

Efficient control pulses for continuous quantum gate families through coordinated re-optimization

Jason D. Chadwick
Department of Computer Science
University of Chicago
Chicago, IL, USA
jchadwick@uchicago.edu

Frederic T. Chong
Department of Computer Science
University of Chicago
Chicago, IL, USA
chong@cs.uchicago.edu

Abstract—We present a general method to quickly generate high-fidelity control pulses for any continuously-parameterized set of quantum gates after calibrating a small number of reference pulses. We find that interpolating between optimized control pulses for different quantum operations does not immediately yield a high-fidelity intermediate operation. To solve this problem, we propose a method to optimize control pulses specifically to provide good interpolations. We pick several reference operations in the gate family of interest and optimize pulses that implement these operations, then iteratively re-optimize the pulses to guide their shapes to be similar for operations that are closely related. Once this set of reference pulses is calibrated, we can use a straightforward linear interpolation method to instantly obtain high-fidelity pulses for arbitrary gates in the continuous operation space.

We demonstrate this procedure on the three-parameter Cartan decomposition of two-qubit gates to obtain control pulses for any arbitrary two-qubit gate (up to single-qubit operations) with consistently high fidelity. Compared to previous neural network approaches, the method is 7.7x more computationally efficient to calibrate the pulse space for the set of all single-qubit gates. Our technique generalizes to any number of gate parameters and could easily be used with advanced pulse optimization algorithms to allow for better translation from simulation to experiment.

I. INTRODUCTION

Quantum circuits, consisting of logical operations on qubits, are typically decomposed into a set of elementary basis operations that are specific to a given hardware device. These basis operations can be individually calibrated through control pulse shaping to achieve high accuracy. However, the space of all quantum operations is much larger than just the hardware basis set, meaning that operations must typically be decomposed into a sequence of basis operations. The capability to perform any operation within a continuous gate set could significantly improve the capabilities of near-term quantum computers, avoiding the runtime and fidelity costs associated with decomposing into basis operations.

Continuous gate sets have been shown to improve fidelity and reduce gate count of quantum circuits [1]–[3]. The implementation of continuous gate sets has been studied on various hardware platforms such as trapped ions [4], [5] and various superconducting platforms [2], [3], [6], [7]. However, each of these studies focused on implementing a specific continuous gate. In this work, we provide a hardware-agnostic method to calibrate any arbitrary continuous gate set. Our approach uses




Fig. 1. a) We seek to efficiently generate high-fidelity control pulses for continuous families of quantum gates. Here, we envision some continuous set of unitary operations between the $CNOT$ and $\sqrt{SWAP}$ operations. The goal is to efficiently obtain control pulses for any arbitrary point in this one-dimensional space of operations while only explicitly calibrating pulses at the two endpoints. b) *Top*: Consider some high-fidelity control pulses that implement $CNOT$ and $\sqrt{SWAP}$ (blue and red). We attempt to obtain a pulse for an intermediate operation (green) through linear interpolation. We find that interpolation yields poor results when the fixed pulses have very different shapes. *Bottom*: However, if we can re-optimize the pulses for $CNOT$ and $\sqrt{SWAP}$ to be more similar to each other (while still performing the correct operations), our simple linear interpolation method can obtain a high-fidelity pulse for the intermediate operation. Our methods generalize to higher-dimensional parameter spaces.

a quantum optimal control solver as a subroutine, iteratively optimizing control pulses to implement a few specific operations from within the gate set, from which we obtain pulses for any other operation.

Quantum control optimization involves shaping hardware control pulses to execute a target operation with high fidelity. Many software packages have been designed to solve quantum control problems for various systems and objectives such as [8]–[16]. If the hardware device of interest is fully controllable, any unitary quantum operation can be implemented using an optimal control solver. Thus, a simple way to support arbitrary continuous gate sets through optimal control is to optimize pulses on an as-needed basis, depending on the operations encountered in the chosen circuit. However, pulse optimizations are computationally expensive, making it infeasible to do this when the input circuit is not necessarily known in advance and long delays in execution time are undesirable.

We address this problem by describing a procedure to pre-calibrate a continuous pulse landscape for a family of quantum operations, from which high-fidelity control pulses for arbitrary operations can be instantly retrieved. We create this landscape by picking a small number of specific operations to directly optimize pulses for, re-optimizing these pulses to be similar to one another, and then defining an interpolation function to retrieve new pulses for any operation in between.

The idea of applying optimal control to continuous spaces is not new. Reference [17] studied the problem of applying a specific operation on a parameterized Hamiltonian to create resilience to varying device errors. Reference [18] examines the single-parameter case in the optimization of a cubic interaction unitary, performing interpolation between two fixed control pulses in a single-parameter space with consistently high fidelity. Reference [19] introduces methods to generate continuous sets of controls to create robustness to experimental deviations in Hamiltonians.

Reference [20] (from which we borrow the term “gate family”) demonstrates the effectiveness of a neural network for generating control pulses for families of parameterized gates. The network takes as inputs the parameters of a specific gate and the time value t , and it outputs the control value(s) $f(t)$. The network is trained over many iterations consisting of batches of randomly-sampled operations from the gate family. Reference [21] uses a similar neural network method, but does not provide detailed data for larger parameter spaces with which to make an objective comparison, so we focus primarily on comparisons with [20] in this work.

