

UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Fisica e Astronomia “Galileo Galilei”

Master’s Degree in Physics of Data

Final dissertation

Closed-Loop Gate Optimisation via Efficient Two-Qubit Tomography for NV Centers

Supervisor

Prof. Simone Montangero

Dr. Matthias Müller

Student

Alessandro Marcomini

Academic Year 2021/2022

Abstract

While moving towards the realisation of commercial quantum computers it has become of paramount importance to be able to implement high fidelity quantum gates. As of today, this is a technical challenge for all hardware platforms because of the difficulty to create theoretical models which reach the level of detail required for designing of high precision gates. To overcome this hurdle, closed-loop quantum optimal control (QOC) provides an excellent tool for the optimization of control pulses by exploiting experimental feedback which does not require modeling and characterisation.

In this work we propose a new protocol for the implementation of an experimentally-convenient figure of merit to perform closed-loop QOC for two-qubit gates. We examine its performances on nitrogen vacancy (NV) centers in diamonds, which have raised as one of the most promising hardware proposals for quantum technologies because of their relatively long decoherence time and their performances at room temperature. In particular, our approach has been tested numerically for the realization of entangling gates for a $^{13}\text{C-NV}^-$ system, with special care in the definition of the pulse duration and the control amplitudes according to typical experimental parameters. Numerical simulations using the dressed Chopped RANdom Basis (dCRAB) algorithm suggest the possibility to build high fidelity CNOT gates within few hundreds of closed-loop optimization iterations, with a significant reduction in the number of required readout operations compared to the state-of-the-art.

Preface

While looking for suitable hardware solutions for the realisation of quantum information processing (QIP) units, past years have made clear that good quantum systems must satisfy a specific set of requirements. Aside from the DiVincenzo criteria [1], the quantum bits must be independently addressable and benefit of long decoherence times in order to allow for controllability and readout. To date, suitable systems are recognised to be superconducting (transmon) qubits [2], trapped atoms and ions [3, 4], photonic circuits [5] and defect centers in diamonds [6]. In particular, the latter consist of a family of technologies based on natural or artificial impurities in Carbon crystals which usually display the aforementioned properties because of their low electron-phonon coupling and their stability (i.e., lack of decoherence) over a wide range of thermal conditions, from cryogenic to room temperature [7]. The choice of diamond as hosting material for the sites allows to exploit its mechanical properties, such as lattice stability, good electrical insulation and thermal conductivity. While diamonds are commonly transparent to light at multiple wavelengths, the presence of these defects can enhance different optical responses which are immensely useful to have clear readouts. Because of this, these sites take also the name of colour centers [8].

While the properties of NVs are undoubtedly good for quantum technologies, concerns about their scalability and their tedious fabrication techniques have kept them so far as a low-profile alternative for quantum computing [9]. Nevertheless, promising research is being carried out to try to overcome these limitations and eventually enable quantum computing with NV centers on a large scale [10]. In this perspective, a new tool to perform high quality quantum state manipulation on the spin states of the NVs is required. Since these centers arise as natural defects in the diamond structure they present individual properties that are different from site to site. In these conditions, a precise modeling of the system to design controlled state evolution becomes a non-trivial challenge.

A solution to this problem is offered by closed-loop quantum optimal control: by performing trial control of the time evolution of the system of interest and properly scoring the results we can learn the best actions to be performed in an iterative way [11]. The goal of this project is therefore the definition of a suitable metric (or "figure of merit") which well combines with the experimental routine to enable and make efficient the feedback process of a closed-loop optimisation for the specific case of two-qubit gates. This procedure has the potential to become a standard to

be applied to all the NVs in the crystal: in this way it will be possible to design tailored quantum gates for each two-qubit system without the necessity of carefully modeling each of them.

Our work is organised as follows: in chapter 1 we review the basics of quantum computing (section 1.1), we describe the experimental details of NV centers (section 1.2) and we introduce the concept of quantum control theory together with the main optimisation algorithms (sections 1.4.3 and 1.5). Chapter 2 contains the full description and analysis of the theoretical protocol we propose to enable efficient control optimisation. In chapter 3 we report simulation details and results that contribute to the validation of our approach. Finally, chapter 4 summarises the fundamental results of this work and introduces the remaining steps to take in order to move to an experimental realisation of the protocol.

This work has been carried out in collaboration with the University of Padova, Italy, and the PGI-8 Institute for Quantum Control at Forschungszentrum Jülich, Germany.

1222·2022
800
ANNI



UNIVERSITÀ
DEGLI STUDI
DI PADOVA



Contents

1	Introduction and theoretical background	1
1.1	Fundamentals of quantum computing	1
1.1.1	Qubits and state representation	2
1.1.2	The Bloch sphere	5
1.1.3	Quantum gates	6
1.1.4	DiVincenzo criteria	10
1.1.5	Time evolution operator	11
1.2	NV centers in diamonds	13
1.2.1	Optical readout	16
1.2.2	Rabi oscillations and optically engineered quantum gates .	18
1.3	System specifics	20
1.4	Control Theory	23
1.4.1	Classical Control Theory	23
1.4.2	Classification of control models	25
1.4.3	Quantum Control	26
1.4.4	Fidelity and Figure of Merit	27
1.5	Optimal control algorithms	30
1.5.1	Gradient Ascend Pulse Engineering (GRAPE)	30
1.5.2	Krotov's method	31
1.5.3	Chopped random basis (CRAB) algorithm	31
1.5.4	The QuOCS library	33
2	Efficient gate fidelity protocol	35
2.1	Application to CNOT and SWAP gates	36
2.1.1	Definition of the best $\hat{\rho}_1$ form	37
2.1.2	Experimental optimisation of the action on the basis . . .	38
2.1.3	Alternative approaches: process tomography and the two- $\hat{\rho}_1$ variant	42
2.1.4	Experimental optimisation of the relative phases	43

3	Numerical validation: simulations	47
3.1	Simulation design	48
3.1.1	States initialisation	48
3.1.2	Time evolution operator and controls	49
3.1.3	Measurement	49
3.1.4	Error propagation	53
3.2	Results	55
3.2.1	Definition of controls and Hamiltonian	55
3.2.2	Reducing the complexity	58
3.2.3	The new Figure of Merit	66
3.2.4	Tuning the algorithm parameters	70
3.2.5	Comparison with the two- $\hat{\rho}_1$ variant	74
3.2.6	Comparison with non-selective e^- spin transitions	76
4	Summary and outlook	79
4.1	Towards the experiment	79

Chapter 1

Introduction and theoretical background

1.1. Fundamentals of quantum computing

In our daily life we have grown to think of technology as an omnipresent, silent partner which accompanies us in every action. From mobile phones to laptops, from autonomous cars to thermostats, we are surrounded by automatised processing units that collect and manipulate data to ease and better our lives.

Also in the field of research technology and computational power have recently been used to enlarge the power of analysis and simulations of natural behaviours in fields ranging from biology to medicine and physics. Focusing on the latter, something extremely peculiar happens while we try to move towards the "infinitesimally small": as discovered by Max Planck over a century ago, the small realm does not follow the classical rules of physics but requires a new description in terms of quantum mechanics. This leads scientists to a plethora of fascinating challenges to be dealt with, including the problem of representing and efficiently studying the dynamics of quantum many body systems.

In 1982, Richard Feynman showed that to simulate the dynamics of a quantum system on classical computer it will be necessary to scale it in an exponential way [12]. This puts an effective bound to the knowledge we can gather of such a system, according to the computational resources at disposal. Nevertheless, looking for a solution to this issue he proposed a completely new idea:

Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy.

Richard Feynman

Indeed, Feynman's idea was the possibility of setting up a "quantum machine", an hardware subject to the laws of quantum mechanics to simulate the behaviour of quantum systems. Moreover, if we are able to encode in such quantum systems a non-quantum task, we could in principle exploit the properties of quantum mechanics to enhance the research of a solution. Following this line of thought, in 1992 Lov Grover proposed an algorithm which is able to introduce a quadratic "quantum advantage" in solving the task of searching for an element in an unstructured database using quantum superpositions [13]. Three years later, a protocol for the efficient factorisation of large numbers with exponential speedup was suggested by Peter Shor, highlighting a possible treat to state-of-the-art cryptographic schemes [14].

As of today, the importance and the utilities of "quantum computers" is widely recognised. In order to build a fully-working machine, some elements need to be designed:

- Fundamental computational units and conventional operations on them
- An efficient and stable quantum hardware
- Software component for the encoding
- A control mechanism that turns software instructions into physical operations on the hardware

In the following section 1.1.1 we introduce the concept of quantum bits and their fundamental representation. Section 1.1.4 describes the requirements for a sustainable quantum hardware, while in section 1.2 we review the details of the specific platform we analyse in detail. Generally speaking, there are no stringent requirements regarding the software: everything reduces to the transpiler module, which is responsible of decomposing the user instructions into a combination of fundamental logic operations (i.e. quantum gates, review in section 1.1.3). Finally, the physical realisation of quantum gates is the deeper focus of this work: the details of the time evolution of quantum states is reported in section 1.1.5, the theory of quantum optimal control is described in section 1.4.3 and the specific proposal for this platform is fully characterised in section 2.

1.1.1. Qubits and state representation

In classical computation theory the fundamental logical units are the bits, i.e. the smallest possible information blocks. A classical bit can assume only two values, 0

and 1 (or "yes" and "no") according to the system's state. In practical realisations bits are often seen as switches (e.g. transistors) which allow or do not allow the flow of electrical current.

As any binary object can in principle represent a classical bit, for analogy we refer to quantum bits (or qubits) speaking about any two level quantum system where we can identify a $|0\rangle$ and a $|1\rangle$ state. Examples are the quantum states associated to fermionic spin states along a direction \hat{z} ($|0\rangle \equiv |\uparrow\rangle$, $|1\rangle \equiv |\downarrow\rangle$), to the polarisation of photons ($|0\rangle \equiv |H\rangle$, $|1\rangle \equiv |V\rangle$) or to an isolated subsystem of ground and excited energy levels ($|0\rangle \equiv |g\rangle$, $|1\rangle \equiv |e\rangle$).

The fundamental difference between bits and qubits lays in the superposition principle: while in classical mechanics the state of the system is either 0 or 1 at any time (and we can only have classical ignorance about it), in quantum mechanics at a given time t the system's state is defined by its quantum wavefunction:

$$|\psi(t)\rangle \equiv \alpha(t)|0\rangle + \beta(t)|1\rangle \equiv \alpha(t) \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \beta(t) \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (1.1)$$

i.e., a superposition of the states $\{|0\rangle, |1\rangle\}$ that form a basis for the Hilbert space that hosts the quantum state $|\psi\rangle$. For a single qubit, the Hilbert space is a 2-dimensional vector space; an n -body state lives in the Hilbert space given by the tensor product of the single-particle spaces, which means that for n qubits the state $|\psi\rangle$ is a vector of size 2^n (and we observe the exponential scaling in the complexity pointed out by Feynman).

The system's state description tells us that the most radical difference between classical and quantum bits is the possibility for the latter to be in the two basis states at the same time. The coefficients $\alpha, \beta \in \mathbb{C}$ are called amplitude probabilities for the basis states and for any time t they satisfy the relation:

$$|\alpha(t)|^2 + |\beta(t)|^2 = P_0(t) + P_1(t) = 1 \quad (1.2)$$

Upon measurement, the system's state will be found in either the $|0\rangle$ or the $|1\rangle$ state with (classical) probabilities given by P_i , $i = 0, 1$: this process is known as the collapse of the quantum wavefunction $|\psi\rangle$. The probability to find the system in a generic state $|\xi\rangle$ is found by the Born's rule as the modulus squared of the projection of the quantum state over $|\xi\rangle$:

$$P(|\xi\rangle) = |\langle \xi | \psi(t) \rangle|^2 \quad (1.3)$$

which is consistent with the previous statements since for $|\xi\rangle = |0\rangle$ it holds:

$$P_0 = P(|0\rangle) = |\langle 0 | \psi(t) \rangle|^2 = |\alpha|^2 \underbrace{|\langle 0 | 0 \rangle|^2}_{=1} + |\beta|^2 \underbrace{|\langle 0 | 1 \rangle|^2}_{=0} = |\alpha|^2 \quad (1.4)$$

and same for P_1 . The contractions between the $|0\rangle$ and $|1\rangle$ states follow the rules for an orthonormal basis.

An alternative representation of the system's state is given by the density matrix formalism: to each state $|\psi\rangle$ is associated a quantum operator defined as:

$$\hat{\rho}(t) \equiv |\psi(t)\rangle\langle\psi(t)| \quad (1.5)$$

which becomes, for a single qubit:

$$\hat{\rho} = |\psi\rangle\langle\psi| = \begin{pmatrix} |\alpha|^2 & \alpha\beta^* \\ \alpha^*\beta & |\beta|^2 \end{pmatrix} = \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{10} & \rho_{11} \end{pmatrix} \quad (1.6)$$

In a more general fashion, being $|\psi\rangle$ a vector of size N , $\hat{\rho}$ has size $N \times N$. It turns out to be a positive semi-definite, Hermitian operator of trace one acting on the Hilbert space of the system. The density operator formalism is especially convenient when dealing with mixed states, i.e. when there is classical uncertainty regarding the quantum state of the system [15]. In fact, assuming that the system can be described by a set of states $\{|\psi_j(t)\rangle\}_j$ with classical probabilities $\{q_j\}_j$, the density operator will be

$$\hat{\rho}(t) = \sum_j q_j |\psi_j(t)\rangle\langle\psi_j(t)| \quad (1.7)$$

Fundamental properties of the density matrix representation of the state are [15]:

- The diagonal elements of the density matrix (populations) are the probabilities associated to measure each basis state. As a result, it holds:

$$\text{Tr}(\hat{\rho}(t)) = 1 \quad \forall t \quad (1.8)$$

- It holds $0 \leq \text{Tr}(\hat{\rho}^2) \leq 1$ and

$$\text{Tr}(\hat{\rho}^2) = 1 \iff \hat{\rho} \text{ is a pure state} \quad (1.9)$$

- Non-zero off-diagonal terms of $\hat{\rho}$ (coherences) indicate the presence of a quantum behaviour: for a fully mixed state (classical mixture) coherences are all zero.
- The Born's rule for density matrices follows from equation (1.7):

$$P(|\xi\rangle) = \sum_j q_j |\langle\xi|\psi_j\rangle|^2 = \sum_j q_j \langle\psi_j|\xi\rangle \langle\xi|\psi_j\rangle = \text{Tr}(\hat{\Pi}_\xi \hat{\rho}) \quad (1.10)$$

being $\hat{\Pi}_\xi = |\xi\rangle\langle\xi|$. This is indeed the expectation value of the projector onto the $|\xi\rangle$ state in a density operator formalism.

Generally speaking, a non-isolated quantum system interacts with the environment over time. We call this an "open quantum system". This interaction leads to a gradual loss of the quantum properties of the system ("quantum dissipation")

and its evolution towards a classical mixture. Since this phenomenon implies the vanishing of off-diagonal terms in the density operators it takes the name of "decoherence".

In this scenario, density matrices are the most suitable tools to investigate the evolution of the quantum system A of interest as a component of the larger, closed quantum system $A + E$ which includes the environment. In fact, the density matrix of the reduced system can be found simply tracing out the complementary components as:

$$\hat{\rho}_A = \text{Tr}_E (\hat{\rho}_{AE}) \quad (1.11)$$

where Tr_E is defined as:

$$\text{Tr}_E (|\xi_u\rangle\langle\xi_v| \otimes |\chi_u\rangle\langle\chi_v|) \equiv |\xi_u\rangle\langle\xi_v| \text{Tr}_E (|\chi_u\rangle\langle\chi_v|) \quad (1.12)$$

being $|\xi_{u,v}\rangle$ arbitrary states for the system A , $|\chi_{u,v}\rangle$ arbitrary states for the system E and $\text{Tr}_E (|\chi_u\rangle\langle\chi_v|) = \langle\chi_u|\chi_v\rangle$ the standard trace operator [16]. This description is extremely convenient since in most real-world cases we deal with open quantum systems where a detailed description of the environment is either not convenient or not possible. Hence, we prefer to focus on the system under investigation.

All these properties, including the practical formalism for the time evolution of an open quantum system described in section 1.1.5, make the density matrix formalism a natural and convenient choice while studying quantum states for the sake of quantum computing.

1.1.2. The Bloch sphere

Equation (1.1) displays the basic form a qubit state in terms of a superposition of the basis states with complex amplitudes subject to the constraint in equation (1.2). An alternative representation of the state in the computational basis is given by the following expression:

$$|\psi\rangle \equiv \cos\left(\frac{\theta}{2}\right) |0\rangle + e^{i\varphi} \sin\left(\frac{\theta}{2}\right) |1\rangle \quad (1.13)$$

where it is exploited the fact that quantum states are defined up to a global phase. Considering now the transformations defined by:

$$\vec{r} = \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} \alpha\beta^* + \alpha^*\beta \\ i(\alpha\beta^* - \alpha^*\beta) \\ |\alpha|^2 - |\beta|^2 \end{pmatrix} = \begin{pmatrix} r \sin(\theta) \cos(\varphi) \\ r \sin(\theta) \sin(\varphi) \\ r \cos(\theta) \end{pmatrix} \quad (1.14)$$

we find a relation between the complex amplitude probabilities $\{\alpha, \beta\}$ and the three real spherical coordinates (r, θ, φ) which uniquely define a point in a sphere of unitary radius. This sphere takes the name of Bloch sphere and is the main visualisation tool for single-qubit states.




Figure 1.1: Example states on the Bloch sphere, corresponding to the X , Y and Z basis states, respectively. As a convention the $|0\rangle$ and $|1\rangle$ states are chosen as the basis states on the Z (vertical) axis. Figure from [17].

Figure 1.1 shows some example states as points on the surface of the Bloch sphere. As a convention, the basis states are assumed to lay on the \hat{z} axis. Any pair of diametrically opposite states forms a basis for the single qubit state. Since:

$$r = \sqrt{\text{Tr}(\hat{\rho}^2)} \quad (1.15)$$

points on the surface of the sphere represent pure states, while points inside the sphere represent mixed states. Finally, the dynamical evolution in time of the qubit state can be seen as a trajectory over the Bloch sphere (where the unitary evolution of pure states remains on the surface of the sphere). An example of a single-qubit state moving from $|0\rangle$ to $|1\rangle$ by means of unitary evolution is reported in figure 1.2.

1.1.3. Quantum gates

In classical computing theory every operation we require our machines to perform is translated into a sequence of fundamental logical operations that are applied to classical bits. Example of such operations are the AND, OR, XOR, NAND gates, and many more. Similarly, in quantum computing we require the possibility to perform some basic operations on qubits that alter their state in a precise way. From a mathematical perspective, these gates are defined as operators (matrices).

Quantum gates are divided in two main classes: single- and multi-qubit gates.

Single qubit gates

In analogy with the classical NOT gate, the quantum version of a bit-flip operator is given by:

$$\hat{X} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = NOT \quad (1.16)$$




Figure 1.2: Quantum trajectory of evolution from the $|0\rangle$ to the $|1\rangle$ state on the surface of the Bloch sphere (π - pulse).

which takes the name "X" since it corresponds to the $\hat{\sigma}_x$ Pauli matrix. From the point of view of the Bloch sphere this corresponds to a rotation of an angle π around the \hat{x} axis and one can easily verify that:

$$|0\rangle \xrightarrow{\hat{X}} |1\rangle \text{ and } |1\rangle \xrightarrow{\hat{X}} |0\rangle \quad (1.17)$$

In a similar fashion, rotation around the \hat{y} and \hat{z} axes are given by the corresponding Pauli matrices:

$$\hat{Y} = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} \quad \hat{Z} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (1.18)$$

In particular, the latter operator has the property of flipping the relative phase between the $|0\rangle$ and $|1\rangle$ states:

$$\alpha |0\rangle + \beta |1\rangle \xrightarrow{\hat{Z}} \alpha |0\rangle - \beta |1\rangle \quad (1.19)$$

These three gates, together with the identity gate:

$$\mathbb{1}_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (1.20)$$

define a basis $\Sigma = \{\mathbb{1}_2, \hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z\}$ for any matrix $M \in \mathbb{C}^{2 \times 2}$.

An arbitrary state rotation of an angle ϕ around a direction \hat{n} in the Bloch sphere can be written as a composition of rotations along the three axes as:

$$R(\phi; \hat{n}) = \cos\left(\frac{\phi}{2}\right) \mathbb{1}_2 - i \sin\left(\frac{\phi}{2}\right) \hat{n} \cdot \vec{\sigma} \quad (1.21)$$

being $\vec{\sigma} = (\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z)^T$.

Other important single-qubit gates are:

$$\hat{H} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad \hat{S} = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\frac{\pi}{2}} \end{bmatrix} \quad \hat{T} = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\frac{\pi}{4}} \end{bmatrix} \quad (1.22)$$

where the Hadamard gate \hat{H} allows to move from the Z to the X basis:

$$|0\rangle \xrightarrow{\hat{H}} |+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \quad \text{and} \quad |1\rangle \xrightarrow{\hat{H}} |-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}} \quad (1.23)$$

while the \hat{S} and \hat{T} gates play the same role as the \hat{Z} gate, shifting the relative phase between the $|0\rangle$ and $|1\rangle$ states of a factor $\pi/2$ and $\pi/4$, respectively. For this reason the S gate is also called *Phase gate*. By convention these gates add a phase between the basis states. To remove the same phase it is sufficient to apply the inverse gates $S^{-1} = S^\dagger$ and $T^{-1} = T^\dagger$.

Multi-qubit gates

To perform quantum computation it is necessary to design a quantum register of multiple qubits and to exploit in full the properties of quantum mechanics (such as entanglement) we must be able to introduce quantum correlations among these quantum logical units. This cannot be done by means of single-qubit, local operations: we need to be able to act simultaneously on the joint state of two (or more) qubits to eventually modify the state of one conditionally upon the state of the other.

The most known example of such a gate is the Controlled-NOT (or *CNOT*) gate [18]. In the computational basis of a two-qubit register it takes the form:

$$CNOT = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad (1.24)$$

and it operates as:

$$\begin{aligned} |0\rangle_A \otimes |0\rangle_B &\xrightarrow{CNOT} |0\rangle_A \otimes |0\rangle_B \\ |0\rangle_A \otimes |1\rangle_B &\xrightarrow{CNOT} |0\rangle_A \otimes |1\rangle_B \\ |1\rangle_A \otimes |0\rangle_B &\xrightarrow{CNOT} |1\rangle_A \otimes |1\rangle_B \\ |1\rangle_A \otimes |1\rangle_B &\xrightarrow{CNOT} |1\rangle_A \otimes |0\rangle_B \end{aligned} \quad (1.25)$$

i.e., it always acts as the identity on the first (control) qubit A , while it acts as $\mathbb{1}_2$ or \hat{X} on the second qubit B whether the state of A is $|0\rangle_A$ or $|1\rangle_A$, respectively. For this reason the *CNOT* gate is also known as *CX* gate.

In a similar way, we can think of applying any arbitrary gate on the second qubit B conditioning to the state of a control, reference qubit A . All these gates take the same form of Controlled-Unitary (CU) gates:

$$CU = \begin{bmatrix} \mathbb{1}_2 & \mathbf{0}_2 \\ \mathbf{0}_2 & \mathbf{U} \end{bmatrix} \quad (1.26)$$

where U is an arbitrary single-qubit gate.

Another important gate is the so-called $SWAP$ gate, which inverts the roles of the two qubits:

$$SWAP = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (1.27)$$

in particular, we observe that it acts as the identity if $|\psi\rangle_A = |\psi\rangle_B$, else as a bit-flip. It can be seen as a $CNOT$ gate up to a redefinition of the states.

The concept of "operation on qubit B conditioned to the state of qubit A " can be generalised to the case of many control qubits. In fact we can condition the action on B to the state of n qubits A_1, \dots, A_n being all simultaneously 1, i.e. $|\psi\rangle_{A_i} = |1\rangle \forall i$. By applying this on the computational basis, this naturally applies to all the possible many-qubit states by means of the superposition principle.

The most common example is the so-called Toffoli gate, or "Controlled-Controlled-NOT" gate, which is a three-qubit gate in the form:

$$\text{Toffoli} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \quad (1.28)$$

the generalisation to n -controls is straightforward.

A schematic summary of the circuital elements associated to the main quantum gates is reported in figure 1.3.

Please notice that, as a fundamental requirement, quantum gates must be unitary operators (i.e. any gate \hat{V} must satisfy $\hat{V}^\dagger \hat{V} = \mathbb{1}_2$). This comes from the normalisation condition of the states, together with the fact that time evolution of quantum mechanical states is naturally unitary (as discussed in section 1.1.5) and therefore quantum gates must reflect this property, being "evolution operator" themselves. All the operators introduced in this section satisfy the aforementioned requirement.




