

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Design and construction of a strontium Fermi gas machine for studies in quantum emulation

Permalink

<https://escholarship.org/uc/item/8wm6k0f1>

Author

Lauria, Paul

Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Design and construction of a strontium Fermi gas machine for studies in quantum
emulation

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Physics

by

Paul Lauria

Committee in charge:

Professor Julio Barreiro, Chair
Professor Daniel Arovas
Professor Richard Averitt
Professor Yeshaiah Fainman
Professor Tenio Popmintchev

2022

Copyright

Paul Lauria, 2022

All rights reserved.

The Dissertation of Paul Lauria is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2022

TABLE OF CONTENTS

Dissertation Approval Page	iii
Table of Contents	iv
List of Figures	vii
List of Tables	ix
Preface	x
Acknowledgements	xi
Vita	xiii
Abstract of the Dissertation	xiv
Chapter 1 Introduction	1
1.1 Quantum emulation	1
1.2 Alkaline-Earth atoms	3
1.3 Laser cooling	4
1.4 Quantum degeneracy	7
1.4.1 Thermal gases	8
1.4.2 Degenerate Fermi gases	9
Chapter 2 Light-matter interaction	13
2.1 Two-level atoms	13
2.2 Stimulated Raman effect	16
2.3 Spin-orbit coupling	19
2.3.1 History	19
2.3.2 SO coupling: physical origin	20
2.3.3 SO coupling in ultracold gases	20
2.4 Quantum Hall systems	21
2.4.1 Classical	21
2.4.2 Integer quantum Hall effect	22
2.4.3 Role of topology	23
2.4.4 Fractional Chern insulators	24
Chapter 3 Experimental setup	27
3.1 Lab conditions	27
3.2 Vacuum chamber	28
3.2.1 Chamber design	29
3.2.2 Chamber bake	30
3.3 Strontium source	35

3.4	Laser systems	39
3.4.1	461 nm laser setup	40
3.4.2	689 nm laser setup	44
3.4.3	SM lattice setup	51
3.4.4	481 nm laser	53
3.4.5	1064 nm dipole trap	54
3.4.6	Optical Stern-Gerlach	56
3.4.7	Lift beam laser	58
3.4.8	Laser shutters	59
3.5	Magnetic field control	62
3.5.1	MOT coils (Bitter coils)	62
3.5.2	Shim coils	65
3.6	RF generation	67
3.6.1	RF sources	68
3.6.2	Amplifiers	70
3.7	Experiment control	70
3.7.1	ADwin hardware	70
3.7.2	Qcontrol	71
3.7.3	Digital buffers	73
3.7.4	Laser intensity control	73
Chapter 4	Achieving quantum degeneracy	76
4.1	From 700 K to 2 μ K	76
4.1.1	Blue MOTs: 700 K to mK	76
4.1.2	Red MOT: mK to μ K	78
4.1.3	MOT benchmark	82
4.2	Evaporative cooling: from μ K to nK	84
4.2.1	Scattering considerations	84
4.2.2	Evaporative cooling guidelines	85
4.2.3	Bose-Einstein condensation	87
4.2.4	Degenerate Fermi gas	87
Chapter 5	Spin-momentum lattices	90
5.1	SM lattice preparation	90
5.1.1	Spin polarization	91
5.1.2	Beam alignment	94
5.1.3	Energy level shifting	95
5.1.4	Raman coupling	98
5.1.5	Spin-resolved imaging	105
5.2	Filling the lattice	106
5.2.1	Visualization	107
5.2.2	Population and atomic motion	107
5.2.3	Model	108
5.3	Experimental realization	110

5.3.1	Introduction	111
5.3.2	Implementation	112
5.3.3	Sequence	114
5.3.4	Conclusion and outlook	118
5.3.5	Acknowledgements	120
Chapter A	Science chamber schematics	121
Chapter B	Recessed viewports	125
Appendix C	DDS Driver	129
Appendix D	Digital buffer schematic	134
Appendix E	Rotating wave approximation	143
Appendix F	Optical shutters	145
F.1	Shutter PCB	145
F.2	Shutters UART server code	151
Appendix G	Switchable gain transimpedance amplifier PCB	170
Bibliography	173

LIST OF FIGURES

Figure 1.1.	Sr energy levels	5
Figure 1.2.	Quantum gas schematic	11
Figure 2.1.	Spin-orbit coupling physics	17
Figure 2.2.	Aharonov-Bohm effect and optical lattices	24
Figure 2.3.	SM Lattice setup	26
Figure 3.1.	Render of experiment	30
Figure 3.2.	Mass spectra evolution during our first bake.	34
Figure 3.3.	Bakeout setup	36
Figure 3.4.	Oven clog	37
Figure 3.5.	461 nm laser setup	40
Figure 3.6.	AOsense optics setup	43
Figure 3.7.	689 nm laser setup	46
Figure 3.8.	Top breadboard optical layout	47
Figure 3.9.	Electronically-switchable-gain transimpedance amplifier.	49
Figure 3.10.	Optical layout at experiment	52
Figure 3.11.	481 nm laser single-mode operation	53
Figure 3.12.	ODT layout	54
Figure 3.13.	ODT control circuit	56
Figure 3.14.	OSG setups	57
Figure 3.15.	OSG images	59
Figure 3.16.	Laser shutters	61
Figure 3.17.	Bitter coil construction and design	64
Figure 3.18.	Bitter coil control circuit	66

Figure 3.19.	Shim coil calibration	68
Figure 3.20.	BeagleBone controller for DDS	69
Figure 4.1.	MOT timing diagram	79
Figure 4.2.	Red MOT time evolution	81
Figure 4.3.	Evaporation trajectory	85
Figure 4.4.	Degenerate Fermi gas	86
Figure 4.5.	Trap frequency measurements	88
Figure 5.1.	Spin polarization schematic	92
Figure 5.2.	Beam alignment	93
Figure 5.3.	Lift optimization	96
Figure 5.4.	Energy level calibration	98
Figure 5.5.	Rabi strength calibration	100
Figure 5.6.	Raman interference	103
Figure 5.7.	Spin-resolved imaging	105
Figure 5.8.	SM lattice result	115
Figure 5.9.	Broken-links figure	116
Figure 5.10.	Cross coupling	118
Figure 5.11.	SM lattice theory comparison	119
Figure E.1.	The rotating wave approximation	144

LIST OF TABLES

Table 1.1.	List of degenerate monoatomic gases	6
Table 3.1.	Blue AOM frequencies	41
Table 3.2.	Red AOM frequencies	45
Table 4.1.	MOT beam intensity list	83
Table 5.1.	Rabi coupling strength list	102

PREFACE

It's hard to write down *everything* you need to know in order to do this experiment, because so much information comes from so many different sources. In this thesis, I've attempted to include as much "original" knowledge as I can—things that I've discovered along the way that are *unique* to this experiment—while calling your attention to the many amazing theses, books, and papers that have guided me along the way. If a section seems lacking in detail, it's (hopefully) because I found that the references cited there are sufficient to explain further. One of the most important of these is Simon Stellmer's thesis [1], which covers the basics of strontium and its cooling. A second important resource is Dan Steck's quantum optics notes [2], from which I learned most of the theory; chapters 1-6 are essential reading! Finally, for very basic or standard AMO techniques, like MOTs or power-broadening, I recommend Foot [3] and Demtroeder [4].

ACKNOWLEDGEMENTS

Although I worked by myself for many years on this experiment, that doesn't mean I did it alone. I had a lot of help from friends and family, past and present, who supported me through this very-trying 9-year PhD.

To my family, who supported me throughout this journey: my mom, who was always so confident I would be able to finish, so sure it was the right thing to do, and always willing to help me in any way she could. To my dad, who was always interested in my work, the experiment, and going to baseball games as often as we could get away. And to Kristin, who kept me going through all the stressful days with her radiant positivity and love.

I was so fortunate to be surrounded by San Diego friends—nerds all, but who also knew how to have a good time, and who got me outside the lab to enjoy life. To protect the innocent, I thank only first names here: in alphabetical order, Alex, Bob, Dan, Devin, Luke, Khang, Giulio, Grady, Jake, Mark, Olivia, Peter, Tucker: thank you all so much—you made this experience worthwhile, and helped me learn about so much more than just science.

Thanks to Julio, who worked so hard to provide everything I needed to succeed—from equipment, to lab space, to ideas and advice, and for the time I needed to grow into a scientist. Thanks to Tyler for taking over my experiment—I await your great results!

Thanks to Chenghui Yu, whose early efforts in the lab helped me survive the ensuing years. Thanks to Sean Mossman and Erik Szwed, who for an all-too-brief time helped out with the SM lattice project, providing many stimulating discussions. Thank you Wei-Ting Kuo, whose assistance on the SML project theory was invaluable!

And finally, thanks to the millions of scientists who came before. The bibliography of this thesis has about 160 citations, encompassing the work of thousands. And each of them has thousands of others to thank, and so on, stringing an exponential web of science which connects people around the world, across time and space, all in search of a little

more knowledge. It is an honor and privilege to have played a part in that.

Chapter 5, in part, has been submitted for publication to Physical Review Letters, 2022. The dissertation author was the primary author of this paper.

VITA

- 2011 Bachelor of Science, University of California Santa Barbara
- 2011-2013 CRAIC Technologies, San Dimas, California
- 2014 Master of Science, University of California San Diego
- 2022 Doctor of Philosophy, University of California San Diego

PUBLICATIONS

Paul Lauria, Wei-Ting Kuo, Nigel R. Cooper, Julio T. Barreiro, Experimental realization of a fermionic spin-momentum lattice, *Phys. Rev. Lett.*, in press, 2022.

P. Ovarthaiyapong, L. M. A. Pascal, B. A. Myers, P. Lauria, and A. C. Bleszynski Jayich, High quality factor single-crystal diamond mechanical resonators, *Appl. Phys. Lett.*, 101, 2012.

ABSTRACT OF THE DISSERTATION

Design and construction of a strontium Fermi gas machine for studies in quantum emulation

by

Paul Lauria

Doctor of Philosophy in Physics

University of California San Diego, 2022

Professor Julio Barreiro, Chair

This work details the design, construction, and use of an apparatus which cools atomic strontium to quantum degeneracy. Degeneracy is a baseline condition necessary for experiments in quantum emulation and quantum gas microscopy. The apparatus is used for the creation and study of an atomic spin-momentum lattice, a step towards the ultracold-atom realization of fractional quantum Hall states.

Chapter 1

Introduction

1.1 Quantum emulation

It's easy to forget that we, as a species, have figured out a lot about the world ¹. For instance, we rarely consider the small, rectangular, glass-metal bricks in our pockets, which talk to each other using invisible waves, and which make sounds at us when someone miles away—or tens of miles, or thousands of miles—wants to talk to us. Such everyday miracles just don't register; we only care about the relayed message, rarely considering the medium itself. I won't dispute this perspective, since starry-eyed wonder never paid anybody's rent. But when the news is particularly troubling and the future looks bleak, I find it helpful to put things in perspective by considering how far we've come as a species. Nobody has given us a *manual* to the universe; nobody told us everything is made of “atoms,” or that invisible waves can interact with those atoms and make them obey our will. These are *our own* accomplishments—made over thousands of years by millions of scientists—and to me, that should give us great pride and hope for an even better future.

But it should also come as no surprise that there is much we don't yet understand. The laws governing our universe as we currently understand them [5] are frequently confounding: although we know (many of) the game's rules, that doesn't mean we can

¹At the time of writing, the English version of Wikipedia.org has 6,436,632 articles with over 4 billion words.

always predict the game’s outcome. Nowhere is this more evident than in the realm of the very small. “Quantum mechanics” governs the interactions between atoms, and while that theory has been essentially unchanged since Schrödinger put forth his equation in 1926 [6], the interactions between more than a few tens of atoms cannot be exactly predicted. To do so at least in principle, one must keep track of all possible configurations of a system—but this number grows exponentially with number of atoms. Even setting aside the requisite computational time, memory becomes a real problem. To figure out the dynamics of $N = 40$ interacting spin-1/2 particles requires 4 TB of memory [7], or about 1% of the capability of a modern supercomputer. Doubling the number of particles would require a *thousand trillion* times that: 10^{12} TB, far exceeding what classical computers are ever likely to achieve. Considering that interacting physical systems have on the order of $N = 10^{23}$ atoms, the amount of needed classical computing power becomes unimaginable.

We could throw up our hands at the sheer complexity of this “many-body” problem, but generations of physicists since Schrödinger have made progress. The field is replete with ingenious approximation schemes, such as perturbation theory, density functional theory, mean field theories, Monte Carlo methods, and Green’s functions [8–10]. But they are all approximations, to be applied within certain limits—and it is at those limits we frequently find ourselves when trying to understand problems like high-temperature superconductivity [11], non-equilibrium dynamics [12], frustrated spin systems [13], topological materials [14, 15], and many others.

Enter quantum simulation. In popular culture, this might be colloquially referred to as ‘quantum computing,’ but it is important to distinguish the two. A quantum computer is a more general device, built to solve particular problems for which it has a “quantum advantage” over classical computers [16]—i.e., that the benefit of using a quantum computer is so profound that it justifies the incredible cost and effort required to make one. In contrast, quantum simulators [17] are more-specialized devices, built to learn about hard physics problems, some of which I mentioned earlier. To be concrete, the

definition of a quantum simulator is a “controllable quantum system[s] that can be used to simulate other quantum systems;” [18] or to put it another way: using atoms which we can control, to emulate the behavior of atoms we can’t—and in so doing, learn about the problem in question.

There are myriad approaches to quantum simulation; the platform of this work is cold neutral atoms, but other possible mediums include trapped ions, superconducting qubits, photons and others, and I refer the interested reader to [7]. For the remainder of this work, I will focus on the lab’s effort to build a quantum simulator using cold atoms, specifically with the element strontium.

1.2 Alkaline-Earth atoms

Among the numerous atoms that have been brought to quantum degeneracy (see § 1.3), why did we choose strontium? This is a sharp question, considering that ready-made, tabletop BEC machines are available for purchase, with options for rubidium, cesium and potassium². These well-studied elements are used at many of the world’s top emulation labs, including those at NIST, Munich, and JILA. The quantum emulation of many-body systems does not explicitly depend on the particular atom; indeed, that is the point! Yet the choice of element is more than just what is convenient, but what experiments are enabled by the particular element.

Alkaline-earth like atoms (AEAs)—of which Sr, Yb, and Ca have been brought to degeneracy—all share a two-electron valence structure. Their outer two spin-1/2 electrons can occupy either the antisymmetric singlet state, labeled 1S_0 (one electron spin-up, and the other spin-down) or the symmetric triplet configuration 3P , which is further split by the fine structure interaction into 3 sublevels, $^3P_0, ^3P_1, ^3P_2$. The dipole selection rule $\Delta s = 0$ [3] forbids transitions between states with differing spin quantum number, so we would expect the 1S_0 and 3P energy manifolds to be isolated. However, due to effects of

²See, e.g., www.coldquanta.com

spin-orbit coupling, these states can be thought of as superpositions which include 1P_1 [19], allowing a decay channel to 1S_0 and thus providing “intercombination” transitions between singlet and triplet states. I show in Fig. 1.1 all Sr transitions relevant to this work.

The dipole-forbidden transitions listed in Fig. 1.1 have narrow linewidths, from kHz down to mHz, providing a main selling point of AEAs. These narrow linewidths are leveraged to make the world’s best optical atomic clocks, achieving fractional frequency uncertainty at the 10^{-18} level [20, 21], two orders of magnitude better than contemporary microwave cesium clocks [22]. One reason we care about ultra-precise clock transitions is because they allow the measurement of anything that can perturb them with commensurate precision. There are thus many applications, including to many-body physics, and I refer the interested reader to [23]. Of particular interest to this experiment, the narrow $^1S_0 \rightarrow ^3P_1$ transition enables our μK MOT discussed in Chapter 2, and free-space spin-orbit coupling scheme, discussed in Chapter 5.

1.3 Laser cooling

First and foremost, cold atom simulation requires that your atoms be just that: cold. Very, very cold. Roughly speaking, any atom simulation experiment shows quantum behavior only when their average kinetic energy is comparable to (or, preferably, much lower than) the energy shifts (forces) you are exerting. And the colder the better; since the first sodium atoms were laser cooled in the 1980s [24], there has been a constant drive to reach ever lower temperatures. Cooling bosons to their ground state, in a Bose-Einstein condensate, was first achieved with ^{87}Rb in 1995 [25], followed by the first degenerate Fermi gas of ^{40}K four years later [26]. In some sense, these landmark experiments enabled the field of cold atom simulation, and in the past 20 years much progress has been made in achieving quantum degeneracy in other elements [27]; see Table 1.1 for a complete list.




Figure 1.1. Atomic energy level diagram for fermionic strontium, showing all transitions used in this thesis, including the laser beams used in the experiment (colored and labeled).

This thesis would not be possible without the combined 40+ year effort made by the entire community, enabling us to cool strontium to degeneracy and perform experiments.

Put simply, laser cooling involves exposing atoms to light at very particular frequencies in order to make them stop moving so much. At room temperature, atoms and molecules floating around in the air move surprisingly fast. If you measure the velocities of a random number of these atoms, you will find that *none* of them are moving in exactly the same way; but if you did this measurement many thousands of times, and then plotted the squares of the measured velocities against the probability of finding the

Table 1.1. Single elements made into a Bose-Einstein condensate, as of January 2022. A star indicates a Fermi gas has also been achieved. In the adjacent references, BEC creation is listed first, and DFG second. Note that the first Sr BECs were achieved simultaneously by the Innsbruck and Rice groups.

Note	Species
Not laser cooled	H [28]
Alkalis	Rb [25], Na [29], Li* [30] [31], K* [32][26], Cs [33]
Noble gas	He* [34][35]
Highly magnetic	Cr [36], Dy* [37],[38], Er* [39][40], Tm [41]
Alkaline-Earth like	Yb* [42][43], Ca [44], Sr* [45, 46][47]

atom at those velocities, you’d start to notice they follow a very particular distribution. This is the Maxwell-Boltzmann probability distribution, and the most probable speed is $v = \sqrt{3k_B T/m}$ where k_B is Boltzmann’s constant, T is the temperature, and m is the particle’s mass [48].

That parameter, T , is of primary importance in cold atom work. While we are all familiar with the colloquial notion of temperature, we must have a more precise understanding in order to meaningfully distinguish -459.6667° F and -459.6668° F. Indeed, one can define T rigorously in terms of thermodynamic quantities [48], but the operational definition that always made the most sense to me, and the one I’ve found most useful, is that it is simply *the parameter that tells you how wide the atomic velocity distribution is*. The smaller the temperature, the narrower the range of possible velocities an atom is likely to have. It is the business of laser cooling to make that distribution as narrow as possible.

While the exact mechanics of the cooling process depend on the involved species, all neutral atom experiments operate some kind of magneto-optical trap (MOT), a standard tool of atomic physics—I refer the reader to [3] for more details. But to make a simple mechanical analogy, we can think of MOT action as simply attaching an atom to six orthogonal springs, each of which damps motion along its axis, pulling a moving atom back to the center. A MOT cannot cool indefinitely, however. When the atoms have cooled

such that their velocity spread is within the driving transition’s linewidth, cooling ceases, and the atoms are then simply making a random walk in momentum space [4]. For a MOT driven with a transition linewidth Γ , a first limit is imposed by the *Doppler temperature*, $k_B T_D = \hbar\Gamma/2$. A second limit—typically much smaller, but more fundamental—relates the amount of momentum acquired when the atom absorbs a single photon. This is the *recoil temperature*, $k_B T_r = (\hbar k)^2/2m$, where $k = 2\pi/\lambda$ is the wavevector of the MOT light and m is the atomic mass.

In the context of trapping atomic Sr, both of these limits are relevant. The maximum MOT capture velocity $v_c \sim \sqrt{T_D}$. In order to capture the (relatively) hot atoms out of an atomic oven, first-stage cooling takes place in the “blue” MOT (bMOT), which operates on the $^1S_0 \rightarrow ^1P_1$ transition at 461 nm, with linewidth $\Gamma = 2\pi \times 30$ MHz ($T_D = 1$ mK). To cool further, we then transfer those atoms to the “red” MOT (rMOT) which operates on the intercombination line $^1S_0 \rightarrow ^3P_1$ at 689 nm, with $\Gamma = 2\pi \times 7.4$ kHz ($T_D = 1\mu\text{K}$), approaching the recoil temperature $T_r = 0.46 \mu\text{K}$. Much has been written about the narrow-line red MOTs of Sr. I encourage the reader to take a close read of the fundamentals [49], the experimental overview [1], and our specific implementation, which takes elements from [50]. I discuss details of our implementation in Section 4.1.2.

1.4 Quantum degeneracy

Cooling to the recoil temperature is not sufficient to perform most quantum simulation. The ultimate limit is achieving *quantum degeneracy*, where the atoms are so cold their wavefunctions begin to overlap, as in the cartoon depiction of Fig. 1.2. It is in this regime where their (in)distinguishability becomes relevant, revealing their bosonic or fermionic nature. Generally speaking, quantum effects become important when the atomic density approaches the de Broglie volume, or $n \approx m/\lambda_T^3 = m(h/p)^{-3}$, where p is the atomic momentum and h is Planck’s constant. For most species cooled so far, the

onset of degeneracy occurs around temperatures of a few tens of nK. Another way of quantifying the onset of degeneracy is when the phase-space density approaches 1. The phase-space density can be expressed as

$$\rho = n\lambda_T^3 = \left(\frac{2\pi\hbar^2}{mk_B T}\right)^{3/2} = N\left(\frac{h\bar{\nu}}{k_B T}\right)^3 \quad (1.1)$$

where the right-most equality holds for a harmonic trap[51] with mean geometric trap frequency $\bar{\nu}$ and atom number N . Whether the gas is classical or quantum-degenerate is then a matter of determining whether $\rho \ll 1$ or $\rho \approx 1$, respectively.

Depending on whether the atom in question obeys bosonic or fermionic statistics, two kinds of degeneracy can be achieved: Bose-Einstein condensation (BEC) for the former, or a degenerate Fermi gas (DFG) for the latter. In this experiment we have achieved both, and I will briefly describe the properties of each.

1.4.1 Thermal gases

To appreciate the nature of quantum gases, it is helpful to first understand the behavior of harmonically trapped thermal gases. In the limit of high temperatures— $\rho \ll 1$, as described above—quantum statistics are unimportant, and the atoms obey the Maxwell-Boltzmann distribution. This says that the probability a particle will occupy a state with energy E_l is

$$f(E_l) = e^{\frac{-E_l - \mu}{k_B T}} \quad (1.2)$$

where μ is the chemical potential, the energy required to add (or remove) a particle from the system. In the Thomas-Fermi (TF) or semi-classical approximation, we regard each atom as occupying a position in phase space (\mathbf{r}, \mathbf{p}) , with density of states $(2\pi\hbar)^{-3}$ [52, 53]. We replace E_l with the system Hamiltonian

$$H(\mathbf{r}, \mathbf{p}) = p^2/2m + V(\mathbf{r}) \quad (1.3)$$

where $V(\mathbf{r}) = 1/2 m\omega^2 r^2$ represents the always-present harmonic trap, m is the particle mass, and ω is the trap frequency. We can then determine the momentum distribution by multiplying by the density of states and integrating over r :

$$n(\mathbf{p}) = \frac{1}{(2\pi\hbar)^3} \int d\mathbf{r} f(H(\mathbf{r}, \mathbf{p})) = \frac{N}{(2\pi\hbar)^{3/2} \sigma^3} e^{-p^2/2\sigma^2} \quad (1.4)$$

with $\sigma = mk_B T$ defining the standard deviation of the resulting Gaussian. The position distribution can be found by a similar integration over \mathbf{p} and is likewise Gaussian (with standard deviation $\sigma_r = k_B T/m\omega^2$).

1.4.2 Degenerate Fermi gases

If the phase-space density $\rho \approx 1$, the quantum statistics of the atoms become important. An excellent in-depth treatment of degenerate Fermi gases (DFGs) is given in Brian DeMarco's thesis [53]; I summarize only key results here. Fermions (such as ^{87}Sr) obey the Fermi-Dirac distribution with chemical potential μ :

$$f(E_l) = \frac{1}{e^{\frac{E_l - \mu}{k_B T}} + 1} \quad (1.5)$$

In the limit of $T \gg \hbar\omega$, with ω given in Eqn. 1.3, the density of states of a 3D harmonic oscillator is [52]

$$g(E_l) = \frac{E_l^2}{2(\hbar\omega)^3} \quad (1.6)$$

As $T \rightarrow 0$, this distribution becomes a step function with a discontinuity at the Fermi energy E_F , which is the energy below which all states are occupied. The Fermi energy can

be found indirectly by determining the total number of particles at $T = 0$:

$$N = \int_0^{E_F} g(\epsilon) f(\epsilon)|_{T=0} d\epsilon \rightarrow E_F = \hbar\omega(6N)^{1/3} \quad (1.7)$$

In benchmarking DFG's, we typically consider the ratio T/T_F , with $T_F \equiv E_F/k_B$ being the Fermi temperature. Most Fermi gas experiments readily reach $T/T_F \approx 0.2 - 0.3$, and with serious effort, very deeply degenerate gases below 0.1 can be achieved; one noteworthy example includes a ${}^6\text{Li}$ system reporting $T/T_F = 0.059(5)$ [54]. With ${}^{87}\text{Sr}$, the best on record include Stellmer reporting 0.10(1) with a 10-state mixture [55], and more recently, Jun Ye's group at JILA reporting 0.07 [56]. At such temperatures, one can find the Pauli exclusion principle in action: atoms occupying states far below the Fermi surface cannot be excited by light, because higher-momentum states are already filled—as recently observed in three landmark experiments [57].

The momentum distribution of a DFG is found by making the same Thomas-Fermi approximation that led to Eqn. 1.4. Integrating the Fermi-Dirac distribution Eqn. 1.5, one finds [53]:

$$n(\mathbf{p}) = -\frac{1}{(2\pi)^{3/2}\hbar^3} \left(\frac{k_B T}{m\omega^2}\right)^{3/2} \text{Li}_{3/2}\left(-\zeta \exp\left(\frac{-p^2}{2mk_B T}\right)\right) \quad (1.8)$$

where Li_s refers to the polylogarithm function of order s , and $\zeta = e^{\mu/k_B T}$ is the fugacity. It is important to point out that real experimental images are acquired by integrating along the direction of an absorption imaging beam, after the atoms have been released and undergo some “time of flight.” We can reconstruct the in-situ momentum distribution, however: after time t of free expansion, each dimension i is simply rescaled as $p_i = p_{i,0} \sqrt{1 + (\omega_i t)^2}$ [58]. At long time of flight ($\omega_i t \gg 1$), the gas appears isotropic³. At earlier times, the aspect ratio reflects any asymmetry in the trapping strengths along the three trapping

³Even at long TOF, the aspect ratio may in fact differ from unity due to the presence of interactions; see for example [56].

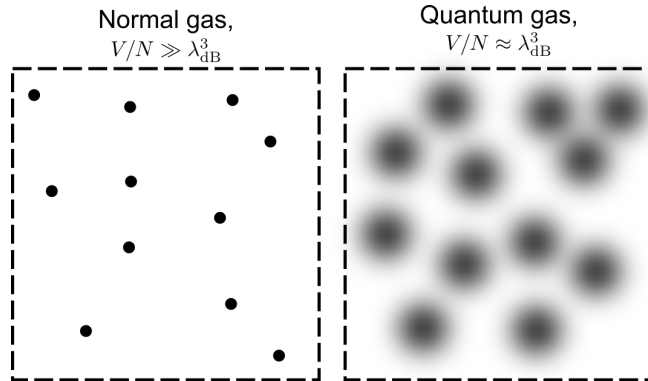


Figure 1.2. The difference between a “normal” gas and a quantum gas, in cartoon form. A quantum gas is so dense the individual wavefunctions begin to overlap.

dimensions, with the most-tightly-trapped dimension expanding fastest—a fact nicely explained by the uncertainty principle (small $\Delta x \rightarrow$ large Δp) [58].

Integrating along the imaging direction (taken to be z) entails expanding the polylog series and integrating term-by-term. For the purposes of fitting, it is actually more convenient to operate with with a density $n(\mathbf{r})$, so that one does not need to know the conversion from pixel-to-momentum (which must be calibrated through some physical process), while the pixel-to-distance conversion is determined by the magnification and pixel size, both of which are usually known. The result is

$$n(x, y) = -\frac{m(k_B T)^2}{(2\pi)^{3/2} \hbar^3 \sqrt{(1 + (\omega_x t)^2)(1 + (\omega_y t)^2)}} \text{Li}_2 \left(-\zeta \exp \left(-\frac{x^2}{r_x(t)^2} - \frac{y^2}{r_y(t)^2} \right) \right) \quad (1.9)$$

with $r_{x,y}(t) = \sqrt{1 + (\omega_{x,y} t)^2}$. In the limit of small fugacity $\zeta \ll 1$, this distribution reduces to the Gaussian distribution described in Eqn. 1.4. In reality, the deviation from a Gaussian is never large; trying a naive fit of a Gaussian to DFG with $T/T_F \sim 0.3$ merely results in a slight overfit at the peak [45, 46]. This overfit results from Pauli blocking limiting the number of atoms that can occupy the velocity class $p \approx 0$.

Since Eqn. 1.9 is a bit cumbersome to work with, I fit instead with the following

expression, following DeMarco’s approach [53]⁴, with fit parameters A, ζ :

$$OD(r) = A \operatorname{Li}_2\left(-\zeta \exp(-r^2/a^2)\right) \quad (1.10)$$

where OD is the measured optical density. This equation can be applied to radially-averaged, time-of-flight absorption images with $p = 0$ atoms aligned on $r = 0$. The fugacity, ζ , is the “shape factor” which determines the degree of deviation from a Gaussian and, thus, the degree of degeneracy, through the implicit relation

$$\operatorname{Li}_3(\zeta) = -\frac{1}{6(T/T_F)^3} \quad (1.11)$$

which can be numerically solved for T/T_F . $\zeta < 1$ is the classical region, and $\zeta \gg 1$ is degenerate, rapidly approaching infinity as $T/T_F \rightarrow 0$.

⁴ Li_2 is usable in Python from the `scipy` library as “`scipy.spence`”

Chapter 2

Light-matter interaction

All our experiments require firing lasers at atoms, so it is imperative we understand the basic physical processes underlying their interaction. This would be most accurately described by *quantizing* the electromagnetic field into operators; see, for instance, [59]. But in the most cold atom experiments, this approach is overkill,¹ and we are justified in approaching these problems “semiclassically:” the electric field is described classically as a plane wave, but the atom is described quantum-mechanically. Following the treatment of [2], I give a brief overview of the processes occurring in the experiment, followed by a description of the quantum Hall physics we are attempting to engineer in the experiment.

2.1 Two-level atoms

The basis for understanding most atom-laser problems is the humble two-level system. Consider an atom that has ground and excited states labeled $|g\rangle$ and $|e\rangle$. Let the ground-state have zero energy, and the upper-state have energy $\hbar\omega_0$. Expose this atom to a monochromatic laser beam; in the long-wavelength limit— $\lambda \gg a_0$, where a_0 is the atom’s Bohr radius—we can ignore the light’s spatial dependence.² We describe the light

¹It becomes important when the number of photons is comparable to the number of particles, such as in quantum cavity experiments.

²Note that standing waves, and Raman coupling, have an important spatial component which we will get to later.

field as a plane wave polarized along direction $\hat{\epsilon}$:

$$\mathbf{E} = \hat{\epsilon} E_0 \cos(\omega t) = \hat{\epsilon} \frac{E_0}{2} (e^{-i\omega t} + e^{i\omega t}) \quad (2.1)$$

The semiclassical Hamiltonian describing the system is then

$$H = H_A + H_{AF} \quad (2.2)$$

where $H_A = \hbar\omega_0|e\rangle\langle e|$ is the free atomic Hamiltonian and $H_{AF} = -\mathbf{d} \cdot \mathbf{E}$ is the atom-field interaction in the dipole approximation. Explicitly, $\mathbf{d} = -e\mathbf{r}$, with position operator \mathbf{r} . The dipole operator is purely off-diagonal in the $\{|g\rangle, |e\rangle\}$ basis due to parity considerations [2]; inserting the identity (with the usual lowering operator notation $a \equiv |g\rangle\langle e|$) we expand H_{AF} as

$$\begin{aligned} H_{AF} &= -\mathbf{d} \cdot \mathbf{E} = \langle g|\mathbf{d}|e\rangle (a + a^\dagger) \cdot \hat{\epsilon} \frac{E_0}{2} (e^{-i\omega t} + e^{i\omega t}) \\ &= -\frac{E_0}{2} \langle g|\mathbf{d} \cdot \hat{\epsilon}|e\rangle (ae^{i\omega t} + a^\dagger e^{-i\omega t} + a^\dagger e^{i\omega t} + ae^{-i\omega t}) \end{aligned} \quad (2.3)$$

The first two terms in parenthesis correspond, respectively, to photon emission with transition to $|g\rangle$, and photon absorption with transition to $|e\rangle$ [59]. These are the “energy conserving” terms. The last two terms describe photon emission with transition to $|e\rangle$, and photon absorption with transition to $|g\rangle$ —these are processes one expects to be highly suppressed when the radiation is near-resonant.³ Dropping these last two is called the *rotating wave approximation* (see Appendix E). And so, we have

$$H_{AF} = \frac{\hbar\Omega}{2} (ae^{i\omega t} + a^\dagger e^{-i\omega t}) \quad (2.4)$$

³These term identifications are by no means obvious; they come from considering the quantized nature of the field. The positively (negatively) signed component of \mathbf{E} is associated with the photon annihilation (creation) operator. [2]

where I have introduced the Rabi frequency, $\Omega \equiv -E_0/\hbar\langle g|\mathbf{d}\cdot\hat{\mathbf{e}}|e\rangle$.

To remove the time-dependence, it is customary to move into a frame rotating around the quantization axis by applying the unitary transformation $U = e^{i\omega t|e\rangle\langle e|}$ to Eqn. 2.4 with the replacement rule $\tilde{H} \rightarrow UHU^\dagger + i\hbar\partial_t U$:

$$H = -\hbar\Delta|e\rangle\langle e| + \frac{\hbar\Omega}{2}(a + a^\dagger) = -\frac{\hbar\Delta}{2}\mathbb{1} + \frac{\hbar\Delta}{2}\sigma_z + \frac{\hbar\Omega}{2}\sigma_x \quad (2.5)$$

where we have defined the detuning from resonance $\Delta \equiv \omega - \omega_0$, and on the right-hand side written the equation in terms of the usual Pauli matrices. From here, we can uncover the two-level dynamics; for the case of resonant light, $\Delta = 0$, and we write the time-dependent Schrödinger equation $i\hbar\partial_t\Psi = H\Psi$ in terms of the amplitudes c_g, c_e :

$$\begin{aligned} \partial_t c_g &= -i\frac{\Omega}{2}c_e \\ \partial_t c_e &= -i\frac{\Omega}{2}c_g \end{aligned} \quad (2.6)$$

We can solve these coupled equations analytically by differentiating one and plugging it into the other. The result, for an initially unexcited atom ($c_g(t=0) = 1, c_e(t=0) = 0$) show *Rabi oscillations* between ground and excited states:

$$\begin{aligned} P_g(t) &= |c_g|^2 = \cos^2(\Omega t/2) \\ P_e(t) &= |c_e|^2 = \sin^2(\Omega t/2) \end{aligned} \quad (2.7)$$

These Rabi oscillations occur in stimulated Raman/spin-orbit coupling experiments, and are an important way of measuring the coupling strengths Ω , as we shall see in Chapter 5.

The last important detail is the realization that the states $|g\rangle, |e\rangle$ are no longer eigenstates of this system—the Hamiltonian is not diagonal in this basis, due to the

presence of coupling. Diagonalizing this Hamiltonian, the new eigenenergies of the system are

$$E_{\pm} = -\frac{\hbar\Delta}{2} \pm \frac{\hbar\Omega}{2} \quad (2.8)$$

with corresponding eigenvectors

$$\begin{aligned} |+\rangle &= \sin(\theta)|g\rangle + \cos(\theta)|e\rangle, \\ |-\rangle &= \cos(\theta)|g\rangle - \sin(\theta)|e\rangle \end{aligned} \quad (2.9)$$

with $\tan(2\theta) = -\Omega/\Delta$. These are the so-called *dressed states*, which are superpositions of the original eigenstates. This dressed-state picture is important in spin-orbit coupling, which we shall soon discuss.

2.2 Stimulated Raman effect

The two-level Hamiltonian Eqn. 2.5 provides a useful conceptual framework we can build on to describe nominally-more-complicated setups—for example, the one in Fig. 2.1(a), describing *two* detuned fields interacting with a three-level “ Λ -system.” It turns out that if the fields cause minimal spontaneous emission, the Hamiltonian reduces to the familiar two-level form.

The free atomic Hamiltonian is analogous to the two-level case, but I drop the at-rest assumption and have

$$H_A = \frac{p^2}{2m} \mathbb{1} - \hbar\omega_1|g_1\rangle\langle g_1| - \hbar\omega_2|g_2\rangle\langle g_2| \quad (2.10)$$




Figure 2.1. The Raman process responsible for spin-orbit coupling of neutral atoms, and which underlies our spin-momentum lattice. (a) Two beams couple two spin states to a shared upper state, with each beam detuned from single-photon resonance by Δ . (b, top) The real-space geometry of a spin exposed to two counter-propagating beams at $t = 0$. (b, bottom) After a π -pulse, the spin is flipped, and the atom acquires a momentum equal to the difference momentum of the two photons. Spin is therefore coupled to the motion—spin-orbit coupling.

Spatial dependence is now relevant, so I preserve it in writing the electric field:

$$\mathbf{E} = \hat{\epsilon}_1 \frac{E_{01}}{2} \left(e^{i\mathbf{k}_1 \cdot \mathbf{r} - i\omega_1 t} + e^{-i\mathbf{k}_1 \cdot \mathbf{r} + i\omega_1 t} \right) + \hat{\epsilon}_2 \frac{E_{02}}{2} \left(e^{i\mathbf{k}_2 \cdot \mathbf{r} - i\omega_2 t} + e^{-i\mathbf{k}_2 \cdot \mathbf{r} + i\omega_2 t} \right) \quad (2.11)$$

The atom-field interaction $H_{AF} = -\mathbf{d} \cdot \mathbf{E}$ is again written in the RWA, and I define the lowering operators as $a_\alpha \equiv |g_\alpha\rangle\langle e|$:

$$H_{AF} = -\frac{\hbar\Omega_1}{2} \left(a_1 e^{-i\mathbf{k}_1 \cdot \mathbf{r}} e^{i\omega_1 t} + c.c. \right) - \frac{\hbar\Omega_2}{2} \left(a_2 e^{-i\mathbf{k}_2 \cdot \mathbf{r}} e^{i\omega_2 t} + c.c. \right) \quad (2.12)$$

where $c.c.$ indicates the complex conjugate and the Rabi frequencies are defined as before, $\Omega_\alpha \equiv -E_0/\hbar\langle g_\alpha|\mathbf{d} \cdot \hat{\epsilon}_\alpha|e\rangle$. To remove the time dependence, move into the rotating frame with the unitary transformation $U = e^{i\omega_1 t|e\rangle\langle e|} + e^{i\omega_2 t|e\rangle\langle e|}$; the full Hamiltonian becomes

$$\begin{aligned}
H &= H_A + H_{AF} \\
&= \frac{p^2}{2m} \mathbb{1} + \hbar\Delta_1|g_1\rangle\langle g_1| + \hbar\Delta_2|g_2\rangle\langle g_2| - \frac{\hbar\Omega_1}{2} \left(a_1 e^{-i\mathbf{k}_1 \cdot \mathbf{r}} + c.c. \right) - \frac{\hbar\Omega_2}{2} \left(a_2 e^{-i\mathbf{k}_2 \cdot \mathbf{r}} + c.c. \right)
\end{aligned} \tag{2.13}$$

For the wavefunction, we make the *ansatz* $|\Psi\rangle = \psi_{g_1}|g_1\rangle + \psi_{g_2}|g_2\rangle + \psi_e|e\rangle$. Before we plug this into the Schrödinger equation $i\hbar\partial_t|\Psi\rangle = H|\Psi\rangle$, let's make one further approximation under the assumption that the lasers are tuned far from resonance, i.e. that $\Delta_\alpha \gg \Gamma$, where $1/\Gamma$ is the lifetime of the excited state. Then the upper state $|e\rangle$ is never significantly populated, so we can set $\partial_t\psi_e = 0$; this is called the *adiabatic elimination* of the upper state.⁴ We then find the equations of motion

$$\begin{aligned}
i\hbar\partial_t\psi_{g_1} &= \frac{p^2}{2m}\psi_{g_1} + \left(\frac{\hbar\delta}{2} + \omega_1^{AC}\right)\psi_{g_1} + \frac{\hbar\Omega_R}{2} e^{i(\mathbf{k}_2 - \mathbf{k}_1) \cdot \mathbf{r}} \psi_{g_2}, \\
i\hbar\partial_t\psi_{g_2} &= \frac{p^2}{2m}\psi_{g_2} + \left(-\frac{\hbar\delta}{2} + \omega_2^{AC}\right)\psi_{g_2} + \frac{\hbar\Omega_R}{2} e^{i(\mathbf{k}_1 - \mathbf{k}_2) \cdot \mathbf{r}} \psi_{g_1}
\end{aligned} \tag{2.14}$$

These equations are formally equivalent to the two-level atom, with $\delta \equiv \Delta_1 - \Delta_2$, $\Delta \equiv 1/2(\Delta_1 + \Delta_2)$, $\omega_i^{AC} \equiv \Omega_i^2/4\Delta$, and Raman Rabi frequency

$$\Omega_R \equiv \frac{\Omega_1\Omega_2^*}{2\Delta} \tag{2.15}$$

Defining $\mathbf{q} \equiv \mathbf{k}_1 - \mathbf{k}_2$, the effective Hamiltonian generating these equations is [62]

⁴If this step seems unjustified, I note that it is not strictly necessary and the same result is derived through more rigorous means in [60]... But for the busy working physicist, I'd recommend [61] for a more-readable and physical explanation.

$$\begin{aligned}
H_{\text{Eff}} &= \frac{p^2}{2m} \mathbb{1} + \frac{\hbar}{2} \begin{pmatrix} \delta + \omega_1^{\text{AC}} & \Omega_R e^{i\mathbf{q}\cdot\mathbf{r}} \\ \Omega_R^* e^{-i\mathbf{q}\cdot\mathbf{r}} & -\delta + \omega_2^{\text{AC}} \end{pmatrix} \\
&= \frac{p^2}{2m} \mathbb{1} + \frac{\hbar\Omega_R}{2} (\cos(\mathbf{q}\cdot\mathbf{r})\sigma_x - \sin(\mathbf{q}\cdot\mathbf{r})\sigma_y) + \frac{\hbar\delta'}{2}\sigma_z
\end{aligned} \tag{2.16}$$

where in the last line I expand in terms of the Pauli matrices, and define $\delta' \equiv \delta + \omega^{\text{AC}}$ (under the assumption that the ac Stark shifts are approximately equal). In the literature, especially in the context of spin-orbit coupling, you will sometimes see this form, which describes a lattice with primitive vector $2\pi/|\mathbf{q}|$. Note that the exponentials $\exp(\pm i\mathbf{q}\cdot\mathbf{r})$ in Eqn. 2.16 are the momentum shift operators [63] which act on a momentum state as $\exp(\pm i\mathbf{q}\cdot\mathbf{r})|\mathbf{p}\rangle = |\mathbf{p} \pm \hbar\mathbf{q}\rangle$. Thus, the off-diagonals connect a state $|g_1, \mathbf{p} = 0\rangle$ to $|g_2, \mathbf{p} = \hbar\mathbf{q}\rangle$. When the atom hops, it absorbs a photon from \mathbf{k}_1 and emits a photon (of energy $c\hbar k_2$) in the \mathbf{k}_2 direction, recoiling along $\mathbf{k}_1 - \mathbf{k}_2$ when it does so, as in Fig. 2.1(b). A final form seen in the literature follows from reducing to the 1D case and counter-propagating beams, so that $\mathbf{q}\cdot\mathbf{r} = 2\hbar k_R x$; then, we make contact with [64]:

$$H_{\text{Eff}} = \frac{p^2}{2m} \mathbb{1} + \frac{\hbar\Omega_R}{2} (F_+ e^{ik_R x} + F_- e^{-ik_R x}) \tag{2.17}$$

where I invoke the usual raising/lowering operators F_+, F_- .

2.3 Spin-orbit coupling

2.3.1 History

The Raman Hamiltonian Eqn. 2.16 describes an artificial spin-orbit (SO) coupling field, which I will show in this section. Artificial SO coupling in ultracold atoms was first implemented by Spielman's group with a Rb BEC [65], followed quickly by Jing Zhang's group with a K Fermi gas [66]. The one-dimensional version has since been extensively

studied in experiments with different atomic species, including the fermionic isotopes of Li [67], Yb [68], Dy [69], and even in a ^{87}Sr atomic clock [70, 71]. Two-dimensional spin-orbit coupling has also been implemented in Rb [72–74] and very recently, ^{87}Sr [75]. For an overview of the field, I recommend [76] and the excellent thesis by K. J. Garcia [77] for properties of spin-orbit coupled Hamiltonians. A more recent review can be found at [78].

2.3.2 SO coupling: physical origin

From an atomic physics perspective, SO coupling is the process that gives rise to fine structure, which is responsible for the J in the familiar 3P_J triplet levels of strontium. In that context, SO coupling is the interaction between the electron’s spin and its orbital angular momentum. This can be written $H_{SO} = -\boldsymbol{\mu} \cdot \mathbf{B}$, where $\boldsymbol{\mu} = e\mathbf{S}/(2m)$, \mathbf{S} is the spin operator, and $\mathbf{B} = -\mathbf{v}/c^2 \times \mathbf{E}$ is the (Maxwell equation describing the) effective magnetic field the electron feels as it orbits the nucleus with velocity \mathbf{v} [63]. This is also called the LS -coupling scheme.

But here we are in the business of quantum emulation, and when we refer to “SO coupling” we generally have in mind the closely-related physics that occur in solid-state materials. Consider a material in which exists a static electric field $\mathbf{E} = E_0\hat{z}$. From the Lorentz-invariant Maxwell equations, a charged particle moving through this field in the xy -plane with velocity $\mathbf{v} = \mathbf{p}/m = \hbar\mathbf{k}/m$ experiences a magnetic field $\mathbf{B}_{SO} = -\mathbf{v}/c^2 \times \mathbf{E} = \frac{1}{c^2} \left(\frac{\hbar\mathbf{k}}{m} \times \mathbf{E} \right) = \frac{E_0\hbar}{mc^2} (k_x\hat{y} - k_y\hat{x})$. The resulting momentum-dependent Zeeman interaction $-\mu B \approx \sigma_x k_y - \sigma_y k_x$ is known as *Rashba SOC*. A similar type of SOC that can emerge in some crystals is *Dresselhaus SOC* with $-\mu B \approx \sigma_x k_x - \sigma_y k_y$ [76].

2.3.3 SO coupling in ultracold gases

We can implement SOC with neutral atoms with the help of Raman coupling. Actually, a kind of SO coupling is already in Eqn. 2.16. Apply the unitary transformation $U = \exp(i(\mathbf{q} \cdot \mathbf{r}/\hbar)\sigma_z)$ to the Hamiltonian with the usual rule $H' = \hat{U}H\hat{U}^\dagger$ and we arrive at

$$H_{\text{Eff}} = \frac{(p\mathbb{1} - \mathbf{q} \cdot \mathbf{r}\sigma_z)^2}{2m} + \frac{\delta'}{2}\sigma_z + \frac{\Omega_R}{2}\sigma_x \quad (2.18)$$

which describes a particle in an equal-parts Rashba and Dresselhaus spin-orbit coupling field [65]. This is a 1D artificial spin-orbit coupling; achieving the fully 2D version requires a third ground state with additional beams, but follows a very similar theory nicely described in [79]. 2D spin-orbit coupling is key to our pursuit of topological matter, and we implement it in the experiment that follows.

