


QUANTUM OPTIMAL CONTROL FOR NV CENTRES USING GATE SET TOMOGRAPHY



David M. Vos

QUANTUM OPTIMAL CONTROL FOR NV CENTRES USING GATE SET TOMOGRAPHY

In partial fulfilment of the requirements for the degree of Master of Science in
Computer Science,
at the Delft University of Technology,
in the Taminiau Lab, Quantum Internet Division of QuTech.
To be defended publicly on Monday June 23, 2025 at 10:30 o'clock.

by

David M. Vos

Student number:	4667050
Project duration:	2 October, 2024 – June 23, 2025
Daily supervisors:	B. van Ommen, MSc. QuTech M. van Riggelen, MSc. QuTech
Defence committee:	Dr. ir. T.H. Taminiau, TU Delft, supervisor Dr. S. Wehner, TU Delft, second examiner Dr. M. Russ, TU Delft, second examiner



*In theory, theory and practice are the same.
In practice, they are not.*

ABSTRACT

We investigate open- and closed-loop quantum optimal control in nitrogen-vacancy (NV) centre systems, using gate set tomography (GST) for pulse characterisation and calibration. Open-loop optimisation with GRAPE revealed a strong dependence of gate fidelity on pulse duration and the need for phase correction in echo-based gates.

Despite promising simulations of advanced pulse designs, experimental performance was consistently higher for a standard weak π pulse, highlighting model limitations. Closed-loop optimisation with dCRAB improved performance but exhibited sensitivity to environmental drift, confirming magnet-induced variations.

GST identified dominant coherent error channels, notably ZZ , attributed to AC Stark-induced phase shifts. A novel gate design suppressed ZZ errors but introduced greater stochastic noise.

These results demonstrate the importance of integrating model-driven and data-driven approaches with advanced tomography protocols to achieve robust, high-fidelity quantum control in noisy intermediate-scale quantum (NISQ) devices.

CONTENTS

1	Introduction	1
2	Theoretical Background	4
2.1	NV Centre in Diamond	4
2.1.1	Optical Spin Polarisation and Read-out	4
2.1.2	System Hamiltonian and Coherent Control	6
2.1.3	Conditional Electron Spin Control with a Weak π Gate	7
2.1.4	Dynamical Decoupling	9
2.2	Quantum Optimal Control	10
2.2.1	GRAPE	11
2.2.2	dCRAB	13
2.3	Gate Set Tomography	15
2.3.1	Mathematical model	16
2.3.2	Experimental Protocol	16
2.3.3	Parameter Estimation	17
2.3.4	Gauge Optimisation and Error Metrics	18
2.3.5	Used Metrics	18
2.3.6	Error Generators	19
3	Methods	21
3.1	QuTiP	22
3.1.1	qutip-qoc	22
3.1.2	Dynamical Decoupling	24
3.1.3	Virtual phase corrections	24
3.2	QuOCS	25
3.2.1	Optimisation Dictionary	25
3.2.2	Construction of the Final Pulse with dCRAB	26
3.2.3	Figure of Merit	28
3.3	pyGSTi	30
3.4	Experimental Setup	31
3.4.1	Optical System	31
3.4.2	Signal Chain	32
4	Results and Discussion	34
4.1	Open-Loop Optimal Control	34
4.1.1	Produced Pulses	35
4.1.2	Phase Evolution and Correction Mechanisms	36
4.1.3	Amplitude Sweeps	39
4.1.4	PL Contrast and Bell State Fidelity	40

4.2	Closed Loop Optimal Control	42
4.2.1	Produced Pulses	43
4.2.2	PL Contrast and Bell State Fidelity	44
4.2.3	Noise Limitations	45
4.3	Gate set tomography	47
4.3.1	Non-Markovianity	47
4.3.2	Average Gate Fidelity	49
4.3.3	Error Generator	49
4.4	Weak $\pi/2$ Sequence	52
4.4.1	Gate Set Tomography	53
5	Conclusion and outlook	57
5.1	Conclusion	57
5.1.1	Open-loop optimisation	57
5.1.2	Closed-loop optimisation	58
5.1.3	Gate set tomography and error analysis	58
5.2	Outlook	58
5.2.1	Refining open-loop approaches	58
5.2.2	Extending closed-loop optimisation	59
5.2.3	Advanced error characterisation	59
5.2.4	Towards fault tolerance and AI-assisted control	59
	Acknowledgements	61
A	Theoretical Background	68
A.1	Uncertainty comparison	68
B	Code Listings	70
B.1	Custom Helper Functions	70
B.2	QuOCS Optimisation dictionary	71
B.3	PyGSTi	72
C	Experimental Results	73

NOMENCLATURE

AOM	Acousto-Optic Modulator
APD	Avalanche Photodiode
AWG	Arbitrary Waveform Generator
CNOT	Controlled-NOT gate
dCRAB	dressed Chopped Random Basis algorithm
DD	Dynamical decoupling
DDRF	Dynamically Decoupled Radio-Frequency
FoM	Figure of Merit
GRAPE	Gradient Ascent Pulse Engineering
GST	Gate Set Tomography
MW	Microwave
NV	Nitrogen–vacancy centre in diamond
PL	Photoluminescence
RF	Radio-frequency
RWA	Rotating-Wave Approximation
pyGSTi	Python Gate Set Tomography Implementation
QuOCS	Quantum Optimal Control Suite
QuTiP	Quantum Toolbox in Python
qutip-qoc	Quantum Optimal Control extension for QuTiP
Δ	Detuning between drive and transition frequency
γ_e	Electron gyromagnetic ratio ($2\pi \times 4.18 \text{ MHz G}^{-1}$)
γ_N	^{14}N gyromagnetic ratio ($2\pi \times 0.308 \text{ kHz G}^{-1}$)
$\Omega(t)$	Time-dependent Rabi frequency of the drive
ϕ	Instantaneous microwave phase

τ	Free-evolution interval in an echo sequence	
A_{\parallel}	Axial hyperfine coupling (≈ -2.18 MHz for Nitrogen)	
B_0	Static magnetic field along the NV axis	
D	Zero-field splitting ($2\pi \times 2.87$ GHz)	<i>Eq. 2.1</i>
F	Gate (or state) fidelity	
P	Parity observable used in Bell-state analysis	

1

INTRODUCTION

QUANTUM COMPUTING AND ITS PROGRESS

Quantum computing harnesses the principles of quantum mechanics to process information in fundamentally new ways. It aims to outperform classical computers in specific tasks by exploiting phenomena such as superposition, entanglement, and quantum interference. The theoretical foundation of quantum computation was first proposed by Paul Benioff in 1980, who introduced a quantum mechanical model of a Turing machine [1]. Shortly thereafter, Richard Feynman and Yuri Manin argued that simulating quantum systems would require quantum computers, as classical computers face exponential overhead in such simulations [2], [3].

A major milestone was reached in 1994 with the introduction of Shor's algorithm, which demonstrated that quantum computers could efficiently factor large integers, threatening classical cryptographic schemes such as RSA [4]. The first physical realisation of a quantum algorithm followed in 1998, implemented using a two-qubit nuclear magnetic resonance (NMR) system [5].

Since then, the field has experienced remarkable progress. Single-qubit gate fidelities have improved from approximately 90% in early demonstrations to 99.99916% in trapped-ion systems [6], 99.998% in superconducting circuits using fluxonium qubits [7] and 99.999% in NV centre systems [8]. Two-qubit gate fidelities, once a significant bottleneck, have also seen substantial progress. State-of-the-art ion trap platforms now exceed 99.9% [9], NV centre systems have reached 99.93% [8], and superconducting systems have achieved comparable performance through optimised control and coupling strategies [10]. Despite these advances, quantum processors remain in the Noisy Intermediate-Scale Quantum (NISQ) regime, limited by noise and decoherence [11]. Achieving fault-tolerant quantum computation requires two-qubit gate fidelities well above 99.9%, in conjunction with efficient error correction protocols [12].

NV CENTRES AS A QUANTUM PLATFORM

Compared to other architectures, nitrogen-vacancy (NV) centres in diamond offer long coherence times and optical addressability. These properties make them strong candi-

dates for scalable quantum computing and networking [13], though they also pose challenges in achieving high-fidelity multiqubit control due to weak couplings and complex spin environments.

QUANTUM OPTIMAL CONTROL

One strategy to meet these challenges is quantum optimal control, which combines classical control theory with quantum dynamics. Since its emergence in the early 2000s [14], this framework has achieved notable success across diverse platforms, including trapped ions [15], superconducting qubits [16], and NV centres [17], [18].

Optimal control techniques are generally classified as either open-loop or closed-loop. Open-loop methods use simulations and accurate system models to design control pulses, while closed-loop control incorporates experimental feedback to iteratively refine the pulse, thereby compensating for model inaccuracies. Algorithms such as GRAPE (Gradient Ascent Pulse Engineering) [14] and dCRAB (dressed Chopped Random Basis) [19] can be employed in either setting. However, GRAPE is most commonly used in open-loop simulations due to its reliance on gradient information, whereas dCRAB is particularly well suited to closed-loop optimisation, where gradients may be unavailable or unreliable.

QUANTUM GATE CHARACTERISATION

Evaluating the performance of an implemented quantum gate is essential for assessing how closely the realised operation approximates the ideal target, and for identifying sources of error that may limit fidelity. While this assessment is distinct from the optimisation procedure itself, it plays a crucial role in validating the outcome and informing future control strategies. In open-loop control, performance evaluation provides experimental confirmation of model-based predictions; in closed-loop control, it helps interpret the effectiveness of feedback-based improvements.

To characterise quantum gate performance, a range of diagnostic tools has been developed, each with different trade-offs. Quantum state and process tomography offer full characterisation of quantum operations but suffer from poor scalability and significant sensitivity to state preparation and measurement (SPAM) errors. Randomised benchmarking (RB), by contrast, is a scalable benchmarking protocol that yields an average error rate over randomised sequences of gates [20]. It is robust against SPAM but does not resolve gate-specific errors or distinguish between coherent and incoherent error sources.

Gate Set Tomography (GST) occupies a complementary role: it is a tomographic method that provides SPAM-robust, self-consistent reconstructions of an entire gate set, including state preparation and measurement processes [21]. Unlike RB, GST yields detailed information about the structure and origin of errors in individual gates, making it particularly useful for diagnosing limitations in optimised control pulses and guiding subsequent improvements.

SCOPE AND STRUCTURE OF THIS THESIS

This thesis demonstrates the implementation of a high-fidelity two-qubit gate in an NV centre system using both open-loop and closed-loop optimal control methods. The resulting operations are characterised using GST, which also guides iterative pulse refine-

ment. Although prior work has applied optimal control to NV systems [17], to our knowledge, this is the first study to combine these techniques with comprehensive GST-based validation.

Our results highlight the synergistic potential of optimal control and advanced characterisation to push gate fidelities closer to thresholds for fault tolerance.

This thesis is structured as follows.

Chapter 2 provides the theoretical background, introducing the NV centre in diamond as our experimental platform, the principles of quantum optimal control—focusing on the GRAPE and dCRAB algorithms—and the GST framework.

Chapter 3 outlines the software and hardware methods used. On the software side, we describe the use of QuTiP for open-loop optimal control, QuOCS for closed-loop optimisation, and pyGSTi for gate characterisation. On the hardware side, we detail the experimental setup and control infrastructure.

Chapter 4 presents and discusses our results. We begin with open-loop optimal control, proceed to closed-loop optimisation, and then describe the GST-based characterisation of our implemented gate. We also explain how these characterisation results inform our pulse design and implementation, leading to the implementation of a weak $\frac{1}{2}\pi$ sequence.

Chapter 5 concludes the thesis by summarising the key findings and outlining potential directions for future research.

2

THEORETICAL BACKGROUND

In this chapter, we present the foundational concepts and models that underpin the experimental and computational work in this thesis. We begin with an overview of the nitrogen–vacancy (NV) centre in diamond, a versatile platform for solid-state quantum technologies. We then introduce the principles of quantum optimal control, focusing on the GRAPE and dCRAB algorithms, and we conclude with gate set tomography—a robust protocol for characterising quantum operations.

2.1. NV CENTRE IN DIAMOND

Diamond’s nitrogen–vacancy (NV) centre is a point defect consisting of a substitutional nitrogen atom next to a nearest-neighbour carbon vacancy (Figure 2.1). We create a controlled density of NVs in high-purity, lab-grown diamond by (i) incorporating nitrogen during chemical vapour deposition, (ii) generating vacancies through electron, ion, or neutron irradiation, and (iii) annealing at $\sim 800\text{--}1000^\circ\text{C}$, which mobilises vacancies so that they pair with substitutional nitrogen atoms [23].

We can stabilise the centre in two charge configurations, NV^0 and NV^- . Throughout this thesis we employ only the spin-triplet ground state of NV^- as a solid-state qubit because it supports optical spin polarisation, coherent microwave control, and spin-dependent fluorescence read-out [22].

2.1.1. OPTICAL SPIN POLARISATION AND READ-OUT

In this work, we use off-resonant excitation at 515 nm to prepare the NV centre in the negatively charged state (NV^-). Spin-selective resonant excitation at 637 nm is then employed to initialise and read out the electron spin state.

The spin-pump laser at 637 nm selectively excites the electron spin if it is in $m_S = \pm 1$, which, through spin-dependent intersystem crossing to the metastable singlet states, efficiently pumps the spin into the $m_S = 0$ ground sublevel [24]. The subsequent resonant readout laser also addresses only the $m_S = 0$ state, enabling projective spin readout via spin-dependent fluorescence.




Figure 2.1: **Crystal structure and molecular orbitals of the NV centre.** (a) Schematic of a substitutional nitrogen atom (green) adjacent to a vacancy (white) in the diamond lattice. Grey spheres denote the ^{12}C lattice; a single ^{13}C isotope (yellow) with nuclear spin $I = 1/2$ naturally occurs with 1.1% abundance. (b) One-electron molecular-orbital diagram of the negatively charged NV^- centre showing the occupation of the a_1 and e defect levels in the triplet ground state. Adapted from Ref. [22].

READOUT UNCERTAINTY

To estimate the number of shots needed to achieve a desired level of statistical certainty in quantum readout, we model the problem as a binomial distribution: each shot yields either a 0 or a 1. Let n be the number of shots, and k the number of observed 1's. The empirical readout frequency is then $f = \frac{k}{n}$, with expectation value p , the true probability of measuring a 1.

Since $k \sim \text{Binomial}(n, p)$, the variance of k is $\text{Var}(k) = np(1 - p)$. The variance of f , the proportion, is therefore:

$$\text{Var}(f) = \frac{p(1 - p)}{n}$$

and the standard deviation is (using f to estimate p):

$$\sigma_f = \sqrt{\frac{p(1 - p)}{n}} \approx \sqrt{\frac{f(1 - f)}{n}}$$

Suppose we wish the uncertainty to be a fraction $1/R$ of the deviation from perfect fidelity, that is, $\sigma_f = \frac{1-f}{R}$. Solving for n gives:

$$\frac{1-f}{R} = \sqrt{\frac{f(1-f)}{n}} \Rightarrow n = \frac{R^2 f}{1-f}$$

For example, if $f = 0.9$ and $R = 100$, we obtain:

$$n = \frac{100^2 \times 0.9}{1 - 0.9} = 900$$

This formulation provides a practical way to plan shot counts based on desired expected fidelity.

In the context of optimisation, the choice of shot count directly affects the ability to find gradients within the control landscape. Too few shots introduce significant stochastic noise into fidelity estimates, obscuring the underlying landscape structure and impeding gradient estimation or convergence of heuristic optimisers. Conversely, excessive shot counts slow down the optimisation loop without proportional benefit once statistical uncertainty is sufficiently low. The trade-off is particularly important in experimental optimisation, where each shot consumes valuable time and resources. By tuning n according to the expected fidelity and required confidence (e.g., using the criterion above), one can maintain a balance between statistical reliability and practical optimisation speed, thereby promoting efficient exploration of the control landscape.

2.1.2. SYSTEM HAMILTONIAN AND COHERENT CONTROL

For a single NV⁻ centre coupled to its host ¹⁴N nucleus ($I = 1$), the secular Hamiltonian in the laboratory frame is [22], [23]

$$\hat{H}_{\text{lab}} = D\hat{S}_z^2 + \gamma_e B_0 \hat{S}_z + \gamma_n B_0 \hat{I}_z + \hat{\mathbf{S}} \cdot \mathbf{A} \cdot \hat{\mathbf{I}} + P\hat{I}_z^2 + 2\Omega(t) \cos(\omega_{\text{MW}}t + \phi) \hat{S}_x, \quad (2.1)$$

where

- $D/2\pi = 2.87 \text{ GHz}$ is the ground-state zero-field splitting;
- $\gamma_e = 2\pi \times 2.802 \text{ MHz G}^{-1}$ and $\gamma_n = 2\pi \times 0.3077 \text{ kHz G}^{-1}$ are the electron and ¹⁴N nuclear gyromagnetic ratios, respectively;
- $\mathbf{A} = \text{diag}(A_{XX}, A_{YY}, A_{\parallel})$ is the hyperfine tensor ($A_{\parallel}/2\pi \approx -2.189 \text{ MHz}$ for ¹⁴N);
- $P/2\pi \approx -4.95 \text{ MHz}$ is the nuclear quadrupole splitting of the ¹⁴N spin;
- B_0 is a static magnetic field aligned with the NV axis (z -axis);
- $\Omega(t)$ and ϕ denote the time-dependent Rabi frequency and phase of the near-resonant microwave drive.

Transforming to the interaction picture with respect to $\hat{H}_0 = D\hat{S}_z^2 + \gamma_e B_0 \hat{S}_z + \gamma_n B_0 \hat{I}_z$ and applying the rotating-wave approximation (RWA) for a drive near the $|m_S = 0\rangle \leftrightarrow |m_S = -1\rangle$ transition yields

$$\hat{H}_{\text{NV}} = \Delta \hat{S}_z + \hat{\mathbf{S}} \cdot \mathbf{A} \cdot \hat{\mathbf{I}} + P\hat{I}_z^2 + \Omega(t) [\cos \phi \hat{S}_x + \sin \phi \hat{S}_y], \quad (2.2)$$

with detuning

$$\Delta = \omega_{\text{MW}} - [D - \gamma_e B_0 + A_{\parallel}]. \quad (2.3)$$

Assuming a fixed nuclear spin state (e.g. $m_I = +1$) and restricting to the electronic subspace $\{|0\rangle, |-1\rangle\}$, Eq. (2.2) reduces to the effective single-qubit Hamiltonian

$$\hat{H}_{\text{RWA}}^{(e)} = -\Delta |-1\rangle\langle -1| + \frac{\Omega(t)}{2} [\cos \phi \hat{\sigma}_x^{(e)} + \sin \phi \hat{\sigma}_y^{(e)}], \quad (2.4)$$

where $\hat{\sigma}_{x,y}^{(e)}$ are Pauli operators acting on $\{|0\rangle, |-1\rangle\}$.

2.1.3. CONDITIONAL ELECTRON SPIN CONTROL WITH A WEAK π GATE

Coherent control of NV-centre qubits is achieved using microwave (MW) and radio-frequency (RF) pulses. This section focuses on the weak π gate—a conditional operation that implements a CROTX gate, where the nitrogen nuclear spin acts as the control and the electron spin as the target.

COMPARISON WITH THE DDRF GATE

The weak π gate applies a conditional rotation to the electron spin, with the rotation occurring only when the nuclear spin is in a specific state. This is achieved through spectral selectivity: by tuning the microwave frequency to match the hyperfine-split transition associated with a particular nuclear spin state, the electron spin undergoes a π rotation if and only if the nuclear spin is in that state. The resulting gate is straightforward to implement and inherently robust against certain types of control errors.

Our group previously developed a high-fidelity two-qubit gate known as the DDRF (Dynamically Decoupled Radio-Frequency) gate [8], which achieved a GST fidelity of 99.93%. The DDRF gate applies a conditional $\pm\pi/2$ rotation to the nuclear spin, conditioned on the electron spin state. Despite its high fidelity, its long duration imposes limitations for certain applications.

A key distinction lies in the direction of conditionality: the DDRF gate conditions a nuclear spin rotation on the electron spin, whereas the weak π gate conditions an electron spin rotation on the nuclear spin state. Moreover, their physical mechanisms differ: the DDRF gate employs differential phase accumulation modulated by dynamical decoupling, while the weak π gate relies purely on spectral selectivity enabled by hyperfine splitting.

NUMERICAL VALIDATION OF CONDITIONALITY

The term “weak π ” refers to a pulse with lower amplitude and longer duration than a conventional strong π pulse. While strong π pulses affect unconditional electron spin flips—typical in dynamical decoupling sequences—the weak π pulse is designed to flip the electron spin only if the nuclear spin is in the $m_I = 0$ subspace. This selectivity enables conditional gate operations such as the controlled- $R_x(\pi)$ (CROTX).

Figure 2.2 illustrates this principle. In the time domain (left), the strong pulse (blue) is brief and intense, while the weak pulse (orange) is longer and weaker.

To quantitatively validate this selectivity, we simulate the driven evolution of a two-level system under a Gaussian-shaped pulse using the time-dependent Rabi Hamiltonian:

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

where Δ is the detuning and $\Omega(t) = \Omega_0 \exp\left(-\frac{(t-t_0)^2}{2\sigma^2}\right)$ is the time-dependent Rabi frequency with a Gaussian envelope centred at t_0 and width σ . Figure 2.2(b) shows the excitation probability as a function of detuning.

