

©Copyright 2026

Victor Sebastian Marcenac

High-Resolution Timing for Nitrogen-Vacancy Spin Memory Control

Victor Sebastian Marcenac

A thesis
submitted in partial fulfillment of the
requirements for the degree of

Master of Science in Electrical and Computer Engineering

University of Washington

2026

Committee:

Kai-Mei C. Fu

Maxwell F. Parsons

Program Authorized to Offer Degree:

Electrical & Computer Engineering

University of Washington

Abstract

High-Resolution Timing for Nitrogen-Vacancy Spin Memory Control

Victor Sebastian Marcenac

Co-Chairs of the Supervisory Committee:

Professor Kai-Mei C. Fu

Physics, and Electrical & Computer Engineering

Professor Maxwell F. Parsons

Electrical & Computer Engineering

Pulse-sequence timing influences the fidelity of electronic-spin operations and the ability to spectrally resolve and control nearby nuclear spins. This thesis presents a method for achieving effective 200 ps timing resolution in nitrogen-vacancy (NV) qubit experiments using the Quantum Instrument Control Kit–Diamond Arbitrary Waveform Generator (QICK-DAWG) implemented on the RFSoc 4×2, without modifying the underlying gateway. By using structured waveform banks and sequencer-level control, this approach enables efficient parameter sweeps while preserving the flexibility of QICK’s pulse-sequencing model.

We apply the methodology to room-temperature electronic-spin characterization and to dynamical-decoupling-based nuclear-spin spectroscopy of a single NV center in diamond. These measurements demonstrate that software-level sequencing strategies can expand the capabilities of existing quantum-control hardware, enabling sub-nanosecond timing control on an inexpensive and flexible open-source platform. This work establishes QICK-DAWG as a practical platform for high-temporal-resolution quantum defect memory control and provides a path toward more precise electron–nuclear spin manipulation.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	v
Chapter 1: Introduction	1
Chapter 2: Theoretical Background	5
2.1 The NV Center	5
2.2 Nuclear Spin Detection and Control	8
2.3 Timing Resolution in Existing Control Hardware	11
Chapter 3: Sub-Clock Timing Method	14
3.1 Method	14
3.2 Limitations	17
Chapter 4: Experimental Demo	21
4.1 Experimental Set-up	21
4.2 Magnetic Alignment Procedure	23
4.3 Electronic Spin Characterization	29
4.4 Nuclear Spin Characterization	33
Chapter 5: Conclusion & Outlook	40

5.1 Conclusion	40
5.2 Future Work	42
Bibliography	44
Appendix A: NV ground-state Hamiltonian for B-field alignment	48

LIST OF FIGURES

Figure Number	Page
2.1 Energy Diagram of the $^{14}\text{NV}^-$ center at room temperature	6
3.1 Waveform bank construction for π and $\frac{\pi}{2}$ pulses, with sub-cycle offset variants	15
3.2 Example delay decomposition between two offset waveforms.	16
4.1 Control schematic of the NV characterization setup, showing two NV centers under solid immersion lenses (SILs): D_o in blue and D_i in orange.	22
4.2 3D orientations and surface-plane projections of the four possible NV axes relative to the sample edges, with the two projection families color coded. . .	22
4.3 NV fluorescence vs Excitation Polarization	24
4.4 Close-up of the magnet positioning stage and sample mount	25
4.5 Pulsed-ODMR sequence (a), and scan (b) showing ^{14}N nuclear spin polarization at ~ 350 G with Gaussian fits.	27
4.6 Rabi sequence (a) and oscillations (b) obtained by driving the electronic spin near the $m_I = +1$ transition.	30
4.7 Ramsey measurement showing a drive detuning of 127.1(1.2) kHz.	32
4.8 CPMG sequence (a) and scan (b) $\Delta\tau - k$ plot with spin assignments (c) and fitted spins overlay (d).	34
4.9 Hahn-echo coherence measurement with exponential decay fit yielding the NV spin coherence time $T_2 = 393(42)$ μs	36
4.10 Population relaxation measurement with exponential recovery fit giving $T_1 = 1011(31)$ μs	37

4.11 CPMG spectroscopy on the Diamond Quantum Processor. (a) Wide-range scan, and 200 ps scans of (b) a strongly coupled ^{13}C spin with $N = 320$ and (c) narrow coherence dips.	39
---	----

LIST OF TABLES

Table Number		Page
2.1	Worst-case timing-induced rotation-angle error for a π pulse.	7

ACKNOWLEDGMENTS

While this journey has been relatively short, it has been profoundly rewarding because of the incredible people who have supported me along the way.

To my advisors, Professors Kei-Mei Fu and Max Parsons, I am deeply grateful for your invaluable guidance, and for your lessons on the importance of questions. Thank you, Professor Fu, for introducing me to my very first qubits, and for teaching me to zoom into the physics of my qubit. Thank you, Professor Parsons, for providing the broader perspective needed to see how my work connects to the wider landscape of quantum computing. Thank you to all my lab mates for the enlightening discussions, and especially to the DQP team: Tommy, Julie, Enrique and Weitao, for all your help.

Finally, thank you to my friends and family for the endless encouragement and for keeping me sane over the last two years. And to my Mom and Dad, thank you for your relentless unconditional support, especially when I was at my most tired.

Chapter 1

INTRODUCTION

Quantum states are abundant in nature, but only a small subset can be isolated, addressed, and controlled well enough to serve as useful qubits. Among the leading platforms, optically addressable spin defects in solids, or color centers, are especially promising for distributed quantum memory nodes. Their electronic spins provide an optical interface to photons, enabling communication between spatially separated systems, while nearby nuclear spins offer long-lived auxiliary qubits for storing quantum information locally.

Quantum memories are broadly useful because they allow quantum information to be stored, synchronized, and retrieved as part of a larger quantum system. In quantum networks, this capability is especially important: long-distance entanglement distribution through direct transmission is limited by photon loss, making repeater-based architectures essential for scaling beyond short distances. Protocols such as DLCZ show that the ability to store entanglement at intermediate nodes can transform an otherwise exponentially lossy entanglement-distribution process into one that scales polynomially with channel length [1]. A useful memory node therefore requires both an optical interface for generating remote entanglement and a local quantum memory that can preserve quantum information while repeated entanglement attempts are made.

Color centers in crystals naturally combine these memory and interface functions in a compact solid-state system. The electronic spin can be optically initialized and read out, entangled with emitted photons, and used to control nearby nuclear spins. The nuclear spins then serve as auxiliary memory qubits, with coherence times that can substantially exceed those of the electronic spin. Recent demonstrations of optically connected remote

color-center nodes further highlight the relevance of these platforms for distributed quantum networks [2, 3].

For certain defects, the crystalline host also provides sufficient isolation from the environment to support coherent single-defect control even at room temperature [4, 5, 6], enabling experimentation that avoids the external trapping and confinement infrastructure required by other qubit platforms such as trapped ions and neutral atoms. Among color centers, the nitrogen-vacancy (NV) center in diamond is the most experimentally mature. Its electronic spin supports optical initialization and readout, microwave control, and coupling to both the NV nitrogen nuclear spin and nearby host ^{13}C nuclear spins in the diamond lattice. Control through the electron–nuclear interactions has enabled selective manipulation of surrounding nuclear spins [7] and multi-qubit electron–nuclear registers under cryogenic conditions [8].

Realizing the full potential of NV-based memory nodes requires not only long-lived nuclear spins, but precise and selective control over them. Nearby ^{13}C spins must be spectrally resolved, individually addressed, and coherently manipulated through pulse sequences applied primarily to the electronic spin. The timing of these microwave pulse sequences determines the spectral selectivity of the interaction and the effective control exerted on each nuclear spin. As a result, limitations in pulse timing resolution directly constrain the ability to resolve, selectively target, and control individual memory qubits.

These requirements make the underlying control hardware central to the experiment. Field-programmable gate arrays (FPGAs) provide a flexible and deterministic platform for implementing such pulse sequences, typically operating at clock rates on the order of hundreds of megahertz. Radio Frequency System-on-Chip (RFSoc) platforms extend this capability by integrating FPGA fabric with high-speed digital-to-analog converters (DACs), enabling direct synthesis of microwave signals with sampling rates approaching 10 GHz. Frameworks such as the Quantum Instrumentation Control Kit (QICK) and its Diamond Arbitrary Waveform Generator extension (QICK-DAWG) provide specialized gateway and software for programming these systems for quantum experiments.

