Optimal control theory for a unitary operation under dissipative evolution

Michael H. Goerz, Daniel M. Reich, and Christiane P. Koch

Theoretische Physik, Universität Kassel, Heinrich-Plett-Str. 40, 34132 Kassel, Germany

(Dated: December 3, 2013)

We show that optimizing a quantum gate for an open quantum system requires the time evolution of only three states irrespective of the dimension of Hilbert space. This represents a significant reduction in computational resources compared to the complete basis of Liouville space that is commonly believed necessary for this task. The reduction is based on two observations: The target is not a general dynamical map but a unitary operation; and the time evolution of two properly chosen states is sufficient to distinguish any two unitaries. We illustrate gate optimization employing a reduced set of states for a controlled phasegate with trapped atoms as qubit carriers and a $\sqrt{i\text{SWAP}}$ gate with superconducting qubits.

PACS numbers: 02.30.Yy,03.65Yz,03.67.Bg

I. INTRODUCTION

Quantum effects such as entanglement and matter interference are predicted cornerstones of future technologies. Their exploitation requires the ability to reliably and accurately control complex quantum systems. A major obstacle is that a quantum system can never completely be isolated from its environment and the interaction with the environment causes decoherence [1]. This is particularly true for condensed phase settings as encountered in, e.g., solid state quantum devices. A number of concepts, such as decoherence-free subspaces [2] and noiseless subsystems [3], dynamical decoupling [4] and spectral engineering [5], have been developed to cope with decoherence. The applicability of these strategies is tied to specific conditions on the interaction between system and environment and, in practice, is often limited to systems that can be described by simple models. For complex quantum systems, numerical optimal control offers an alternative approach. It calculates the external controls that implement a desired target operation by performing an iterative search in the parameter space of the controls [6].

For quantum systems that are subject to decoherence, numerical optimal control was first employed to realize laser cooling of internal degrees of freedom in molecules [7]. Further applications, also utilizing a Markovian master equation to describe the open system dynamics, include controlling coherences [8], automatic protection against noise [9], selective photoexcitation of charge transfer [10], electric current in a molecular junction [11], quantum gates [12], and quantum memories [13]. Due to the formal equivalence between Markovian dissipation and quantum measurements, optimized observations can be determined using the same set of tools [14]. Numerical optimal control can also be applied to non-Markovian quantum systems [15–18] provided the dynamics can be calculated with sufficient efficiency.

The question of numerical effort becomes particularly important in the optimization of high-fidelity quantum gates. High fidelities, or small errors, are best achieved with monotonically convergent optimization algorithms that utilize gradient information and thus require repeated forward and backward propagation [19, 20]. Gate optimization under coherent dynamics implies propagation of a set of states that span the Hilbert (sub)space on which the target is defined [21, 22]. For open system dynamics, this was generalized to a set of states that span the corresponding Liouville (sub)space [9, 12, 18, 23]. It requires not only propagation of density matrices instead of wavefunctions but also a significantly larger number of states since Liouville space dimension is the square of Hilbert space dimension. Realistically, this limits quantum gate optimization to but the simplest examples, i.e., one-qubit and two-qubit operations.

The direct extension from Hilbert to Liouville space [9, 12, 23] overlooks the fact that in quantum gate optimization the target is a unitary operation and not a general dynamical map. The latter would indeed require a basis that spans the full Liouville space. However, much less information is required to assess how well a desired unitary is implemented. This observation is not only relevant for optimal control but also provides the basis for all current attempts at reducing the resources for estimating the average gate error [24–28]. In fact, only two states are necessary to distinguish any two unitaries, irrespective of Hilbert space dimension [29]. We show here that these two states, together with a third state enforcing the dynamical map on the optimization subspace to be contracting and population conserving, can be utilized to construct an optimization functional which attains its optimal value only if the desired gate is implemented with unit fidelity.

The two states that are required for unitary identification are constructed such that the first one consists of non-degenerate contributions from each Hilbert (sub)space direction. This corresponds to choosing a basis, and probing the gate error within this basis. In order to determine the error of gates that are diagonal in the chosen basis, i.e.,

phase errors, the second state is needed. For Hamiltonians which due to their inherent structure allow for nothing but diagonal gates, only the second state together with the third one, enforcing the dynamical map on the optimization subspace to be contracting and norm conserving, is required. In our application, we thus distinguish between gates which are diagonal and those that are non-diagonal in the logical basis.

Our optimization functional is closely related to the gate error. While two states represent the minimal set of states required to distinguish any two unitaries, it is impossible to deduce bounds on the gate error from the two states [29]. This is due to the state corresponding to the choice of basis being a totally mixed thermal state. Numerical bounds on the gate error can be derived when replacing the totally mixed state by a set of d pure states where d is the dimension of Hilbert space, i.e., by choosing a separate basis state for each Hilbert space direction [29]. The resulting set consists of $d+1$ states. Analytical bounds are obtained when also the second state of the minimal set is expanded [30]. The corresponding set is built out of the $2d$ states of two mutually unbiased bases [29]. This observation from process verification motivates the choice of optimization functionals which utilize these extended sets of states. Although the number of states then depends on Hilbert space dimension, this choice still comes with very significant savings in the computational resources. For example, already for a two-qubit gate, both $2d$ and $d+1$ represent a significant reduction in the number of states that need to be propagated, namely a reduction from 16 for the full Liouville space basis to 8 and 5, respectively.

We demonstrate below that two states are sufficient to optimize diagonal gates and three states to optimize non-diagonal two-qubit gates. We also show that, depending on the desired gate error, $d+1$, respectively $2d$ states in the optimization functional correspond to the numerically most efficient choice. We consider a controlled phasegate with neutral trapped atoms that are excited into a Rydberg state and a $\sqrt{iSWAP}$gate with superconducting qubits. In both examples, our optimization identifies gate implementations for which the error is limited by decoherence. This proves that all reduced sets of states are sufficient for determining the fundamental limit to the gate error and thus for quantum gate optimization.

The paper is organized as follows. Section II defines the optimization functional and presents the optimization algorithm. Optimization of a controlled phasegate for neutral atoms is discussed in Sec. III, whereas optimization of a non-diagonal gate for superconducting qubits is studied in Sec. IV. Section V concludes. The algebraic framework and the proofs required for the construction of the three states employed in the optimization functional are presented in Appendix A.

II. OPTIMAL CONTROL THEORY FOR A UNITARY OPERATION UNDER DISSIPATIVE EVOLUTION

A. Optimization functional

In order to employ optimal control theory to determine a high-fidelity implementation of quantum gates, one needs to define a distance measure J_T between the desired unitary $\hat{\mathbf{O}}$ and the actual evolution. We show here that

$$J_T = 1 - \sum_{i=1}^{n} \frac{w_i}{\mathsf{Tr}[\hat{\boldsymbol{\rho}}_i^2(0)]} \, \mathfrak{Re} \left\{ \mathsf{Tr} \left[\hat{\mathbf{O}} \hat{\boldsymbol{\rho}}_i(0) \hat{\mathbf{O}}^\dagger \hat{\boldsymbol{\rho}}_i(T) \right] \right\} , \tag{1}$$

with $n=3$ and specific initial states $\hat{\boldsymbol{\rho}}_i(0)$, represents a suitable choice for J_T. This is in contrast to Refs. [9, 12, 23], where n was taken to be the Liouville space dimension corresponding to $\hat{\mathbf{O}}$, i.e., $n = 2^{2N}$ for N qubits, and $\hat{\boldsymbol{\rho}}_i$ an orthonormal basis (under the Hilbert-Schmidt product) of Liouville space. In Eq. (1), w_i are weights, normalized as $\sum_{i=1}^{n} w_i = 1$. In order to evaluate J_T, the time evolved states $\hat{\boldsymbol{\rho}}_i(T)$ need to be obtained by solving the equation of motion describing the open system's evolution for $\hat{\boldsymbol{\rho}}_i$. While in general the dynamics can be non-Markovian, we will restrict ourselves to a Markovian master equation in the examples below. We assume the coherent part to include coupling to an external control, i.e., the Hamiltonian is of the form $\hat{\mathbf{H}}(t) = \hat{\mathbf{H}}_0 + \epsilon(t)\hat{\mathbf{H}}_1$, and generalization to several controls $\epsilon_i(t)$ is straightforward.

The functional J_T needs to be minimized with respect to $\epsilon(t)$. Further constraints can be added, for example,

$$J = J_T - \lambda_a \int_0^T \left[\epsilon(t) - \epsilon_{ref}(t) \right] / S(t) \, \mathrm{d}t , \tag{2}$$

where $\epsilon_{ref}(t)$ denotes a reference field, $S(t)$ enforces the field to be smoothly switched on and off, and the second term in Eq. (2) ensures a finite pulse fluence [22]. More complex additional constraints, for example restricting the spectral width of the pulse or confining the accessible state space [31, 32], are also conceivable.

Mathematically, our claim that only three states are sufficient to determine proper implementation of the desired unitary $\hat{\mathbf{O}}$ is equivalent to the conjecture that the optimization functional attains its global minimum if and only if

$$\hat{\boldsymbol{\rho}}_i\left(T\right) = \hat{\mathbf{O}}\hat{\boldsymbol{\rho}}_i(0)\hat{\mathbf{O}}^\dagger \tag{3}$$

for the three states $\hat{\boldsymbol{\rho}}_i$. The three states are constructed such that the first one fixes a basis, and the corresponding Hilbert-Schmidt product in Eq. (1) checks whether the gate is correctly implemented in this basis. It misses errors for gates that are diagonal in the basis, i.e., phase errors [29]. The second state is therefore chosen to detect phase errors with its Hilbert-Schmidt product in Eq. (1) [29]. The Hilbert-Schmidt product of the third state determines whether the dynamical map attained at time T conserves the population within the optimization subspace. This is necessary since the time evolution can be non-unitary due to decoherence or due to leakage into states other than the logical basis [41].

In more technical terms, $\hat{\boldsymbol{\rho}}_1(0)$ is a density matrix with N non-degenerate, non-zero eigenvalues. Spanning the d-dimensional Hilbert space ($d = 2^N$ for N qubits) by an arbitrary complete orthonormal basis, $\{|\varphi_i\rangle\}$, $\hat{\boldsymbol{\rho}}_1(0)$ is expressed in terms of a complete set of d one-dimensional orthogonal projectors $\hat{\mathbf{P}}_i = |\varphi_i\rangle\langle\varphi_i|$, i.e., $\hat{\boldsymbol{\rho}}_1(0) = \sum_{i=1}^d \lambda_i \hat{\mathbf{P}}_i$ with $\lambda_i \neq \lambda_j \forall i \neq j$ and $\lambda_i \geq 0$ [29]. The second state, $\hat{\boldsymbol{\rho}}_2(0)$, is constructed to be totally rotated with respect to $\hat{\boldsymbol{\rho}}_1(0)$, i.e., $\hat{\boldsymbol{\rho}}_2(0) = \hat{\mathbf{P}}_{TR}$ where $\hat{\mathbf{P}}_{TR}$ is a one-dimensional projector obeying $\hat{\mathbf{P}}_{TR}\hat{\mathbf{P}}_i \neq 0$ for $i = 1, \ldots, d$ [29]. $\hat{\boldsymbol{\rho}}_3(0)$ is the identity in the optimization subspace. A possible choice for the initial states reads

$$\left(\hat{\boldsymbol{\rho}}_1(0)\right)_{ij} = \frac{2\left(d-i+1\right)}{d\left(d+1\right)}\delta_{ij}\,, \tag{4a}$$

$$\left(\hat{\boldsymbol{\rho}}_2(0)\right)_{ij} = \frac{1}{d}\,, \tag{4b}$$

$$\left(\hat{\boldsymbol{\rho}}_3(0)\right)_{ij} = \frac{1}{d}\delta_{ij}\,, \tag{4c}$$

where the matrix elements are given in the optimization subspace, all other elements are zero. We show in Appendix A that the optimization reaches its target if and only if condition (3) is fulfilled. Specifically, we prove that propagation of three states is sufficient, irrespective of the dimension of the optimization subspace. This represents a significant computational saving compared to the propagation of 2^{2N} initial states deemed necessary in the literature [9, 12, 23] already for a small number of qubits.

