

Introduction to quantum control: From basic concepts to applications in quantum technologies

Christiane P. Koch 

Abstract Quantum control refers to our ability to manipulate quantum systems. This tutorial-style chapter focuses on the use of classical electromagnetic fields to steer the system dynamics. In this approach, the quantum nature of the control stems solely from the underlying dynamics, through the exploitation of destructive and constructive interference to reach the control target. We first discuss two basic control principles — coherent control which uses manipulation in frequency or time to design these interferences, and adiabatic following where access to the control target is enabled by tracking the time-dependent ground state. For complex control targets and system dynamics that exceed the scope of these basic principles, optimal control theory provides a powerful suite of tools to design the necessary protocols. A key consideration for the successful application of optimal control theory is a proper choice of the optimization functional. All concepts are illustrated using recent work from my research group, with a focus on controlling atoms and superconducting qubits. The chapter concludes with an outlook on integrating coherent control with engineered dissipation and a discussion of open questions in the field.

1 Introduction

The superposition principle captures the essence of quantum mechanics and manifests itself in quantum interference. In textbook examples, these interferences are typically taken as given, arising from static interactions. Remember the dihydrogen molecule where the difference between the bonding and antibonding orbitals is just the sign, i.e., a relative phase of π , in the respective superpositions of atomic orbitals. Coherent quantum control involves a change in perspective: Interferences can be created dynamically. The interaction of a quantum system with an external field allows

Christiane P. Koch
Freie Universität Berlin, Dahlem Center for Complex Quantum Systems & Fachbereich Physik,
Arnimallee 14, 14195 Berlin, Germany, e-mail: christiane.koch@fu-berlin.de

for creating superpositions and manipulating the relative phases in these superpositions. Under this angle, quantum control appears as a way to see the superposition principle at work.

Quantum control also refers to our ability to manipulate quantum systems in the desired way. It is a field at the intersection of physics, mathematics, chemistry, and engineering. From a practical viewpoint, it is not necessary to think about superpositions. Rather, quantum control can be seen as an engineering discipline, where the equations of motion are given by physics (or chemistry). Using the language of mathematics, it provides a set of tools to implement tasks that arise in the operation of quantum computers or quantum sensors. In this mindset, quantum control is at its best when it can be used in a black box way. While fully justified when the goal is to design and build a quantum device, the engineering perspective falls short of harnessing the exploratory power of quantum control.

The present lecture attempts a more holistic view, uniting perspectives from physics, chemistry, mathematics, and engineering. Such an interdisciplinary approach establishes a constructive cycle for quantum control: Analyzing the requirements for operating a quantum device leads to the design of new algorithms; implementing these algorithms yields practical protocols; and analyzing the underlying control mechanism improves our microscopic understanding. The lecture is structured as follows: Section 2 will introduce the basic concepts, starting from the statement of the control problem in Sec. 2.1. It provides an overview over the theoretical tools that are nowadays available to solve quantum control problems – from the basic principles of coherent control in frequency and time in Sec. 2.2 and control based on adiabatic evolution in Sec. 2.3 to optimal control theory in Secs. 2.4 and 2.5 and controllability analysis in Sec. 2.6. The application of these tools will be illustrated in Sec. 3 with selected examples, with a focus on the control of open quantum systems and a strong bias towards work from my group (a reader interested in a more comprehensive summary of the state of the art is referred to Refs. [1, 2]). After discussing the design of suitable figures of merit for open quantum systems in Sec. 3.1, the discussion is structured according to the control strategies for open quantum systems. Control avoiding decoherence is covered in Sec. 3.2, followed by a discussion of quantum reservoir engineering as a strategy that exploits the environment in Sec. 3.3. As an intermezzo, Sec. 3.4 presents advances in controllability analysis, as a tool to select suitable interactions for the control of quantum reservoir engineering. A recent addition to the set of control strategies, explained in Sec. 3.5, is the use of auxiliary quantum degrees of freedom in order to engineer desired dissipation. The lecture concludes in Sec. 4 with a discussion of open questions and possible avenues for future research, the selection of which again reflects a personal preference.

2 Basic concepts and methods

2.1 Statement of the control problem

The term “control” is widely used in quantum physics and often with rather diverse meanings. In this lecture, we take the perspective of *control theory* in the mathematical sense; that is, we ask how a dynamical system can be steered from a given initial state to a desired final state using one or several external control knobs. Here, the notion of dynamical system implies that we can write down (and also solve, either analytically or numerically) the system’s equation of motion.

This formalization of the control problem is of course not limited to quantum physics. In fact, optimal control theory (where the control problem is solved using variational calculus to find a minimum or maximum [3]) has been developed to steer the motion of objects that follow the classical laws of motion. A famous example is the landing of a space craft on the moon. For simplicity, we assume that we can describe the space craft as a point particle. In this case, the equation of motion is given by Newton’s law,

$$m \ddot{\mathbf{r}} = \mathbf{F},$$

and the initial and final states are specified in terms of the point particle’s initial, respectively desired, final position and velocity,

$$\mathbf{r}(t = 0) = \mathbf{r}_0, \mathbf{v}(t = 0) = \mathbf{v}_0 \quad \mathbf{r}(t = t_f) = \mathbf{r}_f, \mathbf{v}(t = t_f) = \mathbf{0},$$

where the latter ensures a soft landing of the space craft. Control on the space craft is exerted by burning fuel which can be described by the corresponding force, $\mathbf{F}_{burn}$ and which also leads to a time-dependent mass $m(t)$. In this very simplified picture of a space craft’s motion, the optimal control problem can be solved in closed form [4].

One of the areas of quantum physics, where control theory has been employed early on, is the control of photochemical reactions, i.e., the breaking (and much later also the making [5]) of chemical bonds in molecules using laser light. For molecules in the gas phase, the interaction with the laser typically occurs on such short timescales that the dynamics is well described by Schrödinger’s equation,

$$\frac{\partial |\psi\rangle}{\partial t} = -i\hat{H}(t) |\psi\rangle \quad |\psi(t = 0)\rangle = |\psi_0\rangle, |\psi(t = t_f)\rangle = |\psi_{target}\rangle.$$

Here and throughout we use units such that $\hbar = 1$. $|\psi_0\rangle$ and $|\psi_{target}\rangle$ describe the molecule’s initial and desired final state. The laser being the external control knob, the Hamiltonian comprises both the model for the molecule, $\hat{H}_0$ (including electronic and possibly also vibrational and rotational degrees of freedom), and the molecule’s interaction with the laser. The laser-molecule interaction is usually well described in the electric dipole approximation such that the total Hamiltonian is given by

$$\hat{H}(t) = \hat{H}_0 + \hat{H}_I(t) = \hat{H}_0 + \hat{\mathbf{d}} \cdot \mathbf{E}(t), \quad (1)$$

where $\hat{\mathbf{d}}$ is the molecule's dipole moment and $\mathbf{E}(t)$ the electric field. In many experiments, more than a single initial state is populated; in which case Schrödinger's equation needs to be solved for each of these states. Expectation values are then obtained as an incoherent average over all the states with their respective weights [6], for example the Boltzmann weights for a thermal initial state. This is not just a technicality but important for the observability of quantum control: Typically, control leverages the superposition principle of quantum mechanics to realize the desired dynamics and thus gets reduced the less pure the initial state is [5].

The purity with which initial states can be prepared, together with the accuracy with which we know the Hamiltonian, is what makes the field of quantum information science an ideal playground for quantum control [2]. At the same time, the requirements on the fidelity, i.e., the accuracy with which the desired dynamics is realized, are very stringent. Moreover, there is often no clear separation between the timescales of decoherence and the desired dynamics. In order to find control schemes that realize the best possible fidelity, it is then necessary to include decoherence and describe the system state by $\hat{\rho}$. A common description for such dynamical systems uses the GKLS master equation [7, 8], named after Gorini, Kossakowski, Lindblad and Sudarshan [9],

$$\frac{\partial \hat{\rho}}{\partial t} = -i[\hat{H}, \hat{\rho}]_- + \mathcal{L}_D(\hat{\rho}) = -i[\hat{H}, \hat{\rho}]_- + \sum_j \left(\hat{L}_a \hat{\rho} \hat{L}_a^\dagger - \frac{1}{2} \hat{L}_a^\dagger \hat{L}_a \hat{\rho} - \frac{1}{2} \hat{\rho} \hat{L}_a^\dagger \hat{L}_a \right), \quad (2)$$

where the dissipative part of the evolution is generated by $\mathcal{L}_D$. The jump operators $\hat{L}_a$ describe processes leading to loss of energy and phase, either phenomenologically or derived from a microscopic picture of the system's interaction with its environment [7, 8]. Depending on the control target (as will be discussed in more detail below), we have to consider a single initial state, $\hat{\rho}(t=0) = \hat{\rho}_0$, or a set thereof, $\hat{\rho}_{0,j}$, $j = 1, \dots, M$, and similarly for the final state(s), $\hat{\rho}_{target,j}$, $j = 1, \dots, M$. Unless the external drives are very strong [10, 11], they enter only the Hamiltonian and not $\mathcal{L}_D$.

For the simple example of a qubit, in the most general case

$$\hat{H}(t) = \frac{\omega_0}{2} \hat{\sigma}_z + u_1(t) \hat{\sigma}_z + u_2(t) \hat{\sigma}_x, \quad (3)$$

and the jump operators are $\hat{L}_1 = \sigma^-$, $\hat{L}_2 = \sigma^+$, describing decay, and $\hat{L}_3 = \sigma_z$ for pure dephasing. Recall that the coherent dynamics of a driven two-level system in the rotating-wave approximation (RWA, see Sec. 5.1) can be solved analytically — these are the famous Rabi oscillations [12]: The frequency of the oscillations is determined by the strength of the external control and its detuning from resonant driving. For zero detuning, the populations of the two levels oscillate between zero and one, as long as the external control is applied. This gives rise to control via so-called π - and $\pi/2$ -pulses, the workhorses of many coherent spectroscopies, where the duration of the external control is chosen such that the populations are inverted, respectively brought from a single level into an equal superposition of both levels. The control problem for a two-level system can in many cases be solved analytically

even with explicit account of decoherence [13], using the techniques of optimal control theory [3].

Given the formalization of the control problem in terms of control target and equations of motion, we can now look at ways to solve it. We will start with two intuitive approaches — coherent control which leverages interference to implement the desired dynamics, and adiabatic following where the system state slowly transforms into the desired one, before presenting formal optimal control theory. The latter is generally applicable (provided we can solve the dynamics, analytically or numerically) but may yield solutions which are not always easy to interpret. In contrast, the first two are somewhat limited in their applicability, but provide clear insight into the control mechanism. They can serve as a starting point for further studies but also help interpret numerical optimal control results.

2.2 Coherent control in frequency and time

In Bohr’s famous gedankenexperiment, an electron, when sent through a double-slit, interferes with itself. If a detector were to record the electron, it would yield an interference pattern where constructive and destructive interference take turns, depending on the distance between the slits. Nowadays such experiments are routinely carried out in the lab, with a grating instead of a double-slit and molecules rather than electrons [14]. Even though these experiments are carried out with many molecules, it is always a single molecule that interferes with itself. The idea of a quantum object interfering with itself due to wave-particle duality can be made to work also in a more abstract sense, and coherent control takes it from the real space surrounding us to mathematical Hilbert space. Remarkably, the consequences are measurable, just as the interference pattern behind a double slit (or grating).

In coherent control, interferences are created between different “quantum pathways” connecting initial and final states [15, 16]. These quantum pathways correspond to time evolutions due to the action of an external field. For example, in an atom, electric dipole transitions excited resonantly by the absorption of one photon or non-resonantly by the absorption of three photons, can connect the same initial and final states, and varying the relative phase between these two pathways modifies the population in the final state [15]. Because of the two laser colors involved, this is sometimes referred to as “bichromatic control”. In a molecule without inversion symmetry, i.e., a chiral molecule, electric dipole selection rules also allow for the interference between one-photon and two-photon pathways [17].

The simplest example to see the interference [15], depicted in Fig. 1 (left), assumes the initial state to be given by a coherent superposition of two levels, $|\psi(t)\rangle = c_1^0 |1\rangle + c_2^0 |2\rangle$, $c_{1,2}^0 \in \mathbb{C}$, that are excited to the same final state, $|f\rangle$. A two-color laser field resonantly excites these two transitions, $\mathbf{E}(t) = \mathbf{E}_1 \cos(\omega_1 t) + \mathbf{E}_2 \cos(\omega_2 t + \varphi)$ where φ denotes the relative laser phase. We will describe the process within first order time-dependent perturbation theory [15],




Fig. 1 Coherent control via quantum pathway interference. Left: Spectral interference, starting from an initial superposition state (left). It is the laser phase φ that controls the interference and hence the final population in $|f\rangle$. Right: Temporal interference where a pump pulse creates a superposition state which is subsequently further excited to a common final state by a probe pulse. It is the time delay between pump and probe pulse which controls the interference and hence the final population in $ket f$.

$$|\psi^{(1)}(t)\rangle = \frac{1}{i} \int_{t_0}^t dt' e^{-iH_0(t-t')} H_I(t') e^{-iH_0(t'-t_0)} |\psi(t_0)\rangle . \quad (4)$$

Since the initial state is a superposition of $|1\rangle$ and $|2\rangle$, the only non-zero component of the final state in first order perturbation theory is $\langle f | \psi^{(1)}(t_f) \rangle = c_f^{(1)}(t_f)$. Given the Hamiltonian, Eq. (1), with $H_0 = \text{diag}(E_1, E_2, E_f)$ and expressing the interaction with the two-color laser field in terms of Rabi frequencies $\Omega_{ij} = \mathbf{d}_{fi} \cdot \mathbf{E}_j$ where $\mathbf{d}_{fi}$ denotes the electric dipole transition matrix element from level i to f , Eq. (4) becomes

$$\begin{aligned} c_f^{(1)}(t) = & \frac{\Omega_{11}}{i} \int_{t_0}^t dt' e^{-iE_f(t-t')} \cos(\omega_1 t') e^{-iE_1(t'-t_0)} c_1(t_0) \\ & + \frac{\Omega_{12}}{i} \int_{t_0}^t dt' e^{-iE_f(t-t')} \cos(\omega_2 t' + \varphi) e^{-iE_1(t'-t_0)} c_1(t_0) \\ & + \frac{\Omega_{21}}{i} \int_{t_0}^t dt' e^{-iE_f(t-t')} \cos(\omega_1 t') e^{-iE_2(t'-t_0)} c_2(t_0) \\ & + \frac{\Omega_{22}}{i} \int_{t_0}^t dt' e^{-iE_f(t-t')} \cos(\omega_2 t' + \varphi) e^{-iE_2(t'-t_0)} c_2(t_0) . \end{aligned}$$

Introducing $\tilde{c}_i^0 = e^{-iE_f t} e^{iE_i t_0} c_i(t_0)$ and $\omega_{fi} = E_f - E_i$, we obtain

$$c_f^{(1)}(t) = \frac{\Omega_{11}}{2i} \tilde{c}_1^0 \int_{t_0}^t dt' \left(e^{i(\omega_{f1} + \omega_1)t'} + e^{i(\omega_{f1} - \omega_1)t'} \right) + \dots$$

The terms that oscillate with the sums $\omega_{fi} + \omega_j$ of the transition and laser frequencies describe highly off-resonant transitions and are neglected within the rotating wave approximation,

$$c_f^{(1)}(t) = \frac{\Omega_{11}}{2i} \tilde{c}_1^0 \int_{t_0}^t dt' e^{i(\omega_{f1} - \omega_1)t'} + \frac{\Omega_{12}}{2i} \tilde{c}_1^0 e^{-i\varphi} \int_{t_0}^t dt' e^{i(\omega_{f1} - \omega_2)t'} \\ + \frac{\Omega_{21}}{2i} \tilde{c}_2^0 \int_{t_0}^t dt' e^{i(\omega_{f2} - \omega_1)t'} + \frac{\Omega_{22}}{2i} \tilde{c}_2^0 e^{-i\varphi} \int_{t_0}^t dt' e^{i(\omega_{f2} - \omega_2)t'} .$$

Finally, we assume that the interaction with the laser field is much longer than any other timescale. This allows us to extend the integral boundaries to $\pm\infty$ such that we can replace the integrals by $2\pi\delta(\omega_{fi} - \omega_j)$. We then obtain for the final state amplitude

$$c_f^{(1)} = -i\pi \left(\Omega_{11} \tilde{c}_1^0 + \Omega_{22} \tilde{c}_2^0 e^{i\varphi} \right) ,$$

since we had assumed the laser frequencies $\omega_{1/2}$ to match the two transition frequencies. Note that the result of the transition $|1\rangle \rightarrow |f\rangle$ being excited only by the ω_1 -component and not the ω_2 -component of the field (and conversely for $|2\rangle \rightarrow |f\rangle$) can also be obtained by first going into the interaction picture or a rotating frame and then invoking the (two-photon) rotating wave approximation [18], see App.5.2). The interference shows up in the population, i.e., in the quantity that is measured:

$$|c_f^{(1)}(t \rightarrow \infty)|^2 = \pi^2 \left(\Omega_{11}^2 |\tilde{c}_1^0|^2 + \Omega_{22}^2 |\tilde{c}_2^0|^2 + 2\Re \left(\Omega_{11} \Omega_{22} (\tilde{c}_1^0)^* \tilde{c}_2^0 e^{-i\varphi} \right) \right) . \quad (5)$$

More precisely, it is the last term that can take positive or negative values, corresponding to constructive or destructive interference. The interference depends on the values $\tilde{c}_1^0$, $\tilde{c}_2^0$, i.e., the initial populations and eigenenergies, as well as the relative phase of the laser field φ . It is the latter that can easily be tuned in an experiment and is thus the main control knob.