A key feature of the RFSoc-based control architecture is the separation between timing domains. Waveforms are generated at the DAC sampling rate, providing sub-nanosecond resolution, while control flow and sequencing occur in the FPGA fabric at a much lower clock rate. This results in a mismatch: individual pulse shapes can be specified with high temporal precision, but relative pulse placement remains quantized to the ~ 3.2 ns sequencer period. This limits the control that can be exercised over interpulse delays. A waveform-only approach can recover this resolution by encoding entire sequences at the DAC rate, but this is inefficient due to limited waveform memory and the need to reload waveforms during parameter sweeps.

This thesis addresses the timing-resolution limitation imposed by coarse-grained sequencing in RFSoc-based quantum control systems. Although sub-nanosecond timing is available at the waveform level, it is not readily compatible with the sequencing logic that enables flexible sequence construction and efficient parameter sweeps.

We introduce a hybrid timing strategy that combines sequencer-level scheduling with waveform-level temporal offsets to achieve sub-clock-cycle timing resolution. In this approach, coarse delays are implemented by the sequencer, while fine timing adjustments are realized by selecting from a bank of pre-offset waveforms. This enables precise control of both pulse durations and interpulse delays without relying on waveform-only implementations.

Additionally, this method is implemented entirely at the instruction level within the QUICK framework, requiring no FPGA gateway modifications. We further develop an efficient implementation based on register-level arithmetic and waveform indexing, and analyze the resulting tradeoffs in instruction overhead, waveform memory usage, and sequence scalability.

In the second part of this thesis, we apply this method to an accessible room-temperature testbed based on a single NV center. We characterize the electronic spin and its coupled nuclear-spin environment using pulse sequences with high-resolution timing control under the practical constraints imposed by finite coherence and limited signal-to-noise ratio. We

then use cryogenic measurements to directly demonstrate higher-resolution nuclear-spin spectroscopy enabled by sub-clock-cycle timing control.

Chapter 2

THEORETICAL BACKGROUND

Quantum control of optically addressable spin defects in solids is achieved by applying optical, microwave, and radio-frequency fields that are resonant to transitions in the electronic and nuclear spin states. This chapter introduces the relevant physics for control of the negatively charged nitrogen-vacancy center in diamond and its surrounding nuclear-spin environment, and connects it to the pulse-sequencing hardware used to implement experimental control.

2.1 *The NV Center*

The nitrogen-vacancy (NV) center in diamond is a point defect formed by a substitutional nitrogen atom adjacent to a lattice vacancy. Diamond's low natural abundance of spinful nuclei (primarily ^{13}C at 1.1%), limits the magnetic noise from the surrounding lattice [4]. As a result, NV centers exhibit long coherence times even at room temperature, enabling their use as qubits and as a controllable interface to nearby, longer-lived nuclear spins.

The primary qubit of interest is the electronic spin of the negatively charged NV center. The electronic structure, shown in Fig. 2.1, consists of ground (3A_2) and excited (3E) triplet manifolds, with a zero-field splittings between the $m_s = 0$ and $m_s = \pm 1$ sublevels [4].

Optical initialization into the $m_s = 0$ state is achieved using a 532 nm laser. Following excitation, a (primarily) non-radiative decay pathway through a metastable singlet state preferentially returns population to the $m_s = 0$ ground state. This spin-dependent relaxation pathway also enables optical readout via spin-dependent fluorescence intensity [4].




Figure 2.1: Energy Diagram of the $^{14}\text{NV}^-$ center at room temperature

The discussion in this chapter is framed around room-temperature operation, which is the regime used for most experiments in this thesis. However cryogenic conditions can provide numerous advantages, including longer electronic-spin relaxation times and narrow optical transitions that enable high-fidelity spin-selective readout [9, 10, 11].

2.1.1 Electronic Spin Control

In this work, the electronic qubit is defined by the $m_s = 0$ ($|0\rangle$) and $m_s = -1$ ($|1\rangle$) transition, often preferentially used as it occurs at a lower frequency than the $m_s = 0 \leftrightarrow m_s = +1$ transition and is therefore generally more accessible for microwave control. In the presence of an external magnetic field B_0 aligned with the NV axis, the transition frequency is give approximately by

$$\omega_{0,-1} \approx D - \gamma_e B_0, \quad (2.1)$$

where $D \approx 2.87$ GHz is the zero-field splitting at room-temperature and $\gamma_e \approx 2.802$ MHz/G is the electron gyromagnetic ratio [4, 12].

Control of the NV electronic spin is achieved through resonant microwave (MW) driving. In the rotating frame, the drive produces rotations about axes in the XY plane with angle

$$\theta = \Omega_R t, \quad (2.2)$$

where Ω_R is the Rabi frequency and t is the applied square-pulse duration.

High-fidelity control requires precise implementation of target rotations, including π and $\pi/2$ pulses. In the simplest picture, finite timing resolution δt constrains the set of accessible pulse durations, introducing a corresponding discretization in the achievable rotation angle,

$$\delta\theta = \Omega_R \delta t. \quad (2.3)$$

As a result, rotation errors scale linearly with both the Rabi frequency and the timing resolution, as illustrated in Table 2.1. This estimate shows that finer timing resolution can improve the granularity of electronic-spin rotations, especially at higher Rabi frequencies. In practice, however, electronic rotations can also be calibrated by modifying the Rabi frequency, which is finely tunable through the applied microwave amplitude. In addition, the version of QICK used in this work already supports pulse envelopes with finely tuned durations at the waveform level.

Rabi frequency $\Omega_R/2\pi$	Max angle error (3.2 ns timing)	Max angle error (200 ps timing)
5 MHz	2.88°	0.18°
10 MHz	5.76°	0.36°
20 MHz	11.5°	0.72°

Table 2.1: Worst-case timing-induced rotation-angle error for a π pulse.

Thus, while the electronic-spin rotation angle provide a useful example of how timing granularity enters spin control, it is not the limitation addressed in this thesis. The limitation instead arises from timing resolution in the multi-pulse sequences used for nuclear-spin detection and control, where the interpulse delay itself determines the effective electron–nuclear interaction.

2.2 Nuclear Spin Detection and Control

The interaction between the NV electronic spin and a nearby nuclear spin provides both a mechanism for detection and a pathway toward coherent control of nuclear spins. This can be understood by considering the hyperfine interaction with a single spin- $\frac{1}{2}$ nucleus.

In the secular approximation, the interaction Hamiltonian is given by

$$H = \omega_L I_z + S_z (A_{\parallel} I_z + A_{\perp} I_x), \quad (2.4)$$

where S_z and I_z are the spin-projection operators for the electron and nuclear spin respectively, ω_L is the nuclear Larmor frequency and $A_{\parallel}$, $A_{\perp}$ are the longitudinal and transverse hyperfine coupling components, respectively [12].

2.2.1 Conditional Nuclear Evolution

Because the interaction depends on S_z , the nuclear spin evolves under two different Hamiltonians conditioned on the electronic spin state, with H_0 (H_1) corresponding to $|0\rangle$ ($|1\rangle$):

$$H_0 = \omega_L I_z, \quad (2.5)$$

$$H_{-1} = \omega_L I_z - A_{\parallel} I_z - A_{\perp} I_x. \quad (2.6)$$

Starting from an initial product state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |\phi(0)\rangle$, where $|\phi(0)\rangle$ denotes the

initial nuclear spin state, the system evolves into an entangled state,

$$|\psi(t)\rangle = \frac{1}{\sqrt{2}} (|0\rangle \otimes |\phi_0(t)\rangle + |1\rangle \otimes |\phi_1(t)\rangle), \quad (2.7)$$

where $|\phi_{0,1}(t)\rangle = e^{-iH_{0,1}t} |\phi(0)\rangle$.

The electronic coherence is then determined by the overlap between these two nuclear states,

$$L(t) = \langle \phi_0(t) | \phi_1(t) \rangle. \quad (2.8)$$

This forms the basis for nuclear spin detection, as the conditional nuclear evolution modulates the electronic coherence, which can be assessed via electronic spin readout.

2.2.2 Role of Longitudinal and Transverse Coupling

The longitudinal interaction $A_{\parallel} I_z$ produces a conditional shift in the nuclear precession frequency. For a nuclear spin prepared in a coherent superposition, this leads to relative phase accumulation between the two conditional evolutions.

However, in typical experiments the nuclear spin is not initialized and is well described by a mixed state diagonal in the I_z basis. In this case, longitudinal evolution only generates global phases within each basis state and does not produce a measurable reduction in overlap.