The states $\hat{\boldsymbol{\rho}}_1$ and $\hat{\boldsymbol{\rho}}_2$ of Eq. (4), while sufficient in principle to distinguish any two unitaries, do not allow for stating bounds on the gate error [29]. Numerical, resp. analytical bounds, require $d+1$, resp. $2d$ states instead of just two [29, 30]. Motivated by this fact, we define two additional sets of states that can be employed in Eq. (1). When n in Eq. (1) is taken to be equal to $d+1$, the totally mixed state of Eq. (4a) is replaced by d pure states,

$$\hat{\boldsymbol{\rho}}_j(0) = |\varphi_j\rangle\langle\varphi_j|\,, \tag{5}$$

with $j = 1, \ldots, d$ and $\{|\varphi_j\rangle\}$ the logical basis. $\hat{\boldsymbol{\rho}}_{d+1}(0)$ is simply equal to $\hat{\boldsymbol{\rho}}_2(0)$ of Eq. (4b). In this case, Eq. (4c) is not required since the $d+1$ pure states are sufficient to enforce the dynamical map on the optimization subspace to be contracting and norm conserving. Similarly, the functional (1) employing $n = 2d$ states is constructed by replacing $\hat{\boldsymbol{\rho}}_1(0)$ of Eq. (4a) by $\hat{\boldsymbol{\rho}}_j$, $j = 1, \ldots, d$ of Eq. (5) and $\hat{\boldsymbol{\rho}}_2(0)$ of Eq. (4b) by

$$\hat{\boldsymbol{\rho}}_{d+j}(0) = |\tilde{\varphi}_j\rangle\langle\tilde{\varphi}_j|\,, \tag{6}$$

with $j = 1, \ldots, d$, where the states $|\tilde{\varphi}_j\rangle$ form a mutually unbiased basis with respect to the canonical basis $\{|\varphi_j\rangle\}$. For two qubits ($d = 4$), an example for such a basis is given by

$$|\tilde{\varphi}_1\rangle = \frac{1}{2}\left(|00\rangle + |01\rangle + |10\rangle + |11\rangle\right)\,, \tag{7a}$$

$$|\tilde{\varphi}_2\rangle = \frac{1}{2}\left(|00\rangle - |01\rangle + |10\rangle - |11\rangle\right)\,, \tag{7b}$$

$$|\tilde{\varphi}_3\rangle = \frac{1}{2}\left(|00\rangle + |01\rangle - |10\rangle - |11\rangle\right)\,, \tag{7c}$$

$$|\tilde{\varphi}_4\rangle = \frac{1}{2}\left(|00\rangle - |01\rangle - |10\rangle + |11\rangle\right)\,. \tag{7d}$$

B. Optimization algorithm

We assume in the following a coupling to the external field that is linear in the field and equations of motion that are linear in the states [42]. Moreover, the full optimization functional, Eq. (2), is linear in the states $\hat{\boldsymbol{\rho}}_i(T)$ and does not depend on the states at intermediate times t. In this case, the linear version of Krotov's method is sufficient to yield a monotonically convergent optimization algorithm [33]. It is given in terms of coupled control equations that need to be solved simultaneously. Here, we model the dissipative time evolution by a Markovian master equation,

$$\frac{d\hat{\boldsymbol{\rho}}}{dt} = \mathcal{L}(\hat{\boldsymbol{\rho}}) = -i[\hat{\mathbf{H}}(t), \hat{\boldsymbol{\rho}}] + \mathcal{L}_D(\hat{\boldsymbol{\rho}}) \,. \tag{8}$$

The control equations then read

$$\frac{d\hat{\boldsymbol{\rho}}_i}{dt} = -i[\hat{\mathbf{H}}, \hat{\boldsymbol{\rho}}_i] + \mathcal{L}_D(\hat{\boldsymbol{\rho}}_i) \,, \tag{9a}$$

$$\frac{d\hat{\boldsymbol{\sigma}}_i}{dt} = -i[\hat{\mathbf{H}}, \hat{\boldsymbol{\sigma}}_i] - \mathcal{L}_D(\hat{\boldsymbol{\sigma}}_i) \quad \text{and} \quad \hat{\boldsymbol{\sigma}}_i(t=T) = \frac{w_i}{\mathsf{Tr}[\hat{\boldsymbol{\rho}}_i^2(0)]} \hat{\mathbf{O}} \hat{\boldsymbol{\rho}}_i(0) \hat{\mathbf{O}}^\dagger \,, \tag{9b}$$

$$\Delta\epsilon(t) = \frac{S(t)}{\lambda_a} \sum_{i=1}^n \Im\left\{ \mathsf{Tr}\left(\hat{\boldsymbol{\sigma}}_i^{\text{old}}(t) \frac{\partial \mathcal{L}(\hat{\boldsymbol{\rho}}_i)}{\partial \epsilon}\Big|_{\rho_i^{\text{new}}, \epsilon^{\text{new}}} \right) \right\} \tag{9c}$$

with $i = 1, 2, 3$ when the initial conditions $\hat{\boldsymbol{\rho}}_i(0)$ of Eq. (4) are employed or $i = 1, \ldots d^2$ with d the dimension of Hilbert space when a full basis of Liouville space is propagated. In Eq. (9c), the states $\hat{\boldsymbol{\sigma}}_i^{\text{old}}$ are backward-propagated with the pulse of the previous iteration ('old'), whereas the states $\hat{\boldsymbol{\rho}}_i^{\text{new}}$ are forward-propagated with the updated pulse ('new'). The derivative with respect to the field is given by the commutator

$$\frac{\partial \mathcal{L}(\hat{\boldsymbol{\rho}})}{\partial \epsilon} = -i\left[\frac{\partial \hat{\mathbf{H}}}{\partial \epsilon}, \hat{\boldsymbol{\rho}}\right] \tag{10}$$

and has to be evaluated for the 'new' field and the states $\hat{\boldsymbol{\rho}}$ propagated under the 'new' field. For a complex control, which occurs for example when using the rotating wave approximation (RWA), Eq. (9c) holds for both the real and the imaginary part of $\epsilon(t)$.

The value of the optimization functional in Eq. (1) depends on the number and the specific choice of initial states as well as the choice of weights, and is therefore not suitable to compare the convergence behavior between different sets of states. Instead, we employ the average gate fidelity,

$$F_{\text{avg}} = \int \langle\Psi|\hat{\mathbf{O}}^\dagger \mathcal{D}(|\Psi\rangle\langle\Psi|)\hat{\mathbf{O}}|\Psi\rangle \, d\Psi \,, \tag{11}$$

for the comparison. In Eq. (11), $\mathcal{D}$ denotes the dynamical map describing the time evolution of the open quantum system, i.e., $\hat{\boldsymbol{\rho}}(T) = \mathcal{D}(\hat{\boldsymbol{\rho}}(0))$. The gate fidelity, respectively the gate error $1 - F_{\text{avg}}$, is easily evaluated as [34]

$$F_{\text{avg}} = \frac{1}{d(d+1)} \sum_{i,j=1}^d \left(\langle\varphi_i|\hat{\mathbf{O}}^\dagger \mathcal{D}(|\varphi_i\rangle\langle\varphi_j|)\hat{\mathbf{O}}|\varphi_j\rangle + \mathsf{Tr}\left[\hat{\mathbf{O}}|\varphi_i\rangle\langle\varphi_i|\hat{\mathbf{O}}^\dagger \mathcal{D}(|\varphi_j\rangle\langle\varphi_j|)\right] \right) \,. \tag{12}$$

III. EXAMPLE I: DIAGONAL GATES

It is quite common that a two-qubit Hamiltonian allows only for diagonal gates, such as a controlled phasegate. A prominent example are non-interacting qubit carriers that interact only when excited into an auxiliary state where they accumulate a non-local phase [35]. Neutral trapped atoms with long-range interaction in a Rydberg state present a physical implementation of this setting [35, 36]. Optimal control theory has been employed before to determine the minimum time in which a controlled phasegate can be implemented [37] and the optimum distribution of the single-qubit phases [38]. These optimizations were carried out, however, without explicitly accounting for decoherence. It is thus not clear whether the best solutions to avoid decoherence have indeed been identified. While the logical basis states and the Rydberg state are typically very long-lived, the main source of decoherence is spontaneous decay from an intermediate state which is necessary to access the Rydberg state. Due to experimental feasibility, the excitation
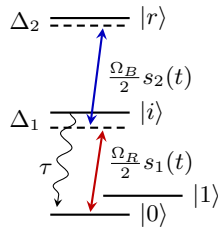
FIG. 1: Atomic levels for two-photon near-resonant excitation to a Rydberg state.

single-photon detuning Δ_1	600 MHz
two-photon detuning Δ_2	0
excitation energy E_1	6.8 GHz
Rabi frequencies Ω_R, Ω_B	300 MHz
interaction energy U	50 MHz
lifetime $\tau = 1/\gamma$	25 ns

TABLE I: Parameters of the Hamiltonian, Eq. (13), for implementing a controlled phasegate with two rubidium atoms.

to the Rydberg state proceeds by a near-resonant two-photon process. The corresponding single atom Hamiltonian in the basis $\{|0\rangle, |1\rangle, |i\rangle, |r\rangle\}$, cf. Fig. 1, and employing a two-color rotating wave approximation is given by

$$\hat{\mathbf{H}}_1 = \begin{pmatrix} 0 & 0 & \frac{1}{2}\Omega_R(t) & 0 \\ 0 & E_1 & 0 & 0 \\ \frac{1}{2}\Omega_R(t) & 0 & \Delta_1 & \frac{1}{2}\Omega_B(t) \\ 0 & 0 & \frac{1}{2}\Omega_B(t) & \Delta_2 \end{pmatrix}. \tag{13a}$$

The total Hamiltonian for two atoms includes an interaction when both atoms are in the Rydberg state,

$$\hat{\mathbf{H}} = \hat{\mathbf{H}}_1 \otimes \mathbb{1} + \mathbb{1} \otimes \hat{\mathbf{H}}_1 - U|rr\rangle\langle rr|. \tag{13b}$$

Spontaneous emission from the intermediate level is accounted for by the dissipator

$$\mathcal{L}_D(\hat{\rho}) = \gamma\left(\hat{\mathbf{A}}\hat{\rho}\hat{\mathbf{A}}^\dagger - \frac{1}{2}\left\{\hat{\mathbf{A}}^\dagger\hat{\mathbf{A}}, \hat{\rho}\right\}\right) \quad \text{with} \quad \hat{\mathbf{A}} = |0\rangle\langle i|, \tag{14}$$

and γ the decay rate, $\gamma = 1/\tau$. The parameters correspond to optically trapped rubidium atoms and are summarized in Table I. Since qubit level $|1\rangle$ remains decoupled throughout the time evolution, cf. Eq. (13a) and Fig. 1, the Hamiltonian (13) admits only diagonal gates. The update equations for real and imaginary part of the red and blue pulses are obtained by evaluating Eq. (9c) for the Hamiltonian given in Eq. (13),

$$\Re\left\{\Delta\Omega_{R,B}(t)\right\} = \frac{S(t)}{\lambda_a}\sum_{i=1}^{n}\Im\left\{\mathsf{Tr}\left(-i\hat{\boldsymbol{\sigma}}_i^{\text{old}}(t)\left[\Re\left\{\hat{\boldsymbol{\Omega}}_{R,B}(t)\right\}, \hat{\boldsymbol{\rho}}_i^{\text{new}}(t)\right]\right)\right\} \tag{15a}$$

$$\Im\left\{\Delta\Omega_{R,B}(t)\right\} = \frac{S(t)}{\lambda_a}\sum_{i=1}^{n}\Im\left\{\mathsf{Tr}\left(\hat{\boldsymbol{\sigma}}_i^{\text{old}}(t)\left[\Im\left\{\hat{\boldsymbol{\Omega}}_{R,B}(t)\right\}, \hat{\boldsymbol{\rho}}_i^{\text{new}}(t)\right]\right)\right\}, \tag{15b}$$

where $\hat{\boldsymbol{\Omega}}_{R,B}$ represents the coupling to the red and blue laser, respectively, in Eq. (13).

