ABSTRACT

Title of dissertation:	ISOTOPE SHIFT SPECTROSCOPY OF ULTRACOLD STRONTIUM
	Neal Carden Pisenti Doctor of Philosophy, 2019
Dissertation directed by:	Dr. G.K. Campbell and Professor S.L. Rolston Joint Quantum Institute, National Institute of Standards and Technology, and Department of Physics, University of Maryland

We describe the design, construction, and performance of a laser system to probe the ultra-narrow ($\Gamma/2\pi \approx \text{mHz}$) clock transition ${}^{1}S_{0} \rightarrow {}^{3}P_{0}$ in strontium. We present the first reported spectroscopy of this transition in two of the bosonic isotopes, ${}^{84}\text{Sr}$ and ${}^{86}\text{Sr}$. Furthermore, we measure the complete set of isotope shifts between all four stable isotopes on the clock line and the narrow intercombination line ${}^{1}S_{0} \rightarrow {}^{3}P_{1}$, permitting a King plot analysis of the isotope shifts. Complications arising from the unambiguous determination of a line center in ${}^{87}\text{Sr} {}^{3}P_{1}$ prevent us from making claims about the King linearity, but we provide a statistical bootstrap analysis of the isotope shifts ${}^{88-84}\text{Sr}$ and ${}^{88-86}\text{Sr}$ to compute a field shift ratio $F_{698}/F_{689} = 0.9979$, with a 95% confidence interval [0.9952, 1.0008]. The intercept term $K_{698} - (F_{698}/F_{689}) K_{689}$ is similarly determined to be -2.0 GHz-amu, with a 95% confidence interval [-3.9, -0.3] GHz-amu. Finally, we describe the design of a next-generation apparatus that will enable improvements on the results described here, as well as other studies that involve coherent manipulation of strontium atoms on the clock line.

ISOTOPE SHIFT SPECTROSCOPY OF ULTRACOLD STRONTIUM

by

Neal Carden Pisenti

Dissertation submitted to the Faculty of the Graduate School of the University of Maryland, College Park in partial fulfillment of the requirements for the degree of Doctor of Philosophy 2019

Advisory Committee: Dr. Gretchen K. Campbell, Advisor Professor Stephen L. Rolston, Chair Professor Christopher Jarzynski, Dean's Representative Dr. James (Trey) Porto Professor James Williams © Copyright by Neal Carden Pisenti 2019

Dedication

For my parents, who have always pushed me to do more and pursue my interest in physics, and for Pascal, who will pull me into the future.

Acknowledgments

It is hard to look back over six years to the person I was entering graduate school; so much has changed since then, and I have grown much as a physicist. In hindsight, I didn't realize quite how lucky I was stumbling into the JQI, with its incredibly talented and diverse set of researchers. The opportunities it provided to work and interact with some of the world's foremost physicists is hard to match elsewhere, and for this I am tremendously grateful.

The people who have directly or indirectly helped me along the way is perhaps innumerable, and I am bound to forget some of them by mistake. A number of undergraduates made valuable contributions to the lab — Eliot Bohr and Peter Zhou in particular helped with everything from vacuum assembly to electronics design. In the early days, Joe Pechkis helped jumpstart the lab for six months as a post-doc while we were still in the building phase. Other graduate students passed through our lab as well, either as visiting students or as a waypoint to other experiments at the JQI and elsewhere. Avinash Kumar, Erin Marshall, Madison Anderson, and Swarnav Banik all helped for some time on the strontium apparatus, either by building electronics or designing optics, but remained close for moral support and the occasional tip when we were facing a challenging technical problem. Alexander Hesse was with us for an entire year from Heidelberg, designing and building a Bitter coil electromagnet with attendant high-current drive electronics. He also provided much-needed manpower in the construction phase of the new strontium lab in the PSC. Junyi Jhang spent a summer with us, and built a very robust 679 nm repump laser to replace the finicky one we had been using, and explored a novel coupledresonance locking scheme for the two repumps. Alessandro Restelli has been a mainstay of the strontium group, and taught me everything I know about electronics design. He always has a clever idea to work around a difficult technical issue, and brought some fantastic outside perspective to the physics monoculture. The new graduate students, Peter Elgee and Ananya Sitaram, and now too Nick Mennona, are bringing new life to the strontium experiment. They've worked tirelessly to bring up the new apparatus, and particularly in the case of Peter, extend the life of the old one. Much of the isotope shift data in this paper wouldn't have been possible without his help. Hiro Miyake joined the lab as a post-doc a little over a year ago, and in that time I've been really grateful for all of his help and expertise, and moral support when everything was breaking and it seemed there was no light at the end of the tunnel. This experiment is as much his now as it is mine. Finally, I don't quite have the words to describe how much I appreciate Dan Barker and Ben Reschovsky, the two previous graduate students in the strontium lab. They raised me, and taught me most of what I know about working in an AMO lab. Dan seemed to have a truly encyclopedic knowledge of the field, and was always happy to answer my dumb questions. Ben had a supernatural talent getting lasers to behave, and often helped me to refocus and just get to work when I had a tendency to overthink before doing. They both have some excellent experimental experimental chops, which I hope to someday be able to match.

Beyond the immediate strontium group, I'm grateful for the camaraderie and collaboration of other members of the JQI, and the "CPR" research group — Matt

Reed, Zach Smith, Aftaab Dewan, Dalia Ornelas, Sandy Craddock, Mary Lyon, Varun Vaidya, Creston Herold, Joe Tiamsuphat, Tsz-Chun Tsui, Sarthak Subhankar, Sean Li, Ana Valdes Curiel, Ryan Price, Dan Campbell, Dimitrios Trypogeorgos, Andika Putra, Prasoon Gupta, Elizabeth Goldschmidt — I could say something about each of them, but Zach Smith in particular was always there for moral and technical support. He has a knack for bridging the gap between physics and the real world. The other members of the Campbell research group made many things possible, both for their friendship and help when the experiment stopped working — Monica Gutierrez, Avinash Kumar, Swarnav Banik, Madison Anderson, and Hector Sosa. Also thanks to the two former sodium post-docs, Fred Jendrzejewski and Steve Eckel, who was always very helpful. Ken Wright in particular helped me through the ups and downs of graduate school, both at Town Hall and sharing the brutal commute to College Park.

I'm very grateful to Gretchen Campbell, who did much more than just pay our salaries. She was always there with advice and help working through a tough problem, but was also happy to let us all experiment with new ideas in the lab without breathing down our necks. I think we are all lucky to have her as an advisor. Steve Rolston and Trey Porto also provided some useful guidance at times as part of the extended research group, as did Luis Orozco who was selflessly there with support both scientific and beyond.

Finally, none of this would've been possible without Sarah's loving and patient support. I am grateful to her in ways that extend beyond graduate school, and am lucky to have her as my partner. This work was supported by the NSF through the PFC at the JQI.

Table of Contents

De	edicat	n i	i
Ac	know	dgements ii	ii
Lis	st of '	ables	X
Lis	st of]	gures	x
Lis	st of A	bbreviations xi	ii
1	Intro 1.1 1.2	luction Interpretation Strontium laser cooling primer Interpretation 1.1.1 Repumping Interpretation I.1.2 Narrow line magneto-optical traps Interpretation Outline of thesis Interpretation Interpretation	1 4 7 8 1
2	Exp 2.1 2.2 2.3 2.4 2.5	imental apparatus13Vacuum systems14Laser Systems172.2.1 Blue 461 nm lasers192.2.2 Red 689 nm lasers292.2.3 Repumping lasers282.2.4 Dipole trap laser282.2.4 Dipole trap laser33Magnetic field coils34Computer control34Custom electronics342.5.1 Laser diode current controller362.5.2 Optical beatnote servo40	337958124560
3	Desi 3.1 3.2	and construction of a narrow-linewidth clock laser 40 Motivation and target specifications 41 Pound-Drever-Hall: noise analysis and limitations 42 3.2.1 Index of refraction 50	5 7 9 0

		3.2.2 Mechanical length stability	51
		3.2.3 Thermal length stability	63
		3.2.4 Servo noise & electronics	68
	3.3	Final design	79
	3.4	Performance	82
4	Pub	lication: An ultra-low noise, high-voltage piezo driver	85
	4.1	Abstract	85
	4.2	Introduction	86
	4.3	Circuit Design	88
		4.3.1 Flyback regulator	89
		4.3.2 Low-noise stabilization and fast modulation	91
		4.3.3 Digital control and slow modulation	93
	4.4	Noise model & analysis	95
	4.5	Results	100
	4.6	Conclusion	103
	4.7	Future developments	104
5	Isote	ope shift spectroscopy	107
	5.1	Isotope shifts: theory and practice	109
		5.1.1 Mass shift \ldots	110
		5.1.2 Field shift \ldots	112
		5.1.3 King plots \ldots	114
		5.1.4 Hyperfine effects	115
	5.2	Precision measurement of strontium isotope shifts	120
		5.2.1 Experimental setup	121
		5.2.1.1 Spectroscopy of 689 nm intercombination lines	122
		5.2.1.2 Spectroscopy of $698 \mathrm{nm}$ clock line $\ldots \ldots \ldots \ldots$	125
		5.2.2 Results	129
		$5.2.2.1$ Systematics \ldots	136
		5.2.2.2 Detailed analysis of the clock transition lineshape	141
	5.3	King plot analysis	146
	5.4	Discussion and future directions	148
6	SrII	and future directions	150
	6.1	Vacuum system	150
		6.1.1 AOSense atomic beam source	153
		6.1.2 Science chamber	153
	6.2	Laser systems	156
	6.3	Magnetic field coils and high-NA imaging	160
	6.4	Computer Control	166
	6.5	Outlook	169
Bi	bliogi	raphy	171

List of Tables

1.1	Summary of strontium lines		
2.1	Dipole trap waists		
2.2	Frequency offsets for 461 nm beam paths 21		
3.1	Summary of ULE mechanical properties		
5.1	Summary of isotope shift measurement		
5.2	Absolute frequency measurements for 698 nm transition		
5.3	Absolute frequency measurements for 689 nm transition		
5.4	Hyperfine A and B coefficients		
5.5	Summary of 689 nm systematics		
5.6	5 Systematic frequency shifts on the clock and intercombination line		
	transitions		
5.7	Atomic and nuclear structure coefficients		
6.1	Zeeman slower simulation data from AOSense		
6.2	Bakeout time constants for various flange sizes		
6.3	Summary of optical access on SrII science chamber		

List of Figures

1.1	Strontium Level Diagram
1.2	Narrow-line cooling force
1.3	Operation of the fermionic red MOT 12
2.1	Schematic of the SrI vacuum system
2.2	SrI chamber-side optics (horizontal)
2.3	SrI chamber-side optics (vertical)
2.4	Spectroscopy for 461 nm 20
2.5	461 nm slave lasers
2.6	689 nm laser system
2.7	Repump laser systems
2.8	Dipole laser system
2.9	Current controller analog subcircuits
2.10	Current controller noise power-spectral-density
2.11	Beatnote lock sample trace
3.1	Clock Zeeman level diagram
3.2	ULE cavity and FEA geometry
3.3	Lab vibrational power spectral density
3.4	Cavity length fluctuation ΔL due to mirror tilt
3.5	Viton rod support cavity FEA
3.6	Final design for Fabry-Pérot cavity mount
3.7	FEA mesh and boundary model comparison
3.8	Summary of clock cavity FEA
3.9	Clock cavity zero-crossing in α_{CTE}
3.10	PDH closed-loop noise analysis
3.11	Schematic of PDH lock
3.12	PDH error signal
3.13	Clock laser system
3.14	Clock laser beatnote
3.15	Atomic spectroscopy of clock transition in Sr87 and Sr84 BEC $\ .$ 84
4.1	Schematic of the high-voltage piezo driver

4.2	Quench circuit for piezo driver
4.3	Noise model for piezo circuit
4.4	Piezo circuit noise model and measured noise power spectral density . 99
4.5	Time-domain traces of high-voltage output from piezo driver 101
4.6	Piezo driver temperature coefficient
4.7	Piezo driver modulation input transfer function
5.1	Finite-size nucleus potential energy
5.2	Magnetic field calibration
5.3	Detailed lineshape curves for 698 nm clock transition
5.4	Bootstrap analysis of 689 nm transition
5.5	Bootstrap analysis of 689 nm isotope shifts
5.6	Bootstrap analysis of 698 nm clock lines
5.7	Thermal line pulling for clock transition
5.8	King Plot (all isotopes)
5.9	Bootstrap distribution of King plot parameters
6.1	Zeeman slower simulation data from AOSense
6.2	CAD rendering of new SrII apparatus
6.3	SrII 461 nm laser system
6.4	SrII 689 nm laser system
6.5	Bitter coil electromagnet
6.6	SrII high-NA objective

List of Abbreviations

AMO	Atomic, Molecular, and Optical physics
AR	Anti-reflection
BEC	Bose-Einstein condensate
bMOT	Blue MOT
CF	ConFlat (vacuum flange)
CPR	"Campbell-Porto-Rolston"
CSS	Computer & Space Sciences Building
ECDL	External cavity diode laser
EMI	Electromagnetic interference
EOM	Electro-optic modulator
FDG	Fermi-degenerate gas
FEA	Finite-element analysis
FMC	FPGA Mezzanine Card
FORT	Far-off-resonance (dipole) trap
FP	Fabry-Perót cavity
FPGA	Field-programmable gate array
GUI	Graphical User Interface
IDE	Integrated developemt environment
JILA	JILA
JQI	Joint Quantum Institute
LAN	Local area network
LDO	Low-dropout voltage regulator
MOT	Magneto-Optical Trap
OSI	Open-systems interconnection model
PDH	Pound-Drever-Hall lock
PHY	Physical layer of the OSI model
PI	Proportional-integral feedback control
PID	Proportional-integral-derivative feedback control
PLL	Phase-locked loop
\mathbf{PM}	Polarization-maintaining fiber
PSC	Physical Sciences Complex
RBW	Resolution Bandwidth
RGA	Residual Gas Analyzer
rMOT	Red MOT
SI	International System of Units (or, Système International d'unités)
SoC	System on Chip
TCP/IP	Transmission Control Protocol/Internet Protocol
UI	User interface
UV	Ultra-violet

VCO Voltage-controlled oscillator

Chapter 1: Introduction

The modern history of atomic physics, in many ways, is intimately intertwined with the history of the laser. Since their conception and first demonstration in the mid-20th century, lasers have enabled us to coherently probe and manipulate the quantum-mechanical nature of atomic structure. With the development of laser cooling techniques, scientists routinely achieve samples of dilute atomic gasses at temperatures many orders of magnitude lower than would otherwise be possible. This ultra-cold regime has opened a window onto new physical phenomena that emerge at such energy scales — Bose-Einstein condensation (BEC) and its fermionic counterpart, synthetic crystal lattice structures, new phase transitions, topological matter — all become tools at the disposal of science.

Recent efforts in the cold-atom community seem to fall into into one of two categories: using atomic systems for precision metrology, or using atomic systems for quantum simulation or information processing. In a sense, the latter category is only possible because of the former — precision measurement enables precision control and the ability to engineer near-defect-free quantum systems for further study. Strontium, which is the element in focus for the entirety of this dissertation, is a particularly interesting candidate for both kinds of scientific investigation. It has a range of optical transitions with natural linewidths spanning over nine orders of magnitude; three of its isotopes are bosons with a wide range of scattering lengths, and one isotope is a fermion with an abnormally high nuclear spin I = 9/2. The narrowest line, connecting the ground ${}^{1}S_{0}$ state to the excited ${}^{3}P_{0}$ "clock" state (see Fig. 1.1), is a leading candidate for the anticipated redefinition of the second based on optical atomic clocks. Fermionic atoms in either the ground or clock state exhibit, to a very high degree, SU(N) symmetry in the nuclear spin subspace [1]. This makes strontium a good system to simulate Hamiltonians which carry this symmetry, where it is now possible to achieve any $N \leq 2I + 1 = 10$. Combining this SU(N) symmetry with the orbital degree of freedom provided by the ${}^{1}S_{0} - {}^{3}P_{0}$ clock transition allows one to simulate other Hamiltonians of interest, such as the Kondo lattice or Kugel-Khomskii models [2].

Beyond simulating these models directly, which often involves atoms pinned in a crystal lattice structure created by a deep three-dimensional (3D) optical lattice, there has been recent interest in constructing lattice models which utilize synthetic dimensions [3–5]. Again, strontium and other alkaline-earth(-like) atoms such as ytterbium feature prominently in many of these proposals because of the coherence offered by their narrow clock transitions, the presence of isotopes with high nuclear spin, and two easily accessible J = 0 orbital states with no magnetic moment. Another set of proposals use the external momentum states of atoms trapped in a uniform potential as a synthetic lattice dimension, realizing an exactly-solvable model with SU(N) symmetry using thermal atoms [6], or a method for estimating the eigenspectrum of the nuclear spin density matrix via Ramsey spectroscopy [7].

In this thesis, we will focus on experiments made possible by the narrow transitions in strontium. In particular, I will discuss the construction of a clock laser to address the ultra-narrow ${}^{1}S_{0} - {}^{3}P_{0}$ transition, and the first reported spectroscopy of this transition in 84 Sr and 86 Sr. We further measure all isotope shifts on both the clock and the narrow intercombination-line transition ${}^{1}S_{0} - {}^{3}P_{1}$; this permits a King plot analysis which is the subject of Chapter 5. However, before we get to that point, we will set the stage by discussing some important experimental details for the laser cooling and trapping of strontium gasses.



Figure 1.1: Low-lying electronic structure of strontium. Linewidths and decay rates are taken from Refs. [8–12]. Hot atoms are first cooled and captured into a magneto-optical trap (MOT) operating on the broad 461 nm transition ${}^{1}S_{0} - {}^{1}P_{1}$ (bMOT). Typical temperatures achieved after this first cooling stage are \sim mK. Atoms are then transferred to a frequency-modulated "broadband" MOT operating on the much narrower intercombination transition ${}^{1}S_{0} - {}^{3}P_{1}$ at 689 nm (rMOT). The intercombination MOT is recoil rather than Doppler limited, and we can achieve temperatures as low as $\lesssim 1 \,\mu$ K. This is cold enough to load directly into a far-off-resonance dipole trap (FORT) and proceed with evaporative cooling to quantum degeneracy. A slow leak from the broad 461 nm transition necessitates repumping during bMOT operation; several repumping schemes exist [13], however we follow the practice at JILA and repump through ${}^{3}S_{1}$ with 707 nm and 679 nm light. Atoms decaying to ${}^{3}P_{1}$ return to the ground state in ~ 21.5 µs, while atoms in ${}^{3}P_{0,2}$ are meta-stable with lifetimes $\gtrsim 100$ s. Repumping in 87 Sr is complicated by the hyperfine splitting of ${}^{3}P_{2}$ and ${}^{3}S_{1}$; full repumping on the 707 nm transition requires at least five tones spread over several GHz. We achieve this by frequency-modulating a fiberized electro-optic modulator (EOM), as described in the text. The "clock" transition at 698 nm is weakly dipole-allowed in ⁸⁷Sr via hyperfine state mixing, while in the bosonic isotopes it must be mixed in, e.g., with a static magnetic field. The broad array of transitions with varying linewidths makes strontium appealing for applications in metrology and quantum simulation.

1.1 Strontium laser cooling primer

There are many good pedagogical introductions to laser cooling and trapping, see for example Ref. [14].¹ Here, I will review the important aspects as they pertain to creating degenerate gasses of strontium. The basic idea behind all laser cooling is to engineer a situation whereby an atomic system preferentially absorbs photons that slow its motion. This is readily achieved via the Doppler effect — by detuning your laser red of a resonant transition, atoms moving towards the light see it blue shifted into resonance, and atoms moving away see it red shifted further from resonance. This produces a velocity dependence to the scattering rate (and thus a velocitydependent force, since each scattering event on average produces a momentum kick $d\mathbf{p}_{scat} = \hbar \mathbf{k}$), and the atomic motion is damped.

To understand this quantitatively, we start with the scattering rate Γ_{scat} for a two-level system. This can be written

$$\Gamma_{\rm scat} = \frac{\Gamma}{2} \left(\frac{s_0}{1 + s_0 + \left(2\Delta/\Gamma\right)^2} \right) \,, \tag{1.1}$$

where Γ is the transition linewidth (related to the excited state lifetime $\Gamma \equiv 1/\tau$ if we can neglect branching to other electronic states), Δ is the detuning from resonance, and s_0 is the on-resonance saturation parameter. For the remainder of this thesis, we adopt a convention (unless otherwise clear from context) that capitalized Greek letters for frequency units are implicitly angular frequency, and carry a factor of 2π .

We relate s_0 to the Rabi rate Ω , which characterizes the strength of the optical drive, via

$$s_0 = \frac{2|\Omega|^2}{\Gamma} \equiv \frac{I}{I_{\text{sat}}}.$$
(1.2)

¹Our discussion here follows the discussion in this reference.

	Transition	$\Gamma/2\pi$ (MHz)	$I_{\rm sat}~({\rm mWcm^{-2}})$	$\lambda \ (nm)$
${}^1S_0 - {}^1P_1$ ${}^1S_0 - {}^3P_1$ ${}^1S_0 - {}^3P_0$	(blue MOT) ("intercombination") ("clock")	30.2 7.48×10^{-3} 10^{-9}	$\begin{array}{c} 40.4 \\ 3.0 \times 10^{-3} \\ 10^{-10} \end{array}$	$460.9 \\ 689.4 \\ 698.4$

Table 1.1: Summary of strontium lines. Linewidths and saturation intensity calculations use recommended values from Ref. [15]. Because the details are isotope-dependent (see discussion in future chapters), we do not list exact values for the clock transition ${}^{1}S_{0} - {}^{3}P_{0}$; however, the order of magnitude is $\Gamma/2\pi \approx \text{mHz}$, $I_{\text{sat}} \approx 10^{-10} \text{ mW cm}^{-2}$.

In Eq. 1.2, we introduce the saturation intensity I_{sat} , which sets the relevant intensity scale for driving a particular transition. It can be conveniently written

$$I_{\rm sat} = \frac{\pi h c \Gamma}{3\lambda^3} \,, \tag{1.3}$$

with h representing Planck's constant, c the speed of light, and λ the wavelength of the transition. Table 1.1 lists these numbers in strontium for easy reference.

Given the scattering rate in Eq. 1.1 and the assumption that spontaneous emission is isotropic, one can write the time-averaged radiation force as $dp/dt = \hbar k \Gamma_{\text{scatt}}$, where each photon imparts a momentum $dp = \hbar k$, with $\hbar = h/2\pi$, and $k = 2\pi/\lambda$. To achieve confinement in addition to cooling, one must engineer a spatial dependence to the scattering rate, such that atoms experience a restoring force towards the center of the trap. This is done with a magnetic field gradient, whereby the detuning Δ in Eq. 1.1 picks up a spatial dependence from the Zeeman effect. A magneto-optical trap (MOT) is generated with six beams (to provide confinement in all three dimensions) in a quadrupole magnetic field, where the beam polarizations are set to address Zeeman lines that are shifted onto resonance as an atom leaves the trapping region, thus providing a restorative force to keep the atom in the MOT. In the appropriate limit, this can be modeled semi-classically as a damped harmonic oscillator, where the equilibrium temperature is set by the Doppler limit $T_D = \hbar \Gamma/2$. Most MOTs operate in this regime.

The exact scheme for laser cooling strontium is informed by its electronic structure, which is shown in Fig. 1.1 [16]. Initial laser cooling is done on the broad $(\Gamma/2\pi = 30.2 \text{ MHz})$ ¹S₀ – ¹P₁ transition at 461 nm. A Zeeman slower [14] decelerates atoms in a thermal beam, and they are captured into a 3D MOT dubbed the "blue MOT" (or "bMOT"). For the 461 nm transition, the Doppler temperature $T_D = 1.1 \,\mathrm{mK}$, which is still too hot to be loaded into a far off-resonant dipole trap (FORT) [17], which for our experiment has a typical trap depth $\approx 15 \,\mu\text{K}$. To further cool beyond the 461 nm Doppler limit, we operate a second stage MOT on the much narrower ($\Gamma/2\pi = 7.5 \text{ kHz}$) intercombination line ${}^{1}S_{0} - {}^{3}P_{1}$ at 689 nm (the "red MOT", or "rMOT"). The MOT dynamics in this regime are very different than the traditional broad-line MOT dynamics, since the recoil shift $\omega_R = \hbar k^2/2m$ is roughly equal to the transition linewidth. In this regime, the relevant energy scale and thus the dynamics are set by the relative size of the detuning Δ , the powerbroadened linewidth $\Gamma' = \Gamma \sqrt{1 + s_0}$, and the recoil shift ω_R [18]. Furthermore, the gravitational force $F_{\text{grav}} = mg$ can provide a non-negligible perturbation to traditional MOT dynamics, which must be accounted for. We discuss these details in more depth in Sec. 1.1.2.

Because the ground state of strontium is non-magnetic, the final stages of cooling are carried out in an optical dipole trap. To achieve good spatial mode matching, we load into a flattened, "pancake" trap formed by a gaussian elliptical beam with a roughly 1:10 aspect ratio, with the tight trap direction oriented along gravity. A second "crossing" dipole beam provides additional confinement along the weaker axes, enhancing the density and thus the elastic scattering rate during evaporative cooling. Our lab has brought ⁸⁴Sr, ⁸⁶Sr, and ⁸⁷Sr to quantum degeneracy [19, 20]; achieving Bose-Einstein Condensation (BEC) in ⁸⁸Sr requires sympathetically cooling with ⁸⁷Sr, because the *s*-wave scattering length in ⁸⁸Sr is near zero ($a_s = -2a_0$) and it does not re-thermalize during evaporative cooling. Additionally, only small BECs of ⁸⁸Sr are possible before the negative interaction energy causes it to collapse.

It is possible to load isotopic mixtures by shelving the first species in the metastable ${}^{3}P_{2}$ state, which is magnetically trapped by the quadrupole field of the MOT, while the second species is loaded. One can then repump to the ground state and proceed with laser cooling on the narrow intercombination-line transition. For a comprehensive review of laser cooling in strontium, see Ref. [13].

1.1.1 Repumping

As indicated in Fig. 1.1, there is a small branching ratio out of the otherwise cycling transition at 461 nm into the ${}^{3}P$ manifold. Atoms entering ${}^{3}P_{1}$ return quickly to the ground state ($\tau = 21.3 \,\mu$ s), but the ${}^{3}P_{2}$ state is metastable and must be depumped. In our lab, we do this by applying a laser at 707 nm, which optically pumps population via ${}^{3}S_{1}$ into both ${}^{3}P_{1}$ and ${}^{3}P_{0}$. Atoms that end up in ${}^{3}P_{0}$ must also be depumped, which is achieved by a second laser at 679 nm. For 87 Sr, this scheme is complicated by the hyperfine structure of both ${}^{3}P_{2}$ and ${}^{3}S_{1}$, which are split into five and three distinct F manifolds, respectively. These hyperfine manifolds are separated by several GHz, and full repumping requires addressing all five ${}^{3}P_{2}$, $F = \{13/2, 11/2, 9/2, 7/2, 5/2\}$ levels with 707 nm light.

The original recipe from JILA [21] involved driving a modulation onto the current and piezo actuator of the 707 nm repump laser to span the whole hyperfine spectrum. This technique has obvious drawbacks – one cannot simultaneously lock the laser and frequency broaden in this manner, and in practice, the mode-hop free tuning range is smaller than the desired modulation bandwidth. The result is a somewhat uncontrolled repumping scheme which requires continuous monitoring and careful tuneup by hand to achieve good results. More recently, groups using the 707 nm/679 nm repump scheme have migrated to using a fiber EOM to generate sideband tones to address each line separately [22, 23]. More details on our implementation can be found in Sec. 2.2.3.

1.1.2 Narrow line magneto-optical traps

The operation of a MOT on a narrow-line transition deviates from the damped harmonic motion discussed above for a broad-line MOT [14]. In the case of strontium, the 689 nm linewidth is $\Gamma/2\pi = 7.5$ kHz. This has a few important consequences. First, the narrow linewidth implies a very low saturation intensity, in our case $3 \mu W \text{ cm}^{-2}$. For typical MOT volumes, the beam waist is several mm, so even μW of optical power will saturate the transition. We thus typically consider the power-broadened linewidth $\Gamma' = \Gamma \sqrt{1 + s_0}$ as the relevant scale parameter. We capture atoms from the bMOT into the rMOT with about 700 μW of power, with rMOT beam waists ≈ 2.5 mm. Thus, $s_0 \approx 2400$ and $\Gamma' \approx 370$ kHz.

Given this saturation-broadened linewidth, the velocity class Δv that is resonant with our MOT beams is only $\approx 40 \text{ mm s}^{-1}$, while the Doppler-limited temperature in the bMOT implies a most-likely velocity $v = \sqrt{2k_BT/m} \approx 370 \text{ mm s}^{-1}$, with k_B Boltzmann's constant. Furthermore, the spatial extent of the MOT cooling force is set by the magnetic field gradient, such that at some distance w_z the atom is Zeeman-shifted by one linewidth. For ${}^{3}P_{1}$, the Zeeman shift is 2.1 MHz G⁻¹; typical rMOTs operate at a field gradient of $\approx 1 \text{ G cm}^{-1}$ to 5 G cm^{-1} , and we find the best performance at 1.5 G cm^{-1} . This implies a resonance "shell" that is at most 1.2 mmwide, which is also smaller than the typical size of the bMOT.

To achieve efficient transfer from the bMOT to the rMOT under these conditions, we modulate the frequency of the rMOT beams. This has the effect of generating a comb structure, where regardless of the velocity or position of an atom in the bMOT, there is some cooling light that is on-resonance with that atom. Once transfer to the rMOT is complete, we increase the magnetic field gradient slightly to compress the rMOT volume, ramp the detuning of the rMOT light closer to res-



Figure 1.2: Illustration of (a) the MOT force and (b) corresponding potential energy (including the effect of gravity) for several different detunings in the narrow-line rMOT, with $s_0 = 36$, v = 0, and a magnetic quadrupole field gradient of $1.5 \,\mathrm{G \, cm^{-1}}$. Unlike a broad-line MOT, which has a dispersive force acting over the entire trapping volume, the narrow-line MOT only affects a force over narrow shells where the atom is Zeeman shifted into resonance. Additionally, gravity adds a non-negligible perturbation to the dynamics, causing atoms to bunch near the bottom of the MOT volume (c). Figures reproduced from Ref. [20].

onance, and then switch from the "broadband" modulated rMOT stage to a singlefrequency MOT. Finally, we decrease the rMOT beam powers to achieve a lower temperature. Empirically, these parameters are set by optimizing the peak density of the cloud (which is typically equivalent to maximizing the total atom number, although we get a better signal-to-noise looking at the fitted peak density).²

Because the gravitational force F_{grav} begins to compete with the cooling force $F_{\text{cool}} = \hbar k \Gamma_{\text{scatt}}$, the location of the rMOT is set by the force-balance between these two effects. This is illustrated in Fig. 1.2.³ Final loading into the dipole trap is achieved by adjusting the magnetic field shim coils to move the rMOT equilibrium position on top of the dipole beam.

The ⁸⁷Sr fermionic rMOT operates on the ${}^{3}P_{1}$ ($F = 9/2 \rightarrow F' = 11/2$) transition, and is complicated by the hyperfine structure in both the ground and excited states. The ground state magnetic field sensitivity is set only by the nuclear magnetic moment, while the excited state field sensitivity is set primarily by the electronic spin magnetic moment — this leads to a large differential Zeeman shift between the ground and excited states, shown schematically in Fig. 1.3. As a result, it is possible to induce a σ^+ (or conversely σ^-) transition anywhere in the MOT volume. This is a terrible situation to find oneself, because we are no longer guaranteed to only absorb photons that give a restorative kick back towards the center of the MOT. To make matters worse, because the linewidth is so small it is possible to exist in a region where the scattering rate is much larger for the "wrong" polarization than for the "right" polarization, depending on the particular m_F level the atom might exist in.

The way around this problem is to add a second laser operating on the

²Ref. [20] gives more details of our exact experimental parameters.

 $^{^{3}}$ As an aside, we would have benefitted greatly from launching our vertical rMOT beams from beneath the chamber, because the end of the rMOT cooling cycle is primarily determined by this beam. Losses going through the optics before retro-reflection decreases the available power

 $F = 9/2 \rightarrow F' = 9/2$ transition, which quickly randomizes the ground state m_F populations. This way, for a particular location in the rMOT, you are guaranteed to end up in a state which scatters favorably towards the MOT center. This process is aided by the Clebsch-Gordon weights for particular σ^+/σ^- transitions. With the addition of this second "stirring" laser, the fermionic rMOT dynamics are stabilized and we get a time-averaged restoring force as is required for the trap to operate. But the performance is degraded somewhat from the performance of the bosonic rMOT, and optimization is more complicated because of the second laser.

1.2 Outline of thesis

The remainder of this thesis is as follows. In the next chapter, I describe the experimental apparatus used to carry out the experiments described here. Chapter 3 describes the construction of a narrow linewidth laser to address the 698 nm clock transition in strontium, and initial tests of its performance. Chapter 4 contains an edited reprint of a paper published in Review of Scientific Instruments, detailing an ultra-low noise, high bandwidth piezoelectric driver. Chapter 5 contains the first reported measurements of the clock transition in ⁸⁶Sr and ⁸⁴Sr, and the results of a detailed isotope shift measurement between all isotopes on the 689 nm intercombination line and 698 nm clock line; these results seed a King plot analysis of strontium to extract details of the nuclear structure. We also discuss improvements to our experiment which might enable enough precision to constrain theories of physics beyond the standard model, which predict deviations from King linearity. Finally, in Chapter 6, we describe the design and construction of a next-generation strontium apparatus, and future directions for the physics we intend to study.



Figure 1.3: Schematic illustration of the behavior of the fermionic rMOT. In (a), we show all possible σ^+ (purple) and σ^- (green) transitions, where the m_F label corresponds to the excited state. Subplots (c) and (d) show the same diagram, but the lines are weighted by the relevant Clebsch-Gordon factors from (b). There is a clear preference from the Clebsch-gordon weights to scatter from the correct MOT beam, but regardless of what ground-state m_F level the atom is in, there is a region where it will scatter a photon from the wrong MOT beam and be pushed out of the trap. To get around this apparent issue, a second laser operating on the $F = 9/2 \rightarrow F = 9/2$ transition randomizes the m_F populations in the ground state. Thus, an atom will be pumped to an m_F level that scatters preferentially from the correct MOT beam, stabilizing the MOT dynamics. Figure reproduced from Ref. [20].

Chapter 2: Experimental apparatus

Most neutral atom experiments require a few key experimental systems to function. First, because the cold atomic gasses we create are fragile, they need to be well isolated from the chaotic laboratory environment. This is achieved by performing experiments in an ultra-high vacuum environment — the exact pressures required depend on the lifetime required of your cold atomic sample, but systems regularly achieve pressures in the low 10^{-11} torr. Second, our only mode for interacting with the atoms comes from electromagnetic fields. Again, the details vary by species and specific experimental requirements, but laser systems (for resonant and off-resonant interactions), a set of magnetic field coils, and possibly electric field control (which can be important for the highly susceptible Rydberg states) are universal. And finally, because the natural timescale for atomic processes is typically in the us regime, real-time computer control systems are necessary to facilitate experimentation, and other classical control systems to stabilize laser frequencies and powers need to be designed and implemented. In the next few sections, I address our solution to this generally posed set of experimental requirements, with references to additional details in previous students' theses.

2.1 Vacuum systems

A detailed description of the SrI apparatus can be found in previous theses [19,20]. Here, I briefly summarize the key design features and reproduce some figures to highlight changes that have been made since Refs. [19] and [20].



Figure 2.1: Schematic of the SrI vacuum system, reproduced from [20]. An effusive oven produces a hot, collimated atomic beam of strontium, which passes through a split stage of transverse cooling before entering a variable-pitch Zeeman slower. The oven, transverse cooling, and main chamber are separated by differential pumping to reduce the impact of the hot oven on vacuum pressures in the main science chamber.

The vacuum system consists of an oven, a transverse cooling stage, a variablepitch "spin-flip" Zeeman slower,¹ and a science chamber (Kimball Physics 8 in "spherical square", MCF800-SphSq-G2E4C4A16). The main science chamber is separated from the oven by two stages of differential pumping, each with a 75 L/s ion pump (Gamma Vacuum 75S), and a pump body adjoins the main chamber opposite the Zeeman slower. This pump body contains a third 75 L/s ion pump and a titanium sublimation pump, along with an ion gauge and sensor package for a residual gas analyzer (RGA). In routine operation, we find the titanium sublimation pump does not help our pressure,² which holds a steady-state value of $\approx 5 \times 10^{-11}$ torr, near the sensitivity limit of the ion gauge (Agilent UHV-24P).

The oven contains two sections: a capped half-nipple "reservoir", which holds

¹This is a misnomer, since strontium has a J = 0 ground state and thus no spin to flip. In this context, we mean the magnitude of the magnetic field goes through zero partway along the axis of the slower.

²However, during the initial pumpdown, the titanium sublimation pump decreased the vacuum pressure by roughly a factor of two.

 $\approx 10 \text{ g}$ strontium, and a nozzle containing 496 hypodermic needles clamped together in a V-shaped groove. Our design is based on the one presented in Ref. [24], and further details can be found in Ref. [20]. The reservoir is heated to 600 °C by an out-of-vacuum clamshell heater (Mellen Company Inc. 11C-1302-TC), and the nozzle is kept 75 °C hotter than the reservoir to prevent clogging. The entire assembly is wrapped in several layers of hardened ceramic insulation (Cotronics #372-2), and the temperature is monitored at several points by high-temperature thermocouples. The temperatures at the reservoir and nozzle are each controlled by a proportionalintegral-derivative (PID) temperature controller from Omega (CNi16D44-EIT), which is programmed via ethernet and actuates a solid state relay to apply or disconnect power from the clamshell heaters.³ The advantage of this design is its mechanical simplicity, however several oven designs exist in the literature that heat invacuum [25], or utilize a 2D MOT and are much more compact [26]. In Chapter 6, we discuss the next-generation strontium apparatus, which uses a compact commercial beam source from AOSense.

The hot atomic beam is blocked by a pneumatic linear actuator (MDC 662006) attached to a sheet-metal flag.⁴ We open the beam shutter only during the bMOT load to reduce strontium buildup on the Zeeman slower viewport. Atoms are decelerated by the Zeeman slower to $\leq 60 \text{ m/s}$ and captured into the bMOT, from which point additional stages of laser cooling bring the dilute gas to ~ nK temperatures or to quantum degeneracy, if desired. A schematic of the vacuum assembly can be

³These temperature controllers seem ubiquitous at the JQI, however we have encountered a rather frequent but intermittent bug in one which causes a programmed temperature ramp to jump immediately to the end value. Power cycling the control box fixes the issue, but it remains disconcerting.

⁴After years of operation, the rubber o-ring which seals the upper and lower pneumatic chambers disintegrated. However, after a few minutes on the phone with MDC, we discovered this is trivial to replace — the top cap unscrews from the pneumatic actuator, and the piston can be unscrewed to access the o-ring. A few months after this failure, the pneumatic solenoid valve also failed, and MDC was only able to replace it with a DC-operated equivalent. This solenoid is driven at 24 V by a monolithic NPN darlington transistor array (Texas Instruments ULN2803A) with an internal free-wheeling diode.