In contrast, the transverse term $A_{\perp} I_x$ drives rotations of the nuclear spin about an axis perpendicular to the quantization axis. This results in a change in the nuclear state, generating distinguishability between $|\phi_0(t)\rangle$ and $|\phi_1(t)\rangle$ even for an initially incoherent state. Consequently, transverse coupling provides the dominant contribution to experimentally observed decoherence features.

2.2.3 Detection via Dynamical Decoupling

Under free evolution, the electronic spin is sensitive to the collective influence of its local spin environment. Because there is no mechanism for selectively enhancing a chosen interaction,

the small perturbations from weakly coupled nuclear spins are easily obscured by stronger interactions and by the overall decay of electronic spin coherence [13]. The resulting lack of spectral selectivity prevents individual weakly coupled nuclei from being reliably isolated or addressed as independent control targets.

Periodic dynamical decoupling sequences such as Carr-Purcell-Meiboom-Gill (CPMG) address this limitation by modulating the electronic spin state with a train of π pulses. This modulation introduces spectral selectivity, concentrating sensitivity near selected nuclear precession frequencies and thereby enabling individual nuclear spins to be resolved and controlled.

The CPMG sequence applies evenly spaced refocusing pulses separated by an interpulse delay 2τ . In the frequency domain, this periodic modulation acts as a band-pass filter whose spectral response is determined by the interpulse spacing and the number of applied pulses [14]. For a weakly coupled nuclear spin, resonances occur when the interpulse delay is synchronized with the effective nuclear precession frequency. This condition can be written approximately as [13]

$$\tau_k \approx \frac{(2k-1)\pi}{2\omega_L + A_{\parallel}}, \quad (2.9)$$

where $k = 1, 2, 3, \dots$ indexes successive resonance orders. Thus, the same nuclear spin can satisfy the CPMG resonance condition at multiple values of τ_k , producing a recurring family of coherence dips as the interpulse delay is swept. With the dip positions being primarily determined by ω_L and $A_{\parallel}$, while the dip widths and contrasts are strongly influenced by the transverse coupling $A_{\perp}$, the number of pulses N , and the background electronic coherence decay [15, 13].

2.2.4 Implications for Control

The same spectral selectivity that enables nuclear-spin detection also provides a route to coherent control. Under resonant dynamical decoupling, the selected nuclear spin evolves conditionally on the electronic spin state. Repetition over many pulse cycles amplifies this conditional evolution, so that even weak hyperfine couplings can produce appreciable nuclear-spin rotations. By choosing the resonance condition and the number of applied pulses, the rotation axis and accumulated rotation angle can therefore be tuned, allowing the nuclear spin to be manipulated rather than only detected [13].

This makes the interpulse delay τ a direct control parameter. Since τ determines how closely the sequence matches a selected resonance order, finite timing resolution produces a detuning from the desired resonance condition. This changes the effective conditional evolution of the nuclear spin, altering both the rotation axis and the accumulated rotation angle per cycle. As a result, coarse timing steps limits, not just how accurately a particular nuclear spin can be targeted, but also how accurately a desired nuclear rotation can be implemented.

2.3 Timing Resolution in Existing Control Hardware

Arbitrary waveform generators (AWGs) are a conventional high-performance approach for microwave pulse generation in NV center experiments. In an AWG-based architecture, pulse shapes are precomputed, loaded into waveform memory, and played back using a fixed sampling clock, either directly as the desired waveform or as an envelope used for subsequent microwave modulation. AWG-based control has been used in state-of-the-art NV quantum processors. For example, Abobeih *et al.* [12] used a Tektronix AWG5014C to control pulse sequences in a fault-tolerant logical qubit experiment in diamond, with real-time feedforward provided by an external ADwin microcontroller. In that experiment, precise timing for quantum gates was based on the AWG clock with 1 ns precision.

The AWG5014C provides a representative NV-control comparison point, with a maximum sampling rate of 1.2 GS/s, corresponding to a waveform sample spacing of approximately 833 ps. Modern state-of-the-art commercial AWGs can operate at substantially faster rates; for example, direct-RF instruments such as the Tektronix AWG70002A and the Active Technologies ARB Rider AWG-7000 reach 25 GS/s and 20 GS/s, respectively, corresponding to sample spacings of 40 ps and 50 ps. These systems provide a well-established route to precise waveform generation, but they are often costly, with high-performance units commonly falling in the tens-of-thousands-of-dollars range or higher. In contrast, the RFSoc span a lower and more accessible price range, with the RFSoc 4×2 development board used in this work, being available through subsidized academic pricing at roughly \$2.5k, while the larger RFSoc boards are priced in the \$15–\$20k range. Beyond cost, AWG-based experimental architectures can also be complex, relying on additional instruments for synchronization, photon counting, and managing real-time feedback.

QICK adopts a different architecture based on the radio-frequency system-on-chip (RFSoC), combining high-speed RF digital-to-analog converters (DACs) with FPGA-based real-time control [16]. A programmable FPGA sequencer is used to schedule pulses and coordinate experiment timing. A numerically controlled oscillator sets the carrier frequency, which, depending on the operating mode, is digitally mixed with I/Q waveforms before conversion to analog signals. Unlike an AWG-style approach, where long or highly varied sequences are often represented by loading many waveform samples into memory, the FPGA-based model represents much of the experiment through compact sequencer instructions, parameter updates, and reusable waveform envelopes. This integrated control model is especially attractive for NV experiments, where microwave control must be synchronized with laser initialization, optical readout, photon counting, and external triggers.

A critical architectural distinction is that QICK exposes two timing scales: the high-speed DAC sample grid, which determines the spacing of waveform samples, and the slower FPGA logic clock, which determines the native granularity of sequenced pulse start times [17]. While

the RFSoc 4×2 DAC used in this thesis is capable of 9.85 GSPS, in our experiments, it is operated at 4.915 GSPS, giving a DAC sample spacing of approximately 200 ps. The FPGA logic operates at one sixteenth of this rate, 307.2 MHz, corresponding to a native sequencer period of about 3.2 ns.

The central limitation addressed in this work is therefore not the RFSoc DAC sampling rate itself, but rather the fact that ordinary sequencer-based pulse placement occurs on the coarser FPGA clock grid. This creates a mismatch between the timing resolution available in the waveform samples and the timing resolution directly exposed to pulse programs. The following chapter describes a method for recovering the finer DAC-scale timing while preserving the QUICK-DAWG FPGA control model.

Chapter 3

SUB-CLOCK TIMING METHOD

3.1 Method

3.1.1 Timing Model

In the QUICK gateway used in this work, pulse programs are centrally sequenced by the tProcessor, and executed by the signal-generator module. In this thesis, the term sequencer is used to refer to this combined FPGA timing path, with both modules operating on a clock period $T_s \approx 3.2$ ns. Pulse sequences are defined by the relative start times and durations of pulses. Instructions such as **sync** advance a global time cursor in integer multiples of T_s , and waveform playback is scheduled relative to this cursor.

The generated waveforms are sampled at the DAC period $T_d \approx 200$ ps, where

$$T_s = 16T_d. \quad (3.1)$$

Pulse envelopes can be specified on the finer DAC sample grid, but their playback start times are natively quantized to the coarser sequencer grid. That is, native scheduling can only place pulses in 3.2 ns increments.

A delay Δ specified at the DAC-sample resolution can be decomposed into a coarse sequencer delay and a fine residual offset,

$$\Delta = nT_s + \delta, \quad (3.2)$$

where $n \in \mathbb{Z}_{\geq 0}$ is the number of sequencer clock cycles and $0 \leq \delta < T_s$. Since $T_s = 16T_d$, the residual offset can be written as

$$\delta = kT_d, \quad k \in \{0, 1, \dots, 15\}. \quad (3.3)$$




Figure 3.1: Waveform bank construction for π and $\frac{\pi}{2}$ pulses, with sub-cycle offset variants

3.1.2 Waveform Construction

For each pulse type, a bank of sixteen waveforms is constructed to represent the possible sub-cycle offsets within one sequencer clock period. The waveform indexed by k contains k leading zero-valued samples followed by the pulse envelope. Additional zero padding is appended after the pulse so that all waveforms in the bank have the same length and the total waveform duration is an integer multiple of T_s . The waveform can therefore be scheduled on the native sequencer grid while the active pulse envelope is shifted by kT_d relative to the scheduled waveform start.

To make waveform selection deterministic, the offset waveforms are initialized sequentially in a contiguous region of block RAM (BRAM). In practice, waveforms are grouped by pulse type (offset index) and ordered by offset index (pulse type), as shown in Fig. 3.1, allowing the correct waveform address to be computed from the pulse type and residual offset.




Figure 3.2: Example delay decomposition between two offset waveforms.

