



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Speedup for quantum optimal control from automatic differentiation based on graphics processing units

Nelson Leung, Mohamed Abdelhafez, Jens Koch, and David Schuster

Phys. Rev. A **95**, 042318 — Published 13 April 2017

DOI: [10.1103/PhysRevA.95.042318](https://doi.org/10.1103/PhysRevA.95.042318)

Speedup for quantum optimal control from GPU-based automatic differentiation

Nelson Leung,^{1,*} Mohamed Abdelhafez,¹ Jens Koch,² and David Schuster¹

¹*The James Franck Institute and Department of Physics, University of Chicago, Chicago, Illinois 60637, USA*

²*Department of Physics and Astronomy, Northwestern University, Evanston, Illinois 60208, USA*

(Dated: March 10, 2017)

We implement a quantum optimal control algorithm based on automatic differentiation and harness the acceleration afforded by graphics processing units (GPUs). Automatic differentiation allows us to specify advanced optimization criteria and incorporate them in the optimization process with ease. We show that the use of GPUs can speed up calculations by more than an order of magnitude. Our strategy facilitates efficient numerical simulations on affordable desktop computers, and exploration of a host of optimization constraints and system parameters relevant to real-life experiments. We demonstrate optimization of quantum evolution based on fine-grained evaluation of performance at each intermediate time step, thus enabling more intricate control on the evolution path, suppression of departures from the truncated model subspace, as well as minimization of the physical time needed to perform high-fidelity state preparation and unitary gates.

I. INTRODUCTION

The techniques and algorithms used to optimize the control of quantum systems [1–17] and those underlying the field of deep neural networks [18, 19] share a number of common elements. Both areas heavily use linear algebra operations combined with gradient descent optimization. Thus, advanced hardware and software technology recently emerging from the rapid development of machine learning also paves the way for a significant boost of optimal quantum control techniques.

A crucial factor for recent impressive progress in machine learning has been the leveraging of massive parallelism native to graphics processing units (GPUs) [20–24]. Similarly GPUs have been used to accelerate computations in many areas of quantum physics and chemistry [25–31]. Specifically, GPUs are extremely efficient in multiplying very large matrices [32, 33]. Such multiplications also form a central step in the simulation and optimal control of quantum systems. Exploiting this advantageous feature of GPUs, we achieve significant speed improvements in optimizing control schemes for systems at the current frontiers of experimental quantum computation. As the number of qubits in these experiments is increasing [34–36], it becomes increasingly important to take advantage of optimal control techniques. Moreover, recent advances in commercially available electronics – e.g., arbitrary waveform generators enabling base-band synthesis of the entire microwave spectrum [37] – afford new capabilities which quantum optimal control is uniquely well-suited to harness.

There have been numerous theoretical developments of numerical and analytical methods for quantum optimal control (see Ref. 1 for a recent review). The algorithms involved are predominantly based on gradient methods, such as realized in gradient ascent pulse engineering (GRAPE) [2, 3], Krotov algorithms [4–10] or rapid monotonically convergent algorithms [11–17], and are available in several open-source packages, including QuTiP [38, 39], DYNAMO [40], Spinach [41], and SIMPSON [42]. Quantum optimal control has been remarkably successful in determining optimized pulse

sequences [43], designing high-fidelity quantum gates [7, 44–53], and preparing entangled states [54–58].

Optimal control is a versatile concept which can be applied to a vast variety of quantum systems. Typically there is a primary goal (e.g. maximizing fidelity to a target state/unitary), as well as additional constraints/costs associated with specific experimental systems. Examples of such constraints include fixed maximum amplitudes of control pulses [59, 60], maximum overall power of control signals [61], and limited time resolution of arbitrary waveform generators [62]. Further, finite coherence of quantum systems motivates minimizing the overall time needed for reaching the intended state or unitary (time-optimal control) [63]. In certain cases, steering the quantum system among an optimal path (time-dependent target) may be desired [64]. Incorporating new constraints in the optimization process often requires the analytical derivation and implementation of additional contributions to the gradient calculation, and may necessitate significant effort to deploy on large computer clusters. This issue can greatly impede the ability to quickly develop control strategies for new problems.

To overcome these obstacles, we have implemented a quantum optimal control scheme that incorporates constraints via automatic differentiation [65, 66] and utilizes GPUs for boosting computational efficiency. Specifically, automatic differentiation handles the updating of gradient calculations in the backward propagation algorithm [19], and thus eliminates the need to hard-code additional gradient contributions from constraints. For the actual optimal control applications we present in this paper, we find that the computational speed-up from utilizing GPUs becomes significant for Hilbert space sizes exceeding dimensions of the order of one hundred, see Fig. 3. Together, these features allow a quick turnaround for varying optimization constraints and system parameters, rendering this approach invaluable for the study of quantum optimal control. In this paper, we describe the implementation of automatic differentiation, demonstrate its application to quantum optimal control of example systems relevant to quantum computing and quantum optics, and discuss the performance gains achieved by utilizing GPUs.

* nelsonleung@uchicago.edu

II. THEORY

We briefly review the essential idea of quantum optimal control and introduce the notation used throughout our paper. We consider the general setting of a quantum system with intrinsic Hamiltonian \mathcal{H}_0 and a set of external control fields $\{u_1(t), \dots, u_M(t)\}$ acting on the system via control operators $\{\mathcal{H}_1, \dots, \mathcal{H}_M\}$. The resulting system Hamiltonian is given by $H(t) = \mathcal{H}_0 + \sum_{k=1}^M u_k(t)\mathcal{H}_k$. Optimal control theory aims to minimize deviations from a target state or target unitary by appropriate adjustments of the control fields $u_k(t)$. To implement this optimization, the time interval of interest is discretized into a large number N of sufficiently small time steps δt . Denoting intermediate times by $t_j = t_0 + j\delta t$, the Hamiltonian at time t_j takes on the form

$$H_j = \mathcal{H}_0 + \sum_{k=1}^M u_{k,j}\mathcal{H}_k. \quad (1)$$

The control fields subject to optimization now form a set $\{u_{k,j}\}$ of $d = M \cdot N$ real numbers.

The quantum evolution from the initial time $t = t_0$ to time t_j is described by a propagator K_j , decomposed according to

$$K_j = U_j U_{j-1} U_{j-2} \dots U_1 U_0 \quad (2)$$

where

$$U_j = \exp(-iH_j\delta t) \quad (3)$$

is the propagator for the short time interval $[t_j, t_j + \delta t]$. (Here and in the following, we set $\hbar = 1$.) Evolution of a select initial state $|\Psi_0\rangle$ from $t = t_0$ to $t = t_j$ then takes the usual form

$$|\Psi_j\rangle = K_j|\Psi_0\rangle. \quad (4)$$

In the decomposition of K_j , each short-time propagator U_i can be evaluated exactly by matrix exponentiation or approximated by an appropriate series expansion. Propagation methods which go beyond the piecewise-constant approximation for the propagation, can further improve speed and accuracy [48].

Optimization of the discretized control fields $\mathbf{u} \in \mathbb{R}^d$ can be formulated as the minimization of a cost function $C(\mathbf{u})$ where $C : \mathbb{R}^d \rightarrow \mathbb{R}^+$. Table I shows some of the most important cost function contributions used for quantum optimal control. The total cost function is a linear combination of these cost functions, $C = \sum_{\mu} \alpha_{\mu} C_{\mu}$. The weight factors α_{μ} must be determined empirically, and depend on the specific problem and experimental realization at hand. In the following, we discuss these relevant cost function contributions.

A. Important types of cost function contributions

The first cost contribution, $C_1(\mathbf{u})$, is the primary tool for realizing a target unitary K_T , such as a single or multi-qubit

| μ | Cost function contribution | $C_{\mu}(\mathbf{u})$ |
|-------|-------------------------------|--|
| 1 | Target gate infidelity | $1 - \text{tr}(K_T^{\dagger}K_N)/D ^2$ |
| 2 | Target state infidelity | $1 - \langle\Psi_T \Psi_N\rangle ^2$ |
| 3 | Control amplitudes | $ \mathbf{u} ^2$ |
| 4 | Control variations | $\sum_{j,k} u_{k,j} - u_{k,j-1} ^2$ |
| 5 | Occupation of forbidden state | $\sum_j \langle\Psi_F \Psi_j\rangle ^2$ |
| 6 | Evolution time (target gate) | $1 - \frac{1}{N} \sum_j \text{tr}(K_T^{\dagger}K_j)/D ^2$ |
| 7 | Evolution time (target state) | $1 - \frac{1}{N} \sum_j \langle\Psi_T \Psi_j\rangle ^2$ |

TABLE I. Relevant contributions to cost functions for quantum optimal control. Names of contributions indicate the quantity to be *minimized*.

gate. Cost is incurred for deviations between the target unitary and the realized unitary K_N at a given final time t_N . For a system with Hilbert space dimension D , its expression $1 - |\text{tr}(K_T^{\dagger}K_N)/D|^2$ [2] represents the infidelity obtained from the trace distance between the target unitary and the realized unitary. Minimizing this cost function is the principle goal of the quantum control problem.

The second cost function, $C_2(\mathbf{u}) = 1 - |\langle\Psi_T|\Psi_N\rangle|^2$ measures the distance between a desired target state $|\Psi_T\rangle$ and the state $|\Psi_N\rangle$ realized at the final time t_N , as obtained from evolution of a given initial state $|\Psi_0\rangle$. In addition, generalizing C_2 to multiple initial and target states is useful for performing a unitary K_T which is only defined on some subspace H_S of the modeled Hilbert space. Such restriction to a selected subspace is of practical importance whenever a desired unitary is to be implemented within some computational subspace only, as is common for quantum computation applications. There, evolution of higher excited states or auxiliary systems outside the computational subspace is immaterial. Optimal control, then, can be achieved by simultaneous evolution of a set of initial states $\{|\Psi_0^s\rangle\}$ ($s = 1, 2, \dots, S$) that forms a basis of H_S . Optimal control fields are obtained from minimizing the composite state infidelity $C_{2\Sigma}(\mathbf{u}) = 1 - |\frac{1}{S} \sum_s \langle\Psi_T^s|P_S|\Psi_N^s\rangle|^2$ relative to the desired target states $|\Psi_T^s\rangle = K_T|\Psi_0^s\rangle$. (Here, P_S is the projector onto subspace H_S .)