Figure 2.2: Horizontal schematic of the SrI chamber-side optics system, reproduced from [20] with modifications reflecting changes since its publication. The clock/689 nm spectroscopy laser is either collimated (waist $w \approx \{1.25 \text{ mm} \times 1.71 \text{ mm}\}$), or focused ($w \approx \{332 \text{ µm} \times 461 \text{ µm}\}$) through the chamber, by removing lens L1 (f = 250 mm) which forms a relay telescope with the final 1" lens (f = 250 mm) before the chamber.

seen in Fig. 2.1.

The chamber-side optics can be seen in Figures 2.2 and 2.3. The rMOT and bMOT beams are combined on dichroic mirrors and enter through the 2.75" CF viewports. Their polarization is set with a two-color $\lambda/4$ waveplate (689 nm and 461 nm), and are retro-reflected on the opposite side of the chamber. Four 1064 nm dipole trapping beams, the "main" pancake, "cross" pancake, vertical beam, and horizontal lattice, each enter along various axes and are combined as indicated with dichroic mirrors or polarizers. The "loose" vertical dipole trap in Fig. 2.3 is not used for the results discussed here, however previous results on the production of ⁸⁶Sr

Boom	Waist (µm)			
Deam	w_x	w_y	w_z	
Main pancake	310		31	
Cross pancake		270	16	
Vertical cross/Lattice	105	105		
Horizontal lattice		100	100	

Table 2.1: Summary of dipole trapping beams and their waists.

BEC required this beam to avoid large three-body collisional losses during evaporation [19,20]. Each dipole beam passes through a Glan polarizer, and is intensity locked through a near-normal-incidence pickoff or backside polished mirror.⁵ The vertical lattice beam (see Fig. 2.3) passes through a liquid-crystal variable waveplate (Meadowlark) on the underside of the chamber. By rotating the polarization, we can dynamically ramp from dipole to lattice configuration by changing the amount of power that is retro-reflected after passing through a polarizing beamsplitter. The waists for each dipole beam can be found in Table 2.1.

2.2 Laser Systems

After the vacuum chamber, the next key ingredient to a cold-atom experiment is the laser system used to trap and manipulate the atoms. Strontium, given its rich electronic structure, is no exception and at times we have needed as many as twelve separate laser systems to run an experimental cycle. In the next sections, I describe some of these laser systems and provide some forward-looking advice for improving the lock stability as the experiment transitions to the next generation (see Chapter 6).

 $^{^{5}}$ We find the Glan polarizer crucial for achieving good intensity lock performance. Due to the small difference in reflectivity of s– versus p–polarized light from the wedged pickoff, polarization rotation can appear to the lock as intensity modulation and thus be written back onto the intensity of the light.



Figure 2.3: Vertical schematic of the SrI chamber-side optics system, reproduced from [20] with modifications reflecting changes since its publication. The "loose" vertical dipole crossing beam is not used for the results discussed here, but is necessary for evaporation of ⁸⁶Sr to avoid large three-body loss.

2.2.1 Blue 461 nm lasers

As we describe in Sec. 1.1, the broad 461 nm transition is used for the first stages of laser cooling, and for imaging. While the exact setup has changed periodically over the history of this apparatus, we currently employ three separate lasers for generating light at 461 nm — a frequency-doubled Toptica system (the TA-SHG Pro), a Toptica ECDL diode laser (the DL Pro), and an injection locked diode from Nichia (NDB4261).

The 461 nm DL Pro is locked to a hollow-cathode lamp (Hamamatsu L2783-38NE-SR⁶) via Doppler-free polarization rotation spectroscopy [28], see Fig. 2.4. Because ⁸⁸Sr is vastly more abundant than the other isotopes, and the 461 nm isotope shifts are small compared to the Doppler-broadened lineshape, there is only one easily-resolved absorption feature for all four isotopes corresponding approximately to the ⁸⁸Sr resonance position. To provide additional flexibility, we reference the Toptica SHG-Pro to this spectroscopy laser via beatnote. This is done with a modified version of the circuit described below in Sec. 2.5.2. External frequency control of the SHG-Pro is done via a slow piezo actuator, making it impossible to phase lock these two lasers. However, we low-pass filter the output of the beatnote circuit at ≈ 1 Hz, such that we slave the Toptica SHG-Pro frequency to the reference laser at long timescales and rely at shorter times on the good frequency stability of the SHG-Pro. Because the linewidth of this transition is 32 MHz, the lock is rather forgiving. We notice performance in the MOT to depend marginally at the level of 5 MHz offsets from the nominal operating frequency.

⁶When making this purchase, we were informed that Hamamatsu will no longer be manufacturing these hollow-cathode lamps. We are hedging against this future on a few fronts: Peter Zhou, an undergrad who has worked with us for the past few years, is building a spectroscopy setup inspired by the dispenser design in [27]. We have also explored transfer cavity locks, whereby we would reference both the 461 nm lasers and 679 nm/707 nm repump lasers to the stable 689 nm system. Finally, the AOSense source described in Chapter 6 has optical access to do spectroscopy transverse to the atomic beam.


Figure 2.4: Schematic of the 461 nm DL Pro spectroscopy laser, reproduced from [20]. Light for the polarization rotation spectroscopy is intensity locked after passing a cleanup cube to convert polarization drift into intensity drift, which can be servoed. We find it crucial to have very clean polarizations for our spectroscopy lock, and place a Glan polarizer after the polarizing beam cube that splits light for the pump and probe arms. Additionally, a μ -metal shield is placed around the hollow cathode lamp to reduce the coupling of magnetic field noise onto the lockpoint.

Beampath	Frequency (MHz)	Notes
bMOT	160	Double-passed from 80 MHz
Transverse Cooling	182	
Zeeman Slower	411	
Imaging	211	

Table 2.2: Summary of the frequency offsets used for each of the main 461 nm beam paths.

This laser locking setup suffers from a few drawbacks. First, while the beatnote provides a great deal of flexibility, bootstrapping a low-bandwidth frequency lock from a circuit nominally designed for high-bandwidth phase locking has two performance drawbacks. First, the 1 Hz low-pass filter between the lock circuit and the laser limits the frequency slew rate. This filter is necessary because of the very high voltage-to-frequency gain at the laser input, but sweeping the laser far enough to lock to a different isotope requires a few seconds. Second, also due to the high voltage-to-frequency gain, we must further attenuate the low-passed feedback signal. This limits the lock capture range, such that we are unable to fully sweep from ⁸⁸Sr to ⁸⁴Sr without manually recentering the lock. Both of these issues could be resolved by implementing a feed-forward, or by moving to a fully digital solution.⁷

Most of the blue light for the experiment is generated by the Toptica SHG-Pro laser system, shown in Fig. 2.5. Acousto-optic modulators (AOMs) provide the necessary frequency offsets from the main laser; see Table 2.2 for the detunings we use. We typically operate with $\approx 20 \text{ mW}$ to 22 mW in the transverse cooling and $\approx 8 \text{ mW}$ to 10 mW per beam in the MOT (derived from a three-way PM fiber splitter from Evanescent Optics), however the power out of this laser has been slowly degrading since its installation. There are three imaging systems, two in

⁷For this kind of lock, a digital solution might provide a lot of benefits. We have discussed building a frequency counter with an FPGA, perhaps using ideas from http://hamsterworks.co. nz/mediawiki/index.php/High_Speed_Frequency_Counter, combined with the PYNQ platform (http://www.pynq.io/) for easy integration into computer control. Briefly, a gray code counter is synthesized in gateware, which is driven by the (digitized) beatnote. The output of this counter sampled at the main FPGA clock frequency, and some arithmetic performed to extract the counted frequency.

the horizontal direction (the low-magnification "Flea", and the high-magnification "Grasshopper") and one in the vertical direction (the high-resolution "PIXIS"). Light for these imaging paths is derived from two AOMs; the Grasshopper exists standalone, but the Flea and PIXIS cameras require a flipper mirror to switch between two fiber launches coupled from the same AOM. All beam paths also have shutters (Stanford Research Systems SR475) for full optical extinction.

Light for the Zeeman slower is generated by injection locking a mediumpower laser diode (Nichia NDB4261). Approximately 10 mW of 461 nm light is fiber-coupled from the main Toptica SHG-Pro system (see Fig. 2.5) and passed through the output rejection port of an optical isolator (Thorlabs IO-5-461-HP).⁸ This injection-locked diode delivers ≈ 50 mW to 60 mW of light to the experiment,⁹ and is located on the vacuum system table to avoid losses due to fiber coupling.

Because the injection current window is small, we employ a crude injection current locking scheme inspired by the system presented in Ref. [29]. Briefly, some of the light is picked off on a polarizing beam cube and sent to a scanning Fabry-Pérot cavity. When the laser is stably injection locked and single mode, the cavity peaks have a characteristic height. As the laser drifts towards unlocking, the peak height decreases and one begins to see secondary frequency modes creep in from below. In Ref. [29], they use a microcontroller to monitor the peak height, and stabilize the current to keep the laser locked.

Our system is essentially the same, however our laser does not exhibit the same gradual decrease in peak height near the high-current edge of the injection locking window. Rather, we find this lasers response to be more symmetric in current tuning,

⁸We had previously used ConOptics isolators for this purpose (711C-3), and still do for our red injection-locked lasers. However, in two separate instances at 461 nm, we had an issue with the Faraday crystal spontaneously cracking. This is quickly discovered, as the output mode turns to speckle, and you can visually see discoloration and hairline fractures extending into the crystal bulk. ConOptics was nice enough to fix our isolators in both instances, but we suspect some issue with this model at shorter wavelengths.

⁹When this laser was first set up, it delivered closer to 70 mW of power.



Figure 2.5: Schematic for the core 461 nm laser systems, reproduced from [20] with modifications.

and quicker to unlock. An Arduino Due, mounted to a Eurocard carrier card and housed in a custom Eurocard rack,¹⁰ handles all aspects of the injection lock. The Arduino Due is based on the SAM3X chipset from Atmel, and has two integrated 12-bit digital-to-analog converters (DACs), as well as several 12-bit analog-to-digital converters (ADCs). Thus, with this single board, we can generate a ramp voltage for the Fabry-Pérot piezo (which passes through an external op-amp to extend the voltage range to 0 V to 15 V) and a feedback voltage for the diode current controller (Thorlabs LD202C), and also read the photodiode voltage to calculate the cavity peak height.¹¹ The user can set the target peak height, the "unlocked" peak height, the number of scans to average before implementing the feedback algorithm, and the feedback stepsize. When the lock is engaged, the Arduino either does nothing if the peak height is above the target threshold, or if this is not the case, it will step the current by some amount. On the next scan, if the peak height decreased, the Arduino will reverse the sign of its current step. This continues until the peak height is above the target threshold. While crude, we find this locking algorithm to be sufficient for most purposes. There is an auto-relock implemented as well (whereby a secondary "unlocked" threshold targets larger current excursions), however we find this to not work particularly well and have not invested the time to improving the performance. Code and Eagle CAD files can be found on GitHub.¹²

Future improvements might include improving the locking algorithm, and integrating the circuit into the planned Eurocard "smart backplane", such that the Arduino talks directly to our custom current controller (see Sec. 2.5.1) rather than via analog voltage to a Thorlabs current controller. Because the injected current

 $^{^{10}}$ This rack currently just supplies $\pm 15\,\mathrm{V}$ to the circuit, but eventually it will enable integration with custom current controller electronics residing in separate modules, see Sec. 2.5

¹¹The photodiode voltage is sampled quickly via the SAM3X Direct Memory Access (DMA) engine, activated on a threshold condition. Thus, during a given scan of the cavity, the photodiode voltages are written directly to a buffer in memory without directly involving cycles from the CPU.

¹²https://github.com/JQIamo/injection-ffwd

window for this diode is so small, the hysteresis in the Thorlabs current controller frontpanel dial can be very frustrating when adjusting the diode to lock. It would also be beneficial to have a mechanism for reporting the "unlocked" or "nearing unlocked" state to the computer control system, such that a cycle could be automatically suspended while corrective action is taken. We currently work around this laser unlocking by closely monitoring the fluorescence of the bMOT, which drops considerably when the Zeeman slower laser is not operating properly.¹³

2.2.2 Red 689 nm lasers

The red 689 nm laser system is centered around a DL-Pro which is locked via the Pound-Drever-Hall technique [30] to a passive Fabry-Pérot cavity made of ULE glass. This laser is close to resonance (up to an AOM shift) with the ⁸⁷Sr F = 9/2 - 11/2 transition, with the offset between the ULE cavity resonance and the atomic frequency provided by double-passed 1 GHz, high-frequency AOM (Brimrose).¹⁴ Light is divided between an injection lock (a few mW) and a network of fiber beamsplitters which provide reference light to several beatnote-locked slave lasers. The primary set of electronics for the PDH lock are provided by Toptica: a FALC analog PID, and a PDD-110 module for demodulating the PDH error signal. We also use a slow integrator from SRS to keep the fast current feedback signal from the FALC centered, by feeding back to the laser piezo. We find this to work better than the built-in "infinite integrator" in the FALC, which seemed to invariably cause a big mode hop whenever the laser happened to come out of lock. The linewidth of this

¹³Other groups have had similar issues keeping this non-AR coated 461 nm diode happily injection locked, and have switched to using the more expensive AR coated version instead [22]. In this configuration, the effective power gain of the diode is much lower, but because the lasing threshold current of the AR-coated diode is so high, the mode is set exclusively by the frequency of the injected light.

¹⁴This AOM is very temperamental and polarization sensitive, which makes double-passing difficult. As discussed in Chapter 6.2, a much more versatile option would be to lock one laser directly to the cavity (or through a phase-modulated sideband generated by a wide-bandwidth EOM, as is done in Chapter 3.3) and beatnote-reference slave lasers as needed.

laser was measured via a heterodyne beatnote with a temporary 689 nm diode locked to the clock laser cavity described in Chapter 3, and found to be ≤ 200 Hz which is more than sufficient for addressing the 7.5 kHz intercombination line transition.

In addition to the injection-locked diode mentioned above, which generates the rMOT light for ⁸⁷Sr, we primarily use light from the "long Steck" laser for rMOT operation (see Fig. 2.6). Specifically, this laser provides the trapping light for all bosonic isotopes, and the stirring light for the fermionic rMOT (see discussion in Sec. 1.1.2). It is a home-built ECDL based on the design from [31]. The diode is a cheap, low-power, single mode diode (HL6738MG) that we have anti-reflection (AR) coated in the University of Maryland FabLab. The coating recipe is the one suggested in Ref. [32]. We remove the hermetic cannister and deposit a layer of Al_2O_3 to bring the facet thickness to $\approx \lambda/2$, followed by a final layer of HfO₂ to bring the thickness to $\lambda/4$ (or an odd multiple thereof). The deposition is monitored in situ by actively scanning the current of the laser diode around threshold, and monitoring the response on the integrated photodiode. As the front facet becomes a better reflector, the threshold current decreases. When we observe the threshold current start to increase again, we halt the deposition of Al_2O_3 and begin depositing HfO_2 until the threshold current is pushed as high as possible. Deposition is done in an electron "E-beam" evaporator, and the diode is held (along with up to four other diodes) in an aluminum mounting clamp affixed to the top of the deposition chamber. Electrical signals for monitoring the photodiode and scanning the laser current are provided via vacuum feedthrough. We find this technique to be rather robust, and has generated several successful AR coatings for diodes at 689 nm and 679 nm.



Figure 2.6: Intercombination line laser systems, reproduced from [20] with modifications. A Toptica DL Pro (the "red master") is locked to a ULE cavity, and is used to reference several slave lasers highlighted in green. The beatnote network (not shown) is a network of butt-coupled fiber splitters which takes reference light from the red master and distributes to the several beatnote locks. It also splits light to the fiber switch (shown in Fig. 2.7) for monitoring on the wavemeter.

2.2.3 Repumping lasers

As discussed in Section 1.1.1, efficient laser cooling of strontium requires repumping out of the metastable ${}^{3}P_{2}$ level populated during operation of the bMOT. Several repump schemes exist for strontium [13,33–35], which use some combination of ${}^{3}S$, ${}^{3}P$, or ${}^{3}D$ states. We follow the scheme used at JILA and elsewhere, pumping out of ${}^{3}P_{2}$ via the ${}^{3}S_{1}$ level at 707 nm. Unfortunately, ${}^{3}S_{1}$ can decay to any of the 5s5p ${}^{3}P$ states, and so a second repump at 679 nm is required to repump out of the metastable ${}^{3}P_{0}$ level as well. Other repump schemes address states which can only decay to ${}^{3}P_{2}$ or ${}^{3}P_{1}$ via selection rules, and so only require one laser. However, the additional complexity of a second repump is somewhat mitigated by the availability of cheap, single-mode diodes at the 679 nm and 707 nm wavelengths.

Both repump lasers are homebuilt Littman-Metcalf ECDL systems. The 707 nm diode was wavelength selected and AR-coated from Toptica, while the 679 nm diode was AR-coated by our lab as described above. Light from the 707 nm ECDL is coupled to a secondary injection-locked system to increase the available power, and passed through a fiber EOM for reasons discussed below. Light from the 679 nm laser passes through a 200 MHz AOM and is brought to the experiment via SM fiber. The reason for the AOM on this repump arm is to eventually gate the repumping of atoms from ${}^{3}P_{0}$ for imaging both ground– and excited-state atoms in experiments with the 698 nm clock transition. Both lasers are locked to a wavemeter (High Finesse, WS7) via custom LabView codes.

Repumping of the fermionic isotope is complicated by the hyperfine interaction, which splits ${}^{3}P_{2}$ into five manifolds and ${}^{3}S_{1}$ into three manifolds, as discussed in Sec. 1.1.1. We initially used a technique inspired by an old JILA method, whereby the 707 nm laser was made multimode by driving a voltage ramp on both the piezo and current modulation ports. By monitoring the bMOT fluorescence while adjust-



Figure 2.7: Repump laser systems, reproduced from [20] with modifications. The 707 nm repump master seeds a secondary injection locked diode, which produces rempump light to the apparatus. For experiments using ⁸⁷Sr, a wide-bandwidth fiber phase modulator adds sidebands to bridge the few-GHz hyperfine splitting in ³P₂. Both the 707 nm and 679 nm lasers are frequency stabilized to a wavemeter.

ing the center of both the laser current and piezo, decent repumping efficiency could be achieved. However, this technique was not very robust or repeatable, and as the laser drifted, so did the repumping efficiency. To circumvent this issue, we now use a wideband fiber EOM (Jenoptik, PM705) to address the repumping transitions in a controlled manner.

Several groups use some variation of the EOM repump scheme for the fermion [22, 23], but we follow the method outlined in Ref. [22] after discussions with the author at DAMOP 2015. A high-frequency DDS evaluation board from Analog Devices (using the AD9914 chip, which is a 12-bit, 3.5 GSPS DDS) is controlled by an Arduino-based microcontroller via SPI. Using the built-in frequency profiles in the AD9914, the microcontroller pre-programs the DDS with a series of frequencies and amplitudes (which can optionally be changed before each experimental cycle). It then toggles the profile-select pins to change the output frequency. We find decent repumping efficiency locking the laser frequency at 423.913 THz, and cycling the EOM frequencies between 655 MHz, 1240 MHz, 1510 MHz, and 1720 MHz with a dwell time of 20 ms per frequency. This is slightly different than the values reported in Ref. [22], which were similarly determined via empirical methods. The efficacy of this technique is measured by steady-state bMOT fluorescence, and we find we are limited primarily by the amount of optical power that can safely be passed through the fiber EOM.¹⁵

¹⁵We initially damaged one of these EOMs, which then had to be sent back for repair. It turns out the damage threshold for CW power is wavelength-dependent, and the value in the datasheet was optimistic. We currently operate with between 4 mW to 5 mW through the fiber. It is also worth noting that the mechanical mate between the PM fiber pigtails attached by Jenoptik and the standard FC connector is not good, and the butt-coupling efficiency is maximized when the connector is slightly unscrewed. This would benefit greatly from a dedicated free-space launch between fibers, or launching directly into the fiber EOM, but we currently have enough CW power to make this less important.

2.2.4 Dipole trap laser

The optical dipole trap, used in the final stages of the experimental cycle, is formed with light from a 30 W fiber laser at 1064 nm (IPG YLR-30-1064-LP-SF). A series of polarizing beam cubes splits light from the laser head into four arms, each of which pass through an AOM (either IntraAction AOM-402AF4 or AOM-302AF4) and are coupled into a large-mode-area (LMA) single mode fiber (Coastal Connections S- FAmkFAmk-10nx/130/3-7.6).¹⁶ The AOM provides power stabilization and fast turn-on/turn-off, while a slower mechanical shutter built off the design in Ref. [36] provides full optical extinction.¹⁷ A schematic of this laser system is shown in Fig. 2.8.¹⁸

The power stabilization servo is the same FPGA-based design in Ref. [37]. Light from a backside-polished mirror or wedged pickoff is incident on a batterypowered photodiode. The transimpedance conversion is achieved with a fixed BNC stub terminator (either $5 k\Omega$ or $10 k\Omega$), and the resulting voltage signal is passed to a custom-built op-amp subtractor circuit. This generates an error signal by subtracting the desired setpoint voltage from the computer-controlled DACs. Finally,

¹⁶Many groups use photonic crystal fibers in this situation, but we've had mostly good luck with these LMA fibers. The advantage is they are very cheap by comparison, however we have burned the fiber tip on two separate occasions. One of these was a very perverse situation — an AOM was turned off, which diverted more power to a beam dump near an upstream mirror of a separate arm. Heat from the beam dump was enough to steer the beam off the fiber core and burn the tip. Care must be taken to align these at low power, and prevent high-power back reflections on lattice arms from burning the output fiber face as well.

 $^{^{17}}$ These shutters worked nicely for about one to two years, before they started sticking intermittently and generally falling apart. We are considering building a replacement based on the very tiny bi-stable solenoid shutter from Takano (BOS10/15). A polished sheet metal flag will replace the anodized flag that comes with this device to avoid excessive heat dissipation in the shutter itself.

¹⁸One piece of lab lore which is worth passing along: typically, the crystal in an AOM is anglecleaved to prevent a reflected acoustic wave. However, this can never be perfect — in our devices, it is ≈ -50 dBc. The result is light which has been shifted to lower frequency appearing in the +1 order of the AOM. In most cases this will have no effect, but if two arms of your dipole trap are driven on opposite AOM orders at close to (but not the same!) frequency, you end up with a running wave which heats atoms quickly out of the trap.



Figure 2.8: 1064 nm dipole laser system, reproduced from [20].

the error signal is passed to one of the digital servo inputs. Using a separate subtraction circuit allows us to run four dynamically-controlled power servos on two of the FPGA-based boxes.

In Chapter 5, we discuss a dithering procedure used to spectroscopically probe the atomic sample in the absence of AC stark shifts from the dipole trap. This is done with an Teensy Arduino microcontroller mounted onto a custom breakout circuit, which provides override logic for the TTL state of the AOM driver on each dipole arm. When triggered by the computer control system, the microcontroller toggles the sample-and-hold input of the FPGA servo in conjunction with the output TTL control of the AOM driver with a programmable duty cycle.

2.3 Magnetic field coils

Beyond the appropriate laser systems, which generate optical fields to probe and manipulate our atomic samples, magnetic fields are another handle we have to interact with the atoms via the Zeeman effect. Indeed, well-controlled magnetic fields are essential to the proper function of a magneto-optical trap. Our apparatus is no exception, and includes wire-wound coils for two field regimes: low-field shim coils, and high field coils that can be switched from (nearly) Helmholtz to (nearly) anti-Helmholtz configuration.

Details on the original coil setup can be found in Ref. [19]. Low-field shim coils are wound in pairs directly onto the vacuum flanges and are driven by an Agilent E3614A current supply, which can be dynamically controlled via analog voltage from the computer control system. These coils can handle a few amperes of current, and generate fields up to a few gauss. Their primary purpose is to cancel stray fields in the lab, and during the final stages of laser cooling, shim the equilibrium position of the red MOT to spatially overlap the atoms with the dipole trap. The verticallyoriented z-shim coil is also used to apply small quantization fields (up to ≈ 6 G) during spectroscopy.

The high-field coil pair is constructed from hollow-core copper tubing wound onto a form and inserted into the recessed vacuum viewports to get closer to the atoms, and produces a field gradient of 0.96 G cm⁻¹ A⁻¹ in the anti-Helmholtz configuration. They are driven by a TDK-Lambda GEN40-125 power supply (in constant voltage mode), and a bank of high-power MOSFETs controls the current flowing to the coils. The gate volate of these MOSFETs is controlled by a simple PI servo developed by other groups at the JQI, and includes a fast turnoff feature which shunts the gate to ground. A snubber circuit shunts the inductive load into 5Ω , which gives an $\approx 75 \,\mu$ s time constant for turnoff. This fast switching is important, because the rMOT and bMOT operate at very different field gradients ($\approx 60 \,\mathrm{G \, cm^{-1}}$ versus $\approx 1.5 \,\mathrm{G \, cm^{-1}}$). A closed-loop chiller forces cooled water through the coils to keep them from overheating.

We have recently added the ability to switch the relative polarity of our two MOT coils, to go from an anti-Helmholtz configuration to a Helmholtz configuration. This allows us to use these coils to apply very large fields, up to $\approx 200 \text{ G}$, when working with the clock line in the bosonic isotopes of strontium (see Chapter 5). This is achieved with an H-bridge built out of high-current Contactor relays (LEV200A4NAF). The relays have a switching time of $\approx 50 \text{ ms}$ to 75 ms, which must be interlocked with the coil drive to avoid switching at high current. The interlock and control of these relays is provided by a chip-programmable logic device (CPLD), which serves two functions: first, it will not change the state of the H-bridge unless the high-current servo output is disabled, and second, it holds the output of that servo to off until the H-bridge has had sufficient time (100 ms, for safety) to fully switch. The CPLD implements this logic in a programmable fabric similar to an FPGA, but which doesn't require re-initializing after a power cycle.

2.4 Computer control

Computer control of the apparatus is achieved by a heterogeneous collection of equipment coordinated by a centralized pulse generator to synchronize the timings. This structure is common amongst neutral atom experiments, since control is done on a per-shot basis and is almost always able to be pre-compiled into a sequence of steps that must happen at a specific time. In our lab, the experimental sequence is programmed in a custom LabView environment that has been built up over the years by various groups at the JQI and NIST, and is collectively known as SetList. Documentation and source code can be found on Github [38].

Each shot of the experiment involves a number of steps, defined per line of the SetList, which specifies the output state of every device connected to the control computer. Examples of devices include National Instruments DAQ cards, which provide analog and digital I/O, Arduino-controlled DDS synthesizers, CONEX linear actuators, scientific CMOS cameras, and SRS function generators, although the advantage of this architecture is any device can (in theory) be added, as long as

it can receive a timing trigger and be programmed to execute something specific on that trigger pulse. The "pseudo-clock" timing device is a PulseBlaster, which provides 24 digital outputs that can be used either as standalone TTL outputs or as triggers for other devices. The rising edge of a trigger pulse can be specified with 10 ns resolution, and the device outputs a 50 ns pulse per trigger event. We optically isolate and buffer all digital outputs, and buffer all analog outputs with BUF634 amplifiers.

An experimental cycle involves programming all devices, starting the pseudoclock triggers, and acquiring an image at the end. Processing of these images is done in IGOR to extract relevant properties of the atomic gas at the end of the experimental cycle.¹⁹ The setlist is parametrized by any number of variables defining the state of a device on a particular line. The top-most VI of SetList wraps the individual shot engine with a sequence generator, allowing randomized or linear scans of any number of variables. It also has the ability to accept a set of variables from some outside program (such as the analysis program) to achieve closed-loop operation, however we have never actually used this feature.

2.5 Custom electronics

Over the course of my degree, I have had occasion to design several custom circuits for laboratory instrumentation. Here, I describe a subset of them, especially as they pertain to laser control. One such design for an ultra-low noise, high-voltage piezoelectric driver, was published in Rev. Sci. Instr. and appears as Chapter 4 of this thesis. Other circuits which I have had considerable involvement with the design and construction, but which are not covered here, can be found on Github.²⁰ We

¹⁹Our analysis scripts are available at https://github.com/JQIamo/sr-scripts, however we have started to migrate some of these to Python for future work.

²⁰These include a full-featured AOM driver (https://github.com/JQIamo/aom-driver) for which a technical note is also available (https://github.com/JQIamo/aom-driver/blob/master/

have also made heavy use of Arduino-based microcontrollers, particularly the Teensy fork of the Arduino project,²¹ for easy integration of new hardware peripherals with our computer control system. Firmware, hardware, and associated design files can be found on the JQI organization GitHub account [39].

2.5.1 Laser diode current controller

External-cavity diode lasers have three frequency control points: diode temperature, diode current, and the length or grating angle of the external cavity [32]. Thus, under typical operating conditions, one needs three active control drivers, one for each frequency knob. Here, I discuss our solution to the second of these controls, namely, the laser diode current. The atomic physics community has gotten substantial mileage out of a design originally due to Libbrecht and Hall [40], core pieces of which have even found their way into commercial drivers (e.g., the one from Vescent Photonics). Recently, several groups have made improvements and updates to the original analog design to incorporate digital control [41] and improve noise performance near the current limits [42]. Inspired by these efforts and the desire to port the design to a Eurocard form-factor, we re-implement the key aspects of the design with some modifications discussed here. The ultimate noise performance is similar to that reported in the literature [41], and outperforms many commercial drivers.

The basic analog design is shown in Fig. 2.9(a). A low-noise op-amp senses the voltage drop across a precision resistor, R_{sense} , and adjusts the gate voltage of a MOSFET such that the voltage drop matches the voltage V_{set} . Thus, the output

aom_driver.pdf), and a linear temperature controller (https://github.com/JQIamo/Linear-Temperature-Controller based on the Wavelength Electronics WTC3243 chip.

²¹See the project website, https://www.pjrc.com/teensy/, for more information.



Figure 2.9: Key analog subcircuits for the low-noise current controller. In (a), we show the basic current controller originating from the Libbrecht-Hall design. The voltage across a precision resistor, R_{sense} , is sensed by an op-amp and controlled by feeding back to the gate of a low-noise MOSFET. The current setpoint is controlled by the voltage V_{set} , which is referenced to the supply voltage V_{supply} . In this way, noise on the supply becomes common-mode, greatly reducing the noise that appears on the output current. In (b), we show our current feed-forward circuit. A voltage from the piezo driver (see Chapter 4) is mirrored across a potentiometer. Tuning the wiper of the potentiometer allows continuous tuning of the gain from G = -1 to 1. This feed-forward current is summed with an external modulation input and the DC current provided by the schematic in (a) to generate the final current output.

current is given by

$$I_{\rm out} = \frac{V_{\rm supply} - V_{\rm set}}{R_{\rm sense}} \,. \tag{2.1}$$

The Libbrecht-Hall design also includes a very clever current modulation circuit, which we will not discuss here, but instead refer the reader to the original paper [40] or to the excellent discussion in [41].

The setpoint voltage, V_{set} , is controlled by a 16-bit DAC (AD5541) with exceptional noise performance. To reduce the differential-mode noise between the DAC and the laser current supply regulator, the DAC is referenced to the supply voltage via an ovenized Zener diode reference (LM399H).²² An op-amp buffers the negative terminal of the Zener reference, and provides a floating virtual ground for the

 $^{^{22}}$ This particular voltage reference appears ubiquitously in similar electronics designs. However, we have noticed it to still be rather sensitive to mechanical and temperature perturbations. In future iterations, I suggest looking into the LTC6655 or LTZ1000 devices from Linear Technology as possible alternatives.

DAC. To handle the reference level translation between this IC and the onboard microcontroller, which digitally controls the voltage setpoint, a pseudo-differential, quad-channel op-amp shifts the microcontroller SPI lines to the floating ground reference. The microcontroller also handles user interaction from the frontpanel, interfacing with a digital rotary encoder, switch, and a low-profile LCD display. The rotary encoder includes a pushbutton, which enables simple menu navigation, e.g., to set the maximum current range.

In contrast to previous designs, we have elected to use a newer, programmable output low-noise voltage regulator from Texas Instruments (TPS7A4700) rather than the old, albeit low-noise, LM317. This new LDO voltage regulator can supply up to 1 A, has a digital enable/disable, and the operating voltage can be adjusted by the jumper array J3 (see the schematic on GitHub [43]). Thus, one can reconfigure the operating voltage with a convenient overhead relative to the laser diode voltage drop to avoid dissipating too much heat through the MOSFET. It has a very low output noise (4 μ V RMS, 10 Hz-100 kHz), which we reduce further by filtering across 50 μ F.

The primary upgrade we have made over previous designs is to implement an adjustable-gain feed forward from the voltage supplied by a piezo driver in the same rack. The lasing frequency of an ECDL is determined by mode competition between the external cavity, the "internal" cavity formed by the (imperfectly anti-reflection coated) diode front facet, and the intrinsic gain profile of the diode itself. Physically, this manifests itself in a mode-hop free tuning range of the laser, that is, the frequency over which the laser can be continuously scanned. Because the "internal" and "external" cavities are controlled by two separate knobs (the former by the diode temperature/current, and the latter by the voltage applied to a piezoelectric actuator controlling the external cavity length), one must scan both concurrently to remain on the same lasing mode. Typically, this is done by adjusting the laser cur-

rent proportional to the external cavity piezo voltage such that the internal cavity resonance tracks the external cavity resonance. For a good pedagogical discussion of this effect, I refer the reader to the recent discussion in [44].

Implementation of this feed forward is complicated in our lab by two effects: first, the feed-forward gain is not necessarily constant across diodes, even of the same type. Second, we have both Littrow and Littman-Metcalf style ECDL lasers, which have an opposite sign in the action of their external cavity.²³ To address both design challenges simultaneously, we implement the schematic shown in Fig. 2.9(b). An op-amp (U1 in Fig. 2.9(b)) inverts the externally-supplied feed-forward voltage and drives the opposite end of a trimpot. By adjusting the trimpot wiper, one can continuously adjust the feed forward gain between ± 1 . The signal from the trimpot is buffered, and sent to the analog modulation input summation node of the current controller, and is added in with an optional external current modulation signal. If one does not wish to use the feed-forward feature, the node can be shunted to ground with a low-profile jumper to avoid introducing additional noise from EMI pickup in the lab. This trimpot is exposed on the frontpanel of our circuit, making it easy to adjust after installation with a laser.

The whole circuit is packaged into a standard 12HP-wide Eurocard module, and is designed to reside in a rack with the piezo driver described in Chapter 4. A divided down monitor of the piezo voltage can be routed via the backplane to the current controller to drive the FFWD node in Fig. 2.9(b). One can either tune the feed-forward gain with it engaged, by slowly scanning the grating piezo of the ECDL and adjusting the gain trimpot to maximize the mode-hop free tuning range, or offline by plotting the necessary current corrections to remain on the proper laser

²³Specifically, we use an old JILA design for the Littman-Metcalf lasers, which at this point shold be considered a "legacy" model and only houses our less-sensitive repump diodes. Newer lasers, and all of our narrow-linewidth systems, implement the design from Daniel Steck's group [31]. Between these two models, the frequency tuning is inverted for a positive voltage change ΔV on the piezo.

mode as the grating piezo voltage is increased. The slope of the resulting line gives the appropriate feed-forward gain; see, for example, Ref. [44].

The final performance of our driver is shown in Fig. 2.10. The spike near 50 kHz is likely spurious, due to pickup from the CRT display of the SR780, however we have not made a thorough investigation of this yet.²⁴ The total RMS-integrated noise from 1 Hz to 10 kHz is 67.2 nA. Below ~ 100 Hz, we see the performance dominated by 1/f flicker noise, with a power spectral density of $\approx 20 \text{ nA}/\sqrt{\text{Hz}}$ at 1 Hz. Above ~ 100 Hz, the noise has a white spectrum with roughly $1 \text{ nA}/\sqrt{\text{Hz}}$ noise spectral density.

2.5.2 Optical beatnote servo

The design for this board originated with the circuit presented in Ref. [45]. Electronics design files, CAD files for rackmount enclosures, and firmware for our implementation can be found on Github [46]. The basic idea is to use a ubiquitous phase-locked loop (PLL) chip from Analog Devices as a phase/frequency discriminator, which generates an error signal for a tunable proportional-integral (PI) loop filter that then feeds back to a laser. In this configuration, the RF signal fed to the PLL is generated by an optical beatnote between a reference laser and the laser to be locked; in effect, we are treating the laser as a VCO which has been demodulated from optical to RF frequencies with a reference laser.

Our design makes several improvements upon the design in Ref. [45]. First, we utilize a limiting amplifier (Texas Instruments ONET1191P) to perform signal conditioning and reduce the sensitivity of the lock to power in the beatnote.²⁵ The

 $^{^{24}}$ In chatting with students from Daniel Steck's group at DAMOP 2016, one of them mentioned they saw a very similar feature when using that spectrum analyzer. We have not had a chance to verify this hypothesis, however.

 $^{^{25}}$ This chip provides up to $\approx 40\,\mathrm{dB}$ gain, while clamping the output at a customizable maximum power, roughly $-5\,\mathrm{dBm}$ in our case, to protect RF components further down the signal chain. It is reasonably quick, up to $\approx 11\,\mathrm{GHz}$, and has the advantage of driving a differential output which is



Figure 2.10: Measured noise power spectral density for our current controller. The output is passed through a 50Ω resistor, R_{meas} , and the voltage drop across R_{meas} is AC-coupled and monitored by a low-noise instrumentation op-amp (AD8421) with a gain of G = 981, see inset. The signal is then sent to a spectrum analyzer (Stanford Research Systems SR780), which records the voltage noise power spectral density. This is converted to a current noise power spectral density by Ohm's law. The large peak around 50 kHz was mysterious, since we had not seen it previously. However, there is some indication it may be due to spurious pickup from the CRT display of the SR780 (see text).

board layout is such that reconfigurable components to change loop gain parameters are soldered on the bottom side, making them easy to swap. To account for longtimescale drifts, a slow integrator is implemented which feeds back to the piezo of an ECDL (as in [45]), however we add the option to reverse the sign of the slow integrator response with respect to the fast current modulation output. This is useful in our lab, as we have a combination of Littrow and Litman-Metcalf ECDLs which have oppositely signed piezo response. A small IDC header adapts to a microcontroller, which handles user input from the frontpanel to reprogram the PLL. Our board is pin-compatible with a line of Analog Devices PLL chips covering 50 MHz to 7 GHz.²⁶ The reference frequency is supplied by a custom-built DDS based off the AD9910 chip.²⁷ In combination with a fixed multiplier from the PLL, the DDS can be used to sweep or hop the frequency of the beatnote between or during experimental cycles.

To generate the optical beatnote, we have traditionally used the fiber-connected biased photodetectors from thorlabs (e.g., DET025AFC), into a 50 Ω termination and then to a Minicircuits low-noise amplifier (e.g., ZX60-P103LN+) to pre-amplify the signal before sending it to the beatnote circuit. A directional coupler located just after the amplifier gives a monitor signal, which is displayed on a spectrum analyzer (either individually or multiplexed with other beatnote signals, see Fig. 2.11). However, we have recently discovered a much superior solution. Hamamatsu sells

the preferred signal input mode for the phase-locked loop chips we use. There is also a loss-of-signal detection feature, which is not currently used, but might be useful in future iterations to detect low beatnote power.