Bichromatic control, and its generalizations to more than two quantum pathways, see e.g. ref. [17], create the desired interference spectrally. A perfectly equivalent way is a temporal perspective [16]. Let us again assume that we start with a coherent superposition of two levels. In its most general form, the state is written

$$|\psi(t)\rangle = c_a(t) e^{-iE_a t} |a\rangle + c_b(t) e^{-iE_b t} |b\rangle .$$

The natural time evolution $e^{-iE_i t}$ of the two levels is irrelevant, when populations are measured. If one can, however, directly probe the superposition, it gives rise to a time-dependent relative phase which results in the so-called "quantum beats". We can again use first order perturbation theory to see it. Let's assume we apply a weak probe field with a finite spectral width, and thus also finite temporal width i.e., a pulse, that excites both $|a\rangle$ and $|b\rangle$ to a common final state $|f\rangle$. According to perturbation theory, the population in $|f\rangle$ is proportional to $|\langle f | \hat{d} | \psi(t) \rangle|^2$. Assuming for simplicity an equal superposition, this results in

$$P_f(t) \sim |\langle f | \hat{d} | \psi(t) \rangle|^2 = \frac{1}{2} |d_{fa}|^2 + \frac{1}{2} |d_{fb}|^2 + |d_{fa}| \cdot |d_{fb}| \cos((E_b - E_a)t).$$

Again, the interference shows up in the last term — due to the sign change in the cosine, the last term can cancel or augment the other two terms, resulting in oscillations between zero and maximal population P_f .

The simplest control scheme realizing coherent control in time is pump-probe spectroscopy, illustrated in Fig. 1 (right): A first pulse, or ”pump”, creates the superposition state $|\psi(t)\rangle$, and a second pulse probes the time evolution of the superposition by excitation to the common final state. Note that this control protocol is necessarily time-dependent: A single pulse (with small to moderate pulse intensity) can create a superposition state only when it has a spectral bandwidth $\Delta\omega > 0$ ensuring that all states of the superposition are resonantly excited. The spectral bandwidth translates into a finite pulse duration τ , for example, $\tau = 1/(\pi\Delta\omega)$ for transform-limited (unshaped) Gaussian pulses. Pump-probe control is widely used across different timescales and various physical systems, from attosecond pulses for electronic motion in molecules [19] to Ramsey interferometry in quantum sensing [20], with the common point being that the pulse duration must be much shorter than the dynamical timescale to be probed.

It is important to note that while we have used perturbation theory to elucidate the control mechanism, coherent control based on spectral or temporal interference is not limited to weak fields. In fact, the stronger the fields, the more pathways will be excited that can interfere with each other [21]. Even outside of its validity regime, perturbation theory can be a very useful tool to rationalize control mechanisms since the book-keeping of the different interactions with the external field allow for easy identification of the interfering quantum pathways. A complementary perspective, relevant for strong external fields, is provided by analyzing the dynamics in terms of field-dressed states. These are at the core of control based on adiabatic following.

2.3 Adiabatic following

The starting point for control via adiabatic following is the observation that the Hamiltonian $\hat{H}(t)$ can be diagonalized at any given instant of time t . The instantaneous eigenstates are also called ”dressed states”, as opposed to the bare states, i.e., the eigenstates of $\hat{H}_0$. While both eigenvalues and eigenvectors of $\hat{H}(t)$ depend on time, one can seek to stay in the same eigenstate, let’s say the ground state. This is possible if the time-dependence of $\hat{H}(t)$ is sufficiently slow, since then the eigenstates for all t are adiabatically connected. The condition for adiabaticity can be derived from the unitary transformation $\hat{U}(t)$ that diagonalizes $\hat{H}(t)$. Moreover, identifying which terms violate adiabaticity allows for designing control protocols that mitigate this violation. They are known as counterdiabatic driving [22] or shortcuts to adiabaticity [23]. We will discuss this using again the example of a (strongly driven) two-level system.

Assume we have found the unitary $\hat{U}(t)$ that diagonalizes the Hamiltonian $\hat{H}(t)$ at time t ,

$$\hat{U}(t)\hat{H}(t)\hat{U}^\dagger(t) = \hat{D}(t),$$

where $\hat{D}(t)$ denotes the diagonal matrix of the time-dependent eigenvalues. We can then use $\hat{U}(t)$ to transform the time-dependent Schrödinger equation into the instantaneous eigenbasis of $\hat{H}(t)$ (see App. 5 for a reminder on frame transformations),

$$i\frac{\partial}{\partial t}|\tilde{\psi}(t)\rangle = \left(\hat{D}(t) - i\hat{U}(t)\frac{\partial\hat{U}^\dagger(t)}{\partial t}\right)|\tilde{\psi}(t)\rangle. \quad (6)$$

If the system starts in an eigenstate $|\phi_n(t=0)\rangle$ of $\hat{H}(t=0)$ and the dynamics is perfectly adiabatic, the state will evolve according to $\exp(-i\lambda_n(t)t)$, where $\lambda_n(t)$ is one of the entries of $\hat{D}(t)$. In other words, the second term on the right-hand side of Eq. (6) is only relevant beyond adiabatic evolution, it captures the non-adiabatic (or "diabatic") transitions! Adiabaticity can be ensured if this term is negligible compared to the first one. We will see below for the example of the two-level system that this requires large Rabi frequencies, i.e., large field amplitudes, and long and smooth pulses. An alternative is provided by counterdiabatic driving which adds further external controls to cancel the second term and thus effectively suppress non-adiabatic transitions [22].

Moving to the two-level system with general Hamiltonian

$$\hat{H}(t) = \begin{pmatrix} E_a & W_{ab}(t) \\ W_{ab}^*(t) & E_b \end{pmatrix}, \quad (7)$$

the unitary that diagonalizes $\hat{H}(t)$ is given by

$$\hat{U}(t) = \begin{pmatrix} \cos\frac{1}{2}\vartheta e^{i\frac{1}{2}\varphi} & \sin\frac{1}{2}\vartheta e^{-i\frac{1}{2}\varphi} \\ \sin\frac{1}{2}\vartheta e^{i\frac{1}{2}\varphi} & \cos\frac{1}{2}\vartheta e^{-i\frac{1}{2}\varphi} \end{pmatrix}, \quad (8)$$

where the two angles are defined via

$$\tan\vartheta = \frac{2|W_{ab}|}{E_a - E_b}, \quad W_{ab} = |W_{ab}|e^{i\varphi}.$$

This yields for the diabatic transition matrix in Eq. (6)

$$i\hat{U}(t)\frac{\partial\hat{U}^\dagger(t)}{\partial t} = \frac{1}{2}\begin{pmatrix} -\frac{\partial\varphi}{\partial t} & i\frac{\partial\vartheta}{\partial t} \\ i\frac{\partial\vartheta}{\partial t} & \frac{\partial\varphi}{\partial t} \end{pmatrix}. \quad (9)$$

The diagonal terms in Eq. (9) add a non-adiabatic phase to the time evolution but do not induce transitions between eigenstates $|\phi_n(t)\rangle, |\phi_m(t)\rangle$. It is the off-diagonal term, i.e., the time derivative of $\vartheta(t)$, that determines whether non-adiabatic transitions will occur. We thus need to compare $\partial\vartheta/\partial t$ with $\hat{D}(t)$. Denoting the two time-dependent eigenvalues by $E_\pm(t)$,




Fig. 2 Adiabatic following in a two-level system, as reflected in the instantaneous eigenenergies $E_{\pm}(t)$ and the time-dependent population of the bare, or field-free, states, i.e., the eigenstates of H_0 , $|\phi_{a/b}\rangle$.

$$E_{\pm}(t) = \frac{1}{2}(E_a + E_b) \pm \frac{1}{2}\sqrt{(E_a - E_b)^2 + 4|W_{ab}|^2},$$

the condition for adiabaticity becomes

$$\frac{1}{2} \left| \frac{\partial \theta}{\partial t} \right| \ll |E_+(t) - E_-(t)|. \quad (10)$$

Expressing the Hamiltonian (7) in terms of time-dependent Rabi frequency and detuning,

$$\hat{H}(t) = -\frac{1}{2}(\Delta_L(t)\hat{\sigma}_z + \Omega_0(t)\hat{\sigma}_x),$$

we find

$$\dot{\theta} = \frac{\Delta_L \dot{\Omega}_0 - \Omega_0 \dot{\Delta}_L}{\Omega(t)} \quad (11)$$

with $\Omega(t) = \sqrt{\Omega_0^2(t) + \Delta_L^2(t)}$ the generalized Rabi frequency. In order to fulfill condition (10), we need

$$\frac{1}{2} |\Delta_L \dot{\Omega}_0 - \Omega_0 \dot{\Delta}_L| \ll \left(\Omega_0^2(t) + \Delta_L^2(t) \right)^{3/2},$$

which translates into the anticipated requirements of long and smooth pulses (small $\dot{\Omega}_0$, small $\dot{\Delta}_L$) and strong fields (large Ω_0). While large detunings Δ_L are also good for adiabaticity, they are typically not helpful for the desired change in the instantaneous eigenstates.

Two choices of control parameters to realize adiabatic following in a two-level system are sketched in Fig. 2. In case (i), the detuning is kept fixed, and a pulse is applied, with its shape reflected in the time-dependence of the eigenvalues $E_{\pm}(t)$. In this case, if the system starts out in one of the eigenstates of $\hat{H}_0$, it will return to the same field-free state. A different result is obtained for a linear time-dependence of the detuning, as in case (ii), where, while adiabatically following one of the dressed

states, $|\phi_{\pm}(t)\rangle$, the population is completely transferred between the bare states, $|\phi_a\rangle$ and $|\phi_b\rangle$. The latter is the famous Landau-Zener model, here realized with a linearly chirped pulse driving a two-level atom. Figure 2 highlights an advantage of adiabatic following – it achieves *robust* population transfer, irrespective of the exact shape or duration of the pulse. This is in contrast to control via a π -pulse which requires very precise pulse amplitude and duration.

A popular and widely used example of adiabatic following is the so-called STIRAP protocol, where STIRAP stands for “stimulated Raman adiabatic passage” [24]. It realizes population transfer in a three-level system (see also App. 5.2), where $|1\rangle$ is connected to $|3\rangle$ only via $|2\rangle$ which is, however, subject to fast decay. If one would apply the first (pump) pulse to drive the population from $|1\rangle$ to $|2\rangle$ and then the second (so-called Stokes) pulse to drive the transition from $|2\rangle$ to $|3\rangle$, the transfer could never reach a fidelity of one due to population decay when in $|2\rangle$. It turns out, however, that, depending on the detunings of the two pulses, one eigenvalue of the time-dependent Hamiltonian may vanish. The corresponding time-dependent eigenstate has non-zero projections only with the bare states $|1\rangle$ and $|3\rangle$ [24]. If one can adiabatically follow this “dark” state, the decaying (and thus “bright”) state $|2\rangle$ never gets populated, allowing for perfect population transfer from $|1\rangle$ to $|3\rangle$. This requires both pulses to overlap in time and, counterintuitively, the Stokes pulse to precede the pump pulse [24]. It also requires strong fields to ensure adiabaticity.

Since large field strengths are not always available, let us inspect the alternative to ensure adiabaticity, counterdiabatic driving. The idea of counterdiabatic driving consists in adding an additional control $\hat{H}_{CD}(t)$ to the Hamiltonian $\hat{H}(t)$ that suppresses non-adiabatic transitions [22]. In order to calculate the counterdiabatic term for a two-level system, it is useful to start in a rotating frame in which $\hat{H}_{RWA}(t) = -\frac{\hbar}{2}(\Omega_0(t)\sigma_x - \Delta_L(t)\sigma_z)$ with $\Delta_L(t) = \omega_0 - \omega_L(t)$, see App. 5.1. With the counterdiabatic term, the Hamiltonian becomes $\hat{H}_{tot}(t) = \hat{H}_{RWA}(t) + \hat{H}_{CD}(t)$. Transforming into the instantaneous eigenbasis of $\hat{H}_{RWA}(t)$, we obtain for the total Hamiltonian

$$\hat{H}'_{tot}(t) = \hat{D}(t) - \hat{U}(t) \frac{\partial \hat{U}^+(t)}{\partial t} + \hat{U}(t) \hat{H}_{CD}(t) \hat{U}^+(t).$$

Now $\hat{H}_{CD}(t)$ is chosen such that the last two terms cancel, that is,

$$\hat{H}_{CD}(t) = i \frac{\partial \hat{U}^+(t)}{\partial t} \hat{U}(t). \quad (12)$$

For the example of the two-level system, we can use Eq. (8) to evaluate the condition (12) and obtain

$$\hat{H}_{CD}(t) = \frac{\dot{\vartheta}}{2} \begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix} = \frac{\dot{\vartheta}}{2} \hat{\sigma}_y.$$

Recalling $\dot{\vartheta}$ given in Eq. (11), the counterdiabatic drive has a Lorentzian shape and, being given by $\hat{\sigma}_y$ compared to $\hat{\sigma}_x$ in $\hat{H}_{RWA}$, it is phase-shifted compared to the original control. Transforming from the rotating frame back into the lab frame,

this implies the counterdiabatic drive to oscillate as $\sin(\omega_L(t)t)$, when the original drive oscillates as $\cos(\omega_L(t)t)$ [22]. When calculating the counterdiabatic term for the STIRAP-Hamiltonian, ironically, the counterdiabatic term is found to drive the (non-existent) $|1\rangle \rightarrow |3\rangle$ transition [22]. This illustrates one of the limitations of counterdiabatic driving, a second limitation being that one needs to be able to diagonalize $\hat{H}(t)$ and determine $\hat{U}(t)$ in order to derive the counterdiabatic control term.

2.4 Optimal control theory: Fundamentals and state-to-state optimization

We now return to formal control theory [3] which is, in essence, an application of variational calculus to the dynamics of a quantum system. The starting point is the ability to solve the equations of motion for $|\psi(t)\rangle$, respectively $\hat{\rho}(t)$, analytically [13] or numerically [1, 2] or on a quantum device [25]. In other words, optimal control theory constructs solutions of the control problem from the knowledge, even if implicit, of the system dynamics. In addition to the equations of motion, the second ingredient is the optimization functional, also referred to as cost functional or target functional. It formalizes the control target in terms of the figure of merit as well as additional constraints, for example on power consumption or smoothness of the controls. Indeed, all information on the physics of the control problem is encoded in these two ingredients — optimization functional and equations of motion — to an otherwise mathematical recipe. It is thus important to design the optimization functional as carefully as one constructs the equations of motion.

Before discussing a first illustrative example, let us consider the general requirements that an optimization functional must meet [26, 27]. First of all, the optimization functional must be real-valued, such that it establishes an order relation that we can use to quantify improvement. It is customary, though not necessary, to normalize the functional to take values between 0 and 1. Second, it is useful to distinguish between so-called final-time costs and intermediate-time costs,

$$J[u(t)] = J_{t_f}[u(t)] + \frac{w}{t_f} \int_0^{t_f} g(\psi(t), u(t), t) dt. \quad (13)$$

Typically, the final-time cost J_{t_f} consists of the figure of merit to be optimized, whereas intermediate-time costs g capture constraints on the dynamics. In Eq. (13), we have assumed that we consider the equation of motion for $|\psi(t)\rangle$, but generalization to $\hat{\rho}(t)$ is straightforward, $w \in \mathbb{R}$ is a weight, and the external controls are denoted by $u(t)$. $J[u(t)]$ is a functional of $u(t)$ in the sense of variational calculus, i.e., asking for an extremum of J , $\delta J = 0$, allows us to determine the function (or “shape” of) $u(t)$. The third requirement that a well-defined functional $J[u(t)]$ must meet is that it takes its optimum value if and only if the optimum is reached. When writing $J[u(t)]$ as in Eq. (13), this implies that the figure of merit takes its minimum

or maximum, depending on the optimization problem, and the second term vanishes, indicating that the constraints are enforced.

