



**Politecnico
di Torino**



POLITECNICO DI TORINO

College of Electronics, Telecommunications and Physics Engineering

Master's Degree in Quantum Engineering

Microwave Up-Conversion for Quantum Sensing with Color Centers in Diamond

Candidate

Hayam Ramadan s327204

Academic Supervisor

Prof. Bartolomeo Montrucchio

Politecnico di Torino

Research supervisor

Dr. Ettore Bernardi

INRiM – Istituto Nazionale
di Ricerca Metrologica

Academic Year 2025–2026

Contents

1	Introduction	1
2	Theoretical Background	6
2.1	Color Centers in Diamond	6
2.1.1	Nitrogen-Vacancy Centers and Charge States	7
2.1.2	Electronic Structure of the NV Center	8
2.1.3	Optical Cycle and Spin-Dependent Readout	9
2.1.4	The NV Center as an Atomic-Scale Magnetic Sensor	11
2.2	ODMR Techniques	16
2.2.1	Continuous wave mode	16
2.2.2	Pulse sequences (Pulsed ODMR)	17
2.3	Coherent Control of the NV Spin	18
2.3.1	Two-Level Approximation	18
2.3.2	Basic Pulsed Control Sequences	19
2.3.3	Microwave Phase and Rotation Axis	24
2.3.4	Pulse Area and Pulse Duration	26
3	Experimental Setup	28
3.1	Photoluminescence Microscopy	28
3.1.1	Light Detector and Imaging Mode	30
3.1.2	Photoluminescence Digitization and Visualization	30
3.1.3	Scanning in the Confocal Microscope	31
3.2	Pulse Generation	31
3.3	Laboratory Set-up	32
3.3.1	Excitation Source and Optical Alignment	32
3.3.2	Confocal Microscope and Optical Filtering	32
3.3.3	Diamond Sample and Antenna Structure	35

3.4	Bias Magnetic Field	38
3.4.1	Nanopositioning System	38
3.4.2	Microwave Control Configuration	39
4	The HDAWG and Qudi Integration	41
4.1	Motivation for Using the HDAWG	41
4.1.1	Capabilities of the Zurich Instruments HDAWG	41
4.1.2	Relevant HDAWG Features for Pulsed ODMR	44
4.1.3	Advantages for Pulse Shaping and Phase Control	45
4.2	Qudi Architecture	46
4.2.1	Graphical User Interface, Logic, and Hardware Layers	47
4.2.2	Configuration and Extension of Qudi	49
4.2.3	Existing Pulse Generation Support and Missing Functionality	51
4.3	Development of the HDAWG Hardware Module	52
4.3.1	Architecture of the Developed Module	52
4.3.2	Communication with the LabOne Data Server	53
4.3.3	Waveform Generation from Qudi Pulse Sequences	54
4.3.4	Generation of SEQC Programs	54
4.3.5	Multi-channel Waveform Management	55
4.3.6	Implementation of Gaussian and DRAG Pulses	56
4.3.7	Sequence Compilation and Upload	57
4.3.8	Instrument Constraints	58
5	Microwave Up-Conversion Setup	60
5.1	Theory of Up-Conversion	60
5.1.1	Need for Up-Conversion in NV-Based Experiments	60
5.1.2	Mixer Operation and Sideband Generation	61
5.1.3	Bandwidth Limitation of the Up-Conversion Chain	63
5.2	Hardware Chain	64
5.2.1	HDAWG Output Stage	65
5.2.2	Local Oscillator Source	65
5.2.3	External IQ Mixer	66
5.2.4	Microwave Amplifier	69
5.2.5	Antenna Coupling to the Diamond Sample	69
5.2.6	Complete Signal Chain	69
5.2.7	IQ Mixer Calibration Procedure	70
5.3	Phase Signal and IQ Control	72
5.3.1	In-Phase and Quadrature Components	72

5.3.2	Amplitude and Phase Control	73
6	Experimental Results	76
6.1	CW ODMR Characterization	76
6.1.1	Rectangular and Gaussian Excitation Profiles	77
6.1.2	Calibration of the π -Pulse Duration Using Rabi Oscillations	79
6.1.3	Pulsed ODMR Comparison	81
6.1.4	Implementation	82
7	Conclusions and Future Work	84
7.1	Summary of the Developed System	84
7.2	Main Engineering Contributions	85
7.3	Limitations	85
7.4	Future Improvements	86
	References	87

1 Introduction

The ability to control individual quantum systems has transformed quantum mechanics from a purely descriptive theory into a technological platform. Modern quantum technologies are built on the possibility of preparing, manipulating, and measuring quantum states with sufficient precision to use them as sensors, information carriers, or probes of microscopic physical processes. Within this broad field, solid-state quantum systems are especially attractive because they combine quantum functionality with the mechanical stability and potential scalability of crystalline materials. Among the different solid-state platforms, color centers in diamond have emerged as one of the most versatile systems for quantum sensing, nanoscale imaging, and quantum information science.

A color center is a point defect in a crystal whose electronic structure introduces optically active states within the bandgap of the host material. In diamond, the most studied and technologically relevant defect is the negatively charged nitrogen-vacancy center, denoted NV^- . It consists of a substitutional nitrogen atom adjacent to a vacancy in the carbon lattice. Although structurally simple, this defect has become a central platform in solid-state quantum physics because it provides an optically addressable spin system embedded in an exceptionally robust material. The large bandgap, high thermal conductivity, chemical stability, and low magnetic noise of diamond make it an excellent host for preserving quantum properties while still allowing optical and microwave access to the defect [1, 2].

The first demonstrations of optically detected magnetic resonance in single NV centers marked an important step in the development of solid-state quantum control. These experiments showed that an individual defect in diamond could be addressed optically, manipulated coherently, and read out through its fluorescence. Since then, NV centers have evolved from a model system for fundamental studies of single defects into a mature platform for quantum sensing and quantum control. Their importance comes from a rare combination of properties: the spin state can be initialized optically, manipulated by microwave fields, and measured through spin-dependent photoluminescence, all under

ambient conditions. This possibility of performing coherent quantum experiments at room temperature distinguishes NV centers from many other quantum platforms and is one of the main reasons for their widespread use [1, 3].

The physical origin of this functionality lies in the electronic ground state of the NV⁻ center. The ground state is a spin triplet, and in the absence of an external magnetic field the $m_s = 0$ and $m_s = \pm 1$ sublevels are separated by a zero-field splitting of approximately $D \approx 2.87$ GHz. This transition lies in the microwave domain. A microwave field resonant with the transition can drive coherent rotations of the spin state, while optical excitation preferentially initializes the system into the $m_s = 0$ state and provides a fluorescence signal that depends on the spin population. This combination of optical initialization, microwave control, and optical readout forms the basis of optically detected magnetic resonance (ODMR), which is the main experimental technique used to probe and control NV-center spin states [1, 4].

Historically, the appeal of NV centers was strengthened by the fact that they bridge two areas that are often experimentally separated: quantum optics and magnetic resonance. From the quantum-optics point of view, an NV center is an optically active emitter that can be detected at the single-defect level. From the magnetic-resonance point of view, it is a spin system whose energy levels can be controlled with microwave radiation. This dual nature makes the NV center a natural interface between photons and spins. As a result, NV centers have been studied not only as sensors but also as single-photon sources, spin qubits, quantum memories, and possible nodes for distributed quantum networks [5, 6].

In quantum sensing, the NV center is especially powerful because its spin transition frequency is sensitive to its local environment. Magnetic fields, temperature variations, strain, electric fields, and nearby magnetic moments can perturb the spin Hamiltonian and shift or modify the ODMR response. By measuring these changes, the NV center can be used as a local quantum sensor. Since the defect has atomic dimensions and can be placed very close to the sample under investigation, it enables spatial resolutions that are difficult to achieve with conventional sensors. This has led to applications in nanoscale magnetometry, magnetic imaging of materials, investigation of current distributions, detection of spin noise, and sensing in biological or chemical environments [3, 2].

The relevance of NV centers therefore lies in their ability to connect several areas of quantum technology: optical readout, microwave spin control, quantum sensing, and solid-state integration. Their room-temperature operation and sensitivity to local perturbations make them especially suited for sensing applications, while their coherent spin dynamics make them a useful platform for testing control techniques originally developed in magnetic resonance and quantum information processing. For this reason, improving the way NV

spins are driven and controlled remains an important part of the development of advanced NV-based sensing platforms.

Microwave control is the central mechanism used to manipulate the NV spin state. The ground-state spin transition of the negatively charged NV center lies near 2.87 GHz, so resonant microwave fields are required to drive coherent rotations between the spin sublevels. In basic experiments, this can be achieved by switching a continuous-wave microwave source on and off, producing rectangular pulses of controlled duration. However, such a method only controls when the microwave field is applied; it does not provide full control over the temporal envelope or the phase of the pulse.

A more flexible approach is to generate the control field with an arbitrary waveform generator. An AWG allows the pulse to be defined as a programmable time-dependent waveform rather than as a simple on-off signal. This makes it possible to control the amplitude envelope, phase, and quadrature components of the microwave drive. These degrees of freedom are important because the shape and phase of a pulse determine how the spin state evolves, how selectively a transition is addressed, and how sensitive the operation is to detuning, amplitude fluctuations, and other experimental imperfections [7, 8].

Pulse shaping provides a direct way to improve microwave control. A rectangular pulse has abrupt temporal edges and therefore contains broad spectral components. Smoother envelopes, such as Gaussian pulses, reduce the abruptness of the switching and can improve spectral selectivity. More advanced pulse forms, such as DRAG corrections, introduce derivative-based quadrature components designed to compensate unwanted dynamics and suppress leakage in weakly nonlinear quantum systems [9]. More generally, shaped pulses expand the available control space and provide a systematic route toward more robust and selective spin manipulation.

The difficulty is that the relevant NV transition frequency is in the microwave range. Direct arbitrary waveform generation at a carrier frequency near 2.87 GHz requires instrumentation with very high sampling rate and large analog bandwidth, which is expensive and often unnecessary if the desired information bandwidth of the pulse is much lower than the carrier frequency itself. A practical alternative is to generate the waveform at a lower intermediate frequency and then translate it to the microwave domain through up-conversion. In this architecture, the AWG defines the programmable intermediate-frequency waveform, while a local oscillator and an IQ mixer shift the signal to the resonance frequency required for NV spin control.

This architecture separates two functions that are otherwise difficult to combine in a single instrument. The AWG provides waveform flexibility: programmable envelopes,

phase control, and quadrature modulation. The up-conversion chain provides access to the gigahertz carrier frequency. The combination therefore enables shaped microwave pulses without requiring a direct multi-gigahertz arbitrary waveform source.

A key point is that up-conversion shifts the carrier frequency but does not increase the bandwidth of the waveform. The mixer translates the spectrum of the intermediate-frequency signal around the local oscillator frequency, but the temporal complexity of the pulse remains limited by the AWG sampling rate, the AWG analog bandwidth, and the bandwidth of the RF chain. Therefore, up-conversion solves the frequency-translation problem, while the achievable pulse shapes remain bounded by the baseband and intermediate-frequency hardware.

Conventional microwave control based on a gated continuous-wave source is limited mainly by its lack of waveform programmability. The microwave carrier is generated externally and then switched on and off to define the pulse duration. The resulting envelope is rectangular, and the phase is not freely assigned for each pulse in the same way as in an IQ-controlled AWG system. This makes the method simple, but it also limits the class of experiments that can be implemented.

The first limitation is spectral. Abrupt switching produces a broad frequency response, which can reduce selectivity when the system contains nearby transitions or inhomogeneous broadening. For NV ensembles or samples with multiple resonances, this lack of selectivity can become relevant. A shaped pulse provides a way to control the spectral content more carefully, whereas a rectangular pulse is constrained by its sharp temporal edges.

The second limitation is robustness. In a real experiment, the microwave amplitude, detuning, and timing are never perfectly ideal. A pulse sequence that works only under exact conditions may become inefficient when the experimental parameters drift. Shaped and phase-controlled pulses can be designed to reduce sensitivity to such imperfections, which is one of the main reasons they are widely used in magnetic resonance and quantum control [7, 8].

The third limitation is compatibility with advanced pulsed protocols. Many experiments benefit from controlled phase shifts, quadrature modulation, or repeated pulses with different phases. These operations are difficult to implement cleanly when the microwave source is only externally gated. An AWG, especially when used through an IQ up-conversion chain, makes these operations natural because the phase and quadrature components are defined directly in the generated waveform.

The up-conversion architecture overcomes the inability of the AWG to directly reach the NV carrier frequency, but it introduces its own practical constraints. The IQ mixer must be calibrated to suppress local oscillator leakage and unwanted sidebands. Any imbalance

between the in-phase and quadrature channels can distort the output spectrum. Moreover, the final microwave waveform is limited by the bandwidth and linearity of the AWG, mixer, amplifiers, cables, and antenna. These constraints must be considered when interpreting experimental results and when comparing different pulse shapes.

The HDAWG provides the digital waveform flexibility, the up-conversion chain enables access to the GHz carrier frequency, and the calibration procedure determines how accurately the desired waveform is delivered to the NV center. The performance of the final system depends on the interaction of all these elements. This complete experimental framework was implemented at the Quantum Optics Laboratory of INRiM and is presented and experimentally validated in this thesis.

The remainder of this thesis is organized as follows.

Chapter 2 introduces the theoretical foundations required to understand NV-center quantum sensing. The electronic structure of the nitrogen-vacancy center, the principles of optically detected magnetic resonance, and the fundamentals of coherent spin control are presented.

Chapter 3 describes the experimental setup employed throughout this work, including the optical excitation and detection system, the microwave instrumentation, the magnetic bias field, and the synchronization between the different components of the setup.

Chapter 4 presents the development of the HDAWG integration within the Qudi framework. The architecture of the implemented hardware module is described together with the communication through the Zurich Instruments Python API, waveform generation procedures, sequence compilation, implementation of Gaussian and DRAG pulse envelopes, and the validation of the developed framework.

Chapter 5 focuses on the microwave up-conversion chain developed for pulsed measurements. The hardware architecture, IQ modulation scheme, and phase and amplitude control procedures are discussed.

Chapter 6 reports the experimental results obtained using the developed platform. The implementation and validation of Gaussian and DRAG pulse sequences are presented, together with oscilloscope measurements and comparisons between different pulse shapes.

Finally, Chapter 7 summarizes the main contributions of this work, discusses its limitations, and outlines possible directions for future developments.

2 Theoretical Background

2.1 Color Centers in Diamond

Diamond has a face-centered cubic crystal structure that can be described as two interpenetrating FCC sublattices displaced along the body diagonal. Each carbon atom forms four covalent bonds with its nearest neighbours in a tetrahedral sp^3 configuration, as shown in Fig. 2.1. The ideal lattice is highly symmetric and has a wide electronic bandgap of approximately 5.5 eV. Point defects or impurity atoms can introduce localized electronic states inside this bandgap. When these defect states allow optical absorption and emission, the defect is referred to as a color center.




Figure 2.1: Diamond crystal structure. (A) Conventional unit cell of diamond, showing the face-centered cubic structure with a lattice constant of approximately $a_0 = 3.57 \text{ \AA}$. (B) Tetrahedral bonding arrangement in diamond, where each carbon atom is bonded to four nearest neighbours in an sp^3 configuration. Adapted from [10].

From a quantum-optical point of view, a color center can be treated as a localized electronic system coupled to the surrounding crystal lattice. The optical and spin properties of the center depend on the symmetry of the defect, its charge state, the local strain environment, and the interaction with nearby nuclear spins. These microscopic param-

eters determine the transition energies, selection rules, fluorescence spectrum, and spin Hamiltonian of the defect.

Several optically active defects exist in diamond, including nitrogen-related, silicon-related, and germanium-related.

2.1.1 Nitrogen-Vacancy Centers and Charge States

The nitrogen-vacancy center is one of the most studied color centers in diamond. It is formed by a substitutional nitrogen atom located next to a vacancy in the carbon lattice. The nitrogen atom replaces one carbon atom, while the vacancy corresponds to a missing nearest-neighbour carbon atom. This pair defines a defect axis oriented along one of the four possible crystallographic directions of the diamond lattice.




Figure 2.2: Atomic structure of the NV defect in diamond. The defect consists of a substitutional nitrogen atom adjacent to a vacancy in the carbon lattice. Adapted from [3].

The NV center can exist in different charge states, mainly the neutral state NV^0 and the negatively charged state NV^- . These charge states have different optical and spin properties. The neutral state NV^0 has its zero-phonon line near 575 nm, while the negatively charged state NV^- has a zero-phonon line near 637 nm, as depicted in Fig. 2.3.

For ODMR-based quantum sensing, the relevant charge state is NV^- because it possesses an electronic spin-triplet ground state that can be initialized optically, manipulated by microwave fields, and read out through spin-dependent photoluminescence [1, 2].

The charge-state stability of the defect is important in experiments. Optical excitation, nearby impurities, surface effects, and the local electrostatic environment can influence whether the center remains in NV^- or switches to another charge state. This is especially relevant for near-surface or nanostructured diamond samples, where surface-related effects can modify the local electronic environment. Maintaining the NV^- charge state is therefore




Figure 2.3: Normalised emission spectra of an ensemble of NV^- and NV^0 centres at 10 K for different excitation (532 nm) powers Adapted from [1].

important because the ODMR contrast and the fluorescence-based spin readout rely on the optical and spin properties of this charge state.

The most important point is that NV^- combines a microwave-addressable spin transition with optical initialization and readout. This combination is what allows the defect to be used as a solid-state spin qubit for pulsed ODMR and quantum sensing experiments.

2.1.2 Electronic Structure of the NV Center

The electronic structure of the NV center determines why the negatively charged state NV^- is suitable for ODMR and microwave-based spin control. While both NV^0 and NV^- are optically active, their electronic level structures are different, as shown in Fig. 2.4. The NV^- center has a triplet ground state, commonly denoted as $^3A_2$, and an optically excited triplet state denoted as ^3E . In addition to these triplet states, intermediate singlet states, usually labelled ^1E and $^1A_1$, participate in the non-radiative decay pathways that make optical spin polarization and spin-dependent readout possible [1, 4].

The ground state of NV^- has total spin $S = 1$, giving three spin projections: $m_s = 0$ and $m_s = \pm 1$. Even without an externally applied magnetic field, these spin sublevels are not degenerate. The $m_s = 0$ state is separated from the $m_s = \pm 1$ states by the zero-field splitting parameter $D \approx 2.87 \text{ GHz}$ ¹.

This energy splitting lies in the microwave frequency range. Therefore, microwave

¹In this thesis the energy couplings are expressed in Hertz. The value in Joule can be derived considering the Planck Law $E = h\nu$ with h expressed in J/Hz.