²⁶These chips are the ADF4110-4113 series, plus the ADF4107. Analog Devices sells PLLs that operate as high as 18 GHz, and one might be able to adapt this design to work at those frequencies with some more careful consideration of RF PCB layout. The distinct advantage of this approach is the feedback loop filter does not need to work at high frequencies; rather, the PLL demodulates the beatnote signal to generate an error signal which can be fed into a standard PID controller.

²⁷For files, see https://github.com/JQIamo/ad9910-dds. This board was designed by Carlos Bracamontes and Alessandro Restelli, and provides digitally-synthesized RF waveforms up to $\approx 400 \text{ MHz}$.



Figure 2.11: Sample trace for optical beatnote lock. On the left, we see the typical performance (power averaged over five traces, 1 kHz RBW, 10 MHz span). The mean-square phase error, an indication of the lock performance, is computed to be $\langle \phi^2 \rangle = 0.100$ (see text). The lock bandwidth, as indicated by the position of the servo bumps, is roughly 1 MHz in this trace. It is possible to push the peak of the servo bump out to a few MHz, however the laser performance degrades as measured by atom number in the rMOT and similarly the mean-square phase error increases. We find the optimal performance corresponds to adjusting the gain such that the servo peaks are flattened but close in to the carrier. On the right is a screenshot from our web-based monitoring program. A small webserver written in Python directs a multiplexed RF switch to alternate between monitor channels, which are then sent to a Rigol spectrum analyzer. The program synchronously pulls traces from the spectrum analyzer over TCP/IP to display in the web browser. This allows us to simultaneously monitor several beatnotes from anywhere in the lab.

a GaAs metal-semiconductor-metal (MSM) photodetector with a 30 ps risetime, in a convenient SMA-connected package (part number G4176-03).²⁸ The bias voltage, ≈ 7 V, is supplied through a bias-tee mounted inside the box with the remaining beatnote circuitry. The only disadvantage is the photodiode is not fiber-coupled; however, one can easily rig a fiber launch assembly and clamp the photodiode using standard optical mounting components. Alternatively, and I suggest this for others wishing to implement this setup, a custom machined mounting plate could be designed and ordered from any of the fast-turn CNC prototype companies for a nominal price.²⁹ Using this photodiode and a beatnote board sporting the speedy

²⁸See their website for details, https://www.hamamatsu.com/us/en/product/category/3100/ 1025/G4176-03/index.html. At the time of our purchase, they were selling for \$174, which is cheaper than the much slower version from Thorlabs.

²⁹We have had good success, on both price and quality/leadtime, with Protolabs. However, their sales representatives are perhaps over-eager and will call you more than necessary about "future projects"!

ADF4113 PLL, we have easily locked a beatnote around 2.6 GHz.³⁰

We show the typical performance of our circuit in Fig. 2.11. The performance figure of merit, in this case, is the the mean-square phase error $\langle \phi^2 \rangle$, defined via the equation [45, 47]

$$\exp\left(-\langle \phi^2 \rangle\right) = \frac{P_{\text{carrier}}}{\int P(\nu)d\nu} \,. \tag{2.2}$$

Physically, it is a measure of how much power is in the delta-function peak, P_{carrier} , relative to the amount of power in the rest of the beatnote spectrum integrated over frequency ν . Ultimately, the lock quality is decided by the atomic response. Here, we primarily use this circuit to offset lock lasers addressing the narrow ($\approx 7.5 \,\mathrm{kHz}$) intercombination transition in strontium. We do notice some degradation of the MOT performance if the beatnote servo bumps are too pronounced; presumably, this corresponds to the onset of an oscillation. This behavior is strongly determined by where your laser is in {piezo, current} space. After locking both the high-frequency current feedback and the low-frequency piezo "infinite integrator", increasing the laser current tends to reduce the effective volt-to-frequency gain (and vice versa) as the piezo is driven to a new operating point to compensate. For the "Long Steck" laser (see Fig. 2.6), we tend to operate with the beatnote gain turned all the way up but then push the servo bumps low and closer in to the carrier by increasing the laser diode current after the system is locked. For the "Short Steck" lasers, depending on the frequency, the beatnote gain does not need to be turned very high, but can also benefit from adjusting the current after the laser is locked.³¹ In an attempt to quantify this heuristic, we have recently added a measure of $\langle \phi^2 \rangle$ to our beatnote

 $^{^{30}}$ This is as high as we've had any reason to go, however, the Rubidium Rydberg lab uses this board with the ADF4107 PLL and a fast commercial photodiode to lock at 6.8 GHz.

³¹While not very scientific or repeatable in an absolute sense, this is just one of many personality quirks one becomes accustomed to in dealing with particular lasers! Enough time spent staring at the beatnote monitor spectrum gives you a sense of what a good lock looks like in terms of MOT atom number and temperature.

spectrum multiplexing software, see Fig. 2.11b. A useful feature to implement in the future is a long-term logger of $\langle \phi^2 \rangle$, which could be correlated back to atom number to say something more quantitative about the lock performance.

While these beatnote circuits are generally pretty robust, there are a few improvements which should be made in any future version. First, integrating the microcontroller on-board (such as in the high-voltage piezo driver, described in [48] and Chapter 4) will greatly simplify the box assembly, and enable a more tightly integrated electronics package. Second, I suggest moving to a tunable PI controller, like the one implemented in [49].³² We have found the "default" PLL loop filter components suggested in Ref. [45] to work well, but depending on the laser, the overall loop gain is easily too low or too high. And the current method, to adjust the final gain by attenuation through a potentiometer, is problematic from a signal-to-noise perspective but also because, in the case of the unlimited piezo integrator, one often wants a very small gain but then is hampered by a very small dynamic range.

There are a few layout improvements which will be useful for diagnosing problems. First, placing probe/testpoints after the limiting amplifier could rule out many signal integrity issues, which can arise especially if the PCB is populated by hand. Second, better attention to how this board will plug together with other RF components in the final box will reduce the headache in assembly. However, both of these are cosmetic modifications.

³²This, in turn, is based off the implementation of Joe Tiamsuphat in the famous Joe-PI, see https://github.com/JQIamo/analog-pi.

Chapter 3: Design and construction of a narrow-linewidth clock laser

As discussed in the introduction, strontium is an appealing system for atomic physics because of its somewhat unique combination of atomic properties. While it doesn't have nearly as many stable isotopes as ytterbium, it does have three bosons with scattering lengths that span almost the entire spectrum from near zero $(a_{88} \simeq -2a_0)$ to near unitarity $(a_{86} \simeq 823)$, and a fermion with an unusually high nuclear spin $I = 9/2^1$. The optical electronic transitions similarly span a wide range, from the broad-line cooling transition at 461 nm ($\Gamma/2\pi \approx 32$ MHz) to the narrow-line intercombination transition at 689 nm ($\Gamma/2\pi \approx 7.5 \,\mathrm{kHz}$) to the mHz-wide clock transition at 698 nm. These disparate transition linewidths provide handles for manipulating strontium with varying levels of coherence and control, and arise from dipole selection rules that are broken by terms in the Hamiltonian coupling the "forbidden" level to levels with fully dipole-allowed transitions. For example, spin-orbit induced state mixing between 1P_1 and 3P_1 creates a small admixture of ${}^{1}P_{1}$ character in the true ${}^{3}P_{1}$ eigenstate, which we've labeled with the Russel-Saunders (LS) coupling term symbol for convenience. Similarly, hyperfine state mixing between states with the same F quantum number creates a small admixture of ${}^{3}P_{1}$ character in the ${}^{3}P_{0}$ eigenstate for 87 Sr, making the strictly dipole-forbidden $J = 0 \rightarrow J' = 0$ transition very weakly allowed. [50] No such hyperfine state mixing exists in the bosonic isotopes, and so we must artificially mix these levels with an

¹This is not so unusual from a nuclear physics perspective, but certainly unusual for species used in laser-cooled atomic physics experiments.

external magnetic field.

In this chapter, I discuss the design and construction of a "clock laser" to address the mHz clock transition in strontium. Target specifications, general considerations for achieving narrow linewidths, and the ultimate performance will be presented in the sections that follow.

3.1 Motivation and target specifications

As its name suggests, the "clock" transition in strontium was initially explored for the purpose of making next-generation atomic clocks, which are now reaching resolutions $\leq 1 \times 10^{-18}$ [51–55].² To get an idea of the precision indicated by this order of magnitude, the age of the universe is roughly 4.3×10^{17} s, meaning a clock with state-of-the-art precision would have lost ≤ 100 ms since the big bang. This level of precision is enabled by recent advances in stable laser systems, which transfer the stability of a Fabry-Pérot cavity in a very carefully engineered environment to the optical domain [57, 58]. The applications for stable laser systems extend far beyond pedestrian atomic clocks, to include spectroscopic observation of quantum many-body phenomena, gravitational wave detection, searches for dark matter, and even deep-space navigation (see [58] and references therein).

For our purposes, we will not need to push the frequency stability of our narrow laser to the cutting edge, although it is worth noting that several companies now offer sub-Hz laser systems for sale, optionally in conjunction with a super-continuum frequency comb to transfer that stability elsewhere in the electromagnetic spectrum. Instead, we wish to ultimately use this laser as a spectroscopic probe of nuclear spin dynamics in ⁸⁷Sr. The differential Zeeman splitting between ${}^{1}S_{0}$ and ${}^{3}P_{0}$ in

²The program to redefine the SI second, of course, relies on both the precise measurement of frequency but also on the careful consideration of systematic frequency shifts. A number of candidate systems are under consideration, from trapped ions to neutral atom lattice clocks, including neutral strontium. For a recent review, which is now already out of date with state-of-the-art stable laser systems, see the 2015 Rev. Mod. Phys. publication from Ludlow *et. al.* [56].



Figure 3.1: Level diagram for the clock transition in 87 Sr. Zeeman sublevels are split by 109 Hz/G, shifting adjacent levels off-resonance from the spectroscopy laser. In this manner, individual ground-state sublevels can be addressed and manipulated for state preparation and readout.

⁸⁷Sr is $m_F \times 109 \text{ Hz/G}$ [59], meaning a laser with $\approx 100 \text{ Hz}$ stability and a modest magnetic field will energetically resolve individual Zeeman transitions between the ground and clock-state manifolds as schematically illustrated in Fig. 3.1. If we can reliably prepare and read out the nuclear spin state of our atomic gas, we can begin to explore some interesting proposals in quantum simulation using alkaline-earth atoms [2,6,7]. With upgrades made on the new strontium experiment described in Chapter 6, we hope to pursue some of these theoretical proposals. In the meantime, we use this laser to perform isotope shift spectroscopy of the clock transition, as described in Chapter 5.

For the remainder of this chapter, we will assume a target linewidth of 100 Hz, which at 698 nm represents a fractional stability

$$\frac{\delta\nu}{\nu} \sim \frac{100 \,\mathrm{Hz}}{429.228 \,\mathrm{THz}}$$

$$\approx 1 \times 10^{-13} \,.$$
(3.1)

With this target specification in mind, we now discuss some of the engineering challenges required to achieve one part in 10^{13} fractional stability.

3.2 Pound-Drever-Hall: noise analysis and limitations

As briefly mentioned in the previous section, narrow-linewidth laser systems achieve their stability by transferring the excellent mechanical stability of a reference cavity to the light via a Pound-Drever-Hall (PDH) lock [30,60]. While care must be taken not to introduce noise via the lock electronics, in principle the PDH technique can achieve shot-noise-limited laser linewidths. Indeed, state-of-the-art laser systems are limited instead by Brownian noise in the cavity spacer, mirror substrate, and mirror coating [58,61]. In Section 3.2.4 we discuss considerations surrounding the PDH lock itself, but first we turn our attention to mechanical considerations in designing an optical resonator capable of delivering a 100 Hz laser.

Inherent in the PDH laser stabilization scheme is the assumption that an optical Fabry-Pérot cavity is capable of supporting a resonant mode whose center frequency has the same fractional stability that we seek to achieve for our laser. Quantitatively, we express the cavity resonance condition $\nu = mc/(2nL)$, where ν is the resonance frequency, c is the speed of light, L is the cavity length, n is the index of refraction, and m is the longitudinal mode number. This comes from requiring that the optical phase be unchanged modulo 2π after one round-trip in the cavity. Considering the response of the cavity resonance, $\delta\nu$, to uncorrelated perturbations δL and δn in the length and index of refraction (respectively), we find

$$(\delta\nu)^{2} = \left(\frac{\partial\nu}{\partial L}\delta L\right)^{2} + \left(\frac{\partial\nu}{\partial n}\delta n\right)^{2}$$

= $\left(\frac{mc}{2nL^{2}}\right)^{2}\delta L^{2} + \left(\frac{mc}{2n^{2}L}\right)^{2}\delta n^{2},$ (3.2)

Taking the linear response to perturbations in each quadrature separately, and ex-

pressing it as a fraction, we find that the fractional stability scales like

$$\frac{\delta\nu}{\nu} \sim \frac{\delta L}{L}, \, \frac{\delta n}{n}.$$
 (3.3)

To achieve $\delta\nu/\nu \sim 10^{-13}$, we must stabilize both $\delta L/L$ and $\delta n/n$ at the same order. For the 100 mm cavity we will use, this means $\delta L \leq 10^{-14} \text{ m.}^3$ This is a rather stringent requirement, and requires careful consideration of perturbations that couple into length changes δL or changes in the index of refraction δn . We will treat each in turn.

3.2.1 Index of refraction

From the Lorenz-Lorentz equation, we know that the index of refraction n is related to the mean polarizability α and density ρ of a gas by

$$\frac{n^2 - 1}{n^2 + 2} = \frac{4\pi}{3} \alpha \rho \,. \tag{3.4}$$

Using the standard equation of state for an ideal gas and solving for n, we find

$$n = \sqrt{\frac{1 + 2\left(\frac{\beta}{k_B T}\right)P}{1 - \left(\frac{\beta}{k_B T}\right)P}}$$
(3.5)

$$\approx 1 + \frac{3}{2} \left(\frac{\beta}{k_B T} \right) P,$$
 (3.6)

where P is the pressure, T is the temperature, k_B is the Boltzmann constant, and $\beta = 4\pi\alpha/3$. The second equality in Eq. 3.6 is valid in the limit $(\beta/k_BT)P \ll 1$.

³As a point of comparison — this is equivalent to stabilizing the length of our cavity to ten times the width of a proton. Clearly, the length isn't well-defined at that level; the light samples a "large" region of the mirror (given by the mode area), and thus the length δL is in some sense averaged over the transverse extent of the optical field.

For air, one can calculate

$$\frac{3}{2} \left(\frac{\beta}{k_B T} \right) \equiv \beta_0 \sim 10^{-9} \,. \tag{3.7}$$

In this limit, we relate $n=1+\beta_0 P$ and take β_0 as our small parameter to obtain

$$\frac{\delta n}{n} = \frac{\beta_0 \delta P}{1 + \beta_0 P} \approx \beta_0 \delta P \,. \tag{3.8}$$

Equation 3.3 requires $\delta n/n \lesssim 10^{-13}$. Taking $\beta_0 \sim 10^{-9}$, this means we must suppress $\delta P \lesssim 10^{-4}$ Pa. The easiest way to reduce pressure fluctuations is to place the cavity under vacuum, which has the added benefit of partially decoupling other environmental perturbations from the cavity (such as temperature and acoustic fluctuations, discussed later). In convenient units, 10^{-4} Pa $\sim 10^{-7}$ torr, which is easily achieved using standard vacuum techniques. The pressure in our vacuum canister is in the mid- 10^{-8} torr. This is higher than expected, given the somewhat lower pressure $\approx 1 \times 10^{-8}$ torr in the vacuum canister housing the 689 nm cavity (see Sec. 2.2.2), but is workable for our purposes. We suspect that either a slightly longer bakeout needed to occur, or that the KF "Kwik" flange attaching the ion pump to the canister has introduced a slow leak.⁴

3.2.2 Mechanical length stability

Equation 3.3 also requires the fractional length stability $\delta L/L \lesssim 10^{-13}$. Our lab had previously purchased a Fabry-Perot cavity from Advanced Thin Films (now part of IDEX) formed by two high-reflectivity mirrors optically contacted to a passive 100 mm spacer ($\Delta \nu_{\rm FSR} = 1.5 \,{\rm GHz}$) made of ULE glass; see Fig. 3.2(d) for

⁴The KF flange is only rated to $\approx 10^{-8}$ torr.

more details.⁵ With this cavity, length fluctuations of the spacer must be suppressed to $\leq 10^{-14}$ m. Damping of high-frequency vibrations can be achieved with mounts made of elasto-polymers such as viton, which absorb mechanical energy in this regime, and low frequency vibrations can be suppressed with a vibration isolation platform. We purchased one such isolation stage, the model 100BM-8 from Minus-K, which provides up to 60 dB of attenuation at frequencies ~ 100 Hz and has a natural frequency ~ 0.5 Hz. Figure 3.3 shows the measured vibrational power spectral density in the lab (before isolation with the Minus-K stage), along with the manufacturer-specified performance of the vibration isolation stage.

In addition to these vibration-damping techniques, a clever mounting configuration can make the cavity length first-order insensitive to elastic deformation due to vibrational acceleration [63]. One can study this problem via numerical simulation, using a strategy known as finite-element analysis (FEA). FEA is a numerical modeling technique used to solve complex boundary-value problems, where analytical solutions are unknown due to the complexity of the shape being modeled. It involves dividing the larger system into smaller "finite elements," approximating the differential equation on that discretized region such that it can be solved, and then stitching together the results to find a global solution. There are many FEA solvers available, both commercial and open-source, with varying degrees of preand post-processing capabilities. We chose to use the MultiPhysics FEA engine by COMSOL to model the Fabry-Pérot cavity under various mounting geometries to aid in the design of a mount which minimizes the sensitivity of the cavity to "bad"

⁵We in fact purchased two such identical cavities, with high-reflectivity mirror line coatings at 689 nm and 698 nm. The plan was to gain experience working with high-finesse cavities for the intercombination-line rMOT, which could be used in the design of the clock laser. This was certainly true, however in hindsight this cavity is overkill for stabilizing the rMOT and perhaps "underkill" for the clock application, as will be discussed below. That said, having the enhanced stability for our 689 nm lasers has worked out well, since we do not need independent spectroscopy on that line and it has facilitated the study of intercombination-line photoassociation resonances, described in [62].





Figure 3.2: (a) Reduced CAD model of the ULE cavity. Appropriate boundary conditions are defined on the planes of symmetry to reduce the computational complexity of the FEA simulation. (b) Meshed geometry. (c) Visualization of how the cavity deforms under a uniform vertical acceleration. Color scale shows displacement along the cavity axis, which is the figure of merit for $\Delta L/L$. (d) Blueprints for the ULE cavity, provided by Advanced Thin Films.



Figure 3.3: At left, the performance curves for 100BM-8 vibration isolation platform from Minus-K. At right, we show the mechanical noise spectral density in our lab, measured with a Wilcox accelerometer. Total integrated RMS noise (1 Hz - 200 Hz) is $< 10^{-3} \text{ m/s}^2$; with the additional damping provided by the Minus-K stage, we expect to remain below the 10^{-3} m/s^2 RMS noise threshold required to achieve our 100 Hz linewidth (see text). It is worth noting the higher mechanical noise floor of the elevated 1/2" breadboard; the assumption that mounting these on rubber feet would further damp vibrations is not borne out in the data.

Mechanical Property	Symbol	Value	
Density Young's Modulus	ho E	2.21×10^{3} 67.6	$ m kg/m^3$ GPa

Table 3.1: Summary of mechanical properties for ultra-low expansion glass (ULE). These are used in performing a finite-element analysis simulation of the Fabry-Pérot cavity, see text.

elastic deformation. The simulations themselves were conducted according to the methods outlined in [63]. A model of the Fabry-Pérot cavity was constructed using the COMSOL model builder, and the material properties of ULE were used (see Table 3.1). We additionally leveraged the cylindrical symmetry of the cavity to reduce the computational complexity of the model. The CAD model for the final design can be seen in Fig. 3.2, along with a visualization of the elastic deformation under uniform vertical acceleration.

As discussed in Ref. [63], we are primarily concerned with analyzing the mechanical response at low frequencies (≤ 100 Hz). Higher Fourier components can be easily attenuated using traditional vibration isolation techniques, such as mounting the cavity vacuum chamber on rubber feet, placing the entire setup on a floated optical table, and enclosing the setup in a box lined with acoustic-damping foam. In general, perturbations ≥ 50 Hz are caused by acoustic noise sources, while those ~ 1 Hz to 50 Hz are usually seismic [64]. With this in mind, we consider designs that minimize the coupling between external mechanical noise and the cavity length. At vibrational frequencies f much smaller than the lowest normal mode of the cavity, we are able to greatly simplify the problem by modeling the mechanical response as a quasi-static deformation; that is, vibrations which are close to DC compared to the cavity's first normal mode can be treated as uniform accelerations. Reference [63] does a back-of-the-envelope calculation to show that we expect the response at frequencies ≤ 35 kHz to be well approximated by a quasi-static treatment, and our FEA model confirms this.



Figure 3.4: Length fluctuation ΔL due to mirror tilt. (a) shows the basic cavity geometry. In (b) and (c), the mechanical axis is denoted with the dashed black line, while the optical cavity axis is shown in red. Tilt modes with even symmetry only affect the optical cavity length at first order if there is a rotation between the mechanical and optical axes, shown in (b). Odd-mode mirror tilt affects the cavity length at first order for an optical axis which is displaced from the mechanical axis, shown in (c). Symmetry considerations predict odd tilting modes for accelerations transverse to the cylindrical axis, while accelerations along the axis of the cylinder induce even tilting modes.
We simulate the cavity response to a uniform vertical acceleration under a variety of mounting configurations and boundary conditions. The figure of merit is the fractional path length change of the optical mode in question (here, the Gaussian TEM₀₀ mode). At lowest order, this is the fractional axial displacement of the mirrors, $\Delta L/L$. Mirror tilt also couples to the optical length of the cavity at first order if the optical and mechanical axes are not coincident, see Fig. 3.4. Since the degree of misalignment between the mechanical and optical axes is hard to estimate *a priori*, we will simply put bounds on the displacement and rotation of the cavity axis relative to the mechanical axis such that the tilt-induced response to a uniform acceleration has the same order of magnitude as a longitudinal shift under the same acceleration.

We first consider the so-called "rod" mounting geometry, used by our group and others to hold a stabilization cavity for the 689 nm intercombination line laser, as well as by some groups for pre-stabilization of a 698 nm clock laser [65]. In this configuration, an aluminum block supports two viton rods at a particular angle from the vertical symmetry plane, which in turn support the ULE cavity. Figure 3.5 shows the resulting (longitudinal) linear response coefficient K_a as a function of θ , the central half-angle. The linear response coefficient is defined via

$$\Delta L/L = K_a \cdot a \,, \tag{3.9}$$

where a is a uniform acceleration. As expected, the closer you are able to support the cavity to a plane of symmetry ($\theta \rightarrow 90^{\circ}$), the less susceptible it becomes to deformation. However, central half-angles greater than ~ 50° to 60° become precarious, as the vertical support is reduced and the cavity is at risk of falling. There is also no true zero-crossing in K_a using this configuration.

Our ultimate choice in mounting geometry, which we discuss next, is inspired



Figure 3.5: FEA analysis for "rod" support method. Fractional length stability under a 1 m/s^2 vertical acceleration is shown at left, for various central half-angles. The geometry is illustrated at right.

by other groups' efforts to find a zero-crossing in K_a [63, 66, 67]. Heuristically, a transverse acceleration compresses the cavity, thereby inducing an elongation along the cavity axis where the magnitude of this effect is governed by Poisson's ratio. However, a competing effect can act to compress along the cavity axis — by supporting it at two points displaced from the center of the cavity, a transverse acceleration causes the middle part to "droop," inducing stresses that point inward. A proper balance of these two effects can in principle cancel first-order longitudinal shifts in $\Delta L/L$. The "cutout" cavity offered by Stable Laser Systems and deployed in several labs takes advantage of this [66,67]. In our case, we are constrained by the existing cavity geometry shown in Figure 3.2(d), which uses an ordinary cylindrical spacer. Our simulation and ultimate design assumes a configuration in which the cavity is supported by four viton pads at a particular angle from the central symmetry plane. The exact angle, width, and position of these pads is varied such that we minimize K_a . The final geometry is shown in Fig. 3.6.

Unfortunately, we found that the exact assumptions made in modelling the boundary conditions at the support points can produce large and qualitative differences in the final value of K_a , as seen in Fig. 3.7. The mounts themselves are specified in COMSOL as small boundary regions, which are constrained to zero dis-



Figure 3.6: Final design for the cavity holder. Viton pads rest in the four pocket cutouts, and provide some additional damping. A base plate, also suspended by viton rods, fixes the cavity axis at a 4° tilt to avoid introducing parasitic etalons between the cavity and the vacuum viewports. All dimensions are in mm.



Figure 3.7: Meshing and boundary model comparison. At left, we show the spread in fractional length stability for different mesh granularities (specified in COMSOL) for a zero-size, fixed-boundary constraint. There is no systematic trend, but the spread gives a sense of the simulation uncertainty at the ~ 10^{-11} level. At right, we show the effect of varying boundary constraints. The color shading, from light to dark, indicates different widths along $\hat{\phi}$ of 0.5 mm, 1 mm, and 2 mm. The specific boundary model used makes a large and qualitative difference in the simulation results, see discussion in text. Plots assume a mounting angle $\theta = 60^{\circ}$, and a pad width w = 10 mm.

placement in different dimensions. Unlike in [67], we did not find much dependence on the exact size of the constrained region, but we do reproduce their observation that K_a depends strongly on the constraints placed on displacement degrees of freedom at the mountpoint. The design in [67] uses four rigid support rods placed under the shoulder of a cutout in the cavity spacer. In that configuration, it is obvious what the constrained displacement is, and indeed they find the best match between simulation and experiment when the model only constrains the vertical displacement to zero.

In our case, the mount is a viton pad, which provides a compliant mechanism to transfer force between the aluminum mounting block and ULE cavity. Rather than deal with the complex physics involved in simulating this boundary constraint, we model the quasi-static force exerted by the viton mounting pad as a body load along the \hat{x} , \hat{y} , or \hat{z} direction while the contact region between viton and ULE remains fixed at zero displacement. We are assuming that the friction between the viton and ULE supplies the necessary force in all dimensions for a no-slip, steady state solution. Under these assumptions, we fix the displacement of the boundary region between the viton pad and ULE spacer. A visualization of these different boundary models and their effect on K_a is summarized in Fig. 3.7. We also investigate the effects of finite mesh size on our simulation. Figure 3.7 shows K_a for various meshing parameters ranging from "Normal" to "Extremely Fine". As can be seen from the plot, meshing errors in the simulation arise at the $\sim 10^{-11}$ level, and can be ignored relative to other systematic errors in the simulation, namely, those arising from ambiguities in the proper boundary constraints.

In terms of mirror tilt, we found that there is a zero crossing in tilt sensitivity at a particular pad position along the cylindrical axis, which was very insensitive to any other varied geometrical parameter. This is unsurprising, as there should exist an Airy point for our geometry. For a straight beam of length L under uniform load, the Airy points exist a distance $L/\sqrt{3}$ apart. Supporting the beam at these points guarantees the two ends to be vertical (*i.e.*, zero tilt of the end faces). For a 100 mm cavity, the Airy points should in theory lie 28.9 mm from the cavity midpoint; we find the optimal central mount position to be between 28 mm and 29 mm, in good agreement with this prediction.

Following Ref. [68], we fit the results of the FEA simulation to interpolate between modeled conditions. The fit is done up to quadratic order, and can be represented as a matrix equation

$$\begin{pmatrix} \Theta \\ K_a \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} \end{pmatrix} \cdot \begin{pmatrix} 1 \\ \theta \\ d \\ \theta \cdot d \\ \theta^2 \\ d^2 \end{pmatrix}, \quad (3.10)$$

where Θ is the fractional tilt sensitivity and K_a is the cavity length susceptibility

discussed previously. The variables θ and d parameterize the mounting pad halfangle and position, respectively, as defined in Fig. 3.8. This can be written more compactly as

$$\mathbf{S} = \mathbf{M} \cdot \mathbf{X}, \qquad (3.11)$$

where now **X** and **S** are both matrices whose rows represent the simulated mounting conditions and associated deformations. While the matrices involved are not strictly invertible, one can solve this equation numerically by least-squares minimization. We do this, given simulation data for $\theta \in [45^{\circ}, 90^{\circ}]$ and $d \in [25 \text{ mm}, 40 \text{ mm}]$. The results can be seen in Fig. 3.8.

Given all the results from the numerical FEA model, we choose the mounting configuration shown above in Figure 3.6. The pads are 10 mm wide, positioned 37 mm from the center axis and at a 60° azimuthal angle. A central half-angle of 60° gives a large enough safety margin for supporting the cavity against gravity while still allowing a zero-crossing in K_a . Furthermore, this zero crossing was preserved for all pad widths investigated, from w = 10 mm to w = 18 mm. Since the exact width does not seem to have a great impact on the final result, we use a pad width of 10 mm for convenience. We expect the longitudinal linear response coefficient $K_a \approx 10^{-11}/(\text{m s}^{-2})$, and the tilt sensitivity to be $\approx 2 \text{ nrad}/(\text{m s}^{-2})$. As discussed above and illustrated in Fig. 3.4, for a tilt sensitivity at this level, we can tolerate a displacement of $\sim 0.25 \text{ mm}$ between the cavity and optical axis before mirror tilt begins to dominate over longitudinal displacement.

Given the expected linear response coefficient $K_a \sim 10^{-11}/(\text{m s}^{-2})$, we must ensure vibrations are damped below 10^{-2} m s^{-2} . In Fig. 3.3, we showed the amplitude spectral density of vibrations in our lab, measured on a Wilcox accelerometer



Figure 3.8: Summary of results from the clock cavity FEA. At top, we show the points and quadratic fit to both tilt and length sensitivity as a function of angle θ and position d. The contour plot shows the same fitted surfaces, with the zero-crossing contours highlighted in red. Dashed lines correspond to the length sensitivity, while solid lines correspond to the tilt sensitivity. The blue dot is the experimentally realized mounting configuration, and the green shaded region is what we consider "safe" mounting angles. The table shows the fit coefficients, and the final sensitivity at the experimentally realized configuration is reported in the last line; see text for further details.

(model 728A, connected through a P703B transducer).⁶ The total RMS noise (1 Hz to 200 Hz) is $1 \times 10^{-4} \text{ m/s}^2$ ($3 \times 10^{-4} \text{ m/s}^2$) measured directly on the optics table (elevated 1/2-inch breadboard), respectively. This is tolerable by itself, however the Minus-K stage will provide an additional 60 dB of attenuation, and floating our optical table will further reduce seismic noise. Finally, a well-designed box around the cavity and laser will provide shielding from acoustic noise sources.

3.2.3 Thermal length stability

In addition to mechanical deformation, temperature fluctuations can couple into the length of the cavity through the coefficient of thermal expansion, α_{CTE} , according to

$$\frac{\delta L}{L} = \alpha_{\rm CTE} \ \delta T \tag{3.12}$$

The cavity is made of Corning ultra-low expansion (ULE) glass, which has $\alpha_{\rm CTE} \lesssim 2 \times 10^{-8}$ /K [70]. In addition, ULE has a zero-crossing in $\alpha_{\rm CTE}$ somewhere near room temperature, at which point the length becomes insensitive to first order changes in temperature. State of the art stable laser systems take great care to tune the temperature to this zero crossing [71], and depending on the complexity, employ multiple temperature-stabilized heat shrouds in vacuum to reduce thermal coupling via blackbody radiation [57]. For our purposes, we actively stabilize the temperature of the vacuum housing with a resistive heater.

To find the zero crossing of α_{CTE} , we locked a 689 nm laser to the "clock" cavity and monitored a beatnote between this laser and our frequency-stabilized

⁶The spectrum was recorded on a PicoScope 4262 digital oscilloscope, using 4096 frequency bins over a 200 Hz bandwidth. The Wilcox has a frequency corner at ~ 1 Hz; additionally, low-frequency spectral components in Figure 3.3 likely suffer from window leakage at DC (spectrum reported with Blackman windowing) and over-estimate the vibrational noise power. We have divided by 1.73 to normalize the power spectral density with respect to the Blackman windowing function [69].

rMOT laser as the temperature setpoint was slowly increased. After each change in setpoint, we either wait sufficiently long for the whole system to thermalize (as measured by the beatnote reaching a steady-state value), or fit the beatnote to an exponential and extracted the steady-state value and time constant. Figure 3.9 shows the results of this analysis. Near the zero-crossing temperature T_0 of $\alpha_{\rm CTE}$, we can express it to first order as

$$\alpha_{\rm CTE} = \beta (T - T_0) \,. \tag{3.13}$$

Plugging this expression into Eq. 3.12, and recalling that $\delta\nu/\nu = -\delta L/L$, we get

$$\delta\nu = \left[\beta \left(\delta T\right)^2\right]\nu \tag{3.14}$$

where $\delta T = T - T_0$, and β is the linear susceptibility of α_{CTE} to changes in temperature about T_0 . From the parabolic fit in Figure 3.9, we extract $T_0 = 30.34(6)$ °C and $\beta = -2.48(7) \times 10^{-10} / \text{K}^2$.

Temperature controllers, such as the WTC3243 from Wavelength Electronics, can stabilize temperatures at ~ 10^{-3} K.⁷ Taking a conservative estimate of the temperature stability, $\delta T \approx 10$ mK, we can easily achieve a fractional stability $\delta \nu / \nu \lesssim 10^{-13}$. The short-term stability is further improved by the effective lowpass filter created by the small thermal conductance between the environment and the cavity itself. Thus, with our setup, we are able to achieve sufficient temperature stability for the target linewidth.

Finally, it is worth mentioning that even if one operates the cavity at the zero-

⁷We have developed a custom temperature controller to use in our lab based on the evaluation board for this chip. It is capable of controlling up to four devices, and can be configured to drive thermo-electric coolers (TEC) or resistive heaters. The feedback gain and integrator time constant, as well as setpoint, temperature, and output current monitor, are exposed via microcontroller interface to a frontpanel display. For Eagle design files, see the GitHub repository https://github.com/JQIamo/Linear-Temperature-Controller.



Figure 3.9: Temperature dependence of the clock laser cavity resonance, used to find the temperature T_0 where $\alpha_{\rm CTE} = 0$. We lock a 689 nm laser to the clock cavity, and track the beatnote between that laser and the narrowed rMOT "master" laser as we adjust the temperature of the clock cavity. Thermalized, steady-state beatnote values are extracted by fitting the beatnote to an exponential and finding the infinite-time asymptote. Here, we neglect the point at 23 °C in the fit because it increases the reduced χ^2 per degree of freedom to $\chi^2/6 \approx 3$, which we can reject with $\alpha \approx 0.006$ in a χ^2 test. We suspect that, far from the zero-crossing temperature, the functional form of $\alpha_{\rm CTE}$ deviates from the linear approximation used here. Errors on each point are taken as the 1σ confidence interval in the steady-state value extracted as described above, and used to weight the parabolic fit to find the zero-crossing temperature, T_0 . Performing a χ^2 test on the remaining datapoints gives a reduced $\chi^2/5 = 0.72$, which is slightly smaller than the ideal value $\chi^2/5 = 0.87$, but can be accepted with an $\alpha = 0.6$. The curvature β (see text) is found to be $\beta = -2.48(7) \times 10^{-10} / \text{K}^2$, and $T_0 = 30.34(6)$ °C. From these parameters, a conservative temperature stability $\delta T \sim 10 \,\mathrm{mK}$ around T_0 is enough to stabilize the fractional length to $\lesssim 10^{-13}$. The very low thermal conductance between the vacuum housing and the cavity itself further serves to low-pass filter any temperature fluctuations, so the short-term thermal stability will be improved even beyond the estimates here.

crossing of α_{CTE} , Brownian motion induced by thermal fluctuations in the cavity spacer, mirror substrates, and mirror coatings lead to perturbations in the cavity length [61]. This thermal noise can be limiting in some systems [72], and has driven efforts to build next-generation cavities operated at cryogenic temperatures [57,58]. Here, we verify that thermal fluctuations are below other noise sources in our system.

Initial treatments of thermal noise in optical cavities were motivated by the long-standing effort to build gravitational wave interferometers [61, 73]. The key physical insight comes from the Fluctuation-Dissipation theorem: thermal fluctuations of a system in equilibrium are related to the dissipative linear response of that system to an external drive [74].⁸ Quantitatively, we can write the fluctuation-dissipation theorem as [61, 75]

$$G_X(\omega) = -\frac{4k_BT}{\omega} \mathbf{Im} \left[H(\omega) \right] \,, \tag{3.15}$$

where $G_X(\omega)$ is the single-sided power spectrum of the fluctuation X at angular frequency $\omega = 2\pi f$, k_B is the Boltzmann constant, T is the temperature, and $H_X(\omega)$ is the (complex) linear-response transfer function of the system, given by

$$H(\omega) \equiv \frac{X(\omega)}{F(\omega)} \tag{3.16}$$

for some generalized force F. The imaginary part of H gives the dissipative portion of the response. What Eq. 3.15 implies is that, by measuring the mechanical linearresponse of a system, we can make a statement about the power spectrum of thermal fluctuations in equilibrium.

Equation 3.15 is a normal-mode representation of the fluctuation-dissipation theorem, however computationally it can be difficult to achieve high precision when

⁸Brownian motion, first theoretically described by Albert Einstein in 1905, and Johnson-Nyquist noise in electrical circuits, are both applications of the more general fluctuation-dissipation theorem.

the elastic material is inhomogeneous [75]. This is true in our case, where the high-reflectivity coating at the surface of the mirror locally damps the mechanical response of the whole cavity assembly. Additionally, we care about the response of X in a spatially localized region, coinciding with the extent of the cavity optical mode (which is $\approx 200 \,\mu\text{m}$). Thus, a normal mode decomposition requires summing many modes for adequate convergence.