Figure 1a shows an example of the problem we solve. We start with two *reference* operations, $CNOT$ and $\sqrt{SWAP}$, with the goal of efficiently obtaining pulses for any intermediate through interpolation. Under the assumption that sufficiently small changes in control pulse values will lead to small changes in the realized quantum operation, linear interpolation between control pulses should yield reasonable results if the initial pulses are similar enough. However, in general, the same quantum operation can be achieved with many different physical pulse sequences; an optimal control solver usually starts from random initial guesses, and, unless heavily constrained, is not expected to reach the same final pulse shape for different initial guesses. Two independently-optimized pulse sequences for different quantum operations could therefore have extremely different pulse shapes (even if the operations themselves are quite similar), meaning that interpolation will generally not provide good results, as observed in the Supplemental Material of [20].

Figure 1b shows the key insight of this work: while optimizing pulses for the two reference operations initially yields poor interpolation accuracy, re-optimizing these pulses to be more similar to each other constrains the pulse solution space and provides significant improvements in average interpolation fidelity. Once the reference pulses are optimized, interpolation can instantly retrieve the pulse for any operation in the parameter space.

In this work, we describe *coordinated re-optimization*, a

method to iteratively optimize reference pulses to be more similar, which enables the use of interpolation to efficiently generate control pulses for continuously-parameterized quantum operations. Our method can achieve equal or better average infidelities as the neural network methods of [20] and [21] using less computation. Additionally, any optimal control algorithm can be used in this framework, allowing for advanced techniques for robust or closed-loop pulse control such as [14]–[16], [22], [23] to directly replace or augment the optimal control unit in our method; these pulse optimization algorithms could significantly improve the translation between simulation and experiment.

We specifically present applications of our method to the set of all single-qubit gates and to the set of two-qubit gate unitaries, which are parameterized by Weyl chamber coordinates, under the belief that these classes of gates are the most important for enabling near-term applications on quantum computers. However, our method can be generalized to any parameterized gate family.

II. OPTIMIZING REFERENCE PULSES

We use the Boulder Opal optimization package [24], [25]. We describe a control pulse by its vector of optimizable variables $\vec{\alpha}$. We use the objective function

$$J = \underbrace{1 - \frac{1}{h^2} \left| \text{Tr} \left(U_{\text{target}}^\dagger U_T(\vec{\alpha}) \right) \right|^2}_{\text{gate infidelity}} + \underbrace{\tilde{\lambda} \sum_k^{n_f} \|\vec{\alpha}_k - \vec{\alpha}_{0,k}\|_2^2}_{\text{Tikhonov regularization}}, \quad (1)$$

with Tikhonov weight

$$\tilde{\lambda} = \frac{\lambda}{n_f \cdot n_p \cdot \alpha_{\text{max}}^2}. \quad (2)$$

Gate infidelity is calculated between the target unitary $U_{\text{target}}^\dagger$ and the unitary $U_T(\vec{\alpha})$ implemented by the pulses. Two-qubit gates have Hilbert space dimension $h = 4$. The Tikhonov regularization term penalizes distance between n_f optimizable pulses, each with n_p optimization variables stored in the vector $\vec{\alpha}_k$, and fixed target vector $\vec{\alpha}_{0,k}$ with user-specified weight λ . It is common to include a similar regularization term in the cost function with $\vec{\alpha}_0 = \vec{0}$ to encourage low-amplitude control pulses, as in [26]. In this work, we found $\lambda = 10^{-2}$ to yield satisfactory results, but this will likely require tuning for different devices and gate families. Choice of $\vec{\alpha}_0$ is key to our method and is discussed in detail in Section II-B.

A. Initial optimization

The first step of our method is to parameterize the gate family of interest and define the volume in parameter space that contains all desired gates, as seen in Figure 2a. To create the interpolation landscape, we first pick reference points on a rectangular grid in this parameter space, where the grid spacing is a tunable parameter that we refer to as the *granularity*. For each point, we obtain the corresponding quantum operation and then use the pulse optimizer to find a valid control pulse sequence that implements this operation.




Fig. 2. Overview of the example of our method provided in this work. More details on this specific example are provided in Section V. a) We use the Cartan decomposition (6) to define a three-dimensional space that contains all two-qubit operations. Our goal is to efficiently generate pulses for any arbitrary point within this continuous space. b) We choose several *reference operations* within the space and generate control pulses that implement these operations via optimal control software. Interpolating between these reference pulses does not necessarily provide a high-quality pulse immediately. c) To solve this problem, we iteratively re-optimize each reference pulse to be more similar to the average of nearby reference pulses, a process we call *neighbor-average re-optimization*. This process is described in Section II. d) Once these reference pulses have been fully calibrated, we can instantly obtain control pulses for any operation in the continuous family by performing linear interpolation between the nearest reference pulses. This process is described in more detail in Section III.

For this initial round of optimization, we choose Tikhonov target $\vec{\alpha}_0 = \vec{0}$ in (1) to encourage low-amplitude pulses.

As can be seen in the examples in later sections, this initial pulse set does not always interpolate well in all parts of parameter space. There is no guarantee that control pulses for similar unitaries will themselves be similar, as observed in [20]; there are often many distinct control sequences that can realize the same unitary operation, yielding near-useless interpolations in between different reference pulses. Tikhonov regularization with $\vec{\alpha}_0 = \vec{0}$ helps by restricting the space of allowed solutions, but does not fully resolve this issue. This motivates a more intelligent approach to finding reference pulses, which we call *neighbor-average re-optimization*.

B. Neighbor-average re-optimization

We generate a simplicial mesh across our reference points using the `scipy.spatial.Delaunay` function [27]. A d -simplex is the simplest d -dimensional polytope, formed by the convex hull of $d + 1$ vertices. For example, a 2-simplex is a triangle and a 3-simplex is a tetrahedron. For a set of $k \geq d + 1$ points in d dimensions, a *simplicial mesh* of contiguous d -simplices can be generated to cover the space between the points¹. Each point is then the vertex of one or more simplices in the mesh. An example of a 3-dimensional simplicial (tetrahedral) mesh is shown in Figure 2c. From this simplicial mesh, we can query the *neighbors* of a given reference point, which are all other points that are connected to the point of interest by an edge.