The simplest example for a control problem is to drive the system from a given initial state $|\psi_0\rangle$ to a desired target state $|\psi_{target}\rangle$. Assuming that the target state shall be reached at time t_f , this can be expressed as

$$J_{t_f}[u(t)] = \left| \langle \psi_0 | \hat{U}^+(t_f, 0; u(t)) | \psi_{target} \rangle \right|^2. \quad (14)$$

Equation (14) is very intuitive since $\hat{U}(t_f, 0; u(t)) |\psi_0\rangle$ is nothing but the system state at the final time that is obtained from time-evolving the initial state $|\psi_0\rangle$ under the external control $u(t)$, whatever $u(t)$ may be. Thus, J_{t_f} measures the overlap of the actually obtained state $|\psi(t_f)\rangle$ with the desired state $|\psi_{target}\rangle$. In Eq. (14), we have taken the square modulus to ensure the real-valuedness of the functional; we could have equally well taken its real (or imaginary) part [28].

A recent application has used Eq. (14) to improve the preparation of circular Rydberg states in rubidium atoms [29, 30]. These are states of the valence electron with a large value of the principal quantum number n and the maximum value that the projection quantum number can take, $m_\ell = \pm(n-1)$. Here, we have made use of the close analogy between the hydrogen atom, where the state of the electron is described by the three quantum numbers n, ℓ, m_ℓ [12], and the valence electron in alkali atoms with the same set of quantum numbers. The name “circular” derives from the shape of the electron’s orbit in that state; the interest in these states is due to their intrinsic protection against decoherence which makes them good candidates for quantum sensing [31] or quantum computing [32, 33]. The protection arises from the very small number of electric dipole-allowed transitions that involve a circular state. This means, however, that it is also not easy to prepare a circular state. The challenge is in the large angular momentum transfer that is required to prepare a circular state.

Typically, one starts with laser excitation from the electronic ground state. Since the Rydberg state is accessed via a two-photon absorption, it results in a state with the desired value of the principal quantum number n and $m_l = 2$. One can then apply a radio-frequency (RF) field which quasi-resonantly drives a sequence of transitions between levels m_l and $m_l + 1$ until the circular state is reached [34]. This protocol can be rationalized in terms of a π -pulse, generalized from a two-level system to a system with $N = n - 3$ levels, and visualized very similarly to the Bloch sphere representation of a two-level system, using the Husimi- Q function [35] for the N -level system. The circular state corresponds to the North pole of the generalized Bloch sphere. All the intermediate levels are not protected against decoherence which is why the generalized π -pulse should be as short as possible. It can, however, not be made arbitrarily short simply by increasing the pulse amplitude, unlike in the case of the two-level system: When the pulse amplitude becomes too large, the pulse drives off-resonantly also the transition from $m_\ell = 2$ to $m_\ell = 1$. When balancing the requirements of avoiding decoherence in the intermediate levels and minimizing population “loss” to $m_\ell = 1$ in the best possible way, the circular state is prepared with a fidelity of about 80% [30].




Fig. 3 Preparation of circular Rydberg state with RF pulse derived with optimal control theory: The pulse (with amplitudes of the two quadratures shown in the top left panel) transfers the population initially in $m_\ell = 2$ to $m_\ell = 51$. The lower left panel compares the experimentally measured populations (crosses) to the theoretical prediction (solid lines), with the colors encoding the m_ℓ states as shown on the right. The inset displays the Husimi- Q functions of the initial state and the state after 15 ns, respectively ≈ 40 ns. The latter, also marked by the vertical line in the left panels, is the time by which a spin-coherent state has been prepared. Adapted from Ref. [30].

With RF fields shaped according to the prediction from optimal control theory, the fidelity measured in the experiment, carried out by the cavity-QED group at Collège de France in Paris, increases to about 97% [30]. This number is mainly limited by the accuracy with which the circular state was detected. Theoretically, the fidelity can be brought much closer to one. Knowing of the limited detection accuracy in the experiment, the calculations were stopped at 99% [29]. Analysis of the optimized pulse revealed a two-stage protocol [29]: In the first 40 ns, the RF field creates a spin-coherent state which subsequently is rotated to the North pole of the Bloch sphere in the second stage. The shape of the RF field reflects the two stages – it consists of modulations during the first stage, whereas constant amplitude is sufficient during the second stage. Note that a constant pulse amplitude will typically not result from numerical optimization unless specifically enforced. In our case, we simply flattened the pulse based on the understanding of the dynamics in order to reduce pulse complexity and ease experimental implementation of the pulse [30].

Encouraged by the seamless interplay of theory and experiment in the circular state preparation, we considered the preparation of a non-classical superposition state $(|m_\ell = 1\rangle + |m_\ell = 51\rangle)/\sqrt{2}$ as another example. For this control task, no protocol was previously known. Similarly to the circular state preparation, a comparatively simple pulse shape was derived with optimal control theory [30]. It consists of modulations in the first 50 ns and a linear chirp in the remaining part and results in an equally high fidelity of about 97%, again limited by the experimental detection capability [30].

One may wonder whether such a seamless interplay of theory and experiment was to be expected. The answer is yes! There are two reasons for this optimism: First of all, the Hamiltonian describing the valence electron in alkali atoms, that is, the energy levels but also all transition matrix elements, is known to high precision.

In other words, our theoretical model of the system is very reliable. Second, the numerical optimization was carried out in a rotating frame. This eliminates all irrelevant timescales from the calculations and helps to avoid numerical artefacts, i.e., physically irrelevant features, in the pulse shape, reducing the overall pulse complexity.

But how does the optimization actually proceed? Let us return to the formal problem of determining the optimum of $J[u(t)]$. There exist two different approaches to solving the optimization problem numerically. Their key difference is whether they are based on J alone or whether they also make use of the gradient of J to construct changes to the controls $u(t)$ [2]. In the gradient-free approach, the controls $u(t)$ are parametrized, for example by expanding them into Fourier components (or any other suitable basis) and treating the frequencies and expansion coefficients as optimization parameters. The actual optimization is then carried out via a non-linear search in the parameter space, with suitable methods found for example in the NLOpt software package [36]. Gradient-free methods come with the advantage that the search is global in the parameter space. However, the search tends to get stuck when the number of parameters is large. In contrast, gradient-based approaches provide fast convergence, almost irrespective of the representation of the pulse. However they are local, performing the search with the information provided by the gradient around the initial “guess” pulse. The best of both worlds can be leveraged in hybrid optimization [37], consisting of a pre-optimization step using a gradient-free method with just a few parameters and a gradient-based optimization that uses the outcome of the pre-optimization as initial “guess”.

The state preparation in Rydberg atoms highlighted above was carried out with a particular version of the gradient-based approach, Krotov’s method [38]. The quantum version of Krotov’s method [39] has been our workhorse for quantum optimal control. Depending on the equations of motion, the figure of merit and additional constraints, Krotov’s method provides us with a constructive way to design an optimization algorithm, i.e., to derive a set of equations that are then solved numerically in an iterative way. The key advantage of Krotov’s method is its guaranteed monotonic convergence, irrespective of the physics of the control problem. Krotov’s method can be applied, in particular, to non-linearities, both with respect to the state in the equations of motion, including dissipation, and with respect to the controls; it also allows for non-convex target functionals. Having this powerful mathematical tool at hand, we are free to focus on the physics of the control problem, that is, on using the most appropriate equations of motion and finding the most suitable control target.

Before providing examples for control targets beyond preparing a certain desired state, let us look at one more level of detail as to what it entails to solve a quantum control problem by numerical optimization with Krotov’s method [40]. As with any gradient-based method, ensuring an extremum of the optimization functional requires both $\delta_u J = 0$ and $\delta_\psi J = 0$ since J depends both explicitly and implicitly, via the states, on the controls $u(t)$. The condition $\delta_\psi J = 0$ yields a dynamical equation for an auxiliary state, typically referred to as adjoint state, where the “initial” condition is given at the final time t_f . When the optimization targets a desired state, the initial

condition is simply this state, $|\psi_{target}\rangle$; in general, it is derived from the figure of merit J_{t_f} . Since the initial condition is given at the final time, the adjoint state needs to be propagated backwards in time. The condition $\delta_u J = 0$ yields the update rule for the external field. For Krotov's method where we minimize, as additional constraint, the change in the field, the rule is given by

$$\Delta u(t) = u^{new}(t) - u^{old}(t) \sim \Im \left[\langle \psi_{target} | \hat{U}^+ (t_f, t; u^{old}) \frac{\partial \hat{H}}{\partial u} \hat{U} (t, 0; u^{new}) | \psi_0 \rangle \right]. \quad (15)$$

Here, $\hat{U} (t, 0; u^{new}) | \psi_0 \rangle$ implies that the initial state is forward propagated from time 0 to t using the new (updated) control u^{new} , whereas $|\psi_{target}\rangle$ is backward propagated from time t_f to time t using the old control u^{old} . The derivative $\partial \hat{H} / \partial u$ yields the operator that couples to the external field, typically the dipole moment, cf. Eqs. (1) or (3). Thus, Eq. (15) implies that, at time t , the dipole moment matches the forward and backward propagated states to yield the change in the control $\Delta u(t)$ at that instant of time. The requirement to both forward and backward propagate is a hallmark of any gradient-based optimization in quantum control. The main difference between gradient-based algorithms is the way in which the update is applied — for Krotov's method, sequentially for each instant in time, as in Eq. (15), whereas for GRAPE [41], another popular algorithm, the update is applied concurrently for all t at once. At first glance, Eq. (15) appears to be implicit in $u^{new}(t)$, but it is sufficient to use two different time grids, shifted by half the numerical time step, for the states and the control [28], in order to solve Eq. (15). Use of optimal control theory is thus equivalent to iteratively solving Eq. (15), or generalizations thereof, starting from a first “guess” for $u^{old}(t)$.

2.5 Optimal control theory: Optimization functionals

How does Eq. (15) change when we use a different functional J in Eq. (13)? For example, a relevant application in the context of quantum technologies is the implementation of a desired unitary evolution [28] such as a gate in quantum computation. At first glance, one may think of it as simultaneous state-to-state transformations, all carried out by the same control $u(t)$. The number of states is given by the size of the Hilbert space on which the desired evolution is defined, i.e., $N = 2$ for a single-qubit gate or $N = 4$ for a two-qubit gate. However, the proper implementation of the desired unitary requires more than N simultaneous state-to-state transitions — also the relevant phase relations between the states must be ensured [28]. This is achieved by generalizing the scalar product between the desired target state $|\psi_{target}\rangle$ and the propagated initial state $\hat{U} (t_f, 0; u(t)) | \psi_0 \rangle$ in Eq. (14) to the Hilbert-Schmidt product for operators. Denoting the desired unitary evolution or quantum gate by $\hat{O}$, the final-time cost thus becomes

$$J_{t_f} = 1 - \frac{1}{N} \Re \left[\text{tr} \{ \hat{O}^+ \hat{U} (t_f, 0; u(t)) \} \right],$$

where we assume, without loss of generality, a minimization problem, i.e., the optimal value that J can take is zero. Often, qubits are carried by physical objects such as atoms or superconducting qubits with a size of Hilbert space much larger than $N = 2$. We then need to project the actually implemented evolution onto the subspace on which $\hat{O}$ is defined,

$$J_{t_f} = 1 - \frac{1}{N} \Re \left[\text{tr} \left\{ \hat{O}^\dagger \hat{P}_N \hat{U}(t_f, 0; u(t)) \hat{P}_N \right\} \right]. \quad (16)$$

Expanding out the trace in Eq. (16), we see the similarity to Eq. (14). The generalization of Eq. (15) is then also straightforward,

$$\Delta u(t) \sim \Im \left[\frac{1}{N} \sum_{k=1}^N \langle \psi_k | \hat{O}^\dagger \hat{U}^\dagger(t_f, t; u^{old}) \frac{\partial \hat{H}}{\partial u} \hat{U}(t, 0; u^{new}) | \psi_k \rangle \right].$$

Here, the backward propagation concerns the target states $\hat{O} | \psi_k \rangle$, and the update $\Delta u(t)$ is given by summing over all N states. In other words, the optimization of a quantum gate requires the forward, respectively backward, propagation of $2N$ states instead of just two. For a small number of qubits, this can easily be done on classical computers, and various software is available, for example GRAPE within QuTiP [42, 43] or Krotov's method as separate package [40], whereas for larger systems, a hybrid quantum-classical approach is necessary [25].

When the control task is to implement an entangling operation for two qubits, it may be advantageous to exploit the concept of local equivalence in the optimization – two two-qubit gates $\hat{A}$ and $\hat{B}$ are said to be locally equivalent when applying single-qubit (“local”) gates before and after $\hat{A}$ results in $\hat{B}$ and vice versa. For example, the controlled phasegate and the controlled NOT gate are locally equivalent. Assume that you want to implement a CNOT gate but the Hamiltonian governing the dynamics of the qubits is diagonal in the logical basis (as is the case, for example, for Rydberg atoms). An optimization targeting a CNOT gate will then fail, whereas an optimization targeting the local equivalence class of CNOT will be successful [44]. The corresponding target functional is expressed in terms of the so-called local invariants, three real numbers that characterize each local equivalence class, more precisely, in terms of the Euclidean distance between the local invariants of the desired and the actually realized gate [44].

The idea of targeting a whole class of quantum gates instead of a single specific one can be generalized further. This is interesting in the context of universal quantum computation where any n -qubit unitary can be realized by a sequence of single qubit gates and one entangling two-qubit gate [47]. These gates form a finite set, termed universal set of gates [47]. Every entangling operation requires an interaction between the qubits; but given an interaction, it is often not immediately obvious which entangling operation is most straightforwardly implemented. This question can be answered using optimal control targeting an arbitrary perfect entangler. Here, the target functional minimizes the distance of the local invariants of the realized gate to the space of perfect entanglers [45, 46], shown as the shaded area in Fig. 4. The result




Fig. 4 The polyhedron of perfect entanglers, W_{PE} , comprising all two-qubit gates $\hat{U} = \hat{k}_1 e^{i(c_1 \sigma_x \sigma_x + c_2 \sigma_y \sigma_y + c_3 \sigma_z \sigma_z)} \hat{k}_2$ with the property of transforming at least one separable state into a maximally entangled state, can be used as an optimization target [45]. Here, $\hat{k}_{1/2}$ are single-qubit operations, and the non-local content is completely characterized by the Weyl chamber coordinates c_1, c_2, c_3 , with the point O corresponding to the identity, and L to the CNOT gate. Optimal control targeting a perfect entangler minimizes the distance to W_{PE} [46].

of the optimization is then not only the control $u(t)$ that implements an entangling gate but also the gate itself.

These control targets — a specific two-qubit gate, a local equivalence class of gates or a two-qubit gate that is a perfect entangler — are expressed by functionals of increasing complexity. Their advantage is that we state the control problem in a more general way which may ease finding a solution. In principle, any physical property can be turned into an optimization functional. For example, one can minimize energy [48] or maximize entropy [49] or squeezing [50], instead of asking to prepare a specific state with minimal energy or maximal entropy or squeezing. This is advantageous as soon as more than one state with the optimal value of the desired property exists, since the optimization can explore all possible target states.

The desired property can be targeted at the final time only or at intermediate times as well. In other words, there is no fundamental difference between targets and constraints in optimal control theory, which is why in the mathematical literature, they are jointly referred to as “costs”. However, targeting a property at several or all intermediate times is a much harder control problem than one with a *single* final-time “cost” J_{tf} ¹. One can rationalize this as follows: In order to solve a control problem, we should first make sure that the problem is solvable. The answer to this question — yes or no — is provided by controllability analysis to which we will get below in Sec. 2.6. In controllability analysis, one assumes to have unlimited resources at the disposal, for example unlimited time (as long as it is finite) or unlimited pulse energy. The result of controllability analysis will be negative if a symmetry precludes

¹ Moreover, the numerical difficulty increases, as the equation of motion for the backward propagation contains a source term for time-dependent targets.

reaching the target; it thus reflects fundamental properties of the system. If the answer is positive, i.e., the quantum system is controllable, then very often many solutions to a given control problem exist — for unlimited resources. As soon as one restricts the resources, the number of solutions will shrink [51], possibly all the way to zero. In the case of the resource time, the transition point to no more solution is referred to as “quantum speed limit” or “minimal evolution time”, and we will see examples of it below in Sec. 3. Including constraints in the optimization functional is another way to limit the resources available for solving the control problem [51]. Conversely, stating the control problem in as general terms as possible will likely increase the search space and the resources. It is therefore important to design the optimization functional in a way that captures all of the desired physics while corresponding to the least constraints on the problem; it will help to make the problem solvable in practice.