3.1.3 Delay Decomposition

The desired delay Δ is defined between the end of the active portion of one pulse and the start of the active portion of the next pulse, expressed in DAC samples. Since each scheduled waveform contains both the active pulse envelope and zero padding, part of this delay is already accumulated before the next sequencer instruction is reached.

Let L_w denote the total waveform length, L_p the active pulse length, and k_{prev} the offset index of the previous pulse. The dead time already included at the end of the previous waveform is

$$D_{\text{dead}} = L_w - (L_p + k_{\text{prev}}). \quad (3.4)$$

The remaining delay is therefore

$$\Delta' = \Delta - D_{\text{dead}}. \quad (3.5)$$

This delay is decomposed as

$$n = \left\lfloor \frac{\Delta'}{16} \right\rfloor, \quad k = \Delta' \bmod 16, \quad (3.6)$$

where n is the sequencer delay in units of T_s , and k selects the offset waveform for the next pulse.

3.1.4 Instruction-Level Implementation

The delay decomposition is implemented using integer arithmetic on FPGA registers. Since $T_s/T_d = 16$, division and modulo can be replaced by bitwise operations:

$$\begin{aligned} n &= \Delta' \gg 4, \\ k &= \Delta' \& 0b1111. \end{aligned} \tag{3.7}$$

The right shift extracts the integer number of sequencer cycles, while the `0b1111` bit mask extracts the residual offset index.

With the waveform layout described above, waveform selection reduces to an address calculation. If S_w is the memory stride between successive waveforms and A_{type} is the starting address of the waveform block for a given pulse type, then

$$addr = kS_w + A_{\text{type}}. \tag{3.8}$$

Here, S_w is determined by the padded waveform duration expressed in sequencer-period units. Grouping waveforms by pulse type and ordering them by offset index allows the FPGA program to select the correct waveform directly from the computed residual k .

Together, the waveform-offset construction and delay decomposition allow pulse timing to be specified at the DAC sample period T_d , while preserving sequencer-based programmability. The approach avoids waveform reloading during parameter sweeps and supports long delays through ordinary sequencer instructions rather than long waveform playback.

3.2 Limitations

The waveform-offset method recovers DAC-scale timing resolution by using additional waveform structure and instruction-level computation. This is a tradeoff rather than a free increase in capability: the practical value of the method depends on how memory usage, instruction count, and sequencing overhead scale with pulse-sequence complexity.

3.2.1 *Waveform Memory*

Waveforms are stored in FPGA block RAM, with finite capacity corresponding to approximately $13.3\ \mu\text{s}$ of samples. Since each pulse type is represented by a bank of offset waveforms, memory usage scales with the number of distinct pulse types, their waveform lengths, and the number of supported offsets. Reducing the number of offsets by a power of two is straightforward, but comes at the cost of a proportional reduction in timing resolution. Each stored waveform also includes zero padding beyond the active pulse duration L_p , slightly increasing the effective storage required per pulse. This constraint becomes most significant for sequences requiring many unique pulse shapes or long-duration waveforms.

3.2.2 *Instruction Memory*

The sequencer has finite instruction memory, limited to 1024 instructions in the present implementation. The waveform-offset method increases the instruction count per pulse because each delay must be decomposed and the appropriate waveform selected before playback.

This cost scales with sequence length. For short and moderate sequences, the additional instructions are manageable, but long pulse trains can approach the instruction-memory limit. Such sequences must therefore be structured carefully, primarily through loops and other program-level optimizations. For example, the CPMGXY8 phase pattern (X, Y, X, Y, Y, X, Y, X) can be compressed into a single loop by encoding the phase pattern in a register and extracting the required phase with bit masking.

3.2.3 *Instruction Overhead*

Each pulse requires additional arithmetic and indexing operations to compute the coarse delay and select the appropriate offset waveform. Although these operations require only a small number of sequencer cycles, they introduce overhead relative to a minimal QICK pulse sequence.

This overhead is usually manageable because QICK separates instruction execution from pulse playback. The tProcessor dispatches pulse commands into timed queues associated with the signal-generator modules, allowing it to continue executing subsequent instructions before a queued pulse is physically played. In this sense, the timed queues provide a buffer between the sequential instruction stream and the scheduled waveform output.

However, the sequencer must still remain sufficiently ahead of the requested pulse times. For sequences with high pulse density or simultaneous use of multiple channels, the added arithmetic and waveform-selection instructions can reduce this timing margin. If the tProcessor does not dispatch a pulse command before its scheduled time is reached, the sequence stalls and the pulse is delayed, apparently silently. Since QICK execution is deterministic, this behavior can be checked experimentally by validating dense pulse programs on an oscilloscope.

3.2.4 Pulse Spacing and Concatenation

Closely spaced pulses cannot always be represented as independent waveform selections. In the waveform-offset construction, each pulse is embedded in a padded waveform containing leading zeros to set the sub-cycle offset and trailing zeros to preserve a fixed waveform stride.

As a result, two pulses can only be treated as independent waveform selections when there is sufficient delay between their active envelopes to accommodate the required padding and any pending sequencer operations that must occur before the next waveform playback. If a shorter separation is required, the pulses must instead be combined into a single composite waveform. This preserves the intended timing, but reduces the modularity of the pulse representation and typically increases waveform memory usage.

3.2.5 Phase Control

Phase control is divided between the stored I/Q envelope and the numerically controlled oscillator (NCO). The I/Q waveform defines any internal phase structure of the pulse, such as the relative phases of multiple segments in a composite waveform. The NCO then applies a uniform phase rotation to the entire waveform at playback.

This separation is useful because a single envelope can be reused for pulses that differ only by a global phase. For example, a $\pi/2$ pulse at phase ϕ can be implemented by updating the NCO phase register before playback, without storing a distinct waveform for each value of ϕ .

The limitation is that NCO phase updates occur at the sequencer time scale and act uniformly on the selected waveform. The NCO therefore cannot modify the internal phase structure of a waveform. Concatenated pulses with different relative phase patterns require distinct precomputed envelopes. For example, (X, Y) and $(Y, -X)$ share the same internal structure up to a global phase and may reuse an envelope, whereas (Y, X) requires a different envelope because the relative phase ordering is changed. Thus, the NCO provides efficient global phase control, while internal phase structure must be encoded directly into waveforms.

3.2.6 Dependence on Clock Ratio

The efficiency of the delay decomposition relies on the clock-domain ratio being a power of two; in this implementation, $T_s/T_d = 16$. This permits bitwise operations to replace division and modulo. If this ratio were not a power of two, more expensive arithmetic operations would be required, increasing instruction overhead.

Chapter 4

EXPERIMENTAL DEMO

4.1 *Experimental Set-up*

4.1.1 *Testbed*

To validate the timing-control methods experimentally, we integrate the RFSoc 4×2 running QICK-DAWG with a confocal microscope platform, enabling synchronized optical and microwave control of individual NV centers (Fig. 4.1). QICK-DAWG performs microwave pulse generation, transistor-transistor logic (TTL) control of the excitation laser through the acousto-optic modulator (AOM), and acquisition of photoluminescence (PL) photon counts from an avalanche photodiode (APD).

Microwave control is delivered through a dedicated RF chain in which the RFSoc-generated pulses are amplified and sent to a wirebond antenna positioned $\sim 25 \mu\text{m}$ above the diamond.

Fluorescence is collected through an APD and used for spin-state readout. The sample consists of NV centers fabricated under solid immersion lenses (SILs), which enhance photon collection efficiency by mitigating total internal reflection at the diamond interface. This is particularly important given the high refractive index of diamond and the angular emission profile of NV centers.

The diamond used in this work has a natural abundance of ^{13}C (1.1%).




Figure 4.1: Control schematic of the NV characterization setup, showing two NV centers under solid immersion lenses (SILs): D_o in blue and D_i in orange.




Figure 4.2: 3D orientations and surface-plane projections of the four possible NV axes relative to the sample edges, with the two projection families color coded.

4.1.2 Sample Geometry and NV Orientation

The diamond is a $[100]$ -oriented plate whose in-plane edges are aligned along the $\langle 110 \rangle$ directions. As illustrated in Fig. 4.2, the projections of the four $\langle 111 \rangle$ NV axes onto (100) plane form a cross aligned with the sample edges, rather than lying along the diagonals.

These projections define two distinguishable orientation families, each consisting of two NV axes with antiparallel in-plane projections. The two axes in each family remain distinct in three dimensions: each NV axis is tilted by 54.74° from the surface normal, and the two axes are separated by the tetrahedral angle 109.47° [4].