Given these computational difficulties, is is standard practice to use finiteelement analysis (FEA) to measure the dissipation, W_{diss} , under a periodic drive delivered to the localized region sampled by our optical field. Specifically, one applies an oscillatory pressure

$$P(\vec{x},t) = F_0 \cos(\omega t) u(\vec{x}) \tag{3.17}$$

to the mirror surface, where F_0 is the amplitude and $u(\vec{x})$ is a gaussian "form factor" which weights the spatial response by the gaussian profile of the laser. The dissipated power W_{diss} , measured via FEA, can be related back to the transfer function $H(\omega)$ by

$$\mathbf{Im}\left[H(\omega)\right] = \frac{2W_{\text{diss}}}{F_0^2\omega} \,. \tag{3.18}$$

Thus, one obtains the power spectrum $G_X(\omega)$ by combining Eqs. 3.15 and 3.18 to get

$$G_X(\omega) = \frac{8k_B T}{\omega^2} \frac{W_{\text{diss}}}{F_0^2} \,. \tag{3.19}$$

In Ref. [61], Numata *et. al.* measure the mechanical quality factor of several common cavity materials by monitoring their mechanical ringdown after delivering an impulse to the material. They then use these results in a finite-element analysis

model to obtain the power spectrum of displacements of the cavity length via the fluctuation-dissipation theorem. Case No. 1 (Table 1 from Ref. [61]) corresponds closely to our our cavity, with ULE spacer and mirror substrate. They find the displacement noise spectral density, $S_x(f) \equiv \sqrt{G_X(2\pi f)}$, to be

$$S_X(f) \approx \frac{6 \times 10^{-17} \,\mathrm{m}/\sqrt{\mathrm{Hz}}}{f}$$
. (3.20)

Thus, for our 100 mm cavity, we can expect thermal fluctuations to contribute to the frequency noise spectral density of our clock laser at the order of $260 \text{ mHz}/\sqrt{\text{Hz}}$ at 1 Hz, falling off with 1/f at higher frequencies. Given the other limitations of our system, we do not expect to hit this thermal limitation to the laser linewidth.⁹

3.2.4 Servo noise & electronics

Having addressed perturbations to the cavity length, we turn to the remaining noise consideration: performance of the PDH lock itself. To understand the closedloop locking behavior of the laser, it is helpful to work in the Laplace domain, where $i\omega \rightarrow s$. Figure 3.10 shows a schematic of the PDH locking setup, where each block (the laser, cavity/PDH discriminator, and servo) has its own transfer function and we also allow each block to introduce additional noise.¹⁰

⁹Converting a frequency noise power spectral density to the traditional laser "linewidth" measured in an experiment is not straightforward. For a white noise power spectrum $|S_{\nu}(f)|^2 = h_0$, which is independent of f, one gets an analytic relationship for the Lorentzian linewidth $\Delta \nu_{\rm FWHM} = \pi h_0$ via the Wiener-Khintchine theorem. In the case of random-walk, or "red" noise, which is the case for the thermal fluctuations discussed here, $|S_{\nu}(f)|^2 = h_0 f^{-2}$ and the linewidth technically diverges. One then introduces a low-frequency cutoff, f_c , and obtains a Gaussian lineshape with FWHM = $\sqrt{(8 \ln 2)h_0 f_c}$. For a more thorough accounting, with approximate formulas for more complicated noise spectra, see the excellent discussion in Ref. [76] (and the references therein).

¹⁰This noise analysis follows the treatment of [70, 77, 78].



Figure 3.10: Block schematic representing the closed-loop behavior of a PDH lock. The laser (aka, the "plant" in control theory literature) outputs some signal $Y(s) = K \cdot U(s)$. Similarly, the PDH cavity acts as a frequency discriminator, converting the laser frequency Y(s) to an electrical error signal E(s) via $E(s) = D \cdot Y(s)$. Finally, the servo closes the loop by computing a feedback signal $U(s) = G \cdot E(s)$. Each stage in the feedback loop has a particular frequency-dependent gain (K, D, G), and might introduce some additional noise with spectral density S_K , S_D , S_G . See text for more details.

Analyzing around the loop, we find

$$Y(s) = K(s)U(s) + S_L \tag{3.21}$$

$$E(s) = D(s) [Y_0(s) - Y(s)] + S_D$$
(3.22)

$$U(s) = G(s)E(s) + S_G, (3.23)$$

where Y, E, and U are (respectively) the output, error, and control signals, and Y_0 is the "command" frequency to which we want to lock (given by the cavity resonance frequency ν_c). The transfer functions for each block are given by D (PDH/cavity discriminator), G (servo), and K (laser), and S_D , S_G , S_L are the associated noise spectral density contributions.

Solving for Y, we find

$$Y(s) = \left(\frac{KGD}{1 + KGD}\right)Y_0 + \frac{KGS_D + KS_G + S_L}{1 + KGD}.$$
(3.24)

The first term is the standard closed-loop transfer function — ie, Y tracks Y_0 as

$$H(s) \equiv \frac{Y}{Y_0} = \frac{G_T}{1 + G_T}, \qquad (3.25)$$

where $G_T(s) \equiv KGD$ is the (total) closed-loop gain. We see that for $G_T \gg 1, H \to 1$ and $Y \approx Y_0$ as desired. The second term in Eq. 3.24, however, is a noise term. This is more properly treated in quadrature, since each noise source is assumed to be uncorrelated. Thus, we have an expression for the noise given by

$$|S_n|^2 = \frac{|KGS_D|^2}{|1 + KGD|^2} + \frac{|KS_G|^2}{|1 + KGD|^2} + \frac{|S_L|^2}{|1 + KGD|^2}, \qquad (3.26)$$

where S_n is the total noise spectral density, related to the final laser linewidth by $\Delta \nu_L = \pi |S_n|^2$ (assuming a white-noise spectrum) [76,77]. We see that by making the servo gain G large enough, we can suppress noise fluctuations S_L of the laser by $1/G \to 0$ in the large-gain limit. Similarly, servo noise S_G is suppressed by $1/G \to 0$. However, noise introduced by the PDH frequency discriminant, S_D , goes as

$$\lim_{G \to \infty} \frac{|KGS_D|^2}{|1 + KGD|^2} = \frac{|S_D|^2}{|D|^2}.$$
(3.27)

Thus, as might be expected from the previous discussion, our final laser linewidth is dominated by the noise of our frequency discriminator — essentially, noise at the output of the discriminator is indistinguishable from noise in the laser itself [77]. Subsequent servo electronics try to compensate for this noise, but ultimately end up writing S_D back onto the laser output. The noise S_D itself has several sources — "fundamental" sources such as shot noise in the photodetector, technical noise in the photocurrent amplification and demodulation to generate the error signal, as well as noise associated with the "true" cavity resonance frequency (induced by vibrations, etc., addressed previously).

We now return to Equation 3.27 for the noise associated with the PDH frequency discriminant. We see there are two "knobs" for reducing the size of this term: we can increase the gain D, or decrease the noise spectral density S_D . Ultimately, the gain D will be limited by stability of the lock, since making it too large will drive the system into oscillation. Similarly, S_D cannot be made arbitrarily small, as we push against fundamental noise sources (like shot noise) which are physically beyond our control. To achieve optimal performance, then, we must ensure that S_D is limited by these fundamental noise sources and not technical noise. To state this another way, we want S_D to be shot-noise limited. For the remainder of this section, we discuss the details required to operate at the shot-noise limit.

The Pound-Drever-Hall locking scheme, described briefly above and shown



Figure 3.11: Schematic illustration of a PDH lock. Optical (phase-modulated) sidebands are imprinted on the laser light with an electro-optic modulator, and the promptly reflected light captured on a fast photodiode. By demodulating this photodiode signal at the phase modulation frequency, one recovers an error signal with a sharp slope around the cavity resonance. This can be used to feed back to the laser to stabilize its frequency to the cavity resonance.

schematically in Figure 3.11, involves the phase-modulation of a coherent laser field coupled into a Fabry-Perot cavity. We can write down the electric field, $E_{\rm inc}$, incident on the cavity as [30]

$$E_{\rm inc} = E_0 e^{i(\omega t + \beta \sin \Omega t)} \tag{3.28}$$

$$\approx E_0 \left[J_0(\beta) e^{i\omega t} + J_1(\beta) \left(e^{i(\omega + \Omega)t} - e^{i(\omega - \Omega)t} \right) \right] , \qquad (3.29)$$

where we have kept the first two terms of the Jacobi-Anger expansion. Here, E_0 is the amplitude of the electric field, β is a parameter describing the modulation depth, $\Omega/2\pi$ is the modulation frequency, and J_0 , J_1 are Bessel functions. To obtain the electric field reflected from the cavity, E_{ref} , we use the transfer function [79]

$$F(\omega) = \frac{E_{\text{ref}}}{E_{\text{inc}}} = \frac{r\left(1 - e^{i\omega/\Delta\nu_{\text{FSR}}}\right)}{1 - r^2 e^{i\omega/\Delta\nu_{\text{FSR}}}},$$
(3.30)

where r is field amplitude reflection coefficient for the mirror and $\Delta \nu_{\rm FSR}$ is the free spectral range, which for our cavity is 1.5 GHz.¹¹ Using Eq. 3.30, we can rewrite the reflected field amplitude as

$$E_{\rm ref} = E_0 e^{i\omega t} \left[J_0(\beta) F(\omega) + J_1(\beta) F(\omega + \Omega) e^{i\Omega t} - J_1(\beta) F(\omega - \Omega) e^{-i\Omega t} \right] .$$
(3.31)

¹¹The cavity, shown in the reduced CAD model of Fig. 3.2, is formed by a flat mirror and a mirror with a 500 mm radius of curvature, optically contacted onto a 100 mm spacer.

The reflected power, proportional to $|E_r|^2$, thus becomes

$$P_{\rm ref} = P_0 \left\{ |J_0 F(\omega)|^2 + |J_1 F(\omega + \Omega)|^2 + |J_1 F(\omega - \Omega)|^2 + J_0 J_1 \left[F(\omega) F^*(\omega + \Omega) e^{-i\Omega t} + F^*(\omega) F(\omega + \Omega) e^{i\Omega t} \right] - J_0 J_1 \left[F(\Omega) F^*(\omega - \Omega) e^{i\Omega t} + F^*(\omega) F(\omega - \Omega) e^{-i\Omega t} \right] + \left[2\Omega \text{ terms} \right] \right\},$$
(3.32)

where we use the shorthand $J_n(\beta) \to J_n$, and $P_0 \propto |E_0|^2$. Rewriting complex exponentials as sines and cosines and collecting like terms, expression (3.32) reduces to

$$P_{\rm ref} = P_0 \left\{ |J_0 F(\omega)|^2 + |J_1 F(\omega + \Omega)|^2 + |J_1 F(\omega - \Omega)|^2 + 2J_0 J_1 \operatorname{Re} \left\{ F^*(\omega) F(\omega + \Omega) - F^*(\omega) F(\omega - \Omega) \right\} \cos \Omega t - 2J_0 J_1 \operatorname{Im} \left\{ F^*(\omega) F(\omega + \Omega) + F^*(\omega) F(\omega - \Omega) \right\} \sin \Omega t + [2\Omega \text{ terms}] \right\}.$$
(3.33)

What is left is a DC term, two terms oscillating with Ω (formed by the beat between the carrier and each sideband), and terms oscillating at 2Ω formed by the sidebands beating with each other. You can see the two Ω terms plotted in Fig. 3.12.

At this point, we make the assumption that $\Delta \nu_{\rm FSR} \gg \Omega/2\pi \gg \delta \nu_c$, which is valid for our cavity (which has a linewidth $\delta \nu_c \sim 6 \,\mathrm{kHz}$, and $\Omega/2\pi \sim 20 \,\mathrm{MHz}$, $\Delta \nu_{\rm FSR} \sim 1.5 \,\mathrm{GHz}$). This lets us write $F(\omega \pm \Omega) \approx 1$ when the laser frequency $\nu = \omega/2\pi$ is near a cavity resonance ν_c (that is, when $\Delta \nu = \nu - \nu_c \rightarrow 0$). In that limit, the cosine term goes to zero while the sine term is maximized and antisymmetric around the cavity resonance — see Figure 3.12. The "lockable" error signal is the quadrature of $P_{\rm ref}$ in phase with $\sin \Omega t$, which can be extracted by



Figure 3.12: Both quadratures of the PDH error signal. We demodulate to obtain the signal in phase with $\sin \Omega t$, finding a steep discriminant at $\Delta \nu = 0$. This plot assumes $\Omega/2\pi = 0.2 \times \Delta \nu_{\rm FSR}$, and r = 0.99.

demodulating at frequency Ω and taking the in-phase component.

To better understand the near-resonance behavior of this error signal, we expand Eq. 3.30 about the resonance condition $\omega_c = 2\pi\nu_c$. To first order in $\Delta\nu$, we find

$$F(\omega)\Big|_{\omega\to\omega_c} \approx -2\pi i \frac{\Delta\nu}{\Delta\nu_{\rm FSR}} \frac{r}{1-r^2},$$
(3.34)

with r the field amplitude reflection coefficient. The cavity finesse, given by

$$\mathcal{F} = \frac{\Delta\nu_{\rm FSR}}{\delta\nu_c} \approx \frac{\pi r}{1 - r^2} \,, \tag{3.35}$$

can be used to reduce Eq. 3.34 to

$$F(\omega)\Big|_{\omega\to\omega_c} \approx -2i\frac{\Delta\nu}{\delta\nu_c}$$
. (3.36)

Plugging this into the sine term of Eq. 3.33 and again taking $F(\omega \pm \Omega) \approx 1$, we arrive at the linearized error signal ϵ , related to $\Delta \nu$ by a cavity transfer function $D_{\rm PDH}$ according to

$$\epsilon = D_{\rm PDH} \Delta \nu \tag{3.37}$$

$$D_{\rm PDH} = \frac{8J_0(\beta)J_1(\beta)}{\delta\nu_c}P_0 \tag{3.38}$$

where again $\delta \nu_c$ is the cavity linewidth and P_0 is the power incident on the cavity.

Returning to our earlier discussion, the dominant technical noise contribution to the final laser linewidth is given by the noise introduced at the PDH discriminant, which goes as $|S_D|^2/|D|^2$. To minimize this term, we want high PDH gain D and low spectral noise density S_D . From Eq. 3.38, we can maximize D by maximizing the coefficient $J_0(\beta)J_1(\beta)$ (which occurs for $\beta = 1.08$) and by increasing the power P_0 incident on the cavity. Clearly narrower cavity linewidths also correspond to higher gain in the PDH transfer function, but this parameter is already fixed by the ULE cavity. Ideally, we want to keep P_0 low because the power buildup in a highfinesse cavity can start to approach the damage threshold for the mirror coatings, and additionally cause heating-induced length fluctuations which negatively impact the stability of the resonator.

Practically speaking, the error signal given in Eq. 3.37 has units of optical power and must undergo further processing in order to obtain a useful signal to pass into the servo block G. Specifically, we must convert ϵ into a voltage error signal E, using a high-speed photodiode and transimpedance amplifier. For simplicity, we use a fixed-gain amplified photodiode (Thorlabs PDA10A).¹² In a locked configuration, $F(\omega_c) = 0$ and from Eq. 3.32 the average power $P_{\rm ref}$ reflected from the cavity is

$$P_{\rm ref} \approx 2P_0 J_1(\beta)^2 \tag{3.39}$$

arising from the two promptly reflected sidebands at $\omega_c \pm \Omega$. This DC optical power gives rise to a fundamental shot noise spectral density in the photocurrent given by [77]

$$S_{\rm sn} = \sqrt{2}\sqrt{2eI}\,,\tag{3.40}$$

where e is the electron charge and I is the photocurrent. The extra factor of $\sqrt{2}$ arises from the demodulation required to generate an error signal, which selects shot noise contributions at both $+\Omega$ and $-\Omega$ fourier frequencies.

The optical power can be converted to a photocurrent via the diode responsiv-

¹²Another possibility, e.g. as done in Ref. [80], would be to build a custom transimpedance amplifier with a nice photodiode such as the S5971 or G4176-03 from Hamamatsu. One could then optimize the signal-to-noise ratio in the desired modulation band. Analog devices has a very nice tool to aid in such designs, and can be found at http://www.analog.com/designtools/en/photodiode/.

ity — for Si-based photodiodes near 698 nm, the responsivity is $R_{\rm PD} \approx 0.45 \,\mathrm{A}\,\mathrm{W}^{-1}$. Combining Equations 3.39 and 3.40 yields

$$S_{\rm sn} = \sqrt{8eR_{\rm pd}J_1(\beta)^2 P_0}$$
. (3.41)

This photocurrent noise spectral density represents the shot-noise limit for a particular input power; in convenient units, it is roughly $S_{\rm sn} \approx 35 \,\mathrm{pA}/\sqrt{\mathrm{Hz}}$ for $1\,\mu\mathrm{W}$ input power, and scales as the square root of optical power. The datasheet for the PDA10A photodiode gives a noise-equivalent power (NEP) of $3.5 \times 10^{-11} \,\mathrm{W}/\sqrt{\mathrm{Hz}}$. This is specified at the peak of the diode's optical response, which happens to be near our wavelength of interest.

It's worth noting that, since setting up the clock laser locking electronics, Thorlabs has discontinued the PDA10A in favor of an upgraded version, the PDA10A2, which features a NEP of $5 \times 10^{-14} \text{ W}/\sqrt{\text{Hz}}$, nearly three orders of magnitude smaller than the PDA10A. One could also use the DET36A2, which is slightly slower (25 MHz bandwidth versus 350 MHz) but has a NEP of only $1.6 \times 10^{-14} \text{ W}/\sqrt{\text{Hz}}$.

The above derivations can be extended to allow for a very versatile, tunableoffset PDH lock [81]. The idea, which we will employ in our system (see the final setup in Fig. 3.13), involves adding a second modulation frequency Ω_{offset} that can be tuned to provide an arbitrary frequency difference between the cavity resonance ν_c and the laser frequency. This is critical, since we have no way to tune the resonance frequency of the ULE cavity and need to bring the laser into resonance with the atomic transition. The incident electric field, then, can be expressed as

$$E_{\rm inc} = E_0 \ e^{i(\omega t + \beta \sin \Omega t + \beta_{\rm offset} \sin \Omega_{\rm offset} t)} \tag{3.42}$$

which contains sidebands at Ω , Ω_{offset} , and $\Omega_{\text{offset}} \pm \Omega$. Requiring $\Omega_{\text{offset}} \gg \Omega$, this expression can be reduced via the same optical transfer function (Eq. 3.30) and simplifications above to give the dual-sideband frequency discriminant, D_{DSB} , as [81]

$$D_{\rm DSB} = J_1(\beta_{\rm offset})^2 \ \frac{8J_0(\beta)J_1(\beta)}{\delta\nu_c}P_0 \tag{3.43}$$

$$= J_1(\beta_{\text{offset}})^2 D_{\text{PDH}}. \qquad (3.44)$$

The intuition here is we are using the modulation frequency Ω_{offset} to generate a sideband which can be tuned onto resonance with the cavity; this sideband in turn has the same PDH sidebands at $\Omega_{\text{offset}} \pm \Omega$ as in the standard PDH setup, and can be used to generate an error signal as before. However, the amplitude of that error signal is decreased by the amount of power we can push into the first-order sideband at Ω_{offset} . It is maximized for a modulation depth $\beta_{\text{offset}} = 1.84$, but is roughly one-third the size of the standard PDH error signal for equivalent laser power. However, the degraded signal-to-noise of the dual-sideband PDH can be compensated for by increasing the optical power P_0 . One must also be careful using this technique, since the additional sidebands can generate spurious features at frequencies far from the desired lockpoint, but with a wavemeter and a little care any confusion can be avoided.

3.3 Final design

The initial layout of this laser included a pre-stabilization cavity, which narrowed the laser linewidth to ≤ 10 kHz before being sent to the high-finesse ULE cavity. The rationale for this was twofold: first, a spectrally narrowed light source would reduce the demands for the second-stage servo responsible for the PDH lock to the ULE cavity, and second, during initial alignment, it would be helpful to see a well-resolved PDH error signal on the ULE cavity. In addition, the final lock to the high-finesse cavity would be done by feeding back to a voltage-controlled-



Figure 3.13: Final layout of the clock laser system. An earlier iteration of the layout included a two stage locking scheme, whereby the laser was locked to a home-built pre-stabilization cavity (finesse $\mathscr{F} \approx 1200$). The final lock to the ULE cavity was achieved via feedback to a voltage-controlled oscillator (VCO), which controlled an AOM to remove residual phase noise from the pre-stabilized light. At longer timescales, the length of the pre-stabilization cavity was served with a piezo such that the VCO control voltage stayed close to its central value. This locking scheme was abandoned in favor of the one presented above for reasons described in the text. Currently, a single Pound-Drever-Hall lock stabilizes the laser to a passive Fabry-Perót cavity made from ULE. The cavity is mounted in such a way to minimize sensitivity to mechanical vibrations, held in a temperaturestabilized vacuum canister, and rests on a vibration isolation stage (BM-8 from Minus-K). The entire ULE breadboard is enclosed in an acoustically damped box to further minimize mechanical noise. Because the ULE cavity resonances do not match up with the atomic transition frequencies, a broadband fiber EOM (Jenoptik, PM705) is used to span the difference; we do not temperature stabilize this device or perform active fiber phase-noise cancellation because of the shortness of the fiber pigtails, but this might be necessary in the future. A variable RF frequency is summed on an RF power splitter/combiner with the 18 MHz dither used to generate the PDH error signal; both tones are then amplified and sent to the EOM. The optical sideband generated by the variable RF drive is made to be resonant with the cavity, and the resulting PDH error signal is passed through the Toptica FALC analog PID and fed back to the current of the laser diode to stabilize its frequency. A final AOM is used for fine frequency scans of the clock light, and for precise pulse shaping. The 1^{st} (+1) order of this AOM is fiber-coupled and sent to the experiment; in the future, it may be necessary to include fiber phase noise cancellation via this AOM. The 0th order is fiber-coupled and sent to the wavemeter for monitoring, and a polarizing beam splitter picks off a few mW before the AOM to be counted against an optical frequency comb.

oscillator (VCO), which has the added advantage of linearity and a lack of mode hops. This was the original design of the clock laser system at JILA [82], which was used successfully for a number of years.

The pre-stabilization cavity was used for the initial alignment and evaluation of the laser system at 689 nm. It consisted of a homebuilt invar spacer (length L = 100 mm) with backside polished mirrors and a custom high-reflectivity line coating from CVI at 698 nm.¹³ The mirrors form a modified hemispherical cavity (the curvatures are $r_1 = 500 \text{ mm}$, $r_2 = \infty$, while the length is L = 100 mm). We measured the finesse to be $\mathscr{F} \approx 1225$ at 689 nm, which yields a cavity linewidth $\gamma_c \approx 1.2 \text{ MHz}$. With a good servo, one can reasonably expect to split the cavity line by a factor of ~ 100 or more, which gets us to something on the order of the ULE cavity linewidth. The pre-stabilization cavity is hermetically sealed in a custom box machined from aluminum, and three Peltier coolers stabilize the temperature of the enclosure.

We ultimately chose to abandon the pre-stabilization cavity, again for several reasons. First, the additional complexity of the servo means more things have the possibility to come unlocked.¹⁴ The linearity and lack of hysteresis in the VCO/AOM frequency control was nice, but even in double-passed configuration it is difficult to completely remove amplitude modulation coupled to changes in drive frequency. We attempted to compensate for this with an intensity lock, but it was unclear if the bandwidth was high enough.

Second, because the laser cooling group had purchased a Toptica DigiLock servo a few years back that was sitting on the shelf, we opted to only buy a single Toptica FALC servo and use the DigiLock for the (nominally less demanding)

¹³The mirrors were quoted with a reflectance $R \gtrsim 0.995$ at 698 nm, but in private communication with CVI, they expect the coating to perform ~ 0.999 even though it could not be guaranteed. The finesse we measured at 689 nm corresponds to a reflectance of 0.997, which seems to be consistent with this statement.

¹⁴And in an ultracold strontium lab, the number of servos is already too damn high!

second-stage lock to the ULE cavity. This digital servo, in retrospect, was a poor choice. The bandwidth performance was never fantastic, the control software was buggy, and it required a computer to be plugged in to the module via USB to relock when needed. We were able to achieve a passable lock with the 689 nm diode used to characterize the system, but when we switched to the 698 nm diode, we could not push the bandwidth larger than ~ 100 kHz without oscillations.

Our experience with the DigiLock, compared to the FALC, made us rethink the wisdom of a pre-stabilization cavity entirely. We had successfully operated a direct PDH lock to an identical cavity for the 689 nm laser systems without trouble, and in conversations with the Menlo field technicians, they hinted that very few labs use pre-stabilization cavities anymore. The reduced complexity meant more laser power for the experiment, which to drive the clock transition in the bosonic isotopes was initially a real concern (see discussion in 5.2.1.2), and after switching to a direct lock with the FALC, the laser would stay locked for several hours at a time compared to the tens of minutes to an hour in the previous scheme.

3.4 Performance

We analyze the performance of the clock laser in two separate ways. First, during construction and initial coupling into the ULE cavity, we used a laser at 689 nm locked to a pre-stabilization cavity¹⁵ and then to the ULE cavity via feedback to a double-passed AOM. This 689 nm laser could then be measured against our master 689 nm laser used for the narrow-line rMOT. We generate a heterodyne beatnote between the two systems, and measure a joint FWHM Lorentzian linewidth of 289(14) Hz, as shown in Fig. 3.14. This puts an upper bound on the linewidth of

¹⁵This cavity, made from an Invar spacer with a piezo-actuator, was originally planned for the final 698 nm system; however, better lock stability was achieved by locking directly to the ULE cavity via PDH. In aligning to a very high finesse cavity, however, it was useful to have the pre-narrowed laser as this allowed us to see an actual error signal as we swept across the cavity resonance.



Figure 3.14: Heterodyne beatnote between two independent lasers at 689 nm. One laser is locked to the "clock" cavity, while the other laser is the master oscillator for the narrow-line rMOT. The joint FWHM Lorentzian linewidth is 289(14) Hz (solid orange line). This is likely dominated by the linewidth of the less well-isolated Fabry Pérot cavity used to lock the rMOT master laser, however from this beatnote measurement one cannot disentangle the relative contributions from each laser.

our 698 nm system, and is likely dominated by the linewidth of the less well-isolated 689 nm Fabry-Pérot cavity.

In addition to the heterodyne beatnote at 689 nm, we can also use the atoms themselves to measure the performance of the laser. Unfortunately, we do not have a magic wavelength dipole trap on this apparatus (see Chapter 6 for discussion of the next-generation system, which will have magic wavelength lattices designed in from the beginning). The resulting lineshape becomes temperature-dependent, and we can use the framework from Ref. [83] to estimate the width of a carrier transition in a deep, non-magic wavelength lattice. Specifically, the full-width-at-half-maximum (FWHM) Γ_c is given by

$$\Gamma_c \approx T\left(\frac{k_B}{h}\right) \frac{\left|\sqrt{\alpha'/\alpha} - 1\right|}{0.295},$$
(3.45)



Figure 3.15: Atomic spectroscopy of the clock transition in ⁸⁷Sr (left) and a ⁸⁴Sr BEC (right). For the ⁸⁷Sr spectroscopy, each point is the mean of five measurements and the error bar taken as the standard deviation; the fit gives a reduced $\chi^2/20 = 0.99$. The FWHM are displayed, and correspond to a 1/e width of 36(2) kHz and 4.4(5) kHz respectively.

where h, k_b are respectively Planck's and Boltzmann's constant, T is the atom temperature, and α'/α is the ratio of polarizabilities in the excited and ground states. At 1064 nm, this polarizability ratio is ≈ 0.7 . This translates to a thermally broadened carrier linewidth of roughly 11.5 kHz µK⁻¹. This is much larger than the ≤ 1 Hz linewidths one expects from an optical clock transition.

Despite this difficulty achieving a truly narrow clock lineshape, we show some spectroscopic loss features out of a dipole trap in Fig. 3.15. Clearly, the spectra are inhomogeneously broadened beyond the ≤ 200 Hz linewidth we might expect from the laser; in the case of 87 Sr, we see a thermal lineshape with a width ≈ 40 kHz. In the case of the 84 Sr BEC, the lineshape is much more complicated but its characteristic width is set by the chemical potential. The narrowest features we observed were in a BEC, on the order of 4 kHz to 5 kHz.

Chapter 4: Publication: An ultra-low noise, high-voltage piezo driver

Over the course of my PhD, I had the opportunity to work on several electronics design projects to enhance the instrumentation and control of our experiment. Here, I include the pre-print arXiv version of a paper that came out of one such project, to create a low-noise, high-bandwidth piezoelectric driver well-suited for controlling ECDL lasers. In conjunction with the several improvements and modifications made to the Durfee/Hall/Librecht current controller and the optical phaselocked-loop beatnote servo (both described previously in Sec. 2.5), this piezo driver completes a set of high-performance laser control electronics that have afforded a great deal of versatility when working with narrow-line transitions in strontium.

The published version of this chapter appears in Review of Scientific Instruments 87, 124702 (2016) [48].¹ What appears below has been edited and reformatted for increased clarity, and includes an additional section at the end discussing aspects which did not make it into the published version, and future improvements that could be made to the design and its firmware.

4.1 Abstract

We present an ultra-low noise, high-voltage driver suited for use with piezoelectric actuators and other low-current applications. The architecture uses a flyback

¹The design was created in discussion and collaboration with Alessandro Restelli, the JQI electronics guru, to whom I am indebted for much of my electronics expertise. I performed the PCB layout and collected the reference data (with occasional assistance from Alessandro), and wrote the manuscript. All authors contributed critical edits to the manuscript, and provided insights and support throughout the process.

switching regulator to generate up to 250V in our current design, with an output of 1 kV or more possible with small modifications. A high slew-rate op-amp suppresses the residual switching noise, yielding a total RMS noise of $\approx 100 \,\mu\text{V}$ (1 Hz–100 kHz). A low-voltage (±10 V), high bandwidth signal can be summed with unity gain directly onto the output, making the driver well-suited for closed-loop feedback applications. Digital control enables both repeatable setpoints and sophisticated control logic, and the circuit consumes less than 150 mA at ±15 V.

4.2 Introduction

Many instrumentation applications in the modern laboratory require agile, low-noise voltage sources capable of supplying hundreds of volts or more. For example, piezo-actuated mirrors and diffraction gratings play an important role in atomic physics experiments (used, e.g., in Fabry-Pérot cavities [84, 85] and external-cavity diode lasers [32]), while avalanche photodiodes and photomultiplier tubes require large bias voltages for proper operation. In the realm of biophysics, electrokinetic separation methods such as free-flow or capillary electrophoresis [86] require large electric field gradients, and the recent push to develop lab-on-a-chip devices could benefit from miniaturized high-voltage sources. [87]

Laboratory devices are often operated in a closed feedback loop, where small voltage changes on top of a large DC voltage are necessary to stabilize the output of a particular system. For example, the frequency of an extended-cavity diode laser can be locked by feeding back to a piezo-actuated diffraction grating or mirror, which in turn supplies optical feedback to the diode. Commercially available piezoelectric drivers typically provide a modulation input for such closed-loop applications, but the input voltage is often gained such that it spans the entire output range of the device. Other designs separate high- and low-voltage control pathways, which can extend the bandwidth to \approx MHz, but the low-voltage control is AC-coupled

to the output. [88] While these designs have advantages, many applications would benefit from an architecture that provides a unity-gain, DC-coupled feedback path to the high-voltage output. This low-gain modulation input could make closed-loop systems less susceptible to noise contributions from the servo controller, which we often find in our laboratory to be a limiting factor in laser lock stability.

Instrumentation electronics capable of supplying high voltages traditionally fall under one of two architectural umbrellas: DC-DC switching converters, and linear-type amplifiers. While DC-DC converters are efficient and can work at very high voltages, they suffer from switching noise and limited control bandwidths. Linear-type devices are typically constructed from a high-voltage operational amplifier (op-amp), powered either from a high-voltage linear regulator or more typically from a secondary switching converter. While the op-amp can provide 100 dB or more of power-supply noise rejection [89], linear regulators must handle any excess voltage by dissipating heat and so may be more cumbersome to deploy in the laboratory.

We present a circuit with a hybrid architecture. The high voltage is generated by a galvanically isolated DC-DC converter, while a low-noise, high-slewrate op-amp simultaneously removes noise at the output and provides a low-gain, high-bandwidth ($\gtrsim 100 \text{ kHz}$) modulation input for closed-loop feedback applications. This architecture is able to achieve extremely low noise ($\approx 100 \text{ µV}_{\text{RMS}}$) over the entire output range, draws very little current, and fits comfortably onto a smallfootprint printed circuit board (PCB). Additionally, the high-voltage output remains single-ended and referenced to ground, allowing it to drive piezo actuators with a grounded terminal. The schematic is presented in Sec. 4.3, with a noise analysis in Sec. 4.4 and characteristic performance data in Sec. 4.5. Complete design files, including the schematic, bill of materials, and board layout, can be found on



Figure 4.1: Schematic of the high voltage supply and stabilization. The voltage $V_{\rm HV}$ is generated using a Texas Instruments DRV2700 high voltage driver in flyback configuration. A high slew-rate op-amp (U2) senses the output voltage across R_1 and R_2 , and controls it by modulating the midground node $V_{\rm MG}$. The DC control signal for this op-amp, $V_{\rm ctl}$, is supplied by a digital-to-analog (DAC) converter, which is passed through a switchable low-pass filter. This design allows for very heavy filtering of the DAC 1/f noise during steady-state operation, but the corner frequency can be increased if the output needs to be scanned more quickly. The $V_{\rm ctl}$ gain is set by $(1 + R_1/R_2)$, while the modulation gain is set by $-R_1/R_{\rm mod}$. The op-amp U2 removes residual switching noise and stabilizes the DC output according to the transfer function given in Eq. (4.3). A MOSFET quench circuit, shown in Fig. 4.2, connects at nodes $V_{\rm out}$ and $V_{\rm MG}$.

GitHub.² The board manufacture and component cost is less than \$200, making it a cost-effective alternative to commercial options.

4.3 Circuit Design

The design principles discussed below show how we leverage the characteristics of a galvanically-isolated switching regulator without sacrificing the low-noise requirements of many laboratory applications. Our design targets a 250 V output, but we discuss straightforward modifications to the schematic that make it possible to tailor the gain and output range to a specific application. The entire electronics package fits into a compact Eurocard rack module (with the high-voltage section taking only a fraction of the PCB), and draws less than 150 mA at 15 V. The highvoltage output current will be limited by the switching regulator and by the LM7171

²https://github.com/JQIamo/hv-piezo-driver

op-amp used for low-noise stabilization (U2 in Fig. 4.1, which can supply at most $\approx 100 \text{ mA} [90]$), but is sufficient for nearly all piezoelectric applications.

Fig. 4.1 shows an overview of the circuit schematic. A flyback regulator (Sec. 4.3.1) controls the potential between the high-voltage (V_{HV}) and floatingground (V_{FG}) circuit nodes, while the low-noise stabilization circuitry (Sec. 4.3.2) controls the output node V_{out} relative to the true circuit ground. A digital-to-analog converter (DAC) generates a voltage setpoint, V_{ctl}, which is sent to the high-voltage flyback regulator and to the low-noise stabilization circuit. The DAC is controlled by an integrated microcontroller, and can be programmed to output slow (≈ 10 Hz) rail-to-rail voltage ramps in addition to setting the DC operating point (Sec. 4.3.3). To improve the large-amplitude slew rate, a MOSFET "quench" circuit (Fig. 4.2) is included to reduce the *RC* time constant of the high-voltage node V_{out} when needed. A low-pass filter with a switchable corner frequency can be engaged during DC operation to reduce 1/f noise from the DAC (discussed in Sec. 4.4). Fast output modulation between ± 10 V can be achieved through the input node V_{mod}. This node is DC-coupled to the high-voltage output, and is useful for closed-loop feedback control.

4.3.1 Flyback regulator

The high-voltage DC-DC converter used here is based on the Texas Instruments DRV2700 piezo driver.³ This integrated circuit can be operated as a boost converter to drive an on-chip differential amplifier up to 100 V, or as a flyback converter up to 1 kV or more. In flyback configuration, the internal-boost switch of the DRV2700 (pin SW in Fig. 4.1) drives a step-up transformer. When the switch closes, current begins to flow through the primary coil of the transformer and induces a

³The identification of commercial products in this paper is for information only and does not imply recommendation or endorsement by the National Institute of Standards and Technology.

corresponding voltage across the secondary coil. In this state, the output diode D1 is reverse-biased, and the capacitor (C_{HV} in Fig. 4.1) holds its charge. When the switch opens, the voltage across the secondary coil is inverted, putting the diode into conduction and charging the capacitor. By changing the switching duty cycle, the DRV2700 is able to regulate the voltage across the galvanically isolated output (nodes V_{FG} and V_{HV} in Fig. 4.1).

The DRV2700 implements output voltage control by comparing the feedback input pin at node V_{FB} with an internal (1.3 V) reference. The resistors R_9 and R_{10} are chosen such that pin FB is at 1.3 V when the output of U1 is at ground: $R_{10}/(R_9 + R_{10}) = 1.3 \text{ V}/5 \text{ V} \approx 0.26$. The op-amp U1 subtracts V_{FG} and $G \cdot V_{ctl}$ from the voltage at node V_{HV} , ensuring the DRV2700 regulates the output voltage such that

$$V_{\rm HV} - V_{\rm FG} = G \cdot V_{\rm ctl} \,, \tag{4.1}$$

where the gain G is set by the resistor ratio $R_3/R_4 \equiv R_5/R_6$, and V_{ctl} is the control voltage set by the DAC. The capacitors C_3 and C_4 are chosen such that $C_3 = 22 \text{ pF}$ and

$$\frac{C_4}{C_3} = \frac{R_3}{R_4 \mid\mid R_7}, \qquad (4.2)$$

as suggested by the DRV2700 datasheet [91], where $R_6 = R_7 = R_8$, and $R_i \parallel R_j \equiv R_i R_j / (R_i + R_j)$. In our implementation, we choose a gain $G \approx 50$ ($R_3 = R_5 = 499 \,\mathrm{k\Omega}$; $R_4 = R_{6,7,8} = 10 \,\mathrm{k\Omega}$), allowing a 5 V control signal V_{ctl} to span 250 V at the output. A different DAC and/or a different gain factor could be chosen to adjust the maximum output voltage.⁴ The transformer (ATB3225, 1:10 step-up winding), diode, and RC feedback network are all based on values suggested in the DRV2700

⁴Note that the passive components must be rated for the chosen output voltage.



Figure 4.2: MOSFET "quench" circuit. When the mid-ground node V_{MG} (also shown in Fig. 4.1) goes below the threshold value set at V_{TH} , the op-amp puts the HV MOSFET Q1 into conduction. When engaged, the quench time constant is given by $\tau \approx R_{MOS}C_{out}$, which for our circuit is set to 1 ms. The capacitor $C_{out} = 1 \,\mu\text{F}$ is shown in Fig. 4.1, and details of this circuit are discussed in Sec. 4.3.3.

datasheet. [91,92]

The output of the flyback regulator is passed through a four-pole, low-pass RC filter. The corner frequency $f_c \approx 30 \,\text{Hz}$ is chosen to be high enough that a slow ($\approx 10 \,\text{Hz}$) rail-to-rail triangle ramp can be applied by the DAC at V_{ctl} (for, e.g., sweeping over a resonance in spectroscopy), but low enough that the $\approx 100 \,\text{kHz}$ switching noise is substantially attenuated. Additional capacitors on both the V_{HV} and V_{FG} resistor networks shunt high frequency noise to ground. This filter topology, modeled on a lossy transmission line, is sufficient for our application, but other corner frequencies or topologies could also be used.

4.3.2 Low-noise stabilization and fast modulation

The low-noise stabilization circuit is crucial to the performance of the design, as it is responsible for removing noise at the output of the flyback converter. To
accomplish this, a high slew-rate op-amp (Texas Instruments LM7171, 4100 V/ μ s; see U2 in Fig. 4.1) drives the galvanically isolated floating ground of the flyback converter. This op-amp senses the voltage V_{out} referenced to true circuit ground, and adjusts its output such that

$$\mathbf{V}_{\text{out}} = \left(1 + \frac{R_1}{R_2 \mid\mid R_{\text{mod}}}\right) \mathbf{V}_{\text{ctl}} - \left(\frac{R_1}{R_{\text{mod}}}\right) \mathbf{V}_{\text{mod}} \,. \tag{4.3}$$

Here, V_{mod} is the applied modulation, which can vary between $\pm 10 \text{ V}$ and is inverted before being summed onto the output. We choose $R_1 = R_{mod} = 1 \text{ M}\Omega$ and $R_2 = 20.5 \text{ k}\Omega$ such that the DC gain $\Delta V_{out}/\Delta V_{ctl}$ is ≈ 50 and the modulation gain $\Delta V_{out}/\Delta V_{mod}$ is unity. Depending on the application, other gain configurations could work equally well provided the non-inverting gain of U2 closely matches the gain of the flyback regulator (since they both derive from V_{ctl}).