2.6 Controllability

Is a given control problem solvable, assuming that all necessary resources can be provided? The mathematical tools for answering questions such as “which states can be prepared?” or “which operations (such as gates) can be implemented?” for a given quantum system are summarized under the term controllability analysis [52]. For closed quantum systems with unitary time evolution, the answer to these questions is determined by the Hamiltonian (1) alone. In other words, we do not need to solve the equations of motion but can infer the answer directly from $\hat{H} = \hat{H}_0 + \hat{H}_I$ ². More precisely, we need to inspect the nested commutators of $\hat{H}_0$ and $\hat{H}_I$, the drift and control Hamiltonians. Why this is sufficient, is readily understood for finite-dimensional systems by visualizing the system state in the (generalized) Bloch sphere [53]: Unitary evolutions are rotations of the state vector $|\psi\rangle$ on the sphere. These rotations are generated by the Hamiltonian. For infinitesimal evolutions $\hat{U}(t + \delta t, t) \approx \mathbb{1} - i(\hat{H}_0 + \hat{H}_I)\delta t$, with $\hat{H}_0$ and $\hat{H}_I$ being two vectors lying in the tangential plane attached to the Bloch sphere at the point corresponding to $|\psi(t)\rangle$. The terms in the Hamiltonian thus define the directions in which the state vector can move (they are indeed elements of a vector space). The ability to reach any point on the sphere then corresponds to the ability to generate motion in all possible directions. For finite evolution times, the directions are not just given by the drift and control Hamiltonians, but also by their (nested) commutators. This can easily be seen when expanding the exponent, $\hat{U}(t, 0) = \exp(i(\hat{H}_0 + \hat{H}_I)t)$ ³. The (nested) commutators are elements of the same vector space as $\hat{H}_0$ and $\hat{H}_I$. The question of controllability has thus been translated into the question of how many linearly independent vectors are generated by expanding the exponent.

² The time-dependence of $\hat{H}(t) = \hat{H}_0 + \hat{H}_I(t)$ is omitted here, since $\hat{H}_I(t) = \sum_j u_j(t)\hat{H}_j$ and what matters for controllability are the properties of the $\hat{H}_j$.

³ More precisely, we should account for time ordering, $\hat{U}(t, 0) = \mathcal{T} \exp(i(\hat{H}_0 + \hat{H}_I(t))t)$, but again, this does not affect the controllability of a system.

In more formal terms, inspecting the nested commutators of $\hat{H}_0$ and $\hat{H}_I$ corresponds to constructing the dynamical Lie algebra of the quantum system⁴. The Lie algebra is said to be of full rank if its dimension is equal to N^2 or $N^2 - 1$, where N is the dimension of Hilbert space [52]. The Lie algebra is then isomorphic to $\mathfrak{u}(N)$, respectively $\mathfrak{su}(N)$, the algebras associated to the unitary group $U(N)$ (containing all $N \times N$ unitary matrices), respectively the special unitary group $SU(N)$ (containing all $N \times N$ unitary matrices with determinant 1). The difference between the two cases concerns merely the physically irrelevant global phase associated with $\mathbb{1}$. If the dynamical Lie algebra is of full rank, any element of the group $U(N)$, respectively $SU(N)$, can be realized, i.e., the quantum system is (completely, or evolution operator-) controllable [52].

Since the dynamical Lie algebra is a vector space, one way to check the full rank-condition is to construct an orthogonal basis of that vector space from the commutators, thus determining its dimension. Alternatively, the Hamiltonian can be represented on a graph. The full rank condition can then be inferred from the connectedness of the graph [52]. While traditionally controllability analysis has been used for given, fixed Hamiltonians, we will see below in Sec. 3.4 that we can turn this perspective around and ask what is the minimal number of terms in $\hat{H}_I$ that yield a completely controllable system. This extends quantum control from the question of how to design the external fields $u(t)$ to that of how to overall engineer the field-matter coupling $\hat{H}_I(t) = \sum_j u_j(t) \hat{H}_j$.

The extension of controllability analysis to open quantum systems is extremely challenging and has largely been confined to systems with Markovian dynamics. The semi-group structure of such dynamics is amenable to Lie algebraic tools which has allowed e.g. for characterizing reachable sets of states [54].

3 Selected applications

Having laid out the basic concepts in Sec. 2, we can now inspect some recent applications of quantum control. The selection is heavily biased towards work from my group but Refs. [1, 2] provide a more comprehensive overview. Indeed, optimal control theory, and quantum control more generally, have proven useful, in the past decades, for applications in quantum information science across all physical platforms from nuclear magnetic resonance to atomic physics or solid state qubits [1, 2]. A key motivation for using optimal control theory is that quantum devices are open quantum systems. While advanced engineering aims to isolate a device's desired quantum features, the need to execute operational and readout protocols implies an inevitable coupling of the device to its environment. Optimal control theory is a tool allowing us to identify the best possible balance between desired control and undesired disturbance, i.e., decoherence.

⁴ In general, an algebra is a vector space with “multiplication”; in case of a Lie algebra, the “multiplication” is given by the commutator.

In addition to its promise to help make quantum technologies practical, the control of open quantum systems tackles questions that are also of fundamental interest and relevant beyond specific applications [26]. For example, what are viable control strategies in the presence of noise and how do they depend on the system-environment coupling? What are fundamental limits to quantum control in terms of attainable targets or errors? While general answers to these questions are still lacking, a set of preliminary useful rules has already been identified [26]:

1. If the desired operation shall keep pure states pure and the system dynamics are Markovian, the effect of the environment is detrimental. Then the best control strategy is to avoid decoherence. We will discuss examples below in Sec. 3.2.
2. If the desired operation changes the system purity, the control target can only be realized thanks to the presence of the environment. Then Markovian dynamics may be desired — the control target is reachable if it is a fixed point of the Liouvillian, and external controls can be used to ensure that this is the case. The corresponding control strategy is sometimes referred to as quantum reservoir engineering [55]. Section 3.3 will show how optimal control theory can be used in this setting.
3. If the realization of the control target requires dissipation which is not naturally available, we can create it by coupling the quantum system of interest to auxiliary degrees of freedom that are then simply discarded or measured. We will discuss this strategy in Sec. 3.5.

When the dynamics of the open quantum system is non-Markovian, the environment may have both beneficial and detrimental effects on control. Examples of beneficial effects are a larger number of implementable gate operations [56] or faster operation [57]. By and large, however, control of systems with non-Markovian dynamics is uncharted territory and beyond the scope of this tutorial.

Before we can examine opportunities for and limitations to control in open quantum systems, we need to take one step back, however, and answer the question how we can quantify success of the control. In other words, what are suitable figures of merit for open quantum systems? While the target functionals discussed in Sec. 2.5 above provide a good starting point, there are also some pitfalls to avoid.

3.1 Control of open quantum systems: How to quantify success

When the target of the control is to reach a given desired state, we can check whether the state of the system after time evolution matches this state by taking the overlap and turning it into a real number, cf. Eq. (14). Moving to open quantum systems, Hilbert space kets become density operators, and naively, we could replace the scalar product between the time-evolved and the target state in Eq. (14) by the Hilbert-Schmidt product for the respective density operators, i.e., $J_T \sim \text{tr} \{ \hat{\rho}(T) \hat{\rho}_{\text{target}} \}$. This works as long as the target $\hat{\rho}_{\text{target}}$ is a pure state but fails if $\hat{\rho}_{\text{target}}$ is mixed [27]




Fig. 5 Which state matches a desired target state better? Bloch sphere representation of matching states $\hat{\rho}_{1/2}$ with Bloch vectors $\mathbf{r}_{1/2}$ to a *mixed* target state with Bloch vector $\mathbf{r}_{trg}$. While a true distance measure quantifies $\mathbf{r}_1$ as closer to $\mathbf{r}_{trg}$ than $\mathbf{r}_2$, the Hilbert Schmidt product measures only the projection onto the axis of $\mathbf{r}_{trg}$ and thus is larger for $\mathbf{r}_2$. This failure of the Hilbert-Schmidt product for a mixed target state can be remedied by a figure of merit that matches both angle and length of the target Bloch vector. Adapted from Ref. [27].

because the Hilbert-Schmidt product $\text{tr} \{ \hat{\rho}_1 \hat{\rho}_2 \}$ is not (unlike the Hilbert-Schmidt distance, $\frac{1}{2} \text{tr} \{ (\hat{\rho}_1 - \hat{\rho}_2)^2 \}$) a distance measure.

Visualizing the states on the Bloch sphere provides an intuitive understanding of the failure [27], cf. Fig. 5. It makes use of the fact that any state $\hat{\rho}$, defined on a finite-dimensional Hilbert space (of dimension N), can be expressed in a complete orthonormal operator basis [53],

$$\hat{\rho} = \frac{\mathbb{1}}{N} + \sum_k r_k \hat{A}_k = \frac{\mathbb{1}}{N} + \mathbf{r} \cdot \hat{\mathbf{A}}.$$

In terms of Bloch vectors, the Hilbert-Schmidt product, or overlap, of two states is given by

$$\text{tr} \{ \hat{\rho}_1 \hat{\rho}_{target} \} = \frac{1}{N} + |\mathbf{r}_1| |\mathbf{r}_{target}| \cos \theta,$$

where θ is the angle between the Bloch vectors, see Fig. 5. Assuming $\mathbf{r}_1 \parallel \mathbf{r}_2$ and $|\mathbf{r}_2| > |\mathbf{r}_1| > |\mathbf{r}_{trg}|$ as shown in Fig. 5, the purer state $\mathbf{r}_2$ has a larger projection onto the target state and thus a larger Hilbert-Schmidt product, although $\hat{\rho}_2$ is “farther” away from the target than $\hat{\rho}_1$. Indeed, optimizing for $\hat{\rho}_{trg}$ with the Hilbert-Schmidt product will yield the pure state with $\theta = 0$ instead of ρ_{trg} . This can be remedied by employing a true distance measure such as the Hilbert-Schmidt distance or the trace distance. This comes, however, with a potential numerical instability due to the discontinuity of the square root. Alternatively, the intuitive picture of Fig. 5 suggests to optimize for the desired length and angle of the target Bloch vector [27]. This has been used in an application of optimal control theory to an optomechanic system where the mixed target state balances purity and squeezing [27].

When targeting a quantum gate, the naive approach would be to lift Eq. (16) from Hilbert space to Liouville space [58]. The trace implies that a complete orthogonal basis of the logical Liouville space, $\{ \hat{\rho}_j \}_{j=N^2}$, needs to be propagated and,

after time evolution, compared with the target operation applied to that basis state, $\hat{\rho}_j(T)\hat{O}\hat{\rho}_j\hat{O}^\dagger$. Here $N = 2^n$ is the dimension of the (logical) Hilbert space for n qubits (the physical Hilbert space of each qubit may comprise more than two levels). This approach involves significantly more effort than necessary because a complete basis of Liouville space needs to be used only if the optimization target is an arbitrary open system evolution. When instead the target is a quantum gate, we can exploit that the desired evolution is unitary, or as close to unitary as possible. In this case, it is sufficient to use three judiciously chosen initial states and compare their time evolution to the desired operation [59–61].

It might be very surprising that the number of states required to verify the desired operation is independent of the system size. This can be rationalized as follows. Assume for a moment that the time evolution *is* unitary. In this case the question is: How many states are necessary to infer from their time evolution whether a desired unitary is implemented? Since all unitaries for a given system form a group, it is equivalent to asking: How many states are needed to distinguish any two unitaries? The answer is two [59], and these two states have a clear geometric interpretation. When assessing whether two unitaries are equal, we need to determine whether they share an eigenbasis and whether their eigenvalues (or eigenphases) are equal. A single state built from one-dimensional orthogonal projectors $\hat{P}_i$, $\hat{\rho}_B = \sum_i \lambda_i \hat{P}_i$ with $\lambda_i \neq \lambda_j$ for all $i \neq j$, is sufficient to “fix” the basis. A second state, that is rotated with respect to all $\hat{P}_i$, $\hat{\rho}_P = \hat{P}_{rot}$ with $\hat{P}_{rot}\hat{P}_i \neq 0 \forall i$, answers the question about the eigenphases. Note that the rotated state is one element of a basis that is mutually unbiased with respect to the basis of the $\hat{P}_i$. In order to generalize this to open system evolution, a third state is necessary to quantify the non-unitarity of the evolution. This quantification is achieved by measuring the unitality on the logical subspace, $\hat{\rho}_3 = \mathbb{1}/N$ [59]. The construction of these three states is useful not just for the figure of merit in gate optimization but also for determining the fidelity of gates or channels of actual quantum devices [60]. In gate optimizations, the most efficient approach is typically to replace $\hat{\rho}_B$, which is mixed, by the N projectors from which $\hat{\rho}_B$ is constructed [61]. In this case, $\hat{\rho}_3$ is not needed since non-unitarity can be inferred from the (non-unitary) time evolution of the $N + 1$ initially pure states. The better numerical performance is rationalized by a more uniform distribution of the relevant information over the propagated states [61]. While the number of states in this variant scales with system size, there is still an exponential saving compared to the naive approach [58]. The same exponential saving carries over to protocols for determining gate or channel fidelities [60].

Besides targeting specific states or gates with quantum optimal control, it may also be useful to optimize towards more general properties. Examples were local equivalence classes of two-qubit gates or the set of all perfect entanglers in Sec. 2.5 above. In order to adapt the corresponding optimization functionals to open system evolution, we have to reconstruct the unitary part of that evolution from knowledge of the propagated states. Similarly to gate optimization, this can be done using $\hat{\rho}_B$ and $\hat{\rho}_P$ or $N + 1$ one-dimensional projectors as initial states, provided the evolution is not too noisy (in which case the question about the unitary part loses its meaning) [62].




Fig. 6 Quantum speed limit for a universal set of gates for transmon qubits (left): The gates are implemented via a shaped microwave drive, designed with quantum optimal control theory, that acts on the cavity to which both superconducting qubits are coupled. Using an optimization functional that targets an arbitrary perfect entangler, cf. Fig. 4, has allowed us to identify $\sqrt{i\text{SWAP}}$ as the “natural” gate for this architecture (right): For long gate times (200 ns in the upper panel), almost all perfect entanglers can be realized, whereas only the $\sqrt{i\text{SWAP}}$ -gate survives when reducing the gate time (10 ns in the lower panel). Adapted from Ref. [63].

The reconstruction algorithm, just as the fidelity estimation discussed above, is useful not just in quantum optimal control theory but also to benchmark actual devices.

Having adapted our figures of merit to open quantum systems, the optimization proceeds as in the closed system case, except for the fact that, instead of the time-dependent Schrödinger equation, the corresponding equation of motion for the open system must be solved (for both forward and backward propagations) [61]. With the necessary tools for studying the control of open quantum systems established, we can now look at examples (taken from the work of my group) which illustrate the three rules identified at the beginning of Sec. 3.

3.2 Control of open quantum systems: Avoiding decoherence

Avoiding decoherence is the best control strategy when the system dynamics are Markovian and the control target is reached without any change of entropy. This is true for coherent state transformations including the implementation of quantum gates. In general, there are two ways to avoid decoherence — one is to “beat” decoherence by operating fast, the other one is to utilize protection from symmetries.

We have already discussed examples of fast operations above, with the examples of Rydberg state preparations [29, 30] in Sec. 2.4. Another example is the implementation of a universal set of gates, defined as the collection of single-qubit gates and one entangling gate from which any desired unitary can be constructed [47]. Optimal control theory is ideally suited to determine the quantum speed limit [48] which can be quite different for different gates. It is therefore crucial to establish the quantum speed limit for the complete universal set, rather than relying on only one or two specific gates. This is particularly true for superconducting qubits where it is *a priori* not

clear, due to residual couplings between the qubits, whether the single-qubit gates or the two-qubit gate determine the speed limit. Moreover, a variety of entangling gates, each with a possibly different speed limit, can be implemented. In order to determine the fastest universal set for superconducting transmon qubits [63], we have used a combination of the tools introduced above, that is, optimization targeting an arbitrary perfect entangler [45, 46] and hybrid optimization [37] combining a global search of the control landscape [36] with Krotov’s method [38–40]. To minimize the impact of external control noise, all gates are implemented via microwave drives that act on the common cavity mode coupling both transmons in the transmission line. For this architecture, manually designed gate protocols typically rely on avoiding cavity excitation which allows the cavity to be eliminated to derive an effective two-qubit Hamiltonian. However, operating in this dispersive limit entails a trade-off — the weak coupling results in slower gate operation. A numerical approach to gate design using quantum optimal control has allowed us to forgo the dispersive limit [63]. Exploring the full landscape of design parameters, specifically couplings, detunings and anharmonicities, we have determined the “quasi-dispersive straddling qutrits” regime to be optimal for the rapid creation and removal of entanglement necessary for a universal set of gates [63]. Optimizing for an arbitrary perfect entangler instead of a specific two-qubit gate has furthermore allowed us to determine the “natural” gate of the architecture — only the $\sqrt{i}$ SWAP can be realized within 10 ns, whereas for gate durations as long as 200 ns almost any perfect entangler can be implemented [63], cf. Fig. 6 (right). We have also found that single-qubit gates require more than 10 ns, putting the quantum speed limit for the universal set at 50 ns, cf. Fig. 6 (left). For such a gate duration, several entangling gates are possible, and we have opted for the B-gate as the perfect entangler that allows for the shortest circuits when decomposing an arbitrary unitary into the gates of universal set [64]. Finally, while operating at the quantum speed limit is optimal for minimizing decoherence, it comes at the expense of spectrally broad and temporally complex pulse shapes which may be hard to benchmark on an actual device. Extending the gate durations may also be beneficial in view of better robustness of the protocols with respect to parameter uncertainties. These points highlight the complexity of optimizing the operation of quantum devices where a number of desiderata need to be balanced.