This composite state-transfer cost function when used over a complete basis is equivalent to the gate fidelity, but has several advantages. Most importantly it is more memory efficient requiring only the current state to be stored rather than the whole unitary. In addition, it is very amenable to distributed computing approaches. However, when the unitary transfer matrix can be stored in memory, propagating the full unitary can take advantage of the parallelism of the GPU for smaller problems (see Fig. 3).

Like many optimization problems, quantum optimal control is typically underconstrained. In order to obtain control fields that are consistent with specific experimental capabilities and limitations, it is often crucial to add further constraints on the optimization. Control fields must be realizable in the lab, should be robust to noise, and avoid large control amplitudes and rapid variations based on signal output specifications of instruments employed in experiments. Exceedingly strong control fields may also be problematic due to heat dissipation which may, for instance, raise the temperature inside a

dilution refrigerator. These points motivate the consideration of additional cost function contributions in the following.

One such contribution, $C_3(\mathbf{u}) = |\mathbf{u}|^2$ suppresses large control-field amplitudes globally, and is commonly employed in quantum optimal control studies [2, 59, 60, 67]. (The generalization to more fine-grained suppression of individual control fields is straightforward to implement as well.) Penalizing the L^2 norm of the control fields favors solutions with low amplitudes. It also tends to spread relevant control fields over the entire allowed time window. While C_3 constitutes a “soft” penalty on control-field amplitudes, one may also apply a trigonometric mapping to the amplitudes to effect a hard constraint strictly enforcing fixed maximum amplitudes [68].

The fourth type of contribution to the cost function, $C_4(\mathbf{u}) = \sum_{j,k} |u_{k,j} - u_{k,j-1}|^2$, penalizes rapid variations of control fields by suppressing their (discretized) time derivatives [67]. The resulting smoothing of signals is of paramount practical importance, since any instrument generating a control field has a finite impulse response. If needed, contributions analogous to C_4 which suppress higher derivatives or other aspects of the time dependence of fields can be constructed. Together, limiting the control amplitudes and their time variation filters out high-frequency “noise” from control fields, which is an otherwise common result of less-constrained optimization. Smoother control fields also have the advantage that essential control patterns can potentially be recognized and given a meaningful interpretation.

The contribution $C_5(\mathbf{u}) = \sum_j |\langle \Psi_F | \Psi_j \rangle|^2$ to the cost function has the effect of suppressing occupation of a select “forbidden” state $|\Psi_F\rangle$ (or a set of such states, upon summation) throughout the evolution. The inclusion of this contribution addresses an important issue ubiquitous for systems with Hilbert spaces of large or infinite dimension. In this situation, truncation of Hilbert space is needed or inevitable due to computer memory limitations. (Note that this need even arises for a single harmonic oscillator.) Whenever the evolution generated by optimal control algorithms explores highly excited states, truncation introduces a false non-linearity which can misguide the optimization. Including additional states can, in principle, mitigate this problem, but is generally computationally very expensive. An independent physics motivation for avoiding occupation of highly-excited states consists of spontaneous relaxation in realistic systems: high-energy states are often more lossy (as is usually the case, e.g., for superconducting qubits), and possibly more difficult to model. Active penalization of such states therefore has the two-fold benefit of keeping Hilbert space size at bay, and reducing unwanted fidelity loss from increased relaxation. To address these challenges, we employ an intermediate-time cost function [64, 69]: the cost function C_5 limits leakage to higher states during the entire evolution, and at the same time prevents optimization to be misinformed by artificial non-linearity due to truncation. We note that the efficacy of this strategy is system dependent: it works well, for example, for harmonic oscillators or transmon qubits [70] which have strong selection rules against direct transitions to more distant states, but may be less effective in systems such as the fluxonium circuit [71] where low-lying states have direct matrix elements to many higher

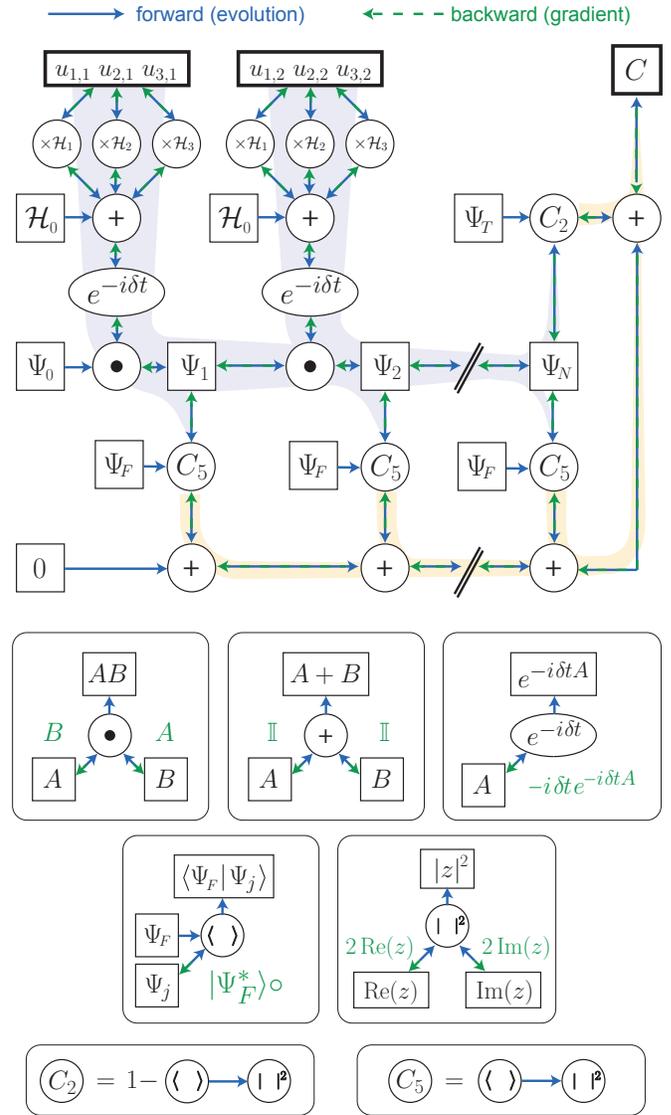


FIG. 1. Computational network graph for quantum optimal control. Circular nodes in the graph depict elementary operations with known derivatives (matrix multiplication, addition, matrix exponential, trace, inner product, and squared absolute value). Backward propagation for matrices proceeds by matrix multiplication, or where specified, by the Hadamard product \circ . In the forward direction, starting from a set of control parameters $u_{k,j}$, the computational graph effects time evolution of a quantum state or unitary, and the simultaneous computation of the cost function C . The subsequent “backward propagation” extracts the gradient $\nabla_{\mathbf{u}} C(\mathbf{u})$ with respect to all control fields by reverse-mode automatic differentiation. This algorithm is directly supported by TensorFlow [72], once such a computational network is specified.

states.

Customarily, algorithms minimizing the cost function $C = \sum_{\mu} \alpha_{\mu} C_{\mu}$ for a given evolution time interval $[t_0, t_N]$ aim to match the desired target unitary or target state at the very end of this time interval. To avoid detrimental effects from decoherence processes during the evolution, it is often beneficial to additionally minimize the gate duration (or state

preparation) time $\Delta t = t_N - t_0$ itself. Instead of running the algorithms multiple times for a set of different Δt , we employ cost function contributions of the form $C_6(\mathbf{u}) = 1 - \frac{1}{N} \sum_j |\text{tr}(K_T^\dagger K_j)/D|^2$ for a target unitary, or $C_7(\mathbf{u}) = 1 - \frac{1}{N} \sum_j |\langle \Psi_T | \Psi_j \rangle|^2$ for a target state, respectively. These expressions penalize deviations from the target gate or target state not only at the final time t_N , but *at every time step*. This contribution to the overall cost function therefore guides the evolution towards a desired unitary or state in as short a time as possible under the conditions set by the other constraints, and thus results in a time-optimal gate.

We will demonstrate the utility of these cost function contributions in the context of quantum information processing in Section V. The versatility of automatic differentiation allows straightforward extension to other contexts such as optimization of quantum observables.

B. Gradient evaluation

The weighted sum of cost functions, $C = \sum_\mu \alpha_\mu C_\mu$, can be minimized through a variety of gradient-based algorithms. Such algorithms are a very popular means of optimization thanks to their good performance and effectiveness in finding optimized solutions for a wide range of problems. At the most basic level, gradient-based algorithms minimize the cost function $C(\mathbf{u})$ by the method of steepest descent, updating the controls \mathbf{u} in the opposite direction of the local cost-function gradient $\nabla_{\mathbf{u}} C(\mathbf{u})$:

$$\mathbf{u}' = \mathbf{u} - \eta \nabla_{\mathbf{u}} C(\mathbf{u}). \quad (5)$$

The choice of the update step size η for the control field parameters \mathbf{u} , plays an important role for the convergence properties of the algorithm. A number of schemes exist which adaptively determine an appropriate step size η in each iteration of the minimization algorithm. Our implementation supports second order methods such as L-BFGS-B [73] as well as gradient descent methods developed for machine learning such as ADAM [74].

For the evaluation of the gradient $\nabla_{\mathbf{u}} C$ we make use of automatic differentiation [65, 66] in reverse-accumulation mode. In brief, this algorithm utilizes the decomposition of the multivariable cost function $C(\mathbf{u})$ into its computational graph of elementary operations (addition, matrix multiplications, trace, etc.), each of which has a known derivative. In reverse-accumulation mode, all partial derivatives of C are evaluated in a recursion from the top level (C) back towards the outermost branches (variables \mathbf{u}) – rather similar to the procedure of obtaining a derivative with pencil and paper. For instance, for the simple function

$$C(\mathbf{u}) = \sin(u_1) + u_1 \cdot \sqrt{u_2} = f_+ \left[\sin(u_1), f_+(u_1, \sqrt{u_2}) \right]$$

one obtains all partial derivatives by a recursion starting with the evaluation of

$$\frac{\partial}{\partial u_j} C = D_1 f_+ \left[\frac{\partial}{\partial u_j} \sin \right] + D_2 f_+ \left[\frac{\partial}{\partial u_j} f_+ \right] = \dots$$

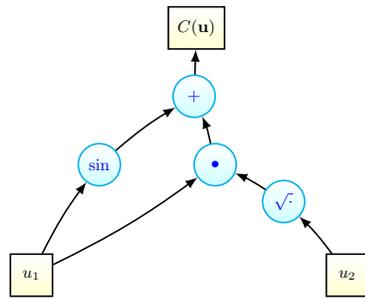


FIG. 2. Sample computational graph for automatic differentiation. Automatic differentiation utilizes the decomposition of the multivariable cost function $C(\mathbf{u})$ into its computational graph of elementary operations, each of which has a known derivative. In reverse-accumulation mode, all partial derivatives of C are evaluated in a recursion from the top level (C) back towards the outermost branches (variables \mathbf{u}).