Figure 2.4: Simplified electronic structure of the NV⁻ and NV⁰ charge states. The negatively charged NV⁻ center has a triplet ground state $^3A_2$ and an excited triplet state ^3E , together with intermediate singlet states involved in spin-dependent relaxation. The neutral NV⁰ center has a different electronic structure and is not the charge state used for ODMR-based spin control in this work. Adapted from [1].

radiation resonant with this transition can drive coherent rotations between the $m_s = 0$ and $m_s = \pm 1$ spin states. This is the physical basis of microwave spin manipulation in NV-center experiments.

The simplified ground-state Hamiltonian, neglecting strain and hyperfine interactions, can be written as

$$H = DS_z^2 + \gamma_e \mathbf{B} \cdot \mathbf{S}, \quad (2.1)$$

where D is the zero-field splitting, γ_e is the electron gyromagnetic ratio, $\mathbf{B}$ is the applied magnetic field, and $\mathbf{S} = (S_x, S_y, S_z)$ are the spin-1 operators. The first term describes the intrinsic splitting between the $m_s = 0$ and $m_s = \pm 1$ levels, while the second term describes the Zeeman interaction with an external magnetic field. This Hamiltonian provides the starting point for understanding ODMR spectra, magnetic-field-dependent resonance shifts, and coherent microwave control of the NV spin [1, 2, 3].

2.1.3 Optical Cycle and Spin-Dependent Readout

The defect can be excited optically, usually with a green laser at 532 nm, and it emits red photoluminescence when it relaxes back toward the ground state. This optical response

makes it possible to detect the NV center using a confocal microscope and to monitor changes in its spin state through variations in the emitted fluorescence intensity.

The optical cycle involves the triplet ground state $^3A_2$ and the triplet excited state ^3E . Under green excitation, population is promoted from the ground state to the excited state. The system can then decay radiatively back to the ground state, emitting red photoluminescence. The zero-phonon line of the NV⁻ center is located near 637 nm, while a broad phonon sideband extends to longer wavelengths. In experiments, the collected fluorescence is therefore usually detected over a broad red spectral range rather than only at the zero-phonon line.

A crucial feature of the NV⁻ center is that the optical cycle is spin-dependent. The $m_s = 0$ and $m_s = \pm 1$ spin projections do not relax in exactly the same way after optical excitation. Population in the $m_s = 0$ state mainly follows a radiative pathway and therefore produces stronger fluorescence. In contrast, population in the $m_s = \pm 1$ states has a higher probability of relaxing through a non-radiative intersystem crossing involving intermediate singlet states. This pathway reduces the detected fluorescence and tends to return population preferentially to the $m_s = 0$ ground-state sublevel [4, 1].




Figure 2.5: Energy-level scheme of the NV⁻ center. Green optical excitation at 532 nm promotes population from the triplet ground state $^3A_2$ to the triplet excited state ^3E . Radiative decay produces red photoluminescence, while spin-dependent intersystem crossing through singlet states enables optical spin polarization and fluorescence-based spin read-out. Adapted from [4].

This spin-dependent relaxation is the physical mechanism behind optical spin initialization. When the NV⁻ center is illuminated repeatedly, population is gradually pumped into the $m_s = 0$ state. Therefore, a laser pulse does not only excite the defect optically; it

also prepares the spin in a well-defined initial state. This property is extremely important in pulsed ODMR experiments, where each measurement sequence typically begins with a laser pulse used to initialize the spin before microwave manipulation.

The same mechanism also enables optical spin readout. Since the $m_s = 0$ state is brighter than the $m_s = \pm 1$ states, the detected photoluminescence intensity contains information about the spin population. If a microwave pulse transfers population from $m_s = 0$ to $m_s = \pm 1$, the fluorescence decreases. By measuring this fluorescence contrast while applying microwave radiation, one can detect the spin resonance optically. This is the principle of optically detected magnetic resonance.

In a typical pulsed measurement, the sequence can be understood in three steps. First, a laser pulse initializes the NV spin mainly into $m_s = 0$. Second, a microwave pulse or pulse sequence manipulates the spin state. Third, a second laser pulse reads out the final spin population through the emitted photoluminescence. The measured signal is therefore not a direct electrical measurement of the spin, but an optical measurement whose contrast originates from the spin-dependent decay pathways of the NV^- center.

This optical interface is one of the main reasons why the NV center is so useful experimentally. It allows a spin state in a solid-state defect to be prepared and measured using light, while coherent control is performed with microwave pulses.

2.1.4 The NV Center as an Atomic-Scale Magnetic Sensor

The magnetic sensing capability of the NV^- center originates from the dependence of its electron spin resonance frequencies on the local magnetic field. When a magnetic field is applied, the energies of the spin sublevels change through the Zeeman interaction. As a result, the microwave resonance frequencies become a direct probe of the magnetic field projected along the NV axis [3, 2].

The NV axis is defined by the direction joining the nitrogen atom and the vacancy. This axis plays the role of the quantization axis of the defect. If the applied magnetic field is denoted by $\mathbf{B}$, it can be decomposed into two components: a component parallel to the NV axis and a transverse component perpendicular to it. The parallel component is usually the most important one for magnetometry because it shifts the two spin transitions in opposite directions.

A useful form of the ground-state Hamiltonian is

$$\mathcal{H} = hDS_z^2 + hE(S_x^2 - S_y^2) + g\mu_B\mathbf{B} \cdot \mathbf{S}, \quad (2.2)$$

where h is Planck's constant, D is the axial zero-field splitting, E is the transverse zero-field splitting parameter, $g \approx 2$ is the electron Landé factor, μ_B is the Bohr magneton, and $\mathbf{S} = (S_x, S_y, S_z)$ are the spin-1 operators. The first term describes the splitting between $m_s = 0$ and $m_s = \pm 1$ in zero magnetic field. The second term accounts for transverse splitting caused mainly by local strain or electric-field effects. The last term is the Zeeman interaction between the NV electronic spin and the applied magnetic field [3].

The parameter D is approximately 2.87 GHz at room temperature and is the main frequency scale of the NV ground state. It originates from spin-spin interaction between the two unpaired electrons of the defect. The parameter E is usually much smaller than D and reflects the fact that local strain or electric fields can slightly break the ideal symmetry of the defect. In high-purity bulk diamond, E can be very small, while in nanodiamonds or nanostructured samples it can reach values of the order of a few megahertz because local strain is stronger [3].

If the magnetic field is mainly aligned with the NV axis, the situation is simple. The $m_s = +1$ and $m_s = -1$ states shift in opposite directions, and the two electron spin resonance frequencies can be approximated as

$$\nu_{\pm} = D \pm \frac{g\mu_B}{h} B_{\text{NV}}, \quad (2.3)$$

where B_{NV} is the projection of the magnetic field along the NV axis. This equation expresses the basic operating principle of NV magnetometry: a magnetic field changes the resonance frequencies, and the field projection can be inferred by measuring the frequency shift.

The term electron spin resonance (ESR) refers to the microwave transitions between the $m_s = 0$ state and the $m_s = \pm 1$ states. In conventional magnetic resonance, such transitions are detected directly through microwave absorption. In NV-center experiments, they are usually detected optically: when the microwave frequency becomes resonant with one of the spin transitions, population is transferred from the brighter $m_s = 0$ state to one of the darker $m_s = \pm 1$ states, and the photoluminescence decreases due to the underlying resonances which are ESR transitions.

Figure 2.6 illustrates this principle. At zero magnetic field, the two transitions are close to the zero-field splitting. As the magnetic field increases, the degeneracy between $m_s = +1$ and $m_s = -1$ is lifted and the two resonance dips separate. The distance between the dips is proportional to the magnetic-field projection along the NV axis in the linear regime. Therefore, the ODMR spectrum provides a direct optical measurement of the local magnetic field.




Figure 2.6: Optically detected ESR spectra of an NV center for different applied magnetic fields. The resonance dips correspond to microwave transitions from $m_s = 0$ to $m_s = \pm 1$. As the magnetic field increases, the Zeeman interaction separates the two resonance frequencies. Adapted from [3].

The behaviour is more subtle when the magnetic field is weak or not perfectly aligned with the NV axis. In the weak-field regime, the transverse splitting parameter E cannot always be neglected. If the transverse magnetic-field component is small and the field projection along the NV axis is denoted by B_{NV} , the ESR frequencies can be written approximately as

$$\nu_{\pm} = D \pm \sqrt{\left(\frac{g\mu_B}{h} B_{\text{NV}}\right)^2 + E^2}. \quad (2.4)$$

This expression shows that, close to zero magnetic field, the ESR frequencies are influenced not only by the magnetic field but also by the transverse splitting E [11]. When the magnetic field is very small, the ESR frequencies may vary quadratically with the field instead of linearly. In that case, the NV spin becomes less sensitive to first-order magnetic-field fluctuations along the NV axis. This behaviour is useful in some contexts because it can protect the spin from magnetic noise, but it is not ideal when the goal is high-sensitivity magnetometry[12].

To recover a linear magnetic response, a bias magnetic field can be applied along the NV axis.

$$\frac{g\mu_B}{h}B_{\text{NV}} \gg E, \quad (2.5)$$

When the effect of E becomes relatively small and Eq. 2.4 reduces approximately to the linear Zeeman relation in Eq. 2.3. In this regime, the ESR frequencies shift almost linearly with the magnetic-field projection, which simplifies the conversion between frequency shift and magnetic field. This is the regime most commonly used for straightforward NV magnetometry [3].

For stronger magnetic fields, especially when the field is not aligned with the NV axis, the transverse component of the magnetic field can no longer be ignored. The spin quantization axis is then no longer strictly fixed by the NV defect axis. The magnetic field mixes the $m_s = 0$ and $m_s = \pm 1$ spin states, and the resonance frequencies depend not only on the field magnitude but also on the angle between the magnetic field and the NV axis. This makes the spectrum more complex than in the simple linear case.

The consequence of this mixing is important experimentally. When the spin states are strongly mixed, optical spin polarization and spin-dependent readout become less efficient. Since ODMR contrast relies on the difference in fluorescence between $m_s = 0$ and $m_s = \pm 1$, a reduction in spin selectivity reduces the visibility of the ESR resonances. Therefore, strong off-axis magnetic fields can make NV magnetometry less efficient, especially in nanodiamonds or nanostructured samples where local magnetic fields and strain can be significant [3].

The magnetic sensitivity of an NV measurement is determined by how precisely a shift in resonance frequency can be measured. In continuous-wave ODMR, the main parameters are the resonance linewidth, the fluorescence contrast, and the detected photon rate. Narrow resonances and high contrast allow smaller frequency shifts to be resolved. However, increasing microwave power can broaden the resonance through power broadening, which can reduce the magnetic-field sensitivity even if the signal becomes stronger [13, 2].

This connection between magnetic sensitivity and microwave control is central for the present work. The microwave signal must be placed at the correct ESR frequency, must be spectrally selective enough to address the desired transition, and must have a controlled temporal shape to avoid unnecessary broadening or unwanted excitation. The Zeeman interaction therefore does not only explain how the NV center senses magnetic fields; it also defines the microwave-control requirements that motivate the use of shaped pulses and an up-conversion architecture.




Figure 2.7: Magnetic-field dependence of the NV^- electron spin resonance frequencies. (a) Definition of the angle θ between the applied magnetic field $\mathbf{B}$ and the NV quantization axis z . (b) Calculated ESR transition frequencies $\nu_{\pm}$ as a function of the magnetic-field amplitude for different field orientations. The curves are obtained from the full ground-state spin Hamiltonian using $D = 2.87$ GHz and $E = 5$ MHz. (c,d) Zooms of the weak-field regime, where the strain-induced transverse splitting E significantly affects the resonance frequencies; the dotted curves in (c) show the weak-field analytical approximation. (e) Example of a zero-field ESR spectrum of a single NV center in a diamond nanocrystal, showing a splitting of approximately $2E$. Adapted from [3].

2.2 ODMR Techniques

2.2.1 Continuous wave mode

In the continuous-wave implementation of ODMR, optical excitation and microwave driving are applied simultaneously. The NV center is continuously illuminated, typically with a green laser, while the microwave frequency is swept across the spin resonance. The detected photoluminescence is then recorded as a function of microwave frequency. When the microwave field is off resonance, the optical cycle keeps the NV population mainly in the brighter $m_s = 0$ state. When the microwave frequency matches one of the $m_s = 0 \rightarrow m_s = \pm 1$ transitions, population is transferred to the darker spin states and the fluorescence decreases.




Figure 2.8: Example of a continuous-wave ODMR measurement. The laser and microwave fields are applied simultaneously while the microwave frequency is swept. Resonant microwave excitation transfers population from the brighter $m_s = 0$ state to the darker $m_s = \pm 1$ states, producing dips in the detected photoluminescence. Adapted from [4].

The resulting ODMR spectrum consists of one or more fluorescence dips. In zero magnetic field, the resonance is located near the zero-field splitting $D \approx 2.87$ GHz. When a magnetic field is applied, the Zeeman interaction separates the $m_s = +1$ and $m_s = -1$ transitions, producing two resonance dips. The position of these dips gives information about the local magnetic-field projection along the NV axis, while their linewidth and contrast determine how accurately the resonance frequency can be extracted [2, 3].

The ODMR contrast is usually defined as the relative decrease in photoluminescence at resonance compared with the off-resonant fluorescence level. A larger contrast improves

the visibility of the resonance. The linewidth describes the spectral width of the resonance dip and is directly related to the precision with which the resonance frequency can be determined. For magnetic sensing, the ideal case is therefore a narrow resonance with high contrast and a high photon collection rate.

However, continuous-wave ODMR involves a compromise between signal strength and linewidth. Increasing the microwave power can increase the population transfer between spin states and improve the observed contrast, but it can also broaden the resonance through power broadening. A broader resonance reduces the precision with which the resonance frequency can be determined and can therefore degrade magnetic-field sensitivity [13]. Similarly, increasing the optical power modifies the balance between initialization, fluorescence collection, and relaxation dynamics. For this reason, CW ODMR is useful and experimentally simple, but it is not always the best configuration for high-sensitivity or coherent-control measurements.

In the context of this work, CW ODMR is mainly used as a characterization tool. It allows the resonance frequency, contrast, linewidth, and magnetic field splitting of the NV center to be identified before implementing pulsed measurements. It therefore provides the frequency reference around which the microwave pulses generated by the AWG and up-conversion chain must be tuned.

2.2.2 Pulse sequences (Pulsed ODMR)

Pulsed ODMR separates the different physical operations of an NV experiment in time. Instead of applying laser and microwave fields continuously, the sequence is divided into distinct steps: optical initialization, microwave manipulation, and optical readout. This separation makes it possible to reduce some of the limitations of continuous-wave ODMR and to study the coherent spin dynamics of the NV center more directly [2, 4].

A basic pulsed ODMR sequence begins with a laser pulse that initializes the NV spin mainly into the $m_s = 0$ state. After the initialization pulse, the laser is turned off and a microwave pulse is applied. The microwave frequency is scanned around the expected spin resonance. If the microwave pulse is resonant with the $m_s = 0 \rightarrow m_s = \pm 1$ transition, it transfers part of the population into the darker spin state. A second laser pulse is then used to read out the final spin population through the emitted photoluminescence. As in continuous-wave ODMR, the resonance appears as a decrease in the detected fluorescence, but the initialization, manipulation, and readout steps are now temporally separated.

This time separation is important because it allows microwave control to be applied without simultaneous optical excitation. During the microwave pulse, the spin evolves coherently under the applied microwave field, rather than being continuously disturbed by

optical pumping. This makes pulsed ODMR more suitable for calibrating coherent control parameters, such as the resonance frequency, π -pulse duration, and Rabi frequency. It is also the basis for more advanced pulsed measurements, including Rabi oscillations, Ramsey interferometry, Hahn echo, and dynamical decoupling sequences [2, 3].

In a pulsed ODMR measurement, the microwave pulse duration is usually fixed while the microwave frequency is swept. When the frequency is resonant, the pulse rotates the spin away from $m_s = 0$, reducing the fluorescence during the readout pulse. The width and contrast of the pulsed ODMR resonance depend on the microwave pulse duration, the pulse shape, decoherence, and the spectral response of the spin transition. Longer pulses generally provide narrower spectral selectivity, while shorter pulses have broader frequency content. This relationship between pulse duration and spectral bandwidth is one reason why pulse shaping becomes important [13].

Compared with continuous-wave ODMR, pulsed ODMR provides more direct access to coherent spin manipulation. It also creates a natural framework for using arbitrary waveform generation. In a conventional setup, the microwave pulse is often rectangular because it is produced by switching a continuous microwave source on and off. With an AWG-based control chain, the microwave pulse can instead be given a programmable envelope and phase. This makes it possible to compare rectangular, Gaussian, and DRAG-like pulses (Derivative Removal by Adiabatic Gate) to study how the temporal shape of the pulse affects spectral selectivity and spin-control performance.

2.3 Coherent Control of the NV Spin

After optical initialization and readout have been introduced, the next key ingredient is coherent microwave control. In the negatively charged NV center, the spin sublevels of the ground state can be driven by microwave radiation near the resonance frequency. By controlling the duration, amplitude, phase, and frequency of the microwave pulse, it is possible to implement controlled rotations of the spin state. These rotations form the basis of pulsed ODMR, Rabi measurements, Ramsey interferometry, Hahn echo sequences, and more advanced sensing protocols.

2.3.1 Two-Level Approximation

The ground state of the NV^- center is a spin-triplet system with three spin projections, $m_s = 0$ and $m_s = \pm 1$. In the most general case, the spin dynamics may involve all three sublevels. However, in many experiments one microwave transition is selected and treated as an effective two-level system. For example, the experiment may address only the transition between $m_s = 0$ and $m_s = -1$, while the $m_s = +1$ state remains far enough

in frequency to be neglected. In this approximation, one defines an effective qubit basis as

$$|0\rangle \equiv |m_s = 0\rangle, \quad |1\rangle \equiv |m_s = -1\rangle, \quad (2.6)$$

or alternatively $|1\rangle \equiv |m_s = +1\rangle$, depending on which transition is driven. The microwave field then couples these two states and produces coherent oscillations of the population between them. This simplified description is extremely useful because it allows the NV spin to be treated using the standard language of qubit control and Bloch-sphere rotations [2, 7]. In the Bloch-sphere representation, the states $|0\rangle$ and $|1\rangle$ are located at opposite poles of the sphere. A resonant microwave pulse rotates the spin state around an axis in the equatorial plane. The angle of rotation depends on the pulse duration and amplitude, while the axis of rotation is set by the phase of the microwave field. Therefore, phase control is not a minor experimental detail: it determines the direction of the rotation in the effective qubit space. The two-level approximation is valid only when the chosen microwave pulse is spectrally selective. If the pulse bandwidth is too broad, or if the two transitions are too close in frequency, the microwave field may also drive unwanted transitions. This can happen at low magnetic field, where the two $m_s = 0 \rightarrow m_s = \pm 1$ transitions may be close, or when very short pulses are used. Since shorter pulses have broader spectral content, there is a direct connection between pulse duration, pulse shape, and the validity of the two-level approximation. This point is important for the present work. The goal of shaped microwave control is not only to create pulses with a desired duration, but also to control their spectral content and phase. A well-designed pulse can address the target transition more selectively and reduce unwanted excitation. This is one of the reasons why an arbitrary waveform generator is useful for NV-center experiments.

