UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Fisica e Astronomia "Galileo Galilei"

Corso di Laurea in Fisica

Tesi di Laurea

Esplorazione del controllo ottimale quantistico mediante diverse basi troncate

Exploration of quantum optimal control using different truncated bases

Relatore

Prof. Simone Montangero

Correlatore

Prof. Marcus Huber
TU Wien

Laureando

Chiara Bettio

Anno Accademico 2023/2024

Abstract

Precise control over quantum systems is essential to exploit their properties effectively, mitigating the effects of external influences such as decoherence and noise. This is crucial for advancing quantum technologies. Quantum optimal control plays a fundamental role in this development, providing a family of techniques to optimise the behaviour of quantum systems through tailored external fields. In this thesis, we introduce how to formulate and solve quantum optimal control problems, focusing on how the selection of the basis function impacts the performance of chopped random basis (CRAB) algorithms. We use the Quantum Optimal Control Suite (QuOCS) Python framework, applying the dressed CRAB algorithm optimising a quantum state transfer in the presence of detuning, caused by inevitable imperfections or fluctuations in the system or control field. We compare the effectiveness of two bases with different properties: the Fourier and Sigmoid bases. Both bases perform well in the absence of detuning. However, as the detuning value increases, the difference in their performance becomes more pronounced, with the Fourier basis demonstrating superior effectiveness due to its ability to exploit the frequency space more effectively. This confirms the importance of choosing an appropriate basis, as certain properties may align better with specific problems. We use this insight to optimise a pulse robust against detuning variations in an ensemble of qubits.

Il controllo preciso dei sistemi quantistici è fondamentale per sfruttare efficacemente le loro proprietà, mitigando gli effetti di influenze esterne come decoerenza e rumore. Questo aspetto è di cruciale importanza per il progresso delle tecnologie quantistiche. Il controllo ottimale quantistico gioca un ruolo fondamentale in questo sviluppo, fornendo una famiglia di tecniche per controllare il comportamento dei sistemi quantistici tramite campi esterni ottimizzati su misura. In questa tesi, presentiamo come formulare e risolvere problemi di controllo ottimale quantistico, concentrandoci sull'impatto della scelta della base di funzioni sulle prestazioni degli algoritmi a base casuale troncata (CRAB). Utilizzando il framework Python Quantum Optimal Control Suite (QuOCS), applichiamo l'algoritmo dressed CRAB per ottimizzare un trasferimento di stato quantistico in presenza di detuning, causato da inevitabili imperfezioni o fluttuazioni del sistema o del campo di controllo. Confrontiamo l'efficacia di due basi con proprietà distinte: la base di Fourier e la base Sigmoidale. Entrambe le basi portano risultati soddisfacenti in assenza di detuning. Con l'aumentare del detuning invece, la differenza nelle loro prestazioni diventa più evidente, con la base di Fourier che dimostra una maggiore efficacia grazie alla sua abilità di sfruttare lo spazio delle frequenze in modo più efficace. Questo conferma l'importanza della scelta di una base appropriata, dato che alcune proprietà possono adattarsi meglio a problemi specifici. Sfruttiamo questa informazione per ottimizzare un impulso robusto alle variazioni di detuning in un insieme di qubit.

Acknowledgements

I want to thank Professor Simone Montangero and Professor Marcus Huber for their willingness to supervise this project, providing me with the opportunity to do this work in collaboration with UniPd and TU Wien. I am especially grateful to Dr. Phila Rembold, whose support and guidance throughout the entire process were fundamental in shaping this thesis. Thank you for the shared insights, careful revisions, and invaluable advice.

Ringrazio la mia famiglia, per aver dimostrato sempre e solo pazienza e supporto nei confronti miei e delle mie scelte di vita.

I am most thankful to the friends I made along the way, for walking with me through different phases of this journey, and to those who have been there for me since long before I started it.

And thanks to you, Marco, for everything else.

Contents

Introduction

Quantum physics is the study of nature at its fundamental level, describing reality at the smallest scales. Since the first developments at the beginning of the 20th century, the theory of quantum mechanics has rapidly evolved following landmark discoveries such as Planck's quantisation of energy [1], Einstein's explanation of the photoelectric effect [2], and Heisenberg, Schrödinger and Von Neumann contributions, among many others, leading to the formulation of a quantum theory [3–5].

The theoretical foundations have since transitioned into practical applications, with interest in quantum technologies steadily increasing over the past decades [6]. Typical quantum properties such as superposition and entanglement offer distinct advantages over classical technologies in various fields. Superposition allows quantum systems to exist in a combination of states before measurement. Entanglement describes the quantum correlation among a group of particles, where each particle's state is intrinsically linked to the others, even across large distances. In quantum computing this translates to the possibility to store and process information differently, allowing significant speedups through quantum algorithms for specific computations, such as Grover's algorithm for search in an unstructured database [7] and Shor's algorithm for number factoring [8]. Quantum communication enables physically secure communication over long distances, employing quantum cryptography protocols for quantum key distribution (QKD) [9] of keys generated with quantum random number generators [10]. Quantum sensing and metrology exploit the high sensitivity of quantum systems to external perturbations, enhancing measurements with quantum clocks [11], spin sensors for magnetic fields [12], and gravity sensors, among many others [13].

However, the development of quantum technologies poses significant challenges: quantum properties are delicate and easily disrupted by external influences, necessitating protocols to mitigate noise, operational imperfections, and decoherence, the loss of information from a system due to interactions with the environment. Overcoming the fragility of quantum systems is essential to transform theoretical knowledge into practical applications. This requires precise control of quantum systems.

Quantum Optimal Control (QOC) [14] plays a critical role in achieving this goal, employing tailored external fields to direct the dynamics of quantum systems. The development of QOC is driven by the desire to fully exploit the potential of quantum systems by optimising a wide set of quantum operations, such as reducing errors [15], speeding up operations [16], preparing and transferring quantum states [17, 18], making systems robust against noise [19], and optimising resource utilisation. This field includes both analytical and numerical pulse-shaping methods [20–22].

We focus on the chopped random basis (CRAB) algorithm [22, 23] and its dressed version dCRAB [24]. CRAB is particularly versatile, allowing both open-loop (numerical) and closed-loop (experimental) approaches, and it can be combined with either gradient-based or gradient-free updating algorithms. In this method, the time-dependent control pulse is parameterised using a finite number of elements randomly selected from a truncated basis set, hence the term chopped random basis. While some bases are more popular and regularly chosen as default, different bases have been implemented and employed for optimal control problems, showing how an educated choice can improve the results of the optimisation [25].

The aim of this project is to explore how the choice of basis function influences the algorithm's effectiveness in finding an optimal solution for the QOC problem of a state transfer in the presence of

constant detuning.

The structure of this thesis is as follows:

- Chapter 1 illustrates the fundamental concepts of quantum optimal control, introducing its theoretical background, the formulation of control problems, and intrinsic limitations. We also discuss numerical algorithms, with a focus on the CRAB method.
- Chapter 2 explores the motivation and approach for using different truncated bases in Quantum Optimal Control, focusing on the Fourier and Sigmoid bases.
- Chapter 3 applies the methods to the control problem of a state transfer for a qubit in the presence of detuning, introducing the concept of robustness to address the issue of an ensemble of qubits with different detunings.

Chapter 1

Quantum Optimal Control Theory

This chapter explores the fundamental concepts of quantum optimal control [14], covering its theoretical background, the formulation of a control problem, and its intrinsic limitations. Additionally, we introduce numerical algorithms employed in quantum optimal control, focusing on the chopped-random basis method (CRAB) [22, 23] which will later be used for the optimisations in Chapter 3.

1.1 Introduction

Control theory provides a mathematical framework to understand and regulate the behaviour of controlled dynamical systems [26]. These are systems whose evolution is influenced by external inputs we can control. The main goal is to determine how to manipulate the inputs to achieve a desired outcome on the measurable outputs or on the final state of the system. Within this framework, optimal control [27] describes a family of strategies to achieve the desired outcome in a way that optimises a specific objective. The objective is defined through optimality criteria, which can involve minimising time or energy consumption, maximising accuracy, or achieving any other measurable goal. The control objective quantifies the effectiveness of the optimised pulse. If the control objective is minimised we call it a cost function. For clarity, in the rest of this thesis, we refer to optimisations involving the minimisation of the cost function.

Optimisation is common in mathematics and engineering and is most prominently addressed using the calculus of variations, through the Euler-Lagrange equations [28]. An early example is the brachistochrone problem, posed by Johann Bernoulli in 1696 [29], which involves finding the curve between two points that minimises the travel time for an object moving under the influence of a uniform gravitational field. Optimal control extends the calculus of variations to include systems with time-dependent controls [27].

Quantum Optimal Control (QOC) [14] applies classical optimal control principles to quantum systems and has been successfully implemented on various platforms, such as superconducting qubits [30], colour centres in diamonds [15], trapped ions [31] and nuclear spin systems in nuclear magnetic resonance (NMR) [32]. Over the past decades, this field has seen rapid development due to the increasing interest in quantum technologies [6]. The developed techniques extend from specialised analytical tools [20] to different numerical algorithms able to adapt to experimental limitations and uncertainties [22, 33].

1.2 Formulation of the control problem

To formulate a control problem, we define a mathematical model for the system and its time evolution. We consider a quantum system evolving under controlled dynamics through an externally applied control field $u(t) = (u_1(t), u_2(t), \ldots, u_n(t))$. Let $\mathcal{H}$ be the Hilbert space of the system and H the

quantum Hamiltonian operator, which can be expressed in the bilinear form:

$$H(t) = H^d + \sum_{i=1}^{n} u_i(t) H_i^c. \tag{1.1}$$

Here H^d represents the time-independent drift Hamiltonian, including the constant terms we are not actively controlling. H_i^c are the control Hamiltonians modulated by the control pulses $u_i(t)$, describing the coupling to the applied control field, e.g. a laser pulse.