Figure 2 shows the gate error of the controlled phasegate versus iteration of the optimization algorithm when using a full basis, i.e., 16 states, or using three, respectively two, states in Eq. (15). The minimum number of states in this example is two since the Hamiltonian admits only diagonal gates, i.e., only phase errors and norm conservation within the logical subspace have to be checked. Therefore, ρ_1 in Eq. (4a) can be omitted, and the two remaining states are ρ_2 (phase errors) and ρ_3 (norm conservation) of Eqs. (4b, 4c). The relative weights w_2 and w_3 in Eq. (1) can be modified to emphasize one of the two aspects. Figure 2 therefore also compares two states with equal and unequal weights in Eq. (1), cf. green dotted and orange solid lines. The fastest convergence was obtained for $w_2/w_3 = 10$. The panels from top to bottom show the optimization without any dissipation, starting from a well-chosen guess pulse; an optimization starting with a bad guess pulse of insufficient fluence; and an optimization taking into account
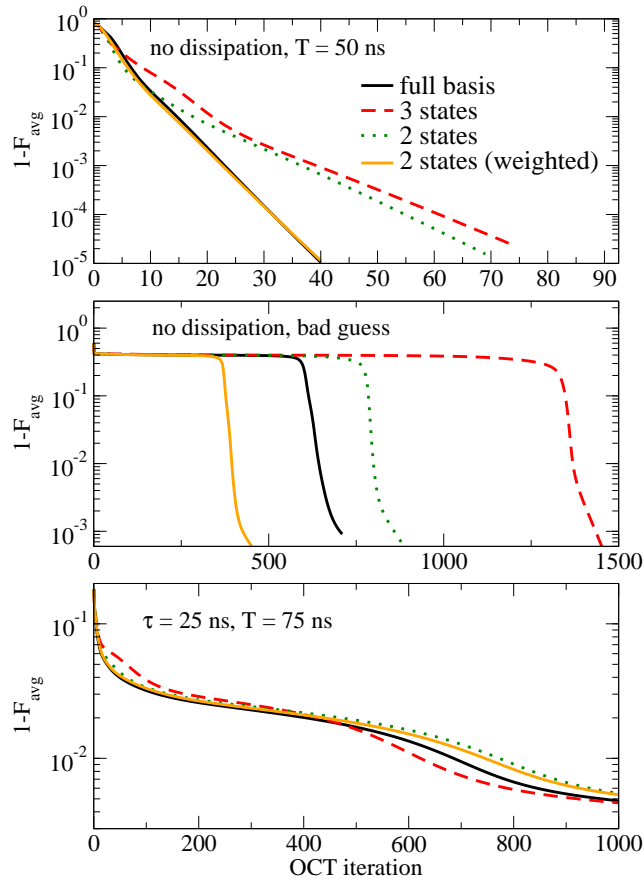
FIG. 2: Optimizing a controlled phasegate for two trapped neutral atoms that are excited to a Rydberg state and subject to spontaneous decay in the intermediate state: Convergence for optimization without and with spontaneous decay. The middle panel shows the same optimization as the top panel, but starting from a badly chosen guess pulse. The calculations employ equal weights of all states, except for those shown in orange where $w_2/w_3 = 10$.

spontaneous decay from the intermediary level. As the main observation, Fig. 2 clearly demonstrates that only two states are sufficient to optimize a quantum gate for a Hamiltonian of this kind. The optimization for coherent time evolution (top panel), shows that while the use of three states converges to gate errors as small as those obtained with the full basis, the convergence rate is only about half that of the full basis. This is due to two factors: (i) For the optimization with three states, there is no bound on the distance between the value J_T and the gate error, such that the path in the optimization landscape may be less direct until an asymptotic value is reached. Since without dissipation, there is no limit to the gate error, the convergence of J_T and that of $1 - F_{avg}$ stay on different trajectories. (ii) The reduced sets of states are constructed specifically to take into account decoherence. In particular, the third state contributes significantly less information that is relevant for reaching the optimization target than the second state. The convergence can be improved dramatically by weighting the three states according to the relevance of the information they carry. In this respect, the use of only two initial states can be seen as choosing $w_1 = 0$. Taking $w_2 > w_3$ addresses the issue of $\hat{\rho}_3$ contributing less to the optimization. Choosing proper weights allows for ensuring the convergence of optimization with a reduced set of states to be as fast as the optimization using the full basis.

The importance of choosing weights appropriate to the optimization problem becomes even more evident when the optimization starts from a bad guess pulse of insufficient fluence, as shown in the center panel of Fig. 2. The features observed in Fig. 2 are typical: The plateau near the beginning corresponds to the optimization increasing the intensity of the pulse without any significant improvement in the gate error, before converging quickly once the pulse is sufficiently intense. The end of the plateau can be significantly influenced by the choice of weights, cf. solid orange and dotted green curves in the middle panel of Fig. 2. Remarkably, the optimal choice of using two properly weighted initial states outperforms the use of the full basis. This might be explained by the fact that each of the three states in the reduced set has a specific physical role to play in the optimization, and this role can be emphasized by choice of the weight. In contrast, all states in the full basis fulfill the same role in the optimization, and thus there is no way
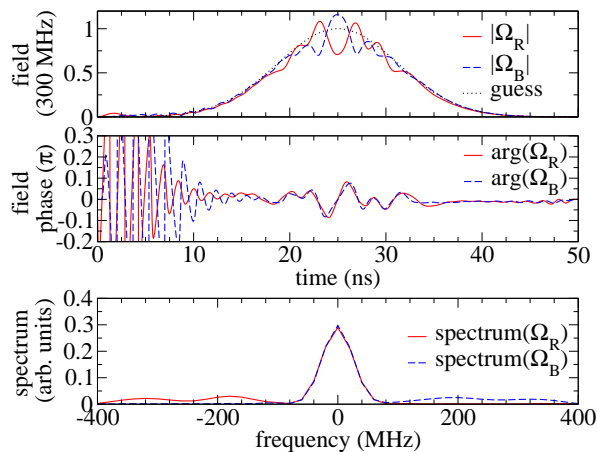
FIG. 3: The optimized pulses resulting from optimization using two states with unequal weights, without spontaneous decay. The pulse amplitudes are shown in the top panel, the complex phase in the center panel, and the pulse spectrum in the bottom panel. In the spectrum, frequency 0 corresponds to the carrier frequency of the laser pulses.

in which different weights on individual states would improve the convergence.

One should point out that even in the cases where the use of two or three states shows a slower convergence than that of the full basis, they still outperform the full basis in terms of numerical resources. Since both CPU time and the required memory scale linearly with the number of initial states in the optimization, using only two states compared to 16 has a 1:8 advantage, which more than offsets the factor of two in the convergence rate in the middle panel of Fig. 2.

Naturally, without the presence of decoherence, there is no reason to perform the optimization in Liouville space. Therefore, the results shown here only serve to illustrate the general convergence behavior of a reduced set of initial states. The more relevant case of non-coherent dynamics is shown in the bottom panel of Fig. 2. The presence of decoherence implies the existence of an asymptotic bound on the gate error. This constraint on the optimization landscape (together with the further constraint that only diagonal gates are reachable) ensures that all sets of reduced states converge at a similar rate, once the asymptotic region is approached. We expect that all choices reach the same asymptotic value; which choice yields the best fidelity after a specific number of iterations cannot be predicted in general. Factoring in all necessary resources, optimization using two states with unequal weights dramatically outperforms optimization using the full basis in this example.

The optimized pulse and spectrum in the case of coherent dynamics is presented in Fig. 3. The result shown here is obtained from the optimization using two initial states with unequal weights. However, the pulse is indistinguishable from the one obtained using the full basis, consistent with the identical convergence behavior for the two sets in the upper panel of Fig. 2. The optimized pulses only show relatively small amplitude modulations compared to the guess pulse (dotted line). These modulations appear as small side-peaks in the spectrum. In the time interval in which there is a significant pulse amplitude, the complex phase only deviates by about $\frac{\pi}{10}$ from zero. This phase evolution is reflected in the asymmetry of the spectrum for the red and the blue pulse (bottom panel). The spectrum nicely illustrates the mechanism of control: while each spectrum by itself is asymmetric, the red pulse showing negative frequencies, the blue pulse showing positive frequencies, the sum of both pulses is again symmetric, i.e., positive and negative frequencies cancel out. This means that the combination of both pulses is two-photon resonant with the transition $|0\rangle \to |r\rangle$, providing multiple pathways for the same transition whose interference might be exploited by the optimization.

The population dynamics induced by the optimized pulses are shown in Fig. 4. The two-photon resonance of the pulse expresses itself in a direct Rabi cycling between $|0\rangle$ and $|r\rangle$ on the left qubit in the propagation of $|01\rangle$ (top panel). The population shows roughly a 4π Rabi flip due to the relatively high pulse intensity. The nearly 25% of the population in the intermediate states in the propagation of $|00\rangle$ (bottom panel) is due to the fact that the decay from these levels was not included in the optimization, and thus the optimization algorithm makes no attempt at suppressing population in these states.