Figure 1.3: Quantum gates for circuitual representation. Each line represents a qubit in the register. Input on the left hand-side. In multi-qubit gates control qubits are marked with a full dot. (a) Hadamard gate (b) X gate (c) Z gate (d) S (*Phase*) gate (e) inverse S gate (f) T gate (g) inverse T gate (h) $CNOT$ gate (i) CZ gate (j) $CPhase$ gate (k) $SWAP$ gate (l) Toffoli gate. Figure from [19]

1.1.4. DiVincenzo criteria

When it comes to quantum computing, a perfectly working quantum machine is expected to satisfied the so-called DiVincenzo's criteria, firstly theorised in [1]:

1. The physical system must be scalable with well-characterized qubits;
2. It must be possible to initialize the state of the qubits to a simple fiducial state;
3. Coherence times must be relevantly long;
4. It must be provided a universal set of quantum gates;
5. Measurements should be able to address single qubits.

Of course, all points 1-5 should be satisfied with high fidelity, i.e. with the smallest possible error probability.

As we will show in section 1.2, NV centers in diamonds intrinsically provide a suitable platform that satisfies points number 3 and 5. Moreover, the possibility to activate optical interaction among different sites would in principle allow to scale the system to a network of quantum registers (point 1). In section 1.2.1 we point out how the optical excitation-relaxation cycle allows to re-initialise the qubit state to the ground level with high fidelity (point 2). As a result, it will be possible to use NVs as a platform for quantum computing when we will manage to design with high fidelity a universal set of quantum gates. Such a universal set is constituted by a collection of single-qubit and two-qubit gates. A popular choice is the so-called "Clifford + T" gate set [20], defined as $\{CNOT, H, S, T\}$ (gates defined in the previous paragraphs).

As better discussed in sections 1.1.5 and 1.4.3, the dynamics of a quantum system is determined by its equation of motion. The latter depends on the system's Hamiltonian that we can control by means of external electromagnetic pulses.

Therefore, to proceed towards the experimental realisation of quantum computers it is necessary to optimise the system's controls in such a way that the time evolution operator defined by the Hamiltonian acts on the basis states as the target gate.

1.1.5. Time evolution operator

In the standard quantum mechanical formalism, the essence of a quantum system is represented by the quantum wavefunction $|\psi(t)\rangle$ which represents his state and lives in an Hilbert space \mathcal{H} . The dynamics of the quantum state is ruled by the Schrödinger equation:

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = \hat{H}(t) |\psi(t)\rangle \quad (1.29)$$

together with some boundary conditions $|\psi(t=t_0)\rangle = |\psi_0\rangle$, where \hat{H} is the Hamiltonian (hermitian operator) of the system.

As a result, it is straightforward to see that the evolution from a starting time $t=0$ to a final time t can be summarised as the action of a time evolution operator $\hat{U}(t)$ which depends on the time and on the Hamiltonian on the system. In many occasions the latter can be found to be time-independent and therefore it holds:

$$|\psi(t)\rangle = \hat{U}(t) |\psi_0\rangle \text{ where } \hat{U}(t) \equiv \exp\left(-\frac{i}{\hbar} \hat{H}t\right) \quad (1.30)$$

If instead the Hamiltonian is time-dependent, the formulation becomes more complex. Please notice how this will be our case, since we will be dealing with a control Hamiltonian with time-dependent electromagnetic pulses (details in section 1.3).

The most complete form of the time evolution operator relies on the introduction of Dyson series and of the time-ordered product of operators \mathcal{T} :

$$\mathcal{T} \left[\hat{H}(t_1) \hat{H}(t_2) \dots \hat{H}(t_n) \right] = \hat{H}(t_n) \dots \hat{H}(t_2) \hat{H}(t_1) \text{ for } t_1 \leq t_2 \leq \dots \leq t_n \quad (1.31)$$

which consists of making the operators act on each other in a time-subsequent way, recollecting that operators act on their right hand-side. In this formalism the time evolution operator becomes:

$$\hat{U}(t) = \mathcal{T} \exp\left(-\frac{i}{\hbar} \int_0^t \hat{H}(t) dt\right) \quad (1.32)$$

where:

$$\begin{aligned} \mathcal{T} \exp\left(-\frac{i}{\hbar} \int_0^t \hat{H}(t) dt\right) = \mathbb{1} + \\ \sum_{n=1}^{\infty} \left(-\frac{i}{\hbar}\right)^n \frac{1}{n!} \int_0^t dt_1 \int_0^{t_1} dt_2 \dots \int_0^{t_{n-1}} dt_n \mathcal{T} \left[\hat{H}(t_1) \hat{H}(t_2) \dots \hat{H}(t_n) \right] \end{aligned} \quad (1.33)$$

While this complex formula can be extremely challenging for analytical analysis, for the sake of numerical simulations some useful tricks can be adopted. First of all, while implementing computational time evolution from $t = 0$ to $t = T$, the total time is split into N steps of size $\Delta t = T/N$ and therefore the Hamiltonian itself is calculated over a finite-size set of timesteps $\{t_j\}_{j=1,\dots,N}$ being $t_j \equiv j\Delta t$. The time evolution operator has the property that composite and direct time evolution must return the same final state:

$$\hat{U}(t_1, t_0) = \hat{U}(t_1, t_2) \hat{U}(t_2, t_0) \quad (1.34)$$

where we specified the initial t_i and final t_f times in $\hat{U}(t_f, t_i)$. Therefore, over the discretisation set it must hold:

$$\hat{U}(T, 0) = \prod_{j=1}^N \hat{U}(t_j, t_{j-1}) \quad (1.35)$$

by paying attention to the fact that the product of operators is build by placing the j -th term on the right of the $(j-1)$ -th (i.e. in a time-ordered way).

If the time step Δt is chosen small enough, the Hamiltonian $\hat{H}(t_j)$ at each time can be thought to be locally time-independent. In this case, the time evolution operator over the small amount of time is given by equation (1.30) and we get:

$$\hat{U}(T, 0) = \prod_{j=1}^N \exp\left(-\frac{i}{\hbar} \hat{H}(t_j) \Delta t\right) \quad (1.36)$$

Further analysis can prove that if $[\hat{H}(t_j), \hat{H}(t_k)] = 0 \forall j, k$ then:

$$\hat{U}(T, 0) = \exp\left(-\frac{i}{\hbar} \sum_{j=1}^N \hat{H}(t_j) \Delta t\right) \quad (1.37)$$

which can be rewritten in the continuous case replacing the sum with the integral from 0 to T . However, since this requirement is not necessarily true and the computational cost of building the time evolution operator is still manageable for relatively large values of N , in this work we compute \hat{U} using the approach of equation (1.36). The theory behind this analysis is fully discussed in [21].

In quantum mechanics it is sometimes more convenient to describe the state of the system by means of the density operator $\hat{\rho}(t)$, instead of the quantum wavefunction. In this case, the state dynamics follows the Liouville - von Neumann equation:

$$i\hbar \frac{d}{dt} \hat{\rho}(t) = [\hat{H}(t), \hat{\rho}(t)] \quad (1.38)$$

With similar reasoning to the previous pages, one can retrieve that an initial state $\hat{\rho}_0$ can be evolved to time T as:

$$\hat{\rho}(T) = \hat{U}(T) \hat{\rho}_0 \hat{U}^\dagger(T) \quad (1.39)$$

being \hat{U} the same time evolution operator as above.

In section 1.1.1 we mentioned the potential of density operator formalism while studying open quantum systems. In particular, we stressed the useful description of these systems in the form of reduced density matrices. When the environment is too difficult or too large to model, it is convenient to study the evolution in time of the reduced density matrix $\hat{\rho}_A$. In a general perspective, the total density matrix $\hat{\rho}_{AE}$ evolves according to the Liouville - von Neumann equation. An analog of equation (1.38) for the reduced matrix of a larger system is given by the Lindblad master equation (or Lindbladian), whose diagonal form is:

$$\dot{\hat{\rho}} = -\frac{i}{\hbar} [\hat{H}, \hat{\rho}] + \sum_{i=1}^{N^2-1} \gamma_i \left(\hat{L}_i \hat{\rho} \hat{L}_i^\dagger - \frac{1}{2} \{ \hat{L}_i^\dagger \hat{L}_i, \hat{\rho} \} \right) \equiv \mathcal{L} \hat{\rho} \quad (1.40)$$

where we set $\hat{\rho} \equiv \hat{\rho}_A$. N is the size of the A system and \hat{H} its Hamiltonian. $\{ \hat{L}_i \}_i$ is the set of Lindblad (or "jump") operators and defines an orthonormal basis for the Hilbert-Schmidt operators on the system's Hilbert space. The coefficients $\{ \gamma_i \}_i$, together with the Hamiltonian, define the reduced system dynamics. It is worth pointing out that for $\gamma_i = 0 \forall i$ we retrieve the Liouville - von Neumann equation.

Please notice that, since this is an open system, the dynamics defined by equation (1.40) can be non-unitary. Nevertheless, the Lindblad equation is Markovian and time-homogeneous. It most importantly preserves density matrix properties.

1.2. NV centers in diamonds

The most common diamond-based technologies exploit Nitrogen-Vacancy (NV) pairs, which are created when a Nitrogen impurity substitutes a Carbon atom in the diamond lattice next to a missing Carbon atom (Vacancy, see figure 1.4). This configuration offers a natural trap for electrons (NV centre) whose spin state can be used to encode quantum information [22, 23, 24, 25, 26, 7, 11]. This solution is especially good since electronic systems interact with photons, which makes them optically controllable [11]. This also potentially enables long-range interactions among different sites (both optical and dipole-dipole), a mandatory requirement for quantum networks [27, 11]. Moreover, the electronic spin of the centre can be coupled to the neighbour nuclear spins of the Nitrogen and Carbon-13 atoms, effectively creating a quantum register which can be exploited for QIP.

From a structural point of view, a diamond is a face-centred cubic Bravais-lattice of Carbon atoms (mainly ^{12}C , spinless nuclei), with a natural abundance of 1.1% ^{13}C isotope (nuclear spin of $1/2$). The characteristic distance of two neighbouring lattice sites is 154 pm with a bond angle of 109.5° . The realisation of an NV center can be artificially achieved by randomly removing Carbon atoms via scattering processes and performing Ion implantation [29]. Then, an annealing procedure is usually exploited to spontaneously move these lattice sites close together to form the actual NV. Alternatively, Nitrogen sites can be added directly




Figure 1.4: Structural picture of NV^- centers in diamonds. (a) The NV center in the crystal lattice of the diamond. The NV axis, i.e. the direction defined by the link between the Nitrogen atom and the vacancy, displays C_{3v} symmetry oriented along the 111 axis of the crystal. (b) Energetic scheme of the effective electronic spin $S = 1$. Transitions can be driven among the $m_s = 0$ and $m_s = \pm 1$ levels via microwave (MW) radiation. Green light ($\lambda \leq 637 \text{ nm}$) can be used to perform off-resonant excitation from the 3A_2 to the 3E level. Spontaneous decay from the latter level happens in the form of photoluminescence (ZPL at $\lambda = 637 \text{ nm}$). Higher energy levels can also display preferential decay via the metastable 1A_1 , 1E singlet levels. (c) Focus on the MW transitions of the lowest energy triplet states: upon Zeeman splitting of the $m_s = \pm 1$ levels, single transitions $\omega_{0\pm}$ can be independently addressed to trigger oscillations among the ground and excited states, with detunings δ_{\pm} . Picture from [28].

in the diamond fabrication process via plasma assisted chemical vapour deposition (CVD, see [30]) or exploiting high-pressure high-temperature techniques (HPHT, see [31]). By tuning the percentage of impurities in the material one can regulate the average density of NV sites, which can be observed singularly via confocal microscopes thanks to the characteristic photoluminescence emission. This centers come in three energetically stable fashions: NV^+ , NV^0 and NV^- .

The implanted Nitrogen atom bonds to three of the surrounding Carbon atoms by sharing three of its electrons. The remaining two valence electrons of the Nitrogen, together with three electrons coming from the remaining Carbon atoms in correspondence of the vacancy add up to a total five-electrons, neutral charge state NV^0 , with effective spin $S = 1/2$. If the crystal presents some impurities, a neighbor atom might act as donor and add a sixth electron to the NV (i.e., NV^-): while generically referring to NV center we will indicate this specific configuration, which is usually the favourite one in QIP protocols because of its optical and magnetic properties [32, 8, 11]. In the NV^- center the effective NV electron and the donor's electron form a triplet state with total spin $S = 1$ and multiplicity $2S + 1 = 3$. The quantum number associated to the spin in the \hat{z} direction can hence take the value $m_s = 0$ (anti-parallel alignment) or $m_s = \pm 1$ (parallel alignment). If no external magnetic field is applied then the latter levels are degenerate, otherwise Zeeman splitting occurs (see figure 1.4 (c)). The energy level scheme is reported in figure 1.4 (b): the characteristic zero-field splitting between $m_s = 0$ and $m_s = \pm 1$ in the ground level 3A_2 triplet state is $D_{gs} = 2.87 \text{ GHz}$, while in the 3E excited state is $D_{ex} = 1.42 \text{ GHz}$ at room temperature (see [33]).

As described and derived in [11] (Appendix A), the full Hamiltonian of the system is:

$$\begin{aligned}
\hat{H} = & \underbrace{\hbar D_{gs} \left[\hat{S}_z^2 - \frac{2}{3} \right]}_{\text{zero-field term}} + \hbar \varepsilon \left(\hat{S}_x^2 - \hat{S}_y^2 \right) + \underbrace{\hbar \gamma_{nv} \vec{B} \cdot \hat{\vec{S}}}_{\text{magnetic interaction}} \\
& + \underbrace{\hbar \delta_{\parallel} E_z \left[\hat{S}_z^2 - \frac{2}{3} \right] - \hbar \delta_{\perp} \left[E_x \left(\hat{S}_x \hat{S}_y + \hat{S}_y \hat{S}_x \right) + E_y \left(\hat{S}_x^2 - \hat{S}_y^2 \right) \right]}_{\text{electric interaction}} \\
& + \hbar \sum_{i=1}^n \left(\underbrace{\hat{S} \hat{N}_i \hat{I}_i}_{\text{hyperfine interactions}} + \underbrace{\gamma_i \vec{B} \cdot \hat{I}_i}_{\text{nuclear Zeeman interactions}} + \underbrace{Q_i \hat{I}_{z,i}^2}_{\text{nuclear quadrupole interactions}} \right)
\end{aligned} \tag{1.41}$$

where it is assumed the NV axis to be the quantisation axis. Here $\hat{\vec{S}} = \left(\hat{S}_x, \hat{S}_y, \hat{S}_z \right)^T$ are the spin operators, ε is the non-axial zero-field parameter and $\gamma_{nv} = 2\pi \times 28 \text{ GHz/T}$ is the gyromagnetic ratio of the NV. $\vec{E} = (E_x, E_y, E_z)$ and $\vec{B} = (B_x, B_y, B_z)$ are the effective electric and magnetic fields accounting for both external and internal (structural) contributes. $\delta_{\parallel} = \mathcal{O}(10^{-1} \text{ Hz} \cdot m/V)$ and $\delta_{\perp} = \mathcal{O}(10^{-3} \text{ Hz} \cdot m/V)$ are the axial and transverse coupling constants, respectively.




Figure 1.5: Fluorescence spectrum of NV centers in diamonds. (a) Schematic optical transitions among the 3A_2 and 3E levels. Excitation is provided by spin-preserving off-resonant radiation, which spontaneously leads to population of the triplet states. Electrons can decay back to the ground state together with fluorescent emission. (b) Spectrum of fluorescent emission for the NV^0 and NV^- species. The zero phonon line (ZPL) is marked at 575 nm and 637 nm , respectively. Figure from [33].

$\{\hat{I}_i\}_i$ are the nuclear spin operators of the neighbouring ${}^{13}\text{C}$ or ${}^{14}\text{N}$ and ${}^{15}\text{N}$ atoms, while \mathcal{N}_i is the hyperfine interaction tensor between the electron and the i -th nuclear spin.

In absence of external fields, the Hamiltonian is reduced to the zero-field term. The electric and magnetic interaction terms account for corrections in presence of these fields. The last line in equation (1.41) accounts for the presence other spins in the lattice: the hyperfine interaction between these and the NV center (comprehensive of isotropic, contact Fermi interactions and anisotropic magnetic dipole interactions which scale as $r_{S,I}^{-3}$), the coupling of the nuclear spin with the external \vec{B} field and the nuclear spin quadrupole splitting represented by Q_i .

1.2.1. Optical readout

Transitions among the 3A_2 and 3E states can be optically triggered by photons of energy above the zero phonon line (ZPL), which is 1.945 eV (637 nm , see figure 1.5 (b)) for the system under investigation. Commonly, such transitions are driven off-resonantly (green light) and are spin preserving because of the selection rules, leading to a conservation of the spin quantum number m_s . Then, the system can return to the ground state either via fluorescent emission (red light, spectrum in figure 1.5 (b)) or following the path involving the additional 1A_1 and 1E singlet states (figure 1.4 (b)), which causes a spin-flip. While the first relaxation is much more probable if the system is in a $m_s = 0$ state (hence called the bright state of the system), the latter path is strongly favoured when the system is in the 3E , $m_s = \pm 1$ states and results in the emission of infrared light which is usually filtered




Figure 1.6: ODMR spectrum of NV center in absence (a) and presence (b) of external aligned magnetic field ($\mathcal{O}(0.1 \text{ mT})$). When no Zeeman splitting occurs, the dip in the spectrum is located at the characteristic frequency of 2.87 GHz . When the magnetic field is turned on the dip splits in two, highlighting the symmetrically distributed resonance peaks. Figure from [36].

out (hence the definition of dark state, see [8]). This allows to perform both state initialisation and readout. The former comes straightforward from the fact that, following the excitation, the relaxation from the 3E level will bring the system to the 3A_2 , $m_s = 0$ state. If such process triggers the emission of red photons then the original pre-excitation state was in $m_s = 0$, while if not then it was in $m_s = \pm 1$. The visibility associated to this readout process, i.e. the maximum qubit population difference observed, is $\sim 70\%$ (more in [34]). In both cases, the resulting state of the measurement is a $m_s = 0$ state: optical pumping can hence be applied to prepare the system in this state.

Assuming now that the system is in the aforementioned state, microwave (MW) radiation can be used to bring it to the 3A_2 , $m_s = \pm 1$ states via Rabi oscillations (details in section 1.2.2). In principle the resonant frequency of this transition can be found by performing a sweep over a range of frequencies and looking for a dip in the fluorescence spectrum. In fact, since the initial state is a bright state, when the incoming light has off-resonant frequency the system remains in a bright state and the fluorescence emission is constant. While the frequency approaches the resonance, the system is more and more likely to move to the dark state, which relaxes via the infrared path and therefore results in a dip in the red fluorescence spectrum [35, 36]. This process takes the name of optically detected magnetic resonance (ODMR, illustration in figure 1.6).

This technique becomes extremely useful in the field of quantum sensing: since the splitting of the $m_s = \pm 1$ levels is linearly proportional to the external magnetic field parallel to the \hat{z} direction as $\Delta\omega = 2\gamma_{nv}B_z$, being $\gamma_{nv} = 2\pi \times 28 \text{ GHz/T}$ the gyromagnetic ratio of the NV, measuring the spectral gap between the two dips allows to estimate the magnetic field [37]. Moreover, by properly tuning the intensity of a constant B field it is possible to lift the degeneracy between these levels up to the point where transitions can be driven between the $m_s = 0$ and $m_s = -1$ being

completely off-resonance with respect to the $m_s = 0$ to $m_s = +1$ transition. This allows to effectively define a two level system that can be used as a qubit for computational purposes with minimal information leakage towards the spectator state $m_s = +1$. Therefore, whenever we refer to the electronic spin in the following sections we address this two-level subsystem that we treat as an electronic qubit in the center, with ground state $|0\rangle_e$ defined by the ${}^3A_2, m_s = 0$ level and excited state $|1\rangle_e$ as the ${}^3A_2, m_s = -1$ level, after removing the degeneracy by application of a constant B_z field.

1.2.2. Rabi oscillations and optically engineered quantum gates

As mentioned above, transitions can be driven across any two levels by means of proper pulses. More in detail, this is usually achieved by employing an oscillating electromagnetic field at frequency ω , while the transition frequency between two energy levels E_a and E_b is defined as

$$\omega_0 \equiv |E_a - E_b|/\hbar \quad (1.42)$$

As reported in [38], if the driving frequency is set on resonance ($\omega = \omega_0$) and a continuous external field is applied, perfect oscillations between the $|a\rangle$ and $|b\rangle$ levels can be performed. This phenomenon is known as Rabi oscillations. Assuming the system to be initially in the $|a\rangle$ state, the probability of measuring the system in the $|b\rangle$ state after a time t is given by Born's rule:

$$P_{a \rightarrow b} = |\langle \psi(t) | b \rangle|^2 = \sin^2 \left(\frac{\Omega t}{2} \right) \quad (1.43)$$

Where the Rabi frequency Ω is proportional to the induced dipole and the intensity of the driving field as $\Omega = dE/\hbar$ (more details for the specific case of NVs in section 1.3). It is worth notice that perfect transitions (population inversion) between the two levels can be achieved for $\Omega t = \pi$: this case is commonly referred as a π -pulse and represents an \hat{X} (or \hat{Y} , according to the phase) single qubit gate (details in section 1.1.3). Similarly, one can notice that for $\Omega t = \pi/2$ balanced quantum superpositions in the form $|\psi(t)\rangle = (|a\rangle + |b\rangle)/\sqrt{2}$ can be prepared (Hadamard gate). In this case we talk about $\pi/2$ -pulses. More in general, by fixing the intensity of the external field the oscillation process can be interrupted at arbitrary time to prepare any superposition of the two states. Moreover, the choice of the phase and the orientation of the field allow to introduce an arbitrary phase in the superposition of the basis states, effectively driving rotations along an equatorial plane of the Bloch sphere (see figure 1.7). An example of experimental Rabi oscillations in an NV center is reported in figure 1.8.

While perfect oscillation can be driven by resonant driving, it is more common to drive slightly off-resonant transitions, for which it holds $\omega = \omega_0 + \Delta$. In this case the theory forbids perfect transitions from a pure state to the other, as the system preserves a non-zero probability of remaining in the initial state. Starting in the $|a\rangle$ state, after a time t it holds:




Figure 1.7: Arbitrary rotation of a single-qubit state by means of a driving field \vec{E} with $|\vec{E}| = E \cos(\omega t)$ and $\vec{E}/|\vec{E}| = (\cos(\theta), \sin(\theta), 0)$. The basis is defined by the spin states in the \hat{z} direction.




Figure 1.8: Experimental data of NV center characterisation. (a) Rabi oscillations of the $m_s = 0$ ($|0\rangle$), $m_s = -1$ ($|1\rangle$) subsystem of the 3A_2 energy level of the NV. After the initialisation the system is in the $|0\rangle$ state. A microwave square pulse of resonant frequency is then applied for different pulse lengths t_p . The data show the perfect population inversion for the first Rabi cycle and subsequent amplitude decay due to decoherence while keeping the oscillatory behaviour. (b) Study of the T_2 coherence time of the system via spin echo measurements (see [39]). Figure from [39].

$$\begin{aligned}
P_{a \rightarrow a} &= |\langle \psi(t) | a \rangle| = \frac{\Omega^2}{\Omega_R^2} \cos^2 \left(\frac{\Omega_R t}{2} \right) + \frac{\Delta^2}{\Omega_R^2} \\
P_{a \rightarrow b} &= |\langle \psi(t) | b \rangle| = \frac{\Omega^2}{\Omega_R^2} \sin^2 \left(\frac{\Omega_R t}{2} \right)
\end{aligned} \tag{1.44}$$

Where now the Rabi frequency Ω_R is redefined in a more general fashion:

$$\Omega_R \equiv \sqrt{\Omega^2 + \Delta^2} \tag{1.45}$$

and for the resonant case ($\Delta = 0$) we retrieve $\Omega_R = \Omega$.

In practical application, instead of tuning the oscillation time it is more convenient to design a variable intensity profile for the driving field while keeping fixed the total pulse duration. Therefore, quantum gates can be designed for the specific system by means of tailored driving fields, defining intensity profiles $\Omega(t)$ and frequency fluctuations $\Delta(t)$ over a fixed gate time T . Since for electromagnetic controls the Rabi frequency is proportional to the electric field, according to the description one wishes to adopt it can take both complex or real values.

The study of the optimal pulse details for gate generation is called Quantum Optimal Control (QOC) and is better described in section 1.4.3.

