

Systems and Control Protocols for Neutral-Atom-Array Quantum Processors

by

Artem Zhutov

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

Waterloo, Ontario, Canada, 2026

© Artem Zhutov 2026

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner: Dvira Segal
 Professor, Department of Chemistry
 University of Toronto

Supervisors: Alexandre Cooper-Roy
 Research Assistant Professor, Department of Physics & Astronomy
 University of Waterloo

David Cory
Research Professor, Department of Chemistry
University of Waterloo

Internal Member: Kazi Rajibul Islam
 Associate Professor, Department of Physics & Astronomy
 University of Waterloo

Internal Member: Raffi Budakian
 Professor, Department of Physics & Astronomy
 University of Waterloo

Internal Member: Dmitry Pushin
Professor, Department of Physics & Astronomy
University of Waterloo

Internal-External Member: Jonathan Baugh
Professor, Department of Chemistry
University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the 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.

Statement of Contributions

Chapter-by-Chapter Contributions

Chapter 2. I designed the microwave control system and, together with Anastasiia Mashko, developed an in-situ magnetic field correction protocol using atoms as quantum magnetometers. We then conducted site-resolved Ramsey interferometry across a $260\text{ }\mu\text{m} \times 160\text{ }\mu\text{m}$ region with $3\text{ }\mu\text{m}$ spatial resolution. I processed this field map to calculate and apply a compensation field. My technical contributions included IQ mixer calibration. This work identified temporal field fluctuations as the next dominant performance limitation and established procedures now standard in my group.

Chapter 3. I performed hardware-aware analysis of Raman laser systems to guide system design. My comparative evaluation of four wavelengths (420 nm, 421 nm, 780 nm, 795 nm) identified the 795 nm D1 transition as optimal. Simulations predicted a single-qubit gate infidelity of 4.4×10^{-4} using BB1 composite pulses to mitigate thermal motion errors. Together with Anastasiia Mashko, I built and characterized the D1 laser system. My specific contribution was engineering a digital servo error signal filter that suppressed a 150 MHz noise peak to stabilize phase noise.

Chapter 4. I built and characterized a cavity-stabilized Rydberg laser system at 420 nm and 1013 nm and developed two approaches for fidelity improvement. First, I implemented servo optimization using linear response theory to compute the predicted Rydberg excitation infidelity from measured phase noise spectra, enabling rapid PID parameter testing without full gate sequences. Second, I implemented quantum optimal control for pulse design. Building on quantum optimal control simulation code from Jeremy Flannery, I converted the simulation to JAX to enable GPU-accelerated optimization. I benchmarked QOC-optimized pulses against analytical protocols, quantifying a sixfold speedup for single-atom excitation. I performed robustness simulations accounting for finite temperature effects.

Abstract

Neutral atom arrays are a leading platform for programmable quantum processors, offering individual qubit addressability, long-lived hyperfine ground states, and strong Rydberg interactions. Recent progress has demonstrated coherent control over thousands of atoms. However, achieving scalable control requires precise mitigation of environmental and hardware imperfections that degrade gate performance.

This thesis presents an integrated neutral-atom array platform built from the ground up that incorporates quantum sensing directly into the processor. Each atom functions both as a qubit and a local magnetometer. We design, build, and characterize from first principles three subsystems: 1) a microwave control system for driving hyperfine transitions in ground-state rubidium atoms; 2) a Raman laser system for site-selective single-qubit gates; and 3) a Rydberg laser system with quantum optimal control for robust two-qubit gates. This work provides a universal gate set and quantifies which error sources limit performance.

First, we develop an in-situ magnetic field imaging technique using the atom array as a quantum sensor. Through site-resolved Ramsey spectroscopy, we image magnetic fields across a $260\text{ }\mu\text{m} \times 160\text{ }\mu\text{m}$ region with $3\text{ }\mu\text{m}$ spatial resolution. We then apply computed corrections that compensate for the bias magnetic fields, producing uniform global microwave single-qubit rotations.

Second, we introduce a hardware-aware simulation framework to evaluate Raman laser systems for hyperfine qubit manipulation. Simulations predict a single-qubit gate infidelity of 4.4×10^{-4} using BB1 composite pulses to mitigate thermal motion errors. We validate the Raman laser system by building and characterizing its phase noise.

Third, we develop a Rydberg laser system for high-fidelity entangling gates. We apply linear response theory to map laser phase noise to single-atom Rydberg excitation fidelity. We then demonstrate fast phase-noise engineering by optimizing laser servo parameters. We employ hardware-aware quantum optimal control to design both Rydberg excitation and two-qubit gate pulses with built-in robustness against physical and control parameter fluctuations, outperforming analytical benchmarks.

This integrated platform demonstrates high-fidelity universal control of neutral-atom registers with hundreds of qubits. By systematically addressing environmental inhomogeneities through integrated sensing and hardware-aware control design, this work provides a validated path for scaling quantum processors while maintaining gate fidelity.

Acknowledgements

I had the opportunity to build a Rydberg atom array quantum simulator from the ground up—a one-of-a-kind platform in Canada. This was the kind of experience I couldn’t get anywhere else, and I am deeply grateful for it.

I have been fortunate to work with Alex, who taught me what it means to think from first principles. Alex’s approach to science is grounded in rigor and precision. When I was applying to IQC six years ago, Alex stood out as the most proactive supervisor, bringing high energy and an unwavering commitment to avoiding assumptions. I vividly remember his principle that “if the laser is broken, it should be repaired the next day,” and his reminder that “physics is precise.” Beyond technical skills, Alex taught me how to design systems, make decisions in uncertain environments, navigate complex problems, and communicate effectively by always asking “who is your audience?” I am also grateful for the experimentalist training—learning to lay down BNC cables, talk to vendors, code properly, and ask the right questions. Alex brought a standard of production quality to everything we did.

David Cory provided a visionary perspective and allowed me significant freedom to explore. His focus on long-term impact shaped the way I approach problems, and I am grateful for his guidance.

I would like to thank my thesis committee—Dvira Segal, Jonathan Baugh, Kazi Rajibul Islam, Dmitry Pushin, and Raffi Budakian—for their thoughtful questions and valuable perspectives throughout this process.

I want to highlight my undergraduate supervisor, Konstantin Kravtsov, who gave me the opportunity to contribute to research and had a great impact on my early development as a scientist.

Progress is not linear. The day we first saw the microwave transition, Anastasiia, Kent, and I ran Rabi, Ramsey, and Spin Echo experiments all in one day, with Soroush’s support. That was the best day of my PhD.

Anastasiia was the most knowledgeable person about atomic physics in our lab, and working closely with her meant I never got lost when building this machine. Her ability to explain complex physics made a significant difference. Through conversations and lunches with the team, I learned about cool places like the sauna—I fell in love with it.

Kent was a great teacher—not just in the lab but also at the gym, where he taught me how to become big. His broad knowledge, great sense of humor, and calm presence brought stability.

Soroush set the highest bar of excellence in the lab. His ability to prototype absolutely anything from scratch in five seconds was mind-blowing. He had a deep understanding of complex topics and was great at explaining them—he’s probably the only person who could actually read Chapter 7 of Steck. I learned how to be a better developer by watching him and through his careful code reviews of my GitHub pull requests. He established the practice of proper code review, and his speed at thinking, writing code, and implementing ideas made a significant impact. The production quality figures we now make reflect the standard he brought.

Those late nights before conferences making posters were some of the best times I had with the team.

Over the last five years, I worked with over 40 undergraduate students. Each person brought something unique—some helped me write better code, Timur taught me skateboarding tricks, and others brought different perspectives and approaches that I learned from.

Luis made the RAC office much less empty after 6 pm. We had many random conversations, and memorably, he taught me how to ride a motorcycle in a parking lot despite the fact that I had never touched a motorcycle before. Looking back, I’m not sure I would do that again—it was terrifying—but it’s a funny memory now.

George seemed to know everything about the building, and his help was always timely.

Beyond the lab, I am grateful to Andrei, Khariton, and Albert, who played significant roles in my growth. Andrei helped me become more organized, teaching me to treat life as a series of experiments—trying new tools, working from different places, keeping a daily log—and building a knowledge base. Khariton and I had brief interactions, but they were incredibly impactful. He taught me to take responsibility for my life and to apply the scientific method not only in the lab but also to understand what I want from life, my thinking, and my emotions. He showed me how to think from first principles about the resources I have. Albert helped me better understand myself through open questions, taught me self-compassion, and guided me through understanding cognitive biases and becoming more self-aware.

I am grateful to Mike for his support. Every interaction with him gave me a significant boost in energy and clarity. Without his encouragement, I wouldn’t have been able to push through personal blocks, especially when I started creating YouTube videos about building personal AI agents. His passion and enthusiasm have had a meaningful impact on my life.

Jovid took me out of my office on weekends to explore the Kitchener-Waterloo area. Anton Borisov and I had the deepest discussions I’ve ever had—three-hour conversations

completely disconnected from the world, fully engaged with a single topic. Bohdan brought a fresh perspective to every challenge I faced—not just in the lab, but in personal life. His grounded advice was eye-opening. Where I would think about complex solutions, Bohdan would say “you should just do that,” and it was always the best decision.

Finally, I’d like to thank my family for their support.

Now, with a new generation of students joining, I am confident the lab is in good hands.

Large language models (LLMs) were used solely to assist with proofreading, the correction of spelling and grammatical errors, and improving word choice. They were not used to generate content.

Table of Contents

Examining Committee	ii
Author's Declaration	iv
Statement of Contributions	v
Abstract	vi
Acknowledgements	vii
List of Figures	xiv
List of Tables	xvi
1 Introduction	1
1.1 Neutral Atom Arrays for Quantum Information Processing	1
1.1.1 Areas of Impact	1
1.1.2 Physical Realizations of Quantum Processors	2
1.2 Scientific Apparatus	5
1.2.1 Atom Array Preparation	5
1.2.2 Universal Control via Microwave, Raman, and Rydberg Systems . .	6
1.3 Spatial and Temporal Uniformity Challenges at Scale	8
1.3.1 Integrated Control Loop: Characterization, Design, and Feedback .	8
1.4 Thesis Organization	9

2	In-Situ Magnetic Field Homogenization with Atomic Magnetometry	14
2.1	Environmental Inhomogeneity Limits Control of Atomic Arrays	15
2.1.1	Sources of Spatial Inhomogeneity	16
2.1.2	Temporal Fluctuations and Homogeneous Dephasing	16
2.2	Magnetic Field Sensing with Atom Arrays	16
2.2.1	Magnetic Field Sensitivity via Zeeman Effect	16
2.2.2	State Preparation and Measurement	17
2.2.3	Ramsey Spectroscopy for Magnetic Field Imaging	18
2.3	High-Fidelity Microwave Control System	18
2.3.1	Microwave Control System Architecture	19
2.3.2	IQ Mixer Calibration	20
2.3.3	Microwave-Driven Coherent Rabi Oscillations	21
2.4	Experimental Demonstration of Field Homogenization	23
2.4.1	Feedback-Based Magnetic Field Correction Protocol	23
2.4.2	Quantifying Improvements in Magnetic Field Uniformity	24
2.5	Discussion and Outlook	25
3	Developing a Raman Laser System for High-Fidelity Single-Qubit Gates	28
3.1	Simulation Framework for Hardware-Aware Raman System Design	29
3.1.1	Design Requirements for High-Fidelity Raman Control	30
3.1.2	Hardware Architectures for Generating Phase-Coherent Fields	33
3.1.3	Simulation Framework Incorporating Realistic Noise Models	37
3.2	Hardware Architecture Evaluation and Selection	39
3.2.1	Transition Wavelength Selection	40
3.2.2	Comparing OPLL and Modulator-Based Architectures	43
3.2.3	Impact of Laser Phase Noise on Gate Fidelity	43
3.2.4	Impact of RF Source Phase Noise on Gate Fidelity	45
3.2.5	Additional Error Mechanisms	46

3.2.6	Design Recommendations and Performance Comparison	48
3.3	Experimental Implementation and Characterization	49
3.3.1	Raman Laser System Design	49
3.3.2	Phase Noise Characterization	51
3.3.3	Summary and Outlook	52
4	Engineering High-Fidelity Rydberg Gates through Laser Stabilization and Quantum Optimal Control	53
4.1	Two-Photon Rydberg Excitation: Physics and Error Sources	54
4.1.1	Two-Photon Rydberg Excitation Scheme	55
4.1.2	Rydberg Laser System Architecture	56
4.1.3	Hamiltonian for Single and Two-Atom Systems	57
4.1.4	Error Channels in Rydberg Gates	58
4.2	Laser Phase Noise Engineering for Rydberg Excitation	60
4.2.1	Rydberg Laser Frequency Stabilization	60
4.2.2	Mapping Phase Noise to Excitation Infidelity via Filter Functions .	61
4.2.3	Servo Parameter Optimization for Rydberg Excitation	63
4.3	Synthesizing Fast, Robust and High-fidelity Adiabatic Control Operations .	66
4.3.1	Quantum Optimal Control	67
4.4	Application I: Single-Atom Rydberg Excitation	70
4.5	Application II: Two-Qubit Entangling Gates	70
4.6	Summary and Outlook	76
5	Conclusion	79
5.1	Summary	79
5.2	Outlook	80
5.3	Moving Forward	83
	References	84

A	System Hamiltonian for Raman Simulations	93
A.1	General Hamiltonian Structure	93
A.1.1	Five-Level System (D1 Transitions: 795 nm, 421 nm)	94
A.1.2	Six-Level System (D2 Transitions: 780 nm, 420 nm)	94
A.2	OPLL Hamiltonian	94
A.3	MZM Hamiltonian	95
B	Raman Gate Simulation Framework	96
B.1	System Hamiltonian	96
B.2	Monte Carlo Simulation of Phase Noise	97
B.3	Noise Incorporation in Optical Phase-Locked Loop Systems	97
B.4	Noise Incorporation in Modulator-Based Systems	98
	Glossary	99

List of Figures

1.1	Quantum Simulation Framework	3
1.2	Magneto-Optical Trap	5
1.3	Single-Atom Trapping and Readout	6
1.4	Control Loop: Characterization, Design, and Feedback	9
1.5	Conceptual Framework for Achieving Scalable Universal Quantum Control	10
1.6	Functional Architecture Mapping Requirements to Operational Capabilities	11
1.7	Physical Systems Implementation of the Neutral Atom Platform	12
2.1	Qubit System and State Preparation in ^{87}Rb	17
2.2	Ramsey Pulse Sequence	18
2.3	Microwave Control System Architecture	20
2.4	IQ Mixer Calibration	21
2.5	Global Microwave Control of Hyperfine Qubit	22
2.6	Iterative Gradient Compensation Protocol	24
2.7	Modulated Ramsey Magnetometry with Diagonal Binning	25
2.8	Magnetic Field Gradient Compensation	26
2.9	Magnetic Field Uniformity Quantification	26
3.1	^{87}Rb Energy Levels and Raman Transition Scheme	31
3.2	Hardware Architectures for Phase-Coherent Raman Fields	34
3.3	Fidelity Map for MZM-Driven Raman Transition	36

3.4	Optical Spectra Comparison: Phase vs Mach-Zehnder Modulators	37
3.5	Simulated Raman Fidelities for Varying RF and Laser Specs	44
3.6	ECDL Frequency Noise Spectrum	47
3.7	OPLL System Architecture for Raman Gates	50
3.8	Raman Laser System Phase Noise	51
4.1	Rydberg Laser System	55
4.2	Rydberg Excitation Energy Level Diagram	56
4.3	Rydberg Laser Servo Control Architecture	62
4.4	Systematic PID Gain Optimization	64
4.5	Impact of PID Gain Fine-Tuning on Laser Noise Spectra	65
4.6	Predicted Rydberg Excitation Infidelity from Optimized Phase Noise	66
4.7	Quantum Optimal Control Framework	67
4.8	High-Fidelity Single-Atom Rydberg Excitation	71
4.9	Robustness of Single-Atom Excitation Protocols	72
4.10	Optimization of Basis Dimension for CRAB Algorithm	72
4.11	CZ Gate Pulse Profiles	74
4.12	Performance Comparison of CZ Gate Implementations	74
4.13	Two-Qubit CZ Gate Robustness	75
4.14	Impact of Finite Temperature on CZ Gate Fidelity	75

List of Tables

3.1	Simulation Operating Points	40
3.2	Simulated Gate Infidelity Comparison by Architecture and Wavelength . .	41
3.3	Candidate Raman Transitions	42

Chapter 1

Introduction

This chapter introduces the neutral atom quantum computing platform and frames the central question this thesis addresses: achieving high-fidelity universal quantum control in arrays with hundreds of individually controlled qubits. Section 1.1 motivates quantum computing and positions neutral atom arrays among competing physical implementations. Section 1.2 describes scientific apparatus used in this thesis to perform the experiments. Section 1.3 frames the scaling problem and outlines our integrated solution. Finally, Section 1.4 previews the thesis structure and contributions.

1.1 Neutral Atom Arrays for Quantum Information Processing

1.1.1 Areas of Impact

Quantum information processing can be broadly classified into three areas: quantum computing, quantum sensing and metrology.

Quantum computing enables efficient solutions to specific problems intractable for classical computers, including integer factorization (Shor’s algorithm [1]) and quantum simulation [2], where classical computational cost scales exponentially with system size. Modeling quantum many-body physics represents one such problem. The computational cost for classical simulation scales exponentially with particle number and entanglement [2, 3].

Quantum simulation [4] uses controlled quantum systems to model these intractable systems, enabling investigation of many-body physics such as quantum magnetism and phase

transitions. Fig. 1.1 illustrates the quantum simulation framework, where an encoding mapping φ establishes equivalence between the target material system and its realization in a programmable neutral atom array platform.

Fault-tolerant computation (FTQC) overcomes the inherent fragility of quantum information and enables reliable algorithms requiring thousands of qubits. FTQC employs Quantum Error Correction (QEC) [6, 7, 8] to encode robust logical qubits from many physical qubits, requiring physical error rates below a threshold (typically 10^{-3} to 10^{-4}).

Quantum sensing and metrology employ controlled quantum systems to measure physical quantities such as time, frequency, and magnetic fields with precision beyond classical limits [9, 10, 11]. By using quantum coherence and entanglement, these approaches can reach or surpass the standard quantum limit [10] that constrains classical sensors. Atomic clocks are a prominent example, where entangled or spin-squeezed atomic ensembles improve time-measurement stability [12, 13].

1.1.2 Physical Realizations of Quantum Processors

Quantum information processing can be realized in multiple physical systems, each with distinct advantages and limitations. This section provides a brief overview of competing platforms and motivates the choice of neutral atom arrays for this work.

It is an exciting time in physics, fueled by the intense interest in quantum information processing and rapid experimental progress. For the first time, we can build and control systems involving hundreds or thousands of quantum particles, allowing us to pursue the realization of quantum advantage, the capability of solving problems intractable for classical computers. This quest has given rise to several competing physical platforms, each with its own set of strengths and limitations.

Superconducting qubits use superconducting resonant circuits operating at cryogenic temperatures. They feature fast gate operations and have demonstrated quantum supremacy in certain computational tasks [14]. However, coherence times are limited and achieving high connectivity between distant qubits requires coupling architectures. Trapped ion systems confine ions in a vacuum and manipulate them with electromagnetic pulses. Trapped ion qubits demonstrate high-fidelity gates and long coherence times [15]. However, scaling to large numbers of qubits is limited by multi-zone trap technologies and managing interactions in long chains. Other platforms include photonic systems using linear optics,




Figure 1.1: Quantum simulation framework. The encoding mapping φ relates initial states between the target material system (left, Hamiltonian H_S) and the neutral atom platform (right, Hamiltonian H_{eff}). After time evolution under their respective Hamiltonians, the decoding mapping φ^{-1} extracts material properties from platform measurements. Adapted from [5].

semiconductor quantum dots using electron spins, and nitrogen-vacancy (NV) centers in diamond. Each of these platforms presents unique advantages and technological constraints.

Among these competing platforms, neutral atom arrays offer a scalable platform for quantum information processing due to several distinct advantages:

Scalability and reconfigurability represent a distinct capability of neutral atom arrays. Arrays up to 256 **qubits** in analog mode are commercially available [16], and experimental systems scaling to 6,100 atoms in approximately 12,000 trapping sites have been realized [17]. Arrays of atoms trapped in optical tweezers, generated by spatial light modulators (SLMs) and moved by acousto-optic deflectors (AODs), provide dynamically reconfigurable geometries with arbitrary connectivity. Real-time rearrangement of atoms compensates for stochastic loading to produce filled arrays in programmable geometries.

Long coherence times represent another distinct advantage. **Qubits** are encoded in atomic hyperfine ground states, which are robust against decoherence and well-isolated from the environment, enabling long-lived quantum memories. Hyperfine **qubit** coherence times of tens of seconds are typical in optical tweezer arrays, with record coherence times of 12.6(1) s demonstrated [17].

Strong, tunable interactions are introduced by coherent coupling to highly excited Rydberg states, which exhibit strong, long-range interactions. Rydberg blockade [18] prevents simultaneous excitation of neighboring atoms, enabling fast entangling gates. Two-qubit controlled phase (CZ) gates have been implemented with fidelities up to 99.5% in parallel on up to 60 atoms [19], approaching the surface code threshold for error correction [8].

Finally, individual optical addressing allows site-selective operations required for multi-qubit gate sequences. Optical techniques allow precise control and readout of individual atoms within the array, providing site-selective rotations and entangling operations required by quantum algorithms.

Despite these advantages, neutral atom arrays face constraints. Rydberg interactions are sensitive to laser frequency noise, and maintaining uniform magnetic field environments across large arrays is demanding. This thesis addresses these limitations through integrated control, stabilization, and sensing protocols.

While neutral atom arrays offer compelling advantages, realizing this potential presents a central question: **How do we achieve high-fidelity quantum control in neutral atom arrays given environmental and control field inhomogeneities?** Magnetic field fluctuations, spatial intensity variations, and phase noise in control fields degrade gate fidelity and coherence. This thesis develops an integrated approach to characterize and mitigate these inhomogeneities. The following sections describe our specific experimental platform and outline the systematic solution.

1.2 Scientific Apparatus

This section introduces the scientific apparatus that serves as the testbed for all subsequent work. We detail the atom array formation, qubit encoding, and the three integrated control systems that enable universal quantum control. All experiments in this thesis employ this custom-built apparatus, which uses ^{87}Rb atoms in optical tweezers to form qubit arrays.

1.2.1 Atom Array Preparation

Single ^{87}Rb atoms are prepared through a multi-stage cooling and trapping process. Atoms are first cooled in a 2D+ magneto-optical trap (MOT), where a push beam directs them into a glass science cell under ultra-high vacuum through a pinhole. A three-dimensional MOT loads 10^7 atoms in 100 ms using two coils in an anti-Helmholtz and anti-Helmholtz coils configuration that generate a magnetic field gradient of 10 G/cm. After loading, PGC is performed for 30 ms, further reducing the atomic temperature to tens of microkelvin. Fig. 1.2 shows the continuously loading 3D MOT in the glass cell.




Figure 1.2: Continuously loading 3D magneto-optical trap in glass science cell. The bright red fluorescence shows trapped ^{87}Rb atoms during MOT loading.

Atoms from this cooled cloud are loaded stochastically into a programmable array of optical tweezers generated by a Ti:Sapphire laser (M Squared Lasers) at 813.4 nm. This wavelength balances two competing requirements. Proximity to the D2 resonance reduces the power needed to achieve a given trap depth, while sufficient detuning minimizes

off-resonant scattering and associated heating. The off-resonant scattering rate at this wavelength is ≈ 75 Hz. Two acousto-optic deflectors (AODs) multiplex the laser beam to provide dynamic spatial control of trap positions, enabling creation of arbitrary square grid geometries and transport of individual atoms. A pair of high-resolution microscope objectives from Mitutoyo ($\text{NA} = 0.5$, $f = 4$ mm) performs high-resolution imaging and beam focusing. The focused laser beam has a waist diameter of $w_0 = 0.82 \mu\text{m}$. Trap depths measured via light-shift spectroscopy are typically 1.0 mK.

Atoms load stochastically into the traps with a loading probability $\eta = 0.53$ (Fig. 1.3). Each trap contains at most one atom. The array is imaged by collecting atomic fluorescence, achieving high-fidelity single-atom readout. Our system demonstrates arrays with up to 1305 traps spanning a $260 \mu\text{m} \times 160 \mu\text{m}$ region (Fig. 1.3).




Figure 1.3: Single-atom trapping and readout. (a) Average fluorescence image of a 45×29 trap array (1305 traps) with a loading probability of 53% showing individual trapped atoms as bright spots. (b) Single-shot image of atom qubit register. This single-shot image demonstrates stochastic loading from optical tweezers.

1.2.2 Universal Control via Microwave, Raman, and Rydberg Systems

Qubits are encoded in long-lived hyperfine ground states of ^{87}Rb $5S_{1/2}$ manifold. The specific hyperfine states depend on the type of experiment performed. For magnetic field sensing (Chapter 2), we use [magnetically sensitive transition](#) states $|F = 1, m_F = -1\rangle$ and

$|F = 2, m_F = -2\rangle$, whose transition frequency shifts linearly with magnetic field, providing a sensitive magnetometer. For control operations (Chapter 3), we instead encode qubits in magnetically insensitive [clock states](#) $|F = 1, m_F = 0\rangle$ and $|F = 2, m_F = 0\rangle$ to minimize unwanted frequency shifts during gate operations and enhance coherence times.