The Hamiltonian represents the total energy of the quantum system and describes its dynamics. We restrict our focus to closed systems, not interacting with the environment via non-Hamiltonian terms. Additionally, we simplify the notation to pure states, described as single vectors in the Hilbert space $\mathcal{H}$. The state of a closed pure quantum system at time t is represented by the vector $|\psi(t)\rangle \in \mathcal{H}$, whose evolution is governed by the solution to the time-dependent Schrödinger equation:

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = H|\psi(t)\rangle, \tag{1.2}$$

where i is the imaginary unit and $\hbar$ is the reduced Planck's constant. The solution to this equation can be expressed in terms of the unitary evolution operator $U(t)$, which defines the evolution of the system from the initial state $|\psi(0)\rangle$ to the state at time t $|\psi(t)\rangle$:

$$|\psi(t)\rangle = U(t)|\psi(0)\rangle = e^{-\frac{i}{\hbar}Ht}|\psi(0)\rangle. \tag{1.3}$$

The control problem consists of optimising $u(t)$ over a given time interval according to a cost function J or figure of merit. A proper cost function is a mathematical expression quantifying the goal we want to achieve. It associates a cost with each possible control field. The general expression of J, known as the Bolza form [34], is:

$$J(u, T) = F(\psi(T), T) + \int_0^T L(t, \psi(t), u(t)) dt. \tag{1.4}$$

The state of the system at the final time T of the evolution is denoted as $|\psi(T)\rangle$, determined according to Eq (1.3). The terminal cost $F(\psi(T), T)$ is evaluated at the final time of the system's evolution and describes the optimisation goal. It can for example quantify the distance between the final state and the target state. Running costs $\int_0^T L(t, \psi(t), u(t)) dt$ are evaluated over the entire pulse duration T and are potentially included to discourage undesired control behaviours, such as surpassing power or frequency limits, by adding penalty terms to the cost function. Stricter boundaries can be enforced through control space restrictions, discussed in Sec. 1.3.2.

For simplicity, in this thesis, we consider Mayer-type [34] cost functions, where $L = 0$.

Various control objectives can be considered and have been studied. The most common control problem is the state-to-state transfer: the goal is driving the system from an initial state $|\psi(0)\rangle$ to the target state $|\psi_T\rangle$ in a finite time T. The overlap between the final state $|\psi(T)\rangle$ and target state can be quantified through the fidelity $\mathcal{F}$:

$$\mathcal{F} = |\langle \psi_T | \psi(T) \rangle|^2, \tag{1.5}$$

which equals 1 if the target state is equal to the final state and 0 if they are orthogonal. As it monotonously increases with the closeness between the target and final state, it represents a terminal cost for state-to-state transfer. The control objective can involve minimising the infidelity $\mathcal{I} = 1 - \mathcal{F}$.

We can extend Eq (1.5) to the optimisation of an observable O expectation value, for example using $J = \langle \psi(T) | O | \psi(T) \rangle$.

Other popular control problems include the optimisation of a unitary gate [35] and the maximisation of entanglement, for instance through the von Neumann entropy [36]. It is also possible to set time minimisation as a control goal [16, 37].

1.3 Restrictions and constraints

There are some limits to when control objectives are theoretically and practically achievable. This section explores three primary restrictions: the controllability of the system, the restrictions introduced by the topology of the control landscape, and the intrinsic limitation imposed by the quantum speed limit.

1.3.1 Controllability

A quantum system is defined as controllable in a given set of configurations if there exists a time-dependent control $u(t)$ capable of driving the system from any initial configuration to any final configuration within a finite amount of time [38]. Various degrees of controllability exist [39], depending on the configurations under consideration (e.g. pure or mixed state controllability, gate controllability). For an N-level closed system a necessary and sufficient condition for complete controllability, the strongest requirement, can be determined with group-theoretical methods [40]. The drift and control Hamiltonians determine whether the full space state is reachable.

It should be noted that a lack of controllability implies that only a subset of configurations can be accessed, allowing some specific control tasks to remain achievable.

1.3.2 Control landscape

Once the system's controllability is established, finding the optimal driving field can be arbitrarily hard. In other words, after confirming the theoretical existence of optimal solutions, one can explore the complexity of finding them.

The control objectives represented as a functional of the control field $u(t)$ determine the control landscape:

$$J(u) = F(|\psi(T)\rangle). \tag{1.6}$$

The optimisation usually leads to the critical points of the landscape, for which the first-order functional derivative with respect to the control field $u(t)$ is zero at all times:

$$\frac{\delta J(u)}{\delta u(t)} = 0 \ \forall t \in [0, T]. \tag{1.7}$$

Therefore, the search complexity of the optimisation task is related to the landscape's topology, influencing the convergence of optimisation algorithms and the likelihood of finding optimal (global) solutions [38].

An example of a generic control landscape is illustrated in Fig. 1.1. For finite-dimensional controllable closed quantum systems with no constraints on the controls, the landscape is free of sub-optimal local minima that could trap the optimisation algorithm, interfering with the search for optimal controls [38]. Instead, any additional sub-optimal points are all saddle points that do not impede finding optimal solutions [42].

However, in practice, control problems are bounded by experimental limitations that can alter the landscape's topology. Significant constraints on the controls, such as limited bandwidth, control power limitations, or noise, can restrict access to optimal solutions isolating portions of the landscape or creating apparent traps that complicate the optimisation process.

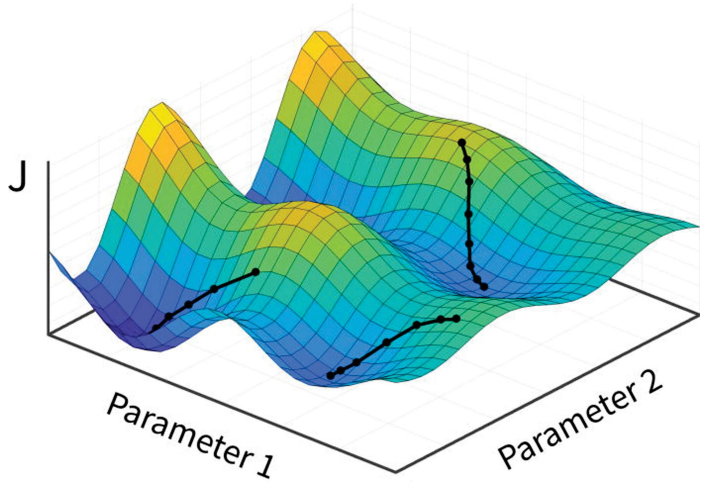

Figure 1.1: Example of control landscape considering two control parameters. In the case of minimisation, the valley represents the locally optimal combinations of control. The black paths represent optimisations starting from different initial guesses, typically moving towards the closest local minima. Figure taken from [41].

1.4 Pulse constraints

We present here some strategies for handling common pulse constraints inevitable in any experiment: amplitude and time.

1.4.1 Amplitude restriction

One common constraint in quantum control is the amplitude of the control pulse. One approach to incorporate amplitude constraints is to impose a cut-off on the control pulses, ensuring their amplitudes remain within a specified range $|u_i(t)| \leq u_i^{\max}$:

$$\widetilde{u}_i(t) = \begin{cases} u_i(t) & \text{if } -u_i^{\max} < u_i(t) < u_i^{\max} \\ u_i^{\max} & \text{if } u_i(t) \geq u_i^{\max} \\ -u_i^{\max} & \text{if } u_i(t) \leq -u_i^{\max} \end{cases} \tag{1.8}$$

The disadvantage of this approach is introducing high-frequency terms to the pulse due to the emerging edges, visible in Fig. 1.2.

Alternatively, we can maintain the original form of the pulse by applying a smooth amplitude restriction. This includes mapping the pulse to fit within the available amplitude space using a continuous, limited function, such as:

$$\widetilde{u}_i(t) = u_i^{\max} \sin u_i(t) \tag{1.9}$$

This method avoids edges, minimising spectral distortions [43].

1.4.2 Time restriction

Limiting the duration of time pulses is another restriction that must be considered when implementing the control pulse. The pulse will start at $t = 0$ and end at $t = T$, leading to $u(t) = 0$ if $t \leq 0$ or $t \geq T$. Similarly to the amplitude restriction, this can be done through cut-off or by multiplying the pulse by a window function, gradually reducing its amplitude to zero at the beginning and end of the pulse duration, as illustrated in Fig. 1.2.

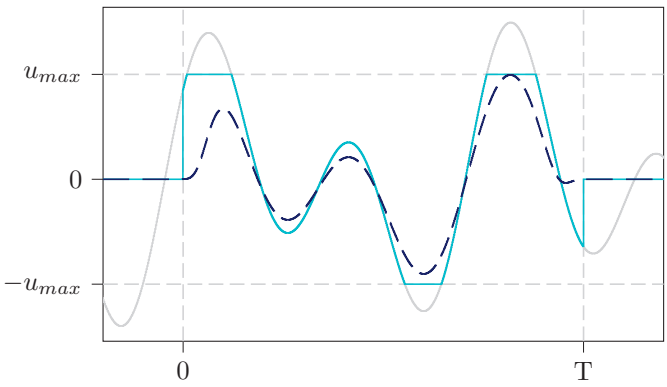

Figure 1.2: Application of restriction approaches. The solid line represents the unconstrained pulse (grey) and the cut-off approach (light blue) in both amplitude and time. The dashed line represents the application of a scaling function and a super-Gaussian [44] window function, effectively reducing the edges and corresponding high-frequency terms.

1.4.3 Quantum speed limit

The quantum speed limit (QSL) is a fundamental constraint rooted in the time-energy uncertainty relation, setting an intrinsic time scale for unitary quantum dynamics [45]. It defines the maximum speed at which a quantum system can evolve in its Hilbert space, setting a lower bound on the time required for a quantum system to evolve from an initial state to a final one. This limit is generally difficult to calculate.

