Actualizing High-Dimensional Qudits on a Trapped-Ion Quantum Computer

by

Gaurav Ashish Tathed

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Science
in
Physics (Quantum Information)

Waterloo, Ontario, Canada, 2025

© Gaurav Ashish Tathed 2025

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Trapped-ion processors offer unrivalled qubit fidelities but encounter heating and spectral crowding limits as chain lengths increase. This thesis mitigates those bottlenecks by encoding high-dimensional qudits in the 25 hyperfine–Zeeman sublevels of a single $^{137}Ba^{+}$ ion, thereby expanding the Hilbert space without adding additional ions. The work presented spans the key challenges of preparing arbitrary qudit basis states, efficiently maintaining calibration of the 80 transitions used to manipulate the ion, and benchmarking coherent control, culminating in performing a full quantum algorithm on a qudit.

A narrow-band optical-pumping sequence was devised to prepare any of the five $6S_{1/2}$ ($F = 2$) sub-levels and then selectively shelve population into the long-lived $5D_{5/2}$ manifold, achieving an average state-preparation-and-measurement (SPAM) fidelity of 99.51% across 25 levels. Two-point Ramsey interferometry, supported by analytic calculations of magnetic-field sensitivities and laser-polarisation geometries, allows for sub-100 Hz transition frequency calibration and Rabi frequency normalization.

Coherence was benchmarked with multilevel Ramsey interferometry on registers of dimension $d = 2 \ldots 24$. Star-topology encodings retained high contrast up to $d = 17$; beyond this, the bus-state architecture required for larger d was limited chiefly by 60 Hz and 180 Hz line noise, as confirmed by a noise-resolved Monte Carlo model that predicted measured contrasts to within $\pm 3\%$ with no free parameters.

Using the same physical toolbox, qudit registers were mapped onto virtual-qubit subspaces and algorithmic primitives executed: the Bernstein–Vazirani algorithm succeeded with probabilities of 95% (two virtual qubits) and 84% (three virtual qubits), while a four-virtual-qubit controlled-Toffoli gate reached 99.5% success—currently limited by SPAM error—using a single Givens rotation.

Collectively, these results establish $^{137}Ba^{+}$ as a versatile 25-level qudit platform and demonstrate that high-fidelity, noise-aware control of large qudit Hilbert spaces is already practical, opening a path toward resource-efficient, fault-tolerant quantum computing with far fewer ions than traditional qubit-based architectures require.

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Professor Crystal Senko, for being an exceptional mentor throughout my Master's program. Her guidance, support, and encouragement gave me the freedom to explore my interests while providing steady direction whenever I needed it.

I am also thankful to my advisory committee members, Professors Alan Jamison and Michal Bajcsy, for their valuable feedback and for keeping our meetings concise and efficient.

A special thanks to my labmates and friends, Pei Jiang Low and Nicholas Zutt , for making long days in the lab both productive and enjoyable. It was truly a pleasure working alongside them.

I am incredibly grateful to my parents, Ashish Tathed and Sonal Tathed, for being a constant pillar of support throughout my undergraduate and master's studies. Their unwavering belief in me made this journey possible.

To Ananya, thank you for calling me out of the blue just to ask random questions — those moments meant more than you realize. To Dolly, thank you for simply being adorable. To Shihochan, thank you for bringing back my sanity after long days of coding and for always appreciating the pretty plots I made.

Lastly, I would like to thank all my gaming friends for providing a much-needed escape and helping me unwind after long days in the lab. Your friendship made this experience more balanced and fun.

Dedication

This thesis is dedicated to everyone's curiosity — the quiet force that drives us to ask, seek, and understand.

Table of Contents

Author's Declaration ii

Abstract iii

Acknowledgements iv

Dedication v

List of Figures ix

List of Tables xi

1 Introduction 1

2 $^{137}Ba^{+}$ — A 25 level workhorse 5

2.1 Energy Level Structure 5

2.1.1 Principal Lasers and modulation Schemes 7

2.2 Experimental Setup 12

2.2.1 Magnetic field Source — Permanent Magnets 13

2.2.2 Beam geometry and Polarizations 16

2.3 Theory tools 17

3 State Preparation and Measurement **24**
3.1 $6S_{1/2}$ Initialization . . . 24
3.1.1 614 nm Repump path . . . 28
3.1.2 Spectral crowding . . . 29
3.1.3 F=1 Pump parameters . . . 31
3.2 $D_{5/2}$ Initialization . . . 33
3.3 State-Selective Readout . . . 34
3.4 SPAM Results . . . 35
3.4.1 Heralding Techniques and Error breakdown . . . 37

4 Calibrations — Transition Frequencies and Strengths **39**
4.1 Qubit Dynamics and Ramsey Interferometry . . . 40
4.2 Ramsey-Based Frequency Calibration — Practical Considerations . . . 45
4.3 Data collection and Fitting . . . 50
4.4 Rabi-Frequency Calibration . . . 53

5 Coherence Characterization — Qudit Ramsey **55**
5.1 Qudit Ramsey Scheme . . . 56
5.2 State Selection and Qudit encoding . . . 58
5.2.1 Qudit encoding . . . 60
5.3 Qudit Ramsey Experiments — Fringe contrast vs. Qudit Dimensions . . . 63

6 Error characterization **66**
6.1 Magnetic field noise . . . 67
6.1.1 Bussed Ramsey sequence and Laser noise cancellation . . . 67
6.1.2 Experimental details . . . 69
6.2 AC Line Noise . . . 71
6.2.1 Experimental Results . . . 73

6.3 Laser Noise . . . 74
6.3.1 Experimental results . . . 75
6.4 Frequency calibration Error . . . 78
6.5 Pulse-angle error . . . 78
6.6 Simulation Model . . . 80

7 Virtual Qubits **83**
7.1 Bernstein–Vazirani Algorithm . . . 84
7.2 Unitary Decomposition under a Star–Topology Constraint . . . 85
7.2.1 Hadamard Decomposition . . . 89
7.3 Experimental Results . . . 91
7.3.1 Bernstein–Vazirani with two and three virtual qubits . . . 91
7.3.2 Controlled-Toffoli Gate . . . 92

8 Conclusion and Outlook **94**

References **96**

APPENDICES **103**

A Quadrupole Geometric Factors **104**

B Intermittent Magnetic Field Noise **109**

C Permanent Magnets Simulations **113**

D SPAM Bright-Dark discrimination threshold **117**

E Effects of Line Noise on Coherent operations **119**

F All Relative Magnetic field Sensitivities **122**

G Bussed Qubit phase scans **125**

List of Figures

2.1 Energy-level structure of $^{137}Ba^{+}$ 6
2.2 Two-photon Ionization Scheme. 8
2.3 Fluorescence laser modulation scheme. 10
2.4 Shelving Laser Modulation Scheme 11
2.5 Permanent Magnet Configuration. 14
2.6 Beam Paths and geometry 15
2.7 Energy Level splitting vs. Magnetic field. 19
2.8 Magnetic field Sensitivity Chart. 22
2.9 Transition Strength Chart. 23

3.1 $6S_{1/2}$ Initialization scheme 25
3.2 Shelving transitions for $6S_{1/2}$ initializations. 26
3.3 Initialization probability vs. Repetitions. 27
3.4 Repump path Map from $6D_{5/2} \rightarrow 6P_{3/2} \rightarrow 6S_{1/2}$ 29
3.5 Off-resonant Scattering probability. 30
3.6 $F = 1$ pump Parameter Chart. 32
3.7 SPAM State encoding. 33
3.8 SPAM Results with average infidelity and data rate. 36

4.1 2-point Calibration Ramsey fringes and Detuning sweep. 47
4.2 π-times and Calibration, reference transitions. 51

5.1 Qudit Ramsey phase scan $d = 5$ and $|0\rangle$ population colormap. 57
5.2 Qudit state encoding for all dimensions. 61
5.3 Qudit Ramsey phase scans. 62
5.4 Qudit Ramsey Contrast vs. Dimension and Operation time. 65

6.1 Bussed Ramsey Contrast vs. Ramsey Wait Time. 70
6.2 Bussed Ramsey T_2^* vs. Magnetic Field Sensitivity κ 72
6.3 AC Line signal for Permanent magnets and EM-coils. 74
6.4 Laser Noise Characteristic Function and Estimated PDF. 76
6.5 Calibration Noise. 77
6.6 Relative transition strength comparison between calibrated theory and measurements. 79

7.1 BVA results on two and three virtual qubits. 91
7.2 Controlled-Toffoli gate on four virtual qubits. 93

A.1 Geometric Factors. 108

B.1 Qubit Ramsey - Inconsistent Contrast. 110
B.2 Voltage Reading from Keysight E3631A Power Supply. 111
B.3 Variation in B-field due to temperature fluctuations. 112

C.1 Permanent Magnet Geometry and Meshing. 115
C.2 Remnant flux density estimation and B-field at the ion position. 116

D.1 SPAM bright/Dark discrimination threshold. 118

E.1 Bussed Qudit Ramsey for d=5. 120
E.2 BVA Simulation on 3-virtual qubits. 121

F.1 All magnetic field sensitivities. 123
F.2 All effective coherence times. 124

G.1 Bussed Qubit Phase scans. 126

List of Tables

2.1 Beam-path summary relative to the quantization axis 16

2.2 Experimental hyperfine constants for selected levels of $^{137}Ba^{+}$. 18

6.1 Fitted coefficients of Eq. (6.15) for the electromagnetic-coil and permanent-magnet field sources. Amplitudes are given in mG; phases (in radians) are referenced to the laboratory 60 Hz trigger. 75

6.2 Noise sources, profiles, and distribution widths. 81

Chapter 1

Introduction

Quantum computing has advanced rapidly over the past decades, with multiple hardware platforms (superconducting circuits, trapped ions, neutral atoms, etc.) now demonstrating control over tens of qubits[1, 2, 3]. Trapped-ion quantum processors are among the leading architectures, offering fully-connected qubit interactions and some of the highest gate fidelities in the field. Trapped-ion quantum computing encodes quantum information in the long-lived hyperfine or Zeeman levels of laser-cooled ions that are confined in electro-magnetic fields. The collective quantized motion of an ion string provides a bus for entangling interactions, while high efficiency state-dependent fluorescence enables projective read-out. Cirac and Zoller were the first to articulate this architecture, showing that universal quantum logic becomes possible when the collective motional modes of an ion string are used to entangle any desired pair of ions—a result that laid the foundation for scalable trapped-ion processors [4]. And subsequent reviews have documented steady progress in experimental control, error sources and scaling strategies [3].

Modern trapped-ion processors satisfy the first five DiVincenzo's criteria for a physical quantum computer outright [5]. For example, techniques like Narrow band optical pumping for state initialization and cabinet shelving procedure for measurement allowed a team at Quantinuum to achieve a state-preparation-and-measurement (SPAM) infidelity of just $(9.0\pm1.3)\times10^{-5}$ [6]. Single qubit rotations implemented with resonant Raman or microwave drives have reached fidelities above 99.99%, while two qubit entangling gates based on Mølmer–Sørensen interaction or geometric phase has surpassed 99.9% [7, 8] Hyperfine "clock" qubits boast coherence times exceeding ten seconds, and cryogenic surface traps further suppress motional heating [9, 10, 11]. Together, these features underpin universal, high fidelity gate sets appropriate for error-corrected operations.

The platform's maturity is evident in recent mid-scale processors. IonQ's Forte system arranges 30 $^{171}Yb^{+}$ ions in a single chain with fully connected two-qubit gates, and comprehensive benchmarking shows a median two-qubit fidelity of 99.5% [12]. Quantinuum's H2-1 processor hosts 56 $^{171}Yb^{+}$ and $^{138}Ba^{+}$ ion pairs in a racetrack-style surface trap; at its launch in June 2024 the company reported record two-qubit-gate fidelities, and in 2025 the very same machine enabled an experiment that generated device-independent certified randomness—a feat that cannot be reproduced on any classical computer [13, 14]. These achievements confirm that trapped-ion hardware has entered the regime where modest numbers of logical qubits and small-scale fault-tolerant protocols are feasible.

As trapped-ion systems scale beyond a few dozen ions, two major effects begin to erode the exquisite control that underpins trapped-ion quantum processors. First, motional heating increases as the ion chain lengthens. The axial centre-of-mass (COM) mode, in which every ion oscillates in phase, acquires a secular frequency that drops roughly as 1 / N once the Coulomb coupling of many ions begins to compete with the external axial confinement, leaving only a narrow energy gap between the COM mode and the broadband electrode noise that permeates the trap. Because that noise is essentially spatially uniform, it couples almost exclusively to the COM motion, so the heating rate rises sharply with chain length. This uniform-field mechanism was identified in the seminal ground-state-heating measurements of Turchette et al in Ref. [15] and reinforced by the extensive survey of Brownnutt et al., which shows that low-frequency modes are universally the most vulnerable to surface noise and heat an order of magnitude faster than their higher-frequency counterparts in Ref. [16]. A second obstacle is spectral crowding. A linear crystal supports 3N collective vibrations; as N grows, adjacent axial frequencies squeeze together. Entangling gates that rely on resolving one or two normal modes must then slow down dramatically to avoid off-resonant excitation due to power broadening or resort to sophisticated pulse shapes that excite several modes at once. Techniques such as amplitude-modulated multi-mode Mølmer–Sørensen gates [17], EASE protocols [18], and power-optimal control sequences [19] have pushed high-fidelity operation into the tens-ion regime, yet the classical optimisation needed to design those pulses scales poorly with N, and theoretical analyses indicate that gate time is ultimately bounded by the narrowest mode spacing—implying at least quadratic slow-down for very long chains[20].

Architectures based on shuttling or dynamical segmentation can mitigate these problems, but they still grow the hardware footprint linearly with the number of physical qubits. This thesis, therefore, turns to a deeper solution: **increase the amount of quantum information carried by each ion instead of adding more ions**. By elevating the two-level qubit to a multilevel qudit encoded in the rich hyperfine or Zeeman structure of a single ion, one expands the computational Hilbert space without lengthening the chain or

forcing the motional spectrum to compress, thereby sidestepping both heating and crowding constraints while opening new avenues for resource-efficient quantum logic.

Trapped-ion qudits were proposed soon after the earliest qubit experiments as a way to exploit the rich multilevel hyperfine and Zeeman structure that naturally exists in ions. Foundational theory established that multilevel trapped-ion systems can support universal quantum logic with a finite, two-body gate set: Muthukrishnan & Stroud proved that any unitary on d-level ions factors into pairwise multivalued gates [21], while Klimov et al. supplied an explicit trapped-ion controlled-NOT for qutrits and showed how to compile arbitrary algorithms from it [22]. Subsequent analytic work by Mischuck & Mølmer extended these constructions to closed-form pulse sequences for one- and two-qudit operations of arbitrary dimension in realistic ion–oscillator Hamiltonians [23] Bridging the gap between these theoretical blueprints and full-scale hardware, Low et al. assembled and numerically validated a complete trapped-ion qudit toolbox—covering high-fidelity state preparation, shelving-based read-out, arbitrary single-qudit rotations, and a generalized Mølmer–Sørensen entangling gate—for dimensions up to $d = 5$, thereby providing the practical recipe book that underpins the recent experimental breakthroughs[24]. Experimental milestones have followed quickly in the past few years. In 2022, Ringbauer et al. demonstrated a universal trapped-ion processor that operates natively in local Hilbert-space dimensions up to $d = 7$, showing single-qudit fidelities and algorithmic performance on par with the best qubit devices [25]. Building on that control, a 2023 experiment realized a native two-qudit entangling gate whose calibration overhead is independent of dimension and produced genuine entanglement for $d \leq 5$ in a single gate application [26]. These results confirm that high-dimensional operations can reach the same ($\approx 99\%$) fidelity frontier as binary gates while carrying more information per ion. The utility of qudits is already visible in application-level tests. Multi-level encoding compresses circuit depth for algorithms with natural high-dimensional structure such as quantum Fourier transforms, lattice gauge simulations, and encoding multiple virtual qubits in a single ion. The recent lattice-gauge-theory simulations leveraged an all-to-all-connected trapped-ion qudit processor to model two-dimensional gauge fields more efficiently than any qubit mapping could manage [27].

Certain ion species lend themselves naturally to multilevel encodings. Practical-protocol surveys, such as Low et al. (2020), highlight $^{43}Ca^{+}$, $^{88}Sr^{+}$, and $^{171}Yb^{+}$ as particularly attractive hosts owing to their long-lived excited manifolds and clean laser-accessible transitions [24]. Among these candidates, $^{137}Ba^{+}$ is especially compelling: its metastable $D_{5/2}$ manifold remains coherent for tens of seconds, providing a spacious and reasonably-isolated arena for qudit logic. Building on this advantage, our group has already demonstrated state preparation and read-out across 13 levels of a single $^{137}Ba^{+}$ ion [28]. This thesis

now advances that work by exploiting the same manifold to approach the full 25-level capacity of $^{137}Ba^{+}$, thereby probing the practical limits of high-dimensional encoding in a single trapped ion. Following this introductory chapter, the thesis opens with a Background section that lays out the atomic and hyperfine energy level structure of $^{137}Ba^{+}$, and reviews the laboratory hardware—the linear Paul trap, laser systems, and control electronics—referencing prior group theses wherever the apparatus is unchanged. The discussion then introduces the theoretical tools that support the experimental work, including calculations of oscillator strengths and Magnetic field sensitivities. Building on this foundation, the next chapter, State Preparation and Measurement, explains how a narrow-band optical-pumping scheme is adapted and engineered to initialize population in the $F = 2$, $6S_{1/2}$ manifold, describes the considerations and optimizations need to get high-fidelity initializations, and then details the fluorescence-based read-out protocol, with particular emphasis on the heralding filters that remove experiments containing easily recognizable errors and thus enhance effective SPAM fidelity. The subsequent chapter, Coherence Characterization, turns to multilevel Ramsey spectroscopy as a metric for coherence charaterization, presenting the pulse sequence, state selection strategy, and an analysis of how the observed Ramsey coherence decays as the qudit dimension grows, thereby benchmarking phase stability across the entire Hilbert space at once. Finally, Error Budget and Fidelity Prediction draws together independently measured noise channels—magnetic-field fluctuations measured in situ, laser phase and intensity noise, and calibration drift in transition frequencies and their Rabi frequencies—and folds them into a Monte-Carlo model that predicts experimental gate-set fidelities to within three percent of observation, providing a quantitative, self-consistent picture of the decoherence mechanisms that limit high-dimensional operations in trapped-ion qudits.

Chapter 2

$^{137}Ba^{+}$ — A 25 level workhorse

This chapter positions $^{137}Ba^{+}$ as the central platform for the 25-level qudit experiments presented in the remainder of the thesis. It first lays out the ion's electronic, hyperfine, and Zeeman structure, focusing on the long-lived $5D_{5/2}$ manifold whose rich ladder of levels enables high-dimensional encoding. The narrative then moves from spectroscopy to hardware, describing the linear Paul trap, laser, and magnetic-field source that together afford precise manipulation and read-out of a single barium ion. The chapter concludes by introducing the theoretical tools — oscillator strength and magnetic field sensitivity calculations from solving the system's Hamiltonian — that underpin the calibrations and noise characterization in subsequent chapters.

2.1 Energy Level Structure

In this section I lay out the electronic, hyperfine and Zeeman structure of $^{137}Ba^{+}$, emphasizing the long-lived $5D_{5/2}$ manifold that underpins high-d encoding. Finally, I conclude with a summary of key hyperfine splittings and radiative pathways, setting up the transitions we will use experimentally.

Singly-ionized barium behaves, to first order, like an alkali atom with a single valence electron outside a xenon-like closed shell. In the odd-mass isotope $^{137}Ba^{+}$, the nuclear spin $I = 3/2$ couples to the electronic angular momentum J via magnetic dipole (and electric-quadrupole) interactions. This splits every fine-structure level into a set of hyperfine states labelled by $F = I + J$. Figure 2.1 shows the Fine structure and laser wavelength used in the experiments, as well as the branching ratios for each orbital. For the ground state, $6S_{1/2}$,

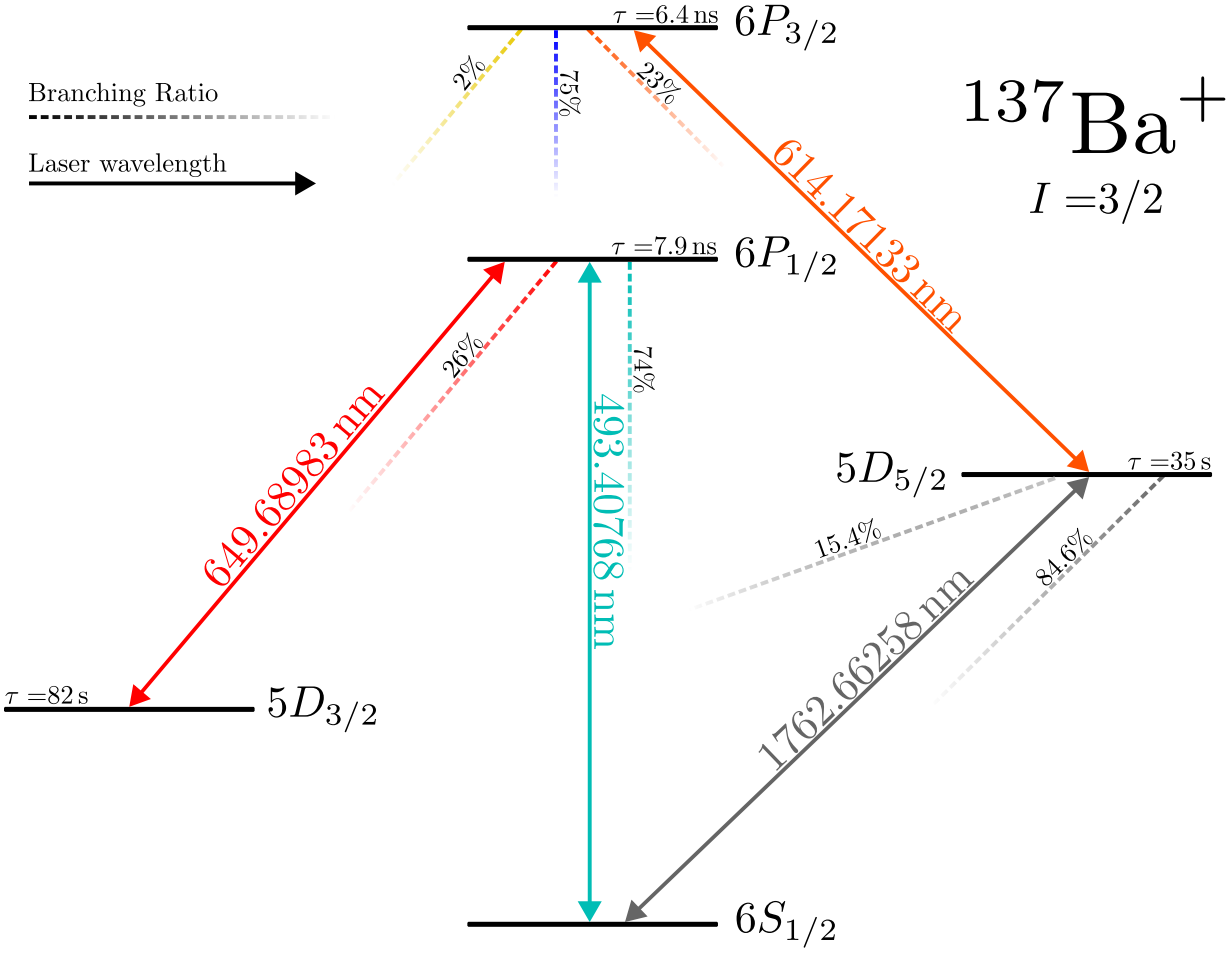


Figure 2.1: **Energy-level structure of $^{137}Ba^+$**, annotated with the laser wavelengths used in our experiment and the branching ratios for each orbital.

there are two such hyperfine states, $F = 1$ and $F = 2$, separated by 8.038 GHz at zero B-field [29]. The first excited level, $6P_{1/2}$, is likewise split into $F = 1$ and $F = 2$ components that lie 930 MHz below and 558 MHz above the fine structure center, respectively, giving total separation of 1.49 GHz [30]. Spontaneous decay from $6P_{1/2}$ returns to the ground state with 74% probability, while 26% of the population leaks into the metastable $5D_{3/2}$ level [31, 32]. The $5D_{3/2}$ itself possesses four hyperfine components ($F = 0$-3) spread over roughly 1.1 GHz and a radiative lifetime of about 80 s, making it a quasi-stable leak channel during fluorescence. The higher-lying $6P_{3/2}$ level exhibits $F = 0$-3 components, spanning

over 670 MHz. Its decay predominantly returns to $6S_{1/2}$, with 75% probability, but also populates $5D_{5/2}$ manifold, with 23% probability, and to a lesser extent $5D_{3/2}$ with 2% probability [33]. The $5D_{5/2}$ manifold, with its four hyperfine states, enjoys a 32 s lifetime, with possible decay channels to $6S_{1/2}$ and $5D_{3/2}$, with 85% and 15% chance respectively [34].

Taken together, three structural features are central to this thesis:

1. The "three" level cycling transitions between $6S_{1/2} \leftrightarrow 6P_{1/2} \leftrightarrow 5D_{3/2}$ used for fluorescence of the ion.
2. The long-lived $5D_{5/2}$ level enables shelving and hosting of high-dimensional qudit encoding.
3. And connectivity between $6S_{1/2} \leftrightarrow 5D_{5/2} \leftrightarrow 6P_{3/2} \rightarrow 6S_{1/2}$ used for pumping the population out of $5D_{5/2}$ into $6P_{3/2}$ from where the population decays to $6S_{1/2}$.

2.1.1 Principal Lasers and modulation Schemes

In a linear Paul trap, oscillating RF quadrupole fields provide radial confinement while static DC endcaps confine axially; a ∼4 G magnetic field sets the quantization axis. Once loaded and Doppler cooled, a single $^{137}Ba^{+}$ ion typically remains trapped for hours to a day, long enough for extended calibration and data runs. When a new ion is needed, neutral barium is released from a $BaCl_2$ salt target and the desired isotope is loaded via isotope-selective photoionization directly into the trap, after which standard cooling prepares the ion for spectroscopy and control and computation.

In daily operation we rely on four key wavelengths for experiments and three lasers for loading and trapping. The underlying optical architecture is essentially unchanged from the system commissioned prior to the start of this work; a complete schematic, parts list and alignment procedure are given in Ref.[35, 36]. Apart from minor updates to EOM drive frequencies, all mirrors, fibres, and dichroics remain as originally installed. In the remainder of this section, I summarize the key instrumentation details and, where applicable, the modulation schemes for lasers employed in our experiments.

532 nm — Pulsed ablation laser

We begin every ion-loading cycle by creating a plume of neutral barium atoms with a *Minilite II*, a Q-switched Nd: YAG laser at 532 nm, provided by *Amplitude Laser Group*.

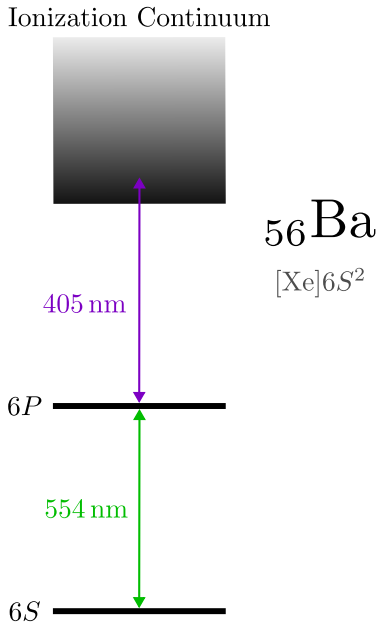


Figure 2.2: **Two-photon Ionization Scheme.** Two photon ionization scheme using 554 nm and 405 nm lasers.

The laser delivers 3-5 ns pulses that pass a half-wave-plate/polarizer pair for energy control, a beam sampler that feeds a power meter, and an ND filter that equalizes the monitored and delivered energies before a 200 mm focal length lens focuses the light onto $BaCl_2$ salt target inside the vacuum chamber. The resulting 196 µm diameter beam gives adjustable fluence (by rotating the HWP); we routinely operate between 0.2 and 0.4 J cm^{-2} depending on the condition of the target spot, which is well below previously record threshold level (>3 J cm^{-2}) to damage the target. Because sub-microsecond synchronization is unnecessary at this stage, pulse triggering is handled by the *Minilite*'s internal sync output, and the beam passes through a mechanical shutter.

554 nm — First–step photo-ionization laser

The neutral atoms liberated by the 532 nm ablation pulse traverse the trap center within a few microseconds. To excite the neutral atoms in an isotope-selective manner, we overlap this neutral plume with a resonant 554 nm (541.433 315 THz) beam that drives the strong $6S \leftrightarrow 6P$ transition — step one of the two-photon resonance-enhanced multiphoton-ionization (REMPI) scheme (Figure 2.2). The exact 554nm frequency is derived from the

previous work done in our group in Ref. [37].

The 554 nm light is generated by frequency-doubling a 1107 nm external-cavity diode laser (ECDL). A 1% of doubled power is sent to the Fizeau-interferometer wavemeter (WS-8) provided by *HighFinesse*, and a PID loop acting on the ECDL grating piezo keeps the absolute frequency parked on the desired isotope line. Because the ionization step is not time-critical, the main beam is gated with a mechanical shutter; downstream, an HWP followed by a polarizer permits stable power adjustment without perturbing the lock.

405 nm — Second-step photo-ionization laser

A laser diode provides the 405 nm photons needed to raise the $6P$ state to the continuum and complete the two-photon REMPI process. After first-step excitation at 554 nm, the atom still needs approximately 3 eV to reach the ionization threshold ($E_{\text{ion}} = 5.21$ eV). A single 405 nm photon carries ~ 3.06 eV, comfortably above this value, so the process is insensitive to THz-level frequency fluctuations. Consequently, the diode is operated open-loop; no piezo or PID servo is required. The beam is mechanically shuttered and timed to overlap the trailing edge of the ablation plume.

493 nm — Doppler-cooling, Fluorescence and Optical-pumping beam

A *Toptica DL-pro* tunable–diode laser is frequency-stabilized with slow feedback from a HighFinesse WS-8 wavemeter and held at 607.431 93 THz ($\lambda \approx$493.406 nm), close to the $\left|6S_{1/2}, F = 1\right\rangle \rightarrow \left|6P_{1/2}, F = 2\right\rangle$ transition. A fibre-coupled phase-modulator (AdvR WPM-K0493) is driven at 8.1 GHz; the first blue-shifted sideband therefore sits near the $\left|6S_{1/2}, F = 2\right\rangle \rightarrow \left|6P_{1/2}, F = 2\right\rangle$ resonance, so that the carrier–sideband pair simultaneously addresses both hyperfine ground states and closes the Doppler-cooling cycle.

A polarising beam-splitter after the EOM diverts $\sim$90% of the optical power to the *fluorescence / Doppler-cooling* arm and the remaining $\sim$10% to an *optical-pumping* arm. Each arm passes a single-pass AOM driven at 200 MHz; the first-order beam acts as a sub-100 ns intensity switch, triggered by the experiment sequencer.

650 nm — Fluorescence Repump beam

Roughly 26% of the population excited by the 493 nm light leaks into the long-lived $\left|5D_{3/2}\right\rangle$ manifold (lifetime $\approx$80 s), so a repumper is required. We use a second *Toptica*

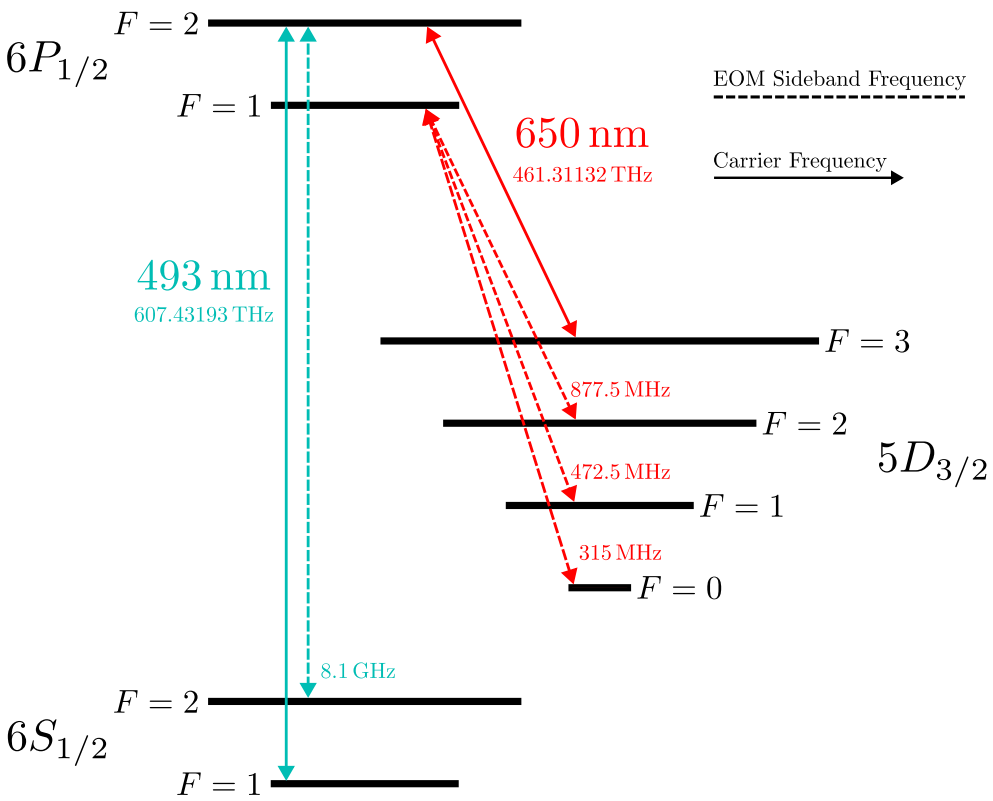


Figure 2.3: **Fluorescence laser modulation scheme.** Modulation Scheme used for Ion Fluorescence using 493 nm and 650 nm. Dashed lines are sidebands generated by the respective EOMs.

DL-pro diode laser, wavemeter-locked near the $\left|5D_{3/2}, F = 3\right\rangle \rightarrow \left|6P_{1/2}, F = 2\right\rangle$ transition at 461.311 32 THz. A free-space phase-modulator adds three additional sidebands at 877 MHz, 473 MHz and 318 MHz; together with the carrier these frequencies connect the $\left|5D_{3/2}, F = 2, 1, 0\right\rangle$ states to $\left|6P_{1/2}, F = 1\right\rangle$, Figure 2.3 shows the EOM sidebands used in ion fluorescence.

The 650 nm beam then traverses a single-pass AOM at 200 MHz, giving identical pulse-picking and frequency agility to the 493 nm paths. After the AOM, the repump light is overlapped with the cooling beam on the same dichroic, shares the reflective delivery telescope, and is focused to a ∼60 µm waist.

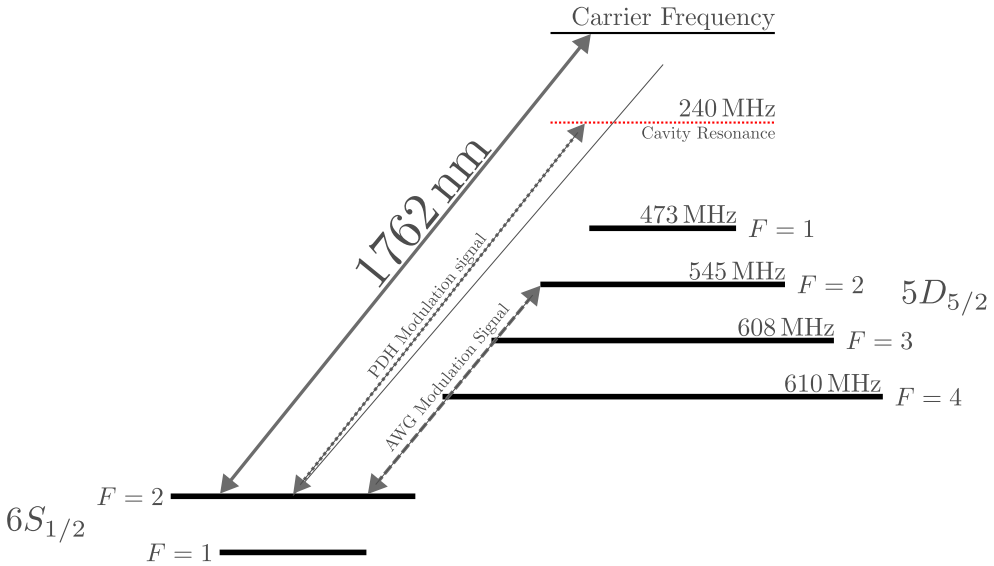


Figure 2.4: **1762 nm Laser modulation schemes** for Locking and to address individual transitions using AWG.

1762 nm — Quadrupole shelving beam

We stabilize our 1762 nm quadrupole laser to a high-finesse ultra-low-expansion (ULE) reference cavity (provided by Stable Laser Systems) using a Pound–Drever–Hall (PDH) scheme that combines a fixed and a tunable radio-frequency (RF) modulation on a single EOM. A permanent 5 MHz modulation is applied to generate the sidebands required for PDH demodulation; by mixing the reflected light with this reference, we obtain the dispersive error signal that is fed back to the laser, ensuring long-term frequency stability. Simultaneously, we drive the same EOM with a second RF tone, Δ_{PDH}, whose frequency we choose to set the carrier's detuning from the nearest cavity resonance. The red sideband ($-\Delta_{\mathrm{PDH}}$) is locked to a cavity mode, so the carrier is held exactly Δ_{PDH} away from that mode while inheriting the cavity's stability. In daily operation we set $\Delta_{\mathrm{PDH}} \approx 240\,\mathrm{MHz}$, a value that keeps the carrier well outside any $6S_{1/2} \Leftrightarrow 5D_{5/2}$ transition yet leaves the first-order sideband of a second, downstream EOM—driven at Δ_{AWG} by an Arbitrary Waveform Generator (AWG) —precisely resonant with the desired transition in the ion shown in Figure 2.4. This arrangement lets us tune or switch the experiment frequency in software without retuning the cavity or disturbing the PDH lock. At the same time, higher-order sidebands remain safely detuned from all spectroscopic features.

614 nm — Shelving-repump beam.

Population shelved in the long-lived $5D_{5/2}$state during quadrupole operations must be returned to the cooling cycle before the next experiment shot. We accomplish this with 614 nm laser, tuned to the $\left|5D_{5/2}, \tilde{F}=4, m_{\tilde{F}}=2\right\rangle \rightarrow \left|6P_{3/2}, \tilde{F}=3\right\rangle$ transition, which decays rapidly to the $6S_{1/2}$ ground state. The 614 nm beam is produced by frequency-doubling a 1228 nm ECDL. A 50:50 beamsplitter sends half the power WS-8, whose PID loop locks the carrier within ±20 MHz of the repump line; the remaining half is routed to the experiment. After the doubler, the experimental arm passes an AOM driven at 200 MHz, giving sub-100 ns intensity switching that is aligned to the experimental sequencer.

Together, the cooling/detection pair closes the fluorescence cycle while dedicated sidebands cover all required ground and leak states. A cavity-stabilized narrow-line drive provides coherent control with software-tunable targeting, and a dedicated depopulation beam returns shelved population to the readout pathway. Fast, synchronized AOM gating keeps pulse sequencing straightforward.

2.2 Experimental Setup

This section consolidates the hardware used to trap and control a single $^{137}Ba^{+}$ ion: a linear Paul trap at the center of a UHV spherical-octagon chamber with dedicated laser access, RF/DC feeds, and integrated pumping.

The experiments are carried out in a conventional four-rod linear Paul-trap that was designed and assembled by earlier members of the group; complete mechanical drawings, wiring diagrams, and alignment protocols are given in Ref. [35, 36, 38, 39]. The essential parameters are collected in this section for ease of reference. The trap sits at the center of a 4.5-inch ConFlat (CF) spherical-octagon UHV chamber whose eight 0.75-inch side ports provide laser access, RF + DC bias feed through. Vacuum is maintained by by a single *SAES NEXTorr Z-200* hybrid pump mounted on the manifold; its integrated 6 Ls^{-1} diode ion-pump handles noble gases and methane, while 200 Ls^{-1} non-evaporable-getter (NEG) cartridge provides high throughput for hydrogen and the other active gases released during ablation shots.

The trap consists of four 0.5 mm diameter tungsten rods arranged at the corners of 1 mm square, so the ion-electrode distance is $r_0 = \sqrt{2}$ mm. One pair of diagonal rods is driven at $\sim +120$ V while the opposite diagonal is driven with the same amplitude, but π-shifted

phase (for a total amplitude $V_{\mathrm{RF}} \sim 240\,\mathrm{V_{pp}}$); the drive frequency is $\Omega_{\mathrm{RF}} = 20.776\,\mathrm{MHz}$. To break the axial symmetry and help compensate stray fields, a static bias of +3 is applied to the pair carrying the +180 RF phase, while the opposite pair remains at dc ground. Two tungsten needle end-caps provide axial confinement, 2.8 mm apart, 11.5 V and 16 V. Under these conditions, a single $^{137}\mathrm{Ba}^{+}$ ion exhibits radial secular frequencies of (1.27, 1.46, and 0.21) MHz. The micromotion is nulled by trimming the rod biases by $\pm 200\,\mathrm{mV}$, reducing the micromotion sidebands—observed via the 1762 nm transition—to $< 5 \times 10^{-3}$ of the carrier.

2.2.1 Magnetic field Source — Permanent Magnets

This subsection motivates replacing the coil-driven bias field—whose line-synchronous and intermittent instabilities limited coherence—with a passive permanent-magnet source. It outlines the design goals (stable quantization axis, preserved optical access, drop-in install) and the modeling-driven configuration that meets them.

A well-defined static magnetic field establishes the quantization axis, lifts Zeeman degeneracies and sets the transition splittings that encode our 25-level qudit. Several of the $\left|5D_{5/2}, \tilde{F}, m_{\tilde{F}}\right\rangle \leftrightarrow \left|6S_{1/2}, F = 2, m_F\right\rangle$ transitions have magnetic field sensitivity of $\sim 3\,\mathrm{MHz\,G^{-1}}$. Consequently, even a drift of a few milligauss produces frequency shifts on the order of our typical Rabi rates, directly limiting coherence times.

In the original apparatus, this field ($B_0 = 4.216\,\mathrm{G}$) was generated by a single copper coil. Although the supply was operated in constant-current (CC) mode, it injected line-synchronized noise at 60 Hz and its third harmonic (180 Hz). We also observed that the magnetic field noise intermittently increased, and this correlated with large voltage excursions—up to 200 mV over a few seconds—on the voltage telemetry of the power supply, as shown in Appendix B. In principle, a constant-current driver should adjust its voltage to compensate for temperature- or vibration-induced resistance changes, or to counteract any stray magnetic field-induced current in the coil loop. Yet, the magnitude and sporadic nature of these excursions pointed to an intrinsic instability of the supply itself or intermittent contact resistance spikes in the wiring. Rather than undertake a detailed fault analysis, we adopted a simpler solution: replacing the electromagnet with a passive array of permanent magnets.