For the optimization with dissipation, the optimized pulse and pulse spectrum is shown in Fig. 5. The characteristics of the pulses are quite different compared to the coherent case. The red pulse remains close to the single Gaussian peak of the guess pulse, except for being slightly narrower. The blue pulse has a more complex structure. It is overall broader than the red pulse and consists of three distinctive features: an initial peak that overlaps but precedes the
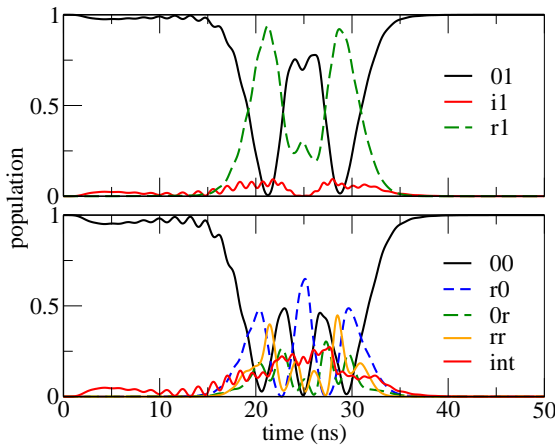
FIG. 4: Population dynamics under the pulse shown in Fig. (3), for the logical basis states $|01\rangle$ (top) and $|00\rangle$ (bottom). The intermediate population ("int") is integrated over all levels with decay, i.e., $|0i\rangle$, $|i0\rangle$, $|ii\rangle$, $|ir\rangle$, and $|ri\rangle$.
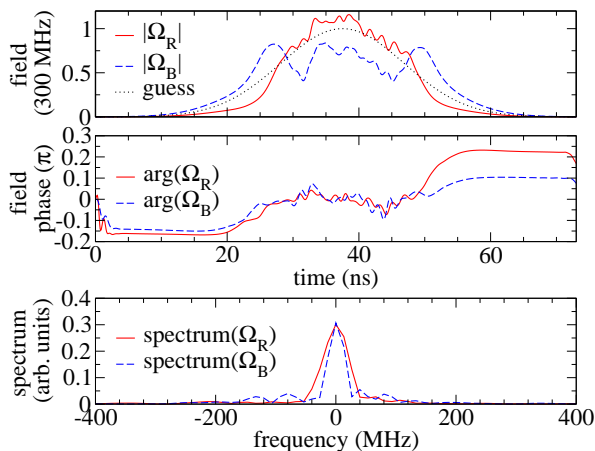

FIG. 5: The optimized pulses resulting from optimization using two weighted states and including spontaneous decay, c.f. Fig. 3.

red pulse, followed by some amplitude oscillations in the center of the pulse, and lastly another peak symmetric to the first, thus following the red laser pulse, with some overlap. For both pulses, the complex phase, shown in the center panel, is close to zero when there is significant pulse amplitude. In the spectrum (bottom panel), the overall narrowing and broadening of the red and blue pulse, respectively, is reflected in a broadening and narrowing of the central peak in the spectrum. The amplitude modulations on the blue pulse appear as side-lobes in the spectrum.

The initial and final peak of the blue pulse, together with the red pulse are reminiscent of the counter-intuitive pulse scheme of STIRAP, with the blue laser acting as the "Stokes" pulse and the red laser as "pump". The STIRAP-like behavior appears also in the population dynamics, shown in Fig. 6, as a population inversion between level $|0\rangle$ and $|r\rangle$, *without* any population in the intermediate decaying state. The amplitude modulations in the central region of both pulses then induce some additional dynamics, generating the entanglement needed for the gate. Note that the pulse duration for the dissipative process ($T = 75$ ns) is longer than that of the coherent process ($T = 50$ ns). This is necessary to allow for an adiabatic time evolution that is essential to the STIRAP-like behavior. Overall, the decaying intermediate state population (red lines in Fig. 6) is almost completely suppressed, which is in contrast to the optimization not taking into account the dissipation, cf. the red lines in Fig. 4. Both Figs. 4 and 6 show a significant population of the $|rr\rangle$ state. This is not surprising since the parameters of Table I are not in the regime of the Rydberg blockade [35, 36].
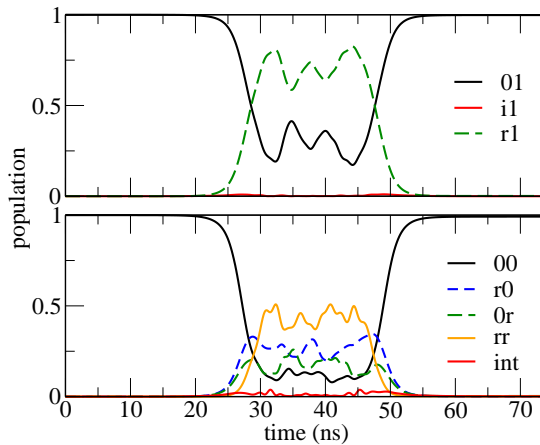
FIG. 6: Dissipative population dynamics under the pulse shown in Fig. 5, for the initial states $\hat{\boldsymbol{\rho}}(0) = |01\rangle\langle 01|$ (top) and $\hat{\boldsymbol{\rho}}(0) = |00\rangle\langle 00|$ (bottom). The intermediate population ("int") is integrated over all levels with decay, i.e. $|0i\rangle$, $|i0\rangle$, $|ii\rangle$, $|ir\rangle$, and $|ri\rangle$.

qubit frequency ω_1	4.3796 GHz
qubit frequency ω_2	4.6137 GHz
drive frequency ω_d	4.4985 GHz
anharmonicity δ_1	-239.3 MHz
anharmonicity δ_2	-242.8 MHz
effective qubit-qubit coupling J	-2.3 MHz
qubit 1 decay time T_1	38.0 μs
qubit 2 decay time T_1	32.0 μs
qubit 1 dephasing time T_2^*	29.5 μs
qubit 2 dephasing time T_2^*	16.0 μs

TABLE II: Parameters of the transmon Hamiltonian, Eq. (16), and Liouvillian, Eq. (17), taken from Ref. [39].

IV. EXAMPLE II: NON-DIAGONAL GATES

Superconducting qubits represent a physical realization of a quantum processor where the Hamiltonian admits both diagonal and non-diagonal entangling gates. In fact, there exist superconducting architectures that admit several two-qubit gates simultaneously [39, 40]. We consider here the example of two transmon qubits coupled via a shared transmission line resonator. In the dispersive limit, the interaction of each qubit with the resonator leads to an effective coupling J between the two qubits, and the cavity can be integrated out [39]. The resulting Hamiltonian reads

$$\hat{\mathbf{H}} = \left(\omega_1 - \frac{\delta_1}{2}\right)\hat{\mathbf{b}}_1^\dagger\hat{\mathbf{b}}_1 + \frac{\delta_1}{2}\left(\hat{\mathbf{b}}_1^\dagger\hat{\mathbf{b}}_1\right)^2 + \left(\omega_2 - \frac{\delta_2}{2}\right)\hat{\mathbf{b}}_2^\dagger\hat{\mathbf{b}}_2 + \frac{\delta_2}{2}\left(\hat{\mathbf{b}}_2^\dagger\hat{\mathbf{b}}_2\right)^2 + J\left(\hat{\mathbf{b}}_1^\dagger\hat{\mathbf{b}}_2 + \hat{\mathbf{b}}_1\hat{\mathbf{b}}_2^\dagger\right) \tag{16}$$
$$+ \Omega(t)\cos\left(\omega_d t\right)\left(\hat{\mathbf{b}}_1 + \hat{\mathbf{b}}_1^\dagger + \hat{\mathbf{b}}_2 + \hat{\mathbf{b}}_2^\dagger\right) ,$$

where $\hat{\mathbf{b}}_{1,2}$, $\hat{\mathbf{b}}_{1,2}^\dagger$ are the ladder operators for the first and second qubit, $\omega_{1,2}$ and $\delta_{1,2}$ represent the frequency and anharmonicity, J is the effective qubit-qubit-interaction, and $\Omega(t)$ and ω_d are amplitude and frequency of the drive, respectively. The two most relevant dissipation channels are energy relaxation and pure dephasing of the qubits, described by the decay rate $\gamma = 1/T_1$ and dephasing rate $\gamma_\phi = 1/T_2^*$ for each qubit. The corresponding dissipator reads

$$\mathcal{L}_D(\hat{\boldsymbol{\rho}}) = \sum_{q=1,2}\left(\gamma_q\sum_{i=1}^{N-1} i D\left[|i-1\rangle\langle i|_q\right]\hat{\boldsymbol{\rho}} + \gamma_{\phi,q}\sum_{i=0}^{N-1}\sqrt{i}D\left[|i\rangle\langle i|_q\right]\hat{\boldsymbol{\rho}}\right) , \tag{17}$$

with

$$D\left[\hat{\mathbf{A}}\right]\hat{\boldsymbol{\rho}} = \hat{\mathbf{A}}\hat{\boldsymbol{\rho}}\hat{\mathbf{A}}^\dagger - \frac{1}{2}\left(\hat{\mathbf{A}}^\dagger\hat{\mathbf{A}}\hat{\boldsymbol{\rho}} + \hat{\boldsymbol{\rho}}\hat{\mathbf{A}}^\dagger\hat{\mathbf{A}}\right) \tag{18}$$
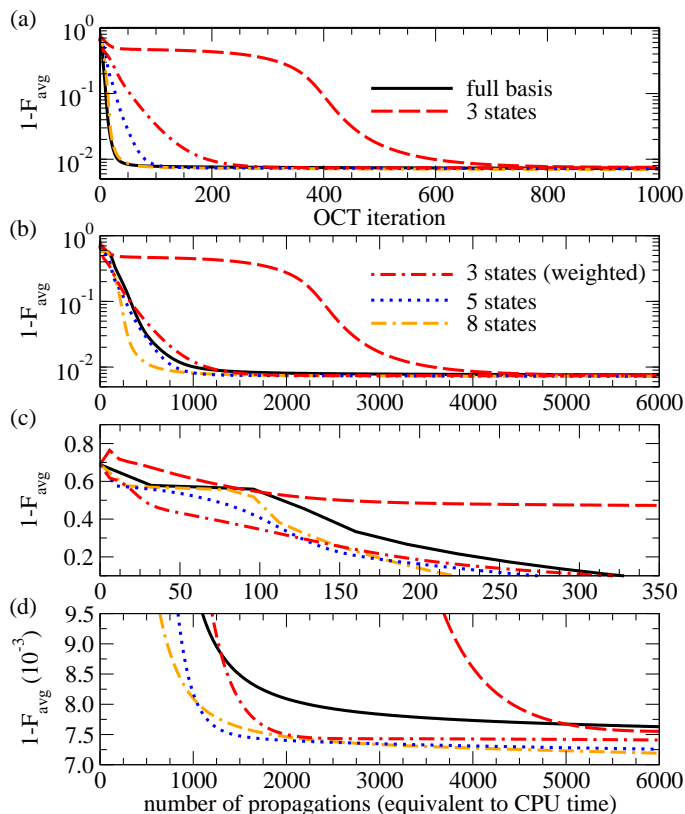
FIG. 7: Optimizing an $\sqrt{i\mathrm{SWAP}}$ gate for two transmons in the presence of energy relaxation and pure dephasing (with the rates given in Table II): Convergence for five choices of sets of initial states, as described in the text. The gate duration is $T = 400\,\mathrm{ns}$. The panels from top to bottom show the gate error over the number of iterations; the gate error over the number of state propagations, indicative of the required CPU time; a zoom on the initial phase of the optimization; and a zoom on the asymptotic convergence.

and each qubit, $q = 1, 2$, truncated at level N. The parameters of the coupled transmon qubits are summarized in Table II. We employ a RWA, centered at the drive frequency ω_d.

The Hamiltonian in Eq. (16) can generate a large number of entangling two-qubit gates; we find $\sqrt{i\mathrm{SWAP}}$ to be a fast converging non-diagonal perfect entangler, and thus choose

$$\hat{\mathbf{O}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{i}{\sqrt{2}} & 0 \\ 0 & \frac{i}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{19}$$

as the optimization target. Figure 7 shows the convergence behavior for several choices of initial states: the 16 canonical states of the full basis of Liouville space; the three states given in Eq. (4) with equal weight and with $w_1/w_2 = w_1/w_3 = 20$; a set of 5 states consisting of $\hat{\rho}_1$ expanded into four pure states, cf. Eq. (5) plus $\hat{\rho}_2$ of Eq. (4b); and lastly a set of eight states, cf. Eqs. (5) and (6), consisting of the expansion of $\hat{\rho}_1$ and the four pure states of a mutually unbiased basis, as explained in Section II. As seen in the top panel, all choices show good convergence. A plateau corresponding to a slowing of convergence is observed only for the three states with equal weights. But even in this case, the same asymptotic value for the gate error is obtained as for the other choices, see also Fig. 7(d). The advantage of employing the reduced sets of states in the optimization functional, Eq. (1), becomes most apparent in Fig. 7(b) which shows the gate error over the number of state propagations. Since optimization requires two propagations per iteration and state, i.e., the backward and forward propagation in Eq. (9), the number of state propagations corresponds directly to the CPU time that is required to obtain a given fidelity. Figure 7(c) and (d) shows a zoom on the same data, once for the initial phase of the optimization and once for the asymptotic behavior. All reduced sets except for the three states with equal weights perform better than the full set during the initial phase. Also, for this specific optimization problem, all reduced sets reach a slightly better asymptotic value than

the full set, although we expect that ultimately all curves will converge to the same value. Figure 7 suggests that the reduced sets have a significant advantage in reaching a good fidelity with a given amount of resources, especially since in practice, an optimization is usually stopped near the beginning of the asymptotic regime. Indeed, the full set shows an advantage only in the intermediate regime between gate errors of 10 and 1 percent, and only over the sets of three states. The choice of 5 or 8 states outperforms the full set in all cases. One should note that the savings in computational resources due to the use of a reduced set of states also extends to the amount of memory required, which is proportional to the number of states. Since in the optimization algorithm, propagated states over the entire time grid need to be stored, these savings can be very substantial.