2.4 Quantum Hall systems

We have seen that we can generate artificial spin-orbit coupling (SOC) fields in neutral atoms. This allows for the simulation/emulation of various solid-state phenomena; but SOC can also be used as a tool for the creation of more complex band structures, specifically those endowed with topology. The goal of this thesis is to engineer a cold atom ensemble which exhibits quantum Hall characteristics. I briefly review these systems below, and explain the SOC connection which is pivotal to our results.

2.4.1 Classical

The Hall effect is a classically-understood process describing the transverse voltage developed across a conductor placed in a magnetic field. The simple Drude model is sufficient to explain this: electrons traveling along the conductor's length experience a Lorentz force, $\mathbf{F} = -e\mathbf{E} - e\mathbf{v} \times \mathbf{B}$, which deflects their motion towards the conductor's boundary. Charge then builds up at one edge, and in the steady-state establishes a voltage with respect to the opposite edge. Materials can be then characterized by their Hall coefficient, $R_H = \rho_{xy}/B$, where ρ_{xy} is the resistivity matrix element which characterizes a resistance along y due to a current along x [80].

2.4.2 Integer quantum Hall effect

In a 1980 experiment [81], it was discovered that—for a certain class of materials under very large magnetic fields—the Hall resistance is quantized in integer steps. It was found to be $R_H = \frac{h}{e^2} \frac{1}{\nu}$ where $\nu \in \mathbb{Z}$, leading to a directly-measurable, highly-precise definition of the fine structure constant. This is a purely quantum effect which can be understood without considering electron-electron interactions; the explanation came a year later by Laughlin [82]. This integer quantum Hall effect (IQHE) has a few subtleties, and a comprehensive guide can be found in David Tong’s lectures on the subject [80]. I will describe here only a few key concepts which will be useful to understand the available cold atom literature.

An electron in an electric and magnetic field is subject to the Hamiltonian $H = (\mathbf{p} + e\mathbf{A})^2/2m + e\phi$, where \mathbf{A} (ϕ) is the vector (electric) potential. This is the classic *Landau level* problem, which can be exactly solved [80] in your favorite gauge. What emerges from the solution are quantized, degenerate energy levels, each of which can accommodate a large (but finite) number of electrons. The number of filled Landau levels corresponds to the integer in the Hall coefficient; the electron density required to fill a given Landau level is $n = B/\phi$, where B the magnetic field and $\phi_0 = h/e$, the flux quantum.

In actual quantum Hall data, we observe the Hall coefficient on “plateaus”—a range of magnetic fields at which only the center satisfies $B = n\phi_0/\nu$. So what happens at other values? The origin of the plateau is material disorder; impurities cause a smearing of the density of states, which are normally δ functions. The states which arise tend to be extended over the entire sample, rather than localized [80]; extended states are the only ones which can carry current, and are guaranteed to be found at the edges. So, as the magnetic field decreases, each Landau level can accommodate fewer electrons; but before jumping to the next Landau level, electrons instead start to fill localized states, which cannot carry current, and thus do not contribute to nor change the conductivity. Hence,

we observe plateaus.

2.4.3 Role of topology

Two years after the discovery of the integer quantum Hall effect, a deep connection to topology was revealed [83]. Materials within the same topological class share “universal” features, regardless of the particular geometry, atom, or impurities present. The classic (mathematical) example is the transformation of a coffee cup to a doughnut: one can be continuously deformed into the other, a relationship encapsulated mathematically by some topological invariant, a quantity which is the same for both. In 2D materials, this invariant is the Chern number, and describes how the wavefunction twists in momentum space; in systems with band structure, a separate Chern number can be assigned for each band, and is given generally by an integration over the first Brillouin zone (BZ):

$$C = \frac{1}{2\pi} \int_{BZ} i\epsilon_{ij} \partial_{q_i} \langle u_q | \partial_{q_j} | u_q \rangle d^2q \quad (2.19)$$

where u_q is the periodic part of the Bloch wavefunction with quasimomentum q ; this is simply the flux of the Berry curvature over the “surface” of the first BZ. The system in which the integer quantum Hall effect was first observed is topological: the total number of filled bands quantized the Hall conductance; that integer can be proven to be the Chern number! These days, systems for which a quantized conductance can associated with a Chern number are said to be “integer quantum Hall” systems. Any band insulator will exhibit the integer quantum Hall effect if the total Chern number of the bands is nonzero.

One way to engineer a topological cold atom system is with an optical lattice that breaks time-reversal symmetry (TRS); breaking TRS is required for Eqn. 2.19 to be nonzero. TRS can be broken if the magnetic flux is a significant fraction of the flux quantum, h/e .⁵ Since the lattice inter-atomic spacing is so small, this would require

⁵We have often heard that the Maxwell equations obey TRS, including the magnetic field. But that’s true under the global operations $\mathbf{B} \rightarrow -\mathbf{B}$, $\mathbf{J} \rightarrow -\mathbf{J}$, and $\mathbf{v} \rightarrow -\mathbf{v}$. If we’re treating the atoms and lattice

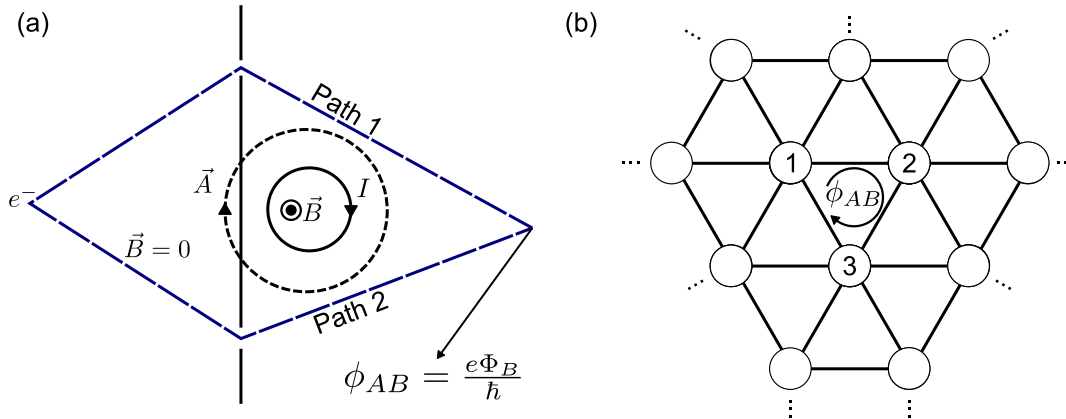


Figure 2.2. Aharonov-Bohm effect and artificial magnetic field generation. (a) An electron beam split in two paths interferes after traversing a region with $\vec{B} = 0$, $\vec{A} \neq 0$, which can be generated by a solenoid with current I . Despite the absence of a magnetic field in the region of electron travel, the electrons still feel the Aharonov-Bohm phase shift due to the vector potential \vec{A} , which is proportional to the enclosed magnetic flux Φ_B . (b) To generate an artificial magnetic field in a lattice potential, it is sufficient to ensure that an atom encircling a plaquette $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$ acquires a net non-zero phase shift that is comparable to the flux quantum e/\hbar .

impractically large magnetic fields. Instead, experimenters have turned to artificial gauge fields [84], which create an SOC that endows an Aharonov-Bohm phase to atoms hopping around a plaquette. This was first implemented in a Harper-Hofstadter Hamiltonian [85].

Generally speaking, if one can engineer the plaquettes such that an atom picks up a net phase as it encircles one, a large magnetic field will be present, breaking TRS; see Fig. 2.2(b). The bands can then be topological. For a recent, comprehensive, and readable review of topology in cold atom systems, see [15].

2.4.4 Fractional Chern insulators

The fractional quantum Hall (FQH) effect describes the quantization of the Hall conductance in rational, instead of integer, steps. The microscopic origin of these states is fundamentally due to electron-electron interaction, unlike the IQHE which is due to Landau level physics alone. Its description is well beyond the scope of this work; I refer the

alone as our system, an external magnetic field breaks time reversal symmetry if we don't simultaneously reverse the (micro)current that originates \mathbf{B} when we reverse t . Sakurai [63] explains this pretty well.

reader to [80] for further details. However, I can say that one key ingredient to the FQH story is the Laughlin wavefunction, describing filling fractions $1/m$ for integer m , and has good overlap with the ground state of some dozen or so interacting particles. Although by itself it breaks down at the macroscopic number of electrons in a real system, it is a good starting point which represents a “universality class” from which more-accurate wavefunctions might be built. More precisely, the Laughlin wavefunctions possess the same topological order as the true wavefunction, and shows the same fractional excitations.

Fractional Chern insulators (FCIs) are the lattice analogues of the FQH states. There is considerable interest in adapting the topological order already established in many cold atom systems by adding the strong correlations necessary to realize FCIs [86–88]. The proposal we are following is given in [89], which engineers flat bands using a spin-momentum-space lattice created by Raman couplings. This is a class of lattices called *optical flux lattices* [90], which generate a large magnetic field via Aharonov-Bohm phase [91]. They can more generally be described as artificial gauge fields [84].

To implement the scheme in [89], the experimental setup is shown in Fig. 2.3. Part (a) shows all present optical and magnetic fields. Three states are coupled together as in part (b), and, when enough links are provided, a lattice is formed in a spin-momentum space; we call this a spin-momentum lattice (SML). Each Zeeman spin state (labeled by the shorthand notation X, Y, Z) is connected to every other spin state with a two-photon Raman transition. The physics behind Raman transitions is discussed in Chapter 2.2, but the main idea is that when an atom hops from one site to another, it changes spin and picks up momentum equal to $\hbar k_R = \hbar(\mathbf{k}_1 - \mathbf{k}_2)$. We simply apply sufficiently many beams such that a lattice is formed in momentum space. If the phases are properly set (as denoted by the α, β, γ in part (c)), the atom experiences an effective Aharonov-Bohm magnetic field as in Fig. 2.2; the system is topological; and, if we can muster enough spin-spin contact interactions, FQH states are predicted to be formed. We present the results of this experiment in Chapter 5.

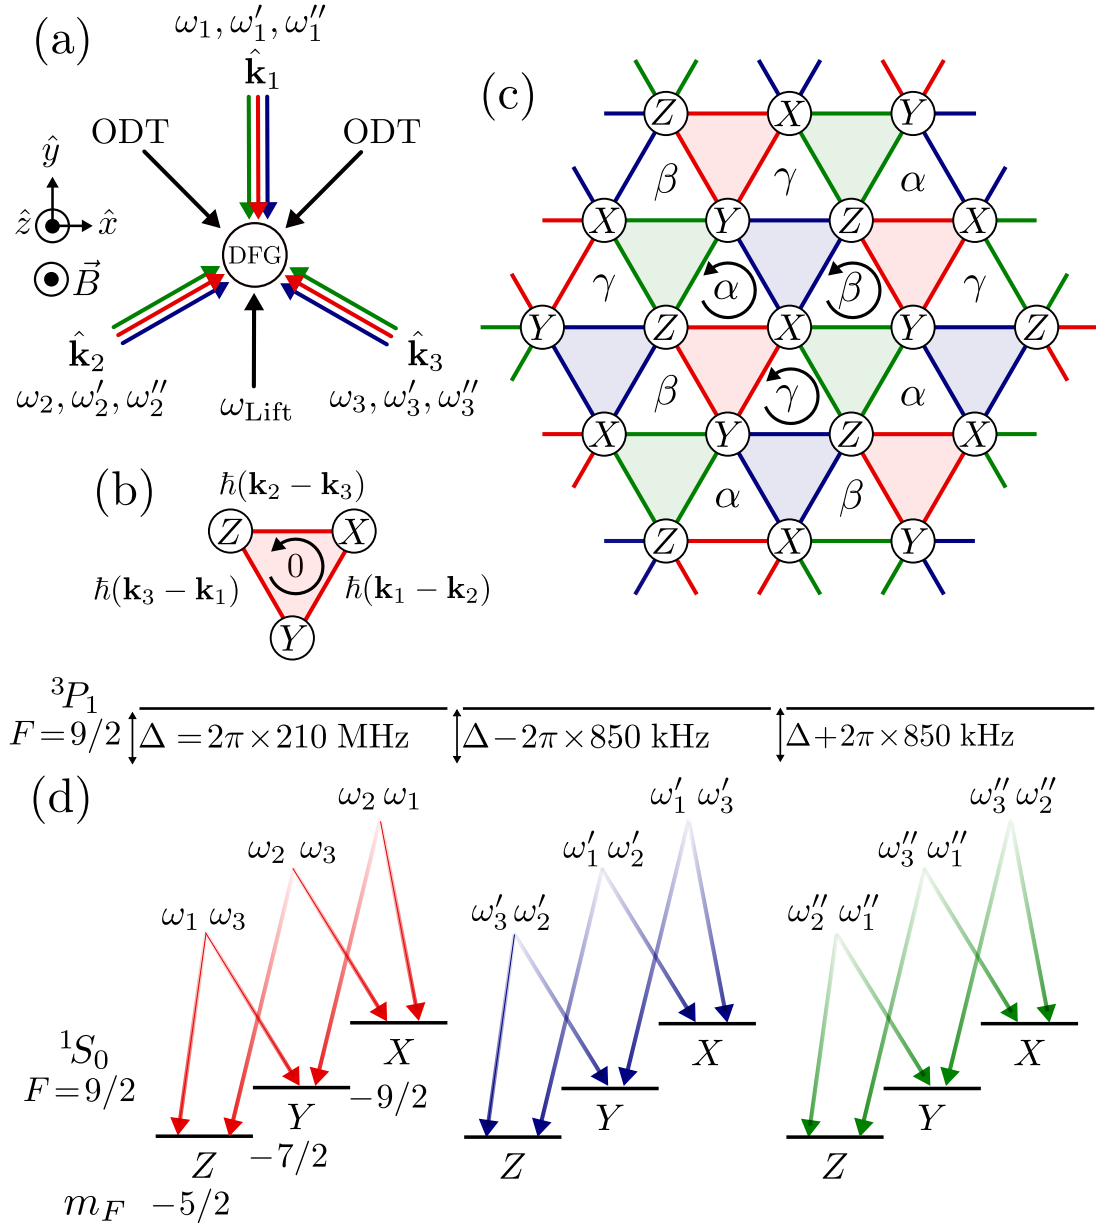


Figure 2.3. SM lattice and coupling details. (a) The DFG is exposed to cyclic Raman couplings between three internal spin states. A beam ω_{Lift} provides a nonlinear energy splitting between the states. (b) States X, Y, Z are connected among themselves through spin-momentum exchange in units of $\hbar\delta\mathbf{k}_{ij} = \hbar(\mathbf{k}_i - \mathbf{k}_j)$. With a single set of beams, no net phase pickup is possible, denoted by the 0. (c) The couplings form a lattice in momentum space. Atoms encircling plaquettes labeled by α, β, γ can pick up a net phase. The link color indicates the frequency set in (d) to which the beams belong. (d) Details of the resonant couplings between the three internal states labeled X, Y, and Z, which represent the nuclear angular momentum projections $m_F = -9/2, -7/2, -5/2$, respectively, in the $^1S_0(F = 9/2)$ ground state. We circularly exchange the roles of the frequencies colored red (ω_i) in the blue (ω'_i) and green (ω''_i) coupling sets.

Chapter 3

Experimental setup

Experiments with neutral-atom SOC described in the previous chapter requires a machine which can produce a quantum gas; making a quantum gas is hard. It requires a wide variety of sometimes-elaborate, always-expensive range of specialty equipment. This chapter aims to describe all the necessary components required for the experiment.

3.1 Lab conditions

The lab where one attempts to make a quantum gas is of vital importance. Although the atoms themselves reside in a vacuum chamber, everything that talks to them is subject to the (in)stability of the ambient temperature. In an experiment with many hundreds of mirrors directing beams all over the place, drifts in temperature on the order of a few degrees can wreck your day: finite coefficients of thermal expansion will cause all your fiber couplings to drift slightly, your AOM efficiencies to change, and cause beams that were perfectly aligned on your atoms to now deviate by a few precious microns. A swing of more than 5 degrees is cause enough to just go home and wait for the temperature to (hopefully) stabilize.

The humidity requirements are less drastic, but changes on the order of 10% cause our strontium lamp setpoint to drift by a few MHz, which can throw off the blue MOT. If it reaches above 70%, the lasers will start to notice and have serious trouble maintaining

lock, owing to increased scattering off water vapor. A humidity swing above 90% is a cause for panic, since condensation in the lasers can be fatal for the doubling crystals, which are extremely hygroscopic.

When I started my PhD, we worked in a lab in Mayer Hall with very poor temperature control. Swings above 5 degrees F were common, and the humidity was poorly controlled, at times producing sauna-like conditions. Turning on too much equipment generated enough heat to throw the AC out of whack. Thankfully, the university noticed, and built us a new lab space in Urey Hall; these facilities are quite amazing. Running 24/7/365, the lab is kept on a tightly-controlled temperature regulation loop of ± 0.1 degrees F at all times at the tables themselves (with the surrounding table doors closed), with a relative humidity set at 37% and swinging within 5%. This results in very good beam pointing and extremely reliable day-to-day operation.

3.2 Vacuum chamber

All ultracold gas experiments take place in vacuum chambers, in order to eliminate collisions with background gases which would otherwise unacceptably limit the duration of experiments. To reach many-second lifetimes—necessary for successful evaporative cooling—ultra-high vacuum (UHV) pressures of lower than 10^{-10} torr must be achieved, preferably around 10^{-11} torr. Reaching these pressures is extremely difficult. Confounding this goal is that elemental strontium is a crumbly, glossy rock, which must be heated to around 500 K to appreciably evaporate [1]. High temperatures are inherently bad for low-pressure work, because that heat is inevitably transferred to the surrounding chamber walls, which exponentially increases outgassing rates [92] and raises the pressure by orders of magnitude.

Several generations of experiments, from MIT [93], Zurich [94], to JILA [95, 96] utilize similar steel vacuum chambers, which are designed to present large pumping

*conductance*¹ at the atoms, while also providing excellent optical access. We use a modified version of this chamber.

3.2.1 Chamber design

Our chamber design is most heavily influenced by the UHV Sr chambers built in Jun Ye’s lab [19, 95, 96]. We made a few design changes from [96]; namely, we include more windows in the xy -plane, including four dedicated 1064 nm-coated ODT windows; separate windows for red and blue horizontal MOT beams; a permanently-attached residual gas analyzer (RGA) to aid with the bake; a discrete NEG (SAES Capacitorr C 2000) in lieu of a chamber NEG coating²; and finally, we don’t use a titanium sublimation pump, opting instead for a 150 L/s ion pump (Agilent Slim Body 150 L/s StarCell).

A render of the experiment is shown in Fig. 3.1, highlighting some of its key features. The entire apparatus is less than 5 feet long, and consists of a commercial atomic beam package from AOsense, which houses a small ion pump and getter to combat the outgassing resulting from the oven’s high temperature. It furthermore includes an integrated 2D MOT and Zeeman slower, which is aligned at an angle to the atomic beam direction, minimizing Sr buildup in the science chamber.

The AOsense is attached to the science chamber via a narrow tube, through which streams Sr atoms azimuthally cooled to about 1 mK. Two recessed viewports (“bucket windows”) allow proximal atom access: there is only 4 mm of separation from the vacuum-side glass surface to chamber center! This allows the design of compact electromagnets and the eventual placement of a high NA objective. The complete science chamber drawing is given in the Appendix A, with the recessed viewport schematics and their anti-reflection coatings in Appendix B.

¹A vacuum term analogous to resistivity: higher is better, and almost all chambers are designed simply to maximize this number.

²The getter coating of [96] started flaking off after some years of use, dirtying their bucket window.




Figure 3.1. Render of the experimental chamber.

3.2.2 Chamber bake

Reaching UHV is at once both an art and science. An excellent, practical guide to all things vacuum-related can be found in the appendix of [97]. I won't repeat the techniques described there, except to distill three key points that apply to all UHV endeavors:

- Cleanliness is paramount.
- Lots of conductance is best.
- Bake until everything is gone; do this a few times.

Having provided by design “lots” of conductance to the atoms in our chamber (hundreds of L/s), and promising ourselves to wear gloves and change them frequently, we simply baked our chamber, and that was that! Haha.

In reality, this was a 6 month ordeal. The largest stumbling block was finding that our conflat viewports, purchased from Lesker and sent to TAKOS for anti-reflection coating, were leaking very badly: TAKOS neglected to mask the conflat knife-edge when they did their coating runs, so the windows never properly sealed. We were unable to remove the coating from the edge, so our entire run of conflat windows had to be replaced.

I cannot recommend this company for any future business. Our second set of viewports was made and coated by MPF Products, and these were excellent. I measured 0.4% (0.6%) loss at 461 nm (689 nm), meeting their specification of 0.25% (0.4%) per surface. Most importantly, they didn't leak.

Bake strategy and technique

The main idea of a bake is to bring the system temperature as hot as possible to drive out water—which is adsorbed onto the steel walls, and is the principal source of pressure in an otherwise-clean steel chamber, limiting the achievable pressure to $10^{-6} - 10^{-8}$ torr. Chambers containing glasses or mechanical valves are extremely sensitive not only to over-heating, but also to temperature gradients, which can crack windows or cause valves to leak. So, we adhere to “backed potato” method described to us by our friends Ruwan, Shankari, and Zach at David Weld’s lab. We loosely wrap everything in a layer of UHV foil³, apply strip heaters over surfaces and band heaters surrounding the window conflat, then wrap everything again in several more layers of foil. We then slowly, carefully raise the temperature of the entire system, striving to keep the temperatures even and rising at a rate far below 1 deg C/s. We aim for 20 degrees below the maximum temperature of the lowest-rated part, which for us was the windows, rated to 200 C.

How long do you bake once you hit setpoint temperature?⁴ That’s where our permanent RGA (SRS RGA100) comes in, which monitors the partial pressures of elements with 1-100 amu (see Fig. 3.3(f)). The idea is to continuously measure mass spectra until the water peak (at 18 amu) and any other junk disappears, except for hydrogen, which is emitted continuously by the steel walls and cannot be depleted entirely. This provides a quantitative measure of vacuum cleanliness, much better than a simple pressure gauge. For \$4000 USD, the RGA also doubles as an inexpensive, effective helium leak checker,

³Also following Weld lab advice, we put steel mesh gratings in front of all our viewports before wrapping with foil, to avoid the foil potentially touching the glass and transferring heat to it.

⁴“Google, how long do I bake my NEG?” - G.K., 2018

which run more than \$20k: it has a mode where it monitors a single element, so if you configure it to look for He (4 amu) you can spray helium around your chamber and look for a peak. We placed the RGA close to the atoms in order to attain higher sensitivity: this proved crucial to understanding the bake progress, since its sensitivity was 2x higher than another RGA we also had at the turbo pump (Agilent TwissTorr 304).

Bake control equipment

We use simple tape and band heaters from Omega as heating elements. For temperature control, we pulse-width modulate (PWM) some solid-state relays (SSRs, Crydom CL240D10RC) using a few 16-channel PWM devices from adafruit (id 815, based on PCA9685). The PWM driver is super useful: you talk to it with a microcontroller (we use BeagleBone Black's) and set a duty cycle; it will output a corresponding square wave, with enough drive current capability to toggle the SSR. I had to slightly modify the adafruit breakout board, since the default PWM clock speed is much too fast to be useful (we need Hz or sub-Hz duty cycles). I gave it a (BeagleBone-supplied) external clock, which I soldered to the EXTCLK pin.

For thermocouple readout, we use the MAX31855 (on a breakout board, adafruit 269), which communicates to a microcontroller via SPI, connected to CAT5 cables going to a central rack. These chips work well enough on their own, but when you need to connect many (> 20) on the same SPI bus, I started to have serious problems with the chips reporting nonsense temperatures. These effects were mitigated by soldering 0.1 uF capacitors across VDD and GND, providing some stability to the power supply. I also had to limit the SPI bus speed to 50 kHz when using long (15 foot) CAT cables. We noticed that these chips are particularly sensitive to RF interference, as we observed a particularly-noisy RF amplifier impacting the chip readouts; this effect has been noticed in various online communities as well, and the only solution is to move the sensors away from sources of RF noise (we used a handheld RF meter to aid us).

High temperature bake

When we first received the chamber, we put on only the valve (DN40 VAT Series 54, all-metal angle valve) and RGA. We bought a set of blank conflat from Lesker, cleaning them in an ultrasonic bath (Elmasonic S900H)⁵ with Alconox and water, then isopropynol, and then finally with methanol⁶. As shown in Fig. 3.3(a), we then blanked off every port and did a “high-temperature” bake to ≈ 230 C for two weeks to clean out the chamber, which is too large to sonicate.

The perils of baking were on full display, even in this relatively-low risk setting without windows. We discovered how easy it was to overheat the system; wrapped with a lot of insulation, the temperature rose long after we stopped increasing the heaters, reaching > 280 C at the top, which (since heat rises) is always the hottest part of a bake. Thus we learned to increase the temperature more slowly, waiting for the system to stabilize before attempting to further increase. I also caused the VAT valve to leak, either through overheating or heating it too fast.

Final bake(s)

We baked the system for two weeks stretches, taking about 2 days for ramp up and ramp down. Special care was taken to optimize our NEG usage: it presents a lot of surface area, so atoms coming off the walls can potentially be adsorbed onto the NEG, wasting its capacity. Thus, during the bake, we kept it in “conditioning” mode, which ran enough current through its heater to keep it much hotter than its surroundings. We gratefully acknowledge assistance from Yevgeniy Lushtakat from SAES for help in determining the optimal bake procedure.

⁵Viewports cannot be cleaned in an ultrasonic bath, since the solvent can etch away the glass-metal seals; valves and bellows also cannot be cleaned, due to their many small crevices which would trap sediment or erode parts.

⁶Methanol is flammable and toxic, and it is a bad idea to fill a 90L bath with it. So, we put our parts in appropriately-sized glass beakers containing the minimum volume of methanol needed to submerge the part(s). The beakers are then placed in the water bath with a shrink wrap cap to minimize evaporation. Regardless of the precautions, this is inherently dangerous and must be closely monitored.

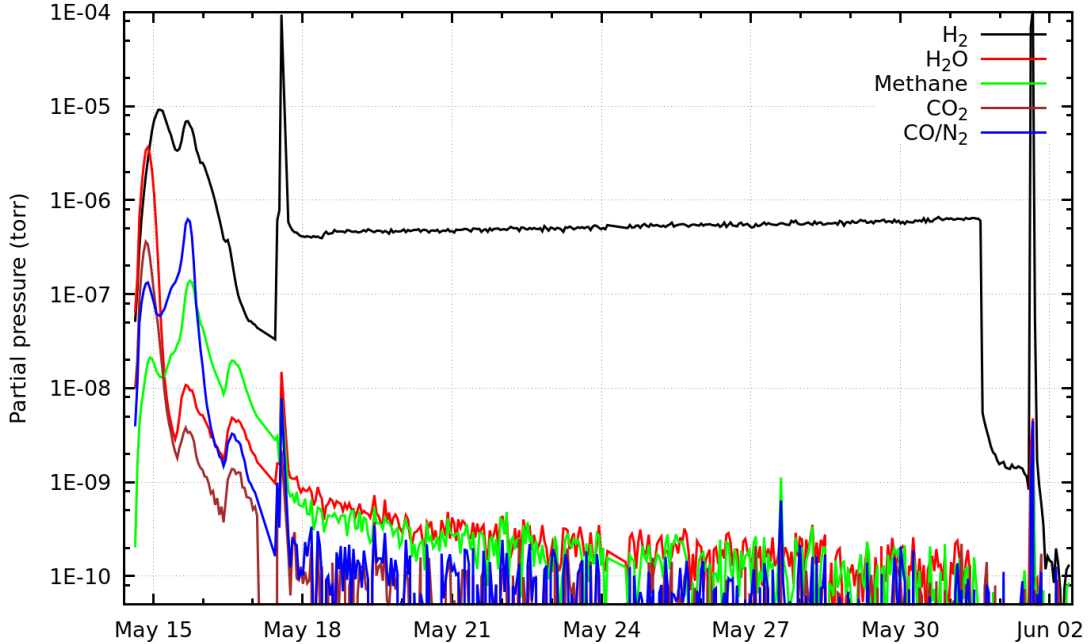


Figure 3.2. Mass spectra evolution of our first bake. This bake reached low 10^{-10} torr. The hydrogen spike on May 18 (June 2) is from NEG conditioning (activation).

Once the water peak reached low 10^{-10} torr (see Fig. 3.2) we started to cool down. When temperatures reached around 100 C, we cycled the ion pump (degassing it) three times for 30 seconds each, and degassed the other gauges including the RGA (which was already mostly clean, having been on throughout the bake). At this point, the NEG was placed into activation mode for two hours.⁷ After letting it cool for 3 hours, we turned on the ion pump, and closed the turbo pump valve at a few hours later, isolating the chamber with temperatures mostly near 70-80 C. The first time we did this, we (heartbreakingly) reached only the low 10^{-10} torr—evidently, our water load was still too high!⁸ So we powered the bake back up and did it all again for another 2 weeks. This time, we reached

⁷We were careful not to turn on the ion pump while the NEG was in activation mode: it releases tons of hydrogen, which can quickly saturate the ion pump.

⁸It's important to note the role of the filament current in our pressure gauge, the Agilent UHV-24P. The default controller current, 0.8 mA, read out 1.5×10^{-10} , but when I increased it to 5 mA, the pressure fell to 0.5×10^{-10} torr. We had two identical UHV-24P's attached to our chamber in different areas, and both exhibited this behavior. The explanation is simply that the increased current provides a better signal: more filament current \rightarrow more ionization \rightarrow more readout current, which is in the pico-amp range and so is extremely challenging for the controller to transduce. The manual specifies 10 mA for pressures below 10^{-10} torr.

3×10^{-11} torr, see Fig. 3.3(e). Corroborating this result is our ion pump: the controller reads “Low Pressure,” indicating operation below 2×10^{-11} torr.

Chamber leak rate

On Yev’s advice, we performed a rate-of-rise test monitoring for Argon permeation. The atmosphere is composed of 1% argon and it is not absorbed by the NEG, so turning off the ion pump and monitoring the partial pressure rise of amu 40 is an excellent measure of the total leak rate of the system. We observed a rise of 4×10^{-10} torr over 1 hour, which, for our ≈ 30 L chamber, is a leak rate of about 3.3×10^{-12} torr l/s, consistent with the MPF window leak rate.

With this leak rate, the system has maintained the same pressure for five years, even after jostling our experiment to a different building in 2019. We shut off the gauges in 2021 to minimize the heat buildup (near the gauge, the chamber walls reach about 60 C), trusting in the ion pump to let us know of any serious problems.

3.3 Strontium source

Building an atomic oven is a difficult, years-long endeavor; so, as at the JILA experiment [96], we purchased a complete commercial atomic beam solution from AOsense. This includes the oven, integrated Zeeman slower, vacuum-side 2D MOT optics, and was shipped to us ready-to-go at UHV pressures. Just plug and play!⁹

Oven clog and ablation attempt

After about a year of operation at 460 C, we began to notice that our flux had decreased. We discovered that strontium was gradually occluding the atomic oven, collecting on one of the heat shields after the collimating tubes. This is an issue known to AOsense, and they said they fixed it in later designs. You can see the problem quite clearly

⁹It also included a window for transverse cooling, but its proximity to the oven quickly coated it with Sr, rendering it unusable. But it did not improve loading rates by more than 10%.




Figure 3.3. (a) The science chamber blanked off and wrapped with heaters and foil for a high-temperature bake. (b) The heaters were controlled by solid-state relays (SSRs), driven by small PWM chips from adafruit. All thermocouples were connected to home-built rack a few feet away, connected to a BeagleBone microcontroller. (c) I wrote a Python GUI for the bake (still in use today in our lab, with some improvements thanks to Grady!). (d) RGA scan over a two-week bake period, showing the gradual pump-out of water and other elements. (e) The final pressure readout after cooling down to room temperature.




Figure 3.4. A view of the oven nozzle showing a clog, as seen through the Zeeman slower window. (a) Strontium deposited directly in front of the collimation tubes onto one of the heat shields, gradually blocking the atomic beam, which propagates directly out of the page. (b) Our first attempt at using laser ablation to try to clear the clog; promising... (c) We have regrets.

in Fig. 3.4(a). Because breaking vacuum is painful, we tried renting a very powerful, ns-pulsed 532 nm laser (Quantel Qsmart 450/850) to attempt to ablate the strontium and clear the blockage. However, as you can see in Fig. 3.4(c), this was a mistake: although it was powerful enough to melt the strontium, much of it melted back onto the tubes themselves, exacerbating the issue. Even worse, the laser put a small burn mark on our Zeeman slower window.¹⁰ We eventually did break vacuum to try to manually clear the blockage, but the ablated Sr was too difficult to remove.

We ended up purchasing a new oven. This was more complicated than simply swapping out the old oven, since they implemented design changes aimed at eliminating the clog issue that made the oven physically longer. This required a replacement back-half of our AOsense beamline; quite an invasive change!

Oven replacement

Strontium is highly reactive in air, quickly forming white, flaky oxide compounds. So the time the oven is spent in atmosphere must be kept to a minimum, or one risks

¹⁰Prof. Rick Averitt suggested next time we try femtosecond lasers, as those apparently ablate more cleanly; but we are unlikely to repeat this exercise!

ruining their new oven. For the replacement, we were extremely paranoid and tried to keep everything under an argon atmosphere (UHP, 5.0 grade) at all times.

In our lab, to vent anything to atmospheric pressure while avoiding a sudden inrush of air and water vapor, we perform an argon purging procedure, as follows. We first connect the part to a turbo pump and pump down the atmosphere-side for a few days. We have an RGA attached somewhere to monitor the water load, and we also have a valved-off molecular sieve (Lesker TSR4MS150QF) which sits in front of a high-purity argon line, which cleans the argon of any residual water (even in high-purity argon, the connections are inherently leaky¹¹). Once we are ready to vent, we open the sieve valve and flood the system with argon to a slight over-pressure¹². Many groups do this by flooding argon through the turbopump itself, but since the turbo blades are particularly sensitive to vibrations, we opt to do it independent of the turbo, valving it off entirely from the purge process.

After flooding with argon, we enclosed the entire oven container in a flexible bag with hands (Sigma Aldrich AtmosBag Z530204¹³), taping the bag as tight as we could over the oven port on the AOsense beam line. All tools and fasteners were cleaned and already inside the bag so that we didn't need to open it. The bag was inflated/deflated at least 3 times with argon before we opened any part to the local Ar atmosphere, at which point we tried to work as quickly as we could to re-seal everything.

After the install, we only baked the oven itself and the newly-installed NEG. With the oven at 250C, slower window at 200C, and the NEG at 2A, after only 2 hours the water peak reached 5e-9 torr, dropping to 1e-9 torr after 2 hours more. Overnight it got below 1e-10 torr. We then ran the oven at 24 hours at 400C, putting it to 500C for 1 hour

¹¹The sieve, in turn, needs baking in order to re-activate it, if it's been sitting out for a while.

¹²No vacuum components are rated for over-pressure, so this can be a risky process—windows are especially vulnerable to leaks caused by too much positive pressure. We carefully watched a high-pressure-capable capacitance gauge, Inficon CDG020D.

¹³Gloves were worn at all times, because these bags make your hands very sweaty, which can make working in them impossible.

(with the slower window at 370C) and observed no visible water peak on an RGA attached to our turbo, indicating that our procedure was a success!

New oven testing

After the oven replacement in April 2019, I measured the absorption through the center of the windows next to the oven, which benchmarks the atom flux. I used a Schafer und Kirchoff 60FC-4-M5-16 collimator (giving a beam $2w \approx 0.9$ mm) with $11 \mu\text{W}$ of power, making $I/I_{\text{sat}} \approx 0.08$, measuring with a Thorlabs PDA100A photodiode with 30 dB gain. The off-resonance frequency used for reference was 921.7148 nm (several GHz away) and the on-resonance frequency was locked with the Sr lamp at 921.72438(3) nm¹⁴. At 465 C I got 60% absorption with its ion pump reading 2.6e-6 torr; at 500 C, I got 91% absorption, ion pump reading 6.7e-6 torr.

After two years of use, this oven has also begun to decline. First, its thermocouple readings one day became unreliable, jumping around 100's of C. I had to resort to open-loop DC control of the oven current, operating it at 19 V, which was 460 C at our lab's ambient temperature. By February 2021, the oven absorption at 460 C was down to half of its former value, to 27%, and that decline has continued. As of this writing in 2022, we are unsure if there is another slow clog, but that seems likely. We currently operate it around 20 V, which is approximately 480 C.

3.4 Laser systems

As shown in Fig. 1.1, strontium has two transitions from its ground state: $^1S_0 \rightarrow ^1P_1$ at 461 nm, and $^1S_0 \rightarrow ^3P_1$ at 689 nm. These are necessary to make the mK blue MOT (bMOT) and μK red MOT (rMOT). There are a few other blue frequencies needed to operate the atomic oven, and for a successful red MOT, we also need a repumping laser

¹⁴These are our wavemeter numbers, but our wavemeter is calibrated to the 689.4 nm transition; so the absolute frequencies listed here are expected to be accurate only up to a few hundred MHz.




Figure 3.5. Simplified 461 nm setup (“to nanotraps” indicates the beam continues to the nanotrap experiment, omitted here), with AOM frequencies in MHz. All beamsplitters are polarizing; half-waveplates are in front of all fibers, which are polarization-maintaining.

at 481 nm. For evaporation, we need 1064 nm, and the SM lattice / OSG beams require some 689 nm light. This section will describe the optical layouts of these setups.

3.4.1 461 nm laser setup

The transition $^1S_0 \rightarrow ^1P_1$ is driven by 461 nm light, and our AOM plan is shown in Table 3.1. We generate 461 nm light using a 532 nm pump laser (Lighthouse Photonics Sprout, 15W) paired with a Ti:Sapph (M Squared SolsTis) to produce 921 nm, which is then frequency-doubled (with an M Squared ECD-X) to 461 nm. The total usable output of the system is about 700 mW when the Sprout is operated at 12 W.

The blue laser setup is on a different table from the experiment and is shown in Fig. 3.5. The AOMs are all from Gooch&Housego¹⁵, with polarizing beam splitters (PBSs)

¹⁵Models: 3080-120, 3110-120, 3200-125, 3350-192.

Table 3.1. Blue AOM frequency table; units in MHz.

Isotope	Isotope shift	Spec. AOM	Overall shift
^{84}Sr	-270.8	405.4	-610.8
^{87}Sr	-124.8	332.4	-391.6
^{88}Sr	-51.6	295.8	-340

from Lens Optics GmbH¹⁶. All beams are fiber coupled using polarization-maintaining APC patch cables (Thorlabs P3-405BPM) coupled via adjustable collimators from Schafter und Kirchoff (60FC-4-M5-33).

The laser sources have had their share of issues over the years. On initial install, the Msquared ECD-X had issues locking due to the lab humidity. The first doubling crystal burned due to this, although we weren't sure of the cause at the time. We installed a new doubling crystal but couldn't get it working; after a few months of back-and-forth with MSquared, we were sent a lid which contained a desiccant, solving the issues. They are now shipping desiccants as standard options on their laser cavities.

The SolsTis is a “black box”, and we were not allowed to open it ourselves; so when internal components drifted significantly, we observed a gradual power decline resulting in eventual loss of ability to run the cavity lock. We had our SolsTis sent back to Scotland once for repair, which took a few months. After two years, it began to lose power again, and we replaced it with another SolsTis.

The Sprout has been mostly reliable; it wasn't until year 5 of operation, around 40,000 hours of use, that we noticed stability problems and gradual power decline. We replaced it in 2021.

¹⁶We operate above the damage threshold on many of these PBSs, but the resulting degradation is still better than the transmission of Thorlabs PBSs at 461 nm.

Spectroscopy

Although the SolsTis has etalon and cavity locks that stabilize the frequency, there still exists a slow drift on the order of a few MHz/minute, as observed on our wavemeter (HighFinesse WSU-8). We spectroscopically correct the drift with a hollow-cathode lamp (Hamamatsu L2783-38NE-SR¹⁷), which is a small glass cell which contains strontium sealed in with an Ne buffer gas. We perform polarization spectroscopy following the setup of [98]. I leave the details to the reference; in short, we use a strong circularly-polarized “pump” beam to saturate the $^1S_0(m_J = 0) \rightarrow ^1P_1(m_J + 1)$ transition of ^{88}Sr . A counter-propagating probe beam¹⁸ tuned near resonance then experiences birefringence, rotating its polarization as it passes through the Sr experiencing the pump light. When scanned through resonance, this produces a nearly-linear error signal on a difference photodiode, which we then lock to using PI control. We use probe and probe beam diameters $2w = 120\mu\text{m}$, and typical probe (pump) powers are $140\ \mu\text{W}$ ($1.0\ \text{mW}$).

There are a few practical details to note. The rated lamp lifetime is 5000 mA-h, and we have operated at the minimum current 8 mA. We are well over the lifetime by now, and have noticed a gradual drop in Sr signal, but the lamp remains usable. Second, any polarization drift will cause a drift in the error signal. Thus, while we do use polarization clean-ups like Wollaston prisms and PBS’s, the lock is especially sensitive to temperature and humidity changes, which modulate the glass’ indices of refraction. Opening the doors surrounding the optical table easily leads to a few-MHz drift, and on the rare days when the AC is malfunctioning, it is better to not even try running experiments. Finally, since we do not lock to the precise zero of the dispersion-like lock curve, the lock point is very sensitive to intensity changes, so this must be actively stabilized as well or the lock point will drift. Satisfying these requirements, we use RedPitaya’s (see Section 3.7.4) to feed a

¹⁷Unfortunately, these are no longer sold by Hamamatsu, nor by any other vendor I know of.

¹⁸Since we use a D-mirror to introduce the pump beam, it is not perfectly counter-propagating, but at an angle of about 10 mRad.

slow correction voltage to the cavity lock on the SolsTis, since we have no need to narrow the (≈ 100 kHz) laser linewidth.

Zeeman slower




Figure 3.6. Simplified AOsense optics setup; the 2D MOTs are at left, Zeeman slower at right. The dashed boxes indicate that these optics are on the under-side of the breadboard; only the last mirrors are top-side, with the beams coming up through large cuts in the tables. The 50/250 mm lenses are cylindrical, with the arrows indicating focusing along \hat{y} , and encircled-dot along \hat{z} .

The Zeeman slower is one of the most important beams in the experiment, and generally determines the loading rate into the 3D MOT. The slope of loading rate vs. power dependence is very linear, until you hit about 60 mW as measured at the slower window, at which point it starts saturating. This is where we typically operate.

The beam is brought to the experiment via SM-maintaining fiber and is shown in Fig. 3.6. It is detuned from the transition by -540MHz, and its diameter at the slower window is $2\omega = 6.5$ mm. Most of its optics are on the underside of the experimental breadboard, and consist very simply of a cleanup PBS, half- and quarter-waveplates, a pick-off, and the shaping telescope, respectively. The polarization of this beam is left-handed circular, as viewed from the polarimeter's perspective (from the perspective of the

beam's k-vector, that makes it right-hand circular).

2D MOTs

These beams propagate almost entirely on the underside of the breadboards, which was quite difficult to align, and reach the top through large holes in the breadboard. We use cylindrical lenses to shape the beam to $2w = 30\text{mm} \times 5\text{mm}$, with the short axis along gravity. We use independent sets of half- and quarter-waveplates in each path, because waveplates at this size are quite pricey. The power we use at each window depends on the available 461 nm power, but is typically on the order of 6-9 mW total per beam. The polarization for each beam is left-handed circular as viewed from the polarimeter's perspective (from the perspective of each beam's k-vector, that makes them right-hand circular). See Fig. 3.6 for the layout.

3D MOT

Similar to our 2D MOT beams, the horizontal blue 3D MOTs are shaped on the underside of the breadboards, reaching the experiment via holes at the window positions. The beams are about $2w \approx 10$ mm in the horizontal plane, and 6 mm in the vertical direction. The horizontal beams at the experiment are shown in Fig. 3.10. The vertical-going beams are shown in Fig. 3.8.

3.4.2 689 nm laser setup

The laser setup for the $^1S_0 \rightarrow ^3P_1$ transition at 689.4 nm is shown in Fig. 3.4.2, and Table 3.2 shows the needed frequencies. The laser source is a littrow-stabilized diode from Toptica operating at 25 mW, amplified by a tapered amplifier (TA pro). These TAs are (evidently) hard to make at this wavelength: our first TA was rated for 400 mW output, but that quickly degraded, and we replaced it when it reached 90 mW about 4 years later. The replacement outputs only 170 mW at best.

Table 3.2. Red AOM frequency table; the units are in MHz.

Isotope	Shift from ^{88}Sr	Laser detuning	Needed frequency
^{88}Sr	0	-160	2×80
^{87}Sr ($F=9/2$)	221.699	-381.699	2×190.85
^{87}Sr ($F=11/2$)	-1241.451	1081.451	4×-270.363
^{84}Sr (removed)	-351.49	191.49	2×95.745

As with the blue setup, we use Schafter und Kirchoff (60FC-4-M5-16) collimators on the laser table, here opting for AOMs from IntraAction¹⁹. The setup is configured to allow operation with either ^{87}Sr or ^{88}Sr . In the past, when we were testing BECs, we also had an AOM path dedicated to ^{84}Sr , but that has since been removed for lack of use.

I did have to replace the diode once at the 1754 hour mark, when I could no longer reach 689.4 nm. But otherwise, the setup has been relatively robust, aside from the gradual TA degradation. The TA beam pointing needs to be corrected for every couple of weeks, as we observe slow drifts in our fiber-coupling efficiency. Most fiber coupling efficiencies are above 60%. The largest source of loss is the quad-pass needed for the fermion (see Table 3.2), which eats 50 mW of power to give only about 9 mW at the experiment.

At the experiment

Fig. 3.10 shows the xy plane of the 689 laser setup at the experiment; the vertical MOT section is shown separately in Fig. 3.8. The stirring beam is split off both to power the SM lattice experiment, and to function as a pump beam resonant with the transition $^1S_0(m_F, F = 9/2) \rightarrow ^3P_1(m_F, F = 9/2)$. This pump beam is useful for kicking out specific m_F states, and is directed onto the camera through achromatic lenses to do red imaging. The 689 MOT and stirring beam intensities are listed in Table 4.1.

¹⁹Models: ATM-3001A1.12 (cavity); ATM-2901A1.12/ATM-2701A1.12 (quad-pass); ATM-801A1.12 (spec-cell and 88-Sr); ATM-1702DA.1 (stirring/SM lattice).

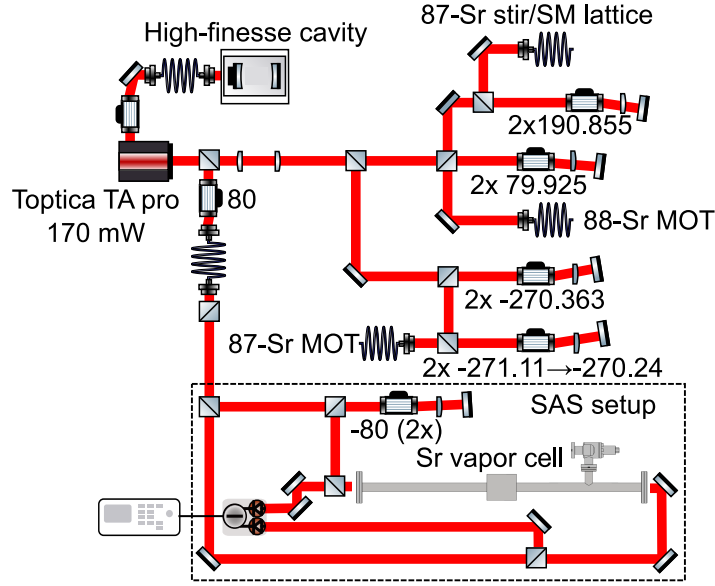


Figure 3.7. Simplified 689 nm laser setup, with AOM frequencies in MHz. All beamsplitters are polarizing, and half-waveplates are in front of all fibers, which are polarization-maintaining. After using it to find the narrow 689 nm transition, the SAS setup has not been used for several years; the SLS cavity has proven stable enough to render it superfluous.