Mandelstam and Tamm derived the first expression of the quantum speed limit time τ_{QSL}, considering systems with a time-independent Hamiltonian [46]. Their derivation is connected to the variance of the Hamiltonian of the quantum system $\Delta H^2 = \langle\psi_0|(H-E)^2|\psi_0\rangle$, $E = \langle\psi_0|H|\psi_0\rangle$. For a process in which the systems evolves from any initial state to an orthogonal final state, they found:

$$\tau_{\mathrm{QSL}} = \frac{\pi\hbar}{2\Delta H} \tag{1.10}$$

Margolus and Levitin [47] later derived an independent similar expression in terms of the Hamiltonian's expectation value, leading to the combined bound [48]:

$$\tau_{\mathrm{QSL}} = \max\left\{\frac{\pi}{2}\frac{\hbar}{\Delta H}, \frac{\pi}{2}\frac{\hbar}{\langle H\rangle}\right\} \tag{1.11}$$

In the context of quantum optimal control, the QSL sets the time scale for an algorithm to converge to an optimal solution [49]. Here τ_{QSL} is given by Battacharyya's bound [50], the generalisation of Eq. (1.10) for arbitrary angles:

$$\tau_{\mathrm{QSL}} = \frac{\arccos|\langle\psi_0|\psi_T\rangle|}{\Delta H} \tag{1.12}$$

This ultimately means that even optimal evolutions cannot be achieved arbitrarily fast for a finite energy of the control pulse. The effects become significant when the time approaches the theoretical QSL.

QOC can be used to determine the controls that drive the system at the quantum speed limit [49].

Additionally, there is a lower limit on the information that needs to be encoded in the control pulse for the optimisation to be successful over a certain error. The dimension of a QOC problem, or the number of independent degrees of freedom in the control pulse, should at least match the dimension of the problem [19, 51].

1.5 Numerical algorithms

Iterative numerical algorithms allow us to handle complex control problems, which are hard or impossible to solve analytically. Numerical methods parameterise the pulse by expanding it into a series of functions. This can include smooth basis functions, such as the Fourier basis, or the discretisation of the pulse in time slices [33]. The coefficients of the expansion are treated as optimisation variables and iteratively optimised. We refer to them as control parameters. The optimisation starting point in the control landscape is defined as the initial guess. The choice of the initial guess can be random or physically motivated. One can also employ a different optimisation with fewer parameters and use the optimised pulse as the initial guess for the original problem [52]. Once the system has evolved, the algorithm checks whether the control objectives have been reached, and either stops or initiates another iteration.

Quantum optimal control algorithms are categorised according to different criteria. A first distinction is whether the cost is derived from a mathematical model (open-loop optimisation) or obtained directly from the system in an experimental setup (closed-loop optimisation) [15, 53]. The closed-loop configuration utilises the direct feedback of the system. The cost function is measured and not computed, allowing remarkable results even for complex (measurable) systems. The external control field is applied and then updated at each iteration according to the measured value of the control function. Open-loop algorithms are typically applied when an accurate theoretical model of the system's Hamiltonian is available. The outcome from open-loop optimal control might serve as an initial guess for subsequent closed-loop optimisation [54].

Gradient-based methods, such as Krotov [55] and GRAPE [56] employ the information about the gradient of the cost function with respect to the control parameters to iteratively refine them. This makes them efficient when the information on the gradient is available and can be easily calculated, as they make use of all the available information. Gradient-free methods are more suitable for complex systems with computationally expensive cost functions [57], or whenever computing the gradient is difficult or impossible. Not needing information on the gradient makes them viable for closed-loop optimisation.

1.5.1 The CRAB algorithm

The chopped random basis algorithm (CRAB [22]) was originally proposed as an optimisation method for complex quantum many-body dynamics as it can be combined with a gradient-free approach. It has since been successfully applied to various control problems and systems, in both closed and open-loop optimisation, and combined with gradient-based and gradient-free methods [23].

The idea behind CRAB is to expand the control pulse with a truncated basis function, also referred to as a chopped basis. This means that only a randomised selection of N_s elements of the original basis function is utilised, restricting the optimisation to a subspace of the original control landscape. Each element $f(s, \vec{A}; t)$ is characterised by a superparameter s and the optimisation parameters $\vec{A}$. An example of superparameters is the frequency for the Fourier basis, described in more detail in Sec. 2.2.1. The number of parameters to optimise depends on the choice and formulation of the basis. The control field is therefore expressed as:

$$u(t) = \sum_{i=1}^{N_s} u_i(t) = \sum_{i=1}^{N_s} f(s_i, \vec{A}_i; t)[+f_0(A_0; t)], \tag{1.13}$$

where $f_0(A_0; t)$ is a term only relevant for the Sigmoid basis, described in Sec. 2.2.2. The parameters $\vec{A}$ are then optimised, via any method of choice, to minimise the figure of merit. The randomness in the choice of the basis elements has been shown to enhance the algorithm convergence compared to employing selected relevant values [22].

Chopped bases are useful in the presence of relevant constraints, such as bandwidth limitations, which

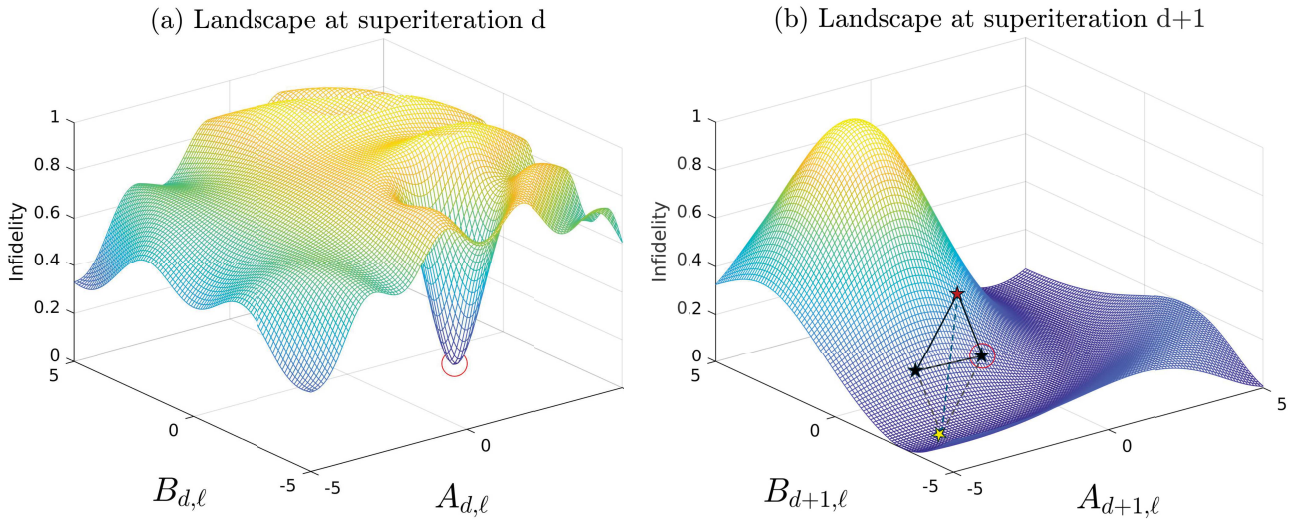

Figure 1.3: Example of how the landscape transforms by changing the basis. In (a) the optimisation ends at the local minimum indicated by the red circle, utilised as the initial guess for the next superiteration. The change of the landscape in (b) allows the dCRAB algorithm to escape the trap. The black triangle indicates the initial test points for Nelder-Mead. The worst vertex (red star) is reflected and substituted with the new improved test point (yellow star) for the next iteration. Figure adapted from [60].

can be included, effectively restricting the optimisation to the accessible sub-space of the total control space.

The optimisation is performed using gradient-based methods or direct search methods. Direct search methods do not require the cost function to be differentiable, as they only rely on the values of the cost function without needing its gradient or attempting to estimate it [58]. One example of a commonly used direct search method is the Nelder-Mead algorithm [59]. This algorithm is simplex-based. A simplex is defined as a group of $n+1$ vertices in an n-dimensional space, e.g. a triangle in a two-dimensional space. During each iteration, the objective function is evaluated at each vertex, which are used as test points. The test points are ordered according to their proximity to the goal. The one further from the goal is referred to as the worst vertex and is reflected through the centroid of the remaining n points to find a new substitute test point candidate. If the new test point is better than the old one, it proceeds to the next iteration starting from the new simplex. Otherwise, the check is repeated for an extended or contracted reflection, and eventually, if no better directions are found, the whole simplex is shrunk to restrict the considered area and eventually converges to the optimal point. Calculations terminate when the stopping criteria are met. This might include exceeding a determined number of iterations or reaching the convergence criteria such as the target control objective or the limit set for the precision of the difference in the evaluation of the cost function. One example of iteration for the Nelder-Mead algorithm in a two-dimensional space is illustrated in Fig. 1.3 (b).

However, confining the search basis to N_s dimensions might lead the algorithm to converge towards sub-optimal points known as false traps.

Here δJ (Eq. (1.7)) vanishes for all the δu in the truncated basis expansion, but not considering all the variations in the original control space. A false-trap is an extremum only in the subspace defined by the truncated basis function. An example of a false-trap is illustrated in Fig. 1.3 (a).

The "dressed" version of CRAB (dCRAB [24]) addresses this issue by initiating a new CRAB optimisation, called superiteration, using the solution from the previous iteration as an initial guess and employing a different random selection of basis elements.

The control pulse at the j^{th} superiteration is formulated as:

$$u^j(t) = u^{j-1}(t) + \sum_{i=1}^{N_s} u_i^j(t), \tag{1.14}$$

where $u^{j-1}(t)$ is the optimised control pulse obtained from the previous superiteration, and $u_i^j(t)$ are the random basis elements introduced in the current superiteration, as calculated in Eq. (1.13). This formulation ensures that the new control pulse consists of one part containing the already optimised old pulse and another part exploring a new direction, enhancing the exploration of the control landscape and reducing the likelihood of getting trapped in local minima.