For the three states with equal weights the gate error shows a non-monotonic behavior in the upper left corner of Fig. 7(c). This is due to the optimization functional, Eq. (1), not being equivalent to the gate error F_{avg}, Eq. (11). Specifically, for a set of three states, no bound on the distance between J_T and $1 - F_{\mathrm{avg}}$ can be derived [29]. Thus, the gate error might increase even though J_T decreases. In fact, the behavior of J_T is fully monotonic as expected (data not shown). With an increasing number of states in the chosen set, the value of the optimization functional is more closely connected to the gate fidelity; and for 5 and 8 states numerical, respectively analytical, bounds can be found [29, 30]. For this reason, we expect the sets of 5 and 8 states to show a faster convergence than the 3 states, when measured in OCT iterations, although not necessarily in CPU time. This expectation is confirmed by Fig. 7. The weak correspondence between the optimization functional and the gate error for three states is most likely also the reason for the plateau observed for the red dashed line in Fig. 7(a) and (b). However, the use of three states can still be a good choice since weighting the states properly improves the convergence significantly. The weights have to be chosen empirically, but the choice can be guided by physical intuition. The three states are responsible for ensuring that the realized gate is diagonal in the correct basis, that the relative phases match the target once the correct basis has been found, and that the gate is unitary on the logical subspace, respectively. The weights should reflect which of these requirements is most difficult to realize. In the present example this is finding the correct basis in which the gate is diagonal. Therefore the choice of $w_1/w_2 = w_1/w_3 = 20$ gave the best convergence rate. This is in contrast to the optimization of the Rydberg gate in Section III, in which the gate was already known to be diagonal, and the first state could be left out of the optimization entirely. Generally, using the set of three states with equal weights is not recommended.

Comparing Fig. 7 with the bottom panel of Fig. 2 for the Rydberg gate shows that the different choices of basis sets show a slightly wider range of the convergence rate. This can be attributed to the fact that, for the Rydberg gate, the optimization landscape is severely constrained since only diagonal gates can be reached. In contrast, the transmon Hamiltonian can generate both diagonal and non-diagonal gates, resulting in a more complex optimization landscape. Different choices of initial states can thus take more strongly varying pathways.

Figure 8 shows the optimization of an $\sqrt{i\mathrm{SWAP}}$ gate for two transmons in the case of weak dissipation, where the decay and dephasing times from Table II have been increased by a factor of 10. A comparison of Fig. 8(a) with Fig. 7(a) shows that the convergence behavior is essentially the same except for the value of the asymptote. We find an asymptotic gate error of approximately 7×10^{-3} with full dissipation, 7×10^{-4} with weak dissipation, and no asymptote without dissipation (data not shown). The value of the asymptote is logarithmically proportional to the decay and dephasing rates. This is as expected since the pulse duration is kept constant at 400 ns and the gate fidelity is solely limited by dissipation. Our claim that the dissipation only affects the asymptotic convergence is supported by a comparison of the initial convergence in Figs. 7(c) and 8(c), which remarkably are completely identical. Furthermore, the crossing between the black solid and red dot-dashed lines for the full basis and the three states with unequal weights near 1000 propagations and that between the blue dotted and orange dash-dash-dotted lines for the sets of 5, respectively 8, states near 1300 propagations in Fig. 8(d) can also be seen in Fig. 7(d). There are however some slight differences in the asymptotically reached values, in that the choice of 3 states (with both equal and unequal weights) reaches a slightly smaller gate error than in the case of full dissipation. Again, we expect that ultimately, all curves will converge to the same value. Which set of states reaches the best gate error at a specific point near the beginning of the asymptotic region seems to depend on the slope of the convergence curve as the limit is approached. This can depend on any number of factors including, e.g., the choice of λ_a in Eq. (2). Again, empirically, the reduced sets of states show a significant numerical advantage over the full basis also for weak dissipation.

As an example, the optimized pulse obtained using a set of three states with unequal weights, taking into account the full dissipation, is presented in Fig. 9, along with the pulse spectrum. The population dynamics that this pulse induces when propagating the logical basis states $\hat{\rho}(t = 0) = |01\rangle\langle 01|$ and $\hat{\rho}(t = 0) = |11\rangle\langle 11|$ is shown in Fig. 10. As can be seen in the top panel of Fig. 9, the optimized pulse shows small oscillations around the guess peak amplitude of 35 MHz. The complex phase, shown in the middle panel, stays relatively close to zero, indicating that the optimization employs mainly amplitude modulation. The pulse amplitude is roughly time-symmetric. The pulse spectrum shown in the bottom panel of Fig. 9 relates easily to the pulse shape. The strongest frequency component remains the driving frequency of the guess pulse (zero in the spectrum). The small oscillations in the pulse shape are approximately 8 ns apart, corresponding to a frequency of ± 125 MHz, which is present in the spectrum. There are peaks with
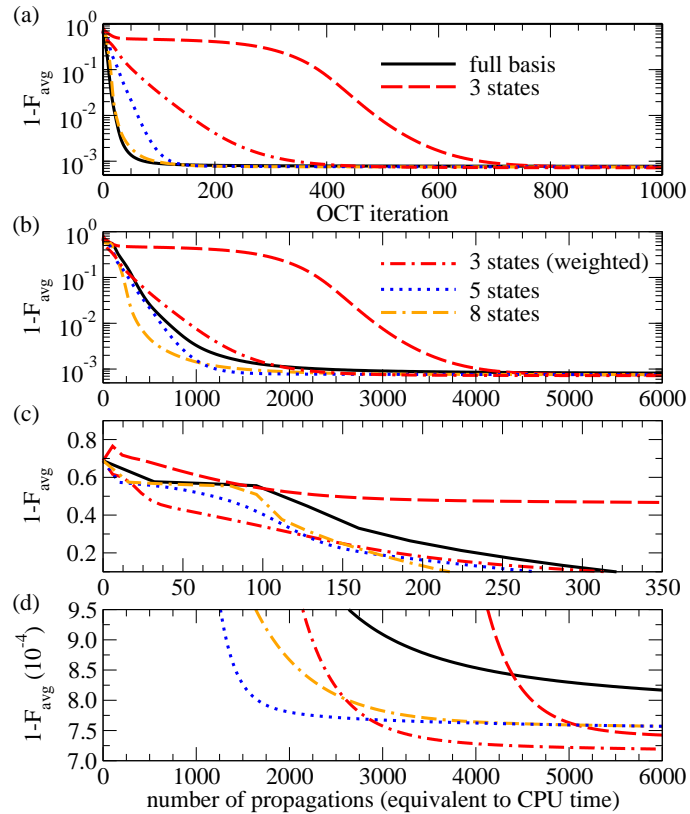
FIG. 8: Optimizing an $\sqrt{i\mathrm{SWAP}}$ gate for two transmons with weak dissipation, using decay and dephasing times increased by a factor of 10 compared to Fig. 7.

exponentially decaying amplitude in the spectrum at multiples of these values. The width of the central peak is due to the 20 ns switch-on and switch-off time of the pulse, and is unchanged from the guess pulse. The fact that there is not a single, but a double peak around ± 125 MHz corresponds the slow beats in the pulse shape. The slight asymmetry of the spectrum is caused by the complex phase of the optimized pulse.

The spectrum of the optimized pulse is very instructive in understanding the population dynamics in Fig. 10. The most relevant transition frequencies from the logical subspace are indicated by vertical lines in the spectrum in the lower panel of Fig. 9. Clearly, the peaks around ± 125 MHz are nearly resonant with the excitation of the left and right qubit, and the excitation to level $|2\rangle$ of the right qubit. There is no significant component in the spectrum that could excite to the level $|2\rangle$ of the left qubit. Consequently, in the population dynamics of both the $|01\rangle\langle 01|$ and $|11\rangle\langle 11|$ state, the right qubit (top panel) leaves the logical subspace (expectation value $\langle j \rangle > 1.0$) to a much more significant extent than the left qubit (middle panel). This behavior is slightly more pronounced for $|11\rangle\langle 11|$, which is the only state for which the total subspace population (gray curve in bottom panel) drops below 80% for a significant amount of time. The fact that for all logical basis states, most of the dynamics occurs within the logical subspace is due to the presence of decoherence, where higher levels have faster decay and faster dephasing due to a stronger coupling to the cavity. In an optimization without dissipation (data not shown), the optimized dynamics would generally veer farther outside the logical subspace. Lastly, the population dynamics show the expected behavior for the $\sqrt{i\mathrm{SWAP}}$ gate: the $|01\rangle$ state ends up in a coherent superposition between $|01\rangle$ and $|10\rangle$, whereas $|11\rangle$ returns to its original state at the end of the gate.

V. CONCLUSIONS

We have utilized the fact that the average error of a quantum gate can be estimated from the time evolution of a reduced set of states [28, 29] to construct a dedicated functional for quantum gate optimization in open quantum systems. Our optimization functional consists of Hilbert-Schmidt products that compare the actual and ideal time-evolved states from the reduced set. The minimal number of states that need to be forward and backward propagated
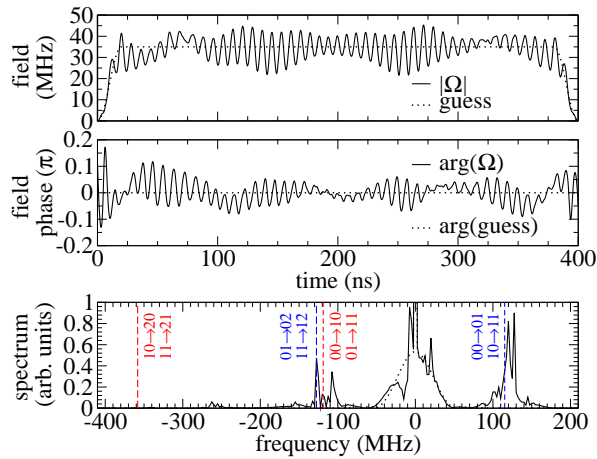
FIG. 9: Shape and spectrum of an optimized pulse, from optimization with 3 weighted states, with strong dissipation. The panels from top to bottom show the amplitude, complex phase, and spectrum of the optimized pulse $\Omega(t)$. The spectrum is shown in the rotating frame, with zero corresponding to the driving frequency w_d of the field. The transition frequencies from the logical subspace are indicated by vertical dashed lines. These are $\Delta_1 = w_1 - w_d = -118.88$ MHz and $\Delta_1 - \delta_1 = -358.18$ MHz in red for the left qubit, and $\Delta_2 = w_2 - w_d = 115.20$ MHz and $\Delta_2 - \delta_2 = -127.58$ MHz in blue for the right qubit. The central peak in the spectrum has been cut off to show the relevant side-peaks, and would extend to a value of approximately 10.0. For all quantities, the values for the guess pulse are shown as a dotted line.

during optimization is two for Hamiltonians that admit only diagonal gates and three for Hamiltonians that allow for both diagonal and non-diagonal gates. Remarkably, the size of the minimal set of states is independent of Hilbert space dimension.