The Gaussian weak π pulse sharply suppresses off-resonant excitation, confirming that only transitions near resonance are driven. This spectral confinement is a key design feature. Compared to a rectangular shape, the Gaussian envelope minimises spectral side lobes and reduces crosstalk with undesired hyperfine transitions. These prop-




Figure 2.2: Illustration of strong and weak π pulses in the NV centre system. **Left:** Time-domain view: the strong pulse (blue) is short and high in amplitude, while the weak pulse (orange) is longer and lower in amplitude. **Right:** Simulated spin excitation probability versus detuning. A strong Gaussian π pulse (blue) excites all three ODMR resonances uniformly, while a weak pulse (orange) is selective, only inducing transitions near zero detuning. This confirms the conditional nature of the weak π gate.

erties make the Gaussian weak π gate particularly effective for conditional control of the electron spin in NV-based quantum gates.

UNITARY AND PHASE CORRECTION

Although the weak π gate induces a full π rotation on the electron spin when the nuclear spin is in the $m_I = 0$ state, its unitary action is subtly different from an ideal CNOT gate. The CNOT gate is defined as

$$U_{\text{CNOT}} = \mathbb{1}_e \otimes |0\rangle\langle 0|_N + X_e \otimes |1\rangle\langle 1|_N,$$

which satisfies $U_{\text{CNOT}}^2 = \mathbb{1}$, since $X^2 = \mathbb{1}$.

In contrast, the weak π gate implements the controlled- $R_x(\pi)$ operation,

$$CR_x(\pi) = \mathbb{1}_e \otimes |0\rangle\langle 0|_N + R_x(\pi) \otimes |1\rangle\langle 1|_N, \quad R_x(\pi) = -iX.$$

This gate differs from U_{CNOT} by a controlled phase, which affects the gate's behaviour under composition.

To explicitly show the relation, consider applying a virtual Z rotation on the control qubit (nitrogen spin):

$$R_z\left(-\frac{\pi}{2}\right) = e^{i\pi/4} |0\rangle\langle 0| + e^{-i\pi/4} |1\rangle\langle 1|.$$

Then,

$$U_{\text{corrected}} = \left(\mathbb{1}_e \otimes R_z\left(-\frac{\pi}{2}\right)\right) CR_x(\pi) = e^{i\pi/4} \mathbb{1}_e \otimes |0\rangle\langle 0| - i e^{-i\pi/4} X_e \otimes |1\rangle\langle 1|.$$

This unitary is equivalent to U_{CNOT} up to a global phase and a known single-qubit phase on the nuclear spin.

In our implementation, this residual phase is corrected via a virtual Z rotation on the nitrogen spin, converting the CROTX gate into an ideal CNOT under composition. This ensures that the gate behaves as expected when embedded within larger quantum circuits.

2.1.4. DYNAMICAL DECOUPLING

To mitigate decoherence caused by continuous coupling between the electron spin and its environment, dynamical decoupling (DD) techniques such as spin echo are essential. In this work, we employ decoupling sequences to preserve electron spin coherence by cancelling the effects of slowly varying environmental noise. The interpulse delay τ is optimised for our system as described in Ref. [22].

To maximise the electron coherence, we avoid interpulse delays τ that satisfy the condition

$$\tau = \frac{(2k-1)\pi}{2\omega_L - A_{\parallel}},$$

for any nuclear spin coupled to the electron, as these values correspond to resonant electron-nuclear interactions that degrade coherence. Here, ω_L is the electron Larmor frequency and A_{\parallel} is the parallel hyperfine coupling to a given nuclear spin. For our system, $\omega_L \approx 4.18$ GHz and $A_{\parallel} \approx 2.18$ MHz for the strongest coupled spin, which is the nitrogen spin. However, the dominant source of decoherence in our system arises from coupling to nearby ^{13}C nuclear spins, which are naturally abundant in the diamond lattice. We determine the optimal τ by sweeping over candidate values and measuring coherence preservation, selecting $\tau = 12.96 \mu\text{s}$ to avoid resonances and maximise coherence.

DD aims to preserve quantum coherence by applying a sequence of precisely timed π -pulses that average out unwanted environmental interactions. When noise varies slowly compared to the pulse interval, each pulse inverts the qubit's evolution, effectively cancelling accumulated phase errors from low-frequency fluctuations. This process can be understood intuitively: during the first free evolution period, noise causes phase accumulation; a π -pulse then flips the qubit frame, so that the subsequent evolution accumulates phase with the opposite sign. If the noise is sufficiently slow, these contributions cancel over each cycle. A more formal perspective is provided by the toggling-frame picture: each pulse switches the sign of the noise Hamiltonian, alternating the effective evolution. Over time, these toggled contributions average to zero, suppressing unwanted couplings and extending the qubit's coherence time.

While basic two-pulse echo sequences can cancel the leading-order effects of static or slowly varying noise, more sophisticated sequences like CPMG, XY4, and XY8 use symmetries to cancel higher-order errors [25].

Figure 2.3 shows two examples: the XY4 sequence, which alternates π -pulses along the X and Y axes, and the XY8 sequence, which symmetrically extends XY4 to suppress more subtle error terms. These sequences are robust to both dephasing noise and pulse imperfections, making them widely used in high-fidelity quantum control.

In this work, we do not employ full XY4 or XY8 decoupling blocks. Instead, we place a single X and $-X$ pulse before and after the pulse being optimised. This choice is motivated by the relatively short duration of the control pulse, during which the environmental noise is assumed to remain quasi-static. In this regime, an echo structure of the CPMG type suffices to cancel low-frequency environmental noise. Although XY4 and XY8 sequences offer enhanced robustness against pulse errors, they provide comparable protection against decoherence from such noise and are therefore not required in this context.

An excellent pedagogical overview of the underlying theory, including a formal deriva-




Figure 2.3: **Dynamical decoupling sequences: XY4 (upper) and XY8 (lower).** Each sequence consists of a series of π pulses applied along specific axes to suppress decoherence due to environmental noise. A free evolution period τ precedes and follows the pulse block, while adjacent pulses are separated by 2τ delays. The XY4 sequence applies X , Y , $-X$, and $-Y$ pulses; XY8 extends this with a symmetric time-reversed mirror to cancel higher-order errors. These sequences are widely used to preserve spin coherence in NV-centre experiments and other qubit systems [25].

tion using average Hamiltonian theory (AHT), can be found in Ref. [26]. While we omit the full derivation here, AHT formalises the idea that the qubit evolves as if governed by an average Hamiltonian over the pulse cycle. A well-designed pulse sequence ensures that this average retains only the desired interactions, while cancelling out noise and unwanted couplings.

2.2. QUANTUM OPTIMAL CONTROL

Optimal control theory originated in the 1950s, primarily through Lev Pontryagin's Maximum Principle and Richard Bellman's dynamic programming [27], [28]. In essence, optimal control determines the control functions that optimise a given objective functional subject to the dynamics of a system. Over the decades, this framework has found extensive application in aerospace engineering [29], robotics [30], chemical process control [31], automotive engineering [32], and economics [33].

To meet the growing need for precise control of quantum states, researchers extended optimal control to the quantum domain in the early 1980s. Early applications included shaped laser pulses for molecular dynamics and advanced NMR techniques [34], [35]. The field expanded in the 1990s with gradient-based algorithms such as Krotov's method [36]. A major milestone came in 2005 with the Gradient Ascent Pulse Engineering (GRAPE) method, which has since become a standard technique for designing high-fidelity control pulses [37].

2.2.1. GRAPE

GRAPE is a widely used numerical optimal-control method introduced by Khaneja *et al.* [37]. It provides a systematic way for us to design piecewise-constant control pulses that drive a quantum system toward a desired evolution with high fidelity. This subsection draws partly on [38].

SETUP AND OBJECTIVE

Consider a quantum system described by the Hamiltonian [37]

$$\hat{H}(t) = \hat{H}_0 + \sum_{i=1}^n u_i(t) \hat{H}_i, \quad (2.5)$$

where:

- \hat{H}_0 is the *drift* Hamiltonian for free evolution,
- $u_i(t)$ are time-dependent control amplitudes, and
- \hat{H}_i are the corresponding control Hamiltonians.

In our system, \hat{H}_0 encodes the hyperfine interaction, while \hat{H}_i generate driving in the x - or y -direction on the Bloch sphere. Specifically, we identify:

$$\hat{H}_0 = \Delta S_z \otimes I_d + A_{\parallel} S_z \otimes I_z, \quad \hat{H}_1 = S_x \otimes I_d, \quad \hat{H}_2 = S_y \otimes I_d.$$

Rather than targeting a specific state, our goal is to implement a target quantum gate \hat{U}_{target} by driving the system with control pulses such that the resulting time-evolution operator $\hat{U}(T)$ approximates \hat{U}_{target} as closely as possible. We quantify success by the gate fidelity

$$F = \frac{1}{d^2} \left| \text{Tr} \left(\hat{U}_{\text{target}}^\dagger \hat{U}(T) \right) \right|^2,$$

where $d = 6$ is the dimension of the system Hilbert space. This fidelity measure, known as the gate unitary overlap, directly captures how well the implemented evolution matches the target gate up to a global phase.

While some implementations of GRAPE optimise for state-transfer fidelity,

$$F_{\text{state}} = \left| \langle \psi_{\text{target}} | \psi(T) \rangle \right|^2,$$

where $\psi(t)$ is the system state evolved under the applied controls and ψ_{target} is a desired target state. Our work uses the gate fidelity F to evaluate performance at the operator level. This is more appropriate when designing universal gate sets or controlling entangling operations.

We maximise the fidelity via gradient ascent while the evolution operator obeys the time-dependent Schrödinger equation,

$$i\hbar \frac{\partial}{\partial t} \hat{U}(t) = \hat{H}(t) \hat{U}(t), \quad \hat{U}(0) = \mathbb{I}.$$

PIECEWISE-CONSTANT APPROXIMATION

In practice, an arbitrary waveform generator (AWG) represents each control amplitude $u_i(t)$ in discrete steps. We adopt the same piecewise-constant approximation:

$$u_i(t) = u_i(j) \quad \text{for } (j-1)\delta t \leq t < j\delta t,$$

where $\delta t = T/N$ is the duration of each slice, and $j \in \{1, \dots, N\}$.

Under this approximation, the time-evolution operator $\hat{U}(T)$ factorizes into a product of operators over each slice:

$$\hat{U}(T) = \hat{U}_N \hat{U}_{N-1} \dots \hat{U}_1,$$

where each \hat{U}_k is given by

$$\hat{U}_k = \exp\left(-\frac{i}{\hbar} \delta t [\hat{H}_0 + \sum_i u_i(k) \hat{H}_i]\right). \quad (2.6)$$

Hence, the fidelity can be written as

$$F = \frac{1}{d^2} \left| \text{Tr} \left(\hat{U}_{\text{target}}^\dagger \hat{U}_N \dots \hat{U}_1 \right) \right|^2.$$

FORWARD AND BACKWARD PROPAGATION

In GRAPE each time step at t_m is defined as:

$$|\rho_m\rangle = \hat{U}_m \dots \hat{U}_1 |\psi_0\rangle, \quad |\lambda_m\rangle = \hat{U}_N \dots \hat{U}_{m+1} |\psi_1\rangle.$$

The state $|\rho_m\rangle$ is the *forward*-propagated initial state up to step m , while $|\lambda_m\rangle$ is the *backward*-propagated target state from step N down to step $m+1$. These partially propagated states meet at t_m , allowing us to compute the fidelity gradient with respect to each control amplitude $u_i(m)$.

GRADIENT DERIVATION

Figure 2.4 illustrates the piecewise-constant control amplitude together with the corresponding gradients that are optimised in GRAPE.

To derive the gradient of F with respect to $u_i(j)$, we need the derivative of the matrix exponential. A well-known identity (see [39]) states that

$$\left. \frac{d}{dt} \right|_{t=0} e^{X+tY} = e^X \int_0^1 d\tau e^{\tau X} Y e^{-\tau X}. \quad (2.7)$$

For sufficiently small δt , the commutator terms in the integral can often be neglected to first order, giving

$$\left. \frac{d}{dt} \right|_{t=0} e^{X+tY} \approx e^X Y.$$

Applying this to one time slice where $X = -\frac{i}{\hbar} \delta t \hat{H}(k)$ and $Y = -\frac{i}{\hbar} \delta t \hat{H}_i$ yields the key simplification in the GRAPE formula.




Figure 2.4: Schematic representation of a control amplitude u_j consisting of N time steps of duration $\delta t = T/N$. During each time step j , the control amplitude $u_i(j)$ is constant. The vertical arrows represent the gradients $\partial F / \partial u_i(j)$, indicating how each amplitude should be adjusted to improve the performance metric F . Image adapted from [38].

After algebra, one obtains the *closed-form* gradient expression:

$$\frac{\partial F}{\partial u_i(j)} = -i \delta t \langle \lambda_j | \hat{H}_i | \rho_j \rangle. \quad (2.8)$$

To *maximize* F , we update $u_i(j)$ via gradient ascent:

$$u_i(j) \mapsto u_i(j) + \epsilon \frac{\partial F}{\partial u_i(j)}, \quad (2.9)$$

where ϵ is the step size (learning rate).

ALGORITHMIC IMPLEMENTATION

Algorithm 1 outlines the basic steps of the GRAPE algorithm. One initialises a guess for the control amplitudes $\{u_i(j)\}$, computes the fidelity and its gradient, then iteratively updates the controls until convergence. The work conducted in this thesis uses QuTiP-qoc for the practical implementation.

2.2.2. dCRAB

dressed Chopped Random Basis (dCRAB) is a gradient-free optimal-control scheme that extends the original CRAB approach of Caneva et al. [40]. Instead of updating each time sample individually (as in GRAPE), dCRAB expands every control amplitude in a low-dimensional, randomly chosen Fourier-like basis and optimises only the expansion coefficients. Whenever the search stalls, we dress the basis with a fresh set of random frequencies, allowing the algorithm to escape local optima while still keeping the number of free parameters small. Because dCRAB never requires gradients, it is particularly attractive for experimental closed-loop optimisation and for systems where evaluating $\partial F / \partial u_i(t)$ is costly or noisy. This subsection follows Refs. [38], [41], [42]. Optimisation proceeds on two complementary control landscapes, as we illustrate in Fig. 2.5.

Algorithm 1 GRAPE Algorithm

Require: **Input:** Initial guess $\{u_i(j)\}$, step size ϵ , convergence tolerance δ , maximum iterations N_{\max}

Ensure: **Output:** Optimised control parameters $\{u_i(j)\}$

```

1: Initialise control parameters  $\{u_i(j)\}$ 
2: Compute initial fidelity  $F(\{u_i(j)\})$ 
3:  $n \leftarrow 0$ 
4: while  $n < N_{\max}$  and  $\Delta F > \delta$  do
5:   Compute forward-propagated states  $|\rho_m\rangle$  and backward-propagated states  $|\lambda_m\rangle$ 
6:   Compute gradient  $\partial F / \partial u_i(j)$ 
7:   Update controls:  $u_i(j) \leftarrow u_i(j) + \epsilon \partial F / \partial u_i(j)$ 
8:   Compute new fidelity  $F_{\text{new}}$ 
9:    $\Delta F = |F_{\text{new}} - F_{\text{old}}|$ 
10:   $F_{\text{old}} \leftarrow F_{\text{new}}$ 
11:   $n \leftarrow n + 1$ 
12: end while
13: return  $\{u_i(j)\}$ 

```

▷ Optimised controls

SETUP AND OBJECTIVE

We again consider the time-dependent Hamiltonian

$$\hat{H}(t) = \hat{H}_0 + \sum_{i=1}^n u_i(t) \hat{H}_i,$$

with fidelity $F = |\langle \psi_1 | \psi(T) \rangle|^2$. Instead of discretising the full waveform, we *parametrise* every control as

$$u_i(t) = u_i^{(0)}(t) + \sum_{k=1}^K c_{ik} f_{ik}(t), \quad (2.10)$$

where

- $u_i^{(0)}(t)$ is an optional bias (e.g. our initial analytic pulse),
- $c_{ik} \in \mathbb{R}$ are the optimisation variables, and
- $f_{ik}(t) = \sin(\omega_{ik}t + \phi_{ik})$ or $\cos(\omega_{ik}t + \phi_{ik})$ with random frequencies $\omega_{ik} \in [\frac{2\pi}{T}, \Omega_{\max}]$ and phases $\phi_{ik} \in [0, 2\pi)$.

The upper frequency limit Ω_{\max} constrains the bandwidth of the control pulse. In practice, we choose Ω_{\max} based on the hardware's analogue bandwidth and the physics of the system, such that $\Omega_{\max} \lesssim \Omega_{\text{phys}}$, where Ω_{phys} is the highest frequency that the system can follow without significant distortion or loss.

OPTIMISATION IN A FIXED BASIS (CRAB CYCLE)

For a fixed random basis the fidelity becomes a smooth function $F(\mathbf{c})$. Because K is small, we efficiently locate a local maximum with derivative-free routines such as CMA-ES, Nelder–Mead, or Powell search.




Figure 2.5: Illustration of control landscapes in dCRAB. **Left:** By varying control parameters c_1 and c_2 we modify the control pulse f , which in turn affects the control objective $J(f)$. The landscape includes local maxima and constrained regions (red barrier), highlighting the complexity of navigating the *dynamic control landscape*. **Right:** The fidelity F is shown as a function of the final quantum state, represented by coordinates ψ_1 and ψ_2 . This typically simpler landscape, with a single maximum, represents the *kinematic control landscape* described by $F(|\psi(T)\rangle)$. Adapted from [42].

DRESSING THE BASIS

If the optimiser stalls, we *dress* the pulse by absorbing the current solution into the bias $u_i^{(0)}(t)$ and drawing a fresh random basis for the next cycle. We repeat this process until the incremental fidelity gain ΔF falls below a preset threshold.

THEORETICAL GUARANTEES

Under mild assumptions an infinite sequence of dressed cycles avoids strict local maxima and converges to the global optimum[41]. Each dressing augments the available search directions, so we eventually follow any remaining gradient of the true landscape.

ALGORITHMIC IMPLEMENTATION

The dCRAB algorithm proceeds in a sequence of optimisation cycles, each performed in a randomly chosen low-dimensional basis. Within each cycle, we optimise only the basis coefficients using a derivative-free method. Once the inner optimiser converges, we update the control waveform and generate a new random basis for the next cycle. This loop continues until the fidelity gain between consecutive cycles falls below a threshold or a maximum number of dressings is reached. The algorithm is outlined in Algorithm 2.

2.3. GATE SET TOMOGRAPHY

Gate Set Tomography (GST) is a self-consistent, high-precision characterisation protocol that simultaneously estimates *all* operations in a quantum processor's native gate library, including state-preparation and measurement (SPAM) processes [43]. Unlike conventional quantum-process tomography, GST remains valid even when SPAM errors are comparable to—or larger than—the gate errors under investigation. We therefore rely on GST to certify devices that aspire to fault tolerance.

Algorithm 2 dCRAB Algorithm

Require: **Input:** Initial waveforms $u_i^{(0)}(t)$; basis size K ; dressing limit L_{\max} ; inner optimiser \mathcal{O} (e.g. CMA-ES); global tolerance δ

Ensure: **Output:** Optimised controls $u_i(t)$

```

1:  $l \leftarrow 0$  ▷ dressing index
2: repeat
3:   Draw random frequencies  $\omega_{ik}^{[l]}$  and phases  $\phi_{ik}^{[l]}$ 
4:   Define basis functions  $f_{ik}^{[l]}(t)$ 
5:   Optimise coefficients:  $\mathbf{c}^{[l]} \leftarrow \mathcal{O}(F(\mathbf{c}))$ 
6:   Update waveform bias:  $u_i^{(0)}(t) \leftarrow u_i^{(0)}(t) + \sum_{k=1}^K c_{ik}^{[l]} f_{ik}^{[l]}(t)$ 
7:    $l \leftarrow l + 1$ 
8:    $\Delta F \leftarrow F_l - F_{l-1}$ 
9: until  $\Delta F < \delta$  or  $l \geq L_{\max}$ 
10: return  $u_i(t) = u_i^{(0)}(t)$ 

```

2.3.1. MATHEMATICAL MODEL

We represent a gate set as $\mathcal{G} = \{\rho, \mathbf{G}, E\}$, where ρ is the prepared initial state, E is the two-element positive operator-valued measure (POVM) describing measurement in the computational basis, and $\mathbf{G} = \{G_1, \dots, G_{|\mathbf{G}|}\}$ is the collection of trace-preserving quantum channels realised by the hardware. We parameterise each channel as a $d^2 \times d^2$ super-operator so that

$$\text{vec}(\rho_{\text{out}}) = G_i \text{vec}(\rho_{\text{in}}).$$

Because the experimental data constrain the gate set only up to a similarity transformation, we fix the gauge *a posteriori*.

2.3.2. EXPERIMENTAL PROTOCOL

We build a GST experiment from three nested building blocks [44]:

1. **Germ sequences** \mathbf{g} are short words in the gate alphabet that amplify coherent errors when repeated.
2. **Fiducials** \mathbf{f}_p (preparation) and \mathbf{f}_m (measurement) map the unknown SPAM processes into an informationally complete basis.
3. **Power iterations:** we repeat each germ g a length L chosen from a compressed exponential grid $L \in \{1, 2, 4, 8, \dots, 2^{k_{\max}}\}$.

We measure all combinations $\mathbf{f}_m g^L \mathbf{f}_p$ and record the binary outcome frequencies $\{p_s, N_s\}$ for $s \in \mathcal{S}$. Long loops ($L \sim 10^4$) convert tiny coherent errors into $\mathcal{O}(1)$ deviations, so GST routinely achieves diamond-norm accuracies below 10^{-4} —orders of magnitude beyond randomised benchmarking.