Electronically switchable gain photodiode amplifier

Servo control of the laser intensity is important in many places in the experiment (for a general description of how we do this, see Section 3.7.4), especially when trying to reach single-digit μK temperatures in the red MOT. On the one hand, the Doppler cooling limit T_D scales with the natural linewidth as $k_B T_D = \hbar \Gamma_{\text{eff}} / 2$, so any power broadening will scale the achievable MOT temperature as $\sqrt{1 + I/I_{\text{sat}}}$, where $I_{\text{sat}} = 3 \mu\text{W}/\text{cm}^2$. On the other hand, at the beginning of the MOT when the atoms are hot, we need large capture velocities, necessitating $I \gg I_{\text{sat}}$. In practice, this means we must ramp the lasers from a few mW to a few tens of μW . Ensuring that the monitoring photodiode doesn't saturate on the high end means it's not going to be able to detect anything on the low-end; in other words, its *dynamic range* is very limited. Ideally, we would have some way to electronically switch the photodiode gain in-sequence, with low gain at the beginning of

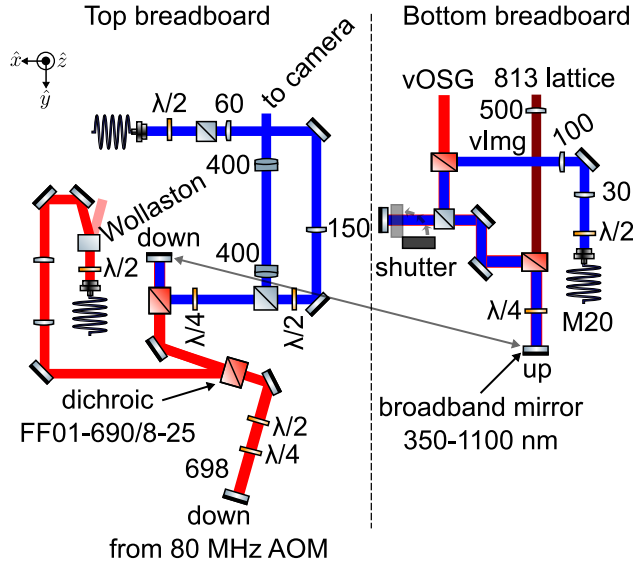


Figure 3.8. Simplified top-most and bottom-most breadboard optical layouts, showing the vertical MOTs and vertical imaging (vImg) beam paths. I did some brief work with the clock transition, hence the path for an 813 nm lattice and 698 nm clock beam.

the rMOT sequence, and higher gain towards the end. Commercial solutions simply do not exist as of 2022; variable-gain photodiodes all have manual (hand) switches.

Enter my switchable-gain transimpedance amplifier, the schematic of which is given in Fig. 3.9, and PCB design in Appendix G. It is based on the Hamamatsu S3072 biased at -12 V (although any photodiode with a similar pinout will work). The design is very simple: it uses a fast op-amp (OPA657²⁰) in a transimpedance configuration on a 2-layer PCB. In the feedback loop are two separate gain paths, set by resistors differing by $30\times$. Control over which path is taken is provided by a high-speed, single-pole double-throw switch (ADG1219).

I isolate the circuit from the rest of the timing hardware with a simple optocoupler (PS2501-1-A), finding it is especially important to avoid ground loops that drive 60 Hz signals on the circuit. The measured bandwidth, which scales like $1/\sqrt{R_F}$, is 1 MHz at high-gain ($R_F = 500k\Omega$), and 8 MHz at low-gain ($R_F = 16k\Omega$), more than sufficient for

²⁰Although the OPA847 is a much faster part in the same family, it is harder to make it stable, and frankly more finicky than it's worth.

our purposes. 3 uW of incident power produces ~ 380 mV, so we have plenty of range at the low-end!

Some care must be taken when the gain switches, because for a brief time ($\ll 5\mu\text{s}$) the op-amp has no feedback and starts to “rail,” or drive its output high or low. So there is a discontinuity and slight power jump at the atoms due until the op-amp recovers; for a MOT, this doesn’t matter and the atoms don’t notice. For more sensitive applications, the simplest solution would a simple design change, pointed out to me by an engineer at AOsense: simply put the RF switch in series with the small resistor, in parallel with the big resistor, so that the op-amp always has some feedback!

I have two ‘fun’ anecdotes about the design/troubleshooting of this circuit. In the first version, the RF switch had a 20% chance to literally catch on fire when I plugged it in. This was...troubling. It turned out to be the long (10 ft) cables I was using to plug it in, which had live 15 V on them going directly to the RF switch. You might not think this is a problem (indeed, I didn’t), but those cables have significant inductance. This means energy is stored in the magnetic field, which gets released when plugged in, causing a large (transient) voltage overshoot at the switch itself. This overshoot reached upwards of 20V (I measured it!), which exceeded the switch’s breakdown voltage, causing a complete short and subsequent smoldering blaze. I simply added some series resistors to dampen the overshoot, and that worked.

What was less fun was troubleshooting a non-functioning voltage regulator. Referring to the schematic, U3 was initially a +5V regulator, the LDL1117S50R. But it was not turning on when the negative supply (sourced by U4, the LT1175) came online first—apparently, enough capacitatively-coupled negative voltage got present on the traces to inhibit proper operation. Having no patience for this, I simply replaced it with the less-elegant but more-robust regulator, LM7805CT. The schematic has been updated to reflect this change, but note that the PCB drawing hasn’t yet, since I haven’t had need of more boards.

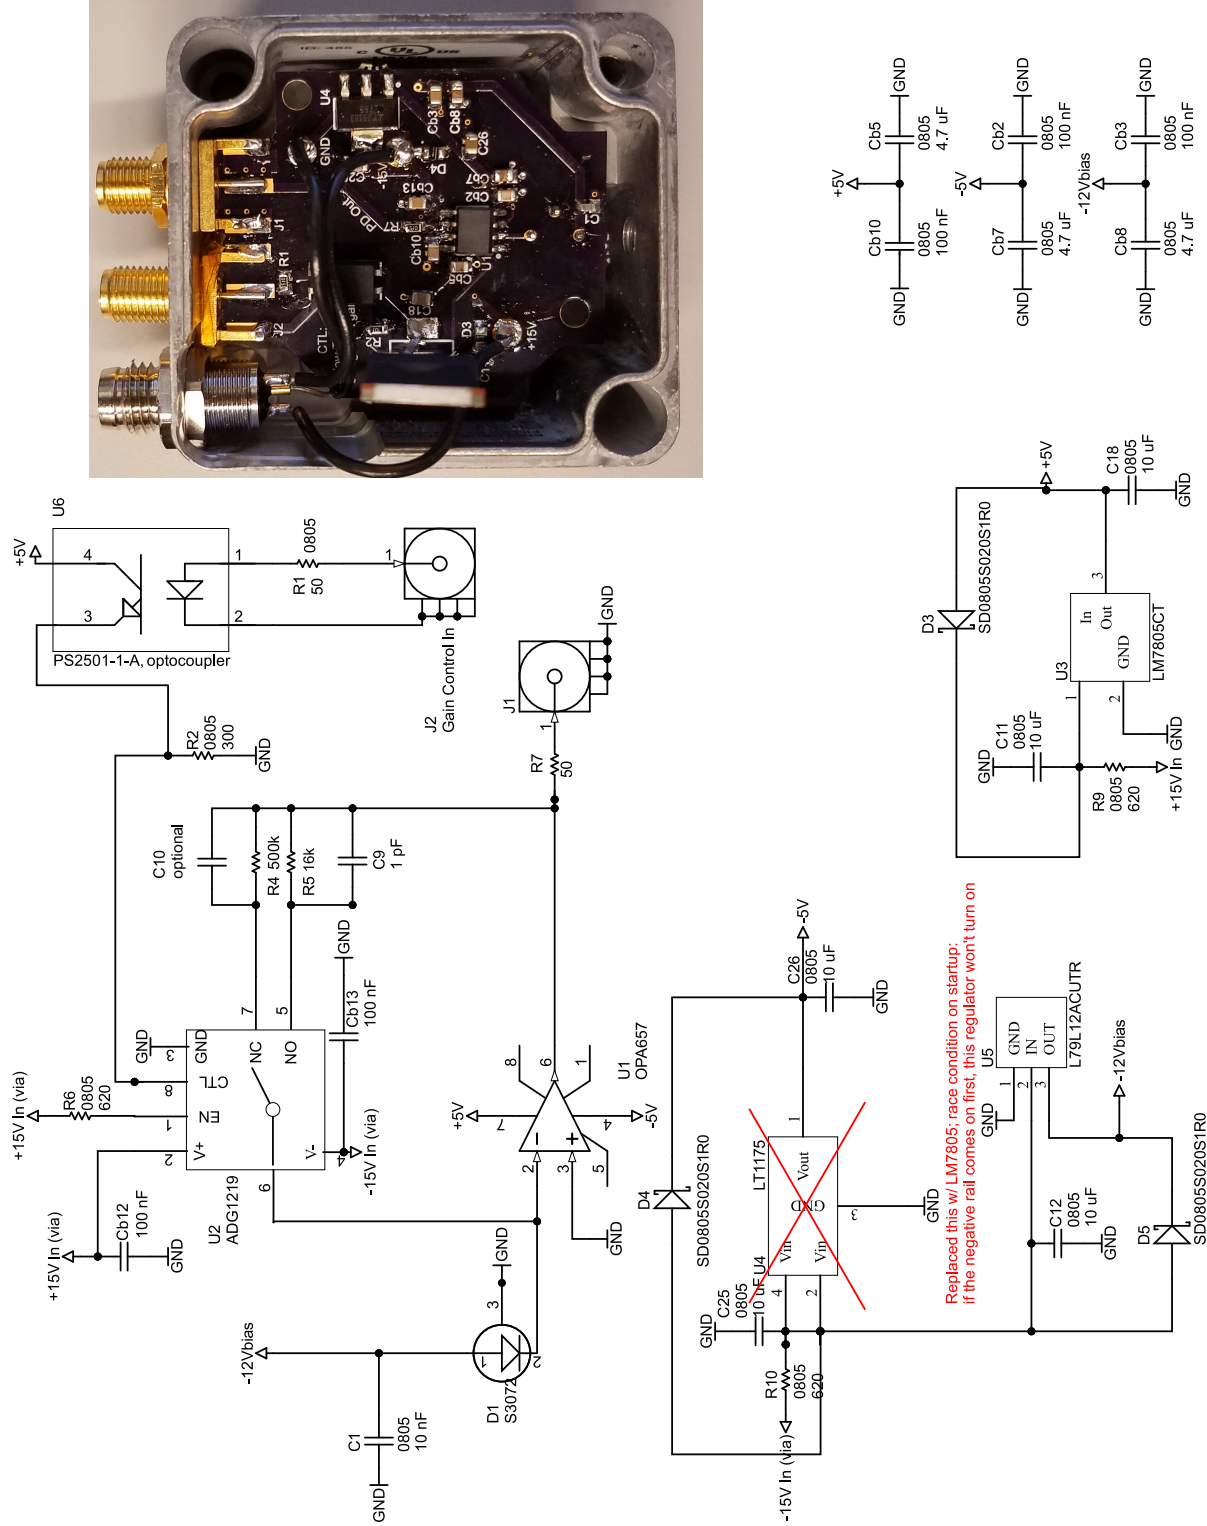


Figure 3.9. Electronically-switchable-gain transimpedance amplifier.

Optical cavity

Owing to its narrow linewidth $\Gamma/2\pi = 7.4$ kHz, the $^1S_0 \rightarrow ^3P_1$ transition at 689 nm has much stiffer requirements than the 461 nm transition ($\Gamma/2\pi = 30$ MHz). With significant effort (and expense), a free-running diode or Ti:Sapph laser can achieve 100 kHz linewidth (as measured over some period of time, usually $5 \mu\text{s}$). This is not even close to good enough for narrow-line spectroscopy: most of the power will be detuned by 10s of kHz, rendering the MOT ineffective and any spectroscopy hopelessly power-broadened and light shifted²¹.

The most common approach to narrow a laser's linewidth is to use a high-finesse optical cavity to provide feedback to the laser current (which controls the frequency); this is called the Pound-Drever-Hall (PDH) technique [99]. We use such an approach here, with a complete solution provided by Stable Laser Systems (SLS). The cavity itself is dual-coated for operation at 689 nm and 698 nm, with a preference on the latter. The specified linewidth at 698 nm is < 1 Hz at 1 s measuring time, and at < 1 kHz at 1 s at 689 nm.

It certainly had its share of problems on install. On the 689 nm side of things, the fiber polarization axis was not aligned with the incoming light, leading to intensity modulation of the laser; I consequently observed a lot of DC drift in the PDH error voltage on thermal timescales. I thus incorporated intensity-stabilization with an AOM, and corrected the polarization alignment. I also observed a lot of 120 Hz noise in their frequency-shifter circuit, so I removed it.

On the 698 side, I had to do a lot of work stabilizing the fiber-noise cancellation interferometer, which kept getting misaligned. I added some waveplates to correct for its polarization drift as well, but did not incorporate intensity stabilization as I did with the 689 path. The 698 fiber was butt-coupled to EOM fiber, which led to a poor polarization

²¹We note that red MOT operation would still be possible, as in David Weld's lab, albeit at somewhat higher temperatures

matching, so I added a couple collimators and a half-waveplate to achieve better matching. With these minor changes, the SLS cavity has performed with exceptional reliability over the years. I used the 698 nm clock transition to measure its linear drift rate of 4.48 ± 0.06 Hz/min, or about 6.45 kHz/day.

3.4.3 SM lattice setup

All beams for the spin-momentum (SM) lattice experiment are sourced from the cavity-stabilized stirring beam, which propagates on the experiment table; see Fig. 3.10 for the layout. At the atoms, these beams are detuned from the $F=9/2$ transition by a net shift of $\Delta/2\pi = 210$ MHz (-290 MHz + 80 MHz). Initially, we were using 80 MHz AOM's to provide a net -160 MHz shift, but to our dismay found that significant amounts of 0th order light was scattered by the first AOM's crystal²² and caused noticeable heating of the atoms. We were unsuccessful in our attempts to spatially filter out this light with a pinhole, instead swapping out the 80 MHz AOM with one at 290 MHz, which has a much larger diffraction angle that naturally separates the crystal-scattered light. This resolved the issue.

The ex-situ measured waists of the SM-lattice beams $k_1, k_2, k_3 \approx 170, 190, 190 \mu\text{m}$, respectively; these are in good agreement with in-situ Raman coupling strength measurements. The beams are picked off and monitored with photodiodes (Thorlabs PDA100A2); since the experiment requires multiple RF tones in each AOM, interfering signals around 100 kHz naturally appear, which must be filtered out to measure the DC light level. I use elliptical RC low-pass filters on these photodiodes, with -3dB corners of around $20\text{-}50$ kHz (Thorlabs EF124²³).

²²This light was due to optical imperfections in the crystal, not due to any residual RF power.

²³Not that you would... but don't even try opening these filters. Thorlabs used a thread-locker on the BNCs, preventing the metal shell from easily opening. While it can be forced, it rips the BNC connector off the PCB.

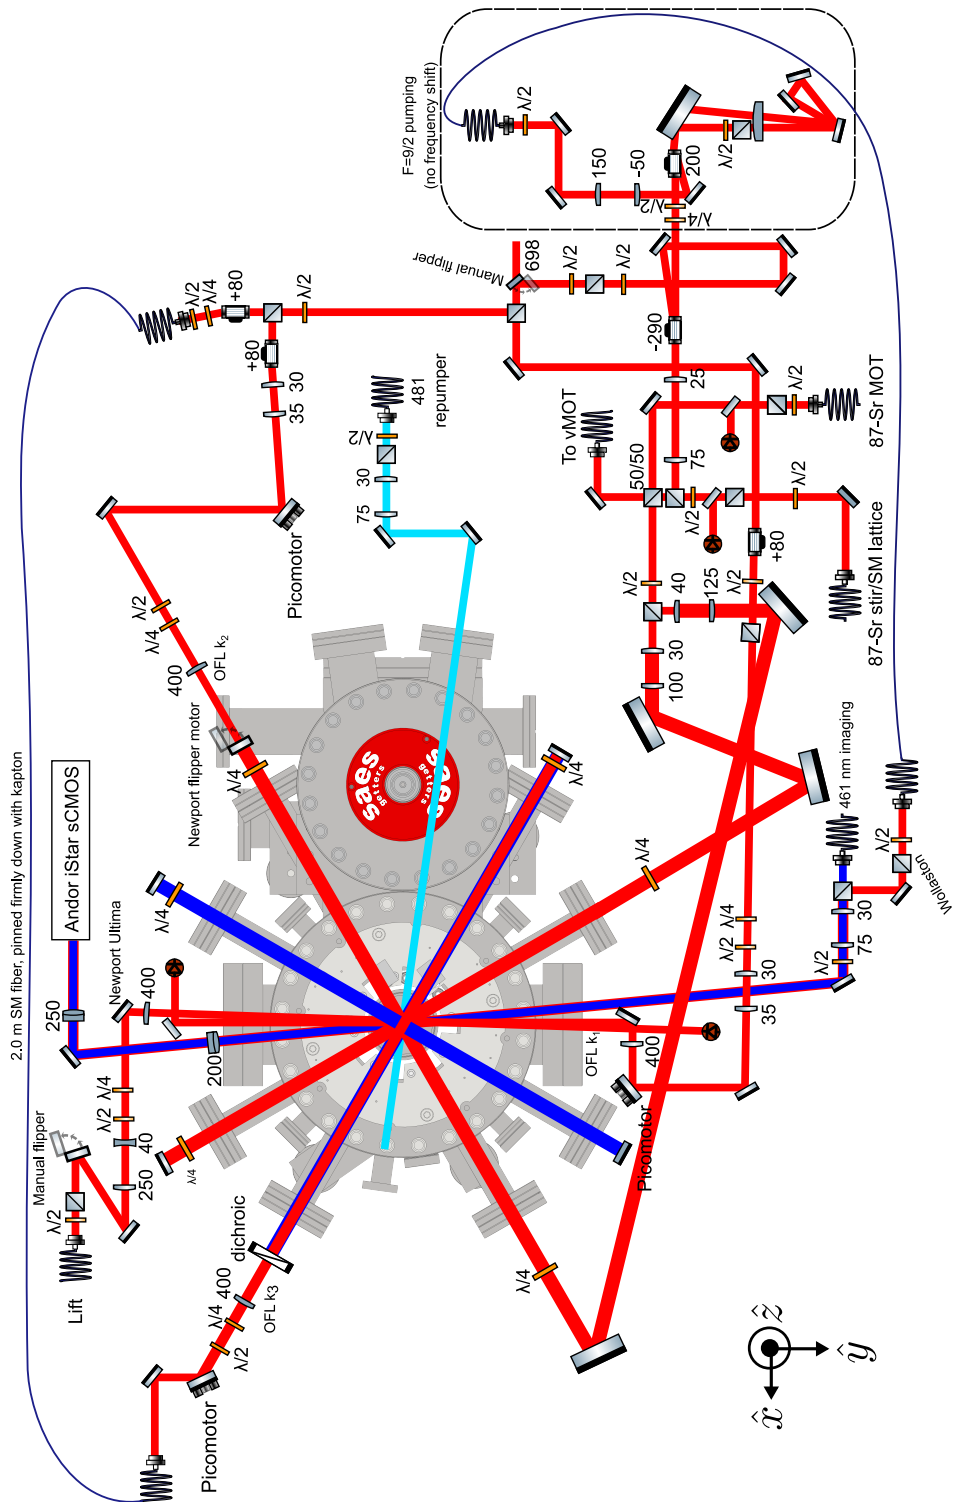


Figure 3.10. Simplified optical layout at the experiment in the xy -plane, without ODT or AOsense optics (see Fig. 3.12, Fig. 3.6).

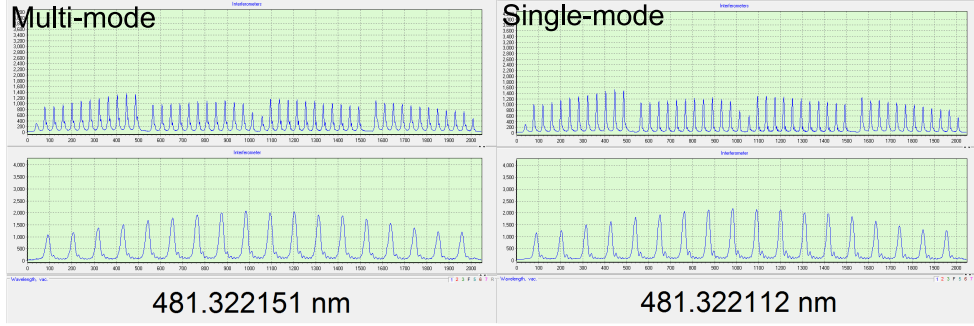


Figure 3.11. The subtle differences between multi- and single-mode operation of the 481 nm diode are sometimes only visible in the wavemeter interferometer displays.

3.4.4 481 nm laser

Atoms in the bMOT have a 1:50000 chance of decaying from 1P_1 into 3P_2 , which is metastable and magnetically-trappable. We take advantage of this fact by accumulating atoms in that state during the MOT loading stage, and then “repump” them back down to the ground state. Most strontium experiments use a combination of 679 nm and 707 nm to do this repumping, but Killian’s group used a single solution at 481 nm [100, 101], which we adopt.

We use a grating-stabilized diode (Toptica DL 100) which outputs about 16 mW of light. We need only a few mW at the experiment. For ^{88}Sr , we keep the vacuum wavelength fixed at 481.322 nm and it works very well. The fermionic isotope has ten spin states, all of which need to be addressed; so, we frequency-modulate the light using the integrated piezo at ≈ 2 kHz (on the SC 110 piezo module, this corresponds to “8” on the frequency range, and “2” on the amplitude knob)²⁴. This light is sent directly to the experiment through a single-mode fiber, expanded to a ≈ 4 mm waist; the (relatively simple) optical layout is shown in Fig. 3.10. We use a laser shutter to gate the light at the atoms, which has been so remarkably repeatable we haven’t had the need to do any AOM switching.

²⁴Killian’s group modulates the current instead, but we had no success doing that—it just caused lots of mode-hops.

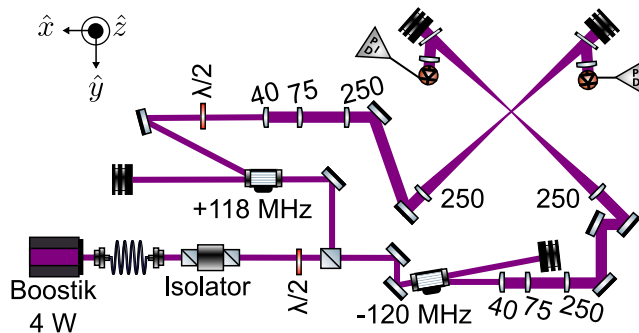


Figure 3.12. The 1064 nm ODT setup. Focal lengths are in mm; the 75 mm (250 mm) lenses are cylindrical lenses focusing along the \hat{y} (\hat{z}) direction. No beam blocks are water-cooled.

This 481 nm diode itself has performed without power drops for years. But, as with most diode lasers, some care must be taken to ensure its operation is single-mode at the wavelength of interest. Evidence of multi-mode operation can be seen directly in the wavemeter; see Fig. 3.11.

3.4.5 1064 nm dipole trap

In order to reach quantum degeneracy, evaporative cooling occurs in an optical dipole trap (ODT) [102]. The laser source here is an NKT Koheras BASIK Y10 seed laser operating at 1064 nm, which is fiber coupled and provides 10 mW of power at an amazing 10 kHz linewidth, making it attractive for dual use as a lattice. This seed light is amplified up to 15 W with an NKT BOOSTIK amplifier²⁵. This laser has been extremely reliable over the years, and the amplifier output is also fiber coupled, providing excellent pointing stability. The setup is shown in Fig. 3.12 (which is not to scale). We do not water cool any beam blocks, but take care to keep them from fibers or other sensitive components, reflecting the to-dump beams across the table about 1 m away²⁶.

Two acoustic deflectors (IntraAction ATD-1202DA6) act as fast optical switches

²⁵In practice, we do not use more than 4 A drive current, providing about 4 W.

²⁶It is very important to physically separate heat sources (such as oscilloscopes or beam dumps) from the output fiber; I have observed beam misalignment (as evidenced by atom loss) as a result of heat-induced drift, although in all cases the alignment returns when the heat source is removed.

and allow intensity control. They shift in opposite directions to avoid interference effects at the atoms. In the same vein, the two beams have orthogonal polarization—although in practice, this latter measure was overkill and didn’t seem to matter. These deflectors were chosen to allow the use of center position modulation (CPM) [103] to dynamically shape the optical trap and improve the evaporation efficiency. However, our early experiments with this were unsuccessful, as we noticed high-frequency noise that caused atom loss. We didn’t spend too much time on it, however, and intend to return to this point in the future. For now, we focus the beams onto the atoms statically, with ex-situ measured waists $w_{xy}, w_z \approx 50, 15\mu\text{m}$. From the trap frequency measurements in Fig. 4.5, the waists are actually closer to $45, 20\mu\text{m}$.

ODT intensity control

Evaporation requires a smoothly-changing intensity profile, fiercely guarded against any noise present at multiples of the trap frequency, which parametrically shakes the atoms out of the trap. In practice, this means intensity noise below 1 kHz must be suppressed, and higher-frequency noise is less important.

We observe the beam power on Thorlabs PDA100A2 photodiodes (bandwidth 10 MHz)²⁷, using pick-off mirrors (Thorlabs BSF20-C) in front of the beam dumps at a small angle (an angle made as small as possible, in order to minimize polarization dependence). The signal is conditioned with a difference circuit (see Fig. 3.13), which subtracts the ADWin’s analog setpoint with the photodiode voltage. Keeping this difference at 0 V is the job of the PID controller, here a RedPitaya. This approach affords enormous flexibility in programming arbitrary intensity ramps, and greatly enhances the RedPitaya’s very limited dynamic range²⁸. To protect the RedPitaya from over-voltage, I used a couple of

²⁷We have heard tell of silicon photodiodes reacting overly-slowly to 1064 nm light, owing to its smaller photon energy; GaAs photodiodes have been recommended in some theses, such as Ruwan’s; but we haven’t noticed problems yet

²⁸For example, an ODT set power of ≈ 2 W would produce 6 V output from the photodiode—but the RedPitaya front-end can handle only ± 1 V on its low-gain setting.

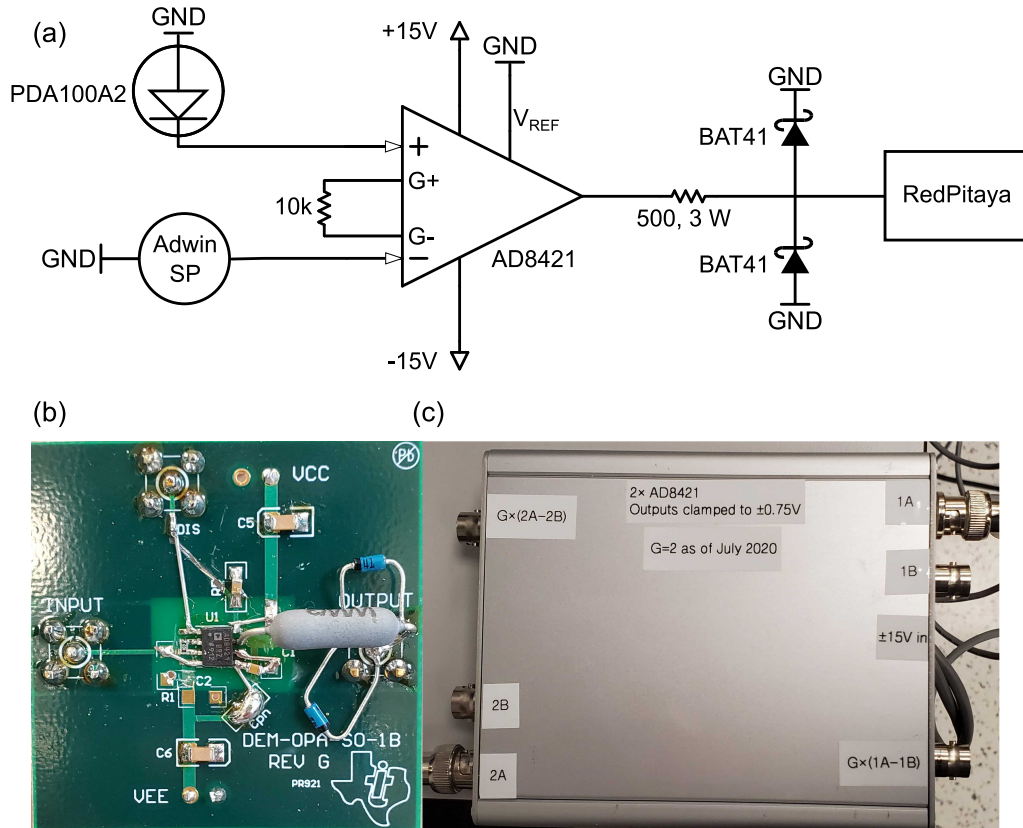


Figure 3.13. Signal conditioning circuit for ODT intensity control. (a) Schematic; a precision op-amp, the AD8421, is configured as a differential amplifier that subtracts an analog setpoint (SP) and the photodiode signal, which is then sent to a RedPitaya for PID control. (b) What it really looks like, soldered onto a handy high-speed demo board from Digikey (TI DEM-OPA-SO-1B). (c) Since we have two ODT beams to control, I made two of the circuits in (b) and put them in a box.

Zener diodes with a beefy 1% series resistor to, in turn, protect the Zeners.

3.4.6 Optical Stern-Gerlach

Inspired by Stellmer’s implementation of the optical Stern-Gerlach effect [104], we built our own version, with a few experimental differences. Since cavity-locked 689 nm power is at a premium and there is no narrow-linewidth requirement to do OSG, we use instead another SolsTis laser pumped by a 532 nm, 15 W MSquared Equinox. This beam goes through a +170 MHz AOM for fast switching and is fiber coupled at the experiment. We have implemented two versions of the OSG beam, one horizontal (hOSG) and one




Figure 3.14. Simplified setup of the vertical optical Stern-Gerlach (vOSG) beam, omitting waveplates and shutters. Two lenses are cylindrical; the arrow beneath the -50 mm lens indicates the focusing direction is \hat{x} , and the encircled dot on the 200 mm lens indicates focusing along \hat{z} . M20 refers to the launch collimator (Schafter und Kirchoff 60FC-4-M20-10).

vertical (vOSG). The setup of the vOSG is shown in Fig. 3.14, and the resulting atom images are exemplified in Fig. 3.15.

The hOSG beam offers better spatial resolution, since it propagates orthogonally to the quantization axis and is horizontally-polarized—and can thus drive σ^+ as well as σ^- transitions. Since we have power to spare, we chose to make a large cylindrical beam, with $1/e^2$ diameters $2\omega_z, \omega_x = 120, 480\mu\text{m}$ at the atoms, requiring 30 mW at 434.827970 THz in a ≈ 7 G bias field²⁹. The beam’s pancake shape minimizes the gradient force along the x -direction, largely preserving the momentum distribution and producing nicer images than in [104].

We also implement the vOSG, with the original intention to observe dynamics in the plane of the optical table—but now it’s used simply because the horizontal path was co-opted by the Raman lift beam (see Section 5.1.3). The vOSG images are less resolved, since it can drive only σ^+ transitions³⁰. This beam is also elliptical, with waists $2\omega_x, \omega_y = 120, 430\mu\text{m}$, requiring about 100 mW to operate in a 9 G bias field.

The OSG has proven instrumental to this work, providing instant feedback on spin-polarization optimization, and for spectroscopy. While one can do narrow-line imaging

²⁹On our wavemeter, this reads 434.827800 THz, since it is picked off before the AOM.

³⁰Its polarization is fixed, since it is introduced with a PBS along the upward-going MOT path.

as in [104], and indeed I do this where appropriate, such imaging has poor signal-to-noise ratio owing to the small saturation intensity.

We have had enormous difficulty obtaining stable lasing from the MSquared Equinox. Originally based on a fiber laser design, their laser was a new product which proved quite unstable; I had issues with it overheating and high-frequency oscillation. Their controllers also had overheating issues, and we had to be issued new ones at several points. Eventually, the Equinox completely failed to turn on; to date, we’ve had 2 complete replacements of the Equinox. The last unit was installed on 2020-12-14; this one seems to have had all its bugs worked out, operating well at the time of writing (May 2022)³¹. Despite the challenges, their customer service has been absolutely stellar, and their willingness to send us free replacement lasers certainly wins them points.

3.4.7 Lift beam laser

In order to lift the degeneracy of the ground $^1S_0(m_F)$ states for the purposes of the flux lattice project, one cannot rely on magnetic field. We require ≈ 100 kHz shifts; given the the nuclear state splitting of 185 Hz/G [19], this would require experimentally-impractical fields of 500-1000 G. Furthermore, the energy levels split in this way would be uniformly spaced; we explicitly require nonuniform spacing.

So, we turned to light shifts and create a “lift” beam, inspired by the Yb spin-orbit coupling experiments [105, 106]; see Section 5.1.3 for more details. Given that we already had a vOSG powered by a several-Watt Ti:Sapph, we co-opted the hOSG path for this. The lift beam focuses ≈ 180 mW at beam waists $(\omega_x, \omega_z) = (350, 330)\mu\text{m}$, which is much greater than the $\approx 30\mu\text{m}$ in-situ DFG size. The beam’s polarization is vertical, and aligned to the vertical axis with a Glan-Taylor prism.

³¹But we are still leary of ramping it all the way off, for fear that it won’t ramp back on!

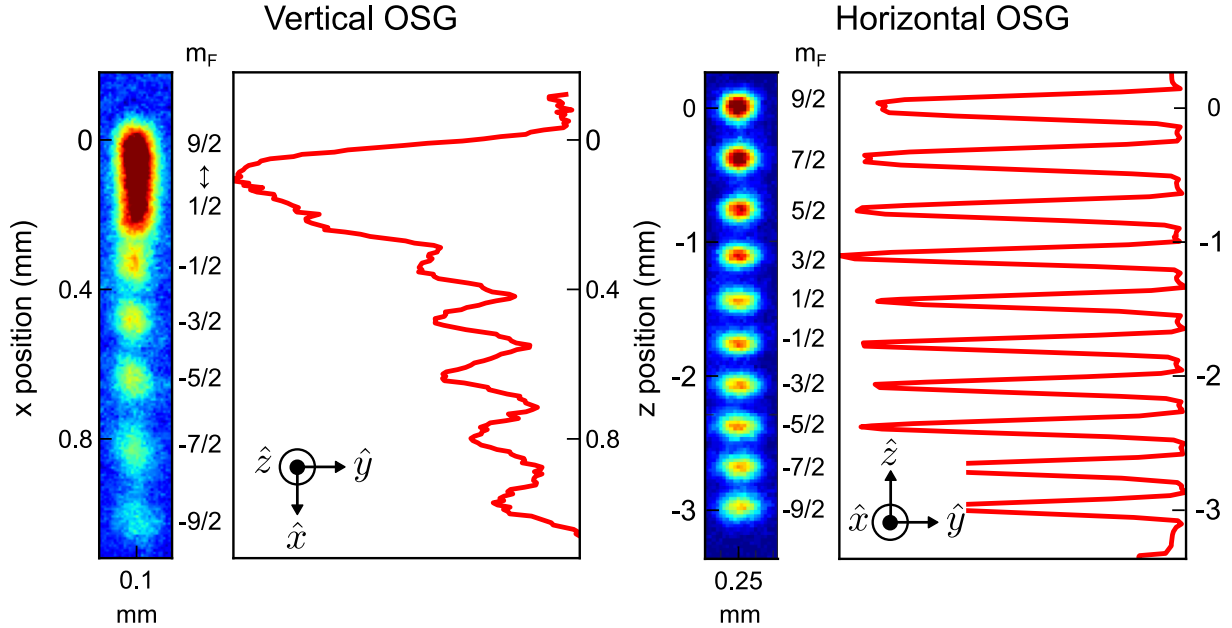


Figure 3.15. Typical single-shot images demonstrating optical Stern-Gerlach (OSG) separation of a DFG ($\approx 50 - 70$ nK), with the traces at right showing the row average of the image. We have two setups for this: the ‘vertical’ OSG (vOSG, left) beam propagates against gravity and along the quantization axis. The ‘horizontal’ OSG (hOSG, right) beam propagates orthogonal to gravity. Neither beam is retro-reflected. The hOSG can drive both σ^+/σ^- transitions, so all m_F states are resolvable. In contrast, the vOSG beam is right-hand circularly polarized and can drive only σ^+ transitions, so it exerts a significant force only on atoms with $m_F < 0$.

3.4.8 Laser shutters

Optical shutters are needed in many places, owing to the finite extinction ratios of AOMs. In particular, the DFG must be shielded from any 461 nm stray light. We build our own shutters using off-the-shelf parts, control them with a custom PCB, and communicate with them via a UART protocol driven by our timing system. Although this is more complex than other solutions³², this approach offers better repeatability and very fast (few-ms) rise times. As measured with a photodiode and the shutter’s approximately 2 in. arm, the close time (90% to 10%) is 1.7 ms, and the fastest optical pulse that can be driven with the shutter itself is about 5 ms. This is competitive with commercial solutions, and is in fact much better than the shutters offered by Thorlabs (the SH05R or SH1),

³²For example, in David Weld’s lab they have small Arduino chips waiting on a trigger.

which offer 8 ms close time at best³³ at a fraction of the price: the stepper motor, driver, and associated hardware cost about \$30 USD, vs. the \$1500 USD cost of the Thorlabs solution.

Hardware

The shutter actuator is a 0.5A/10V stepper motor (Pololu 1208), with an aluminum mounting hub (Pololu 1203) attached to its shaft. This holds a bolted-on anodized aluminum sheet, which acts as the beam block. I used a sheet metal press to cut out 2 in x 0.5 in arms. Their light weight makes it easy for the stepper to swing, while being rigid enough for repeatable movements. The shutter base is epoxied to 1 in. optical posts, with 0.25 in. of vibration-damping sorbothane directly between it and the post body, which in turn is on a square sheet metal cutout to provide structural stability. See Fig. 3.16.

Driver and PCB

The steppers are driven by the Allegro A4988, mounted on a breakout board (Pololu 2986). These chips make it extremely easy to drive the shutter and keep them asleep when they are not used, which is important because the idle power draw can heat them up to 180 F, which would disturb nearby optics. The protocol is simple: a few ms before the intended movement, wake up the shutter with a digital high on the sleep line; then send a series of digital, 1-ms wide pulses. The stepper takes a clockwise (or counter-clockwise, depending on the direction bit) step at every pulse.

I made a PCB (see Appendix F.1 for schematic) to mount 12 of these stepper motor drivers. All connect to a BeagleBone Green (a Black works also), which coordinates sleep and direction settings. The board connects to the individual shutters through CAT5 cabling, with each cable containing the power and control lines for two independent

³³Strictly speaking, this comparison is unfair because their geometry differs in that it's an iris-like shutter. However, I'd argue that geometry is less useful: almost every beam diameter in the actual experiment is about 1 mm or less, and the few that are bigger are easily accommodated with bigger shutter arms.




Figure 3.16. Our lab’s laser shutters (top left). We use simple stepper motors with some clamped-on anodized aluminum sheet metal as a beam block; sorbothane padding mounted on sheet metal base mitigates vibrations. The shutters are controlled by stepper motor drivers on a custom PCB I made (top right), which attaches to a Beaglebone microcontroller. The PCB mounts onto a BNC-breakout front-panel (bottom). The CAT5 cables carry the driver current to the motors around the lab. Our ADwin timing system communicates with the microcontroller through the UART serial protocol.

shutters.

UART design and software

We need to be able to tell the A4988 when to turn on the shutters, when to put them to sleep, and when to switch direction. But our timing system, the ADwin, communicates only via programmable digital (and analog) channels. In order to interface the ADwin with a microcontroller, such as a BeagleBone, we utilize the UART (universal asynchronous receiver/transmitter) protocol³⁴. UART is a serial protocol which can be used to send messages digitally. It is similar to SPI and I2c, except that it requires only Rx (receive) and Tx (transmit) lines. If you only need your device to receive, you need

³⁴While we could, in principle, dedicate a sleep and direction line for each of the 12 shutters, this would result in the use of 2×12 digital channels; we are limited to 48 total digital channels, so this is not an efficient approach.

only Rx and GND - no clock is needed, unlike SPI.

It works because both sender and receiver agree beforehand on a clock speed at which the data line will be sampled. The receiver samples the line at this pre-agreed rate, and when it finds digital high (say, 3.3 V), it interprets that as a 1; when it finds digital low (typically $< 0.2V$), it interprets a 0³⁵. What we do is dedicate a single digital IO channel from the ADWin to send commands to the BeagleBone in the form of a binary string, which nearly-simultaneously configures the sleep and direction states of all 12 attached shutters.

In order to optimize this for speed and reliability, I coded the microcontroller server in C, which interfaces directly with the Beaglebone's Linux kernel. The code is listed in Appendix F.2. Although it looks complicated, the program essentially just receives the UART message and maps the received 0's and 1's onto the output GPIO pins that correspond to each shutter's sleep and direction bits.

3.5 Magnetic field control

Atom manipulation requires careful control of the magnetic field, both to control the energy splittings and to set a well-defined spin quantization axis. In order to run the 461 nm bMOT, field gradients of approximately 60 G/cm must be generated, requiring on the order of 50-100 A from an anti-Helmholz coil. On the other hand, to zero the Earth's magnetic field requires only a few amps in a Helmholtz coil. We have two different coils for these two jobs.

3.5.1 MOT coils (Bitter coils)

We implement a Bitter-coil type electromagnet design, heavily influenced by Cheng Chin's design in [107], making a few modifications and improvements.

³⁵Aside from the obvious advantage of only needing two wires, it is fast - up to 3.7Mbits/s - and better still, the BBB can be controlled via the UART port. It cannot easily be slaved in SPI, and not at all in i2c.

Bitter coil design

As gasket material, we use 15 mil thick, high-temperature silicone rubber sheets of hardness durometer 20 (McMaster-Carr 86435K43, actually Marian HT6220.015-6x6). These were laser-cut for us by a local company (RMS Laser), after our campus machine shop failed at producing reasonable cuts with a water-jet process. The gasket material shrinks when the plastic liner is removed, which can lead to shorts—so the gaskets all had to be oversized (in all dimensions) by about 2.0%.

The gasket material can relax transversely over time. After construction, I tightened all bolts with 1.0 N-m of torque, using a torque wrench³⁶ in a star-pattern, similar to how vacuum windows are installed. When I first measured the resistance of the completed coils, they were about 22 m Ω each, as measured using a 4-wire technique on a 7.5 digit multimeter (Keithly DMM7510). This resistance doubled after a month! I re-tightened to 1.0 N-m and this seemed to fix the issue.

Our initial version, pictured in Fig. 3.17(a, right), used a brass water distribution block. This was too complicated and unwieldy, so after it leaked when I left it on all night without water cooling (oops!), we decided to make a 3D printed water block using the VeroWhite material. The print was manufactured by Fictiv (fictiv.com). This block uses internal channels to route the water. When we received it from Fictiv, a lot of printer material was blocking the internal channels; this was inevitable, due to the layer-by-layer nature of the 3D printing process. The channels had to be cleared by hand, which was tedious but manageable. Zip-ties were the tool of choice: they were at once strong enough clear the blockage, and flexible enough to worm into the internal channels. I relied on the chiller's water pressure to clear the last bits. In the final version of our coils, the connector (McMaster-Carr 5508K171) is epoxied (a two-part boat epoxy from Marine) into a fitted slot on the water block. The VeroWhite has held up well to water pressure. In testing,

³⁶Take some care here; the failure torque, at which point the 8-32 rods will snap, is about 2.5 N-m.




Figure 3.17. Construction (and evolution) of our Bitter coils. (a) Each of the 33 copper turns must be very carefully laid out, with the gasket material neatly flattened. Each subsequent layer is rotated clockwise by $1/8$ of a turn; see inset for an exploded view. Water flows along the rods, and along the gasket openings. Our initial (impractical) water cooling design is shown at right, replaced in favor of a 3D printed water block. (b) Coil with a printed VeroWhite water distribution block. This version eventually leaked at the connector. (c) Render of the final version—running in the experiment for 3 years without issue—using epoxied connectors. The cutaway shows some internal water channels.

it has gone up to 30 PSI without leaking. In daily use, we have both coils connected in parallel, and supply ≈ 17 PSI, which gets us about 0.7-0.8 L/s flow total, which is sufficient for our needs.

Coil control

The current in the coils must be controlled to a precision of about 10 mA (out of 50 A). We were inspired by some coil control circuitry from Alan Stummer at University of Toronto (physics.utoronto.ca/~astummer) called ‘Mag-o-Matic’; the circuit design

we use here is in Fig. 3.18. Our beefy power supply is a rated at 15V/400 A (Delta Elektronika SM15-400), and is operated in constant-voltage mode at 6.60 V. We use two parallel insulated-gate bipolar transistors (IGBTs), keeping their gates in the linear regime, essentially turning them into variable resistors. We use a Hall effect sensor (LEM IT-200S) to replicate the coil current, which avoids dissipating energy with a “sense” resistor. We convert this mirror current (ratio 1:1000) to voltage with a precision, temperature-stable resistor (200 Ohms, $\pm 0.1\%$, 25 ppm/C, 0.4 W)³⁷. Similar to the ODT circuit (see Fig. 3.13), we subtract this transduced voltage from the analog waveform generated by the ADwin. The RedPitaya’s job is to keep that difference at 0 V at all times, operating a PI control loop³⁸ with feedback sent to the IGBT gate. The IGBT’s themselves (IXYS IXYN50N170CV1) are rated for 1700 V / 120 A continuous-duty operation, and sit on large heatsinks cooled passively by some DC fans. An additional (and necessary) feature of our circuit is the ability to switch from anti-Helmholtz (MOT configuration) to Helmholtz, using some solid-state relays (Crydom DC100D100) in an H-bridge configuration.

3.5.2 Shim coils

To offset the Earth’s magnetic and move the red MOT around to co-locate it with the dipole trap, we use ‘shim coils,’ wraps of wire satisfying (very-approximately) the Helmholtz condition. Initially, we had some nice plastic mounts machined for them connected rigidly with cross bars, using square wire that fit pretty well (MWS Microsquare M151241). However, when our first-try Bitter coil leaked and we had to remove them, we also had to remove the shim coils—undoing the painful, week-long process of wrapping the chamber with the square wire. We eventually did a quick-and-dirty wrap job with lots

³⁷From this, we can work out the best achievable regulation: the LEM peak-to-peak noise voltage is 4 mV. That’s $4\text{mV}/200\Omega \times 1000 = 20$ mA on the primary, which, for a current-to-field conversion of 3 G/A, works out to 60 mG.

³⁸The output voltage of the RP is scaled with an analog circuit, which maps the (-1,1) V output to (5,7) V.