This geometry determines both the polarization-dependent optical response and the magnetic-field directions used for alignment.

4.2 Magnetic Alignment Procedure

High-fidelity spin control of the electronic and nuclear spins requires a strong external magnetic field, aligned with the specific NV axis citeAbobeih. For the electronic spin, the axial magnetic field component is lifts the $m_S = \pm 1$ degeneracy, while strong transverse components increase spin mixing. For ^{13}C spins, the external field sets the nuclear Zeeman splitting and defines the bare nuclear quantization axis. Additionally a sufficiently strong magnetic field sets a large bare Larmor frequency, reducing overlap between CPMG resonance orders. Our alignment procedure consists of three stages: orientation identification, coarse alignment, and fine alignment.

4.2.1 Orientation Identification

We first determine the NV orientation family using polarization-dependent photoluminescence (PL). Optical absorption depends on the projection of the excitation polarization onto the optical dipoles transverse to the NV axis [18]. By rotating the excitation polarization




Figure 4.3: NV fluorescence vs Excitation Polarization

and recording the fluorescence intensity, we obtain a sinusoidal photoluminescence PL modulation that identifies the orientation families of the two single NV centers under the SIL, denoted D_o and D_i (Fig. 4.3). The remaining measurements in this thesis focus on D_o .

4.2.2 Coarse Magnetic Alignment

After identifying the NV family, we restrict the external magnetic field to lie within the plane defined by the two candidate NV axes in that family. We then apply a weak (~ 100 G) magnetic field at an oblique angle (approximately $\pm 45^\circ$ relative to the surface normal) corresponding to the two in-plane projections and perform optically detected magnetic resonance (ODMR) spectroscopy on both $m_s = 0 \leftrightarrow \pm 1$ transitions.

The Zeeman splitting depends on the projection of the magnetic field onto the NV axis, $|\mathbf{B}_0 \cdot \hat{n}_{\text{NV}}|$. By comparing the splitting between the two magnet orientations constrained to




Figure 4.4: Close-up of the magnet positioning stage and sample mount

this plane, we determine which of the two candidate NV axes is more closely aligned with the applied field direction. With the NV axis being roughly aligned to the magnet orientation that produces the larger Zeeman splitting.

4.2.3 Fine Alignment

Fine alignment is achieved by iteratively adjusting the magnet position, using the micrometers captured in Fig. 4.4, while monitoring ODMR spectra. The X micrometer is primarily used to set the strength of the magnetic field, while fine angle adjustments are performed by tuning the Z and Y micrometers. By fitting the measured transition frequencies to a spin Hamiltonian model, we estimate both the magnitude and angular alignment of the magnetic field.

At the 300–450 G fields used in this work, the fluorescence becomes increasingly sensitive to small transverse magnetic-field components, as off-axis magnetic fields mix the NV spin

states and reducing both PL intensity and ODMR contrast [19]. Once the ODMR-based alignment is close, we switch to using the field dependence of the PL signal as our fine tuning metric, adjusting the magnet position to maximize fluorescence.

4.2.4 ^{14}N Nuclear Spin Polarization

In the absence of nuclear polarization, each electronic transition exhibits three hyperfine-split resonances corresponding to the $m_I = \{-1, 0, +1\}$ nuclear spin projections, separated by ~ 2.16 MHz. Under robust polarization, the population concentrates into a single hyperfine state, improving both the electronic spin dephasing times (T_2^*) and signal-to-noise ratio (SNR).

When operating under an axially aligned external magnetic field of 300 – 450 G, optical pumping into $m_S = 0$ via the singlet-state pathway simultaneously polarizes the nuclear spin into $m_I = +1$. This polarization arises from residual excited-state mixing between states that share the same total spin angular momentum projection ($m_F = m_S + m_I$), over thousands of excitation-relaxation cycles [20]. While polarization efficiency improves as the field approaches the excited-state level anticrossing (ESLAC, ~ 510 G), the accompanying increase in nuclear spin-flip rates reduces state stability, degrading SNR.

Nuclear spin polarization is verified using pulsed optically detected magnetic resonance (PODMR), performed at low microwave power to suppress power broadening. In PODMR, a fixed-duration microwave pulse is applied before optical readout, and the resulting fluorescence contrast is measured as a function of microwave frequency. At low microwave power the ODMR lineshape is described by inhomogeneous broadening (over repetitions), primarily arising from the magnetic dipolar interactions with the ^{13}C nuclear-spin bath [21]. We therefore fit the PODMR spectrum using a Gaussian profile as opposed to the Lorentzian lineshape characteristic of homogeneous, power-broadened transitions [22].

For the PODMR measurement at ~ 369 G after magnetic alignment, the microwave




Figure 4.5: Pulsed-ODMR sequence (a), and scan (b) showing ^{14}N nuclear spin polarization at ~ 350 G with Gaussian fits.

amplitude was first calibrated on resonance, yielding a Rabi frequency of ~ 190 kHz, corresponding to a $2.6\text{ }\mu\text{s}$ π -pulse. The PODMR scan shown in Fig. 4.5 was then acquired using a $2.3\text{ }\mu\text{s}$ rectangular pulse, producing a controlled under-rotation. The spectrum is fit with a constrained three-Gaussian model with independent amplitudes, a common linewidth, and a shared nearest-neighbor hyperfine splitting. The fitted splitting is $2.15(7)$ MHz, consistent within uncertainty with the axial ^{14}N hyperfine constant [4, 12]. The $m_I = +1$ transition sets the full-width-half-maximum (FWHM) to $0.35(1)$ MHz. The $m_I = 0$ feature is weak relative to the noise floor, while the $m_I = -1$ feature is not resolved.

The weak side features near the dominant $m_I = +1$ dip are consistent with finite-pulse spectral structure from the rectangular microwave pulse. For a $T = 2.3\text{ }\mu\text{s}$ square pulse, the expected Fourier-limited FWHM is approximately

$$\text{FWHM} \approx 0.89/T = 0.39\text{ MHz}, \quad (4.1)$$

comparable to the fitted linewidth of $0.35(1)$ MHz. This agreement suggests that the observed linewidth is likely limited primarily by the finite pulse duration, rather than by inhomogeneous broadening alone.

The relative nuclear-spin populations are extracted from the integrated areas of the three Gaussian peaks and normalized by their total area. We obtain $P_{-1} = 0.0(1.2)\%$, $P_0 = 3.2(1.2)\%$, and $P_{+1} = 96.8(3.8)\%$. Using the experimental spin-1 polarization metric employed by Busaite et al. [23], which in this case simplifies to

$$\mathcal{P} = P_{+1} - P_{-1}, \quad (4.2)$$

we determine a ^{14}N polarization of $\mathcal{P} = 96.8(3.5)\%$. This result is consistent with the field dependence reported by Busaite et al., who measured $\mathcal{P} = 96(2)\%$ near the ESLAC and observed polarization exceeding approximately 90% down to 300, G, below which it declined steeply. Our measurement at 369, G therefore lies within the range over which strong nuclear-spin polarization is expected.

4.3 Electronic Spin Characterization

With the system aligned, we proceed to characterize the NV electronic spin using a sequence of control experiments implemented entirely within the QICK-DAWG framework; performing Rabi oscillations, Ramsey interferometry, and the CPMG dynamical decoupling sequence.

4.3.1 Rabi Oscillations

Driving near the $m_I = +1$ transition identified in the preceding PODMR measurement, at $f = 1838.15$ MHz, we perform a Rabi sweep at high MW power. This sequence is similar to PODMR, but with the MW frequency fixed while the MW pulse duration is varied. The resulting oscillations, shown in Fig. 4.6, yield an electronic-spin Rabi frequency of $\Omega_R = 10.482(1)$ MHz, which is used to set the π - and $\pi/2$ -pulse durations for subsequent experiments. From the same measurement, we obtain a relative photoluminescence contrast of $28.1(1)\%$ between the electronic spin states.

For the remaining experiments, we operate at QICK's maximum MW gain and use a $\pi/2$ -pulse duration of ~ 23.8 ns, approximately 50 ps shorter than the ideal duration inferred from the measured Rabi frequency. This corresponds to an estimated rotation error of 0.18° per $\pi/2$ pulse. The corresponding π pulse is implemented with twice this duration, giving an estimated rotation error of 0.36° per pulse. These errors compound in multi-pulse sequences, reducing the effective coherence through imperfect refocusing and introducing additional modulations. This is particularly relevant for CPMG sequences with large N , and becomes increasingly important when moving from isolated calibration sequences to the concatenated gate operations required for multi-qubit control. Although fine-tuning the MW amplitude can reduce these errors, we do not perform this optimization for the purposes of this demonstration.