1.3. System specifics

As described in section 1.2, NV centers in diamonds allow to define electronic two-level quantum systems with well established energy schemes. Generally speaking, when multiple NVs are found in a restricted area then interaction between different NV spins become possible, generating in fact a multi-qubit system. On the other hand, two-qubit systems can be generated naturally even for an isolated NV center because of the presence of Nitrogen atoms with non-zero nuclear spin, as well as possible neighbour ^{13}C species. In these cases, the nuclear spin of the atom can couple with the electronic spin (e.g. via dipole-dipole interaction).

The physical system we refer to is fully described in [40] and depicted in figure 1.9. In detail, we consider the hyperfine coupling between the electronic spins of a NV centre and the nuclear $1/2$ -spin of a neighbour ^{13}C atom with spin states down ($|0\rangle_n$) and up ($|1\rangle_n$).

The complete system Hamiltonian can be expressed as

$$\hat{H} = \hat{H}_d + \hat{H}_c \tag{1.46}$$

where the drift Hamiltonian defining the self-energy with respect to the ground state $|0\rangle_{en}$ is:

$$\hat{H}_d = \omega_{01} \hat{\sigma}_{11} + \omega_{02} \hat{\sigma}_{22} + \omega_{03} \hat{\sigma}_{33} \tag{1.47}$$




Figure 1.9: Electron spin - nuclear spin system in NV centre. The transition frequencies ω_{ij} are taken from [40] and reported in the text. Controls and detunings ω_k, δ_k ($k = m, r_1$) are subject to optimisation. The control ω_{r_2} is originally introduced for state initialisation only.

where $\hat{\sigma}_{pq} \equiv |p\rangle\langle q|$ and $\omega_{01} \approx 2 \text{ MHz}$, $\omega_{02} \approx 2.88 \text{ GHz}$, $\omega_{03} \approx 130 \text{ MHz}$. The control Hamiltonian \hat{H}_c accounts for the external electromagnetic pulses. Consider an oscillating monochromatic field at frequency ω in the form of a plane wave:

$$\vec{E} = \vec{E}_0(t) e^{-i\omega t} + \vec{E}_0^*(t) e^{i\omega t} \quad (1.48)$$

where the field intensity can be modulated over time. From the light-matter interaction theory within the dipole approximation, it is known that transitions among two energy levels $|p\rangle$ and $|q\rangle$ can occur with probability proportional to the dipole matrix element $\vec{d}_{pq} = \langle p|\hat{\vec{d}}|q\rangle$, being $\hat{\vec{d}}$ the dipole moment of the particle. Being the latter an hermitian operator and considering that $\langle p|\hat{\vec{d}}|p\rangle = \langle q|\hat{\vec{d}}|q\rangle = 0$, we can decompose $\hat{\vec{d}}$ as:

$$\hat{\vec{d}} = \vec{d}_{pq} |p\rangle\langle q| + \vec{d}_{pq}^* |q\rangle\langle p| \quad (1.49)$$

Now, the energy associated to the light-matter interaction is the dipole energy that, applying the definition above, becomes:

$$\begin{aligned} \hat{H}_d &= -\hat{\vec{d}} \cdot \vec{E} \\ &= -\hbar \left[\left(\Omega e^{-i\omega t} + \tilde{\Omega} e^{i\omega t} \right) |p\rangle\langle q| + \left(\Omega^* e^{-i\omega t} + \tilde{\Omega}^* e^{i\omega t} \right) |q\rangle\langle p| \right] \end{aligned} \quad (1.50)$$

where $\Omega = \vec{d}_{pq} \cdot \vec{E}_0 / \hbar$ is the Rabi frequency and $\tilde{\Omega} = \vec{d}_{pq}^* \cdot \vec{E}_0^* / \hbar$ the counter-rotating frequency. By moving to the interaction picture one finds that the terms involving Rabi frequencies oscillate at a frequency $\Delta\omega = \omega - \omega_0$, being ω_0 the resonance frequency between $|p\rangle$ and $|q\rangle$, while the ones with amplitude $\tilde{\Omega}$ oscillate at $\omega + \omega_0$. As a result, if $\omega \approx \omega_0$, we can neglect the fast oscillating terms in

$\tilde{\Omega}$ treating them as constantly equal to their average value in the interaction time window: this operation takes the name of rotating wave approximation (RWA). Alternative approaches suggest the validity of this description also in the case of $\Omega \ll \omega_0$, i.e. the case of weak (or slow) driving which will be the case for this work (see [41] for reference). As a result, in the interaction picture the two-level control Hamiltonian can be written as:

$$\begin{aligned}\hat{H}_d^I &= -\hbar\Omega e^{-i\Delta\omega t} |p\rangle\langle q| - \hbar\Omega^* e^{i\Delta\omega t} |q\rangle\langle p| \\ &= -\hbar\Omega e^{-i\Delta\omega t} |p\rangle\langle q| + h.c.\end{aligned}\quad (1.51)$$

which can now be generalised to the full set of pulses which drive transitions among the states. As a result, the total control Hamiltonian in the Schrödinger picture is given (in units of \hbar for simplicity) by:

$$\hat{H}_c = -\frac{1}{2} (\Omega_m e^{i\omega_m t} \hat{\sigma}_{02} + \Omega_{r_1} e^{i\omega_{r_1} t} \hat{\sigma}_{23} + \Omega_{r_2} e^{i\omega_{r_2} t} \hat{\sigma}_{13} + h.c.) \quad (1.52)$$

being Ω_k the Rabi frequencies controlling the MW transition ($k = m$) and RF transitions ($k = r_1, r_2$) between the energy levels $|p\rangle$ and $|q\rangle$.

It is convenient to move to the interaction picture defined by:

$$\begin{aligned}\hat{U}_0(t) &= \exp(-i\hat{H}_0 t) \\ \text{with } \hat{H}_0 &= a\hat{\sigma}_{22} + b\hat{\sigma}_{33} + c\hat{\sigma}_{11} + d\hat{\sigma}_{00}\end{aligned}\quad (1.53)$$

and with the choice of parameters:

$$\begin{cases} a = \omega_m - \frac{4}{3}\delta_m - \frac{1}{2}\delta_r \\ b = a - \omega_{r_1} \\ c = \omega_{01} + 2\delta_m - \gamma \\ d = a - \omega_m \end{cases} \quad (1.54)$$

we obtain an effective Hamiltonian given by $\hat{H}_{eff} = \hat{U}\hat{H}\hat{U}^\dagger - i\dot{\hat{U}}\hat{U}^\dagger$ that can be written as:

$$\hat{H} = -\frac{1}{2} \begin{pmatrix} \frac{4}{3}\delta_m - \frac{2}{3}\delta_{r_1} & 0 & \Omega_{02} & 0 \\ 0 & \gamma & 0 & \Omega_{13} \\ \Omega_{02}^* & 0 & -\frac{2}{3}\delta_m - \frac{2}{3}\delta_{r_1} & \Omega_{23} \\ 0 & \Omega_{13}^* & \Omega_{23}^* & -\frac{2}{3}\delta_m + \frac{4}{3}\delta_{r_1} \end{pmatrix} \quad (1.55)$$

which in the case $\delta_m = \delta_{r_1} = \delta$ simplifies as:

$$\hat{H} = -\frac{1}{2} \begin{pmatrix} \frac{2}{3}\delta & 0 & \Omega_{02} & 0 \\ 0 & \gamma & 0 & \Omega_{13} \\ \Omega_{02}^* & 0 & -\frac{4}{3}\delta & \Omega_{23} \\ 0 & \Omega_{13}^* & \Omega_{23}^* & \frac{2}{3}\delta \end{pmatrix} \quad (1.56)$$

where $\Omega_{nm} = u_{nm}e^{i\theta_{nm}}$ is the control performing Rabi oscillations between the $|n\rangle$ and $|m\rangle$ levels while δ and γ are detuning parameters.

equation (1.56) defines an effective Hamiltonian for the system that will be used to build the time evolution operator to be optimised for the design of the quantum gate in the simulation of the model. The definition of the Hamiltonian has been done following the reasoning illustrated in [40]. Nevertheless, a different reference frame characterised by a new definition of the a, b, c, d parameters in equation (1.53) has been introduced to overcome some technical issues, as further discussed in section 3.2.1.

1.4. Control Theory

In today's technology most precise high-tech devices exploit some kind of feedback-based approach to automatically calibrate and respond to external stimuli. This is indeed a very natural behaviour: from animals to plants, all creatures spontaneously learn from the environment by performing certain actions, analysing the environment's response and therefore deciding the next action to take. This mental scheme is therefore translated also to the field of high-precision technology: while a whole branch of machine learning exploits such methods to develop a sort of "consciousness" about rules and goals within a specific framework (reinforcement learning, see [42]), every-day applications such as thermostats in refrigerators or laser light emitters make use of this feedback-controlled methods to keep a steady working state [43].

It is important to stress that control theory takes the role of "steering" the dynamics of a specific system. In particular, from a classical perspective it could be defined as the study of the dynamics manipulation given the environment actions on the system itself. When it comes to quantum, control theory is responsible for a coherent manipulation of the evolution of a quantum system which is otherwise subject to the non-coherent action of the environment, which directly affects the system state by means of external fields and measurements effects on the quantum wavefunction. In this "backaction" on the system lays the crucial distinction between classical and quantum devices. According to our goal, it can be either an opportunity to encode environmental information in the system of an unwanted modification that we will want to avoid. Because of all this, a structural theory for controlling the dynamics of quantum systems is required.

In this section we highlight the main features of classical and quantum control theory, defining important concepts and describing how this framework can be useful to enhance the performance of our (quantum) devices. A full description of classical and quantum control specifics can be found in [43].

1.4.1. Classical Control Theory

Let us try to formalise the theoretical framework we are working in: we shall assume that at any arbitrary time t the state of the system under investigation is

defined by a set of variables $\mathbf{x}(t) \in \mathbb{C}^n$, whose values allow to encode all the information about the system. The domain of $\mathbf{x}(t)$ defines the state space and the dynamics of the system will be therefore described by a trajectory in such a space, defined by the values of $\mathbf{x}(t) : 0 \leq t \leq T$ over a certain period of time T . With this picture we are implicitly assuming that the system is subject to an equation of motion (EoM) of the form:

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{a}(\mathbf{x}_{<t}, t) \text{ with } \mathbf{a} : \mathbb{C}^n \times \mathbb{R} \mapsto \mathbb{C}^n \quad (1.57)$$

i.e., we assume that there exists an EoM for the system state determined by the presence of a function \mathbf{a} of the history of the system $\mathbf{x}_{<t} \equiv \{\mathbf{x}(t) : 0 \leq t < T\}$ and the current time t .

While equation (1.57) governs the state evolution, this information is only known to the system itself. Instead, an external observer can gain information about the system state only performing a measurement on it. In real life, perfect continuous monitoring of the system state is impossible to perform (meaning that we cannot access $\mathbf{x}(t)$ per se), but the system can be only partially observed returning outputs

$$\mathbf{y}(t) = \mathbf{c}(\mathbf{x}(t), t) \text{ with } \mathbf{c} : \mathbb{C}^n \times t \mapsto \mathbb{R}^m \quad (1.58)$$

where in the most general case the system description can be complex, while the measurement outputs $\mathbf{y}(t)$ take real values. In this framework we are basically stating that all the information we can retrieve from the system by probing it and the dynamical evolution of the system itself are described by some (possibly non-linear) functions.

The description adopted so far is the one of an open system subject to its own evolution laws that can be only saw by performing a measurement. In a more general fashion, we could also introduce the possibility for us external users to control the dynamics of the system by the means of input (control) signals $\mathbf{u}(t)$. This leads to a modification of the equation of motion for the system in the following form:

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{a}(\mathbf{x}_{<t}, t) + \mathbf{b}(\mathbf{x}_{<t}, \mathbf{u}(<t), t) \text{ with } \mathbf{b} : \mathbb{C}^n \times \mathbb{R}^k \times \mathbb{R} \mapsto \mathbb{C}^n \quad (1.59)$$

while consistently also the measurements could depend on the controls: $\mathbf{y} = \mathbf{c}(\mathbf{x}_{<t}, \mathbf{u}(<t), t)$. This formulation is rather general and covers any possible expression for the evolution function and the controls. Still, to make these problems analytically solvable it is common practice to restrict to the case in which the right-hand terms in equations 1.59 and 1.58 are represented by linear, time-stationary terms: this is the class of the so-called linear time-invariant (LTI) models (see [44]), for which it holds:

$$\begin{aligned} \frac{d}{dt}\mathbf{x}(t) &= \mathbf{a}(\mathbf{x}_{<t}, t) + \mathbf{b}(\mathbf{x}_{<t}, \mathbf{u}(<t), t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \\ \mathbf{y}(t) &= \mathbf{c}(\mathbf{x}(t), \mathbf{u}(t), t) = \mathbf{C}\mathbf{x}(t) + \mathbf{D}\mathbf{u}(t) \end{aligned} \quad (1.60)$$

with $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$ proper linear operators with consistent sizes and domains.

The controls $\mathbf{u}(t)$ act as a steer on the evolution of the states, deviating them from their original trajectory. As a result, they allow to reach points in the state space which would not be naturally be crossed by the dynamics. We can hence perform state preparation in this way. In the most general sense, once a specific target is set, there exist multiple signals $\{\mathbf{u}_i(t)\}_i$ which solve the control problem. According to the situation, one could choose any of this controls or perform further optimisation, for example minimising the total time or the control norm $\|\mathbf{u}_i\|_2$.

1.4.2. Classification of control models

So far we introduced the presence of control signals and their theoretical action on the system dynamics. Nevertheless, even if perfect controls exist it might be non trivial to learn their form. There are two major classes of approaches that allow to infer the optimal control \mathbf{u}^* for a given problem: we review them below, with a particular focus on the closed-loop approach which has been the main focus of this work [45, 43].

Open-loop control

If the evolution of the system is known from some theoretical underlying model, one can proceed to analytically or numerically solve the problem given by:

$$\begin{cases} \dot{\mathbf{x}}(t) = \mathbf{a}(\mathbf{x}_{<t}, t) + \mathbf{b}(\mathbf{x}_{<t}, \mathbf{u}_{<t}^*, t) \\ \mathbf{x}(0) = \mathbf{x}_0 \\ \mathbf{y}(0) = \mathbf{c}(\mathbf{x}_0, \mathbf{0}, 0) \\ \mathbf{y}(T) = \mathbf{c}(\mathbf{x}(T), \mathbf{u}_T^*, T) = \mathbf{y}_{target} \end{cases} \quad (1.61)$$

This approach is especially powerful because it would in principle return an exact solution to the problem and therefore find the truly optimal control of the system. Nevertheless, it presents a major issue: we do not want the system to be too complicated to not be analytically solvable, yet we cannot idealise the system representation up to the point that the description is too simplistic. In fact, a deterministic solution would be possible only by neglecting all the random contributions to the dynamics coming from environmental noise and readout imperfections, which are assumed small. Moreover, if the theoretical background describing the system is not accurate, the optimal signal returned by this optimisation will inevitably not work properly of the real device.

Closed-loop control

To overcome the modeling issues of the open-loop approach, a more trial-and-error based solution can be exploited. As reported in figure 1.10, the aforementioned control problem can be saw as finding the right input to a black box which performs state evolution to retrieve the desired output. Therefore, we can implement




Figure 1.10: Scheme of optimal control inference process. The system is treated as a black box which takes in input the trial control signal and evolves the system via and EoM ruled by $f(x, u, t) \equiv a(x, t) + b(x, u, t)$. The output signal is the measurement result. If the input control is optimal ($\mathbf{u}(t) = \mathbf{u}^*(t)$), then the output signal is the target one. In closed-loop optimisation output signals are used as a feedback to adjust the input one.

an iterative protocol: we propose an input signal $\mathbf{u}(t)$, then probe the system retrieving some measurements $\{\mathbf{y}(t_i)\}$ to infer something about the correct action of the control on the system and ultimately we exploit this information as a feedback to fine tune the details of the control. While being in most cases more challenging due to the difficult technical implementation, this approach has better potential than open-loop control since the real device is directly probed and therefore no precise modeling has to be done.

The closed-loop control class is quite wide and includes several methods and procedures, which share the common feature of performing an iterative signal optimisation and therefore "control learning". Sometimes an excessively fast dynamical evolution or experimental difficulties in performing specific measurements can prevent the user to exploit the real system for the evolution. In these cases it is convenient to substitute the real device with a detailed simulation, which differentiates from the open-loop case because now we are taking into account stochastic processes and noise sources as well. This keeps the fundamental distinction between open- and closed-loop approaches valid: in the first case we seek numerical solutions, while in the second case we infer suitable signals via adaptive learning.

1.4.3. Quantum Control

The theory of classical control depicted in the previous section applies to a large variety of systems, regardless of their size and the physics behind them. However, when it comes to quantum systems some important differences arise: the first and main one is the backaction induced by the measurement operation, described by the projection postulate. This means that we cannot think about retrieving information from the system without perturbing it. Moreover, projective measurements

on the system result in the collapse of the wavefunction onto the respective Hilbert subspace, introducing a stochastic component in the dynamics that is not fundamentally taken into account for the classical case. While trying to interfere as little as possible with the system, weak measurements can be used instead of projective ones: this would in principle reduce the information extracted by the user but it will not make the quantum state fully collapse [46].

Keeping in mind the previous important note, the translation from the classical state-space control model to the domain of quantum systems is quite straightforward. In particular, the intrinsic formalism of quantum mechanics (which is usually described in terms of linear operators and matrices) allows to find direct correspondence with classical LTI models:

- **State vector** \longleftrightarrow **Density matrix**: while for classical systems $\mathbf{x}(t)$ is the state encoding all the useful information about the system, in a quantum mechanical framework the state is defined by the density operator $\hat{\rho}(t)$. Fundamental differences are the probabilistic nature of the quantum state representation and its definition up to a global phase.
- **State space** \longleftrightarrow **Hilbert space**: a consequent analogy lays in the domain of the state, which is no longer the vector space of $\mathbf{x}(t)$ but becomes the Hilbert space which hosts the quantum states. For both spaces the dimension might become infinity.
- **System matrix A and forcing term $Bu(t)$** \longleftrightarrow **Drift and control Hamiltonians \hat{H}_0, \hat{H}_c** : while dealing with quantum states, the evolution of the system is ruled by the Hamiltonian of the system. Therefore, the system matrix describing the free evolution of the system has as counterpart the drift Hamiltonian \hat{H}_0 which determines the energy eigenstates, while the additional term in the EoM which is related to the external controls becomes in a quantum domain an additional term in the total Hamiltonian:

$$\hat{H}(t) = \hat{H}_0(t) + \underbrace{\sum_i u_i(t) \hat{H}_{c,i}}_{\equiv \hat{H}_c(t)} \quad (1.62)$$

According to the previous analogies, if we wish to control quantum systems it is therefore necessary to define a set of controllable actions that contribute to the definition of the system Hamiltonian at each time. This will ultimately impact the time evolution and therefore the dynamics of the quantum states.

1.4.4. Fidelity and Figure of Merit

In the previous sections we stated that the goal of control theory is the manipulation of a system evolution to redirect its trajectories in the state-space and eventually reach a specific target. In order to quantify the quality of a trial solution, it is

necessary to introduce a metric or, as it is usually called in the field, a Fidelity measure \mathcal{F} . The latter must satisfy some basic requirements:

- Upper- and lower-bounded (usually in the interval $[0, 1]$)
- Reaches its maximum if and only if the target is perfectly reached (up to some permitted symmetries, such as a global phase difference in a quantum state preparation)

Once a quality metric has been defined, the control problem can be revisited as a maximisation problem, where the quantity to maximise takes the name of Figure of Merit (FoM). In most cases this coincides with the fidelity, but it can be eventually defined taking into account some other quantities which might affect the optimal outcome, such as the control pulse power $\|\mathbf{u}(t)\|_2$. Practically speaking, it is common to translate the optimisation problem into a minimisation task with loss function \mathcal{J} given by $1 - FoM$, which usually reduces to the simple infidelity measure $1 - \mathcal{F}$.

For the case of quantum states a natural fidelity exists. Consider the time-evolved system state $|\psi(t)\rangle$ which satisfies:

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = \hat{H}(t) |\psi(t)\rangle = \left(\hat{H}_0 + u(t) \hat{H}_c \right) |\psi(t)\rangle \quad (1.63)$$

$$|\psi(t=0)\rangle = |\psi_0\rangle$$

where we restricted to the case of a time-independent drift Hamiltonian and a single control. Then, a Fidelity measure is given by the quantum overlap between the target state $|\xi\rangle$ and $|\psi(t)\rangle$:

$$\mathcal{F}_\xi(\psi(t)) \equiv |\langle \xi | \psi(t) \rangle|^2 = \left| \text{Tr} \left(\hat{\Pi}_\xi \hat{\rho}(t) \right) \right|^2 \quad (1.64)$$

where $\hat{\Pi}_\xi \equiv |\xi\rangle\langle\xi|$ and, for a pure state, $\hat{\rho}(t) \equiv |\psi(t)\rangle\langle\psi(t)|$. One can verify that such a metric satisfies the requirements above. Please notice that:

$$\mathcal{F}_\xi(\psi(t)) = 1 \iff |\psi(t)\rangle = e^{i\varphi} |\xi\rangle \quad (1.65)$$

i.e., maximum fidelity is reached when the time-evolved state is equivalent to the target (equal up to a global phase φ).

As for the optimisation of quantum gates, the principle of superposition in quantum mechanics implies that to ensure the correct implementation of the gate it is sufficient to verify its expected action on the basis states $\{|\phi_i\rangle\}_i$, since each other state can be written as a linear composition of these. Attention should be paid to the global phase added to each basis state upon evolution: if we require it to be the same for all the basis states, a suitable candidate for a fidelity metric for quantum gates is:

$$\mathcal{F} \equiv \frac{1}{N^2} \left| \sum_{i=1}^N \langle \chi_i | \phi_i(t) \rangle \right|^2 \quad (1.66)$$

where the $|\chi_i\rangle$ states are defined as the theoretical application of the target state \hat{V} to the basis states:

$$|\chi_i\rangle \equiv \hat{V} |\phi_i(0)\rangle \quad \forall i \quad (1.67)$$

In the most general case, the time evolution of the quantum states ruled by the system Hamiltonian can be viewed as the action of a time evolution operator $\hat{U}(\mathbf{u}(t), \boldsymbol{\theta}; t)$ from $t_0 = 0$ which depends on the controls $\mathbf{u}(t) = \{u_i(t)\}_i$, the final time t and possibly some additional parameters $\boldsymbol{\theta} = \{\theta_j\}_j$ (see section 1.1.5):

$$|\phi_i(t)\rangle \equiv \hat{U}(\mathbf{u}(t), \boldsymbol{\theta}; t) |\phi_i(t=0)\rangle \quad \forall i \quad (1.68)$$

Therefore, for each basis state it holds:

$$\langle \chi_i | \phi_i(t) \rangle = \langle \phi_i(0) | \hat{V}^\dagger \hat{U} | \phi_i(0) \rangle \quad (1.69)$$

which is the i -th entry of the trace of the $\hat{V}^\dagger \hat{U}$ operator, i.e. its expectation value in the $|\phi_i(0)\rangle$ state. The quantity in equation (1.69) is a complex number with modulus bounded in the interval $[0, 1]$ and the normalisation constraint $\langle \phi_i | \phi_i \rangle = 1$ ensures that the maximum of $|\langle \chi_i | \phi_i(t) \rangle|$ is reached if and only if:

$$\hat{V}^\dagger \hat{U} = e^{i\varphi} \mathbf{1}_N \implies \hat{U} = e^{i\varphi} (\hat{V}^\dagger)^{-1} = e^{i\varphi} \hat{V} \quad (1.70)$$

i.e., the time evolution operator is equivalent to the (unitary) target gate up to a global phase. Taking the normalised sum over all the basis states ensures that maximum gate fidelity can be reached if and only if all states are correctly evolved to their target, eventually up to the same global phase.

These considerations suggest us the possibility to define gate fidelities on the basis of the trace of the $\hat{V}^\dagger \hat{U}$ operator which take into account the amount of symmetries we are willing to accept in our final result. As reported in [45], the control problem can be in fact claimed as solved when the \hat{U} operator is equivalent to the target \hat{V} gate, eventually up to some symmetries such as global or local phases and local operations.

When the goal of the optimisation is to produce $\hat{U} = \hat{V}$, a suitable fidelity measure is given by:

$$\mathcal{F}_{Re} = \frac{1}{N} \operatorname{Re} \left[\operatorname{Tr} \left(\hat{U}^\dagger \hat{V} \right) \right] \quad (1.71)$$

which satisfies $\mathcal{F}_{Re} \in [0, 1]$ and $\mathcal{F}_{Re} = 1 \iff \hat{U} = \hat{V}$.

