

A new record in atomic clock performance

by

Travis L. Nicholson

B.A., University of Colorado, 2006

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2015

This thesis entitled:
A new record in atomic clock performance
written by Travis L. Nicholson
has been approved for the Department of Physics

Jun Ye

Dr. Deborah Jin

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Nicholson, Travis L. (Ph.D., Physics)

A new record in atomic clock performance

Thesis directed by Dr. Jun Ye

The pursuit of better atomic clocks has advanced many fields of research, providing better quantum state control, new insights in quantum science, tighter limits on fundamental constant variation, and improved tests of relativity. This thesis describes the construction and characterization of an ^{87}Sr optical lattice clock with a state-of-the-art stable laser. The performance of an atomic clock is typically gauged by two figures of merit: stability and total systematic uncertainty. Stability is the statistical precision of a clock or frequency standard, and the total systematic uncertainty is the combined uncertainty of all known systematic measurement biases. Several demonstrations of clock stability are presented in this work, one of which was the first to significantly outperform ion clocks. The most recent of these measurements resulted in fractional stability of 2.2×10^{-16} at 1 s, which is the best reported to date. These stability improvements are used for two systematic evaluations of our clock. The first full evaluation at 6.4×10^{-18} total uncertainty took the record for best clock performance. The second evaluation used improved strategies for systematic measurements, achieving a new best total systematic uncertainty of 2.1×10^{-18} . With a combination of accurate radiation thermometry and temperature stabilization of the measurement environment, we demonstrate the first lattice clock to achieve the longstanding goal of 10^{-18} level uncertainty in the formidable blackbody radiation shift. Improvements in the density, lattice ac Stark, and dc Stark shifts were also a result of innovations that are described in this thesis. Due to the low total uncertainty of the Sr clock, timekeeping based on this system would not lose a second in 15 billion years (longer than the age of the Universe), and it is sensitive to a gravitational redshift corresponding to a height change of 2 cm above the Earth's surface.

Dedication

To my kind, brilliant, and loving wife, Dr. Huanqian Loh. And to my caring parents, Jeff and Johannah.

Acknowledgements

I would like to express my gratitude to those who have mentored me, worked with me, educated me, and supported me during my time at JILA. Firstly, I want to thank my advisor, Jun Ye. I have always admired Jun's mastery of a wide range of physics, his intellectual rigor, and his infectious enthusiasm. The time and energy he spends making sure his students meet his high standards have greatly benefited those of us who were fortunate enough to be taken under his wing. Under his mentorship and guiding hand, I have grown immensely as a scientist and a professional. I am indebted to Jun for imparting his wisdom and technical expertise, and I thank him for his dedication to my education and my success.

I worked on Jun's second-generation strontium system (Sr2) from its construction until it was retired, and during that time I interacted with a number of talented graduate students and postdoctoral researchers. I am grateful to Martin Boyd, Gretchen Campbell, and Jan Thomsen for their contributions to the Sr2 system in its early days. Marty was a clever physicist and one of the funniest people I have worked with, and Gretchen impressed me with her ability to successfully work full time on both Sr1 and Sr2. Professor Jan Thomsen is a frequent visitor to our group, and he is a tremendously friendly person and a hardworking labmate. Sebastian Blatt and I were the two graduate students primarily assigned to Sr2, and we spent many long nights together building this experiment. Sebastian is a hardworking and driven scientist who understands that, in the face of long hours and the struggles that come with building a new experiment, people work better together when they make it a point to get out of the lab and have fun with one another.

Shortly after we finished building the Sr2 system (first as a ^{88}Sr machine), Sebastian and I

were joined by graduate student Ben Bloom. As we were wrapping up our first result on Sr2, which was a study of optical Feshbach resonances, we were also joined by postdoctoral researcher Jason Williams. After this first result was published and Sebastian graduated, taking a postdoctoral position at Harvard, Ben, Jason, and I revamped Sr2 to be a ^{87}Sr clock. No description of our trio could be better than to simply state that together the three of us achieved the best atomic clock, which is a feat I will always be proud of. As we were finishing our stability measurements, the lively and talented graduate student Sara Campbell started on Sr2. I am particularly grateful for Sara's work on our Ti:saph laser system. Partway through the first Sr2 full systematic evaluation, Jason accepted a permanent position at Jet Propulsion Laboratories. Around this time Ben decided to graduate early to take up an industry position at Intel.

During the second Sr2 systematic evaluation, graduate student Ross Hutson and undergraduate Rees McNally became part of our team. Ross is a bright student who is always willing to lend a helping hand. Rees is a mature, adept young scientist, and I wish him luck as he heads off to Columbia University for graduate school. As Sara, Ross, Rees, and I were nearing the end of the second evaluation, we were joined by postdoctoral researcher George (Ed) Marti, who quickly proved to be a clever scientist.

With the second evaluation complete, I began writing my thesis, and Sr2 was torn down to make space for Sara, Ross, and Ed's new generation Sr system. Once, while they were clearing away the Sr2 apparatus, they asked me to take a quick break from writing my dissertation and give them a hand disassembling the main vacuum chamber. Although our last result was a fitting end to the system I had worked on for so many years, I confess that it was painful to hear the chamber lose vacuum as I turned the wrench. However, I am looking forward to hearing about the many results that are sure to come from the new Sr apparatus.