Chapter 2

Exploration of different bases

While the Fourier basis is a common choice in implementing the dCRAB algorithm introduced in Sec. 1.5.1, alternative bases with different shapes and properties such as Sigmoid, Chebyshev polynomials [61], or piece-wise constant elements [55, 56] can also be used. This chapter explains the motivation behind exploring different bases and provides details about the considered bases: Fourier and Sigmoid basis. Finally, we will describe our approach to comparing the bases.

2.1 Motivation

As mentioned in Sec. 1.5.1, a chopped basis restricts the search for the control to a subsection of the total control landscape. Different bases confine the search to different sections and may align better with specific control tasks or physical systems. Accessing a smoother part of the landscape with fewer local minima can enhance the chances of finding a global optimum and allow for faster convergence.

The choice of basis also determines which shapes are complex or simple to generate. The properties and shape of a basis can lead to representing an optimal control field using fewer basis elements. For instance, square pulses are hard to produce with Fourier basis elements but easier for Sigmoid functions. This consideration is significant for problems where the expected shape of optimal solutions has predictable or known properties [25, 62].

Furthermore, as discussed in Sec. 1.4, quantum control fields are subject to various physical constraints such as maximum amplitude and bandwidth limits. Some bases naturally incorporate these constraints, making it easier to ensure that the optimised control fields are physically realisable. For example, in the Fourier basis, it is simple to limit the bandwidth by restricting the choice of basis elements. This can also help exclude basis elements that would not contribute to the optimisation, such as high or low-frequency values, impacting the optimisation efficiency.

2.2 Approach

To explore how the choice of truncated basis functions impacts the solution to the control problem, we test them on a simple problem: a state transfer in the presence of detuning (Chapter 3). Specifically, we compare the Fourier and Sigmoid bases, which have distinct properties. We examine their impact on convergence behaviour and solutions quality using the dCRAB algorithm, described in Sec. 1.5.1 and implemented in the QuOCS [63] software suite. To mimic an experimental optimisation process, in Chapter 3 we employ the dCRAB algorithm combined with a gradient-free optimisation, considering a qubit modelled after NV centres in diamonds [64].

We begin this section by describing the two bases and their properties, followed by a discussion on comparability considerations.

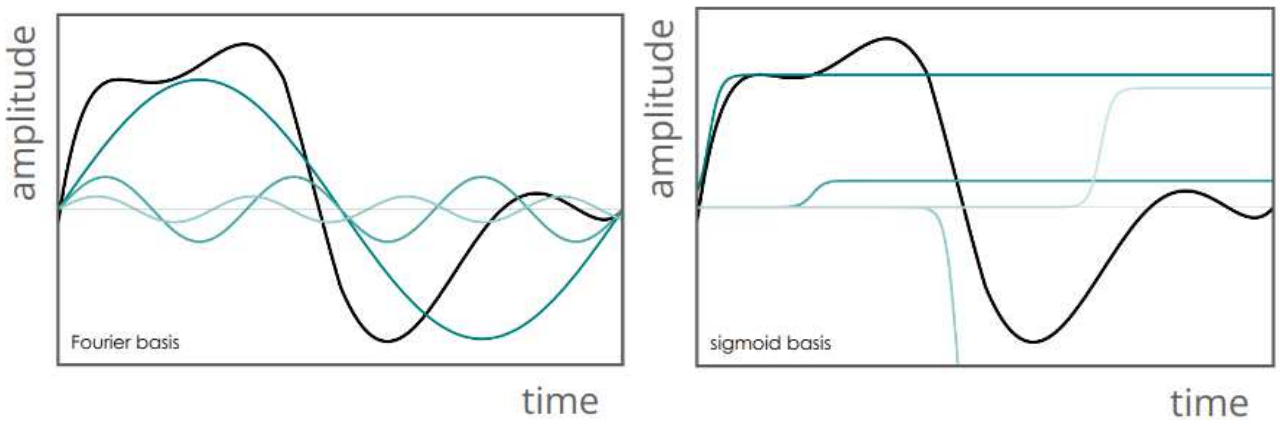

Figure 2.1: Example of decomposition of the same pulse with Fourier and Sigmoid basis. Figure taken from [63].

2.2.1 Fourier basis

The Fourier basis is commonly used with dCRAB. The randomly generated superparameter is the frequency s. Each basis element is composed of a sine and a cosine function at frequency s and is expressed as:

$$f(s, [A^1, A^2]; t) = A^1 \sin st + A^2 \cos st. \tag{2.1}$$

This results in two parameters to be optimised for each basis element. Consequently, the total number of parameters to be optimised in the optimisation process is $N_p = 2N_s$.

In QuOCS [63] the frequencies are determined as $s = \frac{2\pi\omega_i}{T}$, where T is the time duration of the pulse, and ω_i is randomly generated following a uniform distribution within the selected interval. The frequency limits are implemented by directly setting boundaries on the random generation of the superparameters, ensuring $\omega_i \in [\omega_{\min}, \omega_{\max}]$.

2.2.2 Sigmoid basis

Sigmoid functions are curves that smoothly transition between two values. We consider Sigmoid functions defined by the integral of a Gaussian. The basis is composed of smooth time steps, where the randomly generated superparameter s is the time of the step. The width σ influences the rise-time and bandwidth of the pulse. In Sec. 2.2.3 we discuss how to limit the frequency through a restriction on the width.

The Sigmoid basis element is given by:

$$f(s, A; t) = \frac{A}{\sqrt{2\pi}\sigma} \int_{-\infty}^{t} e^{-\frac{1}{2}(\frac{t'-s}{\sigma})} dt' = \frac{A}{2} \left(1 + \operatorname{erf}\left(\frac{t-s}{\sqrt{2}\sigma}\right)\right), \tag{2.2}$$

$$f_0(A_0; t) = f(0, A_0; t). \tag{2.3}$$

This leads to one parameter to be optimised for each basis element, plus the initial amplitude A_0, resulting in $N_p = N_s + 1$ parameters to be optimised.

In QuOCS [63], the width σ is a hyperparameter that is preliminarily set to a constant value (default 0.1). The superparameters $s = \tau$ are randomly generated with a flat distribution in the selected interval $[\tau_{\min}, \tau_{\max}]$.

2.2.3 Comparability considerations

To ensure a fair comparison, the bases should cover comparable ranges of the function space. Therefore, similar restrictions will be applied to each basis.

For the Fourier basis, as mentioned in Sec. 2.2.1, the frequency range can be directly set in the settings as $\omega \in [\omega_{min}, \omega_{max}]$. However, for the Sigmoid basis, there is no such direct method. Since the rise-time of the functions is determined by their width, we can analyse the spectrum of a Sigmoid pulse by calculating its Fourier Transform (FT). This leads to the bandwidth limiting envelope [60]:

$$Y^{\infty}_{max}(\omega) = A_{\max}(\tau_N - \tau_0)e^{-\frac{1}{2}\sigma^2\omega^2}, \tag{2.4}$$

calculated for a square pulse with amplitude $A_{\max}$ and limiting any constructable pulse. The Sigmoid basis is therefore bandwidth-limited through a Gaussian envelope. We can impose a frequency limit by setting [25]:

$$\sigma = \frac{\sqrt{-2\ln\epsilon_\omega}}{\omega_{\max}}. \tag{2.5}$$

This ensures that:

$$\frac{|Y^{\infty}_{\max}(\omega)|}{A_{\max}(\tau_N - \tau_0)} = \epsilon_\omega^{\frac{\omega^2}{\omega_{\max}^2}},$$

effectively reducing the relative amplitude of any component with $\omega \geq \omega_{\max}$ to ϵ_ω at most.

The same amplitude and time restrictions are applied as well. We shrink the pulses by applying the smooth amplitude restriction implemented in QuOCS (see Sec. 3.1) and multiplying them by a flat-top super-Gaussian [44] window function before optimisation, as shown in Fig. 1.2.

In addition, the same number of N_c parameters are optimised, corresponding to $N_c = 2N_s$ for the Fourier basis and $N_c = N_s + 1$ for Sigmoid.

Chapter 3

Application

3.1 Quantum Optimal Control Suite

The Quantum Optimal Control Suite (QuOCS) [63] is an open-source software package implemented in Python. It includes quantum optimal control algorithms such as GRAPE and dCRAB, as well as some direct search algorithms, offering a customisable interface. QuOCS performs both open and closed-loop optimisations, as it is designed to be easily integrated with experimental control setups.

3.1.1 Main settings

The main settings include the choice of the algorithm, whether the objective is minimisation or maximisation, the number of superiterations (SI), restrictions on the maximum number of function evaluations, time limits, and the figure of merit goals for each superiteration. This also includes the choice of the direct search method and its stopping criteria.

Amplitude constraints for the pulse are available, allowing the selection of upper and lower limits and the option to restrict through cut-off or shrinking. Considering the pulse $u_i(t)$ with a maximum (or minimum) amplitude value $u_{\mathrm{max,i}}(t)$ and the amplitude limit u_{max}, the shrink function is implemented as:

$$\widetilde{u}_i(t) = \begin{cases} u_i(t) & \text{if } u_{\mathrm{max,i}}(t) < u_{\mathrm{max}} \\ \frac{u_{\mathrm{max}}}{u_{\mathrm{max,i}}(t)} u_{\mathrm{max,i}}(t) & \text{otherwise} \end{cases} \tag{3.1}$$

A user-defined initial guess and window function can be implemented as well.

We use the dCRAB algorithm combined with the Nelder-Mead method [59], minimising the infidelity $\mathcal{I}$ with a figure of merit goal $J = e^{-10}$ and a maximum number of total evaluations set to 3000. We employ the window function combined with the shrink option.

3.1.2 Bases settings

The bases implemented in QuOCS are Fourier, Chebyshev, Piece-wise constant, and Sigmoid.