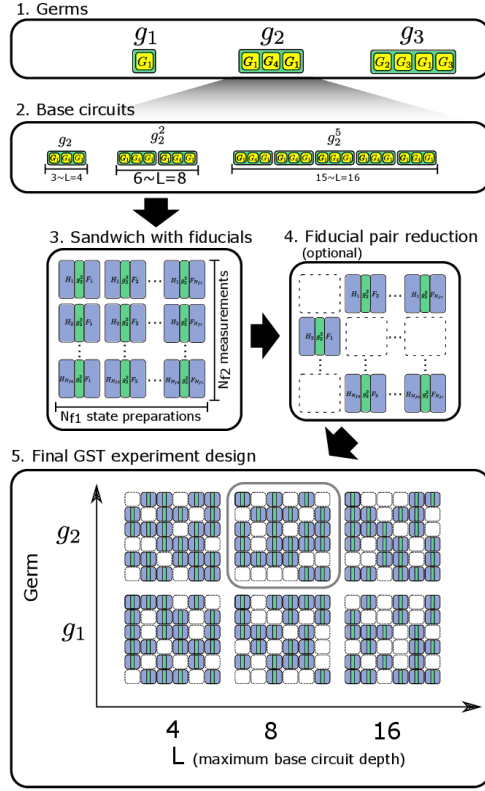


Figure 2.6: Overview of our GST experiment design. (1) We choose a set of *germs* to amplify specific gate errors. (2) We construct *base circuits* by repeating germs to varying depths L . (3) We sandwich each base circuit between sets of *fiducial operations* (state preparations and measurements) to form the full experimental circuits. (4) Optionally, we reduce the number of fiducial pairs to decrease the total number of circuits. (5) The final design is a grid of circuits, indexed by germ and depth L , each including different combinations of fiducials. Adapted from [44].

2.3.3. PARAMETER ESTIMATION

We infer the unknown parameters θ via maximum-likelihood estimation (MLE). In gate set tomography (GST), θ parametrises the quantum gates, the initial state, and the measurement effect that together define the model of the experiment. Writing the model probability for sequence s as $p_s(\theta)$, the log-likelihood is

$$\mathcal{L}(\theta) = \sum_{s \in \mathcal{S}} N_s [n_s \ln p_s(\theta) + (1 - n_s) \ln(1 - p_s(\theta))],$$

where n_s is the observed fraction of ‘1’ outcomes. Because $\mathcal{L}(\theta)$ is non-convex, our GST solvers proceed in two stages [45]:

1. We run *linear-inversion GST (LGST)* to obtain a coarse starting point under the assumption of negligible SPAM error.

2. We refine that estimate with iterative MLE, using Newton or quasi-Newton updates that respect complete positivity and trace preservation.

We certify convergence with the log-likelihood ratio $\lambda = 2[\mathcal{L}_{\max} - \mathcal{L}(\hat{\theta})]$, which follows a χ^2 distribution under Wilks' theorem.

2

2.3.4. GAUGE OPTIMISATION AND ERROR METRICS

Because the experimental data can be explained equally well by many gauge-equivalent gate sets, gauge freedom presents a challenge when interpreting estimated gate sets. In particular, most commonly reported error metrics—such as the average gate fidelity—are not gauge-invariant. To enable meaningful comparison with a target gate set $\mathcal{G}_{\text{ideal}}$, we perform a gauge optimisation: we select the gauge for our estimated gate set $\hat{\mathcal{G}}$ that minimises the Frobenius distance to $\mathcal{G}_{\text{ideal}}$,

$$\min_{T \in \text{GL}(d^2)} \sum_{G \in \hat{\mathcal{G}}} \|T^{-1} \hat{G} T - G_{\text{ideal}}\|_F^2.$$

This optimisation does not resolve the fundamental ambiguity introduced by gauge freedom, but it provides a consistent convention for reporting results. After gauge optimisation, we report metrics such as the diamond-norm error, process infidelity, a stochastic-vs-coherent Pauli decomposition, and the unitarity $u(\hat{G})$, which witnesses decoherence.

2.3.5. USED METRICS

The metrics used in this work quantify gate errors in different operational and statistical regimes, allowing for both high-level and fine-grained benchmarking. We focus on the following diagnostics:

1. N -sigma deviation. The N -sigma (or log-likelihood ratio) test quantifies the statistical plausibility of the estimated gate set under the assumption that the model is correct. Specifically, it measures the discrepancy between the observed data and the predicted probabilities from the best-fit model. The test statistic

$$\lambda = 2[\mathcal{L}_{\max} - \mathcal{L}(\hat{\theta})]$$

follows a chi-squared distribution with k degrees of freedom under the null hypothesis, where $k = |\mathcal{S}| - \dim(\theta)$. The N -sigma score is defined as

$$N\sigma = \frac{\lambda - k}{\sqrt{2k}},$$

which approximates a standard normal deviate and provides an intuitive measure of model adequacy: values of $N\sigma \lesssim 2$ indicate a good fit [43], [45].

2. Average gate fidelity. The average fidelity $F_{\text{avg}}(\hat{G}, G_{\text{ideal}})$ quantifies the overlap between an implemented gate \hat{G} and the ideal gate G_{ideal} under uniform input states:

$$F_{\text{avg}}(\hat{G}, G_{\text{ideal}}) = \int d\psi \langle \psi | G_{\text{ideal}}^\dagger \circ \hat{G}(|\psi\rangle \langle \psi|) |\psi\rangle.$$

This integral simplifies to a closed-form expression involving the Choi matrices of the two channels [46]. The fidelity is closely related to the entanglement fidelity and provides an interpretable single-number summary of performance.

2.3.6. ERROR GENERATORS





Sector	Dimension	Action	Example effect (Bloch sphere)	ε_J (Error Probability)	θ_J (Error Amplitude)
Hamiltonian \mathbb{H}	$d^2 - 1$	$H_P[\rho] = -i[P, \rho]$		0	1
Stochastic (Pauli) \mathbb{S}	$d^2 - 1$	$S_P[\rho] = P\rho P - \mathbb{1}\rho\mathbb{1}$		1	0
Stochastic (Pauli-correlation) \mathbb{C}	$\begin{pmatrix} d^2 - 1 \\ 2 \end{pmatrix}$	$C_{P,Q}[\rho] = P\rho Q + Q\rho P - \frac{1}{2}\{[P, Q], \rho\}$		0	0
					1
Active \mathbb{A}	$\begin{pmatrix} d^2 - 1 \\ 2 \end{pmatrix}$	$A_{P,Q}[\rho] = i\left(P\rho Q - Q\rho P + \frac{1}{2}\{[P, Q], \rho\}\right)$		0	1
					0

Figure 2.7: We represent an imperfect gate by its error generator, defined as $\mathcal{L} = \log(GG^{-1})$, where G is the process matrix of the imperfect gate and G is that of the ideal gate. We construct a basis of elementary error generators spanning the vector space \mathcal{L} , which provides a taxonomy of small Markovian errors. This basis partitions the generators into four sectors (subspaces), listed in this table along with their dimensions and Choi sum representations. For each sector, we consider the single-qubit case and illustrate how an error process generated by a single basis element deforms the Bloch sphere. Adapted from [47].

The error generator \mathcal{L}_G is a linear map that captures deviations from the target gate via

$$\hat{G} = G_{\text{ideal}} e^{\mathcal{L}_G},$$

analogous to the Lindblad generator in open-system dynamics. Expanding \mathcal{L}_G in a Pauli basis allows one to classify errors as stochastic, coherent, or non-Markovian.

A systematic basis for \mathcal{L}_G spans four main types of elementary error generators, each corresponding to a different physical effect:

- **Hamiltonian (H) generators** produce unitary errors. These represent over- or under-rotations about a particular axis of the Bloch sphere and correspond to mis-calibrated control fields.
- **Stochastic (S) generators** represent decoherence processes such as dephasing or amplitude damping. They reduce Bloch vector lengths without introducing rotation, modelling purely incoherent noise.
- **Commutator-Pauli (CP) generators** and **Anticommutator-Pauli (AP) generators** describe non-unitary error processes in the Markovian Lindblad formalism. CP

generators correspond to reversible but non-unitary dynamics, such as those induced by coherent coupling to a fluctuating environment. AP generators model state-dependent decoherence effects, for example, relaxation processes whose rates depend on the system's instantaneous state or the applied control.

Figure 2.7 shows how each type of generator transforms the Bloch sphere, helping to visualise their distinct operational effects. This decomposition enables a more granular understanding of gate errors and supports targeted error mitigation strategies.

3

METHODS

In this chapter, we outline the methodologies employed in conducting the research. We begin by introducing the software tools utilized throughout the study. These include QuTiP for open-loop control and physical simulations, QuOCS for closed-loop optimization, and pyGSTi for gate set tomography.

Next, we detail the experimental hardware. We present the optical systems used for initialisation and readout. Finally, we present the microwave and radio frequency setup used for precise system control.

SOFTWARE

The implementation and validation of quantum control protocols in this thesis required a diverse set of simulation and optimisation tools. These tools enabled the design of high-fidelity control pulses and the quantitative assessment of experimental results. Python-based open-source frameworks were preferred for their flexibility, community support, and extensibility, allowing for the development of custom modules tailored to the specific requirements of NV centre experiments.

This section presents an overview of the primary software packages utilised:

- **QuTiP** [48] and its extension `qutip-qoc` [49] for simulating quantum dynamics and performing quantum optimal control,
- **QuOCS** [50] for flexible, experiment-oriented control pulse optimization, and
- **pyGSTi** [51] for rigorous characterisation of gate performance through gate set tomography.

The interplay between these software packages forms the computational backbone of this work, facilitating both theoretical simulations and the experimental realisation of optimised quantum operations.

3.1. QuTiP

The Quantum Toolbox in Python (QuTiP) is a widely used open-source Python package for simulating quantum systems [52]. QuTiP provides solvers for both the Schrödinger and master equations via the `qutip.sesolve` and `qutip.mesolve` functions, respectively. Its modular and extensible architecture makes it particularly suitable for custom quantum simulation tasks, including the integration of user-defined modules.

In the context of this thesis, several custom helper functions were implemented in QuTiP to aid in the visualisation and interpretation of optimal control results. These functions streamline the analysis of control pulse performance and system dynamics. Brief descriptions and full function signatures are provided in Appendix B.1.

Additionally, some simulations leveraged a specialized Python package called NiVaSiT, developed by the Taminiau group. NiVaSiT builds on top of QuTiP and is tailored to simulations of NV centres in diamond.

3.1.1. QUTIP-QOC

`qutip-qoc` is an extension of QuTiP that provides functionality for quantum optimal control [48]. Originally embedded within the QuTiP codebase, `qutip-qoc` was separated into a standalone package in 2024. It supports both gradient-based and gradient-free optimisation algorithms, including GRAPE, GOAT, CRAB, and reinforcement learning (RL)-based methods.

From version 5 onward, `qutip-qoc` introduces a new API centered around the `Objective` class, which enables multi-objective optimisation.

An example usage of how to run a `qutip-qoc` optimisation is shown below:

```

1 H = [Hsys, tensor(sigmax(), identity(2)), tensor(sigmay(), identity(2))
      ]
2 objective = qoc.Objective(initial=qt.qeye(4), H=H, target=cnot)
3 qoc.optimize_pulses(
4     objectives=objective,
5     control_parameters={
6         "ctrl_x": {"guess": x, "bounds": [None, None]},
7         "ctrl_y": {"guess": y, "bounds": [None, None]},
8     },
9     tlist=tlist,
10    algorithm_kwargs={
11        "fid_err_targ": 1e-20,
12        "min_grad": 1e-20,
13        "alg": "GRAPE",
14        "max_iters": 1000,
15        "dyn_type": "UNIT",
16    },
17    minimizer_kwargs={
18        "method": "l-bfgs-b",
19    }
20 )

```

Listing 3.1: Example of using the `Objective` class and `optimize_pulses` in `qutip-qoc` v5+

Here, the first entry in `H` is the system (drift) Hamiltonian, while the remaining entries define the control Hamiltonians. The `Objective` class also requires an initial unitary,

which is set to the identity in this case, and a target unitary, specified here as a CNOT gate. The `optimize_pulses` function accepts the following key arguments:

- `control_parameters`: A dictionary specifying the control fields. Each entry includes an initial guess for the control amplitudes ("guess") and optional upper and lower bounds for the optimisation ("bounds").
- `tlist`: A list of time points defining the time grid over which the control fields are applied. The number of time bins is implicitly set by the length of `tlist`.
- `algorithm_kwargs`: A dictionary of parameters passed to the underlying optimisation algorithm:
 - `fid_err_targ`: Target fidelity error; the optimization terminates if this threshold is reached.
 - `min_grad`: Minimum allowed gradient magnitude before the optimisation is considered converged.
 - `alg`: The optimisation algorithm to use; in this case, GRAPE.
 - `max_iters`: The maximum number of iterations for the algorithm.
 - `dyn_type`: The type of system dynamics; "UNIT" indicates unitary evolution.
- `minimizer_kwargs`: Parameters passed to the classical minimiser used within the algorithm. Here, the Nelder-Mead simplex method is selected via the "method" key.

The chosen parameters were selected after extensive experimentation. In particular, changing `dyn_type` from the default "GEN_MAT" to "UNIT" proved to be a crucial adjustment, as the optimisation consistently failed to produce meaningful results with "GEN_MAT". This may be because "GEN_MAT" internally expects the Hamiltonian and control operators to be defined in a form compatible with Liouville-space evolution, which is primarily intended for open quantum systems. For purely unitary target operations, this representation introduces unnecessary complexity and can potentially lead to mismatches or numerical instabilities if the problem is defined using a Hilbert-space Hamiltonian, as done here. By restricting the simulation to unitary evolution with "UNIT", the optimiser is aligned with the problem definition, ensuring correct handling of closed-system dynamics and significantly reducing computational overhead.

Furthermore, switching from "l-bfgs-b" to the full BFGS method was empirically observed to accelerate convergence considerably. However, this came at the cost of the optimiser consistently predicting lower final fidelities than those obtained with L-BFGS-B in this specific application. This discrepancy is likely due to BFGS being more prone to converging to sharp local minima in the high-dimensional control landscape, whereas L-BFGS-B, with its limited memory and bound constraints, performs more conservative updates that tend to avoid such traps. While BFGS can be advantageous for quickly exploring the parameter space, L-BFGS-B proved more reliable for achieving higher-fidelity solutions in practice within this work.

3.1.2. DYNAMICAL DECOUPLING

In solid-state quantum systems such as the NV centre, environmental noise leads to rapid decoherence, limiting coherent control times. DD sequences are a well-established method to mitigate these effects by periodically refocusing phase errors through carefully timed π pulses [53].




Figure 3.1: Schematic of the echo sequence used to embed the optimised control pulse. Free evolution periods τ and τ' surround the control pulse of duration τ_{OCP} , with π pulses applied at intermediate points to refocus dephasing errors. In this implementation, $\text{gate_duration} = 1.6 \mu\text{s}$ and $\tau = 12.96 \mu\text{s}$.

Since the electron spin is decoupled throughout the entire experiment, the gates are defined to incorporate the π pulses on the electron within their structure. This ensures correct timing and compatibility with other gates in the sequence. Alternatively, one could implement the weak π pulse independently on the AWG, which would be valid but would complicate integration with other gates, as the timing between decoupling pulses would need to be precisely managed. Each echo in the decoupling sequence includes a period of free evolution (τ) between π pulses, as shown in 3.1.

When incorporating an optimal control pulse (OCP) of duration τ_{OCP} , the timing relationship becomes

$$2\tau' + \tau_{\text{OCP}} = 2\tau.$$

If the intended target gate is a CNOT operation, i.e., $U_{\text{target}} = U_{\text{CNOT}}$, we must account for the surrounding DD dynamics. The first and second echoes are defined as

$$U_{\text{echo1}} = U_{\tau} U_{\pi} U_{\tau'}, \quad U_{\text{echo2}} = U_{\tau'} U_{\pi} U_{\tau}.$$

The composite evolution including the OCP becomes

$$U_{\text{DD, OCP}} = U_{\text{echo1}} U_{\text{OCP}} U_{\text{echo2}}.$$

To ensure that the embedded pulse implements the desired CNOT gate in the DD context, the adjusted target unitary is given by

$$U_{\text{target}} = U_{\text{echo1}}^{\dagger} U_{\text{CNOT}} U_{\text{echo2}}^{\dagger}.$$

3.1.3. VIRTUAL PHASE CORRECTIONS

The weak π gate optimised by GRAPE is a controlled rotation about the X -axis (CROTX), not a canonical CNOT. The Hamiltonian used in our optimisation supports only such CROT-type gates. As a result, the implemented unitary differs from the ideal CNOT by a

known, fixed phase on the control qubit: applying a CROT gate corresponds to a CNOT gate preceded or followed by a $R_z(-\pi/2)$ rotation on the control qubit.

In addition to this structural difference, the experimental implementation requires correcting for dynamical phase accumulation during the gate. Specifically, the electron spin acquires a phase due to coherent evolution under the hyperfine coupling during the pulse. To maintain a consistent rotating frame for the electron spin, this phase must be tracked and compensated by adjusting the phase of subsequent control operations (typically by updating the phase of the local oscillator).

We therefore construct a logically equivalent target unitary that incorporates both corrections:

$$U_{\text{target}} = \left(R_z^{\text{control}}(-\pi/2) \otimes \mathbb{I} \right) \cdot \left(\mathbb{I} \otimes R_z^{\text{target}}(\theta_{\text{hf}})^\dagger \right) \cdot U_{\text{ideal}}, \quad (3.1)$$

where:

- $R_z^{\text{control}}(-\pi/2)$ accounts for the fixed structural phase difference between CROT and CNOT,
- $R_z^{\text{target}}(\theta_{\text{hf}})$ compensates the hyperfine-induced phase accumulation on the electron spin,
- θ_{hf} depends on the total pulse duration.

By applying both corrections—the fixed $R_z^{\text{control}}(-\pi/2)$ and the dynamical $R_z^{\text{target}}(\theta_{\text{hf}})$ —we obtain a target unitary that is logically equivalent to an ideal CNOT in the chosen rotating frame.

3.2. QuOCS

The Quantum Optimal Control Suite (QuOCS) [50] is an open-source Python package designed to facilitate both model-based and experiment-based implementations of quantum optimal control algorithms. It provides a modular framework for deploying a wide range of optimisation techniques, including GRAPE, AD-GRAPE, dCRAB, and various direct search methods.

QuOCS is well-suited for integration with experimental platforms, offering support for Qudi, a modular python suite for experiment control and data processing. Its architecture is built around extensibility and real-time interaction with laboratory equipment, making it a valuable tool for closed-loop optimisation workflows.

In this work, however, all experimental interfacing and control software were developed on top of diamondOS, a custom software package developed by QuTech.

An overview of the QuOCS framework is provided in Figure 3.2.

3.2.1. OPTIMISATION DICTIONARY

The main user interface for defining an experiment is the *optimisation dictionary*. This dictionary contains all relevant configuration details for the optimisation process, except for the figure of merit, which is discussed in the next section. An example of what the optimisation dictionary looks like for a typical experiment can be seen in B.2.

The dictionary first specifies `algorithm_settings`, which define the optimal control algorithm to be used (e.g., dCRAB, GRAPE, Reinforcement Learning, Krotov) and the

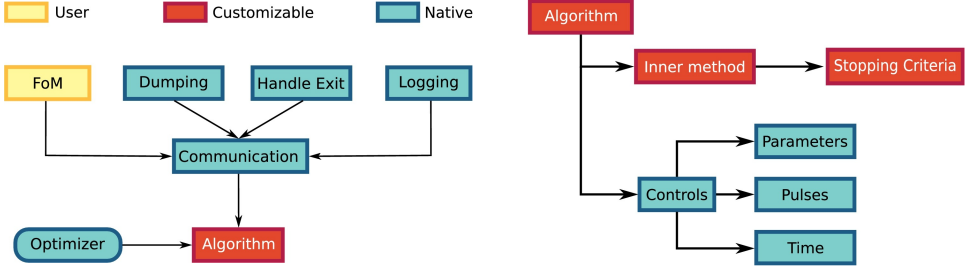


Figure 3.2: Overview of the QuOCS framework. **Left:** The user-defined FoM is integrated into a Communication object, which also connects to native utilities for data dumping, logging, and exit handling. This communication object is then passed to a customizable Algorithm, which is initiated by the Optimiser to begin the optimisation process. **Right:** Structure of the QuOCS optimisation algorithm. The core logic resides in the Inner method, which is customizable and can be replaced with user-defined code. It interfaces with a Stopping Criteria component to manage convergence or timeout conditions. The algorithm also interacts with a Controls object, which manages Parameters, Pulses, and Time—elements critical for defining the optimisation landscape. Figure adapted from [50].

optimisation direction (maximisation or minimisation). It also sets the maximum number of super-iterations and function evaluations allowed.

The direct search method (`dsm_settings`) determines the numerical optimisation algorithm (e.g., Nelder-Mead, Powell, CMA-ES, L-BFGS-B), along with the corresponding stopping criteria based on parameter changes and function evaluations.

Finally, pulse settings are defined. An initial guess can be provided, which in our experiments corresponds to the reference weak π -pulse. This initial guess is specified as an array of piecewise constant values, effectively added to the pulses produced during optimisation. Mathematically, if the pulse created by the optimisation is \mathbf{p}_{opt} and the initial guess is \mathbf{p}_{init} , the final pulse applied is

$$\mathbf{p} = \mathbf{p}_{\text{init}} + \mathbf{p}_{\text{opt}}.$$

Additionally, the pulse basis (e.g., Fourier, Chebyshev, Sinc, Sigmoid) and the number of basis vectors are specified. Time arrays required for pulse definitions are also included.