My thanks are also due to my labmates on Sr1. Graduate student Michael Martin and I overlapped for the majority of my graduate career, and I am grateful for his hard work on the Sr clock laser. Mike is a thoughtful, gifted scientist, and I am certain he will continue to succeed at CalTech and beyond. Matthew Swallows was a postdoctoral researcher on Sr1, and I thank him

for his work on the first Sr clock comparisons. I am sure that Matt's meticulous nature is serving him well at AOSense. Michael Bishof, who has recently accepted a permanent position at Argonne National Labs, is an effective scientist and a great guy to watch movies with. Visiting researcher Yige Lin was a kind labmate and a thoughtful scientist, and I hope his Sr clock in Beijing is working well. Postdoctoral researcher Xibo Zhang is a hardworking and thorough colleague, and I am sure he will be a successful professor. Sarah Bromley is a friendly and dedicated young student who is off to a great start in her career. Tobias Bothwell recently joined Sr1, and I am sure he will continue Sr1's tradition of success.

I am grateful to postdoctoral researcher Wei Zhang for his work on the Sr clock laser, 689 nm cooling laser, and frequency comb. Wei is one of the hardest working and most dedicated people I have worked with. Lindsay Sonderhouse is a bright and friendly new graduate student, and I thank her for important contributions to the Sr laser systems. It has also been fun to interact with the grandfather of our team, Jan Hall. Due to Jan's encyclopedic technical knowledge and enthusiasm for physics, conversations with him were always delightful and immensely educational.

The Sr2 team has been fortunate to collaborate with many great scientists over the years. Paul Julienne provided invaluable theoretical work for the first Sr2 experiment, and since then he and I have collaborated on theoretical studies of intercombination line Feshbach resonances. It has always been an honor and a pleasure to work with Paul. I also want to thank our collaborator Weston Tew from the Sensor Science Division at NIST Gaithersburg. I am grateful for his masterful work on our thermometry project, which was one of the most important pieces of the second Sr2 systematic evaluation. I also wish to thank Gregory Strouse, also from NIST Sensor Science, for his role in our thermometry work. JILA theorist Ana Maria Rey has been a good friend to the Sr team, and I thank her for helping us understand the density shift, which was a central theme of the second publication from Sr2. I thank Murray Barrett for providing useful and enlightening analytic theory of the Sr2 decay measurement. I am grateful to Marianna Safronova for her advice and expert calculations.

I am thankful to my undergraduate advisor, Deborah Jin. Debbie was a great mentor who

made it a point to provide me with a wide range of technical training. My work with her team confirmed my desire to pursue a career in physics, and I am grateful for the many opportunities Debbie gave me to achieve that goal.

My thanks go out to the many skilled teachers who taught my graduate courses. The classes of Leo Radzihovsky, Paul Beale, Andrew Hamilton, and Tom DeGrand were enriching and fun. Thanks also to Paul Beale for his continued friendliness and for always being happy to give me a helping hand. Thanks to Andrew Hamilton for being a fun person to chat physics with.

I am continually impressed at how many people on the JILA technical staff are simultaneously invaluable problem solvers and friends with everyone in the building. This includes instrument shop staff Hans Green, Blaine Horner, Kim Hagen, Todd Asnicar, Kels Detra, and Tracy Keep; electronics shop staff Terry Brown and Carl Sauer; and computing group staff James Raith, Mike Paige, and Jim McKown. Special thanks to Dave Alchenburger for assisting with some of our toughest technical problems over the years. Thanks in particular to Hans Green for his first-rate work on so many of the instrument shop jobs that were critical to this thesis.

I am grateful to my parents for all they have done for me. Since childhood, my mother has been a caring and generous parent, and I have always admired her cleverness and curiosity. My father is a vibrant person, and some of my best memories from the past several years have been of weekend ski days with him and Huanqian. My thanks go to my brother for his friendship and support. He is a smart guy with a strong conscience, and I am grateful to have him in my life. I am thankful to my sister for her wonderful sense of humor and her friendship. I am proud of the person she has become. Thanks to my brother-in-law for being a fun and friendly addition to our family, and thanks to my little nephew for the joy he has brought to my life. My appreciation also goes out to my parents-in-law. I feel lucky to have a kindhearted and reliable mother-in-law and an intelligent and friendly father-in-law. I am grateful to them for welcoming me into their family. I also thank my brother-in-law, his wife, and their children for their kindness and friendship.

Lastly, and most importantly, my thanks go out to my wife, Huanqian. She is an excellent partner for the same reason that she succeeds in so many aspects of life, namely that she approaches

all matters large and small with thoughtfulness, care, dedication, and sincere humbleness. I am grateful for her wonderful sense of humor, which was a welcome respite when the Sr2 lasers were acting up. I hope that over the years, more of her many admirable traits rub off on me, and I look forward to a lifetime of her company.