2.3.2 Basic Pulsed Control Sequences

Pulsed ODMR experiments are constructed from a small number of elementary operations: optical initialization, microwave manipulation, free spin evolution, and optical readout. By arranging these operations in different orders, different physical quantities can be measured. The most common sequences include Rabi oscillations, Ramsey interferometry, and Hahn echo. Each sequence has a different role: Rabi measurements calibrate microwave rotations, Ramsey measurements probe dephasing and frequency detuning, and Hahn echo measurements refocus slowly varying noise to access longer coherence times.




Figure 2.9: Common pulsed ODMR sequences for NV-center spin control. (a) Relaxometry sequence for measuring longitudinal relaxation. (b) Rabi sequence, where the microwave pulse duration is varied to calibrate coherent rotations. (c) Ramsey sequence, where two $\pi/2$ pulses are separated by a free evolution time. (d) Hahn echo sequence, where a refocusing π pulse is inserted between two $\pi/2$ pulses to reduce the effect of slowly varying noise. Adapted from [4].

Relaxometry

The simplest pulsed ODMR protocol is the relaxometry sequence, shown in Fig. 2.9(a). In this measurement, a first laser pulse initializes the NV^- spin mainly into the bright $m_s = 0$ state. The system is then left to evolve freely for a variable delay time τ , without applying a microwave pulse. A second laser pulse is used to read out the remaining spin population through the emitted photoluminescence.

The purpose of relaxometry is to measure the longitudinal relaxation time, usually denoted as T_1 . This time describes how the spin population returns toward thermal equilibrium after optical initialization. In other words, while coherent protocols such as Rabi or Ramsey measurements study controlled spin rotations and phase evolution, relaxometry probes the natural population relaxation of the NV spin. The measured fluorescence signal typically decays as the delay time τ is increased, reflecting the loss of the initialized spin polarization [4, 2].

Relaxometry is important because it provides a basic timescale for the spin system. The value of T_1 sets an upper limit on how long spin population information can be preserved. It is also sensitive to magnetic noise at frequencies close to the NV spin transition, which makes T_1 measurements useful for sensing fluctuating magnetic environments. However, relaxometry does not require coherent microwave rotations and therefore does not calibrate the microwave control field.

This is where the Rabi sequence becomes the natural next step. After the population lifetime has been characterized, a microwave pulse can be inserted between the initialization and readout laser pulses, as shown in Fig. 2.9(b). By varying the duration of this microwave pulse, the experiment moves from observing passive relaxation to actively driving coherent transitions between the spin states. Rabi oscillations therefore provide the first direct calibration of microwave control, while relaxometry provides the reference timescale for spin population decay.

Rabi Oscillations

Rabi oscillations are the basic experiment used to verify and calibrate coherent microwave driving. The sequence begins with a laser pulse that initializes the NV spin mainly into the bright $m_s = 0$ state. A microwave pulse is then applied at a frequency close to the selected spin transition. The duration of this microwave pulse, usually denoted by τ , is varied. Finally, a readout laser pulse measures the final spin population through spin-dependent photoluminescence, as explained in Fig. 2.9(b). If the microwave field is resonant with the selected transition, the spin population oscillates between the two states of the effective qubit. In the ideal resonant case, the probability of remaining in the initial state $|0\rangle$ after

a pulse of duration τ can be written as

$$P_0(\tau) = \cos^2\left(\frac{\Omega_R\tau}{2}\right), \quad (2.7)$$

while the probability of being transferred to the state $|1\rangle$ is

$$P_1(\tau) = \sin^2\left(\frac{\Omega_R\tau}{2}\right), \quad (2.8)$$

where Ω_R is the Rabi angular frequency. Since $|0\rangle$ corresponds to the brighter spin state, these population oscillations appear as oscillations in the detected photoluminescence. The Rabi frequency depends on the amplitude of the microwave magnetic field at the NV center. A stronger microwave field produces faster Rabi oscillations, while a weaker field produces slower oscillations. For this reason, Rabi measurements are used to calibrate the duration of microwave pulses. A pulse that transfers population from $|0\rangle$ to $|1\rangle$ is called a π pulse, while a pulse that creates an equal superposition is called a $\pi/2$ pulse. If the Rabi oscillation period is T_R , the ideal pulse durations are

$$t_\pi = \frac{T_R}{2}, \quad t_{\pi/2} = \frac{T_R}{4}. \quad (2.9)$$

These calibrations are essential before performing more complex sequences such as Ramsey and Hahn echo, because those protocols require accurate $\pi/2$ and π rotations. In real experiments, the observed Rabi oscillations can lose contrast due to several mechanisms, including microwave amplitude inhomogeneity, finite coherence time, and variations in the local environment of different NV centers in an ensemble. Detuning from resonance does not by itself cause damping; instead, it reduces the oscillation contrast and increases the effective Rabi frequency [4, 2].

Ramsey Interferometry

The Ramsey sequence shown in Fig. 2.9(c) is used to measure dephasing and to detect small shifts of the spin resonance frequency. It consists of two $\pi/2$ microwave pulses separated by a free evolution time τ . The first $\pi/2$ pulse transforms the initialized state into a coherent superposition of the two effective qubit states. During the free evolution time, the relative phase between the two states evolves. The second $\pi/2$ pulse converts this accumulated phase into a population difference that can be detected optically. The Ramsey signal depends on the detuning between the applied microwave frequency and the actual spin transition frequency. If the microwave drive is slightly detuned, the relative

phase accumulated during the free evolution time produces oscillations in the final fluorescence signal. These oscillations are called Ramsey fringes. Their frequency is related to the detuning, while their decay envelope reflects the inhomogeneous dephasing time, usually denoted by T_2^* . The Ramsey sequence is powerful because the spin evolves freely between the two microwave pulses. During this interval, the system is sensitive to frequency shifts caused by magnetic fields, strain, temperature, or other local perturbations. This makes Ramsey interferometry useful for DC sensing and for measuring small changes in the transition frequency. Compared with continuous-wave ODMR, the separation between microwave manipulation and optical readout helps avoid some limitations associated with simultaneous optical and microwave excitation, including power broadening [4, 2]. However, the Ramsey signal is sensitive to slowly varying noise and inhomogeneities. Different NV centers, or even the same NV center under a fluctuating environment, may accumulate phase at slightly different rates. This causes the ensemble-averaged signal to decay with a characteristic time T_2^* . Therefore, Ramsey measurements are excellent for probing dephasing and static or slowly varying frequency shifts, but they are not optimal when the goal is to preserve coherence for long times in the presence of low-frequency noise. In the context of microwave control, Ramsey interferometry also highlights the importance of phase stability. The two $\pi/2$ pulses must have a well-defined relative phase, because the measured signal depends on the phase accumulated between them. This requirement is directly connected to the use of programmable waveform generation: an AWG-based system can define the phase of each pulse precisely, which is essential for reliable Ramsey and phase-sensitive sensing protocols.

Hahn Echo

The Hahn echo sequence extends the Ramsey idea by adding a refocusing π pulse in the middle of the free evolution time as shown in Fig. 2.9(d). The sequence is usually written as

$$\frac{\pi}{2} - \tau - \pi - \tau - \frac{\pi}{2}. \quad (2.10)$$

The first $\pi/2$ pulse creates a coherent superposition. During the first free evolution period, the spin accumulates phase due to detuning or local field fluctuations. The central π pulse swaps the two components of the superposition, effectively reversing the phase accumulation caused by quasi-static noise. During the second free evolution period, the phase acquired before the π pulse is partially cancelled. The final $\pi/2$ pulse converts the remaining coherence into a measurable population difference. The Hahn echo sequence is therefore less sensitive to slowly varying dephasing noise than the Ramsey sequence. The decay time extracted from Hahn echo measurements is called T_2 , the spin coherence time. In many systems, T_2 can be significantly longer than T_2^* because the echo pulse removes

the effect of static or slowly fluctuating frequency offsets. This is why Hahn echo is a central tool for characterizing coherence in NV centers and other spin systems [4, 2, 3]. The echo sequence is also useful for sensing alternating magnetic fields. The refocusing pulse makes the sequence sensitive to fields that change sign or oscillate in synchrony with the pulse timing, while rejecting slowly varying noise. More advanced dynamical decoupling sequences extend this principle by using many refocusing pulses to filter specific frequency components of the environment. From the perspective of pulse engineering, Hahn echo is more demanding than a simple Rabi measurement because it requires accurate $\pi/2$ and π pulses, stable phase relationships, and precise timing. Any pulse error can reduce the refocusing efficiency and lower the measured contrast. Therefore, the performance of echo-based protocols depends directly on the quality of the microwave pulse generation system. This provides another motivation for using shaped and phase-controlled pulses in NV-center experiments.

Panel (d) of Fig. 2.9 shows a typical Hahn echo measurement. The horizontal axis is the total free evolution time 2τ , because the spin evolves for a time τ before the refocusing π pulse and for another time τ after it. The measured photoluminescence oscillates and decays as 2τ is increased. The decay envelope reflects the loss of spin coherence and is used to extract the Hahn echo coherence time T_2 .

The revivals visible in the signal arise from the interaction between the NV electronic spin and the surrounding nuclear-spin environment. In diamond, nearby nuclear spins, especially ^{13}C nuclei, can modulate the coherence of the NV spin. At specific evolution times, part of the phase accumulated from the environment is partially rephased, producing echo revivals. The overall decay of the envelope still indicates that the spin coherence is gradually lost, but the revival structure reveals that the environment is not simply a featureless noise bath.

This behaviour shows why the Hahn echo sequence is useful. Compared with a Ramsey experiment, the central π pulse cancels phase errors caused by slowly varying noise, allowing the coherence to survive for longer times. However, the echo signal is still affected by faster fluctuations and by the coupling to nearby nuclear spins. For this reason, the Hahn echo curve gives information not only about the intrinsic coherence time T_2 , but also about the magnetic environment surrounding the NV center.

2.3.3 Microwave Phase and Rotation Axis

In the effective two-level approximation, the selected NV spin transition can be treated as a qubit driven by a microwave field. The role of the microwave field is not only to transfer population between the two spin states, but also to define the axis around which the state

rotates on the Bloch sphere. This is why microwave phase control is essential in coherent spin manipulation.

Consider an effective basis formed by $|0\rangle \equiv |m_s = 0\rangle$ and $|1\rangle \equiv |m_s = -1\rangle$, or equivalently by the transition to $|m_s = +1\rangle$ depending on the resonance being addressed. In a frame rotating at the microwave drive frequency and under the rotating-wave approximation, the driven two-level Hamiltonian can be written as

$$H_{\text{rot}} = \frac{\hbar}{2} [\Delta \sigma_z + \Omega(t) (\cos \phi \sigma_x + \sin \phi \sigma_y)], \quad (2.11)$$

where Δ is the detuning between the microwave frequency and the spin transition frequency, $\Omega(t)$ is the time-dependent Rabi frequency, ϕ is the microwave phase, and σ_x , σ_y , and σ_z are the Pauli matrices acting on the effective two-level system. This expression shows that the microwave amplitude controls the rotation rate, while the microwave phase controls the direction of the rotation axis in the transverse plane of the Bloch sphere. When the microwave drive is detuned from resonance ($\Delta \neq 0$), the rotation axis is no longer confined to the equatorial plane of the Bloch sphere. The detuning term $\Delta \sigma_z$ introduces a longitudinal component along the z axis, causing the effective rotation axis to tilt away from the transverse plane. As a result, the spin evolves around an axis with both transverse and longitudinal components, rather than undergoing a pure rotation around the x or y axis. The larger the detuning, the larger the contribution of the z component to the effective rotation axis.[\[2, 7\]](#).

On resonance, where $\Delta = 0$, the rotation axis lies in the equatorial plane. A microwave pulse with phase $\phi = 0$ drives a rotation around the x axis, while a pulse with phase $\phi = \pi/2$ drives a rotation around the y axis. More generally, changing the microwave phase changes the azimuthal angle of the rotation axis. The corresponding unitary operation can be written as

$$U(\theta, \phi) = \exp \left[-i \frac{\theta}{2} (\cos \phi \sigma_x + \sin \phi \sigma_y) \right], \quad (2.12)$$

where θ is the rotation angle produced by the pulse. Therefore, the phase of the microwave field determines whether the pulse acts as an x -rotation, a y -rotation, or a rotation around another axis in the transverse plane.

This phase dependence is crucial for pulsed protocols. In a Rabi experiment, one can observe coherent oscillations even with a fixed microwave phase. However, Ramsey and Hahn echo sequences require well-defined relative phases between different pulses.

For example, two $\pi/2$ pulses with different phases project different spin components onto the measurement axis. Similarly, phase cycling techniques rely on changing the phase of selected pulses to separate the desired coherent signal from systematic errors or unwanted background contributions.

In a conventional microwave-switching setup, the microwave source is usually gated on and off, but the waveform phase is not freely programmed pulse by pulse. This restricts the set of accessible pulse operations. An arbitrary waveform generator, especially when used with IQ modulation, removes this limitation. The in-phase and quadrature components define the microwave field as

$$s(t) = I(t) \cos(\omega_{\text{IF}} t) - Q(t) \sin(\omega_{\text{IF}} t), \quad (2.13)$$

where $I(t)$ and $Q(t)$ are the programmable quadrature waveforms. By choosing these two components appropriately, both the amplitude and phase of the microwave pulse can be controlled. This quadrature-based description is the basis of the IQ modulation and up-conversion approach used later in the microwave chain [?].

2.3.4 Pulse Area and Pulse Duration

The effect of a microwave pulse is determined by its pulse area. For an effective two-level system driven resonantly, the rotation angle θ is given by the time integral of the Rabi frequency [2, 7],

$$\theta = \int_0^T \Omega(t) dt, \quad (2.14)$$

where T is the pulse duration. This equation shows that a pulse is not defined only by how long it lasts, but by the total area of its amplitude envelope. Two pulses with different temporal shapes can produce the same rotation angle if their integrated Rabi frequency is the same.

For a rectangular pulse with constant Rabi frequency Ω_R , the pulse area becomes

$$\theta = \Omega_R T. \quad (2.15)$$

A π pulse corresponds to a rotation angle $\theta = \pi$, while a $\pi/2$ pulse corresponds to $\theta = \pi/2$. Therefore, for a rectangular pulse,

$$t_\pi = \frac{\pi}{\Omega_R}, \quad t_{\pi/2} = \frac{\pi}{2\Omega_R}. \quad (2.16)$$

Experimentally, these pulse durations are obtained from a Rabi measurement. The microwave pulse duration is varied, and the photoluminescence oscillates as the spin population is driven between the two selected states. The period of the Rabi oscillation gives the calibrated durations for π and $\pi/2$ pulses [2, 4].

For shaped pulses, the same pulse-area principle applies, but $\Omega(t)$ is no longer constant. This pulse-area concept is widely used in magnetic resonance and shaped-pulse quantum control [7, 8]. For example, a Gaussian pulse has a smooth amplitude envelope, so the rotation angle is determined by the area under the Gaussian profile rather than by a simple product of constant amplitude and duration. To make a Gaussian pulse act as a π pulse, its amplitude must be scaled so that

$$\int_0^T \Omega_G(t) dt = \pi. \quad (2.17)$$

Similarly, a shaped $\pi/2$ pulse must satisfy an integrated area of $\pi/2$. This is important because changing the pulse shape without renormalizing its area changes the actual spin rotation.

3 Experimental Setup

This chapter presents the experimental setup used for continuous-wave and pulsed optically detected magnetic resonance measurements. The apparatus is organized around three main requirements: optical excitation and photoluminescence collection, microwave generation and delivery to the diamond sample, and the application of a bias magnetic field to control the NV-center spin resonances. Both the continuous-wave and pulsed operating modes are considered, since the same hardware platform provides the basis for the measurements discussed in the following chapters.

3.1 Photoluminescence Microscopy




Figure 3.1: Schematic of the confocal microscope beam path used in this work. The green ray bundle represents the 532 nm excitation laser, which is focused onto the diamond sample through a high-NA objective lens. Red arrows indicate the Stokes-shifted photoluminescence (PL) collected by the same objective and redirected by the dichroic mirror toward the detection path. Yellow-green rays illustrate out-of-focus emission that is spatially filtered by the confocal pinhole (fibre entrance), ensuring that only signal originating from the focal volume reaches the single-photon detectors. Adapted from [14].

The confocal principle further refines this picture by placing a small aperture (or, equivalently, the entrance face of a single-mode optical fibre) in a plane conjugate to the focus of the objective. Only light originating from the focal spot—the so-called *probe volume*—is transmitted; out-of-focus background is rejected. The result is an axially resolved three-dimensional point-spread function whose lateral extent is diffraction-limited to roughly $\lambda/(2\text{NA})$ as illustrated in 3.1.

The sample is usually mounted on a positioning stage, allowing the optical focus to be aligned with the region of interest in the diamond. For the present setup the objective is positioned *above* the sample (upright configuration), which is well-suited to solid substrates such as the diamond chips used here. An inverted arrangement, in which the objective is placed below the sample, is also possible and is often preferred when immersion objectives must be used or when biological samples in liquid media are studied [4].




Figure 3.2: Schematic representation of a simple photoluminescence-based ODMR setup. A laser beam excites the diamond sample through the microscope objective, while the emitted photoluminescence is collected and directed toward a photodetector. Microwave radiation is delivered to the sample by an antenna placed close to the diamond. Adapted from [4].

In this configuration, the optical excitation is typically provided by a colimated laser beam. The excitation beam is directed toward the diamond sample and focused through a microscope objective. The same objective can also be used to collect the red photoluminescence emitted by the NV centers. Since the excitation and emission wavelengths are different, optical elements such as a dichroic mirror and spectral filters are used to separate the excitation light from the fluorescence signal. The collected photoluminescence is then sent to 2 SPADs (Single Photon Avalanche Diodes) .