The op-amp U2 controls V_{out} via two different feedback pathways. At low frequencies, it modifies the floating ground reference V_{FG} of the flyback converter, and the DRV2700 in turn modifies V_{HV} to maintain a constant voltage between V_{FG} and V_{HV} . At higher frequencies, U2 is decoupled from V_{FG} by the passive filtering network. In this regime, C_{out} provides a low-impedance path between U2 and the output, such that high-frequency switching noise can be directly compensated for by the op-amp. We chose a value $C_{out} = 1 \,\mu\text{F}$, which is a compromise between component size and the desire for a large capacitance. In addition, a small resistance $R_{11} = 50 \,\Omega$ is inserted between U2 and C_{out} to ensure stable operation. Smaller R_{11} and/or larger C_{out} might provide better performance, but this has not been tested.

The choice of components for resistors R_1 and R_2 is crucial for the low-noise performance of the system. Because this resistive divider is responsible for accurately sensing the voltage V_{out} , noise introduced by these resistors cannot be corrected by the op-amp. In general, resistors are fundamentally limited by Johnson noise, in which thermal fluctuations contribute to a white noise power spectrum given by $4k_BTR$, where T is the temperature and k_B is Boltzmann's constant. [93] However, resistors also exhibit 1/f current noise caused by equilibrium fluctuations of the resistance. [94,95] This excess noise depends on the applied voltage, and therefore is an important consideration in a high-voltage circuit. It is also highly dependent on the resistor composition and varies from manufacturer to manufacturer. Seifert, et. al. [96] characterized 1/f noise in a variety of resistors, and found that the Vishay TNPW 0.1%-series resistors showed a noise spectrum almost consistent with Johnson noise down to 1 Hz. Our current design uses this series in a 1206 package, but we noticed low-frequency noise correlated with varying strain on the PCB, potentially due to the relatively large footprint of this package. Future boards might instead use three TNPW 0.1% 0603 resistors in series for both R1 and R2 to minimize strain-induced output noise.

The value of capacitor C_1 is a tradeoff between two competing design considerations. On the one hand, a larger C_1 extends the frequency range where switching noise from the DRV2700 is suppressed. However, large values of C_1 limit the bandwith of V_{ctl} . We empirically settled on $C_1 = 1 \text{ nF}$, which is large enough to saturate the feedback gain in the 40 kHz–100 kHz range where switching noise dominates, but not so large that it limits the bandwidth of V_{ctl} below the corner set by the switchable low-pass filter described in Sec. 4.3.3. Once C_1 was chosen, capacitor C_{mod} was calculated to match the impedances $R_1 \parallel C_1 = R_{mod} \parallel C_{mod}$. For our circuit, this means $C_{mod} = C_1$.

4.3.3 Digital control and slow modulation

The high-voltage setpoint (absent voltages summed in at V_{mod}) is controlled by a low-noise DAC. This has two advantages: digital control enhances setpoint repeatability, and simplifies the integration with computerized control electronics or sophisticated servo loops. While the modulation input V_{mod} has a limited range of $\pm 10 \text{ V}$, larger voltage swings can be achieved by reprogramming the DAC.

As discussed below in Sec. 4.4, without modification the DAC would dominate the noise performance of V_{out} . Therefore we add a single-pole, low-pass filter between V_{ctl} and the non-inverting node of U2 to bring the DAC noise contribution below other noise sources in the circuit. This filter has a switchable corner frequency (between 165 Hz and 0.8 Hz) to optimize noise performance at DC while still permitting AC modulation when needed. It consists of a 20.5 k Ω resistor and 47 nF capacitor, with a secondary 10 µF capacitor that can be switched in to operate with the lower corner frequency.

One downside of the flyback regulator presented above is that while the switched transformer can quickly charge the output capacitors, the discharge time τ is limited to ≈ 1 s by the RC time constant of the circuit. To get around this limitation, we have added an auxiliary MOSFET "quench" circuit⁵ to quickly shunt V_{out} to ground (see Fig. 4.2). This circuit works by monitoring the voltage at V_{MG}, the mid-ground node controlled by op-amp U2. If V_{MG} drops below a threshold set by V_{TH}, the comparator op-amp in Fig. 4.2 changes the gate voltage of the MOSFET to put it into conduction. The time constant for this configuration is given by $\tau \approx R_{MOS}C_{out}$. For our circuit, this changes τ to ≈ 1 ms, allowing C_{out} to be quickly discharged. The threshold is V_{TH} = -10.4 V, but other values could be chosen depending on the design requirements.

The high-voltage design presented here has the flexibility to exist as a standalone circuit or be integrated with other electronics, and we have included several auxiliary features to make this more convenient. For example, the analog modulation input is differentially buffered to break ground loops (not shown in Fig. 4.1),

⁵The quench circuit presented here is based on the pulldown FET discussed in the DRV2700 datasheet, with some additional modifications to suit our purposes.

and the digital portion can be interfaced with other devices in the lab to expand conceivable control scenarios. Secondary features include a divided-down output that can be used as a monitor or fed forward to a low-noise laser diode current controller like the one in Erickson, et. al. [41] Of course many variations are possible, and we refer the reader to our GitHub page for more details on our specific implementation.⁶

4.4 Noise model & analysis

To better understand the circuit performance and the measured output noise reported in Sec. 4.5, we introduce the noise model shown in Fig. 4.3. A summary of each noise contribution (op-amp, DAC, Johnson-Nyquist, and residual ripple from the DRV2700, all calculated at the node V_{out}) is shown in Fig. 4.4, along with the cumulative root-mean-square (RMS) noise estimates in different frequency bands. We will neglect noise appearing at node V_{mod} due to the external modulation because its exact character depends on the external drive.

To facilitate the noise analysis, we calculate the voltage and current (transimpedance) gains from the input nodes of U2 to the output, V_{out} . Starting with the non-inverting node, we find the voltage gain to be

$$G_{+}^{(v)} = 1 + Z_1 \left(\frac{1}{R_2} + \frac{1}{Z_{\text{mod}}}\right),$$
 (4.4)

where Z_1 and Z_{mod} are the equivalent impedances of $R_1 \parallel C_1$ and $R_{\text{mod}} \parallel C_{\text{mod}}$, respectively. The transimpedance gain, $G_+^{(i)}$, follows by multiplying $G_+^{(v)}$ by the impedance, Z_+ , seen from that node. Thus,

$$G_{+}^{(i)} = G_{+}^{(v)} \left[\frac{R_{\rm LP}}{1 + 2\pi i f R_{\rm LP} C_{\rm LP}} \right], \qquad (4.5)$$

⁶https://github.com/JQIamo/hv-piezo-driver



Figure 4.3: Noise model, including contributions from the op-amp, DAC, DRV2700, and modulation input. The element labeled C_{LP} consists of a 47 nF capacitor in parallel with a 10 µF capacitor connected to ground through a switch, such that the corner frequency of the filter can be changed depending on the mode of operation (see text). Though not drawn in the figure, we also consider the Johnson noise contributions from all resistors.

where the bracketed term is Z_+ , f is the Fourier frequency, and i is the imaginary unit.

We now calculate gains from the inverting node of U2. Because any currents appearing at this node will be cancelled by the feedback of U2, the transimpedance gain $G_{-}^{(i)}$ is simply the impedance Z_1 , given by

$$G_{-}^{(i)} = Z_1 \equiv \frac{R_1}{1 + 2\pi i f R_1 C_1} \,. \tag{4.6}$$

With these expressions in hand, we can calculate the output noise contribution from each source in our model.

As shown in Fig. 4.3, the op-amp noise is parametrized by two noise contributions: e_n , the input voltage noise spectral density, and i_n , the input current noise spectral density. For the LM7171 at 10 kHz, $e_n = 14 \text{ nV}/\sqrt{\text{Hz}}$ and $i_n = 1.5 \text{ pA}/\sqrt{\text{Hz}}$ with a 1/f noise character below this frequency. [90] The voltage noise appears at the non-inverting input, while the current noise is present at both inputs. By multiplying each source by the appropriate gain, we obtain the equivalent output-noise power spectral densities (PSD),

$$e_{v,U2}^{2} = |G_{+}^{(v)}|^{2} e_{n}^{2}$$

$$e_{i,U2}^{2} = |G_{-}^{(i)}|^{2} i_{n}^{2} + |G_{+}^{(i)}|^{2} i_{n}^{2},$$
(4.7)

where $e_{v,U2}^2$, $e_{i,U2}^2$ are the output-referred voltage and current noise PSD, respectively (plotted in Fig. 4.4).

Next, we calculate the DAC noise contribution. The voltage gain from the node $V_{\rm ctl}$ is given by

$$G_{\rm ctl}^{(v)} = \left[\frac{1}{1 + 2\pi i f R_{\rm LP} C_{\rm LP}}\right] G_+^{(v)}, \qquad (4.8)$$

where the bracketed term represents the contribution to the transfer function from the switchable low-pass filter. Without the addition of this low-pass filter, the voltage noise of the DAC would dominate both the low- and high-frequency noise performance of the circuit. A simple solution would be to place a fixed-corner filter at this node, but this would severely restrict the AC performance of V_{ctl} . Thus, we use a switchable low-pass filter (as described in Sec. 4.3.3) to achieve low-noise performance during DC operation, while still permitting the DAC to modulate V_{ctl} more quickly when needed. The non-zero resistance of the switch contributes a zero to the transfer function at $\approx 23 \text{ kHz}$, but has negligible effect on the performance. The DAC voltage noise contribution can now be calculated as

$$e_{n,\text{DAC}}^2 = \left| G_{\text{ctl}}^{(v)} \right|^2 v_{n,\text{DAC}}^2 ,$$
 (4.9)

where $v_{n,\text{DAC}}$ is the frequency-dependent output voltage noise of the DAC as reported in the datasheet. [97] In subsequent calculations, we take the DC-mode operation $(f_c = 0.8 \text{ Hz})$ for the switchable low-pass filter.

The Johnson noise contribution can be calculated by modeling each resistor with a parallel current noise given by $i_{\rm R}^2 = 4k_BT/R$. Resistors R_1 , R_2 , and $R_{\rm mod}$ all contribute current noise at the inverting node of U2, which as discussed previously has a transimpedance gain to the output given by $G_{-}^{(i)}$. The resistor $R_{\rm LP}$ contributes current noise at the non-inverting node, which sees a transimpedance gain $G_{+}^{(i)}$. Thus, the total Johnson noise is

$$e_{n,JN}^{2} = 4k_{B}T \times \left[\left| G_{-}^{(i)} \right|^{2} \left(\frac{1}{R_{1}} + \frac{1}{R_{2}} + \frac{1}{R_{\text{mod}}} \right) + \frac{\left| G_{+}^{(i)} \right|^{2}}{R_{\text{LP}}} \right].$$
(4.10)

The low-noise stabilization circuit is limited in its ability to reject residual switching noise from the DRV2700 regulator (after the passive filtering network) by the total loop gain analyzed from node V_{out} . The LM7171 has a reported open-loop gain of 85 dB ($\approx 1.8 \times 10^4$), with a dominant pole at ≈ 10 kHz. [90] We can model the open loop gain, G_{OL} , as

$$G_{\rm OL} = \frac{1.8 \times 10^4}{1 + i(f/10 \,\mathrm{kHz})} \,. \tag{4.11}$$



Figure 4.4: Output voltage noise. The graph on the left shows the calculated noise contributions from each source in our model, along with the total calculated noise. The graph on the right shows the measured output noise at several voltages, with the total noise model from the lefthand graph superimposed for reference. Power spectral density (PSD) is referred to the high-voltage output, and the table shows the integrated RMS noise due to each noise source in different frequency bands. The total RMS noise (summed in quadrature) over the entire $1 \, \text{Hz} - 100 \, \text{kHz}$ range is calculated to be 113 μ V, with the residual DRV2700 switching noise measured at 100 V as described in the text. Frequency-dependent noise spectra for the DAC and op-amp were extracted from the datasheets. The measured noise in the right-hand plot was taken on an SR780 spectrum analyzer, AC-coupled with a corner frequency $\lesssim 1 \, \text{Hz}$. The integrated RMS noise ($1 \, \text{Hz} - 100 \, \text{kHz}$) is $\{138, 80, 101\} \, \mu\text{V}$ measured at $\{50, 100, 200\}$ V, with no output load.

The feedback network contributes a gain

$$G_{\rm FB} = \frac{R_2 || Z_{\rm mod}}{Z_1 + R_2 || Z_{\rm mod}}, \qquad (4.12)$$

arising from the voltage partition between V_{out} and the inverting node of U2. From these, we write down the closed-loop gain seen from V_{out} ,

$$G^{(\text{DRV})} = \frac{1}{1 + G_{\text{OL}}G_{\text{FB}}}.$$
 (4.13)

The contribution to the output from residual switching noise, $v_{n,\text{DRV}}$, is then

$$e_{n,\text{DRV}}^2 = v_{n,\text{DRV}}^2 \left| G^{(\text{DRV})} \right|^2.$$
 (4.14)

For $R_1 = 1 \text{ M}\Omega$ and $C_1 = 1 \text{ nF}$, $v_{n,\text{DRV}}$ is attenuated by as much as 76 dB at 6.3 kHz.

We estimate the noise spectral density $v_{n,\text{DRV}}$ by monitoring the node V_{U2} in Fig. 4.3, since the output of this op-amp represents the control signal required to cancel voltage fluctuations at V_{out} . The trace for $e_{n,\text{DRV}}$ plotted in Fig. 4.4 is derived from the results of this measurement. Because the measured $v_{n,\text{DRV}}$ drops below the noise floor of our spectrum analyzer at $\approx 10 \text{ kHz}$, we only plot the trace out to this frequency.

Given the noise model discussed above, our circuit is dominated by the opamp's intrinsic current noise at lower frequencies, and voltage noise at higher frequencies. The op-amp current noise contribution could be suppressed by using lower resistances R_1 and $R_{\rm mod}$, however one must be careful about power and current limitations when dealing with such high voltages. Each noise source is tabulated and plotted in Fig. 4.4, and the total voltage noise $(1 \, \text{Hz} - 100 \, \text{kHz}, \, V_{\rm out} = 100 \, \text{V})$ is calculated to be $113 \, \mu V_{\rm RMS}$.

4.5 Results

The measured performance of the high-voltage piezo driver is shown in Fig. 4.4, where we plot the noise power spectral density measured at several different output voltages. These traces were taken on a Stanford Research Systems SR780 spectrum analyzer, with the high-voltage output coupled through a 0.5 Hz high-pass filter and without any capacitive load. This represents a worst-case scenario, as larger capacitances at the output will reduce the noise. The integrated noise (1 Hz – 100 kHz) was measured to be { $138 \,\mu V_{\text{RMS}}$, $80 \,\mu V_{\text{RMS}}$, $101 \,\mu V_{\text{RMS}}$ } for { $50 \,\text{V}$, $100 \,\text{V}$,



Figure 4.5: Time-domain traces of the high-voltage output (100 V, color online). (a) Short-time scatterplot, measured on a PicoScope 5442B (AC coupled). The points have been down-sampled for clarity, and colored based on a normally-distributed probability of occurence in each 100-µs timeslice (scaled to the most probable voltage). The color thus provides a visual estimation of the RMS width. (b) Long-term trace, measured on a Keithly 2010 digital multimeter. A 100 V DC offset is subtracted from the plotted values. The top panel shows the lab temperature during the same time period.

200 V} outputs. This matches roughly with the expected total RMS noise calculated in Sec. 4.4. The difference in noise performance at different output voltages can be traced back to the residual ripple of the DRV2700, which is larger for lower output voltages. Indeed, one can see the characteristic shape change in Fig. 4.4 between 100 V and 50 V as the residual ripple $e_{n,\text{DRV}}$ begins to dominate at low frequencies.

Fig. 4.5 shows the performance at both short- and long-time scales. At long times, voltage fluctuations on the order of a few mV can be observed. This is due generically to 1/f noise, but also correlates with the external temperature. A cross-correlation between the measured temperature and output voltage yields an effective temperature coefficient of -24 ppm/° C at 100 V, see Fig. 4.6. The short-term trace was taken on a PicoScope 5442B (AC-coupled, 100 V output), and downsampled for



Figure 4.6: Output voltage fluctuation, measured during a period of laboratory temperature instability. The graph at top shows deviations in the output voltage (centered at ≈ 100 V, plotted in blue), with the temperature plotted in yellow-orange over the same time period. From this data, we compute the covariance matrix, shown at bottom. The diagonal elements represent the variance σ^2 , while the off-diagonal elements represent the covariance $\text{COV}(T, V) = \langle (T - \overline{T})(V - \overline{V}) \rangle$ between the temperature T and voltage V. At ≈ 100 V output, this represents a temperature coefficient of $-21.1 \text{ ppm/}^{\circ}$ C, given by $\text{COV}(T, V)/(\sigma_T^2 \langle V \rangle)$.

clarity. Points are binned into 100-µs slices, and colored based on their normallydistributed probability of occurrence. The color scale is normalized to the most probable voltage in each bin.

Fig. 4.7 shows the measured frequency response under different load conditions. The unloaded bandwidth is as high as a few megahertz, while a 1 µF capacitive load can still be driven at ≈ 100 kHz. Several mechanical resonances can be seen with a 700 nF piezoelectric load, as expected. In a laboratory setting, these resonances can be mitigated by using a digital feedback controller with notch filters tuned to match the exact resonance frequencies observed in the system [98], thereby extending the usable bandwidth out to ≈ 100 kHz.



Figure 4.7: Modulation input transfer function (color online). The solid blue line indicates the unloaded frequency response, with the phase plotted in the lower panel. The dash-dotted orange trace shows the response with a 700 nF piezoelectric actuator (Thorlabs PN AE0505D08F). Several mechanical resonances above $\approx 50 \text{ kHz}$ are clearly visible. The dashed green trace shows the response under a 1 µF capacitive load. Loaded response bandwidth is $\approx 100 \text{ kHz}$, and the unloaded gain is flat within 0.1 dB out to 1 MHz where the phase is -20° .

4.6 Conclusion

We have designed, built, and characterized a high-voltage piezoelectric driver optimized for use in a modern atomic physics laboratory. It is based on a flyback configuration switching regulator, but is able to achieve very low noise performance by active stabilization from a high slew-rate op-amp. This hybrid architecture makes it small and easy to deploy in a variety of situations, without requiring an external high-voltage power supply. The design principles discussed here can be adapted to fit the exact application, and all design files are freely available on GitHub for others to use and modify.

4.7 Future developments

This circuit has functioned nicely in our lab for over a year now, however there are a few design modifications that should be made before fabricating a new version. First, the onboard microcontroller⁷ has an ADC which is connected to a divided-down pickoff of the high-voltage output monitor. The idea here was to allow the microcontroller to read the actual output voltage, and in conjunction with its own knowledge of what the programmed setpoint is, infer the voltage supplied by the analog modulation input. From this reading, the microcontroller can adjust the internal setpoint to keep the external modulation input centered around a particular value.

Unfortunately, when we went to implement this feature, it was discovered that the voltage read by the microcontroller ADC was very non-linear with the output voltage, particularly at the low end, which came as a bit of a surprise! However, on further inspection of the datasheet, we discovered an important footnote: the analog source resistance should be kept as low as possible (less than a few ohms), and the input impedance RC time constant should be ≤ 1 ns. Neither of these is presently true for our design; the equivalent source resistance is ≈ 2.3 k Ω , and given parasitic capacitances can be \sim pF, the source RC time constant will be greater than 1 ns. Luckily, this is quickly resolved by buffering the voltage monitor before sending it to the ADC.

Currently, we work around this nonlinearity by implementing a "calibration" procedure. The microcontroller creates a lookup table at several voltages over the full output range of the device, and thus can map the ADC measured values to the

⁷We are using the MK20DX256VLH7, an ARM Cortex M-4 from Freescale, along with the Teensy bootloader for easy programming and integration with the plethora of Arduino-based library firmware. Not every board has this bootloader populated, however; instead, we have a "programming board" with the bootloader chip mounted, which plugs into the header J5 and performs the same functionality.

actual output voltage, linearly interpolating between calibration points. Unfortunately, this procedure seems to only be good to $\approx 2-3$ V accuracy. Thus, the digital feed-forward functionality is a bit limited. We also have not yet implemented the slow-scan feature, but its addition will be straight forward using the API laid out in src/piezo_api.h [99].

The firmware, which can be found under the code/ directory of the main repository [99], has a good deal of inline documentation. However, I will mention a few high-level design choices here, along with some frustrating workarounds in the Arduino build system. First, piezo.ino is the top-level Arduino sketch. Individual modules, which implement different functionality, are broken into their components in src/*.h. In the main loop() function, calls are made to process input in the individual modules (e.g., to process input from the frontpanel rotary encoder). There is a stub module for implementing the slow scan mentioned above.

Unfortunately, the Arduino build system comes close to being an excellent tool, but misses the mark on several fronts. The IDE, apart from having a difficult UI for any project of reasonable size, does not provide a good way to package code on a per-project basis. We instead use a custom make script, in conjunction with a build options file (build.options.json), to facilitate compilation. This specifies the local lib/ folder as the libraries path. After compilation, one can grab the HEX file generated in bin/, and use the Teensy loader ⁸ to flash the microcontroller. In other projects (e.g., the AOM driver found in [100]), we also use this same code structure, and pull the current git commit hash just before compilation to allow firmware traceability (something which should be ported to the piezo codebase as well). More details about the Arduino builder tool can be found on Github⁹, and a quick google search reveals many helpful StackExchange answers.

⁸https://www.pjrc.com/teensy/loader.html

⁹https://github.com/arduino/arduino-builder

I will conclude by mentioning a long-standing plan to integrate this circuit, along with other Eurocard form-factor circuits, into a local-area network (LAN) ecosystem of controlled devices. One might imagine controlling the current/piezo of a diode laser, either programmatically or by hand from the main computer, without needing to walk to the back of the lab. This could be useful, for example, in our wavemeter-based repump laser locking scheme. Software running on the computer controlling the wavemeter could remotely send packets to the laser control electronics, instructing them to update their output based on some computed frequency error signal.

To facilitate this future integration, all recent Eurocard-based circuits implement an I^2C bus connection to the backplane.¹⁰ A rear-facing module in the Eurocard rack, which could be as simple as an Arduino-based WIZNET module or as fancy as an FPGA or SoC such as the Zynq, exposes rack modules to the lab network. Commands sent to the rack are parsed and passed to individual modules to query or change their behavior. I would've liked to see this planned integration through to completion, but it continually fell down the list of priorities.¹¹

¹⁰For the backplane standard we have settled on, see [101], particularly the in-progress branch at the time of this writing.

¹¹It is worth noting that Alessandro Restelli has been experimenting with some good success with the PYNQ ecosystem, which integrates a python front-end with an FPGA/Linux backend running on a Zynq SoC. Using this as a model, one could imagine an easy way to expose low-level hardware behavior in a convenient pythonic frontend. Xilinx sells a rather cheap but performant evaluation board for the Zynq, the MicroZed (http://zedboard.org/product/microzed), which has built-in ethernet PHY and a high-pincount FMC connector on the bottom. In the future, it could be used to quickly add realtime capability to custom-built electronics, while providing a python interface via the PYNQ project.

Chapter 5: Isotope shift spectroscopy

Recently, there has been a push to use the unprecedented precision offered by cold atomic systems to investigate new physics beyond the standard model. For a good summary of those efforts, including precision measurements to detect variations in fundamental constants, see the recent review by Safronova *et. al.* [102]. Here, however, we focus on measuring isotopic shifts in atomic energy levels as a probe of many-body electron correlations and nuclear effects.

There is a long history of these kinds of measurements, dating back to the mid-1900s [103, 104]. In particular, W.H. King recognized that one could plot the isotope shifts, scaled by a factor related to the difference in mass, of two separate atomic transitions and the points would fall along a line. This relationship is the eponymous "King Linearity," and the plots are known as "King Plots." The slope of the line gives the ratio of the field shifts of the two transitions, and the intercept is related to the mass shift. The theoretical underpinnings of these shifts will be discussed in Sec. 5.1.

Deviations from King Linearity are known in some systems, and hint at the complexity of atomic and nuclear structure calculations [105]. For example, relativistic corrections to electron wavefunctions, which give rise to fine structure splitting, also contribute to a J-dependence of the field shift. Any attempts to attribute deviations from King Linearity to new physics beyond the standard model will require careful theoretical study to rule out any complicated standard model contributions [106]. Furthermore, any attempt to constrain new physics models using

isotope shift spectroscopy requires high-precision measurements establishing King Linearity, and the assumption that no fortuitous cancellations between standard model effects and new physics occur.

In addition to placing constraints on beyond-standard-model theories that predict deviations from King linearity, isotope shift spectroscopy is a very sensitive tool to benchmark state-of-the-art atomic structure calculations. In particular, recent data from photon recoil spectroscopy of the D2 line in Ca⁺ suggest a several-sigma disagreement between experiment and theory [107]. Specifically, the authors of Ref. [107] perform several different *ab initio* calculations of the field shift ratio on the D1 and D2 lines, the closest of which still differs by 3σ from the experimental uncertainty.¹ By contrast, a recent experiment in francium found excellent agreement with the *ab initio* theory [108]. Thus, there is significant utility to performing similar analyses in another independent system to benchmark these *ab initio* methods. Strontium is particularly appealing because it presents several very narrow optical transitions which allow a precise determination of the isotope shifts.

One distinct drawback, however, is the lack of a fourth stable bosonic isotope. As discussed below, complications arising from hyperfine interactions make a clean theoretical analysis of the fermionic isotope, ⁸⁷Sr, challenging. The remaining two isotope shifts of ⁸⁴Sr and ⁸⁶Sr, relative to the most abundant isotope, ⁸⁸Sr, make the King plot overconstrained — we are reduced to fitting a line between two points. However, measurement of all isotope shifts including those involving ⁸⁷Sr will provide high-precision experimental data for theorists to compare in a proper accounting of hyperfine effects. Despite this drawback, one can still extract values for the field shift ratio and other interesting quantities, with uncertainties propagated from the underlying experimental uncertainties in each datapoint.

¹In private communication with M. Safronova, there is reason to think this discrepancy is even larger (closer to 6σ), given known issues with some of the *ab initio* methods.

The rest of this chapter is laid out as follows: in Sec. 5.1, I will describe some theoretical machinery necessary for analyzing and understanding the isotope shift data. In Sec. 5.2, I describe details of the isotope shift measurements we undertook on the 689 nm intercombination transition and the 698 nm clock transition. Results are presented in Sec. 5.2.2, and Sec. 5.3 presents a King plot analysis. We end with a discussion of future extensions to this work.

5.1 Isotope shifts: theory and practice

Atomic systems provide a truly exceptional testbed for theoretical physics. Because they can be studied in near perfect isolation from external perturbations, it is possible to make predictions that are built directly from first principles in quantum theory. Simultaneously, there is an enormous bedrock of experimental techniques that enable us to perform some of the most precise measurements in the history of science on atomic systems. For example, recent results from the Yb lattice clock at NIST report sensitivities approaching a few parts in 10^{-19} [109].²

This clean experimental system notwithstanding, complications in the mathematics of many-body theories make it difficult to do exact calculations; beyond hydrogen, we lack analytical solutions, and numerical methods break down as the computational complexity grows exponentially with the number of interacting particles. This leaves theorists trying to develop approximation methods which can still capture relevant physics, and it again becomes valuable to have clean experimental systems to test those approximations.

For our purposes here, we do not dig into the full complexities of *ab initio* theoretical methods to interpret the experimental data. However, it is necessary to

²As one of my professors said at some point, the only thing we really know how to measure well is frequency. Any precision measurement, at some level, is just a clever way of demodulating the signal of interest in such a way that it becomes equivalent to a frequency measurement. In this Yb lattice clock result, they are able to resolve sub-centimeter differences in the gravitational redshift from earth's gravitational field.

have a basic understanding of the underlying atomic theory. In the next few sections, I present derivations of a few of the key results that will inform the subsequent experimental analysis. I also point out places where further theoretical analysis would add substantially to our understanding.

5.1.1 Mass shift

As with any problem in quantum mechanics, we begin with a Hamiltonian describing our system. In the present case, we have the Hamiltonian of a multielectron atom,

$$\hat{H}_{0} = \frac{\hat{\mathbf{P}}^{2}}{2M} + \sum_{i} \left[\frac{\hat{\mathbf{p}}_{i}^{2}}{2m} - \frac{Ze^{2}}{|\hat{\mathbf{r}}_{i}|} \right] + \sum_{i>j} \frac{e^{2}}{|\hat{\mathbf{r}}_{i} - \hat{\mathbf{r}}_{j}|}, \qquad (5.1)$$

where $\hat{\mathbf{P}}$ is the momentum of the nucleus with atomic number Z, M is the nuclear mass, $\{\hat{\mathbf{p}}_i, \hat{\mathbf{r}}_i\}$ are the momentum and position of the i^{th} electron, m is the electron mass, and e is the elementary charge. In the idealized case of a heavy atom, $M \to \infty$ and the first term of Eq. 5.1 is neglected. However, accounting for the non-infinite mass of the nucleus gives rise to the mass shift electronic energy eigenstates. In Sec. 5.1.2, we will discuss the so-called "field" (or "volume") shift, which arises from changes in the charge distribution of the nucleus as neutrons are added from one isotope to the next.

Taking Eq. 5.1 and moving to the center of mass coordinate system, we can write,

$$\mathbf{P} = -\sum_{i} \mathbf{p}_{i} \tag{5.2}$$

$$\mathbf{P}^2 = \sum_i \mathbf{p}_i^2 + 2 \sum_{i>j} \mathbf{p}_i \cdot \mathbf{p}_j \,. \tag{5.3}$$

Substituting Eq. 5.3 into Eq. 5.1, we can write

$$\hat{H}_{0} = \sum_{i} \left[\frac{\mathbf{p}_{i}^{2}}{2M} + \frac{\mathbf{p}_{i}^{2}}{2m} + \right] + \frac{1}{M} \sum_{i>j} \mathbf{p}_{i} \cdot \mathbf{p}_{j} + \sum_{i>j} V(\mathbf{r}_{i}, \mathbf{r}_{j})$$
$$= \sum_{i} \left[\frac{1}{2} \left(\frac{M+m}{Mm} \right) \mathbf{p}_{i}^{2} \right] + \frac{1}{M} \sum_{i>j} \mathbf{p}_{i} \cdot \mathbf{p}_{j} + \sum_{i>j} V(\mathbf{r}_{i}, \mathbf{r}_{j}), \qquad (5.4)$$

where we have dropped the hats and rewritten the potential energy as $V(\mathbf{r}_i, \mathbf{r}_j)$. The first term in square brackets is quickly recognized as the standard atomic Hamiltonian with reduced mass $\mu = Mm/(M+m)$, while the second term involves correlations between electron momenta. The second term is often referred to as the mass polarization term, and will give rise to the "specific mass shift" discussed below.

Let us now consider what happens when the mass $M \to M'$. We can write the new Hamiltonian

$$\hat{H} = \sum_{i} \left[\frac{1}{2} \left(\frac{M' + m}{M'm} \right) \mathbf{p}_{i}^{2} \right] + \frac{1}{M'} \sum_{i>j} \mathbf{p}_{i} \cdot \mathbf{p}_{j} + \sum_{i>j} V(\mathbf{r}_{i}, \mathbf{r}_{j})$$
(5.5)
= $\hat{H}_{0} + \hat{H}_{1}$,

where \hat{H}_0 is the Hamiltonian from Eq. 5.4. We solve for \hat{H}_1 to find

$$\hat{H}_1 = \hat{H} - \hat{H}_0$$

$$= \left(\frac{1}{M'} - \frac{1}{M}\right) \left[\frac{1}{2}\sum_i \mathbf{p}_i^2 + \sum_{i>j} \mathbf{p}_i \cdot \mathbf{p}_j\right].$$
(5.6)

At first order in perturbation theory, the energy shift associated with this change in nuclear mass is simply

$$\Delta E_{\text{mass}}^{(1)} = \langle \hat{H}_1 \rangle \,. \tag{5.7}$$

This energy shift has two terms. The first term,

$$\left(\frac{1}{M'}-\frac{1}{M}\right)\sum_{i}\frac{\langle \mathbf{p}_{i}^{2}\rangle}{2},$$

is called the "normal mass shift", while the second term,

$$\left(\frac{1}{M'} - \frac{1}{M}\right) \sum_{i>j} \langle \mathbf{p}_i \cdot \mathbf{p}_j \rangle \,,$$

is the "specific mass shift" and arises from electron-electron correlations. For simplicity, one would like to neglect the specific mass shift because it is a difficult many-body calculation to make. However, for multi-electron atoms, this shift can be rather large and does not necessarily even act in the same direction as the normal mass shift. But, the important result here is that both contributions to $\Delta E_{\text{mass}}^{(1)}$ scale in the same way with changing mass, namely, with $\frac{1}{\mu} = (\frac{1}{M'} - \frac{1}{M})$; similarly, when we take the difference in energy shifts between two energy levels to compute the isotope shift of an atomic *transition*, that scales with μ as well. This will be important later when we discuss King Linearity.

5.1.2 Field shift

The other isotopic shift that must be accounted for when calculating atomic spectra arises from an additional term not included in Eq. 5.1. In writing down that equation, we assumed the nucleus to be a point charge with an electrostatic energy

$$V(\hat{\mathbf{r}}) \sim -\frac{1}{\hat{\mathbf{r}}}$$

We know from E&M this is only technically true for an electric monopole, but it reveals an avenue for probing nuclear structure by measuring properties of atomic structure — deviations from the Hamiltonian in Eq. 5.1 will cause shifts which



Figure 5.1: Schematic potential energy for a uniformly charged sphere of radius $r_{\rm nuc}$. In this simple model, electrons with non-zero wavefunction overlap near r = 0 see a potential which deviates from 1/r. The size of the nucleus is approximately assumed to scale $r_{\rm nuc} \sim A^{1/3}$ with mass number A. The field shift is proportional to the electron wave function $|\Psi(0)|^2$ evaluated at the nucleus. See Ref. [110].

depend on the nuclear charge distribution and on higher order electric and magnetic multipoles of the nucleus.

Beyond accounting for higher order terms in the classical multipole expansion of the nuclear charge, isotope shifts can also be sensitive to corrections from relativity and quantum electrodynamics [107]. For example, neglecting these corrections would predict field shift constants that have no J-dependence — that is, the field shift of transitions to the same fine structure multiplet should be identical. If instead one solves the Dirac equation for the electronic wavefunction, the solution is a bispinor which differs in the overlap at $\mathbf{r} = 0$ for fine structure levels with different J. This will manifest as a "splitting isotope shift," $\delta \nu_{\text{SIS}}^{AA'} = \delta \nu_i^{AA'} - \delta \nu_j^{AA'}$, between two fine structure transitions i, j and isotopes with masses A, A' [107, 111]. It will also manifest as a field shift ratio which deviates from unity in the King plot, as discussed in the next section.

5.1.3 King plots

Now that we have established some basic principles, we can discuss techniques for analyzing isotope shifts.³ Perhaps the most common (and model-independent) technique in this field is the King plot analysis. If we assume the mass and field shifts factorize according to

$$\delta\nu_i^{A,A'} = \delta\nu_{i,\text{mass}}^{A,A'} + \delta\nu_{i,\text{field}}^{A,A'}, \qquad (5.8)$$

where $\delta \nu_i^{A,A'}$ is the isotope shift for transition *i* between two isotopes with mass numbers *A*, *A'*, then we can use the results from the previous discussion to rewrite the isotope shift as

$$\delta\nu_i^{A,A'} = K_i \left(\frac{1}{M_{A'}} - \frac{1}{M_A}\right) + F_i \,\,\delta\langle r_c^2 \rangle^{A,A'} \,. \tag{5.9}$$

Here, the coefficient K_i wraps up the physics of the mass shift (normal plus specific mass shift), which are both proportional to

$$\frac{1}{\mu_{A,A'}} = \frac{M_A - M_{A'}}{M_A M_{A'}}$$

$$= \frac{1}{M_{A'}} - \frac{1}{M_A}.$$
(5.10)

The inverse mass factor $\mu_{A,A'}$ is introduced via the definition 5.10 to be consistent with the literature [106, 107, 112]. The second term, $F_i \,\delta\langle r_c^2\rangle^{A,A'}$, represents the field shift $\delta\nu_{i,\text{field}}^{A,A'}$, which is proportional to the change in the mean-square charge radius of the nucleus, $\delta\langle r_c^2\rangle^{A,A'}$.

³This discussion follows the well-presented discussion in Ref. [107].

Taking Eq. 5.9 and multiplying through by $\mu_{A,A'}$, we get

$$\mu_{A,A'}\delta\nu_i^{A,A'} = K_i + F_i \ \mu_{A,A'}\delta\langle r_c^2 \rangle^{A,A'} \,. \tag{5.11}$$

Since the nuclear structure is independent of the electronic configuration, we can eliminate $\delta \langle r_c^2 \rangle^{A,A'}$ if isotope shifts $\{\delta \nu_i^{A,A'}, \delta \nu_j^{A,A'}\}$ are measured on two different electronic transitions $\{i, j\}$. With a little bit of algebra, we find

$$\mu_{A,A'}\delta\nu_i^{A,A'} = K_i - \frac{F_i}{F_j}K_j + \frac{F_i}{F_j}\mu_{A,A'}\delta\nu_j^{A,A'}.$$
(5.12)

Equation 5.12 shows a linear relationship between the modified isotope shifts $\mu_{A,A'}\delta\nu_i^{A,A'}$ and $\mu_{A,A'}\delta\nu_j^{A,A'}$ on two different electronic transitions. This relationship is known as the King Linearity, after the physicist who first noticed this scaling [103]. Plotting the shifts of several (but at least three) isotopes scaled by the inverse mass factor $\mu_{A,A'}$ allows one to fit a line with slope F_i/F_j and intercept $K_i - (F_i/F_j)K_j$.

An alternative analysis is possible if we take previously measured changes in the mean-square nuclear charge radius, $\delta \langle r_c^2 \rangle^{A,A'}$. Here, the field shift constant F_i and mass shift constant K_i can be extracted directly on a single electronic transition via Eq. 5.11. Precise knowledge of the normal mass shift, which is only limited by knowledge of the atomic mass unit, then permits one to extract the much harder to calculate specific mass shift and benchmark theoretical predictions. If data from more transitions are available, one can do a multi-dimensional King analysis [107] and tighten the uncertainties on parameters such as $\delta \langle r_c^2 \rangle^{A,A'}$.

5.1.4 Hyperfine effects

One difficulty with the King plot analysis discussed in the previous section arises when considering isotopes with hyperfine structure. For a given fine-structure transition, the nuclear angular momentum I combines with the electronic angular momentum J to give total angular momentum F = I + J under the traditional L-S coupling rules. For strontium, the bosonic isotopes have I = 0, which greatly simplifies the electronic energy structure; however, the fermion ⁸⁷Sr possesses I = 9/2. For the two transitions considered here, this primarily complicates the ${}^{1}S_{0} - {}^{3}P_{1}$ intercombination line, which splits into three hyperfine manifolds $F = \{11/2, 9/2, 7/2\}$ for the ${}^{3}P_{1}$ fine structure level.