Contents

Chapter		
1	Introduction	1
1.1	Precision, Accuracy, and Systematic Uncertainty	1
1.2	Metrological Standards	2
1.3	Timekeeping	3
1.4	Atomic Time	4
1.5	Optical Clocks	5
1.6	State of the Art Frequency Standards	6
1.7	Thesis Overview	7
2	The Physics of the Strontium Lattice Clock	9
2.1	The Sr Atom	9
2.2	Laser Cooling Sr	10
2.3	The Magic Wavelength Optical Lattice	12
2.3.1	Optical Dipole Traps	12
2.3.2	Retroreflected Lattices	14
2.3.3	Cavity Lattices	15
2.3.4	The Magic Wavelength	18
2.4	Spectroscopy of Narrow Resonances in the Lamb-Dicke Regime	19
2.4.1	System Hamiltonian	19

2.4.2	Sideband Transitions	22
2.4.3	The Lamb-Dicke Regime	24
2.4.4	Anharmonicity and Sideband Broadening	26
2.4.5	The Lattice in the Presence of Gravity	28
2.5	Digital Servo to Atomic Spin Stretched States	31
2.5.1	Spin Polarization	31
2.5.2	Locking to the Clock Transition	33
2.5.3	Digital Lock to Alternate Stretched States	35
2.6	The Magic Wavelength Optical Lattice Revisited	37
2.7	Introduction to Stability and Systematic Effects	37
3	The Sr2 Apparatus	40
3.1	The Sr2 Vacuum Chamber	40
3.1.1	Sr Oven	40
3.1.2	Main Chamber	41
3.2	Cooling on the $^1S_0 \rightarrow ^1P_1$ Transition	43
3.2.1	461 nm Laser System	43
3.2.2	Zeeman Slower and Transverse Cooling	45
3.2.3	Three-Dimensional MOT and Repump Lasers	47
3.3	Cooling on the $^1S_0 \rightarrow ^3P_1$ Transition	48
3.3.1	689 nm Laser System	48
3.3.2	The 689 nm MOT	49
3.4	One-Dimensional Optical Lattice	53
3.4.1	The Cavity Lattice	53
3.4.2	Tapered Amplifier Lattice Laser	56
3.4.3	Titanium Sapphire Lattice Laser	58
3.5	Spin Polarization into the Angular Momentum Stretched States	59

3.5.1	Spin Polarization Bias Field Coils	59
3.5.2	Spin Polarization	61
3.6	The Clock Laser	63
3.6.1	Cavity Stabilization	63
3.6.2	Fiber Phase Noise Cancellation	63
3.7	Digital Atomic Servo	64
3.7.1	Clock Laser Frequency Modulator and Drift Cancellation	64
3.7.2	Camera Imaging and Photomultiplier Tube Detection	65
3.7.3	Clock Transition Scans	66
3.7.4	Digital Servo Software	68
4	Clock Stability	70
4.1	The Allan Deviation	70
4.1.1	Definition	70
4.1.2	Estimators	72
4.2	Effects that Limit Stability	74
4.2.1	Quantum Projection Noise	74
4.2.2	The Dick Effect	75
4.2.3	Long-Term Drift	78
4.3	Stability in Optical Clocks	79
4.3.1	The Stability of Lattice Clocks and Ion Clocks	79
4.3.2	Using Longer Cavities to Reduce Thermal Noise	80
4.3.3	Self-Comparison	82
4.4	Synchronous Interrogation	83
4.4.1	Two Sr Clocks with a Shared Clock Laser Source	83
4.4.2	Excited State Fraction Correlation	85
4.5	Asynchronous Interrogation	88

4.5.1	Using Asynchronous Interrogation to Measure Independent Clock Stability	88
4.5.2	Correlations and Allan Deviation	90
4.6	High Stability Self-Comparison	93
5	Systematic Uncertainty	97
5.1	Systematic Measurements	97
5.1.1	Introduction	97
5.1.2	Digital Lock-In Detection	98
5.1.3	Lever Arm	100
5.1.4	Model Testing	101
5.1.5	Point-String Analysis	104
5.2	Residual First-Order Zeeman Shift	107
5.2.1	Introduction	107
5.2.2	Background Field Servo	108
5.2.3	Shift Analysis	109
5.3	Second-Order Zeeman Shift	111
5.3.1	Introduction	111
5.3.2	The Lattice Tensor Shift	112
5.3.3	First and Second Evaluations	113
5.4	Probe Ac Stark Shift	113
5.4.1	Introduction	113
5.4.2	First and Second Evaluations	115
5.5	Density Shift	118
5.5.1	p-Wave Density Shift	118
5.5.2	Density Shift in a Cavity Lattice	120
5.5.3	Large Lever Arm Measurement	121
5.5.4	Systematic Evaluation	124