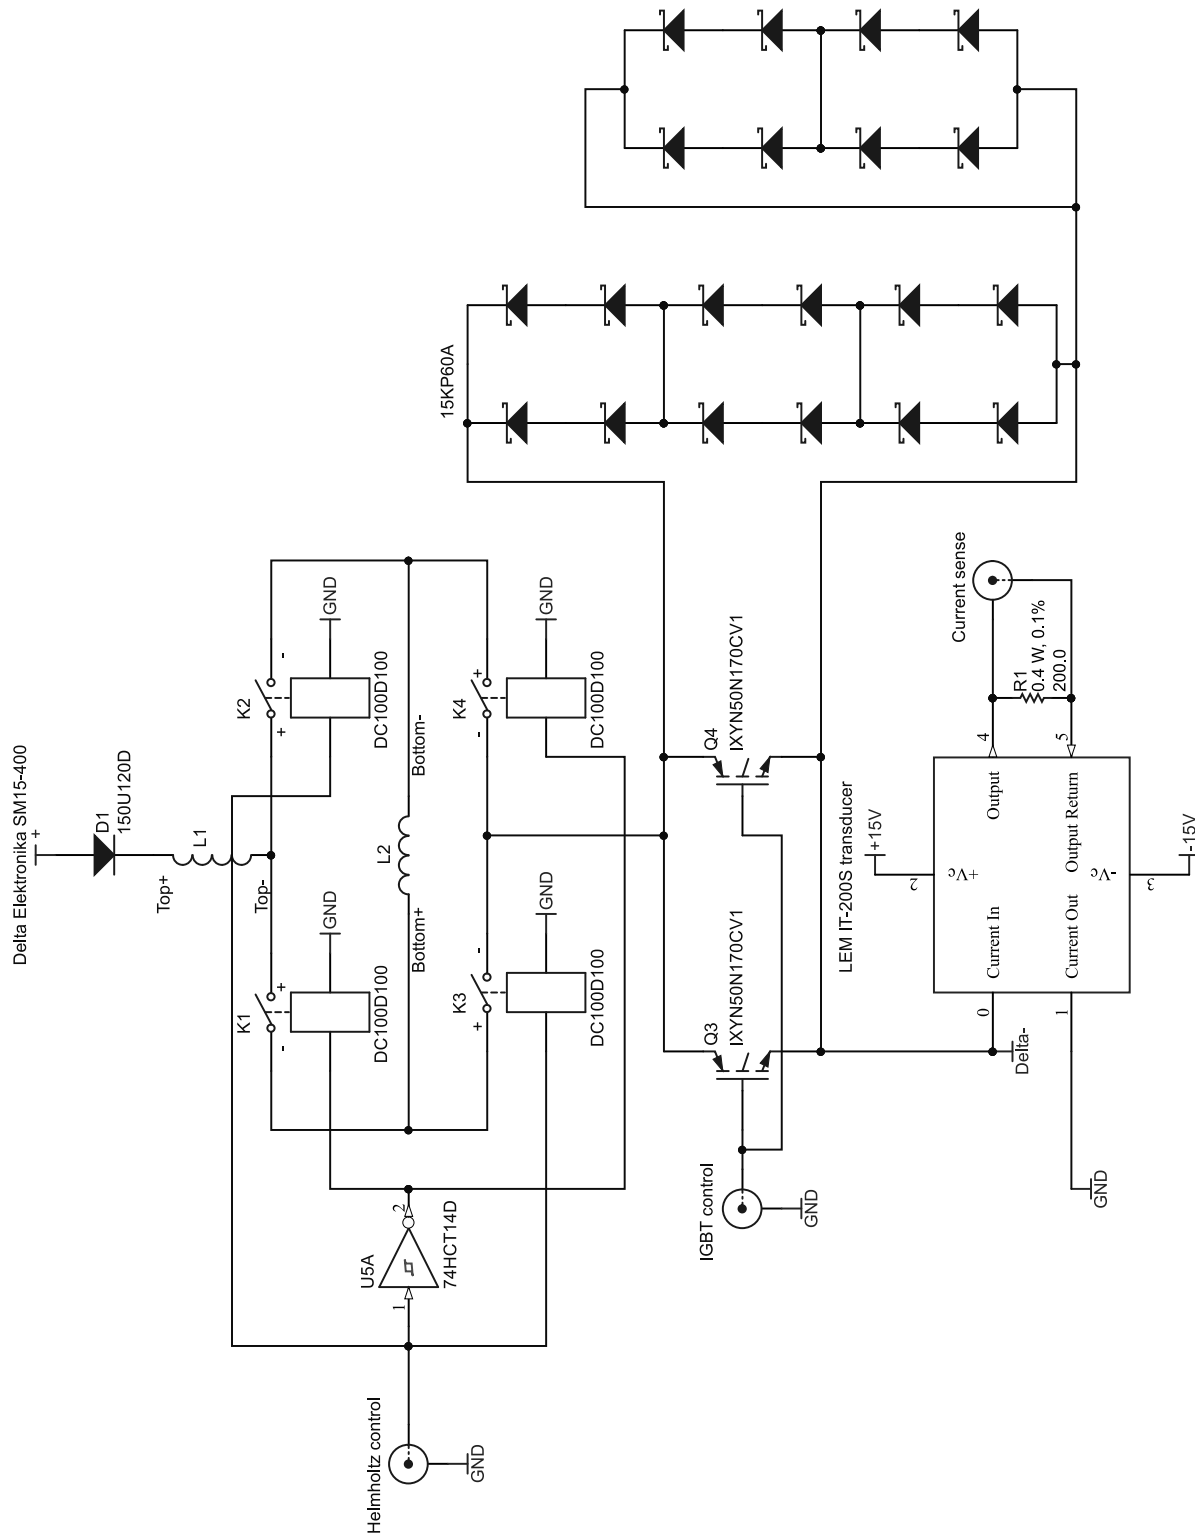


Figure 3.18. Bitter coil current control circuit.

of tape, but I haven't had any problems with them so far³⁹.

The current through these coils is controlled by linear power supplies (GW Instek GPS-1830D) operating in constant-current mode, slaved to an external analog voltage (from the ADWin) in order to allow changing the bias fields in-sequence. The measured supply noise is pretty good; $1\mu A$ at 4 Hz, and about $100\mu A$ at frequencies above 270 Hz. Given the approximately 1 mG/mA conversion ratio for the three coils, the ultimate noise limit is about 0.2 mG.

To zero out the Earth's magnetic field at the atoms, I used narrow-line imaging following JILA's example [96]. After 4 s of evaporation⁴⁰, I took absorption images (with pulse lengths of 50 μs) at 2 ms TOF at various shim coil currents⁴¹, getting the magnetic field from a power-broadened Lorentzian fit to the spacings of the m_F states. The magnetic field component B_i is then fit to the functional form $B_i = \sqrt{B_i^2 + (\alpha(I - I_0))^2}$, where α is the coil's current-to-field ratio and I_0 is the current which minimizes the field along that axis. The results are plotted in Fig. 3.19; we use these currents for all experiments in this work.

3.6 RF generation

Many frequencies are required to address the atoms. Lasers output only a single frequency; and in fact, we go to great lengths to make them do so in a very narrow band! To generate any needed nearby frequencies, we use acousto-optic modulators (AOMs), which take an RF signal and add (subtract) its energy to (from) the light passing through it. This section describes the generation and amplification of these RF signals.

³⁹But I wouldn't run more than a few amps through them. A quick test of 20 A for 1 s showed the wire strongly contracting, making a jerking movement. The wire also heated up at an alarming rate

⁴⁰The power supplies provide positive voltage only. In order to reverse the coil current direction, I had to build a (temporary, and now-removed) H-bridge, similar to the one built for the MOT coil control in Fig. 3.18.

⁴¹The coil currents during the red MOT stage were always the same, because the ODT/MOT overlap is crucial and changing the shim coils changes that overlap. Only after loading into the ODT did I ramp the coil currents.




Figure 3.19. Magnetic field calibration of the shim coils using narrow-line spectroscopy on the $^1S_0 \rightarrow ^3P_1(F = 9/2)$ transition. These measurements were performed after 4 s of evaporation. The fitted $x, y,$ and z shim coil coefficients are, respectively, $0.877 \pm 0.001, 0.446 \pm 0.002, 1.377 \pm 0.003$ G/A. The background magnetic fields at the zero-currents are $0.309 \pm 5, 0.329 \pm 4, 0.1 \pm 0.1$ G, with the large uncertainty in the z -shim calibration arising from lack of m_F -state resolution at those fields.

3.6.1 RF sources

With the exception of the flux lattice experiment, which uses an arbitrary waveform generator, all of our frequencies are generated by direct digital synthesizers (DDS's), which produce sine waves with very low phase noise. We use the Analog Devices AD9914, AD9915, and AD9959. The former two are single-channel, and are used for frequencies > 200 MHz; the 4-channel AD9959 produces all needed lower-frequency waveforms.

All DDS's require a low-noise clock. For this we use relatively low-cost signal sources from RF Consultant (rf-consultant.com), specifically the TPI Synthesizer based on the Analog Devices ADF4351. These are (in turn) phase-locked to a 10 MHz Rb reference (SRS FS725). The advantage of the Rb reference is the low drift; all oscillators, left to themselves, drift over time, which is a problem in atom experiments where we care a lot about the absolute frequency. The reference utilizes Rb for its source generation, providing a very stable locking reference for all signal sources in our lab.




Figure 3.20. One of two DDS-controlling Beaglebones installed in the experiment, with a custom PCB mounted on top.

DDS control

I control all DDS's with two BeagleBones, which have custom PCBs mounted to them; see Fig. 3.20, and the schematics in Appendix C. This is a 4-layer board made using “best-practice” high speed design techniques [108], and it can control up to 6 DDS boards. It buffers the Beaglebone output pins, protecting them from over-current conditions and allows the driving of long-ish transmission lines. With a change of a jumper setting, it can be used to drive the (slow) SPI bus, or it can be used for the programmable real-time unit for faster control projects. In daily practice, I only used it for the SPI bus.

The board has a couple of digital input trigger BNCs, to which you may program it to react. This is extremely useful when doing automated frequency sweeps—so long as you keep in mind the ‘reaction time’ of the board and associated SPI delay is about 300 ms from the time you issue the DDS command.

3.6.2 Amplifiers

The RF sources in the previous section must be amplified to drive AOMs. The source amplitude is typically -10 dBm to -6 dBm; to be useful, AOMs require on the order of +30 dBm, about 1 W. One of my first projects in the lab was modifying an RF amplifier design from Florian Schreck’s lab⁴² to drive all of our AOMs. It is advantageous and economical to build our own RF amps, because the alternative approach is expensive, bulky, and less-featured Mini-Circuits blocks. Our RF amps ended up costing about \$500 each, including part and assembly fees.

I made some simple design changes (updating a few parts), and we had the PCB’s manufactured and soldered by Advanced Circuits, which did a run of 25 amplifiers for us. These power all of our AOMs; in 2019, lab-mate Khang Ton made some really cool tweaks and improvements, including an on-board RF switch, nanosecond attenuator, and general speed improvements to the analog control path.

3.7 Experiment control

The production of quantum gases requires exposing the atoms to different fields—both optical and magnetic—at very precise times, in a repeatable way. This involves the movements of shutters, the ramping up and down of optical and magnetic intensities, and myriad digital signals that switch all these things on and off. All of this must be done with μs precision, and better-than- μs repeatability; essentially, we require a very precise conductor for a very picky electronic orchestra.

3.7.1 ADwin hardware

In our lab, the timing system is the ADwin real-time unit (ADwin Pro II, running the T11 processor). It is “real-time” in the sense that it will, with nanosecond precision

⁴²As of 2022, their schematics are still available on strontiumbec.com.

and accuracy, output a signal when you tell it to. This is in contrast to almost all other computers that exist, which run on regular CPUs; to synch a program to real-world events, you can kind-of use software hacks like python’s `delay()` functions, but you are never guaranteed a response within a certain time interval. If the CPU is otherwise busy, it will triage your request after it e.g. renders the cute cat GIF⁴³ you’re looking at while waiting for the sequence to finish. It is vital in AMO experiments that things happen exactly when you want them to, with precision much faster than the characteristic atomic motion; in most cases, that is milliseconds, so the timing system had better be at least $10^2\times$ faster. This level of accuracy is not possible without dedicated real-time processors.

The ADwin version we bought has 64 digital and 24 analog (± 10 V) outputs, any of which can be toggled at any time. The smallest-allowable pulse time for any channel is $4\mu\text{s}$, sufficient for most purposes⁴⁴. Our ADwin does not have any inputs, nor can it directly control other devices, necessitating some clever protocols for driving e.g. shutters (see Section F.2).

3.7.2 Qcontrol

While the ADwin does have its own programming language, it is a pain to work with; thankfully, we interface with it using Qcontrol, a python-based program developed at the Max Planck Institute for Quantum Optics. We use a fork of the version from 2013, which I have (minimally) maintained since then. The version we use requires python2, an increasingly onerous requirement since python2 reached end-of-life in 2020; all major Linux distros have removed it. Porting the software is frankly not worth the effort, so I have it frozen in a virtual environment running an old version of Fedora.

Qcontrol is extremely easy to use. Its documentation is sufficient to explain its detailed use; here I will give only a brief overview. Each channel on the ADwin, analog or

⁴³Pronounced “GIF,” as in “good”.

⁴⁴Although this is a bit of a pain when trying to do, e.g., Kapitza-Dirac diffraction or electron shelving; in either case, closer to $1\mu\text{s}$ would be nice.

digital, is associated with a named object; for example, you might choose to name one of its digital outputs “SwODT1”, referring to an RF switch controlling one of the ODT beams. To turn this on at in-sequence time 500.000 ms, you simply write `SwODT1(500*ms,1)`; to turn it off 5 ms later, call `SwODT1(505*ms,0)`. In addition to the real-time capabilities, you also enjoy the other features of the language, such as numpy array handling, function calls, and classes. That is it in a nutshell! The entire experiment consists of many such calls, telling what channels to turn on when, for how long, and with what amplitude.

Qcontrol bugs and caveats

There are some bugs that are important to be aware of. First, there is a memory leak that becomes apparent after a few hundred runs during iteration. I haven’t tracked this down, but if you are not careful, somewhere around your 400th iteration run, the OS will kill the timing system and you may not be aware! Secondly, there can be issues with priorities. Even if two channels are not in direct conflict, sometimes the program will obliquely alert you of ‘CONFLICTS!’— but will still allow you to upload the script to the ADwin. And then, in the actual experiment, some channel calls will get silently dropped! This can be maddening to diagnose, and typically occurs with calls to shutters in sub-sequences. It is easily worked-around by including keyword `priority=m`, where m is any integer less than -10 (see the docs for more info). This essentially ups the priority and allows the call to proceed.

Finally, the GUI, although excellent, sometimes experiences rendering errors. For reasons I haven’t been able to pin down, sometimes when you change the timing, the drawing that represents a channel’s activity prior to the update doesn’t get deleted, and the new timing gets drawn over it. This can get very confusing, and the only solution is restarting the GUI. Also, if you weren’t using a channel at all, and then decide to use it, the GUI dictionaries get confused, resulting in uncaught exceptions when you try to plot the new channel. I try to keep all channels I might use activated somewhere, at least

setting them to 0V at some point in the Initialization.pys script, to avoid these kinds of errors.

The only major change I made to Qcontrol, aside from maintenance operations that correct syntax-warnings-cum-errors over the years, was modifying the GUI to show the channel names, and not the internal channel number. In my opinion, it is objectively more useful to display the channel name, instead of e.g. “DIO2.DIO29.”

3.7.3 Digital buffers

The ADwin is very expensive, costing upwards of \$30k USD. It would be a shame if any of its channels were damaged through any of the day-to-day accidents that occur in a busy lab, such as accidental shorts to ground. It would be furthermore beneficial to augment its current-driving capability, and provide an isolated ground in case of high-voltage accidents which could zap the entire unit. For these reasons, and inspired by Alan Stummer’s digital buffers (physics.utoronto.ca/~astummer), I designed a 4-layer digital buffer breakout board to which the digital outputs of the ADwin connect. We use the (now-obsolete) quad-channel ADUM1400CRWZ buffers, using proper high-speed design techniques [108] which completely isolate the ADwin ground from the outputs. The buffered signals go into the line drivers SN64BCT25244, which can drive loads up to 80 mA. The full schematics are given in Appendix D.

3.7.4 Laser intensity control

All laser intensity control in the experiment is handled by the credit-card sized, low-cost, FPGA-based RedPitaya.

Hardware

The RedPitaya has 4 analog channels, 2 inputs and 2 outputs—and since it is an FPGA, can be programmed to perform fast modulation of the outputs based on the analog

inputs. Although the rated bandwidth is 50 MHz, it has an internal delay of about 130 ns from in-to-out, which limits the achievable servo bandwidth to about 500 kHz (see [109] for a good characterization). But this is fast enough for our purposes.

Software

The software that comes with the Red Pitaya is terrible (don't even try it!) and has lately become micro-monetized. We instead use the excellent, open-source PyRPL (pyrpl.readthedocs.io), developed at the quantum optics lab Laboratoire Kastler Brossel in Paris, France. This software has been modified in our lab by nanotraps-experimenter Grady Kestler in order to incorporate an absolutely essential feature: it allows the PID controller to accept an analog waveform as a setpoint. This affords enormous flexibility to the experiment, because the ADwin timing system has an analog arbitrary waveform generator. We can thus make the laser intensity follow any profile we want.

Caveats

There are a few important practical notes regarding their use. First, I have noticed serious switching noise induced on the analog inputs, apparently caused by the cheap power supplies shipped with the units. I therefore always use linear power supplies with them (5V, capable of at least 2 A). Secondly, ground-loops are especially problematic and must be avoided, since 60 Hz noise on either input will get written onto the laser—a fatal situation for an ODT, which has end-stage trap frequencies on the order of 40-100 Hz. To ameliorate ground issues, we plug all RedPitaya's to the lab's orange outlets, which is the 'sensitive' lab ground, separate from the main ground of the lab which can have heavy equipment attached to it. This isn't always enough: even with shielded cables, sometimes one must simply move them away from noise sources. I had to physically move and even shape my ODT-controlling cable, in such a way as to empirically minimize coupled 60 Hz

noise! ⁴⁵

⁴⁵I did this as you might imagine: as I watched the FFT of the input, I moved the cable around to minimize the peaks at 60 Hz and multiples thereof, then taped it in place. Not a perfect solution, but time is always limited in real experiments.

Chapter 4

Achieving quantum degeneracy

Having described the necessary equipment and laser setups, this chapter will explain the process for cooling atoms to quantum degeneracy. This is the initial condition of the SM-lattice experiment.

4.1 From 700 K to 2 μK

Since the vapor pressure of strontium is negligible at room temperature, we have to first heat it up to about 750 K (900 F). These hot atoms propagate through thin micro-tubules which create a “collimated” hot atom beam. We then impose several MOT cooling stages operating on the 461 nm/689 nm transitions. The blue MOT takes us down to a few mK, and the red MOT down to a few μK , reaching the temperatures of the associated Doppler recoil limits (as discussed in Section 1.3).

A list of nominal MOT intensities is given in Table 4.1.

4.1.1 Blue MOTs: 700 K to mK

After the 700 K atoms make it through the AOsense micro-tubules, the atoms are hit with 2D MOT and Zeeman slower light. From measurements provided by Tom Loftus, the atoms are radially slowed to 0.6 m/s, and longitudinally to 40 m/s (90 MPH!). This is what enters our chamber.

Cooling the boson or fermion in the 3D blue MOT is relatively straightforward,

since the fermion’s more-complicated 1P_1 hyperfine structure is not resolved due to the broad transition linewidth of $\Gamma/2\pi = 30$ MHz. The procedure for all species is very simple: after setting the appropriate isotopic frequency shifts, we apply a 47 G/cm magnetic field gradient and illuminate with the 3D MOT beams for some load time t_{load} . During the loading period, atoms have a 1:50000 chance to decay to the metastable state 3P_2 (see Fig. 1.1), which acts as a reservoir for atom accumulation. This is strictly superior than using a blue MOT by itself, which has a very limited lifetime owing to light-assisted collisional losses [110].

At 82% natural abundance, the boson ^{88}Sr is the fastest to load into the reservoir; the fermion, at 7.0%, needs commensurately longer time to achieve similar atom counts. Depending on the oven’s health, I have loaded the boson with $t_{\text{load}} = 1$ s at 460 C, but that has lately gone up to 3 seconds. For the fermion, as of March 2022, I use $t_{\text{load}} = 7$ s at 480 C.

Repumping

To get the atoms out of the metastable state so that we can transfer them to the red MOT, we use a *repumper*. This is a 481 nm laser (see Section 3.4.4 for details), which is mechanically shuttered. At the end of the loading time t_{load} , the 481 is flashed onto the atoms by simply opening the shutter. It is not important when this beam turns off (if ever), since it is low power and far from any ground-state resonances; however, the relative timing between when the Bitter coil current falls (in order to load the red MOT) and the repumper opens is critical, since the reservoir atoms are magnetically-trapped and free-fall as soon as the field drops. If they’re not immediately recaptured by the red MOT, they are lost. Considering this, the shutters have been remarkably repeatable, although I suspect more-optimal results could be obtained using an AOM as an optical shutter, which would allow much better than millisecond timing precision. However, I haven’t found it necessary.

As described in Section 3.4.4, when loading the boson, the repumper wavelength is set statically at 481.3223 nm; for the fermion, the laser is scanned a few GHz to cover the hyperfine structure. This scanning is necessary, and more than triples the resulting red MOT atom count.

4.1.2 Red MOT: mK to μ K

We transfer atoms from the reservoir to the red MOT (rMOT) by simply overlapping the rMOT beams at the reservoir location. The beams operate on the intercombination line $^1S_0(F = 9/2) \rightarrow ^3P_1(F = 11/2)$, which has a narrow linewidth of $\Gamma_{689}/2\pi = 7.4$ kHz. Since the reservoir-trapped atoms are at a few mK, the resulting Doppler-broadening is on the order of a few MHz, reducing the MOT force $F \sim \Gamma_{689}/\Delta \approx 0$ for most velocity classes, where Δ is the Doppler detuning from resonance. People typically work around this problem by artificially broadening the laser linewidth during a ‘capture’ stage of the rMOT [49, 104], driving the laser in a triangle wave in frequency space that simultaneously addresses numerous velocity classes. More recently, another approach, based on a sawtooth wave profile (SWAP MOT), has emerged [50, 111], and this is what we use.

The fermion MOT timing diagram is shown in Fig. 4.1, showing the action of the coils, blue MOT shutter, repumper, rMOT and ODT beams¹. The gray dashed line on the repumper shutter indicates the approximate time at which the 481 light is open to the atoms, within a 1 ms window. The inset shows that the rMOT operates on a sawtooth-wave profile, with the indicated frequency sweep reaching -1.5 MHz (120 kHz) below (above) resonance. The power of the MOT beam is decreased gradually, which lowers the temperature as time proceeds.

The SWAP MOT ends abruptly after about 200 ms from the repump, and we continue cooling with single-frequency light in order to get the lowest temperatures. The stirring beam turns on in the single-frequency stage, which is necessary only for the fermion.

¹The boson timing is very similar.

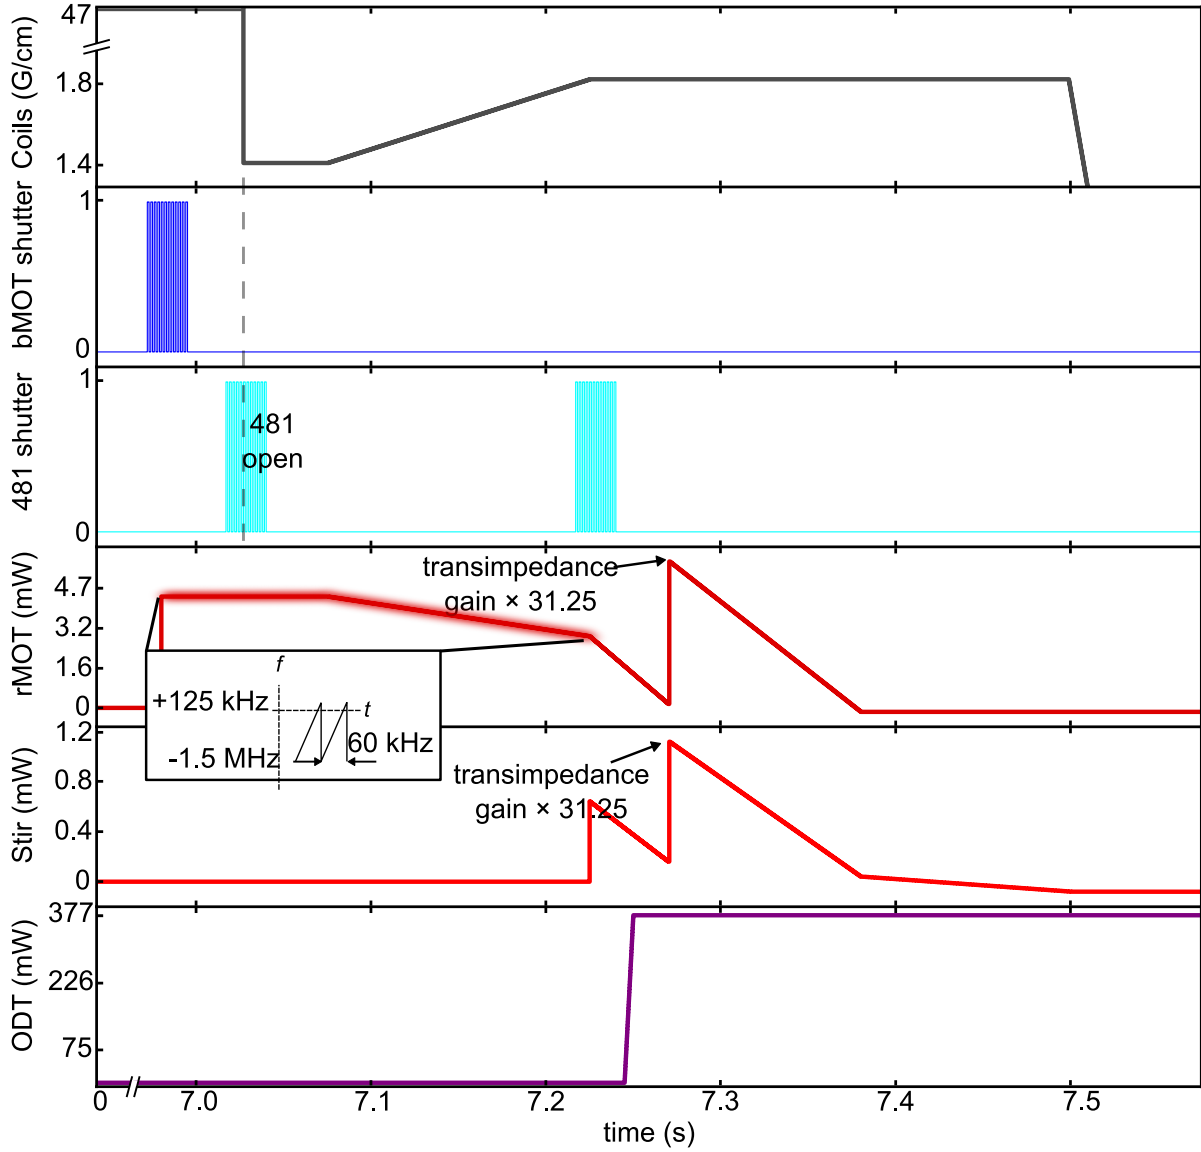


Figure 4.1. Timing diagram of the red and blue MOTs for ^{87}Sr . rMOT powers indicate the total power seen by the atoms; ODT power is per-beam. Only the rMOT beams are frequency-broadened in a sawtooth pattern, as indicated by the blurred line; the scan width is indicated.

This beam, operating on the $^1S_0(F = 9/2) \rightarrow ^3P_1(F = 9/2)$ transition, “stirs” the m_F states, randomizing the population [112]. This is needed because MOT operation on the $F=11/2$ transition alone is unstable, owing to the large spread of g_F factors among the m_F states. This would cause many atoms to be anti-trapped, reducing the achievable atom count by a factor of 20 or so.

The switchable-gain photodiode (see Section 3.4.2) changes gain midway into the single-frequency stage, allowing us to reduce the rMOT power to tens of μW and achieve temperatures close to $2 \mu\text{K}^2$. The ODT turns on during the single-frequency stage, with the timing rationale being to limit its (harmful) ac Stark shift effects on the rMOT; but in practice, it doesn’t seem to matter when it turns on.

The in-situ MOT cooling process for the fermions is pictured in Fig. 4.2, with a 461 nm absorption images taken orthogonal to gravity. This data was taken at an oven temperature of 460 C in October 2019, when the oven was relatively healthy, so the loading time here was only 3 s (now 7 s). We can see the MOT density increase as the atoms cool, with the MOT action forcing the atoms into a smaller volume. During the single-frequency stage, the atoms collapse into more of a pancake shape, resembling the bosonic case. Atoms crossing the dipole trap get trapped by it and are removed from the cooling process, owing to its large ($> 100 \text{ kHz}$) differential ac Stark shift.

MOT alignment

It is sometimes necessary to adjust the rMOT, either due to accidental misalignment, beam drift, or changes to the oven characteristics. Unfortunately, the fermions are extremely sensitive to beam alignment; a blue MOT will be present under just about any circumstances, but the transfer from blue to red is nontrivial since the two are on different

²ODT-trapped atoms are typically hotter by a few μK as compared with atoms in the rMOT alone. This is because atoms entering the ODT are abruptly light-shifted out of the cooling process. One might think to just turn on the ODT later—but that idea is thwarted by the finite hold time in the rMOT: reaching the lowest temperatures requires lowering the MOT powers (and thus MOT force) below that which could hold the atoms against gravity. There are clever ideas which get around this limitation, specifically by using a transparency beam; see [56] for further details.




Figure 4.2. In-situ time evolution of the red MOT cooling process, starting from the repumped atoms at few-mK and ending at few- μ K. The units are optical density; note the scale changes as time evolves. The dipole trap presents as a horizontal line of dense atoms. I omit the stirring beam diagram here, but note that it is the same as in Fig. 4.1.

beam paths. Depending on the degree of misalignment, one might find the rMOT not present at all. What follows are some general problem-solving steps I have developed over the years, ordered in severity. The following guidelines are mostly for the fermion, but it's applicable to the boson also (but the boson should be pretty robust).

1. “rMOT count down $< 10\%$.” Day-to-day drifts and decreases in laser power are frequently at fault; check that the bMOT and rMOT powers, as measured at the windows, are consistent with established values (see Table 4.1 for one example). Next, check the 689 laser lock: if the diode is in a “bad” place, the cavity can't lock it as well, so the linewidth will be broad. Finally, keep a log of the OD over time; if you notice a trend that can't be explained by power drifts, it could very well be the oven health declining.

2. “rMOT down > 50% or more.” Are any beams obstructed? Is the stirring beam on? It could be severely misaligned. Try iteratively walking each rMOT arm, starting from the vertical. Walk it like you would a laser into a fiber: run the sequence in a loop (removing the ODT code to speed things up) and make tweaks to the steering mirrors, watching the resulting OD at -140 ms. Do this for each arm.
3. “rMOT is gone.” Are you sure the frequencies are correct, and lasers locked? Surely you checked this already. Ok, are you sure it’s gone? Check “back in time”: check some of the earlier times, as shown in Fig. 4.2. If you do see atoms at some time and they disappear before the cooling sequence concludes, try sitting at the latest of these visible times and perform the walking as described above.
4. “No, the rMOT is really *gone!* Halp.” Worse comes to worst, if you need a starting point you can always align with ^{88}Sr . This isotope is fairly robust. Retro-reflect each MOT beam as best you can, and it should easily produce an rMOT. Don’t expect to then go to the fermion and have it be perfect, however; it’s just a starting point. But hopefully, at some time shortly after repumping, a workable signal will appear.

4.1.3 MOT benchmark

The laser intensities necessary to run the experiment are given in Table 4.1. These numbers are not intended to be rigidly matched, and certainly evolve over time depending on laser and oven health, but serve as a rough guide to future experimenters. With these intensities—and a relatively healthy oven—in mid 2021, I loaded the fermion for 7 s, at an oven temperature of 470C/19.1 V, producing a red MOT OD 2.6 at -140 ms sequence-time; see Fig. 4.1 for the timing. The resulting peak OD after 12.5 s evaporation, at 20 ms TOF, was 0.55.

Table 4.1. Peak MOT beam intensities measured at the chamber windows. The resulting red MOT OD was 2.6; see text for details. Refer to Fig. 3.10 for the beam locations: bMOT 1 is nearest to OFL k_1 ; bMOT 2 is nearest to OFL k_3 . rMOT 1's retro is nearest the lift beam; rMOT 2 copropogates with OFL k_2 . Ditto with the stirring beams, which copropogate with the rMOT beams. 2D MOT 1 is closest to the oven. -z means the beam propagates against gravity.

Beam	Peak power (mW)	$1/e^2$ diameter (mm)
481 nm repumper	3	4(0.5)
2D MOT 1	3 ¹	30×5
2D MOT 2	2.5	30×5
Zeeman slower	60	6.5
bMOT 1	4.2	10(0.2)
bMOT 2	6.3	9.2(0.2)
bMOT, -z	2.9	6(0.1)
rMOT 1	0.79	11(1)
rMOT 2	0.79	9.1(5)
rMOT,-z	1.08	$7(0.5) \times 8(0.2)$
rStir 1	0.25	10(0.2)
rStir 2	0.26	10(0.2)
rStir, -z	0.48	$7(0.5) \times 8(0.2)$

¹ The 2D MOT powers are measured with a power sensor with a 10 mm diameter, so this indicates peak power only, not total power.

4.2 Evaporative cooling: from μK to nK

Making quantum-degenerate gases requires evaporative cooling. RF evaporation is one older method of doing so, which uses an RF antenna to push hotter atoms into anti-trapped m_F states; but most experiments today use optical dipole traps (details on the laser setup can be found in Section 3.4.5), which is required for Sr due to its weak magnetic moment. Far off-resonant dipole traps (FORTs) [102] provide a strong scalar light shift on the atomic ground states, holding them against gravity with almost no spontaneous emission. The idea is to put the MOT-cooled atoms into the dipole trap, and then gradually lower the trap depth. Hotter atoms escape, taking with them their kinetic energy; the remaining atoms collide with each other and rethermalize, becoming cooler as time proceeds. It's just like a hot coffee set out to cool.

4.2.1 Scattering considerations

The success of evaporative cooling depends on favorable scattering lengths, which in turn depend on the atom, the specific isotope, and even the involved quantum state. For example, evaporation of the boson ^{88}Sr with itself does not work, because its self-scattering length of $-1a_0$ is too low. That scattering length is much larger in the magnetically-trappable 3P_2 state, opening up the intriguing idea of magnetic evaporation; but large two-body collisional losses exceed the trap depth [55]. In practice, a BEC of ^{88}Sr can be achieved only with a mix of another species present to aid thermalization, in a process called ‘sympathetic cooling.’

Only two Sr isotopes can be cooled to degeneracy by themselves: the boson ^{84}Sr , and the fermion ^{87}Sr . In the former case, its scattering length is neither too large nor too small at $124a_0$; a BEC is easily achieved for any reasonable evaporation sequence. For the fermion, evaporation proceeds with an important caveat: identical spin states do not




Figure 4.3. The evaporation trajectory used in experiments, the functional form of which is given. The typical endpoint is highlighted. In the all-spin, $N = 10$ case, this reaches a degeneracy of $T/T_F = 0.2$. The waveform for ODT2 is identical, with a slightly lower peak power of 348 mW. Omitted here is the spin polarization ODT ramp, which is optional and discussed later.

scatter with each other at low temperatures, so fermionic degeneracy can be achieved *only* with a mix of at least two spin states. At the end of the red MOT, we naturally have an even mix of $N = 10$ spin m_F states, with which routine degeneracies of $T/T_F = 0.2$ can be achieved. With two-spin evaporation, however, the achievable degeneracy worsens, rising to about 0.3-0.4 [55], although recent progress has been made on that front at JILA [56]. If a completely spin-polarized Fermi gas is required, one does the spin-polarization right after the red MOT, keeping at least two states; at the end of evaporation, we can take advantage of the narrow-line transitions to kick out the undesired state.

4.2.2 Evaporative cooling guidelines

A natural question to ask is, if evaporative cooling is required, why bother going through the other cooling stages, especially the red MOT? mK trap depths are achievable with enough laser power! While that is true enough, there are scaling laws governing the evaporative cooling process [51]. The nominal requirement for reaching the quantum

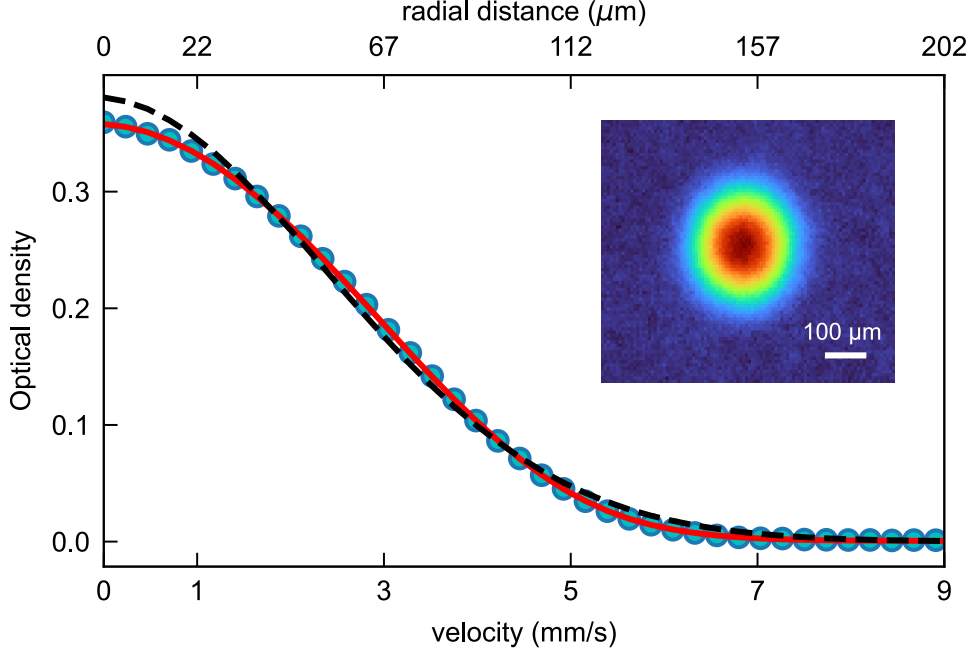


Figure 4.4. Azimuthally-averaged momentum profile of a DFG, averaged over 25 shots and taken at 22 ms TOF. Pauli blocking manifests as a small deviation from a naive Gaussian fit (black, dashed), which over-fits at the peak because the low-momentum states are mostly all occupied. The red line is a fit to the Thomas-Fermi profile, which matches the data well with degeneracy parameter $T/T_F = 0.24$.

regime is to reach a phase-space density of 1, where we recall that $\rho = N(h\nu)^3/(kT)^3$ for a harmonic trap, with ν the geometric mean trap frequency. With the dimensionless ratios $\eta = U/kT$ and $\eta' = \eta + (\eta - 5)/(\eta - 4)$, where U is the trap depth, the ratio of final to initial phase space density is:

$$\left(\frac{\rho}{\rho_i}\right) = \left(\frac{N_i}{N}\right)^{\eta'-4} \quad (4.1)$$

One important consequence of this relation is that if $\dot{N} = 0$, the phase space density remains constant, emphasizing the crucial role of atom loss in the cooling process. To put some numbers to this, for a typical $\eta = 10$ with $\rho_i = 0.003$, lowering the well depth by $84\times$ will reach $\rho = 1$. If we instead started $10^3\times$ hotter at a few mK, then $\rho_i \rightarrow 10^{-9/2}\rho_i$, assuming we can trap the same number of atoms and scale the trap depth as $U \rightarrow 10^3U$, with the trap frequency scaling $\nu \rightarrow 10^{3/2}\nu$. Then the well depth would need to scale down

224300 \times , which is eminently impractical. Thus, we use MOTs, which are by their nature not-conservative, carrying away energy and allowing us to reach the necessarily-large starting phase-space density.

4.2.3 Bose-Einstein condensation

Early in the lab’s construction (mid-2018), we benchmarked our system by cooling the boson ^{84}Sr to BEC³. We loaded for 20 s into the metastable reservoir, accumulating in the rMOT about 5 million atoms at $1.3\mu\text{K}$. Our starting dipole trap frequencies were 116, 107, 587 Hz, with the latter the vertical direction. The evaporation waveform was described by $P(t)/P_0 = \left(1 + \frac{t}{\tau}\right)^{-\beta}$ with $\beta = 1.5, \tau = 0.4$; the BEC transition occurred around $t = 2\tau$.

4.2.4 Degenerate Fermi gas

For the DFG, the evaporation waveform follows the functional form suggested in [47, 51] and is shown explicitly in Fig. 4.3. Since the scattering rate decreases with atom count, I slow the evaporation ramp-down rate significantly after 14 s in order to allow sufficient time for the atoms to thermalize. I’ve tweaked this waveform only a little over the years, but there is certainly room for improvement in light of [56].

The resulting unpolarized Fermi gas is shown in Fig. 4.4 (the fits are described in Section 1.4.2). Producing such a Fermi gas is now a daily part of our experiments. If the red MOT OD is above 2.0 at -140 ms (with respect to the end of the red MOT sequence), one can expect to achieve similar or better results, with the atom count in the DFG depending linearly on the red MOT density.

I note that the overlap of the crossed ODT beams is critical, since the waists in the vertical direction are around 20 μm . Slight misalignment due to tiny bumps to any elements on the path can spoil the evaporation ramp and cause severe atom loss—but

³The ODT (and entire optical setup) has undergone quite an evolution since then; indeed, we have since repurposed the ^{84}Sr rMOT optics.




Figure 4.5. Post-evaporation dipole trap frequency measurements along the directions of ODT1 ($2\pi \times 35$ Hz), ODT2 ($2\pi \times 52$ Hz), and along gravity ($2\pi \times 211$ Hz), respectively. The beam power of ODT1 (ODT2) was 12.5 mW (15.6 mW), corresponding to voltage setpoint 35 mV (45 mV). The $1/e$ decay timescales are 96 ms, 64 ms, and 14 ms, respectively.

fixing this usually only requires micro-adjustments to the vertical knobs on the last Polaris mirrors⁴. I also note that any 461 nm or IR stray light can spoil the DFG by heating the atoms, and any vibrations that are close to the trap frequencies (or multiples thereof) must be avoided and/or moved far away from the ODT mirrors.

Trap frequencies

Although one can measure the ex-situ ODT beam waists, the actual intensity experienced by the atoms depends on the alignment and focal lens positioning, which cannot be directly measured. Fortunately, we can easily get at these values—and the corresponding motional timescales—by inducing dipole oscillations in the ODT: push the atoms and, like a ball on a spring, watch them oscillate. The frequency of these oscillations is the trap frequency. Although the trap is nominally harmonic and thus conservative, we do observe damping on 10 ms timescales (depending on the push direction); this indicates the presence of trap anharmonicity, which is to be expected since the traps are only locally

⁴You’ll want to use an Allen wrench, not your fingers on the knobs; sweep back and forth in steps of much less than 1 degree, as larger changes are likely to overshoot the ‘sweet spot.’

harmonic.

To measure the trap frequency in the direction of gravity, I ramp the ODT beams completely off in 1 ms, before ramping them back on in $100\mu\text{s}$ ⁵. The beams are intensity-servo'd the entire time. This quench and rapid return generates a nice center-of-mass oscillation along the vertical direction. I let the atoms bounce for a varying time before letting them go, imaging them at 10 ms TOF, with spacing of 0.3 ms per shot. The result is the right-most image in Fig. 4.5; clearly, the vertical direction is the most tightly-confined.

For the longitudinal trap frequencies—along the direction of the ODT beams—I suddenly shift (by 500 kHz) one of the ODT AOMs; the atoms then oscillate along the direction of the unperturbed ODT beam⁶. I've verified that shifting a beam and then shifting it back to center gives about the same frequency as simply leaving it shifted. The data in Fig. 4.5, left and center, is for shifting the beam and leaving it there, and we observe trap frequencies along the ODT1, ODT2 directions of 35 and 52 Hz, respectively, at beam powers 12.5 mW and 15.6 mW.

⁵Diabatically switching the beams on and off also works, but the resulting oscillations are too wild, preventing a clear measurement of the center of mass motion.

⁶Since Gaussian beams confine more strongly in the radial direction, and much more weakly along their propagation direction.

Chapter 5

Spin-momentum lattices

Having built the instrument, it's time to do something with it! This chapter details the realization of an atomic spin-momentum lattice using the fermion, ^{87}Sr . This is a step along the way to the creation of optical flux lattices (OFLs) for the study of fractional quantum Hall (FQH) physics.

Part of what follows is the manuscript of our results, to be published in Physical Review Letters [113]. This work is Copyright 2022 by the American Physical Society.

5.1 SM lattice preparation

To prepare for the experiment as diagrammed in Fig. 2.3, several challenges had to be overcome. Firstly, the fermion has $N = 10$ degenerate Zeeman spin states, which populate evenly after evaporation. But the experiment calls for the use of 3 non-degenerate states all linked together, which must furthermore be spaced nonuniformly in energy. Procedurally, then, we need to accomplish the following:

1. Polarize the spins: we've got to push 10 spins into 1, in a manner which minimizes loss
2. Align the lasers onto the atoms
3. Lift the spin state degeneracy in a nonuniform way

4. Link the spins together with a coherent Raman process, and measure the strength of those links
5. Measure the spin populations.

We have already developed one important tool: an optical Stern-Gerlach (OSG) beam, which spatially separates the spins as demonstrated in Fig. 3.15. But since this destroys the momentum-space information, we need another way of imaging to satisfy point 5. We tackle each of these issues in turn.

5.1.1 Spin polarization

The first problem we face is the efficient transfer of population into the desired spin state, while depopulating the undesired ones. While one might initially think to use a coherent process—such as Raman transfer, or adiabatic rapid passage—it is fundamentally impossible to employ unitary dynamics to spin-polarize a sample, based on simple entropy considerations alone [114]: spin polarization reduces the system’s entropy, which cannot happen in a unitary (reversible) process. For instance, applying Raman pulses would simply swap two spin populations, not transfer them. Spontaneous emission *must* be involved, with the spontaneously-emitted photons carrying away the excess entropy.

Spin polarization of strontium is discussed by Stellmer [104] and Killian’s lab [115]. Essentially, it involves using the narrow-linewidth 689 nm transitions to selectively pump atoms out of some spin states and in to others. This optical pumping must occur during the early stages of evaporation to avoid significantly heating (and ruining) a DFG. The transition’s narrow linewidth is necessary in order to selectively excite individual m_F states.

The process is described in Fig. 5.1. Light sourced from the stirring beam and resonant with the transition $^1S_0(F = 9/2, m_F) \rightarrow ^3P_1(F = 9/2, m_F - 1)$ ¹ propagates

¹Since this process is done in the dipole trap, these transitions experience a considerable ac Stark shift as compared to the free-space values resonance values.




Figure 5.1. Fermion spin polarization in a 9 G field. (a) The atoms are exposed to 9 consecutive resonant σ^- pulses. Atoms spontaneously decay from the upper states; the targeted states are (eventually) depleted, and the spin population is driven towards negative m_F values. (b) OSG results, after a few “cleanup” pulses done after evaporation. (c) Polarization purity optimization; light resonant with $m_F = -1/2$ is also resonant for $m_F = -5/2$ if the light has any σ^+ component. The grey line indicates the waveplate angle before this optimization.

downward onto the atoms. The pulse length is determined by $T = t_0/t_{m_F}$, where $t_0 = 1.8$ ms and t_{m_F} refers to the squared branching ratio of the particular transition. Atoms in state m_F are pumped to the excited state with $m_F - 1$, and then spontaneously decay; some return to m_F as well as $m_F - 1$, but the pulse is kept on long enough that eventually m_F is depopulated. A sequence of 9 such pulses fully polarizes the sample; however, it must be emphasized that Fermi degeneracy cannot be achieved with a single spin state: at least two must be present at the start of evaporation. So, the 9th spin pumping pulse is partially applied. In the end, I evaporate with the ratios $N_{-5/2}/N, N_{-7/2}/N, N_{-9/2}/N = 7\%, 23\%, 47\%$, with the rest in the higher m_F states. Further purification occurs approximately 500 ms before the experiment begins. Depending on the desired resulting configuration, this sequence consists of an additional series of σ^- pulses applied with $t_0 = 0.5$ ms. These beams are sufficiently far-detuned in the 9 G field that no heating is observed in the remaining atoms.