To achieve universal quantum control, this thesis develops three integrated physical systems:

Microwave System for Global Manipulation of Hyperfine Qubit. The microwave system performs global single-qubit rotations across the entire array. Microwave pulses at 6.834 GHz drive these rotations. These fields couple directly to the hyperfine transition, enabling uniform array-wide operations that prepare coherent superposition states across all qubits simultaneously. Achieving spatial uniformity across the large array is difficult because magnetic field gradients introduce site-dependent frequency shifts that degrade gate fidelity, as quantified in Chapter 2. Chapter 2 demonstrates in-situ magnetometry using the atoms themselves as distributed magnetometers to map and correct these inhomogeneities.

Raman System for Manipulation of Hyperfine Qubit. In contrast to global microwave control, optical Raman transitions enable site-selective addressing of individual qubits. A two-photon Raman transition at 795 nm drives single-qubit gates. Two laser beams with a frequency difference matching the hyperfine splitting drive coherent transitions through a virtual excited state. Phase coherence between these beams is maintained using an [OPLL](#), an active feedback system that locks the relative phase of two independent lasers by mixing their outputs and stabilizing the beat note. This architecture minimizes phase noise that would otherwise limit gate fidelity. Chapter 3 develops and validates a noise-aware design approach that identifies this [OPLL](#) architecture and characterizes its performance.

Rydberg Laser System for Two-Qubit Entangling Gates. This system performs two-qubit entangling gates. Two-photon Rydberg excitation using lasers at 420 nm and 1013 nm produces strong interactions between atoms. Atoms excited to high-lying Rydberg states have large electric dipole moments, leading to strong van der Waals interactions [18]. To suppress laser frequency noise, both lasers are frequency-locked to a [ULE cavity](#) [20, 21] using the [PDH](#) technique, a reference cavity with minimal thermal expansion that provides a stable optical frequency reference. Chapter 4 achieves robust, high-fidelity control operations through systematic stabilization of this laser system and optimal control protocols.

These systems provide universal quantum control over hundreds of individual qubits. However, maintaining high fidelity as the array scales to hundreds of qubits requires en-

environmental stabilization and noise resilience. The following section examines these challenges and outlines our integrated solution.

1.3 Spatial and Temporal Uniformity Challenges at Scale

Neutral atom arrays offer a clear path to systems with thousands of qubits, but realizing this potential requires managing spatial and temporal inhomogeneities. Control becomes more challenging as neutral atom arrays scale to larger sizes because both spatial and temporal uniformity are fundamentally harder to maintain. Spatially, the optical and magnetic fields that define and manipulate the qubits cannot be perfectly homogeneous over an extended region. Variations in laser intensity and phase lead to position-dependent light shifts, differences in Rabi coupling strengths, and site-dependent detunings, causing non-uniform qubit rotations and frequency shifts across the array.

Temporally, maintaining phase-stable and amplitude-stable control fields becomes more demanding as the number of beams, optical paths, and control channels increases. Laser phase noise, optical path drift, and magnetic field fluctuations introduce correlated errors that affect all qubits simultaneously. These limitations arise because electromagnetic fields vary in space and fluctuate in time, making it increasingly difficult to apply identical, coherent control across thousands of atoms. The resulting spatial and temporal inhomogeneities directly degrade gate fidelity and coherence, constraining the scalability of large neutral atom processors.

The solution developed in this thesis is an **integrated neutral-atom platform** that unifies quantum control and sensing within a single system. This thesis develops magnetic field homogenization, laser phase noise optimization, and hardware-aware control design to enable robust, high-fidelity quantum operations. This integrated approach enables scalable, noise-resilient quantum information processing.

1.3.1 Integrated Control Loop: Characterization, Design, and Feedback

The platform architecture is best understood through multi-level decomposition, illustrated in Figs. 1.4–1.7. We begin with the high-level control feedback loop (Fig. 1.4), which shows how environmental characterization, control design, and quantum operations form a closed-loop system. We then decompose system requirements (Fig. 1.5), map these to functional capabilities (Fig. 1.6), and finally detail physical implementations (Fig. 1.7).




Figure 1.4: The control framework cycles through environmental characterization (measurements on the atom array), hardware-aware control protocol design, and quantum operations. Measurement feedback continuously refines control strategies, closing the loop between understanding the environment and optimizing quantum operations.

This progression from abstract requirements to concrete hardware structures the thesis contributions.

Achieving universal control requires translating platform advantages into robust operational capabilities. While neutral atoms offer long coherence times, maintaining phase coherence across hundreds of spatially distributed qubits under realistic noise remains challenging. The following chapters develop systematic approaches to address these implementation challenges through environmental stabilization and hardware-aware control design.

1.4 Thesis Organization

Each subsequent chapter presents a key stage in the platform’s development, which functions as an integrated quantum processor with built-in quantum sensing capabilities. Each chapter addresses a specific scientific gap and provides experimental evidence to validate the central hypothesis.

Chapter 2: Correcting Magnetic Field Inhomogeneity Using an In-Situ Atomic Magnetometer. This chapter addresses the problem of static magnetic field inhomogeneities limiting global quantum coherence across large arrays. We develop an in-situ magnetic field imaging technique, using the atom array itself as a distributed quantum sensor network. Using site-resolved Ramsey spectroscopy, we image magnetic fields across a $260\ \mu\text{m} \times 160\ \mu\text{m}$ region with $3\ \mu\text{m}$ spatial resolution. We then apply computed corrections to anti-Helmholtz compensation coils (Chapter 2). This produces a spatially uniform magnetic environment that improves ensemble coherence time and the fidelity of control operations. Management of the magnetic environment enables the high-fidelity




Figure 1.5: Conceptual framework for achieving scalable universal quantum control. This diagram presents the conceptual framework for addressing the central thesis question: how to achieve high-fidelity universal control as system size scales. The top-level goal decomposes into two primary requirements: (1) implementing a universal gate set, and (2) ensuring robust, high-fidelity control at scale. The universal gate set requirement is satisfied by developing both single- and two-qubit gates. Single-qubit operations are implemented through a global microwave control system (Chapter 2) and a site-selective Raman laser system (Chapter 3). High-fidelity two-qubit entangling gates are realized using a stabilized Rydberg laser system (Chapter 4). The requirement for robust, high-fidelity control at scale is addressed through environmental stabilization and the development of robust control protocols. Environmental stabilization is achieved by actively compensating for magnetic field gradients (Chapter 2) and stabilizing laser control fields via an optical phase-lock loop (OPLL) architecture (Chapter 3). Robustness is further enhanced by an infidelity-based laser servo optimization method (Chapter 4) and hardware-aware quantum optimal control (QOC) for designing noise-resilient pulses (Chapter 4). Each branch illustrates the specific functional capabilities developed to satisfy these core requirements, mapping directly to the contributions presented in this thesis.




Figure 1.6: Functional architecture mapping requirements to operational capabilities. This diagram maps the conceptual requirements from Fig. 1.5 to the specific operational capabilities developed to satisfy them. The requirement for a universal gate set is addressed through distinct capabilities. Full single-qubit control is achieved through two complementary methods: global single-qubit gates via microwave fields for uniform array-wide operations, and site-selective single-qubit gates via optical addressing for individual qubit manipulation. The requirement for two-qubit entanglement is met by implementing two-qubit entangling gates via Rydberg interaction. The requirement for robust, high-fidelity control at scale is satisfied by engineering a uniform qubit environment and stable control fields. First, qubit environment uniformization is achieved via magnetic field homogenization to correct for static spatial magnetic field variations. Second, control field stabilization is performed through laser phase noise optimization to reduce temporal fluctuations in the laser frequency. Finally, noise-resilient gate protocols are developed using a quantum optimal control (QOC) framework to make operations robust against residual noise.




Figure 1.7: Physical systems implementation of the neutral atom platform. This diagram maps the functional capabilities from Fig. 1.6 to their physical hardware and software implementations. The capability for magnetic field homogenization (Chapter 2) is realized using a microwave control system to perform measurements on the atom array, which acts as a quantum sensor, along with a set of anti-Helmholtz compensation coils. The capability for site-selective single-qubit gates via optical addressing (Chapter 3) is implemented with two optically phase-locked Raman ECDL lasers and acousto-optic deflectors (AODs) for spatial addressing of the individual qubits. The capability for two-qubit entangling gates via Rydberg interaction (Chapter 4) is enabled by a Rydberg laser system at 420 nm and 1013 nm, frequency-locked to a ULE stabilization cavity. Individual two-qubit pairs entanglement is enabled through holographic beam shaping with a spatial light modulator. Finally, the quantum optimal control (QOC) framework (Chapter 4) incorporates measured noise spectra to design hardware-aware pulses. These systems provide global single-qubit gates, site-selective single-qubit gates, and site-selective two-qubit gates. This platform enables the assembly and manipulations of quantum systems with hundreds of individual qubits, enabling exploration of many-body quantum phenomena and discovery of physics previously beyond experimental reach.

site-selective operations and entangling gates developed in subsequent chapters.

Chapter 3: Developing a Noise-Aware Raman Laser System for High-Fidelity Single-Qubit Gates. This chapter addresses achieving high-fidelity, site-selective single-qubit rotations despite laser phase noise and architectural constraints. We introduce a hardware-aware simulation framework to evaluate different hardware architectures for hyperfine qubit manipulation via Raman transition. We quantitatively evaluate each architecture based on laser and RF synthesizer noise spectra. This approach identifies a Raman system with predicted single-qubit gate infidelity of 4.4×10^{-4} using BB1 composite pulses to mitigate thermal motion errors. This result provides quantitative justification for this architecture choice for high-fidelity site-selective control. We then assembled this system in the lab and performed the phase locking of the lasers.

Chapter 4: Engineering High-Fidelity Rydberg Gates through Laser Stabilization and Quantum Optimal Control. This chapter addresses the susceptibility of Rydberg entangling gates and Rydberg excitation protocols to control and physical imperfections. First, we perform systematic optimization of a cavity-stabilized 420/1013 nm Rydberg laser system. This section presents a fast servo optimization method using target control protocol fidelity as a metric. We perform in-loop optimization of the laser phase noise and suppress the noise at frequencies most detrimental to Rydberg excitation fidelity. To guide this hardware optimization, we apply linear response theory to map laser phase noise to Rydberg excitation infidelity, enabling fast phase noise engineering to minimize infidelity for a target excitation protocol. Second, the chapter develops a quantum optimal control (QOC) framework to synthesize hardware-aware, robust control pulses. This framework generates time-dependent control fields that maintain high fidelity under realistic physical parameter fluctuations. We demonstrate this approach on two applications: single-atom Rydberg excitation and two-qubit controlled-Z gates, achieving higher fidelity at similar levels of robustness compared to analytical benchmark protocols.

Chapter 2

In-Situ Magnetic Field Homogenization with Atomic Magnetometry

High-fidelity parallel quantum gates require that control fields address all qubits on-resonance. Environmental magnetic field inhomogeneities produce site-dependent Zeeman shifts that prevent this uniform addressing, degrading gate fidelity and limiting coherence times [9].

As neutral atom arrays scale toward thousands of sites [17], maintaining uniform control grows more challenging. The larger physical size of the register requires distributing laser power uniformly over extended areas, while longer rearrangement times and increased crosstalk between addressing beams further limit operations [22]. Uniform control fields across the entire array are essential for high-fidelity parallel quantum gates.

Residual magnetic field gradients produce site-dependent Zeeman shifts. These shifts degrade the fidelity of quantum operations and limit coherence times. Qubits across the array operate at different resonance frequencies, preventing uniform parallel control. Atomic magnetometry can image magnetic fields [23, 24]. However, previous work has not integrated these measurements into active correction systems.

Here we demonstrate a method for imaging and correcting magnetic field gradients using atomic qubits as sensors. The approach uses the atoms themselves as a spatially-resolved sensor array. The procedure comprises three stages. First, site-resolved Ramsey spectroscopy measures local transition frequencies from which we extract magnetic field values. Second, a linear fit to the spatial field profile extracts the gradient components. Third, iterative feedback suppresses the measured gradients by adjusting gradient coil currents. This homogenizes the magnetic field environment of the qubits.

The protocol employs Ramsey interferometry on the magnetically sensitive $|F = 1, m_F = -1\rangle \rightarrow |F = 2, m_F = -2\rangle$ transition in ^{87}Rb , mapping field distributions across arrays of up to 1305 optical tweezers with $3\text{ }\mu\text{m}$ spatial resolution. We implement this using a custom microwave control system with IQ modulation, achieving carrier suppression of 57.6 dB and image sideband suppression of 50.7 dB.

The compensation reduces field inhomogeneity $\sigma_{\delta B}$ from 0.410 mG to 0.115 mG, improving the ensemble-averaged gate fidelity limit from 99.9% to 99.99%. This establishes a uniform magnetic environment enabling the parallel Rydberg entangling gates demonstrated in Chapter 4.

2.1 Environmental Inhomogeneity Limits Control of Atomic Arrays

Parallel quantum gates require all qubits in the array to experience identical control fields. Spatial magnetic field gradients produce site-dependent Zeeman shifts, creating position-dependent resonance frequencies. This inhomogeneous dephasing degrades gate fidelity and limits coherence times. Temporal fluctuations add homogeneous dephasing affecting all atoms identically, reducing shot-to-shot consistency of quantum evolution.

To quantify the uniformity requirements, we calculate the ensemble-averaged process fidelity [25] for a π -pulse applied to atoms with spatially varying resonance frequencies. We model the spatial field distribution as Gaussian with spread σ_{Δ} in detuning. For the $|F = 1, m_F = -1\rangle \rightarrow |F = 2, m_F = -2\rangle$ transition (Section 2.2.1), the Zeeman sensitivity is 2.10 MHz/G [26]. Achieving $\bar{F} \geq 0.999$ requires magnetic field uniformity $\sigma_{\delta B} \lesssim 0.65\text{ mG}$ (equivalently $\sigma_{\Delta} \lesssim 1.4\text{ kHz}$). For $\bar{F} \geq 0.9999$, $\sigma_{\delta B} \lesssim 0.2\text{ mG}$ ($\sigma_{\Delta} \lesssim 0.4\text{ kHz}$) is required.

Initial measurements revealed spatial variation in qubit resonance frequencies across the atomic array. We systematically characterized potential sources: optical AC Stark shifts from trapping beams, microwave amplitude variations, and magnetic field gradients. We identified magnetic field inhomogeneity as the dominant limitation in our system, motivating the development of the spatially-resolved magnetometry and compensation framework described in this chapter.

The qubit transition exhibits a linear Zeeman shift, making resonance frequency proportional to local magnetic field. Spatial magnetic field variations produce site-dependent frequency offsets. We now examine the physical origins of these spatial and temporal field variations.

2.1.1 Sources of Spatial Inhomogeneity

Static magnetic field inhomogeneities originate from two sources. First, magnetic coils produce gradients. Helmholtz coils for the static quantization field misalign with the atomic array. Anti-Helmholtz coils for magneto-optical trapping generate strong gradients (~ 10 G/cm), leaving residual fields after imperfect compensation.

Second, optical AC Stark shifts from trapping beams contribute to spatial inhomogeneity. Trap intensity variations across the array, combined with differential polarizability between hyperfine states, produce position-dependent frequency offsets. Additionally, imperfect linear polarization can induce effective magnetic fields via vector polarizability [27].

2.1.2 Temporal Fluctuations and Homogeneous Dephasing

Temporal magnetic field fluctuations cause homogeneous dephasing. AC power line noise at 60 Hz and harmonics couples to the magnetically sensitive qubit transition, producing shot-to-shot variation in the magnetic field amplitude.

The magnetic fields in our setup are generated by current-carrying coils powered through supplies powered from the 60 Hz mains. Residual ripple and fluctuations at the mains frequency and its harmonics modulate the coil current amplitude and offset. These modulations produce time-dependent magnetic fields at the atoms' position, causing shot-to-shot variations in the Zeeman shift of the qubit transition frequency. The resulting frequency jitter contributes to ensemble dephasing and limits coherence time T_2^* .

2.2 Magnetic Field Sensing with Atom Arrays

Gradient correction requires spatial field measurements with resolution matching qubit spacing ($3\ \mu\text{m}$). External magnetometers achieve millimeter-scale resolution [23], insufficient for resolving site-to-site variations. We employ the atoms themselves as sensors using site-resolved Ramsey interferometry.

2.2.1 Magnetic Field Sensitivity via Zeeman Effect

We use individual ^{87}Rb atoms trapped in an optical tweezer array as local probes of the magnetic environment. The qubit is defined within the $5S_{1/2}$ ground state manifold of ^{87}Rb , using the hyperfine states $|0\rangle \equiv |F=1, m_F=-1\rangle$ and $|1\rangle \equiv |F=2, m_F=-2\rangle$.

This stretched-state transition contrasts with the clock transition ($|F = 1, m_F = 0\rangle \rightarrow |F = 2, m_F = 0\rangle$), which has negligible first-order Zeeman sensitivity. Comparing measurements on both transitions isolates magnetic from non-magnetic noise sources.

2.2.2 State Preparation and Measurement

We initialize atoms using two-stage optical pumping (Figure 2.1). Coarse pumping with $F = 2 \rightarrow F' = 2$ light drives population toward the $F = 1$ manifold. Fine pumping with σ^+ -polarized $F = 1 \rightarrow F' = 2$ light transfers population into $|0\rangle \equiv |F = 1, m_F = -1\rangle$.




Figure 2.1: Qubit System and State Preparation in ^{87}Rb . Energy levels and optical pumping scheme for the hyperfine qubit. (a) The energy level structure of the $5S_{1/2}$ ground state manifold shows the qubit states $|0\rangle \equiv |F = 1, m_F = -1\rangle$ and $|1\rangle \equiv |F = 2, m_F = -2\rangle$. A static magnetic field creates a linear Zeeman shift, making the transition frequency sensitive to the local field. (b) Coarse optical pumping with light resonant with the $F = 2 \rightarrow F' = 2$ transition drives population towards the $F = 1$ manifold. (c) Fine optical pumping with σ^+ -polarized light resonant with the $F = 1 \rightarrow F' = 2$ transition translates population into the target $|F = 1, m_F = -1\rangle$ state, achieving high-fidelity initialization.

We adiabatically ramp trap depths from ~ 1 mK to $\lesssim 1$ μK , reducing AC Stark shifts from ~ 50 kHz to ~ 500 Hz. We measure atomic temperatures of ~ 24 μK using adiabatic ramp survival probability [28].

After coherent manipulation via microwave pulses, we use state-selective pushout and fluorescence imaging for measurement. The imaging system achieves 99.98% classification fidelity with an atom loss probability of 0.12% per 75 ms imaging cycle. This corresponds to an atom lifetime under imaging of 66.7 ± 0.02 s.

2.2.3 Ramsey Spectroscopy for Magnetic Field Imaging

Magnetic field inhomogeneity creates a spatially-dependent frequency landscape across the array. The linear Zeeman effect shifts each qubit’s resonance frequency in proportion to the local field strength. Ramsey interferometry measures these frequency shifts to construct a magnetic field image.

Ramsey spectroscopy encodes local frequency offsets as measurable phase differences. Two $\pi/2$ pulses bracket a free evolution interval T_{Ramsey} (Fig. 2.2). The first pulse creates superposition $(|0\rangle + |1\rangle)/\sqrt{2}$, after which each qubit precesses at its local resonance frequency. The relative phase between qubit and drive evolves as $2\pi\Delta\nu(r)T_{\text{Ramsey}}$, where $\Delta\nu(r)$ is the detuning at position r . The second pulse applies with phase $2\pi\nu_{\text{mod}}T_{\text{Ramsey}}$ (with $\nu_{\text{mod}} = 15$ kHz), converting accumulated phase to population:

$$P_{|1\rangle}(r, T_{\text{Ramsey}}) = \frac{1}{2} (1 + C \cos [2\pi (\Delta\nu(r) + \nu_{\text{mod}}) T_{\text{Ramsey}}]), \quad (2.1)$$

where C is fringe contrast.

We scan T_{Ramsey} from 0 to 1000 μs . We perform Fourier analysis to extract oscillation frequency at each site: $\nu_{\text{peak}}(r) = \nu_{\text{mod}} + \Delta\nu(r)$. We convert detuning $\Delta\nu(r)$ to local magnetic field using the transition’s Zeeman coefficient [26].




Figure 2.2: Ramsey Pulse Sequence. Two $\pi/2$ pulses separated by free evolution time τ . The second pulse includes phase modulation $2\pi\nu\tau$ to encode frequency information.

Having described the magnetometry protocol, we now present the microwave control system that implements these measurements.

2.3 High-Fidelity Microwave Control System

Ramsey magnetometry requires amplitude uniformity and spectral purity. Amplitude variations create site-dependent Rabi frequencies, preventing simultaneous $\pi/2$ rotations at all

sites introducing the pulse area error. Spectral impurities (carrier leakage, image sidebands) from the imperfect IQ mixing induce off-resonant excitation, causing AC Stark shifts that degrade fidelity.

This section presents the microwave control system in three parts. Section 2.4.1 presents the IQ modulation architecture and shows how baseband I/Q signals implement arbitrary single-qubit rotations. Section 2.4.2 presents the calibration procedure for suppressing spurious spectral components to better than 50 dB. Section 2.4.3 validates the system by demonstrating coherent control of qubits across the array with driving global Rabi oscillations.

We implement IQ modulation combining phase-locked frequency stability with arbitrary pulse shaping. The system achieves carrier suppression of 57.6 dB, lower sideband suppression of 50.7 dB, and spatial Rabi frequency variation $< 1\%$ (35.55 ± 0.13 kHz across the array).

2.3.1 Microwave Control System Architecture

The microwave control system (Figure 2.3) delivers pulses with programmable amplitude and phase to the atoms. IQ modulation upconverts control waveforms from an AWG to the 6.834 GHz hyperfine transition frequency. A bandpass filter suppresses carrier leakage and the image sideband. A switch provides sub-microsecond timing, and a 45 dB amplifier delivers 10 W to a horn antenna.

The system operates phase-coherently at 6.834 GHz, delivering 10 W to the antenna with arbitrary control over amplitude, frequency, and phase.

The control architecture uses IQ modulation scheme to generate arbitrary shaped microwave pulses. A NI QuickSyn FSW-0020 synthesizer generates local oscillator (LO) at 6.734 GHz, phase-locked to an SRS FS725 10 MHz rubidium standard. The LO drives a Marki Microwave MMIQ-0416LS IQ mixer modulated by in-phase (I) and quadrature (Q) signals from an Arbitrary Waveform Generator. Upconversion translates baseband envelopes to microwave domain. A Quntic PMI 5CL6D8-100-CD-SMM cavity bandpass filter suppresses sidebands and LO leakage.

IQ modulation implements arbitrary single-qubit control. In the rotating frame, qubit-field interaction follows

$$H = \frac{\hbar\Omega(t)}{2} [\cos(\phi(t))\sigma_x + \sin(\phi(t))\sigma_y] \quad (2.2)$$

where $\Omega(t)$ is Rabi frequency and $\phi(t)$ is phase. The mixer generates fields with amplitude and phase determined by $I(t) = \Omega(t) \cos(\phi(t))$ and $Q(t) = \Omega(t) \sin(\phi(t))$. Programming $I(t)$ and $Q(t)$ via the AWG performs arbitrary single-qubit rotations on the Bloch sphere.

We use a dual-horn architecture with broadband operation (6-18 GHz), supporting hyperfine qubit control at 6.834 GHz and Rydberg-Rydberg transitions at higher frequencies. The dual-channel design enables future programmable polarization control.

A Pulsar SW2AD-33 SPDT switch provides 100 ns pulse switching. A Mini-Circuits ZCDC10-01203-S+ directional coupler samples signal for monitoring. An RF-Lambda RFLUPA0618GA2 power amplifier (45 dB gain) amplifies before an RF Spin QRH0422 quad-ridged horn antenna radiates into the experimental volume.




Figure 2.3: Microwave Control System Architecture. Schematic of the microwave control and delivery system. A stable local oscillator (LO) signal, phase-locked to a rubidium frequency standard, drives an IQ mixer. The mixer is modulated by I and Q signals from an Arbitrary Waveform Generator (AWG) to create precisely shaped microwave pulses. The signal is then filtered, switched, and amplified. A directional coupler allows for in-situ power monitoring before the signal is delivered to the atoms via a horn antenna. An isolator protects the amplifier from reflected power.

2.3.2 IQ Mixer Calibration

Imperfect IQ mixing generates spurious spectral components degrading gate fidelity [29]. An ideal upconverter produces $V_{\text{RF}}(t) = I(t) \cos(\omega_{\text{LO}}t) - Q(t) \sin(\omega_{\text{LO}}t)$. DC offsets I_0 and Q_0 create carrier leakage $V_{\text{LO,leak}}(t) = I_0 \cos(\omega_{\text{LO}}t) - Q_0 \sin(\omega_{\text{LO}}t)$. The mixer multiplies LO with IF. DC content in IF produces LO-frequency tones.