Because the required field is static and moderate, a permanent-magnet solution promised at least a factor of two reduction in Line noise [40] and zero technical noise. Key design goals were:

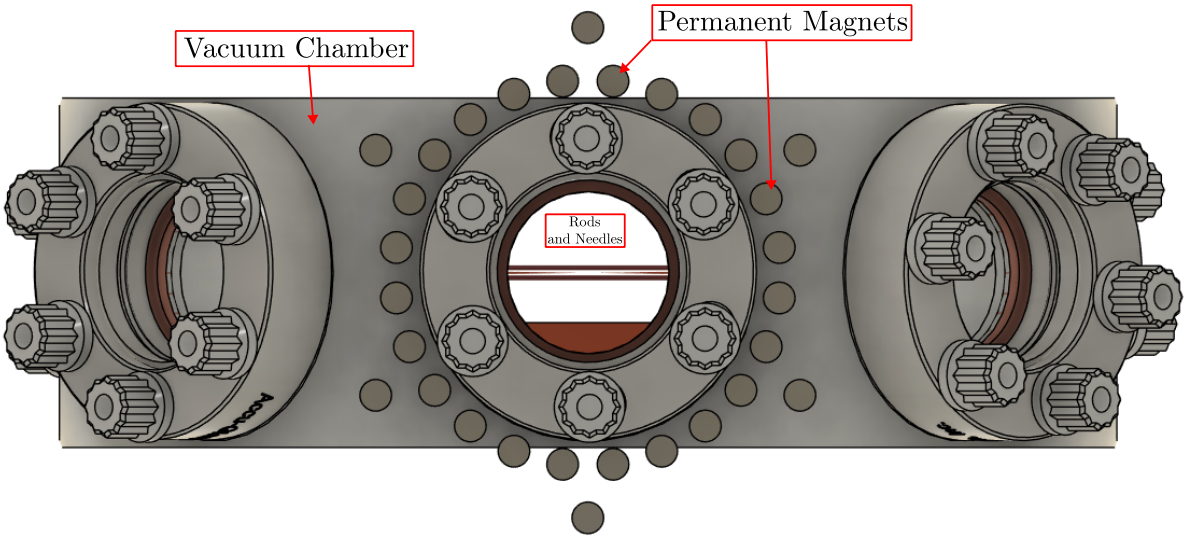


Figure 2.5: **Permanent Magnet Configuration.** A CAD model image of permanent magnets on the view-port of vacuum chamber. This configuration of permanent magnets is used to generate ∼4.2 at the trap center.

1. Deliver an on-axis field of $B_0 = 4.20 \pm 0.05$ G at the trap centre;
2. Preserve full clearance for all eight optical side-ports; and
3. Offer coarse tunability while installing as a drop-in replacement that requires no invasive modifications to the existing apparatus.

Using COMSOL Multiphysics, I built a model of the circular magnet array and stainless-steel vacuum chamber in order to optimize the free parameters like the number of layers, the number of magnets in each layer, the angular separation, the distance between the rings, and the dimensions and power of magnets used. The optimum configuration has two-layers of that hold small cylindrical N52-grade neodymium magnets (nickel-plated, ⌀= 1/8 in, thickness 1/16 in, flux density ≈ 4.5 kG). The inner layer's ⌀=1.03 in, and the outer layer's ⌀=1.5 in. The pair of magnet arrays is clamped to the opposing 4/3 in view-ports, giving axial spacing of ≈ 4.2 in (set by the ⌀4.2 in chamber body). In the final configuration, there are 24 magnets in the inner layer and 6 magnets in the outer layer, shown in figure 2.5, each with its north pole pointing towards the trap center for one array of magnets

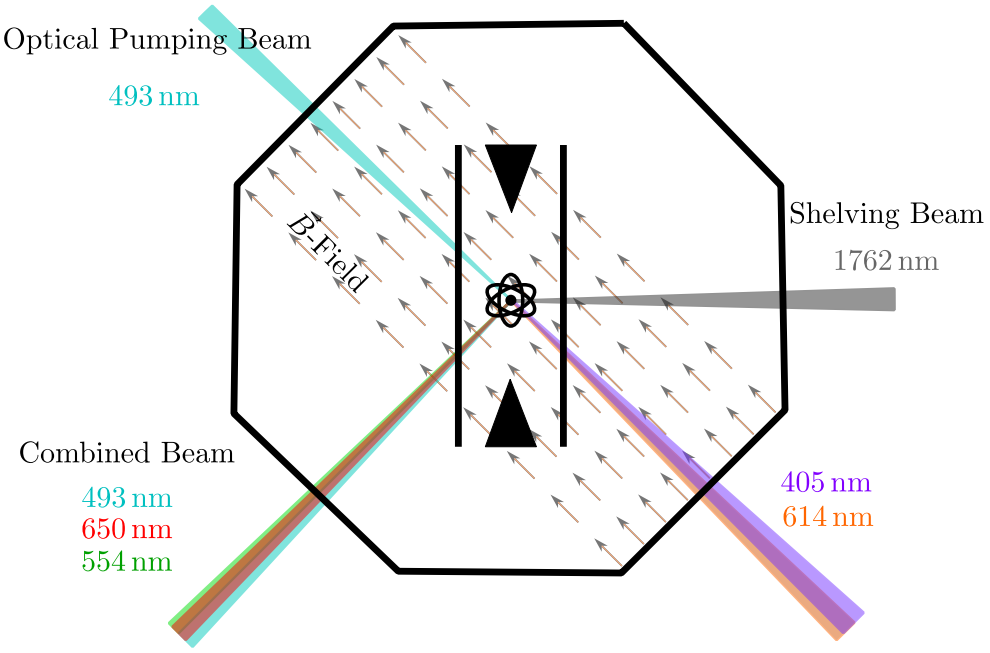


Figure 2.6: **Beam Paths and geometry** with respect to Magnetic field orientation

and (south pole pointing towards the trap center for the other array). The configuration yields a simulated on-axis field of 4.5 G while leaving all eight optical side-ports completely unobstructed (see Appendix C for details on simulations.).

A three-piece PLA collar was then printed with press-fit pockets for *up to* 24 inner-ring magnets and 32 outer-ring magnets, the extra holes allowing future re-configuration without re-printing the fixture. The N52 cylinders were inserted by hand after printing; no magnets were embedded during the print. With the $24+6$ configuration installed we measure an on-axis field of $B_0^{\text{exp}} = 4.2\,\text{G}$ at the ion—which is 1 G discrepancy between simulation and measured magnetic field.

The discrepancy can be traced to two practical factors, (i) the actual magnet–trap distance is slightly larger than the ideal 4.2 in because the collars cannot sit perfectly flush against the viewport flanges, B-field at the center vs. the separation between the rings is shown in Figure C.2 (b) in Appendix C. (ii) the manual press-fit leaves a few magnets tilted by a few degrees, reducing the net axial component. (iii) The quoted surface flux density(and hence the estimated remanent flux density) of the individual magnets might be higher in practice.

2.2.2 Beam geometry and Polarizations

Table 2.1: Beam-path summary relative to the quantization axis

Beam path	**Wavelength (nm)**	$\hat{k}$ **direction**†	**Lab pol.**	**Ion pol.**
Combined Beam	493, 650, 554		Linear	σ^+, π, σ^-
Shelving Repump	614, 405		Linear	σ^+, σ^-
Shelving Laser	1762		Elliptical	$\{\sigma^+, \pi, \sigma^-\}^*$
Optical pump	493		right circular	σ^+

†Black arrows represent the quantization axis
*Details about polarization components of 1762nm in Appendix A

Now that the permanent-magnet array fixes a stable quantization axis, we can turn to the other half of the picture: how each laser beam couples to the ion in that field. The polarisation that actually drives an atomic transition is defined with respect to this axis. To determine it, one projects the laser's electric field onto the spherical basis $\{\sigma^+, \pi, \sigma^-\}$ tied to the magnetic field. The result depends on the angle between the beam's propagation vector $\hat{k}$ and the ellipticity of the field:

- $\hat{k} \parallel \hat{B}$. Circularly polarized light appears as pure σ^+ or σ^-; linear light splits into equal σ^+ or σ^- components.
- $\hat{k} \perp \hat{B}$. Linear light polarized *along* the field is pure π, whereas linear light polarized *across* the field decomposes into equal $\sigma^+ + \sigma^-$.
- **Intermediate angles**. All three components are present, with weights that vary smoothly with the angle and ellipticity.

Thus, the "polarization the ion sees" is set jointly by lab optics and by the beam's direction relative to the quantization axis. With the magnetic geometry now fixed, I list below the propagation directions as seen in Figure 2.6, the power and waist and lab polarization of all operational wavelength, and translate each into the corresponding $\{\sigma^+, \pi, \sigma^-\}$ mix that drives the relevant transitions in our $^{137}Ba^+$ qudit.

2.3 Theory tools

This section documents the mathematical formalism and numerical workflow that allow us to extract practical quantities—most notably transition strengths and magnetic-field sensitivities—for any pair of hyperfine manifolds. Our immediate focus is the quadrupole-allowed $6S_{1/2} \leftrightarrow 5D_{5/2}$ system in Ba^+, but the expressions below hold for any (I, J) combination once the appropriate hyperfine constants are supplied (see Table 2.2).

Hyperfine interactions capture how the nuclear and electronic angular momenta couple within a given fine-structure manifold. This coupling splits each fine-structure level into hyperfine states and, together with the Zeeman interaction from an external magnetic field, determines both the energy spacings and the composition of the eigenstates. In the weak-field regime the total angular momentum F (with projection m_F) is a good quantum number. The Hamiltonian below collects these effects into standard terms so we can compute level energies and eigenvectors used throughout the experiment.

$$H_{\mathrm{hfs}} \;=\; H_{\mathrm{A_{MD}}} \;+\; H_{\mathrm{B_{EQ}}} \;+\; H_{\mathrm{C_{MO}}} \;+\; H_{\mathrm{Z}}, \tag{2.1}$$

The magnetic-dipole term, $H_{\mathrm{A_{MD}}}$ (2.2), is the leading electron-nuclear spin coupling and sets the primary hyperfine splitting pattern with the fine structure level. It describes the coupling between the nuclear magnetic-dipole moment and the magnetic field produced at the nucleus by the electronic cloud. The electric-quadrupole term, $H_{\mathrm{B_{EQ}}}$ (2.3), accounts for the interaction between the nuclear quadrupole moment and the electronic electric-field gradient; it vanishes for states with $J = 1/2$ but contributes in higher-J manifolds. The magnetic-octupole, $H_{\mathrm{C_{MO}}}$ (2.4), term is a higher-order correction from the nuclear magnetic octupole moment; it is typically small but can be included for completeness when extracting precise level energies. Finally, the Zeeman term, H_{Z} (2.5),describes coupling to an external bias field and provides the field-dependent shifts and state mixing.

$$H_{\mathrm{A_{MD}}} = A_{\mathrm{MD}}\,\mathbf{I}\cdot\mathbf{J}, \tag{2.2}$$

$$H_{\mathrm{B_{EQ}}} = \frac{B_{\mathrm{EQ}}}{2I(2I-1)J(2J-1)}\Big[3(\mathbf{I}\cdot\mathbf{J})^2 + \tfrac{3}{2}\,\mathbf{I}\cdot\mathbf{J} - I(I+1)J(J+1)\Big], \tag{2.3}$$

$$\begin{aligned} H_{\mathrm{C_{MO}}} = \frac{C_{\mathrm{MO}}}{I(I-1)(2I-1)J(J-1)(2J-1)}\Big[\,&10(\mathbf{I}\cdot\mathbf{J})^3 + 20(\mathbf{I}\cdot\mathbf{J})^2 \\ &+ 2\big(I(I+1) + J(J+1) + 3 - 3I(I+1)J(J+1)\big)\mathbf{I}\cdot\mathbf{J} \\ &- 5I(I+1)J(J+1)\Big], \end{aligned} \tag{2.4}$$

$$H_{\mathrm{Z}} = \mu_{\mathrm{B}} B_0\big(g_J J_z + g_I I_z\big). \tag{2.5}$$

Table 2.2: Experimental hyperfine constants for selected levels of $^{137}\mathrm{Ba}^{+}$.

Electronic level	A_{MD} (MHz)	B_{EQ} (MHz)	C_{MO} (Hz)	Reference
$6S_{1/2}$	4018.870(5)	-	-	[29]
$6P_{1/2}$	743.7(4)	-	-	[41]
$6P_{3/2}$	126.9(3)	92.5(2)	-	[41]
$5D_{3/2}$	189.731494(17)	44.537594(34)	29.533(86)	[42]
$5D_{5/2}$	-12.029234(11)	59.525520(110)	-12.41(77)	[42]

Here $\mathbf{I}$ and $\mathbf{J}$ are the nuclear- and electronic-angular-momentum operators, μ_{B} is the Bohr magneton, and g_I and g_J are the respective g-factors. For an l-S coupled valence electron, the Landé factor is given by

$$g_J \;=\; \frac{g_L\big[J(J+1) - S(S+1) + L(L+1)\big] + g_S\big[J(J+1) + S(S+1) - L(L+1)\big]}{2J(J+1)}. \tag{2.6}$$

In practice, we work in the uncoupled product basis $|m_I, m_J\rangle$ of dimension $(2I+1)(2J+1)$. Ladder operators

$$S_{\pm}|m\rangle \;=\; \hbar\sqrt{S(S+1) - m(m\pm1)}\,|m\pm1\rangle \tag{2.7}$$

generate the Cartesian components

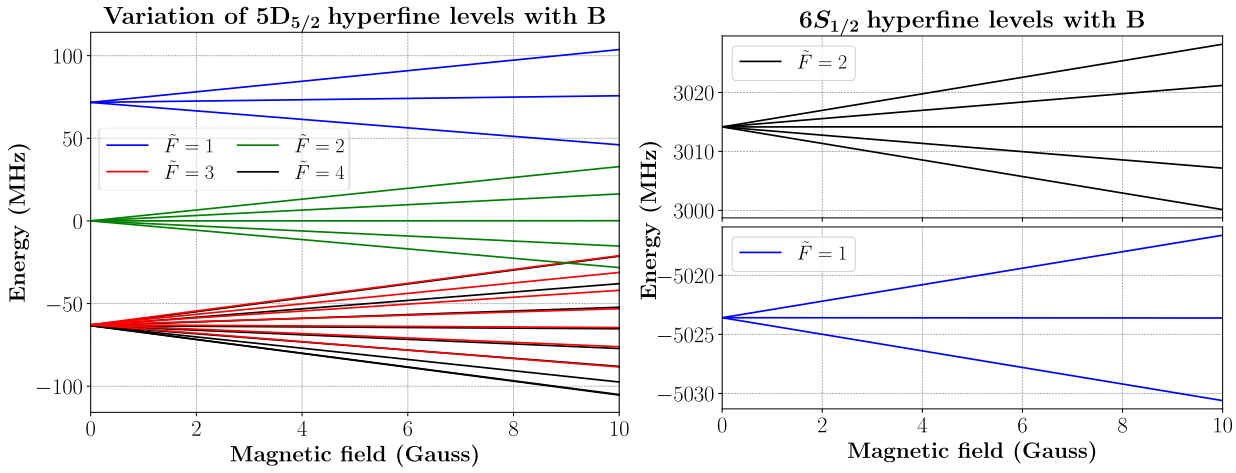


Figure 2.7: **Energy Level splitting vs. Magnetic field.** Variation in Energy levels of $5D_{5/2}$ and $6S_{1/2}$ over 10 G of magnetic field. The zero energy for $5D_{5/2}$ figure is set at magnetic field insensitive state $|5D_{5/2}, \tilde{F}=2, m_{\tilde{F}}=0\rangle$. And the zero energy of $6S_{1/2}$ is at $^{138}Ba^{+}$ energy level (it has no hyperfine splittings).

$$S_x = \tfrac{1}{2}(S_+ + S_-), \quad S_y = \tfrac{1}{2\mathrm{i}}(S_+ - S_-), \quad S_z|m\rangle = \hbar m|m\rangle, \tag{2.8}$$

and replacing $S \to I$ or J yields **I** and **J**. The scalar product that appears throughout Eqs. (2.2),(2.3),(2.4),(2.5) is then

$$\mathbf{I}\cdot\mathbf{J} \;=\; \sum_{k=x,y,z} I_k \otimes J_k. \tag{2.9}$$

After solving this Hamiltonian numerically, we get eight eigenvalues and eigenstates for $6S_{1/2}$ manifold and 24 eigenvalues and eigenstates for $5D_{5/2}$ manifold. The eigenstates will be in the $|I, m_I; J, m_J\rangle$ basis.

$$H_{\mathrm{hfs}}|V_n\rangle = E_n \sum_{m_{I,L}, m_{J,L}} c_{m_{I,L}, m_{J,L}} |I=3/2, m_I; J, m_J\rangle \tag{2.10}$$

where $|V_n\rangle$ is eigenstate in $|I, m_J; J, m_J\rangle$ basis, $c_{m_{I,L}, m_{J,L}}$ are complex coefficients corresponding to respective $|I=3/2, m_I; J, m_J\rangle$ state. And E_n is the eigenvalue or energy shift from the corresponding fine-structure level.

In $^{137}Ba^{+}$ the $5D_{5/2}$ manifold's lowest two hyperfine levels, F=3 and F=4, are split by only ~ 0.5 MHz at $B = 0$. Because this gap is so narrow, even sub-gauss magnetic fields mix pairs of states appreciably, turning the nominal $|F, m_F\rangle$ basis into field-dependent superpositions. Thus, the hyperfine labels are merely convenient references; the actual energy eigenstates we work with are these mixed combinations (denoted by $|\tilde{F}, m_{\tilde{F}}\rangle$). Our numerical routine follows each eigenvector adiabatically in B, permuting it whenever energy levels swap the order, so every energy, transition strength, and sensitivity reported later is tied to a single, continuously tracked physical state.

Algorithm 1 CalculateEnergy – Adiabatic computation of the $^2D_{5/2}$ spectrum

Require: Target field $B_{\max}$; quantum numbers $J = 5/2,\ I = 3/2$

$B_{\text{prev}} \leftarrow 10^{-5}\,\text{G}$ ▷ seed field

$(V_{\text{prev}}, _) \leftarrow \text{Eig}\big[H_{D_{5/2}}(B_{\text{prev}})\big]$

$P \leftarrow I_{24}, \quad P_{Fm_F} \leftarrow I_{24}$ ▷ running permutations

for B from 10^{-6} to $B_{\max}$ in steps $\Delta B = 10^{-4}\,B_{\max}$ **do**

  $(V, \mathcal{E}) \leftarrow \text{Eig}\big[H_{D_{5/2}}(B)\big]$ ▷ diagonalise

  $O \leftarrow \big|V^{\dagger} V_{\text{prev}}\big| \qquad \tilde{O} \leftarrow \text{round}(O)$ ▷ overlap matrix → permutation

  $P \leftarrow \tilde{O}\,P; \quad V \leftarrow V\,P^{\mathsf{T}}; \quad \mathbf{E} \leftarrow P\,\text{diag}(\mathcal{E})$

  if first loop **then** ▷ anchor to F, m_F basis once

    $T \leftarrow$ CG-matrix $\big(|I, m_I; J, m_J\rangle \to |F, m_F\rangle\big)$

    $C \leftarrow \text{diag}\big[\text{round}\big(\mathbf{1}^{\mathsf{T}} T^{\dagger} V\big)\big]$

    $V \leftarrow V\,C; \quad P \leftarrow C\,P$

  end if

  $V_{FmF} \leftarrow T^{\dagger} V$

  $Q \leftarrow \text{round}\big(V_{FmF}^{\dagger} P_{Fm_F}\big)$ ▷ enforce consistent F, m_F labels

  $V \leftarrow V\,Q; \quad V_{FmF} \leftarrow T^{\dagger} V; \quad P_{Fm_F} \leftarrow Q\,P_{Fm_F}$

  set elements of V and V_{FmF} with $|x| < 10^{-6}$ to 0

  $V_{\text{prev}} \leftarrow V$ ▷ store for next field step

end forreturn $\mathbf{E}$ (ordered energies), V (lab-labelled eigenvectors), V_{FmF} (vectors in F, m_F basis), P (final permutation)

Using this method, we can calculate how the energy levels split as the B-field increases, as shown in Figure 2.7. The $F = 3$ and $F = 4$ levels start to cross at $\sim 0.07\,\text{G}$, resulting in eigenstates that are mixes of the pure $|F, m_F\rangle$ states.

The plots in Figure 2.7 can be combined for a given B-field to generate a magnetic field sensitivity chart of transitions between the $6S_{1/2}$ and $5D_{5/2}$ manifolds shown in Figure 2.8.

Furthermore more relative transition strengths can also be calculated using the Clebsch-Gordan (CG) coefficients and polarization-dependent geometric factor as

$$s_{\text{rel}} = g^{(q)} \langle \tilde{F}_S, m_{\tilde{F},S}; k = 2, q = m_{\tilde{F},D} - m_{\tilde{F},S} | \tilde{F}_D, m_{\tilde{F},D} \rangle. \tag{2.11}$$

Since $|\tilde{F}, m_{\tilde{F}}\rangle$ are not the pure hyperfine energy levels, the CG coefficients part of (2.11) can be calculated as:

$$\begin{aligned} &\langle \tilde{F}_S, m_{\tilde{F},S}; k = 2, q = m_{\tilde{F},D} - m_{\tilde{F},S} | \tilde{F}_D, m_{\tilde{F},D} \rangle = \\ &\sum_{m_{I,D}, m_{J,D}} \sum_{m_{I,S}, m_{J,S}} c^*_{m_{I,D}, m_{J,D}} \, c_{m_{I,S}, m_{J,S}} \delta_{m_{I,S}, m_{I,D}} \\ &\times \langle J_S = 1/2, m_{J,S}; k = 2, q = m_{J,D} - m_{J,S} | J_D = 5/2, m_{J,D} \rangle \end{aligned} \tag{2.12}$$

Where $c_{m_{I,L}, m_{J,L}}$ can be numerically calculated as in Eq (2.10). $g^{(q)}$ is the geometric factor corresponding to the polarization and geometry of the laser with respect to the quantization axis. Ref. [35] Section 4.1.2 provides the derivation of these geometric factors for a linearly polarized laser. Appendix A provides a derivation of the geometric factor for a general elliptical polarization. Using (2.11) and (2.12), we obtain a relative transition strengths chart for transitions between $6S_{1/2}$ and $5D_{5/2}$ manifolds. A subset of this chart is shown in the Figure 2.9 with quadrupole laser polarization assumed as linear.

These magnetic field sensitivity values and relative transition strength values are of practical use for calibrating transition frequencies and Rabi frequencies, as discussed in Chapter 4.

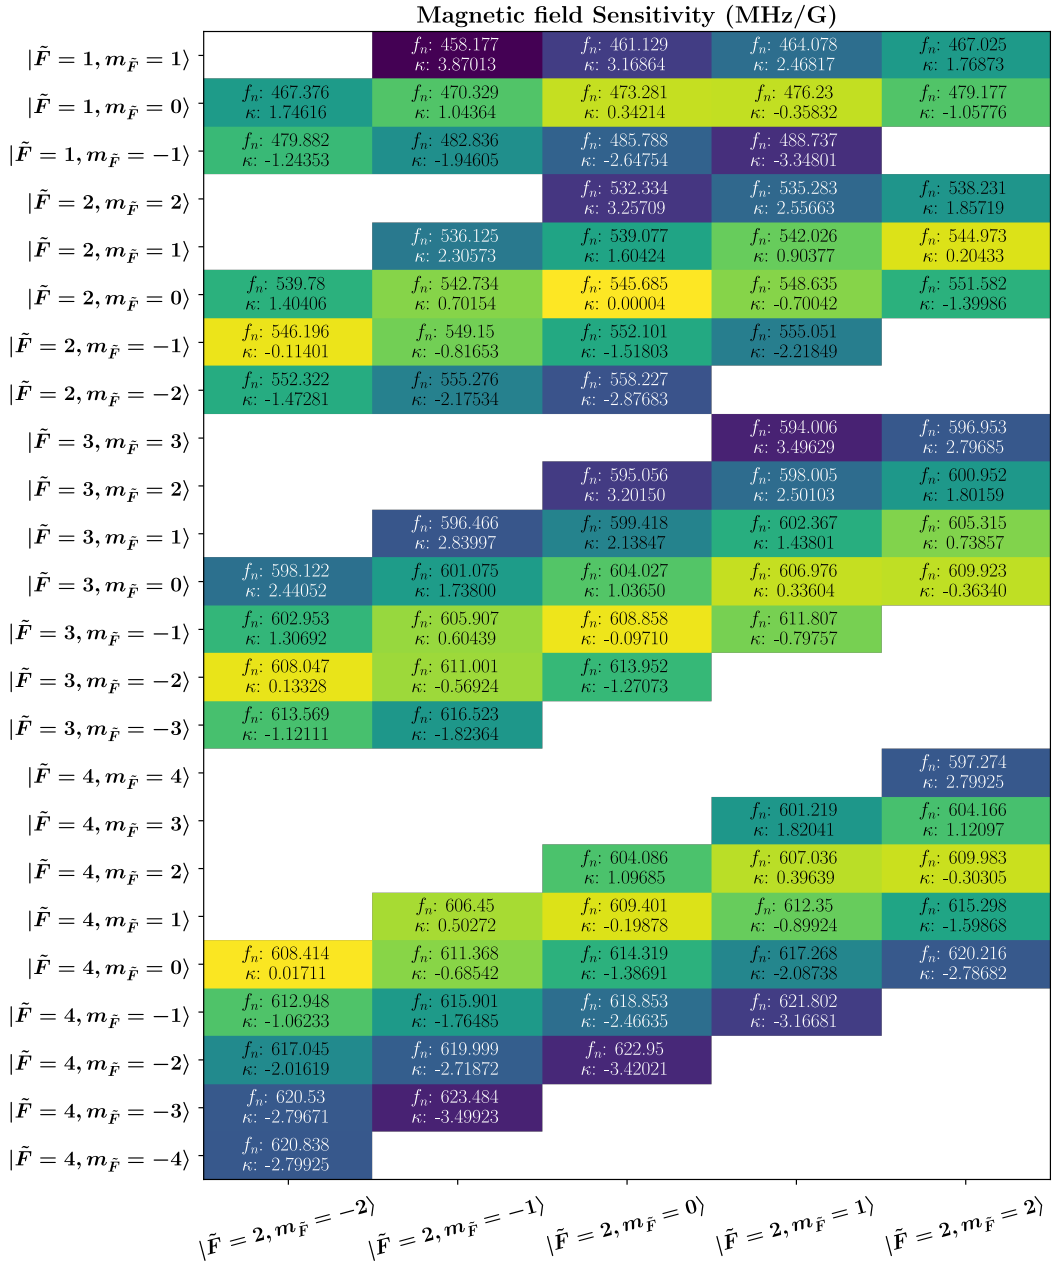


Figure 2.8: **Magnetic field Sensitivity Chart.** Color map of magnetic field sensitivities κ (in MHz/G) for various transitions between the $6S_{1/2}$ $F = 2$ and $5D_{5/2}$ states, calculated at 4.209 G of B-field. Each cell is annotated with the corresponding sensitivity value κ as well as the associated AWG modulation frequency f_n (in MHz) used to drive the transition. The vertical axis denotes the $5D_{5/2}$ states, while the horizontal axis represents the $6S_{1/2}$ states involved in each transition.

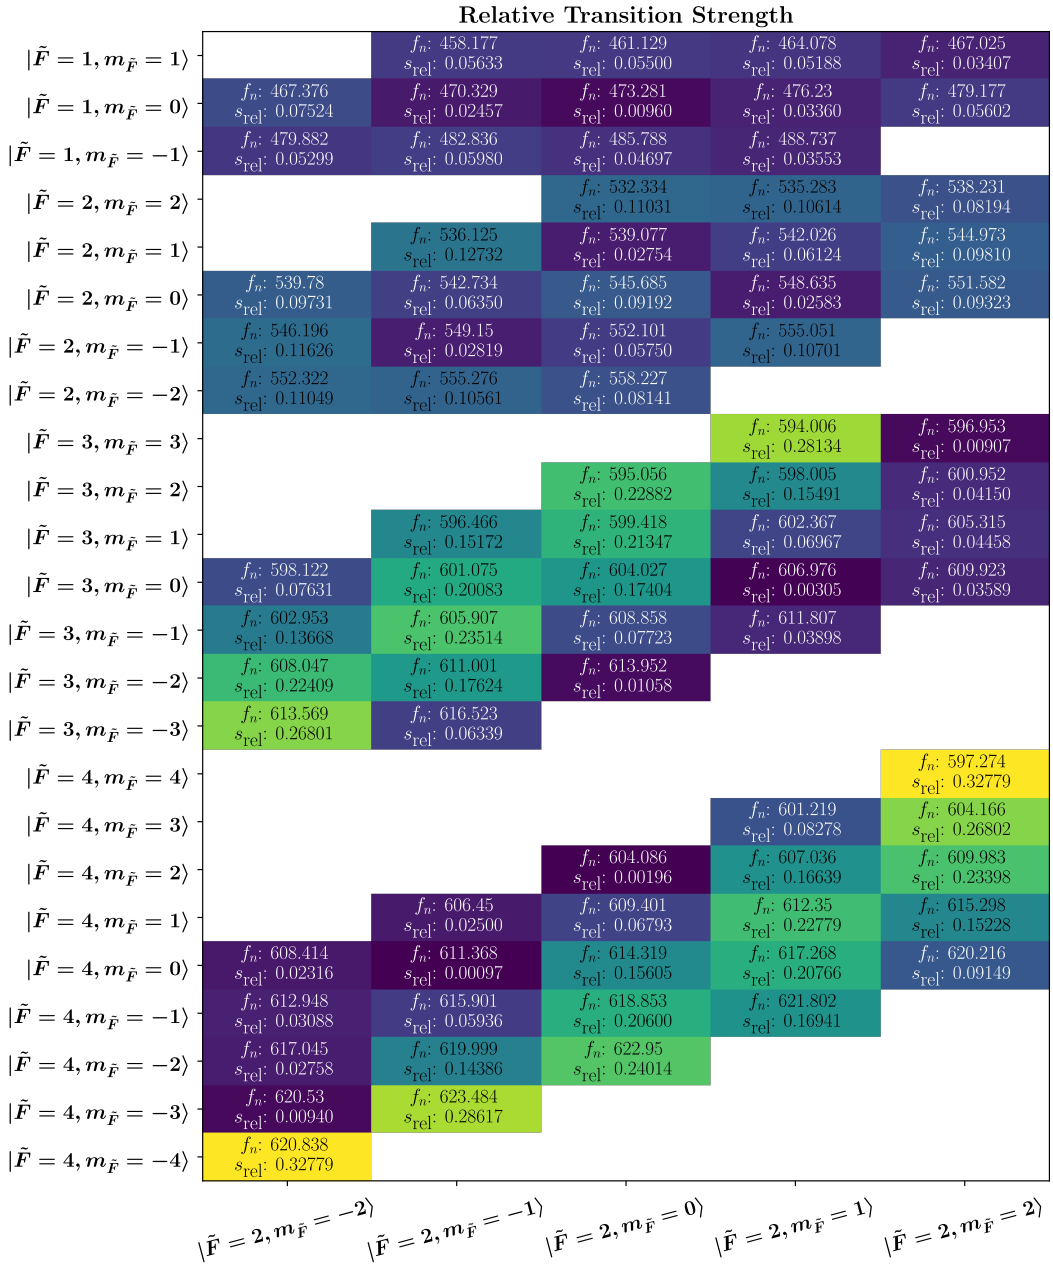


Figure 2.9: **Transition Strength Chart.** Colormap of relative transitions strengths (s_{rel}) between $6S_{1/2}$ $F = 2$ and $5D_{5/2}$ states. The laser polarization is assumed to be linear with a polarization angle of 58°.

Chapter 3

State Preparation and Measurement

A critical aspect of any quantum information experiment is the ability to prepare quantum states and measure them reliably with high fidelity. This chapter presents the techniques and results related to initializing and measuring multi-level qudit states in our trapped $^{137}Ba^{+}$ system. Since the $6S_{1/2}$ manifold contains the ground states of the ion, all experiments must begin with the ion initialized in one of these levels. We begin by discussing the method of state preparation in the $6S_{1/2}$ manifold, including practical considerations that influence fidelity, such as branching ratios, laser polarization, optical pumping duration and off-resonant scattering. We then extend our focus to the $5D_{5/2}$ manifold, where initialization is more straightforward. Having first prepared the ion in a well-defined $6S_{1/2}$ state, we can selectively drive "good" transitions into specific $5D_{5/2}$ levels to prepare arbitrary qudit states in that manifold. Following the discussion on preparation, we detail our measurement protocols, including fluorescence detection and how transitions are chosen for state readout. We then present our SPAM (State Preparation and Measurement) results, quantifying our system's performance across all 25 qudit levels. Finally, we describe heralding techniques employed to boost fidelity and ensure experimental confidence in the presence of imperfections or some initialization and readout errors.

3.1 $6S_{1/2}$ Initialization

This section introduces an NBOP routine that prepares a chosen $6S_{1/2}, F{=}2$ Zeeman state and shelves it to a specified $5D_{5/2}$ sublevel with a single calibrated quadrupole π-pulse. I outline the pump–repump–depletion loop, how the transfer line is selected based on Practical considerations in each step of NBOP sequence.

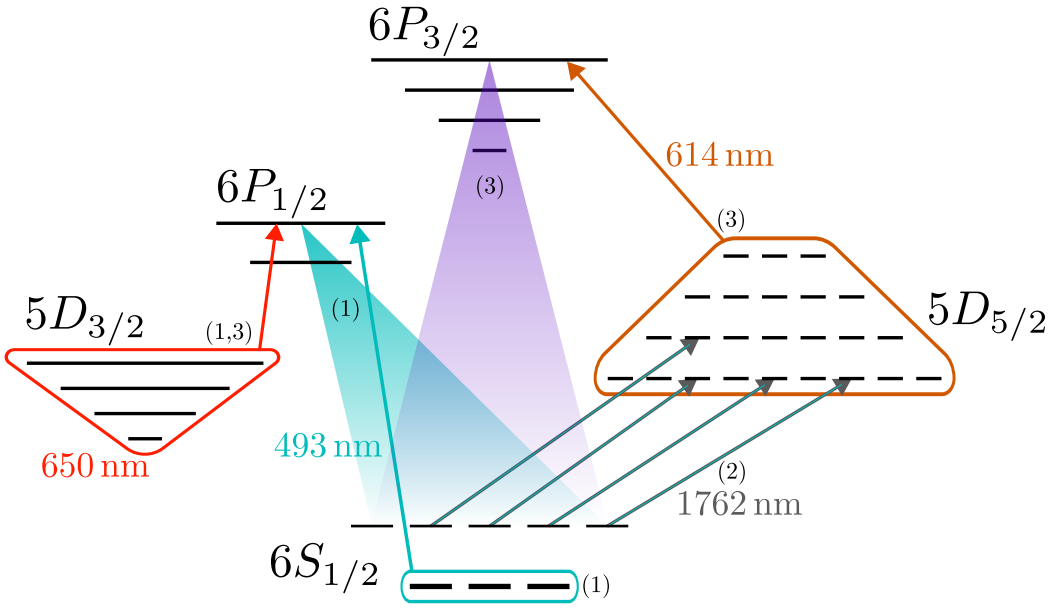


Figure 3.1: **$6S_{1/2}$ Initialization scheme.** In the example shown here $|6S_{1/2}, F = 2, m_F = -2\rangle$ is being initialized. 650 nm laser is used in both steps (1) and (3), since $6P_{1/2}$ and $6P_{3/2}$ have a branching ratio to $5D_{3/2}$.

Polarization-based optical pumping with circularly (or near-circularly) polarized light is the canonical route to preparing a trapped ion in a single Zeeman sub-level. The technique exploits the fact that repeated absorption–spontaneous–emission cycles drive the population toward a "dark" state that is forbidden by the polarization selection rules. High-purity polarization can routinely achieve sub-percent preparation error in qubits [28, 43].

For multi-level qudits, however, polarization-based optical pumping (with circular polarization) alone is insufficient: from the $m_F = +2$ state in $6S_{1/2}$, $F = 2$ manifold, we can access only 12 of 24 sub-levels in $5D_{5/2}$directly. This ceiling was evident in earlier work from our group in Ref. [28], where they demonstrated control and single-shot readout of a 13-level qudit, but deemed it impractical to reach the full 25-level computational space because the remaining $5D_{5/2}$ states were inaccessible from the standard optical-pumping start point. To unlock the entire set of $5D_{5/2}$ levels we adopt, and extend, the Narrow-Band Optical Pumping (NBOP) protocol introduced by An *et al.* in Ref. [6]. This technique combines broadband optical pumping with targeted, narrow-band coherent transitions to achieve high-fidelity initialization in a specific Zeeman sublevel. In each cycle, the undesired $6S_{1/2}$ state population is coherently transferred to $5D_{5/2}$ via a quadrupole

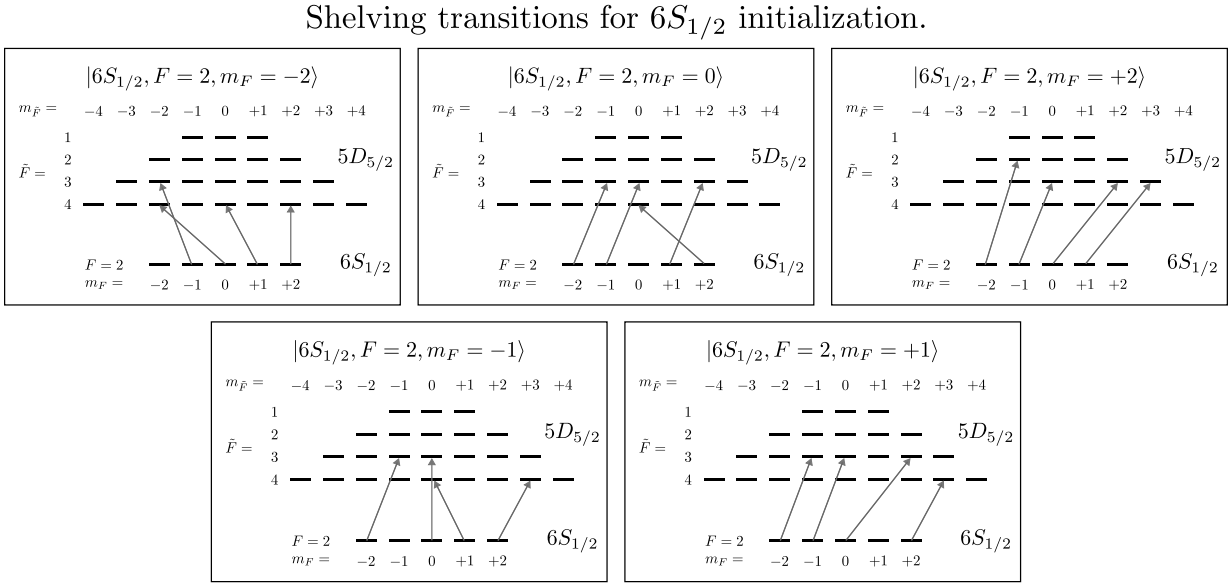


Figure 3.2: **Shelving transitions for $6S_{1/2}$ initializations.** 1762 shelving transitions used in step (2) of Narrowband Optical Pumping (NBOP) scheme. The transitions were chosen keeping the considerations described in

transition, then repumped using broadband dipole light. Repeating this process gradually depopulates all states except the desired one. This method circumvents the limitations of polarization-based pumping and achieves exceptional fidelity even in complex multi-level systems such as $^{137}Ba^{+}$. An *et al.* implemented NBOP by reaping the shelving population out of the two leakage states $|6S_{1/2}, F=2, m_F=\pm1\rangle$ with narrow 1762 nm $\Delta m_F = \pm2$ pulses and repumping, thereby funneling everything into the field-insensitive clock state $|6S_{1/2}, F=1, m_F=0\rangle$; this single-state variant reached a record SPAM infidelity of 9×10^{-5} [6]. We extend the scheme by retuning the narrow-band steps so that *any* of the five $|6S_{1/2}, F=2, m_F\in\{-2,-1,0,+1,+2\}\rangle$ sub-levels can be chosen to be initialized, giving easy access to all 24 magnetic sub-levels of $5D_{5/2}$ manifold.

In the following, I outline the specific steps of our adapted initialization sequence (also shown in figure 3.1), highlighting key considerations at each stage, such as frequency selectivity and timing constraints, that are crucial for achieving state preparation across all target states.

(1) Selective $F=1$ depopulation—Apply 493 nm and 650 nm light with frequency

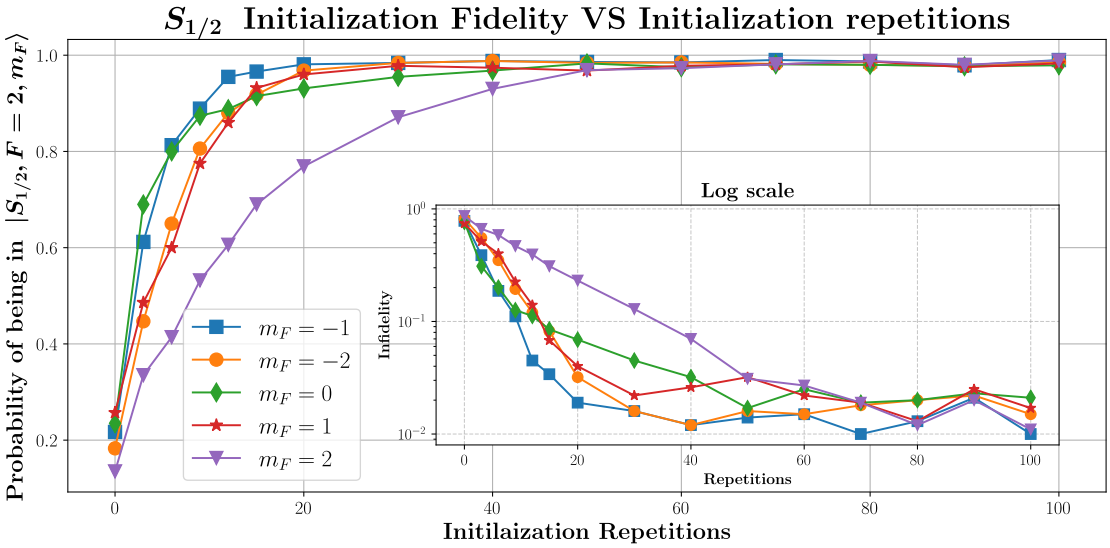


Figure 3.3: **Initialization probability vs. Repetitions.** Plots here show the fidelity and infidelity of initializing in a target $6S_{1/2}$ $F = 2$ level as the initialization repetitions increase.

component tuned specifically to pump out of population from $6S_{1/2}$, $F = 1$ manifold. This ensures that the ion is confined to the $F = 2$ manifold, where state selection will occur.

(2) Coherent depletion of unwanted $F = 2$ states—Apply 1762 nm light with narrow-band frequency component targeting specific transitions from $6S_{1/2}$, $F = 2$ levels—excluding the desired target state. Each component is tuned for a calibrated π-pulse duration to coherently transfer population from each unwanted sublevel to $5D_{5/2}$. The transitions used in this step are shown in figure 3.2.

(3) Repumping via 614 nm light—Apply 614 nm light to depopulate $5D_{5/2}$ by preferentially pumping it to $6P_{3/2}$, $F = 3$, allowing any shelved population to decay back to $6S_{1/2}$, typically redistributing among the available $F = 2$ sublevels.

(4) Iterative refinement—Repeat steps 1-3 in a loop. Over successive cycles, population accumulates in the chosen state —i.e., the only sublevel not targeted by step 2 (1762 nm depletion step)—resulting in preparation of the desired $6S_{1/2}$, $F = 2$ state.

Figure 3.3 shows how the population in the desired $6S_{1/2}$, $F = 2$ state increases with the number of initialization iterations for each of the five sublevels. The population reliably

saturates after approximately 60-80 repetitions, indicating the convergence to the target state. yielding a preparation fidelity of 0.986(8) averaged over the Five target sub-levels. Additionally, the figure includes a state preparation and measurement (SPAM) chart for each $6S_{1/2}$, $F = 2$ sublevel. For these experiments, we initialize the target state and then apply up to six 1762 nm shelving pulses to transfer population to the distinct $5D_{5/2}$ states, mitigating errors due to imperfect shelving. We then detect fluorescence; the absence of photons (a "dark" ion) indicates successful state preparation and shelving, confirming that the population was initialized in the desired sublevel.