There were a few problems with our original spin-pol implementation. The first was polarization impurity: any σ^+ light simultaneously pumps m_F and $m_F - 2$, sabotaging our spin-pol scheme and making it harder to apply targeted blasts. I optimized this away in Fig. 5.1(c): by applying post-evaporative spin blasts on $m_F = -1/2$, I used the OSG




Figure 5.2. Beam alignment onto the atoms. A detuned laser is pulsed on the atoms just as they are dropped, followed by 4 ms TOF; each frame is a different realization of the experiment with the beam at a different position, given in each frame by the horizontal “step” of the Picomotor.

to monitor the relative spin population of $m_F = -5/2$, minimizing its loss as a function of waveplate angle. The starting waveplate position is shown as a gray line—off from optimal by about 5 degrees. It is interesting to note that the red MOT was relatively unaffected in either waveplate position.

The second issue had to do with retaining atoms after the pumping process. We achieved good pump-out results, but weren’t seeing a simultaneous population increase in the desired states. So, operating on the theory that the trap couldn’t contain the photon recoil, my idea was to simply increase the dipole trap frequency to more-tightly contain the atoms. After initially trapping the atoms in the ODT and before we apply spin blasts, we linearly ramp the ODT trap frequencies over 300 ms to (313 Hz, 397 Hz, 2.16 kHz) along the two horizontal and vertical dimensions, respectively, before ramping it back down and proceeding with evaporation. I configured the ramp time to minimize atom loss, and since the compression is mostly adiabatic, don’t notice a change in the resulting degeneracy.

5.1.2 Beam alignment

Lasers need to be aligned onto the atoms; this can be a hard problem. For 10 mm diameter MOT beams, it's pretty easy: simple geometric alignment with the windows is enough. Focused beams are another story. After evaporation, the atoms occupy an oblate spheroid 30-50 μm in diameter, and the nearest chamber windows are 200-400 mm away. Working with beams focused to finer than 1 mm is, as you might imagine, like finding a needle in a haystack. There are different approaches to this problem. In the case of the vOSG beam (see Fig. 3.14), I used the camera: first, I located the atoms on the camera, and then configured the software to image the laser by itself; overlapping the two gives a great starting point². If you don't have a camera (or one won't fit), it can be a lot harder, and you have to rely on the dipole force imparted to the atoms as the beam gets close.

The dipole-pushing strategy consists first of rough-aligning the beam such that entrance and exit spots hit the center of the chamber windows, which are by design centered on the atom location. Next, note that the optical dipole force is $F = -\alpha(\omega)\nabla I(r)$, where $\alpha(\omega)$ is the wavelength-dependant polarizability, and $\nabla I(r)$ is the gradient of the beam intensity. A red- (blue-)detuned beam pulls (pushes) the atoms towards (away from) the beam, predominantly in the *radial* direction—since Gaussian beams have their strongest gradient in the direction transverse to their *k*-vector. You then scan a 2D grid around the atoms until you see the laser push them.

I used Newport Picomotors on the SM lattice beams, which allowed me to scan the beam positions from the control computer. Even so, this can be a tedious process given the 20 s experimental repetition rate. Reasonable step sizes must be used or the process can take days; empirically, 100-200 steps are good step sizes for 10 mW beams focused to 200 μm by lenses positioned 400 mm from the atoms. One then looks for the effect of a push/pull on the atoms in time-of-flight, after diabatically pulsing the beam in-situ for

²Assuming your imaging path has minimal chromatic aberration, or that the imaging and pushing lights are similar in wavelength. Our imaging path uses achromats, so it's not a problem.

some time, typically 1-10 ms.

If it is particularly difficult, one can also use a ‘guide’ beam, which I used on one troublesome arm of the SM lattice. The source was a Ti:Sapph for which I had (essentially) infinite power, and the frequency control to make its effects more dramatic. I show an example of one such experiment in Fig. 5.2. Each frame shows a Fermi gas after being exposed to a 2 mW pulse of light at 434.827632 THz, close to the $^1S_0(F = 9/2) \rightarrow ^3P_1(F = 11/2)$ hyperfine transition. At step 0, the beam is approximately centered on the atoms; the (red-detuned) beam is then moved left, pulling the atoms with it, an effect which becomes quite pronounced. When it gets far enough away, its effects lessen, disappearing by step 725. Note that around step 400 it pulls the atoms down as well as left, indicating that the initial position was a bit low from center (or, that there is some cross-coupling in the Picomotor knob!). The key takeaway is that it is often hard to tell when a detuned beam is well-aligned with the atoms—because when it’s well-aligned, no effect is observable³. It is an unfortunate fact that one must disturb the alignment to check it. To verify that you’re exactly centered, you need to sweep the beam position left, right, up, and down.

5.1.3 Energy level shifting

The ground states are degenerate; we have to break this degeneracy, and in a nonuniform manner. Firstly, what magnitude shift do we need? To resolve lattice dynamics, we require Rabi coupling strengths Ω on the order of the temperature of the atoms, or $\Omega \sim c\kappa$, where $\hbar\kappa$ is the initial spread in atomic momentum of the DFG. On the other hand, to enable a unique correspondence between pairs of Raman beams $\mathbf{k}_i, \mathbf{k}_j$ and three spin states (see Fig. 2.3) requires a nonuniformity ϵ in their energy spacing such that $\epsilon \gg \Omega$. In alkali-like atoms, such shifts could be easily provided by magnetic bias fields. In alkaline-earth-like elements, the ground states are relatively insensitive to magnetic field.

³Except perhaps a decrease in the sample lifetime, owing to spontaneous emission.

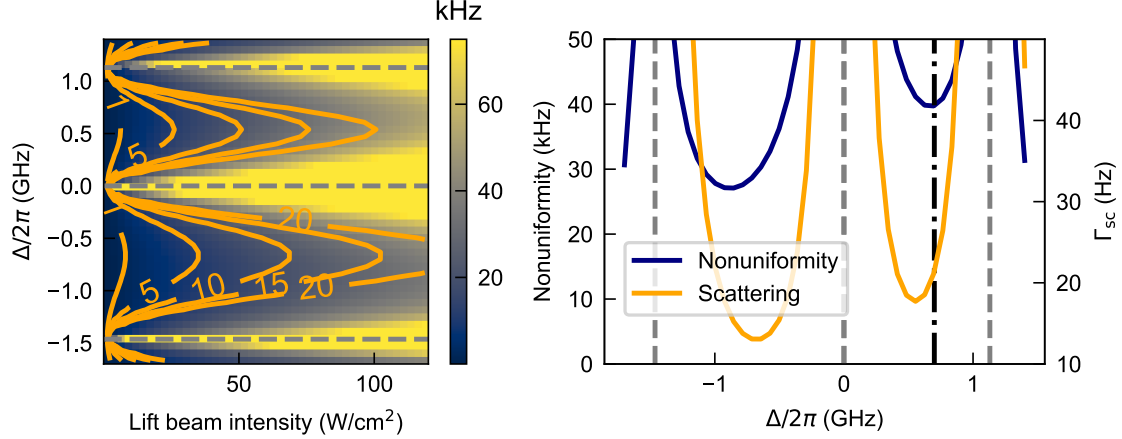


Figure 5.3. The design of the ac Stark shift “lift” beam. Left: Energy level nonuniformity as a function of beam power and detuning. Beam detuning is given relative to the $^1S_0(F = 9/2) \rightarrow ^3P_1(F = 9/2)$ transition. Contours of constant scattering rate are given in orange, in units of Hz. Right: Energy level nonuniformity and scattering against detuning, for a fixed power $I = 95 \text{ W/cm}^2$. Grey dashed lines indicate hyperfine resonances.

The Zeeman shift coefficient of ^{87}Sr is only 185 Hz/G and lacks any significant quadratic component [19], which could provide the nonuniformity.

We instead turn to an ac Stark shift beam called ω_{Lift} , detuned by Δ from the hyperfine resonance line $^1S_0(F = 9/2) \rightarrow ^3P_1(F = 9/2)$. This beam produces a tensor light shift [2] with coefficient α_t proportional to the detuning from the three hyperfine levels in 3P_1 as:

$$\alpha_t(F, \omega) = \sum_{F_e} (-1)^{F+F_e} \sqrt{\frac{40F(2F+1)(2F-1)}{3(F+1)(2F+3)}} \times \left\{ \begin{matrix} 1 & 1 & 2 \\ F & F & F_e \end{matrix} \right\} \frac{|\langle F || \mathbf{d} || F_e \rangle|^2}{2\hbar\Delta_{F_e}} \quad (5.1)$$

where the sum runs over the hyperfine levels $F_e = \{7/2, 9/2, 11/2\}$, $F = 9/2$ is the ground-state total angular momentum, Δ_{F_e} is the detuning from the hyperfine resonance, the curly-braces indicate the Wigner-6j symbol, and $|\langle F || \mathbf{d} || F_e \rangle|$ is the hyperfine dipole

matrix element given by

$$|\langle F || \mathbf{d} || F_e \rangle| = |\langle J || \mathbf{d} || J_e \rangle| (-1)^{F_e + J + 1 + I} \times \sqrt{(2F_e + 1)(2J + 1)} \begin{Bmatrix} J & J_e & 1 \\ F_e & F & I \end{Bmatrix} \quad (5.2)$$

with $J = 0$, $J_e = 1$, nuclear spin $I = 9/2$, and $|\langle J || \mathbf{d} || J_e \rangle|^2 = \Gamma_{689} \frac{2J_e + 1}{2J + 1} \frac{3\epsilon_0 \hbar c^3}{\omega_0^3}$ is the fine-structure dipole matrix element at 689.4 nm. The tensor ac Stark shift $\Delta\nu$ on an atom in state $|F, m_F\rangle$ due to a beam with intensity I_l , polarized linearly along the quantization axis, is given by $h\Delta\nu = -\alpha_t(F, \omega) I_l ((3m_F^2 - F(F + 1))/F(2F - 1))$.

The design criteria of this beam was twofold: we wanted as much nonuniformity as possible, while minimizing unwanted spontaneous emission. These constraints are exemplified in Fig. 5.3 (left), where I analyze a range of possible detunings vs. beam powers, with contours showing the spontaneous emission rates. On the right are the energy level nonuniformity and scattering rates as functions of beam detuning, for a fixed intensity $I = 95 \text{ W/cm}^2$.

We ended up choosing absolute frequency 434.829943(5) THz, with detuning $\Delta = 2\pi \times 700 \text{ MHz}$. The beam's ex-situ measured waists are $(\omega_x, \omega_z) = (350, 330) \mu\text{m}$, much larger than the in-situ cloud size of $\approx 30 \mu\text{m}$. This laser (linewidth 100 kHz) is fiber-coupled at the experiment to improve pointing stability. Its polarization is linear and aligned with the bias magnetic field. For an experimental power of 182 mW and the measured beam waists, we expect energy splittings $\Delta\nu_{XY} = 171 \text{ kHz}$ and $\Delta\nu_{YZ} = 130 \text{ kHz}$.

We verify the energy levels using Raman spectroscopy, as shown in Fig. 5.4. Starting from a DFG polarized mostly into $m_F = -9/2$, we flash on the Raman beams $\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3$ for varying times, with two beams probing a particular transition and the third beam detuned (but still providing its light shift). The frequency at which the population transfer is maximized determines the splitting [106]. The data is fit with a Rabi-like

model with temperature and pulse time as fixed parameters, varying center and Rabi frequency. After compensating for the recoil shift $4E_R/h = 14.49\text{kHz}$, we find the energy differences $\delta\nu_{XY} = 170.1 \pm 0.1\text{kHz}$ and $\delta\nu_{ZX} = 303.2 \pm 0.1\text{kHz}$, in excellent agreement with the expected light shift values.

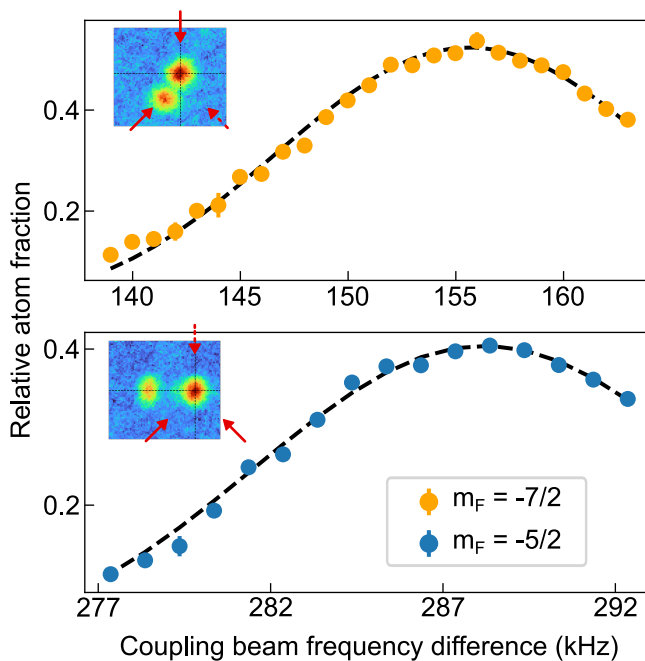


Figure 5.4. Calibration of the ac Stark shift induced by the lift beam with a Fermi gas spin-polarized into $m_F = -9/2$. The beams were flashed on for $50\mu\text{s}$ ($75\mu\text{s}$) when transferring into the $-7/2$ ($-5/2$) state giving fit center 155.6 kHz (288.7 kHz); compensating for the recoil energy, this is 170.1 kHz (303.2 kHz). Insets show exemplary spin-momentum-resolved time-of-flight data, with the crosshairs indicating the location of $\mathbf{p} = 0$. The solid red arrows indicate the involved beams, and the broken red arrows indicate a beam that is far-detuned from Raman resonance but providing its ac Stark shift.

5.1.4 Raman coupling

The three states $X(-9/2)$, $Y(-7/2)$, $Z(-5/2)$ are coherently coupled with Raman transitions (see Section 2.2). Here I discuss some details of the design choices, measurements of the coupling strengths, and some unique effects relating to the beam polarization which cause Raman interference.

Single photon detuning

For the scheme to work, spontaneous emission must be minimized: the couplings must be coherent, so single-photon uptake (and the associated spontaneous emission) must be suppressed. We are fortunate in this case to work with narrow linewidth transitions, and choose specifically ${}^3P_1(F = 9/2)$ as our upper state. We are then assured that coupling to the other fine structure levels, such as 3P_0 , is negligible: the single-photon scattering rate $\Gamma_{sc} \sim (\Gamma/\Delta)^2$; but in this case $\Gamma \sim 10$ kHz and $\Delta \sim 6$ THz, so $(\Gamma/\Delta)^2 \approx 0$. This is not to be taken for granted; other species don't have this luxury. ${}^{87}\text{Rb}$, for instance, has similar fine structure spacing but 1000 times the linewidth, with the resulting destructive interference severely limiting detuning choice and sample lifetime [78].

That being said, we must still choose a reasonable detuning given our available laser power. With the scattering rate $\Gamma_{sc} \sim I(\Gamma/\Delta)^2$ and the coupling strength $\Omega \sim I(\Gamma/\Delta)$, we always win by increasing the detuning—if we can afford the intensity increase necessary to keep the same Rabi frequency. As a compromise between these two goals (and overlapping that compromise with our available AOMs) we chose $\Delta/2\pi = -210$ MHz, although in retrospect, -400 MHz would have been a better choice in terms of lifetime.

Coupling strengths

By observing Rabi oscillations between pairs of coupling beams, we calibrate the coupling strengths Ω_i and quantify the degree to which the SM lattice is limited to the spin projections $m_F = -5/2, -7/2, -9/2$. The Rabi measurement protocol differs from the main experiment, due to our dual use of the Ti:Sapph laser as both lift beam and as a source of optical Stern-Gerlach (OSG) pulses, which are separated by several GHz. While we could bridge the difference with a series of AOMs, we instead apply an open-loop drive voltage to the laser's cavity lock in order to change the laser's function from lift to OSG, as follows.

After applying a Raman pulse of varying duration, both Raman and lift beams are

snapped off while the ODT is kept on. In the next 50 ms, we ramp the laser frequency from the lift beam frequency 434.829943(5) THz to the OSG frequency 434.828370(30) THz. The larger frequency uncertainty induced by the open-loop control translates to small positional variations of the atomic spin populations from shot-to-shot. We account for this in the data analysis by binning the spin locations relative to the location of the locally-maximum atom cloud, which is unambiguously the “spin-up” population since the transfer efficiency is never larger than 0.5 (see Fig. 5.5). Empirically, the 50 ms ramp time dampens the Raman-kicked atomic motion sufficiently that the OSG pulse separates the m_F states along the \hat{y} direction only.

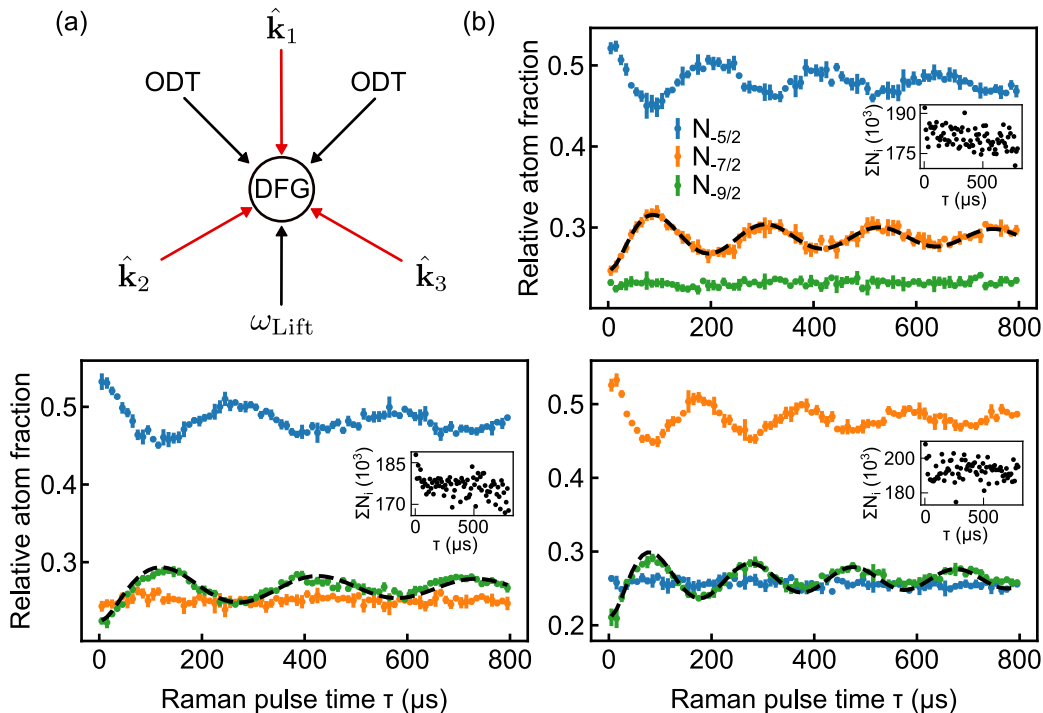


Figure 5.5. Coupling strength calibration and lattice closure measurement from OSG data. (a) Beam setup. For each probed transition, two beams at a time were tuned on-resonance, with the third beam off-resonance but providing its ac Stark shift. (b) Rabi oscillations due to a pulse resonant with (top) $-5/2 \rightarrow -7/2$, (bottom left) $-7/2 \rightarrow -9/2$, (bottom right) $-5/2 \rightarrow -9/2$. Insets show total atom count in the triplet $\sum_i^{-5/2, -7/2, -9/2} N_i$ at each pulse length τ_p . Triplet losses due to coupling to adjacent m_F states (such as $-1/2, -3/2$) would present as atom loss, bounded here to $< 10\%$.

We summarize these measurements in Fig. 5.5. Atoms with quasimomentum \mathbf{q} predominantly populate the “spin up” state $|\uparrow = -5/2, \mathbf{q}\rangle$ in (a) and (c), or $|\uparrow = -7/2, \mathbf{q}\rangle$ in (b). A light pulse of varying length is shone onto the tensor-shifted atoms, the OSG beam is fired as described above, and then all coupling and trapping beams are snapped off. This results in Rabi oscillations between bare states $|\uparrow, \mathbf{q}\rangle$ and $|\downarrow, \mathbf{q} + \sqrt{3}\hbar k_L\rangle$. We fit to the spin-down population with a model [69, 116]

$$P_{\downarrow} = A \int_{-v_0}^{v_0} dv P_e(-\sqrt{3}k_L v, \Omega, t) G_v(\sigma) + N_{\downarrow,0} \quad (5.3)$$

where A is the amplitude, t is the pulse length, $P_e(\delta, \Omega, t) = \frac{\Omega^2}{\Omega^2 + \delta^2} \sin^2(\frac{t}{2}\sqrt{\Omega^2 + \delta^2})$, $G_v(\sigma)$ is the Gaussian velocity distribution with standard deviation $\sigma = \sqrt{2k_B T/m}$, and $N_{\downarrow,0}$ is the fixed amount of atoms populating $|\downarrow\rangle$ at $t = 0$. With only two fit parameters, Ω and A , we find good agreement with the model, observing the decoherence characteristic of finite-temperature SO-coupled fermions [66, 68]. The Rabi coupling strengths between each link in the spin-momentum lattice are obtained by scaling the Rabi frequencies obtained in Fig. 5.5. The scaling considers both the respective branching ratios of each transition, and also the small measured differences in power between each tone in the driving AOMs. The strengths are listed in Table 5.1, with a total uncertainty considering the involved beam waists (measured through Fig. 5.5), polarization angle uncertainty of 1° , and the measured power uncertainty in each tone, which we take to be 5%.

The insets of Fig. 5.5 indicate the total atom number in all three spin states X, Y, Z . Since the timescale of spontaneous emission is small compared to the pulse durations, atom loss is due to atoms coupling to the adjacent dipole-allowed states $m_F = -1/2, -3/2$, which are not counted in the OSG images. The loss into these external states is expected to scale with the Rabi coupling strength and here is at most 10%, showing that the spin-momentum lattice experiment is closed to the spin states $-9/2, -7/2, -5/2$.

Table 5.1. Experimental two-photon coupling strengths Ω_{mn} in units of E_R/\hbar . Some transitions are associated with two coupling strengths, which interfere with relative phase angle $\chi = 2\pi/3$; see text for details.

	$X \rightarrow Y$	$Y \rightarrow Z$	$Z \rightarrow X$
$\mathbf{k}_1 - \mathbf{k}_2$	$0.61(10)e^{-i\chi} +$ $0.47(7)$	$0.70(11)e^{-i\chi} +$ $0.50(8)$	$0.45(7)$
$\mathbf{k}_2 - \mathbf{k}_3$	$0.61(10)e^{i\chi} +$ $0.47(7)e^{-i\chi}$	$0.65(10)e^{i\chi} +$ $0.46(7)e^{-i\chi}$	$0.43(7)$
$\mathbf{k}_3 - \mathbf{k}_1$	$0.63(10) +$ $0.49(8)e^{i\chi}$	$0.58(9) +$ $0.42(6)e^{i\chi}$	$0.42(6)$

Raman interference

In describing the net effective Rabi strengths Ω_{EFF} which form each link in the lattice, we must carefully consider the polarizations. Because each beam \mathbf{k}_i can drive σ^+ , σ^- , and π transitions, “double- Λ ” type couplings [117] are realized between the links $X \leftrightarrow Y$ and $Y \leftrightarrow Z$; see Fig. 5.6 for the specific coupling between $X(-9/2)$ and $Y(-7/2)$. The coupling frequencies may cause a $\pi\sigma^-$ or $\sigma^+\pi$ transition, since the upper-state detunings in each case satisfy $\Delta/\Delta' \approx 1$ and thus have similar coupling strengths (in contrast, the couplings $X \leftrightarrow Z$ may proceed only via $\sigma\sigma$ due to dipole selection rules). For such links, we define a net effective coupling strength $\Omega_{\text{EFF}} = \Omega_1 + e^{i\chi}\Omega_2$, where χ is the relative phase and Ω_i is the two-photon Rabi frequency associated with a specific polarization scheme, as in [72, 73]. In our experiment, the phase between π - and σ -components is fixed, unlike experiments in Refs. [72] and [73], which modify this phase using a tunable path length difference or electro-optic modulator, respectively. As we show, our phases χ are fixed by the beam intersection angle.

Consider ω_1 inside beam \mathbf{k}_1 , propagating along the $-\hat{y}$ direction as in Fig. 5.6(a), with polarization angle $\theta = 33^\circ$ with respect to the xy -plane. The electric field, with real amplitude E_{10} , including a possible phase ψ , which can vary from shot to shot, is

$$\mathbf{E}_1 = E_{10}(\cos(\theta)\hat{x} + \sin(\theta)\hat{z})e^{i\mathbf{k}_1 \cdot \mathbf{r}}e^{i\psi} \quad (5.4)$$




Figure 5.6. The effect of Raman interference in a double- Λ level structure. (a) Two beams incident on a spin with wavevectors k_1, k_2 in the xy -plane, whose polarization vectors are tilted at angle $(\pi/2 - \theta)$ with respect to the z -axis. The beams cross at angle χ . (b) Both beams have components π, σ^+, σ^- , and the upper-state detunings Δ are approximately equal. Thus, Raman transitions can proceed along two equivalent paths with differing coupling strengths Ω_1 and $e^{i\chi}\Omega_2$. This interference must be considered when computing an effective Rabi frequency Ω_{Eff} .

Similarly, the electric field describing ω_2 in \mathbf{k}_2 is

$$\mathbf{E}_2 = E_{20}(-\sin(\phi)\cos(\theta)\hat{x} + \cos(\phi)\cos(\theta)\hat{y} + \sin(\theta)\hat{z})e^{i\mathbf{k}_2\cdot\mathbf{r}} \quad (5.5)$$

with $\phi = 30^\circ$ being the acute angle between \hat{k}_2 and the x -axis. These two fields drive a Raman transition $X \leftrightarrow Y$ when their frequency differences are close to the energy splittings. As shown in Fig. 5.6(b), there are two ways an atom can make the jump, with respective two-photon Rabi frequencies Ω_1, Ω_2 proportional to $E_1^{\pi*}E_2^{\sigma^-}, E_1^{\sigma^{+*}}E_2^{\pi}$. The π -components of the fields are those along the \hat{e}_z direction, and the σ^\pm components are those along the directions $\hat{e}_\pm = \mp 1/\sqrt{2}(\hat{x} \pm i\hat{y})$. We also note that the geometric scaling factor resulting from the branching ratio associated with Ω_1 is $-7\sqrt{2}/99 = \exp(i\pi)7\sqrt{2}/99$ (associated with Ω_2 is $\sqrt{2}/11$). We then find

$$\begin{aligned} \Omega_1 &\propto E_1^{\pi*}E_2^{\sigma^-} = E_{10}E_{20}/\sqrt{2}\sin(\theta)\cos(\theta)e^{-i\psi} \times e^{i(\mathbf{k}_2-\mathbf{k}_1)\cdot\mathbf{r}}e^{i(-\phi-\pi/2)} \\ \Omega_2 &\propto E_1^{\sigma^{+*}}E_2^{\pi} = E_{10}E_{20}/\sqrt{2}\sin(\theta)\cos(\theta)e^{-i\psi} \times e^{i(\mathbf{k}_2-\mathbf{k}_1)\cdot\mathbf{r}} \end{aligned}$$

We find the total phase difference $\chi = -\phi - \pi/2 = -2\pi/3$ between these two paths; we

furthermore see that any phase fluctuation in ψ is common-mode and so does not affect Ω_{Eff} .

To evaluate the impact of the Raman interference phases on the flux lattice scheme, we sum the phases around the plaquette drawn in Fig. 2.3(b), in which an atom makes a complete cycle $X \rightarrow Y \rightarrow Z \rightarrow X$. We write the total Rabi strength of transition mn , where $mn = \{XY, YZ, ZX\}$, as $\Omega_{mn}(\chi_{mn}) = \Omega_{mn,1} + e^{i\chi_{mn}}\Omega_{mn,2} = A_{mn}e^{i\Phi_{mn}}$, with

$$\begin{aligned} A_{mn} &= \sqrt{\Omega_{mn,1}^2 + \Omega_{mn,2}^2 + 2\cos(\chi_{mn})\Omega_{mn,1}\Omega_{mn,2}} \\ \tan(\Phi_{mn}) &= \frac{\sin(\chi_{mn})\Omega_{mn,2}}{\Omega_{mn,1} + \cos(\chi_{mn})\Omega_{mn,2}} \end{aligned} \quad (5.6)$$

Adopting the coordinate system of Eqn. 5.4 and Eqn. 5.5, we note that electric field of \mathbf{k}_3 is $\mathbf{E}_3 = E_{30}(-\sin(\phi)\cos(\theta)\hat{x} - \cos(\phi)\cos(\theta)\hat{y} + \sin(\theta)\hat{z})e^{i\mathbf{k}_3 \cdot \mathbf{r}}$. We then find the phases χ_{mn} to be

$$\begin{aligned} \chi_{XY} &= -\phi - \pi/2, \\ \chi_{YZ} &= \phi + \pi/2, \\ \chi_{ZX} &= 2\pi \end{aligned} \quad (5.7)$$

With no AOM phases set, an atom encircling the plaquette experiences a phase pickup $\exp(i(\Phi_{XY} + \Phi_{YZ} + \Phi_{ZX}))$. Using Eqn. 5.6, we find

$$\begin{aligned} \Phi &= \Phi_{XY} + \Phi_{YZ} + \Phi_{ZX} \\ &= \arctan\left(\frac{-\cos(\phi)\Omega_{XY,1}}{\Omega_{XY,1} - \sin(\phi)\Omega_{XY,2}}\right) + \arctan\left(\frac{\cos(\phi)\Omega_{YZ,1}}{\Omega_{YZ,1} - \sin(\phi)\Omega_{YZ,2}}\right) \end{aligned} \quad (5.8)$$

For equal coupling strength ratios $\Omega_{XY,2}/\Omega_{XY,1} = \Omega_{YZ,2}/\Omega_{YZ,1}$, we see that $\Phi = 0$.




Figure 5.7. Spin-resolved imaging process. (a) Spin-blast Rabi frequency calibrated by electron shelving of the $m_F = -7/2$ (Y) state. The fitted single-photon Rabi frequency is $\Omega = 335(1)$ kHz. (b) Averaged atom shot of the SM lattice at sweep time $8E_R$. (c) Spin-blast beam resonant with $m_F = -7/2$ applied to the SM lattice. (d) Subtraction of images (b) and (c), removing the common-mode background and revealing the locations of $m_F = -7/2$ atoms. A negative-OD region forms at the locations of the blast-scattered atoms.

5.1.5 Spin-resolved imaging

Atoms released from the SM lattice naturally spatially resolve according to their spin and momentum; however, in cases where the starting state is not pure or for verification purposes, it is necessary to have a means to resolve the spins. Directly imaging the spins using narrow-line absorption imaging [104] is impractical due to the low scattering cross-section and relatively low atom density of Fermi gases. Furthermore, since the atoms are moving after they are released from the SM lattice, the use of optical Stern-Gerlach separation is not desirable because it will (necessarily) disturb the momentum distribution.

We instead visualize the spin dynamics with blast pulses [68], propagating in the xy plane, resonant with the narrow $^1S_0(F = 9/2, m_F) \rightarrow ^3P_1(F = 9/2, m_F)$ transition at 689 nm. To be effective, the blast beam strength needs to be similar to the worst-case atomic Doppler shifts in the SM lattice $\Omega_{\text{Doppler}} \approx n \times 2\pi \times 4E_R = n \times 91$ kHz, where n is the number of photons absorbed in the SM lattice. Simultaneously, to avoid exciting neighboring m_F states, the blast strength must be smaller than the typical Zeeman splitting of the upper $^3P_1(F = 9/2)$ states, here $2\pi \times 790$ kHz.

The beams are calibrated as in Fig. 5.7(a) with electron-shelving [101]: atoms

are excited into 3P_1 with a resonant 689 nm pulse for a time t , after which we apply a $4 \mu\text{s}$ pulse of 461 nm MOT light to blow away all remaining ground-state atoms. Atoms so “shelved” in the 3P_1 state do not interact with the 461 nm light; instead, the fraction projected onto the ground state is measured with absorption imaging after a brief 2 ms time of flight. The overall decay curve matches the 3P_1 natural lifetime of $22 \mu\text{s}$. The SM lattice in Fig. 5.7(b) is subjected to this beam in Fig. 5.7(c), yielding the spin-resolved picture in Fig. 5.7(d) after subtracting the two images.

Since the blast process relies on spontaneous emission, an atom can scatter (at most) a few photons before decaying to an adjacent m_F state, which is then transparent to the narrow-linewidth blast beam. This means the momentum acquired by the targeted atoms is comparable to the momentum of the atoms in the SM lattice, so some overlap is inevitable—as the negative optical density region in Fig. 5.7(d) shows. This effect can be particularly noticeable if the scattered atoms have strong geometric overlap with a region of interest, as seen in the X Fig. 5.8, particularly the $8E_R$ panel, which entirely masks the cloud at $\mathbf{p} = -3\hbar k_L \hat{y}$.

Furthermore, although the beam strengths are chosen to be comparable to the Doppler shifts, the efficiency of these single-tone blast pulses still decreases with increasing SM lattice photon absorption. To overcome this, a larger bias magnetic field would enable the use of stronger blast pulses. To overcome the geometric overlap issue, tilting the blast pulses out-of-plane would push targeted atoms out of the imaging focus region, more effectively removing them from the images.

5.2 Filling the lattice

Having set up the necessary tools and configured the atomic energy levels to our liking, we must now ask ourselves (a) how we can visualize and measure a momentum-space lattice, and (b) how we can populate that lattice.

5.2.1 Visualization

In general, real-space cold-atom lattices are visualized indirectly: one typically does not have the imaging resolution to visualize atoms spaced on the order of the beam’s wavelength, typically ≈ 500 nm or less. Instead, we measure the atomic *momentum distribution*, which is dramatically altered by the presence of a periodic potential (and is generally well-described by Bloch wavefunctions). In practice, this means turning off the lattice beams and allowing the atoms to freely expand under gravity—this is called a “time of flight” (TOF) measurement. Adiabatically ramping down the lattice is called a *band mapping* procedure [118], which smoothly deforms the tight-binding (flat) bands into the free-particle parabolas. In such a ramp down, quasimomentum is preserved and the Brillouin zone may be directly mapped.

Another method is simply diabatically turning off the periodic potential. Since the atoms are described by Bloch waves, rapidly turning off the potential amounts to a plane-wave decomposition of the Bloch states. An excellent overview of these two methods is given in [119]. In our experiment, we diabatically turn off the potential; after 12 ms TOF, the atoms are absorption-imaged, and the physical positions of the atoms map directly onto real momentum, reflecting the in-situ momentum distribution.

5.2.2 Population and atomic motion

We need a way of populating the different lattice sites. Atoms starting at $\mathbf{p} \approx 0$ do not have the kinetic energy necessary to tunnel to different lattice sites, which are spaced by $\sqrt{3}\hbar k_L$: they must be given a “push.” This is accomplished by continuously ramping the frequencies of a triplet of beams. In the experiment, we ramp \mathbf{k}_1 , although we could just as well have ramped either of the other two triplets. In so doing, the atoms move preferentially along the \mathbf{k}_1 direction, because $\partial_t(\omega_1, \omega'_1, \omega''_1) > 0$.

The direction of atomic motion can be explained by a conservation of energy

argument. Consider the state $|X, \vec{p} = 0\rangle$; the transition $|X, \vec{p} = 0\rangle \rightarrow |Y, \vec{p} = \sqrt{3}\hbar k_L(\hat{k}_3 - \hat{k}_1)\rangle$ via the blue beams ω'_3, ω'_1 is increasingly off-resonant with increasing ω'_1 . To see this, we note that the effective detuning from Raman resonance δ (for stationary atoms) is given by [61]:

$$\delta = (\omega'_3 - \omega'_1) - (\Delta\omega_{XY} + (\hbar(\mathbf{k}_3 - \mathbf{k}_1))^2/2m) \quad (5.9)$$

where $\Delta\omega_{XY} > 0$ is the Zeeman energy splitting between spin states X and Y , and $(\hbar(\mathbf{k}_3 - \mathbf{k}_1))^2/2m = 4E_R/\hbar$ is the recoil energy. In the experiment, at $t = 0$ we set the laser frequency difference to match the energy splitting ($\omega'_3 - \omega'_1 = \Delta\omega_{XY}$) so that $\delta = -4E_R/\hbar$. If $\partial_t\omega'_1 > 0$, δ becomes more negative, and the Raman process is further suppressed. The upward motion is only favorable if $\partial_t\omega'_1 < 0$, in which case $\delta \rightarrow 0$ as $\omega'_3 - \omega'_1 \rightarrow 4E_R/\hbar$. A similar argument can be applied to the upward blue transition $X \rightarrow Z$.

Thus, atoms hopping in the downward direction are increasingly Doppler-shifted due to their acquired momentum; these shifts are (imperfectly) compensated by the continuous ramping of the beam \mathbf{k}_1 , which propels them further downward. A relatively small number of atoms do venture upward, owing to the finite fermion temperature; those atoms moving away from the beam at $t = 0$ are the ones most likely to make an upward transition.

5.2.3 Model

In order to predict the dynamics, we solve the TDSE. First we define the wave vectors as

$$\mathbf{k}_1 = -k_L\hat{u}_y, \quad \mathbf{k}_2 = \frac{k_L}{2}(\sqrt{3}\hat{u}_x + \hat{u}_y), \quad \mathbf{k}_3 = \frac{k_L}{2}(-\sqrt{3}\hat{u}_x + \hat{u}_y) \quad (5.10)$$

where $k_L = 2\pi/\lambda$. Note that although these beams do not have the same frequency, the difference between wavenumbers is negligible. Since the transitions induced by these

lasers involve two-photon processes, the corresponding momentum transfer is related to differences between $\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3$. We define the single-photon recoil vector $k_R = 2\pi/\lambda \sin(\theta/2)$ where $\theta = 120^\circ$ is the angle between any two beams. The recoil vector becomes $k_R = \sqrt{3}\pi/\lambda = \sqrt{3}k_L/2$, so the wavevectors become

$$\mathbf{k}_1 = -\frac{2}{\sqrt{3}}k_R\hat{u}_y, \quad \mathbf{k}_2 = \frac{k_R}{\sqrt{3}}(\sqrt{3}\hat{u}_x + \hat{u}_y), \quad \mathbf{k}_3 = \frac{k_R}{\sqrt{3}}(-\sqrt{3}\hat{u}_x + \hat{u}_y) \quad (5.11)$$

The relative momentum \mathbf{q}_i is defined as

$$\mathbf{q}_1 = \mathbf{k}_1 - \mathbf{k}_2 = k_R(-\hat{u}_x - \sqrt{3}\hat{u}_y), \quad \mathbf{q}_2 = \mathbf{k}_2 - \mathbf{k}_3 = 2k_R\hat{u}_x, \quad \mathbf{q}_3 = \mathbf{k}_3 - \mathbf{k}_1 = k_R(-\hat{u}_x + \sqrt{3}\hat{u}_y) \quad (5.12)$$

All magnitudes are $2k_R$, consistent with the motivation of defining the single-photon recoil vector. For convenience, we also define

$$\begin{aligned} \omega_{Y \rightarrow X} &= \omega_1 - \omega_2 = \omega_2'' - \omega_3'' = \omega_3' - \omega_1', \\ \omega_{X \rightarrow Z} &= \omega_1''(t) - \omega_2'' = \omega_2' - \omega_3' = \omega_3 - \omega_1, \\ \omega_{Z \rightarrow Y} &= \omega_1'' - \omega_2'' = \omega_2' - \omega_3' = \omega_3 - \omega_1(t) \end{aligned} \quad (5.13)$$

The time dependence only shows up in $\omega_1, \omega_1', \omega_1''$ since only the frequencies in beam \mathbf{k}_1 are swept during the experiment.

To simulate the quantum dynamics in the spin-momentum (SM) lattice, we first set up the Hamiltonian in the Bloch basis, $\psi_{\mathbf{q},\alpha}(\mathbf{r}) = e^{i\mathbf{q}\cdot\mathbf{r}}u_{\mathbf{q},\alpha}(\mathbf{r})$ where \mathbf{q} is the crystal momentum, α is the spin species and $u_{\mathbf{q}}$ is a periodic function. The corresponding Schrödinger equation can be expressed as

$$\left(\frac{(\hat{\mathbf{P}} + \hbar\mathbf{q})^2}{2m} + \epsilon_\alpha + \Omega(\hat{\mathbf{r}}, t) \right) u_{\mathbf{q},\alpha}(\mathbf{r}, t) = i\hbar \frac{\partial}{\partial t} e^{i\mathbf{q}\cdot\mathbf{r}} u_{\mathbf{q},\alpha}(\mathbf{r}, t) \quad (5.14)$$

Note that ϵ_α labels the energy of the internal state. Due to the periodicity of the $u_{\mathbf{q},\alpha}(\mathbf{r}, t)$,

we can decompose $u_{\mathbf{q},\alpha}(\mathbf{r},t)$ into its Fourier components,

$$u_{\mathbf{q},\alpha}(\mathbf{r},t) = \sum_{\mathbf{K}} c_{\mathbf{K},\alpha}(\mathbf{q},t) e^{i\mathbf{K}\cdot\mathbf{r}} \quad (5.15)$$

where $\mathbf{K} = n_1\mathbf{q}_1 + n_2\mathbf{q}_2$ is the reciprocal lattice vector, described by integers n_1, n_2 . The Schrödinger equation becomes

$$\sum_{\mathbf{K}} e^{i\mathbf{K}\cdot\mathbf{r}} \left(\frac{\hbar^2(\mathbf{q} + \mathbf{K})^2}{2m} + \epsilon_\alpha + \Omega(\hat{\mathbf{r}},t) \right) c_{\mathbf{K},\alpha}(\mathbf{q},t) = \sum_{\mathbf{K}} e^{i\mathbf{K}\cdot\mathbf{r}} i\hbar \frac{\partial}{\partial t} c_{\mathbf{K},\alpha}(\mathbf{q},t) \quad (5.16)$$

Next, we formulate the coupling term with frequency sweep velocity v as

$$\begin{aligned} \Omega(\hat{\mathbf{r}},t) &= \sum_n \sum_{\alpha \neq \beta} \Omega_{\alpha\beta} e^{i\mathbf{q}_n \cdot \mathbf{r} - i(\delta_{n,1} - \delta_{n,3})vt^2} e^{-i\omega_{\beta \rightarrow \alpha}t} |\alpha\rangle \langle \beta| \\ &= \sum_n \sum_{\alpha \neq \beta} \Omega_{\alpha'\beta'} e^{i\mathbf{q}_n \cdot \mathbf{r} - i(\delta_{n,1} - \delta_{n,3})vt^2} |\alpha'\rangle \langle \beta'| \end{aligned} \quad (5.17)$$

where $\delta_{m,n}$ is the Kronecker delta, and the last equation comes from using the rotating frame to absorb the usual oscillation term.

In our simulation, we use a radial grid of 832 points in momentum space and use the LSODA differential equation solver [120] to predict the dynamics at each point, time-evolving a Gaussian distribution which closely matches the initial momentum spread. We fix the 9 coupling strengths according to Table 5.1, the values there obtained experimentally. We have empirically found that these coupling strengths are approximately 20% too large to account for the observed dynamics, and thus uniformly scale all Rabi strengths accordingly.

5.3 Experimental realization

What follows is the manuscript of our results, submitted and accepted to Physical Review Letters [113]. The figure fonts have been modified to be consistent with the other figures in this thesis.

5.3.1 Introduction

Ultracold atoms in optical lattices have been established as an important tool for the quantum emulation of condensed matter models [121], especially those with topological features [15, 122]. The inherent tunability afforded by optical lattices provides access to a variety of parameter regimes, which has proved essential in the seminal realizations of topological phases in ultracold matter [85, 123–125]. Since then, efforts to study topology in other systems have led to the exploration of synthetic dimensions [64, 126], which provide internal degrees of freedom beyond those afforded by the trapping geometry and have enabled a new generation of experiments [127].

Several approaches to synthetic dimensions have been experimentally realized. Real-space lattices augmented with spin-orbit coupling (SOC) connect spin “lattice” sites via momentum exchange, creating Hall cylinders pierced by magnetic flux in a synthetic position-spin space [105, 128–134], or creating Hall ribbons in optical clock experiments [70, 71, 135, 136]. Real-space lattices are not always needed; SOC itself can provide synthetic degrees of freedom, which can act as a potent generator of Berry curvature [137–142] or provide control parameters for Hamiltonian engineering [74, 143]. Synthetic lattices entirely in momentum-space [144, 145] have been realized, and, with carefully engineered hopping schemes, have proven topological [146–149]. Recently, a synthetic lattice of Rydberg states has been employed for the study of a Su–Schrieffer–Heeger model [150], and a synthetic dimension of trap states created with patterned light [151].

Lattices composed of spin and momentum states, or spin-momentum (SM) lattices, have been proposed [91] as a platform to exhibit topological features, with some schemes potentially realizing the Laughlin state of the fractional quantum Hall effect [89, 152]. As a step towards this, we realize a fermionic spin-momentum lattice using SOC and three atomic Zeeman spin states. Previous experiments using spin-momentum lattices utilized bosons in a single-dimension [153, 154], and used a real-space lattice with lattice-band

pseudospins [155]. Recently, a two-spin bosonic SM lattice has been implemented in an optical cavity [156]. Here, by providing sufficient links between three spin sites, we build a lattice of fermions in a 2D spin-momentum space, without a traditional scalar optical lattice. This platform increases the flexibility of the synthetic dimension approach. In particular, the use of three spin states in two spatial dimensions allows the simulation of synthetic magnetic fields of high spatial uniformity, which lead to ultra-narrow Chern bands that support robust fractional quantum Hall states [89, 157].

5.3.2 Implementation

The synthetic lattice is composed of three Zeeman spin states in the $^1S_0(F = 9/2)$ ground state of ^{87}Sr , labeled $X \equiv |m_F = -9/2\rangle, Y \equiv |m_F = -7/2\rangle, Z \equiv |m_F = -5/2\rangle$. In a momentum-dependent manner, the spins are cyclically coupled by up to 9 Raman lasers intersecting at 120° . In the rotating-wave approximation, we describe the atom-laser coupling as

$$\hat{V} = \Omega_{mn} e^{i(\mathbf{k}_R \cdot \mathbf{r} + \varphi_i - \varphi_j)} |m\rangle \langle n| + \text{H.c.} \quad (5.18)$$

where $m \neq n$ runs over the states X, Y, Z , $|\mathbf{k}_R| = |\mathbf{k}_i - \mathbf{k}_j| = k_L \sin^2 \frac{\theta}{2}$ is the magnitude of the single-photon recoil wavevector with $i \neq j$ denoting the beams driving a particular m - n coupling, Ω_{mn} is the coupling strength, and $\theta = 120^\circ$ is the angle between any pair of beams. The single-photon recoil energy is $E_R = (\hbar k_R)^2 / 2m = \hbar \times 22.7$ kHz. The phase differences $\varphi_i - \varphi_j$ are set to zero in the experiment, but we note that setting nonzero phases is at the heart of the ultra-narrow-band optical flux lattice experiment [89].