While the minimal number of states allows for determining whether a quantum gate has been implemented, it is insufficient to deduce bounds on the gate error [29]. Numerical and analytical bounds require $d + 1$, respectively $2d$, states in the reduced set, where d is the dimension of the Hilbert space on which the optimization target is defined. Employing the sets of $d + 1$, respectively $2d$, states in quantum gate optimization is still significantly more efficient, both with respect to CPU time and memory requirements, than utilizing a full basis of Liouville space, with d^2 elements [9, 12, 23].

We have demonstrated the power of our approach in the optimization of a diagonal and a non-diagonal two-qubit gate. Specifically, we have optimized a controlled phasegate for trapped neutral atoms that are excited into a Rydberg state and subject to fast spontaneous emission from an intermediate state. The best performance was achieved by two states in the reduced set and a large weight of the Hilbert-Schmidt product for the state responsible for detecting phase errors. In the optimization of a $\sqrt{i\text{SWAP}}$ gate for two transmons coupled to the same transmission line cavity and subject to both energy relaxation and pure dephasing, we have found the best, and roughly identical, performance for the reduced sets consisting of $d + 1$, respectively $2d$, states. In all cases, the final gate error was limited by the decoherence rates. This confirms that employing a reduced set of states in quantum gate optimization is sufficient to determine the physical limit for the gate error.

The significant reduction in computational resources that we report here opens the door for a large-scale, systematic investigation of the fundamental limits of high-fidelity quantum gates in the presence of decoherence. Our approach is not tied to a specific decoherence model. It therefore allows to explore, using optimal control theory, settings for extended Hilbert spaces and beyond Markovian master equations, where a quantum system's complexity may possibly be exploited for control.
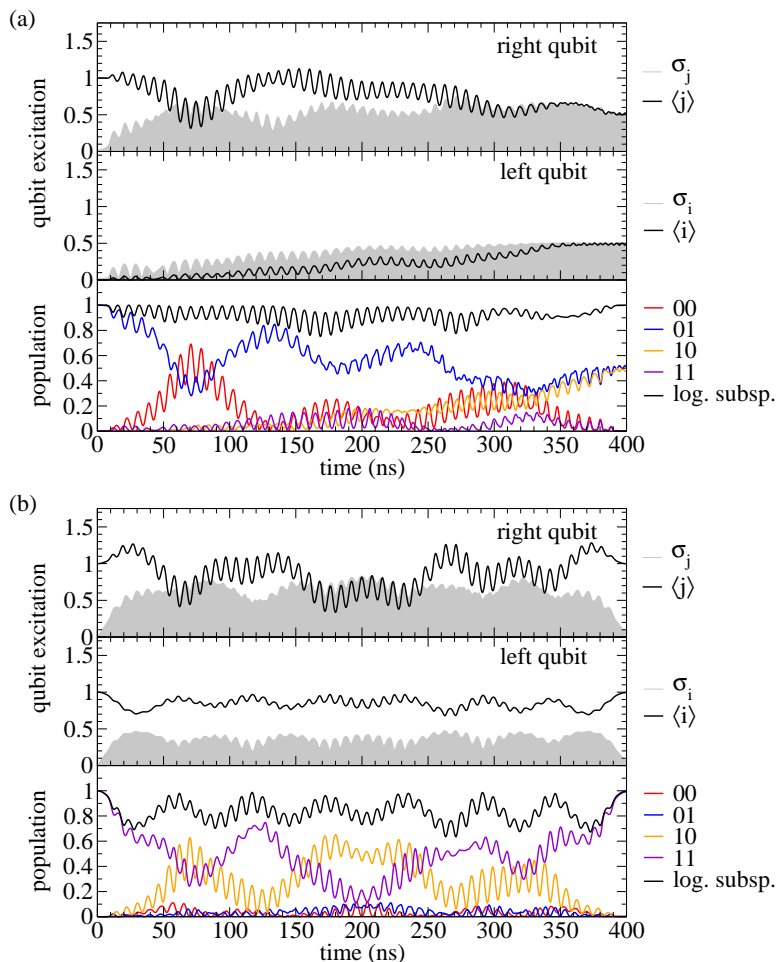
FIG. 10: Population dynamics for $\hat{\boldsymbol{\rho}}(t=0) = |01\rangle\langle01|$ (a) and $\hat{\boldsymbol{\rho}}(t=0) = |11\rangle\langle11|$ (b) under the pulse shown in Fig. 9. For each of the two propagated states, the expectation value of the right qubit excitation quantum number j is shown in the top panel, with the standard deviation in gray, the expectation value for the corresponding quantum number i for the left qubit is shown in the center panel, and the population dynamics for all the logical subspace states is shown in the bottom panel (colored lines), along with the total population in the logical subspace (black line).

Appendix A: Three states are sufficient to assess whether a desired target unitary is implemented

In the following we discuss the functional J_{dist},

$$J_{dist} = \sum_{i=1}^{3} \mathsf{Tr}\left[\left(\hat{\mathbf{O}}\hat{\boldsymbol{\rho}}_i(0)\hat{\mathbf{O}}^\dagger - \hat{\boldsymbol{\rho}}_i(T)\right)^2\right], \tag{A1}$$

which is built on the distance between the ideal and actual states at time T. It attains its global minimum, $J_{dist} = 0$, if and only if the initial states defined in Section II, $\hat{\boldsymbol{\rho}}_i(0)$ for $i = 1, 2, 3$, are mapped to their correct target states, i.e., fulfill condition (3). This functional motivates the use of the optimization functional J_T, Eq. (1), which is also built on only three states, as discussed in Sec. A 1. J_T and J_{dist} differ in that J_T evaluates the Hilbert-Schmidt products, i.e., the projections of the actual onto the ideal states instead of the trace distance. The construction of J_{dist}, and subsequently J_T, is rationalized by a theorem for unital, i.e., identity preserving, dynamical maps. Specifically, the theorem states that a *complete and totally rotating* set of density matrices is sufficient to determine whether a given time evolution is unitary. The functional (A1) exploits the further property of a complete and totally rotating set of density matrices to differentiate any two unitaries [29]. The theorem for unital dynamical maps is proven in Sec. A 2.

It should be stressed that we use J_T, Eq. (1), instead of J_{dist}, Eq. (A1), as optimization functional. This is motivated by the convexity of J_T which implies a much more favorable convergence behavior than would be obtained with a non-convex functional [43]. Mathematically, however, the two functionals are not equivalent. This is illustrated by

rewriting a single summand of J_{dist}, Eq. (A1), and comparing it to the corresponding term in J_T, Eq. (1),

$$\mathsf{Tr}\left[\left(\hat{\mathbf{O}}\hat{\boldsymbol{\rho}}_i(0)\hat{\mathbf{O}}^\dagger - \hat{\boldsymbol{\rho}}_i(T)\right)^2\right] = \mathsf{Tr}\left[\left(\hat{\mathbf{O}}\hat{\boldsymbol{\rho}}_i(0)\hat{\mathbf{O}}^\dagger\right)^2\right] - 2\,\mathsf{Tr}\left[\hat{\mathbf{O}}\hat{\boldsymbol{\rho}}_i(0)\hat{\mathbf{O}}^\dagger\hat{\boldsymbol{\rho}}_i(T)\right] + \mathsf{Tr}\left[(\hat{\boldsymbol{\rho}}_i(T))^2\right]. \tag{A2}$$

The first term on the rhs of Eq. (A2) is constant and thus irrelevant. The second term corresponds to the Hilbert-Schmidt overlap as used in J_T, Eq. (1), up to a prefactor. The main difference between J_T and J_{dist} is due to the third term, the purity of the propagated density matrix. J_T neglects this term. This could potentially disturb convergence, because the functional value of J_T can be decreased by (artificial) purification of the totally mixed states $\hat{\boldsymbol{\rho}}_1$ and $\hat{\boldsymbol{\rho}}_3$, cf. Eq. (4), instead of being decreased due to the desired approach to the target. Note that this problem can only arise for mixed states, i.e., when using the minimal set of states. For the reduced sets consisting of $d+1$, respectively $2d$, states, propagation starts from pure states, and the global minimum of J_T is identical to the global minimum of J_{dist}. Note that the problem of artificial purification is purely hypothetical and was never encountered in our optimizations – 'artificial purification traps' in the optimization landscape of the functional J_T with mixed states are apparently avoided.

1. Construction of the functional

We first define the concept of *complete and total rotation*, which we then use to formulate the required theorem. Let $\mathcal{H}$ be a Hilbert space with dimension N. Let $\mathcal{A}$ be a set of N one-dimensional orthogonal projectors. A one-dimensional projector is a projector with rank one, which means that its spectrum consists of a single eigenvalue equal to one with all remaining eigenvalues being zero.

Definition: A one-dimensional projector $\hat{\mathbf{P}}_{TR}$ is called *totally rotated* with respect to the set $\mathcal{A}$ if $\forall \hat{\mathbf{P}} \in \mathcal{A}$: $\hat{\mathbf{P}}_{TR}\hat{\mathbf{P}} \neq 0$.

Definition: A set of density operators, $\{\hat{\boldsymbol{\rho}}_i\}$ with $\hat{\boldsymbol{\rho}}_i \in \mathcal{H} \otimes \mathcal{H}$, is called *complete* if the set $\mathcal{P}$ of projectors onto the eigenspaces of $\{\hat{\boldsymbol{\rho}}_i\}$ contains exactly N one-dimensional orthogonal projectors.

Definition: A set of density operators, $\{\hat{\boldsymbol{\rho}}_i\}$ with $\hat{\boldsymbol{\rho}}_i \in \mathcal{H} \otimes \mathcal{H}$, is called *complete and totally rotating* if it is complete and there exists a one-dimensional projector in $\mathcal{P}$ that is totally rotated with respect to the orthogonal set of one-dimensional orthogonal projectors necessary for completeness.

Theorem 1: Let $\mathrm{DM}(N)$ be the space of $N \times N$ density matrices and $\mathcal{D} : \mathrm{DM}(N) \mapsto \mathrm{DM}(N)$ a dynamical map. The following three statements are equivalent:

1. $\mathcal{D}$ is unitary, i.e., $\mathcal{D}(\rho) = U\rho U^\dagger \; \forall \rho \in \mathrm{DM}(N)$ and U some element of the projective unitary group, $U \in \mathrm{PU}(N)$.

2. $\mathcal{D}$ maps a set $\mathcal{A}$ of N one-dimensional orthogonal projectors onto a set of N one-dimensional orthogonal projectors as well as a totally rotated projector $\hat{\mathbf{P}}_{TR}$ (with respect to $\mathcal{A}$) onto a one-dimensional projector.

3. $\mathcal{D}$ is unital and leaves the spectrum of a complete and totally rotating set of density matrices invariant.