It is important to note that three main factors influence both the rate at which the population saturates and the final saturation value in our initialization scheme. First, the repump path in step 3—i.e., the decay pathway from $5D_{5/2} \rightarrow 6P_{3/2} \rightarrow 6S_{1/2}$—plays a critical role. Due to the branching ratios involved, certain $5D_{5/2}$ states are more likely to decay into specific $6S_{1/2}$ sublevel, and we can exploit these biases to our advantage. Second, the spectral selectivity of the transitions used in step 2 must be carefully managed. To avoid off-resonant excitation of the targeted state, the chosen 1762 nm transitions must have sufficient frequency separation from the transitions connecting to the targeted state. Third, power and duration of 493 nm (and 650 nm) light in step 1 must be optimized: it should be strong enough to clear out the population from the $F = 1$ manifold without unintentionally exciting population out of the $F = 2$ manifold. These considerations directly inform the choice of transitions and pulse parameters in our initialization sequence. The following sections discuss each of these points in detail.

3.1.1 614 nm Repump path

Using the numerical calculations in Chapter 2 Section 2.3, we map out the repump pathways for population shelved in any of the 24 Zeeman sub-levels of $5D_{5/2}$. Because the 614 nm repump beam contains equally weighted σ^+ and σ^- polarization components (see Table 2.1), these polarizations are taken into account when we compute the branching ratios for the cascade $6D_{5/2} \rightarrow 6P_{3/2} \rightarrow 6S_{1/2}$. This allows us to identify the most advantageous $5D_{5/2}$ states to use in step (2) of NBOP sequence so they preferentially decay to the desired target state.

The resulting repump-probability chart (shown in Figure 3.4) is renormalized so that it only shows decay channels terminating in the $6S_{1/2}, F = 2$ manifold. These probabilities are one of the main factors guiding our choice of shelving states in $6S_{1/2}$: when we wish to initialize a particular m_F sub-level, we deliberately choose shelving states whose repump pathways preferentially funnel population into (or towards) that target m_F state.

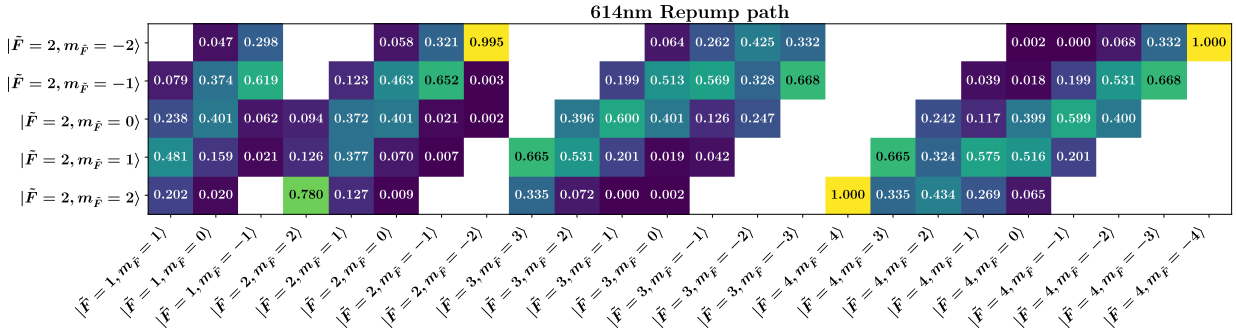


Figure 3.4: **Repump path Map.** Using numerical calculation described in 2 Section 2.3, branching ratios from $6D_{5/2} \rightarrow 6P_{3/2} \rightarrow 6S_{1/2}$ were calculated and normalized for $F = 2$ manifold in $6S_{1/2}$.

3.1.2 Spectral crowding

A second constraint on high-fidelity $6S_{1/2}$ initialization is the spectral isolation of the 1762 nm shelving transitions, particularly from those linking the target $6S_{1/2}\, m_F$ sub-level to the $6D_{5/2}$ manifold. In practice, the spectrum is crowded: with our polarization geometry the quadrupole selection rule ($\Delta m = \pm 2$) connects the five $6S_{1/2}(F = 2)$ Zeeman sub-levels to subsets of the twenty-four $6D_{5/2}$ states, yielding ~ 80 distinct carrier lines. Each of these carriers is further flanked by red and blue motional sidebands at the radial secular frequencies ω_x and ω_y, i.e. at $\omega_0 \pm \omega_x$ and $\omega_0 \pm \omega_y$. These sidebands arise because a trapped ion experiences a harmonic pseudopotential from the RF Paul trap; near the equilibrium point the motion separates into orthogonal normal modes along $\hat{x}$ and $\hat{y}$ with secular frequencies ω_x and ω_y, and the laser couples internal and motional states with strengths set by the Lamb–Dicke parameters $\eta_{x,y} = (\mathbf{k}\cdot\hat{x}, \mathbf{k}\cdot\hat{y})\sqrt{\hbar/(2m\omega_{x,y})}$. With a 1762,nm beam at 45° to both $\hat{x}$ and $\hat{y}$ axes, the projection of $\mathbf{k}$ onto either axis is $k_0/\sqrt{2}$. Using $^{137}\mathrm{Ba}^+$ mass $m = 2.28 \times 10^{-25}$ kg:

$$\eta_x = \frac{k_0}{\sqrt{2}}\sqrt{\frac{\hbar}{2m\omega_x}}, \qquad \eta_y = \frac{k_0}{\sqrt{2}}\sqrt{\frac{\hbar}{2m\omega_y}}. \tag{3.1}$$

Plugging $\omega_x = 2\pi \times 1.46$ MHz and $\omega_y = 2\pi \times 1.27$ MHz:

- $\eta_x \approx 0.0127$
- $\eta_y \approx 0.0136$

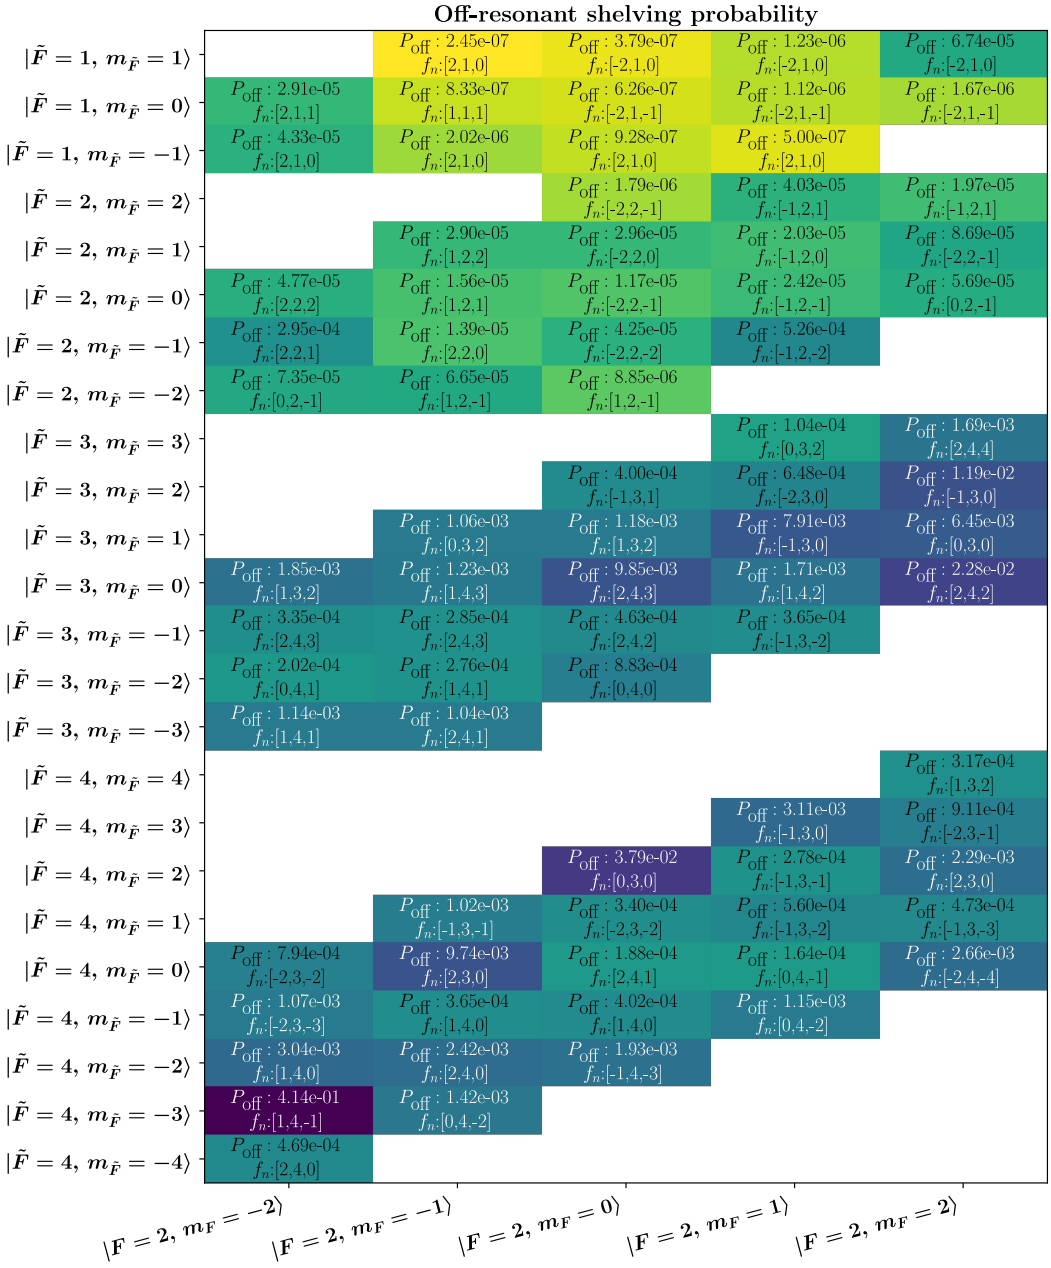


Figure 3.5: **Off-resonant Scattering probability.** A colormap of off resonant scattering probability and most probable transition (denoted as triplet number $[m_{F,S}, \tilde{F}_D, m_{F,D}]$) to be scattered.

The mean phonon number is $\bar{n} \approx 140$ the Lamb–Dicke weighting for the first-order secular side-bands jumps by a factor $\approx \sqrt{140} \approx 12$:

$$\Omega_{\text{blue}} = \eta\sqrt{\bar{n}+1}\,\Omega_0 \;\approx\; 12\,\eta\,\Omega_0, \qquad \Omega_{\text{red}} = \eta\sqrt{\bar{n}}\,\Omega_0 \;\approx\; 12\,\eta\,\Omega_0. \tag{3.2}$$

Now if we drive a transition with frequency f_0 for its corresponding π-time, the probability of off-resonantly driving a transition of frequency f_n and Rabi-frequency Ω_n is given by

$$P_n(t_\pi) \;=\; \frac{\Omega_n^2}{\Omega_n^2+\Delta_n^2}\,\sin^2\!\Big(\tfrac{1}{2}\sqrt{\Omega_n^2+\Delta_n^2}\;t_\pi\Big)\,. \tag{3.3}$$

where $\Delta_n = f_n - f_0$ Total off-resonant shelving probability for the driven line can be found by summing over all the carriers and their $\pm 1.46\,\text{MHz}$ / $\pm 1.27\,\text{MHz}$ side-bands

$$P_{\text{off}} \;=\; \sum_{n\neq 0} P_n(t_\pi), \tag{3.4}$$

Figure 3.5 plots, for every candidate shelving line, the total off-resonant scattering probability obtained by driving a π-pulse, as well as shows the biggest contributor to the off-resonant scattering probability. In step (2) of initialization scheme, the transitions that appreciably couple the target $6S_{1/2}$ state to any of the 24 Zeeman sub-levels of $6D_{5/2}$ are undesirable, because inadvertent population excitation imposes a hard ceiling on the achievable shelving-pulse saturation (and hence on the final state-preparation fidelity). We therefore select, as our primary shelving channels, those transitions that minimize this off-resonant drive while still offering acceptable Rabi rates at the available optical power.

3.1.3 F=1 Pump parameters

We have found that the likelihood of successfully preparing the target state is highly influenced by parameters of the $F = 1$ pumping step (1) that affect off-resonant scatter rates, more specifically, the optical power, modulation configuration, and pulse duration.

In our original modulation configuration, the 493 nm carrier was parked 2.3 GHz above the $F = 1$ manifold, while 2.3 GHz and 5.7 GHz sidebands addressed the $F = 1$ and $F = 2$ manifolds, respectively. The red 2.3 GHz sideband efficiently depopulates the $F = 1$ manifold, but the blue 2.3 GHz sideband sits only ≈ 3.4 GHz from the $F = 2$ manifold, and the carrier itself is 5.7 GHz detuned from $F = 2$. At high laser power, these proximities lead to significant off-resonant excitation of the $F = 2$ manifold during initialization, counteracting the intended shelving and preferential repump sequence. Figure 3.6 illustrates the

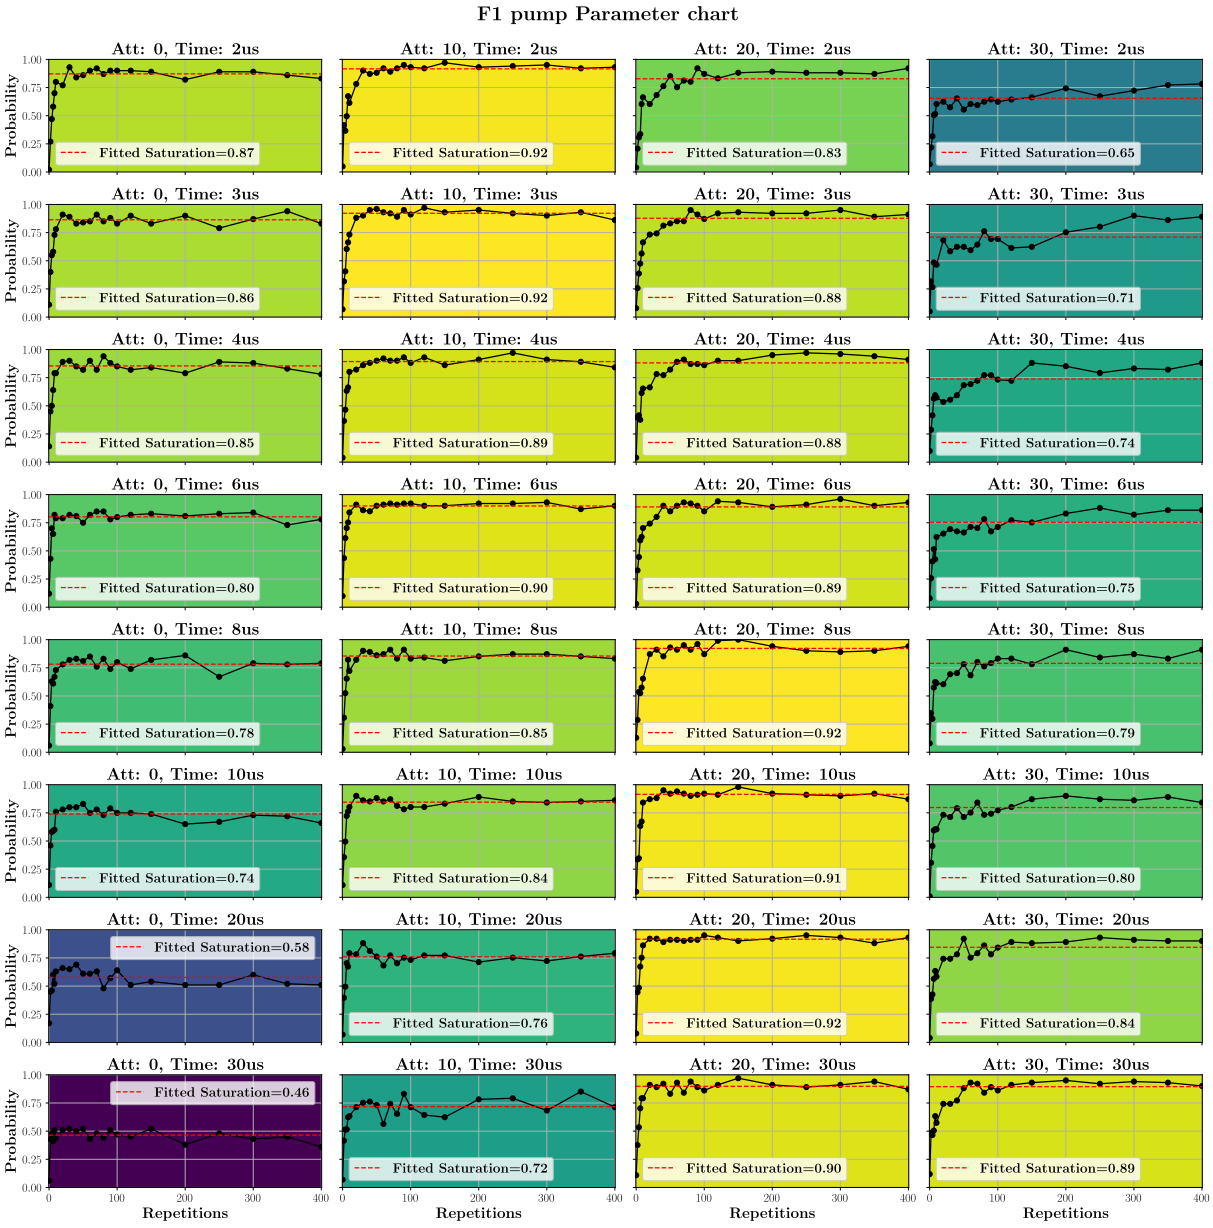


Figure 3.6: ***F* = 1 pump Parameter Chart.** A colormap showing saturation values as we sweep over the 2.3 GHz sideband's attenuation as well as pumping time in step (1) of the initialization scheme.

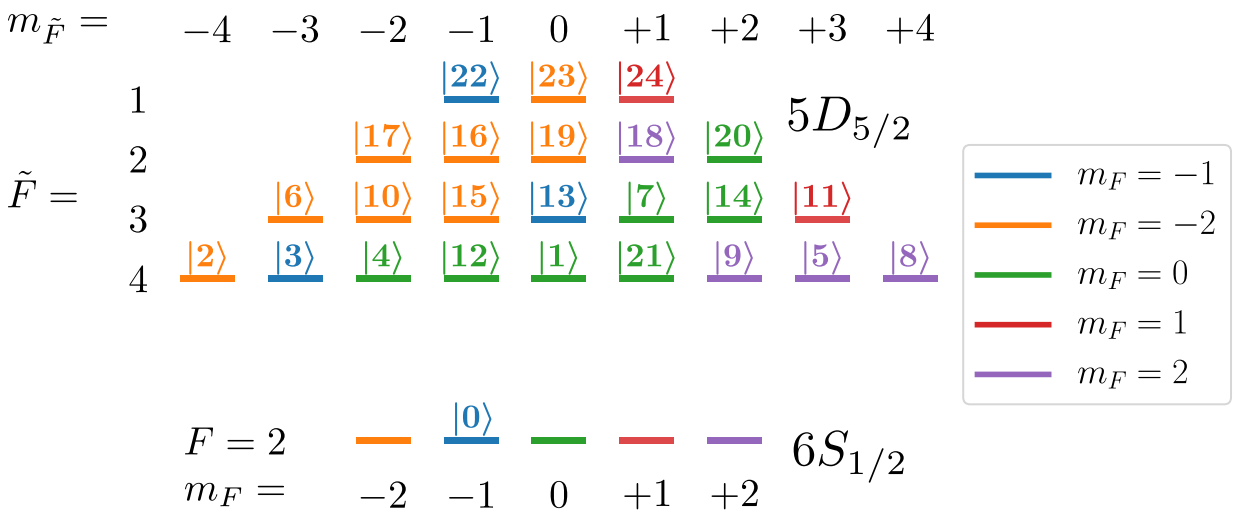


Figure 3.7: **State encoding and shelving transitions for State preparation and Measurement.** The $5D_{5/2}$ levels are colored based on the $6S_{1/2}$ level they originate from.

effect: when we vary the $F = 1$ pump duration and increase the attenuation of the 2.3 GHz sideband, the observed saturation value shifts accordingly.

Building on these results, we have re-optimized the 493 nm beam configuration. The carrier is now tuned directly to the $6S_{1/2}$, $F = 1 \rightarrow 5P_{1/2}$, $F = 2$ resonance, and a single ≈ 8.1 GHz electro-optic sideband addresses the $6S_{1/2}$, $F = 2 \rightarrow 5P_{1/2}$, $F = 2$ transition (as described in Section 2.1.1). By eliminating the former blue sideband that lay only $\approx$ 3.4 GHz from the $F = 2$ manifold and moving the carrier further away, we greatly suppress unintended excitation of $F = 2$ during the $F = 1$ pumping step. Residual off-resonant scattering is still possible, so we continue to tune the pump parameters carefully. A 0.5 µs pulse at full laser power (0 dB attenuation) reliably empties the $F = 1$ manifold without undoing the shelving and preferential repump sequence, and we now use this setting as our standard operating point for high-fidelity initialization.

3.2 $D_{5/2}$ Initialization

To prepare a chosen qudit level within the 24-dimensional $5D_{5/2}$ manifold, we first identify the strongest electric-quadrupole transition linking that $5D_{5/2}$ state to one of the five $6S_{1/2}$ hyperfine levels. Figure 4.2 lists every allowed transition (by quadrupole selection rules)

together with their π-times. Our strategy is to select, for each $5D_{5/2}$ row, the transition with the shortest π-time; this maximizes transfer efficiency and minimizes decoherence and off-resonant driving due to spectral crowding. Figure 3.7 shows our shelving transitions chosen for 24 levels in $5D_{5/2}$. After robustly initializing the ion in the corresponding $6S_{1/2}$ level, extending the procedure to any of the 24 $5D_{5/2}$ levels becomes straightforward: a single calibrated π-pulse promotes the population into the desired $5D_{5/2}$ level. The detailed calibration of the transition frequencies and precise π-times used for each transition is presented in Chapter 4.

3.3 State-Selective Readout

Our measurement scheme needs to distinguish between all 25 levels in our ions. We accomplish this by reducing the problem to a sequence of binary (bright/dark) fluorescence checks, interleaved with calibrated π-pulses that successively "deshelve" individual dark states ($5D_{5/2}$) into the bright manifold ($6S_{1/2}$); the first bright event fixes the state, with a dynamic threshold set from the photon-count histogram. The protocol proceeds as follows:

1. **Detect $|0\rangle$.** Switch on the Doppler–cooling beams at 493 nm and 650 nm for a detection window of duration $t_{\text{det}} = 5\,\text{ms}$.
   - *Bright:* the ion was in $|0\rangle$.
   - *Dark:* population is shelved elsewhere; continue.
2. **Unshelve $|1\rangle$.** Apply a calibrated π-pulse on the quadrupole transition that transfers $|1\rangle$ from the $5D_{5/2}$ manifold to the bright $6S_{1/2}$ manifold.
3. **Detect $|1\rangle$.** Repeat the fluorescence detection as in Step 1.
   - *Bright:* population was in $|1\rangle$.
   - *Dark:* population is shelved elsewhere; continue.
4. **Iterate for the remaining basis states.** Repeat the "unshelve $\rightarrow$ detect" cycle for all remaining $|n\rangle$ ($n \in \{2, d\}$ for d level qudit), each time using the appropriate π-pulse to bring the selected dark state into the bright manifold.

A *bright* event is defined with a *dynamic* threshold: for every block of 1 000 shots we re-fit the bimodal photon-count histogram and set the cutoff at the midpoint between

the bright and dark peaks (see Appendix D). Across the entire data set this threshold converges, on average, to about 10 counts in $t_{\mathrm{det}} = 5\,\mathrm{ms}$; counts below that value are classified as *dark.*

Executing the full sequence generates a Boolean string of up to 25 entries, e.g. `DDDBB...`, where "`B`" denotes a bright detection and "`D`" a dark outcome. Because the act of fluorescence detection projects the ion into the bright manifold, the *first* bright event identifies the state with the population; all subsequent detections are disregarded. This pump–probe sequence relies on a single precisely calibrated π-pulse per state and exploits the high bright/dark discrimination fidelity already demonstrated for ion fluorescence. Its robustness is further supported by the long metastable lifetime of the $5D_{5/2}$ manifold (~ 30 s), which keeps shelved population protected from spontaneous decay during the entire measurement cycle. Care must still be taken in selecting the specific quadrupole transitions that deshelve population from $5D_{5/2}$ to $6S_{1/2}$: each pulse must be spectrally isolated enough to avoid off-resonant coupling that would inadvertently depopulate neighbouring $5D_{5/2}$ levels. With this combination of long-lived shelving and carefully vetted, state-selective π-pulses, we obtain reliable, state-resolved readout across the full 25-level qudit.

3.4 SPAM Results

With the preparation sequence in the $6S_{1/2}$ and $5D_{5/2}$ manifolds fully defined, and a state-selective readout routine established, we are now equipped to evaluate the combined state preparation and measurement (SPAM) performance.

The evaluation procedure is conceptually straightforward: each of the 25 computational basis states is individually prepared, followed by the complete fluorescence-plus-deshelving readout sequence. The observed outcomes are recorded for each input state, resulting in a full detection matrix defined by

$$M_{ij} \equiv P(\text{detect } j \mid \text{prepare } i), \tag{3.5}$$

where M_{ij} represents the conditional probability of detecting state j given that state i was prepared.

The resulting matrix is a 25×25 grid whose diagonal entries reflect the individual SPAM fidelities for each state, while the off-diagonal elements reveal state-to-state leakage, crosstalk, and other preparation or detection errors. Figure 3.8 displays the measured detection matrix M_{ij}, compiled from 5000 experimental shots per prepared state. We see an **average state preparation and measurement fidelity of 99.51(5)%**. Data

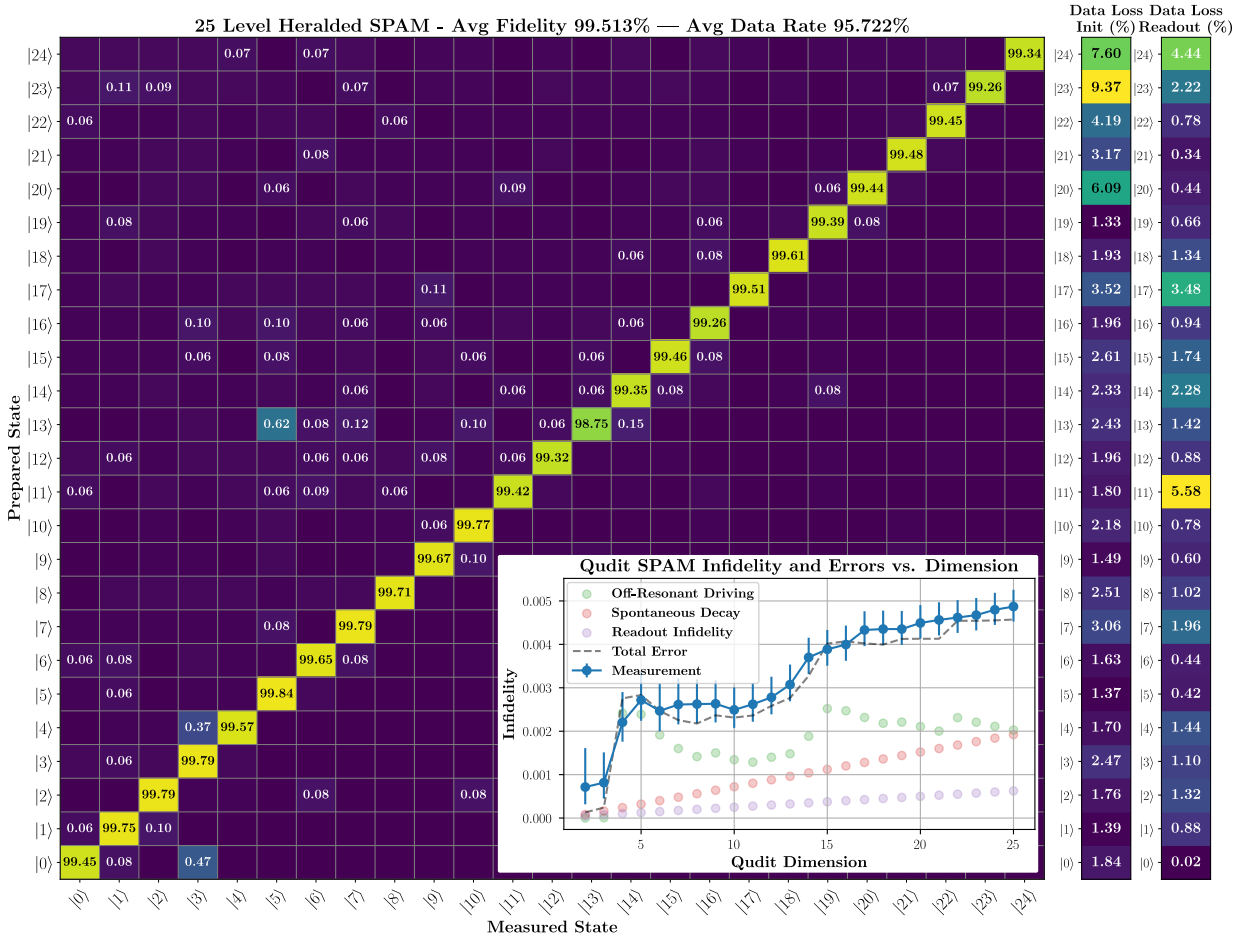


Figure 3.8: **State Preparation And Measurement Results.** A color map showing $M_{i,j}$ elements from (3.5). The inset shows how the error from each channel scales as we add more dimensions to our qudit. The data rate plots on the right show the data loss for each prepared state due to Heralding techniques.

collection was performed in sets of 1000 shots, with frequency calibration of the prepared state carried out before each set. The calibration procedure is detailed in Chapter 4. The states are encoded in descending order of the Rabi frequency of the transition used for shelving and deshelving the states. The encoding for this SPAM experiment is shown in Figure 3.7.

Following data acquisition, the raw outcomes are further refined using a heralding-based post-processing step, which suppresses detectable errors and improves SPAM fidelity. The heralding procedure is described in the subsequent section.

3.4.1 Heralding Techniques and Error breakdown

To suppress trials in which obvious preparation or read-out faults occur, we add two simple heralds around every SPAM cycle and retain only those shots that pass both checks.

Herald #1 — Initialization check. Immediately after the preparation sequence we attempt a *preview* fluorescence measurement of the $6S_{1/2}$ manifold:

1. For target states $|1\rangle \ldots |24\rangle$ (which reside in $5D_{5/2}$) the ion *should* be dark; a bright result signals that population never left $6S_{1/2}$, so the shot is discarded.

2. For the special case $|0\rangle \in 6S_{1/2}$ we first shelve the ion to a reliably isolated $|0\rangle \leftrightarrow 5D_{5/2}$ transition, perform the same fluorescence check, and then deshelve. A bright preview again indicates a failed initialization. And the shot is rejected. In Figure 3.8, we can see the artifact of this heralding technique. The residual population in $M_{0,3}$ is due to this heralding technique and indicates the error in deshelving the pulse.

The fraction of shots that fail Herald #1 is logged as a data rate penalty attributable to imperfect initialization. For the data set presented here, this data loss due to initialization is $\eta_{\text{init}} \approx 2.8\%$ of the raw trials.

Herald #2 — Readout consistency check. A properly functioning readout sequence must yield at least one bright detection at the step that projects the ion onto its prepared state. If *all* 25 detection windows return photon counts below the threshold value, the shot is flagged as a deshelving error and discarded. In the presented data set, this second herald removes, on average, 1.8% of the trials, a loss reported as data loss due to readout error.

Error budget. The state-by-state data losses—and the residual, single-state SPAM Infidelities extracted from the diagonal of M_{ij} are well described by the stochastic model developed by Nicholas Zutt, details can be found in Ref. [44]. In that framework, the total SPAM error splits into three contributions:

1. *Off-resonant scattering* during the shelving and deshelving pulses, which can spuriously transfer population to neighbouring Zeeman sub-levels. This reduces our fidelity by **0.20%.** We can see big off-resonant scattering error in Figure 3.8 $M_{13,5}$. This is also confirmed by looking at the off-resonant scattering chart in Figure 3.5.

2. *Spontaneous decay* from the metastable $5D_{5/2}$ manifold while the sequence is in progress. This reduces our fidelity by **0.19%.**

3. *Bright–dark discrimination errors* arising from the finite overlap of the photon-count histograms. This reduces total fidelity by **0.06%.**

To place our $d = 25$ average SPAM in context, I report a qubit-equivalent figure by normalizing to the information content of the register. If F_d denotes the average correct assignment portability over d basis states, we define

$$F_{\text{qubit-equiv}} = F_d^{1/\log_2 d} \tag{3.6}$$

For our reult $F_{25} = 0.9951$, thi sgives $F_{\text{qubit-equiv}} \approx 0.9989$. For superconducting qubits, a representative state-of-the-art experiment reports an average assignment fidelity of 99.77% (best qubit $> 99.9\%$) using fast multiplexed readout [45]. Against state-of-the-art trapped-ion SPAM on signal qubits, our qubit-equivalent figure is within an order of magnitude in error of the leading records: Quantinuum (99.9904%) [6], IonQ (99.96%) [46], and Oxford Ionics (99.9993%, heralded) [43], noting that those qubit benchmarks typically target one or few favorable states while our figure averages uniformly across 25 basis states.

Furthermore, the main SPAM errors stem from off-resonant scattering during the shelving/deshelving pulses and from spontaneous decay of the $5D_{5/2}$ state while fluorescence is being collected; both are therefore reducible. A higher-numerical-aperture collection optic would allow the detection window to be shortened, cutting the time available for spontaneous decay, while retuning the magnetic field to maximize the isolation of used transitions, suppressing off-resonant excitation, and together pushing the overall SPAM fidelity higher.

Chapter 4

Calibrations — Transition Frequencies and Strengths

High-fidelity state preparation and measurement in the previous chapter hinged on delivering well-calibrated shelving and deshelving pulses. Achieving that fidelity demands that we know the exact transition frequency and Rabi frequency for every line we intend to drive. Yet the $6S_{1/2}\,F = 2 \leftrightarrow 5D_{5/2}$ system offers 80 quadrupole-allowed transitions. A conventional approach—full frequency scans followed by Rabi flops on each line—would occupy the experiment for hours every day, only to be undermined by slow drifts in magnetic field and in the 1762 nm cavity that shift those frequencies on a timescale of few minutes.

Therefore, this chapter reintroduces and adapts the compact calibration toolkit described in Ref. [28] for all 25 levels in $5D_{5/2}$and all 80 possible transitions, which makes exhaustive scans obsolete. Furthermore, instead of traditional way of stepping the 1762 nm laser across each line to hunt for its resonance, we calibrate every transition with a Ramsey-based phase measurement. By letting the ion accumulate phase during a fixed dark time and reading out that phase directly, we pinpoint the frequency detuning in just two data points—achieving finer resolution and far less overhead (in terms of calibration time, AWG programming time, and fitting) than any full sweep could offer.

With the resonant frequencies secured, the next step is to establish reliable transition strengths. Measuring the Rabi frequency for every one of the eighty quadrupole lines would again impose an excessive burden on experiment time. Instead, the calibration procedure determines the π-pulse times of five carefully chosen transitions—one representative for each allowed Δm_F—and then scales the remaining seventy-five using the (relative)

quadrupole matrix-element calculated in Chapter 2. As a result, π-pulse timings for every transition are extracted within minutes rather than hours, keeping the full qudit manifold ready for high-fidelity control without cumbersome overhead.

4.1 Qubit Dynamics and Ramsey Interferometry

Before we delve into the details of the Ramsey experiment on a qubit, it is helpful to first understand how a classical laser field interacts with a two-level atomic system. In the context of trapped ion quantum computing, this interaction forms the basis of coherent control over qubit states. Consider a two-level ion with ground and excited states separated by energy $\hbar\omega_0$, interacting with a classical laser field of frequency ω, Rabi frequency Ω, and a controllable phase ϕ. The system is described by the Hamiltonian:

$$H = -\frac{\hbar\omega_0}{2}\sigma_z + \hbar\Omega\cos(\omega t + \phi)\sigma_x \tag{4.1}$$

Here, σ_z and σ_x are the standard Pauli operators in the two-level basis. The first term represents the internal energy of the qubit, and the second term describes the time-dependent coupling to the laser. To make analytical progress, we move to the interaction picture with respect to a rotating frame defined by:

$$H_0 = -\frac{\hbar\omega}{2}\sigma_z, \quad H_1 = \hbar\Omega\cos(\omega t + \phi)\sigma_x \tag{4.2}$$

The Hamiltonian in the interaction picture becomes:

$$H_{\text{int}} = e^{iH_0 t/\hbar}(H - H_0)e^{-iH_0 t/\hbar} = e^{iH_0 t/\hbar}H_1 e^{-iH_0 t/\hbar} \tag{4.3}$$

To evaluate this, we use the identity for the rotated Pauli operator:

$$e^{i\omega t\sigma_z/2}\sigma_x e^{-i\omega t\sigma_z/2} = \cos(\omega t)\sigma_x + \sin(\omega t)\sigma_y \tag{4.4}$$

Substituting this back in gives:

$$H_{\text{int}} = \frac{\hbar\Delta}{2}\sigma_z + \frac{\hbar\Omega}{2}\left[\cos(\omega t + \phi)\cos(\omega t)\sigma_x + \cos(\omega t + \phi)\sin(\omega t)\sigma_y\right] \tag{4.5}$$

where we have defined the detuning $\Delta = \omega_0 - \omega$. We now simplify the products of trigonometric functions using standard identities:

$$\cos(\omega t + \phi)\cos(\omega t) = \frac{1}{2}[\cos(2\omega t + \phi) + \cos\phi] \tag{4.6}$$

$$\cos(\omega t + \phi)\sin(\omega t) = \frac{1}{2}[\sin(2\omega t + \phi) - \sin\phi] \tag{4.7}$$

Inserting these into the Hamiltonian yields:

$$H_{\text{int}} = \frac{\hbar\Delta}{2}\sigma_z + \frac{\hbar\Omega}{2}\left[\cos(2\omega t + \phi) + \cos\phi\right]\sigma_x + \left[\sin(2\omega t + \phi) - \sin\phi\right]\sigma_y \tag{4.8}$$

Next, we apply the Rotating Wave Approximation (RWA), which involves neglecting the rapidly oscillating terms at frequency 2ω. This leaves us with an effective time-independent interaction Hamiltonian:

$$\Rightarrow H_{\text{int}} = \frac{\hbar\Delta}{2}\sigma_z + \frac{\hbar\Omega}{2}\left(\cos\phi\,\sigma_x - \sin\phi\,\sigma_y\right) \tag{4.9}$$

This expression captures the essential physics of laser-driven qubit rotations with arbitrary detuning and phase. To understand the outcome of a Ramsey experiment, we now solve for the time evolution under this Hamiltonian. The evolution operator is:

$$U(t) = e^{-iH_{\text{int}}t/\hbar} \tag{4.10}$$

Expanding the exponential, we write:

$$U(t) = \exp\left[-i\left(\frac{\Delta}{2}\sigma_z + \frac{\Omega}{2}\cos\phi\,\sigma_x - \frac{\Omega}{2}\sin\phi\,\sigma_y\right)t\right] \tag{4.11}$$

We define the total effective Rabi frequency:

$$\lambda = \sqrt{\Delta^2 + \Omega^2} \tag{4.12}$$

and the normalized axis of rotation:

$$\hat{n} = \frac{1}{\lambda}(\Delta\sigma_z + \Omega\cos\phi\,\sigma_x - \Omega\sin\phi\,\sigma_y) \tag{4.13}$$

Using Euler's identity for exponentials of Pauli operators, we obtain:

$$U(t) = \cos\left(\frac{\lambda t}{2}\right)\mathbb{I} - i\sin\left(\frac{\lambda t}{2}\right)\hat{n} \tag{4.14}$$

In full, this gives:

$$\begin{aligned} U(t) = \cos\left(\frac{\sqrt{\Delta^2+\Omega^2}t}{2}\right)\mathbb{I} \\ - i\sin\left(\frac{\sqrt{\Delta^2+\Omega^2}t}{2}\right)\left(\frac{\Delta}{\sqrt{\Delta^2+\Omega^2}}\sigma_z + \frac{\Omega\cos\phi}{\sqrt{\Delta^2+\Omega^2}}\sigma_x \quad - \frac{\Omega\sin\phi}{\sqrt{\Delta^2+\Omega^2}}\sigma_y\right) \end{aligned} \tag{4.15}$$

This analytical expression governs the qubit state after a pulse of duration t, and will be central to analyzing the phase evolution observed in Ramsey spectroscopy.

In a typical Ramsey experiment, we begin with the qubit in the ground state $|0\rangle$ and apply the following sequence of operations:

1. A $\pi/2$ pulse (duration $t_{\pi/2}$) with phase $\phi = 0$, which creates an equal superposition of $|0\rangle$ and $|1\rangle$.
2. Free evolution for a time T during which the qubit evolves under the detuning $\Delta = \omega_0 - \omega$.
3. A second $\pi/2$ pulse, also of duration $t_{\pi/2}$, but with a variable phase ϕ relative to the first.

The final state is measured in the computational basis, and we are typically interested in the population of the ground state $|0\rangle$ as a function of ϕ or Δ. To derive a compact analytical expression for the probability of measuring the qubit in $|0\rangle$ at the end of the Ramsey sequence, we make the following assumptions:

(1) **Small Detuning:** The detuning Δ is much smaller than the Rabi frequency, i.e., $\Delta \ll \Omega$. This allows us to drop second-order terms like Δ^2 and make the approximation $\sqrt{\Delta^2 + \Omega^2} \approx \Omega$.

(2) **Short Pulses:** The $\pi/2$ pulses are assumed to be short compared to the timescale of detuning-induced phase accumulation, i.e., $t_{\pi/2} \ll T$. This means the free evolution dominates the Ramsey fringe behavior.

(3) **Neglect of Decoherence:** We assume no decoherence or dephasing occurs during the sequence (i.e., $T_2 \to \infty$), so the evolution is entirely unitary.

(4) **Ideal Pulses:** The $\pi/2$ pulses are treated as ideal unitary rotations with full contrast and no amplitude or phase error.