5.6	Lattice Stark Shift	124
5.6.1	Lever Arm and Nonlinear Light Shifts	124
5.6.2	Technical Considerations	126
5.6.3	First and Second Evaluations	127
5.6.4	Hyperpolarizability	130
5.7	Dc Stark Shift	131
5.7.1	Introduction	131
5.7.2	Electrodes and Lock-In Sequence	132
5.7.3	Observation of Dc Stark Shift and Removal of Background Field	134
5.7.4	Active Control of the Dc Stark Shift	136
5.8	Static Blackbody Shift	139
5.8.1	Introduction to the Blackbody Radiation Shift	139
5.8.2	Sources of Uncertainty	140
5.8.3	Radiation Thermometry	142
5.8.4	Kirchhoff's Radiation Law	144
5.8.5	First Evaluation: Silicon Diode Radiation Thermometers	148
5.8.6	First Evaluation: Non-Thermal Heat Bias	150
5.8.7	First Evaluation: Immersion Error	152
5.8.8	First Evaluation: Passive Temperature Control	153
5.8.9	First Evaluation: Other Sources of Uncertainty	154
5.8.10	Second Evaluation: Platinum Resistance Radiation Thermometers	156
5.8.11	Second Evaluation: Sensor Prequalification	157
5.8.12	Second Evaluation: New Radiation Thermometry Hardware	158
5.8.13	Second Evaluation: Sensor Calibration at NIST	159
5.8.14	Second Evaluation: Measurements in the Sr2 Chamber	163
5.9	Dynamic Blackbody Shift	167
5.9.1	A Scheme to Measure the Dynamic Coefficient	167

5.9.2	Decay Measurement	169
5.9.3	Fit Uncertainty	172
5.9.4	Systematic Bias Due to the Finite DFB Laser Pulse	174
5.9.5	Fit Bias Due to Stray Electromagnetic Radiation	176
5.9.6	Density-Dependent Effects	178
5.10	Minor Systematic Shifts	180
6	Outlook and Concluding Remarks	183
	Bibliography	187
	Appendix	
A	Mathematical Relations Not Derived in the Main Text	196
A.1	Two Time Signals Related by a Transfer Function	196
A.2	Relating the Allan Deviation to the Power Spectral Density	197
A.3	Power Spectral Density of Discrete Time Signals	198
A.4	Detection Noise	199

Tables

Table

2.1	Strontium Isotopes	9
5.1	Uncertainty Budget for Both Systematic Evaluations	98
5.2	Diode Sensor Uncertainty Budget	155
5.3	Long PRT Sensor Uncertainty Budget	167
5.4	Decay Measurement Uncertainty Budget	180

Figures

Figure

1.1	The Target Analogy	1
2.1	Strontium Level Diagram	10
2.2	Repumping Diagram	11
2.3	Hyperfine Levels of the 3P_1 State	12
2.4	A retroreflected lattice	14
2.5	A cavity lattice	15
2.6	Cavity Enhancement	17
2.7	Ac Stark shift hyperfine dependence	20
2.8	Motional state quantization	21
2.9	Sideband Rabi frequency	23
2.10	Sideband spectra	24
2.11	Sideband spectra with an improved Lamb-Dicke parameter	25
2.12	Sideband Broadening	27
2.13	The Lattice in the Presence of Gravity	30
2.14	Transitions with Unpolarized and Spin Polarized Atoms	32
2.15	Locking to the Line Center	34
2.16	Spin Polarization Timing	36
3.1	Oven	41

3.2	Vacuum Chamber	42
3.3	Main Chamber	43
3.4	461 nm Laser System	44
3.5	461 nm Saturation Spectroscopy	46
3.6	Zeeman Slower	46
3.7	689 nm Saturation Spectroscopy and Cavity Stabilization	49
3.8	689 nm Slave Lasers	50
3.9	689 nm MOT Timing Diagram	51
3.10	689 nm Laser FM	52
3.11	Bias Coils	54
3.12	Cavity Lattice Mounts	55
3.13	Lattice Mount Cutaway	55
3.14	Tapered-Amplifier-Based Lattice Laser	57
3.15	Titanium Sapphire Lattice Laser	58
3.16	Spin Polarization Bias Coils	60
3.17	Bias Coil Current Servo	61
3.18	Clock Laser Cavity	62
3.19	Sideband Scan	67
3.20	Resolved m_F State Scan	68
3.21	Stretched State Scan After Spin Polarization	69
4.1	Theoretical Allan Deviation	71
4.2	Allan Deviation Estimator	72
4.3	Total Deviation Estimator	73
4.4	Sensitivity Function	77
4.5	Atomic Clock Stability Comparison	79
4.6	7 cm and 40 cm Cavity Stabilization	81

4.7	Optical Clock Stability Comparison	82
4.8	Comparisons with a Shared Clock Laser	84
4.9	Synchronized Raw Data	86
4.10	Simulated Correlation Measurement	86
4.11	Synchronous Interrogation Measurements	87
4.12	Correlation Data with the 40 cm Cavity	90
4.13	Allan Deviations with the 40 cm Cavity	91
4.14	4 hour Allan Deviation	92
4.15	Long-Term Drift Simulation	94
4.16	High Stability Self-Comparison	96
5.1	Lock-In Detection	99
5.2	Probe Stark Simulation	101
5.3	Background Field Servo	110
5.4	Minimizing σ Transitions	112
5.5	Second-Order Zeeman Shift Evaluation	114
5.6	First Probe Stark Shift Evaluation	116
5.7	Second Probe Stark Shift Evaluation	117
5.8	Density Shift Linearity in N	122
5.9	Density Shift Lock-In Data Linearity in N	122
5.10	Density Shift $U_{atom}^{3/2}$ Scaling	123
5.11	Density Shift Evaluation	125
5.12	Lattice Stark Shift Evaluation	128
5.13	Electrodes for Dc Stark Shift Measurement	133
5.14	First Observation of a Dc Stark Shift	134
5.15	Dc Stark Shift Cancellation Servo	136
5.16	Dc Stark Shift Under Active Control	137