Figure 4.6: Rabi sequence (a) and oscillations (b) obtained by driving the electronic spin near the $m_I = +1$ transition.

4.3.2 Ramsey Interferometry

We then perform Ramsey interferometry by sweeping a free-evolution delay τ between two $\pi/2$ pulses. The first $\pi/2$ pulse prepares the electronic spin in an equal superposition, $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, while the second $\pi/2$ pulse maps the accumulated phase back onto the measurement axis for optical readout. In this measurement, the two $\pi/2$ pulses are applied with opposite phases, such that in the absence of detuning and dephasing the spin is returned to $|0\rangle$.

Ramsey interferometry is sensitive to both electronic-spin dephasing and detuning between the MW drive and the electronic-spin transition. Dephasing appears as a decay envelope, while detuning produces oscillations at the detuning frequency. In principle, the ^{14}N hyperfine interaction can produce three distinct Ramsey frequencies; however, with strong nuclear-spin polarization, we expect a single dominant oscillation.

The measured Ramsey signal is fit with a decaying cosine, as shown in Fig. 4.7. A Gaussian decay envelope is used to describe inhomogeneous dephasing from repeated measurements in a quasi-static spin environment [24]. The resulting fit is used to correct the MW drive frequency. Because the experiment is performed at relatively high Rabi frequency, power broadening makes the Rabi frequency only weakly sensitive to small shifts in drive frequency. Nevertheless, after changing the MW frequency, it can still be useful to repeat a Rabi scan to recalibrate the pulse durations.

The fit yields a detuning of 127.1(1.2) kHz and an inhomogeneous dephasing time of $T_2^* = 5.300(67)$ μs . The fit also includes a small phase offset, $\phi = -0.072(21)$ rad, corresponding to approximately -4.1° . A simple finite-pulse estimate, based on the measured detuning and a 24 ns $\pi/2$ -pulse duration, gives a phase accumulation of ~ 0.019 rad, or $\sim 1.1^\circ$, with the opposite sign. We therefore treat ϕ as a nuisance parameter and do not assign physical significance to this offset.

Detuning changes the effective rotation axis and rate, and can therefore appear as an additional calibrated control parameter in pulse-optimization procedures such as bootstrap




Figure 4.7: Ramsey measurement showing a drive detuning of 127.1(1.2) kHz.

pulse tomography [25]. Detuned driving is often employed when multiple nitrogen hyperfine transitions are simultaneously populated, such as when the host nuclear spin is not initialized [26]. In this work, however, our nuclear-spin polarization concentrates the population into a single hyperfine manifold, so we do not exploit detuning as an additional optimization parameter.

4.4 Nuclear Spin Characterization

With precise control over the electronic spin established, we extend sensitivity to surrounding nuclear spins using the CPMG sequence shown in Fig. 4.8(a). As discussed in Section 2.2, CPMG selectively amplifies or suppresses electron–nuclear interactions, acting as a band-pass filter for effective nuclear Larmor frequencies. By sweeping the interpulse delay 2τ , we perform spectroscopy of the local nuclear-spin environment.

We scan interpulse delays over a range of approximately 0.1–9 ns using a CPMG sequence with $N = 32$ refocusing π pulses. The resulting signal shows broad regions of reduced electronic coherence, as well as more isolated coherence dips, as shown in Fig. 4.8(b). These dips recur periodically because the same nuclear spins satisfy successive CPMG resonance conditions at different interpulse delays.

Several methods exist for estimating hyperfine parameters from CPMG spectroscopy. Here, the most prominent features were analyzed by Weitao He using an implementation based on the algorithmic decomposition approach described by Oh [15]. The signal is decomposed into Gaussian features corresponding to electron-coherence dips induced by nearby nuclear spins. Time windowing based on the ^{13}C Larmor precession is then used to assign a period index k . Features associated with the same nuclear spin across multiple periods are grouped, and their positions and widths provide initial guesses for the hyperfine parameters, as shown in the $\Delta\tau - k$ assignment plot in Fig. 4.8(c).

These initial estimates are then refined by fitting an analytical model to the measured




Figure 4.8: CPMG sequence (a) and scan (b) $\Delta\tau - k$ plot with spin assignments (c) and fitted spins overlay (d).

signal, yielding the axial and transverse hyperfine parameters, denoted A and B , respectively [27]. The resulting fit for three nuclear spins is overlaid on the measured CPMG signal in Fig. 4.8(d), and yields

$$(A_{\parallel}^{(1)}, A_{\perp}^{(1)}) = (-18.64(18), 18.36(69)) \text{ kHz},$$

$$(A_{\parallel}^{(2)}, A_{\perp}^{(2)}) = (32.74(15), 28.76(1.08)) \text{ kHz},$$

and

$$(A_{\parallel}^{(3)}, A_{\perp}^{(3)}) = (53.27(27), 34.37(1.65)) \text{ kHz}.$$

This CPMG data was acquired at a magnetic field of approximately 365 G, with a MW drive frequency of 1845.6 MHz, using similar MW power and pulse durations to the preceding room-temperature measurements.

4.4.1 Decoherence and Relaxation

The room-temperature CPMG measurements are limited by both signal-to-noise and finite electronic-spin coherence. As discussed earlier, room-temperature optical readout relies on the same spin-dependent relaxation pathway that also reinitializes the NV into $|m_S = 0\rangle$. Consequently, each shot provides only limited information about the pre-readout spin state, requiring extensive averaging. For long CPMG experiments, where each shot can last hundreds of microseconds, averaging over millions of repetitions extends scans over many hours and makes them increasingly susceptible to slow experimental drift.

A second limitation is finite electronic-spin coherence. As the total CPMG sequence duration increases, decoherence reduces the available contrast, limiting the ability to resolve weak or narrow nuclear-spin features.

To characterize these limitations, we measure both Hahn-echo coherence and longitudinal relaxation. The Hahn-echo sequence mirrors the Ramsey experiment, but includes a π




Figure 4.9: Hahn-echo coherence measurement with exponential decay fit yielding the NV spin coherence time $T_2 = 393(42) \mu s$.




Figure 4.10: Population relaxation measurement with exponential recovery fit giving $T_1 = 1011(31) \mu\text{s}$.

pulse at the midpoint of the 2τ free-evolution period. This refocuses quasi-static frequency shifts and extends the measured coherence time beyond T_2^* , producing an approximately exponential decay envelope [28]. From this measurement, we extract a Hahn-echo coherence time of $T_2 = 393(42) \mu\text{s}$, as shown in Fig. 4.9.

We also measure the longitudinal relaxation time T_1 using a π pulse to prepare the electronic spin in $|m_S = -1\rangle$, followed by a variable delay before optical readout. At room temperature, longitudinal relaxation is expected to arise primarily from spin-phonon processes, with two-phonon Raman processes dominant [9, 4, 24]. At a laser power of approximately 70% saturation 140 μW , we extract $T_1 = 1011(31) \mu\text{s}$, as shown in Fig. 4.10.

These measured coherence and relaxation times constrain the usable duration of room-temperature CPMG sequences, with the CPMG data exhibiting contrast decay on the same

order as the Hahn Echo T_2 . Thus, while the 200 ps timing resolution enables fine sampling of narrow spectral features, the room-temperature signal quality and sequence-dependent contrast loss limit the extent to which this timing advantage can be fully realized.

4.4.2 Cryostatic CPMG validation

To more directly demonstrate the advantage of 200 ps timing resolution, we use data collected on a cryogenic single-NV setup using the same QICK-based sub-clock timing methodology. This cryogenic testbed uses resonant optical readout, providing substantially improved signal-to-noise, and exhibits longer coherence times, with a measured Hahn-echo time of $T_2 = 1.05$ ms. The measurements were collected on the Diamond Quantum Processor testbed by Tommy Nguyen [27].

Figure 4.11 shows CPMG spectroscopy performed on the cryogenic testbed. Panel (a) shows a broad CPMG sweep acquired with 5.09 ns timing increments and $N = 32 \pi$ pulses. Panels (b) and (c) then demonstrate fine sweeps with 203 ps timing increments.

In panel (b), the fine timing resolution resolves sinusoidal modulation within a previously identified ^{13}C feature by using CPMG with $N = 320$ refocusing π pulses. This is then fit to an analytical model enabling a precise estimate for the hyperfine parameters of that ^{13}C nuclear spin, $(A_{\parallel}, A_{\perp}) = (225.01(12), 203.15(44))$ kHz.