The setup implementing the optical couplings in Eqn. (1) is shown in Fig. 1(a). Up to three running-wave triplets of beams are incident on a degenerate Fermi gas (DFG) spin-polarized mostly into state $|X\rangle$ with $T/T_F = 0.36(5)$, where T_F is the Fermi temperature. Each beam $\hat{\mathbf{k}}_i$ contains up to three frequencies $\omega_i, \omega'_i, \omega''_i$, such that the energy difference between any two frequencies $\omega_i('') - \omega_j('')$ matches an energy difference in the X, Y, Z

manifold. These beams provide a Raman coupling between states, as in Fig 1(b). The quantization axis is defined by a \hat{z} -oriented magnetic field $B \approx 9.3$ G, along which we align the linear polarization of a beam ω_{Lift} [105, 106, 158] providing a strong ac Stark shift that lifts the degeneracy of the states X, Y, Z . The coupling beam polarizations are linear and angled at $33(1)^\circ$ with respect to the xy -plane, projecting approximately equal intensity among the possible Raman transition types π, σ^\pm . When using all nine frequencies, the beams form an infinite lattice in spin-momentum space, as in Fig. 1(c). Since the average starting atomic momentum $\langle p \rangle \ll 4\hbar k_R$, the fermions initially occupy only a small spread of states $|X/Y/Z, \mathbf{q} \approx 0\rangle$, where \mathbf{q} is the quasimomentum. After engaging the coupling beams, atoms with differing spin or \mathbf{q} occupy adjacent spin-momentum lattices, shifted by their corresponding quantum numbers. These lattices are not tight-binding, in the sense that in-situ motion is constrained only by the overall harmonic trap, allowing particles of different spin and \mathbf{q} to mix. Unlike a square lattice—which does not readily allow for nonzero Chern number—this triangular lattice naturally breaks inversion symmetry and allows for magnetic flux. We note that nonzero gauge flux is possible only on the upward-pointing triangles, corresponding to momentum transfers involving all three frequency sets.

As shown in Fig. 1(d), the coupling beams utilize the dipole-forbidden transition $^1S_0(F = 9/2) \rightarrow ^3P_1(F = 9/2)$, detuned below resonance by $\Delta/2\pi = 210$ MHz. The transition’s narrow linewidth $\Gamma/2\pi = 7.4$ kHz allows coherent manipulation with minimal spontaneous emission, and no significant destructive interference [78, 159] arises from the THz-separated fine structure states 3P_0 and 3P_2 . In order to make each triplet unique, the upper-state detunings of the blue and green couplings are shifted by $\mp 37 E_R/\hbar$, larger than the $\approx 7.5 E_R/\hbar$ energy splittings. The role of the frequency ω_i in beam \mathbf{k}_i is circularly rotated amongst the three triplets, such that all frequencies resonantly couple all spin states.

In order to realize the lattice, careful attention must be paid to the energy levels of

X, Y, Z , which are naturally degenerate. Since the ground states have $J=0$ —rendering Zeeman shifts insignificant at our bias field—we use an ac Stark shift approach. Further leveraging the narrow-line intercombination transition, the lift beam ω_{Lift} is operated at 434.829943(5) THz, midway between the hyperfine resonance lines $^1S_0(F = 9/2) \rightarrow ^3P_1(F = 7/2)$ and $^1S_0(F = 9/2) \rightarrow ^3P_1(F = 9/2)$. This light was designed to produce a strong tensor shift $\epsilon = 1.76E_R/\hbar$ of the X state, allowing each pair of Raman beams to uniquely couple two spin-momentum states. A necessary condition of the spin-momentum lattice model is that the coupling strengths $\Omega \ll \epsilon$. The coupling strengths here are $\Omega \approx 0.5E_R/\hbar$.

5.3.3 Sequence

In our newly-built apparatus, we source ^{87}Sr from a commercial atomic oven from AOSense, which includes an integrated Zeeman slower and 2D MOT optics. After two MOT loading and cooling stages lasting 7s [1, 50], the atoms are loaded into a crossed 1064 nm optical dipole trap (ODT) with an initial temperature of $\approx 2 \mu\text{K}$. The vertical (horizontal) trapping frequency is ramped up to 2.160(5) kHz (313(1), 397(2) Hz), at which point we spin-polarize the sample with a series of pulses resonant with the different m_F states via the $^1S_0(F = 9/2) \rightarrow ^3P_1(F = 9/2)$ transition [115, 160]. The ODT frequencies are then lowered back to ≈ 1 kHz vertically, and forced evaporation proceeds over the next 10 s, finally reaching a quantum-degenerate sample. Without spin-polarization, we routinely achieve $T/T_F = 0.20$, where T_F is the Fermi temperature, rising to $T/T_F = 0.36(5)$ when spin-polarized. Evaporation ends at mean geometric trap frequency $\bar{\omega} = (\omega_x\omega_y\omega_z)^{\frac{1}{3}} = 71.4(1)$ Hz, yielding a 50 nK Fermi gas. Immediately following evaporation, the sample is spin-polarized in the state $X(80 \pm 7\%)$, and the ac Stark shifting beam ω_{Lift} is ramped on in 0.5 ms. Via optical Stern-Gerlach imaging [104], we verify that this timescale does not alter the spin polarization. We then introduce the coupling beams with a turn-on time of $< 1\mu\text{s}$.

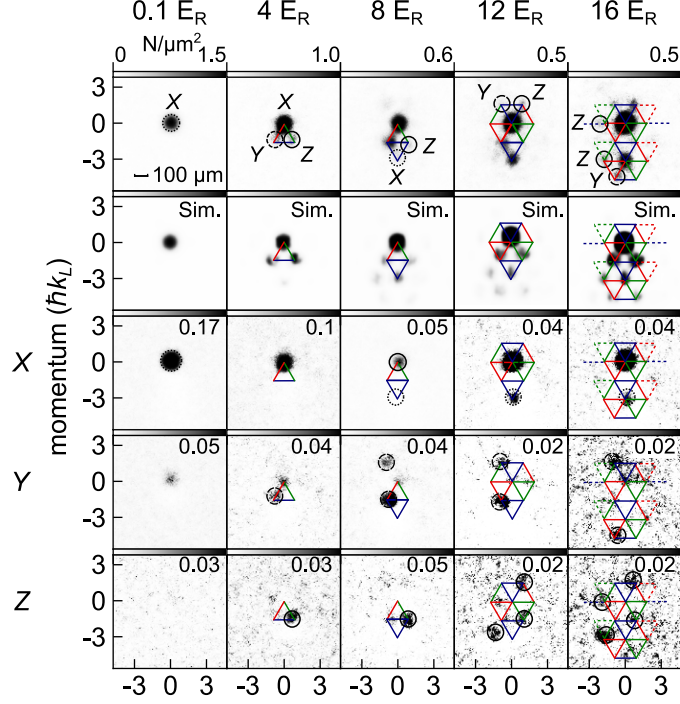


Figure 5.8. Demonstration of the spin-momentum lattice, see text for details. The sweep progress is indicated in units of recoil energy E_R . The top row shows spin-unresolved momentum-space images, the second row the predicted dynamics, and subsequent rows show spin-resolved images. As a guide to the eye, the links are also drawn across all rows with colors corresponding to Fig. 2.3. Each image is an average of ≈ 100 experimental runs taken at 12 ms time of flight.

We demonstrate the spin-momentum lattice in Fig. 5.8. In order to fill more sites, we emulate motion along a single dimension by subjecting the atoms to an inertial force [133] along $\hat{\mathbf{k}}_1$, ramping all three of that beam's frequencies at a rate $\hbar\partial_t(\omega_1, \omega'_1, \omega''_1) = 16.607 E_R/\text{ms}^4$. Hopping to neighboring sites is made favorable when the frequency difference between two coupling beams matches the energy and recoil shifts between states, providing enhanced state transfer between initial state $|X, \mathbf{q}\rangle$ and $|X, \mathbf{q} - \mathbf{K}\rangle$, $|Y, \mathbf{q} - \mathbf{K}\rangle$, $|Z, \mathbf{q} - \mathbf{K}\rangle$ for some reciprocal lattice vectors $\mathbf{K} = n_1\mathbf{q}_1 + n_2\mathbf{q}_2$ with integers n_1, n_2 . After a varying sweep time, all optical fields are quenched off, releasing the atoms from the harmonic trap. Atoms that have tunneled to different lattice sites acquire a concomitant increase in momentum, in discrete units of the two-photon Raman momentum $\hbar k_R =$

⁴We do not expect this sweep rate to be adiabatic with respect to the current Rabi coupling strengths.

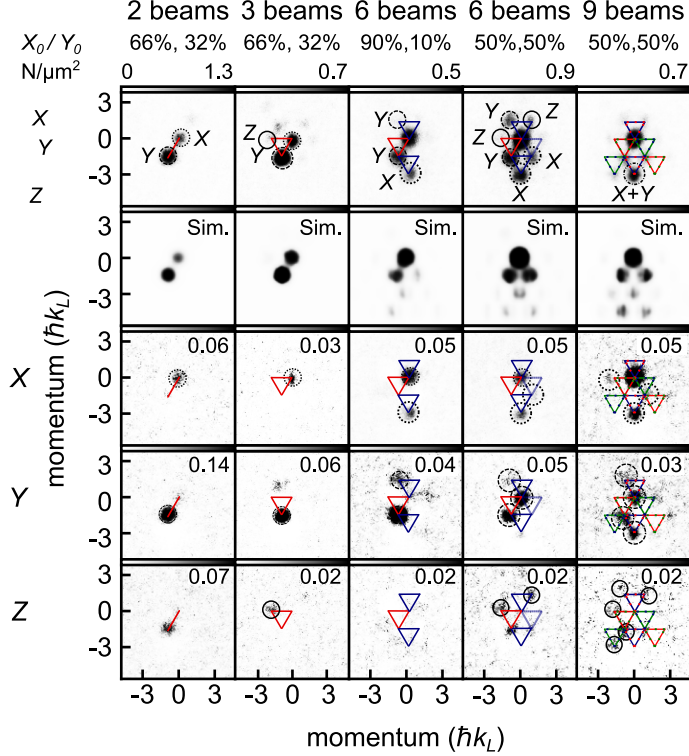


Figure 5.9. Building the spin-momentum lattice; all images taken at a common sweep time of $12E_R$. Each removed frequency corresponds to a missing link in the full lattice setup, allowing the exploration of quasi-1D SOC in the first column, to 2D SOC in the second column, culminating in the full lattice shown in the last column. The multiply-colored links in the last two columns indicate two SM lattices simultaneously overlaid: a two-state configuration, in which the experiment began with equal populations of X, Y .

$\sqrt{3}/2\hbar k_L$. The lattice sites become spatially resolved after 12 ms time of flight, since the starting momentum distribution's full-width half-max width is $1.05(1)\hbar k_L$ and external heating by spontaneous emission from 3P_1 is minimal. The atoms are then absorption-imaged in the xy plane using the $^1S_0 \rightarrow ^1P_1$ transition at 461 nm ($\Gamma_{461}/2\pi = 30\text{MHz}$), which images all spins with approximately equal efficiency [161]. The bias magnetic field is kept on at all times, in order to maintain the spin quantization axis.

The individual columns of Fig. 5.8 demonstrate spin- and momentum-resolved imaging at various quench times. Sweep time is indicated by the final frequency deviation of the swept beam, in units of E_R . Intuitively, one would not expect stationary atoms ($\langle p \rangle \approx 0$) to tunnel before at least overcoming the recoil shift $4E_R$, and we observe

this in the experiment. Denoting transferred momentum by $\hbar\delta k_{ij} = \hbar(|\mathbf{k}_i - \mathbf{k}_j|)$, and referring to beam triplets by their colors in Fig. 1(d), when the sweep reaches $4E_R$ we see beams ω_1, ω_2 from the red triplet driving the corresponding Raman transition $|X, 0\rangle \rightarrow |Y, \hbar\delta k_{12}\rangle$; similarly, the green beams ω'_1, ω'_3 allow $|X, 0\rangle \rightarrow |Z, \hbar\delta k_{13}\rangle$. By $8E_R$, atoms have firmly populated sites $|Y, \hbar\delta k_{12}\rangle, |Z, \hbar\delta k_{13}\rangle$, with initially-faster-moving atoms beginning to populate the site $|X, p = -3\hbar k_L\rangle$, completing a traversal of the first Brillouin zone. By $12E_R$, more atoms have tunneled through the Brillouin zone, and the momentum center-of-mass proceeds downward at $16E_R$; imaging becomes increasingly difficult due to the lower atom density, so we terminate here. As a consistency check, we also demonstrate spin-resolved imaging using spin blasts [68] in order to verify that sites on the SM lattice are of the expected spin projection, m_F . We observe good consistency with the SM lattice model as drawn in Fig. 1(d), although mechanical effects of the spin-blasts can mask some lattice sites; notably, the X -site at $\mathbf{p} = -3\hbar k_L \hat{y}$. Our model shows qualitative agreement with the data for a scaled value of the measured Rabi coupling strengths. Some disagreement is evident, especially at the lattice sites with momenta $\mathbf{p} = \sqrt{3}/2\hbar k_L$, which are predicted to have a stronger amplitude than is observed. We attribute these mismatches to off-resonant effects not included in our effective Hamiltonian.

The lattice scheme presented here is readily tunable. Although the full spin-momentum lattice is composed of nine frequencies, we can remove links between lattice sites at will. We explore this flexibility in Fig. 5.9, where we show the driven dynamics experiment of Fig. 5.8, but now with all images taken at a common sweep time $12E_R$. In the two-beam scheme, composed of a single frequency in each of two beams \hat{k}_1, \hat{k}_2 , we have reduced the system to a 1D SOC model between an effective spin up $|\uparrow\rangle = |X, \mathbf{q}\rangle$ and spin down $|\downarrow\rangle = |Y, \mathbf{q}\rangle$ [66, 68]. In the 3-beam case, with a single frequency in each of the $\hat{k}_1, \hat{k}_2, \hat{k}_3$ beams, we have a 2D spin-orbit coupling [79, 139, 162] cyclically linking the three states X, Y, Z . The 6-beam case consists of beams $\hat{k}_1, \hat{k}_2, \hat{k}_3$ each possessing two frequencies, labeled by their colors red and blue as labeled in Fig. 2.3. In the last

two columns, we explore the dynamics starting from an even spin mix of states X and Y , which, in the 9-beam experiment, can be visualized as two SM lattices overlapped on $\mathbf{p} = 0$.




Figure 5.10. Impact of too-small triplet frequency separation. An earlier version of this experiment separated the triplets by only ± 525 kHz (now ± 850 kHz), resulting in the strong population of positive-momentum lattice sites, a result not expected from the direction of the frequency sweep.

5.3.4 Conclusion and outlook

We have demonstrated a two-dimensional fermionic spin-momentum lattice without the use of standing waves. This adds to the wealth of cold atom synthetic dimension platforms available to study topological materials. The system’s 15 ms lifetime exceeds our current experimental duration by a factor of 10, and could be further improved with increased Raman detuning Δ . The current lift beam strength imposes a 30 ms limit, which can be relaxed under appropriate conditions. The number of visible lattice sites can be increased with larger Rabi coupling strengths, or by slowing the sweep rate, which would couple more atoms out of the $p \approx 0$ momentum class. The spin-resolved imaging presented here could be improved by using stronger blast pulses to overcome the Doppler shifts among the lattice’s numerous momentum states. This work launches a novel platform for exploring topological physics with optical flux lattices. The natural extension of this work would be to load the atoms adiabatically into the lowest band and set nonzero coupling phases such that a gauge flux appears on the plaquettes labeled α, β, γ in Fig. 2.3. The topology of the band structure could then be probed using established anomalous velocity




Figure 5.11. Comparison of the 9-beam SM lattice data from Fig. 5.8 (left column) with the model (middle column), showing the difference image (right column). The coupling strengths used are given in Table 5.1, all of which are uniformly scaled down by 20%. We apply to the model images a small Gaussian blur (std. dev. = 2 pixels) commensurate with the imaging resolution of ≈ 1.5 pixels = $0.1\hbar k_L$. The units are optical density.

techniques [85, 133, 163], which involve accelerating the dressed atoms in the same manner as done here. Demonstrating this topology would enable the exploration of many-body fractional Hall states [89]. In the present experiment, spin-contact interactions are not expected to play a significant role; but these interactions may be increased through the use of a vertical real-space lattice, as detailed further in [89] for ^{87}Rb , which we note has the same scattering length $\approx 100a_0$ as ^{87}Sr .

Some disagreement with the model persists, as can be seen in Fig. 5.11. The model under-predicts the upward-going atomic motion, an effect responsible for the over-predicted population with $\mathbf{p} \approx 0$ at later evolution times. These effects were especially noticeable in an earlier version of the experiment, in which we had a smaller triplet separation;

compare with Fig. 5.10, which shows significant population in modes with $\mathbf{p} > 0$ due to cross-coupling between triplet beams.

We attribute the discrepancy to both the lingering effects of inter-triplet interference, and to off-resonant Raman couplings in the SM lattice—i.e., couplings which reverse the notion of which beam is considered ‘pump’ and which is considered ‘Stokes.’ By increasing the tensor shift and separation between frequency triplets (currently ± 850 kHz), such effects can be further suppressed.


5.3.5 Acknowledgements

This chapter, in part, has been accepted for publication to Physical Review Letters [113], and is copyright 2022 by the American Physical Society. The dissertation author was the primary author of this paper.

Appendix A

Science chamber schematics


Below is the drawing describing the science chamber, exactly as sent to Sharon Vacuum. For the most part, they did an excellent job, with two minor quips; the first time they sent it to us, they rotated port 1a incorrectly, which we had to send to them to fix. We also noticed they missed threading one hole on port 13. Neither turned out to be a big deal, and we commend their professionalism and quality of construction.



DATE	4/11/2016	TITLE	
DESIGNED		BY	
CHECKED		DATE	
APPROVED		USER	D
		PROJECT	UCSD Main Chamber
		SHEET	1 of 3

**All flanges are clearance holes except for port 13.
All flanges fixed (non-rotatable).**



Port Number	Flange OD (in.)	Focal Point	Flange Distance (in.)	Theta (°)	Phi (°)	Notes
1 (a/b)	1 1/3	P0	a: 6.30 b: 5.00	a: 8 b: -8	90	a: Clearance hole must be aligned vertically. Port 1a is atomic beam entrance and needs to be focused on Focal Point P0 to < 0.5° accuracy. b: Through hole CFL33 Flange <1° accuracy
2 (a/b)	2 3/4	P0	8.00	a: 30 b: -30	90	<1° accuracy
3 (a/b)	1 1/3	P0	5.80	a: 45 b: -45	90	<1° accuracy
4 (a/b)	2 3/4	P0	8.0	a: 60 b: -60	90	<1° accuracy
5 (a/b)	4 5/8	P0	6.00	a: 90 b: -90	90	<5° accuracy
6 (a/b)	2 3/4	P0	8.0	a: 120 b: -120	90	<1° accuracy
7 (a/b)	1 1/3	P0	5.80	a: 135 b: -135	90	<1° accuracy
8 (a/b)	2 3/4	P0	10.00	a: 150 b: -150	90	<1° accuracy
9 (a/b)	2 3/4	P0	14.50	a: 172 b: -172	90	<1° accuracy
10 (a/b)	2 3/4		Flange coords. X=+7/8, Y=3.0 (b) Z=-7/8			Centerline of the horizontal tube from port 13. Align sniffer port vertically.
11 (t/b)	10	P0	3.47	t: 0 b: 180		Sniffer port aligned with y axis
12	2 3/4		Flange coords. X=9/5 Y=7/5 Z=-7/6			Centerline of horizontal tube from port 13. Align sniffer port vertically.
13	6	P1	4.78	180	90	Flange must have blind-tapped holes. Sniffer port must be aligned vertically (with the Z-axis) to <1° accuracy
14	8	P1	17.00		0	




DESIGNED	4/11/2016	TITLE	
CHECKED		DATE	
SKETCHED		SCALE	1:1
APPROVED		USER	D
		PROJECT	UCSD Main Chamber
		SHEET	2 OF 3

All mounting tabs: weld top and bottom sides of tabs where they connect to chamber. They must support the chamber's weight. Bottom of all tabs must be coplanar. Tabs should be parallel to XY plane.


Mounting holes all 0.27 thru. All tabs 0.25in thick.



SECTION Top tabs-Top tabs
SCALE 1 / 3



SECTION Bottom tabs-Bottom tabs
SCALE 1 / 2




Side
SCALE 1 / 3

DESIGNED	DATE	SCALE	SHEET
DRAWN	REV		
CHECKED	DATE		
APPROVED	DATE		
PROJECT		UCSD Main Chamber	
JOB NO		D	
SCALE		1 / 3	
SHEET		3	

Appendix B


Recessed viewports

Our “bucket windows” are constructed by the UK Atomic Energy Authority, with much help from Tim Darby (contact Tim.Darby@ccfe.ac.uk). These took about 3 months to construct, after which they were then sent to LaserOptik GmbH in Germany for the application of an anti-reflection coating (contact Heiko Haensel, hhaensel@laseroptik.de). Our 'top' bucket window is flat, while its partner is angled at 2 degrees to avoid etalon effects. The damage threshold for the coating for CW use is “about 1-2 kW per mm beam radius.”



ITEM	QTY	PARTS LIST PART NUMBER	DESCRIPTION
1	1	200CF Flange	SS316LN
2	1	Re Tube	SS316
3	1	Top Plate	SS316
4	1	50KV Ferrule Sleeve	Inconel
5	1	50kv Ferrule	Inconel
6	1	50kv x 5mm THK	Fused Silica Spectrosil 2000 TWE Lambda 10 20-10 Scratch Dig Customer to arrange AR coating


DRAWN 14/06/2016	CHECKED	TITLE	REV
QA	RFG	SIZE	A3084 200CF +80.54 50CV GA V1 TD
APPROVED		SCALE	1
			SHEET 1 OF 1



SECTION A-A
SCALE 1.5:2

ITEM	QTY	PART NUMBER	DESCRIPTION
1	1	200CF Flange	SS316LN
2	1	Re Tube	SS316
3	1	Top Plate	SS316
4	1	50CV Ferrule Sleeve	Inconel
5	1	50cv X 5mm THK	Fused Silica Spectrosil 2000 TWE Lambda 10 20-10 Scratch Dig Customer to arrange AR coating
6	1	50cv Ferrule	Inconel

DRAWN	16/06/2016	TITLE	
CHECKED		REV	
QA		SIZE	C
PMG		SCALE	
APPROVED		DWG NO	A3084 200CF +76.54 2d 50cv GA V1 T
		SHEET 1	OF 1

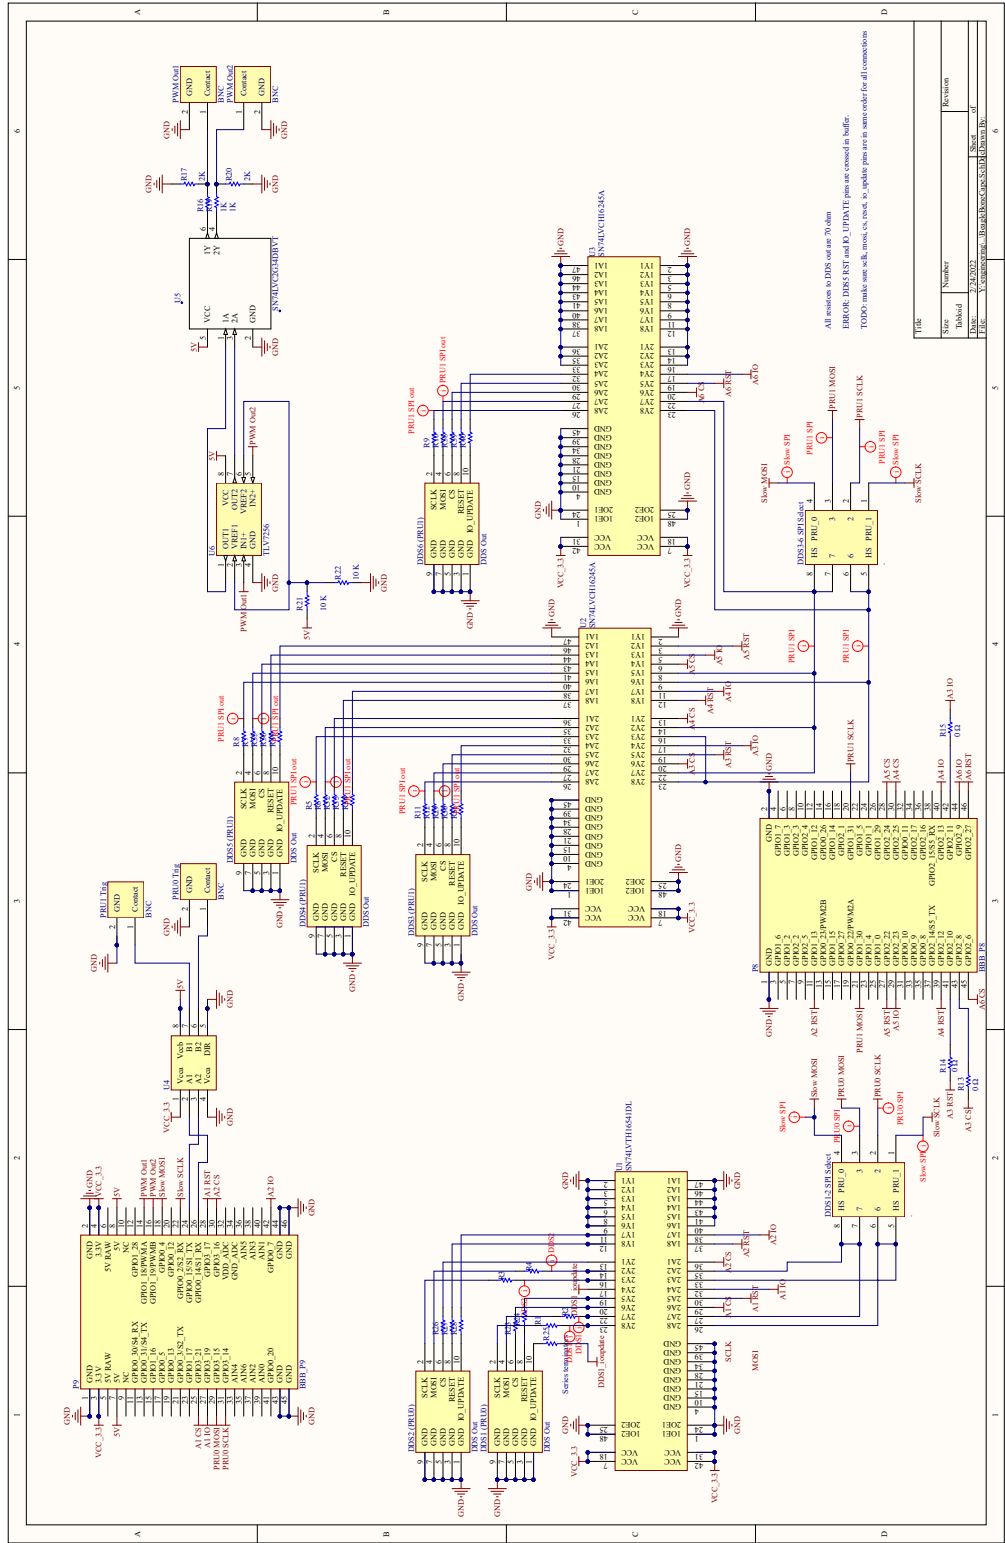


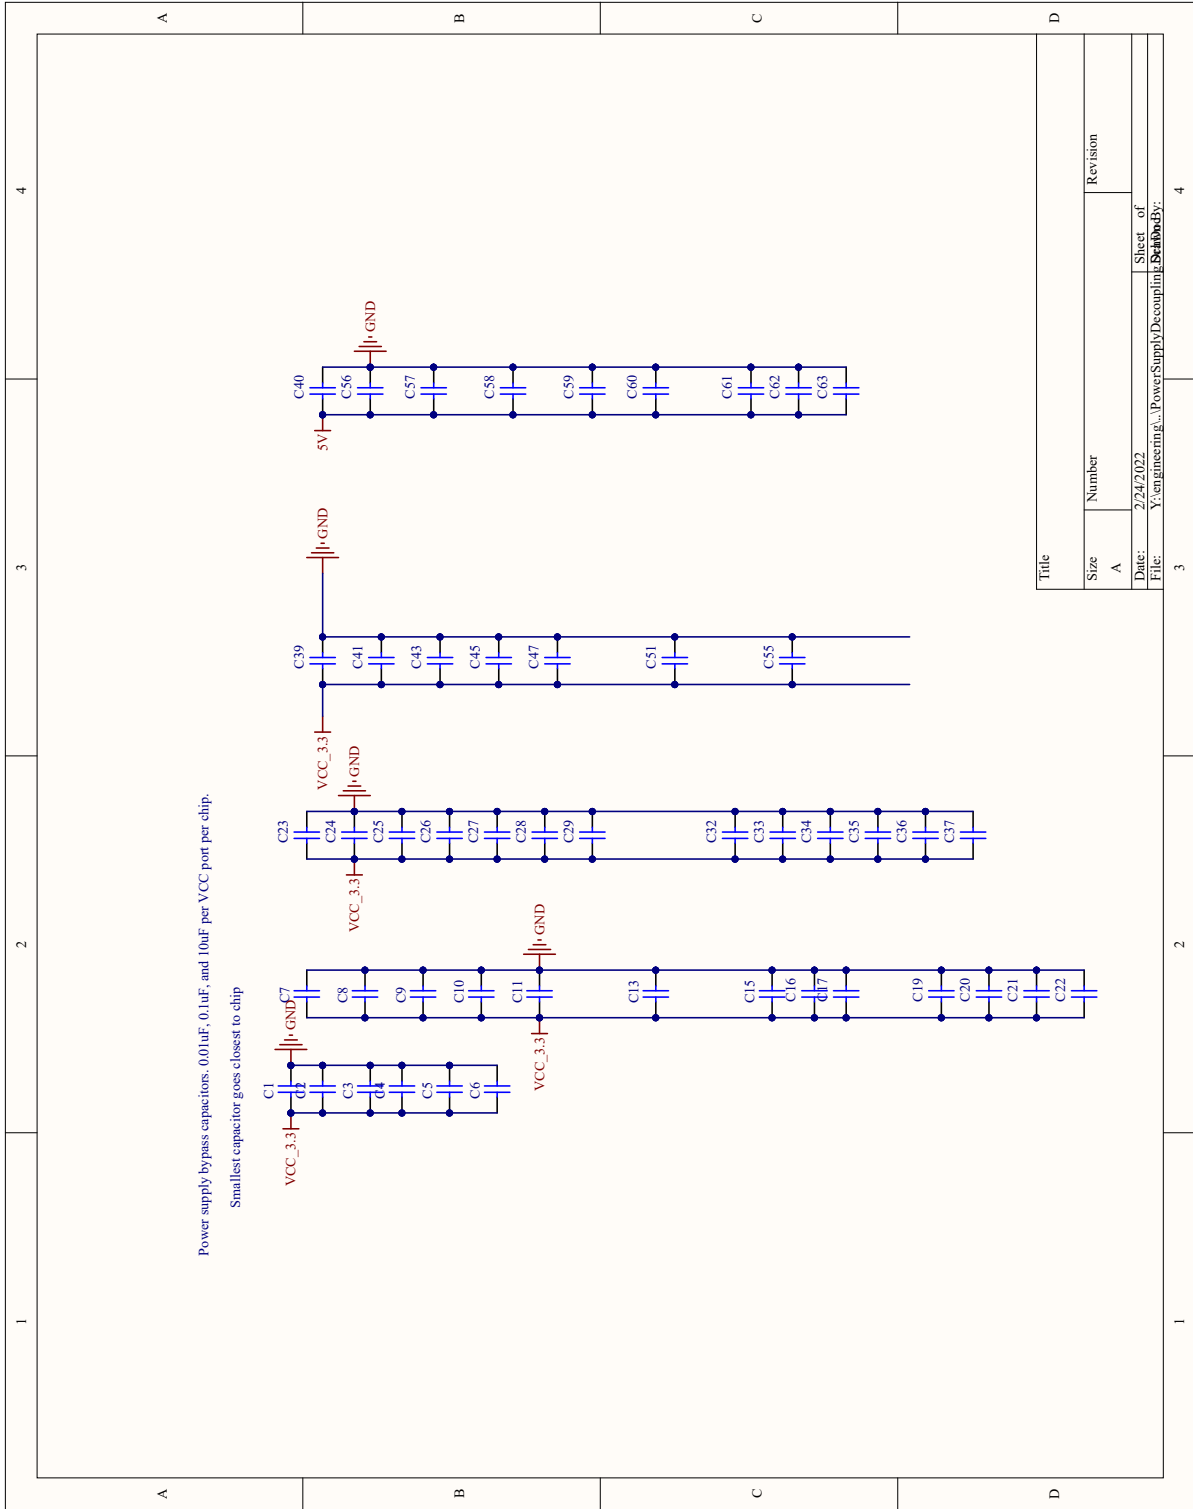
B-13927: AR450-850nm/0° on UHV

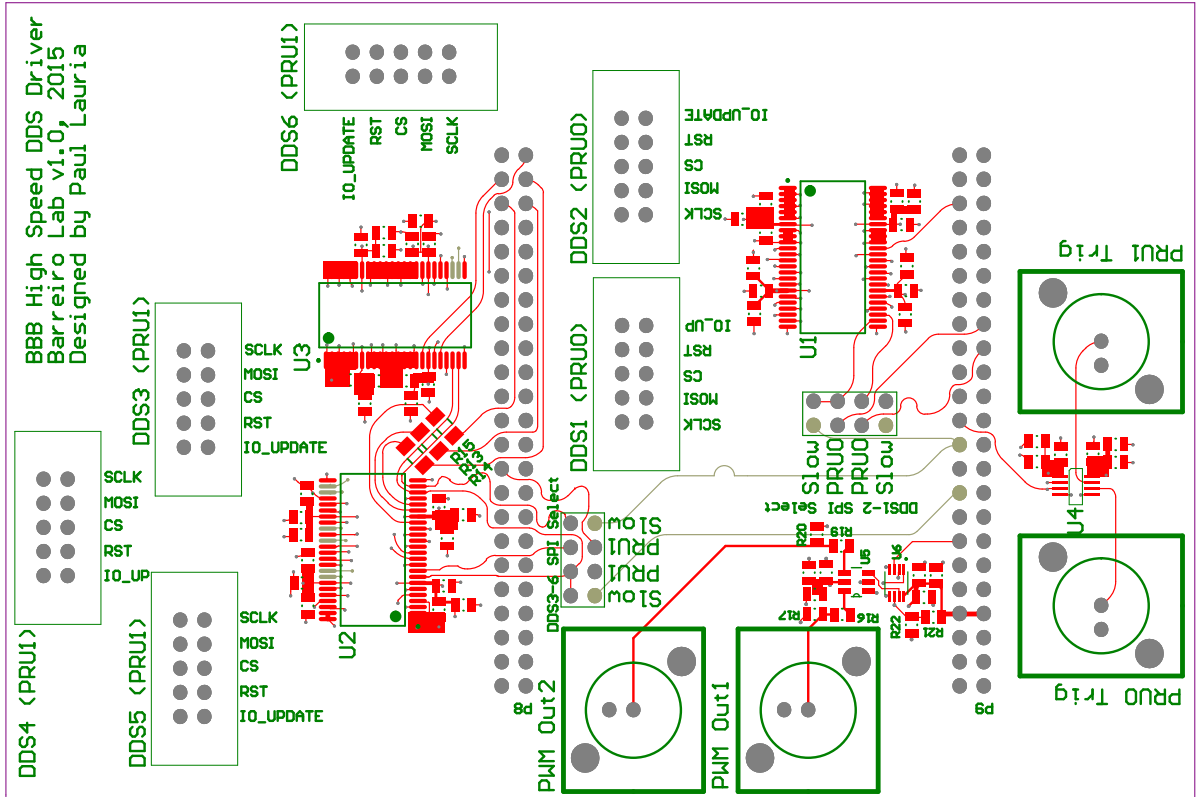
Appendix C

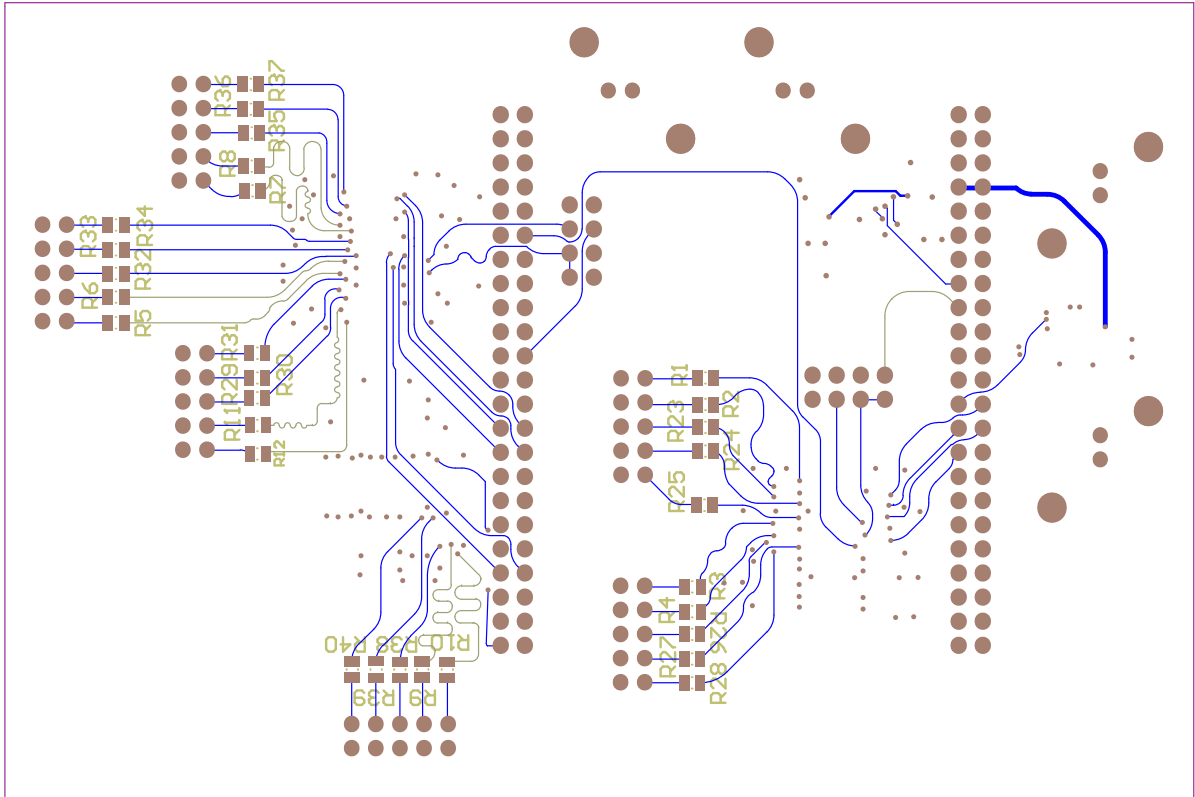
DDS Driver

I designed a custom PCB which can drive any DDS that uses a standard SPI configuration—specifically, one that uses only SCLOCK, MOSI, IOUPDATE, chip select, and RESET. It mounts on top of a BeagleBone Black or Beaglebone Green, as in Fig. 3.20.