The first and most obvious complication arises in how one assigns a single frequency to a transition which is split by several GHz (which is larger than the isotope splitting), such that it can be plotted alongside the bosonic isotopes on a King plot. Fortunately, there is a procedure for assigning a "center of gravity" to transitions split by hyperfine interactions. Unfortunately, as we shall see, this procedure is overconstrained, and does not account for shifts beyond the first-order perturbation theory of the non-shifted level. Without a way to properly identify the line center, we introduce a nonlinearity to the King plot and often studies of isotope shifts will neglect fermionic isotopes due to such complications. Unfortunately in strontium, there are only three stable bosonic isotopes, which means the King plot analysis must either account for the hyperfine complications of ⁸⁷Sr, or accept the limitation that the fitted linear relationship is again overconstrained by insufficiently many points. In the King plot analysis of our results discussed in Sec. 5.3, we shall see that including ⁸⁷Sr does indeed introduce a statistically significant nonlinearity that requires further theoretical study. We will therefore use the over-constrained model and propagate errors from the bosonic isotope shift measurements where noted in calculating properties of atomic and nuclear structure. If one had access to the radioactive ⁹⁰Sr isotope, this complication would be alleviated, and deviations from King linearity (or the lack thereof) might hint at beyond standard-model physics. The high precision offered by narrow optical transitions makes this an intriguing possibility for future investigation.

With that, we will now discuss the effects of hyperfine interactions on the ⁸⁷Sr intercombination line transition.⁴ Given the fine-structure Hamiltonian \hat{H}_0 , which incorporates spin-orbit interactions but not the hyperfine interaction between nuclear and electron spins, one can add a small interaction term $\hat{H}_{\rm HFI}$ which couples the electronic and nuclear angular momentum and consider it in the framework of perturbation theory. Specifically, our Hamiltonian becomes

$$\hat{H} = \hat{H}_0 + \hat{H}_{\rm HFI},$$
 (5.13)

where the hyperfine interaction Hamiltonian is given by

$$\hat{H}_{\rm HFI} = \sum_{k,\mu} (-1)^{\mu} T^{e}_{k,\mu} T^{n}_{k,-\mu} \,.$$
(5.14)

Here, the operators $T_{k,\mu}^e$ and $T_{k,\mu}^n$ are spherical tensors of rank k that act on the electronic and nuclear subspaces, respectively.

Perturbation theory relates the shift of a particular energy eigenstate to matrix elements of the interaction Hamiltonian \hat{H}_{HFI} ; thus, we consider terms of the form

$$\langle \gamma' I J' F' m'_F | \hat{H}_{\rm HFI} | \gamma I J F m_F \rangle \tag{5.15}$$
$$= \delta_{F'F} \delta_{m'_F m_F} (-1)^{I+J+F} \sum_k \begin{cases} F & J & I \\ k & I & J' \end{cases} \langle \gamma' J' || T^e_k || \gamma J \rangle \langle I || T^n_k || I \rangle , \tag{5.16}$$

where the matrix elements are expressed via the Wigner-Eckhart theorem as products of reduced matrix elements $\langle || \cdot || \rangle$ and terms related to the angular momentum. The state labels are the standard angular momentum quantum numbers for nuclear

⁴This discussion and derivation follows the discussion in Ref. [113], which is primarily concerned with the hyperfine structure of ${}^{3}P_{2}$.

spin (I), total electronic angular momentum (J = L + S), and the total angular momentum (F = I + J), with γ encompassing all other quantum numbers. At first order, the energy shift is the diagonal term $\langle \gamma IJFm_f | \hat{H}_{\rm HFI} | \gamma IJFm_F \rangle$, which can be rewritten compactly as a sum over terms of increasing multipole order,

$$\Delta E^{(1)} = \langle \gamma I J F m_f | \hat{H}_{\rm HFI} | \gamma I J F m_F \rangle \tag{5.17}$$

$$=\sum_{k} X_{k} (IJF) \langle T_{k}^{e} \rangle_{J} \langle T_{k}^{n} \rangle_{I}$$
(5.18)

with the function $X_k(IJF)$ wrapping all of the angular dependence

$$X_k(IJF) = (-1)^{I+J+F} \frac{\begin{cases} F & J & I \\ K & I & J \end{cases}}{\begin{pmatrix} J & k & J \\ -J & 0 & J \end{pmatrix} \begin{pmatrix} I & k & I \\ -I & 0 & I \end{pmatrix}}.$$
 (5.19)

The matrix elements are related to the standard hyperfine coupling coefficients A, B by

$$A \equiv \frac{1}{IJ} \langle T_1^n \rangle_I \ \langle T_1^e \rangle_J \tag{5.20}$$

$$B \equiv 4 \langle T_2^n \rangle_I \ \langle T_2^e \rangle_J \tag{5.21}$$

For the ⁸⁷Sr, ³ P_1 fine structure manifold, J = 1 and I = 9/2, so $F \in \{7/2, 9/2, 11/2\}$. In this case, the Wigner six-J symbol in the definition of X_k is only non-zero for $k \leq 2$; the spherical tensor structure coupling I = 9/2 and J = 1 only supports up to a rank two operator. Thus, we do not have higher order multipole coupling coefficients beyond the magnetic dipole term A and the electric quadrupole term B. However, this framework could be extended for other transitions with J > 1, which must include higher-rank tensor operators.

Practically speaking, measuring the hyperfine splittings in the ${}^{3}P_{1}$ manifold of 87 Sr should allow us to calculate the energy shifts $\Delta E^{(1)}$ from Eq. 5.17 and back out the unperturbed energy level $E^{(0)}$, which would be the energy of some hypothetical spin-zero fermionic isotope with mass number 87. However, an obvious problem presents itself: in the derivation discussed thus far, we have only treated the interaction Hamiltonian $\hat{H}_{\rm HFI}$ at first order in perturbation theory. But, we are already discussing fitting a series of three frequency measurements for the ${}^{1}S_{0}(F = 9/2) - {}^{3}P_{1}(F = 11/2, 9/2, 7/2)$ transition to three unknowns — hyperfine coupling constants A and B, and the unperturbed energy level $E^{(0)}$. Thus, we cannot experimentally disentangle the higher-order perturbation theory shifts from what is available in the data, and the number we specify for $E^{(0)}$ is only accurate assuming higher order shifts are negligible.

To see how reasonable this assumption is, we consider the second order perturbative term

$$\Delta E^{(2)} = \sum_{\gamma'J'} \frac{\left| \langle \gamma' I J' F m_F | \hat{H}_{\rm HFI} | \gamma I J F m_F \rangle \right|^2}{E_{\gamma J} - E_{\gamma'J'}}, \qquad (5.22)$$

where the sum extends over all other fine structure levels $|\gamma' J'\rangle \neq |\gamma J\rangle$. A similar Wigner-Eckart decomposition yields an expression again proportional to reduced matrix elements of \hat{H}_{HFI} between fine-structure energy levels, divided by the energy difference $E_{\gamma J} - E_{\gamma' J'}$.

The largest contribution to the shift $\Delta E^{(2)}$ of the ${}^{3}P_{1}$ state will come from couplings to adjacent fine structure levels ${}^{3}P_{0,2}$. Without calculating the reduced matrix elements exactly, we estimate the magnitude of the second order contribution to be smaller by roughly the first order hyperfine splitting in ${}^{3}P_{1}$ divided by the fine structure splitting $E_{\gamma,J=1} - E_{\gamma'J'=0,2}$. We find

$$\Delta E_{\rm HFS} \cdot \left(\frac{\Delta E_{\rm HFS}}{\Delta E_{\rm FS}}\right) \sim \begin{cases} 1\,{\rm GHz} \cdot \frac{1\,{\rm GHz}}{11.8\,{\rm THz}} \approx 85\,{\rm kHz} & \text{for } {}^{3}P_{2} \\ 1\,{\rm GHz} \cdot \frac{1\,{\rm GHz}}{5.6\,{\rm THz}} \approx 180\,{\rm kHz} & \text{for } {}^{3}P_{0} \end{cases},$$
(5.23)

which is non-negligible. A full theoretical treatment of the higher-order perturbation theory shifts of the ⁸⁷Sr fine structure levels is needed to unambiguously incorporate that isotope alongside the bosons in the King plot analysis. Such a treatment is beyond the scope of this thesis; we will instead rely on propagation of errors from the measurement of isotope shifts between the three stable bosons in extracting quantities from the King plot, but cannot say anything about deviations from the predicted King linearity.

5.2 Precision measurement of strontium isotope shifts

We now turn our attention to the experimental measurement of isotope shifts on the strontium "clock" (698 nm) and "intercombination" (689 nm) lines (see level diagram in Fig. 1.1). We present the first reported spectroscopy of the clock transition in ⁸⁶Sr and ⁸⁴Sr, and the first systematic study of all isotope shifts on these two lines. Our measurements of the intercombination line transitions are consistent with previously reported values, in many cases with much smaller errors (see, for example, Refs. [114–118]), and while certain experimental shortcomings discussed below prohibited a measurement of the clock transition with state-of-the-art precision, this work paves the way for more detailed studies in the future. The results discussed in the next few sections are in the process of being written up as a formal publication, which will appear sometime next year.

5.2.1 Experimental setup

The experimental procedure is straight forward, and proceeds as follows. We begin by laser cooling as described in Chapter 1.1. Briefly, we capture each isotope from a Zeeman slower directly into a magneto-optical trap (the "bMOT") using the broad, $461 \text{ nm } {}^{1}S_{0} - {}^{1}P_{1}$ transition. Atoms are then transferred into a MOT operating on the narrow, 689 nm intercombination line (the "rMOT"), where we laser cool to $\approx \mu \text{K}$ temperatures. At this point the rMOT is brought into spatial overlap with the main 1064 nm dipole beam via an adjustable magnetic shim field. Because the isotope splitting is many tens to hundreds of MHz, the transfer from bMOT to rMOT (which has a linewidth of $\approx 7.5 \text{ kHz}$) is isotope selective, and the atom number can be varied by adjusting the time allowed for loading of the bMOT. We typically load a few million to 10^{7} atoms at temperatures $\approx 1 \,\mu \text{K}$ to $3 \,\mu \text{K}$, depending on the isotope.⁵

Measurement of the isotope shift then proceeds via direct spectroscopy, but the details differ slightly for each line due to the technical reasons explained below. In all cases, however, we measure the number of atoms remaining in the dipole trap after performing spectroscopy as a function of the probe laser frequency by on-resonant absorption imaging at 461 nm, either *in situ* or after a short (few ms) time-of-flight expansion to obtain sufficient signal-to-noise in the absorption image. Atom number is extracted by numerically summing the optical density over a region of interest [119], and subtracting a background count rate determined from a nearby region of the image that contains no atoms.

⁵The atom number we can reliably load dropped precipitously towards the end of this work, as discussed elsewhere — we speculate that we have completely coated the Zeeman slower viewport, degrading the overall bMOT performance.

5.2.1.1 Spectroscopy of 689 nm intercombination lines

Light for the intercombination line spectroscopy was derived from the "short Steck" laser, which is referenced to the 689 nm master laser via a dynamically tunable beatnote lock (see Sec. 2.5.2 and Fig. 2.6). The light passes through a 100 MHz AOM and is fiber coupled to the experiment table. A photodiode on the opposite side of the chamber monitors the power and servos the intensity by feeding back to the RF drive amplitude of the AOM. The servo is based on the design from NIST [37], and permits a digital sample and hold during the dithered spectroscopy described below. The absolute frequencies quoted below are obtained by measuring an optical beatnote between the 689 nm master laser and a Menlo Systems frequency comb (FC1500-250-ULN), which is referenced to a 10 MHz Rubidium atomic clock (SRS, FS725).

The spectroscopy beam is collimated at the atoms with a $1/e^2$ waist of approximately {1.25 mm × 1.71 mm}, and propagates at a 45° angle with respect to the dipole trap. Power in this beam is typically 1 µW to 20 µW ($I/I_{\text{sat}} \approx 10$ at 1 µW), where the exact value depends on the transition and is chosen to maximize the signal to noise on resonance without saturating the atom loss signal.

Bosonic intercombination line. Spectroscopy of the 689 nm line in ⁸⁸Sr, ⁸⁶Sr, and ⁸⁴Sr was completed over several days. After loading the dipole trap at high power (≈ 3.2 W in the main dipole trap beam, as described above), we perform a short exponential ramp to the final power (≈ 1.2 W, where the trap depth $U_{\text{dip}} \approx 4.6 \,\mu\text{K}$). At this point, we perform spectroscopy in the absence of any AC stark shifts from the dipole trap by dithering the spectroscopy laser out of phase with the dipole laser, as in Ref. [120]. To maintain constant average power, we ramp the power in the dipole beam to twice its nominal value in 50 µs and dither with a 50% duty

cycle and a 500 µs period (250 µs on followed by 250 µs in the dark, during which the spectroscopy is performed). To ensure (near-)complete extinction of the dipole light, we delay the spectroscopy pulse by 50 µs after the falling edge of the dipole trap. We estimate the residual light shift due to incomplete optical extinction of the dipole trap to be well below the other systematic uncertainties of our measurement. The total duration is fixed by the integrated probe time of the spectroscopy pulses, which is typically 1 ms to 5 ms at powers ranging from 1 µW to 2 µW depending on the isotope. A small 500 mG field is added to define a quantization axis and separate transitions to ${}^{3}P_{1}$ ($m'_{J} = \pm 1$) such that we only address the $m_{J} = 0 \rightarrow m'_{J} = 0$ transition, which is first order field insensitive. After the spectroscopy is over, we hold for several ms to allow lost atoms to fall out of the trap region, and measure the remaining atom number by absorption imaging at 461 nm.

Fermionic intercombination line. The procedure for spectroscopy of the ⁸⁷Sr intercombination lines is similar to that of the bosons with a few changes. Because ⁸⁷Sr has ten m_F levels which are roughly evenly populated, we cannot obtain sufficient signal to noise if they are spectroscopically resolved by an external magnetic field. Thus, we perform this spectroscopy at near zero magnetic field. We adjust the probe power and duration to obtain sufficient contrast; because of the more complex optical pumping dynamics, spectroscopy of the ⁸⁷Sr $F = 9/2 \rightarrow F' = 7/2$ required higher probe powers and durations. For this transition, we typically used $20 \,\mu\text{W}$ and up to 15 ms probe time to achieve reasonable signal to noise. For the $F = 9/2 \rightarrow F' = 9/2$ and $F = 9/2 \rightarrow F' = 11/2$ transitions, we typically used $5 \,\mu\text{W}$ for 5 ms to get similar results.

Data analysis and systematics. To extract useful information from the spectroscopic measurement, we must understand both its accuracy and its precision. The precision is determined by the statistics of the measurement, while the accuracy must take account of various systematic effects that shift the measured frequency from the "true" frequency. Each scan of the resonance is time-tagged and fitted to an inverted gaussian, which allows us to extract a line center and offset corresponding to the average atom number during that scan. We simultaneously count the frequency of the "master" 689 nm laser with the frequency comb in a 1 s gated window, and record the frequency of the beatnote lock reference (which is supplied by a DDS). This allows us to reconstruct the absolute frequency of the spectroscopy laser, f_{laser} , averaged over the duration of the scan. We compute f_{abs} , the absolute frequency of the resonance position, as

$$f_{\rm abs} = f_{\rm laser} + f_{\rm res} \tag{5.24}$$

We take many such measurements of f_{abs} for each isotope, and tabulate them for post-analysis.

To measure the systematic density shift (see Sec. 5.2.2.1), we keep everything else constant but vary the atom number by varying the bMOT load time from ≈ 0.1 s to 10 s depending on the isotope. This generates a scatterplot of {number, f_{abs} }. The slope of a line fit to this data gives the density shift of the resonance. We compute this for each isotope, and apply the systematic offset based on the average atom number across all measurements. The uncertainty of this systematic in the final transition frequency is accounted for by the bootstrap analysis, whereby we draw from the set of all density shift measurements (with replacement) to compute the distribution of final transition frequencies. The density shift (and uncertainty) quoted in Table 5.5 is generated from a fit to the complete set of experimental data, and the uncertainty propagated from the fit uncertainty to the operating density of the measurement.

5.2.1.2 Spectroscopy of 698 nm clock line

Light for spectroscopy of the clock transition is derived from the laser system described in Chapter 3. Briefly, a home-built Littrow-configuration ECDL laser is stabilized to a passive ULE cavity via PDH lock to the first phase-modulated sideband generated by a broadband fiber EOM. Coarse tunability of the laser frequency is provided by the RF drive of the fiber EOM, which provides a few-hundred MHz offset between the atomic transition and the cavity resonance. Changing this drive frequency enables us to span the full isotope shift of the transition. Fine frequency tuning is provided by an 80 MHz AOM, which also stabilizes the probe laser intensity as described above for spectroscopy of the intercombination line. A schematic of this setup was provided in Fig. 3.13. Again, the absolute frequency reference is determined via optical beatnote with a Menlo Systems frequency comb.

Light is fiber coupled from the laser table to the experiment, and follows the same beam path as the spectroscopy light for 689 nm. However, we have removed lens L1 in Fig. 2.2 such that the beam is focused through the center of the chamber with a $1/e^2$ waist of approximately {332 µm × 461 µm}. This is necessary for the bosonic isotopes in particular, because the line we are probing is a doubly-forbidden $J = 0 \rightarrow J' = 0$ transition. To make it allowed under E1 selection rules, we must apply a magnetic field to mix in some ${}^{3}P_{1}$ character [121]. The effective rabi rate Ω then scales as

$$\Omega = \frac{\Omega_L \Omega_B}{\Delta} \,, \tag{5.25}$$

where Ω_L is the Rabi rate coupling ${}^1S_0 \leftrightarrow {}^3P_1$, Ω_B is the magnetic field coupling ${}^3P_0 \leftrightarrow {}^3P_1$, and Δ is the fine structure splitting between 3P_0 and 3P_1 . For light polarized at an angle θ with respect to the magnetic field, the resulting effective

rabi rate can be written

$$\Omega = \alpha \sqrt{I} |\mathbf{B}| \cos \theta \,, \tag{5.26}$$

where $\alpha = 19.8 \text{ mHz}/(G\sqrt{\text{mW/cm}^2})$ [121]. We operate far into the broadened regime, where the linewidth is set by Ω ; given the large inhomogeneous broadening from stark shifts in the 1064 nm dipole trap, this is necessary to achieve sufficient signal to noise on our spectroscopy. To make an order of magnitude estimate, for a field $B \approx 100 \text{ G}$ and 1 mW of probe light, the rabi rate is $\Omega \sim 40 \text{ Hz}$. Given the trap frequencies in our dipole trap are of a similar order of magnitude ($\omega_{x,y,z} \approx$ $2\pi \times \{50, 4, 495\}$ Hz), this is a very atypical regime to operate in. We discuss the effect this has on the spectroscopy lineshape in Sec. 5.2.2.2.

The dithering technique used to probe the 689 nm transition in the absence of AC stark shifts does not work for spectroscopy of the clock transition, because it causes too much heating during the 2s probe time required to achieve sufficient signal to noise in atom loss. Because of this, to reduce the sensitivity to long timescale fluctuations in dipole trap power, we take an entire isotope shift measurement in a single day between ⁸⁸Sr and the target isotope. As discussed below in Sec. 5.2.2.1 with more detail, we account for this AC stark shift systematic.

Fermionic clock line. Spectroscopy of the clock transition in ⁸⁷Sr is done at zero magnetic field. We probe with typical powers $\approx 250 \,\mu\text{W}$ for 2 s, which is enough to achieve reasonable signal to noise on the atom loss spectroscopy.

Bosonic clock line. Spectroscopy of the clock line in the bosonic isotopes requires a large magnetic field. To do this, we reconfigure the MOT quadrupole field coils from an anti-Helmholtz to a Helmholtz configuration; this is done by flipping the



Figure 5.2: Calibration data for the high-current magnetic field. Fit is to a quadratic model, where the magnitude of the magnetic field is measured via spectroscopy of the Zeeman lines in ${}^{3}P_{1}$ of 88 Sr. Vertical green lines show the operation point for 88 Sr, 86 Sr (solid), and 84 Sr (dash dotted). The fit values with statistical uncertainties at the operating points are 109.57(2) G and 197.87(5) G, respectively. A conservative estimate of the systematic error is ≤ 100 mG, which encapsulates a small deviation from the model at high currents, as evidenced by the residuals plotted above. The Landé g-factor for this transition is 2.101 MHz G⁻¹.

polarity of the lower coil with respect to the upper. Control of this is provided by an interlocked H-bridge built out of high-current Contactor relays (LEV200A4NAF). The fully debounced switching time is ≤ 100 ms, and is done during the hold after loading of the dipole trap. The field is then ramped back up to 109.6(2) G for ⁸⁸Sr and ⁸⁶Sr, and 197.9(5) G for ⁸⁴Sr. The field strength is calibrated by Zeeman spectroscopy of the ³P₁ ($m'_J = \pm 1$) levels in ⁸⁸Sr, using the known Landé g-factor for this level (2.101 MHz G⁻¹).⁶ See Fig. 5.2 for more details, and Sec. 2.3 for a more technical discussion of the coil electronics.

⁶Interestingly, at high enough fields, we start to spectroscopically resolve the magnetic field gradient across our cloud. It is a fractionally small effect for spectroscopy of the 698 nm line, but on the more sensitive Zeeman lines of ${}^{3}P_{1}$ we can selectively blow away portions of our trapped cloud.
Data analysis and systematics. Each scan of the resonance is time-tagged and fitted to an inverted gaussian, which allows us to extract a line center and offset corresponding to the average atom number during that scan. We simultaneously count the frequency of the 698 nm clock laser with the frequency comb in a 1 s gated window. As with the intercombination line, we calculate an absolute frequency, $f_{\rm abs}$, for each scan of the transition by averaging the laser frequency over the scan duration and adding the fitted line center. This is repeated for each isotope, with a reference taken in ⁸⁸Sr during the same day. The isotope shift is computed by subtracting the measured frequency in each isotope from the reference frequency measured in ⁸⁸Sr.

To measure the systematic shifts (we explicitly measure density, AC stark, probe power, and probe duration, see Sec. 5.2.2.1), we keep everything fixed but vary the experimental parameter under consideration between a "low" and "high" value. A linear fit to this data allows us to compute an offset, applied to the average value of each parameter during the spectroscopy measurement. For the recoil shift, we compute the systematic directly from the measured frequency and known atomic mass; similarly for the second order Zeeman shift, we compute it from the measured magnetic field strength and the known second-order constant $\beta = -23.8(3)$ MHz T⁻² [122].

Unlike for the density shift measurement in 689 nm, data presented for the clock transition was obtained by interleaving scans with the spectroscopy laser on and off. Points taken with the laser shuttered are used to normalize the atom number during that scan, and apply a density shift correction based on an independently measured density shift of the resonance position.



Figure 5.3: Detailed lineshape curves for the 698 nm clock transition, with fits to a Gaussian model and to the full lineshape (see Sec. 5.2.2.2). Given the atom number shot noise in our data, we cannot resolve statistically significant differences between the two models. In both cases, the fit error on the centroid is roughly 1 kHz, however the full lineshape model fits ω_0 to be ≈ 10 kHz to 20 kHz red of the Gaussian line center. This is attributable to the thermal distribution of atoms in a dipole trap operated far from the magic wavelength.

5.2.2 Results

A summary of the isotope shift measurement can be seen in Table 5.1, and the absolute frequencies extracted from our spectroscopy are shown in Tables 5.2 and 5.3. We also report the hyperfine A and B coefficients for ${}^{3}P_{1}$ in Table 5.4. A breakdown of the systematics for the intercombination and clock-line transitions can be found in Tables 5.5 and 5.6, respectively. These will be discussed in more detail in Sec. 5.2.2.1. We also show a detailed scan of each resonance for the clock transition in Fig. 5.3, including the two newly measured lines in 86 Sr and 84 Sr. For a discussion of the fits displayed in this figure, we defer to Sec. 5.2.2.2.

To generate confidence estimates in each parameter reported in the tables

below, we perform a bootstrap analysis. This is done by aggregating all data for each measurement (including systematics) and resampling with replacement to create an empirical bootstrap distribution. This sampled dataset is used to compute the final parameter (*e.g.*, f_{abs} , the isotope shift, etc.). We repeat this 1000 times to create a histogram of possible values for that parameter, which are shown in Figures 5.4, 5.5 and 5.6. The 95% confidence interval (CI) is taken from this distribution, and the mean is reported (along with the 95% confidence bounds) in each table. It is worth noting that the 95% CI is a wider bound than the typical 1σ error bar, and most of the uncertainty comes from incorporating the systematic shifts.



Figure 5.4: Bootstrap distributions for the (absolute frequency) 689 nm intercombination line transitions, including systematics. The 95% confidence interval (CI) is marked in each subfigure, see text for more details.



Figure 5.5: Bootstrap distributions for isotope shifts on the 689 nm intercombination line, including systematics. The 95% confidence interval (CI) is marked in each subfigure. For ⁸⁷Sr, the isotope shift is taken relative to a hypothetical I = 0 isotope, and the bootstrap distributions for the hyperfine A and B coefficients are shown in the lower panels, see text for more details.



Figure 5.6: Bootstrap distributions for isotope shifts on the 698 nm intercombination line, including systematics. The 95% confidence interval (CI) is marked in each subfigure. Isotope shifts in the lower panel are taken relative to 88 Sr. See text for more details.

Isotope pair (kHz)	${}^1S_0 \rightarrow {}^3P_1$	95% CI	${}^1S_0 \rightarrow {}^3P_0$	95% CI
$\frac{88-87}{88-86}$	$\begin{array}{c} 62187 \\ 163819 \end{array}$	[-19, 19] [-10, 13]	$\frac{62177}{162938}$	[-14, 14] [-17, 16]
88 - 84	351495	[-11, 10]	349656	[-16, 17]

Table 5.1: Summary of isotope shift measurements on each transition relative to ⁸⁸Sr, with uncertainties reported as a 95% confidence interval (CI) generated from the bootstrap distribution shown in Figures 5.5 and 5.6. The isotope shift for the ⁸⁷Sr intercombination line uses the weighted line center from the three hyperfine manifolds in ³P₁. See text for more details.

Isotope	${}^{1}S_{0} \rightarrow {}^{3}P_{0}$ frequency (kHz)			
	(this work)	95% CI	(previous)	
88	429 228 066 386	[-13,13]	$429228066418.0072(2)^{[a]}$	
87	429228004212	[-10, 11]	$429228004229.8732(2)^{[a]}$	
86	429227903444	[-11, 11]		
84	429227716731	$\begin{bmatrix} -11 & , 11 \end{bmatrix}$		

^a Ref. [122]

Table 5.2: Measured absolute frequencies for ${}^{1}S_{0} \rightarrow {}^{3}P_{0}$. Uncertainties in this work are reported as a 95% confidence interval (CI) generated from the bootstrap distribution shown in Fig. 5.6. The discrepancy in absolute frequency for 87 Sr and 88 Sr points to an uncontrolled systematic, likely due to a convolution of thermal line pulling with the measurement of the AC stark shift.

Isc	otope	${}^{1}S_{0} - $	$\rightarrow {}^{3}P_{1}$ frequency	v (kHz)
		(this work)	95% CI	(previous)
88		434 829 121 311.7	[-2.9, 5.0]	$434829121312.30(4)^{[a]}$
	7/2	434830473244.8	[-24.3, 26.4]	434830473270(55) ^[b]
87	9/2	434829342988.3	[-35.1, 38.1]	434829343010(50) ^[b]
	11/2	434827879825.6	[-30.4, 30.0]	434827879860(55) ^[b]
86		434828957492.5	[-10.8, 9.1]	434828957494(10) ^[c]
84		434828769816.4	[-8.1, 10.5]	434828769718(111) ^[d]
^a Re	ef. [116]] ^b Ref. [115] ^c	Ref. [114] d	Ref. [117]

Table 5.3: Measured absolute frequencies for ${}^{1}S_{0} \rightarrow {}^{3}P_{1}$. Uncertainties in this work are reported as a 95% confidence interval (CI) generated from a bootstrap distribution, while for previous works the uncertainty corresponds to a 1σ error bar. Given normally distributed errors, the 1σ error bar is a factor of two smaller than the corresponding 95% confidence bound.

Hyperfine Parameter (kHz)	(this work)	95% CI	(previous)
А	260085	[-4, 4]	260084(2)
В	35667	[-40, 44]	35658(6)

Table 5.4: Measured hyperfine A and B coefficients for ${}^{3}P_{1}$. Uncertainties in this work are reported as a 95% confidence interval (CI) generated from the bootstrap distribution shown in Fig. 5.5, while the previous result (Ref. [123]) reports the uncertainty as a 1σ error bar.

5.2.2.1 Systematics

When performing a (quasi–) precision measurement, a full accounting of statistical and systematic errors is necessary to understand how well you know the quantity you have measured. Statistical errors are easy, in the sense that they are unambiguously calculated from the data.⁷ Systematic errors, on the other hand, can be notoriously hard to calculate or estimate, and tabulating an exhaustive list relies on understanding the physics of the measurement procedure.

We calculated or directly investigated the effect of several systematic shifts to the measurements presented above. Because we are primarily interested in the isotope *shift*, many systematics which have a large effect on the absolute transition frequency become (nearly) common-mode in the isotope shift measurement. However, we will consider each in turn for a full accounting.

Recoil and Doppler shifts. The first and most straight forward systematic shift is the photon recoil shift, which can be treated within the same theoretical framework as the Doppler shift. Consider an atom in the electronic ground state $|g\rangle$ with some initial momentum, $|\mathbf{p}_i\rangle$. Unless one is working in the Lamb-Dicke regime, where the atomic motion is quantized and the photon recoil energy is much less than the gap to the next excited motional energy level, the transition we are driving can be written (in the lab frame) as

$$|g, \mathbf{p}_i\rangle + \hbar\omega \to |e, \mathbf{p}_f\rangle,$$
 (5.27)

where the atom has absorbed a photon with energy $\hbar\omega$, and the electronic energy level difference $\Delta E = E_e - E_g = \hbar\omega_0$.

⁷This is perhaps an oversimplification, because there are many ways one can calculate and report statistical errors, too! See discussion in the previous section, but as always, there is some amount of judgement involved.

By conservation of energy and momentum, we require $E_i = E_f$, where

$$E_i = \frac{\mathbf{p}_i^2}{2m} + \hbar\omega \tag{5.28}$$

$$E_f = \hbar\omega_0 + \frac{\mathbf{P}_f^2}{2m} \tag{5.29}$$

and similarly,

$$\mathbf{p}_i + \hbar \mathbf{k} = \mathbf{p}_f \tag{5.30}$$

where $\hbar \mathbf{k}$ is the photon momentum. Substituting \mathbf{p}_f into the expressions above, we arrive after a few lines of algebra at

$$\hbar(\omega - \omega_0) = \hbar \mathbf{k} \cdot \mathbf{v} + \frac{\hbar^2 k^2}{2m}$$
(5.31)

where now $\mathbf{v} = \mathbf{p}/m$ is the atomic velocity. The energy shift from the "true" resonance, $\hbar\omega_0$, has two contributions as shown in Eq. 5.31. The first term, $\hbar \mathbf{k} \cdot \mathbf{v}$, is the standard first-order Doppler shift, while the second term, $\hbar^2 k^2/2m$ is the recoil shift.

In principle, one might consider higher order Doppler shift corrections, which arise from taking into account relativistic expressions for the energy and momentum. However, these corrections will arise in a Taylor expansion in the (small) quantity $\epsilon \equiv v/c$, which for our µK-temperature samples, is ~ 10⁻¹¹. If we re-write the Doppler term $\hbar \mathbf{k} \cdot \mathbf{v}$ in terms of the photon energy $\hbar \omega$, we get

$$E_{\text{Doppler}} = \hbar\omega\epsilon \tag{5.32}$$

$$\approx 20 \,\mathrm{kHz} \times h$$
. (5.33)

An additional factor of ϵ , to give the second order Doppler shift, will enter at 10^{-8} Hz,

which is far below any of our statistical uncertainties and anything we can detect. Thus, they are neglected in the calculations here.

Density shift. The next systematic we consider is an interaction shift of the resonance, which scales with atomic density and is a consequence of the scattering dynamics parametrized by the (temperature-dependent) scattering length. We do not have an accurate measure of *in situ* density, but we do have a good measure of atom number. At fixed trap frequencies, this becomes proportional to the atomic density. We empirically measure a shift in the resonance for each isotope and transition while varying atom number by a factor of ≈ 2 to 4. The computed correction is applied to the absolute frequency measure before taking the difference to compute the isotope shift.

AC Stark shift. The AC stark shift arises from off-resonant coupling between our levels of interest and other atomic energy levels. We do not tabulate the probeinduced stark shifts because the shift is well below our other systematic uncertainties, but for the clock line interrogation we must consider an AC stark shift arising from the non-magic wavelength 1064 nm dipole trap. We probe this shift empirically, measuring a change in the resonance frequency for each isotope as the trap power is varied. A rough estimate of the differential AC stark shift, given the dipole beam waists from Table 2.1, is 70.5 kHz which is very close to the measured shift in Table 5.6.

2nd order Zeeman shift. Because we probe a $m_J = 0 \rightarrow m'_J = 0$ transition (with the exception of ⁸⁷Sr, which is done at zero magnetic field), there is no linear Zeeman shift to contend with. However, the second order Zeeman shift (proportional to the square of the magnetic field *B*), will in principle give rise to a systematic shift. Because we probe at low fields for the 689 nm spectroscopy, the second order Zeeman shift is below our resolution for that transition. But the high magnetic fields required to induce the clock transition in the bosonic isotopes requires proper accounting. Here, we compute the systematic shift from the known magnetic field, calibrated as discussed above in Fig. 5.2, and the known quadratic susceptibility to magnetic field from the strontium clock literature. Specifically, there is a shift

$$\Delta \nu = \beta \left| \mathbf{B} \right|^2 \,, \tag{5.34}$$

where $\beta = -23.8(3) \text{ MHz T}^{-2}$ [122].

Probe power and duration shifts. We do not tabulate these for the intercombination line transition because we probe for short times and low powers, and any systematic shift is well below the resolution of our spectroscopy. However, for the clock transition, we measure these explicitly again by varying both the power and duration while holding every other parameter constant. These are tabulated in Table 5.6.

	88		87		86	84
Systematic Shift (kHz)		F' = 7/2	F' = 9/2	F' = 11/2		
Density	1.7 ± 1.7	-34.1 ± 14.5	-51.9 ± 26.8	-43.3 ± 15.6	5.1 ± 3.4	-1.4 ± 4.3
Recoil	$4.8 \pm (< 0.1)$	$4.9 \pm (< 0.1)$	$5.0 \pm (< 0.1)$			
Total	6.5 ± 1.7	-29.3 ± 14.5	-47.1 ± 26.8	-38.5 ± 15.6	10.0 ± 3.4	3.6 ± 4.3

Table 5.5: Summary of systematics for the bosonic 689 nm intercombination line. Here we report uncertainties as the uncertainty in the corresponding systematic fit parameter propagated to the operating point during spectroscopy. This allows us to see which systematic errors are likely to dominate the final bootstrapped confidence interval. Statistical uncertainties are all below 1 kHz and do not contribute substantially to the total error budget.

	88-8	87	88-	-86	88-	-84
Systematic Shift (kHz)	88	87	88	86	88	84
Density	0.8 ± 1.6	-3.8 ± 1.2	0.2 ± 0.3	-0.9 ± 0.8	0.4 ± 0.9	-2.3 ± 0.9
Recoil	$4.7 \pm (< 0.1)$	$4.7 \pm (< 0.1)$	$4.7 \pm (< 0.1)$	$4.8 \pm (< 0.1)$	$4.7 \pm (< 0.1)$	$4.9 \pm (< 0.1)$
AC Stark	73.7 ± 2.2	71.2 ± 3.4	76.3 ± 2.2	73.9 ± 2.1	71.9 ± 2.1	73.0 ± 2.1
2nd Order Zeeman	$-2.8 \pm (< 0.1)$	$0.0 \pm (< 0.1)$	$-2.8 \pm (< 0.1)$	$-2.8 \pm (< 0.1)$	$-2.8 \pm (< 0.1)$	$-9.1 \pm (< 0.1)$
Probe Power	3.5 ± 1.6	1.3 ± 0.3	3.5 ± 1.6	3.5 ± 1.6	3.6 ± 1.6	3.6 ± 1.6
Probe Duration	3.4 ± 3.3	3.3 ± 1.3	3.4 ± 3.3	3.4 ± 3.3	3.4 ± 3.3	3.4 ± 3.3
Thermal line pulling	-22 ± 4	-17 ± 4	-22 ± 4	-8 ± 4	-22 ± 4	-21 ± 4
Total	61.3 ± 6.1	59.7 ± 5.6	63.2 ± 5.9	73.9 ± 5.9	59.2 ± 5.9	52.4 ± 5.9

Table 5.6: Systematic frequency shifts and uncertainties for the ${}^{1}S_{0} \rightarrow {}^{3}P_{0}$ transition. The three columns for ⁸⁸Sr correspond to the three independent isotope shift measurements. Here we report uncertainties as the uncertainty in the corresponding systematic fit parameter propagated to the operating point during spectroscopy. This allows us to see which systematic errors are likely to dominate the final bootstrapped confidence interval. The thermal line pulling, discussed in Sec. 5.2.2.2, is not included in the bootstrapping procedure but added as a separate systematic to calculate the final frequencies and isotope shifts. Statistical uncertainties on the absolute frequency are $\leq 1 \text{ kHz}$ for all transitions and do not contribute substantially to the total error budget.

5.2.2.2 Detailed analysis of the clock transition lineshape

As discussed previously in Sec. 3.4, because our dipole trap was not tunable to the magic wavelength for either transition ($\lambda_{\text{magic},698} = 813 \text{ nm}$, $\lambda_{\text{magic},689} =$ 914 nm [59]), we were unable to perform recoil-free spectroscopy in a deep optical lattice, as is done for true precision clock experiments. Given this limitation, it is important to understand the inhomogeneous broadening mechanisms that might affect our spectroscopic signal.

We begin with a simplified, semi-classical treatment of atom loss from the trap, following the line of argumentation from Ref. [124]. Since our signal is a laser-induced one-body loss after some probe time t_p , we start by calculating a loss rate coefficient $K(\delta\omega, I, T, U_{\text{trap}})$. The observable, namely, atom number, is governed by the differential equation

$$\frac{dN}{dt} = -K(\delta\omega, I, T, U_{\rm trap})N.$$
(5.35)

The loss rate coefficient is a function of the laser intensity I and detuning $\delta \omega = \omega_{\text{laser}} - \omega_0$ (relative to the bare atomic resonance ω_0), and and is assumed to be proportional to an ensemble average of the scattering rate over all atoms in the trap. It carries a dependence on theromorphism quantities such as the temperature T, and is parametrized by the confining potential U_{trap} .