In panel (c), fine timing resolution is used with $N = 32$ pulses to resolve narrow coherence dips whose shapes could not be characterized at coarser timing granularity. With 5.09, ns steps, only 3 dips are visible, each represented by a single sample. In contrast, the 203 ps sweep resolves the dip line shapes clearly and reveals two additional dips. The resolved dips all have FWHM values between 1 and 3.5 ns, so the native 3.2 ns timing grid not only under-samples their line shapes, but also risks missing them entirely. Overall, this corresponds to a 16-fold improvement in CPMG spectroscopic resolution over QICK-DAWG’s previous capability.




Figure 4.11: CPMG spectroscopy on the Diamond Quantum Processor. (a) Wide-range scan, and 200 ps scans of (b) a strongly coupled ^{13}C spin with $N = 320$ and (c) narrow coherence dips.

Chapter 5

CONCLUSION & OUTLOOK

5.1 *Conclusion*

This thesis set out to extend the practical timing resolution available for NV-center pulse-sequence experiments using QICK-DAWG. In standard QICK timing, pulse placement is naturally constrained by the sequencer clock, corresponding to a timing granularity of 3.2 ns. The result of this work is the implementation and experimental use of a methodology that realizes 200 ps timing control while preserving QICK’s sequencer-based structure. Because this approach is implemented through sequencer instructions rather than custom gateway, it remains highly flexible. Being applicable across different QICK gateway configurations, adaptable to a wide range of pulse sequences, and fully compatible with QICK’s efficient parameter-sweep framework.

As a result, pulse timings can be specified $16\times$ more finely than in native QICK scheduling. This finer timing control allows greater precision over free-evolution windows, improving the resolution available for nuclear-spin spectroscopy and enabling more precise control of electron–nuclear qubit interactions. More generally, the method makes QICK-DAWG a flexible and inexpensive research platform for sub-nanosecond quantum control, offering capabilities, such as extended-duration CPMG sweeps, that would otherwise require either more specialized hardware or additional gateway.

This methodology was used in pulse sequences throughout this work to characterize an NV center and its surrounding nuclear-spin environment. Many of the room-temperature characterization experiments do not strictly require the full 200 ps timing resolution to ob-

serve the relevant dynamics, but they can still benefit from the extra precision in pulse duration and delay. For CPMG sequences, and in cryogenic conditions in particular, the improvement in spectroscopic resolution is more significant, allowing narrower features and modulations in the frequency response to be resolved. This improvement also extends naturally to nuclear-spin control, where finer control over τ provides a route toward more precise engineering of electron–nuclear spin dynamics using QICK-DAWG.

The implementation also comes with important tradeoffs. Complex sequences using this methodology require careful balancing of waveform memory, instruction memory, sequencer execution time, and the available instruction structure. In some cases, pulses or delays must be grouped in ways that preserve timing while remaining compatible with the sequencer execution model. As a result, designing sequences at 200 ps resolution can require more careful reasoning than designing conventional QICK programs. The added timing flexibility therefore comes at the cost of increased programming complexity, and effective use of the method requires a detailed understanding of how the various constraints interact.

These tradeoffs also clarify how this method should be positioned within the broader QICK toolkit. A useful point of comparison is the recent development of the Real-Time CPMG-XY8n (RT-CPMG) Generator by Dr. Tiamike Dudley at the University of New Mexico [29]. RT-CPMG uses a custom CPMG pulse-generating module to produce CPMG-XY8 sequences in real-time logic, effectively offloading the processing burden from the tProcessor while consuming no waveform memory. This is an especially compelling approach for experiments where additional background processing is needed for active feedback or control of other channels. However, it also carries different constraints: for example, RT-CPMG is limited to square pulses, and its fundamental programmable block is the full XY8 sequence rather than an individual π pulse. The contrast between these approaches illustrates that the optimal implementation depends on the needs of the experiment. In some cases, the two methodologies may even be complementary.

Although this thesis primarily focuses on sequences with fixed pulse lengths and vari-

able interpulse delays, a similar waveform-banking approach can be extended to sequences requiring variable pulse lengths, such as Rabi experiments, provided the pulses contain an extendable constant-amplitude region. In this case, each pulse can be partitioned into starting, middle, and ending segments. The starting and ending segments are selected from waveform banks to achieve fine timing offsets, while the uniform middle portion can be repeated using QICK’s ability to loop over waveform segments. This is the approach used for the Rabi experiments in this work.

Overall, this work demonstrates that high-resolution experimental timing can be achieved on QICK-DAWG without abandoning its central strengths: programmable FPGA sequencing, efficient parameter sweeps, and reusable waveform resources. The result is a practical framework for 200 ps control of NV-center experiments, with immediate applications in electronic-spin calibration, nuclear-spin spectroscopy, and the development of more precise electron–nuclear control protocols.

5.2 Future Work

The most important practical limitation of the present approach is complexity. High-resolution pulse programs require careful management of instruction memory, waveform memory, loop structure, and sequencer timing. As pulse sequences become longer or more adaptive, this complexity becomes a substantial development burden. A natural future direction is therefore the development of a higher-level pulse compiler for QICK-DAWG. Such a compiler would allow users to specify a desired pulse sequence at a high level and automatically generate the corresponding QICK instructions, waveform banks, and timing offsets.

This abstraction would be especially useful because QICK supports multiple control strategies, each with distinct tradeoffs in timing precision, memory usage, and program flexibility. A compiler could map a user-specified sequence onto the appropriate hardware resources while tracking low-level constraints such as waveform-bank size, instruction count, and queueing margin. A model inspired by hardware-aware pulse languages such as QUA

may be useful, since it could preserve the flexibility of QICK while reducing the need for users to reason directly about each constraint.

A second direction is to apply the 200 ps timing methodology to more demanding dynamical-decoupling and nuclear-control sequences. CPMG spectroscopy provides a direct demonstration of the value of fine interpulse-delay control, but more complex sequences may be even more sensitive to timing discretization. For example, Uhrig dynamical decoupling uses nonuniform pulse spacings, making it poorly matched to coarse timing grids [14]. Sub-nanosecond timing resolution would reduce this discretization error and allow more direct comparisons between ideal pulse timings and experimentally realizable sequences.

A final direction, already underway, is the use of this methodology for nuclear-spin control, with the long-term goal of realizing nuclear spins memory registers.

Together, these directions point toward QICK-DAWG as a flexible platform for precision quantum control. The method developed in this thesis demonstrates that sub-clock-cycle control can be implemented within the existing QICK framework, while future work can focus on making this capability easier to program and applying it to more demanding control tasks.

BIBLIOGRAPHY

- [1] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller. Long-distance quantum communication with atomic ensembles and linear optics. *Nature*, 414(6862):413–418, November 2001.
- [2] M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, S. Wehner, and R. Hanson. Realization of a multinode quantum network of remote solid-state qubits. *Science*, 372(6539):259–264, April 2021.
- [3] C. M. Knaut, A. Suleymanzade, Y.-C. Wei, D. R. Assumpcao, P.-J. Stas, Y. Q. Huan, B. Machielse, E. N. Knall, M. Sutula, G. Baranes, N. Sinclair, C. De-Eknamkul, D. S. Levonian, M. K. Bhaskar, H. Park, M. Lončar, and M. D. Lukin. Entanglement of nanophotonic quantum memory nodes in a telecom network. *Nature*, 629(8012):573–578, May 2024.
- [4] Marcus W. Doherty, Neil B. Manson, Paul Delaney, Fedor Jelezko, Jörg Wrachtrup, and Lloyd C. L. Hollenberg. The nitrogen-vacancy colour centre in diamond. *Physics Reports*, 528(1):1–45, July 2013.
- [5] A Lohrmann, B C Johnson, J C McCallum, and S Castelletto. A review on single photon sources in silicon carbide. *Reports on Progress in Physics*, 80(3):034502, January 2017.
- [6] Nai-Jie Guo, Song Li, Wei Liu, Yuan-Ze Yang, Xiao-Dong Zeng, Shang Yu, Yu Meng, Zhi-Peng Li, Zhao-An Wang, Lin-Ke Xie, Rong-Chun Ge, Jun-Feng Wang, Qiang Li, Jin-Shi Xu, Yi-Tao Wang, Jian-Shun Tang, Adam Gali, Chuan-Feng Li, and Guang-Can Guo. Coherent control of an ultrabright single spin in hexagonal boron nitride at room temperature. *Nature Communications*, 14(1):2893, May 2023.
- [7] T. H. Taminiau, J. Cramer, T. Van Der Sar, V. V. Dobrovitski, and R. Hanson. Universal control and error correction in multi-qubit spin registers in diamond. *Nature Nanotechnology*, 9(3):171–176, March 2014.
- [8] C.E. Bradley, J. Randall, M.H. Abobeih, R.C. Berrevoets, M.J. Degen, M.A. Bakker, M. Markham, D.J. Twitchen, and T.H. Taminiau. A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute. *Physical Review X*, 9(3):031045, September 2019.