We now explain how Theorem 1 can be used to prove the claim that J_{dist}, Eq. (A1), attains its global minimum if and only if condition (3) is fulfilled for the three states defined in Sec. II. We first discuss the role of $\hat{\boldsymbol{\rho}}_3 = \frac{1}{N}\mathbb{1}$. It is used to check whether the evolution corresponds to a dynamical map in the optimization subspace and whether it is unital. This dynamical map is obtained by projecting the action of the dynamical map, defined on the total Hilbert space, onto the optimization subspace. The term in the functional (A1) involving $\hat{\boldsymbol{\rho}}_3 = \frac{1}{N}\mathbb{1}$ becomes minimal, and so does the total functional, only if the identity in the optimization subspace is mapped onto itself. Minimization of J_{dist} thus ensures a unital dynamical map on the subsystem such that Theorem 1 is applicable.

We now discuss the role of $\hat{\boldsymbol{\rho}}_1$ and $\hat{\boldsymbol{\rho}}_2$ which by construction form a complete and totally rotating set of density matrices. The functional (A1) becomes zero only if $\mathcal{D}(\hat{\boldsymbol{\rho}}_1) = \hat{\mathbf{O}}\hat{\boldsymbol{\rho}}_1\hat{\mathbf{O}}^\dagger$ and $\mathcal{D}(\hat{\boldsymbol{\rho}}_2) = \hat{\mathbf{O}}\hat{\boldsymbol{\rho}}_2\hat{\mathbf{O}}^\dagger$. This requires the actual evolution to be unitary. Unitary evolution leaves the spectrum of a density matrix invariant. Due to the equivalence relation (1) $\iff$ (3) in Theorem 1, preservation of the spectrum of a complete and totally rotating set of density matrices, i.e., the two states $\hat{\boldsymbol{\rho}}_1$ and $\hat{\boldsymbol{\rho}}_2$, is sufficient to ensure unitarity. Furthermore, it was proven in Ref. [29] that the density matrices $\hat{\boldsymbol{\rho}}_1$ and $\hat{\boldsymbol{\rho}}_2$ are unitary differentiating, i.e., it is possible to distinguish any two unitary evolutions by inspection of $\hat{\boldsymbol{\rho}}_1(T)$ and $\hat{\boldsymbol{\rho}}_2(T)$ only. In particular there is only one unitary dynamical map, $\mathcal{D}(\hat{\boldsymbol{\rho}}) = \hat{\mathbf{U}}\hat{\boldsymbol{\rho}}\hat{\mathbf{U}}^\dagger$, which leads to $\mathcal{D}(\hat{\boldsymbol{\rho}}_i(0)) = \hat{\mathbf{O}}\hat{\boldsymbol{\rho}}_i\hat{\mathbf{O}}^\dagger$ for both $i = 1, 2$, namely the one induced by the target unitary $\hat{\mathbf{O}}$. Therefore the functional (A1) becomes minimal if and only if the target gate $\hat{\mathbf{O}}$ is implemented.

To summarize, J_{dist} is additively composed of three terms, each corresponding to a distance measure between the desired result, $\hat{\mathbf{O}}\hat{\boldsymbol{\rho}}_i\hat{\mathbf{O}}$, and the actually implemented evolution, $\mathcal{D}(\hat{\boldsymbol{\rho}}_i)$. For the total functional to be minimal, the evolutions of all three states have to match. This is the case only if a unital dynamical map on the optimization

subspace is implemented *and* if this is the unitary evolution according to $\hat{\mathsf{O}}$. More explicitly, the distance measure formed by the density matrices $i = 1, 2$ is only meaningful provided the evolution within the optimization subspace corresponds to a unital dynamical map. However, this is ensured by the third density matrix. Consequently, the global minimum of the functional (A1) will only be attained if this condition is fulfilled, too.

Note that the functional (A1) weights all three states equally. This is not a unique choice. In fact, all crucial properties of the functional remain unchanged when scaling the three terms with different positive factors, which has been done in the main text for example when discussing the optimisation using three states with weighting which significantly improved the performance of the optimization.

2. Proof

We utilize in the following the representation of operators by $N \times N$ matrices and therefore omit the operator notation. In order to prove Theorem 1, it is useful to first show the validity of the following lemma.

Lemma 1: Let $\mathcal{D}$ be a unital dynamical map, i.e., $\mathcal{D}$ is completely positive and maps identity onto itself, acting on $N \times N$ density matrices. If and only if there exists a set of N one-dimensional, orthogonal projectors that is mapped by $\mathcal{D}$ onto another set of N one-dimensional orthogonal projectors, there exists a complete set of density matrices whose spectrum is invariant under $\mathcal{D}$.

Proof of Lemma 1: ($\Longrightarrow$ direction) We denote the set of N one-dimensional projectors P_i by $\mathcal{P}$. By assumption, we know that

$$\forall i : \ \mathcal{D}(P_i) = \tilde{P}_i \, ,$$

where the $\tilde{P}_i$ also form a set of N one-dimensional, orthogonal projectors. Clearly, $\mathrm{spec}(P_i) = \mathrm{spec}(\tilde{P}_i)$, hence $\forall P_i \in \mathcal{P}$

$$\mathrm{spec}(\mathcal{D}(P_i)) = \mathrm{spec}(P_i) = (1, 0, \dots, 0) \, .$$

Obviously, $\mathcal{P}$ itself corresponds to a specific complete set of density matrices, $\rho_i = P_i$.

($\Longleftarrow$ direction) This part of the proof proceeds as follows: First we show that the assumption, a dynamical map leaving the spectrum of a given density matrix invariant, implies that $\mathcal{D}$ maps projectors onto the eigenspaces of the initial density matrices into projectors onto the eigenspaces of the resulting density matrix with the same eigenvalue. As a consequence, a one-dimensional projector onto a corresponding one-dimensional eigenspace is mapped into a one-dimensional projector. We then repeat this argument for all density matrices in the complete set. In this set, by definition, there exist density matrices with N one-dimensional, orthogonal projectors onto one-dimensional eigenspaces which, according to the first step of the $\Longleftarrow$ proof, is mapped onto another set of one-dimensional projectors. We show in a second step that the set of the mapped one-dimensional projectors is also orthogonal.

We start by assuming that $\mathcal{D}$ leaves the spectrum of a given density matrix, ρ, invariant,

$$\mathrm{spec}(\mathcal{D}(\rho)) = \mathrm{spec} \sum_k \left(E_k \rho E_k^\dagger \right) = \mathrm{spec}(\rho) \, ,$$

where we have expressed $\mathcal{D}$ in terms of Kraus operators E_k. We can write $\rho = \sum_i \lambda_i P_i'$ where $\mathcal{P}' = \{P_i'\}$ is a set of M orthogonal projectors onto the eigenspaces of ρ with M the number of distinct eigenvalues of ρ. We assume the λ_i to be ordered by magnitude with λ_1 corresponding to the largest eigenvalue. Since we know that the spectrum of $\mathcal{D}(\rho)$ to be identical to that of ρ, we can decompose $\mathcal{D}(\rho)$,

$$\mathcal{D}(\rho) \ = \ \sum_i \lambda_i \tilde{P}_i'$$

with $\{\tilde{P}_j'\}$ another set of M orthogonal projectors. Note that neither the P_i' nor the $\tilde{P}_i'$ have to be one-dimensional but for a given i, $\tilde{P}_i'$ has the same dimensionality as the corresponding P_i'. Specifically,

$$\mathcal{D}(\rho) = \mathcal{D}\left(\sum_i \lambda_i P_i' \right) = \sum_i \lambda_i \mathcal{D}(P_i') = \sum_j \lambda_j \tilde{P}_j' \, .$$

Multiplying by another projector $\tilde{P}_p'$ from the set, where p can take integer values between 1 and M, we obtain

$$\sum_i \lambda_i \mathcal{D}(P_i') \tilde{P}_p' = \sum_j \lambda_j \tilde{P}_j' \tilde{P}_p' = \lambda_p \tilde{P}_p' \, , \tag{A3}$$

since $\tilde{P}'_j$, $\tilde{P}'_p$ are orthogonal. Using proof by (transfinite) induction we now show that

$$\mathcal{D}\left(P'_k\right) = \tilde{P}'_k \quad \forall i = k, \dots, M \,.$$

The idea of the induction is the following: To show that indeed the projectors onto the eigenspaces of ρ, P'_i, are mapped into projectors onto the eigenspaces of $\mathcal{D}(\rho)$ with the same eigenvalue, we start with the projector onto the eigenspace with the largest eigenvalue and then inductively proceed to increasingly smaller eigenvalues. Furthermore, to prevent having to deal with a possible smallest eigenvalue of 0, we treat the lowest eigenvalue case separately. Calling the induction variable k, we have to show that $\mathcal{D}\left(P'_k\right) = \tilde{P}'_k$ follows from the assumption $\mathcal{D}\left(P'_i\right) = \tilde{P}'_i \ \forall i < k$. Note that if $k = M$, i.e., for the smallest eigenvalue,

$$\sum_i \mathcal{D}\left(P'_i\right) = \mathcal{D}\left(\sum_i P'_i\right) = \mathcal{D}(\mathbb{1}) = \mathbb{1} \,, \tag{A4}$$

since, by definition, a unital dynamical map maps identity onto itself. So assume $k \neq M$. Then $\lambda_k > 0$ since it is not yet the smallest eigenvalue because each λk corresponds by construction to a different eigenspace, hence they are different, and we assumed them to be ordered. For $k = p$, we can rewrite Eq. (A3), multiplying by an arbitrary normalized eigenvector $\vec{x}_k \in \mathbb{C}^N$ of $\tilde{P}'_k$ from the left and right,

$$\sum_i \lambda_i \vec{x}_k \cdot \mathcal{D}\left(P'_i\right) \cdot \vec{x}_k = \lambda_k \,. \tag{A5}$$

By assumption of the induction, $\mathcal{D}\left(P'_i\right) = \tilde{P}'_i \ \forall i < k$, therefore

$$\vec{x}_k \cdot \mathcal{D}\left(P'_i\right) \cdot \vec{x}_k = 0 \quad \forall i < k \,.$$

Introducing $d_{kk}^{(i)} \equiv \vec{x}_k \cdot \mathcal{D}\left(P'_i\right) \cdot \vec{x}_k$, Eq. (A5) can be written as

$$\sum_{i \geq k} \lambda_i d_{kk}^{(i)} = \lambda_k \,. \tag{A6}$$

Due to Eq. (A4) and the assumption of the induction,

$$\sum_i d_{kk}^{(i)} = \sum_{i \geq k} d_{kk}^{(i)} = \sum_i \vec{x}_k \cdot \mathcal{D}\left(P'_i\right) \cdot \vec{x}_k = 1 \,,$$

and, since $\mathcal{D}\left(P'_i\right)$ is the image of a positive semidefinite matrix which has to be positive semidefinite itself,

$$d_{kk}^{(i)} = \vec{x}_k \cdot \mathcal{D}\left(P'_i\right) \cdot \vec{x}_k \geq 0 \ \ \forall i \,.$$

Now remember that $\lambda_k \neq 0$ is strictly larger than all the other λ_i with $i > k$ since the eigenvalues are assumed to be ordered. In addition, $d_{kk}^{(i)} \geq 0 \ \ \forall i$ and at least one $d_{kk}^{(i)}$ with $i \geq k$ must be nonzero, otherwise the $d_{kk}^{(i)}$ would not sum up to 1. Then

$$\sum_{i \geq k} \lambda_i d_{kk}^{(i)} \leq \lambda_k \sum_{i \geq k} d_{kk}^{(i)} = \lambda_k \,,$$

with equality if and only if $d_{kk}^{(i)} = 0$ for $i \neq k$. In fact, equality has to hold since otherwise we would contradict Eq. (A6). We conclude that

$$d_{kk}^{(i)} = \vec{x}_k \cdot \mathcal{D}\left(P'_i\right) \cdot \vec{x}_k = \delta_{ik} \,.$$