If instead the optimisation can be considered successful when the time evolution operator is equal to the target gate up to a global phase, it is convenient to use:

$$\mathcal{F}_{sm} = \frac{1}{N^2} \left| \operatorname{Tr} \left(\hat{U}^\dagger \hat{V} \right) \right|^2 \quad (1.72)$$

and one can observe how this is an equivalent formulation of equation (1.66).

Finally, if we wish to allow for phase differences among the individual matrix elements of \hat{U} and \hat{V} a suitable metric is given by:

$$\mathcal{F}_{ss} = \frac{1}{N^2} \sum_{k,l=1}^N |u_{kl}^* v_{kl}|^2 \quad (1.73)$$

For all the aforementioned scenarios, the optimisation problem will be perfectly solved by $\mathcal{F} = 1$.

As for our implementation and simulations, in this work we focus on gate optimisation up to a global phase.

1.5. Optimal control algorithms

In the previous sections we introduced some basic concepts of control theory and we saw how the goal of optimal control is to find the optimal set of controls that steer the system dynamics in a desired way. An optimal solution is therefore found when the state of the system (or the whole basis evolution) is close to the target, given a specific metric. In this section we investigate the state-of-the-art algorithms which allow to learn the best set of controls (a recent full review in [47]).

1.5.1. Gradient Ascend Pulse Engineering (GRAPE)

When the gradient of the control objective can be efficiently computed, a large family of methods are applicable to move in a deterministic and consistent way towards the global minimum. As the self-explanatory name suggests, the GRAdient Ascend Pulse Engineering (GRAPE) method is a gradient-based algorithm widely used in the field of optimal control for state preparation or other sectors where the analytical form of the gradient can be calculated. For the former, let us assume that a target state $\hat{\rho}(T)$ to be prepared from $\hat{\rho}(0)$ by tuning an Hamiltonian as the one in equation (1.62). As discussed in section 1.1.5, the dynamics of the state is ruled by the Liouville-von Neumann equation that, upon discretisation of the total transfer time T into $\{t_i\}_{i=1,\dots,M}$ where $t_i = i \cdot \Delta t$ and $\Delta t = T/M$, becomes:

$$\hat{\rho}(T) = \hat{U}_M \hat{U}_{M-1} \dots \hat{U}_1 \hat{\rho}(0) \hat{U}_1^\dagger \dots \hat{U}_{M-1}^\dagger \hat{U}_M^\dagger \quad (1.74)$$

where:

$$\hat{U}_k = \exp \left[-\frac{i}{\hbar} \left(\hat{H}_0 + \sum_j u_j(k\Delta t) \cdot \hat{H}_j \right) \Delta t \right] \quad (1.75)$$

Since in the discrete case for Δt sufficiently small it holds:

$$\hat{U}(0; T) \simeq \mathcal{T} \left[\left\{ \hat{U}_k \right\}_{k=1,\dots,M} \right] = \hat{U}_M \hat{U}_{M-1} \dots \hat{U}_1 \quad (1.76)$$

where $\mathcal{T}[\cdot]$ the time ordering operator, we can calculate the variation of a fidelity measure as the one in equation (1.64) with respect to each operator \hat{U}_k and consequently, by chain rule, with respect to u_k . The update rule follows then directly from a standard follow-the-gradient approach as:

$$u_j^{i+1}(k\Delta t) = u_j^i(k\Delta t) + \eta \cdot \frac{\delta \mathcal{F}}{\delta u_j(k\Delta t)} \quad (1.77)$$

where now i indicates the iteration number and η the learning rate, quantifying the update momentum.

1.5.2. Krotov's method

Another popular gradient-based technique is given by Krotov's algorithm (firstly introduced in [48], applications to QOC further described in [49], [50]), which works in a very similar way to GRAPE but implements a slightly different update routine. In fact, this approach involves the definition of the control objective in a way that ensures the possibility to sequentially (rather than concurrently) update the controls in a non-linear fashion. In particular, this can be done by discretisation of the control pulse into M time steps which are updated one per iteration according to a prescription rule obtained by derivation of the control objective with respect to the control steps themselves. While this provides a deterministic and precise method for optimisation, its drawback is the slow-down of the FoM convergence while approaching the region of the optimum control due to the vanishing of the gradient.

1.5.3. Chopped random basis (CRAB) algorithm

The chopped random basis (CRAB) algorithm has been firstly introduced in [51] and more recently enhanced in [52]. As widely described in [45], it has become one of the most widely used QOC algorithms because of its capability of work well under experimental constraints, the possibility to access the usually trap-free landscape (and therefore favour convergence) and the adaptability to both open-loop and closed-loop optimisation schemes.

The key idea behind this method is based on the expansion of the control field $u(t)$ into a truncated basis made of N_c functions f_i :

$$u(t) = \sum_{i=1}^{N_c} c_i f_i(t) \quad (1.78)$$

The optimization functional $\mathcal{J}[\{c_i\}_i]$ is minimised by means of gradient-free direct search methods (e.g., the Nelder-Mead (NM) algorithm, see [53]) over a subspace of reduced dimension, aiming to find the best set of coefficients $\{c_i\}_{i=1, \dots, N_c}$. This approach allows to optimise of highly non-linear functionals and complex systems and is especially convenient when gradient-based algorithms are impossible

to apply, either because of analytic constraints or because of low efficiency. The main parameter to fix for this optimisation is the amplitude variation (AV), i.e. the size of the simplex to be investigated at each NM run. Larger AV allows to scan the parameter space faster, getting closer to the minima in less iterations; a smaller AV allows to better map the minima surroundings to find detailed solutions.

In principle, truncating the basis forces a constraint in the search space of the controls that can lead to sub-optimal (local) minima, named false-traps. These are represented by the actual minima of the reduced control problem that might in principle not coincide with the global minima of the original formulation. Therefore, a first step to overcome this phenomenon is the randomisation of the basis functions

While the choice of the basis is in principle arbitrary, a common choice is represented by the (truncated) Fourier basis, for which the randomisation procedure consists of a new random choice of frequencies and phases in the Fourier space. In this case, exploiting trigonometric function equation (1.78) becomes:

$$u(t) = \sum_{i=1}^{N_c/2} c_i \frac{\cos(\omega_i t)}{\Lambda(t)} + \sum_{i=(N_c/2)+1}^{N_c} c_i \frac{\sin(\omega_i t)}{\Lambda(t)} \quad (1.79)$$

where $\omega_i = \tilde{\omega}_i + \delta\omega_i \forall i$, being $\{\tilde{\omega}_i\}_i$ the principal harmonics and $\delta\omega_i$ is randomly chosen in $[-\omega_{max}, \omega_{max}]$.

Each random set of basis functions allows to explore a N_c dimensional subset of $\mathcal{C} \subset L_2$, the set of admissible controls (eventually infinite). The latter is usually defined by the experimental constraints, such as limited bandwidth or power. By iterative randomisation of the basis set we can therefore investigate the more and more regions of \mathcal{C} at a lower computational cost.

The randomisation procedure has been properly formalised in an efficient way in the so-called dressed CRAB (dCRAB) algorithm (see [52]), aiming to overcome at the best the false trap problem. The key idea in here is to proceed with a control expansion on a first set of basis states $\{f_i^0\}$ as in equation (1.78), perform an optimisation in the $span(\{f_i^0\})$ subspace of \mathcal{C} and then "dress" the resulting pulse (which will be either a global or a local optimal control) with new search directions given by a new basis set. The optimisation procedure within the same basis choice takes the name of super-iteration (SI). For the j -th SI it holds:

$$u^j(t) = c_0^j u^{j-1}(t) + \sum_{i=1}^{N_c} c_i^j f_i^j(t) \quad (1.80)$$

where $\{f_i^j(t)\}$ is the new set of basis functions and $u^{j-1}(t)$ is the optimised control pulse returned by the $(j-1)$ -th super-iteration. The coefficient c_0^j allows to move the optimisation on the same direction of the previous pulse, while $\{c_i^j\}_i$ allow to explore new directions.

For each dCRAB search the following parameters must be specified:

- Nature of the basis (e.g. Fourier)
- Dimension of the truncated basis N_c
- Number of super iterations
- Nature of the optimisation (search) algorithm to minimise the control objective (e.g. NM algorithm)
- Number of optimisation iterations within one SI
- Convergence criteria (threshold value for \mathcal{J} or minimal slope for local minima recognition)

In this work dCRAB has been chosen as optimisation algorithm because of its great performances while working in a gradient-free framework under the direct conditioning of the experiment. The NM algorithm has been tested to be the one suggesting best overall optimisation results in a noiseless scenario, while upon introduction of stochastic noise CMA-ES could be preferred [54]. As shown in section 3.2, dCRAB allows to reach good figure of merit results over simulations of the setup in various configurations. This makes the approach promising while moving towards experimental applications.

1.5.4. The QuOCS library

Specifically, for the numerical validation of this work we made use of the Quantum Optimal Control Suite (QuOCS) available for Python [55]. This tool has been designed to favour the implementation of QOC both in numerical, model-based approaches and real experimental optimisation.

QuOCS provides tool for both open-loop and closed-loop QOC, including optimisers as GRAPE and dCRAB as well as multiple search algorithms like Nelder Mead and CMA-ES. It also includes built-in classes for single-qubit simulations that allow to optimise gates in different conditions. As for our case, we independently built a class for two-qubit simulations using the standard of QuOCS and the instructions given in section 3.1. Then, we interface our simulation module with the built-in dCRAB optimiser of the Suite, adjusting the parameters for each run. Default functions have also been used to provide visualisation of the results.

Chapter 2

Efficient gate fidelity protocol

In section 1.4 we introduced the field of control theory and highlighted why it is relevant for (quantum) technologies. Moreover, in 1.4.4 we reported the main fidelity measures that can be exploited to define control objectives for the optimisation of quantum states preparation and quantum gates. Nevertheless, even if the metrics reported in equations 1.71-1.73 provide accurate tools for gate preparation, the experimental estimation of such quantities can be challenging for many-body quantum systems and would require a lot of effort in determining the full set of matrix elements of \hat{U} to compute the fidelity. Therefore, we can exploit alternative fidelity measures which are more experimentally accessible for the gate preparation, while keeping the aforementioned quantities for reference. In this section we describe the theoretical details of the tomography protocol we propose to compute a figure of merit which enables closed-loop QOC starting from experimental data. With the term *tomography* we indicate the characterisation of the quantum gate by means of preparation, evolution and measurement of some well-designed probe states.

A recent work showed how analysing the time evolution of just three well-designed states can suffice to perform open-loop optimisation of quantum gates regardless of the Hilbert space dimension [56]. In detail, the proposed fidelity is presented as follows:

$$\mathcal{F}_J = \sum_{i=1}^n \frac{w_i}{\text{Tr}[\hat{\rho}_i^2(0)]} \text{Re} \left\{ \text{Tr} \left[\hat{V} \hat{\rho}_i(0) \hat{V}^\dagger \hat{\rho}_i(T) \right] \right\} \quad (2.1)$$

where $n = 3$ can be reduced to $n = 2$ if the time evolution is unitary and preserves the population in the optimisation subspace (e.g., absence of decoherence). Restricting to the latter case, the definition of the two states to evolve must take into account the following: the first state fixes a basis $\{\varphi_i\}_i$ and verifies the gate action on the basis vectors. It must satisfy:

$$\hat{\rho}_1(0) = \sum_{i=0}^{d-1} \lambda_i \hat{P}_i \text{ with } \lambda_i \neq \lambda_j \text{ for } i \neq j \text{ and } \lambda_i \geq 0 \forall i \quad (2.2)$$

where $\hat{P}_i = |\varphi_i\rangle\langle\varphi_i|$ and d is the basis size. The second one is a totally rotated pure state used to detect phase errors:

$$\hat{\rho}_2(0) = \hat{P}_{tr} \text{ with } \hat{P}_{tr} = |\phi\rangle\langle\phi| \text{ such that } \langle\phi|\varphi_i\rangle \neq 0 \forall i \quad (2.3)$$

The additional third state to check population conservation in the subspace can be simply taken proportional to the identity in the subspace.

Please notice how the quantity in equation (2.1) reaches 1 when:

$$\hat{\rho}_i(T) = \hat{V} \hat{\rho}_i(0) \hat{V}^\dagger \forall i \quad (2.4)$$

where the left-hand side term represents the realisation of the evolved quantum states under the controls:

$$\hat{\rho}_i(T) \equiv \hat{U} \hat{\rho}_i(0) \hat{U}^\dagger \quad (2.5)$$

and the second defining the theoretical, expected action of the gate:

$$\hat{\rho}'_i \equiv \hat{V} \hat{\rho}_i(0) \hat{V}^\dagger \quad (2.6)$$

It is important to underline that the results achieved in [56], including the definition of equation (2.1), have been developed within the framework of numerical optimisation. In this work we analysed and proposed how to turn this computational tools into experimental routines for closed-loop QOC. In this regard, we recognise the real-part operator in equation (2.1) being a tool for numerical representation of an experimental result, which is indeed always a real number *per se*. Since this is the direction we aim to inspect, we ignore it in the following pages.

2.1. Application to CNOT and SWAP gates

In this section we discuss how to estimate the quantity in equation (2.1) with the experimental realisation, evolution and measurement of specific quantum states. The latter are chosen following the instructions in equations 2.2 and 2.3 in a compatible way with the experiment. Assuming a well defined characterisation of the four basis states and the possibility to perform the control within the decoherence time, we do not investigate information leakage and therefore presume unitary evolution. In this regime, just two quantum states are sufficient to estimate the gate fidelity. While a full analysis of the preparation and readout processes for our system is carried out in sections 3.1.1 and 3.1.3, respectively, here we give a more general definition of the figure of merit based on only three requirements: the possibility of initialising the system in $|0\rangle_{en}$, performing Rabi oscillations between states and accessing the populations of the basis states independently, all with high fidelity.

A suitable choice for the $\hat{\rho}_1(0)$ and $\hat{\rho}_2(0)$ states is:

$$\hat{\rho}_1(0) = \frac{1}{10} \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad \hat{\rho}_2(0) = \frac{1}{4} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \quad (2.7)$$

We focus on the real-case task of optimising a Controlled-NOT gate which flips the nuclear spin conditioned to the NV state. Setting the target gate to $\hat{V} = CNOT$, its action on the initial states is:

$$\begin{aligned}\hat{\rho}_1(0) &\xrightarrow{CNOT} \hat{\rho}'_1 \equiv \hat{V} \hat{\rho}_1(0) \hat{V}^\dagger = \frac{1}{10} \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 4 \end{pmatrix} \\ \hat{\rho}_2(0) &\xrightarrow{CNOT} \hat{\rho}'_2 \equiv \hat{V} \hat{\rho}_2(0) \hat{V}^\dagger = \frac{1}{4} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} = \hat{\rho}_2(0)\end{aligned}\quad (2.8)$$

i.e., the $CNOT$ gate inverts the population of states $|2\rangle$ and $|3\rangle$ while applied to $\hat{\rho}_1$ while acts as the identity on $\hat{\rho}_2$.

2.1.1. Definition of the best $\hat{\rho}_1$ form

The preparation of $\hat{\rho}_1(0)$ as in equation (2.7) has been chosen as a better alternative with respect to the more natural $\hat{\rho}_1(0) = \text{diag}(0.1, 0.2, 0.3, 0.4)$ because of the difficulties of the algorithm in discriminating between $CNOT$ and Identity gates.

Let us quantify this effect. Consider $\hat{V} = CNOT$ and the theoretical case of $\hat{\rho}_1$ being a classical mixture. If the time evolution is perfect then the operator which represents it (which depends on the controls) acts on the state as the target gate:

$$\hat{\rho}(T) = \hat{U} \hat{\rho}(0) \hat{U}^\dagger = \hat{V} \hat{\rho}(0) \hat{V}^\dagger \quad (2.9)$$

and therefore:

$$\frac{\text{Tr} \left[\hat{V} \hat{\rho}_1(0) \hat{V}^\dagger \hat{\rho}_1(T) \right]}{\text{Tr} \left(\hat{\rho}_1^2(0) \right)} = \frac{1}{0.3} \text{Tr} \left[\begin{pmatrix} 0.1 & 0 & 0 & 0 \\ 0 & 0.2 & 0 & 0 \\ 0 & 0 & 0.4 & 0 \\ 0 & 0 & 0 & 0.3 \end{pmatrix} \begin{pmatrix} 0.1 & 0 & 0 & 0 \\ 0 & 0.2 & 0 & 0 \\ 0 & 0 & 0.4 & 0 \\ 0 & 0 & 0 & 0.3 \end{pmatrix} \right] = 100\% \quad (2.10)$$

This applies to both versions of the state:

$$\frac{\text{Tr} \left[\hat{V} \hat{\rho}_1(0) \hat{V}^\dagger \hat{\rho}_1(T) \right]}{\text{Tr} \left(\hat{\rho}_1^2(0) \right)} = \frac{1}{0.3} \text{Tr} \left[\begin{pmatrix} 0.2 & 0 & 0 & 0 \\ 0 & 0.3 & 0 & 0 \\ 0 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0.4 \end{pmatrix} \begin{pmatrix} 0.2 & 0 & 0 & 0 \\ 0 & 0.3 & 0 & 0 \\ 0 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0.4 \end{pmatrix} \right] = 100\% \quad (2.11)$$

since

$$\text{Tr} \left(\hat{\rho}_1^2(0) \right) = \text{Tr} \left(\hat{\rho}_1^2(0) \right) = (0.1)^2 + (0.2)^2 + (0.3)^2 + (0.4)^2 = 0.3 \quad (2.12)$$

But what would we measure for the quantity above if we were to mistakenly find controls that make the time evolution act as the Identity on the system? In this case we would have:

$$\hat{\rho}(T) = \mathbb{1} \hat{\rho}(0) \mathbb{1}^\dagger = \hat{\rho}(0) \quad (2.13)$$

so that:

$$\frac{\text{Tr} \left[\hat{V} \hat{\rho}_1(0) \hat{V}^\dagger \hat{\rho}_1(T) \right]}{\text{Tr} \left(\hat{\rho}_1^2(0) \right)} = \frac{1}{0.3} \text{Tr} \left[\begin{pmatrix} 0.1 & 0 & 0 & 0 \\ 0 & 0.2 & 0 & 0 \\ 0 & 0 & 0.4 & 0 \\ 0 & 0 & 0 & 0.3 \end{pmatrix} \begin{pmatrix} 0.1 & 0 & 0 & 0 \\ 0 & 0.2 & 0 & 0 \\ 0 & 0 & 0.3 & 0 \\ 0 & 0 & 0 & 0.4 \end{pmatrix} \right] = 96.7\% \quad (2.14)$$

while:

$$\frac{\text{Tr} \left[\hat{V} \hat{\rho}_1(0) \hat{V}^\dagger \hat{\rho}_1(T) \right]}{\text{Tr} (\hat{\rho}_1^2(0))} = \frac{1}{0.3} \text{Tr} \left[\begin{pmatrix} 0.2 & 0 & 0 & 0 \\ 0 & 0.3 & 0 & 0 \\ 0 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0.4 \end{pmatrix} \begin{pmatrix} 0.2 & 0 & 0 & 0 \\ 0 & 0.3 & 0 & 0 \\ 0 & 0 & 0.4 & 0 \\ 0 & 0 & 0 & 0.1 \end{pmatrix} \right] = 93\% \quad (2.15)$$

It is immediate to verify that the (relative) difference between the *CNOT* evolution and the Identity evolution is 3.3% in the case of $\hat{\rho}_1$, while it grows to 7% for $\hat{\rho}_1$. Hence, we can conclude that the latter form displays better sensitivity while discerning between these gates. This is important in our case: as illustrated in chapter 3, in many simulations we start from a blind guess for the controls corresponding to turning off the whole control apparatus and in this regime the unitary evolution coincides with the identity. Therefore, it is important to distinguish between the action of *CNOT* and Identity gates to escape the initial, local minimum of the fidelity. In principle, the populations of the states subject to flip can be selected to be maximally different to enhance this effect under the non-degenerate populations requirement.

2.1.2. Experimental optimisation of the action on the basis

Fixing to $n = 2$ the number of states, the fidelity proposed in equation (2.1) is a weighted sum of two terms, each one checking correspondence between the time evolution and the quantum gates under a specific perspective. The first one involves the preparation, evolution and measure of $\hat{\rho}_1$, analysing the action of time evolution on all the basis states at once. For the initial state proposed in equation (2.7) the theoretical, expected evolution reported in equation (2.8) is a diagonal, mixed state. This allows to simplify the term for $\hat{\rho}_1$ in the fidelity measure since:

$$\begin{aligned} \text{Tr} \left[\hat{V} \hat{\rho}_1(0) \hat{V}^\dagger \hat{\rho}_1(T) \right] &= \text{Tr} \left[\hat{\rho}'_1 \hat{\rho}_1(T) \right] \\ &= \sum_{i=0}^3 \langle i | \hat{\rho}'_1 \hat{\rho}_1(T) | i \rangle \\ &= \sum_{i=0}^3 (\hat{\rho}'_1)_{ii} \langle i | \hat{\rho}_1(T) | i \rangle \\ &= \sum_{i=0}^3 (\hat{\rho}'_1)_{ii} (\hat{\rho}_1(T))_{ii} \end{aligned} \quad (2.16)$$

which holds because $\langle i | \hat{\rho}'_1 = (\hat{\rho}'_1)_{ii} \langle i |$ for $i = 0, \dots, 3$ since $\hat{\rho}'_1$ is a diagonal matrix. It is important to highlight that this is true as we consider in here the evolution of the theoretical mixed initial state $\hat{\rho}_1(0)$, which is perfectly diagonal. As we will soon discuss, this is usually not the case for $\hat{\rho}_1(T)$, that represents in this picture the time-evolved (experimental) state of the system. Therefore, equation (2.16) suggests that we can perform a closed-loop estimate of the contribute to the gate fidelity related to $\hat{\rho}_1$ by preparing the system in the initial state, letting it

evolve under the controls action, measure the population of each basis state (i.e., $\{(\hat{\rho}_1(T))_{ii}\}_{i=0,\dots,3}$) and then multiply these experimental values by the theoretical ones ($\{(\hat{\rho}'_1)_{ii}\}_i$).

For the experimental setup under investigation the realisation of the state $\hat{\rho}_1(0)$ cannot be done with ease due to the mixed nature of the state itself. In fact, assuming a default starting state $\hat{\rho}_0 = |0_e 0_n\rangle\langle 0_e 0_n|$ for the two-qubit system, there is no way to reach $\hat{\rho}_1(0)$ by means of unitary transformations. Hence, we propose an approach based on random-phase pulses which turns out to produce our desired state as an average of multiple realisations. For the sake of simplicity, we schematize the protocol in a single-qubit case.

Consider an initial and a final state given by:

$$\hat{\rho}(t_0) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = |0\rangle\langle 0| \quad \hat{\rho}(t_f) = \begin{pmatrix} \lambda_0 & 0 \\ 0 & \lambda_1 \end{pmatrix} \quad (2.17)$$

Generally speaking, we can move from the initial $|\psi_0\rangle = |0\rangle$ towards the $|1\rangle$ state by applying a tailored pulse $\hat{P}_\varphi(\theta)$ of given length which has the effect of rotating the pure state to a configuration with the desired diagonal terms and non-zero coherences:

$$\hat{\rho}(t_0) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = |0\rangle\langle 0| \xrightarrow{\hat{P}_\varphi(\theta)} \hat{\rho}^\theta(t_f) = \begin{pmatrix} \lambda_0 & \kappa \\ \kappa^* & \lambda_1 \end{pmatrix} \quad (2.18)$$

where $\cos(\varphi) = 2\lambda_0 - 1$ and $\kappa = \sin(\varphi)\cos(\theta) - i\sin(\varphi)\sin(\theta) = \sin(\varphi)e^{-i\theta}$, with reference to the notation in figure 2.1 (a). We recall that φ is the altitude coordinate and θ is the pulse phase which determines the azimuth coordinate on the Bloch sphere surface. The approach here is the same underlying Rabi oscillations, fully described in section 1.2.2.

What happens if we consider multiple realisations of the state by starting from $|0\rangle$ and applying pulses $\hat{P}_\varphi(\theta)$ with the same value of φ but random θ ? In section 1.1.2 we introduced the Bloch sphere visualisation tool for single qubit states. In particular, a direct result from equation (1.14) is:

$$\lambda_0 = \frac{1+z}{2} \text{ and } \lambda_1 = \frac{1-z}{2} \quad (2.19)$$

being λ_0 and λ_1 the probabilities associated to the measurement of the $|0\rangle$ or $|1\rangle$ states (diagonal terms of the density matrix), and z is the third component of the Bloch vector \vec{r} , i.e. the state projection onto the vertical axis Z . As a result, since all the states laying on the same XY horizontal plane of the Bloch sphere share the same z coordinate, they also have identical probabilities for $|0\rangle$ and $|1\rangle$. In other words, they are states with the same diagonal.