I/Q amplitude imbalance ε and phase error ϕ generate image sidebands. Ideal upconversion of $s(t) = I(t) + iQ(t)$ creates single sideband at $\omega_{\text{LO}} + \omega_{\text{IF}}$. Imbalance adds component $\propto s^*(t)$, producing unwanted image at $\omega_{\text{LO}} - \omega_{\text{IF}}$ with amplitude $A_{\text{image}} \propto \varepsilon + i\phi/2$.

We calibrate by adjusting DC offsets to null carrier leakage and tuning I/Q amplitude imbalance and quadrature phase to suppress image sidebands (Figure 2.4). We achieve carrier suppression of 57.6 dB and lower sideband suppression of 50.7 dB.

Off-resonant tones induce AC Stark shift $\delta f \approx \Omega_{\text{res}}^2/(4\pi\Delta)$, where Ω_{res} is residual amplitude and Δ is detuning. For $\Omega_R = 2\pi \times 35$ kHz with 50.7 dB suppression, residual amplitude $\Omega_{\text{res}} = 2\pi \times 102$ Hz. At $\Delta \sim 100$ MHz (IF frequency), this predicts an AC Stark shift of $\delta f \sim 0.1$ Hz, negligible compared to the Rabi frequency.




Figure 2.4: IQ Mixer Calibration. **(a)** Lower sideband suppression as a function of amplitude imbalance ε and phase error ϕ showing the optimization landscape. **(b)** Carrier suppression as a function of DC offset voltages $V_{\text{DC},I}$ and $V_{\text{DC},Q}$. The calibration minimizes DC offsets and I/Q imbalance to achieve carrier suppression of 57.6 dB and lower sideband suppression of 50.7 dB.

2.3.3 Microwave-Driven Coherent Rabi Oscillations

We measure Rabi oscillations on the $|F = 1, m_F = -1\rangle \leftrightarrow |F = 2, m_F = -2\rangle$ transition to validate control uniformity. Scanning microwave frequency with fixed-duration pulse reveals resonance lineshape with broadening (Figure 2.5a).

Scanning pulse duration at resonance produces Rabi oscillations (Figure 2.5b), with π -pulse time $14.08 \mu\text{s}$ corresponding to array-averaged Rabi frequency $\Omega_R/(2\pi) = 35.52$ kHz. Site-resolved measurements over 5×5 boxes (Figure 2.5c) show mean Rabi frequency 35.55 ± 0.13 kHz, with spatial variation of 0.36% across the array.




Figure 2.5: Global Microwave Control of Hyperfine Qubit. Array-averaged measurements of microwave-driven qubit dynamics. **(a)** Rabi oscillations: Population as a function of pulse duration. The fit yields π -pulse time of $14.08 \mu\text{s}$ and Rabi frequency $\Omega_R/(2\pi) = 35.52 \text{ kHz}$. **(b)** Rabi spectroscopy: Population transfer to $|1\rangle$ as a function of microwave frequency. The lineshape reveals linewidth broadening. **(c)** Spatial map of Rabi frequency across the array shows uniformity of microwave control with 0.36% spatial variation.

2.4 Experimental Demonstration of Field Homogenization

We demonstrate iterative feedback compensation that suppresses spatial magnetic field gradients across the atomic array. The protocol comprises three stages: site-resolved Ramsey spectroscopy measures the field distribution, linear fitting extracts gradient components, and gradient coil currents adjust to null the gradients. This achieves field uniformity $\sigma_{\delta B} = 0.115$ mG,

2.4.1 Feedback-Based Magnetic Field Correction Protocol

The magnetic environment is controlled using two types of coil systems. Compensation coils consist of three orthogonal Helmholtz pairs that control the field components B_x , B_y , and B_z . These coils compensate the Earth's magnetic field during polarization gradient cooling.

The compensation algorithm:

1. Site-resolved Ramsey spectroscopy measures local frequency shifts $\Delta\nu(y, z)$, which we convert to a spatial field map $B(y, z)$ using the Zeeman coefficient.
2. Linear fit to the spatial profile $B(y, z)$ extracts gradient coefficients G_y and G_z .
3. Gradient coil currents adjust iteratively to minimize gradients.

Ramsey spectroscopy measures the frequency difference between the atomic transition and the microwave drive. The local magnetic field shifts the atomic transition frequency via the Zeeman effect. However, the microwave drive frequency is set by the local oscillator, not by an absolute reference. This introduces a constant offset common to all sites. Comparing frequencies between sites cancels this offset, so spatial gradients are accurately determined. The absolute magnetic field value remains undetermined. We decompose the measured field into a uniform offset and a spatially varying part:

$$B(y, z) = B_0 + G_y(y - y_0) + G_z(z - z_0) \quad (2.3)$$

where B_0 is a fitted reference offset (not an absolute calibration), $G_y = \partial B / \partial y$ and $G_z = \partial B / \partial z$ are the spatial gradients along row and column directions, and (y_0, z_0) is the array center. Defining the spatial variation $\delta B(y, z) \equiv B(y, z) - B_0$, we have

$$\delta B(y, z) = G_y(y - y_0) + G_z(z - z_0). \quad (2.4)$$

Gradient compensation targets the inhomogeneous component δB , reducing the spatial gradients G_y and G_z across the array.

We use iterative feedback to suppress the measured gradients (G_y, G_z) and minimize spatial field variation δB (Figure 2.6).




Figure 2.6: Iterative gradient compensation protocol. Step 1: Site-resolved Ramsey spectroscopy measures the spatial field profile $B(y, z)$ across all atomic sites in the array. Step 2: Linear fit to the spatial distribution extracts gradient components G_y and G_z along the row and column directions. Step 3: Gradient coil currents adjust to counteract the measured gradients. The protocol iterates until gradient suppression plateaus, establishing a spatially uniform magnetic environment.

2.4.2 Quantifying Improvements in Magnetic Field Uniformity

The magnetic field gradient produces tilted fringes in the position-dependent Ramsey signal. Figure 2.7a shows the reconstructed magnetic field map. Figure 2.7b shows the resulting tilted fringes—atoms at different positions accumulate phase at different rates due to spatially varying Zeeman shifts.

Before compensation, linear gradients are $G_y = -0.125$ mG/row and $G_z = -0.135$ mG/col (Figure 2.8a). Iterative feedback suppresses gradients to $G_y = -0.050$ mG/row and




Figure 2.7: **(a)** Magnetic field reconstructed from measurement of local frequency shifts of the modulated Ramsey spectrum ($\nu_{\text{mod}} = 5$ kHz), obtained after averaging the signal over diagonal bins perpendicular to the gradient. Coordinate u denotes position along the gradient field $G_u \hat{e}_u$, rotated by angle θ from the array axes. **(b)** Ramsey fringes after binning along u . Contrast decay at long times reflects incoherent averaging over $M = 15$ shots with fluctuating global offset phase.

$G_z = 0.015$ mG/col (Figure 2.8b). The standard deviation of spatial field variation $\sigma_{\delta B}$ decreases from 0.410 mG to 0.115 mG. Figure 2.9 shows the resulting distribution of δB across the array.

The 99.9% fidelity threshold requires $\sigma_{\delta B} \lesssim 0.65$ mG. The 99.99% threshold requires $\sigma_{\delta B} \lesssim 0.2$ mG. Initial field inhomogeneity ($\sigma_{\delta B} = 0.410$ mG) satisfied the 99.9% requirement. Compensated field uniformity ($\sigma_{\delta B} = 0.115$ mG) meets the 99.99% threshold.

After compensating spatial gradients, temporal magnetic field fluctuations from the 60 Hz mains remain. A zero-crossing detector synchronizes the experimental sequence to a fixed phase of the mains cycle. This line-triggering ensures each shot samples the same instantaneous magnetic field value, reducing shot-to-shot phase variations in the Ramsey signal.

This uniform magnetic environment enables site-selective Raman operations (Chapter 3) and Rydberg two-qubit gates (Chapter 4).

2.5 Discussion and Outlook

This chapter establishes a calibration procedure for the magnetic environment using atomic qubits as spatially-resolved sensors. Site-resolved Ramsey spectroscopy maps the magnetic




Figure 2.8: Spatial field variation δB across the array before and after iterative compensation. **(a)** Before compensation: clear spatial gradient with $G_y = -0.125$ mG/row and $G_z = -0.135$ mG/col. **(b)** After compensation: uniform distribution with suppressed gradients $G_y = -0.050$ mG/row and $G_z = 0.015$ mG/col.




Figure 2.9: Magnetic Field Uniformity Quantification. Histogram of spatial field variation δB measured across all atomic sites before (red) and after (blue) gradient compensation. Before compensation, the distribution has standard deviation $\sigma_{\delta B} = 0.410$ mG. After compensation, the distribution narrows to standard deviation $\sigma_{\delta B} = 0.115$ mG.

field distribution across 1,305 sites with $3\text{ }\mu\text{m}$ resolution. Iterative gradient compensation reduced field inhomogeneity from $\sigma_B = 0.410\text{ mG}$ to 0.115 mG . The full field reconstruction requires approximately 33 minutes of measurement time.

Beyond static field mapping, site-resolved fluorescence detection supports single-shot gradiometry. A magnetic field gradient imposes a position-dependent Ramsey phase $\phi(\mathbf{r}) = \gamma_B \mathbf{G} \cdot \mathbf{r} \tau$, producing a sinusoidal modulation in the survival probability. Fourier analysis of a single image extracts the gradient magnitude and direction from the peak wavevector $\mathbf{k} = \gamma_B \mathbf{G} \tau$. The complex phase of the Fourier peak provides a global phase reference for each shot. Referencing all site phases to this common value suppresses shot-to-shot offset field drift before ensemble averaging.

The same microwave pulses used for Ramsey spectroscopy can implement Quantum Lock-in Amplification (QLIA) [30]. QLIA spatially resolves AC Stark shifts from the Rydberg excitation laser. Amplitude-modulated laser light induces periodic shifts in atomic transition frequencies. Synchronous detection extracts the modulated component, rejecting common-mode noise from laser power or frequency fluctuations. This technique measures the Rydberg beam intensity profile with sub-micron resolution [19].

Chapter 3

Developing a Raman Laser System for High-Fidelity Single-Qubit Gates

Universal quantum computation requires high-fidelity, arbitrary single-qubit rotations with individual qubit addressability. Microwave control systems developed in the previous chapter provide efficient, high-fidelity coherent manipulation across a neutral atom array. The direct manipulation of the hyperfine states using microwave fields is free from photon scattering decoherence. However, microwave fields at the hyperfine frequency have wavelengths of several centimeters. This prevents selective control of atoms spaced by micrometers in optical tweezers. To perform site-selective addressing, an optical approach is therefore required.

There are two main architectures for site-selective addressing. The first uses a global microwave drive combined with a focused laser that induces a differential AC Stark shift on a target atom [31]. The Stark shift tunes the target atom into resonance with the microwave field, while non-target atoms remain spectrally isolated and off-resonant.

This method uses high-fidelity microwave sources, but achieves Rabi frequencies typically limited to tens of kHz. The slow Rabi frequencies require gate times on the order of tens of microseconds, increasing exposure to decoherence and reducing gate operation rate.

The second approach, which we adopt, uses two focused laser beams to drive a stimulated Raman transition between the qubit states. Controlled Raman pulses implement single-qubit gates. A π -pulse performs an X gate, while arbitrary pulse shapes produce general rotations [18, 32]. The Raman approach enables sub-microsecond gates with MHz-scale Rabi frequencies. Recent demonstrations achieve fidelities as high as 99.99% [33, 19].

Fidelity is limited by spontaneous emission from the intermediate state.

Implementing high-fidelity Raman gates requires generating two phase-coherent optical fields separated by the 6.834 GHz qubit splitting. Two-photon Raman transitions require phase coherence between these optical fields over the gate duration, making them sensitive to phase noise.

Multiple hardware architectures exist for generating phase-coherent fields. These include optical phase-locked loops (OPLL) using two offset-locked lasers with active feedback. Modulator-based systems use MZMs or phase modulators (PMs) driven by RF synthesizers to generate optical sidebands from a single laser.

Each architecture has distinct noise characteristics. In OPLL systems, the residual phase error between the two lasers depends on both laser phase noise and RF synthesizer noise, with contributions partially suppressed by the loop transfer function. In modulator-based systems, RF synthesizer noise couples directly to optical sidebands without loop filtering. However, no quantitative comparison existed to determine which architecture achieves lower gate infidelity.

We develop a simulation framework to evaluate Raman gate architectures and select hardware components. We design a system capable of a π -pulse gate (single-qubit X rotation) in $0.25\ \mu\text{s}$ with predicted infidelity below 10^{-4} . We quantify how optical and RF component noise couples to gate errors, enabling systematic architecture comparison and component selection.

First, we establish a simulation framework for hardware-aware system design by defining requirements, surveying architectures, and developing a Monte Carlo method that incorporates realistic component noise (Sec. 3.1). We then apply this framework to systematically evaluate design choices including wavelength selection, architecture comparison, and component specifications (Sec. 3.2). Finally, we implement and validate the chosen OPLL system through experimental characterization (Sec. 3.3).

3.1 Simulation Framework for Hardware-Aware Raman System Design

Designing a Raman laser system with infidelity below 10^{-4} requires comparing hardware architectures and component specifications without building prototypes. Each architecture presents distinct noise characteristics that affect gate fidelity. Optical Phase-Locked

Loops (OPLL) using two lasers differ from modulator-based systems using electro-optic devices. Experimental characterization would require constructing and testing each variant with different laser types and RF sources—a costly and time-intensive approach requiring construction of multiple hardware prototypes.

In this section, we develop a simulation framework for hardware-aware system design. We establish the performance requirements and physical constraints governing Raman gates with scattering probability below 10^{-4} (Sec. 3.1.1). We survey candidate hardware architectures for generating phase-coherent optical fields (Sec. 3.1.2). We implement a Monte Carlo simulation method that incorporates realistic component noise spectra into master-equation dynamics, enabling quantitative architecture comparison (Sec. 3.1.3).

3.1.1 Design Requirements for High-Fidelity Raman Control

Achieving sub- 10^{-4} gate infidelity requires careful selection of the intermediate state transition wavelength, optical power, and single-photon detuning. These parameters must simultaneously satisfy three constraints: (1) fast gate times to minimize decoherence, (2) low photon scattering probability from the intermediate state ($< 10^{-4}$), and (3) phase coherence between optical fields throughout the gate duration. Component specifications (laser linewidth, RF synthesizer phase noise) must be chosen to meet these requirements.

We define the qubit encoding and Raman transition mechanism. We derive the physics constraints including scattering errors and phase coherence requirements. We then formulate the performance targets that guide component selection.

Qubit Encoding and Raman Transition

The qubit is encoded in the $|F = 1, m_F = 0\rangle$ and $|F = 2, m_F = 0\rangle$ hyperfine ground states of ^{87}Rb , separated by 6.834 GHz. These clock states are first-order insensitive to magnetic field fluctuations, providing enhanced coherence. A two-photon stimulated Raman transition couples these states via a far-detuned intermediate state (Fig. 3.1a). Two optical fields with time-dependent Rabi frequencies Ω_1, Ω_2 drive coherent population transfer between the qubit states through the intermediate state. The large detuning Δ from the intermediate level minimizes spontaneous emission while maintaining an effective two-photon coupling at the qubit frequency $\omega_q = 2\pi \times 6.834$ GHz.




Figure 3.1: ^{87}Rb energy levels and stimulated Raman transition scheme at 795 nm (D1 line). (a) General three-level Λ scheme. Two optical fields with Rabi frequencies Ω_1 , Ω_2 drive a two-photon transition between qubit states $|0\rangle \equiv |F=1, m_F=0\rangle$ and $|1\rangle \equiv |F=2, m_F=0\rangle$ via a far-detuned intermediate state $|e\rangle$. The single-photon detuning Δ from the intermediate level suppresses spontaneous emission while enabling coherent coupling. (b) Rubidium-87 D1 line structure showing the $5S_{1/2}$ ground state hyperfine manifolds ($F=1$, $F=2$ with Zeeman sublevels) and the $5P_{1/2}$ excited states ($|e_1\rangle$, $|e_2\rangle$) corresponding to $F'=1$, $F'=2$). The σ^+ polarized Raman beams couple the $m_F = 0$ clock states through the virtual level, detuned by Δ from the $5P_{1/2}$ manifold. The excited-state hyperfine splitting $\omega_{\text{HF}'} = 816.656(30)$ MHz [26] separates the two intermediate states $|e_1\rangle$ and $|e_2\rangle$.

System Hamiltonian

The dynamics of the Raman transition are governed by a time-dependent Hamiltonian describing the qubit states $|g_1\rangle \equiv |F = 1, m_F = 0\rangle$, $|g_2\rangle \equiv |F = 2, m_F = 0\rangle$, and intermediate excited states $|e_i\rangle$:

$$H(t) = H_0 + H_{\text{interaction}}(t) \quad (3.1)$$

The static Hamiltonian H_0 contains the energy offsets of the intermediate states relative to the qubit manifold. For a system with two intermediate hyperfine levels (D1 line at 795 nm):

$$H_0 = -\Delta_1|e_1\rangle\langle e_1| - (\Delta_1 - \omega_{HF2})|e_2\rangle\langle e_2| + \delta_{\text{tot}}|g_2\rangle\langle g_2| \quad (3.2)$$

where Δ_1 is the single-photon detuning from the first intermediate state, ω_{HF2} is the excited-state hyperfine splitting, and δ_{tot} is the two-photon detuning. For D2 transitions (780 nm, 420 nm) with three intermediate hyperfine levels, we include an additional term $-(\Delta_1 - \omega_{HF2} - \omega_{HF3})|e_3\rangle\langle e_3|$.

The interaction Hamiltonian $H_{\text{interaction}}(t)$ describes the laser-atom interaction and depends on the hardware architecture. For OPLL systems using two independent lasers (Laser 1 as reference, Laser 2 as locked), the interaction terms contain electric field amplitudes E_1 , E_2 and time-dependent phase noise $\phi_1(t)$, $\phi_2(t)$. For MZM systems using a single modulated laser, the interaction is expressed as a sum over optical sidebands with amplitudes proportional to Bessel functions $J_n(\beta)$ of the modulation depth. The dipole matrix elements V_{ij} connecting states $|i\rangle$ and $|j\rangle$ determine the interaction strength. We calculate these using the Alkali Rydberg Calculator (ARC) package [34] accounting for hyperfine structure. Complete expressions for both architectures, including intermediate state specifications and Rabi frequency calculations, are provided in Appendix A.

Requirements for high-fidelity, site-selective Raman gates:

Physics requirements

- Scattering error: To achieve scattering probability $P_{\text{scat}} < 10^{-4}$ at the target gate time (Table 3.1), large single-photon detuning ($|\Delta| \gg \Gamma$) is required.
- Transition wavelength: The atomic transition must be accessible using optical frequencies. The dipole matrix element for the transition and the intermediate state lifetime should be maximized.

Engineering requirements

- Optical power: Large detuning requires high optical power to maintain a reasonable two-photon Rabi frequency. This is necessary to meet the scattering error requirement. The system must deliver the required total power to the atomic plane to achieve the target gate time.
- Phase noise suppression: The relative phase between Ω_1 and Ω_2 must be stable. Phase fluctuations translate into dephasing of the qubit and reduce the fidelity of the gate. This requires using low phase noise lasers and control electronics (e.g RF synthesizers that create frequency reference).
- Crosstalk: The addressing laser beams must be tightly focused to excite only the target atom, minimizing unwanted operations on its neighbors.

Performance Targets and Optimization Problem

We consider driving a Raman transition for a single atom. The design target is a single-qubit π -pulse (an X rotation) with predicted infidelity $1 - F < 10^{-4}$ at the target gate time (Table 3.1) corresponding to two-photon Rabi frequency $\Omega_{\text{eff}}/(2\pi) = 2$ MHz.

The infidelity target of 10^{-4} provides substantial margin below typical surface code thresholds of $\sim 10^{-2}$ under depolarizing noise [35, 8]. This allows for reduced error correction overhead and enables longer logical qubit lifetimes.

Performance targets are:

$$1 - F \leq 10^{-4} \quad (\text{fidelity target})$$

$$t_g \leq 0.25 \mu\text{s} \quad (\text{gate speed target})$$

Gate time and optical power are the key resources constrained by available hardware. Table 3.1 specifies operating parameters for each architecture.

3.1.2 Hardware Architectures for Generating Phase-Coherent Fields

Phase coherence between two optical fields separated by 6.834 GHz impacts Raman gate fidelity. Two hardware architectures satisfy this requirement with distinct implementation

complexity, noise characteristics, and performance trade-offs. We survey optical phase-locked loops using two independent lasers with active feedback, and modulator-based systems using a single laser with electro-optic modulators. Quantitative comparison requires a simulation framework that incorporates realistic noise spectra into gate dynamics—the subject of Section 3.1.3.




Figure 3.2: Hardware architectures for generating phase-coherent Raman fields. **(a)** Optical Phase-Locked Loop (OPLL): Two independent lasers are phase-locked using a heterodyne beat note compared to an RF reference. The error signal drives fast current modulation and slow piezo control to maintain a precise frequency difference of 6.834 GHz, producing two pure optical tones with no higher-order sidebands. **(b)** Electro-Optic Modulator (EOM): A single laser is modulated to generate optical sidebands at the qubit frequency, creating the required two-tone spectrum with higher-order sidebands present at amplitudes determined by the modulation depth.

Optical Phase-Locked Loops (OPLL)

An OPLL uses two independent lasers with active feedback to maintain a frequency difference of 6.834 GHz (Fig. 3.2a). The heterodyne beat note between the two lasers is compared to an RF reference. The error signal drives fast current modulation of one laser to phase-lock it to the other. This produces two pure optical tones with no higher-order sidebands, but requires continuous active stabilization through phase-lock electronics. Environmental perturbations such as temperature fluctuations, acoustic vibrations, and mechanical drift

can couple into the feedback loop, requiring careful isolation and loop design to maintain stable lock.

The electric field on atoms is given by:

$$E(t) = E_1(t) + E_2(t) = E_0 e^{i(\omega_1 t + \phi_1(t))} + E_0 e^{i(\omega_2 t + \phi_2(t))}$$

where ω_1, ω_2 are the frequencies of the two lasers and ϕ_1, ϕ_2 are the phases of the two lasers.

Modulator-Based Systems

Modulator-based approaches use a single laser with electro-optic modulators to generate optical sidebands at the 6.834 GHz qubit splitting. Two modulator types are used:

Phase Modulator (PM): A phase modulator generates an optical field $E(t) = E_0 e^{i\omega_0 t} e^{i\beta \sin(\Omega t)}$, where β is the modulation depth. This produces a frequency comb with sideband amplitudes proportional to Bessel functions $J_n(\beta)$. The carrier and higher-order sidebands cannot be fully suppressed, leading to off-resonant couplings that can degrade fidelity through differential AC Stark shifts.

Mach-Zehnder Modulator (MZM): An MZM configured to suppress the carrier frequency drives the two interferometer arms with RF signals of equal strength but opposite phase (Fig. 3.2b). A DC voltage biases the device at an intensity minimum. The modulation depth β is the peak phase shift applied by the modulator, determining the amplitude distribution among optical sidebands via Bessel functions $J_n(\beta)$. The electric field at the atoms for an MZM-modulated system is:

$$E(t) = E_0 e^{i\omega_0 t} \sum_{n=-\infty}^{\infty} J_n(\beta) (e^{-in(\omega_m t + \phi(t))} - e^{+in(\omega_m t + \phi(t))})$$

where ω_0 is the carrier frequency, $J_n(\beta)$ are Bessel functions of the first kind, β is the modulation depth, ω_m is the RF modulation frequency, and $\phi(t)$ is the time-dependent RF phase noise.

We use the iXblue NIR-MX800-LN modulator with the following specifications and operating parameters: 4.5 dB insertion loss reduces input power $P_0 = 2$ mW to effective power $\tilde{P} = P_0/2.82 = 0.71$ mW. At this modulation depth, the power in each sideband is $P_n = \tilde{P} \times J_n^2(\beta)$. Carrier suppression of 27.5 dB yields residual carrier power $P_{\text{carrier}} = \tilde{P}/562.34$. Our simulations include all optical sidebands up to 11th order to accurately model off-resonant coupling and AC Stark shifts from higher-order sidebands.

Optimization of modulation depth

The modulation depth β presents a trade-off between gate speed and spectral purity. When the MZM is biased for carrier suppression, the desired ± 1 sidebands drive the Raman transition with amplitudes $J_1(\beta)$. Parasitic higher-order sidebands ($\pm 3, \pm 5, \dots$) have amplitudes $J_3(\beta), J_5(\beta), \dots$ that induce unwanted AC Stark shifts.

Our simulations predict maximum fidelity at $\beta = 0.36\pi$, not at $\beta = 0.59\pi$ where $J_1(\beta)$ is maximum. At $\beta = 0.36\pi$, power in the desired sidebands is reduced by 15% (from $|J_1(0.59\pi)|^2 = 0.34$ to $|J_1(0.36\pi)|^2 = 0.29$), but power in the dominant parasitic sideband is reduced by 70% (from $|J_3(0.59\pi)|^2 = 0.09$ to $|J_3(0.36\pi)|^2 = 0.03$). This improves the spectral purity ratio $|J_3/J_1|$ by $2.9\times$ (from 0.53 to 0.18). This reduces unwanted differential light shifts from the third-order sideband, which couples off-resonantly to the qubit states and introduces systematic phase errors during gates. We therefore set $\beta = 0.36\pi$ for all MZM simulations.