The settings for the bases include the number of vectors used for the expansion and the limits on the superparameters, which are sampled from a uniform distribution by default. For the Sigmoid basis, the width and offset (default value 0.1) can be selected.

Our optimisations employ Sigmoid and Fourier bases implemented here, with amplitude and superparameter limits discussed in Sec. 3.2.2. Unless stated otherwise, all optimisations are performed with three superiterations and optimising $N_p = 4$ parameters, leading to two basis vectors for Fourier and three for Sigmoid.

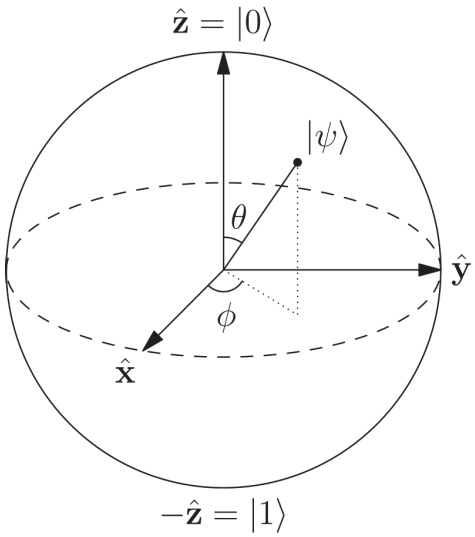

Figure 3.1: Visualisation of qubit states on the Bloch Sphere. The state $|0\rangle$ is represented by the unit vector $\vec{z}$, the state $|1\rangle$ by -$\vec{z}$. $|\psi\rangle$ represents the generic state. Any point on the surface of the sphere is associated with a pure state, points inside of the sphere represent mixed states. Image taken from [66].

3.2 Single qubit with detuning

A qubit, short for quantum bit, is the basic unit of quantum information. Physically a qubit can be any two-level quantum system, such as a system's ground and excited state, a spin system, or two orthogonal polarisation directions of a photon. The two levels are represented in the computational basis as $|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $|1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, spanning the entire Hilbert space of the system $\mathcal{H}^2$ [65]. Generally, a pure qubit state exists in a superposition of the basis states, expressed as:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad \alpha, \beta \in \mathbb{C},\ |\alpha|^2 + |\beta|^2 = 1. \tag{3.2}$$

The Bloch sphere is a three-dimensional unitary sphere that serves as a geometric representation of the qubit space. Here, a pure qubit state can be represented by introducing the notation:

$$|\psi\rangle = \cos\frac{\theta}{2}|0\rangle + \sin\frac{\theta}{2}e^{i\phi}|1\rangle \quad 0 \leq \theta \leq \pi,\ 0 \leq \phi \leq 2\pi, \tag{3.3}$$

with the angles θ and ϕ defining a point on the sphere (Fig. 3.1).

In this section, we consider a state transfer for a single qubit in the presence of detuning. In quantum computing, a state transfer refers to an operation performed on a qubit, typically involving the rotation of its state vector on the Bloch sphere from an initial state to a target state.

We model the qubit based on nitrogen-vacancy (NV) centres in diamonds [64], spin systems used for quantum sensing [41], and other quantum technologies [67]. An NV centre is a point defect in diamonds, where one of the carbon atoms is replaced by a nitrogen atom, accompanied by a neighbouring vacancy site. The NV centre forms a spin triplet system. A static external magnetic field $B_{\|}$ is applied in the $\vec{z}$ direction (the NV axis) to lift the degeneracy of the excited level $m_s = \pm 1$. The two qubits level are the ground state $|0\rangle$ $m_s = 0$ and the exited state $|1\rangle$ associated to one of the two $m_s = \pm 1$.

We can manipulate the qubit's state by applying an external microwave pulse. In the absence of detuning, this manipulation is achieved using a square pulse with the appropriate length and amplitude (Sec. 3.2.1).

When a control pulse is applied to manipulate the qubit state, its frequency ω_c should match the qubit's natural frequency ω_{nv} needed to drive the transition $|0\rangle \leftrightarrow |1\rangle$. However, in practical scenarios,

deviations can occur, leading to detuning. Detuning happens when the frequency of the control pulse ω_c does not precisely match the qubit's natural frequency ω_{nv} due to experimental imperfections, fluctuations in control parameters, or interactions with the qubit's environment. For instance, an inhomogeneous magnetic field $B_\parallel$ along z. As a result, the operation may not have the desired effect on the qubit's state. To mitigate the effects of detuning, time-dependent control pulses can be optimised and implemented using quantum optimal control techniques.

We first analyse and analytically solve the equation for a constant pulse to illustrate how a constant detuning influences the operation. Then, for increasing values of detuning, we explore time-dependent optimal pulses as alternatives. As a final step, we optimise a pulse robust against detuning to address the issue of state transfer for an ensemble of qubits with different detunings. The optimisations are performed using the dCRAB algorithm, employing both the Fourier and Sigmoid bases separately.

3.2.1 Analytical solution for a constant pulse

The Hamiltonian H incorporates the qubit's energy and its interaction with the control field. The specific form of the detuning term depends on the qubit system's implementation and the influences considered. For NV centres it typically involves the qubit's transition frequency $\omega_{nv} = B_\parallel \gamma_{nv}$, depending on the static magnetic field and the gyromagnetic ratio of the NV centre γ_{nv}, and the control field frequency ω_c. We consider a constant detuning value $\Delta = \omega_c - \omega_{nv}$. The Hamiltonian in the rotating frame of $B_\perp$ (using the rotating wave approximation) is then:

$$H = \frac{\hbar}{2}(\Delta\sigma_z + \Omega\sigma_x), \tag{3.4}$$

where $\sigma_{x,z}$ are the Pauli X and Z operators, and Ω is the amplitude of the rectangular pulse applied to the system. The eigenenergies and respective eigenstates of H are [68]:

$$E_{+,-} = \pm\frac{\hbar}{2}\sqrt{\Delta^2+\Omega^2};\quad |\psi_+\rangle = \begin{pmatrix}\cos\left(\frac{\theta}{2}\right)\\ \sin(\frac{\theta}{2})\end{pmatrix},\ |\psi_-\rangle = \begin{pmatrix}-\sin\left(\frac{\theta}{2}\right)\\ \cos(\frac{\theta}{2})\end{pmatrix};\quad \tan\theta = \frac{\Omega}{\Delta}. \tag{3.5}$$

Considering an evolution from the initial state $|\psi(0)\rangle = |0\rangle$, we can determine the fidelity between the state at time t and the target state $|\psi_T\rangle = |1\rangle$:

$$\mathcal{F}(t) = |\langle\psi(t)|1\rangle|^2 = \frac{\Omega^2}{\Delta^2+\Omega^2}\sin^2\left(\sqrt{\Delta^2+\Omega^2}\ \frac{t}{2}\right). \tag{3.6}$$

The fidelity oscillates over time. Since $\sin^2 x \leq 1\ \forall x$, we find that $\mathcal{F}(t) \leq \frac{\Omega^2}{\Delta^2+\Omega^2} < 1$ when $\Delta \neq 0$. This means we can reach the target state only when $\Delta = 0$. Eq. (3.6) then gives the relation:

$$\mathcal{F}_0(t) = \sin^2\left(\Omega\ \frac{t}{2}\right). \tag{3.7}$$

with $\mathcal{F}_0(t) = 1$ for $\Delta = 0$, $t = k\frac{\pi}{\Omega}$, k constant integer value.

For a fixed pulse length T, we can define Ω_π, corresponding to a π rotation on the Bloch Sphere in the absence of detuning, leading to the state transfer we aim to achieve, as:

$$\Omega_\pi = \frac{\pi}{T}. \tag{3.8}$$

In the optimisations of Sec. 3.2.2 we consider this as a unit for both the amplitude constraint of the control field and the values of detuning, setting $T = \pi$ for convenience.

The calculated fidelity in the range $\Delta \in [-10, 10]$ for a constant pulse of length $T = \pi$ and unitary amplitude is represented in Fig. 3.5.

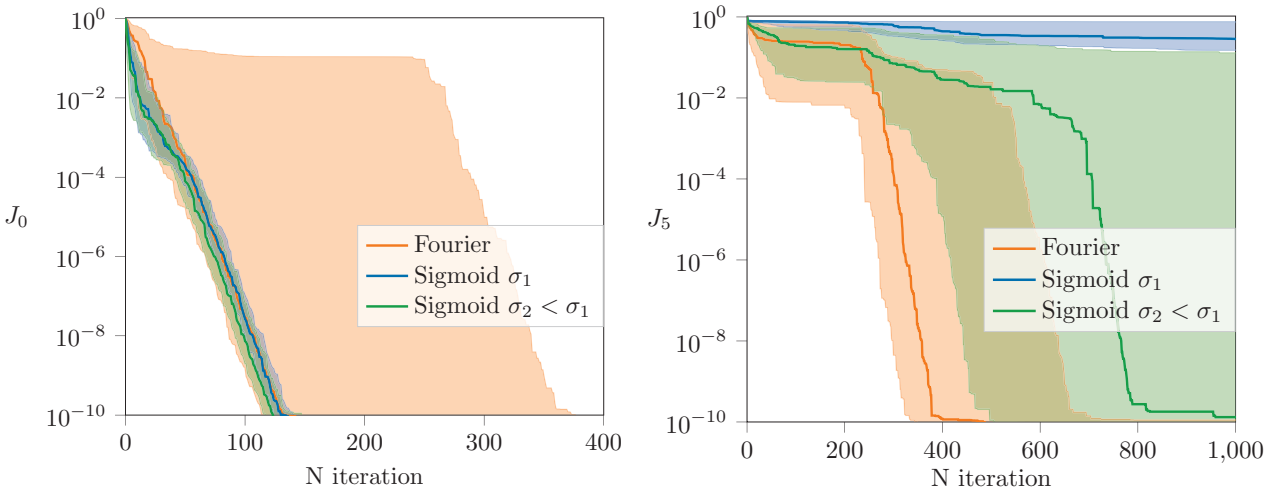

Figure 3.2: Convergence behaviour of 100 optimisations for the three considered bases. The median optimised infidelity is depicted, along with a shaded area representing the 25th and the 75th percentile. On the left, for $\Delta = 0$, the Sigmoid bases perform significantly better in terms of convergence time. On the right, for $\Delta = 5\Omega_\pi$, Fourier performs better, with Sigmoid σ_1 not converging to an optimal solution for some cases and σ_2 not converging at all in the considered percentile range. In both cases $T = \pi$ and $u_{\max} = 5\Omega_\pi$.