Starting with the initial naive reference pulses from the initial optimization round, we generate a simplicial mesh over the reference points. For each reference point p_i , we find the

set of neighboring vertices $\eta(p_i)$ and calculate the *neighbor-average* pulse vector

$$\hat{\vec{\alpha}}_i = \frac{1}{|\eta(p_i)|} \sum_{p_j \in \eta(p_i)} \vec{\alpha}_j, \quad (3)$$

which is the average over the pulse variables $\vec{\alpha}_j$ of all neighboring points p_j . We then calculate the Tikhonov penalty between the optimized pulse $\vec{\alpha}_i$ and the neighbor-average pulse $\hat{\vec{\alpha}}_i$. An example of a neighbor-average pulse calculation is shown in Figure 2c.

For each point (ordered by highest neighbor-average Tikhonov penalty), we recalculate $\hat{\vec{\alpha}}_i$ and then use the optimal control solver to re-optimize $\vec{\alpha}_i$ using cost function (1), with initial guess and $\vec{\alpha}_0$ both set to be $\hat{\vec{\alpha}}_i$. The Tikhonov regularization term encourages the optimized pulse $\vec{\alpha}_i$ to be as similar as possible to $\hat{\vec{\alpha}}_i$ while still yielding low infidelity.

This tune-up procedure can be repeated multiple times, with each round of re-optimization building upon the last and steering the reference pulses to be increasingly similar to each other.

III. INTERPOLATING BETWEEN REFERENCE PULSES

Once the reference points have been calibrated, to calculate the linearly-interpolated control pulse for some new point $\tilde{p}$, we construct a weighted sum of reference pulses $\vec{\alpha}_i$ at nearby reference points $\{p_i\}$. We first locate the point within one of the simplices in the Delaunay mesh (the same mesh we previously used in the calibration process). The interpolated pulse vector is then a linear combination of the reference pulses at the vertices of this simplex, weighted by barycentric coordinates.

Barycentric coordinates are a group of $d + 1$ numbers that uniquely identify a point within a d -dimensional simplex, each corresponding to one vertex of the simplex. Barycentric coordinates are uniquely determined by the requirement that

¹Issues may arise in specific configurations of points. For example, if all points are coplanar in a 3D space, no nonzero-volume tetrahedra can be generated. However, in practice, we can reliably generate a simplicial mesh for any reasonable set of points.

the target point is equal to the coordinate-weighted sum of the vertices,

$$\tilde{p} = \sum_{p_i \in S_{\tilde{p}}} b_i p_i, \quad (4)$$

where b_i is the barycentric coordinate of $\tilde{p}$ with respect to vertex p_i and $S_{\tilde{p}}$ is the set of vertices defining the simplex that contains point $\tilde{p}$. The closer the point is to a given vertex, the larger the corresponding coordinate will be². Each coordinate is bounded by $[0, 1]$ and the coordinates sum to 1.

Given these coordinates, our interpolated pulse vector is determined by

$$\vec{\alpha}_{\tilde{p}} = \sum_{p_i \in S_{\tilde{p}}} b_i \vec{\alpha}_i \quad (5)$$

where the vector $\vec{\alpha}_i$ is the optimized reference pulse at vertex p_i . This barycentric coordinate approach is the generalization of linear interpolation to arbitrary dimension. An example of this interpolation process in three dimensions is shown in Figure 2d.

IV. STEPS OF GENERAL METHOD

In this section, we present our general approach to calibrate a continuous set of quantum control pulses.

- 1) *Setup*. If using a model-based optimizer, obtain a model of the device, such as a Hamiltonian. Choose a pulse description $\vec{\alpha}$ (a finite set of variables used to construct each pulse) and a pulse optimization algorithm. Define the parameters of the gate family and determine the space to interpolate within. Sample a number of parameter points $\{p_i\}$ from this space and obtain the corresponding quantum operations.
- 2) *Initial optimization*. Use the optimizer to generate initial reference pulses $\vec{\alpha}_i$ for each reference point.
- 3) *Re-optimization*. For each reference point p_i :
 - a) Calculate new target pulse $\vec{\alpha}_{0,i}$ based on the set of existing reference pulses.
 - b) Re-optimize the reference pulse $\vec{\alpha}_i$ with this target pulse as the initial guess. Use Tikhonov regularization in the pulse optimization cost function to encourage the final pulse to be close to the target pulse.

Repeat step 3 as needed.

- 4) *Interpolation*. Choose an interpolation function $f : (\tilde{p}, \{p_i, \vec{\alpha}_i\}) \rightarrow \vec{\alpha}_{\tilde{p}}$ that calculates interpolated pulse $\vec{\alpha}_{\tilde{p}}$ at parameter-space point $\tilde{p}$ given the set of optimized reference points and pulses $\{p_i, \vec{\alpha}_i\}$.

²If the point lies on a face of the simplex, the only nonzero barycentric coordinates will be those corresponding to the vertices that define that face. This property ensures that interpolations are consistent along face boundaries of adjacent simplices. Along with the fact that barycentric coordinates vary continuously within a simplex, this means that the interpolation space is continuous throughout the simplicial mesh.