Figure 3.3 shows the fidelity landscape versus detuning and pulse duration for the MZM system, illustrating the operating point selection.




Figure 3.3: Fidelity map for MZM-driven Raman transition limited by intermediate state decay. Simulated infidelity ($1 - F$) for a Raman π -pulse versus intermediate state detuning (Δ) and pulse duration. Darker regions indicate higher fidelity. The simulation models an MZM driving an ^{87}Rb transition at 795 nm with $\beta = 0.36\pi$ and includes contributions from all generated sidebands.

Figure 3.4 compares the optical spectra from phase and Mach-Zehnder modulators, illustrating the difference in spectral purity.




Figure 3.4: Comparison of optical spectra from phase and Mach-Zehnder modulators. **(a)** A PM distributes power among a carrier and multiple sidebands according to Bessel functions. The carrier and higher-order sidebands cannot be eliminated. **(b)** An MZM, biased for carrier suppression, efficiently transfers power into the first-order sidebands.

3.1.3 Simulation Framework Incorporating Realistic Noise Models

Experimentally evaluating each hardware architecture is prohibitively expensive and time-consuming. To systematically design a high-fidelity Raman laser system, we develop a noise-aware simulation framework that incorporates measured noise characteristics from lasers and RF synthesizers into master-equation dynamics. This framework enables quantitative architecture comparison and component selection without building multiple prototypes.

We formulate the gate fidelity metric using master-equation dynamics (Sec. 3.1.3). We describe how phase noise enters the Hamiltonian differently for each architecture (Sec. 3.1.3). We implement the Monte Carlo procedure for averaging over stochastic phase trajectories (Sec. 3.1.3) and generate realistic noise spectra from component specifications (Sec. 3.1.3).

Gate Fidelity Metric

We compute the fidelity of a Raman π -pulse, $F = \langle g_2 | \rho(T_\pi) | g_2 \rangle$, for a two-photon stimulated Raman transition using master-equation simulation.

The system evolution is governed by:

$$\frac{d\rho(t)}{dt} = -\frac{i}{\hbar}[H(t), \rho(t)] + \mathcal{L}_{\text{decay}}[\rho(t)] \quad (3.3)$$

where $H(t)$ is the time-dependent Hamiltonian and $\mathcal{L}_{\text{decay}}$ is the Lindblad operator capturing spontaneous emission from the intermediate state.

The fidelity metric F quantifies the overlap between the final density matrix $\rho(T_\pi)$ at the end of the gate and the target state $|g_2\rangle\langle g_2|$.

Phase Noise Incorporation

Phase noise enters the Hamiltonian differently depending on the hardware architecture.

For OPLL-based systems, the Hamiltonian interaction terms are $\frac{1}{2}(\Omega_1 + \Omega_2 e^{-i(\omega_{HF}t + \Delta\phi(t))})$, where $\Delta\phi(t) = \phi_s(t) - \phi_m(t) - \phi_{RF}(t)$ is the residual phase error—the deviation of the locked laser’s phase from ideally tracking the reference laser and RF synthesizer. The residual phase error between the lasers has a power spectral density $S_\phi^e(f)$ determined by:

$$S_\phi^e(f) = \left| \frac{1}{1 + G_{\text{op}}(f)} \right|^2 \left(S_\phi^m(f) + S_\phi^{(s,fr)}(f) + S_\phi^{RF}(f) \right) \quad (3.4)$$

where $G_{\text{op}}(f)$ is the loop transfer function characterizing the phase-lock loop’s ability to suppress noise [36], with contributions from the loop bandwidth, time delay, and PID controller parameters. S_ϕ^m is reference laser noise, $S_\phi^{(s,fr)}$ is locked laser free-running noise, and S_ϕ^{RF} is RF synthesizer noise.

For modulator-based systems (MZM or PM), phase noise enters through the modulation index. The RF synthesizer driving the modulator has a time-dependent phase $\phi(t)$ sampled from its phase noise power spectral density. Each optical sideband acquires a phase factor $e^{\pm in\phi(t)}$ for the n -th sideband, where the RF phase noise directly modulates the optical sidebands without attenuation by a feedback loop.

Appendix B presents the full Hamiltonian derivation and noise incorporation procedure.

Monte Carlo Procedure

Stochastic phase noise requires statistical analysis to predict average gate performance. Each experimental realization of the gate experiences a different phase noise trajectory

drawn from the same underlying power spectral density. We must average over many realizations to obtain the expected performance.

Performance is reported as the Monte Carlo average over M phase noise realizations:

$$F_{\text{avg}} = \frac{1}{M} \sum_{m=1}^M \text{tr}\{\rho_m(T_\pi)\rho_{\text{target}}\} \quad (3.5)$$

Each stochastic phase trajectory $\phi_j(t)$ is synthesized from measured power spectral densities using:

$$\phi_j(t_i) = \sum_k 2\sqrt{S_\phi(f_k)\Delta f} \cos(2\pi f_k t_i + \phi_{\text{rand},j,k}) \quad (3.6)$$

with random phases $\phi_{\text{rand},j,k} \in U(0, 2\pi)$ [37]. Appendix B presents the complete procedure.

Noise Spectrum Generation

We generate laser phase noise from manufacturer-specified linewidths using the β -line separation model [38]. This allows simulating performance with different laser types: narrow-linewidth [ECDLs](#) ($\Delta\nu \lesssim 50$ kHz) and broader distributed feedback (DFB) diode lasers ($\Delta\nu \sim \text{few MHz}$).

We extract RF synthesizer phase noise from measured specifications. We compare two representative sources: the Vaunix LMS-183DX with phase noise of -83 dBc/Hz at 10 kHz offset from 6 GHz, and the Holzworth HSX9000 with -127 dBc/Hz at 10 kHz offset.

3.2 Hardware Architecture Evaluation and Selection

We apply the simulation framework to identify the architecture, atomic transition, and component specifications that predict sub- 10^{-4} gate infidelity.

We first select the optimal intermediate-state transition for driving the Raman process (Sec. 3.2.1). We then compare OPLL and modulator-based architectures to determine which achieves lower predicted infidelity (Sec. 3.2.2). We systematically evaluate the impact of laser phase noise (Sec. 3.2.3) and RF source phase noise (Sec. 3.2.4) on gate fidelity. Finally, we synthesize these results into design recommendations (Sec. 3.2.6).

Operating point and model settings

Table 3.1 summarizes the operating conditions used for the simulations reported here. All cases target $\Omega_{\text{eff}}/2\pi = 2$ MHz. For OPLL, we model relative phase noise via the loop-filtered pathway described in Eq. 3.4.

Infidelity comparison

Table 3.2 presents the simulated gate infidelities by architecture, laser type, and wavelength.

Table 3.1: Simulation operating points used in the evaluation. MZM power reflects effective power after 4.5 dB insertion loss (2 mW input). All simulations target 0.25 μs gate time ($\Omega_{\text{eff}}/2\pi = 2$ MHz) and use Monte Carlo averaging over 300 stochastic phase trajectories. MZM configuration: iXblue NIR-MX800-LN with $\beta = 0.36\pi$, $\omega_m/2\pi = 3.417$ GHz, carrier suppression 27.5 dB, sidebands modeled up to 11th order.

Configuration	Power	$\Delta/2\pi$ (GHz)
MZM @ 795 nm	0.71 mW	−86.92
MZM @ 780 nm	0.71 mW	−88.45
OPLL @ 795/780 nm	4 mW	−1165
OPLL @ 421 nm	4 mW	−10.52
OPLL @ 420 nm	4 mW	−14.67

3.2.1 Transition Wavelength Selection

We select the intermediate excited-state manifold that minimizes scattering error using the operating points summarized in Table 3.1.

Four excited-state manifolds can mediate the Raman transition in ^{87}Rb : 420 nm ($5S_{1/2} \rightarrow 6P_{3/2}$), 421 nm ($5S_{1/2} \rightarrow 6P_{1/2}$), 780 nm ($5S_{1/2} \rightarrow 5P_{3/2}$, D2), and 795 nm ($5S_{1/2} \rightarrow 5P_{1/2}$, D1). These transitions differ in three properties that affect gate fidelity under power constraints: dipole matrix elements, natural linewidths, and hyperfine pathway interference.

Table 3.2: Simulated Gate Infidelity Comparison by Architecture and Wavelength

Configuration			Wavelength (nm)			
Architecture	Laser	RF Source	420	421	780	795
MZM	ECDL	Vaunix	—	—	3.12×10^{-2}	8.97×10^{-3}
	ECDL	Holzworth	—	—	2.99×10^{-2}	8.44×10^{-3}
	DFB	Vaunix	—	—	2.97×10^{-2}	8.77×10^{-3}
	DFB	Holzworth	—	—	2.98×10^{-2}	8.43×10^{-3}
OPLL	2 × DFB	Holzworth	5.62×10^{-1}	5.30×10^{-1}	2.08×10^{-2}	1.70×10^{-2}
	2 × ECDL	Holzworth	5.75×10^{-1}	5.37×10^{-1}	1.84×10^{-4}	8.80×10^{-5}

Central values only; asymmetric uncertainties omitted for clarity. Cell colors indicate performance:

$< 10^{-3}$, 10^{-3} – 10^{-2} , 10^{-2} – 10^{-1} , $> 10^{-1}$.

Scattering vs. speed tradeoff

The key resources for driving the Raman transition are time and available optical power. The main control knob is the single-photon detuning Δ from the intermediate state. In the large-detuning regime ($|\Delta| \gg \Gamma$), the effective Raman Rabi frequency scales approximately as

$$\Omega_{\text{eff}} \simeq \frac{\Omega_1 \Omega_2}{2\Delta}, \quad (3.7)$$

where $\Omega_{1,2}$ are the single-photon Rabi frequencies set by the beam powers $P_{1,2}$.

Spontaneous emission arises from the small admixture of the excited state and, to leading order, the scattering rate is

$$\Gamma_{\text{sc}} \simeq \Gamma \frac{\Omega_1^2 + \Omega_2^2}{4\Delta^2}, \quad (3.8)$$

so the scattering probability during the gate is $p_{\text{sc}} \simeq \Gamma_{\text{sc}} t_g \propto 1/\Delta^2$ for fixed power and gate time. Increasing $|\Delta|$ suppresses scattering.

Power enters through $\Omega_{1,2} \propto \sqrt{P_{1,2}}$. For power balanced beams ($P_1 \approx P_2$), $\Omega_{\text{eff}} \propto P_{\text{tot}}/|\Delta|$, where $P_{\text{tot}} = P_1 + P_2$. Consequently, to maintain a fixed Ω_{eff} while increasing $|\Delta|$, the required optical power grows linearly with $|\Delta|$. We choose $|\Delta|$ and optical power according to the operating points summarized in Table 3.1, selecting the largest feasible $|\Delta|$ that meets the target two-photon Rabi frequency at the available power.

In the large-detuning regime, the two-photon Raman amplitude is a coherent sum over auxiliary hyperfine levels:

$$\Omega_{\text{eff}} \propto \sum_{F_e, m_{F_e}} \frac{\Omega_{g \rightarrow e} \Omega_{e \rightarrow g'}}{2 \Delta_{F_e}}.$$

Relative signs set by angular-momentum algebra cause pathway interference. For the D2 manifold ($5P_{3/2}$), these pathways interfere destructively, suppressing the achievable Ω_{eff} by approximately 0.07 relative to a single-level estimate [39]. The D1 manifold ($5P_{1/2}$) has simpler hyperfine structure where pathways add constructively, avoiding this cancellation.

The scattering rate scales as $\Gamma_{\text{sc}} \propto (\Gamma_e/\Delta)(\Omega^2/\Delta)$ (Eq. 3.8), where Γ_e is the natural linewidth. Increasing $|\Delta|$ from 10 GHz to 100 GHz reduces scattering by $100\times$ but requires $10\times$ more optical power to maintain the target Ω_{eff} , establishing the power-detuning trade-off.

Table 3.3: Candidate Raman transitions from $5S_{1/2}$ ground state [26]. J-reduced dipole matrix elements and splittings between excited-state hyperfine levels determine the achievable Rabi frequencies and pathway interference. D1 transitions ($j = 1/2$) have 2 intermediate states; D2 transitions ($j = 3/2$) have 3 intermediate states with one forbidden pathway (see Appendix A).

λ (nm)	Transition	$\Gamma_e/2\pi$ (MHz)	$\langle J er J' \rangle$ (ea_0)	$\omega_{HF}/2\pi$ (MHz)	States
420	$6P_{3/2}$	1.347	0.5230	51.7, 87.0	3
421	$6P_{1/2}$	1.231	0.3236	265.2	2
780	$5P_{3/2}$	6.065	6.003	157.0, 266.7	3
795	$5P_{1/2}$	5.746	4.253	816.7	2

The blue 420/421 nm manifolds have narrower linewidths ($\Gamma_e/2\pi = 1.3$ MHz vs 6 MHz for D-lines), potentially reducing scattering at fixed detuning. However, their J-reduced dipole matrix elements are $8\text{--}13\times$ smaller than D1 and $12\text{--}19\times$ smaller than D2 [40, 41, 26]. From Eq. 3.7, the single-photon Rabi frequencies $\Omega_{1,2}$ enter as $\Omega_{1,2} \propto (\text{dipole}) \times \sqrt{P}$. Achieving $\Omega_{\text{eff}}/(2\pi) = 2$ MHz with smaller dipole elements would require $64\text{--}361\times$ more optical power at the same detuning. Since this power level is impractical, we instead compensate by reducing the detuning magnitude: Table 3.1 shows $|\Delta|/2\pi = 10\text{--}15$ GHz for 420/421 nm, compared to 1165 GHz for 795 nm. From Eq. 3.8, the scattering rate scales as $\Gamma_{\text{sc}} \propto 1/\Delta^2$, so reducing $|\Delta|$ by $100\times$ increases scattering by $10,000\times$, proportionally increasing the intermediate state population and decay probability.

Table 3.2 confirms this analysis: the 420 nm and 421 nm transitions yield infidelities of 0.56 and 0.53 (over 50% error), dominated by scattering from the intermediate state.

The 795 nm D1 transition achieves the lowest simulated infidelity at its operating point (Table 3.1). With a large dipole matrix element ($4.253 ea_0$), this transition enables large detuning ($|\Delta|/2\pi = 1165$ GHz) that suppresses scattering while maintaining the target Rabi frequency. The 780 nm D2 transition shows higher infidelity of 1.84×10^{-4} , consistent with reduced effective coupling from hyperfine-path interference.

We select the 795 nm D1 transition because it achieves sub- 10^{-4} simulated infidelity at its operating point. This manifold combines constructive hyperfine interference (unity coupling factor) and a strong dipole matrix element, enabling large detuning that suppresses scattering errors.

3.2.2 Comparing OPLL and Modulator-Based Architectures

Table 3.2 compares simulated infidelity for each architecture at 795 nm. Simulations predict that OPLL with narrow-linewidth ECDLs achieves lower infidelity (8.80×10^{-5}) compared to MZM with low-noise Holzworth RF synthesizer (8.44×10^{-3}), though detailed noise source decomposition analysis is needed to identify the dominant mechanisms responsible for this difference.

The simulations model architecture-specific noise pathways: loop-filtered relative phase error in OPLL (Eq. 3.4), and RF-driven sideband phase modulation in MZM. MZM performance includes contributions from 4.5 dB insertion loss, higher-order sideband coupling, and RF synthesizer phase noise, as described in Sec. 3.1.2.

Figure 3.5 illustrates the Monte Carlo simulations for the MZM and for the OPLL architectures with different laser and RF sources options.

Our simulations predict that OPLL architecture with narrow-linewidth ECDLs achieves sub- 10^{-4} infidelity (8.80×10^{-5}) at the operating point (Table 3.1). In the MZM cases considered, simulated fidelities remain above the 10^{-3} level at this operating point. As described in Sec. 3.1.2, MZM performance is affected by insertion loss reducing available optical power, higher-order sidebands introducing off-resonant coupling, and RF synthesizer phase noise. Simulations also indicate that, with 27.5 dB carrier suppression, residual carrier leakage contributes negligibly to gate infidelity at this operating point, as it couples far off-resonance and primarily contributes higher-order AC Stark shifts.

3.2.3 Impact of Laser Phase Noise on Gate Fidelity

Laser phase noise impacts Raman gate fidelity differently in OPLL versus modulator architectures. In modulator-based systems, both optical sidebands originate from the same




Figure 3.5: Comparison of Raman gate architectures and component specifications through master-equation simulations incorporating realistic phase noise. Each panel shows simulated population transfer $\langle \psi(t)|1 \rangle$ for a $0.25 \mu\text{s}$ π -pulse at 795 nm D1 transition (operating parameters in Table 3.1). Red line shows ideal evolution without noise; blue line shows Monte Carlo average over 300 stochastic phase noise trajectories; shaded region indicates standard deviation. Inset magnifies final fidelity. **(a)** MZM with ECDL laser and Vaunix RF source; **(b)** MZM with ECDL laser and Holzworth RF source; **(c)** OPLL with DFB lasers and Holzworth RF source; **(d)** OPLL with ECDL lasers and Holzworth RF source.

laser. Laser carrier phase noise appears as common-mode noise that cancels in the relative phase between sidebands. However, RF modulation phase noise does not cancel—as described in Sec. 3.1.3, the n -th sideband acquires a phase factor $e^{\pm i n \phi_{\text{RF}}(t)}$, producing opposite phase shifts on the $+1$ and -1 sidebands that drive the Raman transition. Table 3.2 confirms laser noise cancellation: for MZM at 795 nm, ECDL versus DFB lasers show negligible difference in simulated infidelity (all MZM cases yield $8.4\text{--}9.0 \times 10^{-3}$, varying by approximately 0.5×10^{-3}). In OPLL systems, the relative phase noise between two independent lasers depends on both their intrinsic linewidths and the phase-lock loop’s ability to suppress this difference through feedback.

This subsection quantifies how laser specifications affect gate performance in OPLL systems. We compare distributed feedback (DFB) diode lasers with linewidths of a few MHz to narrow-linewidth external-cavity diode lasers (ECDLs) with linewidths below 50 kHz.

Laser linewidth requirements

Within the OPLL architecture, we compare two laser types: distributed feedback (DFB) diode lasers with linewidths of a few MHz, and narrow-linewidth external-cavity diode lasers (ECDLs) with $\Delta\nu \lesssim 50$ kHz. Table 3.2 shows that for 795 nm OPLL systems, two ECDLs meet the target, while two DFB lasers predict 1.70×10^{-2} , failing to meet the target. The table shows a two-order-of-magnitude difference in simulated infidelity between ECDL and DFB lasers in the OPLL configuration, demonstrating the critical importance of narrow-linewidth lasers for achieving sub- 10^{-4} performance.

3.2.4 Impact of RF Source Phase Noise on Gate Fidelity

RF synthesizer phase noise couples to optical fields through different mechanisms in OPLL and modulator architectures. In modulator-based systems, the RF source directly drives the electro-optic modulator, transferring phase fluctuations to the optical sidebands. In OPLL systems, the RF source serves as the local oscillator for the heterodyne phase lock.

Table 3.2 compares two representative RF sources: Vaunix LMS-183DX (phase noise -83 dBc/Hz at 10 kHz offset) and Holzworth HSX9000 (-127 dBc/Hz at 10 kHz offset). For MZM at 795 nm, the table shows small differences between RF sources: ECDL+Vaunix yields 8.97×10^{-3} while ECDL+Holzworth yields 8.44×10^{-3} , a difference of approximately 0.5×10^{-3} .

3.2.5 Additional Error Mechanisms

The phase noise analysis above assumes ideal atomic conditions. Three additional error sources arise beyond phase noise. ECDL relaxation oscillations contribute high-frequency laser noise at frequencies above the servo bandwidth [42, 43]. Thermal atomic motion through the Gaussian beam intensity profile samples position-dependent Rabi frequencies, causing pulse area errors [44]. The two-photon Raman transition imprints a phase $\phi = \mathbf{k}_{\text{eff}} \cdot \mathbf{r}$ determined by the atom’s position within the trap. Since atoms undergo thermal motion within the trapping potential, each experimental shot samples a different position, acquiring a different phase. Averaging over many shots therefore yields dephasing.

ECDL relaxation oscillations

External cavity diode lasers exhibit relaxation oscillations [43]. This resonant feature in the frequency noise spectrum arises from the dynamic exchange between the intracavity photon field and the gain medium population inversion. For ECDLs, this relaxation peak typically occurs at frequencies of order 1 GHz, well above the bandwidth of typical phase-lock servo loops (1–10 MHz) [45, 46]. Consequently, the relaxation oscillation cannot be actively suppressed and remains present in the free-running noise spectrum. Whether this high-frequency noise degrades gate fidelity depends on the qubit’s sensitivity to noise at these frequencies.

To assess the impact of relaxation oscillations on gate fidelity, we apply the χ -separation line criterion introduced by Day et al. [43]. This metric defines a frequency-dependent threshold: noise below the χ -separation line contributes negligibly to gate errors, regardless of its spectral structure. For a π -pulse with Rabi frequency $\Omega/2\pi = 2$ MHz, the χ -separation line at the relaxation frequency $f_{\text{RO}} = 1$ GHz evaluates to $\chi_{\text{sep}} = h_a(2f_{\text{RO}}/\Omega + 1) \approx 1000 \times h_a$, where h_a is the white noise floor.

We use frequency noise data for a Toptica DL pro ECDL [47]. Following the standard decomposition of laser frequency noise into flicker ($1/f$) and white noise components [38], the spectrum converges to a white noise floor $h_a = 50 \text{ Hz}^2/\text{Hz}$ at high offset frequencies. Extrapolating this baseline to 1 GHz gives $\chi_{\text{sep}} = 50,224 \text{ Hz}^2/\text{Hz}$ (Fig. 3.6). A conservative 20 dB relaxation peak produces an amplitude of $5,017 \text{ Hz}^2/\text{Hz}$, well below χ_{sep} .

Monte Carlo simulations validate this analysis: including the 20 dB relaxation oscillation increases the gate infidelity from 8.85×10^{-5} to 9.07×10^{-5} ; a 30 dB peak increases infidelity to 1.09×10^{-4} . We conclude that ECDL relaxation oscillations do not compromise

the predicted gate infidelity at the 10^{-4} level. This analysis assumes a 1 GHz relaxation frequency [43]. Lower-frequency peaks are filtered less effectively and would degrade fidelity more. Measuring the actual peak frequency would confirm this margin.




Figure 3.6: ECDL frequency noise spectrum based on Toptica DL pro specifications [47] (blue points, 45 Hz–100 kHz). The white noise floor $h_a = 50 \text{ Hz}^2/\text{Hz}$ (green dotted) is extrapolated to 10 GHz with a 20 dB relaxation oscillation peak at 1 GHz (red solid). The χ -separation threshold (dashed black) at $50,224 \text{ Hz}^2/\text{Hz}$ lies above the relaxation oscillation peak amplitude of $5,017 \text{ Hz}^2/\text{Hz}$, indicating that relaxation oscillations contribute negligibly to gate infidelity.

Pulse area errors from atomic motion

The simulation assumes uniform Rabi frequency across the atomic wavepacket. In the experiment, the focused Raman beam ($w_0 = 1 \text{ } \mu\text{m}$) produces a position-dependent Rabi frequency $\Omega(r) = \Omega_0 \exp(-r^2/w_0^2)$, where r is the radial displacement from the beam axis. For atoms at temperature $T = 25 \text{ } \mu\text{K}$ with radial trapping frequency $\omega_r/2\pi = 85 \text{ kHz}$, the thermal position spread $\sigma = \sqrt{k_B T / m \omega_r^2} \approx 92 \text{ nm}$ means that atoms at different positions within their thermal distribution experience different Rabi frequencies. A pulse of fixed duration therefore produces a spread of nutation angles. Averaging over the Boltzmann distribution yields an amplitude-error contribution of 1.3×10^{-3} to the gate infidelity, exceeding the phase-noise-limited prediction by an order of magnitude.

This amplitude noise can be suppressed using BB1 composite pulse sequences [48, 49], which cancel pulse-area errors to third order in the fractional Rabi deviation. However, BB1

requires 5π total rotation (versus π for a simple pulse), increasing spontaneous scattering from 1.6×10^{-5} to 7.8×10^{-5} and extending the gate duration fivefold. The phase noise contribution to infidelity for a simple π -pulse is 7.3×10^{-5} (excluding spontaneous emission). Assuming phase noise infidelity scales approximately with gate duration, BB1 accumulates $\sim 5 \times 7.3 \times 10^{-5} \approx 3.6 \times 10^{-4}$ from phase noise. The net predicted infidelity with BB1 is 4.4×10^{-4} , now limited by phase noise rather than amplitude error.