5.17 Dc Stark Data Processing	138
5.18 Simple Model of Radiation Thermometry	142
5.19 Simple Case of Kirchhoff's Radiation Law	145
5.20 Silicon Diode Sensor Mounts	149
5.21 Non-Thermal Heat Simulation	151
5.22 Sr2 Chamber Passive Heat Shield	154
5.23 PRT Electronics	158
5.24 Sensor Calibration Chamber	160
5.25 Sensor Calibration	161
5.26 Sensor Agreement After Transport	164
5.27 Sensor Agreement After Installation	166
5.28 Decay Measurement Scheme	169
5.29 Decay Measurement Calculation	170
5.30 Decay Measurement	171
5.31 Decay Data Noise	172
5.32 Decay Data Simulated Uncertainty	173
5.33 Decay Fit for Finite Pulse Calculation	174
5.34 Decay Fit for 500 ns Pulse Calculation	175
5.35 3D_1 Lifetime as a Function of Atom Number	179

Chapter 1

Introduction

1.1 Precision, Accuracy, and Systematic Uncertainty

Throughout the 20th century, physics research pushed the frontiers of technology and transformed civilization. The quest for better technology continues to motivate numerous physics research endeavors, which are driven by technological goals like renewable energy, quantum information, and biomedical devices. Good technology must be compatible and interoperable across the Earth if it is to meet the demands of the global economy. This means that measurements—which are foundational to the design, construction, characterization, and assessment of technology—must be internationally and precisely consistent.

Two figures of merit that describe the quality of measurements are precision and accuracy, which are exemplified in Figure 1.1. Here I use the famous target analogy for these two concepts. Figure 1.1 depicts different distributions of 70 holes shot into a target. The shooter's precision

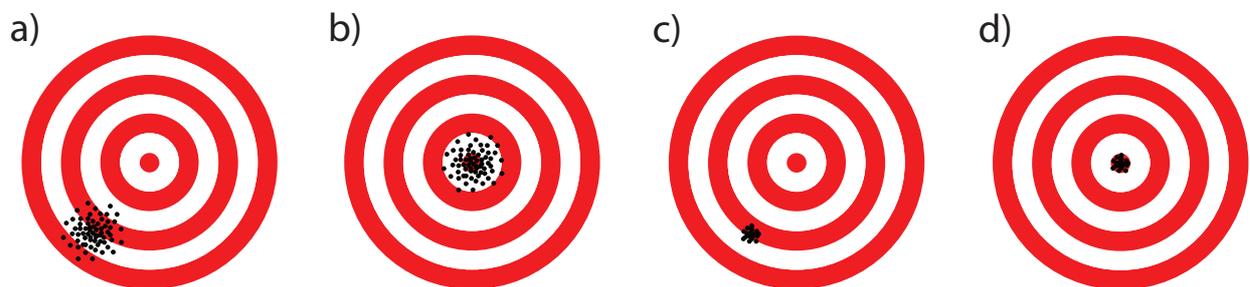


Figure 1.1: a) The shooter is neither precise nor accurate. b) Accurate but not precise. c) Precise but not accurate. d) Precise and accurate.

is the width of the distribution, and the shooter’s accuracy is the distance between the center of the distribution and the bullseye. Similarly, measurement outcomes occur in distributions, where the precision is the width of these distributions, and the accuracy is how much the centers of the distributions differ from some true value.¹

Systematic uncertainty can also be explained with the target analogy. To quantify the shooter’s performance, one must measure the accuracy of the distribution of holes. In Figures 1.1a or 1.1c, if one attempted to measure the difference between the center of the distribution and the bullseye, this measurement would have uncertainty (known as the systematic uncertainty). This is because of the inherent uncertainty in determining the center of the distribution. In Figure 1.1a, the distribution is wider than in Figure 1.1c, so there is more uncertainty in locating the distribution center and therefore more systematic uncertainty. This demonstrates a common feature of measurements: better precision leads to better systematic uncertainty.

1.2 Metrological Standards

Modern, global technology often requires precise and accurate measurement instruments. Precision is needed to resolve small changes in measured quantities. Accuracy is necessary to ensure that measurements are reproducible and agree worldwide. In this case, high accuracy typically entails that measurement instruments have been calibrated to high quality, internationally uniform metrological standards.