3.2.2 Optimisation of a time-dependent pulse

Following Eq. (1.1) the Hamiltonian of the system is:

$$H = \frac{\hbar}{2}(\Delta\sigma_z + u(t)\sigma_x) \tag{3.9}$$

The goal of the optimisation is finding the control field $u(t)$ that minimises the infidelity $\mathcal{I}$.

Once we set the time duration of the pulse T, the amplitude constraint determines whether optimal solutions are achievable for determinate detuning values (see QSL on Sec. 1.4.3).

We use Ref. [69] as orientation, optimising only the σ_x component of the pulse and exploring a looser amplitude restriction $u_{\max} = 5\Omega_\pi$ for all optimisations.

The frequencies of the Fourier basis have also been restricted based on empirical tests, with the optimal range $\omega_i \in [-5, 5]$. Higher frequencies negatively impacted the optimisation convergence. The Sigmoid basis is similarly constrained by setting its width σ_1 to satisfy Eq. (2.5), with $\omega_{\max} = 5$ and $\epsilon_\omega = 0.2$. As the inability of Sigmoid to exploit the frequency space as efficiently as Fourier seems to impede the search for optimal solutions, we compare the results with a different Sigmoid basis with a broader width $\sigma_2 = 0.1$, corresponding to $\omega_{\max} = 18$.

In the absence of detuning, all three bases achieve optimal values. Fig. 3.2 illustrates their convergence behaviour. The two Sigmoid bases have similar convergence patterns. The Fourier basis between the 50th and 75th percentile requires more iterations to converge. The visible plateau indicates that some pulses need more than one superiteration to reach the optimal solution. This does not seem to apply to Sigmoid, suggesting a smoother landscape with fewer false traps for the considered system.

As detuning increases within $\Delta = [-10, 10]$, the fidelity varies as depicted in Fig. 3.5. Considering $\Delta = 5$ as an example, Fourier and Sigmoid with σ_1 still achieve optimal results, while Sigmoid with σ_2 fails to reach an optimal value in most of the optimisations, as illustrated in Fig. 3.2. Since the best Fourier pulses show a more oscillatory behaviour, this might be due to the Sigmoid basis being constrained by its Gaussian spectrum, which tends to produce more low-frequency values. This observation is supported by σ_2, with a broader frequency range, performing significantly better than σ_1. This behaviour becomes more pronounced as the detuning increases. Fig. 3.3 illustrates examples of both the best and worst pulses for Fourier and Fig. 3.4 the best pulses for the two Sigmoid bases, for

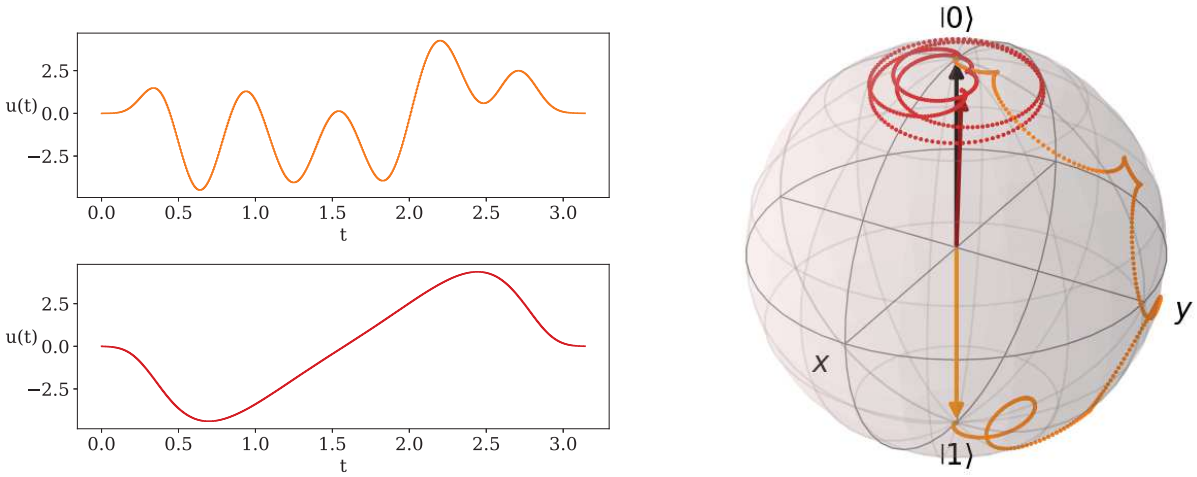

Figure 3.3: Best (orange) and worst (red) optimised pulses for Fourier and their path on the Bloch sphere, for an optimisation with $T = \pi$, $u_{\max} = 5\Omega_\pi$ and $\Delta = 10\Omega_\pi$. The black vector represents the initial state, and the orange and red vectors are the respective final states.

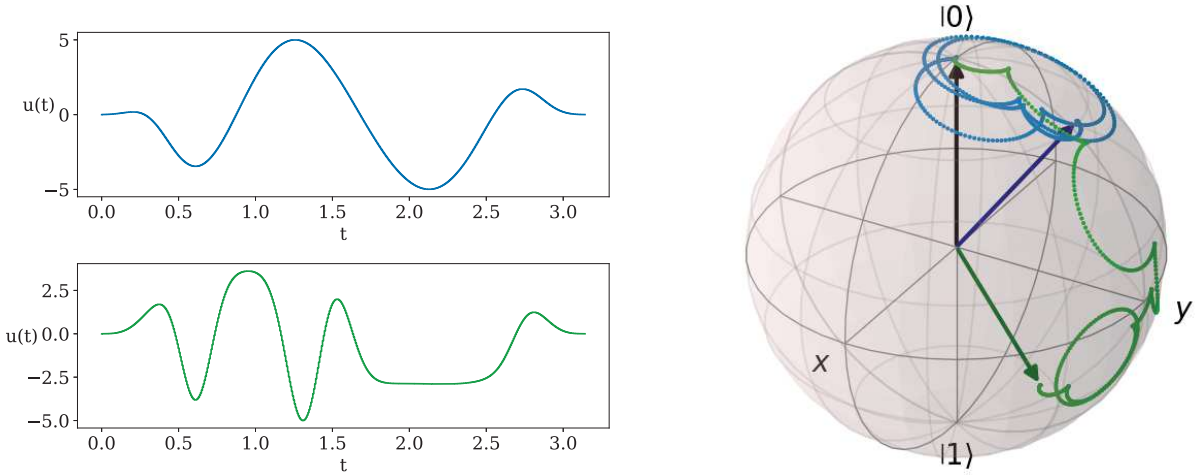

Figure 3.4: Best optimised pulses for the Sigmoid bases with width σ_1 (blue) and σ_2 (green) and their path on the Bloch sphere, for an optimisation with $T = \pi$, $u_{\max} = 5\Omega_\pi$ and $\Delta = 10\Omega_\pi$. The black vector represents the initial state, and the blue and green vectors are the respective final states.

the maximum detuning value considered $\Delta = 10$, along with the evolution of the respective systems on the Bloch Sphere.

To better understand the behaviour of the Sigmoid basis, we conducted additional optimisations with increased numbers of superiterations for $\Delta = 5$, $\sigma = \sigma_1$, yet observed no significant improvement. This suggests that the issue is not related to encountering false traps in the optimisation process. Increasing the number of basis vectors leads to better results, however, it also extends the convergence time considerately. For a total number of optimised parameters $N_p = 10$, the median infidelity approaches that of the Fourier basis for $N_p = 4$.

3.2.3 Robustness for ensemble of qubits with different detunings

In the previous section, the control pulse was optimised for singular values of detunings. This approach only guarantees optimal state transfer when the detuning coincides with the value considered in the optimisation. However, in reality, quantum ensembles of qubits can exhibit variations in the parameters that characterise the dynamics of the systems, including detuning. Since it is generally impractical to employ different control pulses for individual qubits, a pulse that is robust against

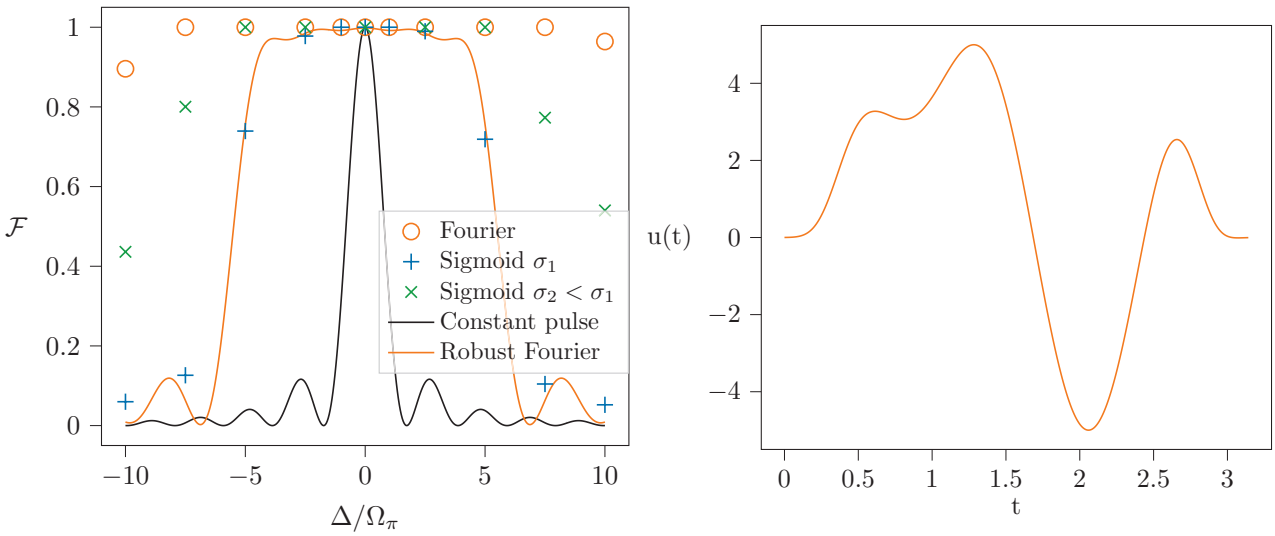

Figure 3.5: On the left, comparison of the median fidelity over 100 optimisations for the three pulses and calculated fidelity for the constant pulse and best optimised robust pulse. On the right, the shape of the best optimised robust pulse over 20 optimisations, run with $T = \pi$, $u_{\max} = 5\Omega_\pi$, $M = 10$, and $\Delta_{\max} = 5$.

detuning is necessary to ensure optimal fidelity for the entire ensemble.