Motional dephasing

The two-photon Raman coupling acquires a position-dependent phase $\phi(\mathbf{r}) = \Delta\mathbf{k} \cdot \mathbf{r}$, where $\Delta\mathbf{k} = \mathbf{k}_1 - \mathbf{k}_2$ is the differential wavevector pointing along the beam propagation direction. For atoms with position uncertainty σ , this phase spread reduces the effective Rabi frequency by the Debye-Waller factor $\exp(-\eta^2/2)$ [50, 51], where $\eta = |\Delta\mathbf{k}|\sigma$ is the Lamb-Dicke parameter. The co-propagating geometry yields $\Delta k = \Delta\omega/c = 2\pi \times 6.834 \text{ GHz}/c \approx 143 \text{ rad/m}$. For thermal atoms at $T = 25 \text{ } \mu\text{K}$ with axial trapping frequency $\omega_z/2\pi = 13 \text{ kHz}$, the position spread $\sigma = \sqrt{k_B T / m\omega_z^2} \approx 600 \text{ nm}$ gives $\eta \approx 9 \times 10^{-5}$. The resulting infidelity contribution $\eta^2/4 \sim 2 \times 10^{-9}$ is negligible compared to the phase-noise-limited prediction.

3.2.6 Design Recommendations and Performance Comparison

The evaluation of hardware architectures identifies a clear configuration in our simulations: an OPLL architecture using narrow-linewidth ECDLs at 795 nm. This design meets the sub- 10^{-4} target with margin at the operating parameters specified in Table 3.1.

Our simulated prediction of 99.9912% fidelity (Table 3.2), if validated experimentally, would be competitive with demonstrated Raman gates. Published Raman gates achieve 99.97% for global rotations [19] and 99.912% for local gates [33]. Microwave gates in neutral-atom arrays reach 99.993% [52]. Experimental validation is required to confirm this prediction.

Since the design of our system, an alternative time-domain approach has been demonstrated for driving optical qubit transitions [53]. This method uses amplitude modulation of a single laser beam to create a time-varying vector light shift that acts as a fictitious magnetic field, driving qubit rotations. This differs from generating two phase-coherent optical fields whose beat note drives a two-photon Raman transition. The Chirped Bragg Grating (CBG) implementation achieves gigahertz-frequency amplitude modulation by first

applying phase modulation via an electro-optic modulator, then converting the phase modulation to amplitude modulation using a dispersive chirped Bragg grating. This approach demonstrates 2 MHz Rabi frequencies with photon-scattering errors below 2×10^{-4} per π -pulse [53]. The CBG architecture offers passive stability without requiring active feedback loops, though the OPLL approach provides independent control of two laser beams that can enable pulse shaping for optimal control applications.

The predicted infidelity (Table 3.2) is a simulation result based on noise spectra generated from component phase noise specifications and atomic properties. We have not validated this prediction experimentally through on-atom gate characterization. Future work requires randomized benchmarking or gate set tomography to measure gate performance and verify the simulation predictions.

3.3 Experimental Implementation and Characterization

Having identified the optimal configuration through quantitative simulation (Section 3.2), we constructed the hardware system and characterized its noise performance. The evaluation identified an OPLL architecture with narrow-linewidth ECDLs at 795 nm as necessary to achieve sub- 10^{-4} predicted gate infidelity. Translating this design into working hardware requires addressing practical challenges in phase locking, active power stabilization, and noise characterization.

We implement the OPLL-based Raman system including laser selection, phase-locking electronics with Toptica FALC Pro controller, and power stabilization using AOM-based servo loops (Sec. 3.3.1). We characterize the system’s phase noise spectrum through out-of-loop beat-note measurement (Sec. 3.3.2).

3.3.1 Raman Laser System Design

We constructed a two-laser OPLL system using narrow-linewidth 795 nm ECDLs (Toptica DLpro).

Laser selection and specifications: The Raman beam pair is generated using two Toptica DL Pro external-cavity diode lasers at 795 nm. Each laser provides 200 mW output power with specified linewidth below 50 kHz. We designate one laser as the reference and the other as the locked laser. We sample small fractions of each beam for beat-note detection and wavemeter monitoring.

Phase-lock electronics. The heterodyne beat from a fast photodiode (bandwidth 2 GHz) is mixed with a microwave local oscillator (National Instruments QuickSyn FSW-0020) set to 6.724 GHz. The signal is filtered, amplified, down-converted, and processed by a Toptica FALC Pro controller acting as phase/frequency discriminator with servo bandwidth 1 MHz. The FALC output corrects the locked laser via fast current modulation (bandwidth 50Hz), phase-locking it to the reference laser at the target Raman splitting of 6.834 GHz.

Intensity stabilization. Each beam passes through an acousto-optic modulator (AOM) driven from a common 80 MHz reference. We use the first-order diffracted beams for the experiment. Photodiodes monitoring the delivered power feed a digital servo that regulates the AOM drive amplitude. Set-points provided by an arbitrary waveform generator (AWG) implement active intensity stabilization and programmable pulse shaping.

Fig. 3.7 presents the system architecture.




Figure 3.7: OPLL System Architecture for Raman Gates. Two Toptica DL Pro external-cavity diode lasers at 795 nm (reference and locked) provide the Raman beams. Small fractions are tapped for beat-note detection and wavemeter monitoring. The heterodyne beat from a fast photodiode is mixed with a microwave local oscillator (NI QuickSyn FSW-0020 at 6.724 GHz), filtered and amplified, then down-converted and error-processed by a Toptica FALC Pro fast PID/servo. The error drives the locked laser (fast current, slow PZT via the DLC Pro) to phase-lock it to the reference at the desired frequency difference. Each beam passes an AOM referenced to a common 80 MHz source; the first-order outputs are delivered to the experiment while the zero-order beams are dumped.

3.3.2 Phase Noise Characterization

Validating the simulation framework requires measuring the constructed system's phase noise spectrum and confirming it matches the assumptions used in the Monte Carlo simulations. The simulation inputs were based on manufacturer-specified laser linewidths and synthesizer phase noise specifications. Direct measurement of the OPLL's relative phase noise provides an independent check of these assumptions.

This subsection presents the phase noise characterization procedure and results. We measure the in-loop phase noise power spectral density by heterodyning the two OPLL lasers on an independent photodiode. We identify technical noise sources that degrade performance and implement mitigation strategies. We compare the measured spectrum to the simulation inputs and assess agreement.

To characterize system performance, we measured the in-loop phase noise power spectral density by beating the two OPLL lasers on an independent photodiode. We calibrate the error-signal voltage to frequency units.

During initial characterization, we observed a technical noise peak at 150 MHz in the frequency noise spectrum. We traced this peak to local-oscillator leakage in the RF mixing stage coupling into the laser fast current modulation input. To suppress this noise, we inserted a low-pass filter before the fast current input.

Fig. 3.8 shows the measured phase noise PSD after suppressing the 150 MHz technical noise peak.




Figure 3.8: Raman laser system phase noise measurement. In-loop beat-note measurement of the two-laser OPLL system. Red: frequency noise PSD without low-pass filter. Blue: PSD with low-pass filter, suppressing the technical noise peak at 150 MHz.

The measured phase noise PSD characterizes the OPLL system’s frequency stability. The 150 MHz peak corresponds to local oscillator leakage into the servo loop.

3.3.3 Summary and Outlook

This chapter develops a Monte Carlo simulation framework designed to predict single-qubit gate fidelity based on component-level noise specifications. By decomposing gate infidelity into its constituent error sources, this framework enables the identification of critical performance bottlenecks prior to hardware construction. This capability facilitates quantitative design decisions, allowing for the optimization of system parameters to meet specific fidelity thresholds.

A comparative analysis of control architectures reveals that the Optical Phase-Locked Loop (OPLL) utilizing narrow-linewidth External Cavity Diode Lasers (ECDLs) at 795 nm achieves a predicted phase-noise-limited infidelity of 8.80×10^{-5} . This represents an improvement of two orders of magnitude over the Mach-Zehnder Modulator (MZM) architecture, which yields an infidelity of 8.44×10^{-3} .

However, the simulations indicate that at typical operating temperatures ($25 \mu\text{K}$), thermal motion through the focused beam’s intensity profile contributes 1.3×10^{-3} to the total infidelity. In this regime, amplitude errors arising from thermal motion become the dominant error source for site-selective operations, exceeding the contribution of laser phase noise. To mitigate this, the implementation of BB1 composite pulses suppresses these amplitude errors, reducing infidelity to 4.4×10^{-4} . This improvement comes at the cost of a fivefold increase in gate duration and a corresponding rise in spontaneous scattering probability (7.8×10^{-5} compared to 1.6×10^{-5}).

These findings guide experimental design: for systems employing site-selective focused beams, the implementation of robust pulse sequences may improve fidelity beyond what reductions in laser phase noise can achieve. To this end, the OPLL system has been constructed and the relative phase noise between the two Raman beams has been characterized. Experimental validation of the full gate performance via randomized benchmarking [54, 55] remains the subject of future work.

Chapter 4

Engineering High-Fidelity Rydberg Gates through Laser Stabilization and Quantum Optimal Control

High-fidelity quantum operations at scale are central for fault-tolerant quantum computing. Strong, tunable dipole-dipole interactions between atoms excited to Rydberg states enable two-qubit entangling gates with demonstrated fidelities of 99.5% in 60-atom arrays [19]. Fault-tolerant error correction thresholds of 10^{-4} infidelity [8] require systematic suppression of multiple error sources: laser phase and amplitude noise, spontaneous emission from intermediate and Rydberg states, and atomic motion-induced Doppler shifts.

Laser phase noise limits gate performance by coupling to the Rydberg excitation via the two-photon detuning, causing dephasing throughout the gate operation. Traditional optimization of laser stabilization requires slow, iterative on-atom experiments. Each adjustment to control parameters requires running full gate sequences to measure fidelity. This approach is impractical for systematic parameter space exploration and yields suboptimal performance.

We integrate fast laser phase noise engineering and hardware-aware pulse design. We characterize the experimental hardware, optimize laser stabilization using filter function analysis to predict Rydberg excitation infidelity from measured noise spectra, and apply quantum optimal control to design robust pulses incorporating realistic constraints such as finite control bandwidth.

The first component focuses on laser stabilization: 420 nm and 1013 nm lasers independently locked to a high-finesse, ultra-low expansion bichromatic cavity using the Pound-

Drever-Hall (PDH) technique. We map measured in-loop frequency noise power spectral densities directly to predicted Rydberg excitation infidelity using filter function analysis based on linear response theory. This physics-based optimization metric enables systematic tuning of proportional-integral-derivative (PID) gains for each actuator, completing optimization in minutes rather than days.

The second component implements quantum optimal control (QOC) to numerically design time-dependent control fields that steer the quantum system toward a desired objective. We use QOC to generate tailored pulse shapes with improved robustness against parameter fluctuations, including laser phase noise, maintaining high fidelity under realistic experimental conditions. The optimization uses the Van Loan formalism for efficient gradient computation and employs a multi-objective cost function targeting final-state fidelity and adiabatic dark-state overlap. Hardware constraints such as finite modulation bandwidth and maximum pulse amplitudes are enforced directly within the optimization loop.

We make three contributions. First, we map measured frequency noise spectra to predicted Rydberg excitation infidelity, avoiding proxy metrics like integrated RMS noise. Second, we develop a multi-objective QOC framework that simultaneously optimizes fidelity and adiabaticity. Third, we apply this framework to design pulses for single-atom Rydberg excitation and two-qubit controlled-Z gates.

Systematic PID tuning yields predicted infidelity from laser phase noise below 10^{-3} for two-photon Rabi frequencies above $2\pi \times 4$ MHz. We achieve additional performance gains through QOC-optimized control pulses. For single-atom Rydberg excitation, we demonstrate 99.7% transfer efficiency with a sixfold reduction in gate time (23.90 ns vs. 143.36 ns) compared to analytical benchmark protocols. For two-qubit controlled-Z gates, we improve simulated fidelity from 99.22% to 99.68% at 209 ns gate duration.

Section 4.1 characterizes the two-photon excitation platform and error sources. Section 4.2 develops hardware optimization through laser phase noise engineering. Section 4.3 develops pulse design using quantum optimal control. Sections 4.4 and 4.5 apply these methods to single-atom Rydberg excitation and two-qubit CZ gates.

4.1 Two-Photon Rydberg Excitation: Physics and Error Sources

High-fidelity gate design requires understanding the physical platform and error mechanisms. This section establishes the two-photon excitation scheme, hardware architecture,

and dominant error sources guiding the optimization strategies in Sections 4.2–4.5.

4.1.1 Two-Photon Rydberg Excitation Scheme

A two-photon excitation scheme drives the $|5S_{1/2}\rangle \rightarrow |6P_{3/2}\rangle \rightarrow |70S_{1/2}\rangle$ transition in ^{87}Rb using 420 nm and 1013 nm lasers. The $70S_{1/2}$ Rydberg state provides strong van der Waals interactions (Section 4.1.3) with radiative decay rate $\Gamma_r = 2\pi \times 2.67$ kHz, enabling gate times up to several hundred nanoseconds with minimal population loss.

The choice of intermediate state detuning Δ_1 balances competing error sources. Small detuning populates the lossy intermediate state, causing decoherence from spontaneous scattering. Large detuning suppresses intermediate state population but reduces the effective two-photon Rabi frequency ($\Omega_{\text{eff}} \propto 1/\Delta_1$) for given laser power, slowing excitation and increasing susceptibility to Rydberg state decay and low-frequency noise.

Figure 4.1 shows the laser system architecture for Rydberg excitation.




Figure 4.1: Rydberg laser system architecture. **Top:** 1013 nm seed laser (NP Photonics) amplified to 10 W, with AOM for frequency control and EOM for PDH locking to stabilization cavity. **Bottom:** SolsTiS Ti:Sapphire produces 840 nm output, sampled to stabilization cavity for locking, then frequency-doubled in resonant cavity to generate 420 nm light for experiments. Both lasers lock to the same ultra-stable cavity (finesse $\approx 50,000$, linewidth 15 kHz).

The excitation scheme uses ground state $|g\rangle = |5S_{1/2}, F = 2, m_F = -2\rangle$, intermediate state $|i\rangle = |6P_{3/2}, F = 3, m_F = -3\rangle$ with natural lifetime $\tau_i = 120.7(1.2)$ ns (decay rate $\Gamma_i = 2\pi \times 1.32$ MHz [56]), and Rydberg state $|r\rangle = |70S_{1/2}, m_J = -1/2\rangle$ with natural radiative lifetime $\tau_r = 375$ μs (decay rate $\Gamma_r = 2\pi \times 2.67$ kHz [57, 34]). The 420 nm laser drives the lower-leg transition ($|g\rangle \rightarrow |i\rangle$) with Rabi frequency Ω_1 and detuning

$\Delta_1 = 2\pi \times 740$ MHz. The 1013 nm laser drives the upper-leg transition ($|i\rangle \rightarrow |r\rangle$) with Rabi frequency Ω_2 and detuning Δ_2 . The two beams are counter-propagating to suppress first-order Doppler shifts. The 420 nm light is generated by frequency doubling an 840 nm laser, which doubles the phase fluctuations: $\phi_{420}(t) = 2\phi_{840}(t)$. Figure 4.2 shows the energy level diagram for this two-photon excitation scheme.




Figure 4.2: Energy level diagram for Rubidium-87 Rydberg excitation. Two-photon scheme excites atoms from ground state $|g\rangle$ to Rydberg state $|r\rangle$ via intermediate state $|i\rangle$ ($6P_{3/2}$, lifetime 120.7 ns) using 420 nm (Ω_1 , blue) and 1013 nm (Ω_2 , red) lasers with single-photon detuning Δ_1 and two-photon detuning $\delta_{tot} = \Delta_1 + \Delta_2$. Spontaneous emission from intermediate state (Γ_{ig}) and Rydberg state (Γ_{ri} , $\Gamma_{rr'}$, $\Gamma_{rg'}$) included in simulations. The dark state, a coherent superposition of $|g\rangle$ and $|r\rangle$ decoupled from the lossy intermediate state, provides a robust pathway for adiabatic population transfer.

4.1.2 Rydberg Laser System Architecture

The laser stabilization system comprises two independent lasers locked to a common ultra-stable optical cavity. The 840 nm light source is a Ti:Sapphire laser (SolsTis, M Squared Lasers). Frequency doubling in a resonant enhancement cavity generates 420 nm light for the lower-leg transition. The 1013 nm light source is a fiber laser (NP Photonics).

Both lasers lock independently to a high-finesse ultra-low expansion (ULE) cavity with finesse $\approx 50,000$ and linewidth 15 kHz. This provides passive filtering of high-frequency noise and a stable frequency reference for active feedback. The cavity operates under vacuum and temperature-stabilized conditions to suppress thermal drift.

Each laser uses Pound-Drever-Hall (PDH) locking for frequency stabilization (detailed in Section 4.2). The PDH error signal drives multiple actuators with different response bandwidths to enable suppression of frequency noise across the full bandwidth relevant for Rydberg gates. Figure 4.1 shows the complete laser system architecture.

4.1.3 Hamiltonian for Single and Two-Atom Systems

The QOC framework (Section 4.3) requires a mathematical model of system dynamics. We define the single-atom and two-atom Hamiltonians below.

Single-Atom Hamiltonian

In the rotating frame, the single-atom Hamiltonian in frequency units is:

$$\mathcal{H}_0(t) = \frac{1}{2} [\Omega_1(t) (|i\rangle\langle g| + \text{h.c.}) + \Omega_2(t) (|r\rangle\langle i| + \text{h.c.})] - \Delta_1 |i\rangle\langle i| - \delta |r\rangle\langle r| \quad (4.1)$$

where $\delta = \Delta_1 + \Delta_2$ is the two-photon detuning. We set $\Delta_2 = -\Delta_1$, giving $\delta = 0$ to suppress spontaneous decay from the Rydberg state [58]. The effective two-photon Rabi frequency is $\Omega_{\text{eff}} = \Omega_1 \Omega_2 / (2\Delta_1)$.

Two-Atom Hamiltonian and Interaction

For two atoms, the Hilbert space expands to four basis states: $|gg\rangle$, $|gr\rangle$, $|rg\rangle$, and $|rr\rangle$. The Rydberg-Rydberg interaction shifts the doubly-excited state energy through the van der Waals potential:

$$V(R) = \frac{C_6}{R^6} \quad (4.2)$$

where C_6 is the van der Waals coefficient and R is the interatomic separation. For Rb-87 in the $70S_{1/2}$ state, $C_6/h \approx 862 \text{ GHz}\cdot\mu\text{m}^6$. At a typical separation $R = 3 \mu\text{m}$, the interaction strength is $V \approx 2\pi \times 1180 \text{ MHz}$.

The two-atom Hamiltonian in frequency units is:

$$\mathcal{H}_2(t) = \mathcal{H}_0^{(1)}(t) \otimes \mathbb{I} + \mathbb{I} \otimes \mathcal{H}_0^{(2)}(t) + V(R) |rr\rangle\langle rr| \quad (4.3)$$

where $\mathcal{H}_0^{(j)}$ is the single-atom Hamiltonian for atom j . This interaction enables entangling gates through conditional phase accumulation.

4.1.4 Error Channels in Rydberg Gates

Error channels from atomic physics, control field imperfections, and thermal motion limit gate fidelity in Rydberg atom arrays. We organize these into four categories by physical origin:

- **A. System-intrinsic & environmental:** Spontaneous emission from intermediate and Rydberg states, black-body radiation-induced transitions.
- **B. Control field fluctuations:** Laser phase noise, laser amplitude noise, inhomogeneous control fields.
- **C. Atomic motion:** Doppler shifts from thermal velocity, position uncertainty affecting interaction strength.
- **D. State preparation and measurement (SPAM):** Optical pumping errors, detection inefficiency.

These categories differ in addressability: control field errors (B) are suppressible through hardware optimization, pulse design (A, B, C) provides partial mitigation, while atomic motion (C) and fundamental radiative processes (A) establish performance limits for current techniques.

Error Analysis

Category A: System-intrinsic & environmental errors. During two-photon excitation, the intermediate state is populated with probability $\rho_{ii} \approx (\Omega_1^2 + \Omega_2^2)/(4\Delta_1^2)$. The scattering rate from spontaneous emission is $\Gamma_{\text{sc}} = \Gamma_i \rho_{ii}$, exhibiting inverse-square scaling with detuning: $\Gamma_{\text{sc}} \propto \Delta_1^{-2}$ [58]. $\Delta_1 = 2\pi \times 740$ MHz provides a suppression factor $\Gamma_i/\Delta_1 \approx 1.78 \times 10^{-3}$ relative to the bare decay rate while maintaining practical laser power requirements. Larger detuning would further reduce scattering but requires higher laser power to achieve a given effective two-photon Rabi frequency [58].

The long Rydberg state lifetime ($\tau_r = 375 \mu\text{s}$) enables gate times of several hundred nanoseconds with minimal population loss. Black-body radiation drives transitions between Rydberg states with rate scaling as T^3 . For Rb-87 $70S_{1/2}$, room temperature (300 K) yields $\Gamma_{\text{BBR}} \approx 2\pi \times 2.7$ kHz, comparable to the natural radiative rate. Cryogenic operation at 4 K suppresses BBR by a factor of 900, reducing Γ_{BBR} to $2\pi \times 0.003$ kHz. Cryogenic systems operating at 40–50 K achieve $8\times$ suppression of BBR-induced transitions for the $70S_{1/2}$ state, with optical tweezers enabling trap lifetimes exceeding 3000 seconds [59].

Category B: Control field fluctuations. Laser frequency fluctuations introduce time-varying errors in the effective two-photon detuning. Intensity fluctuations cause variations in the Rabi frequencies $\Omega_1(t)$ and $\Omega_2(t)$, leading to pulse area errors and AC Stark shift variations. Active intensity stabilization using acousto-optic modulators suppresses relative intensity noise (RIN) to levels below 10^{-3} in the relevant frequency band.

Category C: Atomic motion. Atoms at temperature T have thermal velocity distribution with standard deviation $\sigma_v = \sqrt{k_B T / m}$. At $T = 25 \mu\text{K}$, the Doppler-broadened linewidth FWHM of 51 kHz corresponds to 1% of typical Rabi frequencies ($\Omega_{\text{eff}} \sim 2\pi \times 4 \text{ MHz}$).

For two-atom gates, shot-to-shot fluctuations in initial atom positions cause interaction strength variations through the R^{-6} dependence. Atoms are prepared in optical tweezers at temperature T before release and Rydberg excitation. The thermal distribution gives position uncertainty $\sigma_R = \sqrt{k_B T / (m\omega^2)}$. At $T = 25 \mu\text{K}$ with trap frequency $\omega = 2\pi \times 50 \text{ kHz}$, $\sigma_R = 156 \text{ nm}$. For $R_0 = 3 \mu\text{m}$, fractional position uncertainty $\sigma_R / R_0 = 5.2\%$ yields RMS interaction strength variation $\sigma_V / V_0 = 6(\sigma_R / R_0) = 31\%$ across experimental runs.

Category D: State preparation and measurement. Imperfect optical pumping during state preparation and finite detection efficiency during measurement contribute a baseline infidelity.

Error Mitigation Strategies

Laser stabilization hardware optimization (Section 4.2) suppresses control field phase and amplitude noise (Category B). Pulse shaping via quantum optimal control (Section 4.3) provides robustness to multiple error sources including spontaneous emission (Category A), control field fluctuations (Category B), and low-frequency parameter variations from atomic motion (Category C). We demonstrate this approach for single-atom excitation (Section 4.4) and two-qubit entangling gates (Section 4.5).

4.2 Laser Phase Noise Engineering for Rydberg Excitation

Laser frequency fluctuations introduce time-dependent errors in the two-photon detuning $\delta(t)$, causing dephasing throughout the gate operation. For a gate with two-photon Rabi frequency Ω_{eff} , phase noise in the frequency band near Ω_{eff} couples most strongly to the quantum evolution. Uncontrolled phase noise limits gate fidelity to 99% even for gates faster than all decoherence rates. Fault-tolerant error thresholds demand systematic suppression of this noise.

We develop a physics-based optimization framework for laser stabilization targeting Rydberg excitation fidelity. Section 4.2.1 quantifies the impact of phase noise on excitation performance. Section 4.2.2 introduces filter function analysis, a linear-response theory mapping measured laser frequency noise spectra to predicted excitation infidelity. Section 4.2.3 uses this metric for systematic optimization of servo parameters, yielding measured noise spectra and predicted sub- 10^{-3} infidelity for operating Rabi frequencies $\Omega_{\text{eff}} > 2\pi \times 4$ MHz.

4.2.1 Rydberg Laser Frequency Stabilization

Free-running lasers exhibit frequency instability across a wide range of timescales. Intrinsic linewidths from spontaneous emission, thermal fluctuations in laser cavities, and acoustic vibrations produce frequency noise spanning from sub-Hz to MHz frequencies. For a two-photon Rydberg transition, the effective detuning experienced by the atoms is $\delta(t) = \delta_0 + \delta_{\text{noise}}(t)$, where $\delta_{\text{noise}}(t)$ arises from the sum of phase noise contributions from both lasers: $\delta_{\text{noise}}(t) = 2\nu_{840}(t) + \nu_{1013}(t)$. The factor of 2 accounts for frequency doubling to 420 nm.