Here, $D_j f$ stands for the derivative of a multivariable function f with respect to its j -th argument; square brackets denote subsequent numerical evaluation of the enclosed term. (Function arguments are suppressed for brevity.)

Automatic differentiation has become a central tool in machine learning [75], and equally applies to the problem of optimal control of quantum systems. In this approach, the gradient of a set of elementary operations is defined and more complex functions are built as a graph of these operations. The value of the function is computed by traversing the graph from inputs to the output, while the gradient is computed by traversing the graph in reverse via the gradients. This methodology gives the same numerical accuracy and stability of analytic gradients without requiring one to derive and implement analytical gradients specific to each new trial cost function.

All cost functions summarized in table I can be conveniently expressed in terms of common linear-algebra operations. Figure 1 shows the network graph of operations in our software implementation, realizing quantum optimal control with reverse-mode automatic differentiation. For simplicity, the graph only shows the calculation of the cost functions C_2 and C_5 . The cost function contributions C_1 , C_6 , and C_7 are treated in a similar manner. The suppression of large control amplitudes or rapid variations, achieved by C_3 and C_4 , is simple to include, since the calculation of these cost function contributions is based on the control signals themselves and does not involve the time-evolved state or unitary. The host of steps for gradient evaluation is based on basic matrix operations like summation and multiplication.

Reverse-mode automatic differentiation [19] provides an efficient way to carry out time evolution and cost function evaluation by one forward sweep through the computational graph, and calculation of the full gradient by one backward sweep. In contrast to forward accumulation, each derivative is evaluated only once, thus enhancing computational efficiency. The idea of backward propagation is directly related to the GRAPE algorithm for quantum optimal control pioneered by Khaneja and co-workers [2], see Appendix A. While the

original GRAPE algorithm bases minimization exclusively on the fidelity of the final evolved unitary or state, advanced cost functions (such as C_5 through C_7) require the summation of cost contributions from each intermediate step during time evolution of the system. Such cost functions go beyond the usual GRAPE algorithm, but can be included in the more general backward propagation scheme described above. [Appendix A shows analytical forms for gradients for cost functions that are based on time evolution ($\{C_1, C_2, C_5\}$).]

III. IMPLEMENTATION

Our quantum optimal control implementation utilizes the TensorFlow library developed by Google’s machine intelligence research group [72]. This library is open source, and is being extended and improved upon by an active development community. TensorFlow supports GPU and large-scale parallel learning, critical for high-performance optimization. The simple interface to Python allows non-software professionals to implement high-performance machine learning and optimization applications without excessive overhead.

Typical machine-learning applications require most of the same building blocks needed for quantum optimal control. Predefined operations, along with corresponding gradients, include matrix addition and multiplication; matrix traces; and vector dot products. In addition, we have implemented an efficient kernel for approximate evaluation of the matrix exponential and its gradient. Using these building blocks, we have developed a fast and accurate implementation of quantum optimal control, well-suited to deal with a broad range of engineered quantum systems and realistic treatment of capabilities and limitations of control fields.

In common applications of quantum optimal control, time evolving the system under the Schrödinger equation – more specifically, approximating the matrix exponential for the propagators U_j at each time step t_j – requires the biggest chunk of computational time. Within our matrix-exponentiation kernel, we approximate $e^{-iH_j\delta t}$ by series expansion, taking into account that the order of the expansion plays a crucial role in maintaining accuracy and unitarity. The required order of the matrix-exponential expansion generally depends on the magnitude of the matrix eigenvalues relative to the size of the time step. General-purpose algorithms such as `expm()` in Python’s SciPy framework accept arbitrary matrices M as input, so that the estimation of the spectral radius or matrix norm of M , needed for choosing the appropriate order in the expansion, often costs more computational time than the final evaluation of the series approximation itself. Direct series expansion with only a few terms is sufficient for $H_j\delta$ with spectral radius smaller than 1. In the presence of large eigenvalues, series convergence is slow, and it is more efficient to employ an appropriate form of the “scaling and squaring” strategy, based on the identity

$$\exp M = \left[\exp \left(\frac{M}{2^n} \right) \right]^{2^n}, \quad (6)$$

which reduces the spectral range by a factor of 2^n at the cost

of recursively squaring the matrix n times [76]. Overall, this strategy leads to an approximation of the short-time propagator of the form

$$U_j \approx \left[\sum_{k=0}^p \frac{(-iH_j\delta t/2^n)^k}{k!} \right]^{2^n}, \quad (7)$$

based on a Taylor expansion truncated at order p . Computational performance could be further improved by employing more sophisticated series expansions [77, 78] and integration methods [79].

As opposed to the challenges of general-purpose matrix exponentiation, matrices involved in a specific quantum control application with bounded control field strength ($iH_j\delta t$), will typically exhibit similar spectral radii. Thus, rather than attempting to determine individual truncation levels p_j , and performing scaling-and-squaring at level n_j in each time step t_j , we make a conservative choice for global p and n at the beginning and employ them throughout. This simple heuristic speeds up matrix exponentiation over the default SciPy implementation significantly, primarily due to leaving out the step of spectral radius estimation.

By default, automatic differentiation would compute the gradient of the approximated matrix exponential via back-propagation through the series expansion. However, for sufficiently small spectral radius of M , we may approximate [2]

$$\frac{d}{dx} e^{M(x)} \approx M'(x) e^{M(x)}, \quad (8)$$

neglecting higher-order corrections reflecting that $M'(x)$ and $M(x)$ may not commute. (Higher-order schemes taking into account such additional corrections are discussed in Ref. 3.) Equation (8) simplifies automatic differentiation: within this approximation, only the same matrix exponential is needed for the evaluation of the the gradient. We make use of this in a custom routine for matrix exponentiation and gradient-operator evaluation, further improving the speed and memory performance.

The TensorFlow library currently has one limitation relevant to our implementation of a quantum optimal control algorithm. Operators and states in Hilbert space have natural representations as matrices and vectors which are generically complex-valued. TensorFlow, designed primarily for neural network problems, has currently only limited support for complex matrices. For now, we circumvent this obstacle by mapping complex-valued matrices to real matrices via the isomorphism $H \xrightarrow{\cong} \mathbb{1} \otimes H_{\text{re}} - i\sigma_y \otimes H_{\text{im}}$, and state vectors $\vec{\Psi} \xrightarrow{\cong} (\vec{\Psi}_{\text{re}}, \vec{\Psi}_{\text{im}})^t$. Here, $\mathbb{1}$ is the 2×2 unit matrix and σ_y one of the Pauli matrices. Real and imaginary part of the matrix H are denoted by $H_{\text{re}} = \text{Re } H$ and $H_{\text{im}} = \text{Im } H$, respectively; similarly, real and imaginary parts of state vectors are $\vec{\Psi}_{\text{re}} = \text{Re } \vec{\Psi}$ and $\vec{\Psi}_{\text{im}} = \text{Im } \vec{\Psi}$. Written out in explicit block matrix form, this isomorphism results in

$$H \vec{\Psi} \xrightarrow{\cong} \begin{pmatrix} H_{\text{re}} & -H_{\text{im}} \\ H_{\text{im}} & H_{\text{re}} \end{pmatrix} \begin{pmatrix} \vec{\Psi}_{\text{re}} \\ \vec{\Psi}_{\text{im}} \end{pmatrix}, \quad (9)$$

rendering all matrices and vectors real-valued. For the Hamiltonian matrix, this currently implies a factor two in memory

cost (due to redundancy of real and imaginary part entries). There are promising indications that future TensorFlow releases may improve complex-number support and eliminate the need for a mapping to real-valued matrices and vectors.

IV. PERFORMANCE BENCHMARKING

Obtaining a fair comparison between CPU-based and GPU-based computational performance is notoriously difficult [80]. We attempt to provide a specific comparison under a unified computation framework. TensorFlow allows for straightforward switching from running code on a CPU to a GPU. For each operation (matrix multiplication, trace, etc.), we use the default CPU/GPU kernel offered by TensorFlow. Note that properly configured, TensorFlow automatically utilizes all threads available for a given CPU, and GPU utilization is found to be near 100%. Not surprisingly, we observe that the intrinsic parallelism of GPU-based matrix operations allows much more efficient computation beyond a certain Hilbert space size, see Fig. 3.

In this example, we specifically inspect how the computational speed scales with the Hilbert space dimension when optimizing an n -spin Hadamard transform gate and n -spin GHZ state preparation for a coupled chain of spin-1/2 systems presented in Section V D. (Details of system parameters are described in the same section.) We benchmark the average runtime for a single iteration for various spin-chain sizes and, hence, Hilbert space dimensions. We find that the GPU quickly outperforms the CPU in the unitary gate problem, even for a moderate system size of ~ 100 basis states. For optimization of state transfer, we observe that speedup from GPU usage, relative CPU performance, sets in for slightly larger system sizes of approximately ~ 300 basis states.

The distinct thresholds for the GPU/CPU performance gain stem from the different computational complexities of gate vs. state-transfer optimization. Namely, optimizing unitary gates requires the propagation of a unitary operator (a matrix), involving matrix-matrix multiplications, while optimizing state transfer only requires the propagation of a state (a vector), involving only matrix-vector multiplications:

$$U_j|\Psi\rangle \approx \sum_{k=0}^p \frac{(-i\delta t)^k}{k!} (H_j \dots (H_j(H_j|\Psi))), \quad (10)$$

Computing the matrix-vector multiplication is generally much faster than computing the matrix exponential itself [81]. For an n -dimensional matrix, the computation of the matrix exponential involves matrix-matrix multiplication, which scales as $O(n^3)$. The computation of state transfer only involves matrix-vector multiplication, which scales as $O(n^2)$ [or even $O(n)$ for sufficiently sparse matrices].