We now derive the analytical expression for the probability of measuring the qubit in state $|0\rangle$ after a Ramsey sequence, making use of the approximations listed above. The sequence consists of a $\pi/2$ pulse with phase 0, free evolution under detuning Δ for a time T, and a second $\pi/2$ pulse with phase ϕ. The qubit begins in the state

$$|\psi_0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \tag{4.16}$$

During the first $\pi/2$ pulse, the qubit evolves under the Hamiltonian

$$H = \frac{\hbar\Delta}{2}\sigma_z + \frac{\hbar\Omega}{2}\sigma_x \tag{4.17}$$

for a duration $t_{\pi/2} = \pi/(2\Omega)$. Applying the approximation (1) and (2) from earlier, the unitary evolution operator becomes:

$$U_1 \approx \frac{1}{\sqrt{2}}\left(I - i\left(\frac{\Delta}{\Omega}\sigma_z + \sigma_x\right)\right) \tag{4.18}$$

Applying U_1 to the initial state yields

$$|\psi_1\rangle = U_1|0\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle - i\frac{\Delta}{\Omega}|0\rangle - i|1\rangle\right) = \frac{1}{\sqrt{2}}\left((1 - i\frac{\Delta}{\Omega})|0\rangle - i|1\rangle\right) \tag{4.19}$$

Now for the second step, we allow the system to evolve freely for a time T under the detuning Hamiltonian

$$H = \frac{\hbar\Delta}{2}\sigma_z \quad \Rightarrow \quad U_{\text{wait}} = \begin{bmatrix} e^{-i\Delta T/2} & 0 \\ 0 & e^{i\Delta T/2} \end{bmatrix} \tag{4.20}$$

Applying U_{wait} to $|\psi_1\rangle$, we get

$$|\psi_2\rangle = U_{\text{wait}}|\psi_1\rangle = \frac{1}{\sqrt{2}}\left((1 - i\frac{\Delta}{\Omega})e^{-i\Delta T/2}|0\rangle - ie^{i\Delta T/2}|1\rangle\right) \tag{4.21}$$

For the final step, we apply a second $\pi/2$ pulse with phase ϕ, using the Hamiltonian

$$H = \frac{\hbar\Delta}{2}\sigma_z + \frac{\hbar\Omega}{2}(\cos\phi\,\sigma_x - \sin\phi\,\sigma_y) \tag{4.22}$$

Again using the approximation (1) and (2), the unitary becomes:

$$U_2 \approx \frac{1}{\sqrt{2}}\left(I - i\left(\frac{\Delta}{\Omega}\sigma_z + \cos\phi\,\sigma_x - \sin\phi\,\sigma_y\right)\right) \tag{4.23}$$

We now write

$$|\psi_2\rangle = \frac{1}{\sqrt{2}}\left(a|0\rangle + b|1\rangle\right) \quad \text{where} \quad a = (1 - i\frac{\Delta}{\Omega})e^{-i\Delta T/2}, \quad b = -ie^{i\Delta T/2} \tag{4.24}$$

Applying U_2 to $|\psi_2\rangle$, we compute the amplitude of the state $|0\rangle$:

$$\psi_0 = \frac{1}{2}\left[a - i\left(\frac{\Delta}{\Omega}a + \cos\phi\, b + i\sin\phi\, b\right)\right] = \frac{1}{2}\left[a(1 - i\frac{\Delta}{\Omega}) - ibe^{i\phi}\right] \tag{4.25}$$

Substituting the expressions for a and b, we have:

$$\psi_0 = \frac{1}{2}\left[(1 - i\frac{\Delta}{\Omega})^2 e^{-i\Delta T/2} - e^{i(\Delta T/2+\phi)}\right] \tag{4.26}$$

Now squaring the magnitude to compute the probability:

$$P_0 = |\psi_0|^2 = \frac{1}{4}\left|(1 - i\frac{\Delta}{\Omega})^2 e^{-i\Delta T/2} - e^{i(\Delta T/2+\phi)}\right|^2 \tag{4.27}$$

Factor out $e^{-i\Delta T/2}$:

$$P_0 = \frac{1}{4}\left|(1 - i\frac{\Delta}{\Omega})^2 - e^{i(\Delta T+\phi)}\right|^2 \tag{4.28}$$

Now expand:

$$(1 - i\frac{\Delta}{\Omega})^2 = 1 - 2i\frac{\Delta}{\Omega} - \left(\frac{\Delta}{\Omega}\right)^2 \approx 1 - 2i\frac{\Delta}{\Omega} \tag{4.29}$$

So:

$$P_0 = \frac{1}{4}\left|1 - 2i\frac{\Delta}{\Omega} - e^{i(\Delta T+\phi)}\right|^2 \tag{4.30}$$

Write $z = e^{i(\Delta T+\phi)}$. Then:

$$P_0 = \frac{1}{4}\left[|1 - 2i\frac{\Delta}{\Omega}|^2 + |z|^2 - 2\,\mathrm{Re}\left[(1 + 2i\frac{\Delta}{\Omega})z\right]\right] \tag{4.31}$$

Using $|1 - 2i\frac{\Delta}{\Omega}|^2 = 1 + \mathcal{O}\left((\frac{\Delta}{\Omega})^2\right)$, and $|z| = 1$, we get:

$$P_0 \approx \frac{1}{2}\left[1 + \cos(\Delta T + \phi)\right] \tag{4.32}$$

This is the standard Ramsey interference fringe pattern, expressed in terms of the accumulated phase ΔT and the pulse phase offset ϕ. In practice, we often vary the phase ϕ in Ramsey experiments to trace out the full sinusoidal fringe, allowing for precise characterization of coherence and detuning.

4.2 Ramsey-Based Frequency Calibration — Practical Considerations

In our experimental procedure, we calibrate the transition frequencies by performing Ramsey experiments at two specific phase values, $\pi/2$ and $3\pi/2$. This allows us to probe the detuning of the system. From these measurements, two points are typically sufficient to obtain a good (sub 50 Hz) fit for the transition frequency detuning. However, in practice, we must account for the finite coherence times of the system. As a result, assumptions (2) and (3), which assume perfect coherence and negligible pulse duration effects, are often no longer valid. Specifically, the $t_{\pi/2}$ times of the transitions can be comparable to the Ramsey wait time, thereby affecting the accuracy of our measurements. To mitigate these effects and provide a more accurate estimate of the detuning, we utilize assumption-free unitaries (see 4.2 for their definition), which are computed numerically. This approach ensures that the detuning values are derived from the actual evolution of the system, providing a more robust calibration despite the presence of decoherence.

Algorithm 2 Estimating Detuning from Two Ramsey Points

Input: π-pulse time t_π, wait time T, measured populations $\left(P_{\pi/2}^{(\text{exp})}, P_{3\pi/2}^{(\text{exp})}\right)$
Define Pauli matrices: σ_x, σ_y, σ_z, and identity I
function AnalyticalUnitary(Δ, Ω, ϕ, t)
    $\lambda \leftarrow \sqrt{\Delta^2 + \Omega^2}$
    return $\cos\left(\frac{\lambda t}{2}\right) I - i \sin\left(\frac{\lambda t}{2}\right) \left(\frac{\Delta}{\lambda}\sigma_z + \frac{\Omega\cos\phi}{\lambda}\sigma_x - \frac{\Omega\sin\phi}{\lambda}\sigma_y\right)$
end function
function StatePopulation(t_π, T, Δ, ϕ)
    $\Omega \leftarrow \pi / t_\pi$
    $t_{\pi/2} \leftarrow t_\pi / 2$
    $U_1 \leftarrow$ AnalyticalUnitary$(\Delta, \Omega, 0, t_{\pi/2})$
    $U_{\text{wait}} \leftarrow$ AnalyticalUnitary$(\Delta, 0, 0, T)$
    $U_2 \leftarrow$ AnalyticalUnitary$(\Delta, \Omega, \phi, t_{\pi/2})$
    $U_{\text{total}} \leftarrow U_2 \cdot U_{\text{wait}} \cdot U_1$
    return $|\langle 0|U_{\text{total}}|0\rangle|^2$
end function
function FitError(Δ)
    $P_{\pi/2}^{(\text{model})} \leftarrow$ StatePopulation$(t_\pi, T, \Delta, \pi/2)$
    $P_{3\pi/2}^{(\text{model})} \leftarrow$ StatePopulation$(t_\pi, T, \Delta, 3\pi/2)$
    return Error: $(P_{\pi/2}^{(\text{model})} - P_{\pi/2}^{(\text{exp})})^2 + (P_{3\pi/2}^{(\text{model})} - P_{3\pi/2}^{(\text{exp})})^2$
end function
Use Nelder-Mead method to minimize FitError over Δ
Output: Optimal detuning Δ

To implement this calibration strategy, we use the algorithm described in the code snippet above. This routine fits the measured ground state populations at two Ramsey phase points, $\pi/2$ and $3\pi/2$, to the expected evolution computed via numerically evaluated unitaries. By minimizing the squared error between the model and experimental results, the algorithm returns the optimal detuning Δ for the transition, as illustrated in Figure 4.1 (a).

To validate the accuracy of this method, we performed a controlled experiment in which known detunings were deliberately introduced on a selected transition. The transition chosen for this test is the most magnetically field-insensitive qubit transition in our system, with a coherence time (T_2) of approximately 1200 μs and a π-pulse time of around 128 μs. In this experiment, the Ramsey wait time was 100 μs, making the pulse duration actually

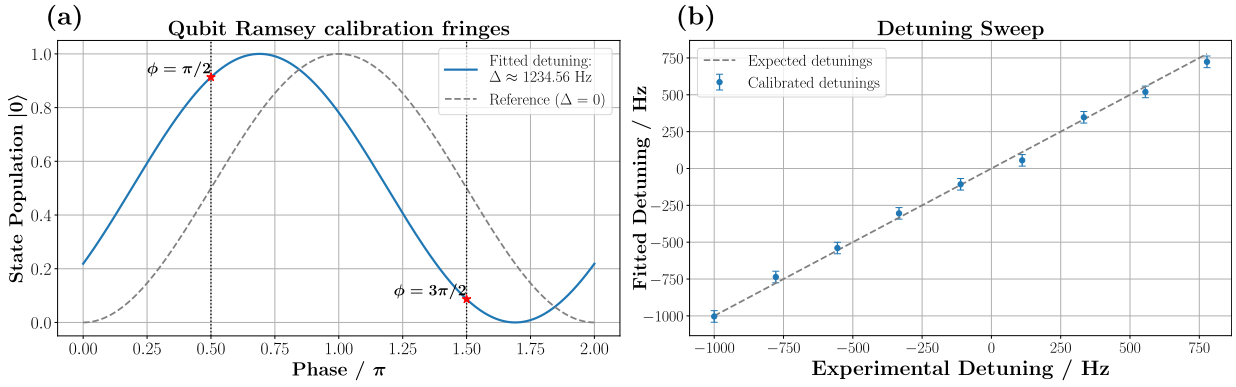


Figure 4.1: **2-point Calibration Ramsey fringes and Detuning sweep. (a)** Illustration of the two-point Ramsey calibration method. The population of the $|0\rangle$ state is measured at phases $\phi = \pi/2$ and $3\pi/2$, and the resulting values are used to fit a Ramsey fringe with the detuning Δ as the only free parameter. The data shown here is from a numerical simulation with a known detuning of $\Delta = 1234.56\,\text{Hz}$, chosen for illustrative purposes. **(b)** Results from a controlled experimental test, in which known detunings were applied to validate the accuracy and robustness of the fitting procedure.

bigger than the free evolution time, conditions under which assumptions (2) and (3) are no longer valid. The calibration script was then used to estimate the detuning from the measured Ramsey populations. The results of this test are shown in Figure 4.1 (b). The detunings predicted by the script closely matched the known offsets, with a standard deviation of just 33 Hz from the expected values.

While the two-point Ramsey method offers efficient (in terms of the number of measurements required) and accurate frequency calibration, it comes with intrinsic limitations. In particular, there is an upper bound on the absolute detuning $|\Delta|$ for which the fitting remains valid. This constraint arises from the requirement that the accumulated phase $|\Delta \cdot T|$ remain within the non-degenerate region of the Ramsey fringe. Empirically, we find that the method becomes unreliable when $|\Delta \cdot T| \gtrsim \pi/2$, as the two points on the sinusoidal curve become ambiguous, leading to degeneracies in the fitted detuning values.

Given the approximate analytical expression for the Ramsey population derived earlier,

$$P_0 \approx \frac{1}{2}\left[1 + \cos(\Delta T + \phi)\right] \tag{4.33}$$

we can estimate the uncertainty in the fitted detuning Δ, assuming the dominant source of uncertainty arises from the finite number of measurement shots n used to estimate the

population P_0. We assume that ϕ and T are known precisely. Inverting the population expression, we obtain

$$\Delta = \frac{\arccos(2P_0 - 1) - \phi}{T}. \tag{4.34}$$

To propagate the uncertainty from P_0 to Δ, we apply standard error propagation:

$$\sigma_\Delta^2 = \left(\frac{d\Delta}{dP_0}\right)^2 \sigma_{P_0}^2. \tag{4.35}$$

Computing the derivative,

$$\frac{d\Delta}{dP_0} = \frac{d}{dP_0}\left(\frac{\arccos(2P_0 - 1) - \phi}{T}\right) = \frac{-2}{T\sqrt{1 - (2P_0 - 1)^2}} \tag{4.36}$$

$$= \frac{-2}{T\sqrt{4P_0(1 - P_0)}} = \frac{-1}{T\sqrt{P_0(1 - P_0)}}. \tag{4.37}$$

The uncertainty in the population P_0 due to projection noise in n measurement shots is given by the standard deviation of a binomial distribution:

$$\sigma_{P_0} = \sqrt{\frac{P_0(1 - P_0)}{n}}. \tag{4.38}$$

Substituting into the expression for σ_Δ, we obtain:

$$\sigma_\Delta = \frac{1}{T\sqrt{P_0(1 - P_0)}} \cdot \sqrt{\frac{P_0(1 - P_0)}{n}} = \frac{1}{T\sqrt{n}}. \tag{4.39}$$

Thus, the final expression for the uncertainty in the predicted detuning is:

$$\boxed{\sigma_\Delta = \frac{1}{T\sqrt{n}}} \tag{4.40}$$

While the analytical result $\sigma_\Delta = 1/(T\sqrt{n})$ is useful for intuition, in our experiment, we estimate Δ by numerically fitting the full model to two or more measurements. In that setting, the analytical expression should be read as a scaling law rather than a strict error bar—it illustrates how the resolution improves with increasing Ramsey wait time T and total number of shots n. The actual uncertainty is determined from a numerical error propagation method based on the first-order Taylor expansion.

Let r_1 and r_2 denote the measured populations at two different phases (e.g., $\phi = \pi/2$ and $\phi = 3\pi/2$ as used in 4.2). These observables are related to the unknown detuning Δ through a nonlinear mapping:

$$\Delta = f(r_1, r_2), \tag{4.41}$$

where f represents the fitting process: a numerical minimization of the squared difference between observed populations and those predicted by the Ramsey model (see 4.2). Since f is a numerical optimizer, we estimate its sensitivity to the inputs using finite differences. By linearizing f using the first-order Taylor expansion around the mean values of r_1 and r_2 and taking variance:

$$\sigma_\Delta^2 \approx \left(\frac{\partial \Delta}{\partial r_1}\right)^2 \sigma_{r_1}^2 + \left(\frac{\partial \Delta}{\partial r_2}\right)^2 \sigma_{r_2}^2. \tag{4.42}$$

This assumes r_1 and r_2 are uncorrelated, which is valid in this case since they're separate experiments. In practice, we approximate the partial derivatives using symmetric finite differences:

$$\frac{\partial \Delta}{\partial r_1} \approx \frac{f(r_1 + \sigma_{r_1}, r_2) - f(r_1 - \sigma_{r_1}, r_2)}{2\sigma_{r_1}}, \qquad \frac{\partial \Delta}{\partial r_2} \approx \frac{f(r_1, r_2 + \sigma_{r_2}) - f(r_1, r_2 - \sigma_{r_2})}{2\sigma_{r_2}}. \tag{4.43}$$

where σ_{r_i} is, standard deviation in binomial sample, given by $\sigma_{r_i} = \sqrt{r_i(1-r_i)/n}$. The final uncertainty is then computed by combining these in quadrature:

$$\sigma_\Delta = \sqrt{\left(\frac{\partial \Delta}{\partial r_1}\sigma_{r_1}\right)^2 + \left(\frac{\partial \Delta}{\partial r_2}\sigma_{r_2}\right)^2}. \tag{4.44}$$

This method is robust and model-agnostic, as it does not require analytic knowledge of the function f, and instead relies solely on the numerical fitting procedure.

The uncertainty or detuning resolution analysis also guides our practical choice of Ramsey wait time T. Since the uncertainty in detuning scales inversely with T (see (4.40)), longer wait times improve precision. However, the accumulated phase $|\Delta T|$ must remain within the non-degenerate region of the Ramsey fringe ($|\Delta T| \lesssim \pi/2$) to ensure the fitting remains unambiguous. In our daily experiments, we typically use a Ramsey wait time of 250,μs, which provides a reliable detuning range of approximately ± 1 KHz. Given that we perform frequency calibrations every few minutes and our system remains stable within

this window, this choice strikes an effective balance between sensitivity and dynamic range. This calibration method enabled us to resonantly drive the transitions used in the SPAM experiment of Chapter 3 with over 99.5% fidelity. For calibrations in those experiments, a Ramsey wait time of $400\,\mu$s was used, providing a detuning range of approximately $\pm 625\,$Hz.

Although this calibration method is significantly faster than performing full frequency sweeps to locate resonance, it remains impractical to apply it individually to every transition when multiple transitions are involved in a single operation. To address this, we implement and adapt the calibration technique described in Ref. [28], which is detailed in the following section.

4.3 Data collection and Fitting

To avoid the need to calibrate every individual transition frequency each time multiple transitions are used in an operation, we adopt a referencing scheme based on two carefully chosen reference transitions. The first transition, which we denote f_0, is selected to be the most magnetically insensitive one: from $S_{1/2}$, $F = 2$, $m_F = 0$ to $D_{5/2}$, $F = 2$, $m_F = 0$ (magnetic field sensitivity $\sim 40\,$Hz/G). This transition is minimally affected by ambient magnetic field fluctuations and primarily experiences frequency drift due to slow changes in laser frequency, such as those arising from optical cavity drift. The second reference transition, denoted f_1, is chosen to be one of the most magnetically sensitive: from $S_{1/2}$, $F = 2$, $m_F = -1$ to $D_{5/2}$, $F = 4$, $m_F = -3$ (magnetic field sensitivity $\sim 3.5\,$MHz/G). By tracking both f_0 and the frequency difference $f_1 - f_0$, we can effectively capture and compensate for the two dominant sources of frequency drift in our system—laser frequency drift and magnetic field fluctuations.

To reference the remaining 78 transitions (denoted f_n) to the two primary transitions f_0 and f_1, we collect a set of "calibration data." This process involves sequentially measuring the frequencies of f_0, f_1, and f_n for 40 selected transitions.These transitions are carefully chosen to collectively span all 24 states in the $6D_{5/2}$ manifold and to provide coverage across the various $6S_{1/2}, F = 2$ Zeeman sublevels. The selected set ensures that for each $D_{5/2}$ level, there is at least one transition measured, and that the frequency separations between the $6S_{1/2}$ $F = 2$ sublevels can be inferred. This allows us to extrapolate the frequencies of the remaining transitions with high accuracy. The chosen transitions used for calibration are shown in Figure 4.2.

To make use of this referencing scheme, we aim to express each transition frequency f_n as a function of f_0 and f_1. From the analysis presented in Chapter 2 Section 2.3, we know

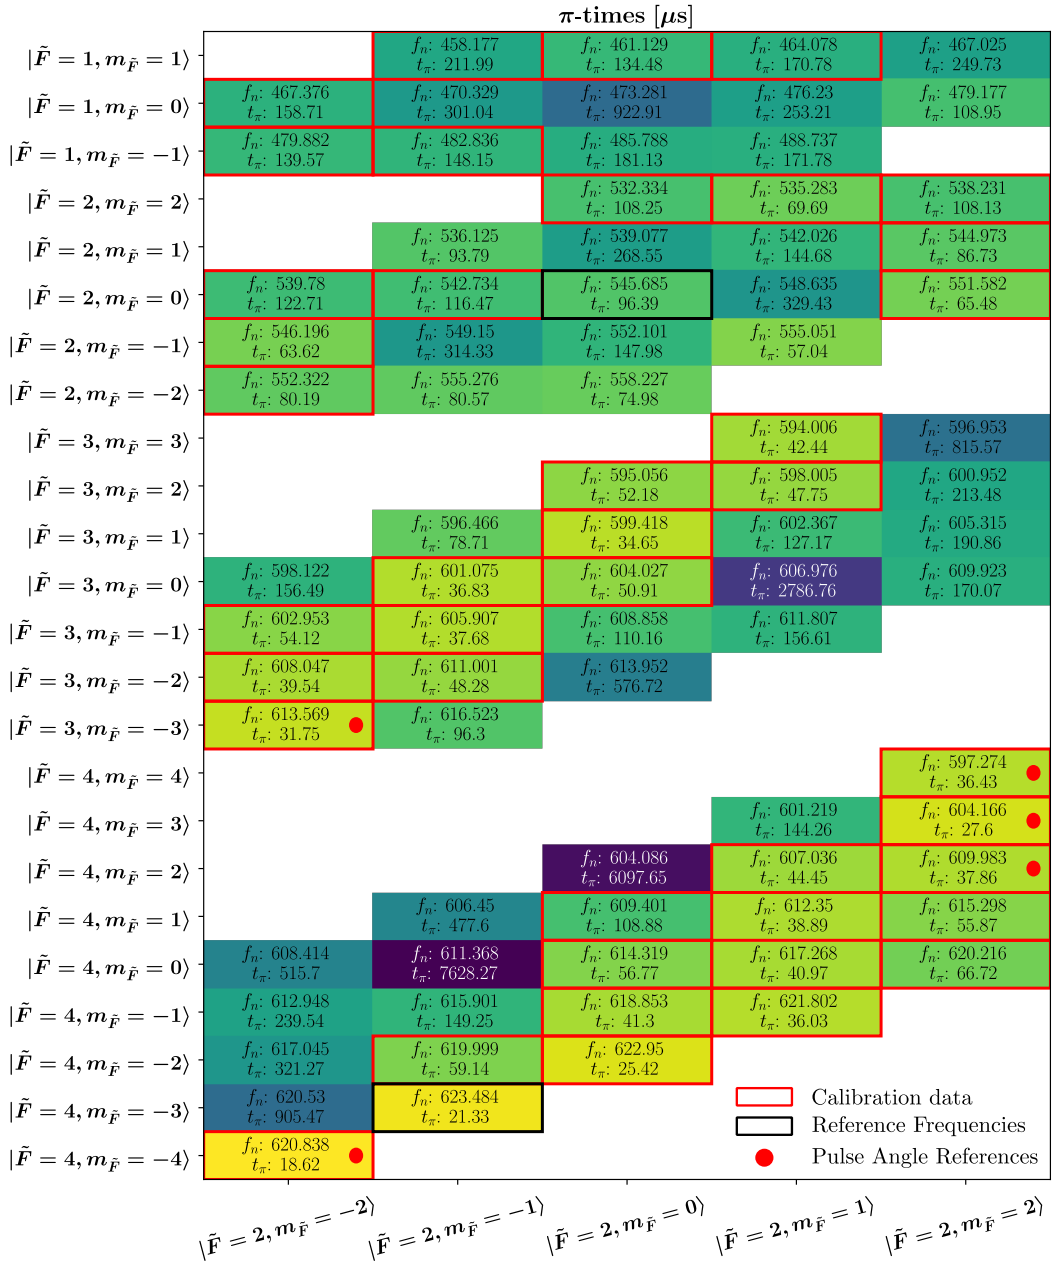


Figure 4.2: **π-times and Calibration, reference transitions.** A colormap of calibrated pi-times of all the allowed transitions (transition strengths are calculated at 4.209 G of B-field). Red boxes show the chosen calibration frequencies for data collection. Black boxes show the two reference transitions for frequency calibrations. And red dots show the five reference transitions for pulse angle calibration.

that the energy levels in the $5D_{5/2}$ manifold exhibit a linear dependence on the magnetic field over small variations in field strength. Consequently, each transition frequency f_n also inherits a linear dependence on the magnetic field, which is captured by the difference $f_1 - f_0$, and a dependence on the laser frequency drift, captured by f_0.

We therefore model the transition frequency f_n using the following form:

$$f_n = \beta_n(f_1 - f_0) + \alpha_n + f_0, \tag{4.45}$$

where β_n encodes the magnetic field sensitivity of the transition relative to the reference, and α_n is a constant offset determined from calibration. We need this calibration offset because the transition frequencies calculated by solving the Hyperfine Hamiltonian (in chapter 2 Section 2.3) have discrepancies when compared to analytically measured transitions. We speculate that state mixing with levels in $5D_{3/2}$ needs to be considered in our calculations to obtain a more complete picture of all perturbations. The parameter β_n can be determined from theory using the calculated magnetic field sensitivities of the transitions:

$$\beta_n = \frac{\kappa_n - \kappa_0}{\kappa_1 - \kappa_0}, \tag{4.46}$$

where κ_0, κ_1, and κ_n are the magnetic field sensitivities of the reference transitions f_0, f_1, and the transition f_n, respectively. These sensitivity values are provided in Figure 2.8. The magnetic field used to calculate these sensitivity values was estimated by back-calculating B_0 in the hyperfine Hamiltonian, based on a fit to the frequency difference between f_0 and f_1. The calibration data described above is then used to determine the best-fit value of α_n for each transition.

Once both α_n and β_n have been determined and the two reference frequencies f_0 and f_1 are measured, the frequency of any desired transition frequency f_n can be calculated using either (4.45) (if the that transition is part of calibration dataset) or by extrapolating its frequency using Zeeman splittings within the $S_{1/2}$, $F = 2$ manifold.

For example, consider the transition from $6S_{1/2}, F = 2, m_F = -1$ to $5D_{5/2}, F = 2, m_F = -2$, which we denote in shorthand as $[-1, 2, -2]$. This transition is not part of the calibration set, but its frequency can be inferred by adding the Zeeman energy difference between the $m_F = -2$ and $m_F = -1$ states in $6S_{1/2}, F = 2$. This frequency spacing can be estimated by averaging the differences observed in transitions that share the same $5D_{5/2}$ state but differ in the $6S_{1/2}$, $F = 2$ sublevel. Specifically, the spacing can be averaged from the differences between:

- $[-2, 3, -2]$ and $[-1, 3, -2]$,

- $[-2, 3, -1]$ and $[-1, 3, -1]$,
- $[-2, 2, 0]$ and $[-1, 2, 0]$,
- $[-2, 1, -1]$ and $[-1, 1, -1]$

And the spacing between any pair of m_F levels within the $F = 2$ manifold of the $6S_{1/2}$ state can be calculated in the same manner, using frequency differences between transitions that share a common upper $5D_{5/2}$ state. Further analysis of the collected calibration data shows that the average deviation between the predicted transition frequencies and the experimentally measured values using this referencing method is $\approx 94\,\mathrm{Hz}$. A more detailed discussion of calibration errors and systematic deviations is presented in Chapter 6.

For the average π-pulse length of $100\,\mu\mathrm{s}$ used throughout this work, the corresponding Rabi frequency is about 16 kHz; the residual detuning therefore represents only $\Delta/\Omega \approx 0.006$. That ratio reduces the maximum population transfer—and thus the contrast in any subsequent Ramsey or Rabi experiment—by less than 10^{-3}, an effect well below other technical noise sources considered later in the thesis. In practice, the calibrated frequencies are therefore sufficiently accurate that detuning-induced contrast loss is negligible for both state preparation and coherence measurements. This accuracy can be improved further by increasing the Ramsey free-evolution time once coherence times in the system are extended, shrinking the statistical uncertainty in the fitted detuning.

4.4 Rabi-Frequency Calibration

While the relative transition strengths for all quadrupole transitions were calculated in Chapter 2 Section 2.3, those calculations assumed a linearly polarized quadrupole laser field. However, based on the observed relative Rabi frequencies across various transitions, it is evident that the laser polarization deviates from being purely linear. A more detailed discussion of this discrepancy is provided in Appendix A. As a result, the geometric factors derived and the relative transition strengths cannot be fully trusted for determining absolute Rabi frequencies.

To overcome this, we experimentally measure the Rabi frequency of five carefully selected transitions, each corresponding to a different value of $\Delta m_F \in \{-2, -1, 0, +1, +2\}$. To generalize the calibration to all transitions, we use the five measured Rabi frequencies as references for their respective Δm_F classes. For any other transition with the same Δm_F, we estimate its π-time based on its theoretical transition strength relative to the

corresponding reference transition. Specifically, if $t_{\pi,n}$ denotes the π-time of the target transition, $t_{\pi,0}$ is the measured π-time of the reference transition with the same Δm_F, and s_n, s_0 are the theoretical transition strengths of the target and reference transitions respectively, then we use the scaling relation:

$$\left\{ t_{\pi,n} = \frac{t_{\pi,0} \cdot s_n}{s_0} \right\}_{\Delta m_F} \tag{4.47}$$

This works because geometric factors for all transitions with the same Δm_F are the same (So essentially we are calibrating over the laser polarization and intensity). Using this approach, we get the π-time values presented in Figure 4.2. The figure also shows the five reference transitions used, one for each Δm_F value. We observe that the π-times estimated using this scaling method deviate from the experimentally measured π-times by approximately 1.7%, indicating good agreement between the model and experimental data. A more detailed discussion of this deviation and its implications is presented in Chapter 6.

For a π-pulse calibrated with a 1.7% error (i.e. the pulse angle is $(1+\varepsilon)\pi$ with $\varepsilon = 0.017$), the shelving probability for an ideal two-level system is

$$P_{\text{shelve}} = \sin^2\!\left[(1+\varepsilon)\tfrac{\pi}{2}\right] = \cos^2\!\left(\tfrac{\varepsilon\pi}{2}\right) \simeq 1 - \left(\tfrac{\varepsilon\pi}{2}\right)^2.$$

Substituting $\varepsilon = 0.017$ gives a fidelity drop of

$$1 - P_{\text{shelve}} \;\approx\; \left(\tfrac{0.017\pi}{2}\right)^2 \;\approx\; 7 \times 10^{-4},$$

so the shelving step still succeeds with about 99.93% probability. This $< 0.1\%$ loss is much smaller than other technical errors discussed in the SPAM budget and becomes negligible once the Rabi rate is re-tuned or the drive amplitude is actively stabilised. This calibration can also be accelerated by adopting a two-point method analogous to the frequency calibration: sample a single Rabi oscillation at two phase-offset times near the expected $\pi/2$ and $3\pi/2$ points, fit with the known Rabi-oscillation model, and extract the Rabi rate directly, cutting the number of data points needed while retaining sub-percent accuracy.

Chapter 5

Coherence Characterization — Qudit Ramsey

Interferometry, in essence, turns phase into an observable by splitting a coherent wave, letting each branch accumulate a phase proportional to some physical quantity, and then recombining the branches so that the relative phase appears as constructive or destructive interference. In the *Ramsey* version applied to internal quantum states, two brief, phase-coherent $\pi/2$ pulses act as the beam-splitter and recombiner: the first pulse creates a superposition of energy eigenstates, the system evolves freely so each component acquires its own phase, and the second pulse mixes the paths so that the final-state populations oscillate (Ramsey fringes) as a function of the accumulated phase. Since Ramsey's original molecular-beam experiment in 1950 [47], this sequence has become a standard ruler for coherence: it anchors precision frequency standards such as auto-balanced Ramsey spectroscopy in $^{171}Yb^{+}$ optical clocks [48], where careful pulse tailoring suppresses probe-induced shifts; it benchmarks coherence and sensing protocols in superconducting qubits, as recently shown by deterministic-control schemes that surpass the conventional Ramsey sensitivity limit [49]; and it gauges spin-ensemble coherence in solid-state systems, exemplified by nitrogen-vacancy centres in diamond whose Ramsey fringes reveal millisecond-scale dephasing times[50].

Multilevel, or qudit, versions of Ramsey interferometry build directly on the two-level scheme by replacing the $\pi/2$ "beam-splitters" with gates that distribute amplitude across several eigenstates and by monitoring a single-state population that re-encodes all accumulated relative phases. In semiconductor optics, Lee *et al.* created coherent three- and four-level superpositions in a charged InGaAs quantum dot and showed that the visibility of the resulting Ramsey fringes tracks phase preservation across the entire manifold [51].

Godfrin *et al.* introduced a formal "generalized Ramsey interferometry" for a nuclear-spin-3/2 system, using Hadamard-like pulses to prepare and recombine equal-weight quartit states so that the single-level fringe contrast reports the combined off-diagonal coherence of the qudit density matrix [52]. Similar ideas have emerged in solid-state colour centres, where Vaňo *et al.* employed multi-frequency microwaves to drive Ramsey sequences on vacancy-related spin sub-levels in silicon carbide, extracting coherence information from a room-temperature qudit platform [53]. On the theoretical side, Feng and Wei analyzed multilevel atomic Ramsey spectroscopy. They showed that equal-superposition sequences can reach quantum-Fisher bounds —the fundamental quantum-mechanical limit on how precisely any method can estimate the desired parameter— underscoring the protocol's metorological power [54]. At the same time, Thenuwara and Sidorov provided closed-form Ramsey envelopes for three- and five-level manifolds in ^{87}Rb condensates, demonstrating how analytic treatments can guide experiments that probe coherence across several hyperfine states [55]. Collectively, these studies establish Ramsey interferometry as a practical, tomography-free metric for assessing how well a multilevel system preserves the relative phases—and therefore the usable coherence—of all the states that participate in a quantum operation.

In this chapter, I describe the Qudit Ramsey scheme, developed by Pei Jiang Low, and the underlying pulse sequences we implement using quadrupole transitions for an arbitrary d-level system. Next, I explain how the computational subspace is chosen: a physically motivated cost function is minimized to select the d basis states that form the qudit.The third section presents the experimental phase scans and the measured fringe contrast as d grows, revealing how coherence scales with Hilbert-space size. The chapter closes with a thorough error budget: each identified noise source is quantified and fed, with No-Free-parameters, into a Monte-Carlo model (details in chapter 7) whose predicted contrasts agree with the data for every dimension within experimental uncertainty.

5.1 Qudit Ramsey Scheme

To interrogate the coherence of the entire d-level manifold, we employ a classical Ramsey-type sequence that was conceived in our group by Pei Jiang Low during his postdoc. The idea is to treat the qudit as a multi-arm interferometer: a first global pulse spreads one unit of population uniformly across all computational states, a second global pulse—with state-dependent phases—attempts to re-focus that amplitude back onto $|0\rangle$. The degree to which the population returns to $|0\rangle$ is our single-number measure of coherence.

The experiment proceeds in three steps:

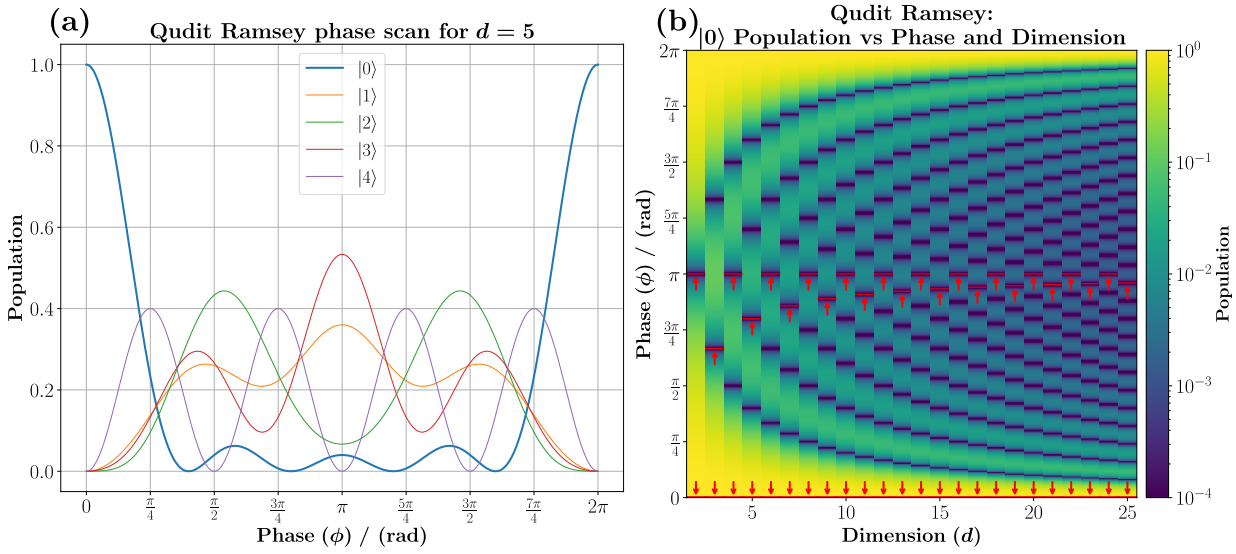


Figure 5.1: **Qudit Ramsey phase scan $d = 5$ and $|0\rangle$ population colormap.(a)** An example of a qudit Ramsey experiment implemented on a $d = 5$ qudit. The plot shows the population dynamics across the five basis states as the phase ϕ is swept from 0 to 2π. Notably, the population returns entirely to the initial state at $\phi = 0$ and $\phi = 2\pi$, indicating constructive interference. At certain intermediate phases, we observe complete destructive interference, where the population is fully transferred to other levels. **(b)** A colormap showing the population of the $|0\rangle$ state as a function of both dimension d and phase ϕ. As the dimensionality increases, interference effects become more complex, resulting in a larger number of phases where the $|0\rangle$ population vanishes. The phases used for contrast measurements in our experiment are indicated in red.

1. **Initialisation** – Prepare the ion in the ground state $|0\rangle$.

2. **Equal-superposition** – Drive a sequence of calibrated 1762 nm quadrupole pulses that, one by one, move a fraction $1/d$ of the population from $|0\rangle$ into $|1\rangle, |2\rangle, \ldots, |d-1\rangle$. For the j^{th} transfer the rotation angle is

$$\theta_j = 2 \arcsin\left(\frac{1}{\sqrt{d-j+1}}\right), \tag{5.1}$$

so that after the final pulse every basis state holds an equal amplitude and the register is in the uniform superposition $\frac{1}{\sqrt{d}} \sum_{k=0}^{d-1} |k\rangle$.

3. **Phase-tagged reversal** – Re-apply the same pulse train in reverse order, but advance the phase of the n^{th} pulse by $n\phi$.

In an ideal noise-free environment, when $\phi = 2\pi\ell$ ($\ell \in \mathbb{Z}$) the phases accumulated by states align to zero and the entire population should return to $|0\rangle$; when $\phi = 2\pi\ell/d$ the interference is maximally destructive and $|0\rangle$ should be empty (except when ℓ is a multiple of d). Analytical derivation of the qudit Ramsey scheme described here can be found in Ref. [44]. Measuring the contrast between these two outcomes provides a direct, tomography-free gauge of how faithfully the qudit preserves relative phases during this operation.

A key departure from textbook Ramsey interferometry is that we insert no free evolution between the superposition step and the recombination step; the two pulse trains follow one another immediately. The contrast we observe, therefore, reflects how well the qudit preserves its relative phases while the relevant superposition components are being driven, rather than during a passive wait. Because the sequence involves every computational state, it effectively exercises the qudit under a congested control pattern we can devise—akin to running a single, maximally complex qudit gate. While a formal gate-complexity proof lies outside the scope of this thesis, the intuition is that any phase error or leakage can appear somewhere in the sequence, so the resulting $|0\rangle$ return probability offers a stringent, system-wide measure of coherence.

5.2 State Selection and Qudit encoding

Because the Qudit-Ramsey experiment probes phase relations among *all* states in the system, the choice of those states must minimize their joint susceptibility to the dominant noise sources in the trap. We frame that choice as an optimization problem: given the measured magnetic-field fluctuations and laser-frequency instabilities (see Chapter 6), we seek the set of d basis states $\{\mathcal{S}\}$ that minimizes a noise-weighted cost function. A Gaussian distribution well describes the magnetic field noise, while the laser noise follows a Voigt profile that accounts for both frequency and phase fluctuations. Decoherence under each noise channel is governed by the characteristic function of its corresponding probability-density function, so the overall vulnerability of a candidate encoding can be expressed as

$$\mathcal{C}(\{\mathcal{S}\}) = \frac{1}{\ell}\sum_{j<k}\left[\left(\frac{\tau_{\{\mathcal{S}\}}}{T_{L,G}^{(j,k)}}\right)^2 + \left(\frac{\tau_{\{\mathcal{S}\}}}{T_{L,L}^{(j,k)}}\right) + \left(\frac{\tau_{\{\mathcal{S}\}}}{T_{B,G}^{(j,k)}}\right)^2\right] \tag{5.2}$$

where the three terms aggregate the pairwise decoherence that arises from Gaussian part of laser noise $T_{L,G}^{(j,k)}$, Lorentzian part of laser noise $T_{L,L}^{(j,k)}$, and Gaussian magnetic-field noise $T_{B,G}^{(j,k)}$, respectively. The parameter

$$\tau_{\{\mathcal{S}\}} = \sum_{j<k} \tau_{\pi}^{(j,k)} \tag{5.3}$$

is the total "π-time budget": the sum of the calibrated π-pulse durations that would be required, in principle, to transfer population between every ordered pair of states in $\{\mathcal{S}\}$. It captures how long the qudit must remain phase-coherent while being actively driven. The normalizing factor ℓ counts the *distinct* states that the full control sequence actually visits; in practice ℓ may exceed the logical dimension d because certain operations route population through auxiliary "bus" states in $D_{5/2}$ or $S_{1/2}$. Framed this way, $\mathcal{C}$ is not tied to the Qudit-Ramsey protocol alone but serves as a platform-agnostic figure of merit for any task that manipulates a selected set of levels in the presence of magnetic field and laser noise.

Applying the cost–function framework to our 29-level spectrum—five Zeeman sub-levels in $S_{1/2}\,(F{=}2)$ and twenty-four in $D_{5/2}$—An *exhaustive search* was carried out for every target dimension d. For each candidate set (i) all $\binom{29}{d}$ possible d-element subsets were compiled, (ii) each subset with any intermediate "bus" states required by the actual control graph were augmented, (iii) the cost $\mathcal{C}$ of the resulting full set was evaluated, and (iv) the subset that produced the global minimum was retained. For $d \leq 9$ it consistently returned a *star-topology* pattern: one state in the $6S_{1/2}\,(F{=}2)$ manifold served as the hub, while the remaining $d-1$ nodes were all drawn from $5D_{5/2}$. At $d = 10$, however, the algorithm introduced another $6S_{1/2}$ Zeeman sublevel to $d = 9$ set, giving the proposed register two states $6S_{1/2}$. Because we encode $|0\rangle$ in the primary $S_{1/2}$ state (see Section 5.2.1), coupling the newly added $6S_{1/2}$ level to $|0\rangle$ would require a two-step route—$|0\rangle \to 5D_{5/2} \to 6S'_{1/2}$—rather than a single quadrupole pulse. This was unexpected, as several $D_{5/2}$ states that couple directly to $|0\rangle$ were still available; yet, the optimizer preferred the more elaborate path. Experimentally, the star-topology register delivered noticeably lower fringe contrast once the second $S_{1/2}$ level was included; replacing that choice with a directly connected $D_{5/2}$ state restored the fidelity. We attribute the discrepancy to a noise source that the present cost function omits—the AC mains line signal component that imparts a systematic phase to each transition in the sequence, decoherence due to this effect increases linearly with total operation time. The influence of this line noise and its impact on coherent operations are also discussed in Appendix E. A revised cost function that includes this additional dephasing channel is left for future work; however, the episode

highlighted the practical value of cross-checking purely numerical state selections against empirical coherence tests.

To make the operation faster, we impose a star-topology constraint: for as large a dimension as possible, the d levels must all couple through a *single* $6S_{1/2}$ hub. Now, for each of the five Zeeman sub-levels in $6S_{1/2}(F = 2)$, the optimizer finds the best set of d-states that connect only to that sub-level; we then compare the five costs and retain the set with the lowest cost. This selection process allows star-topology registers up to $d = 17$. Beyond that point every $5D_{5/2}$ state directly linked to the chosen hub is already in use, so we manually introduce $5D_{5/2}$ level that belong to a different $6S_{1/2}$ hub. From $d = 18$ onward we therefore migrate the $|0\rangle$ population to whichever $6S_{1/2}$ level couples to the $5D_{5/2}$ state being addressed, routing through a single $D_{5/2}$ "bus" state when necessary.

5.2.1 Qudit encoding

After the physical levels are chosen, we must decide how to assign the logical labels $|0\rangle, |1\rangle, \ldots, |d-1\rangle$. The optimal labelling can vary with the task at hand. If an experiment repeatedly couples many levels to a single state in a star-topology graph, it is advantageous to encode that hub in $6S_{1/2}(F{=}2)$; every other level is then accessible with a single quadrupole pulse.