Metrological standards are controlled systems that realize physical units. For example, a temperature standard could be a water bath that is pressure and temperature controlled such that it maintains liquid, solid, and gaseous phases of water all at once (since the Kelvin is defined as exactly $1/273.16$ the temperature of the water triple point). The worldwide uniformity of standards²

¹ In principle, these quantities should be called the “imprecision” and the “inaccuracy”; however, I will be a bit loose in making the distinction between “precision” versus “imprecision,” “stability” versus “instability,” etc. Much of the clock literature is equally loose, and there is a good reason for this. These terms are used often in clock publications, and it is a burden on careful readers to have to be concerned with things like whether “stability” makes sense in its sentence instead of “instability.” In practice, it is often clearer to mostly use “stability,” “precision,” and “accuracy,” and rely on the context to shed light on whether “instability,” “imprecision,” or “inaccuracy” is actually meant.

is ensured through the international Treaty of the Metre (or Metre Convention), which established an international standards bureau as well as metrology conferences and committees of signatory nations. Through these conventions, member states coordinate international metrology and, in 1960, formally established the International System of Units (SI). The SI is an evolving system that updates its definitions when superior standards based on new systems become available.³

Metrology researchers continue to develop more precise standards in anticipation that these improved systems will be needed as technology becomes more precise and more global.

The most precise standards are those of frequency, which define the SI second. Currently, the definition of the SI second is “The duration of 9193631770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the cesium 133 atom.” Many member nations of the Metre Convention maintain one frequency *primary standard*, which is a microwave resonator slaved to the ground state hyperfine transition in a cesium fountain [48]. Primary standards are compared with those of other member nations, and they can also serve as local standards against which other frequencies can be measured. In the United States, the Cs primary standard is located at the National Institute of Standards and Technology (NIST) in Boulder, CO.

1.3 Timekeeping

One of the most critical applications of metrological standards is globally coordinated timekeeping, which provides precise timing and synchronization for modern technology. Some of the technologies that rely on precise timekeeping are GPS, secure communication, space navigation systems, power grid synchronization, and computer network timing. Timekeeping is based on clocks, which are (at least) comprised of *oscillators*, *counters*, and *references* [56]. Examples of oscillators found in clocks are pendulums (grandfather clocks), crystal resonators (quartz clocks), and elec-

² Engineers often use the word “standards” differently. Here I use “standards” to mean only “metrological standards.”

³ For instance, the SI meter was defined in units of wavelengths of a krypton transition. Later this definition was updated to define length as “the distance traveled by light in a vacuum in 1/299792458 second” since the latter could eventually be measured more precisely.

tromagnetic waves (atomic clocks). A clock counts oscillator cycles, converting them to quantities with temporal dimensions by equating a unit of time with a given number of oscillator periods [56]. For instance, the resonators in many quartz clocks are laser trimmed until their fundamental frequencies are 32768 Hz, and these clocks are designed to advance 1 second after 32768 oscillations are counted.

Clock oscillator phase or frequency noise degrades timekeeping precision. For an oscillator that can be expressed as $A \sin(\omega t + \phi)$, where A is a constant amplitude and ω is the (constant) base frequency, allowing the phase offset ϕ to be a stochastic function of time describes both phase and frequency noise [100]. If a clock considers 1 second to be exactly n oscillator periods, the oscillator phase after n cycles is $\omega t + \phi = 2n\pi$. This relation implies that, for a stationary phase noise process, $\delta t = \delta\phi/\omega$ (where δt is the root mean square (RMS) time signal noise, and $\delta\phi$ the RMS phase noise). Therefore, good clock design means employing a phase-stable (small $\delta\phi$) oscillator with the largest possible base frequency.

To select the value of n correctly, one must measure the number of oscillator cycles in 1 s. This is accomplished by referencing or calibrating the clock oscillator frequency to an accurate frequency signal. Historically, this signal was often Earth’s rotational frequency (which could be measured with sundials) [56]. Now, for cutting edge clocks that use large base frequency oscillators, microwave or optical atomic *clock transitions* are used as reference frequencies.⁴

1.4 Atomic Time

A frequency standard is typically a referenced oscillator [100]. Therefore, in principle, a frequency standard plus a counter yields a clock. The best clocks to date are atomic clocks, which are based on frequency standards that use microwave or optical electromagnetic waves as oscillators and atomic transitions as references. High performance atomic clocks exist mostly in academic and government laboratories, such as NIST’s atomic clock that generates official US time (based on their Cs primary standard) [48]. The Cs clock does not lose 1 s in 300 million years, in contrast to

⁴ Any transition used as a frequency reference in a clock is called that system’s clock transition.

the high-quality quartz clocks found in marine chronometers, which lose $\gtrsim 10$ s annually. Rather than use better clocks for consumer timepieces, it is typically sufficient and cost effective to use quartz clocks that are occasionally reset to official US time.

In practice, a Cs primary standard cannot run continuously, so many nations in the Metre Convention utilize *flywheels*. In this case, a flywheel is an array of hydrogen masers, which provides the frequency that is counted to generate the official time of many countries, including the United States [67]. To maintain accuracy and long-term precision, the maser arrays are periodically calibrated to Cs primary standards.