For optimization of the Hadamard transform as well as the GHZ state preparation, we observe a 19-fold GPU speedup for a 10-qubit system (Hilbert space dimension of 1,024) in the former case, and a 6-fold GPU speedup for an 11-qubit system (Hilbert space dimension of 2,048) in the latter case. Since matrix operations are the most computationally intensive task in our software, this speedup is comparable to other

GPU application studies that heavily use matrix operation [20–24, 80, 82]. We emphasize that these numbers are indicative of overall performance trends, but detailed numbers will certainly differ according to the specific system architecture in place. The CPU model we used was an Intel[®] Core[™] i7-6700K CPU @ 4.00 GHz, and the GPU model was an NVIDIA[®] Tesla[®] K40c. In this study, all computations are based on dense matrices. Since most physically relevant Hamiltonians are sparse (evolution generally affects sparsity, though), future incorporation of sparse matrices may further improve computation speed for both CPU and GPU [83, 84].

V. SHOWCASE APPLICATIONS

In this last section, we present a set of example applications of experimental relevance. The first application demonstrates the importance of cost functions suppressing intermediate occupation of higher-lying states during time evolution, as well as cost functions accounting for realistic pulse shaping capabilities. In a second application, we show how the cost function C_6 can yield high-fidelity state transfer within a reduced time interval. Third, we discuss the application of Schrödinger-cat state preparation – an example from the context of quantum optics and of significant interest in recent schemes aiming at quantum information processing based on such states [67, 85, 86]. This application combines considerable system size with a large number of time steps, and utilizes most of the cost functions discussed in Section II A. In the fourth application, we demonstrate the algorithm performance in finding optimal solutions for GHZ state preparation and implementation of a Hadamard transform gate in a chain of qubits with a variable number of qubits. We use either the Adam [74] or L-BFGS-B optimization algorithm [73] for pulse optimization, and achieve a minimum fidelity of 99.9% in all of our following examples.

A. CNOT gate for two transmon qubits

In the first example, we study realization of a 2-qubit CNOT gate in a system of two coupled, weakly anharmonic transmon qubits. For each transmon qubit ($j = 1, 2$) [70], we take into account the lowest two states spanning the qubit computational space, as well as the next three higher levels. The system Hamiltonian, including the control fields $\{\Omega_{x_1}(t), \Omega_{x_2}(t), \Omega_{z_2}(t)\}$, then reads

$$\begin{aligned} H(t) = & \sum_{j=1,2} \left[\omega_j b_j^\dagger b_j + \frac{1}{2} \alpha_j b_j^\dagger b_j (b_j^\dagger b_j - 1) \right] \\ & + J(b_1 + b_1^\dagger)(b_2 + b_2^\dagger) \\ & + \Omega_{x_1}(t)(b_1 + b_1^\dagger) + \Omega_{x_2}(t)(b_2 + b_2^\dagger) + \Omega_{z_2}(t)b_2^\dagger b_2. \end{aligned} \quad (11)$$

Here, the ladder operators b_j , and b_j^\dagger are truncated at the appropriate level. (The qubit frequencies $\omega_j/2\pi$ are chosen as 3.5 and 3.9 GHz, respectively; both transmons have an anharmonicity of $\alpha/2\pi = -225$ MHz; and the qubit-qubit coupling

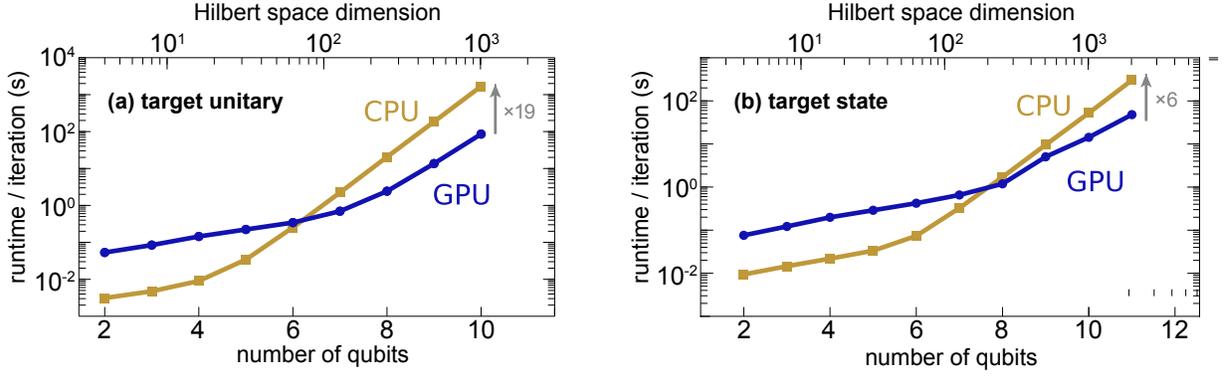


FIG. 3. Benchmarking comparison between GPU and CPU for (a) a unitary gate (Hadamard transform), and (b) state transfer (GHZ state preparation). Total runtime per iteration scales linearly with the number of time steps. For unitary-gate optimization, the GPU outperforms the CPU for Hilbert space dimensions above ~ 100 . For state transfer, GPU benefits set in slightly later, outperforming the CPU-based implementation for Hilbert space dimensions above ~ 300 . The physical system we consider, in this case, is an open chain of N spin-1/2 systems with nearest neighbor $\sigma_z \sigma_z$ coupling, and each qubit is controlled via fields Ω_x and Ω_y .

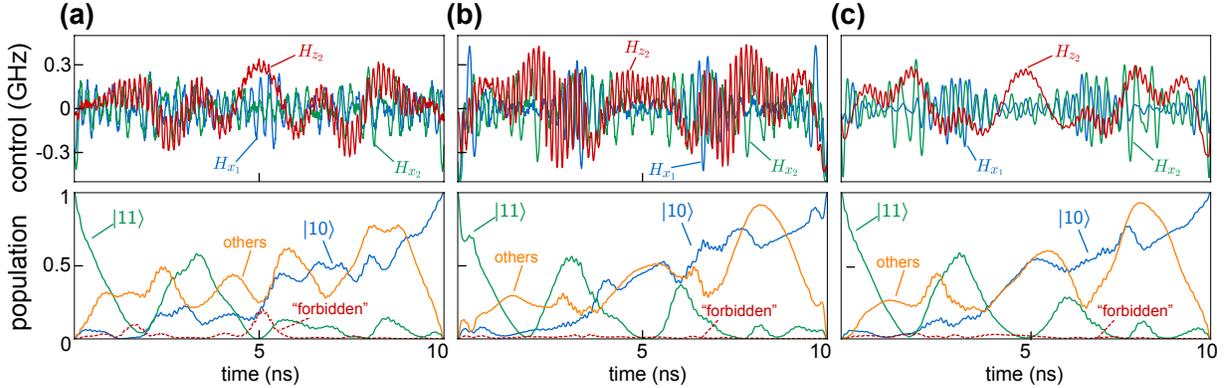


FIG. 4. Control pulses and evolution of quantum state population for a CNOT gate acting on two transmon qubits, (a) only targeting the desired final unitary, (b) employing an additional cost function suppressing occupation of higher-lying states (C_5), and (c) including additional pulse-shape cost functions (C_3, C_4). Here, only the evolution of state $|11\rangle$ is shown, as the evolution of state $|11\rangle$ is most susceptible to the occupation of higher level states. In all three cases, the CNOT gate converged to a fidelity of 99.9%. The results differ in important details: in (a), both high-frequency “noise” on the control signals and significant occupation of “forbidden” states (3rd and 4th excited transmon level), shown as dashed red line, are visible throughout the evolution; in (b), forbidden-state occupation is suppressed at each time step during evolution; in (c), this suppression is maintained and all control signals are smoothed. The maximum occupation of forbidden states is reduced from $\sim 20\%$ in (a) to $\sim 3\%$ in (b) and (c). The population of “others” states (non- $|11\rangle$, $|10\rangle$ or “forbidden”) is also shown for completeness. For demonstration purposes, all three examples use the same gate duration of 10 ns, despite being subject to different constraints. In practice, one would typically increase the gate time for a more constrained problem to achieve the best result in maximizing gate fidelity, minimizing forbidden state occupation, and achieving a realistic control signal.

strength used in the simulation is $J/2\pi = 100$ MHz.) Consistent with recent circuit QED experiments utilizing classical drives as well as parametric modulation, we investigate control fields acting on $H_{x_1} = b_1 + b_1^\dagger$, $H_{x_2} = b_2 + b_2^\dagger$, and $H_{z_2} = b_2^\dagger b_2$.

We next optimize control fields for the realization of a CNOT gate, with transmon qubit $j = 1$ acting as the control qubit. Our control-field optimization reaches a prescribed fidelity of 99.9% for a 10 ns gate duration in all cases, as seen in Fig. 4. Results shown in Fig. 4(a) are obtained with the standard target-gate infidelity cost function (C_1) only. It is evident that the solution encounters two issues: the occupation of the 3rd and 4th excited transmon level (“forbidden”) is

significant, and control fields are polluted by high-frequency components. Including a cost function contribution of type C_5 succeeds in strongly suppressing occupation of higher levels, see Fig. 4(b). This both reduces exposure to increased relaxation rates and ensures that the evolution is minimally influenced by our numerical truncation of Hilbert space. In the final improvement step, shown in Fig. 4(c), our optimization additionally suppresses excessive control amplitudes and derivatives via cost contributions of type C_3 and C_4 . The inclusion of these terms in the overall cost lessens superfluous “noise” in the control signals, and also helps improve convergence of the algorithm – without reducing the achieved target-gate fidelity.

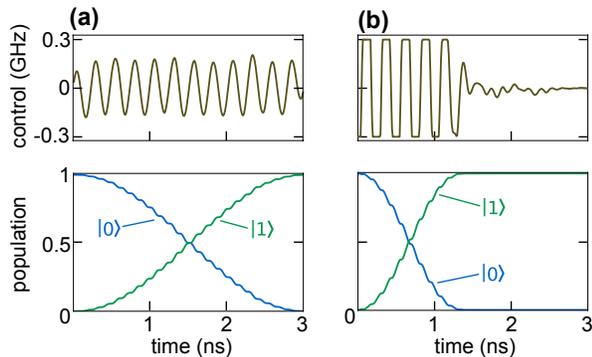


FIG. 5. Minimizing evolution time needed for a high-fidelity state transfer. (a) No time-optimal award function. (b) With time-optimal award function. (a) Without penalty for the time required for the gate, the control field spreads across the entire given time interval. (b) Once evolution over a longer time duration is penalized with a contribution of type C_6 or C_7 (see table I), the optimizer achieves target state preparation in a shorter time, without loss of fidelity.