Fig. 3. Comparison of pulse shapes for two adjacent reference points before and after several rounds of re-optimization. Point A corresponds to Cartan coordinates $(\frac{1}{2}, 0, 0)$ (the coordinates of *CNOT* and *CZ*) and point B corresponds to coordinates $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$. Optimized controls for the two operations (from the same results as shown in Figure 4) are shown for the $\sigma_x \sigma_x$ control (top) and $\sigma_z^{(1)}$ control (bottom) of the Hamiltonian (8). *Inset*: locations of points A and B in the Weyl chamber. *Left*: the pulse shapes are initially significantly different between points A and B. *Right*: the pulses become far more similar after re-optimization, making interpolation easier, but still retain certain differences in their shapes that account for the differences in the resulting operations.

V. EXAMPLE: THE WEYL CHAMBER OF TWO-QUBIT GATES

We demonstrate this re-optimization and interpolation scheme for the three-parameter Cartan decomposition of two-qubit gates. Any two-qubit quantum logic gate $U \in SU(4)$ can be written as

$$U = k_1 \exp \left(-i \frac{\pi}{2} \sum_{j=x,y,z} t_j \sigma_j^{(1)} \sigma_j^{(2)} \right) k_2 \quad (6)$$

in terms of Pauli matrices $\sigma_x, \sigma_y, \sigma_z$ and Cartan coordinates t_x, t_y, t_z [28]. The operations $k_1, k_2 \in SU(2) \otimes SU(2)$ represent single-qubit gates acting on the qubits independently, i.e. $k_1 = U_0 \otimes U_1$ for some single-qubit gates U_0 and U_1 . Distinct two-qubit gates are referred to as “locally equivalent” or “equivalent up to single-qubit gates” if they have the same Cartan coordinates.

The Weyl chamber is a volume in Cartan coordinate space defined by the equations

$$\begin{aligned} 0 &\leq t_x \leq 1 \\ 0 &\leq t_y \leq \min(t_x, 1 - t_x) \\ 0 &\leq t_z \leq t_y. \end{aligned} \quad (7)$$

The Weyl chamber is shown in Figure 2a. It contains the Cartan coordinates of every two-qubit gate, making it a compelling example to demonstrate our interpolation method; the ability to perform any operation in the Weyl chamber with high fidelity using a single pulse could significantly improve the performance of existing quantum computers. For a more

in-depth introduction to Cartan coordinates we refer the reader to [29]. We note that pulse optimization for target operations within the Weyl chamber has been considered before, e.g. [30]–[32]; however, here we demonstrate optimization for *all* operations in the Weyl chamber together, which to the best of our knowledge has only been considered in [20]. In this work, we demonstrate our interpolation scheme for operations of the form (6) with $k_1 = k_2 = I$ for all parameter points that lie within the Weyl chamber (7).

To evaluate the performance of our approach in comparison to [20], we use the same two-qubit Hamiltonian

$$H(t) = f_{xx}^{\vec{\alpha}}(t)\sigma_x^{(1)}\sigma_x^{(2)} + \sum_{j=1}^2 f_{jy}^{\vec{\alpha}}(t)\sigma_y^{(j)} + f_{jz}^{\vec{\alpha}}(t)\sigma_z^{(j)} \quad (8)$$

for specific parameter values $\vec{\alpha}$. The five control functions $f(t)$ are each restricted to values in $[-1, 1]$, and the pulse duration is fixed to π . We optimize reference pulses using the procedures described previously and then test the interpolation quality by evaluating the infidelities of interpolated pulses for new operations in the space. We perform several rounds of neighbor-average re-optimization and track the improvement in average infidelity.

Each control function $f(t)$ is described as a piecewise-constant function of 20 segments, yielding 100 total optimizable parameters. We set the Tikhonov regularization weight $\lambda = 10^{-2}$. The pulse optimizer is run for a maximum of 50 iterations for each individual optimization (or re-optimization), although it often terminates early upon reaching convergence, especially for later rounds of re-optimization where the pulse shapes do not change as drastically. We found this iteration limit to be sufficient for convergence over multiple re-optimization rounds, and it also avoids unnecessary computation during early rounds of optimization (where the pulse shapes need not be finalized). We use the default Q-CTRL convergence criteria [24].

In Figure 3 we examine two specific control pulses for two different reference points in the Weyl chamber. The two points of interest are connected by an edge in the simplicial mesh, and thus are influenced by each other in the neighbor-average re-optimization process. Before the re-optimization steps, the pulse shapes are noticeably different, which leads to lower-quality interpolation in between the pulses, as can be seen in the left side of Figure 4. After re-optimization, the pulse shapes are much more similar, although some differences remain (the pulses cannot be identical because they perform different operations). The infidelity at the test point directly between these two reference points improved from 3.9×10^{-2} initially to 1.5×10^{-4} after the final round of re-optimization. In the right side of Figure 4, we observe a significant improvement in interpolated pulse accuracy across the entire Weyl chamber.

To investigate the relationship between calibration cost and accuracy, we evaluate the performance of our method for several different reference point granularities and varying numbers of re-optimization rounds, which both influence the total computation time. We evaluate reference point granularities




Fig. 4. Interpolation quality after applying the re-optimization scheme. Axes on the top plots correspond to the three Cartan coordinates of two-qubit gates. Reference pulses are generated by optimal control software for 14 reference points in parameter space. Pulses for any point in the chamber can be obtained by linear interpolation between reference pulses. *Left*: interpolated pulse infidelities at 819 test points (granularity 1/24) after initial optimization of reference points. Black lines connect reference points and define the simplicial mesh that is used for interpolation. Each colored point represents a unitary operation defined by its coordinates. The color of each point is the infidelity of the pulse that is obtained via interpolation. Mean infidelity is $1.3 \pm 1.2 \times 10^{-2}$. *Right*: results after repeatedly re-optimizing each reference point to be near the average of its neighbors. Mean infidelity is improved to $1.0 \pm 0.8 \times 10^{-4}$.