Since $\vec{x}_k$ is normalized and arbitrary as long as it lies in the eigenspace $\tilde{\mathcal{E}}_k$ of $\tilde{P}'_k$, $\vec{x}_k$ must be an eigenvector of $\mathcal{D}\left(P'_k\right)$ with eigenvalue 1. Consequently, the operator $\mathcal{D}\left(P'_k\right)$ maps the eigenspace of $\tilde{P}'_k$ onto itself. Now we are almost done with showing that $\mathcal{D}\left(P'_k\right)$ and $\tilde{P}'_k$ are indeed identical. Since $\tilde{\mathcal{E}}_k$ is mapped by $\mathcal{D}$ into itself, $\mathcal{D}\left(P'_k\right)$ has at least $\dim(\tilde{\mathcal{E}}_k)$ eigenvalues equal to 1. The fact that $\mathcal{D}\left(P'_k\right)$ has exactly $\dim(\tilde{\mathcal{E}}_k)$ eigenvalues equal to 1 follows from $\mathcal{D}$ being trace-preserving: $\mathrm{Tr}\left[\mathcal{D}\left(P'_k\right)\right] = \mathrm{Tr}\left[P'_k\right] = \dim(\mathcal{E}_k)$ and $\dim(\mathcal{E}_k) = \dim(\tilde{\mathcal{E}}_k)$, where $\mathrm{Tr}\left[\mathcal{D}\left(P'_k\right)\right]$ is the sum over the eigenvalues of $\mathcal{D}\left(P'_k\right)$. Since all eigenvalues of $\mathcal{D}\left(P'_k\right)$ are non-negative, all other eigenvalues must vanish. Hence

$\mathcal{D}\left(P'_k\right) = \tilde{P}'_k$. This completes the induction and concludes the first step of the $\Longleftarrow$ proof, i.e., we have shown that a unital dynamical map that leaves the spectrum of a given arbitrary density matrix invariant, maps projectors onto the eigenspaces of this density matrix onto projectors of the same rank. This is specifically true for one-dimensional projectors. Iterating the argument for all density matrices in the complete set and selecting a set $\mathcal{P}$ of N orthogonal, one-dimensional projectors, it follows that these projectors will be mapped by $\mathcal{D}$ onto another set of one-dimensional projectors.

In the second step of the $\Longleftarrow$ proof we still need to show that the mapped set is also orthogonal. We denote the complete set of projectors by $\{P_i\}$. From the first step of the $\Longleftarrow$ proof we know that the $\tilde{P}_i$,

$$\mathcal{D}\left(P_i\right) = \tilde{P}_i\,,$$

need to be one-dimensional projectors. Using the unitality of $\mathcal{D}$, we see that

$$\mathbb{1} = \mathcal{D}\left(\mathbb{1}\right) = \mathcal{D}\left(\sum_i P_i\right) = \sum_i \mathcal{D}\left(P_i\right) = \sum_i \tilde{P}_i\,.$$

The unit matrix can only be summed by N one-dimensional projectors if these are orthogonal. Hence we have accomplished the second step, and the lemma follows.

Proof of Theorem 1: The equivalence relation of statement $(1) \Longleftrightarrow (2)$ has already been proven in Ref. [29]. To complete the proof of the more general Theorem 1, we are left with proving $(2) \Longleftrightarrow (3)$.

$(2) \Longrightarrow (3)$: If $\mathcal{D}$ maps a set of N one-dimensional orthogonal projectors onto another set of N one-dimensional orthogonal projectors, it leaves the spectrum of the projectors invariant. This can be seen as follows. Projectors are idempotent and positive semi-definite, hence their spectrum can only consist of zeros and ones. Since the projector is one-dimensional, its image under $\mathcal{D}$ has to be one-dimensional, too, and there can only be one eigenvalue equal to one. Thus any one-dimensional projector has the spectrum $\{1, 0, 0, \ldots\}$ which must be invariant under a mapping between one-dimensional orthogonal projectors. We now use the linearity of dynamical maps to show that $\mathcal{D}$ must be unital. Specifically, let $\{P_i\}$ be the initial set of orthogonal projectors that is mapped to another set of orthogonal projectors, $\{\bar{P}_i\}$. We find for the image of the totally mixed state, $\rho_M = \frac{1}{N}\mathbb{1}$,

$$
\begin{aligned}
\mathcal{D}(\rho_M) &= \mathcal{D}\left(\frac{1}{N}\sum_{i=1}^{N} P_i\right) = \frac{1}{N}\sum_{i=1}^{N}\mathcal{D}(P_i) \\
&= \frac{1}{N}\sum_{i=1}^{N}\bar{P}_i = \rho_M\,,
\end{aligned}
$$

i.e., $\mathcal{D}$ maps identity onto itself, making it unital. We can thus use Lemma 1 to obtain that the spectrum of a complete set of density matrices is invariant under $\mathcal{D}$. We now just have to add $\rho_{TR} = P_{TR}$ to realize a complete and totally rotating set. The spectrum of $\rho_{TR} = P_{TR}$ is also invariant under $\mathcal{D}$ since it is a one-dimensional projector that is mapped onto another one-dimensional projector.

$(3) \Longrightarrow (2)$: From Lemma 1, we obtain that $\mathcal{D}$ maps a set of N one-dimensional, orthogonal projectors onto another set of N one-dimensional orthogonal projectors. We are thus only left with showing that $\mathcal{D}$ maps a totally rotated projector onto a one-dimensional projector: There always exists a density matrix with a one-dimensional eigenspace corresponding to a totally rotated projector P_{TR} whose spectrum is invariant under the action of $\mathcal{D}$. In the proof of Lemma 1, we have shown that a dynamical map that leaves the spectrum of projectors invariant maps these projectors onto projectors of the same rank. Repeating the steps of the proof of Lemma 1, we see that the image of P_{TR} has to be a one-dimensional projector. This completes the proof of Theorem 1.

---

[1] H.-P. Breuer and F. Petruccione, *The theory of open quantum systems* (Oxford University Press, 2007).
[2] D. A. Lidar, I. L. Chuang, and K. B. Whaley, Phys. Rev. Lett. **81**, 2594 (1998).
[3] L. Viola, E. M. Fortunato, M. A. Pravia, E. Knill, R. Laflamme, and D. G. Cory, Science **293**, 2059 (2001).
[4] L. Viola, E. Knill, and S. Lloyd, Phys. Rev. Lett. **82**, 2417 (1999).
[5] J. Clausen, G. Bensky, and G. Kurizki, Phys. Rev. Lett. **104**, 040401 (2010).
[6] S. A. Rice and M. Zhao, *Optical control of molecular dynamics* (John Wiley & Sons, 2000).
[7] A. Bartana, R. Kosloff, and D. J. Tannor, J. Chem. Phys. **106**, 1435 (1997).
[8] Y. Ohtsuki, W. Zhu, and H. Rabitz, J. Chem. Phys. **110**, 9825 (1999).
[9] S. Kallush and R. Kosloff, Phys. Rev. A **73**, 032324 (2006).

[10] J. C. Tremblay and P. Saalfrank, Phys. Rev. A **78**, 063408 (2008).
[11] G.-Q. Li and U. Kleinekathöfer, Eur. Phys. J. B **76**, 309 (2010).
[12] T. Schulte-Herbrüggen, A. Spörl, N. Khaneja, and S. J. Glaser, J. Phys. B **44**, 154013 (2011).
[13] D. J. Gorman, K. C. Young, and K. B. Whaley, Phys. Rev. A **86**, 012317 (2012).
[14] F. Shuang, A. Pechen, T.-S. Ho, and H. Rabitz, J. Chem. Phys. **126**, 134303 (2007).
[15] P. Rebentrost, I. Serban, T. Schulte-Herbrüggen, and F. K. Wilhelm, Phys. Rev. Lett. **102**, 090401 (2009).
[16] E. Asplund and T. Klüner, Phys. Rev. Lett. **106**, 140404 (2011).
[17] R. Schmidt, A. Negretti, J. Ankerhold, T. Calarco, and J. T. Stockburger, Phys. Rev. Lett. **107**, 130404 (2011).
[18] F. F. Floether, P. de Fouquieres, and S. G. Schirmer, New J. Phys. **14**, 073023 (2012).
[19] J. Somlói, V. A. Kazakovski, and D. J. Tannor, Chem. Phys. **172**, 85 (1993).
[20] W. Zhu, J. Botina, and H. Rabitz, J. Chem. Phys. **108**, 1953 (1998).
[21] J. P. Palao and R. Kosloff, Phys. Rev. Lett. **89**, 188301 (2002).
[22] J. P. Palao and R. Kosloff, Phys. Rev. A **68**, 062308 (2003).
[23] Y. Ohtsuki, New J. Phys. **12**, 045002 (2010).
[24] A. Bendersky, F. Pastawski, and J. P. Paz, Phys. Rev. Lett. **100**, 190403 (2008).
[25] S. T. Flammia and Y.-K. Liu, Phys. Rev. Lett. **106**, 230501 (2011).
[26] E. Magesan, J. M. Gambetta, and J. Emerson, Phys. Rev. Lett. **106**, 180504 (2011).
[27] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, Phys. Rev. Lett. **107**, 210404 (2011).
[28] D. M. Reich, G. Gualdi, and C. P. Koch, Phys. Rev. Lett. **111**, 200401 (2013).
[29] D. M. Reich, G. Gualdi, and C. P. Koch, Phys. Rev. A **88**, 042309 (2013), arXiv:1305.3222.
[30] H. F. Hofmann, Phys. Rev. Lett. **94**, 160504 (2005).
[31] D. M. Reich, J. P. Palao, and C. P. Koch, J. Mod. Opt. (2013), URL http://www.tandfonline.com/doi/abs/10.1080/09500340.2013.844866.
[32] J. P. Palao, D. M. Reich, and C. P. Koch, Phys. Rev. A **88**, 053409 (2013).
[33] D. M. Reich, M. Ndong, and C. P. Koch, J. Chem. Phys. **136**, 104103 (2012).
[34] L. H. Pedersen, N. M. Møller, and K. Mølmer, Phys. Lett. A **367**, 47 (2007).
[35] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, Phys. Rev. Lett. **85**, 2208 (2000).
[36] M. Saffman, T. G. Walker, and K. Mølmer, Rev. Mod. Phys. **82**, 2313 (2010).
[37] M. H. Goerz, T. Calarco, and C. P. Koch, J. Phys. B **44**, 154011 (2011).
[38] M. M. Müller, D. M. Reich, M. Murphy, H. Yuan, J. Vala, K. B. Whaley, T. Calarco, and C. P. Koch, Phys. Rev. A **84**, 042315 (2011).
[39] S. Poletto, J. M. Gambetta, S. T. Merkel, J. A. Smolin, J. M. Chow, A. D. Córcoles, G. A. Keefe, M. B. Rothwell, J. R. Rozen, D. W. Abraham, et al., Phys. Rev. Lett. **109**, 240505 (2012).
[40] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, et al., Phys. Rev. Lett. **107**, 080502 (2011).
[41] Strictly speaking, one should enforce the dynamical map on the optimization subspace to be unital, i.e., both norm conserving and contracting. This could only be achieved by employing the trace distance of the ideal and actually time evolved third state, not their Hilbert Schmidt product, in the optimization functional, cf. Appendix A. However, for all practical purposes, the Hilbert Schmidt product turns out to be sufficient.
[42] A generalization to non-linear couplings and equations of motion is straightforward following Ref. [33].
[43] Optimization using non-convex functionals is possible but requires additional terms in the update equation for the field to preserve monotonicity of the convergence [33].