Starting from $|0\rangle$, i.e. the north pole of the Bloch sphere, the angle $\varphi = \Omega t$ is uniquely determined by the pulse amplitude and duration (see equation (1.43) for reference). This is the only parameter responsible for changing the value of z . Hence, adopting the same φ for all the realisations we ensure to be in the previous scenario.




Figure 2.1: (a) State representation on the Bloch sphere. (b) Visualisation of the random realisation of single qubit states sharing the same diagonal on the Bloch sphere. Initial state is $|0\rangle$ (green). Upon evolution with a pulse of tailored intensity and duration, the state is rotated to reach a specific value of z (here $z = -0.5$, orange state). The random phase of the pulse results in a random state position on the horizontal circumference at z (blue dots, here $N = 100$ samples). The average of many uniformly distributed realisations results in a mixed state with Bloch vector $\vec{r} = (0, 0, z)^T$ (blue).

If we now consider N_θ realisations of the state obtained by sampling as many random values for the control field phase ($\theta \sim \mathcal{U}[0; 2\pi]$), we obtain a collection of pure states uniformly distributed on the same horizontal plane of the Bloch sphere. Then, considering an average over all these states:

$$\begin{aligned}
 \left\langle \begin{pmatrix} \lambda_0 & \kappa \\ \kappa^* & \lambda_1 \end{pmatrix} \right\rangle_\theta &= \begin{pmatrix} \langle \lambda_0 \rangle_\theta & \langle \kappa \rangle_\theta \\ \langle \kappa^* \rangle_\theta & \langle \lambda_1 \rangle_\theta \end{pmatrix} = \\
 &= \begin{pmatrix} \lambda_0 & \sin(\varphi) \langle e^{-i\theta} \rangle_\theta \\ \sin(\varphi) \langle e^{i\theta} \rangle_\theta & \lambda_1 \end{pmatrix} \\
 &= \begin{pmatrix} \lambda_0 & 0 \\ 0 & \lambda_1 \end{pmatrix}
 \end{aligned} \tag{2.20}$$

since the diagonal terms do not depend on θ and $\langle e^{\pm i\theta} \rangle_\theta = 0$ for θ uniformly distributed in $[0; 2\pi)$. As a result, one can obtain an equivalent description of the dynamics of the mixed state as an average of the evolved versions of many pure states with randomised coherences. A visualisation of these concepts is proposed in figure 2.1.

The generalisation of this idea to the multi-qubit case is straightforward. Consider the control Hamiltonian in the interaction picture defined in 1.56: assuming that upon initialisation the system is in the $\hat{\rho}_0 = |0_e 0_n\rangle\langle 0_e 0_n| \equiv |0\rangle\langle 0|$ state, it is sufficient to consider sequentially two-states subsystems and apply tailored pulses in the aforementioned way, turning everything off but the desired control. Details of the pulse sequence and initialisation process are described in section 3.1.1. Such

procedure allows us to obtain a two-qubit state

$$\hat{\rho}_1^\theta(0) = \begin{pmatrix} \lambda_0 & \kappa_1 & \kappa_2 & \kappa_3 \\ \kappa_1^* & \lambda_1 & \kappa_4 & \kappa_5 \\ \kappa_2^* & \kappa_4^* & \lambda_2 & \kappa_6 \\ \kappa_3^* & \kappa_5^* & \kappa_6^* & \lambda_3 \end{pmatrix} \quad (2.21)$$

with $\{\kappa_i\}_{i=1\dots 6}$ being the coherences with random phases in the set $\theta \equiv \{\theta_i\}_{i=1\dots 6}$. Again, the average of infinitely many of such realisations results in the desired $\hat{\rho}_1(0)$ state: in practical conditions, the number $N_{samples}$ of realisations required for the off-diagonal terms to average to zero over a finite number of samples will be one of the tunable parameters of the algorithm.

The previous approach allows to experimentally determine the quantity in equation (2.16) simply by averaging the estimation over the whole set of phases θ used in the preparation step. In fact, considering the previously introduced state $\hat{\rho}_1^\theta$ which satisfies $\langle \hat{\rho}_1^\theta(0) \rangle_\theta = \hat{\rho}_1(0)$ it holds:

$$\begin{aligned} \left\langle \text{Tr} \left[\hat{V} \hat{\rho}_1(0) \hat{V}^\dagger \hat{\rho}_1^\theta(T) \right] \right\rangle_\theta &= \left\langle \text{Tr} \left[\hat{\rho}'_1 \hat{\rho}_1^\theta(T) \right] \right\rangle_\theta \\ &= \text{Tr} \left[\hat{\rho}'_1 \left\langle \hat{\rho}_1^\theta(T) \right\rangle_\theta \right] \\ &= \text{Tr} \left[\hat{\rho}'_1 \left\langle U \hat{\rho}_1^\theta(0) U^\dagger \right\rangle_\theta \right] \\ &= \text{Tr} \left[\hat{\rho}'_1 U \left\langle \hat{\rho}_1^\theta(0) \right\rangle_\theta U^\dagger \right] \\ &= \text{Tr} \left[\hat{\rho}'_1 U \hat{\rho}_1(0) U^\dagger \right] \\ &= \text{Tr} \left[\hat{\rho}'_1 \hat{\rho}_1(T) \right] \\ &= \text{Tr} \left[\hat{V} \hat{\rho}_1(0) \hat{V}^\dagger \hat{\rho}_1(T) \right] \end{aligned} \quad (2.22)$$

Again, this holds since the $\hat{\rho}'_1$ term above represents the action of the gate on the theoretical state with exact 0 on the off-diagonal terms which is θ -independent and we are dealing with linear operations only.

We can finally combine the results of equations 2.16 and 2.22 to obtain a theoretically-driven experimental routine:

$$\begin{aligned} \text{Tr} \left[\hat{V} \hat{\rho}_1(0) \hat{V}^\dagger \hat{\rho}_1(T) \right] &\stackrel{(2.22)}{=} \left\langle \text{Tr} \left[\hat{V} \hat{\rho}_1(0) \hat{V}^\dagger \hat{\rho}_1^\theta(T) \right] \right\rangle_\theta \\ &\stackrel{(2.16)}{=} \left\langle \sum_{i=0}^3 (\hat{\rho}'_1)_{ii} \left(\hat{\rho}_1^\theta(T) \right)_{ii} \right\rangle_\theta \\ &\simeq \frac{1}{N_\theta} \sum_\theta \sum_{i=0}^3 (\hat{\rho}'_1)_{ii} \left(\hat{\rho}_1^\theta(T) \right)_{ii} \end{aligned} \quad (2.23)$$

Equation (2.23) summarises the experimental protocol we suggest for the gate fidelity measure over the basis state:

- Define a suitable set of diagonal terms, in accordance with equation (2.2)
- Define a sequence of pulses starting from $|0\rangle_{en}$ which prepares pure states with the desired populations on the diagonal but uses random phases θ_i
- Prepare N_θ several realisations of such states with random coherences and for each of them perform the time evolution and readout the populations at the final time
- Multiply each experimental population with the theoretical value and average over all the N_θ results

This routine allows to access the quantity $\text{Tr} [\hat{\rho}'_1 \hat{\rho}_1 (T)]$, which shall be normalised in the interval $[0, 1]$ by dividing it by a factor $\text{Tr} (\hat{\rho}_1^2 (0))$ (theoretical value).

A final remark: not only this approach allows to experimentally access the quantity of interest, but it also does it in an efficient way. In fact, the output of a qubit measurement is a random variable and to extract any reliable information from the system multiple realisations and measures (at least $N_{rep} = \mathcal{O}(10^4)$ in our case) are required. Since $N_\theta = N_{rep}$ realisations of $\hat{\rho}_1^\theta$ are in most cases more than enough for a reliable representation of $\hat{\rho}_1 (0)$, the random-phase approach has the exact same complexity as the preparation and measure of the mixed $\hat{\rho}_1 (0)$ state, which is extremely convenient because it only requires to work with unitary evolution and pure states. This is potentially a great experimental advantage.

While it would be possible to study the action of the control gate on the N_{basis} individual basis states (process tomography, more on this in section 2.1.3), this would require N_{rep} realisations and measures for each of them, for a total of $N_{basis} \times N_{rep}$ measures, which can become expensive for large systems. In this respect, the protocol illustrated in this section would allow to save a factor N_{basis} of operations.

2.1.3. Alternative approaches: process tomography and the two- $\hat{\rho}_1$ variant

As underlined in [56], the analytical form of $\hat{\rho}_1 (0)$ allows to set the computational basis and verify at once the action of the controlled gate over the basis states $\{|\varphi_i\rangle\}_i$. Of course, the same can be achieved by preparing, evolving and measuring the pure basis states directly. This approach resembles the one adopted in Quantum Process Tomography (QPT): since we can linearly decompose any quantum state in the superposition of basis states, it is sufficient to verify the correct action of the control gate over these and the validation transfers to any quantum state. This approach can be more or less convenient than the preparation and measure of the $\hat{\rho}_1$ state according to the situation. For example, tomography is more suitable for small systems where the basis states can be prepared with ease and the state tomography for the readout is efficient. On the other hand, for many-qubit systems this procedure becomes expansive due to the exponential scaling of the number of basis states.

For the four-states system under investigation both the $\hat{\rho}_1$ state propagation and basis evolution analysis present some important drawbacks, such as the necessity

of adopting a randomised scheme or a large number of measurements for statistics. In here we propose an intermediate approach involving the decomposition of $\hat{\rho}_1(0)$ as the linear combination of two mixed states:

$$\begin{aligned}\hat{\rho}_1(0) &= c_a \hat{\rho}_1^a(0) + c_b \hat{\rho}_1^b(0) \\ &= c_a \begin{pmatrix} \alpha & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1-\alpha & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} + c_b \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \beta & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1-\beta \end{pmatrix}\end{aligned}\quad (2.24)$$

where $c_a + c_b = 1$ ensures $\text{Tr}(\hat{\rho}_1(0)) = 1$ and we can relate this notation to the one in equation (2.2) by setting,

$$c_a \cdot \alpha = \lambda_0 \quad c_b \cdot \beta = \lambda_1 \quad c_a \cdot (1 - \alpha) = \lambda_2 \quad c_b \cdot (1 - \beta) = \lambda_3 \quad (2.25)$$

With respect to the $\hat{\rho}_1$ method described in section 2.1.2, this approach requires a larger experimental effort in the realisation and measurement of an additional state but it reduces the number of samples needed for randomized cancellation of off-diagonal terms of orders of magnitude. In fact, while the preparation of $\hat{\rho}_1(0)$ requires the use of three random-phase pulses which imply the creation of non-zero coherences in the whole density matrix (equation (2.21)), a pure state with diagonal terms as $\hat{\rho}_1^a(0)$ can be prepared by the means of an unique tuned random-phased pulse starting from $|0\rangle\langle 0|$ (and similarly for $\hat{\rho}_1^b(0)$ starting from $|1\rangle\langle 1|$). As a result:

$$\hat{\rho}_0 \xrightarrow{\hat{U}_{t\alpha}^{2\theta}(\theta)} \begin{pmatrix} \alpha & 0 & |\kappa|e^{-i\theta} & 0 \\ 0 & 0 & 0 & 0 \\ |\kappa|e^{i\theta} & 0 & 1-\alpha & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}\quad (2.26)$$

where now ten out of twelve off-diagonal terms are exactly zero. The only θ -dependent terms average at zero with few samples: $N_\theta = \mathcal{O}(10^1)$ is enough to reach the same value of reference fidelity \mathcal{F}_{sm} as the original algorithm with $N_\theta = \mathcal{O}(10^{2+})$, over similar time and in the same optimisation conditions. While this ensures a smaller computational effort, the number of iterations required for convergence of the pulse optimisation might be larger and eventually the choice of the most convenient approach should be investigated for each method in the specific use-case.

2.1.4. Experimental optimisation of the relative phases

As for the second term in equation (2.1), we found that a particularly efficient protocol can be designed for entangling gates on two-qubit systems such as Controlled-NOT (CNOT) and SWAP:

$$CNOT = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad SWAP = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}\quad (2.27)$$

In fact, for all the gates that conditionally invert the populations of two basis states the phase-preservation qualities of the proposed \hat{U} operator can be estimated in

a single measure (up to repetitions for statistics). This also applies to the multi-qubit version of such gates, such as the Toffoli gate, while the efficiency for phase-shifting gates is yet to be studied.

In detail, being $\hat{\rho}_2(0)$ a pure state it exists a unitary operator \hat{M} such that, starting in the initialized $\hat{\rho}_0 \equiv |0_e 0_n\rangle\langle 0_e 0_n| = |0\rangle\langle 0|$ state:

$$\hat{\rho}_2(0) = \hat{M}\hat{\rho}_0\hat{M}^\dagger \quad (2.28)$$

For example, a suitable form for \hat{M} is a chain pulses among states such as:

$$\hat{M} = U_{\pi/4}^{13} \left(\frac{3}{2}\pi\right) U_{\pi/3}^{23} \left(\frac{\pi}{2}\right) U_{\pi/2}^{02} \left(\frac{3}{2}\pi\right) \quad (2.29)$$

where

$$\begin{aligned} U_\phi^{nm}(\theta) &= \exp(-iH^{nm}t_{nm}) \\ H^{nm} &= u_{nm}e^{i\theta}|n\rangle\langle m| + u_{nm}e^{-i\theta}|m\rangle\langle n| \end{aligned} \quad (2.30)$$

and $t_{nm} = \phi u_{nm}^{-1}$.

Therefore, recollecting that $\hat{V}\hat{\rho}_2(0)\hat{V}^\dagger = \hat{\rho}_2(0)$ (equation (2.8)), it holds:

$$\begin{aligned} \text{Tr} \left[\hat{V}\hat{\rho}_2(0)\hat{V}^\dagger\hat{\rho}_2(T) \right] &= \text{Tr} [\hat{\rho}_2(0)\hat{\rho}_2(T)] \\ &= \text{Tr} \left[\hat{M}\hat{\rho}_0\hat{M}^\dagger\hat{\rho}_2(T) \right] \\ &= \text{Tr} \left[\hat{\rho}_0\hat{M}^\dagger\hat{\rho}_2(T)\hat{M} \right] \\ &= \left(\hat{M}^\dagger\hat{\rho}_2(T)\hat{M} \right)_{00} \end{aligned} \quad (2.31)$$

where we exploited the cyclic property of the Trace and recognised that the only non-zero element of $\hat{\rho}_0$ is $(\hat{\rho}_0)_{00} = 1$. Ideally, this corresponds to prepare the $\hat{\rho}_2(0)$ state, evolve it with the trial controlled \hat{U} operator and bring it back to the initial condition: since the action of a *CNOT* gate leaves $\hat{\rho}_2$ untouched, we measure the fidelity as the overlap between the final state and the initial one. It is worth remarking how this works since we are targeting gates that do not introduce phase shifts among states, but just flips. On the other hand, this works for gates in this form on arbitrarily large qubit systems.

To summarise, to verify that the optimisation proposal preserves relative phases between states it is enough to:

- Define a sequence of pulses to move from $|0\rangle\langle 0|$ to $\hat{\rho}_2(0)$ (e.g., \hat{M} operator in equation (2.29))
- Perform time evolution for the $\hat{\rho}_2$ state
- Return to the initial basis (by the means of inverse transformation \hat{M}^{-1})
- Measure the population of the $|0\rangle\langle 0|$ state

and this will allow to access the quantity $\text{Tr} [\hat{\rho}'_2 \hat{\rho}_2 (T)]$ to be used for the computation of \mathcal{F}_J .

To sum up, contributions coming from the two states combine as follows: the second state verifies that no local phases are added between states, allowing only for population swaps; the first one checks whether the population rearrangement caused by the controls matches the one expected from the target gate. Fixing the weights w_i in equation (2.1) allows to balance between these contributions (quantitative analysis in section 3.2.4).

Chapter 3

Numerical validation: simulations

In order to verify numerically the claims above, we create a digital twin of the quantum system and interface it with a Quantum Optimal Control (QOC) algorithm. In particular, we exploit the dCRAB algorithm (details in section 1.5.3) available in the Quantum Optimal Control Suite (QuOCS) library in Python (section 1.5.4, [55]). The following pages describe our approach and the technical details of the code. The control objective is simply set to be $1 - \mathcal{F}$.

To understand whether dCRAB can be considered a good optimiser for our task, we test it on our system model with fidelity measure given by equation (1.72): over few hundreds numerical closed-loop optimisation iterations with mild converging requirements, a final infidelity of $\mathcal{O}(10^{-10})$ has been achieved.

Coming to the real closed-loop simulation, the whole procedure described in the previous sections has been summarised in algorithm 1. Our investigation covered:

- Different choices for $\hat{\rho}_1$ populations $\{\lambda_0, \dots, \lambda_3\}$;
- Various sets of weights $\{w_i\}_{i=1,2}$;
- Changing numbers of dCRAB Super-Iterations (SIs) and stopping criteria;
- Tuning the NM algorithm amplitude variation to change speed and accuracy of the investigation of the parameters space;
- Setting the number N_θ of random-phase realisations of $\hat{\rho}_1$.

As we observed, all these affect greatly the performance of the optimisation process. Therefore, a cautious tuning of these parameters is required in order to reach the best convergence and avoid local minima. Details of the study of $\{\lambda_i\}_i$ are reported in section 2.1.1 while an analysis of all the other hyper-parameters is carried out in section 3.2.4.

3.1. Simulation design

As mentioned, this work relies on the QuOCS open-source library in Python for the optimisation component [55]. The two-qubit framework has been designed specifically for this purpose, as well as the details of the time evolution operators, the controls and the proposed figure of merit.

3.1.1. States initialisation

Here we review the numerical procedure that enables the preparation of $\hat{\rho}_1$ and $\hat{\rho}_1^\theta$ following the instructions in sections 2.1.2 and 2.1.4, respectively. The simulation module assumes that the experiment can be reset to the initial $|0_e 0_n\rangle\langle 0_e 0_n|$ state with relative ease and high fidelity (which is true for the system under investigation). Starting from here, to obtain the $\hat{\rho}_1^\theta(0)$ term in equation (2.21) a total of three pulses with random phases $\{\theta_i\}_i$ is required. Therefore, the simulation draws random values from a uniform distribution $\theta_i \sim \mathcal{U}(0; 2\pi)$ and prepares a randomised state with fixed diagonal terms (0.2, 0.3, 0.4, 0.1) (whose choice has been discussed in section 2.1.1). This can be done in different ways, depending on the possibility to distinguish between transitions $|0\rangle \longleftrightarrow |2\rangle$ and $|1\rangle \longleftrightarrow |3\rangle$. Following the formalism introduced in equation (2.1.4), in the first case it is sufficient to apply three tailored pulses:

$$U_{\phi_1}^{02}(\theta_1), U_{\phi_2}^{23}(\theta_2), U_{\phi_3}^{13}(\theta_3)$$

where $\phi_1 = 0.8 \cdot \pi$, $\phi_2 = 0.375 \cdot \pi$, $\phi_3 = 0.8 \cdot \pi$ (3.1)

where the application order is left to right. In practice, we can tune the product $\phi \equiv \Omega_R t$ between the Rabi frequency and the pulse duration in equation (1.44) for each pulse so to transport the right amount of population among states.

It might be the case that the NV transition is non-selective, i.e. no discrimination can be done between transitions $|0\rangle \longleftrightarrow |2\rangle$ and $|1\rangle \longleftrightarrow |3\rangle$. Here, instead of pulses $U_\phi^{02}(\theta)$, $U_\phi^{13}(\theta)$, we can only apply a single pulse $U_\phi^{NV}(\theta)$ between the $|0\rangle_e$ and $|1\rangle_e$ states of the electrons. Hence, an alternative strategy for $\hat{\rho}_1(0)$ preparation can be:

$$U_{\phi_1}^{NV}(\theta_1), U_{\phi_2}^{23}(\theta_2), U_{\phi_3}^{NV}(\theta_3), U_{\phi_4}^{23}(\theta_4) \quad (3.2)$$

where now we would need four random-phase pulses and the values of $\{\phi_i\}_i$ are to be estimated so to match the required state populations. An alternative solution could be to exploit a surrounding ancillary qubit. Anyways, for the sake of our simulations we assume to be in the first scenario and created controls to move from $|0\rangle\langle 0|$ to $\hat{\rho}_1^\theta(0)$ on a timescale of $\mathcal{O}(100 \mu s)$ and pulse amplitude of $\mathcal{O}(10 \text{ kHz})$.

Intuitively, if instead of $\hat{\rho}_1$ for the standard protocol the goal would be the preparation of $\hat{\rho}_1^a, \hat{\rho}_1^b$ for the two- $\hat{\rho}_1$ variant the reasoning is the very same. As for $\hat{\rho}_1^a$ it is sufficient to

As for the preparation of $\hat{\rho}_2(0)$, a similar scenario arises: in case of selective MW transitions a possible pulse sequence to be applied to $|0\rangle\langle 0|$ is the one reported in equation (2.29) (\hat{M} operator)

$$U_{\pi/2}^{02}(\frac{3}{2}\pi), U_{\pi/3}^{23}(\frac{\pi}{2}), U_{\pi/4}^{13}(\frac{3}{2}\pi) \quad (3.3)$$

while if only non-selective transitions are applicable then we can perform:

$$U_{\pi}^{NV}(0), U_{\pi/2}^{23}(0), U_{\pi/2}^{NV}(0) \quad (3.4)$$

Please notice how in this case no randomised phases are required. Also in this case the error due to imperfect pulses and readout is summarised as a random variable in the final measurement. Nevertheless, the error propagation is actually different in this case (no need to consider average over several phase-randomised states and single measurement for the output).

3.1.2. Time evolution operator and controls

The initialisation procedure described above allows to simulate the preparation of the required states for the control optimisation. Once these are prepared, it is necessary to make them evolve in time to compare the actual evolution with the target gate's one. Therefore, we build the time evolution operator \hat{U} as fully described in section 1.1.5.

Once the time grid $\{t_j\}_{j=1,\dots,N}$ is designed, the optimisation module returns the best set of controls and parameters as a vector of size N . By plugging these values in the system's Hamiltonian (equation (1.56)), we can build a vector of Hamiltonians computed at each timestep $\{H(t_j)\}_j$. Then, by performing iterative exponentialisation and multiplication following the time ordering of the operators, we can create the Δt time evolution operators and ultimately the full time operator $\hat{U}(T)$ as described in equation (1.36).

As mentioned, such approach works if and only if the timestep width Δt is chosen small enough, i.e. if the local approximation of constant controls is valid. In our case, tests have been run to verify this constrain: the approximation works for $\Delta t = \mathcal{O}(10^{-3})$ (a.u.). For the simulations in section 3.2, the total pulse duration is set to $T = 10^{-3}$ s while $N = \mathcal{O}(10^2)$ allows for both small time steps and sustainable computational duration.

After its computation, the \hat{U} operator has to be applied to the test states.

3.1.3. Measurement

Technical considerations and readout alternatives

The most challenging part of the protocol, together with the state preparation, is represented by the measurement procedure. Generally speaking, for NV centers in diamond the most common measure operation would be given by optical readout, as fully described in section 1.2.1. However, this can lead to some issues

when trying to access the nuclear spin states, given the fact that in principle we are optically blind to them and therefore population measurements can be performed only by swapping the state of the electronic and nuclear spin before readout, via a combination of optical readouts and MW/RF π -pulses.

Experimentally, in order to do this we must be able to address selective transitions between states $|0\rangle \longleftrightarrow |2\rangle$ and $|1\rangle \longleftrightarrow |3\rangle$ during the readout phase. In detail, this is possible under specific requirements for the NV center we are dealing with. In section 1.2.1 we described how ODMR can be used to characterise different resonance frequencies (example in figure 1.6). We might exploit this technique to analyse $m_s = 0 - m_s = -1$ transition frequencies in the 3A_2 and 3E levels. If these are spaced in the spectrum enough to be distinguished but not enough not to be able to drive both levels altogether with a wide bandwidth pulse, then we could perform selective π -pulses with a fine-tuned, short bandwidth pulse. In this sense, a weak microwave pulse can flip the NVs spin state selectively on the nuclear spin state, which is experimentally easy to do in the case of large couplings.

The previous method allows the use of optical readout to access the four populations of the basis states, but it also imposes to find a suitable NV candidate. An alternative to optical readout is given by measurements via charge state [57]. In particular, this would allow the user to access directly the single-state population of $|0\rangle_{en}$ and, upon population inversion among the basis states. Nevertheless, also charge state readout is possible only for some systems and in particular it is especially efficient for experiments at cryogenic temperatures.