3.2.2. CONSTRUCTION OF THE FINAL PULSE WITH dCRAB

The final control pulse in dCRAB is built by successively augmenting a base waveform with oscillatory components drawn from randomly sampled bases, as described in Section 2.2.2. Each such augmentation corresponds to a single *super-iteration*, where QuOCS generates a new random basis and optimises the associated coefficients.

Specifically, in super-iteration j , a new set of angular frequencies $\{\omega_{j,1}, \omega_{j,2}, \omega_{j,3}\}$ is drawn, and the optimiser searches for the best amplitudes $\{a_{j,k}, b_{j,k}\}$ for each component. The corresponding update to the pulse is:

$$\Delta p_j(t) = \sum_{k=1}^3 [a_{j,k} \sin(\omega_{j,k} \tau) + b_{j,k} \cos(\omega_{j,k} \tau)],$$

where $\tau = 2\pi t/T$ is a normalised time variable. This representation, equivalent to the basis given in Eq. (2.10), ensures full expressivity via both sine and cosine terms, enabling arbitrary phase shifts at each frequency.

The frequencies $\omega_{j,k}$ in each super-iteration are independently drawn from a uniform distribution over a predefined range [54]. This stochastic sampling strategy ensures that the optimiser explores a wide range of temporal features in the control landscape. By varying the basis at each super-iteration, dCRAB avoids the risk of being confined to a fixed functional subspace, enhancing its ability to escape local optima.

The cumulative pulse after N super-iterations is then:

$$p_{\text{final}}(t) = p_{\text{init}}(t) + \sum_{j=1}^N \Delta p_j(t).$$

Each super-iteration expands the accessible search directions by redefining the basis, thereby avoiding premature convergence to suboptimal solutions. This *basis dressing* mechanism provides a meta-level exploration on top of the local optimisation within a fixed basis, as illustrated in Fig. 2.5.

For reproducibility, the final waveform can be reconstructed from a stored frequency and coefficient lists:

```
1 tau = 2 * np.pi * time / time[-1]
2 pulse = initial_guess.copy()
3 for freqs, params in zip(dcrab_freq_list, dcrab_param_list):
4     for k, w in enumerate(freqs):
5         pulse += params[2*k] * np.sin(w * tau) + params[2*k+1] * np.cos
           (w * tau)
```

Listing 3.2: Reconstruction of the Final Pulse

The table below presents the exact frequencies and amplitudes used by QuOCS in each super-iteration during a dCRAB optimisation of the weak π gate. These values were extracted directly from `optimized_parameters.json` and fully determine the updates applied to the pulse in each super-iteration.

Super-Iter	ω_1	ω_2	ω_3	$a_1 \times 10^3$	$a_2 \times 10^3$	$a_3 \times 10^3$	$b_1 \times 10^3$	$b_2 \times 10^3$	$b_3 \times 10^3$
0	2.0424	6.2940	7.2780	-0.3	0.8	0.4	0.5	0.4	1.0
1	2.9443	4.9170	7.7148	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0618	4.2538	8.1862	0.1	-0.1	0.2	0.2	0.4	1.3
3	1.3234	6.6217	9.9939	-1.3	1.7	-0.4	0.3	1.5	-0.0
4	1.8966	5.3724	9.3921	0.0	0.0	0.0	0.0	0.0	0.0

Super-iterations 1 and 4 resulted in no effective parameter updates, suggesting that the direct search method identified no significant improvement over the current solution when using those particular randomised bases.

This behaviour is characteristic of the stochastic nature of dCRAB, where some super-iterations may fail to yield beneficial directions in the control landscape. However, the alternation of exploration (via random bases) and refinement (via local optimiser) across super-iterations is what enables dCRAB to robustly escape local traps and achieve high-quality control fields.

The corresponding pulse constructed from these parameters is shown in Figure 3.3, confirming the exact reconstruction of the final control field as stored by QuOCS.

More details can be found in the [QuOCS GitHub documentation](#).

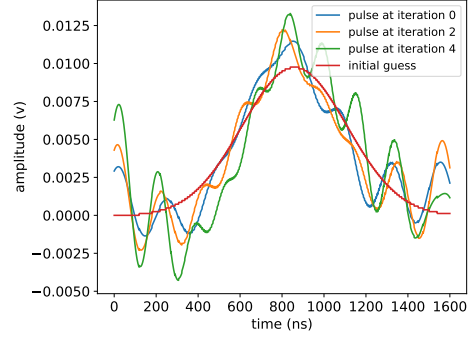


Figure 3.3: Reconstruction of the optimised pulse from the QuOCS parameters listed above. The final pulse (green) matches exactly the stored control field, verifying the additive pulse construction method used by QuOCS.

3.2.3. FIGURE OF MERIT

As shown in Figure 3.2, the component of the QuOCS framework that must be provided by the user is the figure of merit (FoM). The FoM serves as the evaluation metric for the optimal control algorithm, determining how well a given pulse performs. Consequently, the outcome of the optimisation process depends entirely on how the FoM is defined, since the algorithm will attempt to either maximise or minimise it according to the settings specified in the optimisation dictionary. For our research, two FoMs were defined: one based on photoluminescence (PL) contrast, and another incorporating Bell state fidelity.

PL CONTRAST

PL contrast serves as a simple yet informative FoM for evaluating the performance of the optimised weak π gate. The core idea is that a properly functioning CNOT gate will flip the electron spin, conditional on the state of the nitrogen nuclear spin.

The experimental procedure involves two measurements: in the first, the nitrogen spin is initialised to the $|0\rangle$ state; in the second, it is initialised to the $|1\rangle$ state by applying a $R_x(\pi)$ gate. After applying the CNOT operation, the population of the electron spin in the $|0\rangle$ state is measured for both cases. Ideally, the CNOT gate flips the electron spin only when the nitrogen spin is in $|1\rangle$. Hence, the difference in electron $|0\rangle$ population between the two cases quantifies how well the gate performs its conditional logic.

We define the PL contrast as:

$$\text{PL Contrast} = P_{e=0|N=0} - P_{e=0|N=1},$$

where $P_{e=0|N=i}$ is the probability of measuring the electron spin in $|0\rangle$ given that the nitrogen spin was initialized in state $|i\rangle$. A higher contrast indicates a better-performing CNOT gate.

The following algorithm outlines the experimental sequence used to compute the PL contrast:

Algorithm 3 PL Contrast Measurement

```

1: for  $n\_init$  in  $[0, 1]$  do
2:    $init()$   $\triangleright \psi = |00\rangle$ 
3:   if  $n\_init == 1$  then  $\triangleright \psi = |01\rangle$ 
4:      $x180(nitrogen)$ 
5:   end if
6:    $apply\_oc\_weak\_pi(nitrogen, electron)$   $\triangleright$  Apply optimized CNOT
7:    $measure\_electron\_population()$ 
8:    $P[n\_init] \leftarrow$  measured population of electron in  $|0\rangle$ 
9: end for
10:  $PL\ Contrast \leftarrow P[0] - P[1]$ 

```

BELL STATE FIDELITY

The CNOT gate is intended to be an entangling operation, but to fully assess its performance, it is not sufficient to merely verify that it produces some entangled state. A stronger and more informative benchmark is its ability to generate a specific maximally entangled state. In this work, we select the Bell state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ as the target, which allows for an evaluation of both amplitude and phase coherence in the two-qubit system.

To this end, we define a FoM based on the parity of the generated state in three complementary Pauli bases: ZZ , XX , and YY . This approach is motivated by the symmetry of the Bell state, which exhibits strong correlations (or anti-correlations) in each of these bases. Measuring in multiple bases enables the detection of phase errors and leakage that would not be apparent from measurements in the computational (Z) basis alone. Notably, while the earlier contrast measurement in the Z basis verifies population transfer, it is insensitive to relative phases between basis states—phases which are crucial for certifying the creation of coherent superposition states like $|\Phi^+\rangle$.

In each basis, the parity of the measured state is estimated as:

$$P = 2 \cdot (\text{prob}_{00} + \text{prob}_{11}) - 1,$$

where prob_{00} and prob_{11} denote the probabilities of observing the even-parity outcomes in the chosen basis. High parity values indicate strong coherence and correct correlations consistent with the target state.

The final FoM aggregates the parity values from all three bases into a single scalar quantity:

$$\text{FoM} = \text{Bell state fidelity} = \frac{1 + P_{XX} - P_{YY} + P_{ZZ}}{4}.$$

This expression estimates the fidelity with respect to the ideal $|\Phi^+\rangle$ state, assuming ideal measurement and state preparation. The sign conventions are chosen to reflect the expected parities of $|\Phi^+\rangle$: even in Z and X , and odd in Y .

The algorithm below outlines the procedure used to compute this FoM:

Algorithm 4 Bell State Fidelity FoM calculation

```

1: measurements  $\leftarrow [x, y, z]$ 
2: correlations  $\leftarrow \{\}$ 
3: for measurement in measurements do
4:   init()  $\triangleright \psi = |00\rangle$ 
5:   x90(nitrogen)  $\triangleright \psi = |0\rangle|-i\rangle$ 
6:   apply_oc_weak_pi(nitrogen, electron)  $\triangleright \psi = |00\rangle + |11\rangle$ 
7:   if measurement == x then  $\triangleright$  Project on X axis
8:     y90(electron)
9:     y90(nitrogen)
10:  else if measurement == y then  $\triangleright$  Project on Y axis
11:    x90m(electron)
12:    x90m(nitrogen)
13:  end if
14:  correlations[measurement]  $\leftarrow 2 \cdot (\text{prob\_00} + \text{prob\_11}) - 1$ 
15: end for
16: FoM  $\leftarrow \frac{1 + \text{correlations}[x] - \text{correlations}[y] + \text{correlations}[z]}{4}$ 

```

3.3. PYGSTi

pyGSTi is a Python-based implementation of Gate Set Tomography (GST), developed and maintained by Sandia National Laboratories [51]. It provides a comprehensive end-to-end framework for GST, encompassing the construction of target models, experimental design generation, and analysis of experimental data. A typical pyGSTi workflow is illustrated in Figure 3.4.




Figure 3.4: Example of a typical pyGSTi workflow.

The primary objects used to interact with pyGSTi are `Circuit`, `Model`, and `DataSet`:

- **Circuit:** Defines quantum circuits composed of one or more qubits. These objects specify the sequences of operations (circuits) to be executed during the experiment.
- **Model:** Encodes the target gate set of the system under study, including available quantum gates and any state preparation and measurement (SPAM) operations.
- **DataSet:** A dictionary-like container that stores experimental data, enabling feedback of collected data into pyGSTi for further analysis.

pyGSTi offers a range of functionalities, including circuit simulation, data simulation, model testing, model-free characterization, and, most importantly for this work, model estimation. In addition to GST, it provides ready-to-use implementations for randomized benchmarking (RB) and robust phase estimation (RPE).

An example of code implementing PyGSTi can be seen in B.3.

After incorporating the experimental data and performing model fitting, pyGSTi generates a comprehensive HTML report.

In this thesis, we primarily analyse the following quantities:

- **Non-Markovianity indicators** which highlight systematic deviations beyond Markovian assumptions;
- **Average gate infidelity**, to assess the overall quality of implemented gate operations;
- **Error generators**, which give a Lie-algebraic description of coherent and stochastic errors in each gate.

3.4. EXPERIMENTAL SETUP

High-fidelity control of solid-state qubits demands a robust hardware platform that synchronises optical, microwave, and RF components with sub-microsecond precision. We use a previously built cryogenic system integrated with laser, microwave, and timing electronics. This section outlines the experimental setup.

3.4.1. OPTICAL SYSTEM

All experiments are performed on a diamond sample mounted in a Montana cryogenics chamber at 4 K. A static magnetic field of 2518 G is aligned along the NV axis of the sample. We use three distinct lasers which are employed for charge-state and spin control:

- A 515 nm (green) laser is used for charge-resonance reset, ensuring the NV centre is initialised into the negative charge state (NV^-).
- A 637 nm (red) laser is used for optical readout; this laser is resonant with the zero-phonon line transition.
- A second 637 nm (red) laser is used for spin initialisation, preparing the electron spin state in $|0\rangle$ at the start of each experimental sequence.

Acousto-optic modulators (AOMs) act as fast optical switches, providing high extinction ratios (> 100 dB) to prevent laser leakage during experiments.

Photon detection and optical readout are performed using an avalanche photodiode (APD). An ADwin real-time control system orchestrates the experimental timing, ensuring precise synchronisation of control and readout signals. The control signal architecture is detailed further in the following section.

3.4.2. SIGNAL CHAIN

In this section, we discuss the signal chain—that is, the path traversed by the control pulses from their generation in the arbitrary waveform generator (AWG) to their delivery at the sample housed in the cryogenic environment. An overview of the signal chain is provided in Fig. 3.5.

The control pulses are resonant with the electron spin transition at approximately 4.18 GHz. To generate a clean carrier at this frequency, single-sideband (SSB) modulation [55] is employed. The local oscillator (LO) of the signal generator is set to $f_{LO} = 4.43$ GHz, i.e. 250 MHz above the qubit frequency. The AWG produces an intermediate-frequency (IF) waveform at $f_{IF} = 250$ MHz.

This IF signal is mixed with the LO using an IQ mixer. The output contains components at $f_{LO} \pm f_{IF}$ (i.e. 4.68 GHz and 4.18 GHz). By adjusting the relative phase of the in-phase (I) and quadrature (Q) inputs, the upper sideband at 4.68 GHz is suppressed, leaving only the desired lower sideband at 4.18 GHz. This SSB technique also filters low-frequency noise originating from the AWG.

Mathematically, the baseband signal entering the mixer is

$$s_{bb}(t) = I(t) \cos(2\pi f_{IF} t) + Q(t) \sin(2\pi f_{IF} t).$$

For a single-tone pulse of constant amplitude A we choose

$$I(t) = A \cos(2\pi f_{IF} t), \quad Q(t) = A \sin(2\pi f_{IF} t).$$

Substituting into the mixer relations yields

$$s(t) = A \cos[2\pi(f_{LO} - f_{IF})t],$$

demonstrating that only the lower sideband at $f_{LO} - f_{IF} = 4.18$ GHz survives, while the upper sideband is cancelled by destructive interference. This implementation corresponds to the Hartley (or Weaver) method of SSB modulation [56].

The resulting microwave signal is amplified and gated by a switch controlled by the ADwin real-time controller. The switch is open only during pulse application, minimizing leakage and background noise.

For the nitrogen nuclear spin, control pulses at 2.18 MHz are generated directly by the AWG, filtered by a ferrite coil to suppress high-frequency electronic noise, and then amplified to ensure sufficient driving strength at the sample.

The 2.18 MHz nitrogen drive and the 4.18 GHz electron drive are combined using a triplexer (the third port is unused) and delivered to the sample through a DC block. A gold stripline patterned on the diamond surface ensures efficient coupling to the qubits.

Signal distortion can degrade gate fidelity, but in our setup—with microsecond-scale pulses—distortion is modest compared with nanosecond-scale control typical of superconducting qubits [57].

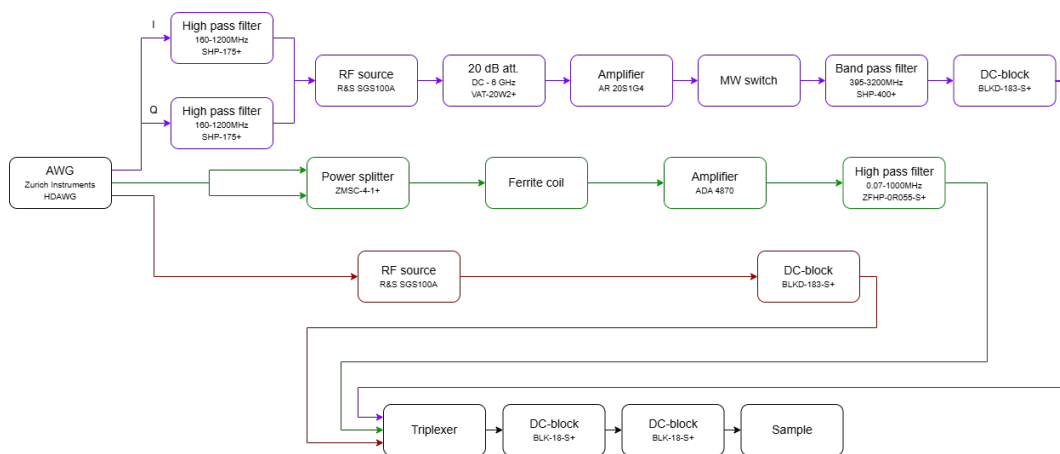


Figure 3.5: Diagram of the microwave and radio-frequency delivery system. For the electron spin (purple), single-sideband modulated pulses at 250 MHz are generated by the AWG and upconverted with an LO at 4.18 GHz using an external RF source. A 20 dB attenuator protects the amplifier, and a microwave (MW) switch suppresses leakage when pulses are inactive. Band-pass and high-pass filters remove out-of-band noise. For the nitrogen nuclear spin (green), the AWG synthesises baseband pulses which are filtered, amplified, and passed through a ferrite coil to reduce high-frequency noise. A second RF source (red) is used for the electron $|+1\rangle$ transition, which is currently not used. All paths are combined via a triplexer and routed through DC blocks to the sample.

4

RESULTS AND DISCUSSION

This chapter presents the results of our control pulse optimisation experiments and their performance evaluations on a two-qubit NV centre system. We begin by exploring open-loop optimal control methods based on idealised system models, using GRAPE to produce high-fidelity pulses under controlled assumptions. We then demonstrate how these pulses, despite their strong simulated performance, degrade in real experimental settings due to model inaccuracies and unaccounted-for noise. This motivates the shift to closed-loop optimisation, where the physical system itself guides the search for optimal control, enabling greater robustness and adaptability to experimental imperfections.

We then apply GST to characterise the implemented gates, identify systematic errors, and quantify performance. These diagnostics reveal key insights into the nature and structure of control errors. Finally, we design and implement a novel gate sequence based on the information learned from GST analysis.

4.1. OPEN-LOOP OPTIMAL CONTROL

Open-loop optimal control seeks to determine the time-dependent control fields that best realise a desired quantum operation, based on a mathematical model of the system dynamics. In this approach, pulses are optimised by simulating the time evolution of the system under a given Hamiltonian and adjusting control parameters to maximise a chosen figure of merit—typically, the fidelity with respect to a target gate or state.

Among the most widely used algorithms in this domain is GRAPE, which leverages gradient-based optimisation under the assumption of a known and differentiable Hamiltonian. This makes it especially well-suited for simulating idealised dynamics where the model closely matches reality.

However, as we demonstrate in later sections, the performance of open-loop pulses is ultimately limited by the accuracy and tractability of the system model. Real-world quantum systems, such as NV centres in diamond, often exhibit non-idealities—such as crosstalk, decoherence, and unknown couplings—that can be challenging to capture

fully. Even when detailed information is available, as in the case of known couplings to many nearby ^{13}C nuclei, it may be computationally infeasible to incorporate this complexity into a model suitable for fast evaluation and optimisation. Consequently, pulses optimised in simulation do not always translate into high-fidelity gates in experiment.

Despite these challenges, open-loop control remains a valuable tool for exploring pulse design under controlled assumptions. It provides a foundation for understanding the system's ideal behaviour and for benchmarking the performance of more robust, data-driven closed-loop techniques introduced later.

4.1.1. PRODUCED PULSES

We produced optimised pulses under three different modelling assumptions. In all cases, the target operation was a CNOT gate with the nitrogen nuclear spin as the control and the electron spin as the target.




Figure 4.1: Optimised control pulses for the CNOT gate under three modelling assumptions. Each panel shows the Rabi frequency of the microwave drive as a function of time. **Left:** *Vanilla* model using a simplified two-qubit Hamiltonian without decoherence or leakage. **Middle:** *Qutrit* model including the full three-level Hilbert space of the nitrogen nuclear spin and penalising leakage from the computational subspace. **Right:** *No Echo* model where the echo pulses are omitted from the target unitary, serving as a coherence control test.

1) Vanilla. This baseline pulse was optimised using the simplified Hamiltonian $H = A_{\parallel} S_z^{(e)} S_z^{(n)}$ with no explicit modelling of decoherence or leakage. It serves as a reference to evaluate the effects of more realistic modelling choices introduced in subsequent scenarios.

2) No Echo. This pulse omits the echo pulses in the target unitary described in Section 3.1.2. It functions as a sanity check to verify that the echo layer is needed and improves coherence. The experiments are performed with echo pulses present, as otherwise the electron decoheres quickly.

3) Nitrogen as a Qutrit. Finally, we modelled the nitrogen nuclear spin as a full qutrit, accounting for the entire $m_I = \{-1, 0, +1\}$ subspace. The Hamiltonian becomes

$$H = A_{\parallel} S_z^{(e)} I_z^{(n)},$$

where $I_z^{(n)}$ is the spin-1 operator for the nitrogen nucleus. The control optimisation explicitly penalised population transfer into the $m_I = -1$ level, effectively constraining the dynamics to the computational subspace $\{m_I = 0, +1\}$. This approach ensures that the resulting gates are robust not only to decoherence but also to leakage.

All four pulses achieved extremely high expected fidelities in simulation, all of them exceeding 99%.