To address the variability in detuning across a quantum ensemble, we adopt a robust optimisation strategy. Following Ref. [69], we assume the detunings to follow a Gaussian distribution centred at zero with standard deviation σ_Δ. The probability density function for a single qubit is given by:

$$p(\Delta_i) = \frac{1}{\sqrt{2\pi}\sigma_\Delta} e^{-\frac{\Delta_i^2}{2\sigma_\Delta^2}}. \tag{3.10}$$

The goal is to optimise a robust cost function J that averages over the infidelity contributions from different detuning values within a specified range:

$$J_{\text{rob}} = 1 - \int d\Delta p(\Delta) |\langle \psi_T | \psi(T, \Delta) \rangle|^2, \tag{3.11}$$

where $|\psi_T\rangle$ is the target state and $|\psi(T, \Delta)\rangle$ is the state evolved under the control pulse $u(t)$ for a time T with detuning Δ. Including the probability distribution $p(\Delta_i)$ ensures the optimisation to favour smaller values of detuning, as they are most likely to occur.

We consider around 98% of the distribution imposing the full width at half maximum (FWHM) as the limit for the detuning values, $|\Delta_i| \leq \Delta_{\max} = 2\sqrt{2\ln 2}\,\sigma_\Delta$.

To simplify computational complexity, we discretise the Gaussian distribution into M evenly spaced detunings values within the considered range. The cost function of the optimisation process is therefore:

$$J_{\text{rob},M} = 1 - \frac{1}{\mathcal{N}} \sum_{i=1}^{M} p(\Delta_i) \mathcal{F}(\Delta_i), \quad \mathcal{N} = \sum_{i=1}^{M} p(\Delta_i), \tag{3.12}$$

Where $\mathcal{F}(\Delta_i)$ represents the fidelity for the detuning Δ_i and $\mathcal{N}$ normalises the probability weights.

We perform 20 optimisations for 10 detuning values in the range $\Delta_i \in [-5, 5]$ employing the Fourier basis, which has proved to be the most efficient among the considered bases. The best optimised pulse is then used to calculate the fidelity over the broader range $\Delta_i \in [-10, 10]$ for comparison with the previous results, shown in Fig. 3.5 together with the pulse shape.

Conclusion and outlook

In this thesis, we applied quantum optimal control to state transfers in the presence of detuning, comparing different truncated bases. We started by establishing the theoretical background for formulating quantum optimal control problems including its inherent limitations and the potential constraints on the optimised pulses. An overview of the main numerical methods employed in QOC was provided, with a particular focus on the dCRAB algorithm. We described two truncated bases, the Fourier and Sigmoid bases, which were applied in the subsequent optimisation. Following a discussion of the employed methods, we compared their efficacy. Here, we briefly introduced NV centres in diamond and the concept of detuning, defining the considered control problem. We analytically determined the ideal amplitude for a constant pulse and optimised time-dependent pulses for increasing detuning values. In our study, we compared the performance of the Fourier basis with two versions of the Sigmoid basis: one with a similar frequency range and one with a broader spectrum. In the absence of detuning both Sigmoid bases showed a smoother control landscape, consistently achieving optimal solutions at the first superiteration. In the presence of detuning, the Fourier basis outperformed both Sigmoid bases. Here, the Sigmoid with a narrower spectrum could only achieve optimal solutions over a limited detuning range compared to the other bases. This outcome, combined with the observed pulse shapes, suggests that Fourier's ability to effectively exploit the frequency space contributed to its better performance. Finally, we employed the Fourier basis to optimise a pulse robust against detuning, i.e. a pulse that would faithfully transfer the state for a range of detunings. This could be applied to an ensemble of qubits with different detunings, for instance in the presence of an inhomogeneous magnetic field.

This study confirms the importance of choosing an appropriate truncated basis to enhance the efficacy of the quantum optimal control algorithm to solve a given problem. We discussed a simple problem involving a single qubit. Recent results indicate that, for specific problems, the relative performance of the bases remains consistent as the number of qubits increases within the same problem class. Therefore testing simplified versions of the problem can help determine the most suitable basis to employ in optimisations involving multiple qubits [25].

While our focus was on pure states and closed systems, the generalisation to mixed states could be done by using density matrices and Lindbladians. Addressing open quantum systems efficiently is a more complex and active area of research, starting from their controllability [70]. Further analysis of the control landscape could provide deeper insight into the convergence behaviour of the basis, guiding the selection of other bases to test. Future work could improve the control over the considered system by adding a σ_y component on the pulse as in Ref. [69], or explore the application of this approach to different optimal control problems and different types of qubits, such as superconducting qubits [71]. The optimisation of state transfer can be extended to produce optimal single qubit gates by replacing the state fidelity with the gate fidelity [37].

In summary, we conclude that the choice of bases affects the convergence speed and quality of the results. Choosing a basis with the right properties can lead to smoother landscapes or optimal solutions even in smaller search spaces, leading to a speed-up in the convergence process and better figures of merit. Quantum optimal control is a valuable tool, further research in this field will continue to advance the capabilities and applications of quantum technologies.

Bibliography

[1] M. Planck, "Über das gesetz der energieverteilung im normalspectrum", Annalen der Physik **309**, 553–563 (1901).

[2] A. Einstein, "Über einen die erzeugung und verwandlung des lichtes betreffenden heuristischen gesichtspunkt", Annalen der Physik **322**, 132–148 (1905).

[3] W. Heisenberg, "Über quantentheoretische Umdeutung kinematischer und mechanischer Beziehungen.", Zeitschrift für Physik **33**, 879–893 (1925).

[4] E. Schrödinger, "An undulatory theory of the mechanics of atoms and molecules", Physical Review **28**, 1049–1070 (1926).

[5] J. v. Neumann, *Mathematical foundations of quantum mechanics* (Princeton University Press, 1955), 462 pp.

[6] 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 Technology **9**, 1–60 (2022).

[7] L. K. Grover, *A fast quantum mechanical algorithm for database search*, Nov. 19, 1996.

[8] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer", SIAM Journal on Computing **26**, 1484–1509 (1997).

[9] S.-K. Liao, W.-Q. Cai, J. Handsteiner, B. Liu, J. Yin, L. Zhang, D. Rauch, M. Fink, J.-G. Ren, W.-Y. Liu, Y. Li, Q. Shen, Y. Cao, F.-Z. Li, J.-F. Wang, Y.-M. Huang, L. Deng, T. Xi, L. Ma, T. Hu, L. Li, N.-L. Liu, F. Koidl, P. Wang, Y.-A. Chen, X.-B. Wang, M. Steindorfer, G. Kirchner, C.-Y. Lu, R. Shu, R. Ursin, T. Scheidl, C.-Z. Peng, J.-Y. Wang, A. Zeilinger, and J.-W. Pan, "Satellite-relayed intercontinental quantum network", Physical Review Letters **120**, 030501 (2018).

[10] V. Mannalatha, S. Mishra, and A. Pathak, "A comprehensive review of quantum random number generators: concepts, classification and the origin of randomness", Quantum Information Processing **22**, 439 (2023).

[11] M. P. Woods, R. Silva, G. Pütz, S. Stupar, and R. Renner, "Quantum clocks are more accurate than classical ones", PRX Quantum **3**, Publisher: American Physical Society, 010319 (2022).

[12] G. Balasubramanian, I. Y. Chan, R. Kolesov, M. Al-Hmoud, J. Tisler, C. Shin, C. Kim, A. Wojcik, P. R. Hemmer, A. Krueger, T. Hanke, A. Leitenstorfer, R. Bratschitsch, F. Jelezko, and J. Wrachtrup, "Nanoscale imaging magnetometry with diamond spins under ambient conditions", Nature **455**, 648–651 (2008).

[13] C. L. Degen, F. Reinhard, and P. Cappellaro, "Quantum sensing", Reviews of Modern Physics **89**, 035002 (2017).

[14] A. P. Peirce, M. A. Dahleh, and H. Rabitz, "Optimal control of quantum-mechanical systems: existence, numerical approximation, and applications", Physical Review A **37**, 4950–4964 (1988).

[15] N. Oshnik, P. Rembold, T. Calarco, S. Montangero, E. Neu, and M. M. Müller, "Robust magnetometry with single nitrogen-vacancy centers via two-step optimization", Physical Review A **106**, 013107 (2022).

[16] A. Carlini, A. Hosoya, T. Koike, and Y. Okudaira, "Time-optimal quantum evolution", Physical Review Letters **96**, 060503 (2006).

[17] 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", Physical Review X **10**, Publisher: American Physical Society, 021058 (2020).