Fig. 5. Different reference point granularities translate to varying amounts of classical computation time needed to optimize all reference pulses. Consecutive points with the same granularity g correspond to subsequent re-optimization rounds, which can further improve average (and maximum) infidelity at the cost of more optimizer iterations. Each optimization iteration corresponds to one system evolution. The points indicated by the plus sign indicate the worst infidelity of any test point.

from $\frac{1}{2}$ to $\frac{1}{8}$ and perform up to 10 rounds of re-optimization. We use the cumulative sum of all iterations used by the optimizer (equivalent to the total number of cost function evaluations) as a measurement of the classical computation time required to calibrate the interpolation landscape. Our results are shown in Figure 5. We find that performance generally improves with more computation, whether this is achieved




Fig. 6. Infidelities at 2197 test points within the parameter space of single-qubit rotations (see Equation (9)). Average and worst-case infidelities are shown in Table I. We observe a region of extremely poor interpolation quality (yellow), indicating that the pulses on either side of this area have very different shapes and actions. Intuitively, the pulses on the two different sides of this boundary take different routes around the Bloch sphere to reach the target unitary. Subsequent rounds of neighbor-average re-optimization reduce the size of this poor interpolation region and eventually eliminate it, resulting in an average interpolation infidelity (across all test points) of $3.5 \pm 6.6 \times 10^{-6}$ after three rounds of re-optimization.

TABLE I
INTERPOLATION INFIDELITIES FOR THE SPACE OF SINGLE-QUBIT ROTATIONS AFTER VARYING ROUNDS OF RE-OPTIMIZATION.

Re-optimization rounds	0	1	2	3	Ref. [20]
Average infidelity	$1.3(2.8) \times 10^{-1}$	$1.2(2.7) \times 10^{-1}$	$2.0(12) \times 10^{-2}$	$3.5(6.6) \times 10^{-6}$	$4(5) \times 10^{-4}$
Maximum infidelity	7.7×10^{-1}	7.8×10^{-1}	7.7×10^{-1}	5.4×10^{-5}	N/A
Total iterations	4680	5445	6072	6654	51200

TABLE II
INTERPOLATION INFIDELITIES FOR THE CARTAN COORDINATE $[0, 1]^3$ BOX AFTER VARYING ROUNDS OF RE-OPTIMIZATION.

Re-optimization rounds	0	1	2	3	4	Ref. [20]
Average infidelity	$1.2(2.5) \times 10^{-1}$	$8.2(21) \times 10^{-2}$	$4.6(16) \times 10^{-2}$	$6.6(60) \times 10^{-3}$	$4.2(6.6) \times 10^{-4}$	$4(4) \times 10^{-4}$
Maximum infidelity	7.1×10^{-1}	6.8×10^{-1}	6.7×10^{-1}	6.1×10^{-1}	3.9×10^{-3}	N/A
Total iterations	6270	11068	15157	18406	21377	51200

by more rounds of re-optimization or a denser sampling of reference points. However, we observe diminishing returns when applying many rounds of re-optimization at the same reference point granularity; simply performing more rounds of re-optimization will not always yield better results, and moving to a higher density of reference points can give a much better performance gain for the same amount of total computation time. Overall, the data suggests that computation time is a relatively accurate predictor of interpolation quality regardless of the specific reference point distribution or number of re-optimization rounds.

VI. COMPARISON TO PREVIOUS METHODS

Finally, we evaluate our approach in comparison to existing methods. We compare to [20] for two gate families. The first is the three-parameter family of single-qubit gates

$$U = \exp \left(-i \frac{\pi}{2} (t_x \sigma_x + t_y \sigma_y + t_z \sigma_z) \right) \quad (9)$$

for $t_x, t_y, t_z \in [0, 1]$. The second family is the three-parameter family of all two-qubit gates that we have studied in previous sections (6). However, instead of restricting to the Weyl chamber, [20] considers a larger volume within this space, specifically $t_x, t_y, t_z \in [0, 1]$. This space is 24 times larger than the Weyl chamber, making it significantly more computationally expensive to calibrate using our method³. We use the same Hamiltonian model (8) (considering only the first qubit for the single-qubit gate case).

A. Results

For the single-qubit gate family, [20] reports average pulse infidelity of $4 \pm 5 \times 10^{-4}$ after using 51,200 system evolutions to train the neural network. Our results for the same space of single-qubit gates are shown in Figure 6 for various numbers of re-optimization rounds. We use a reference point granularity

³The Weyl chamber alone contains the Cartan coordinates of all two-qubit gates, so this larger space merely serves to increase the complexity of the problem without adding any additional benefit.

of 1/4 (125 points). Infidelities are evaluated on a grid in parameter space with granularity 1/12 (2197 test points). Large regions of parameter space appear to interpolate reasonably well after the initial optimization round; we attribute this to the relatively simple operation space, and to the initial Tikhonov regularization that rewards low-amplitude pulses. As shown in Table I, our method reaches an average pulse infidelity of $3.48 \pm 6.63 \times 10^{-6}$ in 6,654 total iterations (where each iteration corresponds to a system evolution), an improvement of over two orders of magnitude in average infidelity with 7.7x lower calibration cost compared to [20].