4.1.2. PHASE EVOLUTION AND CORRECTION MECHANISMS

During open-loop optimal control, the evolution of the nitrogen nuclear spin under the hyperfine coupling A_{\parallel} leads to phase accumulation that must be managed carefully. If left uncorrected, this phase can degrade the fidelity of the implemented gate, even when the pulse correctly performs the desired logical operation. In this section, we show two complementary strategies for addressing this effect: tuning the gate duration and applying an explicit phase correction to the target unitary.

GATE DURATION SWEEPS

An important free variable in open-loop optimal control is the gate duration. We therefore investigated the relationship between gate duration and the resulting expected fidelity, as shown in Figure 4.2.

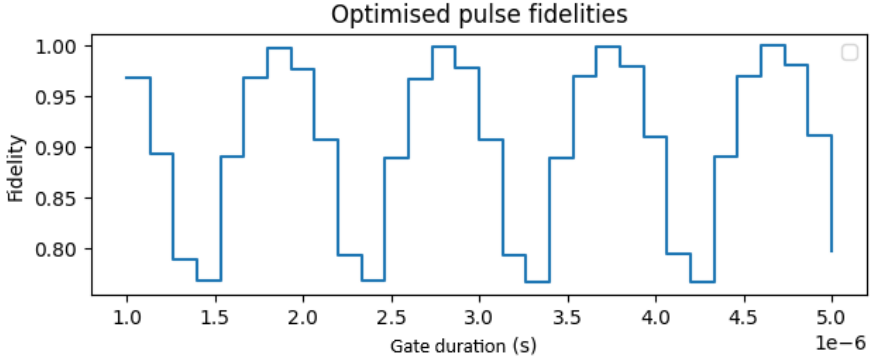


Figure 4.2: Gate duration sweep of GRAPE pulses. The fidelities oscillate as a function of A_{\parallel} , consistent with the phase evolution under the simplified Hamiltonian. The horizontal plateaus arise from the finite step size in gate duration; 60 data points were used in total.

We find that the fidelity exhibits oscillations as a function of A_{\parallel} . This behaviour is expected, as the simplified Hamiltonian used in this optimisation is given by

$$H = A_{\parallel} S_z^{(e)} S_z^{(n)},$$

where $S_z^{(e)}$ and $S_z^{(n)}$ are the Z -operators for the electron and nitrogen nuclear spins, respectively.

The term $A_{\parallel} t$ in the time evolution induces a phase accumulation on the nitrogen spin that depends on the gate duration t . Constructive interference of this phase occurs at times satisfying

$$t = \frac{1}{2A_{\parallel}k},$$

with $k \in \mathbb{N}$.

Given that $A_{\parallel} = 2.18$ MHz, an oscillation period corresponding to A_{\parallel} is expected. An optimal gate duration can therefore be calculated as

$$\text{gate_duration} = \frac{1}{2A_{\parallel}k}.$$

By selecting gate durations that satisfy this condition, the nitrogen spin acquires a phase compatible with the target unitary, thereby improving fidelity.

4

EXPLICIT PHASE CORRECTION

Alternatively, the same phase evolution can be corrected directly by adjusting the phase of the target unitary, rather than tuning the gate duration.

As explained in Section 3.1.3, a phase offset may need to be corrected for when defining the target unitary, especially when echo sequences are used. If this phase is not included in the definition of the target unitary, the overlap between the simulated and target unitaries will appear reduced, leading to an artificially low fidelity even if the implemented operation is otherwise correct.




Figure 4.3: Fidelity as a function of global phase offset applied to the target unitary. A fixed pulse is used for all evaluations. The observed periodicity confirms that the implemented operation differs from the ideal gate only by a global phase, which must be accounted for in fidelity evaluation.

In order to characterise the fidelity landscape with respect to phase correction, we perform a phase sweep using a fixed pulse, as shown in Figure 4.3. The sweep is conducted at a gate duration of $2.5 \mu\text{s}$ with an echo included in the target unitary. The utilised pulse sequence can be seen in Figure 4.4.




Figure 4.4: Pulse sequence used to perform the phase sweep shown in Figure 4.3. The sequence includes a DD layer that introduces a deterministic phase offset in the nitrogen spin, necessitating correction in the target unitary.

4

We observe a clear periodic dependence of fidelity on the applied phase offset, with maxima near $\pi/2$ and $3\pi/2$, and pronounced dips at 0 , π , and 2π . This confirms that the pulse implements the correct logical operation up to a global phase, which must be incorporated into the target unitary to avoid artificially reduced fidelity values.

Using the method described in Section 3.1.3, we determine the optimal phase offset to be

$$\phi_0 \approx 1.08\pi (\approx 194.4^\circ).$$

This value aligns well with the observed fidelity peaks, validating the correction procedure.




Figure 4.5: Effect of phase correction on fidelity as a function of gate duration. **Left:** Without phase correction, fidelity plateaus near 0.77 due to a fixed mismatch between the implemented and target unitaries. **Right:** Applying the optimal phase correction ($\phi_0 \approx 1.08\pi$) restores fidelity to the expected high values near 0.99, confirming that the discrepancy was due solely to a global phase offset.

To further validate the importance of phase correction, we repeat the gate duration sweep using both a fixed phase of 0 and the optimal phase (Figure 4.5). In the left plot, the fidelity remains flat near 0.77, corresponding to the initial mismatch discussed in Section 3.1.3. This can also be observed in Figure 4.2, where at a gate duration of $2.5 \mu\text{s}$ we see a fidelity of approximately 0.77.

In contrast, the right plot shows fidelity stabilising near 0.99, confirming that the optimiser consistently aligns the pulse to a unitary differing only by a global phase. This further demonstrates that uncorrected phases can obscure high-fidelity solutions and lead to misleading optimisation outcomes.

In summary, the oscillations observed in the gate duration sweep and the phase correction explored here stem from the same physical effect: phase accumulation of the nitrogen spin due to its precession under A_{\parallel} . This phase can be managed either by se-

lecting specific gate durations that align the phase appropriately, or by explicitly correcting the phase in the target unitary. Both approaches are complementary and together ensure accurate fidelity evaluation and convergence to physically valid solutions.

Therefore, we explicitly incorporate the echo-induced phase shift into the target unitary for all subsequent optimisations. This ensures meaningful fidelity evaluation and prevents high-fidelity solutions from being obscured by trivial phase mismatches.

4.1.3. AMPLITUDE SWEEPS

After generating the optimised pulses, we implemented them on our experimental setup and applied them to a physical sample containing the NV centre system. To facilitate this, a `NumericalMWPulse` object was defined in *diamondOS*, a custom software package developed by QuTech to streamline experimental procedures. Since the open-loop optimisation returns Rabi frequencies rather than control voltages, we must calibrate the corresponding amplitude scaling factors required for effective hardware implementation. As an initial step, we calculated the Rabi frequency to voltage conversion factor based on the parameters of a square weak π pulse. This provided a theoretical expectation for the correct control amplitude, which led us to apply an initial scaling factor of 1.0 when implementing the optimised pulse. To refine this calibration, we perform amplitude sweeps and monitor the PL contrast to identify the optimal operating point. High PL contrast indicates improved gate fidelity, as discussed in Section 3.2.3.

Figure C.1 shows the results of the amplitude sweeps. Figure 4.6 shows the full amplitude sweep, around multiplication factor 1 and around multiplication factor 19.8 for the qutrit pulse.

The sweep range was normalised by the maximum amplitude present in each pulse. As shown in Figure 4.1, the vanilla pulse features significantly higher peak amplitudes—nearly an order of magnitude larger—than the others.




Figure 4.6: Amplitude sweeps for the GRAPE-optimised qutrit pulse. **Left:** Full sweep over a broad range of scaling factors. **Middle:** Zoomed-in view near scaling factor 1.0, corresponding to the theoretically predicted operating point. **Right:** Zoomed-in view near scaling factor 19.8, which also exhibits high contrast. These two operating points correspond to implementations of the target gate and its logical inverse, respectively.

The vanilla pulse exhibits no clear region of high contrast, suggesting poor control fidelity. The no-echo pulse yields low contrast across the sweep range, with both nitrogen initialisation states having similar lines. This indicates that without echo compensation, the pulse fails to induce meaningful conditional dynamics. In contrast, the qutrit pulse displays two prominent contrast peaks, around scaling factors of 1 and 19.8, suggesting effective conditional gate operation at these amplitudes.

Since the qutrit pulse yielded the highest contrast in the initial amplitude sweeps, it was selected for all subsequent experiments.

As shown in Figure 4.6, both experimental and simulated sweeps reveal two distinct sweet spots near scaling factors of 1 and 19.8. At these points, we observe maximal PL contrast between the readouts conditioned on the nitrogen spin state, indicating strong conditional control fidelity. Notably, although the contrast magnitude is similar at both points, the roles of the $|0\rangle$ and $|1\rangle$ traces are reversed.

This inversion arises from the pulse effectively implementing opposite unitary operations at the two amplitude settings. At a scaling factor of 1, the pulse performs the intended target gate. Near 19.8, the accumulated phase and rotation angle result in an over-rotation that implements the logical inverse of the intended gate. For instance, a CNOT gate targeting the electron spin conditioned on $|1\rangle_N$ becomes an inverse CNOT, flipping the electron spin when the nitrogen spin is in $|0\rangle_N$.




Figure 4.7: Comparison of simulated (left) and experimental (right) amplitude sweeps for the qutrit-optimised pulse. Each trace corresponds to a different initialisation of the nitrogen nuclear spin: blue for $|1\rangle_N$ and orange for $|0\rangle_N$. Vertical axis: final electron spin population. Horizontal axis: amplitude scaling factor. Both experimental and simulated results display distinct sweet spots near scaling factors of 1 and 19.8, confirming conditional control fidelity at these amplitudes.

The existence of sweet spots around an amplification factor of 1 and 19.8 is further validated by simulation, as shown in 4.7.

4.1.4. PL CONTRAST AND BELL STATE FIDELITY

To quantitatively assess pulse performance with minimal experimental overhead, we evaluate the PL contrast and Bell state fidelity figures of merit introduced in subsection 3.2.3. Figure 4.8 compares these metrics for the reference weak π pulse and the GRAPE-optimised pulse, evaluated at the two sweet spots identified in the earlier amplitude sweep—namely, amplitudes 1 and 19.8. The reference pulse, described in subsection 2.1.3, serves as our baseline.

The reference weak π pulse achieves a PL contrast of $97.36\% \pm 0.12\%$, whereas the GRAPE pulses yield $91.11\% \pm 0.33\%$ (amplitude 1) and $89.55\% \pm 0.35\%$ (amplitude 19.8). This reduction is expected, as open-loop optimisation is based on an idealised model that omits experimental imperfections such as decoherence, drift, and pulse distortions.

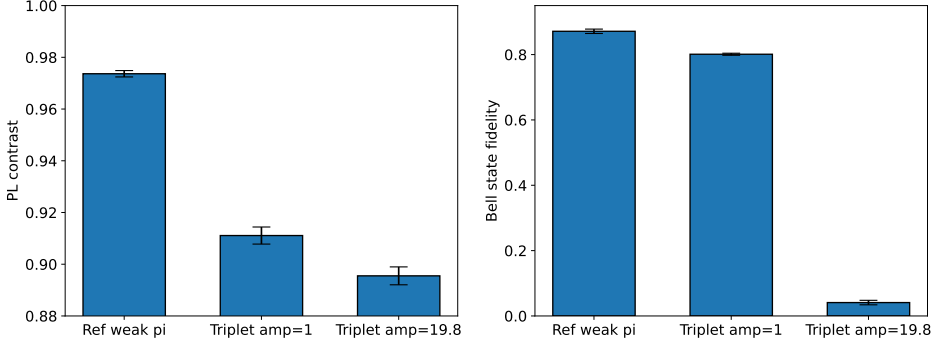


Figure 4.8: Comparison of PL contrast and Bell state fidelity for the reference weak π pulse and the GRAPE-optimised qutrit pulse at amplitude scaling factors 1 and 19.8. For the reference pulse, fidelity is computed with respect to the Bell state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. For the GRAPE pulses, fidelity is computed with respect to $|\Psi^-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$, which is the expected outcome given the inverted control logic. The reference pulse outperforms the GRAPE pulses on both metrics, illustrating the limitations of open-loop control in experimental settings.

To interpret the Bell state fidelity data, we consider the full entanglement circuit. A $R_y(\pi/2)$ gate is first applied to the nitrogen nuclear spin, followed by a conditional weak π pulse on the electron. The reference pulse implements:

$$U_{W-\pi} = \mathbb{1}_e \otimes |0\rangle\langle 0|_N + X_e \otimes |1\rangle\langle 1|_N,$$

where X_e denotes the Pauli X on the electron spin. Acting on the initial state $|0\rangle_e |0\rangle_N$, the sequence yields:

$$\begin{aligned} \text{Initial: } & |0\rangle_e |0\rangle_N \\ & \xrightarrow{Y(\pi/2)_N} |0\rangle_e \frac{1}{\sqrt{2}}(|0\rangle_N + |1\rangle_N) \\ & \xrightarrow{U_{W-\pi}} \frac{1}{\sqrt{2}}(|0\rangle_e |0\rangle_N + |1\rangle_e |1\rangle_N) = |\Phi^+\rangle. \end{aligned}$$

Thus, under ideal conditions, the reference pulse prepares the Bell state $|\Phi^+\rangle$.

The GRAPE pulse, by contrast, was optimised to flip the electron spin if the nitrogen is in $|0\rangle_N$, implementing the opposite control logic:

$$U_{\text{GRAPE}} = X_e \otimes |0\rangle\langle 0|_N + \mathbb{1}_e \otimes |1\rangle\langle 1|_N.$$

Applying this to the same input state gives:

$$\begin{aligned} \text{Initial: } & |0\rangle_e |0\rangle_N \\ & \xrightarrow{Y(\pi/2)_N} |0\rangle_e \frac{1}{\sqrt{2}}(|0\rangle_N + |1\rangle_N) \\ & \xrightarrow{U_{\text{GRAPE}}} \frac{1}{\sqrt{2}}(|1\rangle_e |0\rangle_N + |0\rangle_e |1\rangle_N) = |\Psi^-\rangle. \end{aligned}$$

This is the singlet Bell state $|\Psi^-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$, where the first qubit is the electron.

This prediction is confirmed by correlator measurements shown in Figure 4.9. The reference pulse produces:

$$\langle XX \rangle = 0.847 \pm 0.018, \quad \langle YY \rangle = -0.794 \pm 0.006, \quad \langle ZZ \rangle = 0.844 \pm 0.018,$$

in good agreement with the expected signature of $|\Phi^+\rangle$. In contrast, the GRAPE pulse at amplitude 1 yields:

$$\langle XX \rangle = -0.671 \pm 0.008, \quad \langle YY \rangle = -0.749 \pm 0.007, \quad \langle ZZ \rangle = -0.779 \pm 0.006,$$

consistent with $|\Psi^-\rangle$, whose negative correlations along all three axes are unique among Bell states.




Figure 4.9: Two-qubit correlator measurements ($\langle XX \rangle$, $\langle YY \rangle$, $\langle ZZ \rangle$) for the reference and GRAPE-optimised qutrit pulses. **Left:** Reference weak π pulse yields strong positive correlations, consistent with the target Bell state $|\Phi^+\rangle$. **Middle:** GRAPE pulse at amplitude 1 exhibits strong negative correlations along all axes, characteristic of the singlet state $|\Psi^-\rangle$. **Right:** GRAPE pulse at amplitude 19.8 shows degraded correlations, reflecting reduced fidelity.

The fidelity with respect to the singlet state is:

$$F_{\Psi^-} = \frac{1 - \langle XX \rangle - \langle YY \rangle - \langle ZZ \rangle}{4},$$

yielding $F_{\Psi^-} = 0.801 \pm 0.004$ for the GRAPE pulse at amplitude 1.

At amplitude 19.8, GRAPE pulse performance degrades on both metrics. The state trajectory is effectively longer (approaching ten full rotations), amplifying sensitivity to imperfections. Susceptibility to hyperfine transitions detuned by 2.2 MHz and 4.4 MHz becomes significant at high power, introducing decoherence and state mixing.

These results underscore the limitations of model-based open-loop control under realistic conditions. This motivates a shift to closed-loop control strategies, which incorporate direct hardware feedback to improve robustness and target fidelity.

4.2. CLOSED LOOP OPTIMAL CONTROL

While our GRAPE-optimised pulse performed well within the limits of the simulation, its effectiveness in experiment is fundamentally constrained by the accuracy of the underlying model—which remains relatively simple. To address this limitation, we transitioned to a closed-loop control approach, allowing the system itself to guide the pulse optimisation process in real time.

To this end, we implemented the dCRAB algorithm using the QuOCS library. We used the reference weak π pulse as an initial guess.

Throughout the optimisation runs, we used the PL contrast and Bell state fidelity metrics introduced in 3.2.3 to guide the search for high-performance pulses.

4.2.1. PRODUCED PULSES

Throughout the experimental work in this thesis, control pulses were generated by optimising the PL contrast and Bell state fidelity FoMs, as defined in Section 3.2.3. In addition to full IQ pulses, which include both in-phase (I) and quadrature (Q) components, we also optimised pulses using only the I component. This choice served as an example of a constrained search space, selected to test whether comparable gate performance could be achieved with fewer control parameters, potentially enabling faster convergence and improved robustness. Other possible ways to constrain the search space include reducing the number of basis functions used to parameterise the pulse, limiting the maximum pulse amplitude, or restricting the allowed pulse variation range during optimisation.

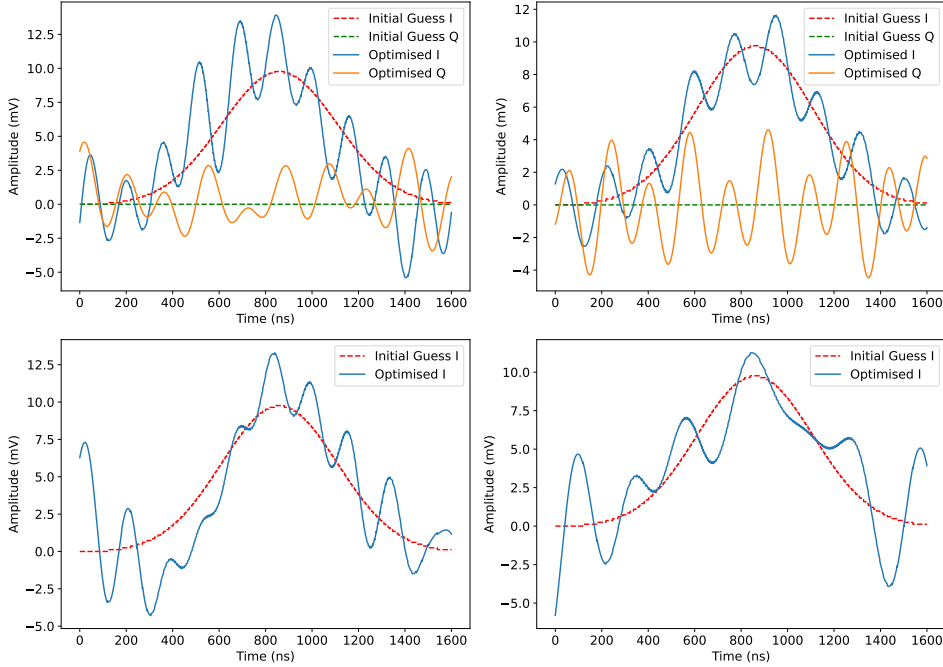


Figure 4.10: Final GRAPE-optimised control pulses for each figure of merit and control parameterisation. **Upper left:** PL contrast optimisation using full IQ control. **Upper right:** Bell state fidelity optimisation using full IQ control. **Lower left:** PL contrast optimisation using I -only control. **Lower right:** Bell state fidelity optimisation using I -only control. Solid lines denote the optimised I and Q components (or I only); dashed lines indicate the initial guesses. These pulses represent the final outcomes of the closed-loop optimisation procedure described in Section 3.2.3.

Pulse Name	Optimisation Objective	Components Used
PL	PL contrast	I and Q
Bell	Bell state fidelity	I and Q
PL only I	PL contrast	I only
Bell only I	Bell state fidelity	I only

Table 4.1: Summary of the optimised pulses, including their objective functions and control parameterisations. Each pulse is labelled according to the figure of merit used during optimisation and whether both in-phase (I) and quadrature (Q) components were employed.

4

Figure 4.2 shows the evolution of the FoM values during optimisation. Characteristic sharp dips are observed, which align with the start of new super-iterations—these mark global reshuffling events where the optimiser samples new directions to escape local optima.

Interestingly, for both the PL contrast and Bell state fidelity objectives, the highest-performing solutions were not found at the end of the optimisation but rather at earlier super-iterations. This may imply that early explorations identified favourable pulse structures that were later perturbed by overfitting to local landscape noise or hardware drift.

Figure 4.10 displays the final control pulses, including both I and Q components.

4.2.2. PL CONTRAST AND BELL STATE FIDELITY

To accurately assess the performance of our control pulses, we evaluate them using the FoMs defined earlier, which also serve as cost functions in our closed-loop optimisation routine. All measurements were performed with 5000 repetitions (shots) per basis setting. To ensure a fair comparison, the reference weak π -pulse data were recollected alongside the other pulses, thereby controlling for environmental variations such as magnetic field fluctuations.

As shown in Figure 4.11, the reference weak π -pulse achieves the highest PL contrast. Interestingly, the PL-optimised pulses yield slightly lower contrast, which is counterintuitive since they were specifically optimised for this metric. One possible explanation is that magnetic field drift or other slow variations occurred during the course of the optimisation process, causing a shift in system parameters that degraded performance relative to the reference.