B. Reducing duration of $|0\rangle$ to $|1\rangle$ state transfer

In this second example, we illustrate the use of cost function contributions (types C_6 , C_7) in minimizing the time needed to perform a specific gate or prepare a desired state. To this end, we consider a two-level spin qubit ($\omega/2\pi$: 3.9 GHz). The system and control Hamiltonians combined are taken to be

$$H = \frac{\omega}{2}\sigma_z + \Omega(t)\sigma_x. \quad (12)$$

We allow for a control field acting on the qubit σ_x degree of freedom, and constrain the maximum control-field strength $\Omega_{\max}/2\pi$ to 300 MHz. When the evolution time needed to perform the state transfer is fixed (rather than subject to optimization itself), we observe that control fields generically spread across the prescribed gate duration time. The desired target state is realized only at the very end of the allowed gate duration. When we incorporate a C_6 or C_7 -type cost contribution, the optimal control algorithm also aims to minimize the overall gate duration, so as to realize the target unitary or state in as short a time as possible, given other active constraints. In our example, this reduces the time for a state transfer from 3 ns to less than 1.5 ns, see Fig. 5. We note that it is further possible to adaptively change the overall simulation time during optimization. For instance, if further optimization was desired in the case of Fig. 5(b), then the simulation time interval could be adaptively reduced to ~ 1.5 ns – resulting in a significant cutback in overall computation time.

C. Generating photonic Schrödinger cat states

As an example of quantum state transfer, we employ our optimal control algorithm to the task of generating a photonic Schrödinger-cat state. The system we consider to this end is a realistic, and recently studied [85, 86] circuit QED setup,

consisting of a transmon qubit capacitively coupled to a three-dimensional microwave cavity. External control fields are restricted to the qubit. Working in a truncated subspace for the transmon (limiting ourselves to levels with energies well below the maximum of the cosine potential), the full Hamiltonian describing the system is

$$H(t) = \omega_q b^\dagger b + \frac{1}{2}\alpha b^\dagger b(b^\dagger b - 1) + \omega_r a^\dagger a + g(a + a^\dagger)(b + b^\dagger) + \Omega_x(t)(b + b^\dagger) + \Omega_z(t)b^\dagger b \quad (13)$$

Here, a and b are the usual lowering operators for photon number and transmon excitation number, respectively. The frequencies $\omega_q/2\pi = 3.5$ GHz and $\alpha/2\pi = -225$ MHz denote the transmon 0-1 splitting and its anharmonicity. The frequency of the relevant cavity mode is taken to be $\omega_r/2\pi = 3.9$ GHz. Qubit and cavity are coupled, with a strength parameterized by $g/2\pi = 100$ MHz. In our simulation, the overall dimension is $154 = (7 \text{ transmon levels}) \times (22 \text{ resonator levels})$. Note that the rotating wave approximation is not applied in order to reflect the capabilities of modern arbitrary waveform generation.

The state-transfer task at hand, now, is to drive the joint system from the zero-excitation state $|0\rangle_q \otimes |0\rangle_r$ (the ground state if counter-rotating terms in the coupling are neglected) to the photonic cat state $|0\rangle_q \otimes |\text{cat}\rangle_r$. Here, the cat state in the resonator corresponds to a superposition of two diametrically displaced coherent states: $|\text{cat}\rangle_r = \frac{1}{\sqrt{2}}(|\lambda\rangle + |-\lambda\rangle)$. Coherent states are defined in the usual way as normalized eigenstates of the photon annihilation operator a , and correspond to displaced vacuum states $|\lambda\rangle = e^{-|\lambda|^2/2} e^{\lambda a^\dagger} |0\rangle$. The cat state $|\text{cat}\rangle_r$ is approximately normalized for sufficiently large λ . As our concrete target state, we choose a cat state with amplitude $\lambda = 2$ (normalization error of $\sim 0.03\%$). The state transfer is to be implemented by control fields $\Omega_x(t)$ and $\Omega_z(t)$ acting on the transverse and longitudinal qubit degrees of freedom, $H_x = (b + b^\dagger)$ and $H_z = b^\dagger b$, respectively. Matching experimental realizations of multi-mode cavity QED systems[87], we do not allow for any direct control of the cavity degrees of freedom.

This state-transfer problem provides an excellent test for an optimal control algorithm. It incorporates the simultaneous challenges of a large number of time steps (8,000), a considerable evolution time (40 ns), and the application of most of the cost functions we discussed in Sect. II A and summarized in Table I. Specifically, in addition to minimizing the target state infidelity (C_2), we penalize occupation of transmon levels 3 to 6 and cavity levels 20 and 21 (C_5) to avoid artifacts from truncation, and penalize control variations (C_4) [88]. Results from the optimization are presented in 6, which shows the control-field sequence, as well as the induced state evolution. At the end of the 40 ns time interval, the control fields generate the desired cat state with a fidelity of 99.9%. The maximum populations at the truncation levels of transmon and cavity are $\sim 6 \times 10^{-6}$ and $\sim 7 \times 10^{-10}$, respectively. We independently confirm convergence with respect to truncation by simulating the obtained optimized pulse for enlarged Hilbert space (8 transmon and 23 cavity levels), and find that the evolution continues to reach the target state with 99.9% fidelity.

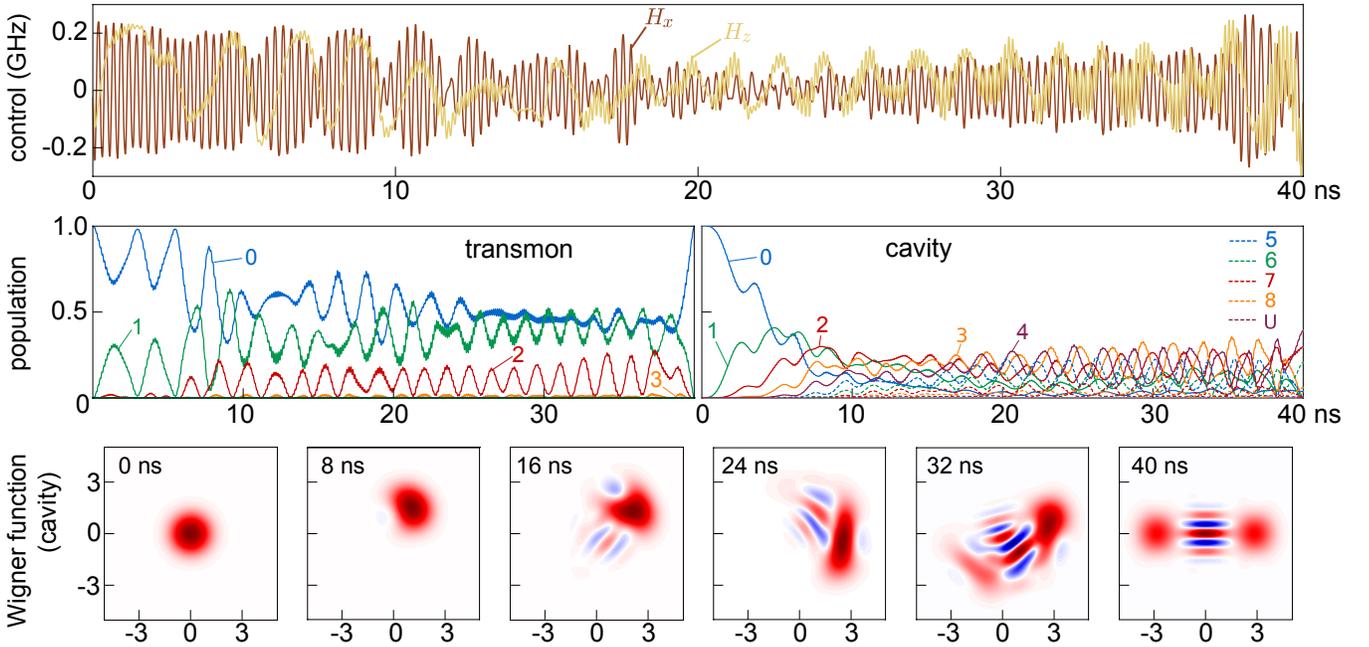


FIG. 6. Cat state generation. Control pulse, state evolution in Fock basis, and Wigner function tomography of the cavity evolution. Photonic cat state generation is shown as a test of state transfer, challenging the quantum control algorithm with a system of considerable size, large number of required time steps, and inclusion of multiple types of cost function. The desired Schrödinger cat state in the resonator is created indirectly, by applying control fields to a transmon qubit coupled to the resonator, and reached within a prescribed evolution time of 40 ns with a fidelity of 99.9%. (Note that occupation of transmon level 4, 5, 6 remains too small to be visible in the graph.)

D. Hadamard transform and GHZ state preparation

We present a final set of examples illustrating the algorithm performance for increasing system size. To that end, we consider a coupled chain of N qubits, or spin-1/2 systems. We assume that all spins are on-resonance in the multiple-rotating frame. This system is described by the Hamiltonian

$$H(t) = \sum_{n=1}^N \left[\Omega_x^{(n)}(t) \sigma_x^{(n)} + \Omega_y^{(n)}(t) \sigma_y^{(n)} + J \sigma_z^{(n)} \sigma_z^{(n+1)} \right], \quad (14)$$

where the coupling term is understood to be dropped for the final summand ($n = N$). The qubit-qubit coupling strength is fixed to $J/2\pi = 100$ MHz. Each qubit (n) is controlled via fields $\Omega_x^{(n)}$ and $\Omega_y^{(n)}$, with a maximum allowed drive strength of $\Omega_{x,y}^{(n)}/2\pi = 500$ MHz.