Besides operating fast, we can avoid decoherence by utilizing protection from symmetries, that is, by keeping the dynamics in decoherence-free subspaces [66]. With the time-dependent “dark state” in STIRAP, we have already encountered a one-dimensional example of a decoherence-free subspace [66] in Sec. 2.3 above. But even though there are several equivalent ways to identify decoherence-free subspaces, being able to do this in practice may be hampered by system complexity. Here is where quantum optimal control can help. Optimizing directly for a decoherence-free subspace results in optimization landscapes with a very large number of local minima. In contrast, targeting the desired operation while accounting for decoherence via the equation of motion is typically sufficient to find decoherence-free subspaces, even approximate ones. For example, we have found that pulses designed to maximize the distinguishability of two quantum states that evolve under slightly different Hamiltonians, a prototypical question in quantum sensing, steer the states close to a




Fig. 7 Keeping the dynamics decoherence-free: Time evolution of two states (ρ_1, ρ_2) that evolve under slightly different Hamiltonians, driven by the same σ_x -drive that was optimized to maximize the final-time distinguishability. The dynamics stays close to a decoherence-free subspace, the state $|0\rangle$ for T_1 decay (left) and the z -axis for T_2 dephasing (right), for a significant part of the protocol. Adapted from Ref. [65].

decoherence-free subspace [65, 67]. The optimized pulses yield a distinguishability that is significantly larger than that obtained with a standard Ramsey protocol. The improvement translates also into a metrological gain in terms of the quantum Fisher information. Remarkably, the distinguishability of the two states can be stabilized to protocol durations that are orders of magnitude larger than the timescale of decoherence, for both T_1 decay and T_2 dephasing [65]. While surprising at first glance, we could rationalize this finding in terms of staying close to a decoherence-free subspace (the ground state for T_1 decay and the z -axis for T_2 dephasing) [65].

A related set of control protocols that seek to avoid decoherence is summarized under the name of “dynamical decoupling” [68]. These protocols use static or time-dependent fields that are designed to average the effect of the interactions with the environment to zero. One way to think of it is as a dynamic way to impose decoherence-free subspaces. Dynamical decoupling can be combined with quantum optimal control to address potentially limiting assumptions on e.g. pulse shapes or pulse timings.

3.3 Control of open quantum systems: Exploiting the environment for control

While beating decoherence may be the most obvious control strategy for open quantum systems, it reaches its limits when the control task requires the presence of an environment. A prominent example is cooling or, more generally, all processes in which the system entropy shall be decreased, see Fig. 8 for illustration. The environment is then needed as an entropy sink. Typical questions for control are how to enhance the cooling rate or how to improve the fidelity of pure state preparation.




Fig. 8 Qubit reset is a paradigmatic example of control problems which can only be solved in the presence of an environment since they require export of entropy. For qubit reset, the protocol should yield the desired target state (light-red arrow) irrespective of the current qubit state (dark red) which may require state rotations (left, middle) and purity increase (right).

Quantum optimal control can be used to answer these questions in non-trivial settings, for example, when the environment is structured, consisting of both strongly and weakly coupled degrees of freedom [49]. The strongly coupled modes allow for fast reset, whereas the remainder of the environment can be used for thermalization on a slower timescale. The simplest model to formalize this picture is a qubit strongly coupled to a single auxiliary two-level system (TLS) that undergoes slow decay. The time evolution of the joint system of qubit and auxiliary TLS is then described by a GKLS master equation (2), and compatibility with the second law of thermodynamics is ensured by imposing a clear separation between the couplings and the energy scale of the system qubit [49]. We have first used numerical optimization and found pulses that drive the qubit into resonance with the auxiliary TLS in such a way that qubit and TLS have their states exchanged at the final time. The reset error for the qubit is determined by the temperature of the TLS which is initially assumed to be in a thermal state, and the minimal time for the reset is given by the qubit-TLS coupling strength [49]. This reset strategy has been implemented experimentally in a slightly modified way for a superconducting qubit, replacing the TLS by a cavity mode [69].

When adding correlations between qubit and TLS to the initial state, we have found, to our surprise, the final error to be smaller and the resonant population swap to be faster than without initial correlations [49]. To understand this result, we have derived an approximate analytical model, which captures only the dynamics of qubit and TLS and neglects the TLS's slow decay. Invoking the rotating-wave approximation, which is relevant for details of the dynamics but not the overall control strategy, we have represented the joint state of qubit and TLS in terms of 16 real variables. Their equations of motion represent a vector flow and can be decoupled by a suitable variable transformation. As a result, the dynamics of the three variables determining the qubit purity, i.e., qubit ground state population and real and imaginary parts of the qubit coherence, can be visualized geometrically [49]. The visualization reveals that the control rotates initial correlations into qubit ground state population, thereby increasing the final purity and reducing the error. The geometric picture also explains the faster reset: The minimal duration for the simple population swap is $\pi/(2J)$, where π is the angle of rotation that the qubit ground state population needs to undergo and J is the strength of the coupling between qubit and auxiliary.

Initial correlations reduce this angle. Unfortunately, in order to take advantage of faster and more accurate qubit reset, one would need to know the initial correlations between the system and environmental degrees of freedom which is typically not the case. In contrast, the simple reset via population swap between qubit and auxiliary works irrespective of the initial state [49].

In our initial model for qubit reset, we have assumed an σ_z -drive on the qubit and a $\sigma_x \otimes \sigma_x$ -coupling between qubit and auxiliary [49]. One may wonder whether a σ_x - or σ_y -drive on the qubit would yield even faster reset times. Moreover, while a $\sigma_x \otimes \sigma_x$ -coupling seems quite natural for population exchange, other couplings could also be engineered, at least for superconducting qubits. We have thus carried out numerical optimizations for a number of couplings [70]. Their results supported the conjecture that, given any qubit-auxiliary coupling excluding those with σ_z for the auxiliary, a specific drive $\epsilon(t)\sigma_i \otimes \mathbb{1}$ exists that enables qubit reset to the maximally attainable purity. Moreover, the minimum reset duration can only be achieved for qubit and auxiliary on resonance and $\epsilon(t)$ constant [70].

The conjecture can be proven combining the tools from controllability analysis introduced in Sec. 2.6 with the so-called Cartan decomposition; this decomposition is fundamental to classifying all two-qubit operations in the Weyl chamber, cf. Fig. 4. The proof formalizes the ability of purifying a qubit via the condition that the qubit dynamics must be non-unital. This, in turn, requires the dimension of the non-local Cartan subalgebra to be equal or larger than two [70]. Provided coupling and drive are chosen to fulfill this condition, there exists an optimal choice for the qubit control σ_i , such that the reset is time-optimal with $T_{min} = \pi/(2\eta)$, where η depends on the choices for the coupling $\sigma_n \otimes \sigma_m$ and the drive σ_i . Several such choices exist that result in $\eta = J$, the overall shortest possible reset time [70]. Interestingly, our result saturates the bound obtained from the so-called Time-Optimal Tori theorem [71].

It is natural to ask whether our results on the minimum duration and maximum purity attainable in qubit reset are modified when the qubit is coupled to an auxiliary degree of freedom with a Hilbert space dimension larger than two. While the minimum reset time stays the same, the maximally attainable purity increases with increasing Hilbert space dimension [70]. Indeed, we were able to show rigorously that the purity error, $1 - \text{tr}\{\hat{\rho}_S^2\}$, can be made smaller than some prespecified error tolerance ε if at least $\lceil d_B(d_S - 1)/d_S \rceil$ eigenvalues of the initial auxiliary state $\hat{\rho}_B$ are below $\varepsilon/(2d_B(d_S - 1))$ where $\lceil \cdot \rceil$ denotes the ceiling function, and d_S and d_B are the Hilbert space dimensions of system (“qubit”) and bath (auxiliary). The better purification can be rationalized in terms of a unitary operation that acts on both system and auxiliary, reshuffling the spectrum of the joint separable state [70]. Unfortunately, this improved bound on the reset fidelity does not easily translate into a better reset protocol, as the reshuffling unitary is likely dependent on the initial state of both system and auxiliary. This leads to the open question whether an approximate unitary can be constructed that surpasses the purity obtained by swapping the population of qubit and auxiliary two-level system, irrespective of the system state.

Qubit reset is the simplest example of quantum reservoir engineering [55] where application of external fields together with the natural interaction of the quantum system with its environment (the “reservoir” or “bath”) steers the system toward its




Fig. 9 Control of quantum reservoir engineering: Making the steady state match the desired target state with quantum optimal control [72, 73]. The original protocol [74] uses the drives shown on the left, plus laser cooling on two magnesium ions. The population flow induced by these drives is shown in the middle panel with the hexagons formed by the $|\uparrow\rangle$, $|\downarrow\rangle$, and $|a\rangle$ states of the two beryllium ions, and the size of the hexagons encoding the vibrational excitation in the ion trap. Orange (black) wiggly lines correspond to transitions that are uni-directional due to excitation by the repump and subsequent spontaneous emission from the electronically excited state (due to laser cooling), whereas the solid black and blue lines represent transitions that are coherently driven by the microwave and sideband laser. The microwave drives population out of the target singlet state (marked in green). A different microwave transition, identified by optimal control, together with a blue and a red sideband and the repump laser results in uni-direction population flow towards the target state (right panel). Adapted from Ref. [72].

ground state. It illustrates the general principle of quantum reservoir engineering to utilize dissipation to drive a system toward a desired steady state. This state can be the ground state but also a non-trivial state such as an entangled state [55]. However, for non-trivial target states, the required combination of external fields and dissipative dynamics is often rather complex, making it challenging to solve the problem manually. This is another avenue for quantum optimal control.

To illustrate the problem and show how ideas from optimal control can be leveraged to solve it, let us consider the example of preparing two beryllium ions in a maximally entangled, or Bell state [74]. The two beryllium ions are co-trapped with two magnesium ions which are laser cooled, ensuring the cooling of all joint vibrations of the four-ion chain. Two hyperfine levels in the electronic ground state of the beryllium ion form the qubit, $|\uparrow\rangle$, $|\downarrow\rangle$, and an additional hyperfine level $|a\rangle$ is also involved in the dynamics, cf. Fig. 9. In order to prepare the two qubits in the spin singlet state, a combination of four external drives and spontaneous emission from a very short-lived electronically excited state have been used [74]. When all four external fields are kept on continuously, the target state is prepared with a fidelity of about 75% [74]. Using the NLOpt package [36], we have optimized the parameters of these four fields, i.e., polarizations, Rabi frequencies and detunings, improving the state fidelity to about 90% [72]. The fidelity is limited because one of the external fields, the microwave drive, couples the target state to another state, as highlighted in the middle panel of Fig. 9. In other words, the steady state of the dynamics is not identical to the desired target state. The original protocol balanced the microwave amplitude with respect to the other drives to minimize this error [74]. But even with




Fig. 10 Performance bounds from quantum optimal control: Ultimately attainable error predicted for dissipative state preparation as a function of the maximal amplitude of the sideband lasers $E_{r/b}$ (bottom x-axis), resp. laser power (top x-axis), for different heating rates. Adapted from Ref. [72].

all drive parameters optimized, the resulting error remains non-negligible [72]. We have therefore used optimization to select the best combination of external drives, before optimizing their parameters [72]. To keep the two-stage optimization manageable, we have classified all transitions that can be driven by an external field, reducing the number of possible combinations. As a result of the optimization, a different microwave drive that does not affect the target state is combined with three lasers, cf. right panel of Fig. 9, and the polarizations, Rabi frequencies and detunings of all fields are optimized. Then the only remaining source of error is anomalous heating in the trap which determines the final state preparation error, predicted to be about 90%, 97%, resp. 98% for heating rates of 100 Hz, 10 Hz and 1 Hz [72].

Optimizing the choice of drives also revealed that laser cooling of the magnesium ions is not essential; instead, a red-detuned sideband transition reduces vibrational excitations in the trap, cf. middle and right panels of Fig. 9. This implies the possibility to work with just the two beryllium ions instead of a four-ion chain — a significant simplification of the experimental setup! The corresponding reduction in the resources required to prepare a Bell state has meanwhile been demonstrated experimentally by the NIST trapped ion group [73], with the best state preparation fidelities at 95%. The remaining error sources in the modified trapped ion experiment are well understood [73], suggesting that our theoretically predicted fidelities are indeed feasible.

Both the gain in performance and the reduction of the resources nicely illustrate the utility of optimal control. We can go one step further and ask: What is the ultimately attainable error? In the estimates quoted above, we had capped the laser amplitudes at values that were typical for the experiment of Ref. [74]. Assuming that

advances in technology eventually will allow for more powerful lasers, even in the UV regime, we can lift this limitation. Large pulse amplitudes allow for operating with larger detunings which is advantageous to minimize undesired photon scattering. Allowing for larger amplitudes, can the state preparation error be pushed below 10^{-4} ? This is the threshold below which error correction is expected to become useful. The question is answered in Fig. 10 which shows the ultimately attainable error as a function of available laser power, accounting also for anomalous heating: For heating rates of 1 Hz and larger, it is not possible to reach the 10^{-4} limit, irrespective of the available power, as seen in the grey and yellow curves levelling off. If anomalous heating can effectively be suppressed on the timescale of the state preparation protocol, the error can be pushed below 10^{-4} , cf. red and blue curves in Fig. 10. The red curve assumes four beams are required to realize the two sideband transitions, whereas the blue curve utilizes a more efficient configuration, achieving the two sideband transitions by combining three beams [72]. Figure 9 illustrates how quantum optimal control determines an ultimate performance bound, specifically the minimally achievable error. Crucially, it simultaneously identifies the required experimental conditions, here constraints on laser power and heating rates.

3.4 A more systematic approach to finding suitable interactions: Controllability analysis

In the previous example, numerical optimization was employed to determine the drives which, in conjunction with spontaneous emission from a short-lived excited state, drive the two-ion system into the desired target state. In formal terms, this corresponds to a search over the space of operators. Such a search incurs a significant numerical cost; even for a system as small as two three-state ions and one vibrational mode, it necessitated a pre-classification of the transition operators. This approach is clearly limited to a few degrees of freedom and cannot be scaled up, raising the question whether alternatives to brute force optimization over operator space exist.

Ensuring that the steady state of a driven-dissipative evolution matches a desired target state is reminiscent of the controllability analysis introduced in Sec. 2.6. Controllability determines whether a certain target state or target operation can be realized, at least in principle (that is, without bounds on resources such as time, pulse power, or pulse bandwidth). Traditionally, this analysis has been applied to a fixed Hamiltonian with a given choice of drives, but it is natural to extend the inquiry to determine which drives are needed to make a system controllable [75]. Unlike numerical optimization, which performs a global search over the operator space, controllability analysis is based on the dynamical Lie algebra, exploiting information from the commutators, and may thus provide a more systematic approach to finding the necessary drives.

Unfortunately, like brute-force numerical optimization, controllability analysis is hampered by the curse of dimensionality. This limitation is immediately apparent as determining evolution-operator controllability requires checking the rank of the




Fig. 11 Controllability analysis via graph representation. Left: Mapping a system to a graph for the example of two coupled qubits subject to one external control. The control connects all eigenstates of the total drift (including the qubit-qubit coupling), represented as the nodes of the graph. However, controllability requires, on top of connectedness, all transitions to be decoupled, such that they can be addressed separately. Right: Once the controllability problem is represented as a graph, commutators can be calculated graphically, avoiding explicit construction of the Lie algebra. Here, the red transition is obtained as commutator of blue and green transitions.

dynamical Lie algebra, cf. Sec. 2.6. The most straightforward way to calculate the rank is by explicitly constructing the dynamical Lie algebra from the commutators. However, this is hampered by both the general double-exponential scaling of the operator space and the numerical difficulty to orthogonalize large matrices. Two recent advances address these problems. The first harnesses graph theoretical methods to avoid constructing an orthogonal basis of the dynamical Lie algebra; ultimately, it remains hampered by the double-exponential scaling of the operator space [76]. For general quantum-mechanical time evolutions, the only path to resolving the scalability issue is to perform at least part of the calculations on a quantum device. This strategy has been adopted for controllability analysis by leveraging a hybrid quantum-classical algorithm [77].