For the two-qubit case, [20] achieves an average infidelity of $4 \pm 4 \times 10^{-4}$ using 51,200 system evolutions. We use a reference point granularity of 1/6 (343 points) and once again test interpolation quality on 2197 points. Our results for the two-qubit Cartan coordinate space are summarized in Table II. We obtain a similar average infidelity to [20] after two rounds of re-optimization with 2.4x lower calibration cost.

We observe a sharp jump in average infidelity between rounds 2 and 3 in the single-qubit case (Table I), whereas the two-qubit example shows a much more consistent downwards trend, improving by around 1 order of magnitude after each round of re-optimization. This is caused by the yellow region that is visible in the testing data in Figure 6; the very large infidelities at these points significantly affect the overall average. Once this yellow region disappears, the maximum test point infidelity and the average infidelity both decrease by four orders of magnitude.

We also observe a relatively small additional iteration cost of performing subsequent rounds of re-optimization. We attribute this to the fact that it takes fewer iterations for the pulse optimizer to find a good solution when it starts from a good initial guess.

These results demonstrate that our method can obtain similar or better results than previous methods with lower computational costs. We expect that calibration cost of our method could be further reduced in several simple ways. For example, the infidelity convergence threshold of the pulse optimizer could be reduced for intermediate re-optimization rounds, since the reference pulses are re-optimized several times and intermediate results are therefore not required to be perfect. Additionally, some reference points may not need to be re-optimized as many times as others if they already have very low neighbor-average Tikhonov penalty.

VII. ALTERNATIVE APPROACHES TO OPTIMIZATION AND INTERPOLATION

In this paper, we demonstrate our procedure using the neighbor-average Tikhonov regularization method and piecewise-linear interpolation. However, the general idea of creating reference pulses to interpolate between could have many possible implementations. Below, we list several possible modifications of our method.

- *Nonuniform reference point distribution.* Reference points need not be distributed on a rectangular grid; the Delaunay mesh methods will work just as well with a

nonuniform distribution of points. It may be effective to use more densely-packed reference points in certain parts of parameter space to give better interpolations, or less dense points in other areas to improve calibration efficiency.

- *Different guess methods for Tikhonov re-optimization.* In this work, we use the simple neighbor-average method to generate the initial guesses for our Tikhonov regularization re-optimization approach. However, our general approach could also be applied using guesses generated in some other way. For example, each new reference point guess could be a weighted sum of all reference pulses, with weights determined by the trace distance between the unitary operations.
- *Selective re-optimization.* Some reference points likely do not need to be re-optimized if they already have a high degree of similarity to neighboring reference pulses; computation time could be easily improved by only re-optimizing the points that differ from nearby pulses by more than some threshold amount. This would also provide a precise stopping criteria for the calibration process.
- *Higher-order interpolation fits.* Regardless of optimization method, different methods can be used for the interpolation step of the procedure. In this paper, we show results using a simple piecewise-linear interpolation. Alternatively, interpolation could be done using a higher-order model such as a spline or other function.

We do not perform these other interpolation methods for this paper due to the effectiveness of the simpler linear model and lack of a clearly-motivated choice of a more intricate method. However, our general framework is flexible and will work with any optimization or interpolation method, and we anticipate that more complex methods could potentially yield better results, depending on the system and gate family of interest.

VIII. DISCUSSION

Our method improves on the work of [20] by achieving similar pulse accuracy with less total computation required. The procedure has the additional benefits of explainability and modularity. The interpolation space contains known reference points with high-fidelity pulses, with the benefit that specific points can be re-optimized without needing to recalibrate the entire space. Modularity is important in the case where simple optimal control pulses do not transfer well to experimental settings (which is generally expected due to significant device variations over time [33]–[35]); more advanced optimization routines such as data-driven or robust methods [14]–[16], [22], [23] can directly replace the optimal control unit without needing to redesign the rest of the procedure.

The Weyl chamber example demonstrates the potential of this calibration routine, allowing the pulses for any two-qubit gates (up to single-qubit operations) to be instantly obtained with infidelities of $1.0 \pm 0.8 \times 10^{-4}$ after only 2828 iterations of the pulse optimizer. On a real device, a similar calibration

routine could allow quantum hardware to natively support any two-qubit operation directly at the pulse level, which could yield significant improvements in both execution time and accuracy by avoiding expensive gate decompositions.

We have described a method to generate a continuous space of control pulses which are constrained to be as similar as possible to each other. In addition to integrating robust pulse optimization methods, it may be beneficial to recalibrate this control pulse space over time to account for device variation over time, which is a significant problem in current quantum hardware [33]–[35]. We suspect that this recalibration could be significantly less expensive than simply recalibrating every reference pulse due to the high degree of similarity among nearby pulses (e.g. Figure 3).

Our neighbor-average re-optimization procedure makes no assumptions about the structure of the parameter space of interest, and thus may not be optimal; we expect that better-motivated approaches could potentially yield higher-accuracy or more efficient interpolations by taking into account the structure of the gate family of interest. The specific method we use in this paper is easily-implemented and conceptually simple, so we provide it as a baseline for future improvements.

Python code used to generate the results in this work is publicly available in a Github repository [36].

We thank Lennart Maximilian Seifert and Andy Goldschmidt for helpful discussions on quantum control and the interpolation scheme, and André Carvalho and the Q-CTRL team for enabling our use of their software tools in this paper. We also thank the anonymous QCE reviewers for their helpful comments. We used the “Bright” colorblind-friendly plot color scheme from [37]. This work is funded in part by EPiQC, an NSF Expedition in Computing, under award CCF-1730449; in part by STAQ under award NSF Phy-1818914; in part by the US Department of Energy Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing Program; and in part by the NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks (NSF Award 2016136) and in part based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. FTC is Chief Scientist for Quantum Software at Infleqion and an advisor to Quantum Circuits, Inc.