Clocks around the world coordinate time by comparing with each other using clock signals broadcast by GPS satellites [67]. Although GPS time signals are noisy, when two nations simultaneously compare their clocks with the same GPS clock, the difference between these comparisons cancels the GPS signal. This yields a comparison between the clocks of both nations that is free of GPS time signal noise.

1.5 Optical Clocks

Cs clocks are based on phase-stable microwave oscillators ($\omega \simeq 2\pi \times 9.2$ GHz). However, with the advent of phase-stable lasers [45], clock oscillators with frequencies of hundreds of THz are now possible. Therefore, state-of-the-art atomic clocks are based on *clock lasers*, which are stable lasers referenced to electronic clock transitions in atoms [93].

Advances in atom cooling and trapping also provide advantages for optical clocks and frequency standards. Cold neutral atoms and atomic ions can be confined in (approximately) harmonic potentials, where $k_B T < \hbar\omega_{tr}$, with T the atomic temperature and $\hbar\omega_{tr}$ the harmonic oscillator energy spacing of the trap. This condition ensures that atoms occupy few motional states of the potential. In this regime, it is possible to interrogate clock transitions in a manner that is free of Doppler and recoil shifts, which cause systematic uncertainties in Cs clocks (since Cs clocks are based on atoms in free flight).

Optical frequency standards are now a mature technology, outperforming standards based

on microwaves. To realize an optical clock, one must also be able to count the cycles of optical clock oscillators. Historically this was a difficult task, but in the 2000s this became possible with self-referenced optical frequency combs [31]. The clock uncertainty introduced by the frequency comb counter is currently negligible [97], so to make better clocks we seek to improve frequency standards.

Two kinds of optical clocks have surpassed Cs: single-ion clocks and optical lattice clocks. Ion clocks are less sensitive than lattice clocks to systematic effects that cause oscillator frequency uncertainty. However, since lattice clocks use $\sim 10^3$ atoms whereas ion clocks use only one, lattice clocks have the advantage in terms of precision. This is because the standard quantum limit to clock precision scales as $1/\sqrt{N}$, where N is the number of atoms in the clock. Until the work described in this thesis, ion clocks held the record for lowest total systematic uncertainty for a number of years.

1.6 State of the Art Frequency Standards

When we talk about one frequency standard outperforming another, it is important to be clear about how performance is gauged. To that end, it is useful to explain precision, systematic uncertainty, and accuracy in the context of atomic frequency standards.

The precision of a frequency standard is also known as its *stability*, which in this case is the uncertainty in the mean frequency due to clock oscillator phase or frequency noise, $\delta\nu$ (also called the “statistical uncertainty”). Although precision gets worse with increasing $\delta\nu$, clock performance improves with larger values of the oscillator base frequency $\nu = \omega/2\pi$; therefore, a good measure of a frequency standard’s performance is the stability in fractional units.

Systematic uncertainty arises because atomic frequency standards are typically based on bare transition frequencies of atoms. Bare transitions are useful because their frequencies are universal, yielding international agreement between frequency standards based on the same atom. In practice, clock researchers must work perturbed transitions, which vary from laboratory to laboratory. Real measurement environments have stray electric and magnetic fields, heat, and atomic interactions

that introduce perturbations; however, it is possible to measure these systematic biases and then correct for them. The uncertainties in measurements of these biases (or *systematic shifts*) are the systematic uncertainties of the frequency standard. Adding all of these uncertainties in quadrature yields the standard’s *total systematic uncertainty*, which is another important measure of a frequency standard’s performance. Systematic shifts are often measured by studying small changes in the clock oscillator frequency, so these measurements are negatively affected by poor stability. As explained in Section 1.1, good precision often lends itself to good systematic uncertainty.

The accuracy of an atomic frequency standard is the degree to which its quoted frequency conforms to the true number of oscillations per SI second of the atom’s bare frequency.⁵ For a standard that is not based on Cs, accuracy can be achieved by comparing the standard’s frequency with that of the Cs transition. This works fine for standards that have comparable or worse performance than that of Cs systems, but for optical standards with total systematic uncertainties that are orders of magnitude better than Cs, this is no longer a viable option. Any attempt to compare optical and Cs standards would be severely limited by the Cs systematic uncertainty.

To improve upon the SI, timekeeping, and modern technology, there has been much discussion in the atomic clock research community of redefining the SI second in terms of an optical frequency [39]. To further this eventual goal, this thesis describes the first lattice clock to surpass all ion clocks and take the record for best clock performance, achieving stability among the best ever reported and demonstrating the lowest total systematic uncertainty to date.