For the Qudit-Ramsey sequence, the requirement is even clearer. The protocol starts by preparing $|0\rangle$, redistributes the population uniformly across the register, and finally reverses the process to return all population to $|0\rangle$. We, therefore, encode $|0\rangle$ in the $S_{1/2}$ state that has the most direct links to the selected $D_{5/2}$ levels. The remaining labels are assigned in *ascending order of coupling strength* to $|0\rangle$. States that couple most weakly become $|1\rangle$, those with progressively stronger couplings become $|2\rangle, |3\rangle, \ldots$, and the level with the largest Rabi frequency is labeled $|d-1\rangle$. This ordering minimizes the total duration of the two pulse sequences in Qudit Ramsey because early pulses only need transfer a small fraction of the initial population, while the final pulse (moving $|d-1\rangle$ to or from $|0\rangle$) uses a $\pi/2$-angle rotation that benefits from the highest available Rabi frequency.

The final level sets chosen for all Qudit-Ramsey runs—together with their logical labels—are collected in Figure 5.2 and serve as the basis for the experimental results presented in the next section.

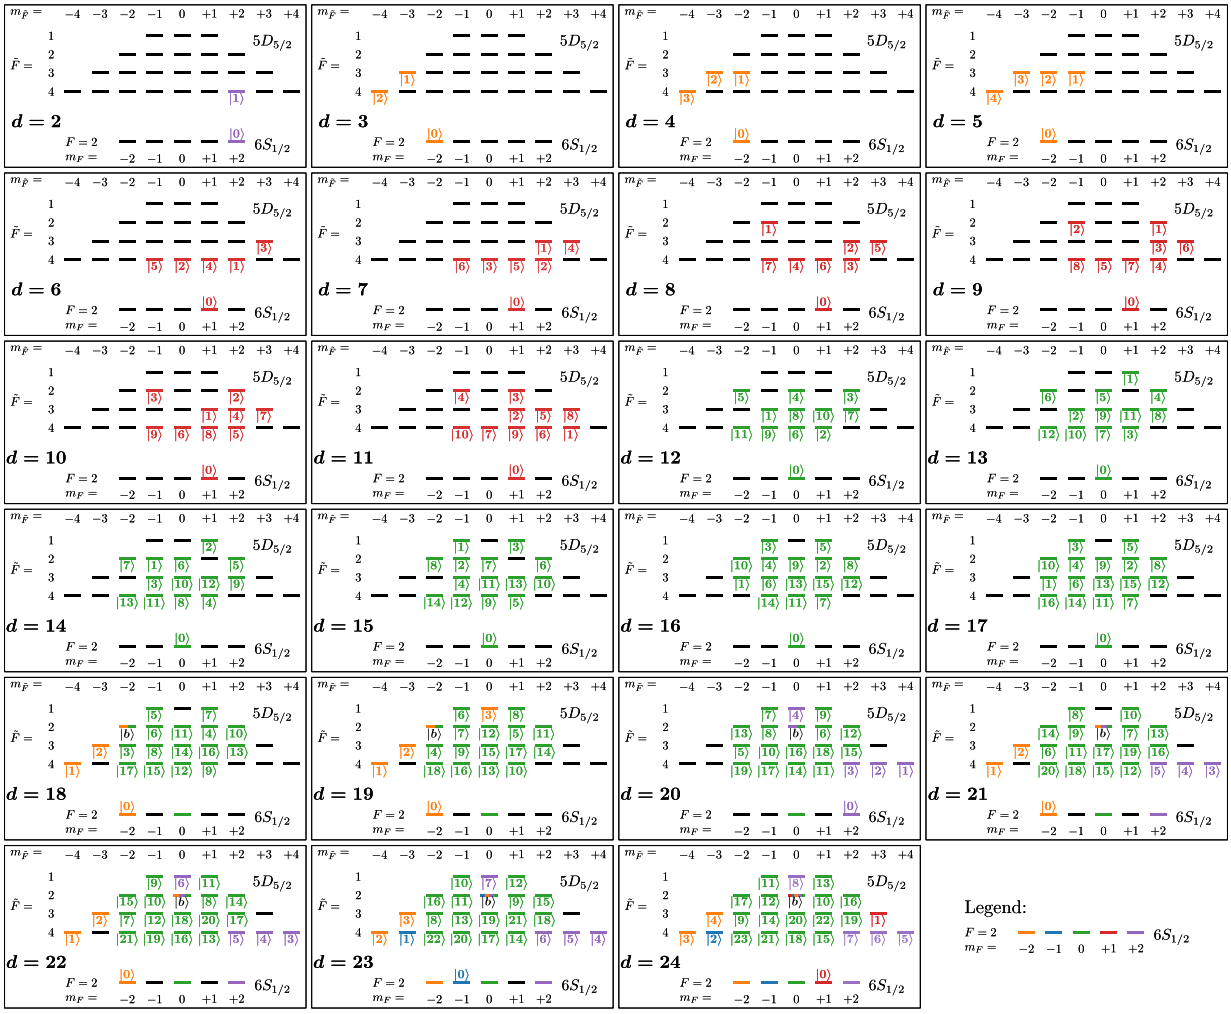


Figure 5.2: **Qudit state encoding for all dimensions.** The plot above illustrates the qudit state encoding for dimensions $d \in \{2, \ldots, 24\}$. Colors indicate the transitions used to address each state in $5D_{5/2}$. Starting from $d = 18$, bus states ($|b\rangle$) are introduced to access additional levels. This requires temporarily shifting the population of $|0\rangle$ to a different $6S_{1/2}$, $F = 2$ state mid-operation to facilitate connectivity.

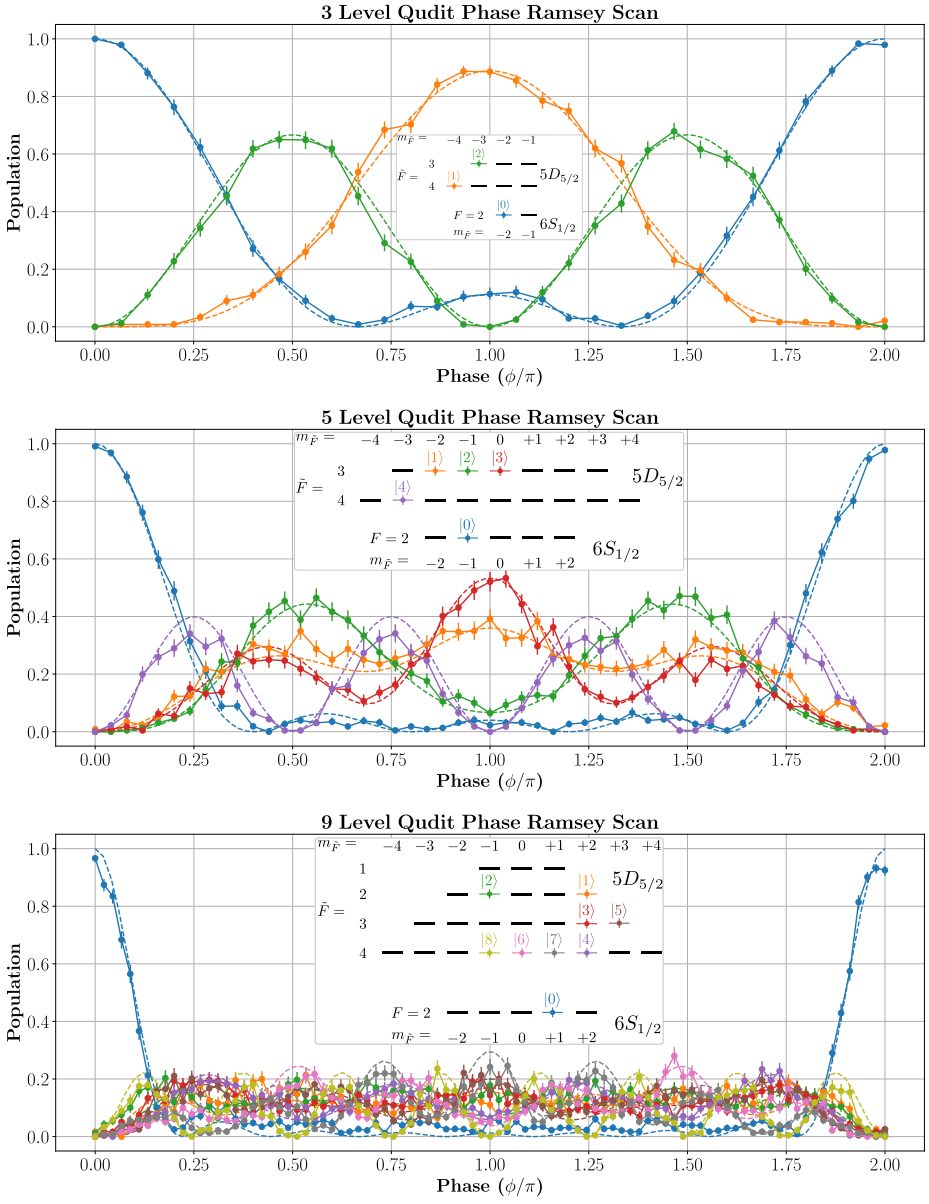


Figure 5.3: **Qudit Ramsey phase scans.** Experimental Ramsey phase scans (solid lines with markers) alongside ideal-case simulation curves (dashed lines) for $d = 3$, 5, and 9. For $d = 3$ and $d = 5$, the experimental data closely follow the theoretical predictions, demonstrating agreement (within $\sim 2\%$). In the $d = 9$ case, the increased number of levels results in significant overlap among population curves, making individual features more difficult to distinguish visually. The legend not only indicates the color-coded mapping between curves and states, but also provides the state encoding used in the experiment.

5.3 Qudit Ramsey Experiments — Fringe contrast vs. Qudit Dimensions

Using the state selection and encoding described in the previous section, we constructed registers of size $d = 2$ through $d = 24$. We stop one level short of the full 25-state manifold because one $D_{5/2}$ level is reserved as a bus for shuttling the $|0\rangle$ population between the different $S_{1/2}(F = 2)$ hubs.

Figure 5.3 plots the complete Ramsey phase scans for $d = 3, 5$, and 9. Solid lines represent the measured populations, dashed lines represent the ideal scans predicted by the theoretical model, and the legend shows the physical–to–logical mapping used in each case. During each scan, the transition frequencies were recalibrated after every five phase points, and each data point is the average of 250 experimental shots. For $d = 3$ and $d = 5$ every branch tracks the theoretical curve very closely (within $\sim$2%). At $d = 9$ the panel becomes too crowded to read individual $|1\rangle \dots |8\rangle$ traces—and the resolution of experimental data only worsens for larger d — so we adopt the $|0\rangle$ fringe as a single, unambiguous performance metric.

Inspection of the $d = 3$ and other dimensions in theory shows that $P_{|0\rangle}$ (population of $|0\rangle$) should be unity at $\phi = 0$ and 2π, and drops to zero at the destructive-interference phases of $2\pi l/d$ for $l \in \mathbb{Z}^+ - \{d\}$ (these phases are $\phi = 2\pi/3$ and $4\pi/3$ for $d = 3$ as seen in the phase scan). We therefore define the *Qudit-Ramsey contrast*

$$C_d = \begin{cases} P_{|0\rangle}(\phi = 0) \;-\; P_{|0\rangle}\big(\phi = \pi(d-1)/d\big) & , \text{ odd } d \\ P_{|0\rangle}(\phi = 0) \;-\; P_{|0\rangle}(\phi = \pi) & , \text{ even } d \end{cases} \tag{5.4}$$

(the phases between which the contrast is measured are shown in figure 5.1 with red arrows) and use this metric to summarise how the system's coherence evolves with the Hilbert-space dimension. A coherent system will have a contrast of 1, and as the system decoheres, the contrast will start trending towards 0. Figure 5.4 shows the measured contrast for each dimension of the qudit up to 24. Each data point represents an average of over 10,000 experimental shots (acquired in ten sets of 1,000), with a full frequency calibration performed immediately before each 1,000-shot batch; the shaded region shows the corresponding $\pm 1\sigma$ statistical uncertainty.

Key observations Qudit Ramsey Contrast vs. Qudit Dimensions (Figure 5.4)

1. **Star Topology vs. Bus Transitions** For $d \leq 17$ every register forms a star-topology connectivity graph, so the increase in implementation time for Ramsey

sequence is about 10s of micro seconds (due to small pulse angles and qudit encoding). Contrast, therefore, decreases very slowly for this dimensional regime. The sudden drop beyond $d = 17$ coincides with the first register that must shuttle $|0\rangle$ through a $5D_{5/2}$ bus to a different $6S_{1/2}$ level; increasing overall implementation time by over a $400\,\mu$s, this is a big jump and hence we see a big decline in contrast.

2. **AC-line crossover.** The simulated error curves reveal that this longer sequence makes the $60 + 180\,$Hz line component (red triangles) dominate from $d \gtrsim 18$ onward, fully accounting for the steep decline. Below that point the AC line noise contribution is four–five orders of magnitude smaller and essentially negligible.

3. **Laser noise contribution.** The laser noise contribution (blue squares) increases very slowly; adding more $D_{5/2}$ states does not enhance dephasing greatly because the common-mode laser phase cancels inside the $D_{5/2}$ manifold (see Chapter 6). Laser noise, therefore, governs only the very lowest dimensions and would be the limiting factor in our fidelity if the AC line signal were compensated for.

4. **Pulse-angle error and laser noise at low d.** For $d \leq 9$ the dominant terms are pulse-angle uncertainty (purple ×) and residual laser-frequency noise. We see that the pulse angle matters less and less as we go to higher dimensions; the reason for this hasn't been investigated yet.

5. **Magnetic-field contribution.** The magnetic-field curve (orange circles) increases gradually as additional $5D_{5/2}$ levels—each with its own Zeeman shift—enter the register, yet it never overtakes the laser-noise floor, showing that permanent magnets are a great choice for magnetic field source for B-field sensitive experiments.

6. **Model agreement.** The total simulated error (red dashed line) tracks the experimental points within the displayed $\pm 1\sigma$ band for every dimension. This agreement is especially impressive, given that the simulated model has no free parameters, indicating that we understand the noise sources in our system very well and know exactly which noise sources to tackle to improve the contrast.

These trends show that different mechanisms limit coherence in different dimensional regimes: laser stability and pulse calibration at low d, and AC-mains pickup once the star topology is broken. Mitigating the line noise by compensating for the phases and detunings imparted (dampening it) should give us the largest coherence gains. The details about each noise source in Figure 5.4 and the simulation method are given in Chapter 6.

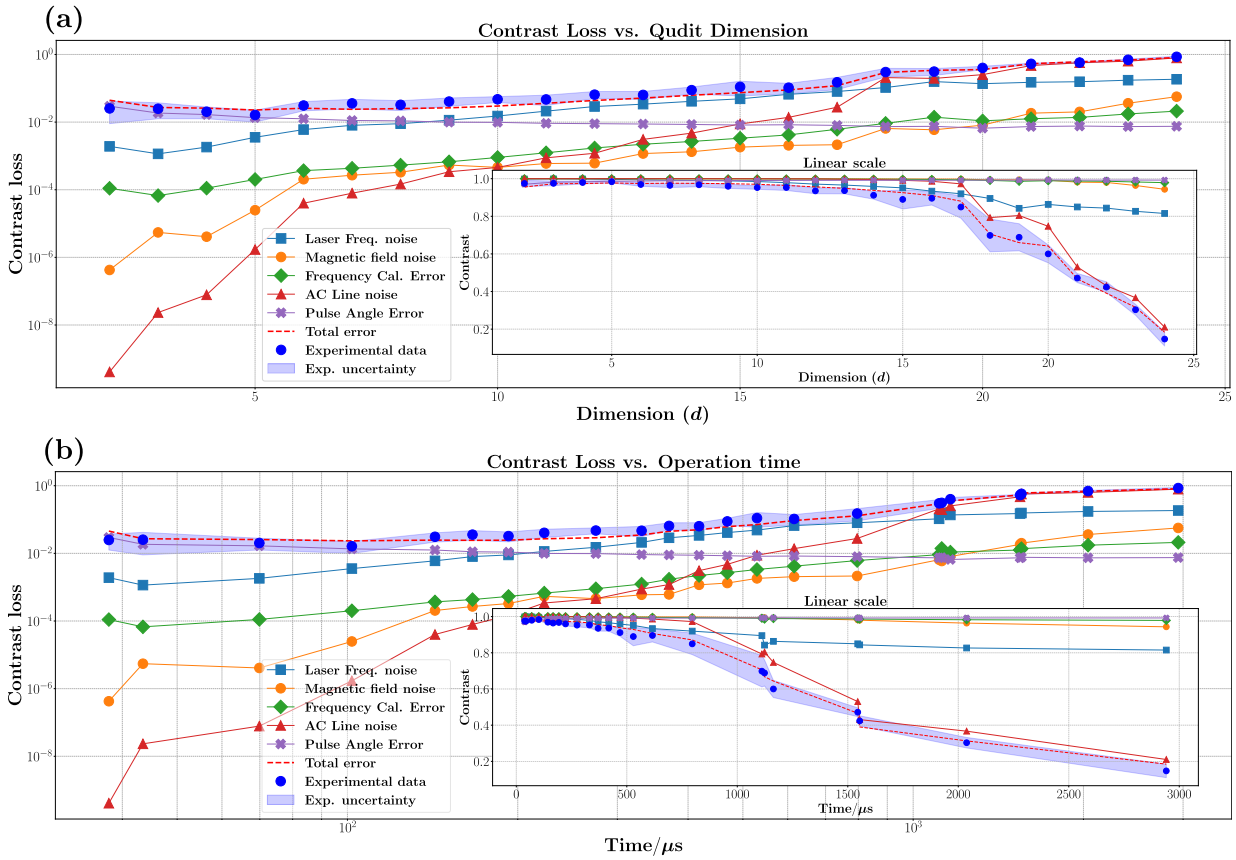


Figure 5.4: **Qudit Ramsey Contrast vs. Dimension and Operation time.** Qudit Ramsey contrast of the $|0\rangle$ population at the phases defined in Equation (5.4), evaluated for each qudit dimension up to $d = 24$.**(a)** main plot shows contrast loss as a function of dimension on log scale, with individual contributions from each noise channel and the total simulated contrast. The inset (linear scale) displays the corresponding measured contrast values. **(b)** presents the same contrast data plotted against the total operation time required to implement the Ramsey sequence for each dimension. In both panels, shaded regions represent one standard deviation (1σ) uncertainty from experimental measurements.

Chapter 6

Error characterization

The simulations in the previous chapter matched the measured contrast–loss data only after five distinct noise channels were incorporated. This chapter presents the experimental evidence supporting each of these channels. For every noise source, we run a tailored variant of Ramsey interferometry (except for pulse angle error, where we use Rabi flops), extract the relevant noise parameters, and feed them into the common error model. In this way the chapter both validates the input parameters of the simulation and showcases Ramsey interferometry as a versatile diagnostic tool for trapped-ion hardware. The error sources considered are

1. **Magnetic-field noise**—fast ambient fluctuations in B-field that modulate the Zeeman splitting during our experiments.
2. **AC-mains line noise**— a combination of 60 Hz and its third harmonic 180 Hz magnetic fields from laboratory mains infrastructure that couple through the electronics around the trap.
3. **Laser noise**—frequency and phase fluctuations of the 1762 nm quadrupole beam, modeled as a Voigt profile. We don't have a working Fiber noise canceling system.
4. **Frequency-calibration error**—residual detuning that remains after the two-point Ramsey calibration routine.
5. **Pulse-angle error**—uncertainty in the Rabi frequency and pulse angle for individual rotations.

Each section that follows is organized in the same way: first the experimental method used to isolate the noise source; then the quantitative results expressed as decoherence times or error distributions; and finally a brief discussion of how the extracted parameters feed into the global error budget used in Chapter **??**. Together, these data form a self-consistent picture of the trapped-ion environment and provide the numerical backbone for all predictive simulations in this thesis.

6.1 Magnetic field noise

Our goal is to measure the *magnetic* contribution to decoherence without contamination from laser–frequency noise. We achieve this by coupling the qubit pair $\{|0\rangle, |1\rangle\}$ through a single bus state (a state that couples to both $|0\rangle$ and $|1\rangle$ through a single quadrupole transition) $|b\rangle$ that sits in a different fine–structure manifold and by doing a "bussed" qubit Ramsey experiment. Because the transition coupling both the computational states are referenced to the *same* $|b\rangle$ the dephasing between $|0\rangle$ and $|1\rangle$ due to noise in the shared laser is equal and cancels out. Below is the pulse sequence of the bussed Ramsey experiment and algebraic proof of the laser noise cancellation phenomenon.

6.1.1 Bussed Ramsey sequence and Laser noise cancellation

First, the ion is prepared in the computational stat $|0\rangle$. We then apply (i) a resonant $\pi/2$ pulse on the $|0\rangle \leftrightarrow |b\rangle$ branch followed immediately by (ii) a resonant π pulse on the $|1\rangle \leftrightarrow |b\rangle$ branch. The corresponding operators can be calculated using equation (4.15), for simplicity assume that transitions are driven very close to resonance, so the unitary becomes

$$U(t) = \cos\left(\frac{\Omega t}{2}\right)\hat{\mathbb{I}} - i\sin\left(\frac{\Omega t}{2}\right)((\cos\phi)\sigma_x - (\sin\phi)\sigma_y) \tag{6.1}$$

where, $\Omega t = \theta$ (pulse angle), and ϕ is the phase imparted on the oscillator from EOM. So for the operators corresponding to the two rotations for bussed Ramseys are

$$\begin{aligned} R_{0b}\Big(\theta = \tfrac{\pi}{2}, \phi = \pi\Big) &= \frac{1}{\sqrt{2}}\Big(|0\rangle\langle 0| + |b\rangle\langle b| + i\ |b\rangle\langle 0| + i\ |0\rangle\langle b|\Big), \\ R_{1b}(\theta = \pi, \phi = 0) &= -i\big(|1\rangle\langle b| + |b\rangle\langle 1|\big). \end{aligned} \tag{6.2}$$

these acting on $|0\rangle$ gives us:

$$R_{0b}\Big(\tfrac{\pi}{2},\pi\Big)\,|0\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle + i\,|b\rangle\big), \tag{6.3}$$

$$R_{1b}(\pi,0)\,R_{0b}\Big(\tfrac{\pi}{2},\pi\Big)\,|0\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big) \;\equiv\; |\psi_1\rangle\,. \tag{6.4}$$

Thus after the first pulse sequence we have a equal super position between computational states and $|b\rangle$ is empty.Now, the system is allowed to evolve freely for a time τ. During this interval each computational state sees a time-dependent detuning

$$\delta_j(t) = \underbrace{\delta_L(t)}_{\text{quasi-static laser noise}} + \underbrace{\kappa_j\,\delta_B(t)}_{\text{magnetic-field drift}}, \qquad j \in \{0,1\}, \tag{6.5}$$

measured *with respect to the common bus state* $|b\rangle$. Because both transitions are driven by side-bands carved from the *same* 1762 nm carrier, the laser term $\delta_L(t)$ is *identical* for $|0\rangle$ and $|1\rangle$. The evolution operator is

$$U_{\text{free}}(\tau) = e^{-i\int_\tau \delta_0(t')dt'}\,|0\rangle\langle 0| + e^{-i\int_\tau \delta_1(t')dt'}\,|1\rangle\langle 1|\,. \tag{6.6}$$

Applying U_{free} to $|\psi_1\rangle$ and discarding the global phase gives

$$|\psi_2(\tau)\rangle = \frac{1}{\sqrt{2}}\Big(|0\rangle + e^{-i\Phi(\tau)}\,|1\rangle\Big), \qquad \Phi(\tau) = \int_\tau \big[\delta_1(t') - \delta_0(t')\big]\,dt'. \tag{6.7}$$

After the free-evolution interval we reverse the first two pulses and flip the phases. These operator's action on $|\psi_2\rangle$ gives

$$R_{1b}\big(\pi,\pi\big)\,|\psi_2\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle + i\,e^{-i\Phi}\,|b\rangle\big)\,, \tag{6.8}$$

$$R_{0b}\Big(\tfrac{\pi}{2},0\Big)R_{1b}\big(\pi,\pi\big)\,|\psi_2\rangle = \frac{1+e^{-i\Phi}}{2}\,|0\rangle - \frac{i\big(1-e^{-i\Phi}\big)}{2}\,|b\rangle\,. \tag{6.9}$$

Hence the population that returns to the computational state $|0\rangle$ is

$$P_{|0\rangle}(\tau) = \left|\tfrac{1+e^{-i\Phi(\tau)}}{2}\right|^2 = \frac{1}{2}\big[1+\cos\Phi(\tau)\big],$$

while all remaining amplitude is parked in the bus state $|b\rangle$.

In the quasi-static limit relevant here the optical detuning is effectively constant during a single shot, so it contributes the *same* phase to both states and cancels in their difference.

$$\Phi(\tau) = \int_{\tau} \big[\delta_1(t') - \delta_0(t')\big]\, dt' = (\kappa_1 - \kappa_0) \int_{\tau} \delta_B(t')\, dt'. \tag{6.10}$$

Because $\Phi(\tau)$ contains only the differential Zeeman term $(\kappa_1 - \kappa_0) \int_\tau \delta_B(t')\, dt'$, the decay of $P_{|0\rangle}(\tau)$ versus τ directly yields the pure magnetic-field dephasing time.

To record the Ramsey fringe we sweep the oscillator phase of the final $0 \leftrightarrow b$ $\pi/2$ pulse from 0 to 2π while keeping the wait time τ fixed; this yields a sinusoidal trace whose peak-to-trough amplitude defines the fringe contrast for that particular τ. The entire phase scan is then repeated at a series of longer wait times, so we obtain a set of contrasts $C(\tau)$. The way $C(\tau)$ decays encodes the characteristic function of the magnetic-field noise; from it we extract a dephasing time $T_{2,B}^*$ for the chosen pair of computational states. Finally, we rerun the complete bussed-Ramsey measurement for several state pairs having different field sensitivities $\Delta\kappa = \kappa_1 - \kappa_0$, full chart of relative sensitivities between the 29 levels is shown in Appendix F. Plotting the resulting $T_{2,B}^*$ values against $\Delta\kappa$ allows us to infer the standard deviation of magnetic-field fluctuations in the trap—providing a quantitative benchmark of the ambient field noise that limits multi-level coherence.

6.1.2 Experimental details

Because the effects of laser dephasing are canceled out, the bussed pairs should display coherence times substantially longer than those obtained from conventional Ramsey on an optical qubit. Generating the long sequences needed to test this limit is, however, constrained by the finite memory in our AWG: a single waveform that includes all four control pulses plus a millisecond-scale dark time would exceed the available buffer. We therefore break the protocol into three separate waveforms:

1. The superposition pulses ($\pi/2$ on $|0\rangle \leftrightarrow |b\rangle$ followed by π on $|1\rangle \leftrightarrow |b\rangle$),
2. a "dummy" waveform that outputs zero RF during the wait time τ, and
3. the reversal pulses (π on $|1\rangle \leftrightarrow |b\rangle$ followed by $\pi/2$ on $|0\rangle \leftrightarrow |b\rangle$).

A digital-I/O line selects each waveform in turn, so τ is set by the spacing between two TTL triggers rather than by a pre-compiled block of zeros inside a waveform (like for

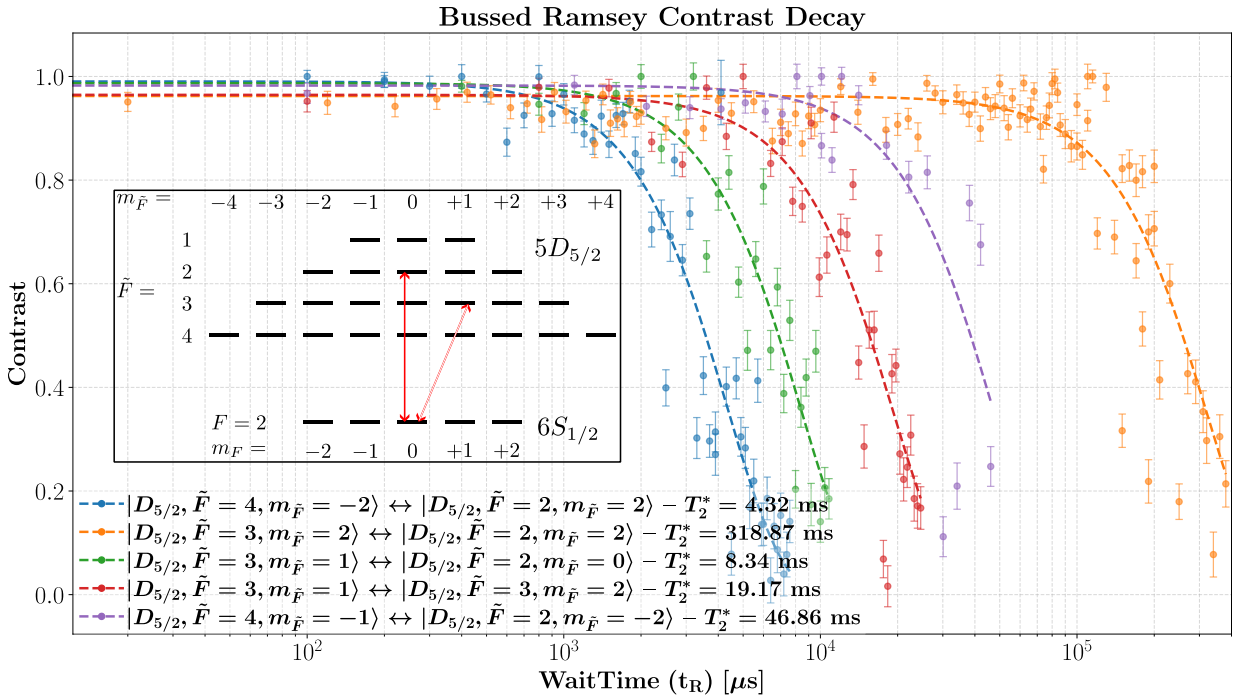


Figure 6.1: **Bussed Ramsey Contrast vs. Ramsey Wait Time.** The figure shows how the contrast of Bussed Ramsey phase scans decays as the Ramsey wait time increases. Contrast curves are plotted for five different pairs of states in the $5D_{5/2}$ manifold, as indicated in the legend. The accompanying schematic illustrates that the computational states are not directly coupled; instead, population transfer occurs via an intermediate bus state.

Ramsey done in frequency calibration). This arrangement causes the absolute fringe phase to drift by an unknown but *shot-to-shot constant* offset so this is of no consequence: for every chosen τ we sweep the phase of the final $\pi/2$ pulse from 0 to 2π, record the resulting sinusoid, and take its peak-to-trough amplitude as the fringe contrast.

Appendix G gives a representative phase scan for a pair of states within the $5D_{5/2}$ manifold. From the peak–to–trough amplitudes of those fringes we obtain the contrast $C(\tau)$ at each wait time τ (plotted in Figure 6.1); fitting the decay yields the magnetic-limited coherence time $T^*_{2,B}$ for the relative Zeeman sensitivity $\Delta\kappa$ of that particular state pair. The fit uses the characteristic function of a Voigt noise profile—the convolution of Lorentzian (full width half max γ) and Gaussian (standard deviation σ) components—whose time-

domain form is

$$C(\tau) = \exp\left(-\frac{\tau^2}{T_G^2} - \frac{|\tau|}{T_L}\right) \quad \text{where,} \quad T_G = \frac{\sqrt{2}}{2\pi\sigma} \quad \text{and} \quad T_L = \frac{1}{\pi\gamma}. \tag{6.11}$$

In our case we are measuring instantaneous detuning caused by the magnetic-field fluctuations weighted by the relative Zeeman sensitivity, $\Delta\omega(t) = \Delta\kappa\,\delta B(t)$. The rms detuning is therefore $\sigma = \Delta\kappa\,\sigma_B$, where σ_B is the rms field noise.

When the data are fitted with the model above we obtain $T_G \lll T_L$; the Gaussian correlation time is at many orders of magnitude shorter than the Lorentzian one. This implies that the field-induced detuning is dominated by the Gaussian component ($T_2^* \approx T_G$). Independent evidence comes from a fast-output magnetometer placed about 2 in from the trap centre: the PDF of its $B_z(t)$ record is well described by a Gaussian profile with no discernible Lorentzian wings. The corresponding noise histogram is plotted in Fig. 6.2.

Finally, the coherence times extracted from each bussed-Ramsey data set are plotted against their corresponding differential sensitivities $\Delta\kappa$. The points follow the expected inverse-trend,

$$T_{2,B}^* = \frac{\sqrt{2}}{2\pi\Delta\kappa\,\sigma_B} \tag{6.12}$$

and a least-squares fit to this expression (dashed line in Fig. 6.2) gives an rms magnetic-field noise. Performing the same analysis for our two field sources shows a clear improvement: with the original electromagnetic coils the fit yields $\sigma_B = 68.1(6)$ μG, whereas replacing them with the permanent-magnet assembly reduces the spread to $\sigma_B = 11.1(1)$ μG—a six-fold suppression of magnetic-field fluctuations at the ion.

6.2 AC Line Noise

Magnetic pickup from the laboratory mains supply is the largest *deterministic* field perturbation seen by the ion. Canada's 120 V distribution operates at 60 Hz, and the wiring that feeds the trap electronics radiates both the fundamental and its harmonics; in our spectra, the third harmonic at 180 Hz is especially prominent. A standard mitigation technique is to *line-trigger* every experimental cycle—i.e. launch the coherent operation sequence at a fixed phase of the 60 Hz waveform. Doing so already extends coherence dramatically, and all routine calibrations and data runs in our lab are phase-locked in this way. Nevertheless, because a typical pulse pattern lasts a finite amount, the qudit still samples a

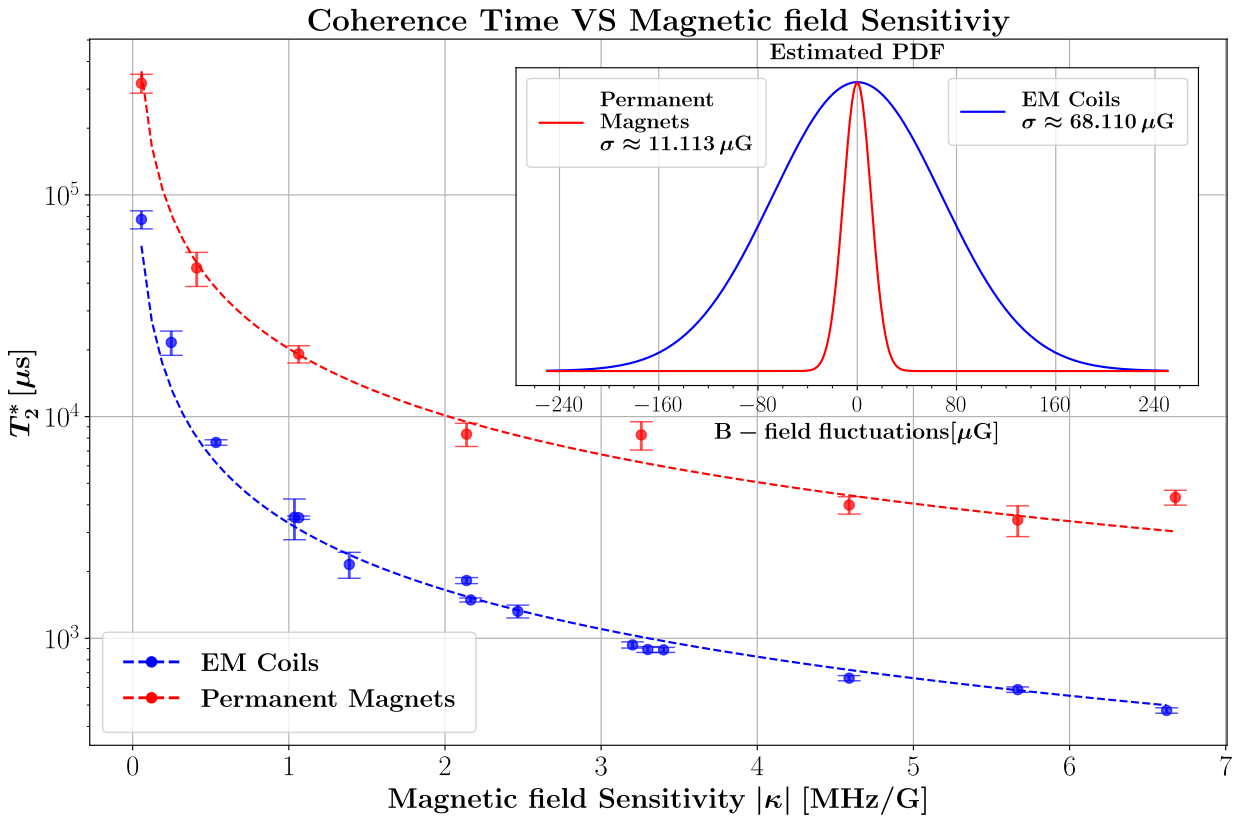


Figure 6.2: **Bussed Ramsey T_2^* vs. Magnetic Field Sensitivity** κ. The plot shows the fitted T_2^* values from Figure 6.1 as a function of magnetic field sensitivity κ between computational states. This relationship allows us to extract the standard deviation of the magnetic field noise. Curves are shown for both the permanent magnet source and the electromagnetic (EM) coil source. The inset compares the extracted magnetic field standard deviations for the two sources, revealing a sixfold reduction in magnetic field noise when switching from EM coils to permanent magnets. The Gaussian envelopes shown in the inset are not normalized and are included solely for visual comparison.

non-negligible fraction of the 60 Hz + 180 Hz period as the sequence plays out, so residual phase slips appear as small detunings and reduce Ramsey contrast.

Although the disturbance is fully deterministic and could, in principle, be pre-compensated pulse-by-pulse, no active compensation techniques have yet been implemented.

We characterize line-synchronous field pickup by repeatedly measuring two reference

transitions: f_0, a magnetically insensitive "clock" transition, and f_1, one of the most sensitive Zeeman transitions available (using the 2-point calibration method described in chapter 4). Their instantaneous frequency difference

$$\Delta f(t) \;=\; f_1(t) - f_0(t) \;=\; (\kappa_1 - \kappa_0)\,\delta B(t) \tag{6.13}$$

provides a direct probe of the field fluctuation $\delta B(t)$, since we already know the sensitivities κ_1 and κ_0 from chapter 2. The measurement protocol is as follows:

1. When the 60 Hz line-trigger fires ($\tau = 0$), we immediately measure the transitions $(f_0(\tau = 0), f_1(\tau = 0))$ and record the pair.
2. After a controlled wait τ, the same two-point calibration is repeated, giving $(f_0(\tau), f_1(\tau))$.
3. Repeat step 1-2 stepping up the delay τ from 0 to ~ 16 ms, mapping the full period of the 60 Hz fundamental (and 3 periods of third harmonic 180 Hz) relative to the line trigger.

The magnetic field change over the interval is

$$\delta B(\tau) = \frac{f_1(\tau) - f_0(\tau) \;-\; f_1(\tau = 0) + f_0(\tau = 0)}{\kappa_1 - \kappa_0}. \tag{6.14}$$

Because the first frequency pair serves as a reference, each $\delta B(\tau)$ data point reflects the deterministic field at a well-defined phase of the 60 Hz+180 Hz waveform, uncontaminated by slower drifts of the absolute field or the laser.

6.2.1 Experimental Results

Figure 6.3 shows the magnetic-field waveform reconstructed from the two-point frequency measurements described above, plotted over a full 16 ms period of the laboratory mains signal. Data for the original electromagnetic (EM) coils and the permanent-magnet field source are overlaid for direct comparison.

The line–synchronised magnetic disturbance is well represented by a dc term plus the first and third harmonics of the mains frequency,

$$B(\tau) = B_0 + A_{60}\sin\!\big(2\pi\cdot 60\,\tau + \phi_{60}\big) + A_{180}\sin\!\big(2\pi\cdot 180\,\tau + \phi_{180}\big), \tag{6.15}$$

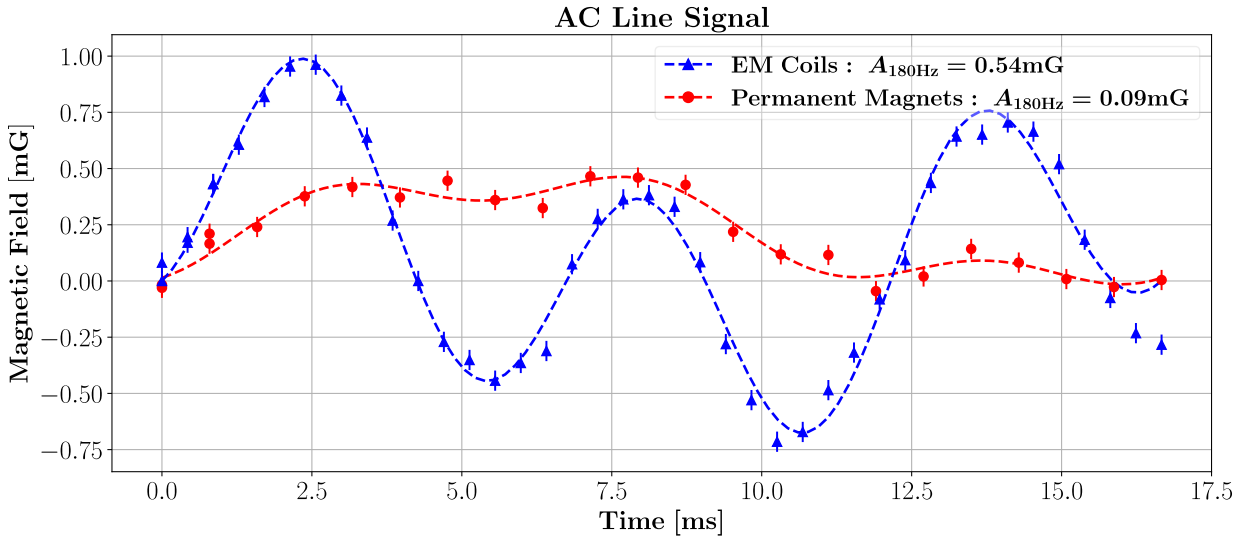


Figure 6.3: **AC Line signal for Permanent magnets and EM-coils.** Magnetic field fluctuations due to the AC line signal, measured using a two-point calibration method. The figure shows both the experimental data and the corresponding best-fit curves, modeled as a combination of 60 Hz and 180 Hz sinusoidal components. Results are shown for two magnetic field sources: a permanent magnet and an electromagnetic coil. Error bars represent one standard deviation (1σ) uncertainty in the measured magnetic field.

where A_{60} and A_{180} are the amplitudes of the 60 Hz fundamental and its 180 Hz third harmonic, ϕ_{60} and ϕ_{180} are their phases at the line trigger point, and B_0 is the static field offset. Least-squares fits of Eq. (6.15) to the data in Fig. 6.3 yield the parameters listed in Table 6.1. As anticipated, replacing the coils with Permanent magnets suppresses both harmonic amplitudes and reduces the static offset, confirming that the coil drive electronics were a major contributor to line-synchronous magnetic noise.

6.3 Laser Noise

Laser-induced dephasing is isolated by probing the *clock* transition between $\left|S_{1/2}, F = 2, m_F = 0\right\rangle$ and $\left|D_{5/2}, F = 2, m_F = 0\right\rangle$, which has magnetic field sensitivity of 40 Hz/G (at operating field of 4.2 G), making it first-order magnetic field insensitive. With magnetic-field noise suppressed, the Ramsey coherence time T_2^* obtained on this transition is therefore a direct

Table 6.1: Fitted coefficients of Eq. (6.15) for the electromagnetic-coil and permanent-magnet field sources. Amplitudes are given in mG; phases (in radians) are referenced to the laboratory 60 Hz trigger.