REFERENCES

- [1] N. Lacroix, C. Hellings, C. K. Andersen, A. Di Paolo, A. Remm, S. Lazar, S. Krinner, G. J. Norris, M. Gabureac, J. Heinsoo, A. Blais, C. Eichler, and A. Wallraff, “Improving the Performance of Deep Quantum Optimization Algorithms with Continuous Gate Sets,” *PRX Quantum*, vol. 1, p. 020304, Oct. 2020.
- [2] P. Gokhale, A. Javadi-Abhari, N. Earnest, Y. Shi, and F. T. Chong, “Optimized Quantum Compilation for Near-Term Algorithms with OpenPulse,” in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 186–200, Oct. 2020.
- [3] N. Earnest, C. Tornow, and D. J. Egger, “Pulse-efficient circuit transpilation for quantum applications on cross-resonance-based hardware,” *Physical Review Research*, vol. 3, p. 043088, Oct. 2021.
- [4] A. Sørensen and K. Mølmer, “Quantum Computation with Ions in Thermal Motion,” *Physical Review Letters*, vol. 82, pp. 1971–1974, Mar. 1999.
- [5] A. Sørensen and K. Mølmer, “Entanglement and quantum computation with ions in thermal motion,” *Physical Review A*, vol. 62, p. 022311, July 2000.
- [6] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Z. Chen, K. Satzinger, R. Barends, F. Arute, K. Arya, R. Babbush, D. Bacon, J. Bardin, S. Boixo, D. Buell, B. Burkett, Y. Chen, R. Collins, E. Farhi, A. Fowler, C. Gidney, M. Giustina, R. Graff, M. Harrigan, T. Huang, S. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, P. Klimov, A. Korotkov, F. Kostritsa, D. Landhuis, E. Lucero, J. McClean, M. McEwen, X. Mi, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, M. Niu, A. Petukhov, C. Quintana, N. Rubin, D. Sank, V. Smelyanskiy, A. Vainsencher, T. White, Z. Yao, P. Yeh, A. Zalcman, H. Neven, J. Martinis, and Google AI Quantum, “Demonstrating a Continuous Set of Two-qubit Gates for Near-term Quantum Algorithms,” *Physical Review Letters*, vol. 125, p. 120504, Sept. 2020.
- [7] M. C. Collodo, J. Herrmann, N. Lacroix, C. K. Andersen, A. Remm, S. Lazar, J.-C. Besse, T. Walter, A. Wallraff, and C. Eichler, “Implementation of Conditional Phase Gates Based on Tunable Z Z Interactions,” *Physical Review Letters*, vol. 125, p. 240502, Dec. 2020.
- [8] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, “Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms,” *Journal of Magnetic Resonance*, vol. 172, pp. 296–305, Feb. 2005.
- [9] P. Doria, T. Calarco, and S. Montangero, “Optimal Control Technique for Many-Body Quantum Dynamics,” *Physical Review Letters*, vol. 106, p. 190501, May 2011.
- [10] T. Caneva, T. Calarco, and S. Montangero, “Chopped random-basis quantum optimization,” *Physical Review A*, vol. 84, p. 022326, Aug. 2011.
- [11] N. Leung, M. Abdelhafez, J. Koch, and D. Schuster, “Speedup for quantum optimal control from automatic differentiation based on graphics processing units,” *Physical Review A*, vol. 95, p. 042318, Apr. 2017.
- [12] N. A. Petersson, F. M. Garcia, A. E. Copeland, Y. L. Rydin, and J. L. DuBois, “Discrete Adjoints for Accurate Numerical Optimization with Application to Quantum Control,” Nov. 2020. Number: arXiv:2001.01013 arXiv:2001.01013 [quant-ph].
- [13] S. Günther, N. A. Petersson, and J. 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)*, pp. 88–98, Nov. 2021.
- [14] R.-B. Wu, B. Chu, D. H. Owens, and H. Rabitz, “Data-driven gradient algorithm for high-precision quantum control,” *Physical Review A*, vol. 97, p. 042122, Apr. 2018.
- [15] C. H. Valahu, I. Apostolatos, S. Weidt, and W. K. Hensinger, “Quantum control methods for robust entanglement of trapped ions,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 55, p. 204003, Oct. 2022.
- [16] A. J. Goldschmidt, J. L. DuBois, S. L. Brunton, and J. N. Kutz, “Model predictive control for robust quantum state preparation,” *Quantum*, vol. 6, p. 837, Oct. 2022. arXiv:2201.05266 [quant-ph].
- [17] P. Luchi, F. Turro, V. Amitrano, F. Pederiva, X. Wu, K. Wendt, J. L. DuBois, and S. Quaglioni, “Control optimization for parametric hamiltonians by pulse reconstruction,” May 2022. arXiv:2102.12316 [quant-ph].
- [18] Y. Shi, A. R. Castelli, X. Wu, I. Joseph, V. Geyko, F. R. Graziani, S. B. Libby, J. B. Parker, Y. J. Rosen, L. A. Martinez, and J. L. DuBois, “Simulating nonnative cubic interactions on noisy quantum machines,” *Physical Review A*, vol. 103, p. 062608, June 2021. arXiv:2004.06885 [physics, physics:quant-ph].
- [19] J.-S. Li, J. Ruths, T.-Y. Yu, H. Arthanari, and G. Wagner, “Optimal pulse design in quantum control: A unified computational method,” *Proceedings of the National Academy of Sciences*, vol. 108, pp. 1879–1884, Feb. 2011.
- [20] F. Sauvage and F. Mintert, “Optimal Control of Families of Quantum Gates,” *Physical Review Letters*, vol. 129, p. 050507, July 2022.
- [21] F. Preti, T. Calarco, and F. Motzoi, “Continuous Quantum Gate Sets and Pulse-Class Meta-Optimization,” *PRX Quantum*, vol. 3, p. 040311, Oct. 2022.
- [22] P. H. Leung, K. A. Landsman, C. Figgatt, N. M. Linke, C. Monroe, and K. R. Brown, “Robust 2-Qubit Gates in a Linear Ion Crystal Using a Frequency-Modulated Driving Force,” *Physical Review Letters*, vol. 120, p. 020501, Jan. 2018.