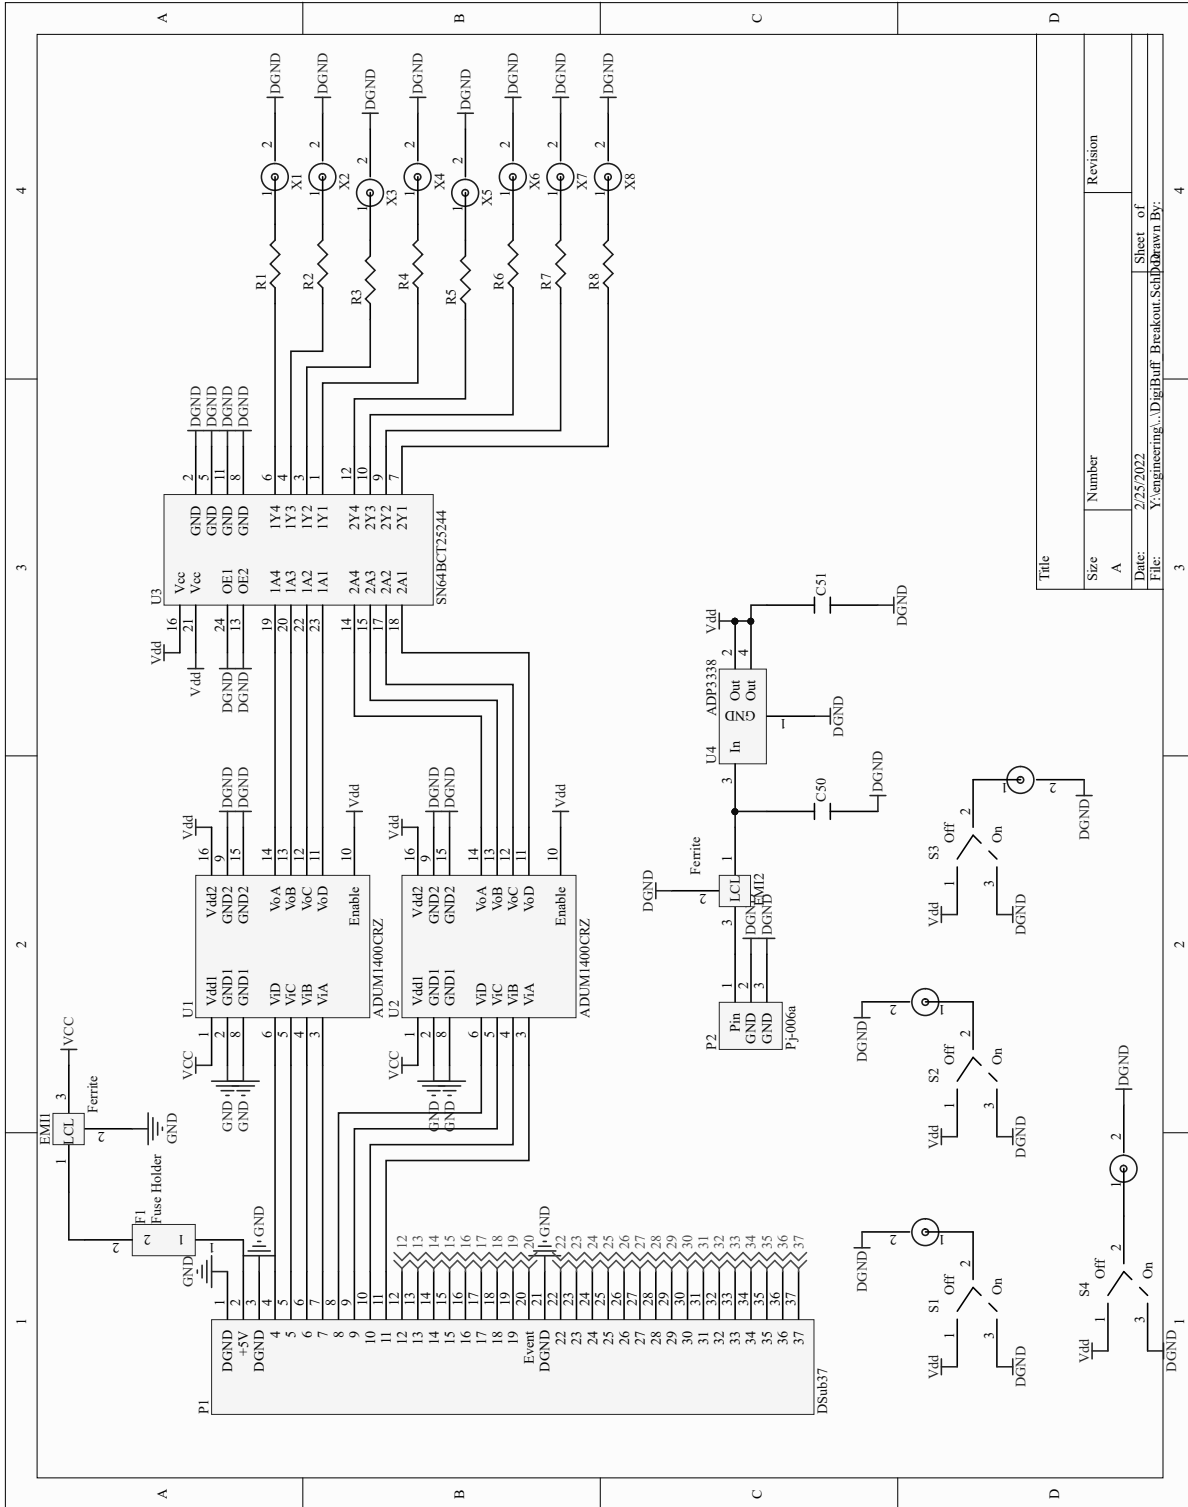




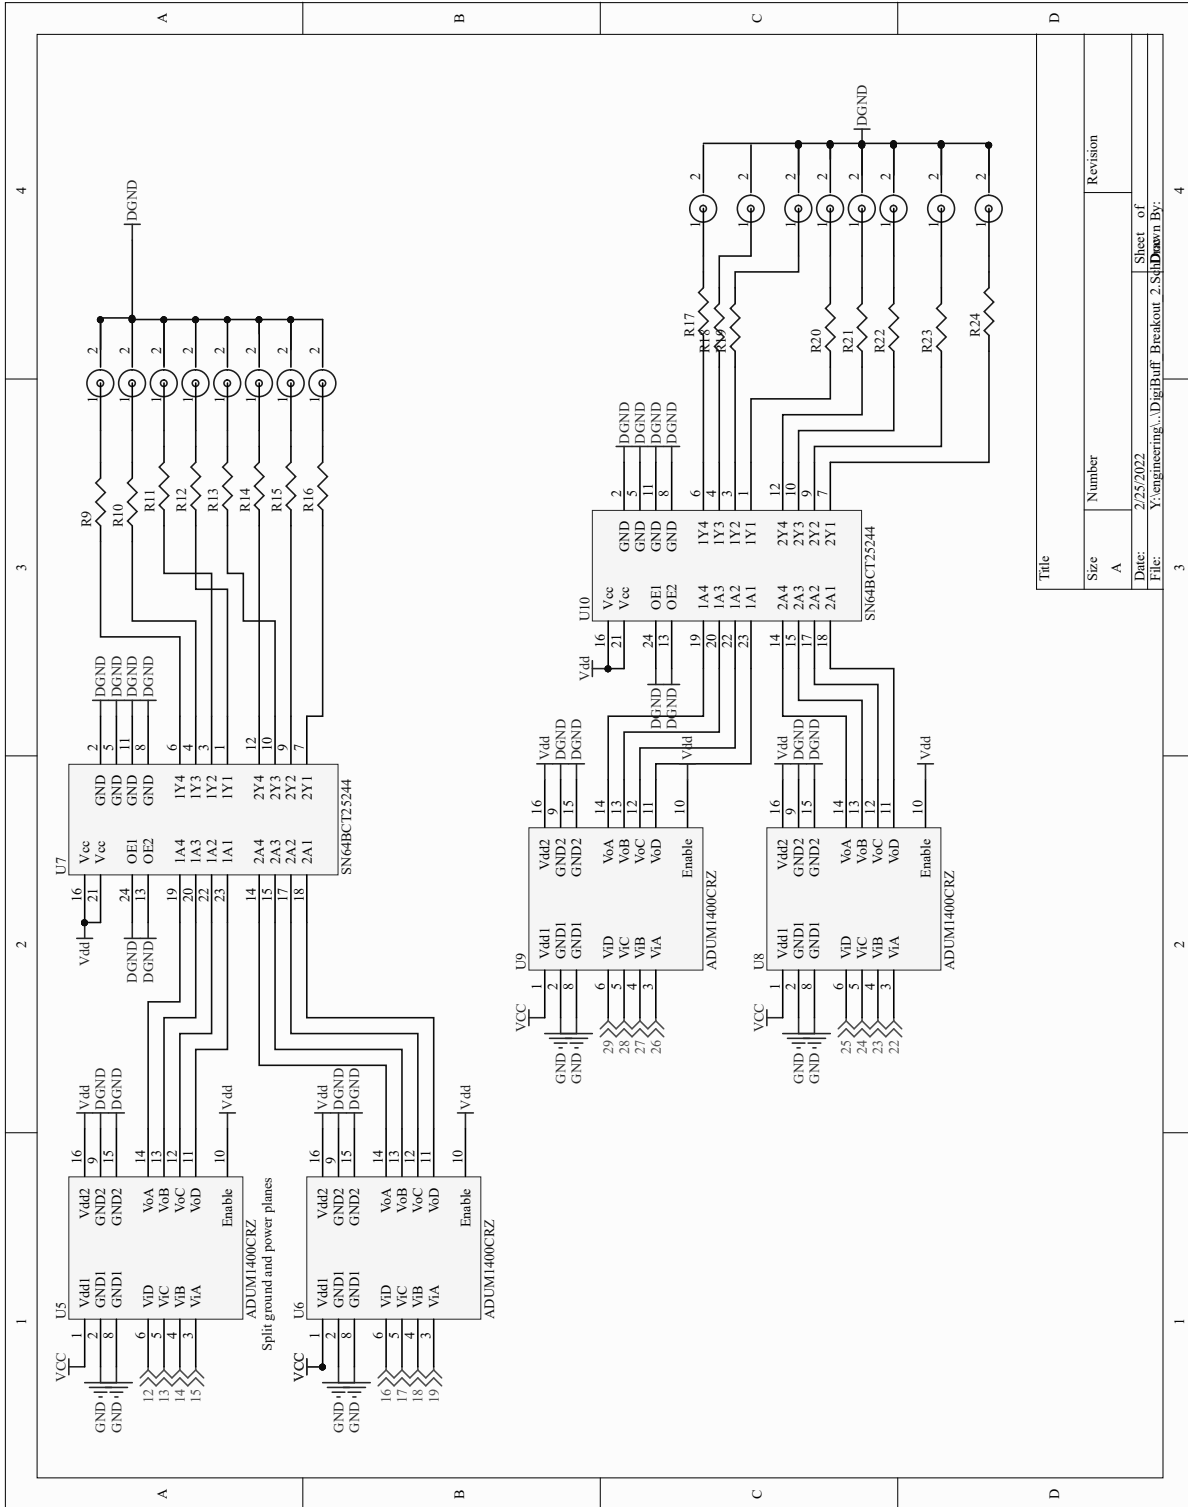
Appendix D

Digital buffer schematic

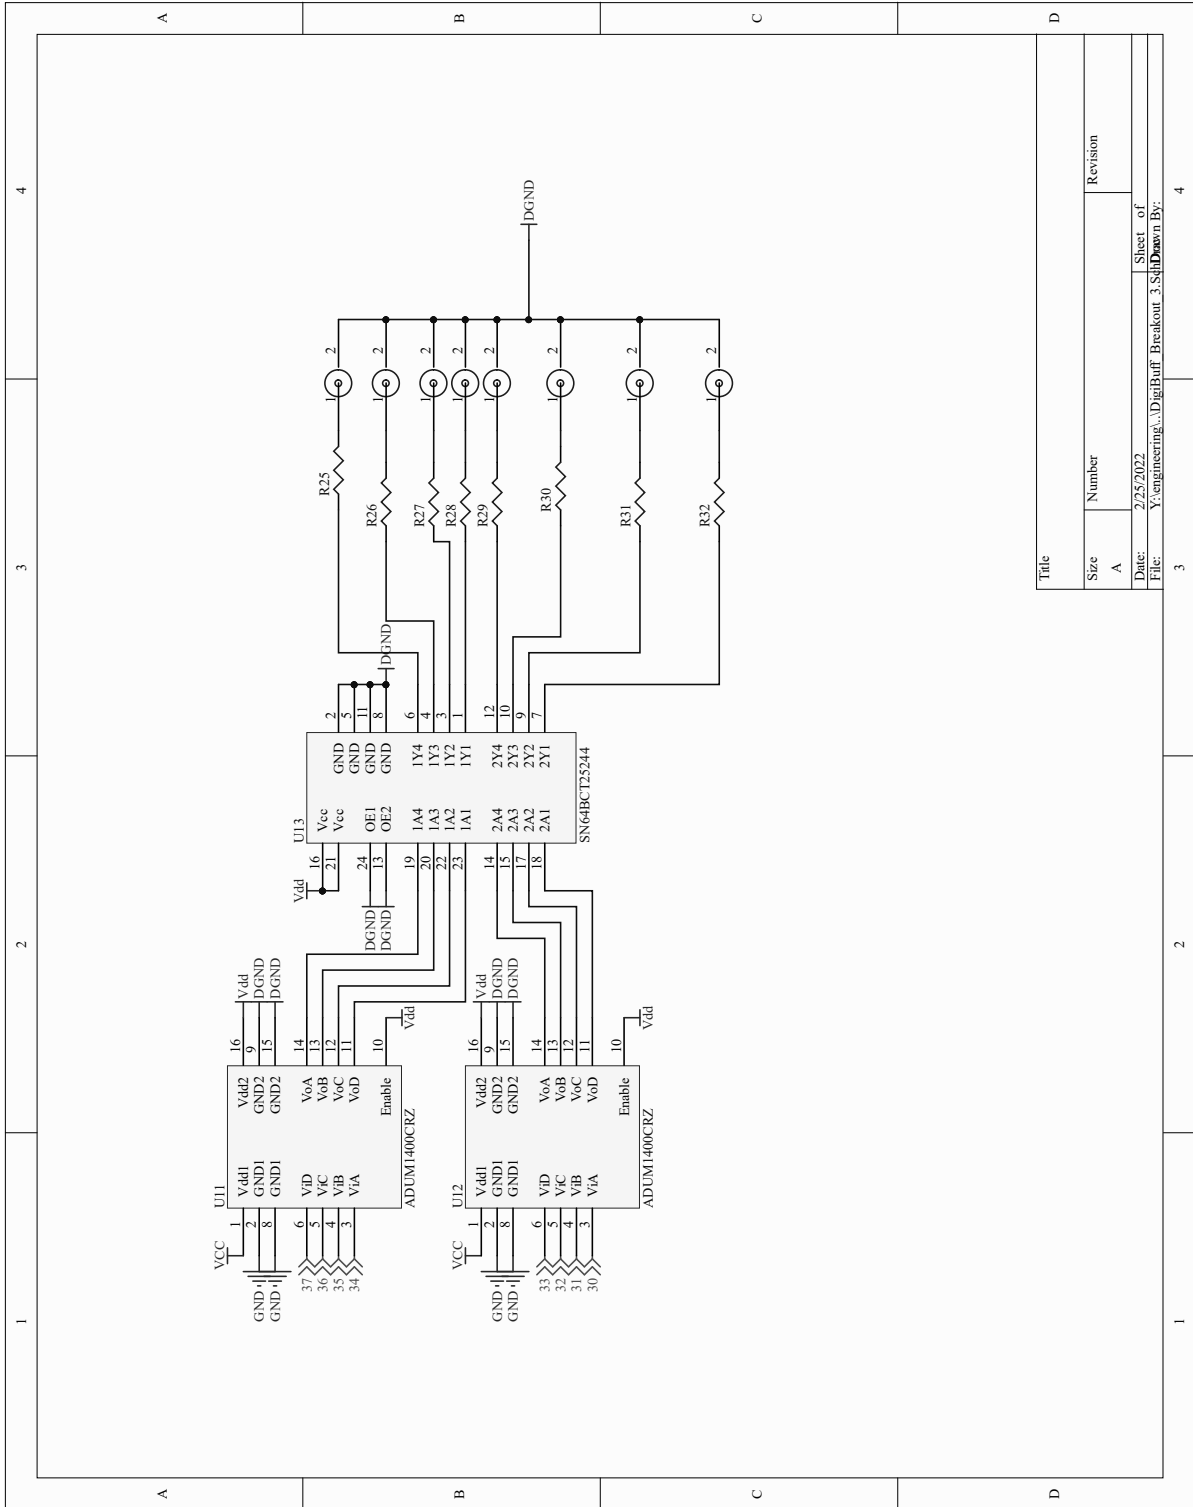
The digital buffers have performed without fail since 2015, with the minor issue of the fuse F1 failing a few times. It seemed to fail at random, so I removed it entirely and jumped that connection.




Title	
Size	Number
A	
Date:	Revision
2/25/2022	
File:	Sheet of
Y:\engineering\...\DigitalBuff_Breakout_Sch.Dwg	1
Drawn By:	Sheet of
	4



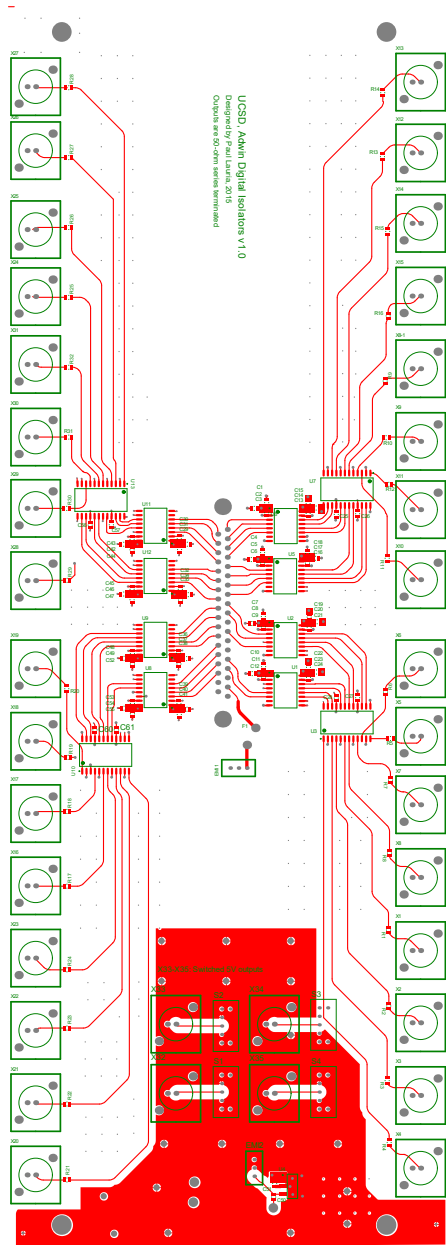
Title	
Size	Number
A	
Revision	
Date:	2/25/2022
File:	Y:\engineering\...\DigiBuff\Breakout_2_SplitDown.rvt
Sheet of	
3	
Drawn By:	




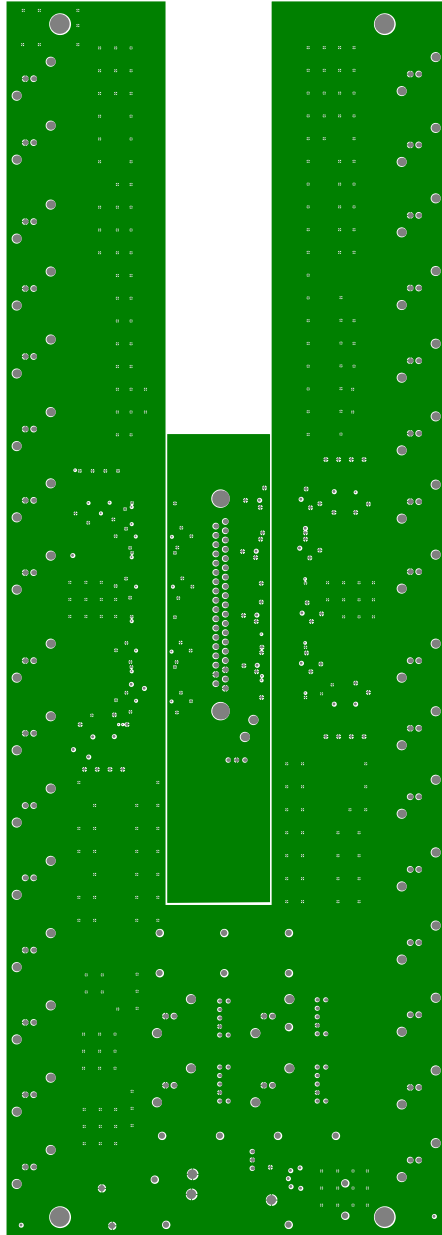
Title	
Size	Revision
A	
Date:	2/25/2022
File:	Y:\engineering\...\DigiBuff_Breakout_3.SchDrawn By:
3	4

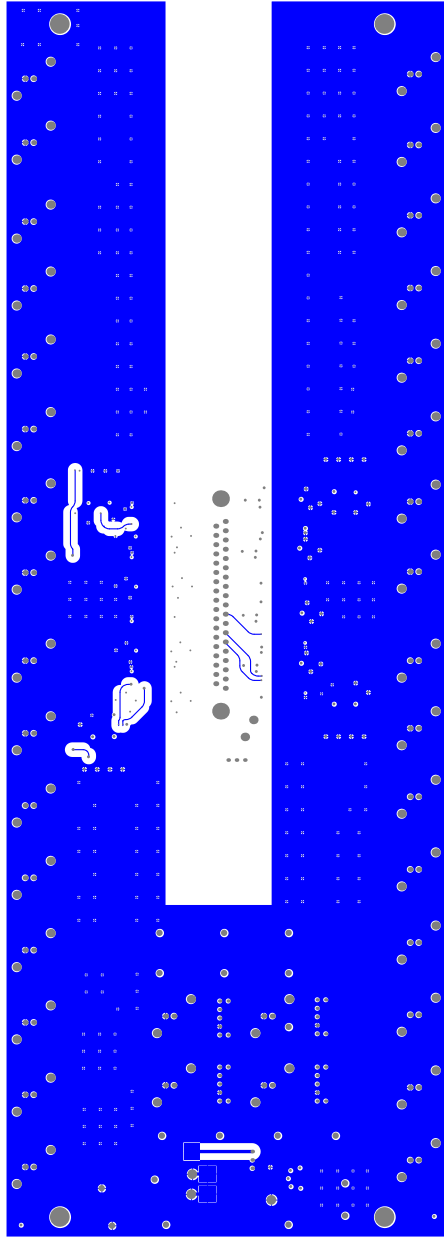


Title		Revision	
Size	A	Number	
Date:	2/25/2022	Sheet of	
File:	Y:\engineering\SupplyBypassing.SchDk1.dwg		
		3	4









Appendix E

Rotating wave approximation

I made the rotating-wave approximation (RWA) in Eqn. 2.4 when describing the dynamics of the 2-level system. It is insightful to consider an alternative derivation that explicitly demonstrates the geometry behind the name, following [59]. I write the dipole approximation as $H_{AF} = -\mathbf{d} \cdot \mathbf{E} = -\mathbf{d} \cdot \hat{\mathbf{e}} E_0 \cos(\omega t)$. Then, as in Eqn. 2.5 but before taking the RWA, I take the two parts of the Hamiltonian H_A and H_{AF} in terms of Pauli matrices as

$$\begin{aligned} H_A &= \frac{-\hbar\omega}{2} \mathbb{1} + \frac{\hbar\omega_0}{2} \sigma_z \\ H_{AF} &= \hbar\Omega \cos(\omega_L t) \sigma_x \end{aligned} \tag{E.1}$$

The Hamiltonian $H = H_A + H_{AF}$ may be thought of as describing a spin-1/2 particle interacting with a crossed magnetic field, with components $\mathbf{B}_0 = \hat{z}\hbar\omega_0/2$ and $\mathbf{B}_1 = \hat{x}\hbar\Omega \cos \omega_L t$ as in Fig. E.1(a). The field \mathbf{B}_1 can be decomposed into two components rotating in opposite directions around the \hat{z} -axis. The counter-clockwise component of \mathbf{B}_1 rotates in the direction of the Larmor spin precession and so can act efficiently on it; the clockwise component rotates at relative frequency $-2\omega_L$, too rapid to have appreciable effect. Dropping this frequency is precisely the rotating-wave approximation.

Taking this a step further, I move into a frame rotating at ω_L as in Fig. E.1(b); we

see that \mathbf{B}_1 has become time-independent, with amplitude reduced by 1/2. The axial field \mathbf{B}_0 sees its field reduced as well, because the Larmor spin precession about \hat{z}' is reduced from ω_0 to $\omega_0 - \omega_L$. Thus, in the rotating frame the spin-1/2 particle experiences a net effective magnetic field \mathbf{B}_{Eff} about which it precesses at frequency

$$\Omega_{\text{Eff}} = \sqrt{\Omega^2 + \delta^2} \quad (\text{E.2})$$

which is otherwise known as the *generalized Rabi frequency* for some laser detuning $\delta \equiv \omega_0 - \omega_L$.




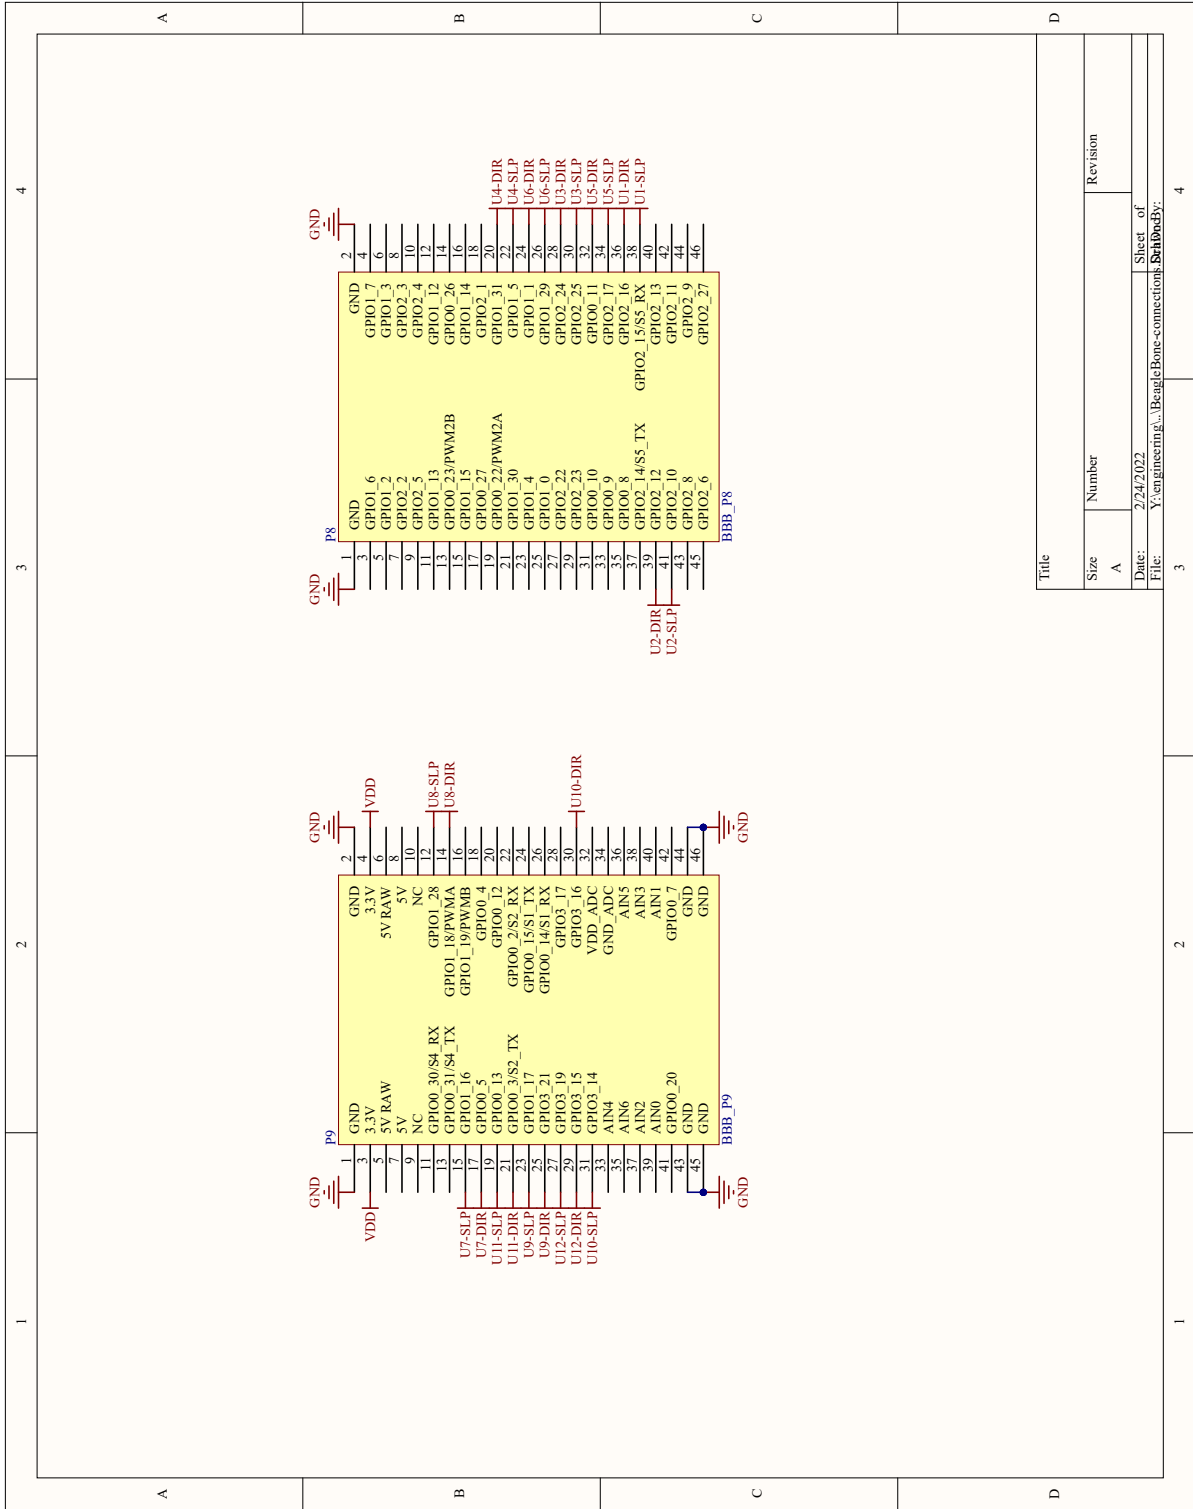
Figure E.1. The rotating wave approximation.

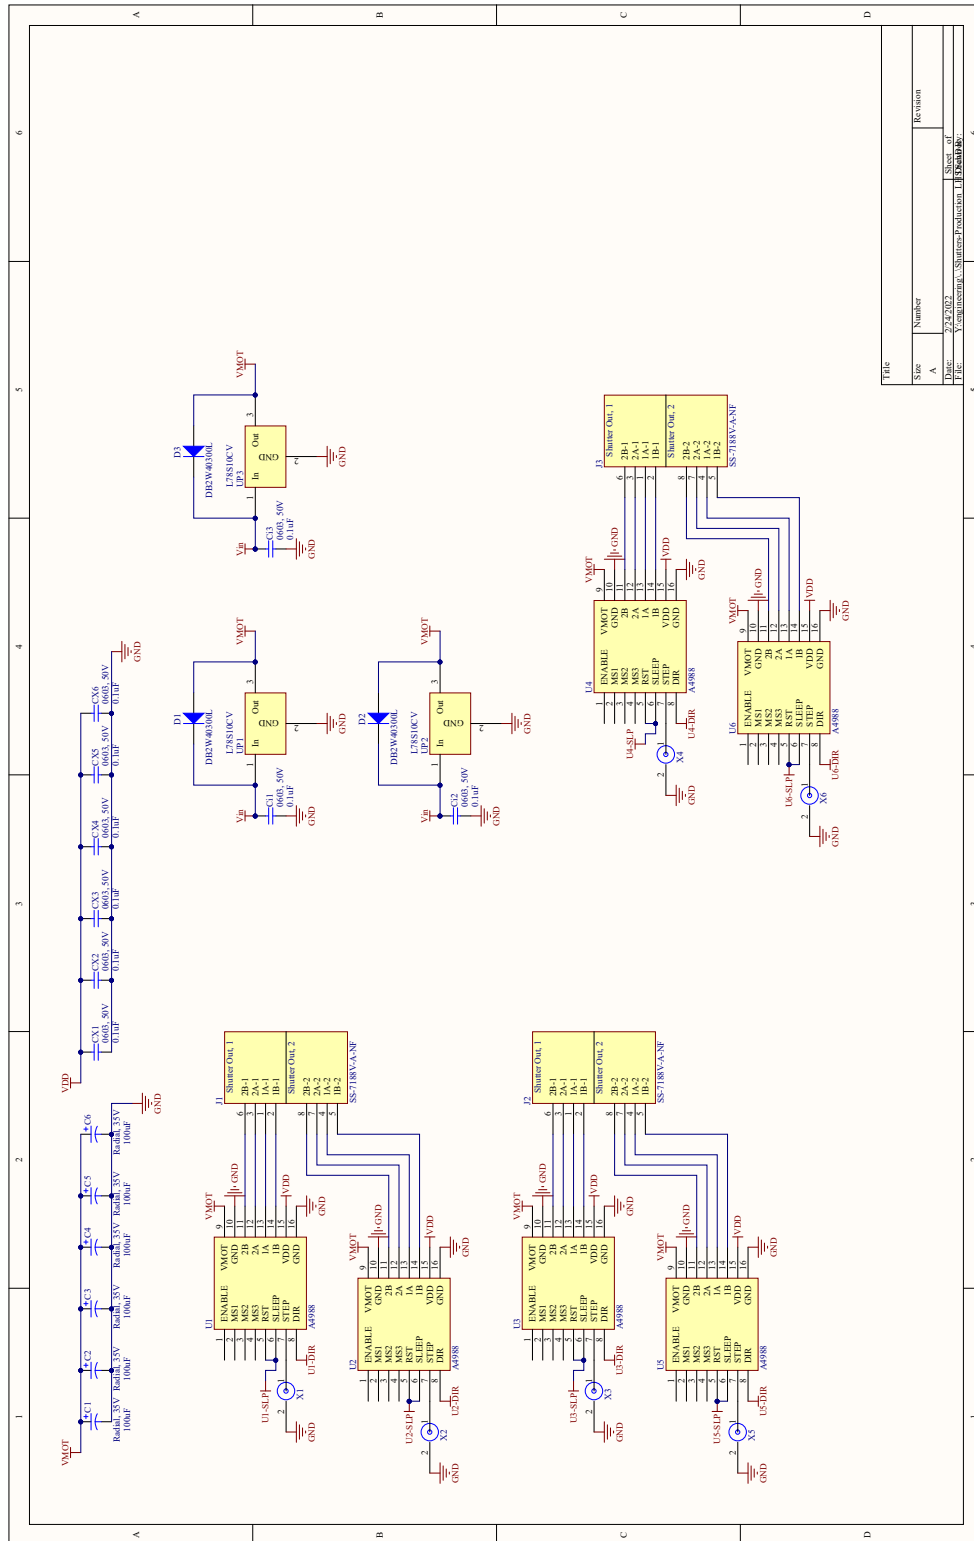
Appendix F

Optical shutters

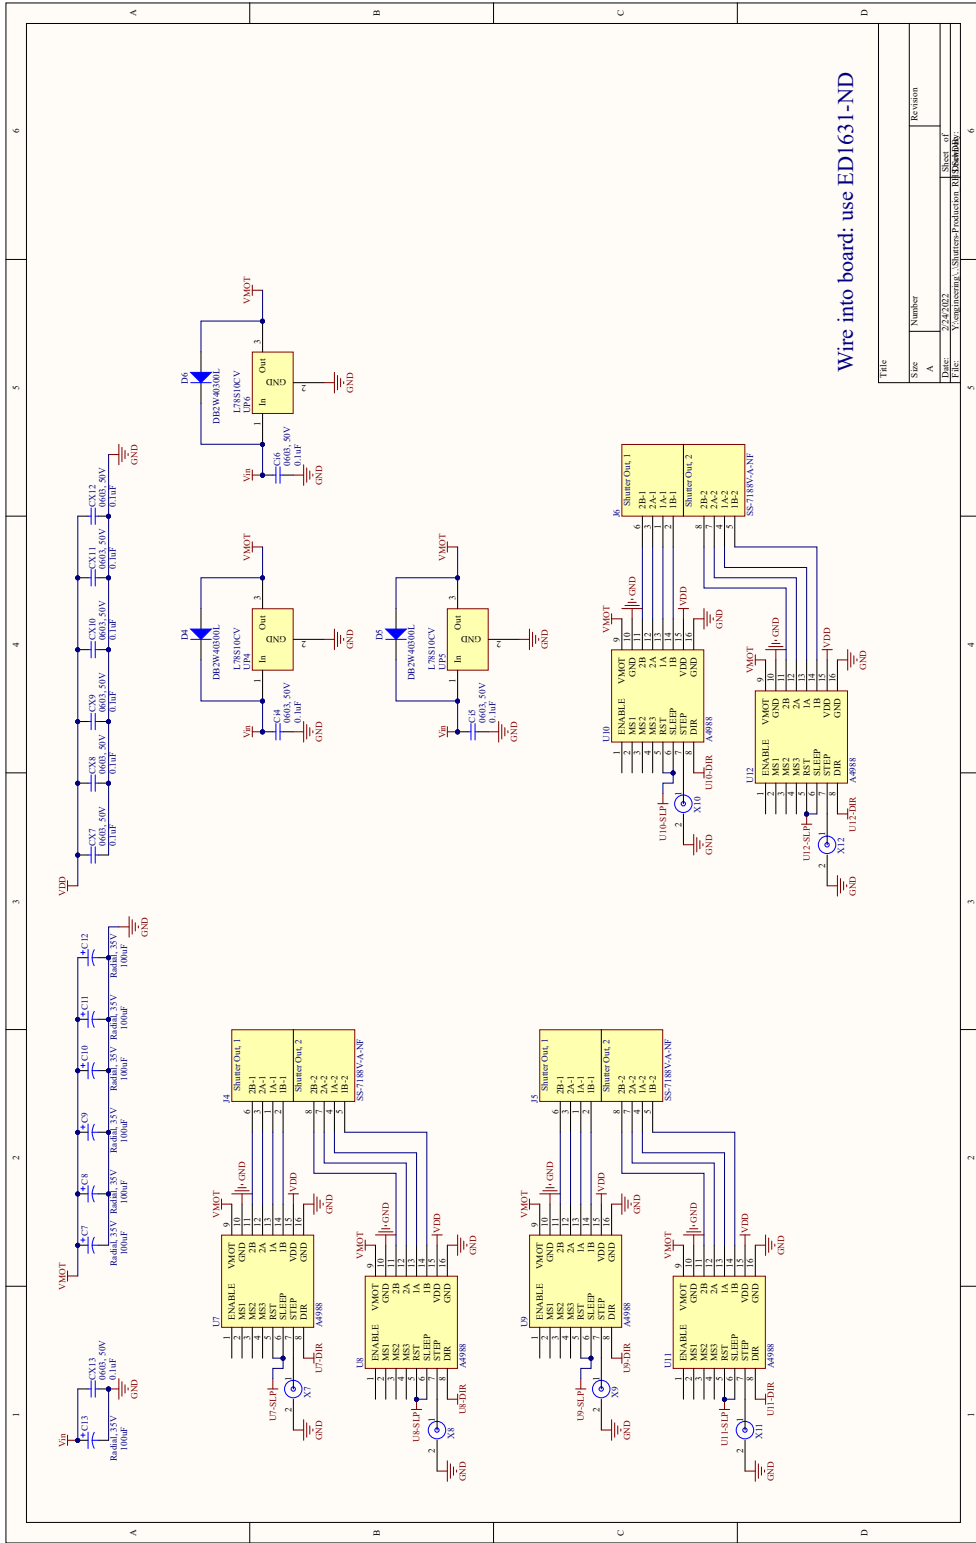
F.1 Shutter PCB


The shutters are controlled by the PCB in this section; for more details, see Section 3.4.8.

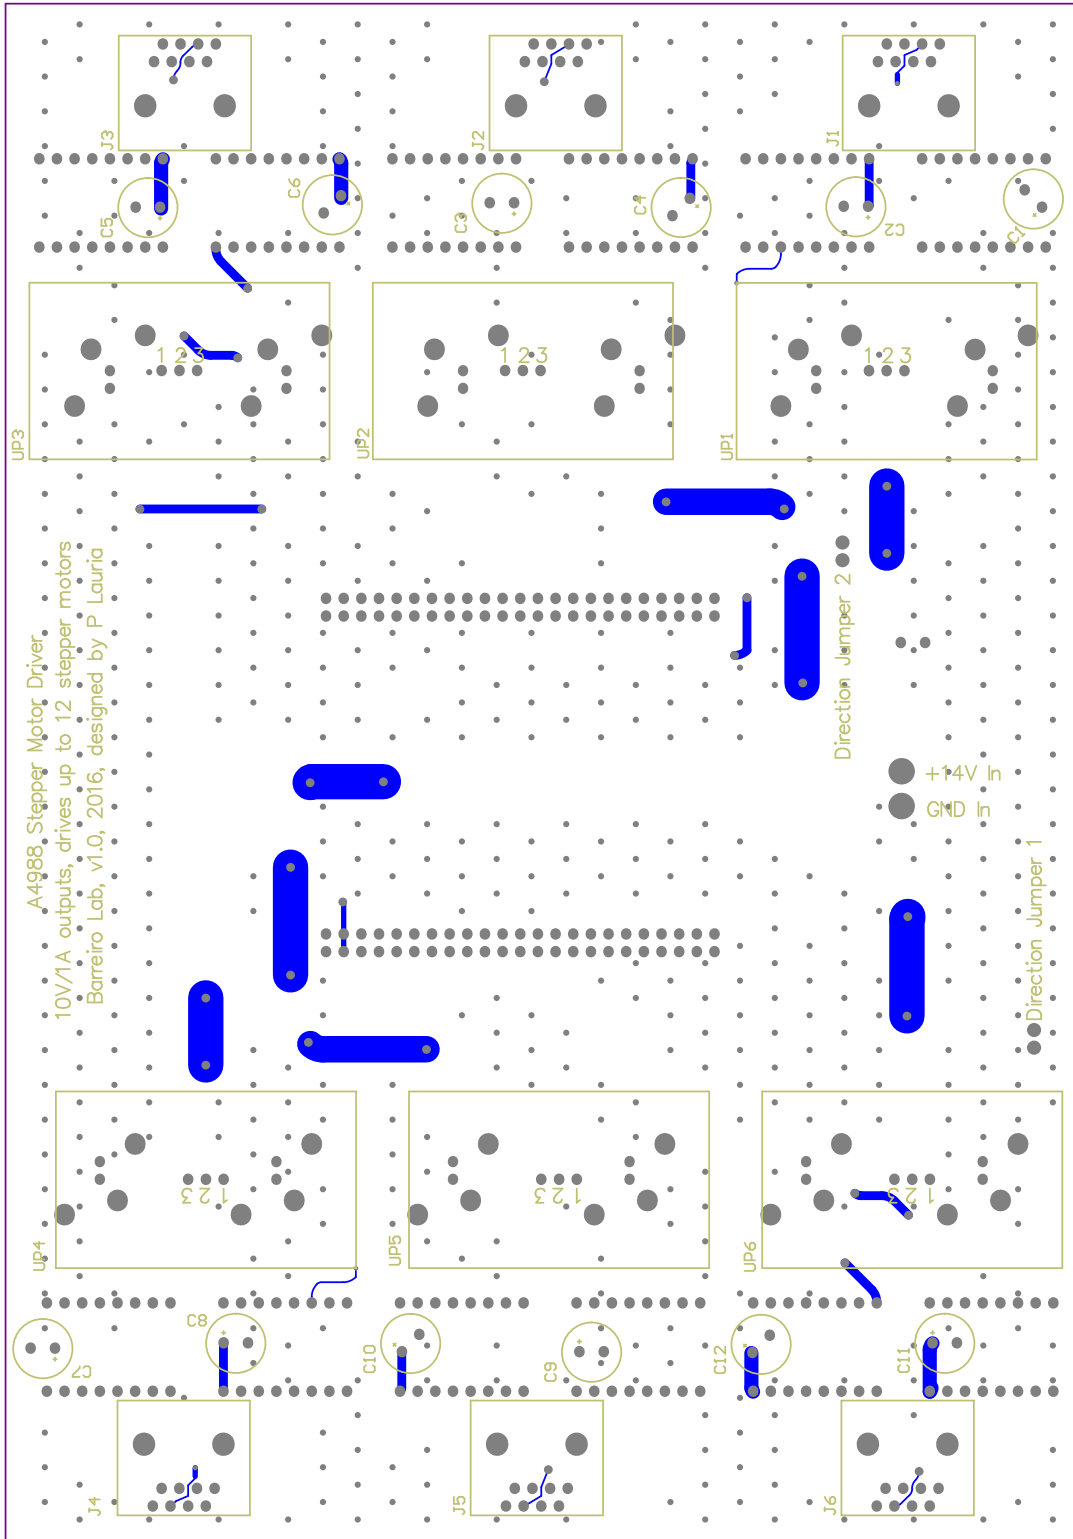




Title	
Size	Number
A	Revision
File	Sheet
C:\Users\user\Documents\Siemens Production\11\SS7188V-A-NF	







F.2 Shutters UART server code

The following has been tested to work on the now-ancient Debian 7.9, using kernel 3.8.13-bone79. My brief attempt at using later kernels and Debian 8.6 was unsuccessful - the UART wasn't working, but I suspect pinmuxing issues with the overlay that might warrant investigation. It is doable, but I haven't found the need (or time) to upgrade.

The core of the `uartserver` is written in C, and essentially just waits (using a kernel event handler) for a UART command, which then updates the memory-mapped GPIO registers. Sending 24 bits (corresponding to 12 servos' sleep and direction bits) to the controlling Beaglebone sets the state of 12 shutters within < 0.5 ms from the end of the UART command. Much credit to [164], chapter 8 in particular, for illuminating the BeagleBone's internals. The GPIO states are accessible via the kernel's memory space. By directly accessing this memory, you can toggle the GPIOs very quickly—much faster than accessing them through `sysfs` (the method of e.g. the `adafruit python` library). You simply need to know the locations in memory of the various GPIOs. Refer to Table 2-3 in the AM335x Technical Reference Manual for the base addresses, and Table 25-5 for the offsets.

In the code, the 'front panel' refers to the shutter front panel, as described in Fig. F.1.

```

/* uartserver.h

* Below defines a mapping between the shutters' front panel numbering
  and
* the GPIO NO. to which sleep and direction correspond on the
  beaglebone.

*
* For example,
*
* #define SLP7 5
*
* Means the sleep pin of Shutter #7 (as marked on the front panel)
  corresponds to GPIO 5
* in the sysfs file mapping - e.g. you'd find the device value file at
*
* /sys/class/gpio/gpio5/value
*
* into which you could echo a 0 or 1 to raise the pin high or low.
*
* gpioX[Y]: X specifies the GPIO bank, of which there are 4, which is
  needed to find the correct register.
*
*         Y specifies the actual bit location in the register which
  corresponds to the specific GPIO.
*
* The correspondance table can be found from Derek Molloy's work online
  , the P8/P9 header tables.
*/

#define DIR1 "112" /* P9_30, gpio3[16]*/
#define SLP1 "110" /* P9_31, gpio3[14]*/
#define DIR2 "111" /* P9_29, gpio3[15]*/
#define SLP2 "115" /* P9_27, gpio3[19]*/
#define DIR3 "117" /* P9_25, gpio3[21]*/
#define SLP3 "49" /* P9_23, gpio1[17]*/

```

```

#define DIR4 "3" /* P9_21, gpio0[3] */
#define SLP4 "4" /* P9_18, gpio0[4] NOTE: on the pcb, soldered is
    P9_19 which is not usable, jump it to P9_18 */
#define DIR5 "5" /* P9_17, gpio0[5] */
#define SLP5 "48" /* P9_15, gpio1[16]*/
#define DIR6 "50" /* P9_14, gpio1[18]*/
#define SLP6 "60" /* P9_12, gpio1[28]*/
#define DIR7 "76" /* P8_39, gpio2[12]*/
#define SLP7 "74" /* P8_41, gpio2[10]*/
#define DIR8 "80" /* P8_36, gpio2[16]*/
#define SLP8 "79" /* P8_38, gpio2[15]*/
#define DIR9 "11" /* P8_32, gpio0[11]*/
#define SLP9 "81" /* P8_34, gpio2[17]*/
#define DIR10 "88" /* P8_28, gpio2[24]*/
#define SLP10 "89" /* P8_30, gpio2[25]*/
#define DIR11 "33" /* P8_24, gpio1[1] */
#define SLP11 "61" /* P8_26, gpio1[29]*/
#define DIR12 "63" /* P8_20, gpio1[31]*/
#define SLP12 "37" /* P8_22, gpio1[5] */

/* Bitmasks for the sent Adwin command. D is direction, S is sleep.
 *
 * A command is a 24-bit number. (command >> DX )&1 gets the
 * direction of shutter X, and (command >> SX)&1 gets the sleep
 * Status of shutter X.
 *
 * No suffix, e.g. D1, denotes the location of the bit
 * in the AdWin UART command.
 *
 * _r denotes the register location of the GPIO.
 *
 */

```



```
#define D6 0
#define D6_r 18
#define S6 1
#define S6_r 28
#define D5 2
#define D5_r 5
#define S5 3
#define S5_r 16
#define D4 4
#define D4_r 3
#define S4 5
#define S4_r 4
#define D3 6
#define D3_r 21
#define S3 7
#define S3_r 17
#define D2 8
#define D2_r 15
#define S2 9
#define S2_r 19
#define D1 10
#define D1_r 16
#define S1 11
#define S1_r 14
#define D12 12
#define D12_r 31
#define S12 13
#define S12_r 5
#define D11 14
#define D11_r 1
#define S11 15
```

```

#define S11_r 29
#define D10 16
#define D10_r 24
#define S10 17
#define S10_r 25
#define D9 18
#define D9_r 11
#define S9 19
#define S9_r 17
#define D8 20
#define D8_r 16
#define S8 21
#define S8_r 15
#define D7 22
#define D7_r 12
#define S7 23
#define S7_r 10

/* Register base addresses for each GPIO bank. */
/* See the AM335x Technical Reference Manual, Table 2-3. */

#define GPIO0 0x44E07000
#define GPIO1 0x4804C000
#define GPIO2 0x481AC000
#define GPIO3 0x481AE000

/* Register offsets to set and clear bits.
 * See AM335x TRM reference Table 25-5. */
#define CLEAR 0x190
#define SET 0x194

/* Each GPIO bank is exactly 4kBytes == 0xFFF */

```

```

#define MAP_SIZE 0xFF

/* This is the fastest rate the admin can support.
 * It corresponds to 1/115200 ~ 8.6uS pulses */
#define BAUDRATE B115200

/* UART 1 on P9_26 */
#define TTY "/dev/tty01"

#define GPIO_EXPORT_PATH "/sys/class/gpio/export"
#define GPIO_UNEXPORT_PATH "/sys/class/gpio/unexport"
#define GPIO_PATH "/sys/class/gpio/gpio"

#define MAX_PATH_SIZE 10000

#define FATAL do { fprintf(stderr, "Error at line %d, file %s (%d) [%s]\n", \
    __LINE__, __FILE__, errno, strerror(errno)); exit(1); } while(0)

#define UART_COMMAND_BYTES 3

/* The below definitions are just used for setting up the GPIOs
 * in sysfs */

char * GPIOs[] = { DIR1, SLP1,
    DIR2, SLP2,
    DIR3, SLP3,
    DIR4, SLP4,
    DIR5, SLP5,
    DIR6, SLP6,

```

```

    DIR7 , SLP7 ,
    DIR8 , SLP8 ,
    DIR9 , SLP9 ,
    DIR10 ,SLP10 ,
    DIR11 ,SLP11 ,
    DIR12 ,SLP12}];

/* File handles to sleep and direction pins - these are unused
 * in the latest versions of uart-interrupt, since I use registers
 * instead, but keep this for backwards compat. for debugging. */
int fD1, fS1;
int fD2, fS2;
int fD3, fS3;
int fD4, fS4;
int fD5, fS5;
int fD6, fS6;
int fD7, fS7;
int fD8, fS8;
int fD9, fS9;
int fD10, fS10;
int fD11, fS11;
int fD12, fS12;

int *fileHandles [] = { &fD1, &fS1,
                        &fD2, &fS2,
                        &fD3, &fS3,
                        &fD4, &fS4,
                        &fD5, &fS5,
                        &fD6, &fS6,
                        &fD7, &fS7,
                        &fD8, &fS8,
                        &fD9, &fS9,

```

```

        &fD10,&fS10,
        &fD11,&fS11,
        &fD12,&fS12});

/* Pointers to the register banks and clr/set registers */
void *gpio0, *gpio0_clr, *gpio0_set;
void *gpio1, *gpio1_clr, *gpio1_set;
void *gpio2, *gpio2_clr, *gpio2_set;
void *gpio3, *gpio3_clr, *gpio3_set;

static volatile int keepRunning = 1;

void UART_command_received (int status);
int unexport_gpio(char * gpio_num);
int export_gpio(char * gpio_num);
void intHandler(int dummy);
void printBinary(int num);

int memfd;
unsigned long read_result, writeval;
int fd;

/* uartserver.c
 *
 * The main UART server. Reacts to serial UART commands, which set the
 * shutter sleep and
 * direction pins.
 * Sets up all sysfs files for all the GPIOs used on shutter box.
 * On receipt of UART1 signal, toggles GPIOs according to the UART
 * command,
 * using direct memory access to the BeagleBone's CPU (AM335x) registers
 * . */

```

```

#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <unistd.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <fcntl.h>
#include <sys/signal.h>
#include <errno.h>
#include <termios.h>
#include <signal.h>
#include <sys/types.h>
#include <sys/mman.h>
#include <limits.h>
#include "uart-interrupt.h"

// Uncomment the below #define to enter debug mode, in which
// the UART commands are printed to the terminal for inspection.
// Note this adds a LOT of latency. Don't use in experiments.
/* #define DEBUG 0 */

/* Handle CTRL+C quitting gracefully */
void intHandler(int dummy) {
    printf("Caught CTRLC, terminating. \n");

    /* Not strictly necessary, as memory is unmapped on termination */
    //close(fd);

    /* if(munmap(gpio0, MAP_SIZE) == -1) FATAL; */
    /* if(munmap(gpio1, MAP_SIZE) == -1) FATAL; */

```

```

/* if(munmap(gpio2, MAP_SIZE) == -1) FATAL; */
/* if(munmap(gpio3, MAP_SIZE) == -1) FATAL; */
/* close(memfd); */

int i;

/* Unexport all gpios, close all opened file handles. */
for(i=0; i < sizeof(GPIOs)/sizeof(GPIOs[0]); i++)
{
    unexport_gpio(GPIOs[i]);
    close(*fileHandles[i]);
}
keepRunning = 0;
exit(0);
}

/* Just to aid in debugging */
void printBinary(int n)
{
    int c, k;
    for (c = 31; c >= 0; c--)
    {
        k = n >> c;

        if (k & 1)
printf("1");
        else
printf("0");
        if(c%4==0)
printf(" ");
    }
}

```

```

printf("\n");
}

int export_gpio(char * gpio_num)
{
    int fd_export,fd_dir, writeresult;

    /* First export the GPIO */
    if ( (fd_export = open(GPIO_EXPORT_PATH, O_WRONLY)) <0) {
        perror("Failed to open GPIO export file, wtf?");
        printf("%s \n",gpio_num);
        exit(1);
    }

    if((writeresult = write(fd_export, gpio_num, sizeof(gpio_num)))<0) {
        perror("Failed to write to output");
        printf("%s \n",gpio_num);
    }
    else {
        printf("Successfully exported GPIO %s \n", gpio_num);
    }

    close(fd_export);

    /* Now set it to output-mode */
    char dirpath[MAX_PATH_SIZE] = {0};
    snprintf(dirpath, sizeof(dirpath), "%s%s/direction",GPIO_PATH, gpio_num)
        ;

    printf("configured dirpath = %s \n",dirpath);

    fd_dir = open(dirpath,O_WRONLY);