The microwave part of the setup is used to drive the ground-state spin transitions of the NV center. Microwave radiation is delivered to the sample through an antenna placed close to the diamond. When the microwave frequency matches one of the NV spin resonances, population is transferred between the spin sublevels, which changes the detected photoluminescence intensity. This fluorescence variation is the measured ODMR signal.

3.1.1 Light Detector and Imaging Mode

Unlike wide-field systems, which require cameras to simultaneously image a large area of the sample, the present setup records the photoluminescence signal point by point while the sample is scanned with a nanopositioning stage. The resulting photoluminescence maps are therefore reconstructed from the photon counts measured at each position. This acquisition method is particularly suitable for confocal investigations of individual NV centers and nanophotonic structures, where high spatial resolution and efficient background rejection are required [4].

Since the optical signal emitted by a single NV center consists of a relatively small number of photons, highly sensitive detectors are necessary. For this reason, the collected photoluminescence was detected using single-photon avalanche diodes (SPADs), which are capable of registering individual photon events with high temporal resolution. After passing through the spectral filtering stage, the fluorescence signal was coupled into an optical beam splitter fiber dividing the signal 50:50. The two outputs were connected to separate SPAD detectors arranged in a Hanbury Brown and Twiss (HBT) configuration.

The HBT configuration enables photon-correlation measurements through the evaluation of the second-order correlation function ($g^{(2)}(\tau)$), which is commonly used to characterize the photon statistics of quantum emitters. In addition to providing the photon-correlation capability required for emitter characterization, the use of fiber-coupled SPAD detectors ensures high detection sensitivity and efficient integration with the data acquisition electronics employed throughout the experiments[15, 16, 17].

3.1.2 Photoluminescence Digitization and Visualization

After spectral filtering and detection by the SPADs, the photoluminescence signal must be converted into a digital format suitable for acquisition, processing, and visualization. Unlike analog photodetectors, which generate a continuous electrical signal proportional to the incident optical power, single-photon avalanche diodes produce discrete transistor-transistor logic (TTL) pulses corresponding to individual photon detection events [4]. As a result, the measured photoluminescence intensity is obtained by counting the number of detected pulses within a specified acquisition interval.

The acquisition electronics recorded the photon-count information and transferred it to the control computer for further processing. The TTL pulses generated by the SPADs are routed to a data-acquisition (DAQ) card, which integrates the photon counts over a defined time window. The resulting count rate is streamed to Qudi in real time, providing a live photoluminescence trace used during alignment, focusing, and confocal scanning. This digital acquisition approach provides a convenient interface between the optical hardware and the software environment used for experiment control and data analysis.

For confocal imaging measurements, the detected photon counts were associated with the corresponding position of the piezoelectric scanning stage. By recording the photoluminescence intensity at each point of the scan, two-dimensional maps of the sample could be reconstructed and displayed in real time. These photoluminescence maps were used to identify individual pillars and locate emitters of interest for subsequent characterization.

The digitized photon-count data were also used during continuous-wave and pulsed ODMR measurements. In the pulsed operating mode, synchronization between the optical excitation, microwave control signals, and acquisition electronics was required to correctly associate the detected photoluminescence with the applied pulse sequence.

3.1.3 Scanning in the Confocal Microscope

Beyond confocal imaging, the experimental setup also supports additional characterization techniques. In particular, the Hanbury Brown and Twiss (HBT) configuration described previously can be used to measure the second-order correlation function $g^{(2)}(\tau)$, providing information on the photon statistics of the investigated emitters and enabling the verification of single-photon emission [4].

3.2 Pulse Generation

Time-domain experiments such as pulsed ODMR require precise control of both the optical and microwave excitation signals. In contrast to continuous-wave measurements, where the laser and microwave fields remain continuously active, pulsed measurements rely on carefully timed sequences that initialize, manipulate, and read out the NV-center spin state. Consequently, a pulse-generation system is required to provide synchronized control signals with sufficient temporal resolution and flexibility [4].

In the experimental setup used in this work, pulse generation was performed using a Zurich Instruments HDAWG. The arbitrary waveform generator provided multiple synchronized output channels, allowing coordinated control of the microwave excitation, optical excitation, and acquisition electronics. Such synchronization is essential for implementing

pulsed ODMR protocols and for ensuring reproducible timing relationships between the different subsystems of the experiment [4].

Optical pulse generation was achieved using an acousto-optic modulator (AOM) placed in the excitation path of the 532 nm laser, AOM-based modulation provides an effective method for producing optical pulses while maintaining stable laser operation. By switching the diffracted optical beam, the AOM enables rapid control of the optical excitation without requiring direct modulation of the laser source itself [4].

The HDAWG served a dual role in the experimental framework: it generated the intermediate-frequency (IF) signals required by the IQ upconversion architecture described in Chapter 4, and its digital marker outputs were used to trigger and synchronize the remaining subsystems, including the AOM gate and the DAQ acquisition window. This allowed the optical excitation, microwave pulses, and photon acquisition to operate within a single synchronized timing framework, which is essential for pulsed ODMR sequences.

3.3 Laboratory Set-up

3.3.1 Excitation Source and Optical Alignment

The optical alignment section therefore served three main purposes: controlling the excitation power, steering the beam into the microscope, and ensuring stable coupling into the objective path. Careful alignment was required because the detected photoluminescence signal strongly depends on the overlap between the focused excitation beam and the NV centers in the diamond sample.

3.3.2 Confocal Microscope and Optical Filtering

Optical excitation of the NV centers was provided by a continuous-wave Coherent Prometheus laser operating at a wavelength of 532 nm. In contrast to many spectroscopic techniques, ODMR does not impose particularly stringent requirements on the spectral linewidth of the excitation source. Instead, stable optical power and efficient delivery of the excitation beam to the sample are generally of greater importance [4].

The laser beam was first directed through an acousto-optic modulator (AOM), which enabled optical pulse generation for pulsed ODMR experiments. After the AOM, the beam passed through neutral density (ND) filters that allowed precise adjustment of the excitation power while maintaining stable laser operation. The beam was then steered through a set of dielectric mirrors and directed toward the microscope through a dichroic mirror, which combined the excitation and collection paths.




Figure 3.3: Main components of the excitation and control setup used during confocal imaging: (a) Coherent Prometheus laser used as the continuous-wave 532 nm excitation source; (b) optical alignment section including the neutral-density filter and dielectric steering mirrors used to control and direct the excitation beam; (c) Keysight PSG microwave generator used during Qudi-controlled operation.

Since ODMR measurements rely on detecting small variations in photoluminescence intensity, fluctuations in laser power can directly affect the measured signal. Careful control of the excitation power and optical alignment was therefore required to ensure stable operation and reproducible measurements throughout the experiments.

The excitation beam entered the confocal microscope, where a 567 nm dichroic mirror reflected the green excitation light toward the objective while transmitting the red-shifted photoluminescence emitted by the NV centers. The microscope objective, with numerical aperture in the range ($NA = 0.6-0.7$), focused the excitation beam onto the diamond sample and collected the emitted fluorescence through the same optical path 3.4.

The collected photoluminescence passed through a filtering chain designed to suppress residual excitation light and isolate the NV emission band. A 532 nm notch filter, as shown in figure 3.5 was used to attenuate scattered laser light, while long-pass filters at 600 nm and 650 nm removed remaining short-wavelength background and transmitted the red NV photoluminescence as shown in figure 3.6. After filtering, the signal was focused using an achromatic doublet and coupled into an optical fiber.

The fiber-coupled signal was then sent to a 50:50 fiber beam splitter connected to two SPAD detectors. This detection path enabled both photon-count acquisition and photon-correlation measurements. During alignment and scanning, the detected count rate was monitored in real time in the control software.




Figure 3.4: Confocal microscope used for NV-center imaging and fluorescence collection. The same objective was used to focus the 532 nm excitation beam onto the diamond sample and to collect the red photoluminescence emitted by the NV centers.




Figure 3.5: NF533-17 notch filter used to suppress residual 532 nm excitation light. (a) Physical appearance of the filter. (b) Transmission spectrum showing strong rejection around 532 nm while maintaining high transmission outside the rejection band. Data adapted from [18].




Figure 3.6: FEL0650 long-pass filter used in the detection path. (a) Physical appearance of the filter. (b) Optical density response showing strong attenuation below the cut-on wavelength. (c) Transmission spectrum demonstrating high transmission above 650 nm. Data adapted from [19].

3.3.3 Diamond Sample and Antenna Structure

The experiments reported in this thesis were performed using a custom diamond sample fabricated by Qnami. The device consists of a $3\text{ mm} \times 3\text{ mm}$ quantum-grade diamond substrate with (111) crystallographic orientation and an initial thickness of approximately $50\text{ }\mu\text{m}$. Nitrogen ions were implanted with an energy of 200 keV, corresponding to an estimated implantation depth of approximately 200–400 nm below the diamond surface. The implantation density was selected to obtain, on average, approximately one NV center per nanophotonic pillar.

The central region of the chip was patterned with nanophotonic pillar arrays designed to improve the collection efficiency of the NV-center photoluminescence. The patterned area occupies approximately $2.2\text{ mm} \times 2.2\text{ mm}$ of the $3\text{ mm} \times 3\text{ mm}$ substrate and is composed of 23×20 replicated unit-cell fields. This large patterned region provides many nominally identical structures, allowing different pillars and emitters to be located and compared during confocal measurements.

Each unit-cell field contains 1044 nanophotonic pillars arranged in a regular lattice, as illustrated in Fig. 3.7. The nominal pillar diameter is approximately 200 nm, while the pillar period is $2.6\text{ }\mu\text{m}$. This spacing is large enough to optically distinguish neighboring pillars in confocal scans, while still providing a dense array of candidate emitters over the patterned area. The protective frame surrounding each field also helps identify the regions during optical navigation.

The use of nanophotonic pillars is important because NV centers embedded in bulk diamond emit photons into a high-refractive-index medium, which limits the fraction of light that can be extracted into the microscope objective. Structuring the diamond into




Figure 3.7: Representative unit-cell field of the nanophotonic pillar array. Each field contains 1044 pillars with a nominal diameter of approximately 200 nm and a period of 2.6 μm .

pillar-like geometries improves the optical extraction of the emitted photoluminescence and therefore increases the detected photon count rate. This is especially useful when working with individual or few NV centers, where the fluorescence signal is intrinsically weak.

The fabricated structures were characterized by scanning electron microscopy. The parabolic geometry visible in the high-magnification image acts as a photonic extraction structure, helping direct more NV fluorescence toward the collection optics.

A photograph of the actual diamond sample used in the experiment is shown in Fig. 3.8. The chip was handled and mounted carefully to preserve the patterned region and to allow optical access during confocal imaging. The nanophotonic pillar array provided the spatially organized platform from which individual bright emitters were located, selected, and subsequently characterized.

The diamond chip was mounted on a custom support incorporating a copper wire positioned in close proximity to the patterned region, as shown in Fig. 3.9. The antenna was designed to generate the oscillating magnetic field required for ODMR measurements while maintaining unobstructed optical access to the nanophotonic pillar array. Its compact size allowed efficient microwave delivery to the implanted NV centers and facilitated integration with the confocal microscope setup used throughout the experiments.




Figure 3.8: Photograph of the Qnami 200 keV nitrogen-implanted diamond sample used in the experiment. The patterned nanophotonic region is located at the center of the chip and contains the pillar arrays investigated during the measurements.




Figure 3.9: 70 μm copper wire used to deliver microwave excitation to the NV centers during ODMR measurements.

3.4 Bias Magnetic Field

Many ODMR experiments require the application of a static magnetic field in addition to the optical and microwave excitation. By lifting the degeneracy of the $m_s = \pm 1$ spin states through the Zeeman effect, a bias magnetic field allows the resonance frequencies of the NV center to be tuned and separated. The magnitude and orientation of the applied field therefore play an important role in determining the observed ODMR spectrum and the sensitivity of the measurement [4].

Several approaches can be used to generate the required magnetic field, common solutions include permanent magnets, Helmholtz coil systems, or hybrid configurations combining both technologies. Permanent magnets provide a simple and compact solution, whereas coil-based systems offer greater flexibility because the field magnitude and direction can be adjusted electronically during the experiment [4].

The experimental setup used in this work employed a permanent magnet positioned in close proximity to the diamond sample. The magnet was oriented such that a static bias field was applied to the NV centers, resulting in the splitting of the spin resonances observed in the ODMR measurements. The position of the magnet could be adjusted to modify the magnetic field experienced by the sample and optimize the experimental conditions.

The use of a permanent magnet offers a practical and stable method for generating the magnetic field required for ODMR measurements while maintaining a compact experimental configuration. Since the magnetic field remained constant during the measurements, a permanent-magnet solution provided sufficient flexibility without the additional complexity associated with coil-based magnetic-field control systems.

3.4.1 Nanopositioning System

The diamond sample was mounted on a three-axis piezoelectric nanopositioning stage, allowing controlled motion along the (x), (y), and (z) directions. The stage used in the setup was a PiezoSystem Jena TRI TOR100SG controlled by an NV40/3CLE controller. This positioning system enabled precise scanning of the diamond sample under the fixed optical focus.

During confocal imaging, the piezo stage was used to scan the sample while the photon count rate was recorded at each position. This allowed two-dimensional photoluminescence maps to be reconstructed and used to locate nanophotonic pillars and identify emitters suitable for subsequent ODMR and photon-correlation measurements.



Figure 3.10: Piezoelectric stage controller used for precise three-dimensional positioning of the diamond sample during confocal scanning.

3.4.2 Microwave Control Configuration

Microwave radiation is required to manipulate the spin states of NV centers and perform ODMR measurements. By applying a microwave field whose frequency is resonant with the ground-state spin transitions, population can be transferred between the spin sublevels, leading to measurable changes in the photoluminescence signal. Consequently, the microwave instrumentation must provide stable frequency generation, sufficient output power, and compatibility with both continuous-wave and pulsed operating modes [4].

In the experimental setup used in this work, microwave signals were generated using a Keysight PSG signal generator. Prior to the developments presented, the microwave chain consisted of the signal generator followed by pulse-generation or switching electronics and the delivery antenna. This configuration enabled conventional continuous-wave and pulsed ODMR measurements.

A central objective of this work was the implementation of an external microwave up-conversion architecture. In the upgraded configuration, the Keysight PSG generator was used as a local oscillator (LO), while an intermediate-frequency (IF) signal was generated by a Zurich Instruments HDAWG. The LO and IF signals were combined using an external IQ mixer, producing an upconverted microwave signal at the desired output frequency. The resulting signal was subsequently amplified and delivered to the sample through the microwave antenna. The design and characterization of this upconversion system are dis-

cussed in detail in Chapter 5.

The microwave radiation was applied to the diamond sample using a custom copper antenna integrated into the sample holder [4], the geometry of the antenna influences both the strength and spatial distribution of the microwave magnetic field. In the present setup, the antenna was positioned in close proximity to the diamond in order to provide efficient microwave coupling to the implanted NV centers while maintaining optical access for confocal measurements.

The microwave control chain evolved during the work. In the initial configuration, microwave excitation was generated using a Keysight N5138B Analog signal generator, followed by switching or pulse-control electronics and then delivered to the diamond sample through the microwave antenna. This configuration was suitable for conventional CW and pulsed ODMR measurements.

The upgraded configuration used the Keysight as a local oscillator, while an external IQ mixer and intermediate-frequency signals were used to generate upconverted microwave signals. The upconverted output was amplified and delivered to the antenna near the diamond sample. The detailed implementation, calibration, and software integration of this upconversion architecture are presented in the following chapters.

4 The HDAWG and Qudi Integration

4.1 Motivation for Using the HDAWG

In CW ODMR, the microwave source is applied continuously while the fluorescence is monitored as a function of microwave frequency. In pulsed ODMR, instead, the microwave excitation must be applied only during well-defined time windows and must be synchronized with laser initialization, optical readout, and photon-counting acquisition. Therefore, the waveform generator must not only produce analog voltages, but must also provide accurate triggering, marker signals, programmable waveform shapes, and stable multi-channel timing.

The Zurich Instruments HDAWG was chosen because it combines high-speed arbitrary waveform generation with precise digital synchronization and software programmability. These properties are essential for generating shaped microwave pulses, controlling the in-phase and quadrature inputs of an IQ mixer, and integrating the pulse-generation process into Qudi.

4.1.1 Capabilities of the Zurich Instruments HDAWG

The front panel of the HDAWG gives direct access to the analog and digital signals used during pulse generation. In the four-channel version used in this work, the instrument provides four main waveform outputs. Each output group contains three relevant connectors: a trigger input, a marker output, and a waveform output[?].




Figure 4.1: Front panel of the Zurich Instruments HDAWG. Each channel group provides a trigger input, a marker output, and an analog waveform output adapted from [20].

The *Wave* connector is the analog output of the arbitrary waveform generator. This is the port from which the programmed voltage waveform is emitted. In the present setup, these outputs are used to generate the baseband signals required for microwave control. For example, two waveform channels can be used as the in-phase and quadrature components of an IQ modulator. The HDAWG has a sampling rate of 2.4 GSa/s and a 16-bit vertical resolution, allowing analog waveforms to be represented by a dense sequence of samples with fine amplitude discretization. The resulting temporal resolution is important for arbitrary waveform generation because the smaller the sampling interval, the more faithfully the discrete waveform approximates the intended continuous envelope. This requirement is especially important for Gaussian and DRAG pulses, where insufficient temporal resolution may distort the smooth envelope and its derivative component [20].

The *Trig* connector is a digital trigger input. A trigger is an external timing signal that tells the HDAWG when to start, wait, branch, or synchronize a programmed waveform sequence. In a pulsed ODMR experiment, triggers are essential because the microwave pulse, laser pulse, and photon-counting window must occur in the correct order. For instance, the HDAWG may wait for an external trigger before playing a microwave pulse, or it may be synchronized with a data-acquisition card that counts photons during a specific readout window. According to the HDAWG manual, the front-panel trigger inputs are managed through the LabOne DIO tools, where trigger levels and trigger states can be configured and monitored [20].

The *Mark* connector is a marker output. A marker is a digital signal generated together with the analog waveform. While the waveform output carries the shaped analog voltage, the marker output carries a high/low digital pulse that can be used to control or synchronize another device. For example, a marker can be used to open a microwave switch, trigger a laser driver, or inform the acquisition system that the microwave pulse has started. This makes markers different from triggers: a trigger is usually an input received by the HDAWG, while a marker is usually an output produced by the HDAWG. The HDAWG manual specifies that marker outputs are available on the front panel and can be assigned to internal marker bits associated with the waveform sequence [20].

The rear panel provides the communication, synchronization, and system-level connections needed to integrate the HDAWG into a larger experimental setup.