Frequency noise couples to the quantum evolution through the Hamiltonian (Eq. 4.1), causing shot-to-shot variations in the accumulated phase. Gate fidelity degrades when the noise-induced phase uncertainty σ_ϕ becomes comparable to 2π . For typical Rydberg gates with durations $T \sim 100\text{--}500$ ns, noise in the frequency range $f \sim 1/T \sim 2\text{--}10$ MHz is most detrimental.

Locking to a high-finesse optical cavity provides two benefits. First, the cavity acts as a narrow-bandwidth frequency reference. The ULE cavity with finesse $\mathcal{F} \approx 50,000$ has linewidth $\Delta\nu_{\text{cav}} \approx 15$ kHz. Second, the cavity passively filters high-frequency noise at $f \gg f_c$ through its optical transfer function. Active feedback via Pound-Drever-Hall locking

extends suppression to lower frequencies, with performance limited by servo bandwidth and actuator noise. We lock both lasers to the high-finesse ULE cavity. The cavity acts as a frequency discriminator with optical transfer function $|H_{cav}(f)|^2 = 1/(1 + (f/f_c)^2)$, where f_c is the cavity pole frequency [21]. High-frequency noise at $f \gg f_c$ is passively suppressed by the cavity's narrow linewidth. Low-frequency drift and noise within the cavity bandwidth require active feedback via Pound-Drever-Hall (PDH) locking [20].

In the PDH scheme, the laser is phase-modulated before entering the cavity. Only light at the cavity resonance frequency is transmitted, while off-resonant light is reflected. The reflected light is detected and demodulated to produce a dispersive error signal proportional to frequency offset from cavity resonance [21]. This error signal drives a feedback loop correcting laser frequency deviations.

The PDH error signal feeds an FPGA-based digital servo controller (Stable Laser Systems) that implements proportional-integral-derivative (PID) filtering. The servo outputs drive frequency actuators with different response bandwidths. For the 840 nm system, a slow piezoelectric transducer (PZT) corrects low-frequency drift, a fast PZT addresses intermediate frequencies, and an electro-optic modulator (EOM) provides high-bandwidth correction. For the 1013 nm system, a PZT handles low-to-mid frequency corrections and an acousto-optic modulator (AOM) provides fast feedback.

Figure 4.3 shows the servo control architecture. Each actuator has independent PID gains (proportional, integral, derivative). Optimal lock performance requires tuning PID gains for each actuator. Suboptimal gains result in either insufficient suppression (too low) or loop instability (too high).

4.2.2 Mapping Phase Noise to Excitation Infidelity via Filter Functions

Linear response theory [60] provides a frequency-domain susceptibility function $\tilde{I}_\nu(f; \Omega_{eff}, T)$ quantifying how frequency noise at frequency f couples to excitation errors. For a square pi-pulse protocol with duration $T = \pi/(2\Omega_{eff})$, the detuning susceptibility function is [60]:

$$\tilde{I}_\nu(f) = \frac{\Omega_{eff}^2 [1 + \cos(2\pi f T)]}{2(\Omega_{eff}^2 - (2\pi f)^2)^2} \quad (4.4)$$

We obtain the phase noise error contribution by numerical integration:

$$\epsilon_\phi(\Omega_{eff}) = (2\pi)^2 \int_0^\infty df S_{\nu, 2\gamma}(f) \tilde{I}_\nu(f; \Omega_{eff}) \quad (4.5)$$




Figure 4.3: Servo control architecture for Rydberg laser stabilization system. Two independent PDH servo loops lock the 840 nm and 1013 nm lasers to a common ULE reference cavity. The 840 nm servo provides feedback via three actuators spanning different bandwidths (slow PZT, fast PZT, EOM) while the 1013 nm servo uses two actuators (PZT, AOM). Each actuator is controlled by an independent PID filter with tunable proportional (P), integral (I), and derivative (D) gains. Optimization of these gain parameters minimizes predicted Rydberg excitation infidelity using filter function analysis (Section 4.2.2).

This framework enables rapid optimization by directly computing predicted infidelity from measured noise spectra, without requiring slow on-atom experiments.

4.2.3 Servo Parameter Optimization for Rydberg Excitation

We minimize predicted excitation infidelity through empirical optimization of PID feedback parameters. We tune each servo gain parameter independently through parameter sweeps while evaluating excitation performance using filter function analysis.

The digital servo controller implements PID filtering with independent gain parameters for each frequency actuator. For each gain parameter, we test a range of values while holding all others constant. We engage the lock, allow stabilization, and measure in-loop frequency noise power spectral density (PSD) from the PDH error signal. Each measured PSD is integrated against the filter function (Eq. 4.4) to compute predicted phase noise infidelity ϵ_ϕ for a representative excitation protocol ($\Omega_{eff} = 2\pi \times 4$ MHz, square pulse). We select the optimal gain as the value minimizing ϵ_ϕ .

Figure 4.4 shows a representative example of the 840 nm EOM gain sweep. It plots predicted ϵ_ϕ vs. gain value, revealing a clear minimum at 49.33. The gain sweep reveals a distinct minimum in predicted infidelity, demonstrating sensitivity of lock performance to precise gain parameter selection. We perform similar optimization for the 1013 nm AOM gain. The non-monotonic optimization curve arises from competing effects: insufficient gain yields inadequate noise suppression, while excessive gain introduces loop instability and actuator-induced noise. This trade-off creates an optimal operating point balancing suppression and stability.

The optimized parameters reduce predicted excitation infidelity. This optimization directly targets excitation fidelity rather than proxy metrics such as integrated RMS frequency noise. For example, noise at frequencies far from the susceptibility function peaks contributes negligibly to excitation error even if it increases total integrated RMS noise.

Optimized Noise Spectra

We extract in-loop frequency noise PSDs for the 840 nm and 1013 nm lasers from the PDH error signal voltage noise. We model the cavity transfer function as a single-pole low-pass filter that suppresses high-frequency noise. We construct the effective two-photon frequency noise PSD as $S_{\nu,2\gamma}(f) = 4S_{\nu,840}^{eff}(f) + S_{\nu,1013}^{eff}(f)$, where the cavity-filtered spectra are denoted by the superscript "eff" and the factor of 4 accounts for frequency doubling to




Figure 4.4: Systematic Optimization of the 840 nm Laser EOM Gain. **(a)** Predicted excitation infidelity from phase noise (ϵ_ϕ) versus EOM proportional gain. The infidelity is calculated for single-atom Rydberg excitation using a square pi-pulse with two-photon Rabi frequency $\Omega_{eff} = 2\pi \times 4$ MHz. The sweep reveals a distinct minimum at gain = 49.33, yielding predicted infidelity of 6.2×10^{-4} . **(b)** Frequency noise power spectral density comparison at optimal gain (49.33, blue) versus slightly sub-optimal gain (45.0, gray), showing noise suppression achieved through fine-tuning. Similar optimization of the 1013 nm AOM proportional gain yielded an optimal value of 0.2.

420 nm. Figure 4.5 shows measured frequency noise PSDs for 840 nm and 1013 nm after cavity filtering, plus the combined effective two-photon $S_{\nu,2\gamma}(f)$.




Figure 4.5: Impact of PID Gain Fine-Tuning on Laser Noise Spectra. **(a)** Frequency noise ASD for the 840 nm laser, comparing the final optimal EOM gain (49.33, blue) to a slightly sub-optimal setting (45.0, gray). The optimization provides suppression in the 100 kHz to 1 MHz band. **(b)** Frequency noise ASD for the 1013 nm laser with optimized AOM gain of 0.2. **(c)** Combined effective two-photon phase noise PSDs for the optimized (blue) and sub-optimal (gray) configurations. The overlay shows the gate filter function for a $\Omega_{eff} = 2\pi \times 4$ MHz gate, highlighting that noise suppression occurs in the frequency band where the gate is most sensitive.

These effective noise PSDs serve as inputs for calculating excitation infidelity via filter function integrals (Eq. 4.5).

Protocol Performance Scaling with Rabi Frequency

Integrating the measured frequency noise PSDs against the filter function (Section 4.2.2) predicts phase noise infidelities $\epsilon_\phi < 10^{-3}$ for two-photon Rabi frequencies in the range $\Omega_{eff} > 2\pi \times 4$ MHz. Figure 4.6 shows (a) susceptibility functions $\tilde{I}_\nu(f)$ for several Rabi frequencies, overlaid with measured noise PSDs to show which frequency bands dominate error contributions, and (b) predicted $\epsilon_\phi(\Omega_{eff})$ vs. Rabi frequency.

At the operating point $\Omega_{eff} = 2\pi \times 4$ MHz used in the PID optimization sweep (Figure 4.4), the optimized system achieves predicted infidelity $\epsilon_\phi = 6.2 \times 10^{-4}$. The infidelity continues to decrease at higher Rabi frequencies, reaching $\epsilon_\phi = 1.1 \times 10^{-4}$ at $\Omega_{eff} = 2\pi \times 20$ MHz. For the practical operating range $\Omega_{eff} > 2\pi \times 4$ MHz, phase noise infidelity remains below 10^{-3} .

Systematic PID optimization reduces laser phase noise to levels where excitation infidelities from this channel are below 10^{-3} for Rabi frequencies above $2\pi \times 4$ MHz. This




Figure 4.6: Predicted Rydberg Excitation Infidelity from Optimized Laser Phase Noise. **(a)** The final, optimized effective two-photon phase noise PSD overlaid with filter functions for several Rabi frequencies. The peaks of the filter functions shift to higher frequencies for faster excitation protocols. **(b)** Predicted Rydberg excitation infidelity from phase noise only, ϵ_ϕ , as a function of the two-photon Rabi frequency Ω_{eff} . The infidelity decreases for faster protocols, reaching a minimum of 1.1×10^{-4} at $\Omega_{eff} = 2\pi \times 20$ MHz for a square pulse. This analysis considers only phase noise and does not include contributions from laser amplitude noise.

optimization directly targets predicted excitation infidelity rather than proxy metrics. The framework identifies the optimal operating point for our hardware, enabling informed experimental choices. Operating at $\Omega_{eff} = 2\pi \times 4$ MHz balances minimizing phase noise infidelity with practical laser power requirements. These predictions provide quantitative targets for pulse optimization (Section 4.3) and identify frequency bands where further hardware improvements yield the greatest fidelity gains.

4.3 Synthesizing Fast, Robust and High-fidelity Adiabatic Control Operations

Laser stabilization (Section 4.2) suppresses phase noise to levels where excitation infidelity from this source falls below 10^{-3} . However, residual phase noise, laser amplitude fluctuations, atomic motion, and spontaneous emission from intermediate and Rydberg states impose additional constraints on achievable gate fidelity. Furthermore, achieving high fidelity across large atomic arrays requires pulse robustness to spatial inhomogeneities in control field intensity and detuning.

Existing pulse protocols face fundamental trade-offs between speed, fidelity, and robustness. Analytical adiabatic protocols such as stimulated Raman adiabatic passage (STIRAP) [61, 62] provide robust population transfer via dark states but require long gate durations that increase sensitivity to decoherence. Shortcut-to-adiabaticity (STA) protocols [63, 64] accelerate adiabatic evolution but require precise control fields engineering.

We introduce a quantum optimal control (QOC) framework [65, 66, 67] for synthesizing Rydberg gate protocols (Fig. 4.7). The dynamics use the Hamiltonians defined in Section 4.1.3. The optimization employs a cost function combining final-state fidelity and adiabatic pathway adherence (Section 4.3.2). Hardware constraints including finite modulation bandwidth and maximum laser power are enforced within the optimization (Section 4.3.3). Gradient computation uses the Van Loan formalism [68, 69, 70, 71].

4.3.1 Quantum Optimal Control




Figure 4.7: Quantum optimal control framework for synthesizing Rydberg gate protocols. The iterative loop parametrizes control pulses using Chopped Random Basis (CRAB) expansion with Gaussian basis functions characterized by amplitudes $a_{k,n}$, centers $b_{k,n}$, and widths $c_{k,n}$. Regularization enforces hardware amplitude limits through $\Omega_{\max} \tanh(|\Omega|)$. The Van Loan propagator computes time evolution $U(T)$ and adiabatic metric $A(T)$ from the time-dependent Hamiltonian $H(t)$. The cost function $\Phi = w_F(1 - F) + w_A\Phi_A$ balances final-state fidelity F and adiabatic pathway adherence. Automatic differentiation provides gradients $\nabla_{a,b,c}\Phi$, and the Adam optimizer updates parameters for N iterations until convergence.

The QOC framework consists of four components in an iterative loop: a system model, a numerical propagator, a multi-objective cost function, and an optimizer. The system dynamics are governed by the single-atom and two-atom Hamiltonians defined in Section

4.1.3 (Fig. 4.2). For the single-atom case, coherent evolution follows Eq. 4.1. For two atoms, the interaction term (Eq. 4.2) shifts the doubly-excited state $|rr\rangle$ via the two-atom Hamiltonian (Eq. 4.3).

We incorporate decoherence via a Lindblad master equation with spontaneous emission rates defined in Section 4.1.4. The full dynamics are governed by:

$$\frac{d\rho}{dt} = -i[H(t), \rho] + \Gamma_i \mathcal{L}[|g\rangle\langle i|]\rho + \Gamma_r \mathcal{L}[|g\rangle\langle r|]\rho \quad (4.6)$$

where $\mathcal{L}[L]\rho = L\rho L^\dagger - \frac{1}{2}\{L^\dagger L, \rho\}$ is the Lindblad superoperator and $H(t)$ is the time-dependent Hamiltonian with control fields $\Omega_1(t)$ and $\Omega_2(t)$. The instantaneous dark state is the coherent superposition

$$|D(t)\rangle = \frac{\Omega_2(t)|g\rangle - \Omega_1(t)|r\rangle}{\sqrt{\Omega_1^2(t) + \Omega_2^2(t)}} \quad (4.7)$$

which is decoupled from the lossy intermediate state $|i\rangle$ for all control field amplitudes.

Cost Function

The cost function simultaneously optimizes final-state fidelity and adherence to a robust evolution pathway:

$$\Phi_{\text{total}} = w_F \cdot \Phi_{\text{infidelity}} + w_A \cdot \Phi_{\text{adiabaticity}} \quad (4.8)$$

where weights w_F and w_A set the relative importance of fidelity and adiabaticity. The infidelity term is $\Phi_{\text{infidelity}} = 1 - F$, where $F = |\langle\psi_{\text{target}}|U(T)|\psi_{\text{initial}}\rangle|^2$ for state transfer or $F = |\text{Tr}(U_{\text{target}}^\dagger U(T))|^2/d^2$ for process fidelity.

The adiabaticity term quantifies adherence to the instantaneous dark state throughout evolution. We define the adiabatic evolution integral

$$A(T) = \int_0^T dt U^\dagger(t) |D(t)\rangle \langle D(t)| U(t) \quad (4.9)$$

measuring the accumulated overlap with the dark state trajectory. The adiabaticity cost is

$$\Phi_{\text{adiabaticity}} = 1 - \frac{1}{T} |\langle\psi_{\text{target}}|A(T)|\psi_{\text{initial}}\rangle|^2 \quad (4.10)$$

This term penalizes diabatic transitions out of the dark state. Maximizing dark state overlap throughout evolution provides robustness to slow parameter variations, as systems adiabatically track instantaneous eigenstates when parameters drift slowly compared to the instantaneous energy gap.

Hardware Constraint Enforcement

We integrate hardware constraints into the optimization. We model the finite bandwidth of control hardware by applying a digital low-pass filter (50 MHz cutoff, Gaussian profile) to control pulses within each optimization iteration. This filtering is part of forward propagation, so the optimizer adapts pulse shapes to the bandwidth constraint.

We enforce maximum pulse amplitudes $\Omega_{\max}$ through hyperbolic tangent regularization:

$$\Omega_j^{\text{phys}}(t) = \Omega_{\max} \frac{\Omega_j(t)}{|\Omega_j(t)|} \tanh(|\Omega_j(t)|) \quad (4.11)$$

This smooth constraint function provides continuous gradients for efficient gradient-based optimization.

Numerical Implementation

The optimization employs key techniques for efficient pulse design. The Van Loan auxiliary matrix formalism [71] enables simultaneous evaluation of the time evolution operator $U(T)$ and adiabatic integral $A(T)$. The Chopped Random Basis (CRAB) algorithm [72, 73] parametrizes control pulses as superpositions of Gaussian basis functions:

$$\Omega_j(t) = \sum_{n=1}^{N_j} A_{j,n} \exp\left(-\frac{(t - b_{j,n})^2}{c_{j,n}^2}\right) \quad (4.12)$$

where $A_{j,n}$, $b_{j,n}$, and $c_{j,n}$ are the amplitudes, centers, and widths. The amplitudes are optimization parameters, while centers and widths are initialized to mimic adiabatic protocols. For single-atom STIRAP-like transfer, Ω_2 centers are concentrated in the first half of the pulse ($b_{2k} < T/2$) and Ω_1 centers in the second half ($b_{1k} > T/2$), matching the counter-intuitive pulse ordering where the Stokes field precedes the pump to maintain dark state adiabaticity. For two-qubit gates, initialization follows the Rao analytical protocol [74]. This physically-inspired initialization improves convergence. We optimize real and imaginary parts of each Rabi frequency independently for full amplitude and phase control.

Automatic differentiation combined with the Adam optimizer (learning rate 0.01, 1000 iterations) provides efficient gradient-based optimization.

4.4 Application I: Single-Atom Rydberg Excitation

We demonstrate the QOC framework on single-atom Rydberg excitation. High-fidelity population transfer from ground state $|g\rangle$ to Rydberg state $|r\rangle$ through a lossy intermediate state $|i\rangle$ requires suppressing spontaneous emission while maintaining speed sufficient to minimize decoherence.

The optimization problem is defined by the Hamiltonian (Eq. 4.1), with time-dependent Rabi frequencies $\Omega_1(t)$, $\Omega_2(t)$ and single-photon detuning Δ_1 as control parameters. The objective is to maximize state transfer fidelity $F_{\text{state}} = |\langle r|U(T)|g\rangle|^2$, where $U(T)$ is the unitary propagator at gate time T . We compare QOC-optimized pulses against STIRAP and SATD benchmarks, evaluating transfer efficiency, gate duration, and robustness to parameter variations.

Figure 4.8 compares QOC pulses against SATD and STIRAP protocols. Operating parameters: $\Omega_1 = \Omega_2 = 2\pi \times 127$ MHz on two-photon resonance ($\Delta_1 = 0$). The temporal profiles in panel (a) show smooth waveforms satisfying the 50 MHz AOM bandwidth limit. Population dynamics in panel (b) demonstrate efficient transfer to the Rydberg state with minimal time in the lossy intermediate state. The QOC protocol reduces the time-integrated intermediate-state population $T_i = \int_0^T P_i(t)dt$ to 0.165 ns, compared to 0.347 ns for SATD and 0.602 ns for STIRAP.

Panel (d) shows transfer efficiency as a function of pulse duration. QOC achieves 99.7% transfer efficiency at 23.90 ns, compared to 143.36 ns for STIRAP at the same efficiency, representing a sixfold speed improvement.

Figure 4.9 quantifies pulse robustness. QOC pulses demonstrate robustness to parameter variations, maintaining high fidelity over broad ranges of detuning and Rabi frequency fluctuations compared to diabatic and SATD protocols.

Using $N_j = 10$ basis functions per field balances expressiveness and computational efficiency (Fig. 4.10). Too few basis functions limit achievable fidelity, while too many increase computational cost and variability. Figure 4.10 shows infidelity decreases from $\sim 10^{-2}$ with 5 basis functions to $\sim 2 \times 10^{-3}$ with 10 basis functions, then plateaus.

4.5 Application II: Two-Qubit Entangling Gates

Universal quantum computation requires entangling gates in addition to single-qubit rotations. In neutral atom arrays, the controlled-Z (CZ) gate uses the Rydberg blockade




Figure 4.8: Optimized pulses achieve sixfold speed improvement for single-atom Rydberg excitation. **(a)** Temporal pulse profiles for superadiabatic transitionless driving (SATD, dashed) and quantum optimal control (QOC, solid). Red: Ω_1 (420 nm), Blue: Ω_2 (1013 nm). QOC generates smooth control waveforms satisfying the 50 MHz acousto-optic modulator (AOM) bandwidth. **(b)** Population evolution for (ground state: red, solid) and Rydberg state (blue, solid). **(c)** Dark state overlap $|\langle D|\psi \rangle|^2$ throughout evolution. QOC (red) and SATD (blue) maintain higher overlap than STIRAP (green). **(d)** Transfer efficiency as a function of pulse duration. QOC achieves 99.7% at 23.90 ns, compared to 143.36 ns for STIRAP (sixfold speed-up).




Figure 4.9: Robustness of single-atom Rydberg excitation to parameter variations. (a) Infidelity as a function of two-photon detuning fluctuations δ_{tot} . QOC (green circles) demonstrates robustness near resonance compared to SATD (red) and STIRAP (blue). (b) Infidelity as a function of Rabi frequency fluctuations $\delta\Omega_1$.




Figure 4.10: Infidelity versus number of Gaussian basis functions for single-atom Rydberg excitation (gate time 100 ns, $\Omega_1 = \Omega_2 = 2\pi \times 127$ MHz, $\Delta_1 = 0$). Error bars show variability across multiple optimization runs with different random initializations.

interaction to implement a conditional phase shift: the state $|11\rangle$ accumulates a geometric phase π while other computational basis states remain unchanged [75]. Experimental demonstrations in neutral atom systems achieve CZ gate fidelities exceeding 99.5% [19]. Achieving sub- 10^{-4} infidelity requires suppressing multiple error channels simultaneously.

We apply the QOC framework to design CZ gate pulses operating in the intermediate blockade regime, where the interaction strength V is comparable to the two-photon Rabi frequency Ω_{eff} . The optimization employs the two-atom Hamiltonian (Eq. 4.3) including the van der Waals interaction term $V|rr\rangle\langle rr|$. The control parameters are the Rabi frequencies $\Omega_1(t)$, $\Omega_2(t)$ and their relative phases $\varphi_1(t)$, $\varphi_2(t)$, along with detuning Δ_1 . We quantify performance using process fidelity $F_{\text{process}} = |\text{Tr}(U_{\text{CZ}}^\dagger U(T))|^2/d^2$, where $d = 16$ is the Hilbert space dimension and $U_{\text{CZ}} = \text{diag}(1, 1, 1, -1)$ is the target CZ gate in the computational basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$.

We compare optimized pulses against two benchmark protocols: the diabatic π -gap- π sequence [75] and the analytical adiabatic protocol by Rao and Mølmer [74]. The Rao protocol uses sinusoidal pulse envelopes

$$\Omega_1(t) = \Omega \sin\left(\frac{\pi t}{2\tau}\right), \quad \Omega_2(t) = \Omega \left| \cos\left(\frac{\pi t}{2\tau}\right) \right| \quad (4.13)$$

for $0 \leq t \leq 2\tau$, with optimized detuning $\Delta \sim \Omega \sim 2V$ creating an adiabatic pathway to the doubly-excited state $|rr\rangle$ in the intermediate regime.

Figure 4.11 shows the pulse shapes for the Rao analytical protocol and QOC-optimized pulses. Operating parameters for both protocols: $\Omega_1 = \Omega_2 = 2\pi \times 127$ MHz, single-photon detuning $\Delta_1 = 2\pi \times 36$ MHz. The van der Waals interaction strength differs: Rao protocol uses variable V (numerically optimized for each gate duration), while QOC operates at $V = 2\pi \times 6.3$ MHz. At 209 ns gate duration and zero temperature, QOC achieves 99.68% fidelity compared to 99.22% for the Rao protocol, demonstrating improved performance through optimized pulse shaping.

Figure 4.12a quantifies the performance advantage. At all gate times, the QOC protocol achieves higher fidelity than both the diabatic π -gap- π protocol and the analytical adiabatic protocol. The QOC pulse maintains consistently higher overlap with the two-atom dark state (Figure 4.12b).

Figure 4.13 demonstrates robustness of QOC pulses to parameter variations, maintaining high fidelity over broad ranges of detuning and Rabi frequency fluctuations.

The van der Waals interaction $V(R) \propto R^{-6}$ shifts the doubly-excited Rydberg state $|rr\rangle$, where R is the interatomic separation. The CZ gate's accumulated phase depends linearly




Figure 4.11: Pulse shaping for adiabatic entangling CZ gates. **(a)** Normalized amplitude profiles for Rao analytical protocol (dashed) and QOC-optimized pulses (solid). Blue: Ω_1 , Red: Ω_2 . QOC pulses show smoother evolution compatible with 50 MHz hardware bandwidth. **(b)** Phase evolution φ/π for Ω_1 and Ω_2




Figure 4.12: High-fidelity CZ gate implementation using QOC-optimized pulses. **(a)** Infidelity as a function of gate duration. QOC (green) outperforms Rao analytical adiabatic protocol (blue) and diabatic π -gap- π gate (red) at all pulse durations. QOC achieves 99.68% fidelity at 209 ns compared to 99.22% for Rao protocol. **(b)** Dark state overlap $|\langle D|\psi\rangle|^2$




Figure 4.13: Robustness of two-qubit CZ gates to parameter variations. (a) Infidelity as a function of single-photon detuning fluctuations $\delta\Delta_1$. (b) Infidelity as a function of Rabi frequency $\delta\Omega_1$. QOC (green) maintains lower baseline infidelity compared to diabatic π -gap- π protocol (red) and Rao analytical adiabatic protocol (blue).