Representing the controllability problem as a graph makes its solution very intuitive. Separating the Hamiltonian into drift $\hat{H}_0$ and interactions with external controls, $\hat{H}(t) = \hat{H}_0 + \sum_j u_j(t) \hat{H}_j(t)$, the eigenstates of the drift are taken to be the nodes of the graph, whereas the controls become edges connecting the nodes. As an example, the left part of Fig. 11 shows the graph for two coupled qubits, one of which is driven by an external control. A system is controllable if a connected subgraph exists that contains all nodes of the graph and only decoupled⁵ transitions [76] (and refs. therein). The two examples shown in Fig. 11 are *not* controllable: In the example on the left, the subgraph contains all nodes but the transitions are coupled (indicated by the color). In contrast, the right-most subgraph contains only decoupled transitions but not all nodes of that example are contained in the subgraph. The examples in Fig. 11 have mainly an illustrative purpose. For the important case of coupled qubit arrays, Ref. [76] provides algorithm flowcharts for setting up the graph, constructing all commutators graphically and checking the subgraph controllability condition. The utility of the graphical controllability test was shown for qubit arrays consisting of five qubits connected in T-shape as in the ibmqquito architecture [76]. In

⁵ Two transitions are coupled if they are connected to the same control $u_j(t) \hat{H}_j$ and have the same transition frequency.




Fig. 12 Controllability analysis via a hybrid quantum-classical algorithm: Parametrized quantum circuit to determine pure state controllability for three qubits. Adapted from Ref. [77].

the original design, each qubit is driven by a $\hat{\sigma}_x$ -control, and neighboring qubits are coupled via $\hat{\sigma}_x\hat{\sigma}_x + \hat{\sigma}_y\hat{\sigma}_y$. Controllability analysis revealed that the number of external controls can be reduced from five to two, keeping the qubit-qubit coupling fixed. Using instead qubit-qubit couplings of the type $\hat{\sigma}_x\hat{\sigma}_x + \hat{\sigma}_y\hat{\sigma}_y + \hat{\sigma}_z\hat{\sigma}_z$, a single external control is sufficient for evolution-operator controllability, i.e. for implementing arbitrary unitaries on the five qubits [76]. Since external controls come with considerable overhead in quantum chip design, reducing their number may be beneficial in view of scaling up the size of those chips. For larger and larger devices, however, running controllability tests on classical computers becomes impractical due to the exponential scaling of the Hilbert space.

To enable controllability tests for large qubit arrays, a hybrid quantum-classical algorithm based on a parametric quantum circuit was devised [77]. The algorithm hinges on the proof that controllability is linked to the number of independent parameters θ_i of the circuit [77]. These parameters are, in turn, obtained by dimensional expressivity analysis, which was originally developed for obtaining a maximally expressive circuit ansatz with a minimum number of parameters [78]. Crucially, dimensional expressivity analysis can efficiently be implemented on quantum hardware [78].

The hybrid quantum-classical algorithm to determine controllability proceeds as follows. The controllability problem is presented as a parametric quantum circuit by associating each control plus the drift with rotation angles θ_j . One layer of the circuit is shown in Fig. 12, where $R_j(\alpha) = \exp(i\alpha\hat{H}_j)$ with $j = 0, \dots, m$ for m controls. The complete circuit consists of n_l layers with different values of the parameters θ_j in each layer. Dimensional expressivity analysis for the circuit [78] is carried out. If the circuit reaches maximal expressivity, the system is controllable (in the sense that any pure state of the qubit array can be reached), otherwise another layer is added and the dimensional expressivity test is repeated [77]. In case the controllability test fails even with a large number of layers, different initial states and random choices of the parameters θ_i need to be tested to confirm the negative controllability result.

Typically, evolution operator controllability is the focus (rather than pure state controllability), as it is a requirement for universal quantum computation. Similarly to extending quantum state tomography to quantum process tomography, this extension of the controllability test is achieved via the Choi-Jamio lkowski isomorphism [77].

It implies using twice as many qubits which must be initially prepared in an entangled state. Aside from this overhead, the algorithm proceeds as before.

Implementing the algorithm on an actual device will enable controllability tests to be used for resource-efficient design of quantum chips. The controllability analysis provides the minimal number of local controls and qubit couplings required for controllability and, consequently, universal quantum computation [77]. It allows for obtaining this information before a device is actually built, provided the associated quantum circuit can be implemented on a different device. However, reducing the number of controls comes at the expense of longer gate durations. Therefore, balancing the number of external controls and the gate durations will be essential. To achieve this balance, it will be important to understand how the removal of redundant controls affects the minimum gate duration.

3.5 Control of open quantum systems: Creating the desired dissipation

Let us now return to the control of open quantum systems. Below, we will examine examples demonstrating how the tools of controllability analysis identify the necessary operators, as conjectured in Sec. 3.3 in the context of quantum reservoir engineering. Before proceeding with these examples, we first recap the known control strategies for open quantum systems. Rephrasing the preliminary set of rules discussed on p. 21, we can classify them as follows:

1. Detrimental: Avoiding decoherence

As discussed above in Sec. 3.2, the impact of the environment is detrimental for Markovian dynamics⁶ and for targets that do not require any entropy export. The available strategies are then to operate fast compared to the decoherence timescale or to engineer protection from decoherence via decoherence-free subspaces or dynamical decoupling. Quantum optimal control theory is useful to implement these strategies, since it allows for designing control protocols with durations at the quantum speed limit or leveraging decoherence-free subspaces, as we have seen in the examples in Sec. 2.4 and Sec. 3.2.

2. Beneficial-1: Exploiting dissipation

For control targets that require entropy export, the presence of the environment is crucial and, consequently, beneficial. The key paradigm is quantum reservoir engineering [55], which involves choosing external drives such that the steady state of the driven-dissipative evolution matches the desired target state. Section 3.3 presented the Bell state preparation for trapped ions as an example, demonstrating how optimal control theory can simplify the experimental design, improve state preparation errors, and identify ultimate performance bounds.

⁶ For systems with non-Markovian dynamics, the question of available control strategies is still open. It is more complex than for systems with Markovian dynamics because of additional ways to exploit the environment [56].




Fig. 13 The principle of measurement-induced dynamics: Coupling a system to one or more auxiliary degrees of freedom (M for “meters”) which are discarded after their interaction with the system S , for example by projective measurements. Adapted with permission from Ref. [80].

3. Beneficial-2: Creating desired dephasing and dissipation

When the required dissipation channels are not naturally available, they can be effectively created by coupling the system to auxiliary degrees of freedom. Following their interaction with the system, these auxiliaries are either simply discarded or subjected to projective measurements, see Fig. 13. This dissipation engineering can be viewed as quantum reservoir engineering 2.0. It offers the advantage of being readily applicable to many-body systems and has, consequently, become a crucial component of the quantum control toolbox, with applications in quantum simulation and quantum error correction [79].

The remainder of this section will present three examples of engineered decoherence and dissipation, illustrating the basic concepts and showcasing where tools from quantum control can be leveraged.

The first example exploits dephasing to endorse adiabaticity [81, 82]. The basic intuition can be rationalized with the Landau-Zener model for an avoided crossing, see also Fig. 2 (bottom left),

$$\hat{H}_S(t) = \frac{\epsilon t}{2} \hat{\sigma}_z + \frac{g}{2} \hat{\sigma}_x.$$

The goal is to avoid diabatic transitions. One strategy is to adiabatically traverse the crossing, but often such slow dynamics cannot be realized. Diabatic transitions would also be avoided if the coherence between the two levels is suppressed. This corresponds to pure dephasing of the Landau-Zener qubit. One possibility to engineer this desired dephasing is by coupling the qubit to an auxiliary degree of freedom, e.g. a cavity coupled to a thermal bath, which we term the “meter” to emphasize the analogy with quantum measurements [81]. Pure dephasing is achieved via a system-meter coupling of the same form as quantum non-demolition (QND) detection. This requires the system-meter coupling to commute with the system Hamiltonian. Since $\hat{H}_S(t)$ is time-dependent, the system-meter coupling must also be time-dependent. Assuming this coupling can be engineered and made sufficiently strong, diabatic transitions in the Landau-Zener qubit can be strongly suppressed [81]. A more realistic scenario, however, involves a stroboscopic, rather than fully time-dependent implementation of the system-meter coupling; in this case the effectiveness of suppressing diabatic transitions is governed by the coupling rate [81].

Applying engineered dephasing to many qubits, such as those found in quantum annealers, requires the system-meter coupling to be both time-dependent and global [82]. Such coupling provides speedups in the time-to-solution of the anneal-

ing that are linearly proportional to the coupling strength. Surprisingly, replacing the cavity “meter” with a two-level system — which allows for calculating the exact reduced system dynamics — reveals that a fully coherent mechanism (effective energy rescaling) is the more effective approach for enhancing the system’s adiabaticity, rather than pure dephasing [82]. Implementing the required system-meter couplings in quantum annealing is daunting and beyond what is readily available in existing devices. This difficulty can be mitigated because the actual schedule is often irrelevant, provided it yields the final Hamiltonian that encodes the computational problem. This fact can be exploited to simplify the time-dependence of the system-meter coupling, albeit at the cost of smaller speedups [82]. In summary, while conceptually very appealing, the implementation of engineered dephasing protocols is rather difficult, at least in time-dependent scenarios such as quantum annealing. Therefore, applications that require only static system-meter interactions may represent a more promising way forward.

Static system-meter interactions can also be used to engineer dissipation, or energy relaxation, rather than pure dephasing. The driven-dissipative dynamics then effectively “cools” the system into the target state. An important application of this approach is to overcome the limitations of quantum reservoir engineering protocols that are based on natural decay processes. Overcoming these limitations would be crucial for realizing theoretical proposals targeting the preparation of non-trivial many-body quantum states as steady states of Markovian master equations [83]. To understand the limitations of quantum reservoir engineering for many-body systems, consider the example of trapped ions, discussed Sec. 3.3. There, the dissipative process is spontaneous emission from an electronically excited state, which presents several drawbacks. (i) Spontaneous emission is essentially local and cannot, by itself, not introduce entanglement. (ii) It leads to leakage to undesired states. Even for just two ions, this necessitated a complex combination of several fields to repump all population into the cooling cycle. (iii) Finally, a natural decay process such as spontaneous emission, which can be coupled to in a controlled way, may simply not exist for a given target state or space.

Instead of relying on natural decay processes, measurement-induced dynamics [79] exploits the fact that quantum measurements induce non-unitary evolution in the measured system [84]. Continuous measurements, in particular, can be described by a GKLS master equation, cf. Eq. (2), where the system part of the system-meter interaction Hamiltonian defines the jump operators [84]. In other words, quantum measurements provide a constructive recipe to engineer jump operators at will [79]! A complementary perspective on measurement-induced dynamics is offered by treating the meter degrees of freedom as “colliders” in so-called collision models [85], which highlights the effective “simulation” of the desired open system dynamics [86].

To illustrate the approach with a specific example, consider the so-called AKLT model⁷ describing a chain of spin-1 particles with only nearest-neighbor interactions. The model is frustration-free; with periodic boundary conditions, it has a non-degenerate ground state that is a paradigmatic example of both a matrix product

⁷ The model is named after Affleck, Lieb, Kennedy and Tasaki who introduced it [87, 88].




Fig. 14 Dilute cooling of the AKLT chain: Weak measurements of a single link of N spin-1 particles, coupled by nearest-neighbor interactions, allow one to drive the system into its ground state. The couplings to the five meter qubits give rise to jump operators $L_{1,\dots,5}^{(1,2)}$ which ensure that all population in the local hot subspace is cooled, whereas the interactions within the chain move all excitations through the chain until they reach the cooled link. Adapted from Ref. [91].

state and a symmetry-protected topological phase [87, 88]. This ground state can be prepared dissipatively, without any action of the Hamiltonian [89]. To this end, each link in the chain of spin-1 particles is repeatedly connected to five meter qubits. Following their interaction with the chain, they are discarded and replaced by freshly initialized qubits. The five jump operators resulting from the system-meter interactions deplete any population in excited states [89]. A second possibility to prepare the AKLT ground state is by alternating unitary evolution and fusion measurements [90]. In both cases, the number of meters scales with the size of the AKLT chain. Remarkably, control theory allows for modifying the protocol of Ref. [89], reducing the number of links that are cooled to a single one [91]. This requires the action of the AKLT Hamiltonian, in addition to coupling one link to five meter qubits. It comes at the expense of slower convergence to the target state with increasing size of the chain. Meter qubits are a precious resource, making the modification very appealing in terms of resource efficiency — the repeated preparation of a large number of meter qubits implies a significant operational cost, whereas engineering the two-body couplings of the AKLT Hamiltonian is of a similar difficulty as engineering the coupling to the meters.

The possibility of such "dilute" cooling can be rationalized by combining two concepts, analysis of the state space structure and the notion of controllability (cf. Section 2.6) [91]. To this end, consider the reduced state on the cooled link (sites 1,2 in Fig. 14); it is obtained by integrating out all other sites (3, ..., N in Fig. 14). Denoting the target state defined on all of the chain by $\rho_{\oplus}$, we refer to the support of $\rho_{\oplus}$ on the cooled link as the cold subspace, and its complement in the local Hilbert space (of sites 1,2) as the hot subspace, cf. Fig. 14. The interaction with the meter qubits is designed in such a way that all population in the hot subspace is transferred into the cold subspace, without affecting any population in the cold subspace. Dilute cooling works because all excited states of the AKLT chain have a non-zero overlap with the local hot subspace on the cooled link: Cooling depletes the population of the local hot subspace, but the interactions within the chain refill it with excitations from other parts of the chain, until all excitations have been cooled. Therefore, the minimal overlap of the excited states with the hot subspace defines the cooling rate [91].

The non-vanishing overlap of all excited states with the hot subspace is a beneficial characteristic of the AKLT model but this intuition can be formalized into necessary and sufficient conditions [91]. This allows dilute cooling to be generalized beyond the AKLT model. A first necessary condition is that the target state, when reduced to the cooled link, must be distinguishable from the reduced excited states; in essence, the local hot subspace cannot be empty, $V_{\text{hot}}^{(i,i+1)} \neq \emptyset$. A second necessary condition concerns the coherent interactions within the system. To formulate this condition, we construct a vector space containing all quasi-local operators that respect the target state as elements. This space has been dubbed the “kernelizer” since the target state is an eigenstate of these operators with eigenvalue zero [91]. The condition then dictates that the target state must necessarily be the only state invariant to both the cooling mechanism (i.e., the jump operators) and all coherent interactions originating from the kernelizer. This second necessary condition ensures that excitations on links that are not directly cooled propagate through the system until they reach the cooled link. Since the elements of the kernelizer are not necessarily part of the system Hamiltonian, the condition is only necessary, but not yet sufficient. It can be turned into a sufficient condition by ensuring controllability on the complement of the target subspace. In other words, if the Lie algebra generated by the kernelizer has full rank, coherent interactions exist that, when added to the system Hamiltonian, ensure the necessary population flow towards the cooled link [91].

The necessary condition of local distinguishability of the target state is not fulfilled for certain highly entangled states, such as GHZ states. Indeed, these states can only be cooled with a number of meters that is extensive in system size. Conversely, the state space structure of the AKLT model is expected for any frustration-free spin chain with Hamiltonian given as a sum over noncommuting terms. These systems should therefore be amenable to dilute cooling. This was demonstrated for the example of the Majumdar-Ghosh model, a chain of spin-1/2 particles with nearest- and next-to-nearest neighbor interactions. Remarkably, a suitable combination of jump operators proved sufficient to selectively prepare one out of two degenerate ground states [91].

The key limitation of dilute cooling follows from the necessary conditions: As seen in the example of the GHZ states above, not every target state allows for the required partitioning of the state space. Moreover, the protocol requires the implementation of three-body interactions between the two sites of the cooled link and the meters. Three-body interactions are often engineered in an effective way from two-body interactions; this incurs an additional operational overhead. These limitations can be overcome by stochastic cooling which utilizes only two-body interactions between system sites and meters (the parameters of which are chosen randomly). Stochastic cooling does not make any assumptions on the target state except for being sufficiently gapped [92]. The approach is universal in the sense that it requires no detailed knowledge of the Hamiltonian or its spectrum. This generality requires a trade-off in efficiency: it necessitates many meter qubits prepared in their ground state, resulting in a large resource overhead [92].

The examples for engineered dissipation discussed above are based on meters that are simply discarded after their interaction with the system. Protocols that read out the state of the meter before disposing them are referred to as active. The corresponding

variant of quantum control is quantum feedback [84]. While beyond the scope of this lecture, the combined use of coherent interactions, quantum measurements, and quantum feedback, also referred to as “quantum interactive matter”, is an exciting emerging avenue of quantum control.

4 Open questions

With this lecture I have attempted to forge a link from the very basic concepts of quantum control theory, as established in the 1980s and 1990s, to present-day applications in quantum information science. The sheer number of the latter is growing at a breathtaking speed [1, 2], making quantum optimal control an indispensable tool in the practical implementation of quantum technologies. Moreover, the range of applications keeps expanding, as highlighted by the most recent examples from dissipation engineering in many-body systems [79]. It would therefore be presumptuous to predict the future directions of quantum control. Instead, I will conclude this lecture with a brief overview of open questions that I personally find interesting, challenging, or otherwise relevant. The discussion so far as well as my concluding remarks have deliberately adopted a focused perspective illustrating key concepts of quantum control and providing intuition with examples drawn primarily from the work of my research group.