The choice of the best readout process in the experiment is then to be decided according to the apparatus specifics. In our simulation we assume to be dealing with an NV center which allows to readout the whole set of populations via optical readout. Indeed, we propose a measurement routine which enables to infer the four diagonal terms of the density matrix in a single run, with a factor 3 speedup with respect to schemes which address each term individually.

Modeling of the readout scheme

Formally, the full system's state before the measurement is given by the 4×4 density matrix:

$$\hat{\rho} = \begin{pmatrix} \rho_{00} & \rho_{01} & \rho_{02} & \rho_{03} \\ \rho_{10} & \rho_{11} & \rho_{12} & \rho_{13} \\ \rho_{20} & \rho_{21} & \rho_{22} & \rho_{23} \\ \rho_{30} & \rho_{31} & \rho_{32} & \rho_{33} \end{pmatrix} \quad (3.5)$$

and we remember that for the sake of our protocol we must learn the diagonal terms of $\hat{\rho}$ individually.

The reduced density matrix of the electronic spin is given by:

$$\hat{\rho}_e = \text{Tr}_n(\hat{\rho}) = \begin{pmatrix} \rho_{00} + \rho_{11} & \rho_{02} + \rho_{13} \\ \rho_{20} + \rho_{31} & \rho_{22} + \rho_{33} \end{pmatrix} \quad (3.6)$$

and the nuclear spin one is given by:

$$\hat{\rho}_n = \text{Tr}_e(\hat{\rho}) = \begin{pmatrix} \rho_{00} + \rho_{22} & \rho_{01} + \rho_{23} \\ \rho_{10} + \rho_{32} & \rho_{11} + \rho_{33} \end{pmatrix} \quad (3.7)$$

Performing optical readout on the electronic subsystem allows to access the diagonal values of its reduced density matrix:

$$\begin{aligned} \alpha &\equiv \rho_{00} + \rho_{11} \\ \beta &\equiv \rho_{22} + \rho_{33} \end{aligned} \quad (3.8)$$

After the measurement and the consequent reset of the electronic component, its state collapses into:

$$\hat{\rho}_e = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad (3.9)$$

while the nuclear component remains the same. The full quantum state is now:

$$\hat{\rho} = \begin{pmatrix} \mathbf{M} & \mathbf{0}_2 \\ \mathbf{0}_2 & \mathbf{0}_2 \end{pmatrix} \quad (3.10)$$

where:

$$\mathbf{M} = \begin{pmatrix} \rho_{00} + \rho_{22} & \rho_{01} + \rho_{23} \\ \rho_{10} + \rho_{32} & \rho_{11} + \rho_{33} \end{pmatrix} \equiv \begin{pmatrix} m_{00} & m_{01} \\ m_{10} & m_{11} \end{pmatrix} \quad (3.11)$$

Now, by applying a *CNOT* gate with control qubit the nuclear one and target qubit the electronic one:

$$C_nNOT_e = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \quad (3.12)$$

the resulting state would be:

$$\hat{\rho}' = \begin{pmatrix} m_{00} & 0 & 0 & m_{01} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ m_{10} & 0 & 0 & m_{11} \end{pmatrix} \quad (3.13)$$

leading to reduced density matrices:

$$\hat{\rho}'_e = \begin{pmatrix} m_{00} & 0 \\ 0 & m_{11} \end{pmatrix} = \hat{\rho}'_n \quad (3.14)$$

As a result, performing again optical readout on the electronic spin allows to access:

$$\begin{aligned} \gamma &\equiv m_{00} = \rho_{00} + \rho_{22} \\ \delta &\equiv m_{11} = \rho_{11} + \rho_{33} \end{aligned} \quad (3.15)$$




Figure 3.1: Scheme of the routine to access the population of the four basis states via Optical Readout (OR) on the NV spin. First, OR is performed to measure the populations of the NV states and simultaneously reset it to the $|0\rangle_n$ state (green pulse on the left-hand side). Then, a selective π -pulse on the NV conditioned to state of the nuclear spin (effective C_nNOT_e gate) transfers information regarding the ^{13}C spin state to the NV (orange pulse). Finally, another OR reads again the populations of the NV (second green pulse). Combining the four measurement outcomes allows to reconstruct the density matrix diagonal.

Now, within the formalism of density matrices, in case of unitary evolution of the system (i.e., no decoherence) it must hold:

$$\text{Tr}(\hat{\rho}) = \sum_{i=0}^3 \rho_{ii} = 1 \quad (3.16)$$

Given this information, after some algebraic combination of $\alpha, \beta, \gamma, \delta$ we can therefore extrapolate the original diagonal values of the full state density matrix, which are the ones we are interested in for our protocol:

$$\begin{cases} \rho_{00} = (+\alpha - \beta - \delta + 1) / 2 \\ \rho_{11} = (+\alpha + \beta + \delta - 1) / 2 \\ \rho_{22} = (-\alpha + \beta - \delta + 1) / 2 \\ \rho_{33} = (-\alpha - \beta + \delta + 1) / 2 \end{cases} \quad (3.17)$$

The aforementioned procedure is schematised in figure 3.1. In our simulations we implement it together with the statistics required by the average over random-phase states for $\hat{\rho}_1$.

Once the measurement takes place, we are in possess of the information regarding the populations of the evolved states and we can therefore proceed to compute the fidelity following equation (2.1) to rate the quality of the proposed controls and then proceed with the optimisation.

The full protocol is therefore summarised in algorithm 1.

Algorithm 1 CNOT FoM estimation protocol

Select $3 \cdot N_\theta$ ▷ 3 values of θ required for each $\hat{\rho}_1^\theta$ preparation
 Draw $[\theta_1, \dots, \theta_{3N_\theta}]$, $\theta_i \sim \mathcal{U}(0, 2\pi) \forall i$
 Decide populations of $\hat{\rho}_1(0)$ (and of $\hat{\rho}_1'$ accordingly)
 Define \hat{M} operator (e.g. equation (2.29))
 Collect trial controls Ω and parameters θ

for $k = 0, \dots, 3$ **do** ▷ $\hat{\rho}_1$ contribute to the FoM
 $i = 1$
 while $i < N_\theta$ **do**
 Prepare $\hat{\rho} = |0\rangle\langle 0|$
 Apply $U^{02}(\theta_i)$ for a suitable time^(*)
 Apply $U^{23}(\theta_{i+1})$ for a suitable time^(*)
 Apply $\hat{U}^{13}(\theta_{i+2})$ for a suitable time^(*)
 Apply controlled operator $\hat{U}(\Omega)$
 Measure $\langle k|\hat{\rho}|k\rangle$ and storage
 $i += 3$
 end while
 Average previous measurements
 Multiply for $\langle k|\hat{\rho}_1'|k\rangle$ and storage
end for
 Sum contributes of k , normalise and multiply for w_1

(*) to be selected according to equation (1.44) and $\hat{\rho}_1$ populations

Prepare $\hat{\rho} = |0\rangle\langle 0|$ ▷ $\hat{\rho}_2$ contribute to the FoM
 Apply \hat{M} to move to $\hat{\rho}_2(0)$
 Apply controlled operator $\hat{U}(\Omega)$
 Apply \hat{M}^\dagger to come back to initial basis
 Measure $\langle 0|\hat{\rho}|0\rangle$, normalise and multiply for w_2

return Sum of the two contributions

3.1.4. Error propagation

To build a simulation which represents the experiment with high precision we must design a full, detailed model which includes stochastic errors for each pulse. Assuming to be able to start in the reset state $|0_e 0_n\rangle\langle 0_e 0_n|$ with negligible error, the uncertainty concerning each applied pulse will eventually propagate, leading to a cumulative error on the $\hat{\rho}_i(0)$ states and, ultimately, on their evolution $\hat{\rho}_1^\theta(T)$. After this, measurement errors are to be taken into account. Finally, since mul-

multiple realisations of $\hat{\rho}_i$ are required to gather statistics for the computation of the FoM, the total uncertainty on the fidelity estimation will be given by further error propagation.

For the sake of our implementation, we summarised this behaviour considering only a Gaussian random error on the average fidelity outcome with tunable standard deviation, to verify to which extent our solution can be considered robust. Of course, uncertainties will affect the optimisation algorithm: components and details of the protocol will have to be selected so to have a global error below the working threshold of the protocol.

As for the readout uncertainty, assuming that the estimates of $\kappa = \alpha, \beta, \delta$ come together with experimental errors σ_k , one can retrieve that the error associated to the estimates of $\{\rho_{ii}\}_i$ by simple error propagation. The result is $\sigma_\rho = \sqrt{\sum_k \sigma_k^2}/2$. While readout error is included in the code, for most simulations this has been tuned to a minimum to verify the correct evolution of the protocol.

3.2. Results

Here we present a report of the most interesting numerical tests that have been carried out. For the sake of completeness, we highlight both the working and non-working features of the simulations. As for the optimisation procedure, results in sections 3.2.1 - 3.2.2 have been achieved minimising directly the theoretical infidelity $1 - \mathcal{F}_{sm}$ (equation (1.72)), while in sections 3.2.3 - 3.2.6 we use our protocol for the optimisation (details in algorithm 1) and we report $1 - \mathcal{F}_{sm}$ just for reference.

3.2.1. Definition of controls and Hamiltonian

Before proceeding to the optimisation of the quantum gate it is necessary to identify the pulses we are willing to develop in order to perform quantum control of the system. In principle this breaks down to the definition of the Hamiltonian and the classification of controls between parameters (i.e., constant values) and actual pulses in the form of time series. While our ultimate goal is the definition of a protocol that allows for efficient closed-loop optimisation, at this stage we care about proving that the system is actually controllable and individuate the best set of controls. For this reason, first optimisations have been performed aiming to minimise \mathcal{F}_{sm} (equation (1.72)), the theoretical FoM which is not directly accessible in the experiment but provides mathematical validation of the model.

Of course, having many controls to optimise implies the necessity to investigate a large parameter space: hence, we wish to reduce the complexity of the control mechanism to a minimum without sacrificing the quality of the solution. The first question arises naturally: is there a straightforward way to create a *CNOT* gate by the means of a constant pulse over time? As a matter of fact, looking at figure 1.9 we observe that an intuitive way to flip the nuclear spin according to the electron spin state as in a *CNOT* gate is given by driving resonant transitions between the states $|2\rangle$ and $|3\rangle$ with the r_1 radio-frequency pulse. In fact, since $\omega_{01} \ll \omega_{23} = \omega_{03} - \omega_{02}$, a system prepared in the $|0\rangle$ state will not be affected by EM radiation oscillating at ω_{23} , while a system in $|2\rangle$ will be subject to Rabi oscillations and eventually move to state $|3\rangle$ after a time $T_\pi = \pi |\Omega_{23}|^{-1}$ (with Ω_{23} the complex Rabi frequency of the transition). While this is in practice a selective flip on the nuclear spin, this does not correspond to a precise *CNOT* gate: in fact, the aforementioned transition introduces a local phase difference between the $\{|0\rangle, |1\rangle\}$ and $\{|2\rangle, |3\rangle\}$ subsets of the basis. More in detail a single, constant pulse of intensity $\Omega_{23} = e^{i\varphi}\alpha$ with fixed phase $\varphi = 0$ will lead to a time-independent system Hamiltonian in the interaction picture in the form:

$$\hat{H} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \alpha \\ 0 & 0 & \alpha & 0 \end{pmatrix} \quad (3.18)$$

from which we can retrieve the time evolution operator:

$$\hat{U}_\alpha(t) = \exp(-i\hat{H}t) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos(\alpha t) & -i \sin(\alpha t) \\ 0 & 0 & -i \sin(\alpha t) & \cos(\alpha t) \end{pmatrix} \quad (3.19)$$

Hence, by tuning the intensity and the pulse duration such that $\alpha t = \pi/2$ it is possible to set to zero the matrix elements related to the $|2\rangle$ and $|3\rangle$ states preservation:

$$\langle 2|\hat{U}_\alpha(\frac{\pi}{2\alpha})|2\rangle = \langle 3|\hat{U}_\alpha(\frac{\pi}{2\alpha})|3\rangle = 0 \quad (3.20)$$

but we cannot remove the $\pi/2$ -phase difference introduced between the first and last two states. As a result, we shall conclude that a single constant pulse is not sufficient to obtain a *CNOT* gate which is exact up to global phase differences and therefore we must include higher complexity optimisation.

Of course, while trying to add controls to our model we aim to be minimally intrusive and to increase the complexity gradually: hence, we propose to drive slightly off-resonant RF and MW transitions to introduce detuning parameters and to let the optimiser manipulate the phases of the Rabi frequencies Ω_{ij} . In the original paper which describes the system under investigation (see [40]), the two-qubit effective Hamiltonian is defined with different choice of the a, b, c, d terms while moving to interaction picture (see section 1.3 for reference). This will eventually result in a different representation of the Hamiltonian \hat{H}' which completely decouples the level $|1\rangle$ of the system from the others, leading to an effective 3-state subsystem:

$$\hat{H}' = -\frac{1}{2} \begin{pmatrix} \frac{2}{3}\delta & 0 & \Omega_{02} & 0 \\ 0 & 0 & 0 & 0 \\ \Omega_{02}^* & 0 & -\frac{4}{3}\delta & \Omega_{23} \\ 0 & 0 & \Omega_{23}^* & \frac{2}{3}\delta \end{pmatrix} \quad (3.21)$$

where $\delta = \delta_{r_1} = \delta_m$. The introduction of the detunings could in principle help adjusting the local phase induced by the constant Ω_{23} drive. Nevertheless, equation 3.21 shows how the unperturbed state $|1\rangle$ fixes the global gate phase to 0: as a result, there is still no direct way to fix the phase difference between the state $|1\rangle$ and the states $|2\rangle$ and $|3\rangle$.

To reach this conclusion, multiple variations have been investigated: changing the total time scale of the pulse, reducing the time-step size and ultimately adding new controls. In detail, simulations have been done with constant Ω_{02} and Ω_{23} Rabi frequencies (where in the syntax of QOC constant controls take the name of parameters), with constant Rabi frequencies and detunings, then with fixed detunings and time-dependent intensity of the pulses $\Omega(t)$ and ultimately time-varying control amplitudes and relative phases (for a total of 4 pulses) with detunings as parameters. In all these scenarios repetitive analysis with different dCRAB settings has been carried out to look for the global minimum. Nevertheless, even over

$\mathcal{O}(10^4)$ iterations, 5 super-iterations with strong convergence requirements (average FoM slope smaller than 10^{-4} over 400 iterations) and large amplitude variation parameter to investigate a larger simplex, the minimum infidelity reaches an unsatisfactory bound at $1 - \mathcal{F}_{sm} = 0.1875$, for all simulations. The robustness of this result suggests a fundamental limit which we believe to be due to the local phase difference induced by optimisation in the aforementioned picture. Therefore, we feel the necessity to rethink the problem and to overcome some possible issues in the design of the model, making sure to have the possibility to operate over all the four basis states.

Please note that it is also possible that an optimised gate with constant pulses can be combined with local operations on the states to rectify the phase imbalance and eventually lead to a perfect *CNOT* gate. Nevertheless, this analysis would require higher level considerations regarding the single-state addressability of the controls and would inevitably lead to longer gate preparation due to the additional local operations.

To be able to manipulate the evolution of all the basis states, we concluded to be convenient a change of framework and therefore move to a different interaction picture defined by a new set of a, b, c, d values (namely, the set of parameters reported in equation (1.54)). The result is the one in equation (1.56): the introduction of the γ term allows to both set the picture and fix a reference phase for the optimisation. This action combines with the relaxation of the detuning constraint (so that $\delta_m \neq \delta_{r_1}$) to allow for a multi-parametric optimisation of the gate. In the end, the full control Hamiltonian we use to build the time evolution operator is (reporting from equation (1.55)):

$$\hat{H} = -\frac{1}{2} \begin{pmatrix} \frac{4}{3}\delta_m - \frac{2}{3}\delta_{r_1} & 0 & \Omega_{02} & 0 \\ 0 & \gamma & 0 & \Omega_{13} \\ \Omega_{02}^* & 0 & -\frac{2}{3}\delta_m - \frac{2}{3}\delta_{r_1} & \Omega_{23} \\ 0 & \Omega_{13}^* & \Omega_{23}^* & -\frac{2}{3}\delta_m + \frac{4}{3}\delta_{r_1} \end{pmatrix} \quad (3.22)$$

Unless explicitly said, we perform optimisations assuming the possibility to drive selective NV transitions using only the MW control $\Omega_{02}(t) = \Omega_m(t) = u_m(t) e^{i\theta_m(t)}$, which means that we set $\Omega_{13} = 0$ in equation (3.22). The other control is given by the RF pulse on the nuclear spin: $\Omega_{23}(t) = \Omega_r(t) = u_r(t) e^{i\theta_r(t)}$. For all the quantities subject to optimisation it is necessary to specify boundaries within which the search takes place. Defaults settings have been chosen to be:

- $u_m, u_r \in [-15 \text{ kHz}, 15 \text{ kHz}]$
- $\theta_m, \theta_r \in [0, 2\pi]$
- $\delta_m, \delta_r, \gamma \in [-2 \text{ kHz}, 2 \text{ kHz}]$

Where thresholds have been chosen performing various trials and following the results proposed in [40]. It is worth mentioning that the 15 kHz threshold for the pulse amplitudes seems to be problematic for some particular cases: in these cases, usually a saturation effect becomes visible (see figure 3.5 for an example). When

this happens the pulses can generally be smoothed to make their implementation easier.

3.2.2. Reducing the complexity

Two pulses, five parameters

Studying the system Hamiltonian in the new picture, the optimisation finally works properly and the theoretical fidelity can reach values $\geq 99\%$. For the first trials, we leave complete freedom to the algorithm to optimise all the details of the controls, namely two time-varying pulses (the real amplitudes $u_m(t)$, $u_r(t)$ of the microwave and the radio-frequency pulses between the $|0\rangle$, $|2\rangle$ and the $|2\rangle$, $|3\rangle$ states, respectively) and the five parameters defined by the two independent detunings δ_m and $\delta_r \equiv \delta_{r_1}$, the two constant pulse phases θ_m and θ_r and reference coefficient for the $|1\rangle$ state γ . The results are reported in figure 3.2: over $\mathcal{O}(10^3)$ optimisation iterations the theoretical infidelity for the gate reaches $\mathcal{O}(10^{-6})$.

The best parameters found by the optimiser are:

$$\begin{cases} \theta_m = 0.029735 \\ \theta_r = 0.005620 \\ \delta_m = 0.551751 \text{ kHz} \\ \delta_r = 0.009098 \text{ kHz} \\ \gamma = -0.349105 \text{ kHz} \end{cases} \quad (3.23)$$

while instructions for dCRAB are as follow (section 1.5.3 for reference):

- Max number of SIs: 5
- Target infidelity: 10^{-7}
- Search algorithm: Nelder-Mead
- Number of Fourier basis vectors: 4
- Stopping criteria for SI: linear fit slope below 10^{-4} over 400 iterations
- Initial guess: blind (0 for all values)
- Time bins: 300
- Amplitude variation: $\sim 3\%$ of max pulse intensity

Looking at the results in figure 3.2 we observe that the convergence towards the minimum of the FoM is approximately linear in semi-log scale (top panel). The trend is repeated for all SIs: with this target of FoM which is reached with relative ease, 5 super-iterations are more than what it is necessary. The control amplitudes $u_m(t)$, $u_r(t)$ oscillate in the range $[-3, 5]$ kHz (bottom panel). These corresponds to oscillations of the electromagnetic fields in the direction of the induce dipoles in the spins.




Figure 3.2: Simulation results for a run with two controls ($u_m(t)$, $u_r(t)$) and five parameters ($\theta_m, \theta_r, \delta_m, \delta_r, \gamma$). Top panel: infidelity measure for each function evaluation. Bottom panel: best set of controls returned by dCRAB.

Two pulses, five parameters - fine tuning

By reducing the amplitude variation parameter (i.e., sampling a smaller simplex size with the NM algorithm) for a better focus over this rough pulse estimate we can see that dCRAB reaches very low FoM values within a single super-iteration. Figure 3.3 displays results in these conditions: using the best solution of a previous run (figure 3.2) as initial guess, the figure of merit starts at a lower value and continuous minimisation is persistent through the whole optimisation process. The best parameters found by the optimiser are:

$$\begin{cases} \theta_m = 0.078384 \\ \theta_r = 0.006862 \\ \delta_m = 0.552352 \text{ kHz} \\ \delta_r = 0.009237 \text{ kHz} \\ \gamma = -0.349066 \text{ kHz} \end{cases} \quad (3.24)$$

which are a fine tuning of the ones in equation (3.23), as we would expect. Also, the control shapes reported in figure 3.3 (bottom panel) are similar in shape to their analogue in figure 3.2.

The time evolution operator build with the optimal controls and parameters of this second optimisation is finally:

$$\hat{U} = e^{i\theta} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (3.25)$$

with $\theta = \pi/3$. All the terms in \hat{U} with $|\hat{U}_{ij}| \leq 10^{-6}$ have been set to 0.




Figure 3.3: Simulation results for a run with two controls ($u_m(t)$, $u_r(t)$) and five parameters ($\theta_m, \theta_r, \delta_m, \delta_r, \gamma$) as a fine search over the results of figure 3.2. Top panel: infidelity measure for each function evaluation. Bottom panel: best set of controls returned by dCRAB.

Two pulses, three parameters

The previous analysis confirms that very good *CNOT* approximations can be obtained with just two controls and a set of parameters. The next question to answer is: are there parameters which are superfluous and that can be fixed directly in the implementation? To work in this direction, we first remove the optimisation over the pulse phases θ_m and θ_r , effectively reducing to two pulses and three parameters to find. This corresponds to fix $\theta_m = \theta_r = 0$ in the experimental design. An approach similar to the previous case has been followed, with a fine-tuning run built on the top of the results of figure 3.2 as initial guess. Results are reported in figure 3.4.

The best parameters found by the optimiser are:

$$\begin{cases} \delta_m = 0.547087 \text{ kHz} \\ \delta_r = 0.006443 \text{ kHz} \\ \gamma = -0.349065 \text{ kHz} \end{cases} \quad (3.26)$$

while the initial instruction for dCRAB are as follow:

- Max number of SIs: 5
- Target infidelity: 10^{-10}
- Search algorithm: Nelder-Mead
- Number of Fourier basis vectors: 4
- Stopping criteria for SI: linear fit slope below 10^{-4} over 1000 iterations
- Initial guess: best of pre-optimisation
- Time bins: 300
- Amplitude variation: $\sim 1\%$ of max pulse intensity

The time evolution operator returned by the analysis is indeed very similar to the one in equation (3.25):

$$\hat{U} = e^{i\theta} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (3.27)$$

with $\theta = \pi/3$. All the terms in \hat{U} with $|\hat{U}_{ij}| \leq 10^6$ have been set to 0. This is indeed to be expected since between the best parameters found for the previous case it holds $\theta_m \approx 0 \approx \theta_r$.




Figure 3.4: Simulation results for a run with two controls ($u_m(t)$, $u_r(t)$) and three parameters ($\delta_m, \delta_r, \gamma$) as a fine search over the results of figure 3.2. Top panel: infidelity measure for each function evaluation. Bottom panel: best set of controls returned by dCRAB. Both the control shape and the infidelity trend are very similar to the case of two controls and five parameters (figure 3.3), showing how the complexity reduction does not affect the quality of the optimisation.

Two pulses, two parameters

Ultimately, we investigate the optimisation performances in the case $\delta_m = \delta_r = \delta$ as proposed in the paper introducing the original control Hamiltonian [40]. We modify the setup accordingly and run an optimisation with the same approach of the previous section. In figure 3.5 we report the results of a fine-tuning run which follows a more wide search in the parameter space: as we can notice, very small FoM values can be reached also in this configuration. It is worth mentioning that in this case the threshold on the controls amplitude is reached and in figure 3.5, bottom panel we can observe this behaviour at ~ 1.6 ms for the RF pulse: in this case it would be up to the lab team to verify if this is a hard constraint (and hence an alternative, more regular pulse shape must be found) or if it can be relaxed.

The best parameters found by the optimiser are:

$$\begin{cases} \delta = 1.397678 \text{ kHz} \\ \gamma = -0.349066 \text{ kHz} \end{cases} \quad (3.28)$$

and the optimised gate is:

$$\hat{U} = e^{i\theta} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (3.29)$$

with $\theta = \pi/3$.

Regardless of the fidelity values, these tests allow us to conclude that optimal control can be efficiently applied to our system by considering the optimisation of two pulses and two parameters. In the next stage of the work, we hence implement the figure of merit estimation described in section 3.1 and aim to retrieve similar results for a *CNOT* gate design with closed-loop optimisation.




Figure 3.5: Simulation results for a run with two controls ($u_m(t)$, $u_r(t)$) and two parameters ($\delta_m = \delta_r = \delta$, γ). Top panel: infidelity measure for each function evaluation. Bottom panel: best set of controls returned by dCRAB.