As a first optimization task, we search for control fields to implement the unitary realizing a Hadamard transform, commonly used in various quantum algorithms. The gate time we allow for the Hadamard transform is $(2N)$ ns, simulated with $10N$ time steps. Figure 7(a) shows the number of iterations and wall-clock time required to converge to the desired 99.9% process fidelity. For the same spin-chain system, we have also employed our code to optimize control fields for transferring the system ground state to a maximally entangled GHZ state. The overall time and time steps we allow for the GHZ state preparation is identical to that used for the Hadamard transform gate. Figure 7(b) shows the number of iterations neces-

sary and the total wall-clock time spent for reaching convergence to a result with 99.9% state fidelity. For both examples, we employed computation on either CPU or GPU, depending which one is faster. (This performance benchmarking data was shown in Section IV). We note that, when using a modest desktop PC with graphics card, optimal control problems for small Hilbert space size converge within seconds. For a 10-qubit Hadamard gate (Hilbert space dimension of 1,024) or 11-qubit GHZ state (Hilbert space dimension of 2048), it takes ~ 1 day to obtain a solution meeting the 99.9% fidelity threshold. The total wall-clock time could likely have been reduced significantly by appropriate choice of optimizer, hyperparameters, and/or initial control fields. In the case of spin-chain system, like many quantum information systems, as the number of elements increase, not only does the Hilbert space grow exponentially, the number of control fields and the required number of time steps also get larger. This further increases the complexity of the problem.

VI. CONCLUSION

In conclusion, we have presented a quantum optimal control algorithm harnessing two key technologies that enable fast and low-overhead numerical exploration of control signal optimization. First, we have demonstrated that automatic differentiation can be leveraged to facilitate effortless inclusion of diverse optimization constraints, needed to obtain realistic control signals tailored for the specific experimental capabili-

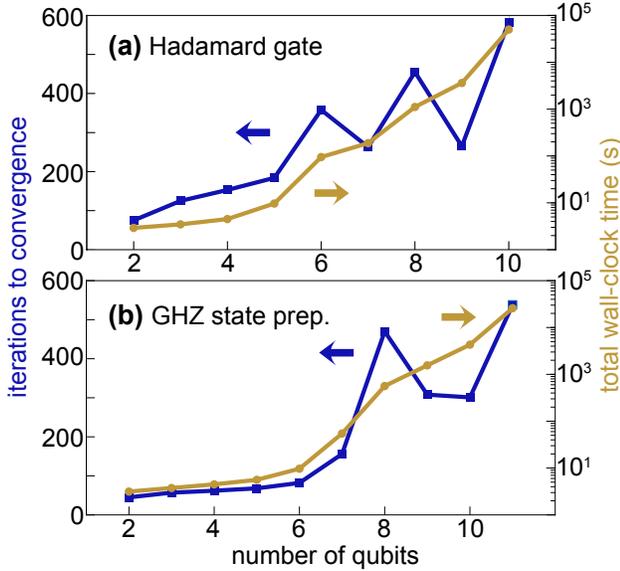


FIG. 7. Performance of optimal control algorithm as a function of qubit number for (a) a Hadamard transform gate, and (b) GHZ state preparation. As system size increases, total time and number of iterations for the algorithm grow rapidly. The larger number of control parameters and complexity of the target state add to the challenge of quantum optimal control for systems with many degrees of freedom.

ties at hand. Automatic differentiation dramatically lowers the overhead for adding new cost functions, as it renders analytical derivations of gradients unnecessary. For illustration, we have presented concrete examples of optimized unitary gates and state transfer, using cost functions relevant for applications in superconducting circuits and circuit QED. We emphasize that this is but one instance within a much larger class of quantum systems for which optimal control is instrumental, and the methods described here are not limited to the specific examples shown in this paper.

The second key technology we have incorporated is the implementation of GPU-based numerical computations, which offers a significant speedup relative to conventional CPU-based code. The use of the TensorFlow library [72] hides the low-level details of GPU acceleration, allowing implementation of new cost functions at a high level. The reduction in computational time will generally depend on a number of factors including system type, Hilbert space size, and the specific hardware employed by the user. We observe that runtime speedup by an order of magnitude is not unusual when using a standard desktop PC, enabling the development of sophisticated quantum control without enormous investments into powerful computing equipment. The underlying libraries also have support for high-performance distributed computing systems for larger optimizations. Our software implementation is open source and can be downloaded at: github.com/SchusterLab/quantum-optimal-control.

The increased efficiency and ease of optimal quantum control due to the employment of GPUs and automatic differentiation makes our work valuable to a broad range of research.

Future work will address sparse-matrix implementations, as well as the deployment of adaptive step size and Runge-Kutta methods for time evolution.

ACKNOWLEDGMENTS

The authors thank J. Kutasov for help in comparing optimization performance with existing frameworks as well as T. Berkelbach and D. Mazziotti for discussions. We gratefully acknowledge the support of NVIDIA[®] Corporation through the donation of the Tesla[®] K40 GPU used for this research, and support from the David and Lucille Packard Foundation. This material is based upon work supported by the Department of Defense under contracts H98230-15-C0453 and W911NF-15-1-0421.

Appendix A: Analytical gradients and algorithms

In the following, we outline the analytical calculation of gradients for cost functions such as those summarized in Table I. We stress that our automatic-differentiation implementation evaluates these gradients autonomously, without the need of these analytical derivations or hard-coding any new gradients. The following derivations are thus merely intended as illustrations of the gradients calculated without user input by means of automatic differentiation.

For a systematic treatment of the different types of cost functions, we note that most cost functions involve an absolute value squared of an inner product between target and final states or target and final unitaries (Hilbert-Schmidt inner product). To obtain the gradients of expressions such as $C_1(\mathbf{u}) = 1 - |\text{tr}(K_T^\dagger K_N)|^2$ with respect to the control parameters, we note that control parameters enter via the final states or unitaries through the evolution operators, $K_N = U_N(\mathbf{u})U_{N-1}(\mathbf{u}) \cdots U_1(\mathbf{u})U_0$. To streamline our exposition, we first summarize multiple matrix-calculus relations of relevance.

Consider two complex-valued matrices A and B , compatible in row/column format such that the matrix product AB is defined. Then, one finds

$$\frac{\partial \text{tr}(AB)}{\partial B_{ji}} = \frac{\partial (A_{nm} B_{mn})}{\partial B_{ji}} = A_{ij}. \quad (\text{A1})$$

Throughout this appendix, we use Einstein convention for summation, and follow the Jacobian formulation (also known as numerator layout) for derivatives with respect to matrices. We will further encounter expressions of the following form, involving a third matrix C of the same dimensions as B^t :

$$\begin{aligned} \text{tr} \left[\frac{\partial [|\text{tr}(AB)|^2]}{\partial B} C \right] &= \frac{\partial [\text{tr}(AB) \text{tr}(AB)^*]}{\partial B_{ji}} C_{ji} \\ &= \frac{\partial \text{tr}(AB)}{\partial B_{ji}} \text{tr}(AB)^* C_{ji} = A_{ij} \text{tr}(AB)^* C_{ji} \\ &= \text{tr}(AC) \text{tr}(AB)^*. \end{aligned} \quad (\text{A2})$$

In the framework of Wirtinger derivatives in complex analysis, derivatives treat quantities X and X^* as independent variables, and Eq. (A1) is used in the step from line 1 to line 2.

The evaluation of cost-function gradients requires the application of the chain rule to expressions of the type $\frac{\partial}{\partial u_i} c(M(\mathbf{u}))$. Here, c maps a complex-valued $\ell \times \ell$ matrix M (e.g., the propagator K_N with ℓ denoting the Hilbert space dimension) to a real number (the cost). The matrix $M = (M_{mn})$ itself depends on the real-valued control parameters $\mathbf{u} \in \mathbb{R}^d$. The subscript in u_i is understood as a multi-index $i = (k, j)$ encoding the control-field label k and discretized-time index j . The matrix-calculus result

$$\begin{aligned} \frac{\partial}{\partial u_i} c(M(\mathbf{u})) &= \frac{\partial c}{\partial M_{mn}} \frac{\partial M_{mn}}{\partial u_i} + \frac{\partial c}{\partial M_{mn}^*} \frac{\partial M_{mn}^*}{\partial u_i} \\ &= \text{tr} \left(\frac{\partial c}{\partial M} \frac{\partial M}{\partial u_i} \right) + \text{c.c.} \end{aligned} \quad (\text{A3})$$

is straightforward to derive with the ‘‘regular’’ chain rule by re-interpreting the functions involved as $c: \mathbb{C}^{\ell^2} \rightarrow \mathbb{R}$ and $M: \mathbb{R}^d \rightarrow \mathbb{C}^{\ell^2}$. In the following, Eqs. (A2) and (A3) are used to obtain the analytical expressions for several examples of cost-function gradients.

1. Gradient for C_1 : target-gate infidelity

The cost function $C_1 = 1 - |\text{tr}[K_T^\dagger K_N(\mathbf{u})]|^2$, penalizes the infidelity of the realized unitary $K_N = U_N U_{N-1} \dots U_1 U_0$ with respect to the target propagator K_T . It has the gradient

$$\begin{aligned} \frac{\partial C_1}{\partial u_{k,j}} &\stackrel{(\text{A3})}{=} \text{tr} \frac{\partial C_1}{\partial K_N} \frac{\partial K_N}{\partial u_{k,j}} + \text{c.c.} \\ &= -\text{tr} \left[\frac{\partial [|\text{tr}(K_T^\dagger K_N)|^2]}{\partial K_N} \frac{\partial K_N}{\partial u_{k,j}} \right] + \text{c.c.} \\ &\stackrel{(\text{A2})}{=} -\text{tr} \left(K_T^\dagger \frac{\partial K_N}{\partial u_{k,j}} \right) \text{tr}(K_T^\dagger K_N)^* + \text{c.c.} \\ &= \text{tr} \left(K_T^\dagger \left[\prod_{j'>j} U_{j'} \right] i \delta t \mathcal{H}_k K_j \right) \text{tr}(K_T^\dagger K_N)^* + \text{c.c.} \\ &= -2 \delta t \text{Im} \left\{ \text{tr} \left(K_T^\dagger \left[\prod_{j'>j} U_{j'} \right] \mathcal{H}_k K_j \right) \text{tr}(K_T^\dagger K_N)^* \right\} \end{aligned} \quad (\text{A4})$$

where \prod is understood to produce a time-ordered product.

This expression shows that automatic reverse-mode differentiation requires the propagators K_j from every time step. Within TensorFlow, the set of intermediate propagators $\{K_j\}$ is stored in memory during the forward evolution. The resulting memory demand therefore scales as $O(\ell^2 \times N)$.