Turning to the Bell state fidelity metric, we observe a more level distribution of values across different pulses, though the uncertainties are larger. This increased uncertainty is primarily due to how the fidelity is computed as detailed in section 4.1.

To further dissect the Bell state fidelity results, Figure 4.12 presents the individual two-qubit correlators used in its calculation. The correlators show consistent structure across all pulses, with XX and ZZ generally yielding higher absolute values than YY . This may reflect phases being accrued in certain bases. The uncertainty in each correlator is substantial, as discussed in 4.1, which propagates into the total Bell fidelity metric.

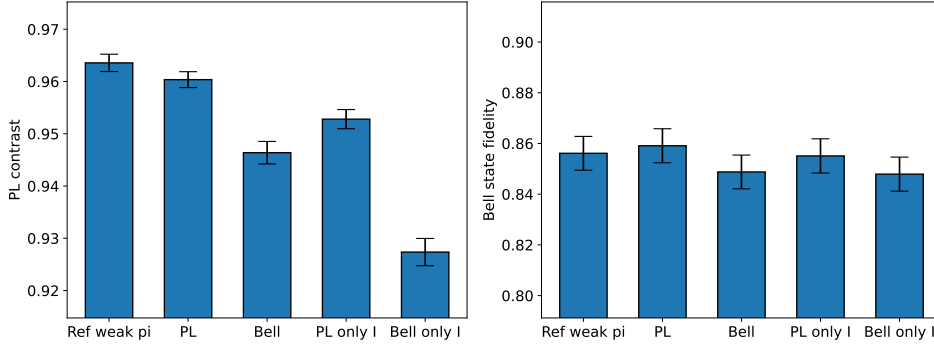


Figure 4.11: Comparison of all optimised pulses and the reference weak π pulse using the two primary figures of merit: PL contrast and Bell state fidelity. All data were collected under consistent experimental conditions with 5000 repetitions per basis setting. The reference pulse achieves the highest PL contrast, while Bell state fidelities are more uniformly distributed across pulses with higher uncertainties.




Figure 4.12: Measured two-qubit correlators ($\langle XX \rangle$, $\langle YY \rangle$, $\langle ZZ \rangle$) for each pulse, used in computing Bell state fidelities. The correlator structure is broadly consistent across all pulses, though variation in magnitude and uncertainty—particularly in $\langle YY \rangle$ —affects the overall fidelity calculation. All correlators were estimated using 5000 shots per basis setting.

4.2.3. NOISE LIMITATIONS

A valuable dataset for characterising the system's intrinsic noise was obtained during runs involving only I pulses. Due to a bug in the sequencer code, the initial guess pulse was replayed at every iteration rather than being updated. Although unintended, this behaviour provided a unique opportunity to systematically probe the system's noise characteristics under repeated application of a fixed control pulse.

Each function evaluation involved 5000 measurement shots. Assuming binary outcomes and a typical population contrast of $p = 0.96$, as discussed in Section 2.1.1, the expected standard deviation from quantum projection noise is

$$\sigma_{\text{stat}} = \sqrt{\frac{p(1-p)}{N}} = \sqrt{\frac{0.96 \cdot 0.04}{5000}} \approx 0.00277.$$

This defines a theoretical noise floor for comparison with observed fluctuations.

The recorded PL contrast data exhibit both random variation from shot noise and structured drifts, particularly during the early hours of the experiment. Across the entire dataset, the mean standard deviation in PL contrast is $\bar{\sigma}_{\text{PL}} = 0.00843 \pm 0.00028$, approximately three times the projection noise limit. Excluding the first 8 hours—where drift was most pronounced—this reduces to $\bar{\sigma}_{\text{PL}} = 0.00545 \pm 0.00021$, indicating improved system stability thereafter.

Magnetic field calibrations were performed hourly, and objective lens repositioning scans were carried out every 4 hours. However, during the initial phase of data collection, neither the magnet nor the lens had been realigned for some time. This misalignment led to pronounced drift in the PL signal and an exceptionally large repositioning step in the first lens scan. The magnetic field calibrations during this period show strong temporal correlation with abrupt corrections in the PL signal, marking a transition to a more stable regime for the remainder of the experiment.

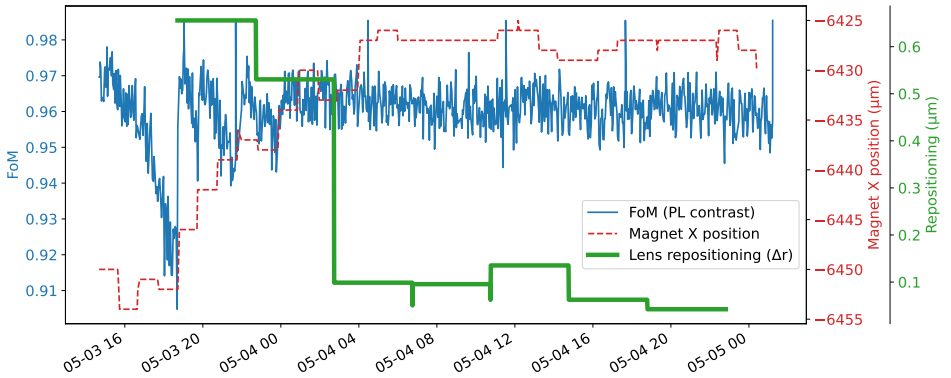


Figure 4.13: Correlation between PL contrast (blue), magnet position (red dashed line), and objective lens repositioning magnitude (green) over time. The PL contrast exhibits both gradual drifts and sudden jumps, many of which align with manual adjustments to the magnet or optical system. The largest contrast change coincides with the first objective lens realignment, indicating that optical focus has a dominant effect on signal quality. Smaller contrast variations are also associated with magnetic field recalibrations, suggesting that magnetic drift is a non-negligible noise source.

While the goal was to maintain a stable magnetic field and optical alignment, this required periodic manual adjustments. The absolute magnet position is not directly meaningful, but sudden changes often coincide with jumps in the PL signal, indicating magnetic field drift as a significant noise source.

The lens repositioning trace reflects regular calibration events, performed every 4 hours via full scans along the X, Y, and Z axes. At each of these time points, the net mechanical displacement is computed as the Euclidean distance between the current and previous scan centres:

$$\Delta r = \sqrt{(x_i - x_{i-1})^2 + (y_i - y_{i-1})^2 + (z_i - z_{i-1})^2},$$

where x_i, y_i, z_i denote the scan centre coordinates at time t_i . The largest displacement occurs after the initial two alignments, while subsequent refocusing steps are smaller

but still non-negligible.

Notably, the first objective lens repositioning produces an immediate and pronounced jump in the FoM, from approximately 0.90 back to 0.97. In contrast, magnetic field calibrations tend to induce smaller, gradual corrections. This highlights the dominant impact of optical alignment on readout quality, particularly when significantly out of focus.

4.3. GATE SET TOMOGRAPHY

GST is a comprehensive protocol for characterising quantum gates, state preparations, and measurements in a self-consistent and device-agnostic manner. Unlike traditional quantum process tomography, which assumes perfect SPAM, GST reconstructs the entire gate set from experimentally accessible data, making no such assumptions. This allows it to isolate and quantify errors not just in individual gates, but in the full logical gate set implemented on a quantum device.

In this section, we apply GST to the produced pulses in order to evaluate the performance of our optimised gate, identify the dominant sources of error, and determine the degree to which the system dynamics can be described by a time-independent, Markovian model. We begin by assessing global deviations from this model using a non-Markovianity metric, then quantify per-gate performance via average infidelities and error generators.

4.3.1. NON-MARKOVIANITY

The non-Markovianity metric, denoted N_σ , quantifies how strongly the experimental data deviates from the assumptions of a Markovian quantum process—specifically, that each gate acts independently and identically at every invocation. This metric is computed from the log-likelihood ratio between the best-fit gateset model and the raw data, scaled in units of standard deviation. In essence, it measures how well the stationary CPTP (completely positive trace-preserving) model assumed by GST captures the observed behaviour.

An explicit error bar is not reported for N_σ , it expresses statistical deviation in units of standard deviation under the assumption that the fitted model is correct. As such, it already incorporates expected sampling fluctuations, and its interpretation is similar to that of a goodness-of-fit statistic. Values of $N_\sigma \lesssim 1$ are consistent with the model within statistical noise, whereas $N_\sigma \gg 1$ indicates an inconsistency between model and data, suggesting the presence of non-Markovian effects [58].

The way it is constructed, N_σ aggregates contributions across all circuits and is sensitive to non-Markovian effects such as temporal correlations, drift, or gate-context dependence that violate the assumptions of time-independent noise. Therefore, it is an excellent metric for diagnosing whether the overall GST model is fundamentally inadequate to describe the true experimental system dynamics.

Looking at Figure 4.14, we observe that non-Markovianity is present in all experiments and follows similar patterns across the different base lengths. For $L = 1$, N_σ is around 45 for all pulses. For $L = 2$, it drops slightly to values between 38 and 42. For $L = 4$, a similar range of 39 to 43 is observed. Finally, for $L = 8$, N_σ increases again to values between 44 and 53.

It is important to interpret these trends carefully. N_σ does not directly quantify the

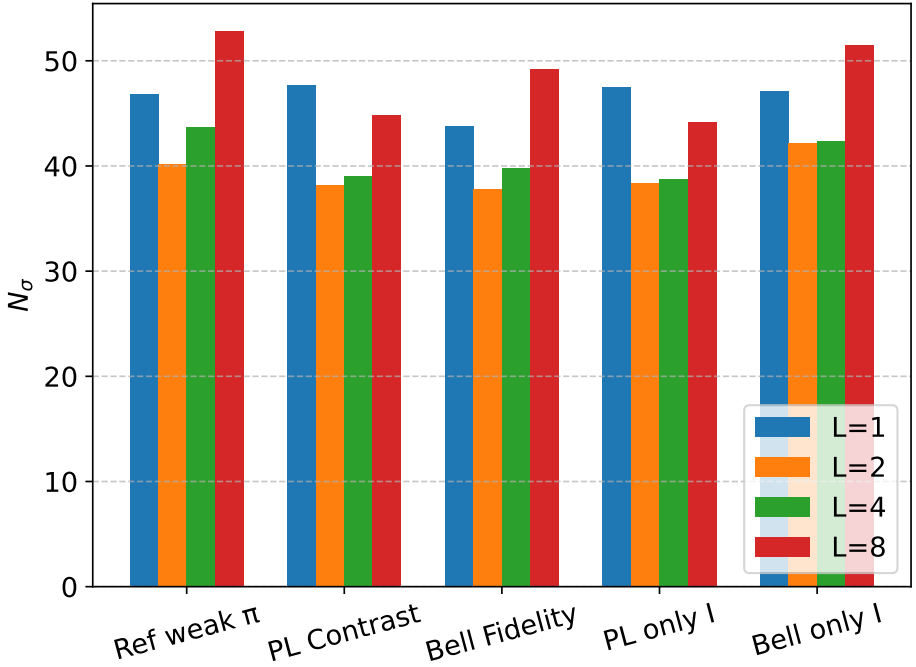


Figure 4.14: Non-Markovianity metric N_σ for five experiments, shown across base lengths $L = 1, 2, 4$, and 8 . Higher values of N_σ indicate stronger statistical evidence against the Markovian noise model assumed by GST. While N_σ increases slightly with base length—partly due to the cumulative nature of deviations over deeper circuits—all experiments exhibit significant non-Markovianity. No single pulse deviates markedly from the others, suggesting that non-Markovian effects are broadly distributed across the system.

amount of non-Markovianity present in the system, nor does a larger value imply a stronger non-Markovian effect per se. Rather, it quantifies the statistical confidence with which the data rejects the Markovian model [43], and it grows roughly linearly with the total number of circuit repetitions. Thus, changes in N_σ with base length L partly reflect the fact that deeper circuits comprise more repeated operations, leading to larger aggregate deviations even if the underlying noise processes remain unchanged. In this light, the observed variations across L should not be overinterpreted as evidence for qualitatively different non-Markovian behaviour at different depths.

A more robust observation is that no single optimised control pulse stands out as an outlier in N_σ across the different base lengths. This suggests that none of the pulses introduces disproportionately strong violations of the Markovian model, and that non-Markovian effects are instead broadly distributed across the system and circuits.

Finally, it must be emphasised that N_σ is a global metric: it aggregates evidence over the entire GST dataset and does not indicate which specific gates or circuits contribute most to the observed deviations. More targeted diagnostics would be required to localise the sources of non-Markovianity.

4.3.2. AVERAGE GATE FIDELITY

To evaluate the performance of individual gates, we turn to the average gate infidelity estimates provided by GST. These values are derived from the reconstructed process matrices and offer insight into how closely each implemented gate approximates its ideal counterpart.

As the CNOT is the gate that we optimised, we visualise the CNOT gate results as a bar chart in Figure 4.15.




Figure 4.15: Average gate infidelity for the CNOT gate across five GST reconstructions, each corresponding to a different control pulse. The reference weak π pulse achieves the lowest infidelity, indicating robust performance under the GST model. Optimised pulses yield slightly higher infidelities, ranging from 3.64% to 4.21%, but remain within a comparable performance range. Error bars reflect the statistical uncertainty in the GST estimate.

We notice that the reference weak π pulse achieves the lowest average gate infidelity of $3.36 \pm 0.14\%$, suggesting that this pulse remains well-suited for implementing the CNOT gate. The other pulses exhibit slightly higher but comparable infidelities, ranging from $3.64 \pm 0.14\%$ to $4.21 \pm 0.14\%$.

To better understand the source of these errors, we now turn to one of the most informative metrics provided by GST: the error generator.

4.3.3. ERROR GENERATOR

The error generator is one of the most informative diagnostics provided by GST, offering a detailed decomposition of gate errors into coherent and stochastic contributions. It is constructed by comparing the estimated superoperator to its ideal form within a Lie-algebraic framework, where small deviations are expressed in the Pauli basis as elements of the corresponding error algebra [47].

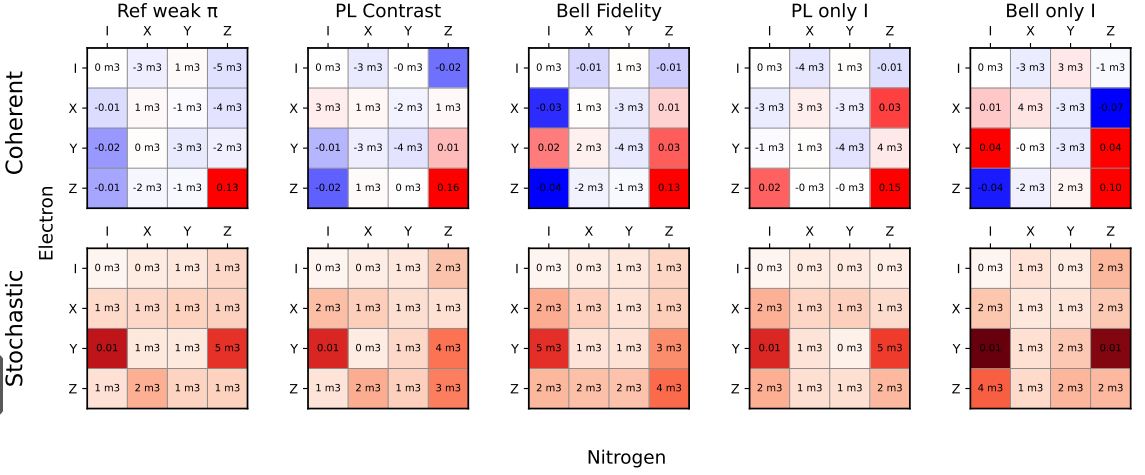


Figure 4.16: Coherent (upper) and stochastic (lower) error generators for the CNOT gate under each optimised pulse, expressed in the Pauli basis. Coherent errors reflect systematic, unitary deviations from the ideal gate and contribute quadratically to gate infidelity, whereas stochastic errors represent decoherent noise processes and contribute linearly. All pulses exhibit a dominant ZZ coherent error component, indicating residual two-qubit coupling as the primary source of infidelity. Stochastic contributions are comparatively small but show consistent structure along YI and YZ terms, possibly reflecting system-specific dephasing channels.

In Figure 4.16, we present the coherent and stochastic error generators for the optimised pulses. It is important to note that coherent errors contribute quadratically to gate infidelity, whereas stochastic errors contribute linearly.

Coherent error generators. We observe a clear and consistent pattern across all configurations: a large ZZ component dominates the coherent error generator, with values ranging from 0.10 to 0.16. Additional observations for each pulse are as follows:

- **Ref weak π pulse:** Moderate single-qubit error terms are present, with -0.01 XI, -0.02 YI, and -0.01 ZI.
- **PL pulse:** Similar structure is observed with 0.02 IZ, -0.01 YI, and -0.02 ZI.
- **Bell pulse:** This pulse exhibits larger single-qubit coherent terms: -0.03 XI, 0.02 YI, and -0.04 ZI. Interestingly, despite these larger terms, the Bell pulse achieves a lower average gate infidelity than the PL pulse, which may be attributed to its larger ZZ term (-0.16 compared to -0.13).
- **PL only I pulse:** This pulse exhibits significant -0.03 XZ and 0.02 ZI terms, as well as a large ZZ component (0.15).
- **Bell only I pulse:** While this pulse achieves the lowest ZZ term (0.10), it suffers from large other coherent terms: -0.07 XZ, 0.04 YZ, 0.04 YI, and -0.04 ZI. The

presence of multiple large coherent components in this case leads to the worst overall gate infidelity among the pulses studied.

Stochastic error generators. For the stochastic error generators, we observe small contributions across all configurations, consistent with the low stochastic noise level of the system. Nevertheless, a pattern of elevated YI and YZ components is apparent. This may indicate a susceptibility to stochastic noise along the Y axis. The fact that these stochastic terms remain small compared to the coherent errors suggests that further improvements in gate fidelity should prioritise addressing the coherent contributions, particularly the ZZ interaction.

A clear pattern emerges: all optimised pulses exhibit a large ZZ coherent error component, which dominates the overall error budget. Single-qubit coherent terms vary between pulses, and their impact on gate fidelity depends sensitively on their magnitude and combination with ZZ errors. Stochastic errors remain minor, with some consistent structure in Y -axis terms. These results underscore the importance of targeting ZZ error suppression.

4

ZZ TERM DUE TO AC-STARK SHIFT

The persistence of ZZ indicates that it originates from a mechanism invariant under a global reference-frame update—specifically, an AC-Stark (Bloch-Siegert) shift of the electron that is conditional on the state of the nitrogen nuclear spin.

The microwave carrier is tuned to the bare electron resonance, $\omega_e/2\pi \simeq 4.18$ GHz. To select the $|m_I = 0\rangle$ manifold we impose a numerical envelope modulation at the parallel hyperfine splitting, $A_{\parallel}/2\pi = 2.18$ MHz. Consequently, pulses applied to the $|m_I = +1\rangle$ manifold are detuned by $\Delta/2\pi = 2.18$ MHz.

For a constant detuning Δ , the effective Hamiltonian for the electron spin (first qubit) is

$$H = \frac{\hbar}{2} (\Omega X_e + \Delta Z_e),$$

where Ω is the Rabi frequency of the drive, and I_x, I_z are the usual electron spin- $\frac{1}{2}$ operators. The effective rotation frequency is then

$$\Omega_{\text{eff}} = \sqrt{\Omega^2 + \Delta^2}.$$

During the weak- π pulse the electron therefore acquires an additional phase

$$\phi = \Delta\tau,$$

where τ is the pulse duration, and this phase accrues only when $m_I = +1$.

The resulting two-spin unitary can be written

$$U \approx (X_e \otimes |0\rangle\langle 0|_N) + (e^{i\phi Z_e} \otimes |1\rangle\langle 1|_N),$$

where X_e is the weak π -rotation applied to the electron when $m_I = 0$, and $e^{i\phi Z_e}$ is a conditional phase acquired in the $m_I = +1$ manifold due to the off-resonant drive.

Expanding $e^{i\phi Z_e} = \cos(\phi)I_e + i \sin(\phi)Z_e$, we can rewrite the unitary as

$$U \approx X_e \otimes |0\rangle\langle 0|_N + [\cos(\phi)I_e + i \sin(\phi)Z_e] \otimes |1\rangle\langle 1|_N.$$

From this form it is evident that a term proportional to $Z_e Z_N$ is present:

$$U \approx X_e \otimes |0\rangle\langle 0|_N + \cos(\phi) I_e \otimes |1\rangle\langle 1|_N + i \sin(\phi) Z_e \otimes Z_N,$$

where we used the identity

$$|1\rangle\langle 1|_N = \frac{1}{2}(I_N + Z_N), \quad |0\rangle\langle 0|_N = \frac{1}{2}(I_N - Z_N),$$

Crucially, this conditional phase is dependent on the nitrogen nuclear spin state: the electron spin acquires an additional phase only in the $m_I = +1$ branch. It is not simply the fact that the phase arises within the weak- π pulse that is problematic, but rather its conditionality: there is no global single-qubit gate that can correct this term once it has been imprinted.

This explains why the ZZ component remains unchanged after applying single-qubit corrections.

The gate infidelity corroborates this picture: it remains 0.0426 ± 0.0075 before and after correction. Since the dominant single-qubit error is mitigated while the infidelity is unaffected, we conclude that the residual ZZ coupling is now the principal error mechanism.

To counteract this we designed an improved gate scheme that suppresses the coherent ZZ interaction while preserving both coherence and fluorescence contrast.