- [23] K. R. Brown, A. W. Harrow, and I. L. Chuang, "Arbitrarily accurate composite pulse sequences," *Physical Review A*, vol. 70, p. 052318, Nov. 2004.
- [24] H. Ball, M. J. Biercuk, A. R. R. Carvalho, J. Chen, M. Hush, L. A. De Castro, L. Li, P. J. Liebermann, H. J. Slatyer, C. Edmunds, V. Frey, C. Hempel, and A. Milne, "Software tools for quantum control: improving quantum computer performance through noise and error suppression," *Quantum Science and Technology*, vol. 6, p. 044011, Oct. 2021.
- [25] Q-CTRL, "Boulder Opal," 2022. <https://q-ctrl.com/boulder-opal>.
- [26] N. Anders Petersson and F. Garcia, "Optimal Control of Closed Quantum Systems via B-Splines with Carrier Waves," *SIAM Journal on Scientific Computing*, vol. 44, pp. A3592–A3616, Dec. 2022.
- [27] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, I. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, A. Vijaykumar, A. P. Bardelli, A. Rothberg, A. Hilboll, A. Kloeckner, A. Scopatz, A. Lee, A. Rokem, C. N. Woods, C. Fulton, C. Masson, C. Häggström, C. Fitzgerald, D. A. Nicholson, D. R. Hagen, D. V. Pasechnik, E. Olivetti, E. Martin, E. Wieser, F. Silva, F. Lenders, F. Wilhelm, G. Young, G. A. Price, G.-L. Ingold, G. E. Allen, G. R. Lee, H. Audren, I. Probst, J. P. Dietrich, J. Silterra, J. T. Webber, J. Slavič, J. Nothman, J. Buchner, J. Kulick, J. L. Schönberger, J. V. de Miranda Cardoso, J. Reimer, J. Harrington, J. L. C. Rodríguez, J. Nunez-Iglesias, J. Kuczynski, K. Tritz, M. Thoma, M. Newville, M. Kümmerer, M. Bolingbroke, M. Tartre, M. Pak, N. J. Smith, N. Nowaczyk, N. Shebanov, O. Pavlyk, P. A. Brodtkorb, P. Lee, R. T. McGibbon, R. Feldbauer, S. Lewis, S. Tygier, S. Sievert, S. Vigna, S. Peterson, S. More, T. Pudlik, T. Oshima, T. J. Pingel, T. P. Robitaille, T. Spura, T. R. Jones, T. Cera, T. Leslie, T. Zito, T. Krauss, U. Upadhyay, Y. O. Halchenko, and Y. Vázquez-Baeza, "SciPy 1.0: fundamental algorithms for scientific computing in Python," *Nature Methods*, vol. 17, pp. 261–272, Mar. 2020.
- [28] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, "Geometric theory of nonlocal two-qubit operations," *Physical Review A*, vol. 67, p. 042313, Apr. 2003.
- [29] G. E. Crooks, "Gates, States, and Circuits," Jan. 2022.
- [30] M. M. Müller, D. M. Reich, M. Murphy, H. Yuan, J. Vala, K. B. Whaley, T. Calarco, and C. P. Koch, "Optimizing entangling quantum gates for physical systems," *Physical Review A*, vol. 84, p. 042315, Oct. 2011.
- [31] P. Watts, J. Vala, M. M. Müller, T. Calarco, K. B. Whaley, D. M. Reich, M. H. Goerz, and C. P. Koch, "Optimizing for an arbitrary perfect entangler. I. Functionals," *Physical Review A*, vol. 91, p. 062306, June 2015.
- [32] M. H. Goerz, G. Gualdi, D. M. Reich, C. P. Koch, F. Motzoi, K. B. Whaley, J. Vala, M. M. Müller, S. Montangero, and T. Calarco, "Optimizing for an arbitrary perfect entangler. II. Application," *Physical Review A*, vol. 91, p. 062307, June 2015.
- [33] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver, "Noise spectroscopy through dynamical decoupling with a superconducting flux qubit," *Nature Physics*, vol. 7, pp. 565–570, July 2011.
- [34] K. W. Chan, W. Huang, C. H. Yang, J. C. C. Hwang, B. Hensen, T. Tanttu, F. E. Hudson, K. M. Itoh, A. Laucht, A. Morello, and A. S. Dzurak, "Assessment of a Silicon Quantum Dot Spin Qubit Environment via Noise Spectroscopy," *Physical Review Applied*, vol. 10, p. 044017, Oct. 2018.
- [35] P. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Y. Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, E. Lucero, J. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. White, S. Boixo, R. Babbush, V. Smelyanskiy, H. Neven, and J. Martinis, "Fluctuations of Energy-Relaxation Times in Superconducting Qubits," *Physical Review Letters*, vol. 121, p. 090502, Aug. 2018.
- [36] J. D. Chadwick, "pulse-interpolation," 2023.
- [37] P. Tol, "Qualitative Color Schemes," Apr. 2021.