Memory-efficient algorithm.— We note that storage of $\{K_j\}$ can be avoided by applying the strategy introduced in the original GRAPE paper [2]: since the evolution is unitary, one may time-reverse the evolution step by step, and recalculate the intermediate propagator via $K_j = U_{j+1}^\dagger K_{j+1}$. Here, each short-time propagator U_j is re-generated locally in time, using only the control fields at time t_j . Such a

backwards-propagation algorithm leads to an increase in computation time by roughly a factor of 2 (each U_j is then calculated twice), but has a memory demand of only $O(\ell^2)$ – which does not scale with N , the number of time steps. Thus for large problems the memory efficient algorithm is superior. This memory-efficient algorithm, currently not realized in this implementation, is given by

Algorithm 1 C_1 gradient via backwards propagation

```

1:  $P = \text{tr}(K_T^\dagger K_N)^* K_T^\dagger$ 
2:  $X = K_N$ 
3: for  $j = N$  to 0 do
4:   for all  $k$  do
5:      $\partial C_1 / \partial u_{k,j} = -2\delta t \text{Im}[\text{tr}(P \mathcal{H}_k X)]$ 
6:   end for
7:    $X = U_j^\dagger X$ 
8:    $P = P U_j$ 
9: end for
10: return  $\partial C_1 / \partial \mathbf{u}$ 

```

2. Gradient for C_2 : target-state infidelity

For state preparation or unitaries specified only in a subspace, it is sufficient to optimize the evolution for only a few initial states, rather than for the complete basis. This is achieved by minimizing a cost function based on $C_2(\mathbf{u}) = 1 - |\langle \Psi_T | \Psi_N \rangle|^2$, where the realized final state $|\Psi_N\rangle$ depends on the control parameters \mathbf{u} . Again applying equations (A3) followed by (A2) (and using that the trace of a number results in that number), we obtain

$$\frac{\partial C_2}{\partial u_{k,j}} = -2 \delta t \text{Im} \left[\langle \Psi_T | \left[\prod_{j'>j} U_{j'} \right] \mathcal{H}_k | \Psi_j \rangle \langle \Psi_T | \Psi_N \rangle^* \right]$$

Memory-efficient algorithm.— In TensorFlow-based automatic differentiation algorithm here, the intermediate states $\{|\Psi_j\rangle\}$ are stored, leading to a memory requirement of $O(\ell \times N)$, rather than $O(\ell^2 \times N)$ for the full propagators. By using the same backward propagation strategy as above, a more memory-efficient algorithm with memory requirement $O(\ell)$ independent of the time-step number is possible:

Algorithm 2 C_2 gradient via backwards propagation

```

1:  $P = \langle \Psi_T | \Psi_N \rangle^* \langle \Psi_T |$ 
2:  $X = |\Psi_N\rangle$ 
3: for  $j = N$  to 0 do
4:   for all  $k$  do
5:      $\partial C_2 / \partial u_{k,j} = -2\delta t \text{Im}[P \mathcal{H}_k X]$ 
6:   end for
7:    $X = U_j^\dagger X$ 
8:    $P = P U_j$ 
9: end for
10: return  $\partial C_2 / \partial \mathbf{u}$ 

```

3. Gradient for C_5 : occupation of forbidden state

Occupation of a ‘‘forbidden’’ state is discouraged by the cost function $C_5 = \sum_j |\text{tr}(\Psi_F^\dagger \Psi_j)|^2$. This cost function differs qualitatively from the gate and state infidelity cost functions: the latter are evaluated based on the result at the final time, while forbidden-state occupation involves intermediate states at every time step. Accordingly, the corresponding gradient takes a different form. First, Eq. (A3) is replaced by

$$\begin{aligned} & \frac{\partial}{\partial u_i} c(\Psi_0(\mathbf{u}), \Psi_1(\mathbf{u}), \dots, \Psi_N(\mathbf{u})) \\ &= \text{tr} \frac{\partial c}{\partial \Psi_j} \frac{\partial \Psi_j}{\partial u_i} + \text{c.c.} \end{aligned} \quad (\text{A5})$$

where introduction of the trace of a c -number is convenient for direct application of Eq. (A2). We then obtain

$$\begin{aligned} \frac{\partial C_5}{\partial u_{k,j}} &\stackrel{(\text{A5})}{=} \sum_J \text{tr} \frac{\partial C_5}{\partial \Psi_J} \frac{\partial \Psi_J}{\partial u_{k,j}} + \text{c.c.} \\ &= \sum_{J \geq j} \sum_{j'} \text{tr} \left[\frac{\partial [|\text{tr}(\Psi_F^\dagger \Psi_{j'})|^2]}{\partial \Psi_J} \frac{\partial \Psi_J}{\partial u_{k,j}} \right] + \text{c.c.} \\ &\stackrel{(\text{A2})}{=} \sum_{J \geq j} \text{tr} \left(\Psi_F^\dagger \frac{\partial \Psi_J}{\partial u_{k,j}} \right) \text{tr} (\Psi_F^\dagger \Psi_J)^* + \text{c.c.} \\ &= 2 \delta t \sum_{J \geq j} \text{Im} \left[\langle \Psi_F | \left[\prod_{j'=j+1}^J U_{j'} \right] \mathcal{H}_k | \Psi_j \rangle \langle \Psi_J | \Psi_F \rangle \right] \end{aligned} \quad (\text{A6})$$

The double sum of eq. (A6) makes it appear as though the computation of this gradient would take $O(N^2)$, however after simplification, the relationship between the limits of the sum and product allow it to be calculated in $O(N)$ time. The corresponding backward propagation algorithm then takes the following form:

Algorithm 3 C_5 gradient via backwards propagation

```

1:  $P = \langle \Psi_N | \Psi_F \rangle \langle \Psi_F |$ 
2:  $X = | \Psi_N \rangle$ 
3: for  $j = N$  to 0 do
4:   for all  $k$  do
5:      $\partial C_5 / \partial u_{k,j} = 2 \delta t \text{Im}[P \mathcal{H}_k X]$ 
6:   end for
7:    $X = U_j^\dagger X$ 
8:    $P = P U_j + \langle X | \Psi_F \rangle \langle \Psi_F |$ 
9: end for
10: return  $\partial C_5 / \partial \mathbf{u}$ 

```

This cost function and gradient is also used as the time-optimal award function, using a negative cost to reward rather than penalize the target state at every time step (rather than just at the end). The gradients of cost functions involving only control fields do not involve the time propagation, so we also omit their derivation.

4. Summary

Algorithms for each cost function along with their computation and memory costs have been presented. The computation time of the algorithms all scale linearly with the number N of time steps. Automatic gradient calculation which requires caching of each step causes memory to scale like N , while reducing the run time by a constant factor of 2. By contrast, algorithms which directly exploit the unitary structure of quantum evolution can have memory requirements do not scale with the number of time steps. Hence, it may be worth implementing analytic gradients for very long computations which otherwise would not fit in memory.

Computing the fidelity and gradient for the whole unitary evolution as in algorithm 1, requires $O(\ell^2)$, whereas state transfer requires $O(\ell)$ memory. It should be noted that full unitary evolution fidelity can also be calculated as ℓ^2 state transfer computations over a complete basis. This has the memory requirements of state transfer, and the same computation requirements as algorithm 1, though is less efficient by a constant factor. In principle, each state transfer can be performed in parallel and assembled to compute the total cost and gradient. In addition, the Hamiltonians of many physical problems can be represented sparsely allowing a significant speedup in computation as well. For practical problems, the number time steps required may scale with the size of the problem, as more complex quantum gates/algorithms require more time than simple ones.

[1] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, *Eur. Phys. J. D* **69**, 1 (2015).
[2] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, *Journal of Magnetic Resonance* **172**, 296 (2005).
[3] P. de Fouquieres, S. G. Schirmer, S. J. Glaser, and I. Kuprov,

Journal of Magnetic Resonance **212**, 412 (2011).
[4] S. E. Sklarz and D. J. Tannor, *Phys. Rev. A* **66**, 053619 (2002).
[5] R. Eitan, M. Mundt, and D. J. Tannor, *Phys. Rev. A* **83**, 053426 (2011).
[6] C. Gollub, M. Kowalewski, and R. de Vivie-Riedle, *Phys. Rev. Lett.* **101**, 073002 (2008).
[7] R. Nigmatullin and S. G. Schirmer, *New Journal of Physics* **11**,