4.4. WEAK $\pi/2$ SEQUENCE

As shown in Fig. 2.7 and discussed in Section 4.3.3, the dominant coherent two-qubit error in our system is an unwanted $Z_e Z_N$ interaction. It originates from an AC-Stark shift accumulated during periods when the microwave drive, resonant with the $m_I = 0$ electron spin transition, is applied while the nitrogen nuclear spin is in the $m_I = +1$ state. In this case, the drive is detuned by approximately 2.18 MHz, leading to a differential phase accumulation.

To suppress this asymmetry, we restored symmetry in the pulse sequence by replacing the solitary weak π pulse with two weak $\pi/2$ pulses placed symmetrically between dynamical decoupling pulses (Fig. 4.17). While the nitrogen spin state remains fixed throughout the sequence, the time-symmetric structure ensures that any phase accumulated due to off-resonant driving is refocused. This is because the effective Hamiltonian during the second $\pi/2$ pulse has a flipped sign compared to the first, cancelling out Stark-induced evolution. Thus, the differential phase between the $m_I = 0$ and $m_I = +1$ manifolds is eliminated, suppressing the coherent $Z_e Z_N$ interaction.

To leading order in χt , the net evolution is

$$U_{+1} \approx \exp[-i H_0 t] + \mathcal{O}((\chi t)^2).$$

This cancellation arises from the symmetry of the sequence, but its effectiveness depends on the form of H_0 . In the $m_I = 0$ manifold, the drive is on-resonance, so H_0 is dominated by a large X_e term. Although this does not commute with the $Z_e Z_N$ interaction, the fast rotation it induces suppresses the effect of the weaker χZ_e perturbation. In the $m_I = +1$ manifold, the drive is off-resonant, and H_0 is approximately a Z_e rotation.

Since this commutes with $Z_e Z_N$, the two terms effectively evolve independently and the time-symmetric structure of the sequence cancels the net effect of the χZ_e term.

In both cases, the symmetric structure of the sequence eliminates the leading-order contribution of the unwanted $Z_e Z_N$ interaction, suppressing the coherent error.




Figure 4.17: Pulse sequence for the symmetric weak- $\pi/2$ dynamical decoupling scheme. The original central weak π pulse is replaced by two weak $\pi/2$ pulses (black Gaussian envelopes), placed symmetrically between DD pulses. π and $-\pi$ DD pulses are shown as blue rectangles. This symmetric construction ensures equal time spent in resonant and off-resonant driving configurations for both nuclear spin states, cancelling deterministic $Z_e Z_N$ phase accumulation via echo-based refocusing.

We verified the effectiveness of this sequence with the same error-amplification circuit used to benchmark the original weak- π gate: initial local $R_x(\pi/2)$ rotations on both qubits, followed by n applications of the optimised control pulse (nominally a CNOT), and final $R_x(\pi/2)$ rotations. This maps Bell-state preparation and read-out onto a parity measurement that is extremely sensitive to residual ZZ phases.




Figure 4.18: Measured population dynamics under repeated application of the control gate, using the original weak- π DD sequence (left) and the symmetric weak- $\pi/2$ DD sequence (right). Solid markers indicate measured electron (blue) and nitrogen (orange) spin populations; hollow markers represent the ideal evolution. The symmetric sequence exhibits significantly reduced deviation from the ideal populations over multiple gate applications, confirming effective suppression of coherent $Z_e Z_N$ error via symmetric phase cancellation.

Figure 4.18 shows that, with the original weak- π sequence, ZZ errors introduced by the AC Stark shift grow rapidly: the electron population shifts after only a few gate applications. In contrast, the symmetric weak- $\pi/2$ sequence keeps both electron and nitrogen populations almost constant over the same repetition range, demonstrating effective cancellation of the unwanted ZZ interaction. Although minor residual errors remain, the overall dynamics are much closer to the ideal and the gate fidelity is markedly improved.

4.4.1. GATE SET TOMOGRAPHY

To benchmark the new $\pi/2$ gate sequence, we once again employ GST. Since the reference pulse used in this gate has approximately half the amplitude of a standard π pulse,

we also benchmark a sequence based on a $5\pi/2$ pulse for comparison.

The rationale for this is that the reduced amplitude of the $\pi/2$ pulse may lead to increased decoherence, as it becomes more susceptible to stochastic noise arising from environmental couplings.

Additionally, we benchmark a $\pi/2$ pulse that was optimised using our closed-loop optimal control scheme, targeting the Bell state fidelity as the figure of merit.

NON-MARKOVIANITY

Figure 4.19 shows the extracted non-Markovianity metric N_σ for the three pulse types—reference weak $\pi/2$, reference $5\pi/2$, and the optimised Bell-state-targeted weak $\pi/2$ —across various sequence lengths $L = 1, 2, 4, 8$.

We observe similar trends as in Figure 4.14. For the reference weak $\pi/2$ pulse, N_σ values are 44, 39, 46, and 68 for $L = 1, 2, 4, 8$, respectively. The $5\pi/2$ pulse shows a consistently lower N_σ , with values of 36, 31, 33, and 44. The Bell-optimised pulse lies in between: 48, 44, 52, and 64.

Notably, the $5\pi/2$ pulse has the lowest N_σ across all L , suggesting a reduced degree of temporally correlated errors. The weak $\pi/2$ pulse exhibits the highest non-Markovianity at $L = 8$, indicating that its extended susceptibility to low-frequency noise may be more pronounced over longer gate sequences. Meanwhile, the optimised pulse appears to mitigate this somewhat, but still shows elevated N_σ at higher L .




Figure 4.19: Non-Markovianity metric N_σ for three pulse types—reference weak $\pi/2$, reference $5\pi/2$, and the Bell-state-optimised weak $\pi/2$ pulse—evaluated at sequence base lengths $L = 1, 2, 4, 8$. Higher N_σ values indicate stronger statistical deviation from a Markovian noise model.

AVERAGE GATE FIDELITY

Figure 4.20 shows the average gate infidelity for the reference weak $\pi/2$, reference $5\pi/2$, and the optimised Bell state fidelity weak $\pi/2$ pulse.

The reference weak $\pi/2$ pulse exhibits an average gate infidelity of 5.91 ± 0.19 , while the optimised weak $\pi/2$ pulse improves on this with 4.71 ± 0.16 . The weak $5\pi/2$ pulse performs slightly better still, with an infidelity of 4.16 ± 0.15 .

All three values remain significantly higher than the infidelities reported in Figure 4.15, where values range from 3.35 ± 0.14 for the reference weak π pulse to 4.20 ± 0.14 for the Bell only I pulse. Although the weak $5\pi/2$ pulse has the lowest infidelity among the three shown here, it would still rank among the poorest performers in the weak π GST analysis.

To investigate whether this performance degradation arises from insufficient suppression of the ZZ interaction term or an increase in stochastic noise, we next examine the corresponding error generators.




Figure 4.20: Average gate infidelity for the CNOT gate, as estimated by GST, for three control sequences: reference weak $\pi/2$, reference $5\pi/2$, and the Bell-state-optimised weak $\pi/2$ pulse. The optimised pulse achieves lower infidelity than the standard weak $\pi/2$, indicating improved performance, but is still outperformed by the weak $5\pi/2$ sequence. All three infidelities are notably higher than those reported for weak π pulses in Figure 4.15, suggesting that weak- $\pi/2$ -based sequences are more susceptible to stochastic errors.

ERROR GENERATORS

Figure 4.21 shows the coherent and stochastic error generators extracted from GST for the three weak $\pi/2$ gate implementations.

Coherent error generators. As expected, the coherent ZZ term—prominent in the original weak π sequence—is absent across all three variants. This confirms that the symmetric pulse design effectively cancels the deterministic $Z_e Z_N$ interaction.

For the reference weak $\pi/2$ pulse, the dominant coherent terms are XI (-0.05) and XZ ($+0.04$), indicating a small overrotation and conditional phase error. The optimised weak $\pi/2$ pulse shows similar magnitudes but with rotated axes: YI ($+0.03$) and YZ (-0.03).

In contrast, the weak $5\pi/2$ pulse exhibits negligible coherent errors, with no dominant terms, suggesting it implements the target unitary more cleanly.

Stochastic error generators. The stochastic error generators show amplified noise compared to the weak π case. All three pulses exhibit strong contributions in YI, YZ, and IZ, matching the noise structure observed previously (Fig. 4.16) but with amplitudes increased by nearly an order of magnitude.

Among the three, the reference weak $\pi/2$ pulse exhibits the largest stochastic errors, followed by the optimised weak $\pi/2$, which closely mirrors the reference due to its shared starting point in the optimisation. The weak $5\pi/2$ pulse shows the lowest overall noise,




Figure 4.21: Coherent (upper) and stochastic (lower) error generators for the CNOT gate under three weak $\pi/2$ gate implementations: reference weak $\pi/2$, reference $5\pi/2$, and the Bell-state-optimised weak $\pi/2$ pulse. All three sequences successfully suppress the coherent ZZ component, confirming effective cancellation of deterministic two-qubit phase accumulation. However, significant stochastic errors remain—particularly in YI, YZ, and IZ—with amplitudes nearly an order of magnitude larger than in the original weak π configurations. Among the three, the $5\pi/2$ pulse exhibits the lowest overall error, suggesting that stronger driving mitigates stochastic noise.

further supporting the hypothesis that moderately stronger driving improves noise resilience.

Another contributing factor to the increased stochastic noise is the effective doubling in gate sequence length compared to the weak π implementation. Longer gate durations lead to greater exposure to decoherence, which manifests as enhanced terms in the Stochastic error generator.

5

CONCLUSION AND OUTLOOK

5.1. CONCLUSION

In this work, we have investigated both open- and closed-loop quantum optimal control for NV centre systems and characterised the resulting optimised pulses using GST to inform calibration routines.

In summary, this work demonstrates the interplay between theoretical control design and experimental feedback in quantum systems. While open-loop methods offer rapid prototyping, their performance is constrained by model correctness and unaccounted environmental effects. Closed-loop optimisation provides a more robust pathway to high-fidelity control, particularly when paired with diagnostic tools such as GST and drift-aware analysis.

5.1.1. OPEN-LOOP OPTIMISATION

Open-loop control was implemented using the GRAPE algorithm via `qutip-qoc`. In a simplified model, we observed that the maximum achievable fidelity depends strongly on the gate duration relative to characteristic frequencies of the Hamiltonian. Furthermore, when echo sequences are incorporated into the target unitary, any resulting phase must be corrected via an explicit phase gate; otherwise, the optimiser fails to identify meaningful solutions. We generated three distinct pulses using different ansätze and assessed their experimental performance through PL contrast as a function of pulse amplitude.

Of the three pulses, two failed to produce significant PL contrast. However, the qutrit-based pulse exhibited high contrast at both $\text{amp}=1$ and $\text{amp}=19.8$. Characterisation using PL contrast and Bell-state fidelity FoMs revealed that this pulse was nonetheless outperformed by the reference weak π pulse. The superior performance of the reference weak π pulse highlights the limitations of open-loop optimisation in the presence of model inaccuracies and motivates the shift toward closed-loop approaches.

5.1.2. CLOSED-LOOP OPTIMISATION

To this end, we employed dCRAB-based closed-loop optimisation using the QuOCS framework. Four pulses were optimised using either PL contrast or Bell-state fidelity as the FoM, with and without Q components, to assess whether reduced control dimensionality could still yield high-performance pulses. We compared their performance across FoMs and found that the Bell-state fidelity FoM incurred greater uncertainty. The PL contrast results suggest that the reference weak π pulse again outperforms others, likely due to environmental drift during the optimisation process.

To better understand such drift, we analysed a dataset in which the same control pulse was repeatedly applied across iterations. We discovered a strong correlation between PL contrast, magnet position and lens position, indicating that magnetic field drift and optical alignment play a significant role in performance variation and highlighting the need for more stable experimental conditions or drift-robust FoMs.

5.1.3. GATE SET TOMOGRAPHY AND ERROR ANALYSIS

Subsequently, we characterised the closed-loop pulses using GST. We closely examined the CNOT gate infidelities. By analysing the error generators, we identified that most infidelity arises from coherent errors. In particular strong ZZ terms were observed.

We attributed these to phases acquired via the AC Stark effect during off-resonant driving. To compensate for this, we implemented a dynamical decoupling sequence composed of two half- π pulses designed to cancel the accumulated phases. In the resulting coherent error generators, the ZZ term vanished; however, the overall gate fidelity worsened. Further inspection of the stochastic error generators revealed increased stochastic contributions, indicating that the error profile had shifted from predominantly coherent to predominantly stochastic. This degradation is likely due to the reduced amplitude of the half- π pulses, which weakened the control sufficiently that ambient stochastic noise began to dominate the dynamics.

5.2. OUTLOOK

While this work provides a broad exploration of both open- and closed-loop quantum optimal control techniques for NV centre systems, several promising avenues remain open for further investigation.

5.2.1. REFINING OPEN-LOOP APPROACHES

The open-loop GRAPE-based approach, though ultimately sidelined in favour of closed-loop methods, remains valuable for exploring system mechanics and pulse intuition. Future work could explore sophisticated noise models incorporating, for example, spectral density characteristics, correlated noise, or anisotropic dephasing.

Additionally, our preliminary experiments hinted at a dependence of optimiser performance on the underlying Hamiltonian structure. A systematic comparison of how different classical optimisers (e.g., Nelder-Mead, CMA-ES, or gradient-based methods) behave across a range of system models could yield insights into their suitability for various control landscapes.

5.2.2. EXTENDING CLOSED-LOOP OPTIMISATION

For closed-loop optimisation, the outlook is even more expansive. Despite the high experimental overhead, the strategy of reducing the search space through low-dimensional parameterisations (e.g., few-frequency dCRAB [59]) and embedding the system directly into the feedback loop shows considerable promise. Our results indicate that even with limited evaluation bandwidth, meaningful improvements can be obtained. A key limitation, however, remains the simplicity of the employed FoMs, which fail to fully account for drifts. Developing more robust and informative FoMs, potentially incorporating tomographic or phase-sensitive elements, could significantly improve control fidelity.

The use of more comprehensive fidelity metrics, such as RB, also stands out as an exciting prospect [60]. While current experimental overheads made RB impractical in this study, future efforts aimed at optimising experimental throughput could make such metrics accessible. Compared to simple figures of merit based on, for example, population contrast or state fidelity, RB offers a more robust characterisation of average gate performance and noise resilience, without requiring full tomographic reconstruction.

5.2.3. ADVANCED ERROR CHARACTERISATION

The GST-based analysis presented in this work also opens several promising directions for improved characterisation and mitigation of gate errors. The presence of coherent errors suggests that systematic control miscalibrations or persistent Hamiltonian terms may be contributing to gate infidelity. A detailed error budget that links these terms to specific physical mechanisms, such as pulse misalignment, crosstalk, or residual couplings, would be a valuable next step.

In addition, the consistent appearance of stochastic contributions, notably in the YI , YZ , and IZ channels, is of particular concern. These were especially pronounced in GST sequences involving the weak $\pi/2$ pulse. If these noise channels stem from experimental artefacts such as AWG signal distortion, amplifier nonlinearities, or environmental instabilities, identifying their sources could be crucial for achieving gate fidelities exceeding the 99% threshold common in leading NV centre platforms.

Further insight may also be gained by investigating the system's control transfer function, the effective mapping from the programmed waveform to the actual field experienced by the qubit, which can introduce nontrivial distortions even in nominally calibrated systems [61]. Incorporating an experimentally determined transfer function into both open- and closed-loop control schemes could enable more accurate compensation for such distortions.

Ultimately, bridging the gap between pulse design, experimental implementation, and gate performance will require a unified framework that combines GST, transfer function analysis, and control optimisation. Developing such a framework represents a compelling challenge for future work.

5.2.4. TOWARDS FAULT TOLERANCE AND AI-ASSISTED CONTROL

Looking ahead, as the field moves closer to the threshold of fault tolerance, the role of high-precision quantum control will become increasingly pivotal. Optimal control theory, grounded in decades of cross-disciplinary development, is poised to contribute substantially to this frontier. Finally, we note the growing impact of machine learning

in quantum control. Recently, convolutional neural nets have been used in a closed feedback loop to tune quantum dots with a 95% success rate [62]. Deep reinforcement learning has been utilised for superconducting qubits to optimise readout [63] and gate design [64]. In the context of NV centres, deep learning has been used to detect surrounding nuclear spins [65].

Other work using reinforcement learning [66] and neural networks [67] have demonstrated similarly encouraging results, and as the capabilities of AI systems expand, they may become essential tools for real-time pulse design and qubit calibration in increasingly complex quantum devices.

ACKNOWLEDGEMENTS

Science is a collaborative endeavour, and luckily so, because otherwise it would be boring.

Ben and Margriet, thank you for being my daily supervisors. You guys were always available for help or answering all the stupid questions I could come up with. Good luck with the final push for your PhD Ben, and I wish you a lot of happy climbing sessions Margriet.

Jiwon, thanks for the cool ideas and the time spent in the basement. I enjoyed our conversations while waiting for measurements a lot.

Tim thank you for welcoming me into your lab with open arms and for providing much needed feedback and guidance. Furthermore, I want to thank the other members of the Taminiau labs for everything. A lab is only as fun as its members, and this is a fun one.

I want to thank my parents, without my birth none of this would be possible, and my sister for picking on me as a kid leading to character development (according to my therapist).

Lastly, I want to thank Kirsten. Long distance sucks, but not for long.

BIBLIOGRAPHY

- [1] P. Benioff, "The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines," *Journal of Statistical Physics*, vol. 22, no. 5, pp. 563–591, May 1980. DOI: [10.1007/bf01011339](https://doi.org/10.1007/bf01011339). [Online]. Available: <https://doi.org/10.1007/bf01011339>.
- [2] R. P. Feynman, "Simulating physics with computers," *International Journal of Theoretical Physics*, vol. 21, no. 6-7, pp. 467–488, Jun. 1982. DOI: [10.1007/bf02650179](https://doi.org/10.1007/bf02650179). [Online]. Available: <https://doi.org/10.1007/bf02650179>.
- [3] I. I. Manin, *Vychislimoe i nevychislimoe / I.I. Manin*, Russian. "Sov. radio," Moskva, 1980, 127 p.
- [4] P. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proceedings 35th Annual Symposium on Foundations of Computer Science*, IEEE Comput. Soc. Press, 1994. DOI: [10.1109/sfcs.1994.365700](https://doi.org/10.1109/sfcs.1994.365700). [Online]. Available: <https://doi.org/10.1109/sfcs.1994.365700>.
- [5] I. L. Chuang, N. Gershenfeld, and M. Kubinec, "Experimental implementation of fast quantum searching," *Phys. Rev. Lett.*, vol. 80, pp. 3408–3411, 15 Apr. 1998. DOI: [10.1103/PhysRevLett.80.3408](https://doi.org/10.1103/PhysRevLett.80.3408). [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.80.3408>.
- [6] C. M. Löschner, J. Mosca Toba, A. C. Hughes, *et al.*, "Scalable, high-fidelity all-electronic control of trapped-ion qubits," *arXiv preprint arXiv:2407.07694*, 2024. [Online]. Available: <https://arxiv.org/abs/2407.07694>.
- [7] D. A. Rower, L. Ding, H. Zhang, *et al.*, "Suppressing counter-rotating errors for fast single-qubit gates with fluxonium," *PRX Quantum*, vol. 5, p. 040342, 4 Dec. 2024. DOI: [10.1103/PRXQuantum.5.040342](https://doi.org/10.1103/PRXQuantum.5.040342). [Online]. Available: <https://link.aps.org/doi/10.1103/PRXQuantum.5.040342>.
- [8] H. P. Bartling, J. Yun, K. N. Schymik, *et al.*, *Universal high-fidelity quantum gates for spin-qubits in diamond*, 2024. arXiv: [2403.10633](https://arxiv.org/abs/2403.10633) [quant-ph]. [Online]. Available: <https://arxiv.org/abs/2403.10633>.
- [9] *Ionq achieves industry breakthrough – first trapped ion quantum system to surpass 99.9% fidelity on barium*, <https://investors.ionq.com/news/news-details/2024/IonQ-Achieves-Industry-Breakthrough--First-Trapped-Ion-Quantum-System-to-Surpass-99.9-Fidelity-on-Barium/>, Accessed: 2025-05-25, 2024.
- [10] X. Li, Y. Zhang, L. Wang, *et al.*, "Realization of high-fidelity cz gate based on a double-transmon coupler," *Physical Review X*, vol. 14, no. 4, p. 041050, 2024. DOI: [10.1103/PhysRevX.14.041050](https://doi.org/10.1103/PhysRevX.14.041050). [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.14.041050>.