```



```

char * direction = "high";
write(fd_dir, direction, sizeof(direction));
close(fd_dir);

/* have a handle to 'value' open at all times */
char gpiopath[MAX_PATH_SIZE] = {0};
snprintf(gpiopath, sizeof(gpiopath), "%s%s/value", GPIO_PATH, gpio_num);
printf("configured giopath = %s \n", gpiopath);
return open(gpiopath, O_WRONLY); // return an open file handle for
    later use
}

int unexport_gpio(char * gpio_num) {
    /* Unexports the GPIO */
    int fd_unexport, res;
    if ( (fd_unexport = open(GPIO_UNEXPORT_PATH, O_WRONLY)) <0) {
        perror("Failed to open GPIO unexport file, wtf?");
        printf("%s \n", gpio_num);
        exit(1);
    }

    if ( (res = write(fd_unexport, gpio_num, sizeof(gpio_num))) <0) {
        perror("Failed to unexport GPIO:");
        printf("Failed unexporting GPIO # %s \n", gpio_num);
    }
}

int main(int argc, char *argv[])
{

```

```

/* Catch CTRL+C */
signal(SIGINT, intHandler);

/* Open file handles to all the GPIOs we're going to
   be using. */
size_t i;
for(i=0; i < sizeof(GPIOs)/sizeof(GPIOs[0]); i++)
{
    printf("GPIOs[i] = %s \n", GPIOs[i]);
    *fileHandles[i] = export_gpio(GPIOs[i]);
}

/* Open the system's memory table */
if((memfd = open("/dev/mem", O_RDWR | O_SYNC)) == -1) FATAL;
printf("/dev/mem opened.\n");
fflush(stdout);

/* Map GPIO address spaces */
gpio0 = mmap(0, MAP_SIZE, PROT_READ | PROT_WRITE, MAP_SHARED, memfd,
    GPIO0);
gpio1 = mmap(0, MAP_SIZE, PROT_READ | PROT_WRITE, MAP_SHARED, memfd,
    GPIO1);
gpio2 = mmap(0, MAP_SIZE, PROT_READ | PROT_WRITE, MAP_SHARED, memfd,
    GPIO2);
gpio3 = mmap(0, MAP_SIZE, PROT_READ | PROT_WRITE, MAP_SHARED, memfd,
    GPIO3);
if(gpio0 == (void *) -1) FATAL;
if(gpio1 == (void *) -1) FATAL;
if(gpio2 == (void *) -1) FATAL;
if(gpio3 == (void *) -1) FATAL;
fflush(stdout);

```

```

/* Compute the memory locations of the clear and set registers */
gpio0_clr = gpio0 + CLEAR;
gpio0_set = gpio0 + SET;
gpio1_clr = gpio1 + CLEAR;
gpio1_set = gpio1 + SET;
gpio2_clr = gpio2 + CLEAR;
gpio2_set = gpio2 + SET;
gpio3_clr = gpio3 + CLEAR;
gpio3_set = gpio3 + SET;

/* Open the UART Rx port - we use UART1 on the BeagleBone, Rx is pin
P9_26 */
fd = open(TTY, O_RDONLY | O_NOCTTY);
if (fd == -1)
{
    perror("open_port: Unable to open /dev/tty01\n");
    exit(1);
}

/* Configure some default serial port settings. */
struct termios termAttr;
struct sigaction saio;

saio.sa_handler = UART_command_received; // This function is called
    when a UART signal is rec'd
saio.sa_flags = 0;
saio.sa_restorer = NULL;
sigaction(SIGIO, &saio, NULL);
fcntl(fd, F_SETFL, FASYNC);
fcntl(fd, F_SETOWN, getpid());
tcgetattr(fd, &termAttr);
cfsetispeed(&termAttr, BAUDRATE);

```

```

cfsetospeed(&termAttr,BAUDRATE);
// termAttr.c_cc[VMIN]=sizeof(int);
termAttr.c_cc[VMIN]=UART_COMMAND_BYTES; // minimum number of bytes per
    read() - will block until satisfied
termAttr.c_cc[VTIME]=1;// time to wait in deciseconds, 0 = infinite
termAttr.c_cflag &= ~PARENB; // disable parity check
termAttr.c_cflag &= ~CSTOPB; // one stop bit
termAttr.c_cflag &= ~CSIZE;
termAttr.c_cflag |= CS8;
termAttr.c_cflag |= (CLOCAL | CREAD); // ignore modem control lines,
    enable receiver
termAttr.c_lflag &= ~(ICANON | ECHO | ECHOE | ISIG); // noncanonical
    mode
termAttr.c_iflag &= ~(IXON | IXOFF | IXANY);
termAttr.c_oflag &= ~OPOST;
tcsetattr(fd,TCSANOW,&termAttr);
printf("UART1 configured. Server is running.\n");

while(keepRunning == 1){/* Loop forever until CTRL+C */}

printf("Terminating...\n");

}

int UART_command;
size_t szCommand = sizeof(UART_command);

/* Event handler called when a UART string is received.
 * */
void UART_command_received (int status)
{

```

```

#ifdef DEBUG
printf("UART command received; waiting on read()...\n");
#endif

/* This sysfs read() adds ~40 us measured latency - negligble at AdWin
's baud rate. */
/* read(fd,&UART_command,szCommand); // TODO: can we replace szCommand
with UART_COMMAND_BYTES? */
read(fd,&UART_command,UART_COMMAND_BYTES);

#ifdef DEBUG
printf("Command is: 0x%x, dec %d\n",UART_command,UART_command);
printf("Command in binary: ");
printBinary(UART_command);
#endif

/* Parse the UART_command: get the relevant bits from the UART command
* and set the corresponding registers.
*
* See uart-interrupt.h for details.
*
* The 0xFFFFF^... in the various *gpioX_clr inverts the sent command
,
* because you counter-intuitively write a 1 to actually clear the bit
(see the AM335x TRM) */

*((unsigned long *) gpio0_clr) = (((((0xFFFFF^UART_command)>>D4)&1)
<< D4_r) +
    (((0xFFFFF^UART_command)>>S4)&1) << S4_r) +
    (((0xFFFFF^UART_command)>>D5)&1) << D5_r) +
    (((0xFFFFF^UART_command)>>D9)&1) << D9_r));

*((unsigned long *) gpio0_set) = (((UART_command>>D4)&1) << D4_r) +

```

```

        (((UART_command>>S4)&1) << S4_r) +
        (((UART_command>>D5)&1) << D5_r) +
        (((UART_command>>D9)&1) << D9_r));
#ifdef DEBUG
int cmd = (((UART_command>>D4)&1) << D4_r) +
        (((UART_command>>S4)&1) << S4_r) +
        (((UART_command>>D5)&1) << D5_r) +
        (((UART_command>>D9)&1) << D9_r));
printf("gpio0 hex: %x\n",cmd);
printBinary(cmd);
#endif
*((unsigned long *) gpio1_clr) = (((((0xFFFFFFFF^UART_command)>>S3)&1)
<< S3_r) +
        (((0xFFFFFFFF^UART_command)>>S5)&1) << S5_r) +
        (((0xFFFFFFFF^UART_command)>>D6)&1) << D6_r) +
        (((0xFFFFFFFF^UART_command)>>S6)&1) << S6_r) +
        (((0xFFFFFFFF^UART_command)>>D11)&1) << D11_r) +
        (((0xFFFFFFFF^UART_command)>>S11)&1) << S11_r) +
        (((0xFFFFFFFF^UART_command)>>D12)&1) << D12_r) +
        (((0xFFFFFFFF^UART_command)>>S12)&1) << S12_r));
*((unsigned long *) gpio1_set) = (((UART_command>>S3)&1) << S3_r) +
        (((UART_command>>S5)&1) << S5_r) +
        (((UART_command>>D6)&1) << D6_r) +
        (((UART_command>>S6)&1) << S6_r) +
        (((UART_command>>D11)&1) << D11_r) +
        (((UART_command>>S11)&1) << S11_r) +
        (((UART_command>>D12)&1) << D12_r) +
        (((UART_command>>S12)&1) << S12_r));
#ifdef DEBUG
cmd = (((UART_command>>S3)&1) << S3_r) +
        (((UART_command>>S5)&1) << S5_r) +
        (((UART_command>>D6)&1) << D6_r) +

```

```

(((UART_command>>S6)&1) << S6_r) +
(((UART_command>>D11)&1) << D11_r) +
(((UART_command>>S11)&1) << S11_r) +
(((UART_command>>D12)&1) << D12_r) +
(((UART_command>>S12)&1) << S12_r));

printf("gpio1: %x\n", cmd);
printBinary(cmd);
#endif
*((unsigned long *) gpio2_clr) = (((((0xFFFFFFFF^UART_command)>>D9)&1)
<< D9_r) +
    (((0xFFFFFFFF^UART_command)>>S9)&1) << S9_r) +
    (((0xFFFFFFFF^UART_command)>>S10)&1) << S10_r) +
    (((0xFFFFFFFF^UART_command)>>D10)&1) << D10_r) +
    (((0xFFFFFFFF^UART_command)>>D7)&1) << D7_r) +
    (((0xFFFFFFFF^UART_command)>>S7)&1) << S7_r) +
    (((0xFFFFFFFF^UART_command)>>D8)&1) << D8_r) +
    (((0xFFFFFFFF^UART_command)>>S8)&1) << S8_r));
*((unsigned long *) gpio2_set) = (((UART_command>>D9)&1) << D9_r) +
    ((UART_command>>S9)&1) << S9_r) +
    ((UART_command>>S10)&1) << S10_r) +
    ((UART_command>>D10)&1) << D10_r) +
    ((UART_command>>D7)&1) << D7_r) +
    ((UART_command>>S7)&1) << S7_r) +
    ((UART_command>>D8)&1) << D8_r) +
    ((UART_command>>S8)&1) << S8_r));
#ifdef DEBUG
cmd = (((UART_command>>D9)&1) << D9_r) +
    (((UART_command>>S9)&1) << S9_r) +
    (((UART_command>>S10)&1) << S10_r) +
    (((UART_command>>D10)&1) << D10_r) +
    (((UART_command>>D7)&1) << D7_r) +

```

```

    (((UART_command>>S7)&1) << S7_r) +
    (((UART_command>>D8)&1) << D8_r) +
    (((UART_command>>S8)&1) << S8_r));
printf("gpio2: %x\n", cmd);
printBinary(cmd);
#endif
*((unsigned long *) gpio3_clr) = (((((0xFFFFFFFF^UART_command)>>D1)&1)
    << D1_r) +
    (((0xFFFFFFFF^UART_command)>>S1)&1) << S1_r) +
    (((0xFFFFFFFF^UART_command)>>D2)&1) << D2_r) +
    (((0xFFFFFFFF^UART_command)>>S2)&1) << S2_r) +
    (((0xFFFFFFFF^UART_command)>>D3)&1) << D3_r));

*((unsigned long *) gpio3_set) = (((UART_command>>D1)&1) << D1_r) +
    (((UART_command>>S1)&1) << S1_r) +
    (((UART_command>>D2)&1) << D2_r) +
    (((UART_command>>S2)&1) << S2_r) +
    (((UART_command>>D3)&1) << D3_r));


#ifdef DEBUG
cmd = (((UART_command>>D1)&1) << D1_r) +
    (((UART_command>>S1)&1) << S1_r) +
    (((UART_command>>D2)&1) << D2_r) +
    (((UART_command>>S2)&1) << S2_r) +
    (((UART_command>>D3)&1) << D3_r));
printf("gpio3: %x\n", cmd);
printBinary(cmd);
printf("-----\n");
#endif
}


```


Appendix G

Switchable gain transimpedance amplifier PCB

Below is the 2-layer PCB art of the electronically-switchable-gain amplifier, whose description and schematic is in Section 3.4.2.





Bibliography

- [1] Simon Stellmer, *Degenerate Quantum Gases of Strontium*, Ph.D. thesis, University of Innsbruck (2013).
- [2] Daniel A. Steck, *Quantum and Atom Optics* (University of Colorado, 2021).
- [3] Christopher J. Foot, *Atomic Physics* (OXFORD UNIV PR, 2005).
- [4] W. Demtroeder, *Laser spectroscopy* (Springer, Berlin, 2008).
- [5] Richard P. Feynman, Robert B. Leighton, Matthew Sands, and E. M. Hafner, The Feynman lectures on physics Vol. I, *American Journal of Physics* **33**, 750–752 (1965).
- [6] E. Schrödinger, An undulatory theory of the mechanics of atoms and molecules, *Phys. Rev.* **28**, 1049–1070 (1926).
- [7] I. M. Georgescu, S. Ashhab, and Franco Nori, Quantum simulation, *Rev. Mod. Phys.* **86**, 153–185 (2014).
- [8] Alexander L. Fetter and John Dirk Walecka, *Quantum Theory of Many-Particle Systems* (Dover Publications Inc., 2003).
- [9] D. J. Thouless, *The quantum mechanics of many-body systems* (Dover Publications, Inc, Mineola, New York, 2014).
- [10] R. O. Jones, Density functional theory: Its origins, rise to prominence, and future, *Rev. Mod. Phys.* **87**, 897–923 (2015).
- [11] Patrick A. Lee, Naoto Nagaosa, and Xiao-Gang Wen, Doping a Mott insulator: Physics of high-temperature superconductivity, *Rev. Mod. Phys.* **78**, 17–85 (2006).
- [12] Anatoli Polkovnikov, Krishnendu Sengupta, Alessandro Silva, and Mukund Vengalattore, Colloquium: Nonequilibrium dynamics of closed interacting quantum systems, *Rev. Mod. Phys.* **83**, 863–883 (2011).

- [13] Leon Balents, Spin liquids in frustrated magnets, *Nature* **464**, 199–208 (2010).
- [14] M. Z. Hasan and C. L. Kane, Colloquium: Topological insulators, *Rev. Mod. Phys.* **82**, 3045–3067 (2010).
- [15] N. R. Cooper, J. Dalibard, and I. B. Spielman, Topological bands for ultracold atoms, *Rev. Mod. Phys.* **91**, 015005 (2019).
- [16] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien, Quantum computers, *Nature* **464**, 45–53 (2010).
- [17] Richard P. Feynman, Simulating physics with computers, *Int. J. of Theo. Phys.* **21**, 467–488 (1982).
- [18] Iulia Buluta and Franco Nori, Quantum simulators, *Science* **326**, 108–111 (2009).
- [19] Martin M. Boyd, *High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard for Frequency and Time*, Ph.D. thesis, University of Colorado (2007).
- [20] Tobias Bothwell, Dhruv Kedar, Eric Oelker, John M Robinson, Sarah L Bromley, Weston L Tew, Jun Ye, and Colin J Kennedy, JILA SrI optical lattice clock with uncertainty of 2.0×10^{-18} , *Metrologia* **56**, 065004 (2019).
- [21] E. Oelker, R. B. Hutson, C. J. Kennedy, L. Sonderhouse, T. Bothwell, A. Goban, D. Kedar, C. Sanner, J. M. Robinson, G. E. Marti, D. G. Matei, T. Legero, M. Giunta, R. Holzwarth, F. Riehle, U. Sterr, and J. Ye, Demonstration of 4.8×10^{-17} stability at 1 s for two independent optical clocks, *Nature Photonics* **13**, 714–719 (2019).
- [22] David R. Leibbrandt, Samuel M. Brewer, Jwo-Sy Chen, Chin-Wen Chou, Aaron M. Hankin, David B. Hume, and David J. Wineland, Trapped-ion optical atomic clocks at the quantum limits, in *The Precise Time and Time Interval Systems and Applications Meeting* (Institute of Navigation, 2017).
- [23] Andrew D. Ludlow, Martin M. Boyd, Jun Ye, E. Peik, and P. O. Schmidt, Optical atomic clocks, *Rev. Mod. Phys.* **87**, 637–701 (2015).
- [24] William D. Phillips and Harold Metcalf, Laser deceleration of an atomic beam, *Phys. Rev. Lett.* **48**, 596–599 (1982).
- [25] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, Observation of Bose-Einstein condensation in a dilute atomic vapor, *Science* **269**, 198–201 (1995).

- [26] B. DeMarco and D. S. Jin, Onset of Fermi degeneracy in a trapped atomic gas, *Science* **285**, 1703–1706 (1999).
- [27] Florian Schreck and Klaasjan van Druten, Laser cooling for quantum gases, *Nat. Phys.* **17**, 1296–1304 (2021).
- [28] Dale G. Fried, Thomas C. Killian, Lorenz Willmann, David Landhuis, Stephen C. Moss, Daniel Kleppner, and Thomas J. Greytak, Bose-Einstein condensation of atomic hydrogen, *Phys. Rev. Lett.* **81**, 3811–3814 (1998).
- [29] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, Bose-Einstein condensation in a gas of sodium atoms, *Phys. Rev. Lett.* **75**, 3969–3973 (1995).
- [30] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet, Evidence of Bose-Einstein condensation in an atomic gas with attractive interactions, *Phys. Rev. Lett.* **75**, 1687–1690 (1995).
- [31] S. R. Granade, M. E. Gehm, K. M. O'Hara, and J. E. Thomas, All-optical production of a degenerate Fermi gas, *Phys. Rev. Lett.* **88**, 120405 (2002).
- [32] G. Modugno, G. Ferrari, G. Roati, R. J. Brecha, A. Simoni, and M. Inguscio, Bose-Einstein condensation of potassium atoms by sympathetic cooling, *Science* **294**, 1320–1322 (2001).
- [33] Tino Weber, Jens Herbig, Michael Mark, Hanns-Christoph Nagerl, and Rudolf Grimm, Bose-Einstein condensation of cesium, *Science* **299**, 232–235 (2003).
- [34] F. Pereira Dos Santos, J. Léonard, Junmin Wang, C. J. Barrelet, F. Perales, E. Rasel, C. S. Unnikrishnan, M. Leduc, and C. Cohen-Tannoudji, Bose-Einstein condensation of metastable helium, *Phys. Rev. Lett.* **86**, 3459–3462 (2001).
- [35] J. M. McNamara, T. Jeltsov, A. S. Tychkov, W. Hogervorst, and W. Vassen, Degenerate Bose-Fermi mixture of metastable atoms, *Phys. Rev. Lett.* **97**, 080404 (2006).
- [36] Axel Griesmaier, Jörg Werner, Sven Hensler, Jürgen Stuhler, and Tilman Pfau, Bose-Einstein condensation of chromium, *Phys. Rev. Lett.* **94**, 160401 (2005).
- [37] Mingwu Lu, Nathaniel Q. Burdick, Seo Ho Youn, and Benjamin L. Lev, Strongly dipolar Bose-Einstein condensate of dysprosium, *Phys. Rev. Lett.* **107**, 190401 (2011).
- [38] Mingwu Lu, Nathaniel Q. Burdick, and Benjamin L. Lev, Quantum degenerate dipolar Fermi gas, *Phys. Rev. Lett.* **108**, 215301 (2012).

- [39] K. Aikawa, A. Frisch, M. Mark, S. Baier, A. Rietzler, R. Grimm, and F. Ferlaino, Bose-Einstein condensation of erbium, *Phys. Rev. Lett.* **108**, 210401 (2012).
- [40] K. Aikawa, A. Frisch, M. Mark, S. Baier, R. Grimm, and F. Ferlaino, Reaching Fermi degeneracy via universal dipolar scattering, *Phys. Rev. Lett.* **112**, 010404 (2014).
- [41] E. T. Davletov, V. V. Tsyganok, V. A. Khlebnikov, D. A. Pershin, D. V. Shaykin, and A. V. Akimov, Machine learning for achieving Bose-Einstein condensation of thulium atoms, *Phys. Rev. A* **102**, 011302 (2020).
- [42] Yosuke Takasu, Kenichi Maki, Kaduki Komori, Tetsushi Takano, Kazuhito Honda, Mitsutaka Kumakura, Tsutomu Yabuzaki, and Yoshiro Takahashi, Spin-singlet Bose-Einstein condensation of two-electron atoms, *Phys. Rev. Lett.* **91**, 040404 (2003).
- [43] Takeshi Fukuhara, Yosuke Takasu, Mitsutaka Kumakura, and Yoshiro Takahashi, Degenerate Fermi gases of ytterbium, *Phys. Rev. Lett.* **98**, 030401 (2007).
- [44] Sebastian Kraft, Felix Vogt, Oliver Appel, Fritz Riehle, and Uwe Sterr, Bose-Einstein condensation of alkaline earth atoms: ^{40}Ca , *Phys. Rev. Lett.* **103**, 130401 (2009).
- [45] Y. N. Martinez de Escobar, P. G. Mickelson, M. Yan, B. J. DeSalvo, S. B. Nagel, and T. C. Killian, Bose-Einstein condensation of ^{84}Sr , *Phys. Rev. Lett.* **103**, 200402 (2009).
- [46] Simon Stellmer, Meng Khoon Tey, Bo Huang, Rudolf Grimm, and Florian Schreck, Bose-Einstein condensation of strontium, *Phys. Rev. Lett.* **103**, 200401 (2009).
- [47] B. J. DeSalvo, M. Yan, P. G. Mickelson, Y. N. Martinez de Escobar, and T. C. Killian, Degenerate Fermi gas of ^{87}Sr , *Phys. Rev. Lett.* **105**, 030402 (2010).
- [48] Daniel V. Schroeder, *An Introduction to Thermal Physics* (Oxford University Press, 2021).
- [49] Thomas H. Loftus, Tetsuya Ido, Martin M. Boyd, Andrew D. Ludlow, and Jun Ye, Narrow line cooling and momentum-space crystals, *Phys. Rev. A* **70**, 063413 (2004).
- [50] S. Snigirev, A. J. Park, A. Heinz, I. Bloch, and S. Blatt, Fast and dense magneto-optical traps for strontium, *Phys. Rev. A* **99**, 063421 (2019).
- [51] K. M. O'Hara, M. E. Gehm, S. R. Granade, and J. E. Thomas, Scaling laws for evaporative cooling in time-dependent optical traps, *Phys. Rev. A* **64**, 051403 (2001).

- [52] D. A. Butts and D. S. Rokhsar, Trapped Fermi gases, *Phys. Rev. A* **55**, 4346–4350 (1997).
- [53] Brian DeMarco, *Quantum Behavior of an Atomic Fermi Gas*, Ph.D. thesis, University of Colorado (2001).
- [54] Rianne S. Lous, Isabella Fritsche, Michael Jag, Bo Huang, and Rudolf Grimm, Thermometry of a deeply degenerate Fermi gas with a Bose-Einstein condensate, *Phys. Rev. A* **95**, 053627 (2017).
- [55] Simon Stellmer, Rudolf Grimm, and Florian Schreck, Production of quantum-degenerate strontium gases, *Phys. Rev. A* **87**, 013611 (2013).
- [56] Lindsay Sonderhouse, Christian Sanner, Ross B. Hutson, Akihisa Goban, Thomas Bilitewski, Lingfeng Yan, William R. Milner, Ana M. Rey, and Jun Ye, Thermodynamics of a deeply degenerate SU(n)-symmetric Fermi gas, *Nat. Phys.* **16**, 1216–1221 (2020).
- [57] Brian DeMarco and Joseph H. Thywissen, No vacancy in the Fermi sea, *Science* **374**, 936–937 (2021).
- [58] Georg M. Bruun and Charles W. Clark, Ideal gases in time-dependent traps, *Phys. Rev. A* **61**, 061601 (2000).
- [59] Claude Tannoudji, *Atom-photon interactions: basic processes and applications* (Wiley, New York, 1998).
- [60] James Bateman, André Xuereb, and Tim Freegarde, Stimulated Raman transitions via multiple atomic levels, *Phys. Rev. A* **81**, 043808 (2010).
- [61] M. Kasevich and S. Chu, Measurement of the gravitational acceleration of an atom with a light-pulse atom interferometer, *Appl. Phys. B* **54**, 321–332 (1992).
- [62] Marco Mancini, *Quantum Simulation with Ytterbium Atoms in Synthetic Dimensions*, Ph.D. thesis, University of Florence (2015).
- [63] Jim Napolitano J. J. Sakurai, *Modern Quantum Mechanics* (Cambridge University Pr., 2020).
- [64] A. Celi, P. Massignan, J. Ruseckas, N. Goldman, I. B. Spielman, G. Juzeliūnas, and M. Lewenstein, Synthetic gauge fields in synthetic dimensions, *Phys. Rev. Lett.* **112**, 043001 (2014).
- [65] Y.-J. Lin, K. Jiménez-García, and I. B. Spielman, Spin-orbit-coupled Bose-Einstein

- condensates, *Nature* **471**, 83–86 (2011).
- [66] Pengjun Wang, Zeng-Qiang Yu, Zhengkun Fu, Jiao Miao, Lianghai Huang, Shijie Chai, Hui Zhai, and Jing Zhang, Spin-orbit coupled degenerate Fermi gases, *Phys. Rev. Lett.* **109**, 095301 (2012).
- [67] Lawrence W. Cheuk, Ariel T. Sommer, Zoran Hadzibabic, Tarik Yefsah, Waseem S. Bakr, and Martin W. Zwierlein, Spin-injection spectroscopy of a spin-orbit coupled Fermi gas, *Phys. Rev. Lett.* **109**, 095302 (2012).
- [68] Bo Song, Chengdong He, Shanchao Zhang, Elnur Hajiyev, Wei Huang, Xiong-Jun Liu, and Gyu-Boong Jo, Spin-orbit-coupled two-electron Fermi gases of ytterbium atoms, *Phys. Rev. A* **94**, 061604 (2016).
- [69] Nathaniel Q. Burdick, Yijun Tang, and Benjamin L. Lev, Long-lived spin-orbit-coupled degenerate dipolar Fermi gas, *Phys. Rev. X* **6**, 031022 (2016).
- [70] S. Kolkowitz, S. L. Bromley, T. Bothwell, M. L. Wall, G. E. Marti, A. P. Koller, X. Zhang, A. M. Rey, and J. Ye, Spin-orbit-coupled fermions in an optical lattice clock, *Nature* **542**, 66–70 (2016).
- [71] S. L. Bromley, S. Kolkowitz, T. Bothwell, D. Kedar, A. Safavi-Naini, M. L. Wall, C. Salomon, A. M. Rey, and J. Ye, Dynamics of interacting fermions under spin-orbit coupling in an optical lattice clock, *Nat. Phys.* **14**, 399–404 (2018).
- [72] Zhan Wu, Long Zhang, Wei Sun, Xiao-Tian Xu, Bao-Zong Wang, Si-Cong Ji, Youjin Deng, Shuai Chen, Xiong-Jun Liu, and Jian-Wei Pan, Realization of two-dimensional spin-orbit coupling for Bose-Einstein condensates, *Science* **354**, 83–88 (2016).
- [73] Wei Sun, Bao-Zong Wang, Xiao-Tian Xu, Chang-Rui Yi, Long Zhang, Zhan Wu, Youjin Deng, Xiong-Jun Liu, Shuai Chen, and Jian-Wei Pan, Highly controllable and robust 2D spin-orbit coupling for quantum gases, *Phys. Rev. Lett.* **121**, 150401 (2018).
- [74] Zong-Yao Wang, Xiang-Can Cheng, Bao-Zong Wang, Jin-Yi Zhang, Yue-Hui Lu, Chang-Rui Yi, Sen Niu, Youjin Deng, Xiong-Jun Liu, Shuai Chen, and Jian-Wei Pan, Realization of an ideal Weyl semimetal band in a quantum gas with 3d spin-orbit coupling, *Science* **372**, 271–276 (2021).
- [75] Ming-Cheng Liang, Yu-Dong Wei, Long Zhang, Xu-Jie Wang, Han Zhang, Wen-Wei Wang, Wei Qi, Xiong-Jun Liu, and Xibo Zhang, Realization of Qi-Wu-Zhang model in spin-orbit-coupled ultracold fermions, arXiv:2109.08885 [cond-mat.quant-gas] (2021).

- [76] Victor Galitski and Ian B. Spielman, Spin-orbit coupling in quantum gases, *Nature* **494**, 49–54 (2013).
- [77] Karina Jimenez Garcia, *Artificial Gauge Fields for Ultracold Neutral Atoms*, Ph.D. thesis, University of Maryland (2012).
- [78] Shanchao Zhang and Gyu-Boong Jo, Recent advances in spin-orbit coupled quantum gases, *J. Phys. Chem. Solids* **128**, 75–86 (2019).
- [79] D L Campbell and I B Spielman, Rashba realization: Raman with RF, *New J. Phys.* **18**, 033035 (2016).
- [80] David Tong, *Lectures on the quantum Hall effect* (2016).
- [81] K. v. Klitzing, G. Dorda, and M. Pepper, New method for high-accuracy determination of the fine-structure constant based on quantized Hall resistance, *Phys. Rev. Lett.* **45**, 494–497 (1980).
- [82] R. B. Laughlin, Quantized Hall conductivity in two dimensions, *Phys. Rev. B* **23**, 5632–5633 (1981).
- [83] D. J. Thouless, M. Kohmoto, M. P. Nightingale, and M. den Nijs, Quantized Hall conductance in a two-dimensional periodic potential, *Phys. Rev. Lett.* **49**, 405–408 (1982).
- [84] Jean Dalibard, Fabrice Gerbier, Gediminas Juzeliūnas, and Patrik Öhberg, Artificial gauge potentials for neutral atoms, *Rev. Mod. Phys.* **83**, 1523–1543 (2011).
- [85] M. Aidelsburger, M. Lohse, C. Schweizer, M. Atala, J. T. Barreiro, S. Nascimbène, N. R. Cooper, I. Bloch, and N. Goldman, Measuring the Chern number of Hofstadter bands with ultracold bosonic atoms, *Nat. Phys.* **11**, 162–166 (2014).
- [86] Siddharth A. Parameswaran, Rahul Roy, and Shivaji L. Sondhi, Fractional quantum Hall physics in topological flat bands, *Comptes Rendus Physique* **14**, 816–839 (2013).
- [87] Emil J. Bergholtz and Zhao Liu, Topological flat band models and fractional Chern insulators, *Int. J. of Mod. Phys. B* **27**, 1330017 (2013).
- [88] Botao Wang, Xiaoyu Dong, and André Eckardt, Measurable signatures of bosonic fractional Chern insulator states and their fractional excitations in a quantum-gas microscope, *SciPost Physics* **12**, 10.21468/scipostphys.12.3.095 (2022).
- [89] Nigel R. Cooper and Jean Dalibard, Reaching fractional quantum Hall states with optical flux lattices, *Phys. Rev. Lett.* **110**, 185301 (2013).

- [90] N. R. Cooper and J. Dalibard, Optical flux lattices for two-photon dressed states, *Europhys. Lett.* **95**, 66004 (2011).
- [91] N. R. Cooper and R. Moessner, Designing topological bands in reciprocal space, *Phys. Rev. Lett.* **109**, 215302 (2012).
- [92] J. L. De Segovia, *Physics of Outgassing*, Tech. Rep. (CERN, 1999).
- [93] Avid Keshet, *A Next-Generation Apparatus for Lithium Optical Lattice Experiments*, Ph.D. thesis, Massachusetts Institute of Technology (2012).
- [94] Christian Thomas Joachim Gersdorf, *A Quantum Gas Microscope for Fermionic Atoms*, Ph.D. thesis, ETH Zurich (2012).
- [95] Benjamin J. Bloom, *Building a Better Atomic Clock*, Ph.D. thesis, University of Colorado (2014).
- [96] Sara L. Campbell, *A Fermi-degenerate three-dimensional optical lattice clock*, Ph.D. thesis, University of Colorado (2017).
- [97] Kevin Michael Birnbaum, *Cavity QED with multilevel atoms*, Ph.D. thesis, California Institute of Technology (2005).
- [98] Yosuke Shimada, Yuko Chida, Nozomi Ohtsubo, Takatoshi Aoki, Makoto Takeuchi, Takahiro Kuga, and Yoshio Torii, A simplified 461-nm laser system using blue laser diodes and a hollow cathode lamp for laser cooling of sr, *Review of Scientific Instruments* **84**, 063101 (2013).
- [99] Eric D. Black, An introduction to Pound–Drever–Hall laser frequency stabilization, *American Journal of Physics* **69**, 79–87 (2001).
- [100] Francisco Camargo, *Strontium Laser cooling and trapping apparatus*, Master’s thesis, Rice University (2015).
- [101] James Aman, *Two-photon photoassociative spectroscopy of strontium-86*, Ph.D. thesis, Rice University (2019).
- [102] Rudolf Grimm, Matthias Weidemüller, and Yurii B. Ovchinnikov, Optical dipole traps for neutral atoms, in *Advances In Atomic, Molecular, and Optical Physics* (Elsevier, 2000) pp. 95–170.
- [103] Richard Roy, Alaina Green, Ryan Bowler, and Subhadeep Gupta, Rapid cooling to quantum degeneracy in dynamically shaped atom traps, *Phys. Rev. A* **93**, 043403 (2016).

- [104] Simon Stellmer, Rudolf Grimm, and Florian Schreck, Detection and manipulation of nuclear spin states in fermionic strontium, *Phys. Rev. A* **84**, 043611 (2011).
- [105] Bo Song, Long Zhang, Chengdong He, Ting Fung Jeffrey Poon, Elnur Hajiyev, Shanchao Zhang, Xiong-Jun Liu, and Gyu-Boong Jo, Observation of symmetry-protected topological band with ultracold fermions, *Sci. Adv.* **4**, eaao4748 (2018).
- [106] Jeong Ho Han, Jin Hyoun Kang, and Y. Shin, Band gap closing in a synthetic Hall tube of neutral fermions, *Phys. Rev. Lett.* **122**, 065303 (2019).
- [107] Dylan O. Sabulsky, Colin V. Parker, Nathan D. Gemelke, and Cheng Chin, Efficient continuous-duty bitter-type electromagnets for cold atom experiments, *Review of Scientific Instruments* **84**, 104706 (2013).
- [108] Alexander Weiler, Alexander Pakosta, and Ankur Verma, *High-Speed Layout Guidelines*, Tech. Rep. (Texas Instruments, 2017).
- [109] Florian Seidler, *Digital high bandwidth feedback controller*, Master's thesis, University of Bonn (2015).
- [110] S. B. Nagel, C. E. Simien, S. Laha, P. Gupta, V. S. Ashoka, and T. C. Killian, Magnetic trapping of metastable 3P_2 atomic strontium, *Phys. Rev. A* **67**, 011401 (2003).
- [111] Juan A. Muniz, Matthew A. Norcia, Julia R. K. Cline, and James K. Thompson, A robust narrow-line magneto-optical trap using adiabatic transfer (2018).
- [112] Takashi Mukaiyama, Hidetoshi Katori, Tetsuya Ido, Ying Li, and Makoto Kuwata-Gonokami, Recoil-limited laser cooling of 87-Sr atoms near the Fermi temperature, *Phys. Rev. Lett.* **90**, 113002 (2003).
- [113] Paul Lauria, Wei-Ting Kuo, Nigel R. Cooper, and Julio T. Barreiro, Experimental realization of a fermionic spin-momentum lattice, *Phys. Rev. Lett.* (2022, in press).
- [114] Simon M. Rochester, Konrad Szymański, Mark Raizen, Szymon Pustelny, Marcis Auzinsh, and Dmitry Budker, Efficient polarization of high-angular-momentum systems, *Phys. Rev. A* **94**, 043416 (2016).
- [115] J. D. Whalen, S. K. Kanungo, R. Ding, M. Wagner, R. Schmidt, H. R. Sadeghpour, S. Yoshida, J. Burgdörfer, F. B. Dunning, and T. C. Killian, Probing nonlocal spatial correlations in quantum gases with ultra-long-range Rydberg molecules, *Phys. Rev. A* **100**, 011402 (2019).
- [116] Liang Hu, Enlong Wang, Leonardo Salvi, Jonathan N Tinsley, Guglielmo M Tino, and

- Nicola Poli, Sr atom interferometry with the optical clock transition as a gravimeter and a gravity gradiometer, *Class. Quant. Grav.* **37**, 014001 (2019).
- [117] Bao-Zong Wang, Yue-Hui Lu, Wei Sun, Shuai Chen, Youjin Deng, and Xiong-Jun Liu, Dirac-, Rashba-, and Weyl-type spin-orbit couplings: Toward experimental realization in ultracold atoms, *Phys. Rev. A* **97**, 011605 (2018).
- [118] Markus Greiner, Immanuel Bloch, Olaf Mandel, Theodor W. Hänsch, and Tilman Esslinger, Exploring phase coherence in a 2d lattice of Bose-Einstein condensates, *Phys. Rev. Lett.* **87**, 160405 (2001).
- [119] Tracy Li, *Probing Bloch band geometry with ultracold atoms in optical lattices*, Ph.D. thesis, Ludwig-Maximilians-Universität München (2016).
- [120] Linda Petzold, Automatic selection of methods for solving stiff and nonstiff systems of ordinary differential equations, *SIAM J Sci. Stat. Comp.* **4**, 136–148 (1983).
- [121] Florian Schäfer, Takeshi Fukuhara, Seiji Sugawa, Yosuke Takasu, and Yoshiro Takahashi, Tools for quantum simulation with ultracold atoms in optical lattices, *Nat. Rev. Phys.* **2**, 411–425 (2020).
- [122] N. Goldman, J. C. Budich, and P. Zoller, Topological quantum matter with ultracold gases in optical lattices, *Nat. Phys.* **12**, 639–645 (2016).
- [123] Hirokazu Miyake, Georgios A. Siviloglou, Colin J. Kennedy, William Cody Burton, and Wolfgang Ketterle, Realizing the Harper Hamiltonian with laser-assisted tunneling in optical lattices, *Phys. Rev. Lett.* **111**, 185302 (2013).
- [124] M. Aidelsburger, M. Atala, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch, Realization of the Hofstadter Hamiltonian with ultracold atoms in optical lattices, *Phys. Rev. Lett.* **111**, 185301 (2013).
- [125] Gregor Jotzu, Michael Messer, Rémi Desbuquois, Martin Lebrat, Thomas Uehlinger, Daniel Greif, and Tilman Esslinger, Experimental realization of the topological Haldane model with ultracold fermions, *Nature* **515**, 237–240 (2014).
- [126] O. Boada, A. Celi, J. I. Latorre, and M. Lewenstein, Quantum simulation of an extra dimension, *Phys. Rev. Lett.* **108**, 133001 (2012).
- [127] Tomoki Ozawa and Hannah M. Price, Topological quantum matter in synthetic dimensions, *Nat. Rev. Phys.* **1**, 349–357 (2019).
- [128] Shao-Liang Zhang and Qi Zhou, Shaping topological properties of the band structures in a shaken optical lattice, *Phys. Rev. A* **90**, 051601(R) (2014).

- [129] M. Mancini, G. Pagano, G. Cappellini, L. Livi, M. Rider, J. Catani, C. Sias, P. Zoller, M. Inguscio, M. Dalmonte, and L. Fallani, Observation of chiral edge states with neutral fermions in synthetic Hall ribbons, *Science* **349**, 1510–1513 (2015).
- [130] B. K. Stuhl, H.-I. Lu, L. M. Aycock, D. Genkina, and I. B. Spielman, Visualizing edge states with an atomic Bose gas in the quantum Hall regime, *Science* **349**, 1514–1518 (2015).
- [131] Jin Hyoun Kang, Jeong Ho Han, and Y. Shin, Realization of a cross-linked chiral ladder with neutral fermions in a 1D optical lattice by orbital-momentum coupling, *Phys. Rev. Lett.* **121**, 150403 (2018).
- [132] Dina Genkina, Lauren M Aycock, Hsin-I Lu, Mingwu Lu, Alina M Pineiro, and I B Spielman, Imaging topology of Hofstadter ribbons, *New J. Phys.* **21**, 053021 (2019).
- [133] Thomas Chalopin, Tanish Satoor, Alexandre Evrard, Vasiliy Makhalov, Jean Dalibard, Raphael Lopes, and Sylvain Nascimbene, Probing chiral edge dynamics and bulk topology of a synthetic Hall system, *Nat. Phys.* **16**, 1017–1021 (2020).
- [134] Chuan-Hsun Li, Yangqian Yan, Shih-Wen Feng, Sayan Choudhury, David B. Blasing, Qi Zhou, and Yong P. Chen, Bose-Einstein condensate on a synthetic topological Hall cylinder, *PRX Quantum* **3**, 010316 (2022).
- [135] 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, *Phys. Rev. Lett.* **117**, 220401 (2016).
- [136] Xiao-Tong Lu, Tao Wang, Ting Li, Chi-Hua Zhou, Mo-Juan Yin, Ye-Bing Wang, Xue-Feng Zhang, and Hong Chang, Doubly modulated optical lattice clock: Interference and topology, *Phys. Rev. Lett.* **127**, 033601 (2021).
- [137] Zengming Meng, Lianghai Huang, Peng Peng, Donghao Li, Liangchao Chen, Yong Xu, Chuanwei Zhang, Pengjun Wang, and Jing Zhang, Experimental observation of a topological band gap opening in ultracold Fermi gases with two-dimensional spin-orbit coupling, *Phys. Rev. Lett.* **117**, 235304 (2016).
- [138] Chang-Rui Yi, Long Zhang, Lin Zhang, Rui-Heng Jiao, Xiang-Can Cheng, Zong-Yao Wang, Xiao-Tian Xu, Wei Sun, Xiong-Jun Liu, Shuai Chen, and Jian-Wei Pan, Observing topological charges and dynamical bulk-surface correspondence with ultracold atoms, *Phys. Rev. Lett.* **123**, 190603 (2019).
- [139] A. Valdés-Curiel, D. Trypogeorgos, Q.-Y. Liang, R. P. Anderson, and I. B. Spielman, Topological features without a lattice in Rashba spin-orbit coupled atoms, *Nat.*

Commun. **12**, 593 (2021).

- [140] Q.-Y. Liang, D. Trypogeorgos, A. Valdés-Curiel, J. Tao, M. Zhao, and I. B. Spielman, Coherence and decoherence in the Harper-Hofstadter model, *Phys. Rev. Research* **3**, 023058 (2021).
- [141] Ren Zhang, Yangqian Yan, and Qi Zhou, Localization on a synthetic Hall cylinder, *Phys. Rev. Lett.* **126**, 193001 (2021).
- [142] Aurélien Fabre, Jean-Baptiste Bouhiron, Tanish Satoor, Raphael Lopes, and Sylvain Nascimbene, Laughlin’s topological charge pump in an atomic Hall cylinder, *Phys. Rev. Lett.* **128**, 173202 (2022).
- [143] Zejian Ren, Dong Liu, Entong Zhao, Chengdong He, Ka Kwan Pak, Jensen Li, and Gyu-Boong Jo, Chiral control of quantum states in non-hermitian spin-orbit-coupled Fermions, *Nat. Phys.* 10.1038/s41567-021-01491-x (2022).
- [144] Eric J. Meier, Fangzhao Alex An, and Bryce Gadway, Atom-optics simulator of lattice transport phenomena, *Phys. Rev. A* **93**, 051602 (2016).
- [145] Fangzhao Alex An, Bhuvanesh Sundar, Junpeng Hou, Xi-Wang Luo, Eric J. Meier, Chuanwei Zhang, Kaden R. A. Hazzard, and Bryce Gadway, Nonlinear dynamics in a synthetic momentum-state lattice, *Phys. Rev. Lett.* **127**, 130401 (2021).
- [146] Eric J. Meier, Fangzhao Alex An, and Bryce Gadway, Observation of the topological soliton state in the Su–Schrieffer–Heeger model, *Nat. Commun.* **7**, 13986 (2016).
- [147] Fangzhao Alex An, Eric J. Meier, and Bryce Gadway, Direct observation of chiral currents and magnetic reflection in atomic flux lattices, *Sci. Adv.* **3**, e1602685 (2017).
- [148] Dizhou Xie, Wei Gou, Teng Xiao, Bryce Gadway, and Bo Yan, Topological characterizations of an extended Su–Schrieffer–Heeger model, *npj Quantum Inf.* **5**, 55 (2019).
- [149] Dizhou Xie, Tian-Shu Deng, Teng Xiao, Wei Gou, Tao Chen, Wei Yi, and Bo Yan, Topological quantum walks in momentum space with a Bose-Einstein condensate, *Phys. Rev. Lett.* **124**, 050502 (2020).
- [150] S. K. Kanungo, J. D. Whalen, Y. Lu, M. Yuan, S. Dasgupta, F. B. Dunning, K. R. A. Hazzard, and T. C. Killian, Realizing topological edge states with Rydberg-atom synthetic dimensions, *Nat. Commun.* **13**, 972 (2022).
- [151] Christopher Oliver, Aaron Smith, Thomas Easton, Grazia Salerno, Vera Guarrera, Nathan Goldman, Giovanni Barontini, and Hannah M. Price, Bloch oscillations along

- a synthetic dimension of atomic trap states, arXiv:2112.10648 [cond-mat.quant-gas] (2021).
- [152] Horst L. Stormer, Daniel C. Tsui, and Arthur C. Gossard, The fractional quantum Hall effect, *Rev. Mod. Phys.* **71**, S298–S305 (1999).
- [153] Yangqian Yan, Shao-Liang Zhang, Sayan Choudhury, and Qi Zhou, Emergent periodic and quasiperiodic lattices on surfaces of synthetic Hall tori and synthetic Hall cylinders, *Phys. Rev. Lett.* **123**, 260405 (2019).
- [154] R. P. Anderson, D. Trypogeorgos, A. Valdés-Curiel, Q.-Y. Liang, J. Tao, M. Zhao, T. Andrijauskas, G. Juzeliūnas, and I. B. Spielman, Realization of a deeply subwavelength adiabatic optical lattice, *Phys. Rev. Research* **2**, 013149 (2020).
- [155] M. A. Khomehchi, Chunlei Qu, M. E. Mossman, Chuanwei Zhang, and P. Engels, Spin-momentum coupled Bose-Einstein condensates with lattice band pseudospins, *Nat. Commun.* **7**, 10867 (2016).
- [156] Rodrigo Rosa-Medina, Francesco Ferri, Fabian Finger, Nishant Dogra, Katrin Kroeger, Rui Lin, R. Chitra, Tobias Donner, and Tilman Esslinger, Observing dynamical currents in a non-Hermitian momentum lattice, *Phys. Rev. Lett.* **128**, 143602 (2022).
- [157] N. R. Cooper, Optical flux lattices for ultracold atomic gases, *Phys. Rev. Lett.* **106**, 175301 (2011).
- [158] Sivaprasad Omanakuttan, Anupam Mitra, Michael J. Martin, and Ivan H. Deutsch, Quantum optimal control of ten-level nuclear spin qubits in 87-Sr, *Phys. Rev. A* **104**, L060401 (2021).
- [159] J. B. Naber, L. Torralbo-Campo, T. Hubert, and R. J. C. Spreeuw, Raman transitions between hyperfine clock states in a magnetic trap, *Phys. Rev. A* **94**, 013427 (2016).
- [160] Francesco Scazza, *Probing $SU(N)$ -symmetric orbital interactions with ytterbium Fermi gases in optical lattices*, Ph.D. thesis, Ludwig Maximilian University of Munich (2015).
- [161] Daniel Schaefer Barker, *Degenerate Gases of Strontium for Studies of Quantum Magnetism*, Ph.D. thesis, University of Maryland, College Park (2016).
- [162] D. L. Campbell, G. Juzeliūnas, and I. B. Spielman, Realistic Rashba and Dresselhaus spin-orbit coupling for neutral atoms, *Phys. Rev. A* **84**, 025602 (2011).
- [163] H. M. Price and N. R. Cooper, Mapping the berry curvature from semiclassical

dynamics in optical lattices, Phys. Rev. A **85**, 033620 (2012).

- [164] Derek Molloy, *Exploring BeagleBone: tools and techniques for building with embedded Linux* (Wiley, Indianapolis, IN, 2015).