[18] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi, J. Cui, M. Rossignolo, P. Rembold, S. Montangero, T. Calarco, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, "Generation and manipulation of schrödinger cat states in rydberg atom arrays", Science **365**, Publisher: American Association for the Advancement of Science, 570–574 (2019).

[19] T. Caneva, A. Silva, R. Fazio, S. Lloyd, T. Calarco, and S. Montangero, "Complexity of controlling quantum many-body dynamics", Physical Review A **89**, 042322 (2014).

[20] U. Boscain, M. Sigalotti, and D. Sugny, "Introduction to the pontryagin maximum principle for quantum optimal control", PRX Quantum **2**, Publisher: American Physical Society, 030203 (2021).

[21] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, "Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework", Physical Review A **84**, Publisher: American Physical Society, 022305 (2011).

[22] T. Caneva, T. Calarco, and S. Montangero, "Chopped random-basis quantum optimization", Physical Review A **84**, 022326 (2011).

[23] M. M. Müller, R. S. Said, F. Jelezko, T. Calarco, and S. Montangero, "One decade of quantum optimal control in the chopped random basis", Reports on Progress in Physics **85**, 076001 (2022).

[24] N. Rach, M. M. Müller, T. Calarco, and S. Montangero, "Dressing the chopped-random-basis optimization: a bandwidth-limited access to the trap-free landscape", Physical Review A **92**, 062343 (2015).

[25] A. Pagano, M. M. Müller, T. Calarco, S. Montangero, and P. Rembold, *The role of bases in quantum optimal control*, May 31, 2024.

[26] E. Fernandez-Car and E. Zuazua, "Control theory: history, mathematical achievements and perspectives", (2003).

[27] E. R. Pinch, *Optimal control and the calculus of variations* (Oxford University Press, Oxford, New York, Sept. 7, 1995), 242 pp.

[28] L. N. Hand and J. D. Finch, *Analytical mechanics* (Cambridge University Press, Nov. 13, 1998), 598 pp.

[29] J. Bernoulli, *Acta eruditorum* (Christoph Günther, 1696), 638 pp.

[30] W. Kalfus, D. Lee, G. Ribeill, S. Fallek, A. Wagner, B. Donovan, D. Riste, and T. Ohki, "High-fidelity control of superconducting qubits using direct microwave synthesis in higher nyquist zones", IEEE Transactions on Quantum Engineering **1**, 1–12 (2020).

[31] C. Figgatt, A. Ostrander, N. M. Linke, K. A. Landsman, D. Zhu, D. Maslov, and C. Monroe, "Parallel entangling operations on a universal ion-trap quantum computer", Nature **572**, Publisher: Nature Publishing Group, 368–372 (2019).

[32] G. Zhang, F. Schilling, S. J. Glaser, and C. Hilty, "Reaction monitoring using hyperpolarized NMR with scaling of heteronuclear couplings by optimal tracking", Journal of Magnetic Resonance **272**, 123–128 (2016).

[33] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, "Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework", Physical Review A **84**, 022305 (2011).

[34] R. Cao, "Stengel r.f. - optimal control and estimation", (1994).

[35] J. P. Palao and R. Kosloff, "Quantum computing by an optimal control algorithm for unitary transformations", Physical Review Letters **89**, 188301 (2002).

[36] S. L. Braunstein and P. van Loock, "Quantum information with continuous variables", Reviews of Modern Physics **77**, 513–577 (2005).

[37] A. Carlini, A. Hosoya, T. Koike, and Y. Okudaira, "Time-optimal unitary operations", Physical Review A **75**, 042308 (2007).

[38] C. Brif, R. Chakrabarti, and H. Rabitz, "Control of quantum phenomena: past, present and future", New Journal of Physics **12**, 075008 (2010).

[39] F. Albertini and D. D'Alessandro, "Notions of controllability for bilinear multilevel quantum systems", IEEE Transactions on Automatic Control **48**, 1399–1403 (2003).

[40] S. G. Schirmer, H. Fu, and A. I. Solomon, "Complete controllability of quantum systems", Physical Review A **63**, 063410 (2001).

[41] P. Rembold, N. Oshnik, M. M. Müller, S. Montangero, T. Calarco, and E. Neu, "Introduction to quantum optimal control for quantum sensing with nitrogen-vacancy centers in diamond", AVS Quantum Science **2**, 024701 (2020).

[42] T.-S. Ho and H. Rabitz, "Why do effective quantum controls appear easy to find?", Journal of Photochemistry and Photobiology A: Chemistry, Coherent Control of Photochemical and Photobiological Processes **180**, 226–240 (2006).

[43] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, "Tunable, flexible, and efficient optimization of control pulses for practical qubits", Physical Review Letters **120**, 150401 (2018).

[44] Z. Hui, Z. Guiliang, and D. Wang, "Super-gaussian window function and its applications", in China., 1991 international conference on circuits and systems (June 1991), 595–598 vol.2.

[45] S. Deffner and S. Campbell, "Quantum speed limits: from heisenberg's uncertainty principle to optimal quantum control", Journal of Physics A: Mathematical and Theoretical **50**, 453001 (2017).

[46] L. Mandelstam and I. Tamm, "The uncertainty relation between energy and time in non-relativistic quantum mechanics", in *Selected papers*, edited by I. E. Tamm, B. M. Bolotovskii, V. Y. Frenkel, and R. Peierls (Springer, Berlin, Heidelberg, 1991), pp. 115–123.

[47] N. Margolus and L. B. Levitin, "The maximum speed of dynamical evolution", Physica D: Nonlinear Phenomena, Proceedings of the Fourth Workshop on Physics and Consumption **120**, 188–195 (1998).

[48] L. B. Levitin and T. Toffoli, "Fundamental limit on the rate of quantum dynamics: the unified bound is tight", Physical Review Letters **103**, 160502 (2009).

[49] T. Caneva, M. Murphy, T. Calarco, R. Fazio, S. Montangero, V. Giovannetti, and G. E. Santoro, "Optimal control at the quantum speed limit", Physical Review Letters **103**, 240501 (2009).

[50] K. Bhattacharyya, "Quantum decay and the mandelstam-tamm-energy inequality", Journal of Physics A: Mathematical and General **16**, 2993 (1983).

[51] S. Lloyd and S. Montangero, "Information theoretical analysis of quantum optimal control", Physical Review Letters **113**, 010502 (2014).

[52] D. Basilewitsch, Y. Zhang, S. M. Girvin, and C. P. Koch, "Engineering strong beamsplitter interaction between bosonic modes via quantum optimal control theory", Physical Review Research **4**, Publisher: American Physical Society, 023054 (2022).

[53] S. Rosi, A. Bernard, N. Fabbri, L. Fallani, C. Fort, M. Inguscio, T. Calarco, and S. Montangero, "Fast closed-loop optimal control of ultracold atoms in an optical lattice", Physical Review A **88**, 021601 (2013).

[54] D. J. Egger and F. K. Wilhelm, "Adaptive hybrid optimal quantum control for imprecisely characterized systems", Physical Review Letters **112**, 240503 (2014).

[55] D. M. Reich, M. Ndong, and C. P. Koch, "Monotonically convergent optimization in quantum control using krotov's method", The Journal of Chemical Physics **136**, 104103 (2012).

[56] 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).

[57] P. Doria, T. Calarco, and S. Montangero, "Optimal control technique for many-body quantum dynamics", Physical Review Letters **106**, 190501 (2011).

[58] R. M. Lewis, V. Torczon, and M. W. Trosset, "Direct search methods: then and now", Journal of Computational and Applied Mathematics, Numerical Analysis 2000. Vol. IV: Optimization and Nonlinear Equations **124**, 191–207 (2000).

[59] J. A. Nelder and R. Mead, "A simplex method for function minimization", The Computer Journal **7**, 308–313 (1965).

[60] P. Rembold, "Quantum optimal control of spin systems and trapped atoms", PhD thesis (University of Cologne, 2022).

[61] M. O. Rayes, V. Trevisan, and P. S. Wang, "Factorization properties of chebyshev polynomials", Computers & Mathematics with Applications **50**, 1231–1240 (2005).

[62] J. H. M. Jensen, F. S. Møller, J. J. Sørensen, and J. F. Sherson, "Achieving fast high-fidelity optimal control of many-body quantum dynamics", Physical Review A **104**, 052210 (2021).

[63] M. Rossignolo, T. Reisser, A. Marshall, P. Rembold, A. Pagano, P. J. Vetter, R. S. Said, M. M. Müller, F. Motzoi, T. Calarco, F. Jelezko, and S. Montangero, "QuOCS: the quantum optimal control suite", Computer Physics Communications **291**, 108782 (2023).

[64] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. L. Hollenberg, "The nitrogen-vacancy colour centre in diamond", Physics Reports, The nitrogen-vacancy colour centre in diamond **528**, 1–45 (2013).

[65] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information: 10th anniversary edition*, Higher Education from Cambridge University Press, (Dec. 9, 2010)

[66] Smite-Meister, *Bloch sphere, a geometrical representation of a two-level quantum system.* Jan. 30, 2009.

[67] F. Jelezko and J. Wrachtrup, "Single defect centres in diamond: a review", physica status solidi (a) **203**, 3207–3225 (2006).

[68] D. J. Griffiths and D. F. Schroeter, *Introduction to quantum mechanics* (Cambridge University Press, Aug. 16, 2018), 511 pp.

[69] J. Tian, H. Liu, Y. Liu, P. Yang, R. Betzholz, R. S. Said, F. Jelezko, and J. Cai, "Quantum optimal control using phase-modulated driving fields", Physical Review A **102**, 043707 (2020).

[70] C. P. Koch, "Controlling open quantum systems: tools, achievements, and limitations", Journal of Physics: Condensed Matter **28**, 213001 (2016).

[71] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "A quantum engineer's guide to superconducting qubits", Applied Physics Reviews **6**, 021318 (2019).