Parameter	**EM coils**	**Permanent magnets**
A_{60} (60 Hz amplitude)	0.22(2)	0.36(1)
ϕ_{60} (60 Hz phase)	-0.56(6)	1.30(3)
A_{180} (180 Hz amplitude)	0.09(1)	0.54(1)
ϕ_{180} (180 Hz phase)	-1.4(2)	-1.20(2)
B_0 (static offset)	0.22(1)	0.000(1)

measure of the laser's frequency-phase stability. The Ramsey protocol for this characterization is as follows:

1. The ion prepared in state $|0\rangle \equiv \left|S_{1/2}, F = 2, m_F = 0\right\rangle$, using NBOP method.
2. A resonant $\pi/2$ pulse on the clock transition transfers half the population to $|1\rangle \equiv \left|D_{5/2}, F = 2, m_F = 0\right\rangle$, creating the superposition $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.
3. The superposition evolves freely for a wait time τ.
4. A second $\pi/2$ pulse, phase-shifted by ϕ, converts accumulated phase into population difference.
5. The phase ϕ is scanned from 0 to 2π to record a Ramsey fringe for that particular τ.

The peak-to-trough amplitude of each fringe gives the contrast $C(\tau)$. Repeating the phase scan for progressively longer wait times yields $C(\tau)$ over the range of interest. The decay is fitted with the Voigt characteristic function introduced in Eq.(6.11):

6.3.1 Experimental results

Figure 6.4 displays a representative Ramsey fringe for the clock transition and the resulting contrast-versus-delay curve $C(\tau)$. A least-squares fit of the Voigt characteristic function yields $T_G = 1521(77)\,\text{ms}$ and $T_L = 3038(471)\,\text{ms}$, and overall $T_2^* = 1187(64)\,\mu\text{s}$. The

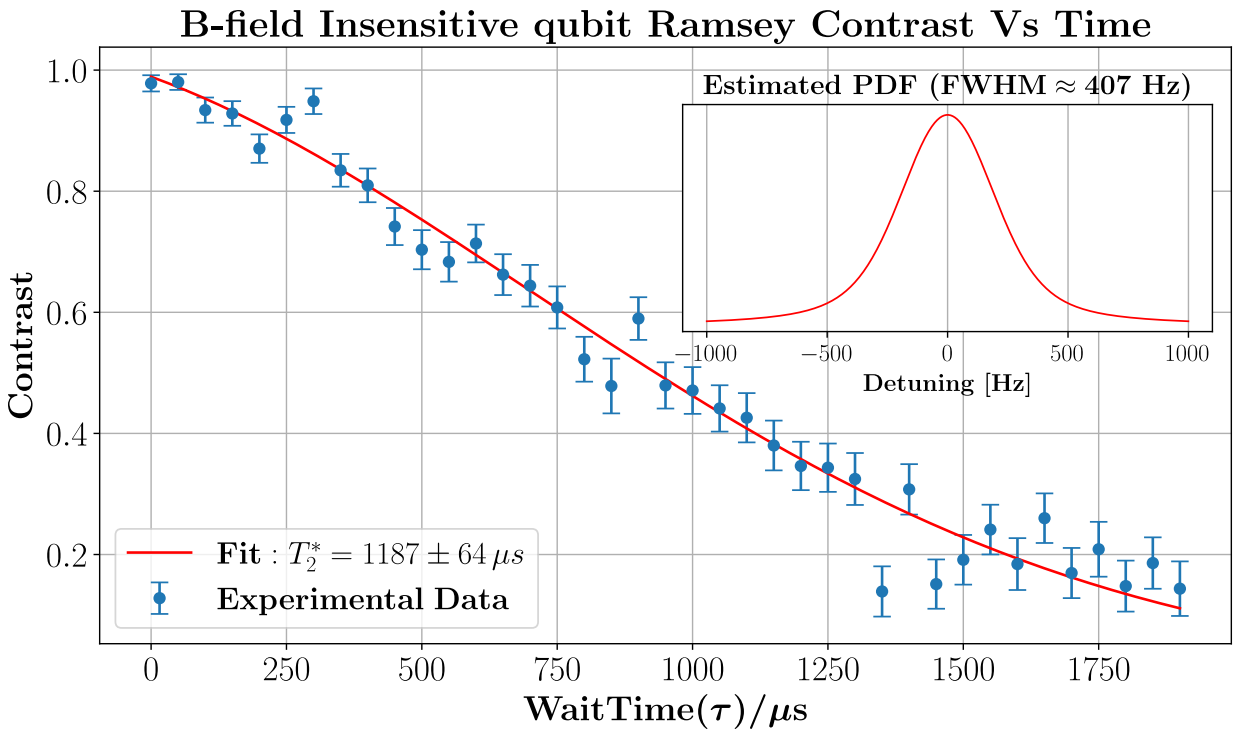


Figure 6.4: **Laser Noise Characteristic Function and Estimated PDF.** The figure shows the Ramsey contrast as a function of wait time (τ) for the most magnetically insensitive transition. The contrast data is fitted using the Voigt characteristic function, from which the probability density function (PDF) of the laser frequency noise is estimated. The inset displays the resulting PDF, illustrating the underlying noise profile.

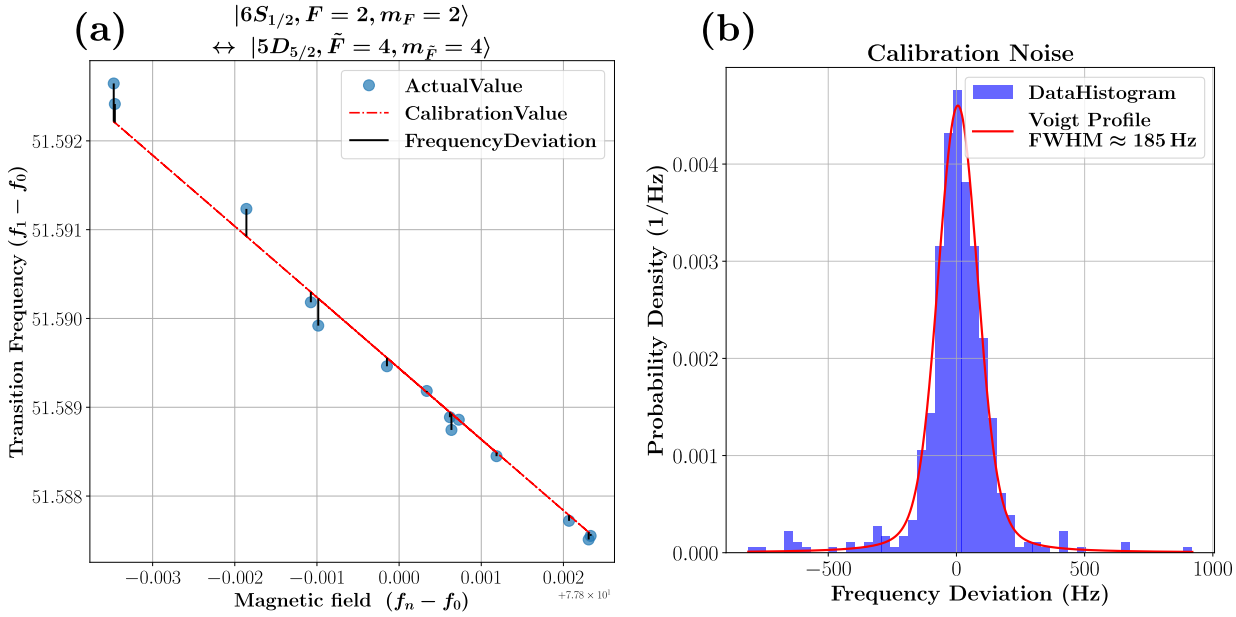


Figure 6.5: **Calibration Noise. (a)** An example comparison between the calibrated transition frequency and the corresponding measured value. The difference between these two defines the frequency deviation attributed to calibration noise.**(b)** A histogram of measured frequency deviations collected across multiple calibrations. The distribution is fitted with a Voigt profile to capture both Gaussian and Lorentzian components.

inset in Figure 6.4 plots the corresponding Voigt profile shape, constructed from the fitted parameters. Its full width at half-maximum (FWHM) is

$$\mathrm{FWHM_V} \simeq 0.5346\,\Gamma_L + \sqrt{0.2166\,\Gamma_L^2 + \Gamma_G^2}, \quad \Gamma_L = \frac{1}{\pi T_L}, \quad \Gamma_G = 2\sigma\sqrt{2\ln 2} = \frac{2\sqrt{\ln 2}}{\pi T_G}, \tag{6.16}$$

giving $\mathrm{FWHM_V} = 407(20)$ Hz [56],[57].

Laser frequency noise is often idealized as purely Lorentzian, yet here the best-fit values indicate Gaussian and Lorentzian contributions of comparable magnitude. Both components must therefore be retained for an accurate prediction of laser-limited coherence times elsewhere in this thesis.

6.4 Frequency calibration Error

The two–point frequency calibration introduced in Chapter 4 predicts the resonance of any target transition from a measurement of the reference pair (f_0, f_1). To quantify the residual error of this procedure, we analyze the calibration data itself. For every shot of calibration data set described in Chapter 4 data we

1. use the measured (f_0, f_1) to compute the *predicted* transition frequency f_{pred} for the line of interest, and
2. record the difference $\Delta f = f_{\text{meas}} - f_{\text{pred}}$ between this prediction and the frequency actually obtained in that shot.

The collection of all Δf values is then binned into a histogram, which shows the empirical probability density function (PDF) of the calibration error. Unlike the calibration uncertainty from a single measurement, this PDF reflects how uncertainties in the two references propagate through the calibration model to the target frequency.

Figure 6.5 displays this empirical PDF for the permanent-magnet data set. The histogram is well-described by a Voigt profile (this profile is expected, given that we observed laser noise to be a Voigt profile); the best-fit full width at half-maximum is 185(1) Hz. This value represents the combined, shot-to-shot frequency calibration error that must be accounted for in any realistic simulation of qudit operations.

6.5 Pulse-angle error

In Chapter 4, Section 4.4, we described the procedure used to calibrate the pulse angles (i.e., effective Rabi frequencies and π-times) for all quadrupole transitions in the system. This calibration is based on measurements from only five reference transitions, from which the pulse angles of all other transitions are inferred using theoretical calculations of their relative transition strengths.

To evaluate the accuracy of this calibration method, we now compare the measured pulse angles (obtained from Rabi oscillation scans) to the calculated values from calibrations. Specifically, we analyze the relative transition strengths—that is, the pulse angles normalized to their corresponding reference transitions. These comparisons are shown in Figure 6.6, where open circles represent the measured relative strengths, and the overlaid bars indicate the predicted values from our calibration.

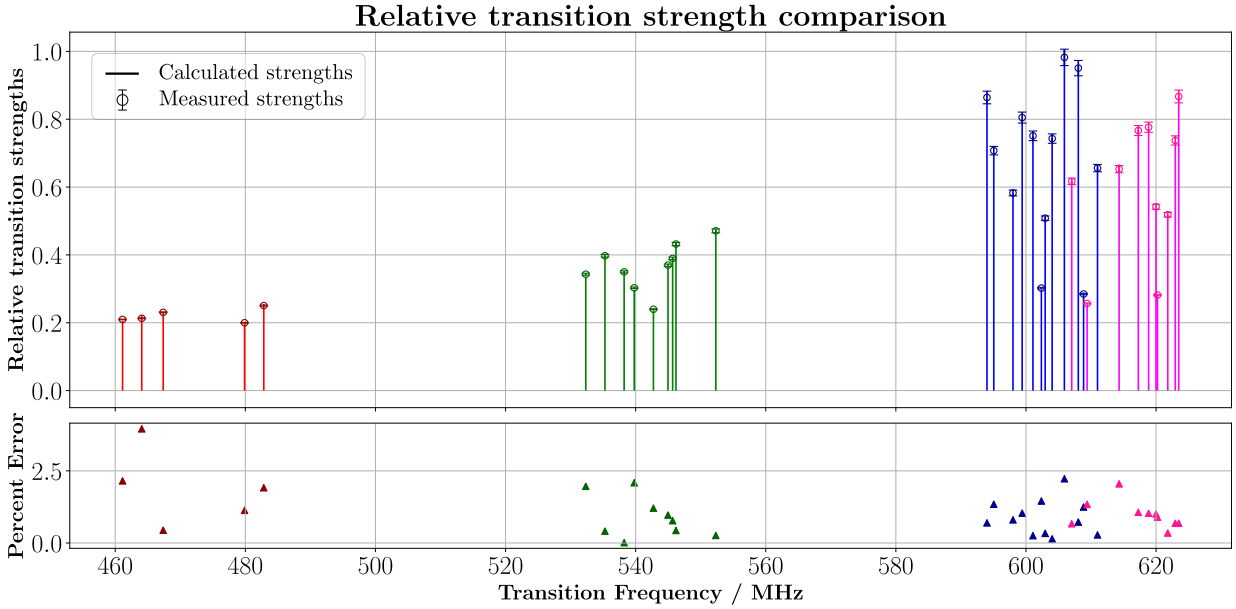


Figure 6.6: **Relative transition strength comparison between calibrated theory and measurements.**The vertical lines in the top panel represent the calculated relative transition strengths (with respect to their corresponding reference transitions), obtained from Chapter 2, Section 2.3, and rescaled using Eq. (4.47). The hollow circles indicate the experimentally measured relative transition strengths. The error bars correspond to one-sigma uncertainties in the measured values. The bottom panel displays the percent deviation between the theoretical predictions and experimental results.

.

The lower panel of the figure displays the percent error between measurement and prediction for each transition. Across the set of transitions considered, the mean absolute percent error, denoted δ_τ, is found to be 1.77 %. For context, this error is significantly larger than the typical uncertainty of 0.23 % associated with directly fitting Rabi oscillations to extract frequency. In fact, we directly measure slow drifts in calibrated pulse durations over a time scale of several hours and find a relative drift of $\delta_d = 2.61\,\%$, this could be due to slow drifts in laser intensity and beam pointing, both of which can cause fluctuations in the effective Rabi frequency during the measurement sequence. Taken together ($\delta_p = 4.38\,\%$), these results suggest that while the theory-based calibration scheme provides reasonably accurate estimates of relative transition strengths, its accuracy is fundamentally limited by experimental instabilities in laser amplitude and alignment as well

as theoretical calculations.

6.6 Simulation Model

We model the experiment with a Monte-Carlo procedure. For every "shot" we draw one random value from the probability-density function of each characterized error source, add those contributions to the pulse detunings, propagate the complete Ramsey sequence, and extract the resulting fringe contrast. Repeating this sample-and-propagate cycle a large number of times—exactly analogous to the shot-to-shot averaging in the laboratory—yields an ensemble-averaged contrast for the chosen Ramsey sequence. In this section the details about this simulation model described.

All coherent control in this work is executed with a single 1762 nm quadrupole laser, so *every pulse couples one* $6S_{1/2}$ *level to one* $5D_{5/2}$ *level* and no two transitions are driven in parallel. This one-at-a-time strategy converts the full-dimensional problem into a succession of independent two-level evolutions, making numerical simulation straightforward.

1. Indexing the Hilbert space. Each physical Zeeman sub-level is assigned a fixed column–row index $n, m \in \{0, \ldots, d-1\}$ in the simulation basis $\{|0\rangle, |1\rangle, \ldots, |d-1\rangle\}$. When a pulse addresses the transition $\left|S_{1/2}, F = 2, m_F\right\rangle \leftrightarrow \left|D_{5/2}, F', m'_F\right\rangle$ the pair (n, m) corresponding to those two states is passed to the propagator.

2. Two-level Hamiltonian and detuning sampling With only one transition illuminated, the system Hamiltonian reduces to

$$\hat{H}_p^{(n,m)} = \frac{\Omega'}{2}\big(e^{+i\phi_{n,m}}\,|m\rangle\langle n| + e^{-i\phi_{n,m}}\,|n\rangle\langle m|\big) + \Delta n, m\big(|m\rangle\langle m| - |n\rangle\langle n|\big), \tag{6.17}$$

where Ω' is the Rabi rate after accounting for a shot-constant pulse angle error:

$$\Omega \;=\; \frac{\Omega_{n,m}}{(1+\delta_p)}. \tag{6.18}$$

and Ω_{nm} is the correct Rabi frequency of states corresponding to n and m, $\phi_{n,m}$ and $\Delta_{n,m}$ are the programmed phase and the instantaneous detuning given by

$$\Delta_{n,m} = 2\pi\big[\kappa_{n,m}\big(\delta B + B_{\text{line}}(t)\big) + \delta\nu_L + \delta\nu_f\big]\,, \tag{6.19}$$

$$\phi_{n,m} = \phi_{\text{Ramsey}}^{(n,m)} + 2\pi t\big[\kappa_{n,m}\,\delta B + \delta\nu_L + \delta\nu_f\big] + 2\pi\kappa_{n,m}\int_0^t B_{\text{line}}(t')\,dt', \tag{6.20}$$

where t is the pulse start time relative to the line trigger. And the following table tells us what each of the remaining notations represent the individual noise samples from their respective profiles discussed in the earlier sections. The notation mapping to each noise source as well as their profiles and FWHM are summarized in the table below:

Noise/Error Source	Notation	Profile	FWHM
Magnetic Field	δB	Gaussian	26 µG
Laser Frequency	δ_L	Voigt	407 Hz
Calibration Frequency	δ_f	Voigt	185 Hz
AC Line Signal	$B_{\text{line}}(t)$	(6.15)	6.1
Pulse angle	δ_p	Gaussian	4.38 %

Table 6.2: Noise sources, profiles, and distribution widths.

Once the Hamiltonian is fixed, the unitary for each coupling is

$$\hat{U}_{(n,m)}(t_p) = e^{-i\hat{H}_p^{(n,m)}\, t_p}, \tag{6.21}$$

with

$$\begin{aligned} \theta &= 2\sin^{-1}\left(\sqrt{f}\right), \\ t_p &= \frac{\theta}{\Omega_{n,m}}. \end{aligned} \tag{6.22}$$

These unitaries are calculated separately for every coupling (n, m) in the Ramsey pulse sequence $\mathcal{C}$ and applied sequentially to the state.

$$|\psi_{\text{final}}\rangle = \prod_{(n,m)\in\mathcal{C}} \hat{U}_{(n,m)}(t_p)\ |\psi(0)\rangle\,. \tag{6.23}$$

3. Repetitions and averaging The full procedure is repeated for N simulation shots (typically $N = 1000$). For each shot j we record the binary outcome $p_k^{(j)} \in \{0, 1\}$ of measuring every basis state $|k\rangle$. The reported population of $|k\rangle$ is the sample mean

$$P_k = \frac{1}{N}\sum_{j=1}^{N} p_k^{(j)}. \tag{6.24}$$

The noise characterization and simulation methods developed in this section allow us to *predict the contrast of any operation* carried out on the system and to identify the

dominant error sources for each type of operation. The agreement between prediction and experiment is illustrated in Figure 5.4 (Chapter 5): every simulated data point lies within the corresponding experimental uncertainty band, highlighting the predictive power of the model.

Chapter 7

Virtual Qubits

The preceding chapters of this thesis demonstrated the feasibility of high-dimensional quantum control in a single trapped $^{137}Ba^{+}$ ion. We developed and implemented a 25-level state preparation and measurement (SPAM) protocol with a fidelity exceeding 99.5%, developed calibration and various error characterization techniques, and characterized coherence properties across a range of multi-level superpositions. These results establish a robust platform for exploring quantum information processing beyond the qubit paradigm. With this foundation in place, we now turn to the potential applications of such high-dimensional control. The expanded state space enables a variety of applications, including native multi-level entangling gates with reduced calibration overhead and circuit depth [26], more compact and tolerant quantum error-correcting codes [58], efficient quantum compilation schemes, and direct simulation of complex quantum systems such as high-spin particles and lattice gauge models [27].

One particularly promising application is the encoding of *virtual qubits* within a single qudit. In this approach, 2^n internal levels of a qudit are mapped to the computational basis states of n logical qubits, allowing multiple qubits to be hosted within a single ion. For example, in the $^{137}Ba^{+}$ system used in this work—with 25 distinguishable states—up to four virtual qubits can be encoded in a single ion. This approach not only reduces the number of ions required for quantum computation but also opens the door to novel architectures with reduced overhead in ion-trap arrays, compilation, and gate synthesis.

This chapter builds directly on the calibration, control, initialization, and readout techniques developed in the preceding chapters to explore practical applications of virtual qubit encoding. In particular, we implement the Bernstein–Vazirani algorithm using two and three virtual qubits encoded within a single $^{137}Ba^{+}$ ion. In addition to demonstrating

standard algorithmic primitives, we also showcase the use of virtual qubits to implement operations that are intensive in standard qubit systems. Specifically, we demonstrate a controlled-controlled-controlled-NOT (CCCNOT) gate—also known as a controlled Toffoli gate—within a four-qubit encoding. While this gate typically requires at least five two-qubit gates on standard architectures, the native access to a 25-dimensional Hilbert space enables a more direct and efficient realization using a single Givens rotation. This highlights a key advantage of virtual qubit encodings: certain high-weight operations that are resource-intensive on qubit systems can be natively or more compactly implemented in the qudit formalism.

7.1 Bernstein–Vazirani Algorithm

The Bernstein–Vazirani (BV) algorithm is a foundational quantum algorithm that demonstrates how quantum parallelism can provide exponential speedup in identifying hidden information encoded in a deterministic Boolean function. Given a black-box function $f_s : \{0,1\}^n \to \{0,1\}$, defined as

$$f_s(x) = s \cdot x \mod 2,$$

where $s \in \{0,1\}^n$ is an unknown bit string and $\cdot$ denotes the bitwise inner product modulo 2, the goal is to determine the string s using as few function queries as possible. Classically, determining s requires n queries to f_s. In contrast, the BV algorithm recovers s with a single query using quantum interference.

The algorithm begins by initializing the system in the all-zero computational basis state $|0\rangle^{\otimes n}$. Applying Hadamard gates in parallel on all qubits creates an equal superposition of all computational states:

$$|0\rangle^{\otimes n} \xrightarrow{H^{\otimes n}} \frac{1}{\sqrt{2^n}} \sum_{x=0}^{2^n-1} |x\rangle \,. \tag{7.1}$$

The oracle $\mathcal{O}$ is then applied, introducing a relative phase determined by $f(x)$, i.e.,

$$|x\rangle \to (-1)^{f(x)} |x\rangle = (-1)^{x \cdot s} |x\rangle \,. \tag{7.2}$$

A second Hadamard transform is then applied to interfere the phase-encoded amplitudes. Due to the structure of the phase, the resulting state becomes $|s\rangle$ concentrating all amplitude on the bit string s. Measuring the final state in the computational basis, therefore, reveals the hidden key with a single query to the oracle.

To implement the Bernstein–Vazirani algorithm using virtual qubits encoded in the $6S_{1/2}$ and $5D_{5/2}$levels of a single $^{137}Ba^{+}$ ion, we must first translate the standard quantum circuit into one compatible with the experimental constraints of our system. Unlike conventional architectures where direct qubit–qubit gates and universal single-qubit rotations are assumed, our system restricts coherent control to electric quadrupole transitions between the $S_{1/2}$ and $D_{5/2}$ manifolds. This induces a *star topology* due to our state selection criteria (described in Chapter 6), where all control operations must pass through a common "hub" state in $S_{1/2}$. Transitions between two states in $D_{5/2}$ are mediated indirectly. As a result, even standard gates such as the Hadamard must be decomposed into physically allowed sequences involving the accessible transitions.

7.2 Unitary Decomposition under a Star–Topology Constraint

Based on the state selection and encoding describe din Chapter 5, one computational state is in $6S_{1/2}$manifold which we label as $|0\rangle$ and the remaining $d-1$ states are encoded in $5D_{5/2}$. Since the coherent operations are driven by 1762 nm electric quadrupole transition, the resulting connectivity graph has a star topology with a central node at $|0\rangle$ (this connectivity is enough to give us universal control over all the levels [59]).

This restricted connectivity poses a fundamental challenge for gate decomposition strategies: standard QR-based protocols [60], which rely on unrestricted pairwise Givens rotations, are not directly applicable. To ensure physical implementability within our system, we adapt the decomposition procedure such that all unitary operations are constructed from rotations confined to two-dimensional subspaces of the form $\mathrm{span}\{|0\rangle, |i\rangle\}$, with $i \in \{1, \ldots, d-1\}$. This constraint preserves compatibility with the underlying topology and sets the stage for pulse-level implementations tailored to our control hardware.

Decomposition Strategy

Let U denote the target logical operation expressed in the basis $\{|0\rangle, |1\rangle, \ldots, |d-1\rangle\}$. For any index $i > 0$ we define a *star-topology Givens rotation* that couples the hub state

$|0\rangle$ to $|i\rangle$ as

$$
\begin{aligned}
G_{0,i}(\theta) \;&=\; \exp\Big[-\theta\big(|0\rangle\langle i| - |i\rangle\langle 0|\big)\Big] \\
&= \begin{array}{c} \begin{array}{cccc} \quad|0\rangle & \cdots & \quad|i\rangle & \cdots \end{array} \\ \begin{pmatrix} \cos\theta & 0 & \sin\theta & 0 \\ 0 & \mathbb{I}_{i-2} & 0 & 0 \\ -\sin\theta & 0 & \cos\theta & 0 \\ 0 & 0 & 0 & \mathbb{I}_{d-1-i} \end{pmatrix}_{0,i} \end{array}
\end{aligned} \tag{7.3}
$$

Our goal is to synthesize U as a time-ordered product of such physically realizable pulses:

$$
U \;=\; G_0^\dagger\, G_1^\dagger \cdots G_{N-1}^\dagger, \qquad G_j = G_{0,i_j}(\theta_j),\; i_j \in \{1,\ldots,d-1\}. \tag{7.4}
$$

so that playing the pulses $\{G_{N-1},\ldots,G_0\}$ on the ion implements U exactly within the star connectivity. The algorithm starts by setting a running matrix $V_0 = U$ and pre-multiplies by Givens rotations until the running matrix is the identity, $V_N = I$. Each step preserves one pivot element (leading diagonal entry) while annihilating a single off-diagonal entry. We follow the following steps:

1. **Initial sweep.** Clear the first column by applying $G_{0,i}\big(\theta_i^{(0)}\big)$ for $i = \{1,2,\ldots,d-2,d-1\}$ in some order $\mathcal{O}$.

$$
\theta_i^{(0)} \;=\; \begin{cases} \tan^{-1}\!\left(\dfrac{V_{i0}}{V_{00}}\right), & V_{00} > 0, \\ \tan^{-1}\!\left(\dfrac{V_{i0}}{V_{00}}\right) + \pi, & V_{00} < 0, \\ \frac{\pi}{2}\,\mathrm{sgn}(V_{i0}), & V_{00} = 0, \end{cases}
$$

   When finished, column 0 equals the basis vector $|0\rangle$.

2. **Swap cycle** (for columns $k = 1,\ldots,d-1$ in reversed order $\mathcal{O}$).

   (a) *Swap.* A fixed $\frac{\pi}{2}$ pulse $S_{0k} = G_{0,k}(\frac{\pi}{2})$ exchanges rows 0 and k, moving the would-be diagonal entry into the pivot position V_{0k}.

   (b) *Column elimination.* For each row $i > k$ (processed in order $\mathcal{O}$) apply $G_{0,i}(\theta_{ik})$ to zero V_{ik} while leaving the pivot V_{0k} unchanged.

(c) *Swap-back.* Re-apply S_{0k} to restore the original row ordering.

Repeating the swap cycle for every column drives V to the identity: $G_{N-1}\cdots G_0\,U = I$. Therefore, reversing the list and taking adjoints yields the pulse sequence for U. The routine uses N Givens rotations.

$$N \;=\; (d-1) \;+\; \sum_{k=1}^{d-1}\big[\,2+(d-1-k)\big] \;=\; \frac{(d-1)(d+4)}{2} \tag{7.5}$$

Algorithm 3 Unitary decomposition (Star topology)

Require: $U \in U(d)$, ordering $\mathcal{O} = (1,\ldots,d-1)$, tolerance $\varepsilon = 10^{-12}$

$V \leftarrow U$
$\mathcal{R} \leftarrow [\,]$ ▷ rotation record
▷ — **Initial sweep clears column 0** —
for $i \in \mathrm{reverse}(\mathcal{O})$ **do**
  if $|V_{i0}| > \varepsilon$ **then**
    $\theta \leftarrow \arctan 2(V_{i0},\, V_{00})$
    $G \leftarrow G_{0,i}(\theta)$; $V \leftarrow G\,V$
    $\mathcal{R}$.append(G)
  end if
end for
▷ — **Swap-cycle for each remaining column** —
for $k \in \mathcal{O}$ **do**
  $S \leftarrow G_{0,k}(\frac{\pi}{2})$ ▷ swap rows $0, k$
  $V \leftarrow S\,V$; $\mathcal{R}$.append(S)
  for $i \in \mathrm{reverse}(\mathcal{O}[k{:}])$ **do**
    if $|V_{ik}| > \varepsilon$ **then**
      $\theta \leftarrow \arctan 2(V_{ik},\, V_{0k})$
      $G \leftarrow G_{0,i}(\theta)$; $V \leftarrow G\,V$
      $\mathcal{R}$.append(G)
    end if
  end for
  $V \leftarrow S\,V$; $\mathcal{R}$.append(S)
end for
return $\mathcal{R}$ (sequence of pulses), V

Each entry of $\mathcal{R}$ is a laser pulse on the 1762 nm quadrupole line.A rotation $G_{0,i}(\theta)$ is effected by a resonant pulse of duration $\tau = \theta/\Omega_{0,i}$, where $\Omega_{0,i}$ is the calibrated Rabi frequency for the $|0\rangle \leftrightarrow |i\rangle$ transition.Swap operations S are simply $\frac{\pi}{2}$ pulses on the same link. Concatenating the pulses in $\mathcal{R}$ (time-ordered left to right) implements the desired unitary U exactly within the star topology.

Iterative Sequence Compression

While the above method produces a pulse sequence to implement the exact target unitary, it is rarely optimal in practice. The algorithm yields a pulse list $\left\{ i_m, \theta_m^{(0)}, \phi^{(0)} \right\}_{m=1}^{M}$, here M is total length of sequence. $\theta_m^{(0)}$ is m^{th} rotation angle. And every pulse carries the same phase $\phi^{(0)}$ (a consequence of the fixed sign convention in 7.3). This uniformity lets us fuse back-to-back rotations that address the same transitions $(0, i)$:

$$R_{0,i}\big(\pi\theta_p^{(0)}\big)\, R_{0,i}\big(\pi\theta_{p+1}^{(0)}\big) \;=\; R_{0,i}\big(\pi\,(\theta_p^{(0)} + \theta_{p+1}^{(0)})\big),$$

The new sequence will have $M_0(\leq M)$ pulses.

We then treat the two continuous parameters per pulse - angle of rotation $\theta_m \in [0, 2]$ and phase $\phi_m \in [0, 2]$ - as optimization variables and minimize the Frobenius-norm cost

$$\mathcal{L}(\theta, \boldsymbol{\phi}) \;=\; \left\| U_{\theta,\boldsymbol{\phi}} \;-\; U_{\text{tgt}} \right\|_F^2,$$

using a Limited-memory Broyden-Fletcher-Goldfarb-Shanno Bound (L-BFGS-B) routine that enforces box constraints. Phase updates are implemented experimentally as instantaneous frame-shifts of the local oscillator; they are realized cost-free *virtual Z gates*.

Finally, we perform a sequential pulse-elimination sweep:

1. For each pulse index m, temporarily delete pulse m, re-optimize the reduced parameter set, and evaluate the residual error $\mathcal{L}$
2. if $\mathcal{L} < \epsilon(\simeq 10^{-3})$ accept the deletion and restart the scan, otherwise keep the pulse.

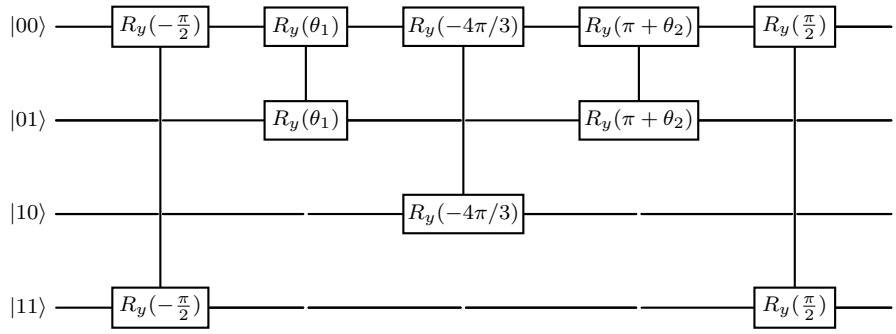
The process terminates when no single-pulse deletion keeps the error below ϵ. On average, the process reduces the depth of the pulse sequence by 35-40% (based on $\mathcal{O}$), in the case of Hadamard decomposition. The shorter execution window reduces exposure to laser-amplitude noise and magnetic-field drift, offering a practical path toward high-fidelity multi-level gates within the native star topology.

7.2.1 Hadamard Decomposition

Applying the star–topology decomposition to the two-virtual-qubit Hadamard,

$$H^{\otimes 2} = \frac{1}{2}\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix}. \tag{7.6}$$

produces the pulse sequence shown here:

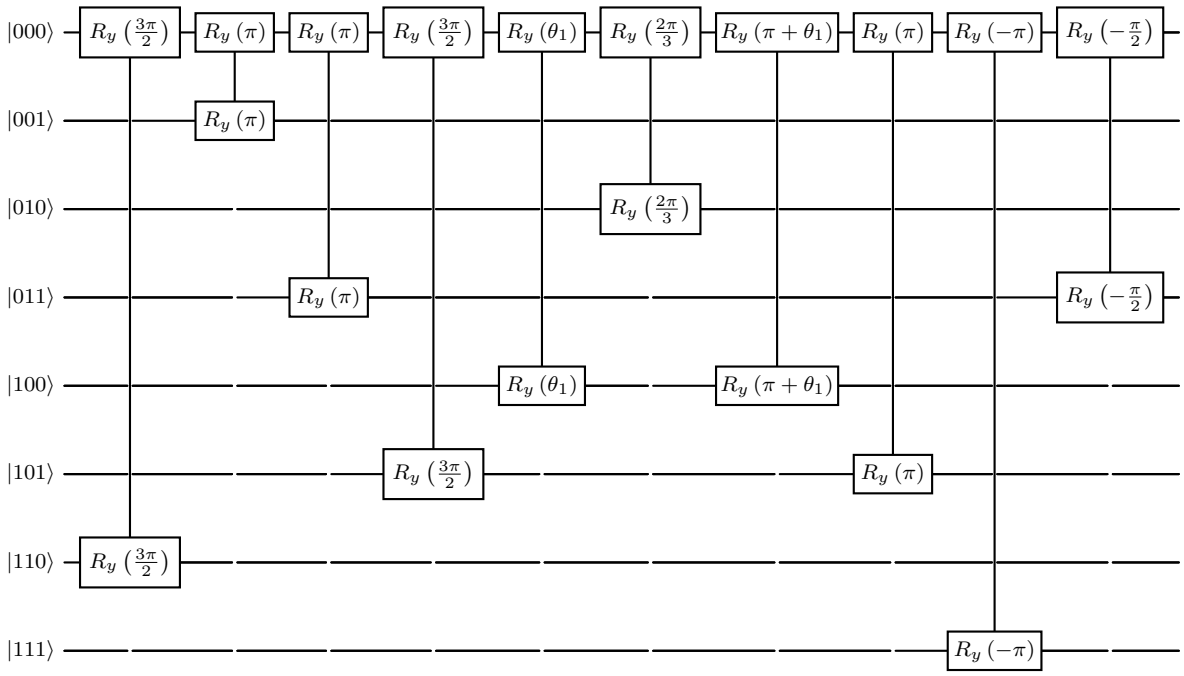


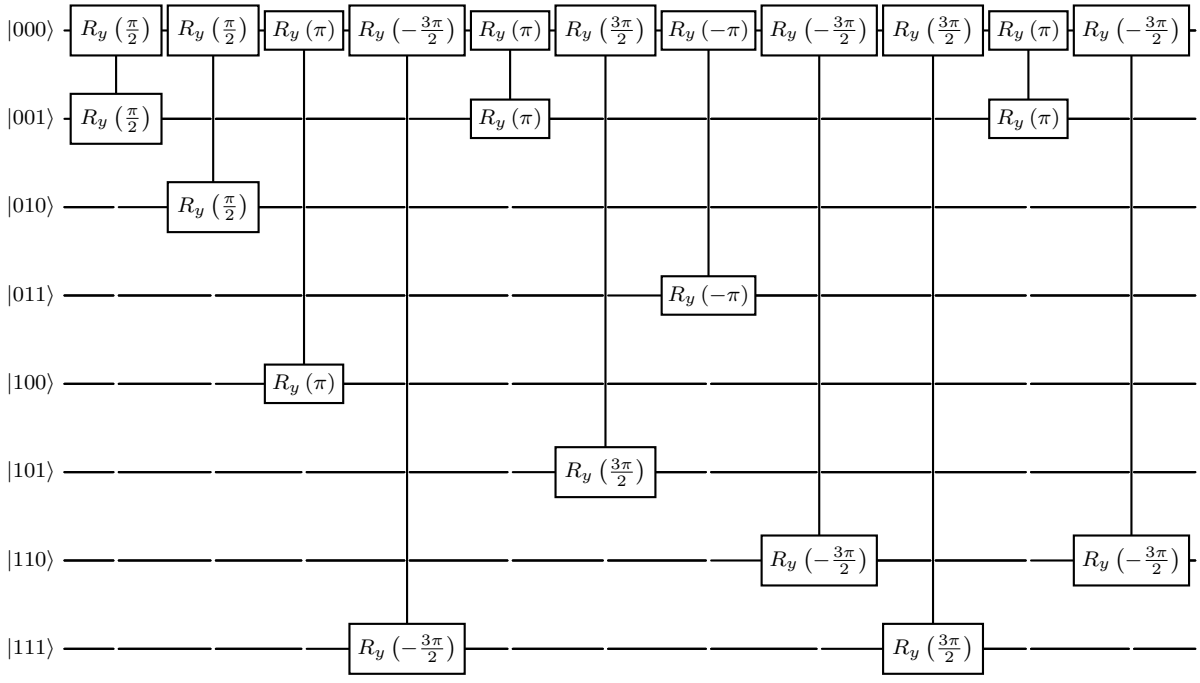
where, $\theta_i = 2\arcsin\left(\sqrt{i/3}\right)$

And applying the decomposition to the three virtual qubits

$$H^{\otimes 3} = \frac{1}{2\sqrt{2}}\begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 \\ 1 & -1 & 1 & -1 & -1 & 1 & -1 & 1 \\ 1 & 1 & -1 & -1 & -1 & -1 & 1 & 1 \\ 1 & -1 & -1 & 1 & -1 & 1 & 1 & -1 \end{bmatrix}. \tag{7.7}$$

Yields a 21-pulse long sequence shown here:


|000⟩
|001⟩
|010⟩
|011⟩
|100⟩
|101⟩
|110⟩
|111⟩
$R_y\left(\frac{3\pi}{2}\right)$
$R_y(\pi)$
$R_y(\pi)$
$R_y\left(\frac{3\pi}{2}\right)$
$R_y(\theta_1)$
$R_y\left(\frac{2\pi}{3}\right)$
$R_y(\pi+\theta_1)$
$R_y(\pi)$
$R_y(-\pi)$
$R_y\left(-\frac{\pi}{2}\right)$
$R_y(\pi)$
$R_y\left(\frac{2\pi}{3}\right)$
$R_y\left(-\frac{\pi}{2}\right)$
$R_y(\theta_1)$
$R_y(\pi+\theta_1)$
$R_y\left(\frac{3\pi}{2}\right)$
$R_y(\pi)$
$R_y\left(\frac{3\pi}{2}\right)$
$R_y(-\pi)$


|000⟩
|001⟩
|010⟩
|011⟩
|100⟩
|101⟩
|110⟩
|111⟩
$R_y\left(\frac{\pi}{2}\right)$
$R_y\left(\frac{\pi}{2}\right)$
$R_y(\pi)$
$R_y\left(-\frac{3\pi}{2}\right)$
$R_y(\pi)$
$R_y\left(\frac{3\pi}{2}\right)$
$R_y(-\pi)$
$R_y\left(-\frac{3\pi}{2}\right)$
$R_y\left(\frac{3\pi}{2}\right)$
$R_y(\pi)$
$R_y\left(-\frac{3\pi}{2}\right)$
$R_y\left(\frac{\pi}{2}\right)$
$R_y(\pi)$
$R_y(\pi)$
$R_y\left(\frac{\pi}{2}\right)$
$R_y(-\pi)$
$R_y(\pi)$
$R_y\left(\frac{3\pi}{2}\right)$
$R_y\left(-\frac{3\pi}{2}\right)$
$R_y\left(-\frac{3\pi}{2}\right)$
$R_y\left(-\frac{3\pi}{2}\right)$
$R_y\left(\frac{3\pi}{2}\right)$

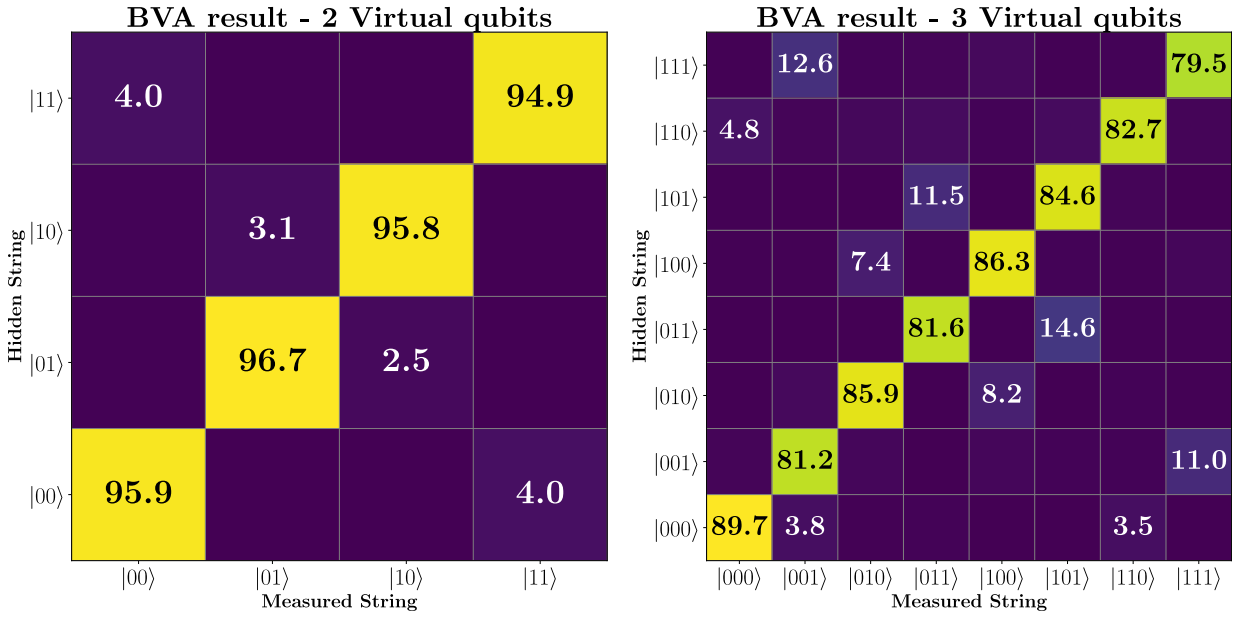


Figure 7.1: **BVA results on two and three virtual qubits.** The plots show the results of the BVA implementation for two virtual qubits (left) and three virtual qubits (right). The average success probabilities of the algorithm are 95% and 84%, respectively. The observed XX and XXI errors in the outcomes are attributed to residual noise from the AC line signal.