Figure 4.2: Rear panel of the Zurich Instruments HDAWG, showing clock, trigger, digital I/O, Ethernet, USB, and power connections, adapted from [20].

The *Reference Clock* connector allows the HDAWG to share a common frequency reference with other instruments. This is important when several devices must remain phase-stable over time. If the waveform generator, microwave source, and acquisition electronics are all referenced to the same clock, their relative timing and phase drift are reduced.

The *Sample Clock* connectors are used for clock synchronization of the waveform-generation engine. The sample clock determines the rate at which waveform samples are

converted into analog voltages. Since pulse duration is ultimately represented by a number of samples, stable sample-clock control is essential for accurate pulse timing.

The *DIO* connector provides digital input/output communication. This can be used for more complex synchronization patterns, digital control, or communication with external electronics. The HDAWG manual describes the DIO tab in LabOne as the interface used to monitor and control digital I/O connectors, marker outputs, and trigger inputs [20].

The *1GbE* Ethernet port is the main communication interface used to connect the HDAWG to the control computer. Through this connection, LabOne and the Zurich Instruments Python API can configure the instrument, upload waveforms, compile SEQC programs, and control execution. This is the interface that makes software integration with Qudi possible[20].

The *USB* connection provides an additional communication interface. In practice, Ethernet is typically preferred for remote control and integration into automated experimental software, while USB remains available for direct communication.

The *MDS* connector is used for multi-device synchronization. This is relevant when several Zurich Instruments devices must operate as a synchronized system. Although the present work focuses on a single HDAWG, this feature is important for future scalability.

Finally, the rear panel contains the power switch and AC power input, which supply the instrument. Although these are not part of the pulse sequence itself, they are part of the complete hardware environment required for stable operation[20].

4.1.2 Relevant HDAWG Features for Pulsed ODMR

Several HDAWG specifications are directly relevant to the experiment. The 2.4 GSa/s sampling rate corresponds to a sampling interval of approximately 416 ps, allowing even short microwave pulses with durations of only a few tens of nanoseconds to be represented by a large number of samples and therefore generated with smooth, high-fidelity envelopes [?].[20].

The 16-bit vertical resolution determines how finely the voltage amplitude can be discretized. A higher resolution is useful when the pulse envelope changes gradually, because small amplitude steps can be represented without strong quantization errors. This is especially relevant for shaped pulses, where the exact envelope shape affects the spectral response of the microwave excitation[20].

The 750 MHz signal bandwidth determines how rapidly the output waveform can change without being strongly distorted by the analog output chain. This is important for fast

pulse edges and for intermediate-frequency I/Q signals. Even if the final microwave frequency is produced after upconversion, the baseband signals must still be generated cleanly by the HDAWG.

The 5 V_{pp} maximum output amplitude gives flexibility in matching the voltage required by external components, such as the IQ modulator. In practice, the maximum amplitude is not always used, because the correct drive level depends on mixer linearity, attenuation, and the desired microwave power. However, having a large available range makes the setup more adaptable.

The trigger-to-output delay below 50 ns is important for synchronization. It means that after the HDAWG receives a trigger, the first waveform sample can be produced with a very small delay. In pulsed ODMR, this reduces timing uncertainty between the waveform generator and the rest of the experiment.

The HDAWG also supports digital modulation at multiple frequencies. This is relevant for generating phase-controlled signals and for experiments where the waveform must include controlled frequency or phase components. In this work, the most important use of phase control is the generation of I and Q signals for IQ modulation.

The LabOne AWG Sequencer and Compiler are central to the operation of the HDAWG. Instead of only uploading a static waveform, the user can write a sequence program in the SEQC language. This program defines when waveforms are played, when the instrument waits for triggers, and how repetitions or branches are handled. The compiler then converts this sequence into instructions executable by the HDAWG hardware [20].

Finally, the possibility of controlling the HDAWG through LabOne and through application programming interfaces is essential for Qudi integration. Manual control through LabOne is useful during testing and debugging, while API control is required for automation. In the final implementation, Qudi must be able to generate pulse sequences, translate them into HDAWG-compatible waveforms and SEQC code, upload them to the instrument, and start the programmed sequence.

4.1.3 Advantages for Pulse Shaping and Phase Control

The main advantage of using the HDAWG is that the microwave control signal can be programmed as a shaped waveform rather than as a simple rectangular pulse[20]. Rectangular pulses are easy to generate, but their abrupt edges introduce additional spectral components. This can be undesirable in spin-control experiments because it may excite unwanted transitions or reduce spectral selectivity[13].

Gaussian pulses provide smoother amplitude variation. Their gradual rising and falling

edges reduce the spectral broadening associated with sharp rectangular pulses. This makes them more suitable for controlled microwave excitation of the NV-center spin transition[13].

DRAG pulses extend this idea by adding a quadrature correction component. In practice, this requires independent control of two synchronized channels: one channel carries the main Gaussian component, while the other carries a derivative-shaped quadrature component. The HDAWG is suitable for this because its channels can be programmed coherently and controlled from the same sequencer environment[21].

Phase control is equally important for IQ modulation. When two HDAWG outputs are connected to the I and Q inputs of an IQ mixer, the relative phase between them determines the generated sideband structure. A correct 90° phase relationship is crucial for single-sideband generation, while phase or amplitude imbalance leads to unwanted sideband leakage. Therefore, the HDAWG is not only a pulse generator, but also a phase-control tool for microwave upconversion.

For these reasons, the HDAWG was selected as the pulse-generation platform for the experiment. Its waveform outputs provide shaped analog pulses, its trigger inputs and marker outputs provide synchronization, its rear-panel connections allow integration with the wider setup, and its LabOne/API control makes automated Qudi integration possible.

4.2 Qudi Architecture

Modern optically detected magnetic resonance (ODMR) experiments require the coordinated operation of several laboratory instruments, such as laser sources, microwave generators, photon-counting detectors, piezoelectric positioning stages, arbitrary waveform generators, oscilloscopes, and data-acquisition systems[4].

Qudi is an open-source software framework developed specifically for ODMR experiments and other laboratory measurements requiring coordinated instrument control. Implemented entirely in Python, it adopts a modular architecture in which the graphical user interface, experimental logic, and hardware communication are separated into independent software layers. This organization improves software maintainability, facilitates the integration of new instruments, and allows the same experimental procedures to be executed using different hardware implementations with minimal modifications [22].

Unlike conventional laboratory software that directly couples experimental routines with individual instruments, Qudi introduces an abstraction layer between the experiment and the physical hardware. Experimental procedures interact only with standardized software interfaces, while the communication with the actual instruments is delegated to dedicated hardware modules. This design considerably simplifies the development of new

instrument drivers and allows the software to remain independent of manufacturer-specific communication protocols [4].

4.2.1 Graphical User Interface, Logic, and Hardware Layers

The Qudi software architecture is organized into three principal software layers: the Graphical User Interface (GUI), the Logic layer, and the Hardware layer. Each layer performs a well-defined task and communicates with the other layers only through standardized interfaces, thereby reducing software complexity and improving code reusability [4].




Figure 4.3: Architecture of the Qudi framework. During initialization, `start.py` loads the core modules and reads the experiment configuration file, which specifies the GUI, Logic, and Hardware modules required for a particular setup. Communication between the Logic and Hardware layers is performed through standardized interfaces, allowing different hardware implementations to be integrated without modifying the experimental logic. Adapted from [4].

Graphical User Interface (GUI)

The GUI constitutes the layer directly accessed by the experimenter. It provides the graphical windows used to configure experimental parameters, launch measurements, monitor the acquisition in real time, and visualize the collected data. Typical parameters that can be modified through the GUI include microwave frequency, microwave power, scanning ranges, pulse durations, acquisition time, averaging parameters, and detector settings.

Besides configuring experiments, the GUI also provides real-time visualization of the acquired data. Confocal images, ODMR spectra, pulse sequences, fluorescence counts, and other experimental results can be displayed while the measurement is running, allowing the user to immediately evaluate the quality of the acquisition. Since the GUI communicates exclusively with the Logic layer, modifications to the graphical interface do not affect the hardware drivers, thereby improving software maintainability and simplifying future developments [4].

Another important characteristic of the GUI is that it remains optional. Although graphical interfaces offer a convenient method for interacting with the instruments, Qudi can also be operated directly from the integrated Python console or through external Python scripts. This flexibility enables automated measurements and facilitates the integration of Qudi into larger experimental workflows [4].

Logic Layer

The Logic layer represents the core of the experimental software. Rather than communicating directly with laboratory instruments, it coordinates the execution of measurements by receiving commands from the GUI, configuring the appropriate hardware modules, processing the acquired data, and returning the results to the graphical interface for visualization.

Each experimental technique implemented in Qudi is associated with one or more dedicated logic modules. Examples include confocal microscopy, continuous-wave ODMR, pulsed ODMR, Rabi oscillations, Ramsey measurements, and Hahn echo experiments. These modules define the sequence of operations required to perform each measurement while remaining completely independent of the specific instruments connected to the setup [4].

The separation between experiment logic and hardware communication provides significant advantages. Since the experimental algorithms interact only with standardized interfaces, the same logic module can operate with different laboratory instruments provided that they implement the required interface. Consequently, replacing a microwave generator, data-acquisition system, or arbitrary waveform generator generally requires only changing the corresponding hardware driver without modifying the measurement algorithms themselves [4].

Logic modules also communicate with one another whenever an experiment requires several coordinated operations. For example, a pulsed ODMR experiment may simultaneously control the microwave source, laser excitation, photon-counting electronics, and positioning stages while collecting and processing fluorescence data. The modular organization of the Logic layer allows these operations to be distributed among several specialized modules that cooperate during the execution of the experiment [4].

Hardware Layer

The Hardware layer provides the communication between Qudi and the physical laboratory instruments. Each hardware module implements the commands necessary to configure and operate a specific device while translating generic software instructions into manufacturer-dependent communication protocols.

This abstraction isolates all device-specific operations inside dedicated hardware drivers. Consequently, communication libraries, instrument commands, data formatting, and synchronization procedures remain hidden from the rest of the software. From the perspective of the Logic layer, every instrument simply provides a predefined set of functions independently of its manufacturer or communication protocol [4].

Hardware modules are responsible for initializing the instruments, configuring their operating parameters, executing measurement commands, acquiring experimental data, and returning the acquired information to the Logic layer. In addition to controlling real laboratory equipment, Qudi also supports virtual or dummy hardware modules that emulate the behaviour of physical instruments. These simulated devices are particularly useful during software development because they allow new measurement routines to be tested before the corresponding hardware becomes available [4].

Communication between the Logic and Hardware layers is achieved through standardized hardware interfaces. Rather than interacting with individual devices directly, the Logic layer communicates only with interface classes defining the operations expected from each category of instruments, such as pulse generators, microwave sources, cameras, data-acquisition systems, or positioning stages. Each hardware driver must implement these interfaces while internally handling the communication with the actual instrument. This object-oriented approach ensures hardware independence and considerably simplifies the integration of new laboratory equipment into the Qudi framework [4].

The work presented in this thesis follows exactly this architectural philosophy. Instead of modifying the existing measurement logic implemented in Qudi, a dedicated hardware module was developed for the Zurich Instruments HDAWG. By implementing the pulser interface expected by the Logic layer, the HDAWG can be controlled using the same generic pulse-generation routines already available within Qudi while encapsulating all communication through the Zurich Instruments LabOne API inside the hardware driver. This approach preserves the modularity of the framework and allows the new arbitrary waveform generator to be integrated without affecting the remaining software components.

4.2.2 Configuration and Extension of Qudi

Qudi is designed to be easily adapted to different laboratory setups through a flexible software configuration process. Before an experiment can be executed, the framework must be installed together with its software dependencies and configured to match the available instrumentation. Rather than embedding the characteristics of a specific experimental setup into the source code, Qudi relies on configuration files that define the software modules to be loaded and the communication links between them. This design greatly simplifies the

deployment of the framework across different laboratories and enables experiments to be reproduced with different hardware configurations [4].

The software is typically installed inside a dedicated Python environment, allowing all required libraries and dependencies to be managed independently of other Python projects. Once the environment has been created, Qudi is launched through the `start.py` script, which initializes the framework, loads the core modules, parses the selected configuration file, and establishes communication between the graphical interface, logic modules, and hardware drivers[4].

The configuration file constitutes the central element of the framework. Written in YAML format, it specifies every software module required by an experiment together with the relationships between them. The file contains the graphical user interfaces, logic modules, hardware modules, communication addresses, and device-specific parameters necessary to initialize the experimental setup. In addition to defining which instruments are available, the configuration file establishes how information flows between the different software layers, allowing the Logic modules to communicate with the appropriate hardware drivers through standardized interfaces [4].

One of the principal advantages of this architecture is its extensibility. When a laboratory instrument is not already supported by Qudi, a dedicated hardware module can be implemented by inheriting the appropriate hardware interface. This interface defines a collection of abstract methods describing the operations expected from a given category of instruments, while the hardware module provides the device-specific implementation using the communication protocol supplied by the manufacturer. Once the new driver has been developed, it can be referenced in the configuration file and becomes available to the rest of the software without requiring modifications to the existing measurement logic or graphical interface [4].

After all required modules have been configured, experiments can be executed directly through the graphical user interface. Experimental parameters may be adjusted interactively, measurement sequences can be selected or customized, and the acquired data are displayed in real time using the visualization and analysis tools integrated within the framework. This workflow enables automated execution of complex ODMR measurements while maintaining a clear separation between user interaction, experiment control, and hardware communication [4].

The integration of the Zurich Instruments HDAWG presented in this thesis follows this design philosophy. A new hardware module was developed to implement the pulser interface required by Qudi, while the corresponding device configuration was added to the YAML configuration file. As a result, the HDAWG can be initialized automatically

during software startup and accessed by the existing pulsed-ODMR logic modules without modifying the overall architecture of the framework.

4.2.3 Existing Pulse Generation Support and Missing Functionality

The original release of Qudi was designed to provide a complete software framework for ODMR experiments by integrating instrument control, data acquisition, measurement automation, and online data processing within a single platform [4]. The framework initially focused on experiments involving single quantum emitters and included software modules for confocal microscopy, continuous-wave ODMR, and pulsed spin-control measurements.

For confocal microscopy, Qudi provides automated control of piezoelectric positioning stages together with synchronized photon-counting acquisition. These capabilities enable raster scanning of the sample, depth profiling through z -stack measurements, and multidimensional image acquisition. In addition, the software supports automatic relocation and tracking of individual emitters during long measurements, reducing the effects of mechanical or thermal drift [4].

For magnetic resonance experiments, the platform includes dedicated measurement routines for continuous-wave ODMR, allowing automated microwave-frequency sweeps, real-time fluorescence acquisition, and online fitting of resonance spectra. These tools simplify the optimization of experimental parameters while improving the reproducibility of the measurements [4].

Qudi also incorporates a pulse-generation framework for coherent spin manipulation. Standard pulsed ODMR protocols, including Rabi oscillations, Ramsey interferometry, and Hahn echo, are already implemented within the Logic layer. During these experiments, the software coordinates the pulse generator, microwave source, laser excitation, and photon-counting hardware while providing utilities for sequence configuration, pulse visualization, and automated parameter management [4].

Despite these capabilities, the original pulse-generation framework was primarily developed for conventional pulse generators and did not provide native support for the Zurich Instruments HDAWG employed in this work. In particular, advanced arbitrary waveform generation, synchronized multi-channel output, direct communication through the LabOne Python API, and the generation of shaped microwave envelopes such as Gaussian and DRAG pulses were not available within the existing hardware modules. Furthermore, no interface existed for compiling pulse sequences into the Zurich Instruments SEQC language required by the HDAWG.

These limitations motivated the development of a dedicated HDAWG hardware module. The module extends the existing Qudi pulse-generation framework while preserving its modular architecture, allowing the new arbitrary waveform generator to be controlled through the same Logic layer without modifying the higher-level experimental routines. The implementation of this module is described in the following section.

4.3 Development of the HDAWG Hardware Module

The central software contribution of this work is a dedicated Qudi hardware module, `HDAWGToolkitPulser`, which integrates the Zurich Instruments HDAWG arbitrary waveform generator into the Qudi pulse-generation framework. The module implements the complete Qudi `PulserInterface`, allowing the HDAWG to be used interchangeably with any other pulser device supported by Qudi, while internally translating Qudi’s abstract pulse representations into instrument-specific waveform data and sequencer programs.

4.3.1 Architecture of the Developed Module

The module is implemented as a Python class inheriting from Qudi’s abstract `PulserInterface` base class, which defines the contract that any pulse-generation hardware must fulfil: methods for device activation, waveform writing, asset loading, channel configuration, and run-state control. By adhering strictly to this interface, the HDAWG becomes a drop-in replacement for previously supported pulsers, and all higher-level Qudi logic modules (pulsed measurement logic, sequence editors, graphical interfaces) operate on it without modification.

Internally, the module is organised around three functional layers:

1. **Connection layer** — establishes and manages the session with the LabOne Data Server and exposes the two AWG cores of the HDAWG as controllable objects.
2. **Buffering layer** — accumulates the sampled analog and digital waveform data delivered by Qudi in chunks, normalises the total length to hardware granularity constraints, and applies channel activation masks.
3. **Upload and execution layer** — generates a deterministic SEQC sequencer program matched to the waveform memory layout, compiles and uploads it to each AWG core, transfers the waveform data, and controls sequence execution.

The lifecycle of the module follows the standard Qudi pattern: the `on_activate()` method establishes the instrument connection, reads the hardware sample clock, applies

the static device configuration, and constructs the constraints object; `on_deactivate()` safely disables all outputs and closes the session. A typical measurement workflow proceeds as `write_waveform() → load_waveform() → pulser_on()`, mirroring the sequence of calls issued by Qudi’s pulsed measurement logic[23].

Channel grouping. The HDAWG contains four analog output channels and four digital marker outputs, organised in two independent AWG cores operating in a 2×2 grouping: core 0 drives outputs OUT1/OUT2 together with markers MARK0/MARK1, while core 1 drives OUT3/OUT4 with MARK2/MARK3. The module addresses both cores symmetrically, uploading to each core its own waveform pair and marker data, so that all four analog channels and all four marker channels are available to Qudi as `a_ch1–a_ch4` and `d_ch1–d_ch4` respectively.

4.3.2 Communication with the LabOne Data Server

All communication with the instrument is mediated by the LabOne Data Server, the vendor-provided middleware that exposes every instrument setting and data stream as a node in a hierarchical tree. The module employs the modern `zhinst.toolkit` Python API, which wraps the low-level node tree in an object-oriented interface [20].