- [9] A. Jarmola, V. M. Acosta, K. Jensen, S. Chemerisov, and D. Budker. Temperature- and Magnetic-Field-Dependent Longitudinal Spin Relaxation in Nitrogen-Vacancy Ensembles in Diamond. *Physical Review Letters*, 108(19):197601, May 2012.
- [10] Lucio Robledo, Lilian Childress, Hannes Bernien, Bas Hensen, Paul F. A. Alkemade, and Ronald Hanson. High-fidelity projective read-out of a solid-state spin quantum register. *Nature*, 477(7366):574–578, September 2011.
- [11] Richard Monge, Tom Delord, Gergő Thiering, Ádám Gali, and Carlos A. Meriles. Resonant Versus Nonresonant Spin Readout of a Nitrogen-Vacancy Center in Diamond under Cryogenic Conditions. *Physical Review Letters*, 131(23):236901, December 2023.
- [12] M.H.M.A. Abobeih. *From atomic-scale imaging to quantum fault-tolerance with spins in diamond*. PhD thesis, Delft University of Technology, 2021.
- [13] T. H. Taminiau, J. J. T. Wagenaar, T. van der Sar, F. Jelezko, V. V. Dobrovitski, and R. Hanson. Detection and Control of Individual Nuclear Spins Using a Weakly Coupled Electron Spin. *Physical Review Letters*, 109(13):137602, September 2012.
- [14] M J Biercuk, A C Doherty, and H Uys. Dynamical decoupling sequence construction as a filter-design problem. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 44(15):154002, July 2011.
- [15] Hyunseok Oh, Jiwon Yun, M. H. Abobeih, Kyung-Hoon Jung, Kiho Kim, T. H. Taminiau, and Dohun Kim. Algorithmic decomposition for efficient multiple nuclear spin detection in diamond. *Scientific Reports*, 10(1):14884, September 2020.
- [16] Leandro Stefanazzi, Kenneth Treptow, Neal Wilcer, Chris Stoughton, Collin Bradford, Sho Uemura, Silvia Zorzetti, Salvatore Montella, Gustavo Cancelo, Sara Sussman, Andrew Houck, Shefali Saxena, Horacio Arnaldi, Ankur Agrawal, Helin Zhang, Chunyang Ding, and David I. Schuster. The QICK (Quantum Instrumentation Control Kit): Read-out and control for qubits and detectors. *Review of Scientific Instruments*, 93(4):044709, April 2022.
- [17] Chunyang Ding, Martin Di Federico, Michael Hatridge, Andrew Houck, Sebastien Leger, Jeronimo Martinez, Connie Miao, David Schuster I, Leandro Stefanazzi, Chris Stoughton, Sara Sussman, Ken Treptow, Sho Uemura, Neal Wilcer, Helin Zhang, Chao Zhou, and Gustavo Cancelo. Experimental advances with the QICK (Quantum Instrumentation Control Kit) for superconducting quantum hardware. *Physical Review Research*, 6(1):013305, March 2024.

- [18] Thiago P. Mayer Alegre, Charles Santori, Gilberto Medeiros-Ribeiro, and Raymond G. Beausoleil. Polarization-selective excitation of nitrogen vacancy centers in diamond. *Physical Review B*, 76(16):165205, October 2007.
- [19] J-P Tetienne, L Rondin, P Spinicelli, M Chipaux, T Debuisschert, J-F Roch, and V Jacques. Magnetic-field-dependent photodynamics of single NV defects in diamond: an application to qualitative all-optical magnetic imaging. *New Journal of Physics*, 14(10):103033, October 2012.
- [20] Antra Asare, Reinis Lazda, Florian Gahbauer, Volker Cimalla, and Marcis Auzinsh. Theoretical analysis and experimental measurement of N 14 nuclear spin polarization in nitrogen-vacancy centers in diamond. *Physical Review B*, 111(6):064113, February 2025.
- [21] 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*, 84(19):195204, November 2011.
- [22] Dennis Lönard, Isabel Cardoso Barbosa, Stefan Johansson, Jonas Gutsche, and Artur Widera. Limits of absolute vector magnetometry with NV centers in diamond, 2025. Version Number: 3.
- [23] Laima Busaite, Reinis Lazda, Andris Berzins, Marcis Auzinsh, Ruvin Ferber, and Florian Gahbauer. Dynamic ^{14}N nuclear spin polarization in nitrogen-vacancy centers in diamond. *Physical Review B*, 102(22):224101, December 2020.
- [24] Mykyta Onizhuk and Giulia Galli. Colloquium: Decoherence of solid-state spin qubits: A computational perspective. *Reviews of Modern Physics*, 97(2):021001, April 2025.
- [25] V. V. Dobrovitski, G. de Lange, D. Ristè, and R. Hanson. Bootstrap Tomography of the Pulses for Quantum Control. *Physical Review Letters*, 105(7):077601, August 2010.
- [26] Hyma H. Vallabhapurapu, Ingvild Hansen, Chris Adambukulam, Rainer Stöhr, Andrej Denisenko, Chih Hwan Yang, and Arne Laucht. High-fidelity control of a nitrogen-vacancy-center spin qubit at room temperature using the sinusoidally modulated, always rotating, and tailored protocol. *Physical Review A*, 108(2):022606, August 2023.
- [27] Victor Marcenac, Tommy Nguyen, Julie Chen, Weitao He, Enrique Garcia, Yuyang Han, Bethany E. Matthews, Tiamike Dudley, Andrew Mounce, Kai-Mei C. Fu, and Maxwell F. Parsons. Sub-nanosecond control for spin-defect quantum memories with a low-cost, compact FPGA platform, April 2026.

- [28] J. R. Maze, J. M. Taylor, and M. D. Lukin. Electron spin decoherence of single nitrogen-vacancy defects in diamond. *Physical Review B*, 78(9):094303, September 2008.
- [29] Tiamike Dudley. *Real-Time Waveform Synthesis for RFSOC-based Quantum Control Systems*. PhD thesis, University of New Mexico, Albuquerque, New Mexico, June 2026.
- [30] X.-D. Chen, C.-H. Dong, F.-W. Sun, C.-L. Zou, J.-M. Cui, Z.-F. Han, and G.-C. Guo. Temperature dependent energy level shifts of nitrogen-vacancy centers in diamond. *Applied Physics Letters*, 99(16):161903, October 2011.

Appendix A

NV GROUND-STATE HAMILTONIAN FOR B-FIELD ALIGNMENT

We use the NV^- ground-state Hamiltonian defined in Doherty et al. [4] to estimate the magnitude and alignment of the external magnetic field from the resonance dips. The Hamiltonian is given by

$$\begin{aligned} \hat{H}_{\text{gs}} = & D(T) \left(\hat{S}_z^2 - \frac{S(S+1)}{3} \right) + A_{\parallel} \hat{S}_z \hat{I}_z \\ & + A_{\perp} \left(\hat{S}_x \hat{I}_x + \hat{S}_y \hat{I}_y \right) + P \left(\hat{I}_z^2 - \frac{I(I+1)}{3} \right) + \hat{H}_Z, \end{aligned} \quad (\text{A.1})$$

where $D(T)$ is the temperature-dependent zero-field splitting. We model $D(T)$ using the sixth-order polynomial fit to experimental data from Chen et al. [30], for which $D \approx 2.87$ GHz at room temperature. The spin numbers are $S = 1$ and $I = 1$ for the ^{14}N NV ground state. We use $A_{\parallel} = -2.14$ MHz ($A_{\perp} = -2.7$ MHz) for the axial (transverse) nitrogen hyperfine coupling, and $P = -5.01$ MHz for the nuclear quadrupole parameter.

The Zeeman term is given by

$$\hat{H}_Z = \mu_B g_e (\mathbf{B} \cdot \hat{\mathbf{S}}) + \gamma_n (\mathbf{B} \cdot \hat{\mathbf{I}}), \quad (\text{A.2})$$

where μ_B is the Bohr magneton in frequency units. The electron g-factor is $g_e = 2.0028$, and the ^{14}N nuclear gyromagnetic ratio is set to $\gamma_n \approx 3.076$ MHz/T.