3.2.3. The new Figure of Merit

What we have done so far is a calibration of the optimisation module, which allows us to verify that the system is controllable, that optimal solutions (in principle) exist and that the time evolution operator is well designed. Nevertheless, the ultimate goal of this work is to prove the efficiency of an experiment-friendly FoM, since the quantity \mathcal{F}_{sm} is not directly relatable to measurements. Therefore, at this stage we proceed to implement the protocol described in section 2. It is actually fair to say that the multiple issues faced during the code implementation helped in the development of the procedure we previously described. We report here some of the most fundamental observations taken from the simulations.

General discussion

The initial proposal for the $\hat{\rho}_1$ state comes from the work of C. Koch *et al.* ([56]): $\hat{\rho}_1 = \text{diag}(0.1, 0.2, 0.3, 0.4)$. In this case, while performing optimisation for a *CNOT* gate starting from a blind guess (all controls set to 0), the solution returned by CRAB has consistently been an identity gate $\mathbb{1}$, rather than the target gate. The analysis we conducted to solve this problem is fully summarised in section 2.1.1 and it led to the redefinition of the state as $\hat{\rho}_1 = (0.2, 0.3, 0.4, 0.1)$. After this update in the protocol the optimiser did not mistake the identity as a global minimum anymore.

At this level the optimisation works as predicted: the convergence with the new FoM is reached and the final output is indeed close to a *CNOT* gate. Still, by comparing the results of the new fidelity measure \mathcal{F}_J with the reference one \mathcal{F}_{sm} (computed on the numerical model at each iteration with the new proposal controls) we observe that the two share the same trends but reach convergence at different values. An example is reported in figure 3.6: as we observe, the blue curve (the optimised one) and the red one (the reference) are about one order of magnitude different but their trends are very similar.

Figure 3.6 confirms that proceeding to minimise \mathcal{F}_J we are indeed performing a minimisation of \mathcal{F}_{sm} as well. From the data we can notice how in the final part of each SI the two curves diverge: as the FoM subject to minimisation reaches lower values, the convergence of the reference one slows down. This could be a signal of overfitting over the realisations $\{\hat{\rho}_1^\theta\}_\theta$ (or other specifics of the single run). This view is supported also by the noticeable fluctuations close to the minimum. Practically speaking this is not a problem: first of all, in a real experiment we cannot measure the reference gate fidelity in real-time, but we can perform full process tomography over the optimised pulses to get a single reference measurement \mathcal{F}_{sm} of the final result. This translates into using \mathcal{F}_J for optimisation and \mathcal{F}_{sm} for quality certification. Secondly, we observe that the divergence between the red and blue curves in figure 3.6 begins to be substantial for values of $\mathcal{F}_{sm} \leq 10^{-3}$ which are of the same order of magnitude of state-of-the-art solutions [58]. This means that even if after some time the fine-tuning minimisation of \mathcal{F}_J does not improve the




Figure 3.6: Simulation results for a run with two controls ($u_m(t)$, $u_r(t)$) and two parameters ($\delta_m = \delta_r = \delta$, γ). Top panel: reference (\mathcal{F}_{sm}) and actually minimised (\mathcal{F}_J) infidelity measures for each function evaluation. Bottom panel: best set of controls returned by dCRAB.

reference fidelity this is not a problem: with this approach we can still push \mathcal{F}_{sm} to very high values within few hundreds of optimisation iterations. Of course, the discrepancy between the two curves could be better studied to figure out a scaling parameter which might lead to a longer-term accordance of the trends. In any case, we believe the compatibility of the FoM estimation with the readout process is a strong point this proposal, since it allows to reduce the total number of readout operations required (details in section 2.1.2). As a result, we estimate the total time required by the optimisation of a two-qubit gate to be of the order of hours, given the typical number of dCRAB iterations required for convergence and assuming a single measurement to last for a time of the order of the minute. This means our approach is well suited for applications where multiple NV sites are to be optimised in a reasonable time, such as quantum computers and quantum networks.

A final consideration: it is worth mentioning that the randomised coherences in the $\hat{\rho}_1^\theta$ states have been produced with a constant, random set of phases to be used for each iteration (meaning that the i -th realisation of $\hat{\rho}_1^\theta$ is the same at each optimisation step). As a matter of fact, the initial choice was the generation of completely random phases at each run, to simulate the behaviour of real lasers without additional constraints. Nevertheless, this has turned out to affect the estimate of the FoM in a non negligible way: close to the optimum the fluctuations of the fidelity are extremely small and therefore the effect on the final estimate of changing the phases might be comparable to the one of the controls update proposed by dCRAB, mining the optimisation process. Because of this, we feel it is more wise to sample a set of pulse phases to be used for each SI. This operation does not affect the experimental feasibility since phases can be stored and applied as preferred in modern devices.

Simulation results

Coming to the details of the simulation, figure 3.6 displays the trends of the figures of merit together with the optimiser outputs for the control amplitudes for a 5-SIs optimisation run. The best parameters found by the optimiser are:

$$\begin{cases} \delta = 1.016831 \text{ kHz} \\ \gamma = -0.357037 \text{ kHz} \end{cases} \quad (3.30)$$

while the initial settings are:

- Max number of SIs: 5
- Target infidelity: 10^{-15}
- Search algorithm: Nelder-Mead
- Number of Fourier basis vectors: 5
- Stopping criteria for SI: max 400 iterations per SI
- Initial guess: blind (0 for all values)

- Time bins: 300
- Amplitude variation: $\sim 10\%$ of max pulse intensity
- Number of random samples for $\hat{\rho}_1$: $N_\theta = 100$
- Weights for \mathcal{F}_J : $\vec{w} = [0.75, 0.25, 0.00]$

The best controls allow to build the time evolution operator $\hat{U} =$

$$\begin{pmatrix} 0.537 + 0.843j & 0. + 0.j & 0.011 + 0.032j & -0.012 + 0.020j \\ 0. + 0.j & 0.479 + 0.878j & 0. + 0.j & 0. + 0.j \\ -0.023 + 0.001j & 0. + 0.j & 0.001 - 0.016j & 0.493 + 0.869j \\ -0.023 - 0.025j & 0. + 0.j & 0.468 + 0.883j & 0.014 + 0.007j \end{pmatrix} \quad (3.31)$$

where we observe that the operator in (3.31) has similar terms with non-zero phase in correspondence of the 1 terms in a CNOT while other matrix elements have small absolute value in comparison.

3.2.4. Tuning the algorithm parameters

We already mentioned how several setting must be adjusted for each run to work: the number of iterations and the convergence requirements impact the minimum reachable FoM, the amplitude variation affects the convergence speed and the fluctuation intensities, a reasoned initial guess can save about 100 iterations to break the $\mathcal{F}_{sm} = 1\%$ threshold.

At this point we ask ourselves: is there any other parameter which has a large impact on the performances?

Indeed, the definition of \mathcal{F}_J requires the choice of the two weights to balance the contributes of the gate action on the basis ($\hat{\rho}_1$ optimisation) and the respect of relative phases ($\hat{\rho}_2$ optimisation). As mentioned, the third weight is simply set to $w_3 = 0$ since we focus of the no information leakage case. Investigation of the best ratio among them, together with the considerations about their role held in section 2, answers that an initial choice of $w_1 > w_2$ is preferable. An example of the importance of good weights is reported in figure 3.7, while trying to better the fidelity of the run depicted in figure 3.6 which we use as initial guess. Blue and yellow dots show the output infidelities reached over three SIs for \mathcal{F}_J , \mathcal{F}_{sm} respectively (again, \mathcal{F}_J is the quantity actually subject to minimisation while \mathcal{F}_{sm} is computed for reference). Green dots report the overall best \mathcal{F}_{sm} fidelity reached during the whole optimisation process, which might not be in correspondence of the minimum control objective found (i.e., generally $\text{argmin}(\mathcal{F}_{sm}) \neq \text{argmin}(\mathcal{F}_J)$, where the minimisation is over the controls). It is clear how the best performances on the reference level can be achieved for $w_1 \ll w_2$ (in contrast with the initial run) or with $w_1 \approx w_2$.

By performing a similar study starting from a blind guess, we realise it is instead more convenient to give larger weight to the optimisation of the gate action first, rather than focusing on the cleaning of the local phases that the controls might induce. To summarise, we believe the best approach would need to be run over multiple SIs in an adaptive way such that:

$$\begin{aligned} AV_1 &< AV_2 < \dots < AV_n \\ (w_1)_1 &> (w_1)_2 < \dots > (w_1)_n \end{aligned} \tag{3.32}$$

i.e., successive super-iterations should use reduced amplitude variation for a better investigation of the neighborhood of the previously-found minimum and alternate the major importance of the two components of \mathcal{F}_J , eventually performing the last optimisation with $w_1 = w_2 = 50\%$. This has turned out to enhance the convergence in the first stages while keeping it close to the minimum and improving it in the following steps. A successful example of this approach is shown in figure 3.9.




Figure 3.7: Minimum value reached by the infidelity measures by changing the weights $w_{1,2}$ which rule the contributes of the $\hat{\rho}_1$ and $\hat{\rho}_2$ states evolution. Initial guesses are pre-optimised pulses favouring the $\hat{\rho}_1$ component. Optimisation settings aside the weights are left the same for all simulations. Blue and yellow dots represent the minimum value of \mathcal{F}_J and \mathcal{F}_{sm} returned after for the controls minimising \mathcal{F}_{sm} over the whole optimisation. Green dots represent the overall minimum value of \mathcal{F}_J reached at some point during the search.




Figure 3.8: Analysis of the protocol performance by changing the number of θ samples used for mixed state preparation of $\hat{\rho}_1$ in fixed optimisation conditions.

Another factor that impacts heavily the computation and the results is the number N_θ of random-phase realisation of the state $\hat{\rho}_1$. Both from an experimental and a numerical perspective, it is convenient to find a value of N_θ such that:

$$\frac{1}{N_\theta} \sum_{i=1}^{N_\theta} \hat{\rho}_1^\theta \simeq \hat{\rho}_1 \quad (3.33)$$

with good approximation but limited workload.

Figure 3.8 reports optimisation runs carried out with identical settings and varying N_θ . Details are:

- Max number of SIs: 3
- Target infidelity: 10^{-15}
- Search algorithm: Nelder-Mead
- Number of Fourier basis vectors: 5
- Stopping criteria for SI: linear fit slope below 10^{-5} over 100 iterations
- Initial guess: blind (0 for all values)
- Time bins: 300
- Amplitude variation: $\sim 10\%$ of max pulse intensity
- Weights for \mathcal{F}_J : $\vec{w} = [0.75, 0.25, 0.00]$

Looking at figure 3.8 we can observe how $\mathcal{O}(10)$ random phases are not enough to grant the requirement given by equation (3.33) and as a result both fidelities fail to reach good values. Moreover, it is noticeable how noisy \mathcal{F}_J is, compared with the other plots. On the other hand, exaggerating the number of phases ($\mathcal{O}(10^3)$) brings limited advantages compared to the computational cost. We conclude that an average number of realisations $N_\theta = \mathcal{O}(10^2)$ is the best solution.

Finally, part of the settings for the optimiser is the number of Fourier basis vectors (namely N_c in equation (1.79)). This tuning does not affect heavily the complexity of the algorithm, since a single SI over 4 basis vectors is in practice equivalent to two SIs over 2 basis vectors, as for the sample space investigated. Hence, we must tune N_c together with the number of SIs we plan on performing, eventually taking care of adaptive AV and \vec{w} parameters, as mentioned above. Over a single SI for the same case under study mentioned above we find the following results:

- $N_c = 2$: $\max \mathcal{F}_{sm} < 80\%$
- $N_c = 4$: $\max \mathcal{F}_{sm} = 95\%$
- $N_c = 5$: $\max \mathcal{F}_{sm} = 99.8\%$

which suggest that a single SI with a large enough number of basis vectors is sufficient to reach high quality solutions.

3.2.5. Comparison with the two- $\hat{\rho}_1$ variant

In section 2.1.3 we introduced an alternative idea for the study of the action of the time evolution operator on the basis states. Here, we perform numerical validation of the aforementioned two- $\hat{\rho}_1$ variant designing the code and running simulations similar to the standard $\hat{\rho}_1$ case. The dCRAB settings are:

- Max number of SIs: 5
- Target infidelity: 10^{-15}
- Search algorithm: Nelder-Mead
- Number of Fourier basis vectors: 5
- Stopping criteria for SI: linear fit slope below 10^{-5} over 100 iterations
- Initial guess: adaptive (SI₁: blind, SI₂: best of SI₁, SI₃₋₅: best of SI₁)
- Time bins: 300
- Amplitude variation: adaptive (SI₁: 10%, SI₂: 0.5%, SI₃₋₅: 0.05%)
- Weights for \mathcal{F}_J : adaptive (SI₁: $w_1 = 75\%$, SI₂: $w_1 = 10\%$, SI₃₋₅: $w_1 = 50\%$)

Figure 3.9 reports results for the two- $\hat{\rho}_1$ variant build as follows: first, a single SI with large amplitude variation (to favour the $\hat{\rho}_1$ contribute), then a single SI with medium amplitude variation favouring the $\hat{\rho}_2$ contribute and finally three SIs with balanced weights and small amplitude variation. The number of random-phase realisation is reduced to $N_\theta = 10$. It is noticeable how at the beginning of each SI the trends of the two figures of merit are quite similar. With respect to the standard case, the number of FoM evaluation increases largely, but the number of realisations and measurements to be performed for each of them decreases of orders of magnitude. Overall, the convergence of the infidelity is reached at $\mathcal{F}_{sm} = 99.95\%$, best overall for the simulations so far. The best parameters are found to be:

$$\begin{cases} \delta = 1.041042 \text{ kHz} \\ \gamma = -0.351753 \text{ kHz} \end{cases} \quad (3.34)$$

while the time evolution operator is $\hat{U} =$

$$\begin{pmatrix} 0.508 + 0.861j & 0. + 0.j & -0.014 + 0.014j & 0.006 + 0.009j \\ 0. + 0.j & 0.493 + 0.87j & 0. + 0.j & 0. + 0.j \\ -0.005 - 0.01j & 0. + 0.j & 0.004 + 0.011j & 0.500 + 0.866j \\ -0.019 + 0.006j & 0. + 0.j & 0.491 + 0.871j & -0.007 - 0.009j \end{pmatrix} \quad (3.35)$$

We can hence conclude that the two- $\hat{\rho}_1$ approach is indeed convenient since it works very well with an order of magnitude less random realisations and it naturally displays less noisy curves. Nevertheless, the total number of function evaluations to reach the optimum is relatively large: in a practical implementation of




Figure 3.9: Simulation results for a run with two controls ($u_m(t)$, $u_r(t)$) and two parameters ($\delta_m = \delta_r = \delta$, γ) in the two- $\hat{\rho}_1$ variant. Top panel: reference (\mathcal{F}_{sm}) and actually minimised \mathcal{F}_J infidelity measures for each function evaluation. Bottom panel: best set of controls returned by dCRAB.

the protocol this can potentially lead to some issues. Looking at the differences between the standard protocol and this variant, it is possible that testing the gate action on the basis state-per-state (i.e., tomography by basis states) will lead to even more accurate results, since the ones displayed are naturally affected by the random realisations. Nevertheless, we believe that this approach can be most convenient because of the better suitability with the natural measurement process of the spins.

3.2.6. Comparison with non-selective e^- spin transitions

In all the analysis carried out in the previous pages a fundamental detail has been assumed, i.e. the possibility to selectively drive the $|0\rangle - |2\rangle$ or $|1\rangle - |3\rangle$ transitions. In some situations it is either convenient or necessary to have non-selective electron excitation, either via bi-chromatic pulses or strong illumination with large bandwidth. Hence, we verify the possibility of building the gate with the non-selective excitation $\Omega_{nv} = \Omega_{02} = \Omega_{13}$. The approach is the same of section 3.2.5: multiple SIs with adaptive amplitude variation and different balance for the two contributes to \mathcal{F}_J . Best results are reported in figure 3.10.

The best parameters found by the optimisation are

$$\begin{cases} \delta = 1.040245 \text{ kHz} \\ \gamma = -0.372271 \text{ kHz} \end{cases} \quad (3.36)$$

while the time evolution operator is: $\hat{U} =$

$$\begin{pmatrix} 0.590 + 0.796j & -0.094 + 0.068j & -0.011 - 0.021j & -0.009 + 0.059j \\ -0.100 + 0.057j & 0.442 + 0.885j & 0.045 + 0.007j & 0.025 - 0.068j \\ -0.059 - 0.009j & 0.072 + 0.014j & -0.015 - 0.011j & 0.463 + 0.881j \\ 0.017 + 0.021j & 0.021 - 0.037j & 0.425 + 0.903j & -0.004 + 0.021j \end{pmatrix} \quad (3.37)$$

While the final fidelity is slightly worse for the non-selective case with respect to the selective one (studied in the previous sections), we believe that a deeper optimisation procedure might enforce the results also in this situation. In any case, we must remember that the model of the system is not meant to be used for actual optimisation but simply to prove the validity of the FoM introduced in equation (2.1). Once again, we see in the evolution of the control objective (figure 3.10 above) the similar trends of the two infidelity measures. This makes us confident that application of this technique to a real system could help building high-fidelity gates regardless of the selectivity of the NV transitions.




Figure 3.10: Simulation results for a run with two controls ($u_{nv}(t)$, $u_r(t)$) and two parameters (δ , γ) with phase-random $\hat{\rho}_1$ approach and non-selective electronic excitation. Top panel: reference (\mathcal{F}_{sm}) and actually minimised \mathcal{F}_J infidelity measures for each function evaluation. Bottom panel: best set of controls returned by dCRAB.

Chapter 4

Summary and outlook

In this work we addressed the problem of designing an efficient protocol to perform closed-loop two-qubit gate optimization on NV centers in diamonds. Our approach includes a review of modern quantum optimal control techniques, a deep understanding of the quantum hardware specifics, the layout of the optimisation procedure and its numerical validation by development of a digital twin of the system.

On the basis of the simulations carried out, we can conclude that our optimisation routine reaches the goal of shaping electromagnetic control pulses that steer the system time evolution as a CNOT gate application with a reference fidelity $\mathcal{F}_{sm} > 99\%$, comparable with the state of the art. Most importantly, being the process fully designed on the basis of experimental details, it enables a straightforward implementation on real devices. This will allow to perform the control search directly based on the system specifics, looking at the actual response and saving the effort of developing a model (which is usually either too complex or too approximate to be a perfect representation).

We managed to propose the protocol with some degrees of freedom that can be tuned according to the user requirements. If a fast convergence of the figure of merit is needed, the recommended approach would be the investigation of the gate action on all the basis states at once via the preparation and average of multiple pure states with random-coherence and non-degenerate populations. Instead, if a slow yet precise infidelity minimisation is preferred, then it is convenient to monitor the evolution of the whole set of basis vectors or to exploit an hybrid approach. In any case, the random-phase methods appear to be more experiment-friendly, since they are well suited for the multiple realisations required to gather statistics in the readout process and therefore require no extra experimental effort.

4.1. Towards the experiment

The next interesting development of this work will be the experimental validation of the protocol. To move in this direction it is first of all necessary to find an ad-

equate system: while ^{13}C , ^{14}N and ^{15}N atoms are all suitable candidates, they present different challenges together with different academic interest. For example, to initialise the nuclear spin to $|0\rangle_n$ it is necessary to prepare the electronic spin in $|0\rangle_e$ and then transfer this state to the nuclear one: this procedure is technically faster on Nitrogen atoms [59] and it is a partial motivation to the fact that experimental efforts towards two-qubit gates on a single NV center have been especially focused on this platform so far [60]. Moreover, in section 3.1.3 we highlighted how the ability to independently readout the populations of the four basis states is strongly conditioned to the possibility to drive selective transitions between the $m_s = 0$ and $m_s = -1$ states of the 3A_2 and 3E levels. Of course, this is only feasible if the respective transition frequencies are sufficiently far apart in the spectrum. Since the energy scheme of the NV center depends on its internal structure together with the configuration of the neighbouring atoms, it is important to individuate a suitable candidate for the application of the protocol.

Another huge factor is given by the decoherence time of the two qubits. We discussed the specifics of a suitable experiment for our proposal: typical decoherence scales are given by $t_2^* = \mathcal{O}(20 \mu s)$ for the NV (with possible extension of one order of magnitude if spin echo is performed) and $t_2^* = \mathcal{O}(1 ms)$ for a ^{13}C nuclear spin [61, 62]. Hence, the total gate time must be much shorter than the one fixed in the previous simulations. We can aim for optimisation in two possible regimes: the first one requires $\Omega_m \approx \Omega_r$ with a total gate time T of order $\mathcal{O}(1 \mu s)$, while the second scenario employs $\Omega_m = \mathcal{O}(50 MHz)$, $\Omega_r = \mathcal{O}(40 kHz)$ and $T = \mathcal{O}(100 \mu s)$.

While performing open-loop analysis to certify the quality of the protocol before moving to the experimental setup, some issues arise. The first one refers to the simulation details illustrated in section 1.3, where we stated that the system Hamiltonian we control is defined in the interaction picture within the rotating wave approximation. The latter is valid in the case $\Omega \ll \omega$, i.e. Rabi frequency much smaller than the transition frequency, for both the MW and RF pulses. Nevertheless, since $\omega_{13} = \mathcal{O}(100 MHz)$ as well as $\omega_{23} = \mathcal{O}(100 MHz)$ (where $\omega_{ij} = |E_i - E_j|/\hbar$), for the RWA to be valid the control amplitudes should be at most of the order of $1 MHz$, which is incompatible with the requirement of a gate time $T \sim 1 \mu s$. Hence, to test the protocol with the first set of conditions mentioned above it will be necessary to prepare a more complex model which does not exploit approximations for fast oscillating terms.

This is not the case for the second scenario, since over a longer gate time we can reduce the control intensity and remain in a (borderline) valid approximation regime. With this new settings, we observe that simulations fail to find a reasonable solution. Following the approach indicated in [63], we believe that the use of two controls such that $\Omega_m \gg \Omega_r$ introduces in the system a Zeno subspace in which the survival probability for any initial state within the subspace is given by:

$$P_s(t) = \left(1 - \frac{2\Omega_r^2}{\Omega_r^2 + \Omega_m^2} \sin^2 \left(\frac{\sqrt{\Omega_r^2 + \Omega_m^2} t}{2} \right) \right)^2 \xrightarrow{\Omega_m \gg \Omega_r} 1 \quad (4.1)$$

which means that the action of the RF pulse becomes useless.

To summarise, numerical validation of this protocol for setup parameters comparable with the ones of the real lab we are starting a collaboration with is still to be verified and will require a more detailed characterisation of the system. Another solution would be the design of controls with a pulse duration of $T \geq \mathcal{O}(1 \text{ ms})$, taking into account that for this pulse length the system is subject to decoherence and therefore additional, correction terms in the figure of merit should be considered. In any case, these results simply indicate that the modeling we used for the numerical validation is not suitable to represent the real-world setup: the power of closed-loop control lays exactly in the possibility of directly accessing the system without the need of a model. In this sense, we are confident that a direct implementation of the protocol will lead to some quality results, making another step forward towards the realisation of two-qubit gates to enhance the future of quantum computers.

Acknowledgements

Many years ago, when I first got in touch with the realm of quantum technologies, I found myself excited and enthusiast about the idea of bringing a concrete contribute to the field. After all this time I am proud of what happened in my quantum career, with this project representing just the last yet probably most meaningful achievement of mine. In this regard, I want to thank all the people who I crossed paths with and that made my passion for this world grow day by day.

First and foremost, I feel it is rightful to thank all my family for the unconditional support and trust they gave me, not only during the period I dedicated to this work but through my whole university experience: immensely thanks to **Deme, Paola, Federico** and **Giacomo** for being here for me and pushing me to reach my best. A special thanks to **Angela** for being such an amazing person, for caring the way she does, bringing out the best in me by celebrating when I succeed and helping me in my dark moments. Thank you very much for believing in me, Fufi.

The work I summarised in these pages is the fruit of a six-months effort I carried out at the PGI-8 Institute of the Forschungszentrum Jülich, in Germany in collaboration with the University of Padova. In this regard, I want to thank Prof. **Tommaso Calarco** for welcoming me to his group, as well as all the people who helped me going through my German days. I want to say thank you to **Monika** for her assistance through the bureaucracy of FZJ. I thank the other seniors of the Institute (**Matteo, Felix, Robert, Zsolt, Alessandro**) for stimulating continuously the ideas of us yourger researchers. I am grateful to all my colleagues, PhD and MSc students, who shared the office with me everyday: thanks to **Thomas** for lending me his desk and sharing his knowledge about NVs to clear all my doubts, to Juhi and Cristina for their positive energy, to **Michael, Francesco, Akshay** and **Martino** for the insightful thoughts over coffee and lunch breaks, to **Florian** for the beers discussing research on the Rhine and to **Bora** for hacking in London together, designing new applications for quantum computing. A huge thank to my friend **Sandeep** for all his help, his insights, and his continuous motivation. Thanks also for the chats about sports and the football tournament: I hope your eyeglasses are now ok.