Frequent questions that are raised after my lectures concern the robustness of control protocols as well as the comparison of machine learning to optimal control. While a detailed discussion is found in Ref. [2], let me summarize the key points. Robustness, as any property of a control, needs to be accounted for in the design of the protocol. In quantum optimal control, the most straightforward way is to include a desired feature in the target functional or to optimize over an ensemble of systems representing parameter fluctuations. However, these techniques are quite costly and often do not work “out of the box”. More practical approaches to include robustness would significantly improve the applicability of quantum optimal control. Moreover, it would be important to quantify whether robustness can be achieved for operations at the quantum speed limit [93] or whether one needs to balance robustness and operation speed.

Quantum optimal control and machine learning are closely linked [2]. For example, both need to quantify their targets in terms of suitable figures of merit, with a lot of potential for cross-fertilization. Both are hampered by the cost of solving the quantum mechanical equation of motion. They differ in their strengths and weaknesses. For example, quantum optimal control often has superior convergence properties but its control protocols are not readily applicable to slightly different settings. Machine learning is designed to be transferrable but is hampered by the cost of quantum measurements — quantum physics is not big data. A comprehensive and fair comparison of the two approaches is another open challenge.

An attentive reader will have wondered why controllability has been discussed primarily for unitary evolution whereas control protocols focused on open quantum

systems. Controllability of open quantum systems is largely uncharted territory! A key limitation for analyzing the controllability of open quantum systems is that traditional controllability analysis does not allow for resource constraints, including time. But open quantum systems are characterized by several competing timescales. If these need to be ignored, only asymptotic statements are possible. This explains why controllability results for open quantum systems have largely been confined to the characterization of reachable sets of states in Markovian dynamics [54]. In order to overcome this limitation, one would need to find a statement of the controllability problem that accounts for competing timescales. A second difficulty is given by the lack of a general equation of motion for open quantum systems. It is thus not clear how to rigorously state the controllability problem for systems with non-Markovian dynamics. These challenges call for new ideas. They will likely emerge from an overall better understanding of open systems control and the fast-paced development of quantum engineering.

Our current understanding of control strategies for open quantum systems clearly is incomplete. An attentive reader will have noticed a missing discussion of control strategies for non-Markovian dynamics in the preliminary balance in Sec. 3.5. Memory effects or strong coupling to environmental degrees of freedom may serve as a resource for control, as showcased in examples such as the implementation of quantum gates thanks to coupling to environmental modes [56]. But control design typically requires knowledge of the dynamics, and the difficulty to compute non-Markovian dynamics has so far impeded a comprehensive analysis.

Another avenue that is being explored only recently is the interplay of strong driving and dissipation. Take the example of dynamics described by Markovian GKLS master equations. If the system is strongly driven, the master equation must be derived in the dressed basis (with respect to the drive) [10]. The resulting drive-dependence of the jump operators can be leveraged in quantum optimal control, for example to speed up qubit reset [94]. Exploring the full potential of this type of control will be an important contribution to the field of dissipation engineering and quantum interactive matter.

Finally, a number of intriguing questions concern the “classical-quantum boundary” in the control of quantum systems. Whereas the external fields used in most of this discussion were classical, the auxiliary degrees of freedom in Sec. 3.5 can be viewed as quantum controllers. This perspective becomes particularly relevant when combining coherent control and quantum feedback [95]. More generally, whether an external field should be modeled as classical or quantum depends on its state. But is a given task better carried out with a classical or a quantum control? Naively, one would assume classical controls to be more readily available. At the same time, operating a quantum device with classical controls incurs both calibration overhead and noise. This problem is exacerbated in quantum feedback control which is based on classical external fields and quantum measurements (with additional costs). In contrast, quantum controllers can operate autonomously and thus minimize the costs associated with the classical-quantum interface. While already exploited in autonomous quantum error correction [96], its potential for quantum control more generally has not yet been tapped.

Acknowledgements I am indebted to my former and present students for dedicating their time, energy, and efforts to advancing quantum control theory. This lecture is largely based on the master's and doctoral theses by Daniel Reich, Michael Goerz, Sabrina Patsch, Daniel Basilewitsch, Karl Horn, Fernando Gago, and Josias Langbehn. I would like to thank Simone Sergi for inspiration and encouragement.

Competing Interests The author declares no competing interests.

References

- [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, “Training Schrödinger’s cat: Quantum optimal control. Strategic report on current status, visions and goals for research in Europe”, *Eur. Phys. J. D* **69**, 279 (2015).
- [2] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, “Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe”, *EPJ Quantum Technol.* **9**, 19 (2022).
- [3] L. S. Pontryagin, *Mathematical theory of optimal processes*, Classics of Soviet Mathematics (Routledge, 1987).
- [4] J. L. Speyer and D. H. Jacobson, *Primer on optimal control theory* (Society for Industrial and Applied Mathematics, 2010).
- [5] L. Levin, W. Skomorowski, L. Rybak, R. Kosloff, C. P. Koch, and Z. Amitay, “Coherent control of bond making”, *Phys. Rev. Lett.* **114**, 233003 (2015).
- [6] M. W. Heger, C. P. Koch, and D. M. Reich, “Optimized sampling of mixed-state observables”, *Phys. Rev. E* **100**, 052105 (2019).
- [7] H.-P. Breuer and F. Petruccione, *The theory of open quantum systems* (Oxford University Press, 2007).
- [8] D. A. Lidar, *Lecture Notes on the Theory of Open Quantum Systems*, arXiv:1902.00967, 2019.
- [9] D. Chruściński and S. Pascazio, “A Brief History of the GKLS Equation”, *Open Syst. Inf. Dyn.* **24**, 1740001 (2017).
- [10] T. Albash, S. Boixo, D. A. Lidar, and P. Zanardi, “Quantum adiabatic Markovian master equations”, *New J. Phys.* **14**, 123016 (2012).
- [11] R. Dann, A. Levy, and R. Kosloff, “Time-dependent Markovian quantum master equation”, *Phys. Rev. A* **98**, 052129 (2018).
- [12] C. Cohen-Tannoudji, B. Diu, and F. Laloë, *Quantum Mechanics, Volume 2: Angular Momentum, Spin, and Approximation Methods*, 2nd (Wiley-VCH, June 2020).
- [13] U. Boscain, M. Sigalotti, and D. Sugny, “Introduction to the Pontryagin Maximum Principle for Quantum Optimal Control”, *PRX Quantum* **2**, 030203 (2021).

- [14] K. Hornberger, S. Gerlich, P. Haslinger, S. Nimmrichter, and M. Arndt, “Colloquium: quantum interference of clusters and molecules”, [Rev. Mod. Phys.](#) **84**, 157–173 (2012).
- [15] M. Shapiro and P. Brumer, *Quantum Control of Molecular Processes*, 2nd, revised and enlarged (Wiley-VCH, Jan. 2012).
- [16] D. Tannor, *Introduction to Quantum Mechanics: A Time-Dependent Perspective* (University Science Books, Sausalito, 2007).
- [17] R. E. Goetz, C. P. Koch, and L. Greenman, “Quantum control of photoelectron circular dichroism”, [Phys. Rev. Lett.](#) **122**, 013204 (2019).
- [18] B. W. Shore, *Manipulating quantum structures using laser pulses* (Cambridge University Press, 2011).
- [19] F. Krausz, “Nobel lecture: sub-atomic motions”, [Rev. Mod. Phys.](#) **96**, 030502 (2024).
- [20] C. L. Degen, F. Reinhard, and P. Cappellaro, “Quantum sensing”, [Rev. Mod. Phys.](#) **89**, 035002 (2017).
- [21] Z. Amitay, A. Gandman, L. Chuntonov, and L. Rybak, “Multichannel selective femtosecond coherent control based on symmetry properties”, [Phys. Rev. Lett.](#) **100**, 193002 (2008).
- [22] M. Demirplak and S. A. Rice, “Adiabatic population transfer with control fields”, [J. Phys. Chem. A](#) **107**, 9937–9945 (2003).
- [23] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, “Shortcuts to adiabaticity: concepts, methods, and applications”, [Rev. Mod. Phys.](#) **91**, 045001 (2019).
- [24] N. V. Vitanov, A. A. Rangelov, B. W. Shore, and K. Bergmann, “Stimulated Raman adiabatic passage in physics, chemistry, and beyond”, [Rev. Mod. Phys.](#) **89**, 015006 (2017).
- [25] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, and M. Sarovar, “From pulses to circuits and back again: a quantum optimal control perspective on variational quantum algorithms”, [PRX Quantum](#) **2**, 010101 (2021).
- [26] C. P. Koch, “Controlling open quantum systems: tools, achievements, and limitations”, [J. Phys.: Condens. Matter](#) **28**, 213001 (2016).
- [27] D. Basilewitsch, C. P. Koch, and D. M. Reich, “Quantum optimal control for mixed state squeezing in cavity optomechanics”, [Adv. Quantum Technol.](#) **2**, 1800110 (2019).
- [28] J. P. Palao and R. Kosloff, “Optimal control theory for unitary transformations”, [Phys. Rev. A](#) **68**, 062308 (2003).
- [29] S. Patsch, D. M. Reich, J.-M. Raimond, M. Brune, S. Gleyzes, and C. P. Koch, “Fast and accurate circularization of a rydberg atom”, [Phys. Rev. A](#) **97**, 053418 (2018).
- [30] A. Larrouy, S. Patsch, R. Richaud, J.-M. Raimond, M. Brune, C. P. Koch, and S. Gleyzes, “Fast navigation in a large hilbert space using quantum optimal control”, [Phys. Rev. X](#) **10**, 021058 (2020).

- [31] A. Facon, E.-K. Dietsche, D. Grosso, S. Haroche, J.-M. Raimond, M. Brune, and S. Gleyzes, “A sensitive electrometer based on a Rydberg atom in a Schrödinger-cat state”, *Nature* **535**, 262–265 (2016).
- [32] T. Xia, X. L. Zhang, and M. Saffman, “Analysis of a controlled phase gate using circular Rydberg states”, *Phys. Rev. A* **88**, 062337 (2013).
- [33] S. R. Cohen and J. D. Thompson, “Quantum Computing with Circular Rydberg Atoms”, *PRX Quantum* **2**, 030322 (2021).
- [34] A. Signoles, E. K. Dietsche, A. Facon, D. Grosso, S. Haroche, J. M. Raimond, M. Brune, and S. Gleyzes, “Coherent Transfer between Low-Angular-Momentum and Circular Rydberg States”, *Phys. Rev. Lett.* **118**, 253603 (2017).
- [35] G. S. Agarwal, “State reconstruction for a collection of two-level systems”, *Phys. Rev. A* **57**, 671–673 (1998).
- [36] S. G. Johnson, *The NLOpt nonlinear-optimization package*, <https://github.com/stevengj/nlopt>, 2007.
- [37] M. H. Goerz, K. B. Whaley, and C. P. Koch, “Hybrid optimization schemes for quantum control”, *EPJ Quantum Technol.* **2**, arXiv:1505.05331, 21 (2015).
- [38] V. Krotov, *Global Methods in Optimal Control Theory* (CRC Press, Oct. 1995).
- [39] D. M. Reich, M. Ndong, and C. P. Koch, “Monotonically convergent optimization in quantum control using krotov’s method”, *J. Chem. Phys.* **136**, 104103 (2012).
- [40] M. H. Goerz, D. Basilewitsch, F. Gago-Encinas, M. G. Krauss, K. P. Horn, D. M. Reich, and C. P. Koch, “Krotov: A Python implementation of Krotov’s method for quantum optimal control”, *SciPost Phys.* **7**, 80 (2019).
- [41] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, “Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms”, *Journal of Magnetic Resonance* **172**, 296–305 (2005).
- [42] J. Johansson, P. Nation, and F. Nori, “QuTiP: An open-source Python framework for the dynamics of open quantum systems”, *Computer Physics Communications* **183**, 1760–1772 (2012).
- [43] J. Johansson, P. Nation, and F. Nori, “Qutip 2: a python framework for the dynamics of open quantum systems”, *Computer Physics Communications* **184**, 1234–1240 (2013).
- [44] M. M. Müller, D. M. Reich, M. Murphy, H. Yuan, J. Vala, K. B. Whaley, T. Calarco, and C. P. Koch, “Optimizing entangling quantum gates for physical systems”, *Phys. Rev. A* **84**, 042315 (2011).
- [45] P. Watts, J. Vala, M. M. Müller, T. Calarco, K. B. Whaley, D. M. Reich, M. H. Goerz, and C. P. Koch, “Optimizing for an arbitrary perfect entangler. i. functionals”, *Phys. Rev. A* **91**, 062306 (2015).
- [46] M. H. Goerz, G. Gualdi, D. M. Reich, C. P. Koch, F. Motzoi, K. B. Whaley, J. Vala, M. M. Müller, S. Montangero, and T. Calarco, “Optimizing for an arbitrary perfect entangler. ii. application”, *Phys. Rev. A* **91**, 062307 (2015).

- [47] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge University Press, 2000).
- [48] P. Doria, T. Calarco, and S. Montangero, “Optimal control technique for many-body quantum dynamics”, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [49] D. Basilewitsch, R. Schmidt, D. Sugny, S. Maniscalco, and C. P. Koch, “Beating the limits with initial correlations”, *New J. Phys.* **19**, 113042 (2017).
- [50] A. Halaski, M. G. Krauss, D. Basilewitsch, and C. P. Koch, “Quantum optimal control of squeezing in cavity optomechanics”, *Phys. Rev. A* **110**, 013512 (2024).
- [51] K. W. Moore and H. Rabitz, “Exploring constrained quantum control landscapes”, *J. Chem. Phys.* **137**, 134113 (2012).
- [52] D. D’Alessandro, *Quantum control and dynamics* (Chapman and Hall (CRC), 2008).
- [53] R. A. Bertlmann and P. Krammer, “Bloch vectors for qudits”, *Journal of Physics A: Mathematical and Theoretical* **41**, 235303 (2008).
- [54] G. Dirr, U. Helmke, I. Kurniawan, and T. Schulte-Herbrüggen, “Lie-semigroup structures for reachability and control of open quantum systems: kossakowski-lindblad generators form lie wedge to markovian channels”, *Rep. Math. Phys.* **64**, 93–121 (2009).
- [55] J. F. Poyatos, J. I. Cirac, and P. Zoller, “Quantum reservoir engineering with laser cooled trapped ions”, *Phys. Rev. Lett.* **77**, 4728–4731 (1996).
- [56] D. M. Reich, N. Katz, and C. P. Koch, “Exploiting Non-Markovianity for Quantum Control”, *Sci. Rep.* **5**, 12430 (2015).
- [57] S. Deffner and S. Campbell, “Quantum speed limits: from heisenberg’s uncertainty principle to optimal quantum control”, *J. Phys. A: Math. Theo.* **50**, 453001 (2017).
- [58] T. Schulte-Herbrüggen, A. Spörl, N. Khaneja, and S. J. Glaser, “Optimal control for generating quantum gates in open dissipative systems”, *J. Phys. B* **44**, 154013 (2011).
- [59] D. M. Reich, G. Gualdi, and C. P. Koch, “Minimum number of input states required for quantum gate characterization”, *Phys. Rev. A* **88**, 042309 (2013).
- [60] D. M. Reich, G. Gualdi, and C. P. Koch, “Optimal strategies for estimating the average fidelity of quantum gates”, *Phys. Rev. Lett.* **111**, 200401 (2013).
- [61] M. H. Goerz, D. M. Reich, and C. P. Koch., “Optimal control theory for a unitary operation under dissipative evolution”, *New J. Phys.* **16**, 055012 (2014).
- [62] A. Romer, D. M. Reich, and C. P. Koch, “Reconstructing the unitary part of a noisy quantum channel”, *arXiv:2507.17648* (2025).
- [63] M. H. Goerz, F. Motzoi, K. B. Whaley, and C. P. Koch, “Charting the circuit-qed design landscape using optimal control theory”, *npj Quantum Inf.* **3**, 37 (2017).
- [64] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, “Minimum Construction of Two-Qubit Quantum Operations”, *Phys. Rev. Lett.* **93**, 020502 (2004).
- [65] D. Basilewitsch, H. Yuan, and C. P. Koch, “Optimally controlled quantum discrimination and estimation”, *Phys. Rev. Research* **2**, 033396 (2020).