- [11] J. Preskill, “Quantum computing in the nisq era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
- [12] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Physical Review A*, vol. 86, no. 3, p. 032 324, 2012.
- [13] M. W. Doherty, N. B. Manson, P. Delaney, *et al.*, “The nitrogen-vacancy colour centre in diamond,” *Physics Reports*, vol. 528, no. 1, pp. 1–45, 2013.
- [14] N. Khaneja, T. Reiss, C. Kehlet, *et al.*, “Optimal control of coupled spin dynamics: Design of nmr pulse sequences by gradient ascent algorithms,” *Journal of Magnetic Resonance*, vol. 172, no. 2, pp. 296–305, 2005.
- [15] P. Schindler, M. Müller, D. Nigg, *et al.*, “Quantum simulation of dynamical maps with trapped ions,” *Nature Physics*, vol. 9, no. 6, pp. 361–367, 2013.
- [16] D. J. Egger and F. K. Wilhelm, “Optimized pulses for the control of superconducting qubits,” *Superconductor Science and Technology*, vol. 27, no. 1, p. 014 001, 2014.
- [17] F. Dolde, I. Jakobi, B. Naydenov, *et al.*, “High-fidelity spin entanglement using optimal control,” *Nature Communications*, vol. 5, p. 3371, 2014.
- [18] S. J. Glaser, U. Boscain, T. Calarco, *et al.*, “Training schrödinger’s cat: Quantum optimal control,” *The European Physical Journal D*, vol. 69, no. 12, p. 279, 2015.
- [19] T. Caneva, T. Calarco, and S. Montangero, “Chopped random-basis quantum optimization,” *Physical Review A*, vol. 84, no. 2, p. 022 326, 2011.
- [20] E. Knill, D. Leibfried, R. Reichle, *et al.*, “Randomized benchmarking of quantum gates,” *Physical Review A*, vol. 77, no. 1, p. 012 307, 2008.
- [21] R. Blume-Kohout, J. K. Gamble, E. Nielsen, *et al.*, “Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit,” *arXiv preprint arXiv:1310.4492*, 2013.
- [22] C. E. Bradley, “Order from disorder: Control of multi-qubit spin registers in diamond,” Final published version, PhD thesis, Delft University of Technology, 2021. DOI: [10 . 4233 / uuid : acafe18b - 3345 - 4692 - 9c9b - 05e970ffbe40](https://doi.org/10.4233/uuid:acafe18b-3345-4692-9c9b-05e970ffbe40). [Online]. Available: <https://doi.org/10.4233/uuid:acafe18b-3345-4692-9c9b-05e970ffbe40>.
- [23] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. L. Hollenberg, “The nitrogen-vacancy colour centre in diamond,” *Physics Reports*, vol. 528, no. 1, pp. 1–45, 2013. DOI: [10.1016/j.physrep.2013.02.001](https://doi.org/10.1016/j.physrep.2013.02.001).
- [24] F. Jelezko and J. Wrachtrup, “Single defect centres in diamond: A review,” *physica status solidi (a)*, vol. 203, no. 13, pp. 3207–3225, 2006. DOI: [10 . 1002 / pssa . 200671403](https://doi.org/10.1002/pssa.200671403).
- [25] K. Khodjasteh and D. A. Lidar, “Dynamical decoupling in the presence of realistic pulse errors,” *arXiv*, vol. 1011.1903, 2010, Discusses how two-axis DD sequences like XY4/XY8 correct higher-order errors.

- [26] L. M. K. Vandersypen and I. L. Chuang, "Nmr techniques for quantum control and computation," *Reviews of Modern Physics*, vol. 76, no. 4, pp. 1037–1069, Jan. 2005, ISSN: 1539-0756. DOI: [10.1103/revmodphys.76.1037](https://doi.org/10.1103/revmodphys.76.1037). [Online]. Available: <http://dx.doi.org/10.1103/RevModPhys.76.1037>.
- [27] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko, *The Mathematical Theory of Optimal Processes*. Interscience Publishers, 1962.
- [28] R. Bellman, *Dynamic Programming*. Princeton University Press, 1957.
- [29] A. E. Bryson and Y.-C. Ho, *Applied Optimal Control: Optimization, Estimation and Control*. Taylor & Francis, 1975.
- [30] B. Siciliano and O. Khatib, *Springer Handbook of Robotics*. Springer, 2016. DOI: [10.1007/978-3-319-32552-1](https://doi.org/10.1007/978-3-319-32552-1).
- [31] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle III, *Process Dynamics and Control*, 3rd. Wiley, 2010.
- [32] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, 2017.
- [33] M. I. Kamien and N. L. Schwartz, *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*. Dover Publications, 2012.
- [34] R. S. Judson and H. Rabitz, "Teaching lasers to control molecules," *Physical Review Letters*, vol. 68, no. 10, pp. 1500–1503, 1992.
- [35] D. J. Tannor, R. Kosloff, and S. A. Rice, "Control of selectivity of chemical reaction via control of wave packet evolution," *The Journal of Chemical Physics*, vol. 83, no. 10, pp. 5013–5018, 1985.
- [36] V. F. Krotov, *Global Methods in Optimal Control Theory*. Marcel Dekker, 1996.
- [37] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, "Optimal control of coupled spin dynamics: Design of nmr pulse sequences by gradient ascent algorithms," *Journal of Magnetic Resonance*, vol. 172, no. 2, pp. 296–305, 2005.
- [38] F. K. Wilhelm, S. Kirchhoff, S. Machnes, N. Wittler, and D. Sugny, *An introduction into optimal control for quantum technologies*, 2020. arXiv: [2003.10132](https://arxiv.org/abs/2003.10132) [quant-ph]. [Online]. Available: <https://arxiv.org/abs/2003.10132>.
- [39] B. C. Hall, *An Elementary Introduction to Groups and Representations*. American Mathematical Society, 2000, Available at <https://arxiv.org/abs/math-ph/0005032>.
- [40] T. Caneva, T. Calarco, and S. Montangero, "Chopped random-basis quantum optimization," *Physical Review A*, vol. 84, p. 022 326, 2011. DOI: [10.1103/PhysRevA.84.022326](https://doi.org/10.1103/PhysRevA.84.022326).
- [41] N. Rach, M. M. Müller, T. Calarco, and S. Montangero, "Dressing the chopped-random-basis optimization: A bandwidth-limited access to the trap-free landscape," *Physical Review A*, vol. 92, p. 062 343, 2015. DOI: [10.1103/PhysRevA.92.062343](https://doi.org/10.1103/PhysRevA.92.062343).

- [42] M. M. Müller, R. S. Said, F. Jelezko, T. Calarco, and S. Montangero, “One decade of quantum optimal control in the chopped random basis,” *Reports on Progress in Physics*, vol. 85, no. 7, p. 076 001, 2022, Preprint arXiv:2104.07687 (2021). DOI: [10.1088/1361-6633/ac723c](https://doi.org/10.1088/1361-6633/ac723c).
- [43] D. Greenbaum, “Introduction to quantum gate set tomography,” *arXiv preprint arXiv:1509.02921*, 2015, Accessed 17 Apr 2025.
- [44] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout, “Gate set tomography,” *Quantum*, vol. 5, p. 557, 2021. DOI: [10.22331/q-2021-10-05-557](https://doi.org/10.22331/q-2021-10-05-557).
- [45] PYGSTI Developers, *PYGSTI user manual (v0.9.x)*, <https://pygsti.readthedocs.io>, Accessed 17 Apr 2025, 2022.
- [46] M. A. Nielsen, “A simple formula for the average gate fidelity of a quantum dynamical operation,” *Physics Letters A*, vol. 303, no. 4, pp. 249–252, Oct. 2002, ISSN: 0375-9601. DOI: [10.1016/S0375-9601\(02\)01272-0](https://doi.org/10.1016/S0375-9601(02)01272-0). [Online]. Available: [http://dx.doi.org/10.1016/S0375-9601\(02\)01272-0](http://dx.doi.org/10.1016/S0375-9601(02)01272-0).
- [47] R. Blume-Kohout, M. P. da Silva, E. Nielsen, *et al.*, “A taxonomy of small markovian errors,” *PRX Quantum*, vol. 3, no. 2, May 2022, ISSN: 2691-3399. DOI: [10.1103/prxquantum.3.020335](https://doi.org/10.1103/prxquantum.3.020335). [Online]. Available: <http://dx.doi.org/10.1103/PRXQuantum.3.020335>.
- [48] N. Lambert, E. Giguère, P. Menczel, *et al.*, *Qutip 5: The quantum toolbox in python*, 2024. arXiv: [2412.04705](https://arxiv.org/abs/2412.04705) [quant-ph]. [Online]. Available: <https://arxiv.org/abs/2412.04705>.
- [49] T. Schuster, R. MacDonell, N. K. Singh, *et al.*, *QuTiP-QOC: Quantum Optimal Control Module for QuTiP*, <https://github.com/qutip/qutip-qoc>, Version 0.2.0, 2022.
- [50] M. Rossignolo, T. Reisser, A. Marshall, *et al.*, “Quocs: The quantum optimal control suite,” *Computer Physics Communications*, vol. 291, p. 108 782, Oct. 2023, ISSN: 0010-4655. DOI: [10.1016/j.cpc.2023.108782](https://doi.org/10.1016/j.cpc.2023.108782). [Online]. Available: <http://dx.doi.org/10.1016/j.cpc.2023.108782>.
- [51] D. Greenbaum, K. Rudinger, E. Nielsen, T. Proctor, K. Rudinger, and R. Blume-Kohout, *Pygsti: Python gate set tomography implementation*, version v0.9.9.1, 2015. DOI: [10.5281/zenodo.46777](https://doi.org/10.5281/zenodo.46777). [Online]. Available: <https://doi.org/10.5281/zenodo.46777>.
- [52] J. Johansson, P. Nation, and F. Nori, “Qutip: An open-source python framework for the dynamics of open quantum systems,” *Computer Physics Communications*, vol. 183, no. 8, pp. 1760–1772, Aug. 2012, ISSN: 0010-4655. DOI: [10.1016/j.cpc.2012.02.021](https://doi.org/10.1016/j.cpc.2012.02.021). [Online]. Available: <http://dx.doi.org/10.1016/j.cpc.2012.02.021>.
- [53] L. Viola, E. Knill, and S. Lloyd, “Dynamical decoupling of open quantum systems,” *Physical Review Letters*, vol. 82, no. 12, pp. 2417–2421, 1999. DOI: [10.1103/PhysRevLett.82.2417](https://doi.org/10.1103/PhysRevLett.82.2417).

- [54] H. M. Lim, G. T. Genov, R. Sailer, *et al.*, *Efficiency of optimal control for noisy spin qubits in diamond*, 2024. arXiv: [2411 . 05078 \[quant-ph\]](https://arxiv.org/abs/2411.05078). [Online]. Available: <https://arxiv.org/abs/2411.05078>.
- [55] J. G. Proakis and M. Salehi, *Digital Communications*, 5th. New York: McGraw-Hill, 2007, ISBN: 9780072957167.
- [56] D. K. Weaver, “A third method of generating single-sideband signals,” *Proceedings of the IRE*, vol. 44, no. 12, pp. 1703–1705, 1956. DOI: [10 . 1109 / JRPROC . 1956 . 275144](https://doi.org/10.1109/JRPROC.1956.275144).
- [57] F. Arute *et al.*, “Mitigating transients in flux-control signals in a superconducting quantum processor,” *arXiv preprint arXiv:2503.08645*, 2024. [Online]. Available: <https://arxiv.org/abs/2503.08645>.
- [58] R. Blume-Kohout, J. K. Gamble, E. Nielsen, *et al.*, “Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography,” *Nature Communications*, vol. 8, no. 1, Feb. 2017, ISSN: 2041-1723. DOI: [10 . 1038 / ncomms14485](https://doi.org/10.1038/ncomms14485). [Online]. Available: <http://dx.doi.org/10.1038/ncomms14485>.
- [59] M. M. Müller, S. Gherardini, T. Calarco, S. Montangero, and F. Caruso, “Information theoretical limits for quantum optimal control solutions: Error scaling of noisy channels,” *arXiv preprint*, vol. arXiv:2006.16113, 2020. [Online]. Available: <https://arxiv.org/abs/2006.16113>.
- [60] P. J. Vetter, T. Reisser, M. G. Hirsch, *et al.*, “Gate-set evaluation metrics for closed-loop optimal control on nitrogen-vacancy center ensembles in diamond,” *npj Quantum Information*, vol. 10, no. 1, Oct. 2024, ISSN: 2056-6387. DOI: [10 . 1038 / s41534 - 024 - 00893 - y](https://doi.org/10.1038/s41534-024-00893-y). [Online]. Available: <http://dx.doi.org/10.1038/s41534-024-00893-y>.
- [61] M. A. Rol, L. Ciorciaro, F. K. Malinowski, *et al.*, “Time-domain characterization and correction of on-chip distortion of control pulses in a quantum processor,” *Applied Physics Letters*, vol. 116, no. 5, Feb. 2020, ISSN: 1077-3118. DOI: [10 . 1063 / 1 . 5133894](https://doi.org/10.1063/1.5133894). [Online]. Available: <http://dx.doi.org/10.1063/1.5133894>.
- [62] V. Yon, B. Galaup, C. Rohrbacher, *et al.*, “Experimental online quantum dots charge autotuning using neural networks,” *Nano Letters*, vol. 25, no. 10, pp. 3717–3725, Feb. 2025, ISSN: 1530-6992. DOI: [10 . 1021 / acs . nanolett . 4c04889](https://doi.org/10.1021/acs.nanolett.4c04889). [Online]. Available: <http://dx.doi.org/10.1021/acs.nanolett.4c04889>.
- [63] Y. Baum, M. Amico, S. Howell, *et al.*, “Experimental deep reinforcement learning for error-robust gate-set design on a superconducting quantum computer,” *PRX Quantum*, vol. 2, p. 040324, 4 Nov. 2021. DOI: [10 . 1103 / PRXQuantum . 2 . 040324](https://doi.org/10.1103/PRXQuantum.2.040324). [Online]. Available: <https://link.aps.org/doi/10.1103/PRXQuantum.2.040324>.
- [64] E. Wright and R. De Sousa, “Fast quantum gate design with deep reinforcement learning using real-time feedback on readout signals,” in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, IEEE, Sep. 2023, pp. 1295–1303. DOI: [10 . 1109 / qce57702 . 2023 . 00146](https://doi.org/10.1109/qce57702.2023.00146). [Online]. Available: <http://dx.doi.org/10.1109/QCE57702.2023.00146>.

- [65] K. Jung, M. H. Abobeih, J. Yun, *et al.*, “Deep learning enhanced individual nuclear-spin detection,” *npj Quantum Information*, vol. 7, no. 1, Feb. 2021, ISSN: 2056-6387. DOI: [10.1038/s41534-021-00377-3](https://doi.org/10.1038/s41534-021-00377-3). [Online]. Available: <http://dx.doi.org/10.1038/s41534-021-00377-3>.
- [66] L. Giannelli, S. Sgroi, J. Brown, *et al.*, “A tutorial on optimal control and reinforcement learning methods for quantum technologies,” *Physics Letters A*, vol. 434, p. 128 054, May 2022, ISSN: 0375-9601. DOI: [10.1016/j.physleta.2022.128054](https://doi.org/10.1016/j.physleta.2022.128054). [Online]. Available: <http://dx.doi.org/10.1016/j.physleta.2022.128054>.
- [67] M. Y. Niu, S. Boixo, V. Smelyanskiy, and H. Neven, *Universal quantum control through deep reinforcement learning*, 2018. arXiv: [1803.01857](https://arxiv.org/abs/1803.01857) [quant-ph]. [Online]. Available: <https://arxiv.org/abs/1803.01857>.



THEORETICAL BACKGROUND

A.1. UNCERTAINTY COMPARISON

Since both the PL contrast and Bell state fidelity metrics are derived from measured populations, it is important to consider how their respective uncertainties compare. Specifically, the parity in each basis is given by

$$P = 2(p_{00} + p_{11}) - 1,$$

with uncertainty

$$\sigma_P = 2 \cdot \sqrt{\sigma_{00}^2 + \sigma_{11}^2}.$$

By contrast, the PL contrast metric involves two directly measured populations:

$$\text{PL Contrast} = P_{e=0|N=0} - P_{e=0|N=1},$$

with associated uncertainty

$$\sigma_{\text{PL}} = \sqrt{\sigma_0^2 + \sigma_1^2}.$$

The final Bell state fidelity is computed by combining the parity measurements in the three complementary bases:

$$F = \frac{1 + P_{XX} - P_{YY} + P_{ZZ}}{4},$$

with propagated uncertainty

$$\sigma_F = \frac{1}{4} \sqrt{\sigma_{XX}^2 + \sigma_{YY}^2 + \sigma_{ZZ}^2}.$$

If all underlying measurement uncertainties were equal, one would expect a smaller total uncertainty for the Bell state fidelity ($\sigma_F \approx 1.22\sigma$) than for the PL contrast ($\sigma_{\text{PL}} \approx 1.41\sigma$), since the fidelity averages over three terms. However, in practice the parity measurements involve joint readout of two qubits, which introduces larger uncertainties

than single-qubit readout. This results in the larger uncertainties typically observed in the Bell state fidelity.

This increased measurement uncertainty does not reflect a larger uncertainty in the intrinsic quality of the gate itself, but rather a larger uncertainty in the estimate of that quality provided by the parity-based metric.

B

CODE LISTINGS

B.1. CUSTOM HELPER FUNCTIONS

This appendix details the custom helper functions used throughout the thesis for analysing optimal control results.

- `plot_expectation_values(result: qutip.Result, times: numpy.array)`: Takes a `qutip.Result` object and a `numpy` array of timesteps, and plots the x , y , and z expectation values for all qubits over time. This visualisation helps assess the state trajectories during the control sequence.
- `plot_bloch_spheres(result: qutip.Result, show_all_states: bool, show_initial_and_final_states: bool)`: Visualises the state evolution of qubits on the Bloch sphere. The `show_all_states` flag controls whether all intermediate states are plotted, while `show_initial_and_final_states` highlights only the initial and final states. This provides an intuitive geometric interpretation of the qubit dynamics.
- `plot_concurrence(result: qutip.Result, times: numpy.array)`: Computes and plots the concurrence, a measure of bipartite entanglement, as a function of time. This is particularly useful for evaluating the generation and preservation of entanglement during gate operations.
- `plot_optimised_amps(result: qutip_qoc.Result, show_initial: bool)`: Plots the optimised control amplitudes obtained from the optimal control routine. If `show_initial` is set to `True`, the initial guess amplitudes used in the optimisation are also displayed for comparison.
- `plot_fft_of_result(result: qutip_qoc.Result)`: Computes and plots the Fourier transform of the optimised control pulses. This analysis provides insight into the spectral content of the control fields and helps identify whether undesired frequency components are present.

B.2. QUOCS OPTIMISATION DICTIONARY

An example of an optimisation dictionary used in this thesis is shown below:

```

1 optimization_dictionary = {"optimization_client_name": "
    CNOT_BELL_PARITY"}
2 optimization_dictionary["algorithm_settings"] = {
3     "algorithm_name": "dCRAB",
4     "optimization_direction": "maximization",
5     "super_iteration_number": 5,
6     "max_eval_total": 30000,
7     "dsm_settings": {
8         "general_settings": {
9             "dsm_algorithm_name": "NelderMead",
10            "is_adaptive": False
11        },
12        "stopping_criteria": {
13            "xatol": 1e-6,
14            "fatol": 1e-6,
15            "change_based_stop": {
16                "cbs_funcnt_evals": 500,
17                "cbs_change": 1e-4
18            }
19        }
20    }
21 }
22
23 pulse_x = {
24     "pulse_name": "Pulse_0",
25     "upper_limit": 0.1,
26     "lower_limit": -0.1,
27     "bins_number": n_bins,
28     "amplitude_variation": 0.01,
29     "time_name": "time_1",
30     "initial_guess": {
31         "function_type": "list_function",
32         "list_function": x_array
33     },
34     "basis": {
35         "basis_name": "Fourier",
36         "basis_vector_number": 3,
37         "random_super_parameter_distribution": {
38             "distribution_name": "Uniform",
39             "lower_limit": 0.01,
40             "upper_limit": 10.0
41         }
42     }
43 }
44
45 pulse_y = pulse_x.copy()
46 pulse_y["pulse_name"] = "Pulse_1"
47 pulse_y["initial_guess"]["list_function"] = y_array
48
49 time1 = {"time_name": "time_1", "initial_value": n_bins}
50
51 optimization_dictionary["pulses"] = [pulse_x, pulse_y]
52 optimization_dictionary["parameters"] = []

```

```
53 optimization_dictionary["times"] = [time1, time1]
```

Listing B.1: Example of an optimisation dictionary

B

B.3. PyGSTi

An example of using pyGSTi is shown below:

```
1 import pygsti
2 from pygsti.modelpacks import smq1Q_XYI
3
4 # 1) Obtain the ideal "target" Model (a predefined stock model)
5 mdl_ideal = smq1Q_XYI.target_model()
6
7 # 2) Generate a GST experiment design
8 edesign = smq1Q_XYI.create_gst_experiment_design(4) # Define max
   circuit length
9
10 # 3) Create a data-set template
11 pygsti.io.write_empty_dataset("MyData.txt", edesign.
   all_circuits_needing_data,
12                               "## Columns = 0 count, 1 count")
13
14 # STOP! "MyData.txt" now contains placeholders (zeros) for data.
15 # Replace the zeros with actual experimental data, then continue:
16
17 ds = pygsti.io.load_dataset("MyData.txt") # Load data into a DataSet
   object
18
19 # Alternatively, simulate experimental data:
20 # ds = pygsti.data.simulate_data(mdl_ideal, edesign, num_samples=1000)
21
22 # 4) Run GST using the modern object-based interface
23 data = pygsti.protocols.ProtocolData(edesign, ds) # Bundle dataset and
   circuits
24 protocol = pygsti.protocols.StandardGST() # Select the GST
   protocol
25 results = protocol.run(data) # Execute the
   protocol
26
27 # 5) Generate an HTML report with detailed results
28 report = pygsti.report.construct_standard_report(results, title="My
   Report", verbosity=1)
29 report.write_html("myReport", auto_open=True, verbosity=1) #
   Optionally write Jupyter notebooks as well
```

Listing B.2: Example usage of pyGSTi

C

EXPERIMENTAL RESULTS




Figure C.1: Amplitude sweep results. Top row: full sweep range; bottom row: zoomed-in view. Left: vanilla (echo-compensated) pulse; centre: no-echo pulse; right: Qutrit (echo-compensated) pulse.

C




Figure C.2: FoM value versus function evaluation for various pulse parameterisations. Vertical dotted lines indicate the start of new super-iterations.