I want then to thank all people that shared with me the everyday life, making my stay abroad pleasant and engaging. To my roommates: thanks to **Govind** for sharing with me the passion for football and Italian cousine, to **Maidier** and **Atuka** for the daily meetings playing basketball and discussing life, to **Josh, Fabrizio**

and **Maria** for the trips on weekend, discovering the beauty of the surrounding countries. Thanks for my Spanish friends **Luis** and **Araceli** for bringing so much energy and fun into the Gästehaus, with your passion for parties and gatherings. Thanks to **Nawar**, to **Victor**, **Julio** and **Richie** for being always there with their positive attitude and their laughs. A special thanks to **Henrique** for caring about me like a little brother, with his wise suggestions and his responsibility, and to **Norbert** for being a genuine friend, always helpful and mindful of others. You all made me feel at home in my second home.

A thought goes to all the people of **Collegio Mazza - Job Campus**, who I had the pleasure to live with during my stay in Padova: thanks for making me grow as a person in these years, for accepting me and my ideas while sharing yours. From a coffee break to a late night chat, every morning I would never knew what awaited me, and I am most grateful for this.

A final thanks to Prof. **Simone Montangero**, my supervisor from UniPd, for his teachings and his promotion of extra-academic activities. Then, I am incredibly grateful to **Matthias** (which you might call Dr. Matthias Müller) for being an amazing person, besides an amazing researcher and supervisor. I enjoyed working on his side and I really want to say thanks to him for believing in me, supporting my proposals, promoting activities to enhance my learning and making me grow in the academic field.

Ultimately, I want to thank my friends from University: thanks to **Marco**, **Michele**, **Antonio**, **Luca**, **Lorenzo**, **Giulia**, **Cristina**, **Filippo**, and all the ones I shared my best and worst moments with. A special mention to **Tommaso** for the experiences we shared, the enthusiasm and the drive to bring new ideas to life, and to **Samuele** for having been my companion in this journey in Germany and for the mutual help to face the most troublesome days. The last of my thanks goes to the younger **Alessandro**, for choosing to pursue his dreams and believing in himself enough to let me be, today, the person that I am.

To you all, to whom I forgot and to whoever shared a moment and a smile with me, I say the most profound and sincere of my thanks.

*Ad maiora,
Alessandro Marcomini*

List of Figures

1.1	Example states on the Bloch sphere, corresponding to the X , Y and Z basis states, respectively. As a convention the $ 0\rangle$ and $ 1\rangle$ states are chosen as the basis states on the Z (vertical) axis. Figure from [17].	6
1.2	Quantum trajectory of evolution from the $ 0\rangle$ to the $ 1\rangle$ state on the surface of the Bloch sphere (π - pulse).	7
1.3	Quantum gates for circuitual representation. Each line represents a qubit in the register. Input on the left hand-side. In multi-qubit gates control qubits are marked with a full dot. (a) Hadamard gate (b) X gate (c) Z gate (d) S (<i>Phase</i>) gate (e) inverse S gate (f) T gate (g) inverse T gate (h) $CNOT$ gate (i) CZ gate (j) $CPhase$ gate (k) $SWAP$ gate (l) Toffoli gate. Figure from [19]	10
1.4	Structural picture of NV^- centers in diamonds. (a) The NV center in the crystal lattice of the diamond. The NV axis, i.e. the direction defined by the link between the Nitrogen atom and the vacancy, displays C_{3v} symmetry oriented along the 111 axis of the crystal. (b) Energetic scheme of the effective electronic spin $S = 1$. Transitions can be driven among the $m_s = 0$ and $m_s = \pm 1$ levels via microwave (MW) radiation. Green light ($\lambda \leq 637 \text{ nm}$) can be used to perform off-resonant excitation from the 3A_2 to the 3E level. Spontaneous decay from the latter level happens in the form of photoluminescence (ZPL at $\lambda = 637 \text{ nm}$). Higher energy levels can also display preferential decay via the metastable 1A_1 , 1E singlet levels. (c) Focus on the MW transitions of the lowest energy triplet states: upon Zeeman splitting of the $m_s = \pm 1$ levels, single transitions $\omega_{0\pm}$ can be independently addressed to trigger oscillations among the ground and excited states, with detunings δ_{\pm} . Picture from [28].	14

- 1.5 Fluorescence spectrum of NV centers in diamonds. (a) Schematic optical transitions among the 3A_2 and 3E levels. Excitation is provided by spin-preserving off-resonant radiation, which spontaneously leads to population of the triplet states. Electrons can decay back to the ground state together with fluorescent emission. (b) Spectrum of fluorescent emission for the NV^0 and NV^- species. The zero phonon line (ZPL) is marked at 575 nm and 637 nm , respectively. Figure from [33]. 16
- 1.6 ODMR spectrum of NV center in absence (a) and presence (b) of external aligned magnetic field ($\mathcal{O}(0.1\text{ mT})$). When no Zeeman splitting occurs, the dip in the spectrum is located at the characteristic frequency of 2.87 GHz . When the magnetic field is turned on the dip splits in two, highlighting the symmetrically distributed resonance peaks. Figure from [36]. 17
- 1.7 Arbitrary rotation of a single-qubit state by means of a driving field \vec{E} with $|\vec{E}| = E \cos(\omega t)$ and $\vec{E}/|\vec{E}| = (\cos(\theta), \sin(\theta), 0)$. The basis is defined by the spin states in the \hat{z} direction. 19
- 1.8 Experimental data of NV center characterisation. (a) Rabi oscillations of the $m_s = 0$ ($|0\rangle$), $m_s = -1$ ($|1\rangle$) subsystem of the 3A_2 energy level of the NV. After the initialisation the system is in the $|0\rangle$ state. A microwave square pulse of resonant frequency is then applied for different pulse lengths t_p . The data show the perfect population inversion for the first Rabi cycle and subsequent amplitude decay due to decoherence while keeping the oscillatory behaviour. (b) Study of the T_2 coherence time of the system via spin echo measurements (see [39]). Figure from [39]. 19
- 1.9 Electron spin - nuclear spin system in NV centre. The transition frequencies ω_{ij} are taken from [40] and reported in the text. Controls and detunings ω_k, δ_k ($k = m, r_1$) are subject to optimisation. The control ω_{r_2} is originally introduced for state initialisation only. 21
- 1.10 Scheme of optimal control inference process. The system is treated as a black box which takes in input the trial control signal and evolves the system via and EoM ruled by $f(x, u, t) \equiv a(x, t) + b(x, u, t)$. The output signal is the measurement result. If the input control is optimal ($\mathbf{u}(t) = \mathbf{u}^*(t)$), then the output signal is the target one. In closed-loop optimisation output signals are used as a feedback to adjust the input one. 26

2.1	(a) State representation on the Bloch sphere. (b) Visualisation of the random realisation of single qubit states sharing the same diagonal on the Bloch sphere. Initial state is $ 0\rangle$ (green). Upon evolution with a pulse of tailored intensity and duration, the state is rotated to reach a specific value of z (here $z = -0.5$, orange state). The random phase of the pulse results in a random state position on the horizontal circumference at z (blue dots, here $N = 100$ samples). The average of many uniformly distributed realisations results in a mixed state with Bloch vector $\vec{r} = (0, 0, z)^T$ (blue). . .	40
3.1	Scheme of the routine to access the population of the four basis states via Optical Readout (OR) on the NV spin. First, OR is performed to measure the populations of the NV states and simultaneously reset it to the $ 0\rangle_n$ state (green pulse on the left-hand side). Then, a selective π -pulse on the NV conditioned to state of the nuclear spin (effective C_nNOT_e gate) transfers information regarding the ^{13}C spin state to the NV (orange pulse). Finally, another OR reads again the populations of the NV (second green pulse). Combining the four measurement outcomes allows to reconstruct the density matrix diagonal.	52
3.2	Simulation results for a run with two controls ($u_m(t), u_r(t)$) and five parameters ($\theta_m, \theta_r, \delta_m, \delta_r, \gamma$). Top panel: infidelity measure for each function evaluation. Bottom panel: best set of controls returned by dCRAB.	59
3.3	Simulation results for a run with two controls ($u_m(t), u_r(t)$) and five parameters ($\theta_m, \theta_r, \delta_m, \delta_r, \gamma$) as a fine search over the results of figure 3.2. Top panel: infidelity measure for each function evaluation. Bottom panel: best set of controls returned by dCRAB. . .	61
3.4	Simulation results for a run with two controls ($u_m(t), u_r(t)$) and three parameters ($\delta_m, \delta_r, \gamma$) as a fine search over the results of figure 3.2. Top panel: infidelity measure for each function evaluation. Bottom panel: best set of controls returned by dCRAB. Both the control shape and the infidelity trend are very similar to the case of two controls and five parameters (figure 3.3), showing how the complexity reduction does not affect the quality of the optimisation.	63
3.5	Simulation results for a run with two controls ($u_m(t), u_r(t)$) and two parameters ($\delta_m = \delta_r = \delta, \gamma$). Top panel: infidelity measure for each function evaluation. Bottom panel: best set of controls returned by dCRAB.	65
3.6	Simulation results for a run with two controls ($u_m(t), u_r(t)$) and two parameters ($\delta_m = \delta_r = \delta, \gamma$). Top panel: reference (\mathcal{F}_{sm}) and actually minimised \mathcal{F}_J infidelity measures for each function evaluation. Bottom panel: best set of controls returned by dCRAB.	67

3.7	Minimum value reached by the infidelity measures by changing the weights $w_{1,2}$ which rule the contributes of the $\hat{\rho}_1$ and $\hat{\rho}_2$ states evolution. Initial guesses are pre-optimised pulses favouring the $\hat{\rho}_1$ component. Optimisation settings aside the weights are left the same for all simulations. Blue and yellow dots represent the minimum value of \mathcal{F}_J and \mathcal{F}_{sm} returned after for the controls minimising \mathcal{F}_{sm} over the whole optimisation. Green dots represent the overall minimum value of \mathcal{F}_J reached at some point during the search.	71
3.8	Analysis of the protocol performance by changing the number of θ samples used for mixed state preparation of $\hat{\rho}_1$ in fixed optimisation conditions.	72
3.9	Simulation results for a run with two controls ($u_m(t)$, $u_r(t)$) and two parameters ($\delta_m = \delta_r = \delta$, γ) in the two- $\hat{\rho}_1$ variant. Top panel: reference (\mathcal{F}_{sm}) and actually minimised \mathcal{F}_J infidelity measures for each function evaluation. Bottom panel: best set of controls returned by dCRAB.	75
3.10	Simulation results for a run with two controls ($u_{nv}(t)$, $u_r(t)$) and two parameters (δ , γ) with phase-random $\hat{\rho}_1$ approach and non-selective electronic excitation. Top panel: reference (\mathcal{F}_{sm}) and actually minimised \mathcal{F}_J infidelity measures for each function evaluation. Bottom panel: best set of controls returned by dCRAB. . .	77

Bibliography

- [1] David P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000.
- [2] Morten Kjaergaard, Mollie E. Schwartz, Jochen Braumüller, Philip Krantz, Joel I.-J. Wang, Simon Gustavsson, and William D. Oliver. Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics*, 11(1):369–395, 2020.
- [3] D. Schrader, I. Dotsenko, M. Khudaverdyan, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede. Neutral atom quantum register. *Phys. Rev. Lett.*, 93:150501, Oct 2004.
- [4] M. Johanning, Andrés Varón, and Christof Wunderlich. Review: Quantum simulations with cold trapped ions. *Journal of Physics B-atomic Molecular and Optical Physics - J PHYS-B-AT MOL OPT PHYS*, 42, 08 2009.
- [5] Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian Qin, Dian Wu, Xing Ding, Yi Hu, Peng Hu, Xiao-Yan Yang, Wei-Jun Zhang, Hao Li, Yuxuan Li, Xiao Jiang, Lin Gan, Guangwen Yang, Lixing You, Zhen Wang, Li Li, Nai-Le Liu, Chao-Yang Lu, and Jian-Wei Pan. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020.
- [6] Romana Schirhagl, Kevin Chang, Michael Loretz, and Christian L. Degen. Nitrogen-vacancy centers in diamond: Nanoscale sensors for physics and biology. *Annual Review of Physical Chemistry*, 65(1):83–105, 2014. PMID: 24274702.
- [7] Victor Acosta and Philip Hemmer. Nitrogen-vacancy centers: Physics and applications. *MRS Bulletin*, 38, 02 2013.
- [8] Marcus W. Doherty, Neil B. Manson, Paul Delaney, Fedor Jelezko, Jörg Wrachtrup, and Lloyd C.L. Hollenberg. The nitrogen-vacancy colour centre in diamond. *Physics Reports*, 528(1):1–45, jul 2013.
- [9] William Koehl, Hosung Seo, Giulia Galli, and David Awschalom. Designing defect spins for wafer-scale quantum technologies. *MRS Bulletin*, 40:1146–1153, 12 2015.
- [10] Matteo Pompili, Sophie Hermans, Simon Baier, Hans Beukers, Peter Humphreys, Raymond Schouten, Raymond Vermeulen, Marijn Tiggelman, Laura Martins, Bas Dirkse, Stephanie Wehner, and Ronald Hanson. Quantum networks based on diamond nv centers. page 9, 03 2021.
- [11] Phila Rembold, Nimba Oshnik, Matthias M. Müller, Simone Montangero, Tommaso Calarco, and Elke Neu. Introduction to quantum optimal control for quantum sensing with nitrogen-vacancy centers in diamond. *AVS Quantum Science*, 2(2):024701, 2020.

- [12] Richard P Feynman. Simulating physics with computers. *International journal of theoretical physics*, 21(6/7):467–488, 1982.
- [13] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *STOC '96*, 1996.
- [14] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26:1484–1509, 1999.
- [15] U. Fano. Description of states in quantum mechanics by density matrix and operator techniques. *Rev. Mod. Phys.*, 29:74–93, Jan 1957.
- [16] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.
- [17] Anton Frisk Kockum and Franco Nori. *Quantum Bits with Josephson Junctions*, pages 703–741. 09 2019.
- [18] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a fundamental quantum logic gate. *Phys. Rev. Lett.*, 75:4714–4717, Dec 1995.
- [19] Chia-Hung Chien, Rodney Van Meter, and Sy-Yen Kuo. Fault-tolerant operations for universal blind quantum computation. *ACM Journal on Emerging Technologies in Computing Systems*, 12, 06 2013.
- [20] Sergey Bravyi and Dmitri Maslov. Hadamard-free circuits expose the structure of the clifford group. *IEEE Transactions on Information Theory*, PP:1–1, 05 2021.
- [21] *Introduction to Quantum Mechanics*, page 1340. Springer US, Boston, MA, 2006.
- [22] F. Jelezko and J. Wrachtrup. Single defect centres in diamond: A review. *physica status solidi (a)*, 203(13):3207–3225, 2006.
- [23] I Aharonovich, S Castelletto, D A Simpson, C-H Su, A D Greentree, and S Praver. Diamond-based single-photon emitters. *Reports on Progress in Physics*, 74(7):076501, jun 2011.
- [24] C. D. Clark, H. Kanda, I. Kiflawi, and G. Sittas. Silicon defects in diamond. *Phys. Rev. B*, 51:16681–16688, Jun 1995.
- [25] Takayuki Iwasaki, Fumitaka Ishibashi, Yoshiyuki Miyamoto, Yuki Doi, Satoshi Kobayashi, Takehide Miyazaki, Kosuke Tahara, Kay Jahnke, Lachlan Rogers, Boris Naydenov, Fedor Jelezko, Satoshi Yamasaki, Shinji Nagamachi, Toshiro Inubushi, Norikazu Mizuochi, and Mutsuko Hatano. Germanium-vacancy single color centers in diamond. *Scientific reports*, 5:12882, 08 2015.
- [26] Takayuki Iwasaki, Yoshiyuki Miyamoto, Takashi Taniguchi, Petr Siyushev, Mathias H. Metsch, Fedor Jelezko, and Mutsuko Hatano. Tin-vacancy quantum emitters in diamond. *Phys. Rev. Lett.*, 119:253601, Dec 2017.
- [27] B. Hensen, N. Kalb, M. S. Blok, A. E. Dréau, A. Reiserer, R. F. L. Vermeulen, R. N. Schouten, M. Markham, D. J. Twitchen, K. Goodenough, D. Elkouss, S. Wehner, T. H. Taminau, and R. Hanson. Loophole-free bell test using electron spins in diamond: second experiment and additional analysis. *Scientific Reports*, 6(1), aug 2016.

- [28] Silvia Arroyo Camejo, Andrii Lazariiev, Stefan Hell, and Gopalakrishnan Balasubramanian. Room temperature high-fidelity holonomic single-qubit gate on a solid-state spin. *Nature Communications*, 5:4870, 09 2014.
- [29] Vogel T. et al. Meijer J., Pezzagna S. Towards the implanting of ions and positioning of nanoparticles with nm spatial resolution. *Appl. Phys. A*, pages 491, 567–571, 2008.
- [30] A. Tallaire, M. Lesik, V. Jacques, S. Pezzagna, V. Mille, O. Brinza, J. Meijer, B. Abel, J. F. Roch, A. Gicquel, and J. Achard. Temperature dependent creation of nitrogen-vacancy centers in cvd diamond layers. 2014.
- [31] Minoru Akaishi, Hisao Kanda, and Shinobu Yamaoka. Synthesis of diamond from graphite-carbonate system under very high temperature and pressure. *Journal of Crystal Growth*, 104(2):578–581, 1990.
- [32] M. W. Doherty, F. Dolde, H. Fedder, F. Jelezko, J. Wrachtrup, N. B. Manson, and L. C. L. Hollenberg. Theory of the ground-state spin of the nv^- center in diamond. *Phys. Rev. B*, 85:205203, May 2012.
- [33] Ariful Haque and Sharaf Sumaiya. An overview on the formation and processing of nitrogen-vacancy photonic centers in diamond by ion implantation. *Journal of Manufacturing and Materials Processing*, 1:6, 08 2017.
- [34] J. Harrison, M.J. Sellars, and Neil Manson. Measurement of the optically induced spin polarisation of nv centres in diamond. *Diamond and Related Materials - DIAM RELAT MATER*, 15:586–588, 04 2006.
- [35] A. Dréau, M. Lesik, L. Rondin, P. Spinicelli, O. Arcizet, J.-F. Roch, and V. Jacques. Avoiding power broadening in optically detected magnetic resonance of single nv defects for enhanced dc magnetic field sensitivity. *Phys. Rev. B*, 84:195204, Nov 2011.
- [36] S. Becker, Nicole Raatz, St Jankuhn, R. John, and J. Meijer. Nitrogen implantation with a scanning electron microscope. *Scientific Reports*, 8, 01 2018.
- [37] Jeronimo Maze, Paul Stanwix, Jonathan Hodges, Sangjun Hong, Jacob Taylor, Paola Cappellaro, Liang Jiang, Gurudev Dutt, E Togan, A. Zibrov, Amir Yacoby, R. Walsworth, and M Lukin. Nanoscale magnetic sensing with an individual electronic spin in diamond. *Nature*, 455:644–7, 11 2008.
- [38] Gilbert Grynberg, Alain Aspect, Claude Fabre, and Claude Cohen-Tannoudji. *Introduction to Quantum Optics: From the Semi-classical Approach to Quantized Light*. Cambridge University Press, 2010.
- [39] Chathuranga Abeywardana, Viktor Stepanov, Franklin Cho, and Susumu Takahashi. Magnetic resonance spectroscopy using a single nitrogen-vacancy center in diamond. *Proc SPIE*, 9269:92690K, 11 2014.
- [40] R. S. Said and J. Twamley. Robust control of entanglement in a nitrogen-vacancy center coupled to a ^{13}C nuclear spin in diamond. *Phys. Rev. A*, 80:032303, Sep 2009.
- [41] Ying Wu and Xiaoxue Yang. Strong-coupling theory of periodically driven two-level systems. *Phys. Rev. Lett.*, 98:013601, Jan 2007.
- [42] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. *J. Artif. Int. Res.*, 4(1):237–285, may 1996.

- [43] JM Geremia. An introduction to control theory from classical to quantum applications (lecture notes), 2003.
- [44] JunBum Kwon, Xiongfei Wang, Frede Blaabjerg, and Claus Leth Bak. Comparison of lti and ltp models for stability analysis of grid converters. In *2016 IEEE 17th Workshop on Control and Modeling for Power Electronics (COMPEL)*, pages 1–8, 2016.
- [45] Matthias M. Müller, Ressa S. Said, Fedor Jelezko, Tommaso Calarco, and Simone Montangero. One decade of quantum optimal control in the chopped random basis, 2021.
- [46] H. J. Carmichael. *An open systems approach to quantum optics*. Springer-Verlag, Berlin Heidelberg, 1993.
- [47] Christiane P. Koch, Ugo Boscain, Tommaso Calarco, Gunther Dirr, Stefan Filipp, Steffen J. Glaser, Ronnie Kosloff, Simone Montangero, Thomas Schulte-Herbrüggen, Dominique Sugny, and Frank K. Wilhelm. Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe. *arXiv e-prints*, page arXiv:2205.12110, May 2022.
- [48] A. I. Konnov and V. F. Krotov. On global methods for the successive improvement of control processes. *Automation and Remote Control*, 60(10):1427–1436, 1999.
- [49] Morzhin Oleg and Alexander Pechen. Krotov method in optimal quantum control. i. closed systems, 09 2018.
- [50] S G Schirmer and Pierre de Fouquieres. Efficient algorithms for optimal control of quantum dynamics: the krotov method unencumbered. *New Journal of Physics*, 13(7):073029, jul 2011.
- [51] Tommaso Caneva, Tommaso Calarco, and Simone Montangero. Chopped random-basis quantum optimization. *Phys. Rev. A*, 84:022326, Aug 2011.
- [52] 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. *Phys. Rev. A*, 92:062343, Dec 2015.
- [53] John A. Nelder and Roger Mead. A simplex method for function minimization. *Comput. J.*, 7:308–313, 1965.
- [54] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195, 2001.
- [55] Rossignolo Marco, Marshall Alastair, Reisser Thomas, Rembold Phila, and Pagano Alice. Quantum optimal control suite. <https://github.com/Quantum-OCS/QuOCS>, 2022.
- [56] Michael H. Goerz, Daniel M. Reich, and Christiane P. Koch. Optimal control theory for a unitary operation under dissipative evolution. 2014.
- [57] Sinan Karaveli, Ophir Gaathon, Abraham Wolcott, Rey Sakakibara, Or A. Shemesh, Darcy S. Peterka, Edward S. Boyden, Jonathan S. Owen, Rafael Yuste, and Dirk Englund. Modulation of nitrogen vacancy charge state and fluorescence in nanodiamonds using electrochemical potential. *Proceedings of the National Academy of Sciences*, 113(15):3938–3943, 2016.

- [58] Wenzheng Dong, F A Calderon-Vargas, and Sophia E Economou. Precise high-fidelity electron–nuclear spin entangling gates in NV centers via hybrid dynamical decoupling sequences. *New Journal of Physics*, 22(7):073059, jul 2020.
- [59] Jianming Cai, Alex Retzker, Fedor Jelezko, and M. Plenio. A large-scale quantum simulator on a diamond surface at room temperature. *Nature Physics*, 9, 08 2012.
- [60] Yang Wu, Ya Wang, Xi Qin, Xing Rong, and Jiangfeng Du. A programmable two-qubit solid-state quantum processor under ambient conditions. 08 2018.
- [61] J. F. Haase, P. J. Vetter, T. Uden, A. Smirne, J. Roskopf, B. Naydenov, A. Stacey, F. Jelezko, M. B. Plenio, and S. F. Huelga. Controllable non-markovianity for a spin qubit in diamond. *Phys. Rev. Lett.*, 121:060401, Aug 2018.
- [62] Yuliya Mindarava, Rémi Blinder, Yan Liu, Jochen Scheuer, Johannes Lang, Viatcheslav Agafonov, Valery A. Davydov, Christian Laube, Wolfgang Knolle, Bernd Abel, Boris Naydenov, and Fedor Jelezko. Synthesis and coherent properties of ^{13}C -enriched sub-micron diamond particles with nitrogen vacancy color centers. *Carbon*, 165:395–403, 2020.
- [63] Matthias Müller, Stefano Gherardini, and Filippo Caruso. Quantum zeno dynamics through stochastic protocols. *Annalen der Physik*, 529, 07 2016.