- [66] D. A. Lidar and K. Birgitta Whaley, “Decoherence-free subspaces and subsystems”, in *Irreversible quantum dynamics*, edited by F. Benatti and R. Floreanini, arXiv:quant-ph/0301032 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2003), pp. 83–120.
- [67] S. Qin, M. Cramer, C. P. Koch, and A. Serafini, “Optimal control for Hamiltonian parameter estimation in non-commuting and bipartite quantum dynamics”, *SciPost Phys.* **13**, 121 (2022).
- [68] D. Suter and G. A. Álvarez, “Colloquium: protecting quantum information against environmental noise”, *Rev. Mod. Phys.* **88**, 041001 (2016).
- [69] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J.-C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, and A. Wallraff, “Fast and unconditional all-microwave reset of a superconducting qubit”, *Phys. Rev. Lett.* **121**, 060502 (2018).
- [70] D. Basilewitsch, J. Fischer, D. M. Reich, D. Sugny, and C. P. Koch, “Fundamental bounds on qubit reset”, *Phys. Rev. Research* **3**, 013110 (2021).
- [71] N. Khaneja, R. Brockett, and S. J. Glaser, “Time optimal control in spin systems”, *Phys. Rev. A* **63**, 032308 (2001).
- [72] K. P. Horn, F. Reiter, Y. Lin, D. Leibfried, and C. P. Koch, “Quantum optimal control of the dissipative production of a maximally entangled state”, *New J. Phys.* **20**, 123010 (2018).
- [73] D. C. Cole, S. D. Erickson, G. Zarantonello, K. P. Horn, P.-Y. Hou, J. J. Wu, D. H. Slichter, F. Reiter, C. P. Koch, and D. Leibfried, “Resource-efficient dissipative entanglement of two trapped-ion qubits”, *Phys. Rev. Lett.* **128**, 080502 (2022).
- [74] Y. Lin, J. P. Gaebler, F. Reiter, T. R. Tan, R. Bowler, A. S. Sørensen, D. Leibfried, and D. J. Wineland, “Dissipative production of a maximally entangled steady state of two quantum bits”, *Nature* **504**, 415–418 (2013).
- [75] M. Leibscher, E. Pozzoli, C. Pérez, M. Schnell, M. Sigalotti, U. Boscain, and C. P. Koch, “Full quantum control of enantiomer-selective state transfer in chiral molecules despite degeneracy”, *Commun. Phys.* **5**, 110 (2022).
- [76] F. Gago-Encinas, M. Leibscher, and C. P. Koch, “Graph test of controllability in qubit arrays: a systematic way to determine the minimum number of external controls”, *Quantum Sci. Technol.* **8**, 045002 (2023).
- [77] F. Gago-Encinas, T. Hartung, D. M. Reich, K. Jansen, and C. P. Koch, “Determining the ability for universal quantum computing: testing controllability via dimensional expressivity”, *Quantum* **7**, 1214 (2023).
- [78] L. Funcke, T. Hartung, K. Jansen, S. Kühn, and P. Stornati, “Dimensional Expressivity Analysis of Parametric Quantum Circuits”, *Quantum* **5**, 422 (2021).
- [79] P. M. Harrington, E. J. Mueller, and K. W. Murch, “Engineered dissipation for quantum information science”, *Nature Reviews Physics* **4**, 660–671 (2022).
- [80] J. Langbehn, PhD thesis (Freie Universität Berlin, 2026).
- [81] R. Menu, J. Langbehn, C. P. Koch, and G. Morigi, “Reservoir-engineering shortcuts to adiabaticity”, *Phys. Rev. Research* **4**, 033005 (2022).

- [82] M. Sveistrys, J. Langbehn, R. Menu, S. Campbell, G. Morigi, and C. P. Koch, “Speeding up quantum annealing with engineered dephasing”, [Quantum](#) **9**, 1731 (2025).
- [83] B. Kraus, H. P. Büchler, S. Diehl, A. Kantian, A. Micheli, and P. Zoller, “Preparation of entangled states by quantum Markov processes”, [Physical Review A](#) **78**, 042307 (2008).
- [84] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control* (Cambridge University Press, Cambridge, 2009).
- [85] F. Ciccarello, S. Lorenzo, V. Giovannetti, and G. M. Palma, “Quantum collision models: open system dynamics from repeated interactions”, [Phys. Rep.](#) **954**, Quantum collision models: Open system dynamics from repeated interactions, 1–70 (2022).
- [86] M. Cattaneo, G. De Chiara, S. Maniscalco, R. Zambrini, and G. L. Giorgi, “Collision models can efficiently simulate any multipartite markovian quantum dynamics”, [Phys. Rev. Lett.](#) **126**, 130403 (2021).
- [87] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, “Rigorous results on valence-bond ground states in antiferromagnets”, [Phys. Rev. Lett.](#) **59**, 799–802 (1987).
- [88] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, “Valence bond ground states in isotropic quantum antiferromagnets”, [Commun.Math. Phys.](#) **115**, 477–528 (1988).
- [89] S. Roy, J. T. Chalker, I. V. Gornyi, and Y. Gefen, “Measurement-induced steering of quantum systems”, [Physical Review Research](#) **2**, 033347 (2020).
- [90] K. C. Smith, E. Crane, N. Wiebe, and S. Girvin, “Deterministic Constant-Depth Preparation of the AKLT State on a Quantum Processor Using Fusion Measurements”, [PRX Quantum](#) **4**, 020315 (2023).
- [91] J. Langbehn, K. Snizhko, I. Gornyi, G. Morigi, Y. Gefen, and C. P. Koch, “Dilute measurement-induced cooling into many-body ground states”, [PRX Quantum](#) **5**, 030301 (2024).
- [92] J. Langbehn, G. Mouloudakis, E. King, R. Menu, I. Gornyi, G. Morigi, Y. Gefen, and C. P. Koch, “Universal cooling of quantum systems via randomized measurements”, [arXiv:2506.11964](#) (2025).
- [93] D. Basilewitsch, S.-D. Börner, C. Berke, A. Altland, S. Trebst, and C. P. Koch, “Chaotic fluctuations in a universal set of transmon qubit gates”, [arXiv:2311.14592](#) (2023).
- [94] D. Basilewitsch, F. Cosco, N. L. Gullo, M. Möttönen, T. Ala-Nissilä, C. P. Koch, and S. Maniscalco, “Reservoir engineering using quantum optimal control for qubit reset”, [New J. Phys.](#) **21**, 093054 (2019).
- [95] P. Rouchon, “A tutorial introduction to quantum stochastic master equations based on the qubit/photon system”, [Annual Reviews in Control](#) **54**, 252–261 (2022).
- [96] M. Mirrahimi, Z. Leghtas, V. V. Albert, S. Touzard, R. J. Schoelkopf, L. Jiang, and M. H. Devoret, “Dynamically protected cat-qubits: a new paradigm for universal quantum computation”, [New Journal of Physics](#) **16**, 045014 (2014).

- [97] S. M. Girvin, “Circuit QED: superconducting qubits coupled to microwave photons”, in Proceedings of the 2011 Les Houches Summer School on Quantum Machines, edited by M. Devoret, R. Schoelkopf, B. Huard, and L. F. Cugliandolo (2014).

Appendix

5 Prerequisites

5.1 Frame transformations and rotating-wave approximation (RWA): Two-level system

Assuming that we need to consider only the ground and first excited state of an atom, the atomic Hamiltonian, or drift, is given by

$$\hat{H}_0 = \frac{\omega_0}{2} \hat{\sigma}_z,$$

and the electric dipole interaction with a laser becomes

$$\hat{H}_I(t) = \mathbf{d} \cdot \mathbf{E}(t) \hat{\sigma}_x,$$

where $\mathbf{E}(t) = E_0 S(t) \cos(\omega_L t)$. Here, a shape function $S(t)$ has been introduced in order to model the switching on and off of the electric field. Since the electric field in $\hat{H}_I(t)$ appears as it is realized in the lab, this representation is referred to as “lab frame”.

When we consider resonant excitation of the atom (which is a prerequisite for truncating the large atomic Hilbert space to just two levels), ω_0 and ω_L differ only by a small amount, the detuning $\Delta_L = \omega_0 - \omega_L$. When oscillations with $\omega_{0/L}$ are much faster than all other relevant timescales, it is useful to transform the time-dependent Schrödinger equation,

$$\frac{\partial}{\partial t} |\psi(t)\rangle = -i\hat{H}(t) |\psi(t)\rangle,$$

into a reference frame that rotates with one of the large frequencies, ω_0 or ω_L . This is achieved by a unitary transformation $\hat{U}$ in the usual way, that is, we insert $|\psi\rangle = \hat{U} |\psi'\rangle$ and identity into the time-dependent Schrödinger equation,

$$\frac{\partial}{\partial t} (\hat{U}(t) |\psi'(t)\rangle) = -i\hat{U}(t)\hat{U}^\dagger(t)\hat{H}(t) (\hat{U}(t) |\psi'(t)\rangle),$$

which can be rewritten for $|\psi'(t)\rangle$,

$$\frac{\partial}{\partial t} |\psi'(t)\rangle = -i \left(\hat{U}^\dagger(t) \hat{H}(t) \hat{U}(t) - \hat{U}^\dagger(t) \frac{\partial \hat{U}(t)}{\partial t} \right) |\psi'(t)\rangle = -i\hat{H}'(t) |\psi'(t)\rangle.$$

The frame rotating with ω_0 is reached when taking $\hat{U} = e^{-i\hat{H}_0 t}$. In this case $\frac{\partial \hat{U}}{\partial t} = -i\hat{H}_0 \hat{U}(t)$ and $\hat{H}'(t)$ becomes

$$\hat{H}'(t) = \begin{pmatrix} 0 & -\frac{1}{2}\Omega_0 S(t) (e^{-i(\omega_L - \omega_0)t} + e^{-i(\omega_L + \omega_0)t}) \\ -\frac{1}{2}\Omega_0 S(t) (e^{i(\omega_L - \omega_0)t} + e^{i(\omega_L + \omega_0)t}) & 0 \end{pmatrix},$$

where we have introduced the Rabi frequency $\Omega_0 = \mathbf{d} \cdot \mathbf{E}_0$. The rotating wave approximation consists in neglecting the terms that oscillate with twice the large frequency,

$$\hat{H}'_{RWA}(t) = \begin{pmatrix} 0 & -\frac{1}{2}\Omega_0 S(t) e^{i\Delta_L t} \\ -\frac{1}{2}\Omega_0 S(t) e^{-i\Delta_L t} & 0 \end{pmatrix}.$$

This is justified when the fast oscillations average to zero on the timescale of interest. Indeed, all remaining time-dependencies in $\hat{H}'_{RWA}(t)$ are slow since $\Delta_L \ll \omega_L, \omega_0$ and $S(t)$ has been introduced as the envelope of the electric field amplitude.

In this simple example, the transformation into the rotating frame is identical to the transformation into the interaction picture. For N -level systems, the two usually do not coincide any longer, as discussed below. In general, several choices for a rotating frame are possible, while there is only one interaction picture. This can be seen already for a two-level system when considering a more general expression for the electric field, $\mathbf{E}(t) = \mathbf{E}_0 S(t) \cos(\omega_L t + \varphi(t))$ including a time-dependent phase $\varphi(t)$. Such a phase is generated for example when "chirping" the pulse. In order to understand this expression, it is instructive to sketch the electric field as a function of time for $\varphi(t) = 0$ and $\varphi(t) = \alpha t^2$ where α is the so-called chirp rate. The general form of the unitary transformation into a rotating frame for a two-level system is given by $\hat{U}(t) = \exp(i/2\vartheta(t)\hat{\sigma}_z)$. In the transformation above, $\vartheta(t)$ was taken to be $\vartheta = \omega_0 t$ but equally valid choices are $\vartheta(t) = \omega_L t$ or $\vartheta(t) = \omega_L t + \varphi(t)$. With these choices, the frame rotates with the central, respectively the instantaneous, laser frequency. It is instructive to derive the corresponding Hamiltonians $\hat{H}''_{RWA}(t)$, $\hat{H}'''_{RWA}(t)$. It is also instructive to visualize the difference between the frames on the Bloch sphere: The equation of motion for a two-level system can be rewritten in terms of the Bloch vector $\mathbf{r}(t) \in \mathbb{R}^3$ precessing around a vector $\mathbf{\Omega}(t)$ representing the Hamiltonian,

$$\frac{d}{dt} \mathbf{r}(t) = \mathbf{r}(t) \times \mathbf{\Omega}(t) \quad \text{with} \quad \hat{H}(t) = \mathbf{\Omega}(t) \cdot \boldsymbol{\sigma}.$$

5.2 Frame transformations and rotating-wave approximation for N -level systems

A good choice of rotating frame is essential. First of all, it simplifies the equations and provides a more intuitive understanding of the (essential) dynamics. Further, and even more importantly, for N -level systems with $N > 2$ it gives rise to different

ways to take the rotating-wave approximation. This is exploited for example when deriving effective Hamiltonians to model superconducting qubits [97].

To see how one can generalize the concept of frame transformation, let us consider a three-level system where levels 1, 2 and levels 2, 3 are connected by a dipole transition (i.e., the corresponding matrix elements in the Hamiltonian are non-zero) [18]. The system interacts with a two-color electric field $\mathbf{E}(t) = \mathbf{E}_1 S_1(t) \cos(\omega_1 t) + \mathbf{E}_2 S_2(t) \cos(\omega_2 t)$, such that the total Hamiltonian, $\hat{H}(t) = \hat{H}_0 + \hat{H}_I(t)$, is given by

$$\hat{H}(t) = \begin{pmatrix} \mathcal{E}_1 & \hat{\mathbf{d}}_{12} \cdot \mathbf{E}(t) & 0 \\ \hat{\mathbf{d}}_{21} \cdot \mathbf{E}(t) & \mathcal{E}_2 & \hat{\mathbf{d}}_{23} \cdot \mathbf{E}(t) \\ 0 & \hat{\mathbf{d}}_{32} \cdot \mathbf{E}(t) & \mathcal{E}_3 \end{pmatrix}$$

Applying the following unitary transformation to the time-dependent Schrödinger equation,

$$\hat{U} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \exp(-i\omega_1 t) & 0 \\ 0 & 0 & \exp(-i(\omega_1 - \omega_2)t) \end{pmatrix},$$

results in

$$\hat{H}' = \begin{pmatrix} 0 & \frac{1}{2} \hat{\mathbf{d}}_{12} (\mathbf{E}_1 (e^{i\Delta_1 t} + e^{i(\omega_1 + \omega_{21})t}) + \mathbf{E}_2 (e^{i(\omega_2 - \omega_{21})t} + e^{i(\omega_2 + \omega_{21})t})) & 0 \\ c.c. & \Delta_1 & c.c. \\ 0 & \frac{1}{2} \hat{\mathbf{d}}_{32} (\mathbf{E}_1 (e^{i(\omega_1 - \omega_{32})t} + e^{i(\omega_1 + \omega_{32})t}) + \mathbf{E}_2 (e^{i\Delta_2 t} + e^{i(\omega_2 + \omega_{32})t})) & -\Delta_{2P} \end{pmatrix},$$

where $\omega_{ij} = \mathcal{E}_i - \mathcal{E}_j$ are the transition frequencies, $\Delta_1 = \omega_1 - \omega_{21}$, $\Delta_2 = \omega_2 - \omega_{32}$ are the one-photon detunings, and $\Delta_{2P} = \omega_{32} - (\omega_1 + \omega_2)$ is the two-photon detuning. It is instructive to compare this Hamiltonian to the one obtained from transforming into the interaction picture, i.e., for $\hat{U} = \exp(i\hat{H}_0 t)$ where the diagonal becomes zero and the detunings appear only in the off-diagonal matrix elements.

Now we assume that ω_1, ω_2 are near-resonant to the two transitions in the three-level system. This allows us to take the so-called two-photon rotating wave approximation where we neglect terms that oscillate with sum frequencies as well as those oscillating with the large difference frequencies $\omega_1 - \omega_{32}$, $\omega_2 - \omega_{21}$, while keeping those with the small difference frequencies, i.e., the one-photon detunings Δ_1, Δ_2 . This reflects that while the electric field contains both spectral components, the probability for the first (second) component to excite transitions between $|2\rangle$ and $|3\rangle$ ($|1\rangle$ and $|2\rangle$) is very small. In other words, within the RWA each transition "sees" only the near-resonant component of the electric field. The Hamiltonian then becomes

$$H'_{RWA} = \begin{pmatrix} 0 & \frac{1}{2} \Omega_1(t) e^{i\Delta_1 t} & 0 \\ \frac{1}{2} \Omega_1(t) e^{-i\Delta_1 t} & \Delta_1 & \frac{1}{2} \Omega_2(t) e^{i\Delta_2 t} \\ 0 & \frac{1}{2} \Omega_2(t) e^{-i\Delta_2 t} & \Delta_{2P} \end{pmatrix},$$

which is very often taken as the starting point [18, 24], but one should keep in mind that it is the result of an approximation.

Comparing the unitary transformations for two- and three-level systems, we can devise a general strategy for frame transformations and rotating-wave approximations for N -level systems: The unitary is diagonal and the complex phases determine by how much the corresponding lab-frame energy is shifted. In comparison to the interaction picture, where all lab-frame energies are shifted to zero and all detunings appear in the off-diagonal matrix elements of the Hamiltonian, one can also choose frame transformations where the detunings appear on the diagonal, and the off-diagonal matrix elements become real.