Figure 4.14: Thermal position fluctuations dominate CZ gate infidelity at finite temperature. Gate infidelity as a function of atomic temperature for simulations including thermal position fluctuations (blue) or Doppler shifts from thermal velocities (red). Position fluctuations cause infidelity to rise from $\sim 3 \times 10^{-3}$ at 0 K to ~ 0.15 at 25 μK (99.68% to 91.47% fidelity).

on interaction strength, so thermal position fluctuations cause shot-to-shot variations in R that translate into phase errors.

Figure 4.14 quantifies these effects. Simulated fidelity of the QOC-optimized CZ gate degrades from 99.68% at 0 K to 91.47% at 25 μ K due to position fluctuations (blue curve). Doppler shifts from thermal velocities contribute less to infidelity (red curve). Alternative gate protocols can mitigate this position-induced phase error. Echoing rapid adiabatic passage (ERAP) uses sequential RAP pulses to cancel the interaction-dependent dynamical phase while accumulating a robust geometric phase of π [76]. ERAP requires adiabatic following of the instantaneous eigenstate throughout the pulse sequence. Our quantum optimal control framework generates pulse shapes that minimize diabatic transitions without extending gate duration. The optimization approach can be applied to ERAP and similar adiabatic gate protocols.

Experimental demonstrations using time-optimal control have achieved 99.5% CZ fidelity with ^{87}Rb [19] and 99.71% with ^{88}Sr [60]. Time-optimal control pulses minimize integrated Rydberg population, reducing the scattering error probability.

Recent analyses have established a roadmap for Rydberg gate fidelity exceeding 99.9% [77]. Scaling to higher principal quantum numbers extends the radiative lifetime. Alkaline-earth atoms (Sr, Yb) eliminate intermediate state scattering via single-photon Rydberg excitation [60, 78]. To close the gap between current demonstrations and the 99.9% threshold, recent studies propose scaling laser power to triple the Rabi frequency and double the intermediate detuning [19].

4.6 Summary and Outlook

This chapter develops two components for high-fidelity two-qubit gates: servo optimization for Rydberg laser stabilization (Section 4.2) and quantum optimal control pulse design (Section 4.3). The filter function framework maps laser phase noise to Rydberg excitation infidelity, providing a quantitative guide for selecting servo parameters. Simultaneously, the QOC framework generates pulses robust to physical parameter variations while suppressing spontaneous decay.

The filter function framework establishes the necessity of protocol-aware servo design. As demonstrated through ab initio error modeling [79], different quantum operations exhibit varying sensitivities to specific noise sources. Laser frequency noise limits CZ gate fidelity, whereas intensity noise limits many-body quench dynamics. Therefore, optimal stabilization requires shaping the relevant noise parameter’s power spectral density (PSD).

Specifically, we must suppress high-frequency intensity noise for quenches and phase noise for gates. By utilizing information about the phase noise structure, we adjust hardware parameters to optimize the laser noise profile specifically for our experimental system and physical parameters.

In parallel, we utilize adiabatic pulses for their inherent robustness against parameter variations. We then apply quantum optimal control to increase adiabatic following, enabling us to design faster pulses while maintaining higher fidelity. This suppresses spontaneous decay—the primary limitation of standard adiabatic protocols—while maintaining their robust character.

This strategy closes the loop between hardware stabilization and control protocol design. While the current work optimizes hardware based on target protocols and optimizes pulses based on physical parameters, future work can extend this integration. For example, measured phase noise spectra could be incorporated directly into the optimal control cost function, allowing for the design of pulses explicitly robust to specific spectral features of the hardware.

Validating these contributions requires confirming that our noise-aware pulse design accurately predicts experimental performance. We propose a definitive measurement to confirm the model’s accuracy and validate the robustness of our optimal control pulses: characterizing single-atom Rydberg excitation infidelity as a function of two-photon detuning.

Prediction: Our model predicts distinct robustness windows for three specific protocols: primitive π -pulses, STIRAP, and the QOC pulses developed in Section 4.3. Figure 4.9 illustrates the expected Rydberg excitation infidelity profiles for STIRAP and QOC pulses, showing broad, high-fidelity plateaus compared to the narrow response of a standard π -pulse.

Measurement: We will measure the excitation probability vs. detuning for each protocol.

Validation: If the measured fidelity curves match the specific shapes predicted by our model, it confirms that the QOC pulses successfully suppress spontaneous decay through adiabatic following while retaining the robustness of adiabatic protocols.

This measurement demonstrates that our optimal control strategy efficiently navigates the trade-off between gate speed and robustness. The broader high-fidelity window provided by these pulses allows the system to tolerate experimental drifts without performance degradation. This relaxes the requirements for frequent interleaved recalibration, allowing the machine to operate for longer periods while maintaining high performance.

Regarding two-qubit gates, our simulations currently predict a fidelity of $\approx 91\%$. This limit arises from variations in interaction energy and atom position fluctuations. Consequently, the contribution of this section to two-qubit gate design is methodological. We developed a framework that improves upon existing adiabatic protocols by increasing the dark state following speed while maintaining their inherent robustness. This optimization strategy is generalizable and can be applied to other adiabatic protocols.

Chapter 5

Conclusion

5.1 Summary

This thesis addresses a central question: how to achieve high-fidelity quantum control in the presence of noise and fluctuating parameters. Bridging atomic physics and control engineering, we develop a full-stack approach that spans from optical hardware design and magnetic field stabilization to control pulse optimization. We contribute four distinct strategies: a self-diagnostic sensing framework for the magnetic environment, predictive gate fidelity modeling, laser phase noise optimization for Rydberg lasers, and robust control pulse design.

First, we demonstrate parallel magnetometry across 1,305 optical tweezer sites using Ramsey interferometry to map the magnetic field with $3\text{ }\mu\text{m}$ spatial resolution. We implement single-shot gradiometry to extract magnetic gradients in one measurement, replacing a 33-minute scanning procedure. By treating each atom as both a qubit and a sensor, we establish a self-diagnostic approach to environmental monitoring, enabling the realization of a self-calibrating quantum system.

Second, we develop a Monte Carlo simulation framework that predicts single-qubit gate fidelity by propagating laser and RF synthesizer phase noise power spectral densities through master equation simulations. The framework decomposes gate infidelity into constituent error sources, allowing us to compare control architectures quantitatively based on fundamental design principles before hardware construction. Our analysis reveals that the OPLL architecture achieves a predicted phase-noise-limited infidelity of 8.80×10^{-5} , which improves upon the MZM architecture (8.44×10^{-3}) by two orders of magnitude.

While thermal motion at 25 μK dominates the error budget for site-selective operations, implementing BB1 composite pulses suppresses these amplitude errors to achieve a total predicted infidelity of 4.4×10^{-4} .

The predicted single-qubit infidelity depends on operating conditions: 8.80×10^{-5} for phase-noise-limited operation with stationary atoms, 1.3×10^{-3} including thermal motion at 25 μK with a π -pulse, and $\sim 4 \times 10^{-4}$ with BB1 composite pulses that suppress amplitude errors at the cost of fivefold longer gate duration. At typical operating temperatures, thermal motion contributes more to infidelity than phase noise. This can be compared to reported values on other platforms. Superconducting qubits achieve single-qubit fidelities of 99.9% with gate durations of 13–50 ns [80]. Trapped-ion systems report single-qubit fidelities approaching 99.99% [81]. Recent neutral-atom demonstrations achieve 99.9% single-qubit fidelity [33]. The contribution is methodological: the simulation framework identifies which hardware configurations achieve competitive performance before experimental implementation.

Third, we optimized servo parameters for both the 420 nm and 1013 nm Rydberg lasers. A filter function analysis maps the laser phase noise power spectral density to Rydberg excitation infidelity. This approach enables protocol-specific optimization of the locked laser phase noise profile by shaping the noise spectrum to minimize its impact on the target quantum operation.

Fourth, we design quantum optimal control pulses that suppress spontaneous decay while achieving robustness to slow parameter variations comparable to conventional adiabatic protocols. The QOC framework achieves sixfold speed improvement over analytical protocols for single-atom Rydberg excitation. Experimental validation of the predicted gate fidelities remains future work.

5.2 Outlook

Scaling control architectures for Rydberg atom arrays presents a trade-off between gate connectivity and cycle time. The primary challenge for large-scale processors is achieving high-fidelity control across the entire array in parallel.

Two distinct paradigms have emerged for executing two-qubit gates. In the fixed array approach with beam-steering demonstrated by Graham et al. [82], atoms remain stationary while control beams are steered to address specific qubit pairs. This architecture enables fast gate cycles limited only by the microsecond-scale Rydberg pulse duration.

However, its connectivity is constrained by the beam geometry and the acousto-optic deflector’s field of view. As the number of qubits (N) grows, the calibration overhead scales as N^2 for all possible pairs, leading to prohibitive calibration times for large processors. Maintaining high-fidelity parallel operations requires increasingly sophisticated compensation for spatial inhomogeneities and beam-pointing stability across the full addressing range.

Conversely, the transport-based architecture demonstrated by the Harvard group enables gates between arbitrary qubit pairs by physically moving atoms mid-circuit via dynamic optical tweezers [49]. This approach provides the connectivity required for error-corrected computation, such as syndrome extraction in surface codes [33]. The scaling strategy relies on a dedicated control region illuminated by a global Rydberg beam, calibrated once. By transporting atoms to this localized region for interaction, the system avoids the N^2 calibration bottleneck, as the beam only needs to be characterized once regardless of total atom count.

This calibration efficiency comes at the cost of operation speed and atom retention. The transport-based approach, which requires implementing atom reconfiguration algorithms [83, 84, 85, 86], introduces overhead of several hundred microseconds per gate layer, significantly slowing the cycle time compared to beam-steering architectures. Furthermore, each transport operation carries a finite probability of atom loss [84, 85, 86], which compounds over the many cycles required for deep circuits. Despite these constraints, the transport-based approach has proven more beneficial in the near term for scaling, achieving 99.5% parallel two-qubit gate fidelity [19] and recently demonstrating state-of-the-art performance in logical error-corrected qubits [33].

Trap depth uniformity becomes harder to maintain as arrays scale. For tweezer arrays generated by crossed acousto-optic deflectors, each trap depth is the product of two RF amplitudes, one from each deflector axis. An $N \times M$ array has $N \times M$ traps but only $N + M$ independent amplitude values. A 100×100 array has 10,000 traps constrained by 200 values. This under-actuation prevents independent equalization of trap depths. Non-uniform trap depths cause site-dependent light shifts [87].

Future work will implement adaptive field stabilization using the magnetometry capability. The microwave oscillator provides a stable frequency reference. By interleaving calibration shots with sensing shots, the system can track magnetic field drift and adjust the microwave drive frequency accordingly. The calibration rate must exceed the magnetic field drift rate. This would maintain optimal Ramsey evolution time for continuous sensing applications.

The noise model predicts gate fidelity from component specifications. Randomized

benchmarking would test these predictions experimentally. Discrepancies between predicted and measured fidelity would identify error sources not captured by the current model. This iterative process of noise measurement, fidelity prediction, and experimental validation enables co-development of control protocols and noise characterization.

Future work will focus on homogenizing Rydberg coupling strengths across the array through light-shift spectroscopy at 1013 nm. Quantum lock-in detection [30] enables high-resolution imaging of the Rydberg beam intensity profile with sub-micron resolution, robust against laser power and frequency fluctuations. This technique will characterize spatial variations in the Rydberg laser intensity and guide beam alignment, addressing a key source of site-to-site gate fidelity variation in scaled neutral-atom processors.

Our system has microscope objectives in crossed configuration, enabling tight focusing of Rydberg beams. Holographic beam shaping provides site-selective addressing of atom pairs. This configuration opens questions about which Hamiltonians become accessible with site-selective control of Hamiltonian parameters.

Extending optimal control to multi-qubit gates (CCZ, Toffoli) could reduce overhead from gate decomposition. A single optimized pulse sequence would implement a three-qubit gate without decomposing it into multiple two-qubit operations.

The calibrated atomic array functions as a spatially-resolved magnetometer. Beyond characterizing the qubit environment, this capability suggests applications to imaging magnetic structures in external samples. Spatially-resolved magnetization measurements are of particular interest for two-dimensional magnetic materials—van der Waals crystals that exhibit magnetic order even at the monolayer limit [88]. However, direct proximity-based sensing of such materials poses challenges: magnetic samples produce stray fields that would interfere with the bias field control and qubit frequencies of atoms confined in vacuum.

Magneto-optical techniques offer an alternative that circumvents the proximity requirement. In the Faraday effect, light transmitted through a magnetized sample acquires polarization rotation proportional to the local magnetization. Placing a thin magnetic sample outside the vacuum chamber, the transmitted probe light would carry a spatially-varying polarization pattern encoding the magnetic domain structure. Since atoms are intrinsically sensitive to light polarization, site-resolved detection of this pattern could enable imaging of magnetic domains without requiring the sample to be in close proximity to the sensor array.

Simulating arbitrary Hamiltonians remains an active research challenge. Given a target material or model, how does one map it to a neutral atom array with specific control parameters and interaction physics?

5.3 Moving Forward

As the field pushes the limits of coherence and control, commercial hardware specifications increasingly determine the limits of experimental performance. Scaling to larger arrays or faster gates requires improvements such as higher optical power handling or the development of custom optics to meet specific thermal and optical transmission requirements. Correcting beam astigmatism and ellipticity becomes critical as power requirements increase with array size [89].

To accelerate progress, AMO laboratories must establish tighter feedback loops with laser suppliers and optics manufacturers. By clearly articulating specific performance limitations and engaging in hardware co-development, we can drive the creation of products that specifically address the bottlenecks of neutral-atom quantum computing.

Due to the complexity of the field, onboarding new members takes a long time. Our lab has hosted over 40 brilliant undergraduate students, but it takes time before they can contribute to the experiment. For four-month undergraduate internships, onboarding takes two months, leaving one month of productive work before knowledge transfer begins.

I see an opportunity to rethink how AMO labs approach documentation. Modern software engineering teams utilize large language models (LLMs) to accelerate onboarding, leveraging the models’ ability to explain code and answer questions within the context of a specific codebase. Adopting a similar approach could compress the onboarding timeline in experimental physics labs. Establishing a “single source of truth” with LLM-friendly documentation becomes an important, high-leverage investment for the long-term productivity of the groups.

References

- [1] Peter W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994.
- [2] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, 1982.
- [3] Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma. Exponential improvement in precision for simulating sparse Hamiltonians. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 283–292, 2014.
- [4] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.
- [5] S. Somaroo, C. H. Tseng, T. F. Havel, R. Laflamme, and D. G. Cory. Quantum simulations on a quantum computer. *Physical Review Letters*, 82(26):5381–5384, 1999.
- [6] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52:R2493–R2496, 1995.
- [7] John Preskill. Reliable quantum computers. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 454(1969):385–410, 1998.
- [8] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86:032324, 2012.
- [9] C. L. Degen, F. Reinhard, and P. Cappellaro. Quantum sensing. *Reviews of Modern Physics*, 89:035002, 2017.

- [10] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Advances in quantum metrology. *Nature Photonics*, 5(4):222–229, 2011.
- [11] Jiehui Huang, Ming Zhuang, and Chaohong Lee. Entanglement-enhanced quantum metrology: From standard quantum limit to heisenberg limit. *Applied Physics Reviews*, 11(3):031302, 2024.
- [12] Andrew D. Ludlow, Martin M. Boyd, Jun Ye, Ekkehard Peik, and Piet O. Schmidt. Optical atomic clocks. *Reviews of Modern Physics*, 87(2):637–701, 2015.
- [13] Edgar Pedrozo-Peñafiel, Stefano Colombo, Chao Shu, Albert F. Adiyatullin, Zeyang Li, Esteban Mendez, Boris Braverman, Akihisa Kawasaki, Daisuke Akamatsu, Yun Xiao, and Vladan Vuletić. Entanglement-enhanced optical atomic clock. *Nature*, 588:414–418, 2020.
- [14] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574:505–510, 2019.
- [15] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709–711, 2002.
- [16] Jonathan Wurtz, Alexei Bylinskii, Boris Braverman, Jesse Amato-Grill, Sergio H. Cantu, Florian Huber, Alexander Lukin, Fangli Liu, Phillip Weinberg, John Long, Sheng-Tao Wang, Nathan Gemelke, and Alexander Keesling. Aquila: QuEra’s 256-qubit neutral-atom quantum computer. *arXiv preprint arXiv:2306.11727*, 2023.
- [17] Hannah J. Manetsch, Gyohei Nomura, Elie Bataille, Kon H. Leung, Xudong Ly, and Manuel Endres. A tweezer array with 6100 highly coherent atomic qubits. *Nature*, 2025.
- [18] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with rydberg atoms. *Rev. Mod. Phys.*, 82:2313–2363, Aug 2010.
- [19] Simon J. Evered, Dolev Bluvstein, Marcin Kalinowski, Sepehr Ebadi, Tom Manovitz, Hengyun Zhou, Sophie H. Li, Alexandra A. Geim, Tout T. Wang, Nishad Maskara, Harry Levine, Giulia Semeghini, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. High-fidelity parallel entangling gates on a neutral atom quantum computer. *Nature*, 622:268–272, 2023.

- [20] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B*, 31:97–105, 1983.
- [21] Eric D. Black. An introduction to pound-drever-hall laser frequency stabilization. *American Journal of Physics*, 69(1):79–87, 2001.
- [22] Flavien Gyger, Maximilian Ammenwerth, Renhao Tao, Hendrik Timme, Stepan Snigirev, Immanuel Bloch, and Johannes Zeiher. Continuous operation of large-scale atom arrays in optical lattices. *Physical Review Research*, 6:033104, 2024.
- [23] D. Budker and M. Romalis. Optical magnetometry. *Nature Physics*, 3(4):227–234, 2007.
- [24] Dominik Schäffner, Tobias Schreiber, Fabian Lenz, Malte Schlosser, and Gerhard Birkel. Quantum sensing in tweezer arrays: Optical magnetometry on an individual-atom sensor grid. *PRX Quantum*, 5:010311, 2024.
- [25] M. D. Bowdrey, D. K. L. Oi, A. J. Short, K. Banaszek, and J. A. Jones. Fidelity of single qubit maps. *Physics Letters A*, 294:258–260, 2002.
- [26] Daniel A. Steck. Rubidium 87 d line data, 2021.
- [27] Jean Dalibard, Fabrice Gerbier, Gediminas Juzeliūnas, and Patrik Öhberg. Colloquium: Artificial gauge potentials for neutral atoms. *Reviews of Modern Physics*, 83:1523–1543, 2011.
- [28] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais, and P. Grangier. Energy distribution and cooling of a single atom in an optical tweezer. *Physical Review A*, 78:033425, 2008.
- [29] Philip Krantz, Morten Kjaergaard, Fei Yan, Terry P. Orlando, Simon Gustavsson, and William D. Oliver. A quantum engineer’s guide to superconducting qubits. *Applied Physics Reviews*, 6:021318, 2019.
- [30] Sijie Chen, Min Zhuang, Ruihuang Fang, Yun Chen, Chengyin Han, Bo Lu, Jiahao Huang, and Chaohong Lee. Quantum double lock-in amplifier. *Communications Physics*, 7:189, 2024.
- [31] T. Xia, M. Lichtman, K. Maller, A. W. Carr, M. J. Piotrowicz, L. Isenhower, and M. Saffman. Randomized benchmarking of single-qubit gates in a 2d array of neutral-atom qubits. *Phys. Rev. Lett.*, 114:100503, Mar 2015.

- [32] K. Biswas, L. Saha, J. Denschlag, G. Smirne, S. F. Huelga, and M. B. Plenio. Efficient two-photon excitation of rydberg states in cesium. *Physical Review A*, 69:062506, 2004.
- [33] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin. Logical quantum processor based on reconfigurable atom arrays. *Nature*, 626:58–65, 2024.
- [34] Nikola Šibalić, Jonathan D. Pritchard, Charles S. Adams, and Kevin J. Weatherill. Arc: An open-source library for calculating properties of alkali rydberg atoms. *Computer Physics Communications*, 220:319–331, 2017.
- [35] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and John M. Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014.
- [36] Naresh Satyan. *Optoelectronic Control of the Phase and Frequency of Semiconductor Lasers*. PhD thesis, California Institute of Technology, 2011.
- [37] Sylvain de Léséleuc, Daniel Barredo, Vincent Lienhard, Antoine Browaeys, and Thierry Lahaye. Analysis of imperfections in the coherent optical excitation of single atoms to rydberg states. *Physical Review A*, 97:053803, 2018.
- [38] Gianni Di Domenico, Stéphane Schilt, and Pierre Thomann. Simple approach to the relation between laser frequency noise and laser line shape. *Applied Optics*, 49:4801–4807, 2010.
- [39] J. B. Naber, L. Torralbo-Campo, T. Hubert, and R. J. C. Spreeuw. Raman transitions between hyperfine clock states in a magnetic trap. *arXiv:1605.05230v2*, 2016.
- [40] C. D. Herold, V. D. Vaidya, X. Li, S. L. Rolston, J. V. Porto, and M. S. Safronova. Precision measurement of transition matrix elements via light shift cancellation. *Physical Review Letters*, 109:243003, 2012.
- [41] M. S. Safronova and U. I. Safronova. Critically evaluated theoretical energies, lifetimes, hyperfine constants, and multipole polarizabilities in ^{87}Rb . *Physical Review A*, 83:052508, 2011.

- [42] Tobias J. Kippenberg, Albert Schliesser, and Michael Gorodetsky. Phase noise measurement of external cavity diode lasers and implications for optomechanical sideband cooling of GHz mechanical modes. *arXiv preprint*, 2011.
- [43] Matthew L. Day, Pei Jiang Low, Brendan White, Rajibul Islam, and Crystal Senko. Limits on atomic qubit control from laser noise. *npj Quantum Information*, 8(1):72, 2022.
- [44] Katharina Gillen-Christandl, Glen D. Gillen, M. J. Piotrowicz, and M. Saffman. Comparison of Gaussian and super Gaussian laser beams for addressing atomic qubits. *Physical Review A*, 94(1):013406, 2016.
- [45] Sin Hyuk Yim, Sang-Bum Lee, Taeg Yong Kwon, Kyu Min Shim, and Sang Eon Park. Optical phase-locking of two extended-cavity diode lasers by serrodyne modulation. *Applied Optics*, 58(10):2481–2484, 2019.
- [46] Niklas Kolodzie, Ivan Mirgorodskiy, Christian Nölleke, and Piet O. Schmidt. Ultra-low frequency noise external cavity diode laser systems for quantum applications. *Optics Express*, 32(16):29781–29794, 2024.
- [47] TOPTICA Photonics SE. DL pro – Tunable Diode Laser. <https://www.toptica.com/products/tunable-diode-lasers/ecdl-dfb-lasers/dl-pro/>, 2024. Accessed: 2024.
- [48] S. Wimperis. Broadband, narrowband, and passband composite pulses for use in advanced NMR experiments. *Journal of Magnetic Resonance, Series A*, 109(2):221–231, 1994.
- [49] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletić, and M. D. Lukin. A quantum processor based on coherent transport of entangled atom arrays. *Nature*, 604:451–456, 2022.
- [50] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103(3):259–328, 1998.
- [51] Marko Cetina. *Hybrid Approaches to Quantum Information Using Ions, Atoms and Photons*. PhD thesis, Massachusetts Institute of Technology, 2011.

- [52] B. Nikolov, E. Diamond-Hitchcock, J. Bass, N. L. R. Spong, and J. D. Pritchard. Randomized benchmarking using non-destructive readout in a 2d atom array. *Physical Review Letters*, 131:030602, 2023.
- [53] Harry Levine, Dolev Bluvstein, Alexander Keesling, Tout T. Wang, Sepehr Ebadi, Giulia Semeghini, Ahmed Omran, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. Dispersive optical systems for scalable raman driving of hyperfine qubits. *Physical Review A*, 105:032618, 2022.
- [54] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. *Physical Review A*, 77:012307, 2008.
- [55] Easwar Magesan, Jay M. Gambetta, and Joseph Emerson. Scalable and robust randomized benchmarking of quantum processes. *Physical Review Letters*, 106:180504, 2011.
- [56] U. D. Rapol, A. Krishna, A. Wasan, and V. Natarajan. Lifetime and hyperfine splitting measurements on the 7s and 6p levels in rubidium. *Journal of the Optical Society of America B*, 21(11):2058–2062, 2004.
- [57] I. I. Beterov, I. I. Ryabtsev, D. B. Tretyakov, and V. M. Entin. Radiative lifetime measurements of rubidium Rydberg states. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 43:015002, 2010.
- [58] Caroline Laure Mbakob Tchouawou. Simulation and mitigation of noise in rydberg atom systems using quantum optimal control. Master’s thesis, University of Waterloo, Waterloo, Ontario, Canada, 2021. Master of Science in Physics (Quantum Information).
- [59] Zhenpu Zhang, Ting-Wei Hsu, Ting You Tan, Daniel H. Slichter, Adam M. Kaufman, Matteo Marinelli, and Cindy A. Regal. A high optical access cryogenic system for rydberg atom arrays with a 3000-second trap lifetime. *PRX Quantum*, 6:020337, 2025.
- [60] Richard Bing-Shiun Tsai, Xiangkai Sun, Adam L. Shaw, Ran Finkelstein, and Manuel Endres. Benchmarking and fidelity response theory of high-fidelity rydberg entangling gates. *PRX Quantum*, 6:010331, 2025.