From Eq. 1.1, the scattering rate is

$$\Gamma_{\rm scat} = \frac{\Gamma}{2} \left(\frac{s_0}{1 + s_0 + \left(2\Delta/\Gamma\right)^2} \right) \,. \tag{5.36}$$

We rearrange this expression, pulling out constant terms to write

$$\Gamma_{\rm scat} \propto \frac{1}{\left(\frac{\Gamma'}{2}\right)^2 + \Delta^2}$$
(5.37)

where $\Gamma' = \Gamma \sqrt{1 + s_0}$ is the saturation-broadened linewidth. For an atom in the trap with a given phase space coordinate (\mathbf{r}, \mathbf{p}) , the effective detuning Δ can be written

$$\Delta = \delta \omega - \frac{\mathbf{p} \cdot \mathbf{k}}{m} - \Delta U_{\text{trap}}(\mathbf{r}) \,. \tag{5.38}$$

The term $\delta \omega - \mathbf{p} \cdot \mathbf{k}/m$ is the Doppler-shifted laser frequency, and $\Delta U_{\text{trap}}(\mathbf{r})$ is the differential stark shift which arises from operating the dipole trap away from the magic wavelength. Specifically, we can write

$$\Delta U_{\rm trap}(\mathbf{r}) = \frac{1}{2} m \left(\bar{\omega}_e^2 - \bar{\omega}_g^2 \right) \mathbf{r}^2 \tag{5.39}$$

$$=\frac{1}{2}m\bar{\omega}_g^2\mathbf{r}^2\left(\frac{\bar{\omega}_e^2}{\bar{\omega}_g^2}-1\right)\,,\tag{5.40}$$

where $\omega_g(\omega_e)$ is the ground (excited) state trap frequency, and $\bar{\omega}$ is the geometric mean of the trap frequencies in all three dimensions. Since $\omega_i \propto \sqrt{\alpha_i}$, where α_i is the AC polarizability of state $i \in \{g, e\}$, we find

$$\Delta U_{\rm trap}(\mathbf{r}) = U_{\rm trap}(\mathbf{r}) \left(\frac{\alpha_e}{\alpha_g} - 1\right) \,. \tag{5.41}$$

At 1064 nm, assuming g is the ground ${}^{1}S_{0}$ state and e is the clock state ${}^{3}P_{0}$, we can compute $\alpha_{e}/\alpha_{g} \approx 0.7$.⁸ This has the effect of re-scaling the trap potential, such that

$$\Delta U_{\rm trap}(\mathbf{r}) = \alpha U_{\rm trap}(\mathbf{r}) \tag{5.42}$$

⁸We calculate the polarizability by taking a sum-over-states [125], with dipole matrix elements computed predominantly from values in the NIST database. Where better line strength data is available, we have substituted those instead. Finally, we add a small ad-hoc offset to make the static polarizability match the value from Ref. [15], which is calculated with much more complicated theoretical machinery. We justify this by noting it makes the predicted magic wavelengths match much more closely to their empirical values, and can be understood as compensating for incorrect or incomplete transition dipole matrix data.

with $\alpha = (\alpha_e/\alpha_g - 1) \approx -0.295$. In the limit of operating at the magic wavelength, $\alpha_e = \alpha_g$ such that $\alpha = 0$ and the spatial dependence drops out of the equation.

We now turn our attention to solving for the loss rate coefficient K, by taking an ensemble average over the scattering rate expressed in Eq. 5.37. Because we are interested in deriving a shape function which can be fit to real atom loss data, we ignore normalizations and overall constants which will be brought into a final fit parameter. The ensemble average of Eq. 5.37 is proportional to

$$\int d^{3}\mathbf{r} \ e^{-U_{\rm trap}(\mathbf{r})/k_{B}T} \ \int d^{3}\mathbf{p} \ e^{-p^{2}/2mk_{B}T} \left[\frac{1}{\left(\frac{\Gamma'}{2}\right)^{2} + \left(\delta\omega - \frac{\mathbf{p}\cdot\mathbf{k}}{m} - \alpha U_{\rm trap}(\mathbf{r})\right)^{2}}\right],\tag{5.43}$$

where we have taken an integral over phase space $\{\mathbf{r}, \mathbf{p}\}$ weighted by the appropriate Boltzmann factor. Here, k_B is the Boltzmann constant, T is the temperature, and m is the atomic mass.

We wish to make this dimensionless to easily work in a numerical fitting routine with experimental data. Focusing on the integral $d^3\mathbf{p} = dp_x dp_y dp_z$ first, we can choose \hat{p}_z to point along **k**. Thus, $\mathbf{p} \cdot \mathbf{k} = p_z k$, and the Boltzmann factor must be rewritten

$$e^{-p^2/2mk_BT} = e^{-(p_x^2 + p_y^2)/2mk_BT} e^{-p_z^2/2mk_BT}$$
(5.44)

The integral over p_x and p_y now factors, and can be brought into an overall scale factor. We define the dimensionless variable $y \equiv p_z/\sqrt{2mk_BT}$, $dp_z = \sqrt{2mk_BT}dy$, where now $p_z k/m = \beta y$, with $\beta \equiv k\sqrt{2k_BT/m}$. In convenient units, for ⁸⁸Sr and $2\pi/k = 698$ nm, $\beta/2\pi = 19.7$ kHz \sqrt{T} with T measured in µK. This parameterization of y serves to scale the momentum p_z to the most probable momentum at a given temperature. Putting it all together, the integral from Eq. 5.43 becomes

$$\int d^3 \mathbf{r} \ e^{-U_{\text{trap}}(\mathbf{r})/k_B T} \ \int dy \ e^{-y^2} \left[\frac{1}{\left(\frac{\Gamma'}{2}\right)^2 + \left(\delta\omega - \beta y - \alpha U_{\text{trap}}(\mathbf{r})\right)^2} \right]$$
(5.45)

where we've neglected overall constants. Turning back to the integral over \mathbf{r} , we know $U_{\text{trap}}(\mathbf{r}) = f(r^2)$. Thus, we can pull the angular integral from $d^3\mathbf{r} \equiv r^2 \sin\theta dr \ d\theta \ d\phi$ into an overall constant, leaving just the integral in r given by

$$\int r^2 dr \ e^{-r^2 m \bar{\omega}^2 / 2k_B T} \ \int dy \ e^{-y^2} \left[\frac{1}{\left(\frac{\Gamma'}{2}\right)^2 + \left(\delta\omega - \beta y - \alpha \frac{1}{2} m \bar{\omega}^2 r^2\right)^2} \right], \quad (5.46)$$

where we have also replaced $U_{\text{trap}}(\mathbf{r})$ with its explicit form $\frac{1}{2}m\bar{\omega}^2 r^2$.

Defining the dimensionless variable $x \equiv r\sqrt{m\bar{\omega}^2/2k_BT} = r \ (\bar{\omega}k/\beta)$ (which scales r by the ratio of the trap potential energy to the thermal energy k_BT), we can rewrite the integral as

$$\int dx \ x^2 e^{-x^2} \ \int dy \ e^{-y^2} \left[\frac{1}{\left(\frac{\Gamma'}{2}\right)^2 + \left(\delta\omega - \beta y - \left(\frac{\alpha m}{2k^2}\right)\beta^2 x^2\right)^2} \right]. \tag{5.47}$$

The term $\alpha m/2k^2$ can be evaluated exactly now (note, there is a factor of \hbar to convert energy to angular frequency units, which I've dropped previously by setting it to unity). Again, for m = 88 amu and $k = 2\pi/698$ nm, we find

$$\frac{\alpha m}{2k^2} \approx -2.52 \times 10^{-6} \,\mathrm{s}\,.$$
 (5.48)

What is left is to put the angular frequencies $\delta \omega$, Γ' into units commensurate with β , such that the numerical integration is more friendly. Recall $\beta/2\pi \sim 10^4$ at $T = 1 \,\mu\text{K}$, and our empirical linewidth is of a similar order. So let us express frequency in units of 10 kHz, such that

$$\tilde{\omega} \equiv 2\pi \times 10 \,\mathrm{kHz} \tag{5.49}$$

$$\delta\omega = \tilde{\omega}\ \tilde{\delta} \tag{5.50}$$

$$\Gamma' = \tilde{\omega} \ \tilde{\Gamma} \tag{5.51}$$

$$\beta = \tilde{\omega} \ \tilde{\beta} \tag{5.52}$$

$$\tilde{\omega}\left(\frac{\alpha m}{2k^2}\right) \equiv \tilde{\alpha} \approx -0.158.$$
 (5.53)

We re-write the integral finally as

$$K \propto \int_0^\infty dx \ x^2 e^{-x^2} \ \int_{-\infty}^\infty dy \ e^{-y^2} \left[\frac{1}{\left(\frac{\tilde{\Gamma}}{2}\right)^2 + \left(\tilde{\delta} - \tilde{\beta}y - \tilde{\alpha}\tilde{\beta}^2 x^2\right)^2} \right], \qquad (5.54)$$

having pulled $1/\tilde{\omega}^2$ out into the overall normalization constant.

Returning to Eq. 5.35, we use our expression Eq. 5.54 for K to solve for atom number as a function of the renormalized detuning $\tilde{\delta}$ and intensity $\tilde{\Gamma}$. We get a simple exponential,

$$N(\tau) \propto e^{-aK(\delta, \dot{\Gamma}, \beta)\tau}$$
, (5.55)

which can be used as an integral fit function to vary the parameters $\{a, (\tilde{\delta}_0 \equiv \tilde{\omega}_{\text{laser}} - \tilde{\delta}), \tilde{\Gamma}, \tilde{\beta}\}$ in a least-squares minimization routine, where τ is wrapped in to the overall normalization of K, given by the fit parameter a in Eq. 5.55.

We perform a fit using Eq. 5.55 to the loss spectra shown in Fig. 5.3. The large atom number shot noise makes it difficult to differentiate the quality of the fit between the full integral lineshape and a simple Gaussian model, but there is significant thermal line pulling from a full accounting of the lineshape. To account for this systematic, we numerically simulate the systematic Gaussian fit offset as a



Figure 5.7: Effects of thermal line pulling on the clock transition. We simulate the lineshape model discussed in the text as a function of temperature (left panel), and extract a systematic offset to the Gaussian fitted center (right panel). This systematic shift is -7.6(3) kHz µK⁻¹.

function of temperature, and find it to be -7.6(3) kHz μ K⁻¹, as shown in Fig. 5.7.

Typical temperatures for each isotope are $\{2.9 \,\mu\text{K}, 2.2 \,\mu\text{K}, 1.1 \,\mu\text{K}, 2.7 \,\mu\text{K}\}$ for $\{^{88}\text{Sr}, ^{87}\text{Sr}, ^{86}\text{Sr}, ^{84}\text{Sr}\}$, with uncertainties $\approx \pm 500 \,\text{nK}$. Given this, we add an additional systematic to Table 5.6. This is used to seed a parametric bootstrap distribution of the systematic shift due to thermal line pulling, which is included in the final bootstrap distributions discussed above.

5.3 King plot analysis

Armed with the results from Sec. 5.2.2, we conduct a King plot analysis to extract properties of strontium nuclear structure. As we can see from Fig. 5.8, including ⁸⁷Sr introduces a strong non-linearity to the King plot. The likely origin of this was discussed previously in Sec. 5.1.4 — from the experimental data, it is ambiguous what the proper "spin-zero" ⁸⁷Sr intercombination line center should be. If we add a 400 kHz offset to that ${}^{3}P_{1}$ isotope shift, we can bring all three isotopes into alignment. This is roughly the order of magnitude we expected second-order perturbative shifts to appear from couplings between the ${}^{3}P$ levels. Clearly, however, a more careful theoretical analysis of these shifts is necessary to include ⁸⁷Sr on the



Figure 5.8: King plot, fit to all isotopes. Error bars, where not visible, are smaller than the point. The line does not pass through any of the points, indicating a clear non-linearity likely due to misidentification of the hyperfine triplet linecenter in ⁸⁷Sr ³P₁. A shift of roughly 400 kHz is sufficient to bring that isotope into line with the bosonic isotope shifts. Clearly a more detailed theoretical accounting of higher order corrections is necessary to properly include ⁸⁷Sr in the King plot.

King plot.

Instead, we proceed with our bootstrapped isotope shift distributions for ⁸⁶Sr and ⁸⁴Sr, and compute the field shift ratio F_{698}/F_{689} as well as the intercept term $K_{698} - (F_{698}/F_{689})$ K_{689} from these shifts alone. This is shown in Fig. 5.9, and tabulated in Table 5.7.

Parameter	Experiment	95% CI
$\frac{F_{698}/F_{689}}{K_{698} - (F_{698}/F_{689})} K_{689}$	$0.9979 \\ -2.0$	$\begin{bmatrix} 0.9952, & 1.0008 \\ [-3.9, & -0.3 \end{bmatrix}$

Table 5.7: Summary of atomic and nuclear structure parameters extracted from the King plot analysis. The field shift ratio F_{698}/F_{689} and the difference $K_{698} - (F_{698}/F_{689}) K_{689}$ (GHz amu) are taken from the mean of the bootstrap distribution calculated over ⁸⁶Sr and ⁸⁴Sr only.



Figure 5.9: Bootstrap distributions of the field shift ratio F_{698}/F_{689} (left) and the intercept term $K_{698} - (F_{698}/F_{689}) K_{689}$ (right). These are generated from the ⁸⁶Sr and ⁸⁴Sr isotope shifts only.

5.4 Discussion and future directions

Here, we have presented the first measurement of the 698 nm clock transition in ⁸⁶Sr and ⁸⁴Sr, and the first systematic measurement of strontium isotope shifts between all isotopes on the narrow 689 nm intercombination and 698 nm clock transitions. While our measurement achieves a fractional uncertainty $\approx 10^{-11}$, there are several aspects which could be markedly improved to reduce both systematic and statistical errors.

The first and most substantial improvement would be to operate both transitions in a magic-wavelength lattice. The lack of a magic-wavelength trap manifests itself directly in the lineshape one measures for these nominally ultra-narrow optical transitions, and in the inability to perform coherent spectroscopy with reasonable signal-to-noise — the dominant linewidth is set by inhomogeneous broadening instead of the laser or (ultimately) the natural width of the atomic transition. Truly narrow features would allow one to integrate down noise much more quickly, and a magic-wavelength lattice would remove much of the uncertainty due to the AC stark shift. It would also allow us to take advantage of our narrow laser to address the entire ensemble coherently, reducing the currently necessary long probe time, larger magnetic fields, and higher probe intensities used in the 698 nm spectroscopy. Our lab recently purchased a Ti:Sapphire laser, which can be tuned between 813 nm and 914 nm (the magic wavelengths for the clock and intercombination lines, respectively). This laser will be incorporated in the next generation strontium system described in Chapter 6. We briefly considered bringing magic-wavelength light from this system into the current apparatus, but due to legacy optics there is no simple way to combine dipole beams at these wavelengths with the existing beam paths.

Chapter 6: SrII and future directions

When I joined the strontium team in 2012, the system described in Chapter 2 was already under vacuum and ready to start trapping. One of the first things I did as a graduate student was help wrap shim coils onto the chamber, and a few weeks later we achieved our first bMOT. However, the current experiment was scheduled to ultimately move to the new Physical Sciences Complex (PSC) building, which was nearing the end of construction.¹ The timing and details of that move have been rather fluid, but when we had money in the budget to re-work the vacuum system and buy new lasers, the decision was made to build up "SrII" in tandem with the operation of the "SrI" apparatus in the Computer & Space Sciences (CSS) building.

Lessons learned building the first apparatus produced a laundry list of ideas to improve the operation and performance of some future system. In this chapter, I describe the design and construction of the new strontium machine. As I write this, the vacuum system is fully assembled and holding a pressure $\leq 1 \times 10^{-11}$ torr, and is ready for the installation of lasers.

6.1 Vacuum system

The SrI machine was plagued by several issues related back to vacuum design choices. First, the high optical-access "spherical square" vacuum chamber from Kimball Physics, which formed the main science chamber, in reality ended up with very restricted optical access. This was a result of two factors: first, the 1.33" mini-CF

¹Of course, ahead of schedule and under budget!

viewports had no line of sight to the center of the chamber due to a miscommunication with the manufacturers of the recessed bucket windows, which are installed on the 8" flanges on the top and bottom of the chamber. Second, the Zeeman slower enters the chamber on one of the large 4.5" CF ports, which precluded using that axis for more optically demanding laser beam paths. The limited optical access necessitated some creative optical contortions. In addition, the Zeeman slower window was not heated adequately or placed on a short bellows as is done in other experiments, and so over time we have coated that viewport with strontium and drastically degraded the performance of our Zeeman slower. We are nearing the point in early 2018 where we'll either need to pull the plug and move to the new lab, or open vacuum to replace that viewport.

For some reason, the atomic flux of our experiment has always suffered slightly relative to other strontium groups. We speculate this is related to the aggressive differential pumping between our oven and main chamber. The first iteration of our oven used a complex, slightly-focusing array of microcapliaries manufactured out of a vacuum-compatible epoxy. However, it was later discovered to not be rated to the 650 °C temperatures we heat it to, and we suspect outgassing from the epoxy was responsible for poisoning the nearest ion pump twice in roughly five years of operation. The upgraded oven and nozzle, which was installed over winter break in 2015, is based on the design in Ref. [24] and involves a clamped array of stainless steel hypodermic needles.² The new design roughly doubled our steady-state atom number, and allowed quick progress towards quantum degeneracy in ⁸⁷Sr. However, as mentioned previously, the increased atomic flux also increased the rate at which the Zeeman slower viewport was being coated, and the atom number has drifted down continuously since the initial oven upgrade. This has a very real impact on productivity — load times for the bMOT have increased, which increases the

²For details of this design, see Ben Reschovsky's thesis [20] or the discussion in Sec. 2.1.



Figure 6.1: Zeeman slower simulation data from AOSense. The figure on the left is an example simulation of initial and final atomic velocity distributions, assuming a Zeeman beam with a $1/e^2$ beam diameter of 8 mm and 120 mW of power. On the right, we show a summary of the simulations. For each beam diameter and power, the detuning was optimized to give the highest flux below 60 m s^{-1} . From the simulations, we see that at a 5 mm diameter there is reduced spatial overlap between the atomic beam and the Zeeman slower light, which degrades performance. There is little difference between 7 mm and 8 mm beam diameters, although the optimal detuning will vary slightly (see Table 6.1).

duty cycle for an experimental shot; we are strained for power, particularly in the Zeeman slower beam, which necessitates daily peakup of fiber and AO alignments and the use of the rather touchy injection-locked 461 nm diode to provide light for the Zeeman slower. And lower atom number makes each stage of optimization from bMOT through evaporation more crucial to maintain high enough numbers to reach quantum degeneracy.

I mention some of this history as a motivation for design choices made on SrII. First, to avoid the headache of poor atomic flux, we chose to purchase a strontium beam source from AOSense (see Sec. 6.1.1). These commercial atomic beam sources have been used successfully in a few labs at JILA, and our decision to buy rather than design our own drove much of the remaining vacuum design. The final system described in the next section attempts to remedy many of the design flaws in the current system.

6.1.1 AOSense atomic beam source

The AOSense "Beam II" is designed to be a compact, high-flux atomic beam source ready to install on a standard cold-atom apparatus³. At its core, it is a short, permanent-magnet Zeeman slower attached to a dual 2D MOT which provides transverse beam collimation and 20° deflection from the Zeeman slower axis. The oven is controlled by an in-vacuum heating element, which reduces power consumption and eliminates the need for water cooling or careful thermal load management between the oven and the rest of the chamber. The Zeeman slower viewport has a secondary in-vacuum hot window which is kept at a high enough temperature to avoid strontium buildup. All other optics for beam shaping and splitting between axes on the 2D MOT stages are also housed in-vacuum, making the entire package easy to integrate with another system. One simply has to provide a beam of 461 nm light at the appropriate detuning and cold atoms emerge out the other end.

6.1.2 Science chamber

The science chamber itself is a modified spherical square from Kimball Physics, made from 316 stainless steel. We had one of the 4.5 in conflat ports enlarged to accept a 6 in minimum-length nipple, which connects the science chamber to a large pump body. The pump body is effectively a shrunken reducer cross, which adapts the 6 in nipple from the chamber into an 8 in flange. The pump body contains two pumps — first, a Nextorr-1000 non-evaporable getter (NEG)/ion pump from SAES provides the bulk of the pumping speed on hydrogen and nitrogen. We originally planned to just use this single pump, but the NaEr group at the JQI found their base pressure was argon-limited. Thus, we place a second 75 L s⁻¹ ion pump from Gamma Vacuum on the other port of the pump body. The ion pump has a secondary 2.75 in

³"Just add lasers!"

$1/e^2$ Diameter (mm)	Power (mW)	Detuning (MHz)	Captured frac. (%)
	20	-580	5.2
	40	-585	9.3
	60	-595	11.2
F	80	-595	12.7
9	100	-605	12.9
	120	-610	12.6
	150	-620	12
	180	-620	11.9
	20	-560	3.5
	40	-575	9.9
	60	-585	13.5
7	80	-585	14.1
1	100	-590	15.3
	120	-590	17
	150	-595	17.7
	180	-610	17.7
	20	-565	2.1
	40	-565	9.4
	60	-575	12.9
	70	-585	14.2
	80	-590	15.1
0	90	-595	16.1
0	120	-595	17.1
	150	-605	18
	180	-605	18.7
	200	-605	19
	220	-605	19.4
	250	-615	19.5

Table 6.1: Summary of Zeeman slower simulations from AOSense, including optimal detunings. There is little difference in performance between 7 mm and 8 mm beam diameters, but depending on the power available, the detuning should be adjusted for peak performance. See Figure 6.1 for a convenient graphical representation.

CF Size (in.)	Thickness (cm)	5 Time 350 °C	Constan 400 °C	ts (days) $425^{\circ}\mathrm{C}$
1.33	1.9	48.6	33.9	24.2
2.75	1.3	21.6	15.1	10.8
4.5	1.7	40.0	27.9	19.9
6	2.0	52.6	36.7	26.2
8	2.2	67.0	46.7	33.4
10	2.5	81.3	56.7	40.5

Table 6.2: Summary of H_2 outgassing rates for various flange sizes. The time constant scales quadratically with the flange thickness. Data from [126].

port which can be used to attach a turbo-molecular pump for the initial pump-down and bakeout. There is also a small half-nipple on the pump body which contains an ion pump (Agilent UHV-24p) for pressure monitoring if needed. The AOSense beam source attaches to the science chamber via one of the 1.33" ports next to the pump body, and care was taken to avoid any obstruction of optical access to the center of the chamber.

Before assembly, the science chamber and pump body were both baked under vacuum at 400 °C for several days to reduce the outgassing rate of hydrogen from the stainless steel bulk. The rationale for this procedure over other similar bakeout schemes comes from the vacuum sciences division at NIST [127]. This paper found that a high-temperature bake (≥ 400 °C) under vacuum is responsible for the bulk of gain in H₂ outgassing. They also found that a second high-temperature air bake could further reduce outgassing by a factor of two, likely due to the formation of an oxide layer that acts as a diffusion barrier. However, we decided not to risk that second bake due to risk of contamination.⁴ The bakeout chamber itself was built from a stock 14 in nipple, with a custom 14 in to 2.75 in reducer on one flange to

⁴Interestingly, they've also found (although perhaps anecdotally) that the benefits of a high temperature bakeout don't accrue until you reach above $400 \,^{\circ}$ C, and in fact might increase the base pressure for bakeouts between 250 °C and 400 °C. The thinking here is that an intermediate-temperature bake is enough to bring much of the hydrogen to the surface, but not enough to increase the desorption rate from the steel.

attach the turbo pump. A chicken-wire frame wrapped with heater tape and many layers of fiber glass insulation drops over the entire assembly to provide even heating, and several thermocouples monitor the temperatures and adjust power to the heater tape accordingly. After the chamber was assembled, we performed a much gentler bake at slightly above 100 °C to remove any water and residual contaminants from inside the vacuum chamber.

The supports for the chamber are made from thick-walled 1.5 in OD stainless steel pipe, which are filled with lead shot. While not as good as the dynamicallydamped posts you can buy from Thorlabs or Newport, they have the advantage of not being nickel-plated and thus are non-magnetic. However, the damping is slightly improved from a standard stainless steel post. We also purchased large, 3/4" custom breadboards from BaseLab Tools, which feature cutouts at convenient locations for optical access to the chamber. One breadboard is supported at mid-level, such that the beam height for standard Thorlabs optomechanical components is level with the center of the chamber, and a second breadboard is mounted above the chamber for beam paths that must enter vertically. Each viewport on the main chamber is attached with a custom adapter washer that allows cage components or other mechanical assemblies to be mounted directly to the chamber. A CAD rendering of the whole system can be seen in Fig. 6.2. Table 6.3 also lists each port with optical access to the atoms, the corresponding distance, and the highest-NA stock lens that can be used along that path.

6.2 Laser systems

Here, I describe the planned laser systems, beam paths, and proposed power budgets for the SrII experiment. This should serve as a reference for future students, based on some of our experiences in SrI, but may change with the facts on the ground. The main cooling light at 461 nm is generated from a Toptica SHG system,



Figure 6.2: CAD rendering of new SrII apparatus. The main chamber is a modified spherical square from Kimball Physics, with one 4.5 in CF port replaced by a 6 in minimum-length nipple to attach to the pump body. The top and bottom ports accept custom recessed viewports from MDC/Insulator Seal, which have a 2.69 in clear aperture. About the axis of symmetry, there are three 4.5 in viewports, four 2.75 in viewports, fifteen 1.33 in mini-viewports, and one 6 in viewport which attaches on the opposite side of the pump body. All viewports except for two, which are reserved for the ultra-violet (UV) Rydberg dressing laser, have a broadband anti-reflection (AR) coating covering visible through $\sim 1.2\,\mu\text{m}$. The pump body is a 8 in to 6 in four-way reducer cross with a custom 2.75 in port welded on for an ion gauge (Agilent UHV-24p). We employ both a 75 Ls^{-1} ion pump (Gamma Vacuum 75S) and an integrated non-evaporable getter (NEG)/ion pump from SAES (NexTorr-1000). The NEG provides $1000 \, L \, s^{-1}$ pumping on N₂, which we expect to be conductance-limited. The integrated ion pump is a $20 \,\mathrm{L\,s^{-1}}$ pump, which enhances the pumping speed on Noble gasses and other non-reactive species. The Gamma Vacuum ion pump has a secondary 2.75 in CF port, which we attach to an all-metal angle valve (VAT, 54132-GE02) and use for initial pumping with a turbo-molecular pump. All vacuum parts were baked for several weeks at $\gtrsim 400$ °C in an evacuated bakeout chamber (see text) to reduce the outgassing rate of hydrogen. A secondary in situ bake at slightly above $100\,^{\circ}\mathrm{C}$ was performed while attached to the turbo pump, and pressures $\leq 1 \times 10^{-8}$ torr were achieved before cooling down and valving off the all-metal angle value. Final pressures are at the detection floor of our ion gauge, $\lesssim 1 \times 10^{-11}$ torr, although we see a high degree of correlation between pressure reading and humidity in the laboratory. We suspect this is due to hydroscopic plastic in the ion gauge connector, which causes small changes in the leakage current used to indicate pressure.

Viewport	Distance (washer face to atoms)	Lens / max N.A.
4.5"	$132\mathrm{mm}$	$f = 150 \mathrm{mm}, (2" \mathrm{dia.})$ $\mathrm{NA}_{\mathrm{max}} = 0.17$
2.75"	$134\mathrm{mm}$	f = 150 mm, (1" dia.) $\text{NA}_{\text{max}} = 0.084$
1.33"	$135\mathrm{mm}$	$f = 150 \mathrm{mm}, (0.5" \mathrm{dia.})$ $\mathrm{NA}_{\mathrm{max}} = 0.042$
6.0"	$438\mathrm{mm}$	$f = 500 \mathrm{mm}, (2" \mathrm{dia.})$ $\mathrm{NA}_{\mathrm{max}} = 0.025$
Recessed bucket	$134\mathrm{mm}$	$f_{\rm eff} = 66 \mathrm{mm}, (2" \mathrm{dia.})$ $\mathrm{NA}_{\mathrm{max}} = 0.29$

Table 6.3: Summary of viewports available on SrII science chamber, along with the working distance to the center of the chamber and maximum clear aperture/numerical aperture (NA) using catalog lenses. The numbers quoted for the recessed bucket window assumes a custom-designed objective made from catalog lenses but aberration-corrected at 461 nm. Each viewport is mounted to the vacuum chamber with a custom stainless-steel washer which adapts to 15, 30, or 60 mm cage systems, depending on the viewport. Alternatively, 4-40 threaded holes allow for easy mounting of custom assemblies to the vacuum viewports.

and is split into several arms (see Fig. 6.3). Each arm represents a beam path to the experiment, and is equipped with an AOM for frequency and power stabilization, and a shutter where high optical extinction is desired. Unused light from the zeroth order of the Zeeman slower AOM can be used to monitor on the wavemeter and generate a beatnote to lock the SHG laser to a secondary spectroscopy laser. As described previously, this beatnote locking scheme allows flexibility to adjust the frequency quickly to address different isotopes within a single experimental shot. In the SrI system, the optomechanics of this laser system were very unstable; care was taken in the sketched layout to keep beam paths short and direct; since we do not require large frequency sweeps from any of our AOMs, we avoid using them in a double-passed configuration.

The 689 nm laser systems sketched in Fig. 6.4 follow a similar design philosophy, privileging compact and simple layouts over infinite flexibility. A single master laser is locked to a stable Fabry-Pérot cavity via the Pound-Drever-Hall method.



Figure 6.3: Planned layout for the SrII 461 nm laser system. Note, better AOM diffraction efficiency might be achieved (especially for the Zeeman arm) if the light coming out of the laser is detuned further and AOM frequencies are adjusted accordingly. All AOMs are 200 MHz models from IntraAction (ATM-2001A1), except for the Zeeman arm AOM which is a 350 MHz model (ATM-3501A1). Half-wave plates are omitted for clarity, but should be installed before each beam cube to appropriately set the power splitting between each arm. Polarization-maintaining (PM) fiber carries light from the laser table to the apparatus. Light for the 3D MOT is split in a 3-1 fiber coupler array (Evanescent Optics) for each arm, which is retro-reflected on the apparatus. The 2D MOT AOM can be used as an optical beam shutter for atoms emerging from the AOSense source; shutters are required on the 3D MOT arms and on the probe arms to achieve full extinction of 461 nm light during the rMOT and dipole trapping stages of the experiment.

Light for this lock is intensity stabilized by noise-eating on the zeroth order of an AOM; the rejected first order can be used as a wavemeter monitor. Remaining light is split four ways (or more, as needed) on polarizing beam cubes, and fiber coupled to a beatnote lock with a slave laser. All fibers and fiber splitter/combiners in the beatnote network should be FC/APC connectorized; in the past, back-reflections from the front facet of a fiber have destabilized the lasing mode and degraded laser performance in strange ways. Note here we do not require any high-frequency AOMs to offset the laser light to the cavity resonance, since light for the experiment is generated from the system of beatnote slave lasers.

Each slave laser is similarly simple – light from the laser passes an optical isolator, then the rejected port of a polarizing beam cube is fiber coupled to the beatnote lock. The remaining light passes an AOM, where the zeroth order is fiber coupled for a wavemeter monitor, and the deflected order is fiber coupled to the experiment. In the case of light for the red MOT, we forego the AOM and shutter and instead couple light directly to a tapered amplifier (TA).⁵ The TA current should be interlocked on the optical power from the slave seed laser, and amplified light should at this point be coupled through an AOM/shutter and into a four-way fiber splitter (Evanescent Optics) to the apparatus.

6.3 Magnetic field coils and high-NA imaging

Beyond basic laser and vacuum systems, the final necessary pieces for a working neutral-atom quantum gas experiment include a way to generate magnetic fields, a way to image your gasses (preferably with high resolution), and a computer control system to orchestrate experimental sequences. The first two of these will be addressed in this section, and we will reserve the next Section 6.4 to discuss plans for an upgraded computer control architecture.

 $^{^5\}mathrm{This}$ is the hilariously named "boosTA", from Toptica.



Figure 6.4: Planned layout for the SrII 689 nm laser system. The "Master" laser is locked to a convenient mode of a stable Fabry-Pérot cavity via the PDH technique. The remaining light is split into four (or more) arms, which reference beatnote locks for slaved lasers. The beatnote circuit is described in more detail in Sec. 2.5.2. For the beatnote laser used to reference the red MOT, the final AOM is omitted and instead light is first passed to a Tapered Amplifier (TA). Remaining beatnote slaves include stirring laser for ⁸⁷Sr, optical pumping, and optical Stern-Gerlach (OSG). Halfwave plates are omitted for clarity.

Most neutral-atom experiments at the JQI have built high-current magnetic field coils using hollow square copper tubing wound onto a form. To remove heat dissipated in the coil by the high current, a closed-loop cooling system runs chilled water through the hollowed center. During design discussions for the SrII system, we became aware of the Bitter-type electromagnet. Instead of running chilled water serially through the whole coil, a Bitter coil is constructed from interlocking copper disks with water cooling channels between each layer. In this design, the cooling occurs in parallel, drastically increasing the coil's ability to dissipate heat.

Recently, there has been interest in so-called "orbital" Feshbach resonances in alkaline-earth(-like) atoms [128–130]. Because the ground state of alkaline-earth species is an electron spin singlet, the magnetic moment is (nearly) zero and thus affords no magnetically-tunable Feshbach interaction like there is in alkali species. However, given spin exchange interactions and a slight differential Landé g-factor between ${}^{1}S_{0}$ and ${}^{3}P_{0}$, one can generate a Feshbach-type resonance using the clock transition [128].

To do this, one operates in a mixture of ground $|g\rangle$ and excited $|e\rangle$ orbital states. Because of the differential Landé g-factor between the two orbital levels, there is a small energy difference between an atom in $|g\uparrow\rangle$ and $|e\uparrow\rangle$ (with the spin state designated by $\uparrow\downarrow$), which is tunable over some range with a magnetic field. Consider the (s-wave) scattering interaction of two such atoms. We can write two scattering channels, designated $|o\rangle$ for "open" (i.e., energetically allowed) and $|c\rangle$ for "closed" (energetically forbidden), which in properly symmetrized form are

$$|o\rangle = \frac{1}{\sqrt{2}} \left[|g\uparrow\rangle_1 |e\downarrow\rangle_2 - |e\downarrow\rangle_1 |g\uparrow\rangle_2 \right]$$
(6.1)

$$|c\rangle = \frac{1}{\sqrt{2}} \left[|e\uparrow\rangle_1 |g\downarrow\rangle_2 - |g\downarrow\rangle_1 |e\uparrow\rangle_2 \right]$$
(6.2)

At infinite separation, $|g \uparrow\rangle_1 |e \downarrow\rangle_2$ and $|g \downarrow\rangle_1 |e \uparrow\rangle_2$ have an energy offset δ propor-

tional to $\Delta m_F = m_F^{\downarrow} - m_F^{\uparrow}$ and given by

$$\delta = (g_e - g_g) \Delta m_F \mu_B B \tag{6.3}$$

where the differential Landé g-factor $\delta g_F = g_e - g_g$ arises from hyperfine mixing between ${}^{3}P_0$ and ${}^{3}P_1$. This energy difference now provides a handle for tuning the energy of bound states in the open and closed channels by changing the magnetic field; given some spin exchange interaction, this now behaves analogously to a traditional magnetic Feshbach resonance.

S-wave orbital Feshbach resonances have been observed in Ytterbium [129, 130], and should exist in strontium as well but only at very large fields. To see why this is, we consider the pseudo-potential model from Ref. [128], which predicts a scattering length a_s given by

$$a_s = \frac{-a_{s0} + \sqrt{m\delta/\hbar^2 (a_{s0}^2 - a_{s1}^2)}}{a_{s0}\sqrt{m\delta/\hbar^2} - 1}$$
(6.4)

where $a_{s0} = (a_s^+ + a_s^-)/2$ is the average of s-wave scattering lengths in the symmetric (a_s^+) and anti-symmetric (a_s^-) spin channels, $a_{s1} = (a_s^- - a_s^+)/2$, m is the mass (in this case 87 amu), and δ is the magnetically-tunable energy difference between the open and closed channels. The resonance position predicted by Eq. 6.4 occurs when the scattering length a_s diverges, that is, when $\delta \to \hbar^2/ma_{s0}^2$. Crucially, the magnetic field at which this occurs is a function of the differential Landé g-factor, but also on the scattering lengths for the symmetric and anti-symmetric spin channels, contained in a_{s0} .

Using scattering numbers from Ref. [1], the known differential Zeeman shift $m_F \cdot 109 \text{ Hz/G}$ [59], and assuming a spin mixture of $m_F = \pm 9/2$, we find the scattering resonance should occur near 3000 G. This field is likely to be unreachable. The reason this technique works in ytterbium is the spin-symmetric scattering channel


Figure 6.5: CAD renderings of the bitter coil and a photo of everything fully assembled. Because the current-carrying brass rods snapped easily, we instead welded a rectangular-profile copper bar to the brass bottom layer, which can be seen in (c) attached to the low-gauge welding cable. The entire coil is potted in epoxy to reduce the risk of water leaks.

has a very large s-wave scattering length $a_s^+ = 3300 a_0$, much larger than strontium, for which both a_s^+ and a_s^- are near 100 a_0 . Stated another way, in ytterbium there is a near-threshold bound state supported by the closed channel scattering potential which can be pushed around slightly to yield a scattering resonance.

That all said, in discussions with theory collaborators, there may exist a nearthreshold p-wave resonance in strontium.⁶ If this indeed the case, we want the ability to still apply moderately high magnetic fields to access this resonance. This motivated the decision to build Bitter coil electromagnets for SrII.

The design for our coils was based loosely on the design in [131,132], with additional tips from [133] and discussions with the CAVS group at NIST. A CAD rendering is shown in Fig. 6.5. The coil consists of two concentric stacks of copper disks with insulating polyester spacers sandwiched between each layer. Cooling water passes through eight vertical cutouts passing through all layers, and then horizontally between each layer in space cutout from the polyester spacer. The final layer at the bottom of the coil is slightly thicker and machined from brass, which provides better mechanical rigidity than copper for clamping the coil together, but has a higher conductivity than steel. The high-current input and return happens on

⁶Paul Julienne, *private comm*.

vertical copper bars welded to the bottom of the brass clamp. The original design followed ideas from Refs. [131, 132], where current was brought on a threaded brass rod passing through one of the vertical water cooling channels. However, we had difficulty tightening the coils to the point where they would seal without snapping the brass rod.⁷

Our primary modification from the design in Refs. [131, 132] is to add a second inner coil, and simplify the water distribution manifold. We constructed a dualchamber stainless steel reservoir which attaches to the copper coils; water enters the bottom chamber, passes through the coil, and returns to the top chamber before returning to the chilled water closed loop. The inner structure of this can be seen in the exploded view of Fig. 6.5(c).

Imaging system. In an imaging system, aberrations arise when light deviates significantly from the paraxial approximation, as is the case for high-NA optics. They are classified into different categories depending on the physical origin (defocus, spherical aberration, coma, etc.), and the goal of a good optical design is to engineer these aberrations to cancel to a high degree. If this can be accomplished, the resulting image will be limited by diffractive effects only.