- 105032 (2009).
- [8] D. M. Reich, M. Ndong, and C. P. Koch, *The Journal of Chemical Physics* **136**, 104103 (2012), arXiv:1008.5126.
- [9] J. P. Palao and R. Kosloff, *Phys. Rev. A* **68**, 062308+ (2003).
- [10] D. Tannor, V. Kazakov, and V. Orlov, in *Time-Dependent Quantum Molecular Dynamics*, Nato ASI Series, Vol. 299, edited by J. Broeckhove and L. Lathouwers (Springer US, 1992) pp. 347–360.
- [11] W. Zhu, J. Botina, and H. Rabitz, *The Journal of Chemical Physics* **108**, 1953 (1998).
- [12] W. Zhu and H. Rabitz, *The Journal of Chemical Physics* **109**, 385 (1998).
- [13] Y. Ohtsuki, G. Turinici, and H. Rabitz, *The Journal of Chemical Physics* **120**, 5509 (2004).
- [14] Y. Maday and G. Turinici, *The Journal of Chemical Physics* **118**, 8191 (2003).
- [15] Y. Ohtsuki, Y. Teranishi, P. Saalfrank, G. Turinici, and H. Rabitz, *Phys. Rev. A* **75**, 033407 (2007).
- [16] A. Borzi[†], J. Salomon, and S. Volkwein, *Journal of Computational and Applied Mathematics* **216**, 170 (2008).
- [17] P. Ditz and A. Borzi[†], *Computer Physics Communications* **178**, 393 (2008).
- [18] S. Haykin, *Neural Networks: A Comprehensive Foundation*, 1st ed. (Prentice Hall PTR, Upper Saddle River, NJ, USA, 1994).
- [19] R. Hecht-Nielsen, in *Neural Networks, 1989. IJCNN., International Joint Conference on* (IEEE, 1989) pp. 593–605 vol.1.
- [20] K.-S. Oh and K. Jung, *Pattern Recognition* **37**, 1311 (2004).
- [21] B. Catanzaro, N. Sundaram, and K. Keutzer, in *Proceedings of the 25th International Conference on Machine Learning, ICML '08* (ACM, New York, NY, USA, 2008) pp. 104–111.
- [22] T. Sharp, in *Computer Vision ECCV 2008*, Lecture Notes in Computer Science, Vol. 5305, edited by D. Forsyth, P. Torr, and A. Zisserman (Springer Berlin Heidelberg, Berlin, Heidelberg, 2008) Chap. 44, pp. 595–608.
- [23] R. Raina, A. Madhavan, and A. Y. Ng, in *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09* (ACM, New York, NY, USA, 2009) pp. 873–880.
- [24] D. Steinkraus, I. Buck, and P. Y. Simard, in *Eighth International Conference on Document Analysis and Recognition (ICDAR'05)*, ICDAR '05 (IEEE, Washington, DC, USA, 2005) pp. 1115–1120 Vol. 2.
- [25] B. Block, P. Virnau, and T. Preis, *Computer Physics Communications* **181**, 1549 (2010).
- [26] M. A. Clark, R. Babich, K. Barros, R. C. Brower, and C. Rebbi, *Computer Physics Communications* **181**, 1517 (2010).
- [27] K. A. Wilkinson, P. Sherwood, M. F. Guest, and K. J. Naidoo, *Journal of computational chemistry* **32**, 2313 (2011).
- [28] I. S. Ufimtsev and T. J. Martinez, *Computing in Science & Engineering* **10**, 26 (2008).
- [29] L. Vogt, R. Olivares-Amaya, S. Kermes, Y. Shao, C. Amador-Bedolla, and A. Aspuru-Guzik, *The Journal of Physical Chemistry A* **112**, 2049 (2008).
- [30] R. Olivares-Amaya, M. A. Watson, R. G. Edgar, L. Vogt, Y. Shao, and A. Aspuru-Guzik, *Journal of chemical theory and computation* **6**, 135 (2009).
- [31] Z. Hou, H.-S. Zhong, Y. Tian, D. Dong, B. Qi, L. Li, Y. Wang, F. Nori, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, *New Journal of Physics* **18**, 083036 (2016).
- [32] X. Cui, Y. Chen, and H. Mei, in *Parallel and Distributed Systems (ICPADS), 2009 15th International Conference on* (IEEE, Washington, DC, USA, 2009) pp. 42–48.
- [33] K. Fatahalian, J. Sugerman, and P. Hanrahan, in *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware*, HWSW '04 (ACM, New York, NY, USA, 2004) pp. 133–137.
- [34] A. D. Córcoles, E. Magesan, S. J. Srinivasan, A. W. Cross, M. Steffen, J. M. Gambetta, and J. M. Chow, *Nature Communications* **6**, 6979 (2015).
- [35] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, *Nature* **536**, 63 (2016).
- [36] J. Kelly *et al.*, *Nature* **519**, 66 (2015).
- [37] *M8196A 92 GSa/s Arbitrary Waveform Generators — Keysight (formerly Agilent's Electronic Measurement)* (2016).
- [38] J. R. Johansson, P. D. Nation, and F. Nori, *Computer Physics Communications* **183**, 1760 (2012), arXiv:1110.0573.
- [39] J. R. Johansson, P. D. Nation, and F. Nori, *Computer Physics Communications* **184**, 1234 (2013).
- [40] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, *Phys. Rev. A* **84**, 022305 (2011).
- [41] H. J. Hogben, M. Krzystyniak, G. T. P. Charnock, P. J. Hore, and I. Kuprov, *Journal of Magnetic Resonance* **208**, 179 (2011).
- [42] Z. Tošner, T. Vosegaard, C. Kehlet, N. Khaneja, S. J. Glaser, and N. C. Nielsen, *Journal of Magnetic Resonance* **197**, 120 (2009).
- [43] T. W. Borneman, M. D. Hürlimann, and D. G. Cory, *Journal of Magnetic Resonance* **207**, 220 (2010).
- [44] A. Spörl, T. Schulte-Herbrüggen, S. J. Glaser, V. Bergholm, M. J. Storcz, J. Ferber, and F. K. Wilhelm, *Phys. Rev. A* **75**, 012302 (2007).
- [45] P. Rebentrost and F. K. Wilhelm, *Phys. Rev. B* **79**, 060507 (2009).
- [46] V. Nebendahl, H. Häffner, and C. F. Roos, *Phys. Rev. A* **79**, 012312 (2009).
- [47] J. Kelly *et al.*, *Phys. Rev. Lett.* **112**, 240504 (2014).
- [48] S. Machnes, D. J. Tannor, F. K. Wilhelm, and E. Assémat, “Gradient optimization of analytic controls: the route to high accuracy quantum optimal control,” (2015), arXiv:1507.04261.
- [49] D. J. Egger and F. K. Wilhelm, *Superconductor Science and Technology* **27**, 014001+ (2014).
- [50] P. J. Liebermann and F. K. Wilhelm, *Phys. Rev. Applied* **6**, 024022+ (2016).
- [51] P. Rebentrost, I. Serban, T. Schulte-Herbrüggen, and F. K. Wilhelm, *Phys. Rev. Lett.* **102**, 090401 (2009).
- [52] D. J. Egger and F. K. Wilhelm, *Phys. Rev. Lett.* **112**, 240503+ (2014).
- [53] R. L. Kosut, M. D. Grace, and C. Brif, *Phys. Rev. A* **88**, 052326 (2013).
- [54] F. Dolde *et al.*, *Nature Communications* **5** (2014), 10.1038/ncomms4371.
- [55] F. Platzer, F. Mintert, and A. Buchleitner, *Phys. Rev. Lett.* **105**, 020501 (2010).
- [56] P. Watts, J. Vala, M. M. Müller, T. Calarco, K. B. Whaley, D. M. Reich, M. H. Goerz, and C. P. Koch, *Phys. Rev. A* **91**, 062306+ (2015).
- [57] 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, *Phys. Rev. A* **91**, 062307+ (2015).
- [58] M. H. Goerz, F. Motzoi, K. B. Whaley, and C. P. Koch, “Charting the circuit QED design landscape using optimal control theory,” (2016), arXiv:1606.08825.
- [59] T. E. Skinner, T. O. Reiss, B. Luy, N. Khaneja, and S. J. Glaser, *Journal of magnetic resonance (San Diego, Calif. : 1997)* **167**, 68 (2004).
- [60] K. Kobzar, T. E. Skinner, N. Khaneja, S. J. Glaser, and B. Luy, *Journal of Magnetic Resonance* **170**, 236 (2004).
- [61] K. Kobzar, T. E. Skinner, N. Khaneja, S. J. Glaser, and B. Luy, *Journal of Magnetic Resonance* **194**, 58 (2008).

- [62] F. Motzoi, J. M. Gambetta, S. T. Merkel, and F. K. Wilhelm, Phys. Rev. A **84**, 022307 (2011), arXiv:1102.0584.
- [63] Q. M. Chen, R. B. Wu, T. M. Zhang, and H. Rabitz, Phys. Rev. A **92**, 063415 (2015).
- [64] I. Serban, J. Werschnik, and E. K. U. Gross, Phys. Rev. A **71**, 053810 (2005).
- [65] M. Bartholomew-Biggs, S. Brown, B. Christianson, and L. Dixon, Journal of Computational and Applied Mathematics **124**, 171 (2000).
- [66] R. E. Wengert, Commun. ACM **7**, 463 (1964).
- [67] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, “Implementing a Universal Gate Set on a Logical Qubit Encoded in an Oscillator,” (2016), arXiv:1608.02430.
- [68] J. D. Farnum and D. A. Mazziotti, Chemical Physics Letters **416**, 142 (2005).
- [69] J. P. Palao, R. Kosloff, and C. P. Koch, Phys. Rev. A **77**, 063412 (2008).
- [70] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Phys. Rev. A **76**, 042319 (2007), arXiv:cond-mat/0703002.pdf.
- [71] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret, Science **326**, 113 (2009), arXiv:0906.0831.pdf.
- [72] M. Abadi *et al.*, “TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems,” (2016), arXiv:1603.04467.
- [73] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, SIAM J. Sci. Comput. **16**, 1190 (1995).
- [74] D. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” (2015), arXiv:1412.6980.
- [75] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, “Automatic differentiation in machine learning: a survey,” (2015), arXiv:1502.05767.
- [76] C. Moler and C. Van Loan, SIAM Review **45**, 3 (2006).
- [77] M. Arioli, B. Codenotti, and C. Fassino, Linear Algebra and its Applications **240**, 111 (1996).
- [78] W. J. Cody, G. Meinardus, and R. S. Varga, Journal of Approximation Theory **2**, 50 (1969).
- [79] A. Jameson, W. Schmidt, and E. Turkel, in *14th fluid and plasma dynamics conference* (1981) p. 1259.
- [80] V. W. Lee, C. Kim, J. Chhugani, M. Deisher, D. Kim, A. D. Nguyen, N. Satish, M. Smelyanskiy, S. Chennupaty, P. Hammarlund, R. Singhal, and P. Dubey, SIGARCH Comput. Archit. News **38**, 451 (2010).
- [81] R. B. Sidje, ACM Trans. Math. Softw. **24**, 130 (1998).
- [82] Theano Development Team, arXiv e-prints **abs/1605.02688** (2016).
- [83] N. Bell and M. Garland, *Efficient Sparse Matrix-Vector Multiplication on CUDA*, NVIDIA Technical Report NVR-2008-004 (NVIDIA Corporation, 2008).
- [84] W. Liu and B. Vinter, in *2014 IEEE 28th International Parallel and Distributed Processing Symposium* (IEEE, 2014) pp. 370–381.
- [85] M. Mirrahimi, Z. Leghtas, V. V. Albert, S. Touzard, R. J. Schoelkopf, L. Jiang, and M. H. Devoret, New Journal of Physics **16**, 045014+ (2013), arXiv:1312.2017.
- [86] B. Vlastakis, G. Kirchmair, Z. Leghtas, S. E. Nigg, L. Frunzio, S. M. Girvin, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, Science **342**, 607 (2013).
- [87] D. C. McKay, R. Naik, P. Reinhold, L. S. Bishop, and D. I. Schuster, Phys. Rev. Lett. **114**, 080501+ (2015).
- [88] A cost function for reducing evolution time (C_7) was not included in this example.