Upon activation, the module opens a `Session` to the Data Server and connects to the instrument by its device identifier:

```
self.session = Session("localhost")
self.hdawg   = self.session.connect_device("DEV8244")
self.awg0    = self.hdawg.awgs[0]
self.awg1    = self.hdawg.awgs[1]
```

Instrument parameters are then accessed as attribute-style nodes; for example, the hardware sample clock is read from `system.clocks.sampleclock.freq`, the output ranges and offsets are configured through the `sigouts[n]` nodes, and the routing of the front-panel trigger outputs to internal marker signals is established via the `triggers.out[n].source` nodes. This node-based access pattern makes every configuration action explicit and verifiable: after each write, the module reads the value back and logs it, providing a traceable record of the instrument state at activation time.

Waveform uploads use the dedicated `Waveforms` container of the toolkit, which serialises the analog sample arrays and packed marker bits into the vector format expected by the instrument’s waveform memory, and transfers them through the session with `write_to_waveform_memory()` [20].

4.3.3 Waveform Generation from Qudi Pulse Sequences

Within Qudi, pulse sequences are described hierarchically. The elementary unit is the `PulseBlock`, an ordered list of block elements each specifying a duration and, per channel, an analog envelope function and a digital state. Blocks are combined with repetition counts into a `PulseBlockEnsemble`, and ensembles can in turn be arranged into a `PulseSequence` with branching and looping directives.

Qudi’s sampling logic converts this abstract description into concrete sample arrays at the device sample rate, and delivers them to the hardware module through the chunked `write_waveform()` interface. Each call provides a dictionary of analog sample arrays (`a_ch1...a_ch4`, floating-point values normalised to the output range) and digital sample arrays (`d_ch1...d_ch4`, boolean), together with flags marking the first and last chunk and the total expected sample count.

The module accumulates these chunks into two internal buffers: an $(N \times 4)$ single-precision floating-point matrix for the analog channels and an $(N \times 4)$ integer matrix for the digital channels. On receipt of the final chunk, the buffers are truncated or zero-padded to exactly the declared total length, guaranteeing consistency between the sampled data and the sequencer program generated subsequently.

Two hardware granularity constraints are enforced at this stage: the waveform length must be at least 32 samples and a multiple of 16 samples. A dedicated normalisation routine rounds the buffer length up to the next valid value, zero-padding the appended samples, so that arbitrary user-defined sequence durations are accepted transparently[20].

4.3.4 Generation of SEQC Programs

The HDAWG executes waveforms under the control of a sequencer program written in SEQC, a C-like domain-specific language compiled on the instrument. A subtle but critical aspect of the implementation is that the memory layout expected by the compiled sequencer must match exactly the layout of the uploaded waveform data: a mismatch between the two — for instance, a program compiled without marker placeholders while the uploaded vectors include marker bits — results in a vector transfer error and a failed upload[20].

To eliminate this class of failure, the module generates the SEQC program deterministically from the known properties of the waveform to be uploaded. For a waveform of N samples with markers in use, the generated program is:

```
wave w1 = placeholder(N, true, true);  
wave w2 = placeholder(N, true, true);
```

```

assignWaveIndex(w1, w2, 0);

while (true) {
    playWave(1, 2, w1, 1, 2, w2);
}

```

The `placeholder()` declarations reserve waveform memory of the correct length and marker configuration without embedding the sample data in the program itself; `assignWaveIndex()` binds the placeholders to index 0 of the waveform table; and the infinite `playWave()` loop replays the waveform continuously once the core is enabled. Because the program text depends only on the length N and the marker flag, the compiled memory layout is guaranteed to be consistent with the subsequently uploaded `Waveforms` container, which is assigned to the same table index[20].

After loading the program with `load_sequencer_program()`, the module explicitly waits for the core’s `ready` node to assert before writing the waveform memory, ensuring the compiler has finished before data transfer begins.

Trigger and marker routing. The four front-panel trigger outputs of the instrument are routed to the marker bits of the two cores during device configuration. This makes the digital channels `d_ch1`–`d_ch4` physically available as TTL signals, which in the present setup are used to gate the acousto-optic modulator `d_ch1` and to trigger the photon-counting acquisition `d_ch2`, thereby placing laser excitation, microwave generation, and detection in a single synchronised timing frame. In addition to the digital channels used to gate the acousto-optic modulator and trigger the data-acquisition system, the digital output `d_ch3` was configured to control the microwave switch. This ensured that microwave radiation was applied only during the intended pulse window, providing precise synchronization between the optical excitation, microwave pulses, and photon-counting sequence.

4.3.5 Multi-channel Waveform Management

The mapping between Qudi’s logical channels and the physical instrument resources is summarised in Table 4.1.

For each core, the two digital channels are packed into a single 16-bit marker word per sample, with the first channel occupying bit 0 and the second bit 1:

$$m[n] = (d_1[n] \ll 0) \vee (d_2[n] \ll 1), \quad (4.1)$$

Table 4.1: Mapping between Qudi logical channels and HDAWG hardware resources as implemented in the developed module.

Qudi channel	HDAWG resource	AWG core
a_ch1	Analog output 1	core 0
a_ch2	Analog output 2	core 0
a_ch3	Analog output 3	core 1
a_ch4	Analog output 4	core 1
d_ch1	Marker bit 0 (MARK0)	core 0
d_ch2	Marker bit 1 (MARK1)	core 0
d_ch3	Marker bit 2 (MARK2)	core 1
d_ch4	Marker bit 3 (MARK3)	core 1

where $d_1, d_2 \in \{0, 1\}$ are the sampled digital states. The packed marker array is attached to the corresponding analog waveform pair in the **Waveforms** container, so that analog and digital signals are stored in the same memory record and are replayed with intrinsic sample-level synchronisation[20].

This two-channel-per-core organisation maps naturally onto IQ modulation: the in-phase (I) and quadrature (Q) components of a microwave pulse are assigned to the two outputs of a single core, guaranteeing deterministic relative timing and phase between them. In the microwave upconversion architecture described in Chapter 5, these two outputs drive the I and Q ports of the external mixer, while the marker channels of the same core provide the synchronised gating signals for the rest of the experiment.

Channel activation states requested by Qudi are applied at upload time: samples belonging to deactivated analog channels are zeroed, and markers of deactivated digital channels are cleared, so that the physical outputs faithfully reflect the activation configuration selected in the user interface.

4.3.6 Implementation of Gaussian and DRAG Pulses

Shaped pulses enter the framework at the level of Qudi’s analog envelope functions and are carried through the module transparently as sampled arrays. Two envelope families were implemented and validated in this work.

Gaussian envelope. A Gaussian pulse of peak amplitude A , centre t_0 , and standard deviation σ is described by

$$\Omega_I(t) = A \exp\left[-\frac{(t - t_0)^2}{2\sigma^2}\right]. \quad (4.2)$$

Compared with a rectangular pulse of equal area, the smooth rise and fall of the Gaussian strongly suppress the high-frequency spectral tails associated with sharp edges, reducing off-resonant excitation of neighbouring transitions. Because the spin rotation angle is set by the time-integrated Rabi frequency, the amplitude of a shaped π or $\pi/2$ pulse must be renormalised so that $\int \Omega_I(t) dt$ equals the required rotation area; this normalisation is applied when the envelope parameters are defined, before sampling[13].

DRAG correction. The DRAG (Derivative Removal by Adiabatic Gate) scheme supplements the in-phase Gaussian envelope with a quadrature component proportional to its time derivative:

$$\Omega_Q(t) = \beta \frac{d\Omega_I(t)}{dt} = -\beta \frac{(t - t_0)}{\sigma^2} \Omega_I(t), \quad (4.3)$$

where β is a dimensionless calibration parameter. The derivative component cancels, to first order, the leakage and phase errors induced by off-resonant coupling to spectator levels. In the implementation, the I and Q envelopes of Eqs. (4.2) and (4.3) are sampled onto the two analog channels of one AWG core; the shared-core architecture described in Section 4.3.5 then guarantees that the derivative channel is replayed with exact sample-level alignment to the primary envelope, which is essential since any relative delay converts the intended quadrature correction into an amplitude distortion[21].

4.3.7 Sequence Compilation and Upload

The complete upload procedure executed by `load_waveform()` proceeds as follows:

1. The buffered ($N \times 4$) analog and digital matrices are length-normalised to the hardware granularity (minimum 32 samples, multiple of 16) by zero-padding.
2. Channel activation masks are applied, zeroing inactive analog channels and clearing inactive markers.
3. For each AWG core, the corresponding analog pair is extracted and the two digital channels are packed into the per-sample marker word.

4. A deterministic SEQC program matching the waveform length and marker configuration is generated and compiled on the core via `load_sequencer_program()`; the module blocks until the core reports ready.
5. The waveform data and markers are assembled into a `Waveforms` container bound to table index 0 and transferred with `write_to_waveform_memory()`.
6. The core is left disabled; execution starts only when `pulser_on()` subsequently enables the signal outputs and the sequencers of both cores.

Every step is wrapped in defensive error handling: failures are logged with their originating operation, invalid buffer shapes are rejected with explicit diagnostics, and the module falls back to a safe idle sequencer program (an infinite `wait()` loop) whenever assets are cleared. This design ensures that malformed uploads cannot leave the instrument in an undefined state, and that error conditions are reported to the Qudi log in a form directly usable for debugging[20].

The result is a robust, fully automated pipeline in which a pulse sequence designed graphically in Qudi is sampled, normalised, compiled, uploaded, and executed on the HDAWG without any manual interaction with the vendor software.

4.3.8 Instrument Constraints

Qudi’s higher-level logic queries each pulser for a constraints object that declares the admissible ranges of every operating parameter, allowing the graphical interface and sampling logic to validate sequences before they are generated. The developed module constructs this object in `_build_constraints()`, reporting the values summarised in Table 4.2.

Table 4.2: Principal constraints reported by the developed module to the Qudi framework.

Parameter	Value / range	Notes
Sample rate	fixed at device clock	read from hardware
Analog amplitude	0.1 – 5.0 V _{pp}	per channel
Analog offset	–3.0 – +3.0 V	per channel
Digital high level	3.3 V	TTL markers
Waveform length	32 – 10 ⁶ samples	step of 16
Waveform number	1 – 32000	

Reporting the sample rate as a fixed value pinned to the hardware clock (rather than a user-settable range) reflects the fact that the module reads the sample clock directly

from the instrument at activation and does not attempt to reprogram it, ensuring that the timing assumed by Qudi always matches the timing actually realised by the hardware[20].

5 Microwave Up-Conversion Setup

This chapter describes the theoretical principles and experimental implementation of the microwave up-conversion system developed in this work. Since the arbitrary waveform generator (AWG) employed in the experimental setup cannot directly synthesize microwave signals at the frequencies required to manipulate NV centre spins, an external frequency conversion stage was implemented. The combination of a local oscillator (LO), an IQ mixer, a microwave amplifier, and the associated control electronics enables the generation of arbitrary microwave pulses with independently controllable amplitude and phase. The following sections first introduce the theoretical concepts underlying microwave up-conversion before describing the hardware architecture and the implementation adopted in this work.

5.1 Theory of Up-Conversion

Microwave up-conversion is a widely used technique in communication systems and quantum control experiments whenever low-frequency waveform generators must produce signals at much higher carrier frequencies[24]. Rather than generating the microwave signal directly, the desired waveform is first synthesized at baseband and then translated to the target frequency through frequency mixing with a stable local oscillator. This architecture preserves the flexibility of digital waveform generation while providing access to the gigahertz frequencies required for coherent spin manipulation[1, 3].

5.1.1 Need for Up-Conversion in NV-Based Experiments

Coherent control of the negatively charged nitrogen-vacancy (NV) centre relies on the application of resonant microwave fields that induce transitions between the ground-state spin sublevels. At zero magnetic field, the electronic spin triplet exhibits a zero-field splitting of approximately 2.87 GHz between the $m_s = 0$ and $m_s = \pm 1$ states. When an external magnetic field is applied along the NV axis, the Zeeman interaction lifts the

degeneracy of the $m_s = \pm 1$ states, shifting the transition frequencies according to the magnetic field strength. Consequently, microwave sources capable of generating stable signals in the GHz range are required for pulsed ODMR, Rabi oscillations, spin echo sequences, and other coherent control experiments [1, 3].

Modern arbitrary waveform generators provide excellent timing resolution together with precise digital control of waveform amplitude, phase and pulse shape. In this work, the microwave pulse envelopes were generated using a Zurich Instruments HDAWG, which supports the implementation of both conventional rectangular pulses and shaped envelopes such as Gaussian and DRAG pulses. However, the analog output of the HDAWG is limited to baseband frequencies and cannot directly generate microwave carriers around 2.87 GHz [20].

For this reason, an external frequency translation stage is required. Instead of attempting to synthesize the microwave carrier directly, the HDAWG generates two low-frequency baseband signals that contain the desired pulse information. These signals are subsequently mixed with a continuous-wave microwave signal provided by a local oscillator. The frequency conversion process transfers the information contained in the baseband waveforms onto the microwave carrier while preserving the programmed amplitude and phase modulation. As a result, complex microwave pulse sequences can be generated without sacrificing the flexibility offered by digital waveform synthesis[24, 20].

Beyond extending the accessible frequency range, the up-conversion approach also enables independent control of the microwave field amplitude and phase through the two baseband channels. This capability is particularly important for high-fidelity quantum control, where shaped microwave pulses are used to reduce spectral leakage, minimise off-resonant excitation, and improve gate performance. The implementation adopted in this work therefore combines the versatility of digital waveform generation with the frequency stability of a microwave local oscillator, providing a flexible platform for coherent manipulation of NV centre spins[24].

5.1.2 Mixer Operation and Sideband Generation

The frequency translation is performed using a microwave mixer, a nonlinear device that combines two input signals to produce new frequency components. If two sinusoidal signals with angular frequencies ω_{LO} and ω_{IF} are applied to an ideal multiplier, the output can be expressed as

$$\cos(\omega_{\text{LO}}t) \cos(\omega_{\text{IF}}t) = \frac{1}{2} [\cos((\omega_{\text{LO}} + \omega_{\text{IF}})t) + \cos((\omega_{\text{LO}} - \omega_{\text{IF}})t)]. \quad (5.1)$$

Equation 5.1 shows that the mixing process produces two new frequency components located above and below the local oscillator frequency. These are commonly referred to as the upper sideband (USB)

$$f_{\text{USB}} = f_{\text{LO}} + f_{\text{IF}},$$

and the lower sideband (LSB)

$$f_{\text{LSB}} = f_{\text{LO}} - f_{\text{IF}}.$$

In a conventional mixer, both sidebands are generated simultaneously, together with a residual component at the local oscillator frequency caused by imperfect isolation. Depending on the application, unwanted sidebands are usually removed using microwave filters. However, filtering becomes increasingly difficult when the intermediate frequency is relatively small compared with the carrier frequency[24].

To overcome this limitation, IQ mixers employ two independent baseband inputs referred to as the in-phase (I) and quadrature (Q) channels. The local oscillator signal is internally divided into two paths with a relative phase difference of 90° , while the I and Q signals are also maintained in quadrature. By properly adjusting the relative amplitudes and phases of these two channels, one of the sidebands can be made to interfere destructively while the other adds constructively. This technique, known as single-sideband (SSB) modulation, enables the generation of a microwave signal containing only the desired frequency component [24].

In practice, perfect sideband suppression is limited by amplitude imbalance, phase errors, DC offsets, finite mixer isolation and component tolerances. Consequently, experimental calibration is generally required to minimise local oscillator leakage and maximise unwanted sideband rejection. The optimisation procedure adopted for the experimental setup developed in this thesis is presented in the next Chapter, where the influence of phase imbalance and DC offset on the output spectrum is experimentally characterised.

Experimental Illustration of Sideband Generation

Figure 5.1 presents a representative spectrum measured during the development of the microwave up-conversion chain. The spectrum was acquired using the spectrum analyser while the IQ mixer was driven by the local oscillator and the baseband signals generated by the HDAWG. Three main spectral components can be distinguished: the local oscillator leakage located at the carrier frequency, together with the upper and lower sidebands generated by the mixing process.

The upper sideband corresponds to the desired microwave signal that is used to excite the NV centres, whereas the lower sideband represents the image frequency produced by the frequency conversion process. In an ideal single-sideband modulation scheme, the unwanted sideband would be completely suppressed as already said. In practice, however, finite amplitude imbalance, phase mismatch and DC offsets between the I and Q channels lead to residual sideband components, which can be observed in the measured spectrum.

The measurements shown in Fig. 5.1 illustrate the physical principles discussed above and motivate the calibration procedure presented later in this thesis. By optimising the amplitude balance, relative phase and DC offsets of the IQ channels, both the local oscillator leakage and the unwanted sideband can be significantly reduced, resulting in a spectrally pure microwave output suitable for coherent manipulation of NV centre spins.




Figure 5.1: Spectrum measured at the output of the IQ mixer illustrating the frequency components generated during microwave up-conversion. The dominant peak corresponds to the desired upper sideband, while the residual local oscillator leakage and lower sideband arise from the finite isolation and imbalance of the practical IQ mixer.

5.1.3 Bandwidth Limitation of the Up-Conversion Chain

The performance of the complete microwave generation system is determined not only by the mixer itself but also by the bandwidth limitations of every component in the signal chain. Since the final microwave pulse is generated through the combination of several independent devices, the achievable frequency range and waveform fidelity are constrained by the narrowest bandwidth within the system[24].

The first limitation originates from the arbitrary waveform generator. Although the HDAWG offers a sampling rate of up to 2.4 GSa/s, its analog outputs operate at baseband

and are intended to generate intermediate-frequency (IF) signals rather than microwave carriers. Consequently, the waveform bandwidth that can be faithfully synthesized is limited by the analog output stage of the instrument[20].

The second limitation is introduced by the IQ mixer. Commercial microwave mixers are designed to operate over specific ranges of local oscillator, radio-frequency (RF), and intermediate-frequency (IF) signals. Efficient frequency conversion requires that all operating frequencies remain within the specifications provided by the manufacturer. Outside these ranges, conversion loss increases while sideband suppression and amplitude balance deteriorate[25].

Additional constraints arise from the microwave amplifier, coaxial cables, connectors and antenna. Each component exhibits a finite frequency response together with insertion losses that influence the power ultimately delivered to the diamond sample. At microwave frequencies, cable attenuation, impedance mismatches and connector reflections become increasingly important and may distort the generated waveform if not properly considered during system design.

Despite these practical limitations, the bandwidth of the up-conversion chain implemented in this work is fully compatible with the frequency range required for NV-centre experiments. The local oscillator frequency, intermediate-frequency waveform and microwave amplification stage were selected such that arbitrary pulse envelopes generated by the HDAWG could be translated efficiently to the desired microwave carrier while maintaining sufficient spectral purity and output power for coherent spin manipulation. The individual hardware components used to realise this architecture are described in the following section.