The high-resolution imaging system designed for SrII is based off the principles in [134]. Using Zemax or other ray-tracing design software, one creates a "best-form" lens for diffraction-limited imaging at a particular wavelength. An iterative process follows: replace one lens with the closest stock-catalog lens available, and re-optimize the design with the remaining surfaces. Repeat this process until you are left with an

⁷Our Bitter coil design was made by Alex Hesse, a visiting student from Heidelberg, and the painstaking construction was taken over by Peter after Alex's departure. Getting these coils to seal has been an ongoing nightmare, as we first thought to use teflon gaskets to both seal and provide electrical insulation. This plan did not work, and we eventually reverted back to the polyester spacers with laser-cut rubber gaskets as in [131]. We also found that machining a clamp plate, and using the 1/4-20 through-hole mounts in the water manifold to pre-tension the coil and check for leaks, has been crucial to preventing the threaded rods from snapping.



Figure 6.6: High-NA imaging objective for SrII, corrected at 461 nm for aberrations introduced by the vacuum viewport. It has an image-space NA of 0.29, and a diffraction-limited spot size of $1.9 \,\mu\text{m}$. It is constructed out of catalog 2" lenses, and optimized with Zemax as described in the text.

objective made from stock lenses. Surprisingly, this seems to converge to the same set of lenses for a variety of wavelengths, with the precise lens spacing set by details of the imaging wavelength and the optical thickness of the final vacuum viewport. The final design for our high-resolution, infinity-corrected objective can be found in Fig. 6.6. It has an effective focal length of 66 mm, and an image space NA of 0.29. At 461 nm, this corresponds to a diffraction-limited spotsize of $\approx 1.9 \,\mu\text{m}^8$ The entire objective is set in a 2" lens tube, with cylindrical spacers machined from aluminum.

6.4 Computer Control

Historically, labs that are a part of the laser cooling and trapping group in the JQI have used an amalgamation of LabView VIs to visualize and program experi-

⁸There are, of course, many metrics one can use to define resolution. Here, we are taking the characteristic spotsize given by the first null of the airy disk, the so-called Raleigh criteria. The airy radius, then is half of this.

mental sequences. Collectively known as SetList, this computer control system has been previously described in Sec. 2.4 and in Ref. [19]. The original SetList was rewritten circa 2014 to utilize some of LabView's object-oriented programming structures, and the newer version was much improved. However, we find that we have pushed SetList far beyond the intended capabilities of LabView. The code has grown unwieldy and undocumented, the GUI is frustratingly slow to respond to user input, and there are intermittent bugs which can be time-consuming to address and solve. In recent years, with the growing popularity of Python, there has been movement in the AMO community to adopt control and analysis software written in Python. On the analysis side, there is a wealth of well-developed fitting, plotting, and image analysis packages. On the experimental control front, the community seems to have coalesced around three candidates: the "Sandia" control system⁹, LabScript Suite [135], and ARTIQ [136].

The Sandia system was first developed at Sandia by Peter Maunz, and later extended by ion trapping groups at the University of Maryland. The core of the system is based on an Opall Kelly FPGA; a GUI written in Python/Qt allows the user to compose "pulse programs," which utilize a subset of the Python language, and instruct the FPGA on deterministic timing. The pulse programs support data acquisition and conditional branching, which makes the construction of complex, real-time behaviors possible. The FPGA is directly connected to several "fast" peripherals such as DACs, ADCs, and DDSs, as well as digital I/O. All data is reported back to the computer and plotted in real-time for easy analysis. For a more detailed description, see the thesis by Kenneth Wright [137].

LabScript Suite is another Python-based system developed at Monash University. It utilizes a very similar timing architecture to the one employed by SetList,¹⁰

 $^{^{9}\}mathrm{To}$ my knowledge, this system has not appeared in published form, but has been adopted by several ion trapping groups at the University of Maryland

¹⁰More specifically, it uses a PulseBlaster or its generic cousin the "PineBlaster" to deliver pseudo-

but uses a Pythonic scripting language instead of LabView for programming of sequences. LabScript also includes an image analysis pipeline, which allows the user to write custom analysis scripts which interact seamlessly with images acquired during the course of an experiment. Files are stored in the popular HDF5 format, which can then be interfaced with other analysis programs if desired. LabScript Suite has been adopted by a few neutral atom labs at the JQI and NIST, and is being actively developed by those groups.

Finally, ARTIQ is a "next generation control system for quantum information experiments."¹¹ It is being actively and heavily developed by a professional software company, and the accompanying hardware is being developed by a group at Warsaw University. Like the two previous solutions, ARTIQ is Python-based but aims to be more agnostic to particular hardware toolchains. For comparison, the Sandia system is tied to the Opall Kelly hardware and frontpanel software drivers for communication, and LabScript Suite assumes a particular control architecture. The stated goals of the ARTIQ ecosystem are to serve as a backbone "operating system" for quantum information and other disciplines with stringent timing requirements. Careful thought and planning has gone into making it robust, well-tested, and extensible. And the fact that it is developed by professional software and hardware engineers means we physicists can leverage their expertise in design architectures while focusing on the end physics applications.

While each system has its upsides and pitfalls, we have decided to deploy Labscript Suite in the SrII laboratory.¹² We have also been working to develop a

clock triggers to a heterogeneous array of peripherals. These peripherals are programmed at the top of a given experimental cycle, and update their outputs to the next state when triggered to do so by the PulseBlaster.

¹¹See the project GitHub page for more information: https://github.com/m-labs/artiq

¹²Our original intention was to deploy ARTIQ, but I seemed to be the only person pushing this path. Given the collective experience with Labscript Suite at the JQI, and the relative newness of ARTIQ, keeping with a known solution seemed prudent. ARTIQ certainly has a vastly more powerful real-time control infrastructure, but as it stands we are without a compelling use case.

replacement for the PulseBlaster, using the PYNQ development platform. Eventually, this might be extended into a larger ecosystem of custom instrumentation electronics.

6.5 Outlook

As the experiment transitions fully to the new lab space, we have begun to build the capability for incorporating Rydberg-dressed interactions in the clock state. Based on the proposal in Ref. [138], the idea is to generate spin squeezing on the clock transition to reduce the quantum projection noise, which will soon limit atomic clock performance.

Another idea which we would like to pursue is to generate uniform 1D potentials with a blue-detuned dipole trap imaged onto the atoms with a digital micromirror array (DMD). There are a number of interesting proposals to study spin physics in a 1D box potential [6,7]. These proposals all rely on the highly groupsymmetric Hamiltonian accessible in a quantum simulation with strontium. Because the momentum eigenstates of a particle in a box all have identical overlap integrals, the interaction energy does not depend on the exact levels that are occupied. This adds an additional permutation symmetry between momentum eigenstates in a uniform square-well potential.

If one thinks of the external energy eigenstates of the box potential as a synthetic lattice dimension, we have changed the spin-spin interaction energy scale from nearest-neighbor J^2/U (with J the hopping energy and U the onsite interaction energy) to all-to-all coupling with constant strength. This model is exactly solvable and can be implemented with thermal atoms, which is interesting from the point of

And because the development thus far has targeted ion trapping groups (who have supplied the majority of the funding for both ARTIQ hardware and software), concise programming examples tailored to the way neutral-atom experiments are structured have not been created. Without some level of institutional experience, or a very high level of python-fu, the learning curve was going to be steep.

view of benchmarking a strontium quantum simulator.

One other proposal for the 1D box potential enables efficient estimation of the eigenspectrum of the nuclear spin density matrix $\hat{\rho}$. The experiment is remarkably simple — a Ramsey pulse sequence is performed on the clock transition, and the dynamics depend on the state of the nuclear spin subspace. Careful analysis of the Ramsey dynamics allows one to extract this information.

Finally, we might revisit the isotope shift measurements discussed in Chapter 5, this time with a magic-wavelength lattice. Norbert Linke is building a Sr^+ experiment at the JQI, and there might be the possibility for combining precision measurements in both neutral and ion systems for a more complete test of the King linearity. We have had some discussions with theory collaborators about how to handle the nonlinearities in ⁸⁷Sr, which will be necessary before moving forward with this research program.

In all, some very exciting times lie ahead for the new strontium experiment. Improvements in the design of the apparatus will enable a richer exploration of the physics accessible with ultracold strontium, and I look forward to seeing what is possible.

Bibliography

- X. Zhang, M. Bishof, S. L. Bromley, C. V. Kraus, M. S. Safronova, P. Zoller, A. M. Rey, and J. Ye, "Spectroscopic observation of su(n)-symmetric interactions in sr orbital magnetism," *Science*, vol. 345, pp. 1467–1473, Sept. 2014.
- [2] A. V. Gorshkov, M. Hermele, V. Gurarie, C. Xu, P. S. Julienne, J. Ye, P. Zoller, E. Demler, M. D. Lukin, and A. M. Rey, "Two-orbital su(n) magnetism with ultracold alkaline-earth atoms," *Nature Physics*, vol. 6, pp. 289–295, Apr. 2010.
- [3] O. Boada, A. Celi, J. I. Latorre, and M. Lewenstein, "Quantum simulation of an extra dimension," *Physical Review Letters*, vol. 108, p. 133001, Mar. 2012.
- [4] A. Celi, P. Massignan, J. Ruseckas, N. Goldman, I. B. Spielman, G. Juzeliunas, and M. Lewenstein, "Synthetic gauge fields in synthetic dimensions," *Physical Review Letters*, vol. 112, p. 043001, Jan. 2014.
- [5] L. F. Livi, G. Cappellini, M. Diem, L. Franchi, C. Clivati, M. Frittelli, F. Levi, D. Calonico, J. Catani, M. Inguscio, and L. Fallani, "Synthetic dimensions and spin-orbit coupling with an optical clock transition," *Physical Review Letters*, vol. 117, p. 220401, Nov. 2016.
- [6] M. E. Beverland, G. Alagic, M. J. Martin, A. P. Koller, A. M. Rey, and A. V. Gorshkov, "Realizing exactly solvable su (n) magnets with thermal atoms," *Phys. Rev. A*, vol. 93, p. 051601, May 2016.
- [7] M. E. Beverland, J. Haah, G. Alagic, G. K. Campbell, A. M. Rey, and A. V. Gorshkov, "Spectrum estimation of density operators with alkaline-earth atoms," *Physical Review Letters*, vol. 120, p. 025301, Jan. 2018.
- [8] C. W. Bauschlicher, Jr., S. R. Langhoff, R. L. Jaffe, and H. Partridge, "Theoretical electric quadrupole transition probabilities for ca, sr and ba," *Journal* of *Physics B Atomic Molecular Physics*, vol. 17, pp. L427–L431, July 1984.
- [9] L. R. Hunter, W. A. Walker, and D. S. Weiss, "Observation of an atomic starkelectric-quadrupole interference," *Physical Review Letters*, vol. 56, pp. 823– 826, Feb. 1986.

- [10] S. B. Nagel, P. G. Mickelson, A. D. Saenz, Y. N. Martinez, Y. C. Chen, T. C. Killian, P. Pellegrini, and R. Côté, "Photoassociative spectroscopy at long range in ultracold strontium," *Physical Review Letters*, vol. 94, p. 083004, Mar. 2005.
- [11] M. Yasuda, T. Kishimoto, M. Takamoto, and H. Katori, "Photoassociation spectroscopy of ⁸⁸sr : Reconstruction of the wave function near the last node," *Phys. Rev. A*, vol. 73, p. 011403, Jan. 2006.
- [12] J. E. Sansonetti and G. Nave, "Wavelengths, transition probabilities, and energy levels for the spectrum of neutral strontium (sr i)," *Journal of Physical* and Chemical Reference Data, vol. 39, pp. 033103–033103, Sept. 2010.
- [13] S. Stellmer, F. Schreck, and T. C. Killian, DEGENERATE QUANTUM GASES OF STRONTIUM, ch. CHAPTER 1, pp. 1–80. 2014.
- [14] H. J. Metcalf and P. van der Straten, Laser Cooling and Trapping. Springer New York, 1999.
- [15] M. S. Safronova, S. G. Porsev, U. I. Safronova, M. G. Kozlov, and C. W. Clark, "Blackbody-radiation shift in the sr optical atomic clock," *Phys. Rev. A*, vol. 87, p. 012509, Jan. 2013.
- [16] S. Stellmer, F. Schreck, and T. C. Killian, Degenerate Quantum Gases of Strontium, pp. 1–80. World Scientific Publishing Co, Mar. 2014.
- [17] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, "Optical dipole traps for neutral atoms," Advances in Atomic Molecular and Optical Physics, vol. 42, pp. 95–170, 2000.
- [18] T. H. Loftus, T. Ido, M. M. Boyd, A. D. Ludlow, and J. Ye, "Narrow line cooling and momentum-space crystals," *Phys. Rev. A*, vol. 70, p. 063413, Dec. 2004.
- [19] D. S. Barker, Degenerate Gases of Strontium for Studies of Quantum Magnetism. PhD thesis, University of Maryland, 2016.
- [20] B. J. Reschovsky, Studies of Ultracold Strontium Gasses. PhD thesis, University of Maryland, 2017.
- [21] X. Xu, T. H. Loftus, J. L. Hall, A. Gallagher, and J. Ye, "Cooling and trapping of atomic strontium," *Journal of the Optical Society of America B Optical Physics*, vol. 20, pp. 968–976, May 2003.
- [22] M. Norcia, New tools for precision measurement and quantum science with narrow linewidth optical transitions. PhD thesis, University of Colorado, 2017.
- [23] T. Yang, K. Pandey, M. Srinivas Pramod, F. Leroux, C. C. Kwong, E. Hajiyev, Z. Y. Chia, B. Fang, and D. Wilkowski, "A high flux source of cold strontium atoms," *European Physical Journal D*, vol. 69, p. 226, Oct. 2015.

- [24] R. Senaratne, S. V. Rajagopal, Z. A. Geiger, K. M. Fujiwara, V. Lebedev, and D. M. Weld, "Effusive atomic oven nozzle design using an aligned microcapillary array," *Review of Scientific Instruments*, vol. 86, p. 023105, Feb. 2015.
- [25] M. Schioppo, N. Poli, M. Prevedelli, S. Falke, C. Lisdat, U. Sterr, and G. M. Tino, "A compact and efficient strontium oven for laser-cooling experiments," *Review of Scientific Instruments*, vol. 83, pp. 103101–103101–6, Oct. 2012.
- [26] I. Nosske, L. Couturier, F. Hu, C. Tan, C. Qiao, J. Blume, Y. Jiang, P. Chen, and M. Weidemüller, "Two-dimensional magneto-optical trap as a source for cold strontium atoms," *Physical Review A*, vol. 96, no. 5, p. 053415, 2017.
- [27] E. M. Bridge, J. Millen, C. S. Adams, and M. P. A. Jones, "A vapor cell based on dispensers for laser spectroscopy," *Review of Scientific Instruments*, vol. 80, pp. 013101–013101–3, Jan. 2009.
- [28] C. P. Pearman, C. S. Adams, S. G. Cox, P. F. Griffin, D. A. Smith, and I. G. Hughes, "Polarization spectroscopy of a closed atomic transition: applications to laser frequency locking," *Journal of Physics B Atomic Molecular Physics*, vol. 35, pp. 5141–5151, Dec. 2002.
- [29] B. Saxberg, B. Plotkin-Swing, and S. Gupta, "Active stabilization of a diode laser injection lock," *Review of Scientific Instruments*, vol. 87, p. 063109, June 2016.
- [30] E. D. Black, "An introduction to pound-drever-hall laser frequency stabilization," American Journal of Physics, vol. 69, pp. 79–87, Jan. 2001.
- [31] E. C. Cook, P. J. Martin, T. L. Brown-Heft, J. C. Garman, and D. A. Steck, "High passive-stability diode-laser design for use in atomic-physics experiments," *Review of Scientific Instruments*, vol. 83, pp. 043101–043101–9, Apr. 2012.
- [32] C. E. Wieman and L. Hollberg, "Using diode lasers for atomic physics," *Review of Scientific Instruments*, vol. 62, pp. 1–20, Jan. 1991.
- [33] K. R. Vogel, Laser cooling on a narrow atomic transition and measurement of the two-body cold collision loss rate in a strontium magneto-optical trap. PhD thesis, UNIVERSITY OF COLORADO AT BOULDER, 1999.
- [34] S. Stellmer and F. Schreck, "Reservoir spectroscopy of 5s5p ${}^{3}p_{2}$ 5snd ${}^{3}d_{1,2,3}$ transitions in strontium," *Phys. Rev. A*, vol. 90, p. 022512, Aug. 2014.
- [35] P. G. Mickelson, Y. N. Martinez de Escobar, P. Anzel, B. J. De Salvo, S. B. Nagel, A. J. Traverso, M. Yan, and T. C. Killian, "Repumping and spectroscopy of laser-cooled sr atoms using the (5s5p)³p₂-(5s4d)³d₂ transition," *Journal of Physics B Atomic Molecular Physics*, vol. 42, p. 235001, Dec. 2009.

- [36] K. Singer, S. Jochim, M. Mudrich, A. Mosk, and M. Weidemüller, "Low-cost mechanical shutter for light beams," *Review of Scientific Instruments*, vol. 73, pp. 4402–4404, Dec. 2002.
- [37] D. R. Leibrandt and J. Heidecker, "An open source digital servo for atomic, molecular, and optical physics experiments," *Review of Scientific Instruments*, vol. 86, p. 123115, Dec. 2015.
- [38] "Setlist: Labview control for bec experiments." Available at https://github. com/JQIamo/SetList.
- [39] Open-source projects from the JQI, hosted by Github. Avalable at https: //github.com/JQIamo/eurocard-backplane/.
- [40] K. G. Libbrecht and J. L. Hall, "A low-noise high-speed diode laser current controller," *Review of Scientific Instruments*, vol. 64, pp. 2133–2135, Aug. 1993.
- [41] C. J. Erickson, M. Van Zijll, G. Doermann, and D. S. Durfee, "An ultrahigh stability, low-noise laser current driver with digital control," *Review of Scientific Instruments*, vol. 79, pp. 073107–073107, July 2008.
- [42] C. M. Seck, P. J. Martin, E. C. Cook, B. C. Odom, and D. A. Steck, "Noise reduction of a libbrecht-hall style current driver," *Review of Scientific Instruments*, vol. 87, p. 064703, June 2016.
- [43] "Laser current controller." Avalable at https://github.com/JQIamo/ current-controller/.
- [44] S. C. Doret, "Simple, low-noise piezo driver with feed-forward for broad tuning of external cavity diode lasers," *Review of Scientific Instruments*, vol. 89, no. 2, p. 023102, 2018.
- [45] J. Appel, A. MacRae, and A. I. Lvovsky, "A versatile digital ghz phase lock for external cavity diode lasers," *Measurement Science and Technology*, vol. 20, p. 055302, May 2009.
- [46] Avalable at https://github.com/JQIamo/beatnote-pll.
- [47] M. Prevedelli, T. Freegarde, and T. Haensch, "Phase locking of grating tuned diode lasers," vol. 61, pp. S241–S248, 01 1995.
- [48] N. C. Pisenti, A. Restelli, B. J. Reschovsky, D. S. Barker, and G. K. Campbell, "An ultra-low noise, high-voltage piezo-driver," *Review of Scientific In*struments, vol. 87, p. 124702, Dec. 2016.
- [49] "Bipolar mosfet controller." Available at https://github.com/JQIamo/ bipolar-mosfet-board.

- [50] M. M. Boyd, T. Zelevinsky, A. D. Ludlow, S. Blatt, T. Zanon-Willette, S. M. Foreman, and J. Ye, "Nuclear spin effects in optical lattice clocks," *Phys. Rev. A*, vol. 76, p. 022510, Aug. 2007.
- [51] T. L. Nicholson, S. L. Campbell, R. B. Hutson, G. E. Marti, B. J. Bloom, R. L. McNally, W. Zhang, M. D. Barrett, M. S. Safronova, G. F. Strouse, W. L. Tew, and J. Ye, "Systematic evaluation of an atomic clock at 2 × 10⁻¹⁸ total uncertainty," *Nature Communications*, vol. 6, p. 6896, Apr. 2015.
- [52] I. Ushijima, M. Takamoto, M. Das, T. Ohkubo, and H. Katori, "Cryogenic optical lattice clocks," *Nature Photonics*, vol. 9, pp. 185–189, Mar. 2015.
- [53] S. L. Campbell, R. B. Hutson, G. E. Marti, A. Goban, N. Darkwah Oppong, R. L. McNally, L. Sonderhouse, J. M. Robinson, W. Zhang, B. J. Bloom, and J. Ye, "A fermi-degenerate three-dimensional optical lattice clock," *Science*, vol. 358, pp. 90–94, Oct. 2017.
- [54] M. Schioppo, R. C. Brown, W. F. McGrew, N. Hinkley, R. J. Fasano, K. Beloy, T. H. Yoon, G. Milani, D. Nicolodi, J. A. Sherman, N. B. Phillips, C. W. Oates, and A. D. Ludlow, "Ultrastable optical clock with two cold-atom ensembles," *Nature Photonics*, vol. 11, pp. 48–52, Jan. 2017.
- [55] N. Huntemann, C. Sanner, B. Lipphardt, C. Tamm, and E. Peik, "Single-ion atomic clock with 3×10^{-18} systematic uncertainty," *Physical Review Letters*, vol. 116, p. 063001, Feb. 2016.
- [56] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt, "Optical atomic clocks," *Reviews of Modern Physics*, vol. 87, pp. 637–701, Apr. 2015.
- [57] T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M. J. Martin, L. Chen, and J. Ye, "A sub-40-mhz-linewidth laser based on a silicon single-crystal optical cavity," *Nature Photonics*, vol. 6, pp. 687–692, Oct. 2012.
- [58] D. G. Matei, T. Legero, S. Häfner, C. Grebing, R. Weyrich, W. Zhang, L. Sonderhouse, J. M. Robinson, J. Ye, F. Riehle, and U. Sterr, "1.5 μ m lasers with sub-10 mhz linewidth," *Physical Review Letters*, vol. 118, p. 263202, June 2017.
- [59] M. M. Boyd, High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard for Frequency and Forbidden. PhD thesis, University of Colorado, 2007.
- [60] 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: Lasers and Optics*, vol. 31, pp. 97–105, June 1983.

- [61] K. Numata, A. Kemery, and J. Camp, "Thermal-noise limit in the frequency stabilization of lasers with rigid cavities," *Physical Review Letters*, vol. 93, p. 250602, Dec. 2004.
- [62] B. J. Reschovsky, B. P. Ruzic, H. Miyake, N. C. Pisenti, P. S. Julienne, and G. K. Campbell, "Narrow-line photoassociation spectroscopy and mass-scaling of bosonic strontium," arXiv e-prints, Aug. 2018.
- [63] L. Chen, J. L. Hall, J. Ye, T. Yang, E. Zang, and T. Li, "Vibration-induced elastic deformation of fabry-perot cavities," *Phys. Rev. A*, vol. 74, p. 053801, Nov. 2006.
- [64] M. Notcutt, L.-S. Ma, J. Ye, and J. L. Hall, "Simple and compact 1-hz laser system via an improved mounting configuration of a reference cavity," *Optics Letters*, vol. 30, pp. 1815–1817, July 2005.
- [65] F. Vogl, "Building a strontium clock laser," Master's thesis, University of Innsbruck, 2013.
- [66] S. A. Webster, M. Oxborrow, and P. Gill, "Vibration insensitive optical cavity," *Phys. Rev. A*, vol. 75, p. 011801, Jan. 2007.
- [67] J. Millo, D. V. Magalhães, C. Mandache, Y. Le Coq, E. M. L. English, P. G. Westergaard, J. Lodewyck, S. Bize, P. Lemonde, and G. Santarelli, "Ultra-stable lasers based on vibration insensitive cavities," *Phys. Rev. A*, vol. 79, p. 053829, May 2009.
- [68] M. J. Martin. PhD thesis, University of Colorado, 2013.
- [69] "The fundamentals of fft-based signal analysis and measurement in labview and labwindows/cvi." http://www.ni.com/white-paper/4278/en/. Accessed: 2018-05-15.
- [70] K. Möhle, *Piezoelectrically Tunable Optical Cavities for the Gravitational Wave Detector LISA*. PhD thesis, Humboldt-Universität zu Berlin, 2013.
- [71] C. Hagemann, C. Grebing, C. Lisdat, S. Falke, T. Legero, U. Sterr, F. Riehle, M. J. Martin, and J. Ye, "Ultrastable laser with average fractional frequency drift rate below 5×10⁻¹⁹/s," *Optics Letters*, vol. 39, p. 5102, Sept. 2014.
- [72] Y. Y. Jiang, A. D. Ludlow, N. D. Lemke, R. W. Fox, J. A. Sherman, L.-S. Ma, and C. W. Oates, "Making optical atomic clocks more stable with 10⁻¹⁶-level laser stabilization," *Nature Photonics*, vol. 5, pp. 158–161, Mar. 2011.
- [73] P. R. Saulson, "Thermal noise in mechanical experiments," Phys. Rev. D, vol. 42, pp. 2437–2445, Oct. 1990.
- [74] H. B. Callen and R. F. Greene, "On a theorem of irreversible thermodynamics," *Physical Review*, vol. 86, pp. 702–710, June 1952.

- [75] Y. Levin, "Internal thermal noise in the ligo test masses: A direct approach," *Phys. Rev. D*, vol. 57, pp. 659–663, Jan. 1998.
- [76] G. Di Domenico, S. Schilt, and P. Thomann, "Simple approach to the relation between laser frequency noise and laser line shape," *Appl. Opt.*, vol. 49, p. 4801, Sept. 2010.
- [77] T. Day, E. K. Gustafson, and R. L. Byer, "Sub-hertz relative frequency stabilization of two-diode laser-pumped nd:yag lasers locked to a fabry-perot interferometer," *IEEE Journal of Quantum Electronics*, vol. 28, pp. 1106–1117, Apr. 1992.
- [78] J. L. Hall, Stabilizing Lasers for Applications in Quantum Optics, vol. 12, pp. 273 – 284. Berlin, Heidelberg: Springer Berlin Heidelberg, 1986.
- [79] J. Ye and M. Martin, "High Precision Laser Stabilization via optical cavities," in Optical Coatings and Thermal Noise in Precision Measurement (G. Harry, T. Bodiya, and R. DeSalvo, eds.), ch. 15, pp. 237–258, Cambridge University Press, Feb. 2012.
- [80] J.-H. Carstens, Stabilization of a Clock Laser for High Precision Spectroscopy on Ultracold Quantum Gases. PhD thesis, Universität Hamburg, May 2011.
- [81] J. I. Thorpe, K. Numata, and J. Livas, "Laser frequency stabilization and control through offset sideband locking to optical cavities," *Optics Express*, vol. 16, p. 15980, Sept. 2008.
- [82] A. D. Ludlow, The Strontium Optical Lattice Clock: Optical Spectroscopy with Sub-Hertz Accuracy. PhD thesis, University of Colorado, 2008.
- [83] M. McDonald, B. H. McGuyer, G. Z. Iwata, and T. Zelevinsky, "Thermometry via light shifts in optical lattices," *Phys. Rev. Lett.*, vol. 114, p. 023001, Jan 2015.
- [84] E. Riedle, S. H. Ashworth, J. T. Farrell, Jr., and D. J. Nesbitt, "Stabilization and precise calibration of a continuous-wave difference frequency spectrometer by use of a simple transfer cavity," *Review of Scientific Instruments*, vol. 65, pp. 42–48, Jan. 1994.
- [85] P. Bohlouli-Zanjani, K. Afrousheh, and J. D. D. Martin, "Optical transfer cavity stabilization using current-modulated injection-locked diode lasers," *Re*view of Scientific Instruments, vol. 77, pp. 093105–093105, Sept. 2006.
- [86] D. Kohlheyer, J. C. T. Eijkel, A. van den Berg, and R. B. M. Schasfoort, "Miniaturizing free-flow electrophoresis – a critical review," *ELEC-TROPHORESIS*, vol. 29, pp. 977–993, mar 2008.

- [87] Y. Temiz, R. D. Lovchik, G. V. Kaigala, and E. Delamarche, "Lab-on-achip devices: How to close and plug the lab?," *Microelectronic Engineering*, vol. 132, pp. 156–175, jan 2015.
- [88] A. J. Fleming, "A megahertz bandwidth dual amplifier for driving piezoelectric actuators and other highly capacitive loads," *Review of Scientific Instruments*, vol. 80, pp. 104701–104701, Oct. 2009.
- [89] Apex Microtechnology, PA84 power operational amplifier. Rev. V.
- [90] Texas Instruments, LM7171 Very High Speed, High Output Current, Voltage Feedback Amplifier, sep 2014.
- [91] Texas Instruments, DRV2700 Industrial Piezo Driver With Integrated Boost Converter, Apr. 2015.
- [92] Texas Instruments, DRV2700EVM High Voltage Piezo Driver Evaluation Kit, Apr. 2015.
- [93] P. Horowitz and W. Hill, The Art of Electronics, ch. 8.1.1, p. 474. Cambridge University Press, 3 ed., 2015.
- [94] J. Clarke and R. F. Voss, "1/f Noise from Thermal Fluctuations in Metal Films," *Physical Review Letters*, vol. 33, pp. 24–27, July 1974.
- [95] R. F. Voss and J. Clarke, "Flicker (1/f) noise: Equilibrium temperature and resistance fluctuations," Phys. Rev. B, vol. 13, pp. 556–573, Jan. 1976.
- [96] F. Seifert, "Resistor current noise measurements," Tech. Rep. LIGO-T0900200, LIGO, 2009.
- [97] Analog Devices, Dual 12-/14-/16-Bit nanoDAC with 5 ppm/°C On-Chip Reference. Rev. G.
- [98] A. Ryou and J. Simon, "Active Cancellation of Acoustical Resonances with an FPGA FIR Filter," *ArXiv e-prints*, Apr. 2016.
- [99] Design files avalable at https://github.com/JQIamo/hv-piezo-driver.
- [100] "Open-source aom driver." Available at https://github.com/JQIamo/aomdriver.
- [101] "Eurocard backplane." Avalable at https://github.com/JQIamo/eurocardbackplane/.
- [102] M. S. Safronova, D. Budker, D. DeMille, D. F. J. Kimball, A. Derevianko, and C. W. Clark, "Search for new physics with atoms and molecules," *Reviews of Modern Physics*, vol. 90, p. 025008, Apr. 2018.

- [103] W. King, "Comments on the article "peculiarities of the isotope shift in the samarium spectrum"," JOSA, vol. 53, no. 5, pp. 638–639, 1963.
- [104] W. King, *Isotope shifts in atomic spectra*. Springer, 1984.
- [105] V. V. Flambaum, A. J. Geddes, and A. V. Viatkina, "Isotope shift, nonlinearity of king plot and search for nuclear island of stability and new particles," ArXiv e-prints, Sept. 2017.
- [106] J. C. Berengut, D. Budker, C. Delaunay, V. V. Flambaum, C. Frugiuele, E. Fuchs, C. Grojean, R. Harnik, R. Ozeri, G. Perez, and Y. Soreq, "Probing new light force-mediators by isotope shift spectroscopy," *ArXiv e-prints*, Apr. 2017.
- [107] C. Shi, F. Gebert, C. Gorges, S. Kaufmann, W. Nörtershäuser, B. K. Sahoo, A. Surzhykov, V. A. Yerokhin, J. C. Berengut, F. Wolf, J. C. Heip, and P. O. Schmidt, "Unexpectedly large difference of the electron density at the nucleus in the 4p ²P_{1/2,3/2} fine-structure doublet of Ca⁺," *Applied Physics B: Lasers and Optics*, vol. 123, p. 2, Jan. 2017.
- [108] M. R. Kalita, J. A. Behr, A. Gorelov, M. R. Pearson, A. C. DeHart, G. Gwinner, M. J. Kossin, L. A. Orozco, S. Aubin, E. Gomez, M. S. Safronova, V. A. Dzuba, and V. V. Flambaum, "Isotope shifts in the 7 s 8 s transition of francium: Measurements and comparison to ab initio theory," *Phys. Rev. A*, vol. 97, p. 042507, Apr. 2018.
- [109] W. McGrew, X. Zhang, R. Fasano, S. Schäffer, K. Beloy, D. Nicolodi, R. Brown, N. Hinkley, G. Milani, M. Schioppo, *et al.*, "Atomic clock performance enabling geodesy below the centimetre level," *Nature*, p. 1, 2018.
- [110] G. Breit, "Theory of isotope shift," Reviews of Modern Physics, vol. 30, pp. 507–516, Apr. 1958.
- [111] E. Riis, A. G. Sinclair, O. Poulsen, G. W. F. Drake, W. R. C. Rowley, and A. P. Levick, "Lamb shifts and hyperfine structure in ⁶Li⁺ and ⁷Li⁺: Theory and experiment," *Phys. Rev. A*, vol. 49, pp. 207–220, Jan. 1994.
- [112] F. Gebert, Y. Wan, F. Wolf, C. N. Angstmann, J. C. Berengut, and P. O. Schmidt, "Precision isotope shift measurements in calcium ions using quantum logic detection schemes," *Physical Review Letters*, vol. 115, p. 053003, July 2015.
- [113] K. Beloy, A. Derevianko, and W. R. Johnson, "Hyperfine structure of the metastable p3₂ state of alkaline-earth-metal atoms as an accurate probe of nuclear magnetic octupole moments," *Phys. Rev. A*, vol. 77, p. 012512, Jan. 2008.

- [114] G. Ferrari, P. Cancio, R. Drullinger, G. Giusfredi, N. Poli, M. Prevedelli, C. Toninelli, and G. M. Tino, "Precision frequency measurement of visible intercombination lines of strontium," *Physical Review Letters*, vol. 91, p. 243002, Dec. 2003.
- [115] I. Courtillot, A. Quessada-Vial, A. Brusch, D. Kolker, G. D. Rovera, and P. Lemonde, "Accurate spectroscopy of Sr atoms," *European Physical Journal* D, vol. 33, pp. 161–171, May 2005.
- [116] T. Ido, T. H. Loftus, M. M. Boyd, A. D. Ludlow, K. W. Holman, and J. Ye, "Precision spectroscopy and density-dependent frequency shifts in ultracold Sr," *Physical Review Letters*, vol. 94, p. 153001, Apr. 2005.
- [117] H. Liu, F. Gao, W. Ye-Bing, T. Xiao, R. Jie, L. Ben-Quan, X. Qin-Fang, X. Yu-Lin, and C. Hong, "Precision frequency measurement of ¹S₀ — ³P₁ intercombination lines of Sr isotopes," *Chinese Physics B*, vol. 24, p. 013201, Jan. 2015.
- [118] S. Stellmer, *Degenerate quantum gases of strontium*. PhD thesis, University of Innsbruck, 2013.
- [119] W. Ketterle, D.S. Durfee, and D.M. Stamper-Kurn, "Making, probing and understanding bose-einstein condensates," *Proceedings of the International School of Physics "Enrico Fermi"*, vol. 140, no. Bose-Einstein Condensation in Atomic Gases, pp. 67–176, 1999.
- [120] M. Borkowski, P. Morzyński, R. Ciuryło, P. S. Julienne, M. Yan, B. J. DeSalvo, and T. C. Killian, "Mass scaling and nonadiabatic effects in photoassociation spectroscopy of ultracold strontium atoms," *Phys. Rev. A*, vol. 90, p. 032713, Sept. 2014.
- [121] A. V. Taichenachev, V. I. Yudin, C. W. Oates, C. W. Hoyt, Z. W. Barber, and L. Hollberg, "Magnetic field-induced spectroscopy of forbidden optical transitions with application to lattice-based optical atomic clocks," *Physical Review Letters*, vol. 96, p. 083001, Mar. 2006.
- [122] T. Takano, R. Mizushima, and H. Katori, "Precise determination of the isotope shift of ⁸⁸Sr — ⁸⁷Sr optical lattice clock by sharing perturbations," *Applied Physics Express*, vol. 10, p. 072801, July 2017.
- [123] G. Zu Putlitz, "Bestimmung des elektrischen kernquadrupolmomentes des ungeraden stabilen strontium-87-kerns," Zeitschrift fur Physik, vol. 175, pp. 543–552, Oct. 1963.
- [124] R. Ciuryło, E. Tiesinga, S. Kotochigova, and P. S. Julienne, "Photoassociation spectroscopy of cold alkaline-earth-metal atoms near the intercombination line," *Phys. Rev. A*, vol. 70, p. 062710, Dec. 2004.

- [125] J. Mitroy, M. S. Safronova, and C. W. Clark, "Topical review: Theory and applications of atomic and ionic polarizabilities," *Journal of Physics B Atomic Molecular Physics*, vol. 43, p. 202001, Oct. 2010.
- [126] Daniel Barker, private comm., Nov. 2016.
- [127] M. Sefa, J. A. Fedchak, and J. Scherschligt, "Investigations of mediumtemperature heat treatments to achieve low outgassing rates in stainless steel ultrahigh vacuum chambers," *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, vol. 35, p. 041601, jul 2017.
- [128] R. Zhang, Y. Cheng, H. Zhai, and P. Zhang, "Orbital feshbach resonance in alkali-earth atoms," *Physical Review Letters*, vol. 115, p. 135301, Sept. 2015.
- [129] G. Pagano, M. Mancini, G. Cappellini, L. Livi, C. Sias, J. Catani, M. Inguscio, and L. Fallani, "Strongly interacting gas of two-electron fermions at an orbital feshbach resonance," *Physical Review Letters*, vol. 115, p. 265301, Dec. 2015.
- [130] M. Höfer, L. Riegger, F. Scazza, C. Hofrichter, D. R. Fernandes, M. M. Parish, J. Levinsen, I. Bloch, and S. Fölling, "Observation of an orbital interaction-induced feshbach resonance in ¹⁷³yb," *Physical Review Letters*, vol. 115, p. 265302, Dec. 2015.
- [131] D. O. Sabulsky, C. V. Parker, N. D. Gemelke, and C. Chin, "Efficient continuous-duty bitter-type electromagnets for cold atom experiments," *Re*view of Scientific Instruments, vol. 84, pp. 104706–104706–5, Oct. 2013.
- [132] T. Luan, T. Zhou, X. Chen, and Z. Ma, "A modified bitter-type electromagnet and control system for cold atom experiments," *Review of Scientific Instruments*, vol. 85, p. 024701, Feb. 2014.
- [133] J. Johansen, Feshbach and Efimov resonances in a ⁶Li⁻¹³³Cs atomic mixture. PhD thesis, University of Chicago, 2017.
- [134] J. D. Pritchard, J. A. Isaacs, and M. Saffman, "Long working distance objective lenses for single atom trapping and imaging," *Review of Scientific Instruments*, vol. 87, p. 073107, July 2016.
- [135] P. T. Starkey, C. J. Billington, S. P. Johnstone, M. Jasperse, K. Helmerson, L. D. Turner, and R. P. Anderson, "A scripted control system for autonomous hardware-timed experiments," *Review of Scientific Instruments*, vol. 84, pp. 085111–085111–10, Aug. 2013.
- [136] S. Bourdeauducq, R. Jördens, P. Zotov, J. Britton, D. Slichter, D. Leibrandt, D. Allcock, A. Hankin, F. Kermarrec, Y. Sionneau, R. Srinivas, T. R. Tan, and J. Bohnet, "Artiq," 2016.
- [137] K. Wright, Manipulation of the Quantum Motion of Trapped Atomic Ions via Stimulated Raman Transitions. PhD thesis, University of Maryland, 2017.

[138] L. I. R. Gil, R. Mukherjee, E. M. Bridge, M. P. A. Jones, and T. Pohl, "Spin squeezing in a rydberg lattice clock," *Physical Review Letters*, vol. 112, p. 103601, Mar. 2014.