7.3 Experimental Results

7.3.1 Bernstein–Vazirani with two and three virtual qubits

Leveraging the composite Hadamard sequences developed above, we implemented the Bernstein–Vazirani algorithm on a single $^{137}Ba^{+}$ ion encoded as either two or three virtual qubits. The oracle is compiled entirely into virtual Z operations, so its action reduces to an instantaneous frame update and adds no physical pulse or decoherence overhead.

Figures 7.1 display the outcome distributions. For every hidden bit string s, we plot the measured probability $P(m \mid s)$ of registering the computational basis string m at the end of the sequence. In all cases, the correct string emerges as the dominant peak, with about 95% success probability in two virtual qubits and 84% success probability in three virtual qubits. The leading error channel in both experiments was traced to magnetic-field

fluctuations synchronous with the 60 + 180 Hz AC mains, see Appendix E for simulation results. These field drifts imprint correlated phase shifts that appear as XX-type errors in the two-qubit register and as XXI errors in the three-qubit register. The mechanism by which line noise couples preferentially to these specific types of errors is still under investigation.

Earlier trapped-ion studies have benchmarked the Bernstein–Vazirani algorithm at the two- and three-qubit scale. Fallek *et al.* used a micro-fabricated surface trap with shuttled ion chains, achieving 97.6% success for the 2-qubit case and 80.9% for three ions [61]. These results provide useful ion-platform reference points for the 2- and 3-virtual-qubit BV experiments reported here.

7.3.2 Controlled-Toffoli Gate

The Hadamard layer required by the Bernstein–Vazirani algorithm exemplifies a *unfavourable* operation for our virtual-qubit architecture. Because concurrent pulses are not yet available, the amplitude and phase redistribution must be carried out serially through the hub state $|0\rangle$. In comparison, a standard qubit-based processor can do this by applying n parallel single-qubit gates as on a conventional register. By contrast, the encoding excels when two-qubit gates are involved since all the gates are mapped to intra-ion gates. Multi-controlled NOT gates are a prime example: a conventional trapped-ion implementation of a CCCNOT (Controlled-Toffoli) demands a minimum of five CNOTs, whereas a virtual-qubit register realizes the same logic in a single $|1110\rangle \leftrightarrow |1111\rangle$ Givens rotation.

Figure 7.2 shows the four-virtual-qubit controlled-Toffoli executed with a mean success probability (over 1000 shots) of 99.5%. The residual error budget is dominated by state preparation and measurement (SPAM) infidelity rather than by the gate itself.

Earlier trapped-ion experiments provide useful reference points for multi-control Toffoli gates. Figgatt *et al.* implemented Toffoli-3 and Toffoli-4 gates in a five-ion chain, reporting process fidelities of 89.6% and 70.5%, respectively, after SPAM correction [62]. More recently, Fang *et al.* extended a Cirac–Zoller protocol to longer chains and achieved truth-table fidelities of 95.4% for Toffoli-3, 92.2% for Toffoli-4, and 90.4% for Toffoli-5 [63]. These results set a trapped-ion baseline against which the 99.5% success of our 4-virtual-qubit CCCNOT can be directly compared.

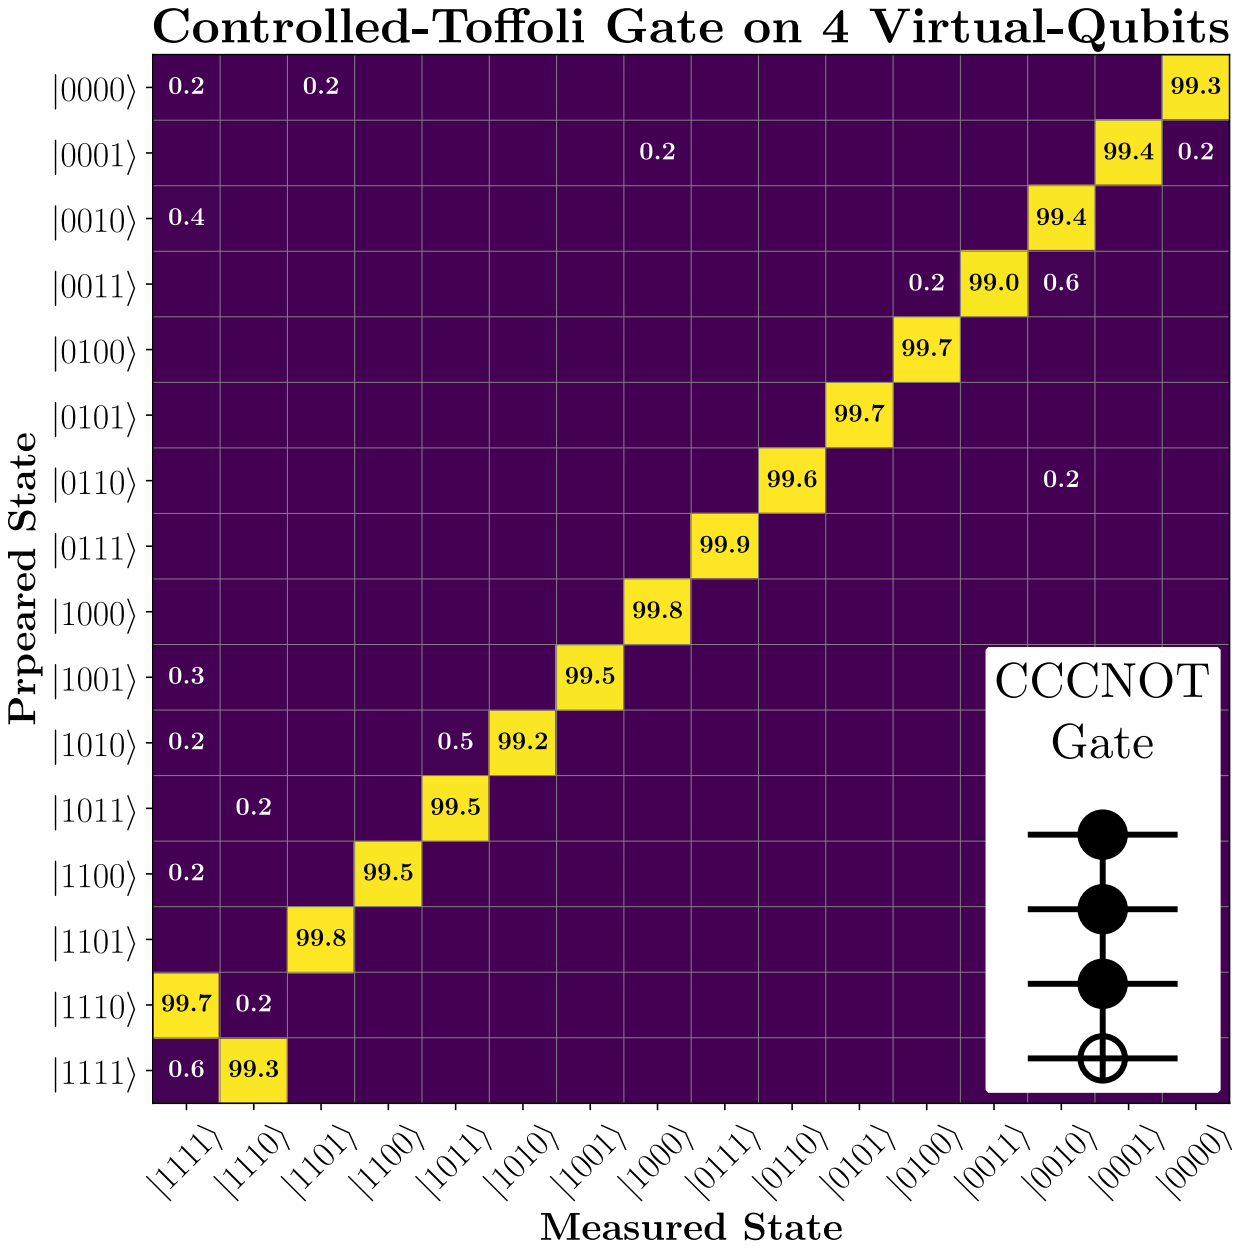


Figure 7.2: **Controlled-Toffoli gate on four virtual qubits.** The figure shows the results of the Controlled Toffoli gate (shown in the inset circuit) on four virtual qubits. The average success probability is 99.5% (currently limited by SPAM fidelity).

Chapter 8

Conclusion and Outlook

The experiments reported here demonstrate $^{137}Ba^{+}$ as a versatile high-dimensional trapped-ion qudit platform that reliably controls 25 computational basis states within a single ion, reaching error rates suitable for proof-of-principle algorithms while significantly reducing hardware overhead compared with multi-ion chains. A narrow-band optical-pumping scheme followed by shelving read-out delivers an average SPAM fidelity of 99.51% across all levels. A two-point Ramsey protocol actively compensates laser-phase and detuning drift to keep shot-to-shot errors below 10^{-2}. Multilevel Ramsey interferometry shows that star-topology encodings retain high fringe contrast up to dimension $d = 17$, with residual 60 Hz/180 Hz magnetic pick-up accounting for the contrast loss at larger d. Mapping the 25-level register onto virtual-qubit subspaces enables algorithmic primitives, including a Bernstein–Vazirani search that succeeds with 95% (two virtual qubits) and 84% (three virtual qubits) probability, and a four-virtual-qubit controlled-Toffoli gate that achieves 99.5% success using a single Givens rotation. A self-consistent Monte Carlo error model reproduces these results within a few percent, confirming that high-fidelity, noise-limited control of a 25-level Hilbert space is already practical and dramatically reduces the number of physical ions required for large-scale quantum computation.

A clear upgrade path now targets the dominant technical limits: applying Fiber noise-cancellation to the 1762 nm quadrupole beam and surrounding the trap with multi-layer Mu-metal shielding will suppress laser-phase and 60 Hz/180 Hz magnetic field fluctuations. Now that it has been shown that the AC line signal is deterministic, one could compensate for the phase and detunings due to 60 Hz and 180 Hz. Experiment speed, and hence calibration value, can be improved by inserting a broadband electro-optic modulator that generates sidebands addressing every $6P_{3/2}$ hyperfine line simultaneously, and by replacing long, memory-hungry AWG waveforms with direct-digital-synthesis phase updates

whose resource cost grows only linearly with the 25 states. Although Doppler cooling suffices today, ground-state preparation can be accelerated without adding a second ion by implementing EIT cooling on the $6S_{1/2}$ and $6P_{1/2}$ [64] manifold. Together, these noise-reduction, bandwidth, and cooling upgrades pave the way to faster and higher-fidelity universal SU(25) gates, error-corrected logical qudits within a single ion, and ultimately fault-tolerant, resource-efficient quantum processors whose scale is set by Hilbert-space density rather than ion count.

References

[1] Mark Saffman. Quantum computing with neutral atoms. *National Science Review*, 6(1):24–25, January 2019.

[2] He-Liang Huang, Dachao Wu, Daojin Fan, and Xiaobo Zhu. Superconducting quantum computing: A review. *Science China Information Sciences*, 63(8):1–32, August 2020.

[3] Colin D. Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M. Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2), June 2019.

[4] J. I. Cirac and P. Zoller. Quantum Computations with Cold Trapped Ions. *Physical Review Letters*, 74(20):4091, May 1995.

[5] David P. DiVincenzo. The Physical Implementation of Quantum Computation.

[6] Fangzhao Alex An, Anthony Ransford, Andrew Schaffer, Lucas R. Sletten, John Gaebler, James Hostetter, and Grahame Vittorini. High Fidelity State Preparation and Measurement of Ion Hyperfine Qubits with I>1/2. *Physical Review Letters*, 129(13):130501, September 2022.

[7] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits. *Physical Review Letters*, 117(6):060504, August 2016.

[8] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland. High-Fidelity Universal Gate Set for 9Be+ Ion Qubits. *Physical Review Letters*, 117(6):060505, August 2016.

[9] C. Langer, R. Ozeri, J. D. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. B. Blakestad, J. Britton, D. B. Hume, W. M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland. Long-Lived Qubit Memory Using Atomic Ions. *Physical Review Letters*, 95(6):060502, August 2005.

[10] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas. High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit. *Physical Review Letters*, 113(22):220501, November 2014.

[11] Jaroslaw Labaziewicz, Yufei Ge, Paul Antohi, David Leibrandt, Kenneth R. Brown, and Isaac L. Chuang. Suppression of Heating Rates in Cryogenic Surface-Electrode Ion Traps. *Physical Review Letters*, 100(1):013001, January 2008.

[12] Jwo-Sy Chen, Erik Nielsen, Matthew Ebert, Volkan Inlek, Kenneth Wright, Vandiver Chaplin, Andrii Maksymov, Eduardo Páez, Amrit Poudel, Peter Maunz, and John Gamble. Benchmarking a trapped-ion quantum computer with 30 qubits. *Quantum*, 8:1516, November 2024.

[13] Minzhao Liu, Ruslan Shaydulin, Pradeep Niroula, Matthew DeCross, Shih-Han Hung, Wen Yu Kon, Enrique Cervero-Martín, Kaushik Chakraborty, Omar Amer, Scott Aaronson, Atithi Acharya, Yuri Alexeev, K. Jordan Berg, Shouvanik Chakrabarti, Florian J. Curchod, Joan M. Dreiling, Neal Erickson, Cameron Foltz, Michael Foss-Feig, David Hayes, Travis S. Humble, Niraj Kumar, Jeffrey Larson, Danylo Lykov, Michael Mills, Steven A. Moses, Brian Neyenhuis, Shaltiel Eloul, Peter Siegfried, James Walker, Charles Lim, and Marco Pistoia. Certified randomness using a trapped-ion quantum processor. *Nature*, 640(8058):343–348, April 2025.

[14] M. DeCross, R. Haghshenas, M. Liu, E. Rinaldi, J. Gray, Y. Alexeev, C. H. Baldwin, J. P. Bartolotta, M. Bohn, E. Chertkov, J. Cline, J. Colina, D. DelVento, J. M. Dreiling, C. Foltz, J. P. Gaebler, T. M. Gatterman, C. N. Gilbreth, J. Giles, D. Gresh, A. Hall, A. Hankin, A. Hansen, N. Hewitt, I. Hoffman, C. Holliman, R. B. Hutson, T. Jacobs, J. Johansen, P. J. Lee, E. Lehman, D. Lucchetti, D. Lykov, I. S. Madjarov, B. Mathewson, K. Mayer, M. Mills, P. Niroula, J. M. Pino, C. Roman, M. Schecter, P. E. Siegfried, B. G. Tiemann, C. Volin, J. Walker, R. Shaydulin, M. Pistoia, S. A. Moses, D. Hayes, B. Neyenhuis, R. P. Stutz, and M. Foss-Feig. Computational Power of Random Quantum Circuits in Arbitrary Geometries. *Physical Review X*, 15(2):021052, May 2025.

[15] Q. A. Turchette, Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Physical Review A*, 61(6):063418, May 2000.

[16] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt. Ion-trap measurements of electric-field noise near surfaces. *Reviews of Modern Physics*, 87(4):1419, December 2015.

[17] K. A. Landsman, Y. Wu, P. H. Leung, D. Zhu, N. M. Linke, K. R. Brown, L. Duan, and C. Monroe. Two-qubit entangling gates within arbitrarily long chains of trapped ions. *Physical Review A*, 100(2):022332, August 2019.

[18] Nikodem Grzesiak, Reinhold Blümel, Kenneth Wright, Kristin M. Beck, Neal C. Pisenti, Ming Li, Vandiver Chaplin, Jason M. Amini, Shantanu Debnath, Jwo-Sy Chen, and Yunseong Nam. Efficient arbitrary simultaneously entangling gates on a trapped-ion quantum computer. *Nature Communications*, 11(1):1–6, June 2020.

[19] Reinhold Blümel, Nikodem Grzesiak, Neal Pisenti, Kenneth Wright, and Yunseong Nam. Power-optimal, stabilized entangling gate between trapped-ion qubits. *npj Quantum Information*, 7(1):1–8, October 2021.

[20] David Schwerdt, Lee Peleg, Yotam Shapira, Nadav Priel, Yanay Florshaim, Avram Gross, Ayelet Zalic, Gadi Afek, Nitzan Akerman, Ady Stern, Amit Ben Kish, and Roee Ozeri. Scalable Architecture for Trapped-Ion Quantum Computing Using rf Traps and Dynamic Optical Potentials. *Physical Review X*, 14(4):041017, October 2024.

[21] Ashok Muthukrishnan and Jr C. R. Stroud. Multivalued logic gates for quantum computation. *Physical Review A*, 62(5):052309, October 2000.

[22] A. B. Klimov, R. Guzmán, J. C. Retamal, and C. Saavedra. Qutrit quantum computer with trapped ions. *Physical Review A*, 67(6):062313, June 2003.

[23] Brian Mischuck and Klaus Mølmer. Qudit quantum computation in the Jaynes-Cummings model. *Physical Review A*, 87(2):022341, February 2013.

[24] Pei Jiang Low, Brendan M. White, Andrew A. Cox, Matthew L. Day, and Crystal Senko. Practical trapped-ion protocols for universal qudit-based quantum computing. *Physical Review Research*, 2(3):033128, July 2020.

[25] Martin Ringbauer, Michael Meth, Lukas Postler, Roman Stricker, Rainer Blatt, Philipp Schindler, and Thomas Monz. A universal qudit quantum processor with trapped ions. *Nature Physics*, 18(9):1053–1057, September 2022.

[26] Pavel Hrmo, Benjamin Wilhelm, Lukas Gerster, Martin W. van Mourik, Marcus Huber, Rainer Blatt, Philipp Schindler, Thomas Monz, and Martin Ringbauer. Native qudit entanglement in a trapped ion quantum processor. *Nature Communications*, 14(1):1–6, April 2023.

[27] Michael Meth, Jinglei Zhang, Jan F. Haase, Claire Edmunds, Lukas Postler, Andrew J. Jena, Alex Steiner, Luca Dellantonio, Rainer Blatt, Peter Zoller, Thomas Monz, Philipp Schindler, Christine Muschik, and Martin Ringbauer. Simulating two-dimensional lattice gauge theories on a qudit quantum computer. *Nature Physics*, 21(4):570–576, April 2025.

[28] Pei Jiang Low, Brendan White, and Crystal Senko. Control and readout of a 13-level trapped ion qudit. *npj Quantum Information*, 11(1):1–10, May 2025.

[29] R. Blatt and G. Werth. Precision determination of the ground-state hyperfine splitting in Barium 137 ion using the ion-storage technique. *Physical Review A*, 25(3):1476, March 1982.

[30] Matthew Hoffman. Observation of the Nuclear Magnetic Octupole Moment of 137Ba+.

[31] M. D. Davidson, L. C. Snoek, H. Volten, and A. Doenszelmann. Oscillator strengths and branching ratios of transitions between low-lying levels in the barium II spectrum. *Astronomy and Astrophysics*, 255:457–458, February 1992.

[32] D. De Munshi, T. Dutta, R. Rebhi, and M. Mukherjee. Precision measurement of branching fractions of \Ba{138}: Testing many-body theories below the 1% level. *Physical Review A*, 91(4):040501, April 2015.

[33] Zhiqiang Zhang, K. J. Arnold, S. R. Chanu, R. Kaewuam, M. S. Safronova, and M. D. Barrett. Branching fractions for P_{3/2} decays in $Baˆ+$. *Physical Review A*, 101(6):062515, June 2020.

[34] Carolyn Auchter, Thomas W. Noel, Matthew R. Hoffman, Spencer R. Williams, and Boris B. Blinov. Measurement of the branching fractions and lifetime of the 5D52 level of Ba+. *Physical Review A*, 90(6):060501, December 2014.

[35] Pei Jiang Low. Control and Readout of High-Dimensional Trapped Ion Qudits. September 2023.

[36] Brendan Bramman. Ablation loading and qudit measurements with barium ions. August 2023.

[37] Brendan M. White, Pei Jiang Low, Yvette de Sereville, Matthew L. Day, Noah Greenberg, Richard Rademacher, and Crystal Senko. Isotope-selective laser ablation ion-trap loading of \Ba{137} using a \ce{BaCl_2} target. *Physical Review A*, 105(3):033102, March 2022.

[38] Pei Jiang Low. Tolerable Experimental Imperfections for a Quadrupole Blade Ion Trap and Practical Qudit Gates with Trapped Ions. September 2019.

[39] Brendan Bramman. Measuring Trapped Ion Qudits. September 2019.

[40] T. Ruster, C. T. Schmiegelow, H. Kaufmann, C. Warschburger, F. Schmidt-Kaler, and U. G. Poschinger. A long-lived Zeeman trapped-ion qubit. *Applied Physics B*, 122(10):1–7, October 2016.

[41] P. Villemoes, A. Arnesen, F. Heijkenskjold, and A. Wannstrom. Isotope shifts and hyperfine structure of 134-138Ba II by fast ion beam-laser spectroscopy. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 26(22):4289, November 1993.

[42] Nicholas C. Lewty, Boon Leng Chuah, Radu Cazan, Murray D. Barrett, and B. K. Sahoo. Experimental determination of the nuclear magnetic octupole moment of 137Ba+ ion. *Physical Review A*, 88(1):012518, July 2013.

[43] A. S. Sotirova, J. D. Leppard, A. Vazquez-Brennan, S. M. Decoppet, F. Pokorny, M. Malinowski, and C. J. Ballance. High-fidelity heralded quantum state preparation and measurement. https://arxiv.org/abs/2409.05805v1, September 2024.

[44] Pei Jiang Low, Nicholas C. F. Zutt, Gaurav A. Tathed, and Crystal Senko. Quantum logic operations and algorithms in a single 25-level atomic qudit. https://arxiv.org/abs/2507.15799v1, July 2025.

[45] Peter A. Spring, Luka Milanovic, Yoshiki Sunada, Shiyu Wang, Arjan F. van Loo, Shuhei Tamate, and Yasunobu Nakamura. Fast Multiplexed Superconducting-Qubit Readout with Intrinsic Purcell Filtering Using a Multiconductor Transmission Line. *PRX Quantum*, 6(2):020345, June 2025.

[46] State Preparation & Measurement with Barium Qubits. https://ionq.com/resources/state-preparation-and-measurement-with-barium-qubits.

[47] Norman F. Ramsey. A Molecular Beam Resonance Method with Separated Oscillating Fields. *Physical Review*, 78(6):695, June 1950.

[48] Christian Sanner, Nils Huntemann, Richard Lange, Christian Tamm, and Ekkehard Peik. Autobalanced Ramsey Spectroscopy. *Physical Review Letters*, 120(5):053602, January 2018.

[49] M. O. Hecht, Kumar Saurav, Evangelos Vlachos, Daniel A. Lidar, and Eli M. Levenson-Falk. Beating the Ramsey limit on sensing with deterministic qubit control. *Nature Communications*, 16(1):1–8, April 2025.

[50] N. Bar-Gill, L. M. Pham, A. Jarmola, D. Budker, and R. L. Walsworth. Solid-state electronic spin coherence time approaching one second. *Nature Communications*, 4(1):1–6, April 2013.

[51] J. P. Lee, A. J. Bennett, J. Skiba-Szymanska, D. J. P. Ellis, I. Farrer, D. A. Ritchie, and A. J. Shields. Ramsey interference in a multilevel quantum system. *Physical Review B*, 93(8):085407, February 2016.

[52] Clément Godfrin, Rafik Ballou, Edgar Bonet, Mario Ruben, Svetlana Klyatskaya, Wolfgang Wernsdorfer, and Franck Balestro. Generalized Ramsey interferometry explored with a single nuclear spin qudit. *npj Quantum Information*, 4(1):1–7, October 2018.

[53] Soltamov Va, Kasper C, Poshakinskiy Av, Anisimov An, Mokhov En, Sperlich A, Tarasenko Sa, Baranov Pg, Astakhov Gv, and Dyakonov V. Excitation and coherent control of spin qudit modes in silicon carbide at room temperature. *Nature communications*, 10(1), November 2019.

[54] X. N. Feng and L. F. Wei. Multilevel atomic Ramsey interferometry for precise parameter estimations*. *Chinese Physics B*, 30(12):120601, December 2021.

[55] Anushka Thenuwara and Andrei Sidorov. Ramsey interferometry in three-level and five-level systems of $ˆ{87}Rb$ Bose-Einstein condensates. https://arxiv.org/abs/2308.11095v1, August 2023.

[56] John F. Kielkopf. New approximation to the Voigt function with applications to spectral-line profile analysis. *JOSA*, 63(8):987–995, August 1973.

[57] Empirical fits to the Voigt line width: A brief review. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 17(2):233–236, February 1977.

[58] James Keppens, Quinten Eggerickx, Vukan Levajac, George Simion, and Bart Sorée. Qudit vs. Qubit: Simulated performance of error correction codes in higher dimensions. https://arxiv.org/abs/2502.05992v1, February 2025.

[59] Gavin K. Brennen, Dianne P. O'Leary, and Stephen S. Bullock. Criteria for exact qudit universality. *Physical Review A*, 71(5):052318, May 2005.

[60] Gavin K. Brennen, Stephen S. Bullock, and Dianne P. O'Leary. Efficient Circuits for Exact-Universal Computations with Qudits. https://arxiv.org/abs/quant-ph/0509161v1, September 2005.

[61] S. D. Fallek, C. D. Herold, B. J. McMahon, K. M. Maller, K. R. Brown, and J. M. Amini. Transport implementation of the Bernstein–Vazirani algorithm with ion qubits. *New Journal of Physics*, 18(8):083030, August 2016.

[62] C. Figgatt, D. Maslov, K. A. Landsman, N. M. Linke, S. Debnath, and C. Monroe. Complete 3-Qubit Grover search on a programmable quantum computer. *Nature Communications*, 8(1):1–9, December 2017.

[63] Chao Fang, Ye Wang, Ke Sun, and Jungsang Kim. Realization of Scalable Cirac-Zoller Multi-Qubit Gates. https://arxiv.org/abs/2301.07564v1, January 2023.

[64] Chuanxin Huang, Chenxi Wang, Hongxuan Zhang, Hongyuan Hu, Zuqing Wang, Zhichao Mao, Shijiao Li, Panyu Hou, Yukai Wu, Zichao Zhou, and Luming Duan. Electromagnetically Induced Transparency Cooling of High-Nuclear-Spin Ions. *Physical Review Letters*, 133(11):113204, September 2024.

[65] D. F. V. James. Quantum dynamics of cold trapped ions with application to quantum computation. *Applied Physics B*, 66(2):181–190, February 1998.

APPENDICES

Appendix A

Quadrupole Geometric Factors

In any electric-quadrupole transition, the strength of coupling depends not only on the internal matrix element of the ion,$\langle J_D|Q|J_S\rangle$, but also on how the laser is pointed and polarized. Following Low (2023) [35] § 4.1.2, this extrinsic contribution is collected into a *geometric factor*

$$g^{(q)}(\theta_p, \theta_k) = \sum_{\alpha,\beta} \left|\epsilon_\alpha k_\beta c_{\alpha\beta,q}\right|, \tag{A.1}$$

where $q = 0, \pm 1, \pm 2$ labels the Δm component that is being driven, θ_k is the angle between the laser wave-vector k and the quantitation ($\hat{z}$) axis, and θ_p is the angle between the electric-field polarization and the xz-plane. For a linearly-polarized beam propagating in the xz-plane,it is shown that

$$g^{(0)}(\theta_p, \theta_k) = \frac{\sqrt{6}}{4} \left|\cos\theta_p \sin 2\theta_k\right|, \tag{A.2}$$

$$g^{(\pm 1)}(\theta_p, \theta_k) = \frac{1}{2} \left| \mp \cos\theta_p \cos 2\theta_k + i \sin\theta_p \cos\theta_k\right|, \tag{A.3}$$

$$g^{(\pm 2)}(\theta_p, \theta_k) = \frac{1}{4} \left| \cos\theta_p \sin 2\theta_k \mp 2i \sin\theta_p \sin\theta_k\right| , \tag{A.4}$$

In the interaction Hamiltonian, the full Rabi frequency for a given Δm channel is therefore

$$\Omega_q \approx E_0 \left|\langle J_D||\hat{Q}||J_S\rangle\right| g^{(q)}(\theta_p, \theta_k) \langle \tilde{F}_S\, m_{\tilde{F}_S}; k = 2, q|\tilde{F}_D\, m_{\tilde{F}_D}\rangle, \tag{A.5}$$

so g acts as a dimensionless weighting factor that converts the internal matrix element into the observable transition strength. Physically, it encodes two simple geometric facts:

1. **Directionality** —Rotating the beam away from the quantization axis changes the mixture of q-components that the field can drive.
2. **Polarization** — adjusting the polarization ellipse redistributes optical power among those q-components.

Because g depends only on the beam geometry, its form is completely general: once derived, it can be multiplied into any quadrupole matrix element or Clebsch–Gordan coefficient to predict relative Rabi frequencies across a manifold of states.

The expressions above were derived under the assumption of linearly polarized light. But this assumption is not true, as we see in figure 4.2, the pi-times for [-2,4,-4] and [2,4,4] are not the same, but from figure 2.9 they should be the same under linear polarization of the laser. In the work that follows, we extend the treatment to an elliptically polarized beam. By allowing both the ellipticity and orientation of the polarization ellipse to vary, we obtain a more versatile geometric factor that is directly applicable to experiments that rely on controlled ellipticity, such as the selective suppression of $\Delta m = \pm 1$ transitions while maintaining balanced $\Delta m = 0, \pm 2$ couplings. The derivation closely mirrors [35] method, differing only in the explicit form of ϵ; the reduced matrix elements and Clebsch–Gordan structure remain unchanged.

We will start with the electric field:

$$E(\vec{r}, t) = E_a e^{i(kz-\omega t)}\hat{x} + E_b e^{i(kz-\omega t+\phi)}\hat{y} \tag{A.6}$$

Where E_a and E_b are the semi-major and semi-minor axis of the polarization ellipse. k is the wavenumber, ω is the frequency of the electric field, and ϕ is the phase difference between the orthogonal components of the elliptically polarized electric field. This electric field propagates in $\hat{z}$ direction. Now, we will change the direction and propagate in $\hat{x} - \hat{z}$ plane, making an angle θ_k with $\hat{z}$ axis.

The wavevector defining the direction of propagation will be:

$$\frac{\hat{k}}{k} = \sin(\theta_K)\hat{x} + 0 + \cos(\theta_K)\hat{z} \tag{A.7}$$

And since electric fields oscillate perpendicular to the direction of propagation,

$$E(\vec{r}, t) = \left(E_a \cos\theta_K \hat{x} + E_b e^{i\phi}\hat{y} - E_a \sin\theta_K \hat{z}\right) e^{i(\vec{k}\vec{r}-\omega t)} \tag{A.8}$$

Now let the ellipticity of polarization ellipse be χ, we define it as $\tan\chi = E_b/E_a$. So we can write the amplitudes of electric fields along the major and minor axes as $E_a = E_0\cos\chi$ and $E_b = E_0\sin\chi$. Now, substituting this in the above equation, we get:

$$E(\vec{r},t) = \left(E_0\cos\chi\cos\theta_K\hat{x} + E_0\sin\chi e^{i\phi}\hat{y} - E_0\cos\chi\sin\theta_K\hat{z}\right)e^{i(\vec{k}\vec{r}-\omega t)} \tag{A.9}$$

$$= E_0\underbrace{\left(\cos\chi\cos\theta_K\hat{x} + \sin\chi e^{i\phi}\hat{y} - \cos\chi\sin\theta_K\hat{z}\right)}_{\vec{\epsilon}}e^{i(\vec{k}\vec{r}-\omega t)} \tag{A.10}$$

The physical electric field is the real part of this expression.

$$E(\vec{r},t) = E_0\Big[\underbrace{(\cos\chi\,\cos\theta_K\,\hat{x} + \sin\chi\,\cos\phi\,\hat{y} - \cos\chi\,\sin\theta_K\,\hat{z})}_{\boldsymbol{\epsilon}_1}\cos(\mathbf{k}\cdot\mathbf{r}-\omega t)$$

$$-\underbrace{(\sin\chi\,\sin\phi\,\hat{y})}_{\boldsymbol{\epsilon}_2}\sin(\mathbf{k}\cdot\mathbf{r}-\omega t)\Big]. \tag{A.11}$$

That is the standard real-valued representation of the elliptically polarized wave propagating in the $\hat{k}$ direction given by the angle θ_K in the xz-plane.

The geometric factor is defined as $g^{(q)} = \sum c_{i,j}^{(q)}\epsilon_i k_j$, where $c_{i,j}^{(q)}$ are Racah Tensors described in Ref. [65]. Belpow is the Mathematica code use to calculate the analytical expression of these geometric factors.

```
cMatrices = {Sqrt[1/6] {{-1, 0, 0}, #Racah Tensors
                        {0, -1, 0},
                       {0, 0, 2}},
                  (1/2)*{{0, 0, -1},
                             {0, 0, I},
                            {-1, I, 0}},
                (1/2)*{{0, 0, 1},
                         {0, 0, I},
                         {1, I, 0}},
                 (1/2)*{{1, -I, 0},
                           {-I, -1, 0},
                          {0, 0, 0}},
                        (1/2)*{{1, I, 0},
                             {I, -1, 0},
                             {0, 0, 0}}};
```

```
thetaK = \[Theta]k;
thetaP = \[Theta]p;
chi = \[Chi];

epsilon = {Cos[chi]  Cos[thetaK], Sin[chi]  Exp[I  d], -Cos[chi]  Sin[thetaK]};

k = {Sin[thetaK], 0, Cos[thetaK]};

calculateG[c_] := Sum[c[[i, j]]*epsilon[[i]]*k[[j]], {i, 1, 3}, {j, 1, 3}];

gList = FullSimplify[calculateG /@ cMatrices];
```

This gives the following geometric factors

$$
\begin{aligned}
g^{(q=0)} &= -\sqrt{\frac{3}{2}}\cos(\theta_k)\cos(\chi)\sin(\theta_k) \\
g^{(q=\pm1)} &= \frac{1}{2}(\pm\cos(2\theta_k)\cos(\chi) + ie^{i\phi}\cos(\theta_k)\sin(\chi)) \\
g^{(q=\pm2)} &= \frac{1}{2}\sin(\theta_k)\left(\cos(\theta_k)\cos(\chi) \pm ie^{i\phi}\sin(\chi)\right)
\end{aligned}
\tag{A.12}
$$

Figure A.1 shows the colormap of geometric factors at a fixed $\theta_k = 48.35°$ with respect to the quantization axis. We use the measured values of relative transition strengths of the five reference transitions used in Chapter 4 Section 4.4 to fit the θ_k, χ and ϕ values in equations (A.12). It is important to note that the calculations and equations presented in this appendix have not yet been tested experimentally. While the relative transition strengths show reasonable agreement with empirical values, the model includes three free parameters and is fit to only four data points. This results in a marginally underconstrained fit, meaning that multiple parameter combinations could potentially yield similar agreement with the data.

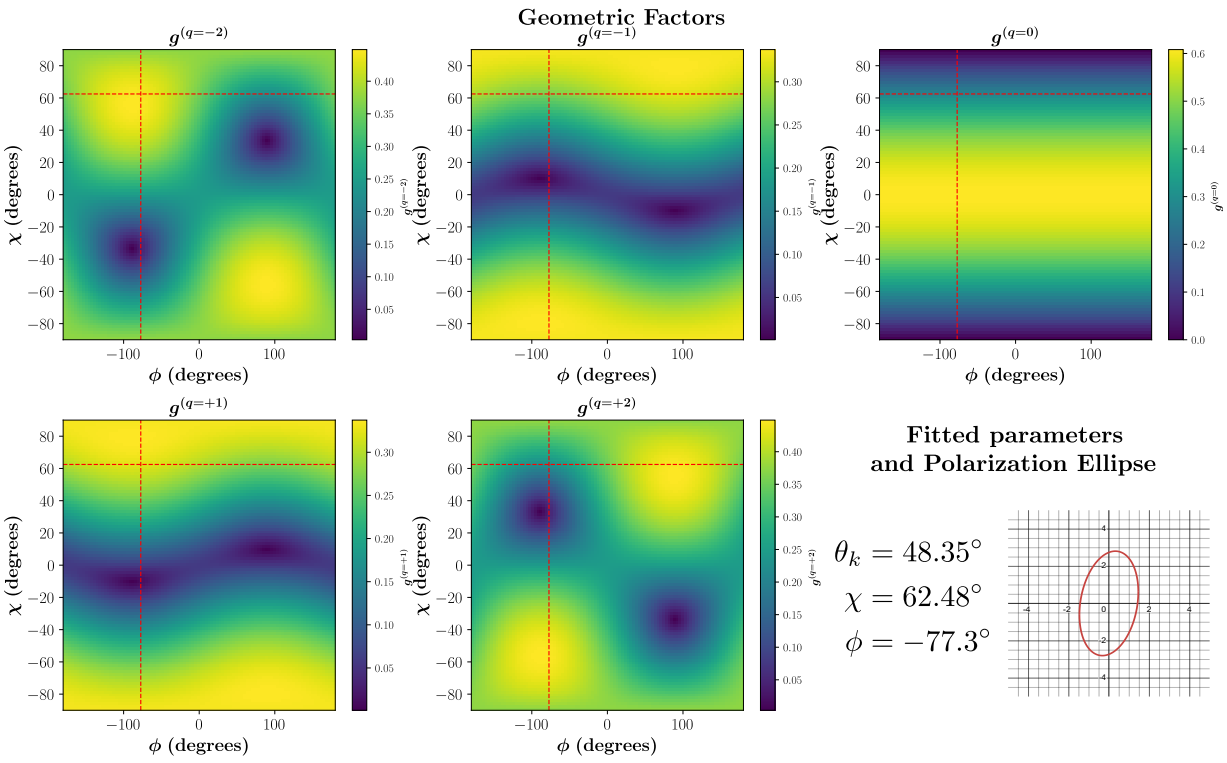


Figure A.1: **Geometric factors.**This figure shows the five geometric factors corresponding to each allowed value of Δm in the quadrupole transition. The parameter θ_k is held fixed throughout. The color plots represent the magnitude of each geometric factor as a function of the polarization angle χ and the phase angle ϕ. The intersection of the dashed red lines indicates the fitted values of χ and ϕ that are used consistently in all experiments presented in this thesis.

Appendix B

Intermittent Magnetic Field Noise

During data collection for a qubit Ramsey phase scan, we observed inconsistencies in the contrast of the phase scans, even when the wait time between pulses was held constant (see Fig. B.2). Notably, this contrast instability was absent when performing Ramsey experiments on a magnetic field-insensitive transition, suggesting that the observed effect originated from intermittent increases in magnetic field noise.

At the time, the quantization magnetic field was generated using electromagnetic coils powered by a Keysight E3631A power supply operating in constant current mode. Under ideal conditions, this configuration should produce a stable magnetic field by adjusting the output voltage to compensate for changes in temperature or ambient magnetic fields. Typically, voltage adjustments fluctuate by about 20 mV. However, during periods of elevated magnetic field noise, we observed unusually large voltage fluctuations—on the order of 100 mV (see Figure B.2). While minor voltage variations are expected to maintain a constant current, fluctuations of this magnitude are atypical and suggest potential instability in the electromagnetic environment, a faulty connection to the coils, or erratic compensation behavior by the power supply. The exact cause of these fluctuations was not investigated further.

To mitigate this issue and improve magnetic field stability, we replaced the electromagnetic coils with permanent magnets to serve as the source of the static B-field. This change effectively eliminated the observed contrast fluctuations in Ramsey scans and significantly improved the consistency of our measurements.

Magnetic field fluctuations with temperature, measured using two-point Ramsey calibration on our two reference transitions. Figure B.3 shows the measured frequencies and the corresponding calculated magnetic field values. Since the drift due to temperature

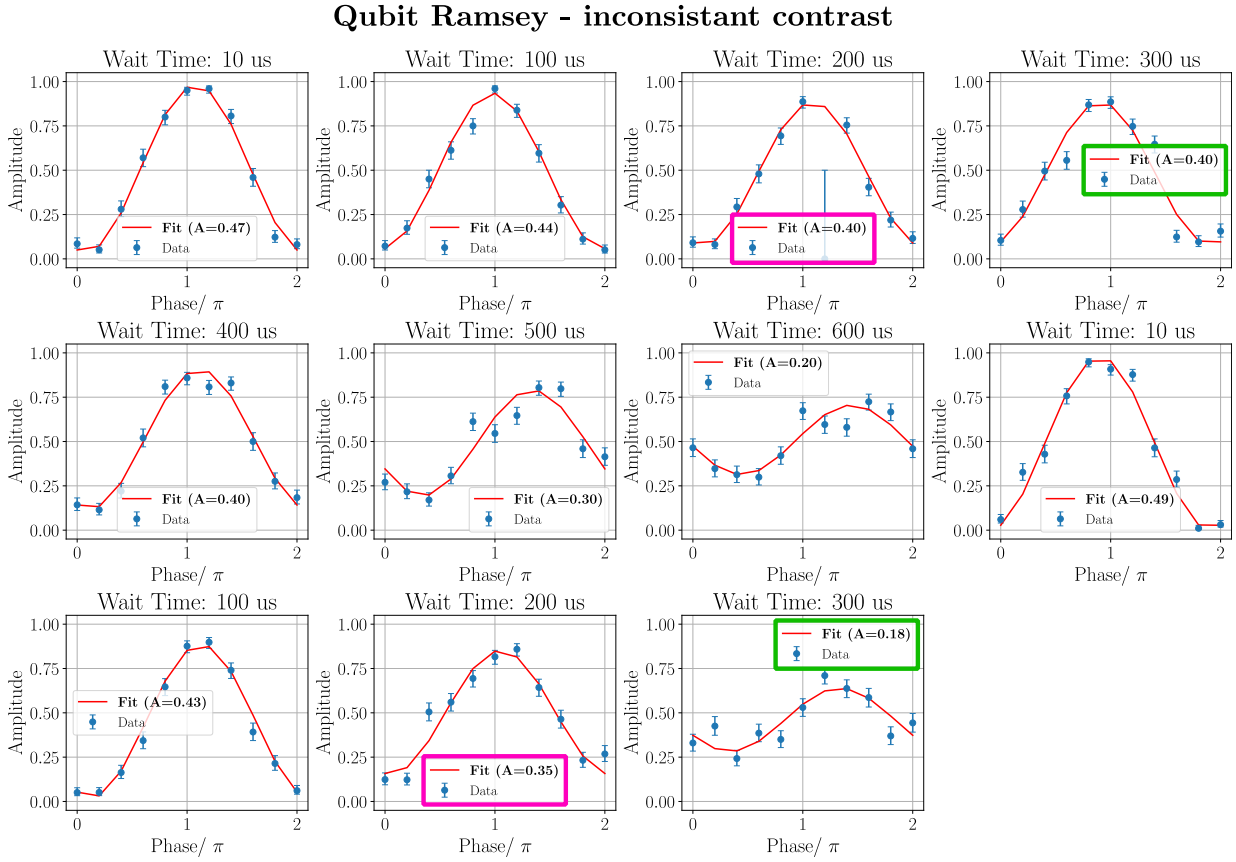


Figure B.1: **Qubit Ramsey - Inconsistent Contrast.** The panels show qubit Ramsey phase scans taken sequentially from left to right. Increased magnetic field noise becomes apparent starting from the second scan with a 200μs wait time, where the contrast drops from 80% to 70%. The effect is even more pronounced in the scan with a 300μs wait time, where the contrast decreases from 80% to 36%, indicating significant magnetic field fluctuations.

changes is relatively slow, it can be tracked using two-point calibration, which typically takes about 3 minutes per frequency measurement.