5.2 Hardware Chain

The microwave up-conversion chain implemented in this work was designed to translate the low-frequency pulse envelopes generated by the HDAWG to the microwave frequency range required for driving NV-centre spin transitions. The system consists of five main stages: the HDAWG baseband output, the microwave local oscillator, the external IQ mixer, the microwave amplification stage, and the antenna used to deliver the microwave field to the diamond sample. This architecture separates pulse shaping from carrier generation: the temporal envelope and phase are programmed digitally at baseband, while the microwave carrier is supplied by a stable external source. The resulting signal is then amplified and coupled to the sample.

5.2.1 HDAWG Output Stage

The first stage of the up-conversion chain is the Zurich Instruments HDAWG, which generates the low-frequency signals applied to the I and Q inputs of the external IQ mixer. In the implemented configuration, two analog output channels of the HDAWG were used as the in-phase and quadrature components of the baseband waveform. The HDAWG provides programmable arbitrary waveforms with high timing resolution, making it suitable for producing arbitrary, rectangular, Gaussian, and DRAG-like microwave pulse envelopes before frequency translation [20].

The role of the HDAWG in this setup is not to generate the microwave carrier directly, but to define the time-dependent envelope of the microwave control pulse. This distinction is important because the NV spin transition lies close to 2.87 GHz, whereas the useful analog output bandwidth of the arbitrary waveform generator is intended for baseband and intermediate-frequency waveform generation rather than direct microwave synthesis. Therefore, the HDAWG outputs encode the desired pulse amplitude, phase, and shape at low frequency, while the following IQ mixer stage translates this information to the microwave domain [20, 24].

During the experiments, the HDAWG outputs were connected to the I and Q ports of the mixer through coaxial cables and attenuators. The attenuation reduced the baseband drive level delivered to the mixer and helped keep the I/Q inputs within a controlled operating range. This was particularly important during the calibration of local-oscillator leakage and sideband suppression, since excessive drive levels or channel imbalance can increase unwanted spectral components at the mixer output.

The use of two synchronized AWG channels also allowed direct control of the relative amplitude and phase of the I and Q components. By programming the two channels with a nominal 90° phase relationship, the generated baseband waveform could be used for single-sideband up-conversion. Small corrections to the channel amplitude, phase, and DC offset were later used to compensate practical imperfections of the mixer and cables.

5.2.2 Local Oscillator Source

The local oscillator provides the microwave reference tone required by the IQ mixer. In this work, the LO was supplied by a microwave signal generator operated in continuous-wave mode. Its frequency was selected such that, after mixing with the intermediate-frequency waveform generated by the HDAWG, the desired sideband was positioned at the NV spin transition frequency. For an intermediate frequency f_{IF} , the two possible output sidebands are located at $f_{\text{LO}} + f_{\text{IF}}$ and $f_{\text{LO}} - f_{\text{IF}}$, as discussed in the previous section [24].

A stable local oscillator is essential because its phase noise and frequency stability are directly transferred to the up-converted microwave signal. In coherent spin-control experiments, fluctuations of the microwave carrier can reduce pulse reproducibility and degrade the phase coherence of the applied control field. For this reason, the LO was treated as the frequency reference of the microwave chain, while the HDAWG supplied only the programmable baseband modulation.

In the experimental configuration, the LO power was adjusted to satisfy the drive requirement of the IQ mixer.

5.2.3 External IQ Mixer

Frequency translation was performed using a Polyphase Microwave AM0350A quadrature modulator. This component was selected because it allows direct conversion of baseband I/Q modulation to an RF signal over a frequency range compatible with the NV microwave transition. According to the manufacturer specifications, the AM0350A operates over an LO/RF frequency range of 300–5000 MHz, accepts a nominal LO power of +5 dBm, and provides typical sideband suppression and LO leakage values of approximately -40 dBc and -35 dBm, respectively [25].



Figure 5.2: Polyphase Microwave AM0350A quadrature modulator used as the external IQ mixer in the up-conversion chain. The device converts baseband I/Q signals to an RF output and is specified for operation from 300 MHz to 5000 MHz. Figure adapted from the manufacturer’s datasheet [25].

The mixer receives three input signals: the LO tone, the I-channel baseband waveform, and the Q-channel baseband waveform. Internally, the LO is split into two paths with a quadrature phase relation, and the baseband signals modulate these two LO components. The resulting RF output is the vector sum of the two modulation paths. Ideally, this

configuration allows one sideband to be enhanced while the opposite sideband is cancelled by destructive interference. In practice, the degree of cancellation is limited by finite amplitude imbalance, quadrature phase error, DC offsets, and imperfect port isolation [24, 25].

The manufacturer performance data shown in Fig. 5.3 summarize the main imperfections relevant to this work. The conversion loss determines the reduction in power between the baseband/LO drive and the generated RF output. The amplitude imbalance and quadrature phase error indicate how closely the two modulation paths satisfy the ideal I/Q condition. The LO leakage and sideband suppression curves provide the expected level of residual unwanted components at the RF output. These quantities are important because they directly affect the spectral purity of the microwave pulse delivered to the NV sample [25].




Figure 5.3: Typical performance characteristics of the AM0350A quadrature modulator, including conversion loss, input IP3, LO leakage, amplitude imbalance, quadrature phase error, sideband suppression, I/Q baseband response, and output 1-dB compression point. Figure adapted from the manufacturer’s datasheet [25].

In the laboratory setup, the HDAWG I and Q outputs were connected to the corresponding baseband ports of the AM0350A. The RF output of the mixer contained the desired sideband, the residual image sideband, and a finite amount of LO leakage. The relative strength of these components depended on the programmed I/Q phase relation, the relative channel amplitudes, and the DC offsets applied to the two baseband inputs. The calibration of these parameters was therefore necessary to improve the spectral purity of the generated microwave signal.

5.2.4 Microwave Amplifier

After up-conversion, the RF output of the IQ mixer was routed to a microwave amplifier. The amplifier was required because the mixer output power was not sufficient to efficiently drive coherent spin transitions at the diamond sample after accounting for conversion loss, cable losses, attenuators, and imperfect antenna coupling. The amplifier increased the microwave power before delivery to the antenna while preserving the pulse timing imposed by the HDAWG waveform.

5.2.5 Antenna Coupling to the Diamond Sample

The amplified microwave signal was delivered to the diamond sample through a microwave antenna positioned close to the region addressed by the confocal microscope. The purpose of the antenna is to generate an oscillating magnetic field at the NV position, with a field component transverse to the NV quantization axis. This transverse microwave magnetic field drives transitions between the $m_s = 0$ and $m_s = \pm 1$ spin states when the microwave frequency is resonant with the corresponding transition.

The antenna placement was chosen as a compromise between microwave coupling efficiency and optical access to the sample. Bringing the antenna closer to the diamond increases the microwave magnetic field at the NV centre, but the antenna must not obstruct the optical excitation and fluorescence collection paths. In addition, the antenna geometry and its distance from the sample influence the microwave field distribution and therefore the effective Rabi frequency observed in pulsed measurements.

Since the antenna and sample geometry were part of the experimental implementation, their final alignment was optimized empirically. The optimization criterion was the observation of a clear ODMR contrast and, later, coherent Rabi oscillations under pulsed microwave excitation. This procedure ensured that the up-converted microwave signal was effectively coupled to the NV centres under study.

5.2.6 Complete Signal Chain

The complete microwave up-conversion chain is summarized in Fig. 5.4. The HDAWG generates two synchronized low-frequency waveforms corresponding to the I and Q components. These are applied to the baseband inputs of the AM0350A IQ mixer. A continuous-wave microwave signal from the local oscillator drives the LO port of the mixer. The mixer output contains the up-converted microwave pulse, which is subsequently amplified and delivered to the diamond sample through the microwave antenna.

This configuration provides a flexible microwave-control platform for NV-centre experiments. The HDAWG determines the temporal structure of the pulse, the local oscillator




Figure 5.4: Schematic representation of the complete microwave up-conversion chain.

defines the carrier frequency, and the IQ mixer combines both inputs to generate the final microwave waveform. As a result, the setup can produce rectangular pulses as well as shaped envelopes while maintaining phase control through the I/Q representation. The following section describes how the phase and amplitude of the microwave signal are controlled through the baseband components.

5.2.7 IQ Mixer Calibration Procedure

The performance of an IQ up-conversion system strongly depends on the precise balance between the in-phase (I) and quadrature (Q) baseband signals. In an ideal quadrature modulator, equal amplitudes, an exact phase difference of 90° , and zero DC offsets would completely suppress the unwanted sideband and eliminate local oscillator (LO) leakage. In practice, however, imperfections in the waveform generator, mixer, cables and connectors introduce small amplitude and phase mismatches that degrade the spectral purity of the generated microwave signal. As a result, an experimental calibration procedure is required before performing any coherent spin manipulation experiments [24, 25].

The calibration was carried out using the Zurich Instruments LabOne software environment, which provides complete control over the waveform generation and output parameters of the HDAWG. Figure 5.5 shows the interface used throughout the calibration procedure. The upper part of the window contains the configuration panels for the waveform generation process. The oscillator and sine generator modules define the baseband frequencies and their relative phase, while the waveform generator section controls the uploaded pulse sequence executed by the HDAWG. In addition, the wave output panel allows independent adjustment of the output amplitudes, DC offsets, delays and voltage ranges of each analog channel. These parameters were modified iteratively during the calibration in order to compensate for imperfections in the analog signal chain.

The lower panel of the interface displays the waveform viewer, where the generated baseband signals can be inspected before being transmitted to the IQ mixer. In the example shown in Fig. 5.5, two sinusoidal waveforms with a relative phase difference of approximately 90° are generated on the two analog output channels. This visualization provides an immediate verification that both channels are correctly synchronized and that the programmed amplitude and phase relationships correspond to the desired IQ modulation scheme before the signals reach the mixer.



Figure 5.5: LabOne interface used during the calibration of the HDAWG and the AM0350A IQ mixer. The upper panels provide access to the oscillator configuration, waveform generation, output amplitudes, DC offsets and voltage ranges of the analog channels, while the waveform viewer at the bottom displays the synchronized I and Q baseband signals before frequency up-conversion. These controls were used to iteratively optimize the IQ mixer calibration by minimizing the local oscillator leakage and the unwanted image sideband observed on the spectrum analyser.

During the calibration procedure, the microwave output of the IQ mixer was continuously monitored using a spectrum analyser while the waveform parameters were adjusted within the LabOne interface. The amplitudes of the I and Q channels, their relative phase, and the DC offsets were modified iteratively until the local oscillator leakage and the unwanted image sideband were minimized. Because these quantities are highly sensitive to even small mismatches between the two channels, several iterations were generally required before reaching the optimum operating point. The final calibration parameters obtained through this process were subsequently used for all microwave pulse sequences presented in this thesis.

This calibration step is fundamental for high-fidelity microwave generation. Without proper balancing of the I and Q channels, residual local oscillator leakage and image frequencies would introduce unwanted spectral components that could drive off-resonant

transitions in the NV centre or reduce the fidelity of the applied pulse sequence. The effectiveness of the calibration is experimentally demonstrated in the following chapter, where the measured suppression of the local oscillator leakage and unwanted sideband is quantified using the spectrum analyser.

5.3 Phase Signal and IQ Control

5.3.1 In-Phase and Quadrature Components

Quadrature modulation relies on representing a microwave signal as the combination of two orthogonal baseband components known as the in-phase (I) and quadrature (Q) signals. The I component is multiplied by a cosine carrier, while the Q component is multiplied by a sine carrier shifted by 90° . The superposition of these two signals allows independent control of both the amplitude and the phase of the resulting microwave field after frequency up-conversion. This principle forms the basis of IQ modulation and enables the generation of arbitrary microwave waveforms without mechanically or electronically modifying the local oscillator frequency [24].

In the experimental setup developed in this work, the I and Q baseband waveforms are generated by two synchronized analog output channels of the Zurich Instruments HDAWG. These low-frequency signals are applied directly to the corresponding I and Q ports of the external AM0350A IQ mixer, while the microwave carrier is simultaneously provided through the local oscillator input. The mixer combines the three inputs to produce the desired microwave signal at the radio-frequency output. Since both baseband channels originate from the same arbitrary waveform generator, their relative timing and phase are inherently synchronized, providing excellent reproducibility during pulse generation [20, 25].

The flexibility of this architecture makes it possible to generate not only conventional rectangular microwave pulses but also shaped pulse envelopes that improve coherent spin control. In particular, the implementation developed in this work supports Gaussian excitation pulses together with Derivative Removal by Adiabatic Gate (DRAG) correction. While the Gaussian envelope is applied to the in-phase channel, the DRAG correction is generated on the quadrature channel. The quadrature component is proportional to the time derivative of the Gaussian envelope and compensates for unwanted spectral components that may otherwise induce off-resonant excitation. Although originally introduced for superconducting qubits, the DRAG formalism has become a widely adopted pulse-shaping technique whenever high-fidelity coherent control is required [21, 8].

The use of synchronized I and Q channels therefore provides considerably greater flex-




Figure 5.6: Oscilloscope measurement of the baseband waveforms generated by the HDAWG. Channel 1 (yellow) carries the Gaussian pulse applied to the in-phase (I) input of the IQ mixer, while Channel 2 (green) contains the corresponding DRAG waveform applied to the quadrature (Q) input. The marker output generated by the HDAWG provides a synchronization signal for external instruments during pulse sequence execution.

ibility than conventional single-channel pulse generation. Instead of controlling only the microwave amplitude, the experimental system is capable of independently programming the amplitude, relative phase and temporal shape of the microwave field. This capability is essential for implementing advanced pulse sequences and for performing the calibration procedures presented later in this thesis, where the relative phase, amplitude balance and DC offsets of the two channels are optimized to suppress local oscillator leakage and unwanted sidebands.

5.3.2 Amplitude and Phase Control

One of the principal advantages of IQ modulation is the possibility of independently controlling both the amplitude and the phase of the generated microwave signal. Unlike conventional microwave generation, where the output phase is fixed by the local oscillator and only the signal amplitude can be varied, an IQ modulator combines two orthogonal baseband signals to synthesize a microwave field with an arbitrary complex amplitude. This additional degree of freedom is particularly valuable in quantum control experiments, where accurate manipulation of the microwave phase is essential for implementing coherent pulse sequences and minimizing systematic control errors [24].

The microwave signal at the output of an ideal IQ mixer can be expressed as

$$V_{\text{RF}}(t) = I(t) \cos(\omega_{\text{LO}}t) - Q(t) \sin(\omega_{\text{LO}}t), \quad (5.2)$$

where $I(t)$ and $Q(t)$ denote the in-phase and quadrature baseband waveforms generated by the HDAWG, while ω_{LO} represents the angular frequency of the local oscillator. The resulting microwave signal can alternatively be written in polar form as

$$V_{\text{RF}}(t) = A(t) \cos(\omega_{\text{LO}}t + \phi(t)), \quad (5.3)$$

where the instantaneous amplitude and phase are given by

$$A(t) = \sqrt{I^2(t) + Q^2(t)}, \quad (5.4)$$

and

$$\phi(t) = \tan^{-1} \left(\frac{Q(t)}{I(t)} \right). \quad (5.5)$$

These relations show that the microwave amplitude is determined by the vector magnitude of the two baseband signals, whereas the relative amplitudes of the I and Q channels define the microwave phase. Consequently, modifying either waveform immediately changes the properties of the generated RF signal without altering the frequency of the local oscillator [24].

The capability to independently adjust the amplitude and phase of the microwave field constitutes one of the main advantages of the adopted up-conversion architecture. Besides enabling efficient suppression of unwanted spectral components, it provides the flexibility required for implementing advanced pulse-shaping techniques and coherent control protocols, which form the basis of the experimental investigations presented in the following chapters.

The independent control of the microwave amplitude and phase provided by the IQ modulation architecture also enables the implementation of arbitrary spin rotations on the Bloch sphere. In coherent spin control, the pulse amplitude determines the Rabi frequency and therefore the rotation angle applied to the spin state, while the microwave phase defines the rotation axis within the equatorial plane of the Bloch sphere. By appropriately programming the I and Q baseband signals, rotations about different axes can be generated without modifying the experimental hardware or the local oscillator frequency. This capability forms the physical basis for implementing single-qubit quantum gates and advanced pulse sequences used in quantum sensing and quantum information processing [7, 8, 26].

6 Experimental Results

This chapter presents the experimental validation of the microwave control system developed in this work. After identifying an isolated NV centre through confocal imaging and second-order autocorrelation measurements, pulsed ODMR experiments were performed using different microwave pulse envelopes generated by the HDAWG. The objective was to compare conventional rectangular microwave pulses with shaped Gaussian-DRAG and DRAG pulses while keeping the remaining experimental conditions unchanged. The measurements include pulsed ODMR spectra, Rabi oscillations, and the corresponding numerical fits used to extract resonance frequencies, linewidths and coherent spin dynamics.

The experimental methodology follows the pulsed ODMR approach commonly employed for single NV centres, where microwave excitation is separated from optical initialization and readout to minimize power broadening and obtain accurate measurements of the spin transition properties [3, 13].

6.1 CW ODMR Characterization

Before performing pulsed measurements, continuous-wave optically detected magnetic resonance was used to identify the electron-spin resonance frequencies of the selected NV centre. During the CW ODMR measurement, continuous optical excitation polarizes the electronic spin preferentially into the $m_s = 0$ state, while a microwave field is swept across the expected ground-state transition frequencies. When the microwave frequency becomes resonant with an allowed spin transition, population is transferred to the $m_s = \pm 1$ manifold, which exhibits a lower fluorescence rate and therefore appears as a dip in the detected photoluminescence signal [1, 3, 13].

Figure 6.1 shows the CW ODMR spectrum acquired from the single NV centre selected for the subsequent pulse-shaping experiments. The upper panel displays the averaged fluorescence signal as a function of microwave frequency together with the fitted resonance

profile. The lower panel presents the individual frequency scans as a two-dimensional intensity map, where each row corresponds to one acquisition cycle. The persistence of the resonance features across the repeated scans confirms that the selected transition frequencies remained stable during the measurement.

The CW ODMR spectrum was used to determine the approximate resonance region for the subsequent pulsed ODMR and Rabi measurements. However, because the laser and microwave fields are applied simultaneously in the CW protocol, the measured linewidth can be broadened by both optical pumping and microwave driving. For this reason, the more quantitative comparison between rectangular, Gaussian and DRAG pulse envelopes was carried out using pulsed excitation, where optical initialization, microwave manipulation and optical readout are temporally separated [13].