- [61] N. V. Vitanov, A. A. Rangelov, B. W. Shore, and K. Bergmann. Stimulated raman adiabatic passage in physics, chemistry, and beyond. *Rev. Mod. Phys.*, 89:015006, 2017.
- [62] K. Bergmann, N. V. Vitanov, and B. W. Shore. Perspective: Stimulated raman adiabatic passage: The status after 25 years. *J. Chem. Phys.*, 142:170901, 2015.
- [63] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga. Shortcuts to adiabaticity: Concepts, methods, and applications. *Rev. Mod. Phys.*, 91:045001, 2019.
- [64] M. Nakahara. Counterdiabatic formalism of shortcuts to adiabaticity. *Phil. Trans. R. Soc. A*, 380:20210272, 2022.
- [65] Christiane P. Koch, Ugo Boscain, Tommaso Calarco, Gunther Dirr, Stefan Filipp, Steffen J. Glaser, Ronnie Kosloff, Simone Montangero, Thomas Schulte-Herbrüggen, Dominique Sugny, and Frank K. Wilhelm. Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe. *EPJ Quantum Technology*, 9:19, 2022.
- [66] Callum W. Duncan, Pablo M. Poggi, Marin Bukov, Nikolaj Thomas Zinner, and Steve Campbell. Taming quantum systems: A tutorial for using shortcuts-to-adiabaticity, quantum optimal control, and reinforcement learning. *PRX Quantum*, 6:040201, 2025.
- [67] Q. Ansel, E. Dionis, F. Arrouas, B. Peaudecerf, S. Guérin, D. Guéry-Odelin, and D. Sugny. Introduction to theoretical and experimental aspects of quantum optimal control. *J. Phys. B: At. Mol. Opt. Phys.*, 57:133001, 2024.
- [68] C. Van Loan. Computing integrals involving the matrix exponential. *IEEE Transactions on Automatic Control*, 23(3):395–404, 1978.
- [69] David L. Goodwin and Ilya Kuprov. Auxiliary matrix formalism for interaction representation transformations, optimal control and spin relaxation theories. *Journal of Chemical Physics*, 143:084113, 2015.
- [70] Holger Haas, Daniel Puzzuoli, Feihao Zhang, and David G. Cory. Engineering effective hamiltonians: a general framework for time-dependent perturbation and control. *New J. Phys.*, 21:103011, 2019.
- [71] Sahand Tabatabaei, Holger Haas, William Rose, Ben Yager, Michèle Piscitelli, Pardis Sahafi, Andrew Jordan, Philip J. Poole, Dan Dalacu, and David G. Cory. Numerical engineering of robust adiabatic operations. *Physical Review Applied*, 15:044043, 2021.

- [72] Tommaso Caneva, Tommaso Calarco, and Simone Montangero. Chopped random-basis quantum optimization. *Physical Review A*, 84:022326, 2011.
- [73] M. M. Müller, R. S. Said, F. Jelezko, T. Calarco, and S. Montangero. One decade of quantum optimal control in the chopped random basis. *Rep. Prog. Phys.*, 85:076001, 2022.
- [74] D. D. Bhaktavatsala Rao and Klaus Mølmer. Robust rydberg-interaction gates with adiabatic passage. *Physical Review A*, 89:030301, 2014.
- [75] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin. Fast quantum gates for neutral atoms. *Physical Review Letters*, 85:2208–2211, 2000.
- [76] Ming Xue, Shijie Xu, Xinwei Li, and Xiangliang Li. High-fidelity and robust controlled- z gates implemented with rydberg atoms via echoing rapid adiabatic passage. *Physical Review A*, 110:032619, Sep 2024.
- [77] M. Saffman. Quantum computing with atomic qubit arrays: confronting the cost of connectivity. *arXiv preprint*, 2025.
- [78] Michael Peper, Yiyi Li, Daniel Tock, Mariia Popova, Jingjing Jiang, Pritish Nair, Eric Megidish, and Hannes Bernien. Spectroscopy and modeling of ^{171}Yb Rydberg states for high-fidelity two-qubit gates. *Physical Review X*, 15:011009, 2025.
- [79] Adam L. Shaw, Zhuo Chen, Joonhee Choi, Daniel K. Mark, Pascal Scholl, Ran Finkelstein, Andreas Elben, Soonwon Choi, and Manuel Endres. Benchmarking highly entangled states on a 60-atom analogue quantum simulator. *Nature*, 628:71–77, 2024.
- [80] Muhammad AbuGhanem. Superconducting quantum computers: Who is leading the future? *EPJ Quantum Technology*, 12:102, 2025.
- [81] Thomas Strohm, Karen Wintersperger, Florian Dommert, Daniel Basilewitsch, Georg Reuber, Andrey Hoursanov, Thomas Ehmer, Davide Vodola, and Sebastian Luber. Ion-based quantum computing hardware: Performance and end-user perspective. *arXiv preprint*, 2024.
- [82] T. M. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer, M. Kwon, M. Ebert, J. Cherek, M. T. Lichtman, M. Gillette, J. Gilbert, D. Bowman, T. Ballance, C. Campbell, E. D. Dahl, O. Crawford, N. S. Blunt, B. Rogers, T. Noel, and M. Saffman. Multi-qubit entanglement and algorithms on a neutral-atom quantum computer. *Nature*, 604:457–462, 2022.

- [83] A. Cooper, S. Maaz, A. E. Mouawad, and N. Nishimura. Parameterized complexity of reconfiguration of atoms. *Algorithmica*, 86(10):3284–3308, 2024.
- [84] B. Cimring, R. El Sabeh, M. Bacvanski, S. Maaz, I. El Hajj, N. Nishimura, A. E. Mouawad, and A. Cooper. Efficient algorithms to solve atom reconfiguration problems. i. redistribution-reconfiguration algorithm. *Phys. Rev. A*, 108:023107, 2023.
- [85] R. El Sabeh, J. Bohm, Z. Ding, S. Maaz, N. Nishimura, I. El Hajj, A. E. Mouawad, and A. Cooper. Efficient algorithms to solve atom reconfiguration problems. ii. assignment-rerouting-ordering algorithm. *Phys. Rev. A*, 108:023108, 2023.
- [86] F. Afiouni, R. El Sabeh, N. Nishimura, I. El Hajj, A. E. Mouawad, and A. Cooper. Efficient algorithms to solve atom reconfiguration problems. iii. the bird and batching algorithms and other parallel implementations on gpus. *Phys. Rev. A*, 112:023109, 2025.
- [87] Soroush Khoubyarian. Systems and methods for generating large arrays of optical traps in neutral atom array quantum processors. Master’s thesis, University of Waterloo, 2025.
- [88] B. Huang, G. Clark, E. Navarro-Moratalla, D. R. Klein, R. Cheng, K. L. Seyler, D. Zhong, E. Schmidgall, M. A. McGuire, D. H. Cobden, W. Yao, D. Xiao, P. Jarillo-Herrero, and X. Xu. Layer-dependent ferromagnetism in a van der waals crystal down to the monolayer limit. *Nature*, 546:270–273, 2017.
- [89] Soroush Khoubyarian, Anastasiia Mashko, and Alexandre Cooper. Correcting astigmatism and ellipticity in Gaussian beams using a cylindrical lens pair with tunable focal lengths. *Optics Express*, 33(22):47292–47320, Nov 2025.

Appendix A

System Hamiltonian for Raman Simulations

This appendix provides complete expressions for the Hamiltonians used in Chapter 3 simulations. The formulas document the quantum model, intermediate state specifications, and calculation methods enabling reproduction of the simulation results.

A.1 General Hamiltonian Structure

The system Hamiltonian models the qubit states $|g_1\rangle$ and $|g_2\rangle$ and the intermediate excited-state manifold. The qubit states are $|g_1\rangle \equiv |5S_{1/2}, F=1, m_F=0\rangle$ and $|g_2\rangle \equiv |5S_{1/2}, F=2, m_F=0\rangle$, separated by the ground-state hyperfine splitting $\omega_{HF1}/2\pi = 6.8347$ GHz (Rb-87 clock transition). Intermediate states $|e_i\rangle$ correspond to hyperfine levels in the excited P manifold, with $i = 1, 2$ for D1 transitions (795 nm, 421 nm) and $i = 1, 2, 3$ for D2 transitions (780 nm, 420 nm).

The Hamiltonian structure is:

$$H(t) = H_0 + H_{\text{interaction}}(t) \tag{A.1}$$

A.1.1 Five-Level System (D1 Transitions: 795 nm, 421 nm)

For transitions to $j = 1/2$ excited states, the intermediate manifold has two hyperfine levels corresponding to $F = 1$ and $F = 2$. The static Hamiltonian is:

$$H_0 = -\Delta_1|e_1\rangle\langle e_1| - (\Delta_1 - \omega_{HF2})|e_2\rangle\langle e_2| + \delta_{\text{tot}}|g_2\rangle\langle g_2| \quad (\text{A.2})$$

where Δ_1 is the single-photon detuning from $|e_1\rangle$, ω_{HF2} is the splitting between excited-state hyperfine levels (Table 3.3), and δ_{tot} is the two-photon detuning from Raman resonance. The five basis states are $\{|g_1\rangle, |e_1\rangle, |e_2\rangle, |d\rangle, |g_2\rangle\}$, where $|d\rangle$ represents dark states.

A.1.2 Six-Level System (D2 Transitions: 780 nm, 420 nm)

For transitions to $j = 3/2$ excited states, the intermediate manifold has three hyperfine levels ($F = 1, 2, 3$). The static Hamiltonian is:

$$H_0 = -\Delta_1|e_1\rangle\langle e_1| - (\Delta_1 - \omega_{HF2})|e_2\rangle\langle e_2| - (\Delta_1 - \omega_{HF2} - \omega_{HF3})|e_3\rangle\langle e_3| + \delta_{\text{tot}}|g_2\rangle\langle g_2| \quad (\text{A.3})$$

where ω_{HF2} and ω_{HF3} are the splittings between consecutive hyperfine levels (Table 3.3). The six basis states are $\{|g_1\rangle, |e_1\rangle, |e_2\rangle, |e_3\rangle, |d\rangle, |g_2\rangle\}$. The transition $|g_1\rangle \rightarrow |e_3\rangle$ is forbidden by angular momentum conservation ($\Delta F = 3 - 1 = 2 > 1$ for single-photon σ^+ polarization), resulting in zero dipole matrix element $V_{31} = 0$.

A.2 OPLL Hamiltonian

For OPLL architecture using two independent lasers (Laser 1 as reference, Laser 2 as locked), the electric fields at the atoms are $E_1(t) = E_1 e^{i(\omega_1 t + \phi_1(t))}$ and $E_2(t) = E_2 e^{i(\omega_2 t + \phi_2(t))}$, where E_1 and E_2 are field amplitudes, ω_1 and ω_2 are optical frequencies separated by $\omega_2 - \omega_1 \approx \omega_{HF1}$, and $\phi_1(t)$ and $\phi_2(t)$ are time-dependent phase noise traces. The field amplitudes relate to optical powers via $E_i = \sqrt{2P_i/(\epsilon_0 c \pi w_0^2)}$, where P_i is the power in beam i , w_0 is the beam waist, ϵ_0 is the permittivity of free space, and c is the speed of light.

The time-dependent interaction Hamiltonian for a five-level system (795 nm, 421 nm) is $H_{\text{interaction}}(t) = H_{g_1 e_1}(t) + H_{g_2 e_1}(t) + H_{g_1 e_2}(t) + H_{g_2 e_2}(t) + \text{h.c.}$, where each term couples a pair of states via one of the laser fields. For the $|g_1\rangle \leftrightarrow |e_1\rangle$ transition:

$$H_{g_1 e_1}(t) = -\frac{V_{11}}{2\hbar} [E_1 e^{i\phi_1(t)} + E_2 e^{-i\delta\omega t} e^{i\phi_2(t)}] |g_1\rangle\langle e_1| \quad (\text{A.4})$$

where $\delta\omega = \omega_2 - \omega_1 = \omega_{HF1}$ is the frequency difference between lasers and V_{11} is the dipole matrix element connecting $|g_1\rangle$ and $|e_1\rangle$. For the $|g_2\rangle \leftrightarrow |e_1\rangle$ transition:

$$H_{g_2e_1}(t) = -\frac{V_{12}}{2\hbar} [E_2 e^{i\phi_2(t)} + E_1 e^{i\delta\omega t} e^{i\phi_1(t)}] |g_2\rangle\langle e_1| \quad (\text{A.5})$$

Similar expressions apply for $|g_1\rangle \leftrightarrow |e_2\rangle$ (with V_{21}) and $|g_2\rangle \leftrightarrow |e_2\rangle$ (with V_{22}). For six-level systems (780 nm, 420 nm), additional interaction terms $H_{g_1e_3}(t)$ and $H_{g_2e_3}(t)$ are included for the third intermediate state, using dipole matrix elements $V_{31} = 0$ (forbidden) and V_{32} (allowed).

A.3 MZM Hamiltonian

For MZM architecture using a single laser with amplitude modulation, the modulated electric field at the atoms is:

$$E(t) = E_0 e^{i\omega_0 t} e^{i\phi_{\text{laser}}(t)} \sum_{n=-\infty}^{\infty} J_n(\beta) [e^{-in(\omega_m t + \phi_{\text{RF}}(t))} - e^{+in(\omega_m t + \phi_{\text{RF}}(t))}] \quad (\text{A.6})$$

where E_0 is the electric field amplitude (accounting for insertion loss), ω_0 is the optical carrier frequency, $\phi_{\text{laser}}(t)$ is the laser phase noise, $J_n(\beta)$ is the Bessel function of the first kind of order n , β is the modulation depth, ω_m is the RF modulation frequency ($\omega_m \approx \omega_{HF1}/2$), and $\phi_{\text{RF}}(t)$ is the RF phase noise. The factor $[e^{-in\omega_m t} - e^{+in\omega_m t}] = -2i \sin(n\omega_m t)$ arises from the push-pull MZM configuration with carrier suppression, which suppresses even-order sidebands ($n = 0, \pm 2, \pm 4, \dots$) while retaining odd-order sidebands ($n = \pm 1, \pm 3, \pm 5, \dots$).

For the $|g_1\rangle \leftrightarrow |e_1\rangle$ transition:

$$H_{g_1e_1}(t) = \frac{E_0 V_{11}}{\hbar} e^{i\phi_{\text{laser}}(t)} \sum_{n=1,3,5,\dots} J_n(\beta) \sin[n(\omega_m t + \phi_{\text{RF}}(t))] |g_1\rangle\langle e_1| + \text{h.c.} \quad (\text{A.7})$$

For the $|g_2\rangle \leftrightarrow |e_1\rangle$ transition, an additional phase factor $e^{i2\omega_m t}$ accounts for the frequency offset $\omega_{HF1} \approx 2\omega_m$:

$$H_{g_2e_1}(t) = \frac{E_0 V_{12}}{\hbar} e^{i\phi_{\text{laser}}(t)} e^{i2\omega_m t} \sum_{n=1,3,5,\dots} J_n(\beta) \sin[n(\omega_m t + \phi_{\text{RF}}(t))] |g_2\rangle\langle e_1| + \text{h.c.} \quad (\text{A.8})$$

The Raman transition is driven by pairs of sidebands where sideband $+n$ couples $|g_1\rangle \rightarrow |e\rangle$ while sideband $-n$ couples $|e\rangle \rightarrow |g_2\rangle$, and the effective two-photon Rabi frequency is a coherent sum over all contributing sideband pairs.

Appendix B

Raman Gate Simulation Framework

This appendix provides the complete simulation framework for evaluating Raman gate fidelity under realistic noise conditions. We present the quantum model (Hamiltonian structure and intermediate state specifications), the method for incorporating stochastic phase noise from measured power spectral densities, and the Monte Carlo sampling procedure enabling performance prediction before hardware construction.

B.1 System Hamiltonian

The system Hamiltonian models the qubit states $|g_1\rangle \equiv |5S_{1/2}, F=1, m_F=0\rangle$ and $|g_2\rangle \equiv |5S_{1/2}, F=2, m_F=0\rangle$, separated by the ground-state hyperfine splitting $\omega_{HF1}/2\pi = 6.8347$ GHz (Rb-87 clock transition), and the intermediate excited-state manifold. Intermediate states $|e_i\rangle$ correspond to hyperfine levels in the excited P manifold, with $i=1, 2$ for D1 transitions (795 nm, 421 nm) and $i=1, 2, 3$ for D2 transitions (780 nm, 420 nm).

The Hamiltonian structure is $H(t) = H_0 + H_{\text{interaction}}(t)$, where H_0 contains the energy offsets of the intermediate states and $H_{\text{interaction}}(t)$ describes the time-dependent laser-atom interaction. For transitions to $j=1/2$ excited states (D1 lines), the intermediate manifold has two hyperfine levels corresponding to $F=1$ and $F=2$. The static Hamiltonian is:

$$H_0 = -\Delta_1|e_1\rangle\langle e_1| - (\Delta_1 - \omega_{HF2})|e_2\rangle\langle e_2| + \delta_{\text{tot}}|g_2\rangle\langle g_2| \quad (\text{B.1})$$

where Δ_1 is the single-photon detuning from $|e_1\rangle$, ω_{HF2} is the splitting between excited-state hyperfine levels, and δ_{tot} is the two-photon detuning from Raman resonance. The

five basis states are $\{|g_1\rangle, |e_1\rangle, |e_2\rangle, |d\rangle, |g_2\rangle\}$, where $|d\rangle$ represents dark states populated by spontaneous emission.

For transitions to $j = 3/2$ excited states (D2 lines), the intermediate manifold has three hyperfine levels ($F = 1, 2, 3$), giving six basis states $\{|g_1\rangle, |e_1\rangle, |e_2\rangle, |e_3\rangle, |d\rangle, |g_2\rangle\}$. The transition $|g_1\rangle \rightarrow |e_3\rangle$ is forbidden by angular momentum conservation ($\Delta F = 3 - 1 = 2 > 1$ for single-photon σ^+ polarization), resulting in zero dipole matrix element $V_{31} = 0$.

B.2 Monte Carlo Simulation of Phase Noise

We simulate the effect of phase noise on gate fidelity via a Monte Carlo method [37]. An ensemble of time-domain noise realizations is synthesized from the measured power spectral density (PSD), and for each realization, the system's master equation is solved to determine the final state fidelity. The simulation requires converting the measured single-sideband (SSB) noise density, $L_{\text{dBc}}(f)$, to a double-sided phase PSD: $S_\phi(f) = 2 \cdot 10^{L_{\text{dBc}}(f)/10}$, where the factor of two accounts for both positive and negative sidebands.

Each time-domain noise trace $\phi_j(t)$ is then synthesized via:

$$\phi_j(t_i) = \sum_k 2\sqrt{S_\phi(f_k) \cdot \Delta f} \cdot \cos(2\pi f_k t_i + \phi_{\text{rand}_{j,k}}) \quad (\text{B.2})$$

where Δf is the frequency resolution and a random spectral phase $\phi_{\text{rand}_{j,k}} \in U(0, 2\pi)$ for each frequency component k and run j ensures the statistical independence of each trace. The noise is incorporated into the time-dependent Hamiltonian $\tilde{H}_j(t)$, either directly through phase terms $e^{\pm i\phi_j(t)}$ or as an instantaneous frequency fluctuation. We solve the master equation $\dot{\rho}_j = -i[\tilde{H}_j(t), \rho_j] + \mathcal{L}\rho_j$ for each trace, where $\mathcal{L}$ contains noise-independent dissipation. The final noisy fidelity is the ensemble average $F_{\text{noisy}} = \langle F_j \rangle$, and the performance degradation is the infidelity $\Delta F = F_{\text{ideal}} - F_{\text{noisy}}$.

B.3 Noise Incorporation in Optical Phase-Locked Loop Systems

For OPLL architecture using two independent lasers (reference laser and locked laser), the electric fields at the atoms are $E_1(t) = E_1 e^{i(\omega_1 t + \phi_1(t))}$ and $E_2(t) = E_2 e^{i(\omega_2 t + \phi_2(t))}$, where E_1 and E_2 are field amplitudes, ω_1 and ω_2 are optical frequencies separated by $\omega_2 - \omega_1 \approx \omega_{HF1}$,

and $\phi_1(t)$ and $\phi_2(t)$ are time-dependent phase noise traces. The phase of the locked laser $\phi_s(t)$ is related to the reference phase $\phi_m(t)$, the RF phase $\phi_{RF}(t)$, and the residual loop error $\phi_e(t)$: $\phi_s(t) = \phi_m(t) + \phi_{RF}(t) + \phi_e(t)$. The total noise $\Delta\phi_{total}(t) = \phi_{RF}(t) + \phi_e(t)$ determines the relative phase error. The reference laser noise $\phi_m(t)$ cancels in the difference, except for its contribution to $\phi_e(t)$ through the loop transfer function.

The OPLL Hamiltonian is derived in a rotating frame defined by the reference laser frequency (ω_1) for the excited state and the beat note frequency (ω_{HF1}) for the ground state $|b\rangle$. The interaction terms acquire the phase noise difference, and the sampled total phase error $\Delta\phi_{total}(t)$ is applied as a relative phase shift $e^{\pm i\Delta\phi_{total}(t)}$ to the Rabi frequency term of the locked laser (Ω_2) in the off-diagonal matrix elements.

B.4 Noise Incorporation in Modulator-Based Systems

In the MZM approach, a single input laser beam is modulated by an RF synthesizer. The noise originates from the phase fluctuation $\phi(t)$ of the RF synthesizer, which simultaneously affects the phase and the instantaneous frequency of the sidebands. The modulated electric field at the atoms is:

$$E(t) = E_0 e^{i\omega_0 t} e^{i\phi_{\text{laser}}(t)} \sum_{n=-\infty}^{\infty} J_n(\beta) [e^{-in(\omega_m t + \phi_{RF}(t))} - e^{+in(\omega_m t + \phi_{RF}(t))}] \quad (\text{B.3})$$

where E_0 is the electric field amplitude (accounting for insertion loss), ω_0 is the optical carrier frequency, $\phi_{\text{laser}}(t)$ is the laser phase noise, $J_n(\beta)$ is the Bessel function of the first kind of order n , β is the modulation depth, ω_m is the RF modulation frequency ($\omega_m \approx \omega_{HF1}/2$), and $\phi_{RF}(t)$ is the RF phase noise.

The Raman transition is driven by pairs of sidebands where sideband $+n$ couples $|g_1\rangle \rightarrow |e\rangle$ while sideband $-n$ couples $|e\rangle \rightarrow |g_2\rangle$. The RF phase noise $\phi_{RF}(t)$ enters with opposite signs for the $+n$ and $-n$ sidebands, producing a relative phase error $\Delta\phi_{\text{rel}}(t) = 2n\phi_{RF}(t)$ for the n -th sideband pair. This relative phase error couples directly to the two-photon Raman transition without the filtering provided by an optical phase-lock loop.

Glossary

clock states Hyperfine qubit states with $m_F = 0$, exhibiting zero first-order magnetic field sensitivity. In ^{87}Rb , the $|F = 1, m_F = 0\rangle \rightarrow |F = 2, m_F = 0\rangle$ transition frequency is stable against magnetic field fluctuations to first order [7](#)

ECDL External-Cavity Diode Laser: a semiconductor laser with optical feedback from an external cavity containing a diffraction grating [39](#)

Helmholtz and anti-Helmholtz coils Coil configurations for magnetic field generation and control. Helmholtz coils (two coaxial coils separated by their radius, currents in same direction) produce static magnetic fields; anti-Helmholtz coils (currents in opposite directions) create magnetic field gradients with zero field at the center, used in magneto-optical traps [5](#)

magnetically sensitive transition An atomic transition between states with different magnetic quantum numbers, exhibiting linear Zeeman shift with applied magnetic field [6](#)

MZM Mach-Zehnder Modulator: an electro-optic device that converts a single laser into multiple optical sidebands by modulating the phase of light in two interferometer arms [29](#)

OPLL Optical Phase-Locked Loop: an active feedback system that stabilizes the frequency difference between two independent lasers [7](#), [29](#)

PDH Pound-Drever-Hall: a laser stabilization technique that phase-modulates light before cavity reflection, then demodulates the reflected signal to generate a dispersive error signal proportional to the frequency offset from cavity resonance [7](#)

PGC Polarization Gradient Cooling: a laser sub-Doppler cooling technique exploiting spatial variation of light polarization in counter-propagating beams with orthogonal circular polarizations [5](#)

qubit A quantum mechanical system whose wavefunction is constrained to a two-dimensional subspace of the total Hilbert space. Physical implementations include atomic hyper-fine states, photon polarizations, electron spins, and superconducting circuits [2](#), [4](#)

ULE cavity Ultra-Low Expansion cavity: an optical reference cavity made from glass with near-zero thermal expansion coefficient [7](#)