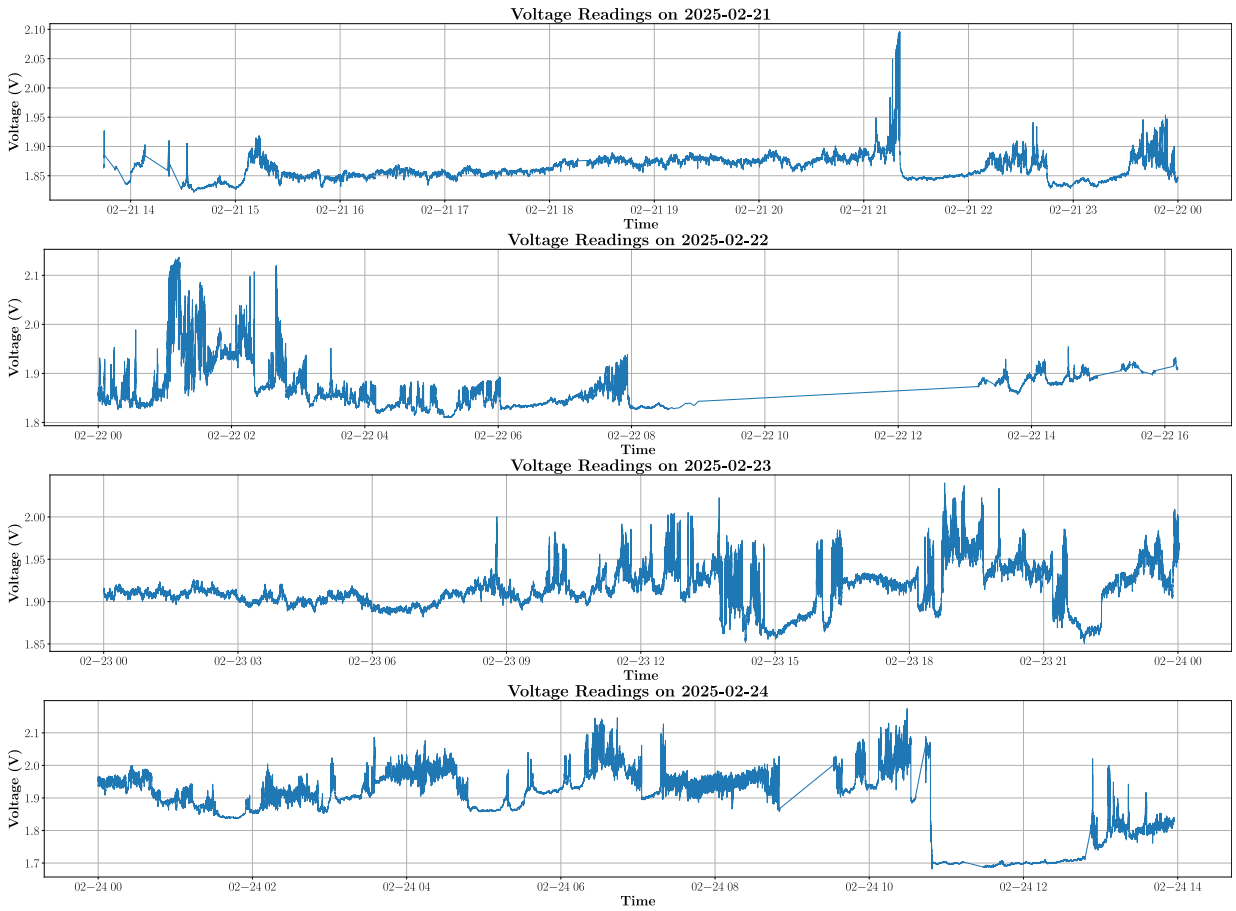


Figure B.2: **Voltage Reading from Keysight E3631A Power Supply.** The panels show the output voltage readings from the E3631A power supply over a four-day period while operating in constant current mode. A clear change in behavior is observed, with noticeable variations in the standard deviation over time, indicating intermittent instability.

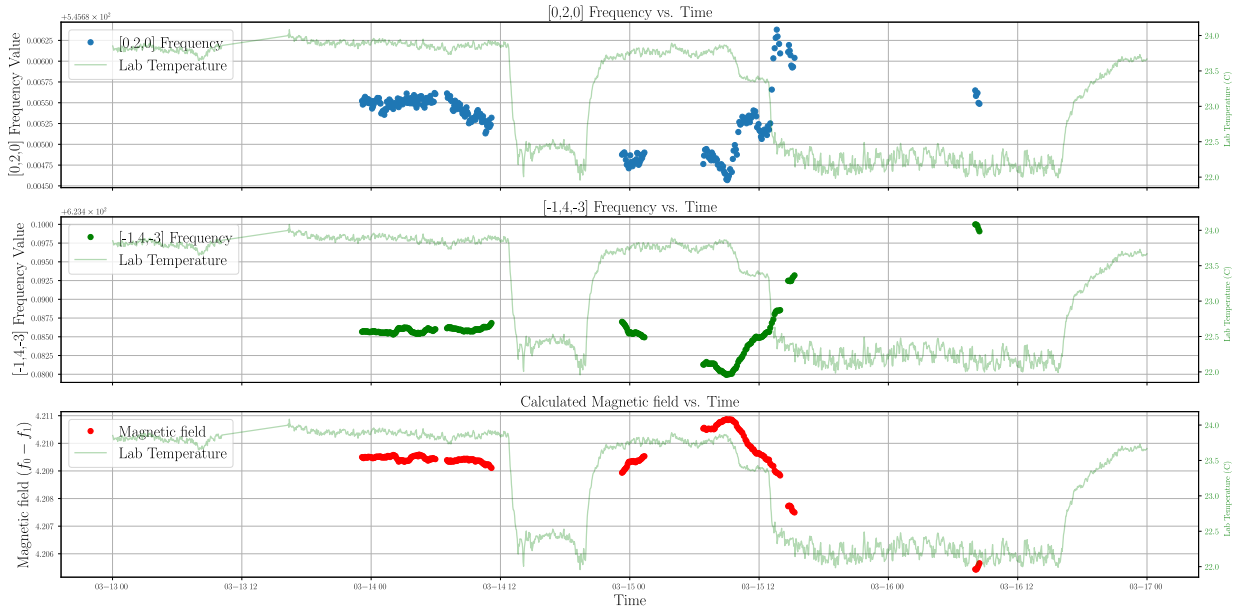


Figure B.3: **Variation in B-field due to temperature fluctuations.** The top two panels show the two reference frequencies: the $[0, 2, 0]$ transition, which is largely insensitive to magnetic fields and thus primarily tracks laser cavity drifts; and the $[-1, 4, -3]$ transition, which is sensitive to both magnetic field fluctuations and laser frequency drifts. To extract the magnetic field change, we compute the frequency difference between these two transitions and divide by the known magnetic field sensitivity of the $[-1, 4, -3]$ transition.

Appendix C

Permanent Magnets Simulations

In this Appendix, I describe the magnetostatic model exactly as it is implemented in COMSOL Multiphysics. The geometry consists of two coaxial, coplanar ring pairs of identical cylindrical Nd-Fe-B magnets (see Figure C.1 (a)). All dimensions—the ring radius R, the magnet diameter $\varnothing$, the axial thickness t, the radial gap g between the two layers, and the total number of magnets n in each layer, and separation between the two rings—are declared as global parameters so that a single change automatically regenerates the entire array. Each ring is populated by a rotate operation on cylinder geometry (representing single magnet) operation in which the angular pitch is computed internally as $\Delta\phi = 360^\circ/n$. Now, in the material data,a we need Remnant Magnetic flux density B_r of a single permanent magnet, which can be found in two ways: First is to use the quoted *maximum energy product*

$$(BH)_{\max} = 414 \text{ kJ m}^{-3} \tag{C.1}$$

For a hard magnet with a nearly straight demagnetization line (like for N52 grade magnets) ($\mu_r \approx 1$) the relationship

$$(BH)_{\max} \approx \frac{B_r^2}{4\mu_0} \tag{C.2}$$

holds. Rearranging,

$$B_r = \sqrt{4\mu_0 (BH)_{\max}} = \sqrt{4\,(4\pi \times 10^{-7})\,(4.14 \times 10^5)} \approx 1.45 \text{ T}. \tag{C.3}$$

We can also calculate the B_r using the quoted surface flux density of $B_s = 4430\,\mathrm{G}$. We do this by simulating a single permanent magnet and varying the remnant flux density in material properties to get the desired B_s (See figure C.2 (a)). From this method we find the $B_r = 1.305$ which we use in the simulation of the full array of magnets. COMSOL converts value to a fixed magnetization $\mathbf{M}_0 = \mathbf{B}_r/\mu_0 \approx 1.03 \times 10^6$ A m^{-1} while the relative permeability is left at $\mu_r = 1.05$.

The computation itself is performed with the *Magnetic Fields (mf)* interface in a *stationary* study, i.e. Ampère's law is solved without the displacement-current term. In vector-potential form, the governing equation is

$$\nabla \times (\nu \nabla \times \mathbf{A}) = -\nabla \times \mathbf{M}_0, \tag{C.4}$$

where $\nu = 1/\mu$ is the magnetic reluctivity and $\mathbf{A}$ is the magnetic vector potential. Once $\mathbf{A}$ is known, the flux density and field strength follow from

$$\mathbf{B} = \nabla \times \mathbf{A}, \qquad \mathbf{H} = \frac{1}{\mu_0}\mathbf{B} - \mathbf{M}_0.$$

No free-current sources appear, so the remnant magnetization alone drives the field.

The computational domain is enclosed by a cubical air region whose outer wall is set to the "magnetic insulation" boundary condition ($\mathbf{n}\cdot\mathbf{A} = 0$) and placed at least one ring diameter away to suppress spurious edge effects. A Free Tetrahedral mesh is used inside each magnet (Shown in Figure C.1 (b)), and the same kind of mesh with a larger size fills the surrounding air.

After the solver converges—typically within seconds—the post-processing step probes the magnitude of the Magnetic flux density of $\mathbf{B}$ at the center of the two rings (representing the ion position). This B-field is simulated to be 4.5 G. In contrast, empirical data show the field is about 4.2 G. This discrepancy almost certainly comes from the distance between the two rings not being the same as in the simulation. This is because we then 3-D print a holder for these magnets that need to sit comfortably on the viewports of the chamber, and this fit is challenging to engineer well due to space constraints. But from figure C.2 (b) we can see that the B-field changes by one whole gauss for a change in the separation between the arrays of 1 cm. It is quite likely that we are a few millimetres off from the simulated distance.

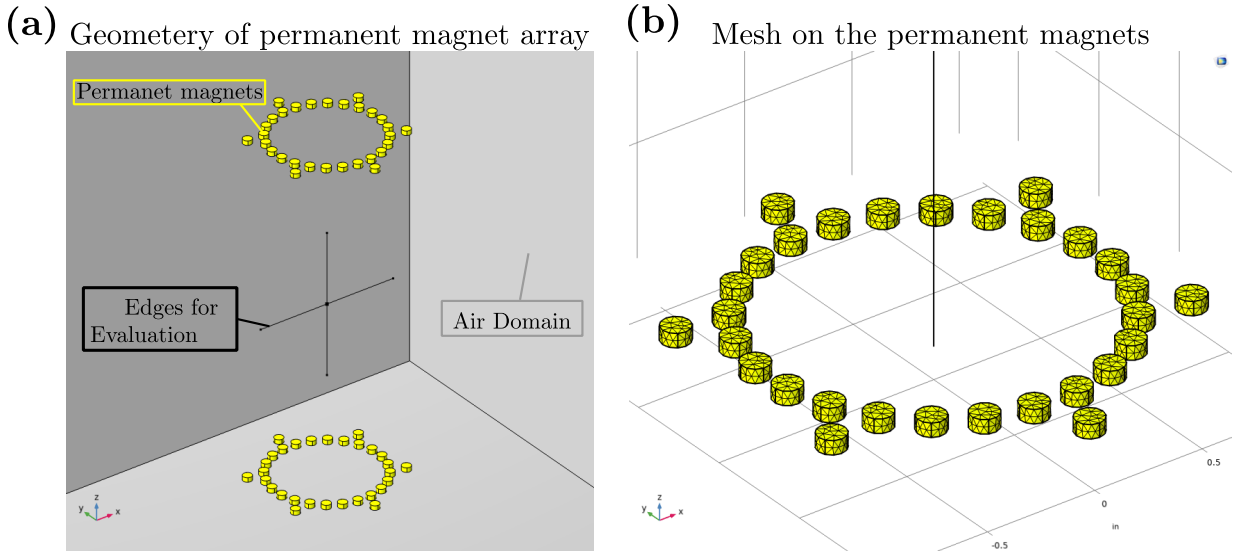


Figure C.1: **Permanent Magnet Geometry and Meshing**. **(a)** The figure shows the Permanent Magnet Configuration - two rings with 24 magnets in the inner layer and six magnets in the outer layer - separated by 4.8 in (separation set by vacuum chamber diameter). An air domain is added as computational space, and it is large enough that the edge effects are negligible. **(b)** Free tetrahedral mesh on permanent magnets. Each one has at least 40 elements in each magnet.

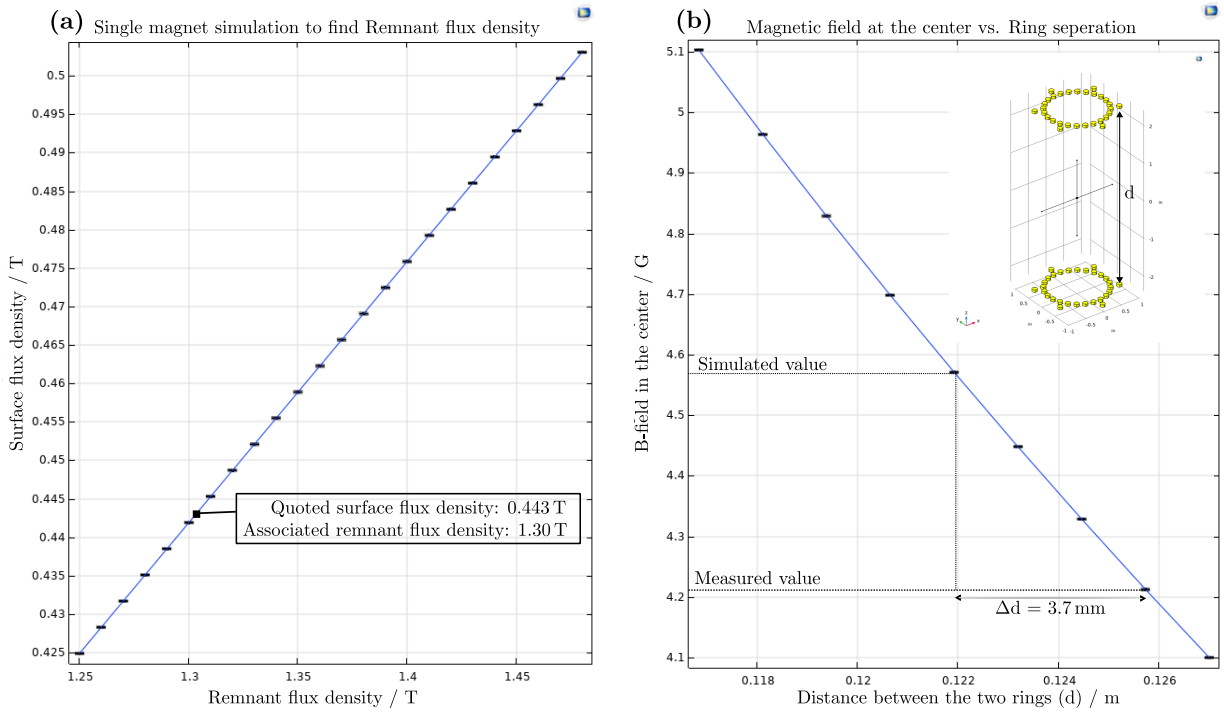


Figure C.2: **Remnant flux density estimation and B-field at the ion position. (a)** Shows the single magnet simulation result in the form of measured magnetic flux density on the surface vs. swept remnant flux density. We observe that the dependence is linear, and at the quoted value of the manufacturer, $B_r = 1.30\,\mathrm{T}$. **(b)** Shows how sensitive the B-field at the center of he trap is with respect to the separation between the rings. The total swept distance is 1 cm and for that range, the B-field changes by 1 G. The discrepancy in measured vs. the simulated B-field can be explained by $\sim 3.7\,\mathrm{mm}$ difference in ring separation.

Appendix D

SPAM Bright-Dark discrimination threshold

Figure 3.7 shows the bright/dark discrimination thresholds for each state in the SPAM experiment, as discussed in Chapter 3. These thresholds are determined by identifying the minimum between the two peaks in the fluorescence histogram.

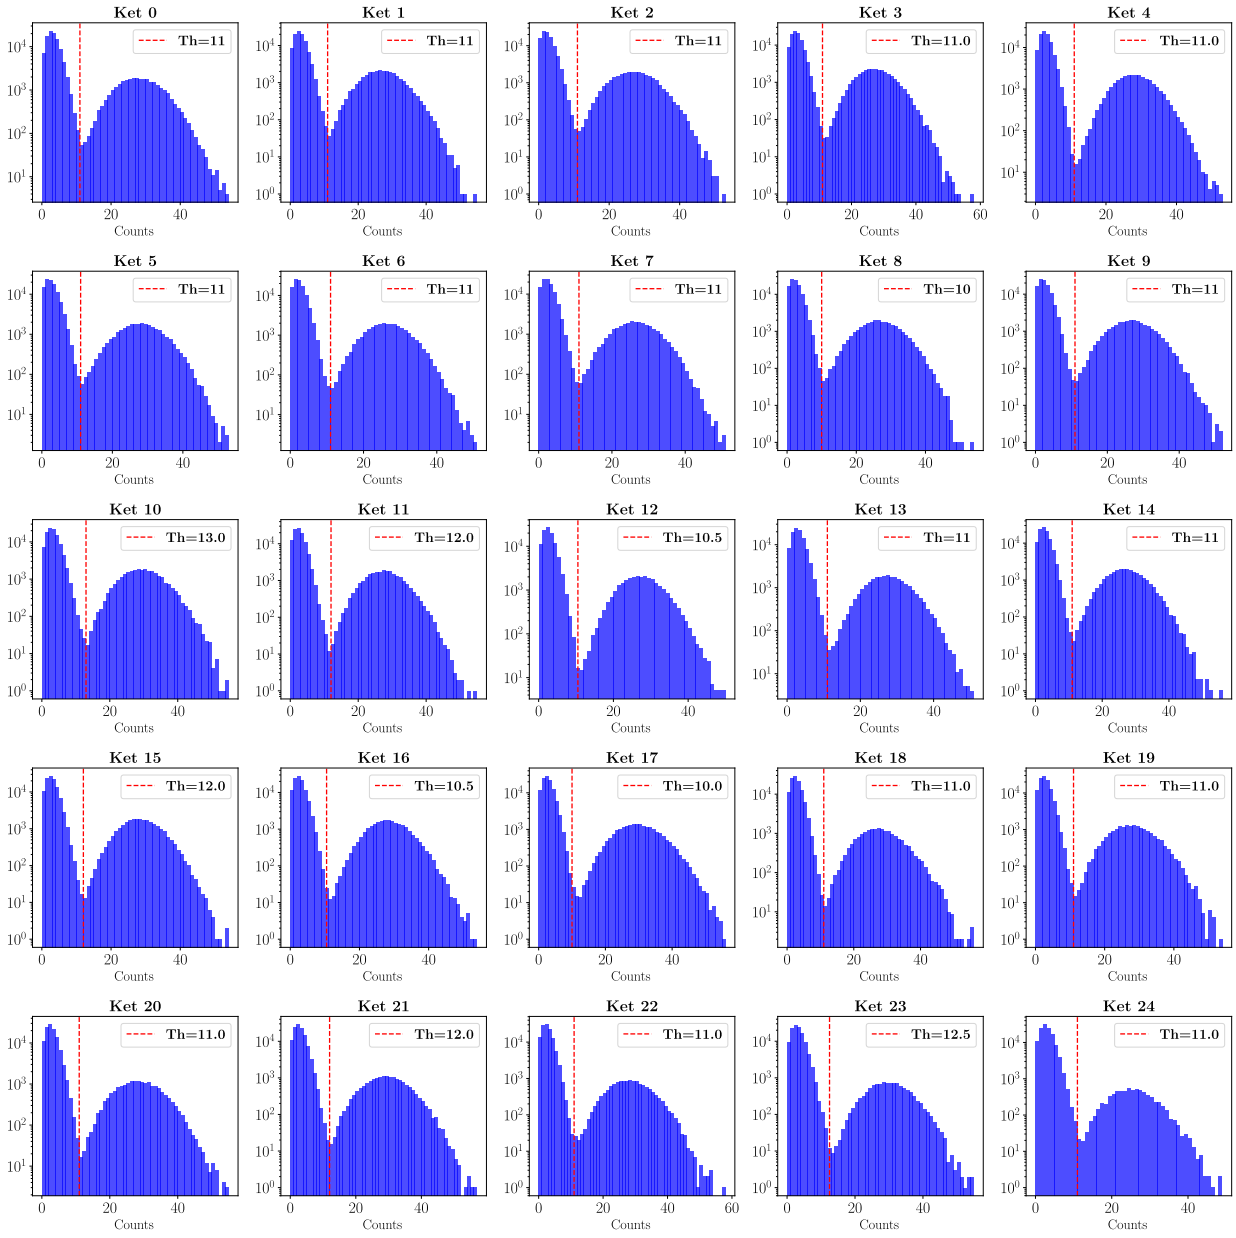


Figure D.1: **SPAM bright/Dark discrimination threshold.** Each of the above panels has 125000 shots. The average threshold for the SPAM experiment presented here is 11 counts.

Appendix E

Effects of Line Noise on Coherent operations

We have seen from Figure 5.4 that line noise (magnetic field changes due to AC line signal of 60 Hz and 180 Hz) is a major source of decoherence for qudits with large dimension. This is partly due to the increased operation time required for these larger qudits. In this appendix, I present some effects of the AC line signal observed during coherent operations, such as the Qudit Ramsey and Bernstein-Vazirani Algorithm.

Bussed Qudit Ramsey. In this experiment, we encode $|0\rangle$ in $5D_{5/2}$ (to take advantage of laser noise cancellation). Now to couple $|0\rangle$ to other states in $5D_{5/2}$ we will need two pulses, the first one will be a pulse that will take a fraction of the population from $|0\rangle$ to one of the states in $6S_{1/2}$ and second pulse will push that population to the desired state $|j\rangle$ in $5D_{5/2}$. So we more than double the operation time of the first part of Qudit Ramsey (since all the transitions now require a fractional pulse as well as a π-pulse). Then to bring the population back to $|0\rangle$ from each state in $5D_{5/2}$ we need three extra pulses — first, a π-pulse to deshelve the population from $|j\rangle$ into $6S_{1/2}$ level, second, a fractional pulse with a phase $\phi_j = j\phi$ for $|j\rangle$. And finally, another π-pulse back to $|j\rangle$ make sure that the bus state is empty at the end of the operation. So the total operation time is increased by more than 400% on average. The results of the experiment show an interesting feature in the Ramsey phase scans (shown in Figure E.1 (b)). We see an unexpected peak in an intermediate phase between 0 and 2π. This feature can also be seen in the Bussed-Qudit Ramsey simulation (shown in Figure E.1 (a)). Since this feature is deterministic, we should be able to compensate for the effects of line noise and improve the contrast by eliminating this feature.

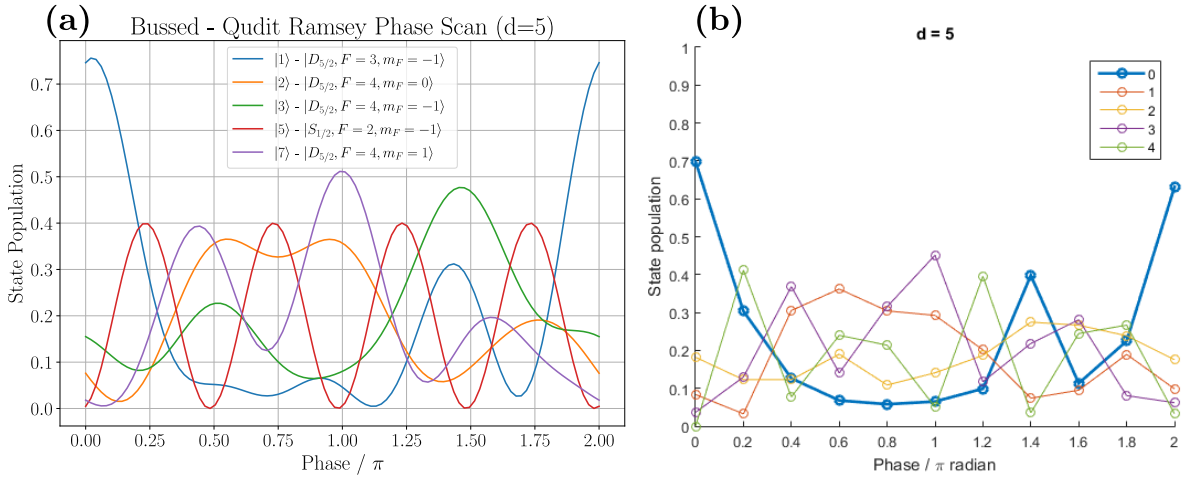


Figure E.1: **Bussed Qudit Ramsey for d=5.** **(a)** shows the simulated result for a qudit Ramsey where $|0\rangle$ is in $5D_{5/2}$ this simulation only has magnetic field changes due to AC line signal and no other noise source. **(b)** shows the experimental result of the same qudit ramsey for d=5. In both the we can clearly see a $|0\rangle$ population peak at phase = 1.4π.

Bernstein-Vazirani Algorithm As seen in Chapter 7, we see XXI-type errors in 3-virtual qubit BVA implementation. This errors are also predicted by simulating the BVA with AC line noise. Figure E.2 displays the simulated results of these experiments on 3 virtual qubits (with and without AC line signal), confirming that the effects of the AC line signal are deterministic and can be compensated for.

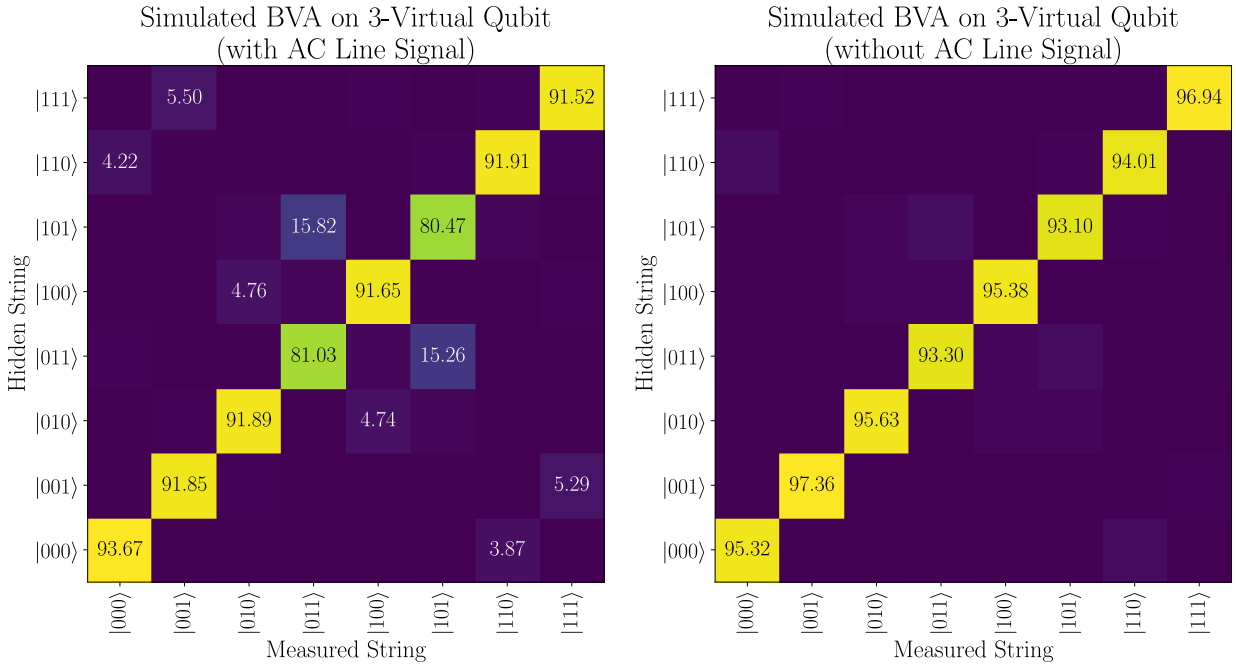


Figure E.2: **BVA Simulation on 3-virtual qubits.** Comparison of BVA results with and without the line signal. We clearly see that majority of XXI error disappears with AC line signal.

Appendix F

All Relative Magnetic field Sensitivities

Figure F.1 shows the relative magnetic field sensitivities between each of the 29 coherently controllable levels. These values are calculated by solving the Hamiltonian described in Chapter 2 Section 2.3. And We can also extrapolate the T_2^* times for all qubit pairs, using the error characterization described in 6. The coherence time between two levels in $5D_{5/2}$ can be calculated using equation (6.12)

$$T_2^* = \frac{1}{\sqrt{2}\pi\Delta\kappa\sigma_B} \tag{F.1}$$

And coherence between direct a state in $6S_{1/2}$ and $5D_{5/2}$ can be calculated using

$$T_2^* = \frac{-L_{\mathrm{L}} + \sqrt{L_{\mathrm{L}}^2 + 4\left(|\kappa| \cdot \sigma_B\right)^2 + 4L_{\mathrm{G}}^2}}{2\left(\left(|\kappa| \cdot \sigma_B\right)^2 + L_{\mathrm{G}}^2\right)} \tag{F.2}$$

Where $L_{\mathrm{L}} = 1/T_L$ and $L_{\mathrm{G}} = 1/T_G$, where T_L and T_G are from Chapter 6 Section 6.3.

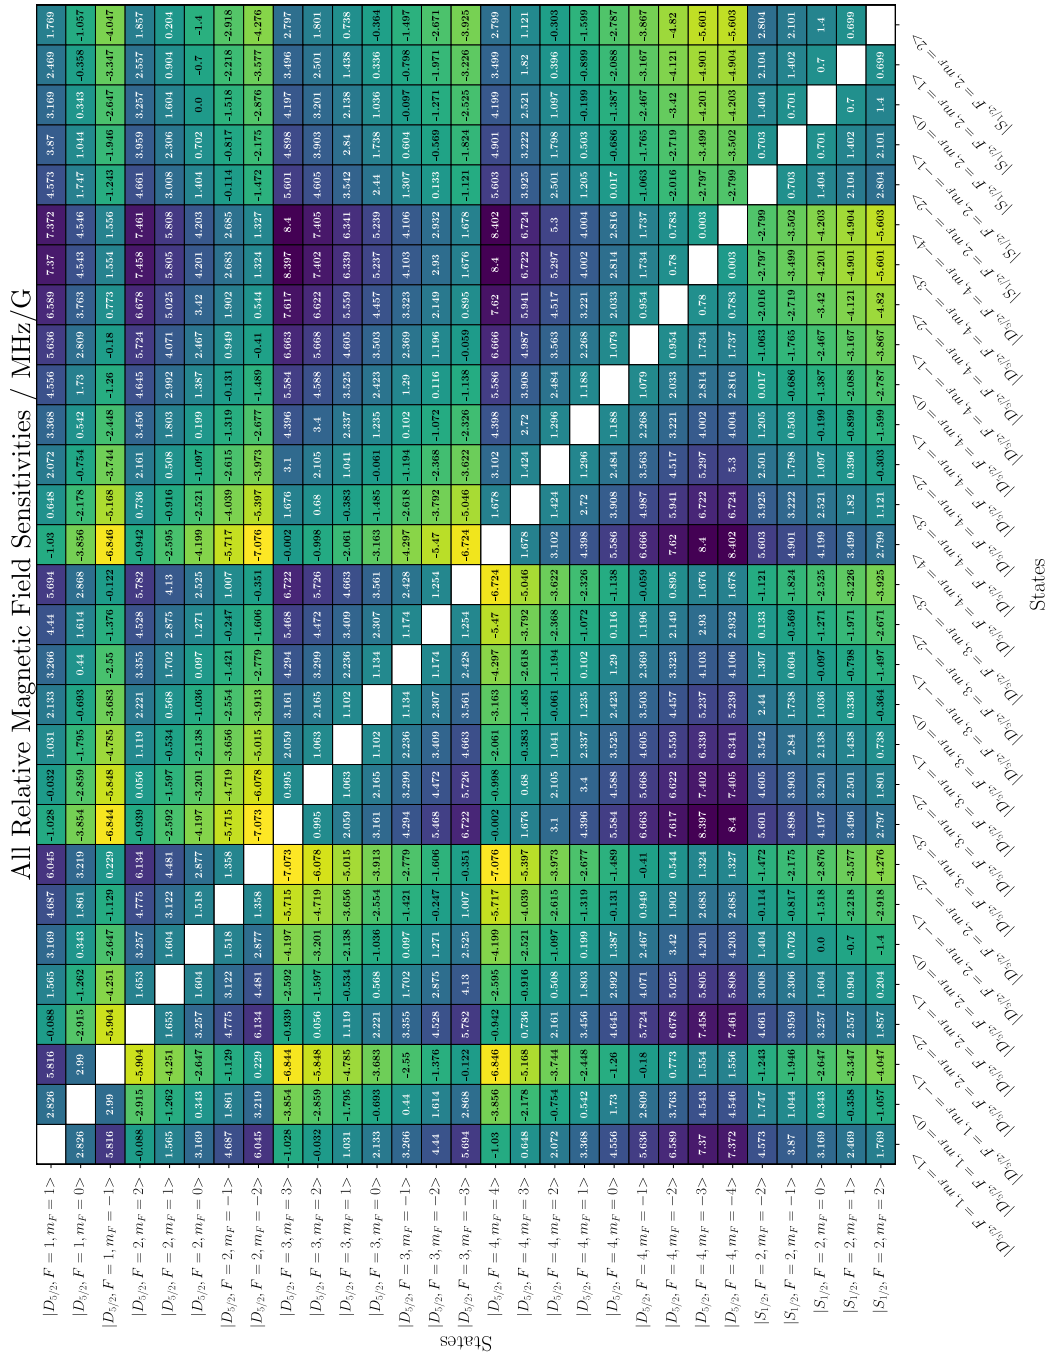


Figure F.1: **All magnetic field sensitivities.** The Colormap shows the absolute magnetic field sensitivity in MHz/G.

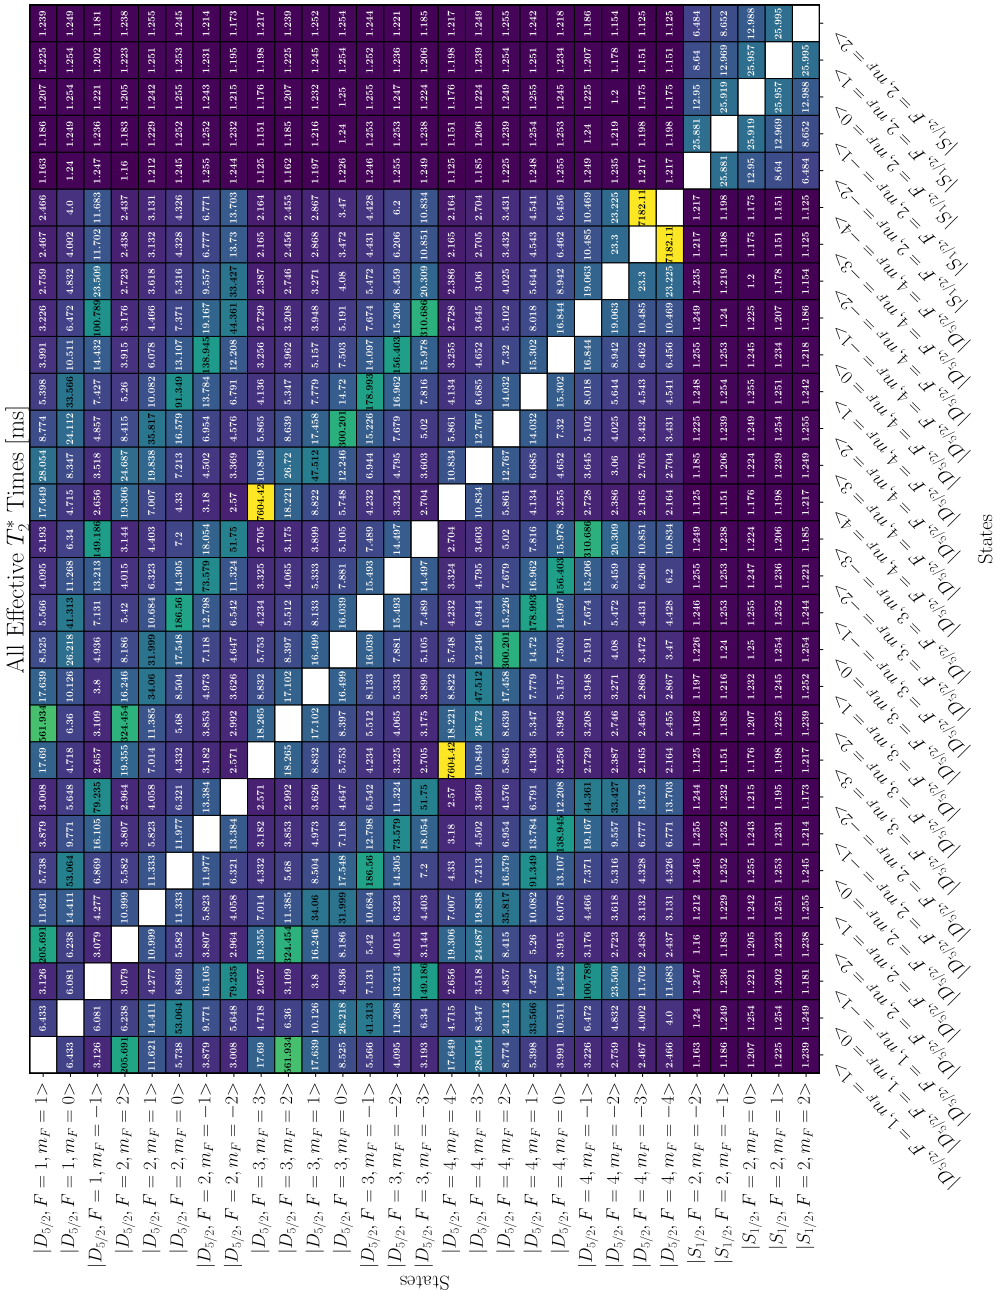


Figure F.2: **All effective coherence times.** The color plot above shows the coherence times between all pairs of qubits, extrapolated from the noise characterization presented in Chapter 6. Notably, direct transitions between $6S_{1/2}$ and $5D_{5/2}$ exhibit significantly shorter coherence times, primarily due to their sensitivity to laser noise.

Appendix G

Bussed qubit Phase scans

The following plots show bussed-qubit phase scans for one of our highest-performing state pairs: $|D_{5/2}, \tilde{F} = 2, m_{\tilde{F}} = 2\rangle \leftrightarrow |D_{5/2}, \tilde{F} = 3, m_{\tilde{F}} = 2\rangle$. These scans also highlight imperfections in the AWG timing, evidenced by a consistent phase shift observed across different wait times. This phase drift cannot be attributed to other noise sources, as the qubit pair is magnetically insensitive and both states reside within the $5D_{5/2}$ manifold, ensuring common-mode rejection of laser phase noise.

Bussed Qubit Ramsey phase scans

$|D_{5/2}, \tilde{F} = 3, m_{\tilde{F}} = 2\rangle \leftrightarrow |D_{5/2}, \tilde{F} = 2, m_{\tilde{F}} = 2\rangle$

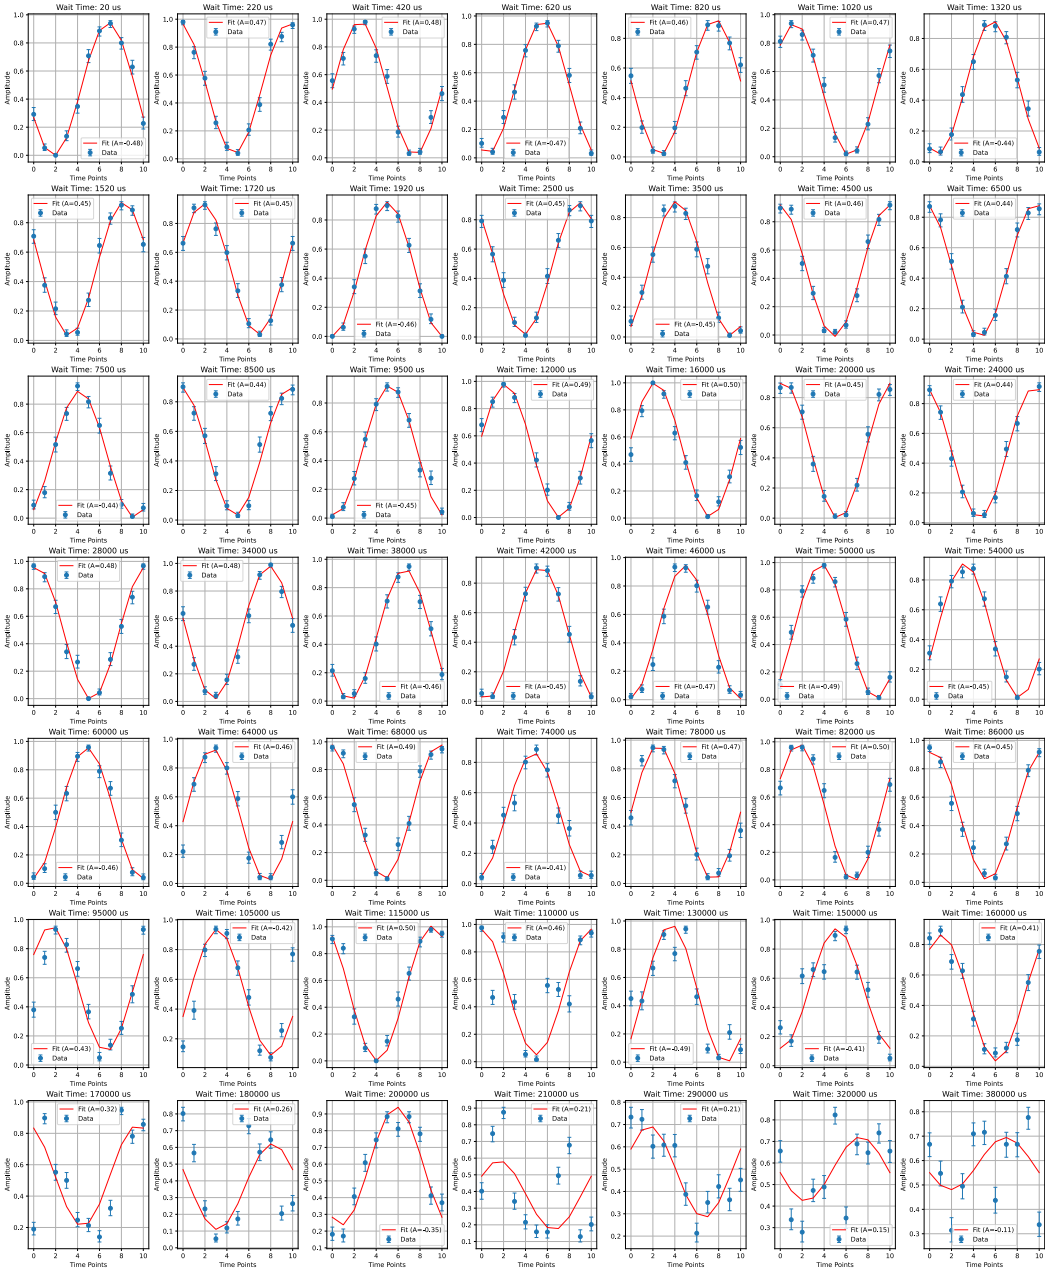


Figure G.1: **Bussed Qubit Phase scans.** The scans shown here are a subset of the many bussed Ramsey scans used to characterize magnetic field noise in our system.