Figure 6.1: Continuous-wave ODMR characterization of the selected single NV centre. The upper panel shows the averaged fluorescence signal and fitted resonance profile as a function of microwave frequency. The lower panel displays the individual repeated scans, demonstrating the stability of the resonance features throughout the acquisition.

6.1.1 Rectangular and Gaussian Excitation Profiles

The spectral response of a microwave pulse is determined by its temporal envelope through the Fourier transform. Consequently, different pulse shapes exhibit different excitation

profiles, which directly affect their ability to selectively manipulate spin transitions. Before comparing the experimental performance of rectangular and Gaussian pulses, it is therefore useful to examine their theoretical excitation characteristics.




Figure 6.2: Comparison of rectangular and Gaussian excitation profiles.

Figure 6.2 compares the simulated excitation profiles of a rectangular pulse and a Gaussian pulse reported by Vandersypen and Chuang [7]. In both cases, the horizontal axis represents the frequency detuning from resonance, while the vertical axis corresponds to the normalized spin response after application of a π pulse.

The rectangular pulse, shown in Fig. 6.2(a), exhibits the well-known sinc-shaped excitation profile resulting from the Fourier transform of a finite-duration constant-amplitude waveform. Although the central lobe produces the desired π rotation at resonance, several secondary side lobes are also present on both sides of the resonance frequency. These side lobes indicate that spins with frequencies detuned from the target resonance can still undergo partial excitation. Consequently, rectangular pulses possess limited frequency selectivity and may unintentionally excite neighbouring transitions when multiple resonances are closely spaced.

In contrast, the Gaussian pulse shown in Fig. 6.2(b) produces a significantly smoother excitation profile. The excitation is concentrated around the resonance frequency, while the secondary side lobes observed for the rectangular pulse are strongly suppressed. This reduction of the side lobes decreases off-resonant excitation. As a result, Gaussian envelopes are widely employed in coherent spin control, where selective excitation is required without perturbing nearby transitions.

The comparison between these excitation profiles illustrates the principal advantage of Gaussian pulse shaping. Rather than modifying the excitation at the resonance frequency itself, the Gaussian envelope redistributes the spectral content of the microwave pulse, suppressing unwanted spectral components that give rise to off-resonant excitation. This property provides the motivation for investigating Gaussian pulse envelopes in the experimental measurements presented in the following sections.

6.1.2 Calibration of the π -Pulse Duration Using Rabi Oscillations

Before performing pulsed ODMR measurements, the microwave pulse duration must be calibrated to produce a π rotation of the NV electronic spin. This calibration was carried out experimentally by performing Rabi oscillation measurements for each pulse envelope separately.

A Rabi experiment consists of repeatedly varying the microwave pulse duration while keeping all other parameters fixed and recording the resulting photoluminescence (PL). The oscillatory behaviour originates from coherent rotations of the two-level spin system under resonant microwave excitation. The oscillation frequency, commonly referred to as the Rabi frequency Ω_R , determines the pulse duration required to perform specific quantum rotations. In particular, the duration corresponding to the first population inversion is selected as the π pulse, while half of this value corresponds to a $\pi/2$ pulse used in more advanced pulse sequences [3, 13].

Since the pulse envelope changes the temporal distribution of the microwave amplitude, the effective driving strength is also modified. Consequently, the rectangular and Gaussian–DRAG pulses require independent Rabi calibrations and generally produce different π -pulse durations even when generated using similar output amplitudes.

Figure 6.3a presents the Rabi oscillations obtained using the conventional rectangular microwave pulse. A damped sinusoidal model was fitted to the experimental data in order to extract the oscillation frequency, from which the corresponding π -pulse duration was calculated.

Similarly, Figure 6.3b shows the Rabi oscillations measured using the Gaussian–DRAG pulse configuration. In this implementation, the Gaussian envelope was applied to the in-phase (I) channel, while the derivative Gaussian (DRAG) component was simultaneously applied to the quadrature (Q) channel using the HDAWG waveform generator. The extracted oscillation frequency differs from that obtained using the rectangular pulse because the Gaussian envelope distributes the microwave power over the pulse duration instead of

maintaining a constant amplitude. The resulting π -pulse duration was subsequently used for all Gaussian–DRAG pulsed ODMR measurements presented in the following section.

These measurements establish a consistent calibration procedure for both pulse implementations, ensuring that the comparison between rectangular and Gaussian–DRAG pulsed ODMR experiments is performed under equivalent spin-rotation conditions.




Figure 6.3: Comparison of the Rabi oscillation measurements used to calibrate the microwave π -pulse duration for the two pulse implementations investigated in this work. For the Gaussian–DRAG implementation, the Gaussian envelope was generated on the in-phase (I) channel and the DRAG component on the quadrature (Q) channel of the HDAWG. The result is an Upconverted gaussian envelop. The oscillation frequency extracted from each fitted curve was used to determine the appropriate π -pulse duration prior to performing the corresponding pulsed ODMR measurements.

To calculate the duration, we got

$$I(\tau) = I_0 + A \cos(2\pi f_R \tau + \phi) e^{-\tau/T_R}, \quad (6.1)$$

where I_0 is the signal offset, A is the oscillation amplitude, f_R is the Rabi frequency, ϕ is the phase offset and T_R is the decay time accounting for dephasing during the measurement. The fitted Rabi frequency determines the period of the coherent oscillation,

$$T_R = \frac{1}{f_R}. \quad (6.2)$$

Since a complete Rabi period corresponds to a 2π rotation of the Bloch vector, the duration required to produce a π rotation is obtained directly as

$$\tau_\pi = \frac{T_R}{2} = \frac{1}{2f_R}, \quad (6.3)$$

while the corresponding $\pi/2$ pulse duration is

$$\tau_{\pi/2} = \frac{1}{4f_R}. \quad (6.4)$$

The extracted π -pulse duration was then used as the microwave pulse length in all subsequent pulsed ODMR experiments.

6.1.3 Pulsed ODMR Comparison

Following the calibration of the microwave π -pulse duration from the Rabi oscillation experiments, pulsed ODMR measurements were performed using the standard pulse sequence described by Dréau *et al.* [13]. Each measurement cycle begins with a laser pulse that optically initializes the NV centre into the $m_s = 0$ spin state. This is followed by a resonant microwave pulse of fixed duration T_π , which rotates the spin population. Finally, a second laser pulse simultaneously reads out the spin state through spin-dependent photoluminescence and reinitializes the NV centre for the next measurement cycle.

During the experiment, the microwave pulse duration remained fixed at the calibrated π -pulse value obtained from the corresponding Rabi oscillation measurement. The microwave frequency was then swept across the expected resonance while recording the fluorescence intensity after each sequence, producing the pulsed ODMR spectrum. This procedure was repeated for both the rectangular microwave pulse and the Gaussian–DRAG pulse implementation, allowing a direct comparison of the two pulse envelopes under identical experimental conditions.

Figure 6.4(a) shows the pulsed ODMR spectrum obtained using conventional rectangular microwave pulses. As expected from the excitation profile discussed in the previous section, the rectangular pulse excites not only the resonant transition but also nearby frequency components owing to the presence of sinc side lobes in the frequency domain. The three hyperfine transitions associated with the ^{14}N nuclear spin are clearly visible and can be accurately described using a three-Lorentzian model.

The corresponding measurement obtained using the Gaussian–DRAG pulse implementation is presented in Fig. 6.4(b). In this case, the Gaussian envelope was applied to the in-phase (I) channel while the derivative Gaussian (DRAG) waveform was simultaneously applied to the quadrature (Q) channel of the HDAWG. The smoother temporal envelope

suppresses the side lobes present in the rectangular excitation profile, thereby reducing out-of-resonance excitation. As a result, the spectrum exhibits a smoother resonance without the pronounced off-resonant features observed for the rectangular pulse.

These measurements demonstrate the expected trade-off between excitation selectivity and spectral resolution: rectangular pulses preserve the hyperfine structure at the expense of stronger spectral side lobes, whereas Gaussian–DRAG pulses reduce off-resonant excitation but produce a smoother and less structured resonance, consistent with the excitation profiles discussed previously [7, 13].




Figure 6.4: Comparison of pulsed ODMR measurements obtained using (a) rectangular microwave pulses and (b) Gaussian–DRAG pulses. The rectangular pulse resolves the three hyperfine transitions of the NV centre, whereas the Gaussian–DRAG pulse suppresses out-of-resonance excitation caused by spectral side lobes but does not sufficiently resolve the hyperfine structure to permit reliable three-Lorentzian fitting.

6.1.4 Implementation

The Gaussian pulse experiments were performed using the waveform generation framework developed in this work and integrated into the Zurich Instruments HDAWG through the Qudi pulse module.

For each microwave pulse, the in-phase (I) channel generated a Gaussian envelope, while the quadrature (Q) channel simultaneously generated the first derivative of the Gaussian waveform, and as already explained the generated pulse is an upconverted gaussian.

Figure 6.5 shows the waveform generated by the HDAWG inside LabOne. The blue trace corresponds to the Gaussian pulse transmitted through the I channel, whereas the red trace represents its derivative applied to the Q channel. The pulse duration and sampling

were automatically generated by the developed software module without requiring manual waveform construction.




Figure 6.5: Waveform generated by the HDAWG through the developed Qudi module. The I channel contains the Gaussian envelope, while the Q channel simultaneously generates its derivative for DRAG pulse generation.

7 Conclusions and Future Work

7.1 Summary of the Developed System

This thesis presented the design, implementation and experimental validation of a waveform generation framework enabling advanced microwave pulse control for nitrogen-vacancy (NV) centre experiments using the Zurich Instruments HDAWG within the Qudi measurement environment.

The developed software extends the existing pulse module by introducing support for arbitrary pulse envelopes while maintaining compatibility with the original Qudi pulse sequence abstraction. Starting from high-level pulse descriptions, the implemented framework automatically generates sampled waveforms, compiles them, uploads them to the HDAWG and executes the corresponding microwave pulse sequences without requiring manual intervention from the user.

The implementation supports conventional rectangular microwave pulses together with Gaussian and DRAG pulse envelopes. The generated waveforms can be assigned to independent output channels, allowing simultaneous generation of the Gaussian in-phase component and its derivative on the quadrature channel. This provides a flexible platform for coherent microwave control while remaining fully integrated into the existing experimental workflow.

The developed system was experimentally validated using a single NV centre. Continuous-wave ODMR, Rabi oscillations and pulsed ODMR measurements confirmed the correct operation of the implemented software and demonstrated that the generated waveforms can be successfully employed for coherent spin manipulation. In particular, the pulsed ODMR measurements highlight the different frequency selectivity of Gaussian and rectangular microwave pulses.

7.2 Main Engineering Contributions

The principal contribution of this work lies in the development of a complete software framework that integrates arbitrary waveform generation into Qudi while providing direct compatibility with the Zurich Instruments HDAWG.

Unlike the original implementation, which relied primarily on rectangular microwave pulses, the developed framework introduces configurable envelope functions without modifying the experimental pulse sequence description. The waveform generation, sampling, compilation and hardware upload processes are performed automatically, allowing users to switch between pulse shapes through software configuration rather than hardware-specific programming.

A second contribution is the implementation of multi-channel waveform generation, enabling synchronized output on the I and Q channels of the HDAWG. This functionality was required for implementing Gaussian and DRAG pulse generation and provides the foundation for more advanced microwave control techniques.

Finally, the developed framework was experimentally validated through coherent spin-control measurements on a single NV centre. The successful observation of Rabi oscillations and pulsed ODMR spectra confirms the correctness of both the software implementation and its integration with the experimental hardware.

7.3 Limitations

Although the developed framework successfully enables arbitrary waveform generation, several limitations remain.

The experimental validation presented in this work was restricted to Gaussian and DRAG pulse envelopes applied to relatively simple pulse sequences. More sophisticated pulse optimization techniques were not investigated.

The current implementation also assumes predefined analytical pulse shapes. Users wishing to implement custom envelopes must currently extend the software source code, limiting flexibility for rapid prototyping of new pulse designs.

Finally, the experimental characterization focused on single-qubit microwave control. The framework has not yet been evaluated for multi qubit libraries, optimal control sequences or more complex quantum algorithms.

7.4 Future Improvements

Several directions can be pursued to further extend the capabilities of the developed framework.

From a software perspective, additional analytical pulse envelopes such as Blackman, Hamming, raised-cosine and adiabatic pulse shapes could be incorporated into the waveform generation module. Support for user-defined arbitrary waveforms imported from external files would further increase the flexibility of the framework.

Future versions of the software could also include automatic pulse calibration routines, allowing π and $\pi/2$ pulse durations to be determined directly from experimental Rabi measurements. Such calibration procedures would simplify experimental operation and improve measurement reproducibility.

On the experimental side, the implemented infrastructure provides a suitable platform for investigating more advanced coherent control techniques. Systematic comparisons between rectangular, Gaussian and DRAG pulses could be performed using measurements of gate fidelity, pulse robustness and coherence preservation rather than solely linewidth comparisons.

Finally, the modular architecture developed in this work facilitates the implementation of more advanced quantum control protocols, including composite pulse sequences, dynamical decoupling methods and optimal control techniques. These extensions would broaden the applicability of the framework to quantum sensing and quantum information processing experiments involving NV centres.

References

- [1] 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.
- [2] C. L. Degen, F. Reinhard, and P. Cappellaro, “Quantum sensing,” *Reviews of Modern Physics*, vol. 89, p. 035002, 2017.
- [3] L. Rondin, J.-P. Tetienne, T. Hingant, J.-F. Roch, P. Maletinsky, and V. Jacques, “Magnetometry with nitrogen-vacancy defects in diamond,” *Reports on Progress in Physics*, vol. 77, no. 5, p. 056503, 2014.
- [4] H. Babashah, H. Shirzad, E. Losero, V. Goblots, C. Galland, and M. Chipaux, “Optically detected magnetic resonance with an open source platform,” *arXiv preprint arXiv:2205.00005*, 2022.
- [5] F. Jelezko and J. Wrachtrup, “Quantum information processing in diamond,” *Physica Status Solidi A*, vol. 203, no. 13, pp. 3207–3225, 2006.
- [6] I. Aharonovich, A. D. Greentree, and S. Prawer, “Diamond photonics,” *Nature Photonics*, vol. 5, no. 7, pp. 397–405, 2011.
- [7] L. M. K. Vandersypen and I. L. Chuang, “NMR techniques for quantum control and computation,” *Reviews of Modern Physics*, vol. 76, p. 1037, 2005.
- [8] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, “Training Schrödinger’s cat: quantum optimal control,” *European Physical Journal D*, vol. 69, p. 279, 2015.

- [9] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, “Simple pulses for elimination of leakage in weakly nonlinear qubits,” *Physical Review Letters*, vol. 103, p. 110501, 2009.
- [10] D. Das, “Nanocrystalline diamond,” in *Carbon-Based Nanofillers and Their Rubber Nanocomposites*, pp. 125–177, Elsevier, 2019.
- [11] G. Zanelli, E. Moreva, E. Bernardi, E. Losero, S. Ditalia Tchernij, J. Forneris, Ž. Pastuović, P. Traina, I. P. Degiovanni, and M. Genovese, “Interplay between dressed and strong-axial-field states in nitrogen vacancy centers for quantum sensing and computation,” *Physical Review B*, vol. 112, no. 23, p. 235201, 2025.
- [12] E. Moreva, E. Bernardi, P. Traina, A. Sosso, S. Ditalia Tchernij, J. Forneris, D. Gatto Monticone, F. Picollo, F. Bosia, P. Olivero, and M. Genovese, “Practical applications of quantum sensing: A simple method to enhance the sensitivity of nitrogen-vacancy-based temperature sensors,” *Physical Review Applied*, vol. 13, no. 5, p. 054057, 2020.
- [13] A. Dréau, M. Lesik, L. Rondin, P. Spinicelli, O. Arcizet, J.-F. Roch, and V. Jacques, “Avoiding power broadening in optically detected magnetic resonance of single nv defects for enhanced dc magnetic field sensitivity,” *Physical Review B*, vol. 84, no. 19, p. 195204, 2011.
- [14] Wikipedia contributors, “Confocal microscopy.” https://en.wikipedia.org/wiki/Confocal_microscopy, 2025. Accessed: October 29, 2025.
- [15] R. Brouri, A. Beveratos, J.-P. Poizat, and P. Grangier, “Photon antibunching in the fluorescence of individual color centers in diamond,” *Optics Letters*, vol. 25, no. 17, pp. 1294–1296, 2000.
- [16] C. Kurtsiefer, S. Mayer, P. Zarda, and H. Weinfurter, “Stable solid-state source of single photons,” *Physical Review Letters*, vol. 85, no. 2, pp. 290–293, 2000.
- [17] P. Grangier, G. Roger, and A. Aspect, “Experimental evidence for a photon anti-correlation effect on a beam splitter: A new light on single-photon interferences,” *Europhysics Letters*, vol. 1, no. 4, pp. 173–179, 1986.
- [18] Thorlabs, “Nf533-17 notch filter, 532 nm center wavelength.” <https://www.thorlabs.com>, 2025. Accessed: 2026-06-09.

- [19] Thorlabs, “Fel0650 longpass filter, 650 nm cut-on.” <https://www.thorlabs.com>, 2025. Accessed: 2026-06-09.
- [20] Zurich Instruments, *HDAWG User Manual*. Accessed: 2026-06-10.
- [21] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, “Simple pulses for elimination of leakage in weakly nonlinear qubits,” *Physical Review Letters*, vol. 103, no. 11, p. 110501, 2009.
- [22] J. M. Binder, A. Stark, N. Tomek, J. Scheuer, F. Frank, K. D. Jahnke, C. Müller, S. Schmitt, M. H. Metsch, T. Unden, T. Gehring, A. Huck, U. L. Andersen, L. J. Rogers, and F. Jelezko, “Qudi: A modular python suite for experiment control and data processing,” *SoftwareX*, vol. 6, pp. 85–90, 2017.
- [23] Qudi Developers, “Qudi: Modular python suite for experiment control and data processing.” <https://github.com/Ulm-IQO/qudi>, 2025. Accessed: July 2026.
- [24] D. M. Pozar, *Microwave Engineering*. Hoboken, NJ: John Wiley & Sons, 4 ed., 2012.
- [25] Polyphase Microwave, *AM0350A IQ Mixer Datasheet*, 2024. Accessed July 2026.
- [26] F. Dolde, V. Bergholm, Y. Wang, I. Jakobi, B. Naydenov, S. Pezzagna, J. Meijer, F. Jelezko, P. Neumann, T. Schulte-Herbrüggen, J. Biamonte, and J. Wrachtrup, “High-fidelity spin entanglement using optimal control,” *Nature Communications*, vol. 5, p. 3371, 2014.