1.7 Thesis Overview

This thesis describes the stability characterization and systematic evaluation of an optical lattice clock based on a stable laser referenced to a transition in atomic strontium.⁶ This is the second Sr lattice clock built by our team at JILA. I begin by discussing the physics of the

⁵ A corollary to this is that the Cs primary frequency standard has perfect accuracy within its systematic uncertainty

⁶ This work will not discuss our frequency comb counter in detail, focusing instead on the strontium frequency standard. This is because the counter does not limit clock performance, so any improvements to the clock must come from the frequency standard. For this reason, this thesis often uses the word “clock” to refer to the strontium apparatus even though a discussion of the counter will be notably absent.

strontium lattice clock, covering cooling, trapping, and spectroscopy of this atom. Next, I describe the second-generation lattice clock apparatus (Sr2) that my labmates and I built during my time on the JILA strontium team. Then, I discuss the stability of the Sr2 frequency standard, which was improved upon over previous results from Sr1 using a more stable laser. Lastly, I detail the results of two systematic evaluations of Sr2. The initial evaluation was the first instance of a lattice clock surpassing an ion clock, and the second evaluation reached a longstanding goal in the optical clock community of low 10^{-18} level total systematic uncertainty.⁷

⁷ Fractional frequency units will be used throughout this thesis for both stability and total uncertainty. “Stability” will refer to the statistical frequency uncertainty divided by the Sr clock transition frequency of 429.228 THz. “Total systematic uncertainty” (or “total uncertainty” for short) will refer to the combined systematic frequency uncertainty divided by the clock transition frequency.

Chapter 2

The Physics of the Strontium Lattice Clock

2.1 The Sr Atom

The strontium lattice clock is based on laser cooled Sr atoms trapped in an optical lattice, which is a periodic potential generated by laser light. The Sr atom has four stable isotopes (Table 2.1), one of which is a fermion and the rest are bosons. Like other alkaline earth (Group II in the Periodic Table) or alkaline-earth-like atoms, the bosonic isotopes all have zero nuclear spin.

Sr has two valence electrons, resulting in spin singlet and triplet angular momentum states. This work often utilizes transitions between the $5s^2 \ ^1S_0$ ground state and the $5s5p \ ^3P_0$ or $5s5p \ ^3P_1$ levels¹ (Figure 2.1), both of which are forbidden by electric dipole selection rules. In the case of the $^1S_0 \rightarrow ^3P_1$ line, which violates the $\Delta S = 0$ selection rule (where S is the total electron spin quantum number), angular momentum state mixing due to the spin-orbit interaction allows for a weak electric dipole *intercombination* transition [16].

In the bosonic isotopes of Sr, the $^1S_0 \rightarrow ^3P_0$ transition is still electric dipole forbidden,

¹ Here I use Russell-Saunders notation, in which angular momentum states are written as $^{2S+1}L_J$, where S is the total electron spin, L is the orbital angular momentum, and J is the total angular momentum. Since there are two valence electrons, $S = 0$ or 1 , where the $S = 0$ case is referred to as a “singlet” state and $S = 1$ is the “triplet.”

Isotope	Abundance	Nuclear Spin
⁸⁴ Sr	0.56%	0
⁸⁶ Sr	9.86%	0
⁸⁷ Sr	7.0%	9/2
⁸⁸ Sr	82.58%	0

Table 2.1: Stable isotopes of strontium.

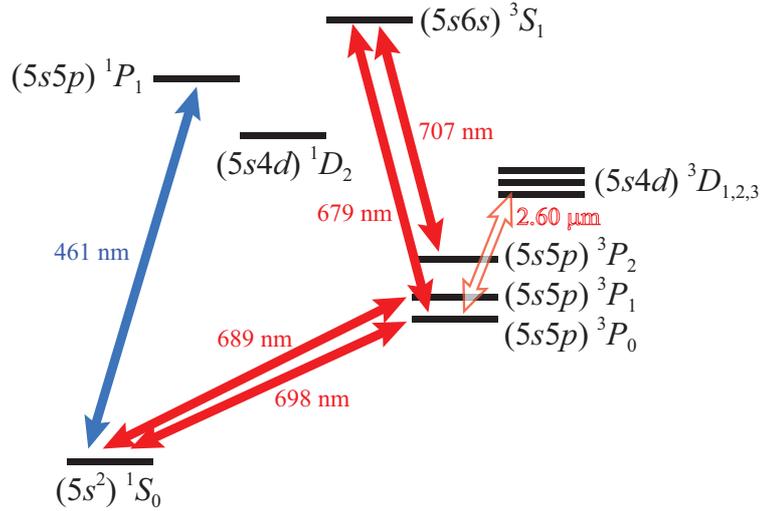


Figure 2.1: The Sr energy levels discussed in this thesis. All of the lasers used to access these levels are represented by arrows.

and theory predicts that, for all practical purposes, it never decays [107]. In ^{87}Sr , however, the hyperfine interaction (that occurs because of the nuclear spin of $9/2$) provides another state mixing mechanism that allows the $^1S_0 \rightarrow ^3P_0$ transition to take place [16]. The result is an ultranarrow ~ 1 mHz linewidth electric dipole resonance, which is used as the strontium clock transition.

The $^1S_0 \rightarrow ^3P_0$ line is a good clock transition because it is naturally insensitive to stray magnetic fields [16], which can cause unwanted systematic shifts of the clock transition. For $J = 0$ states, the transition's first-order Zeeman shift is given by the nuclear g -factor g_I , which is smaller than the Landé g -factor g_J (that determines the magnetic sensitivity of $J \neq 0$ states) by a factor roughly equal to the proton-to-electron mass ratio. The $^1S_0 \rightarrow ^3P_0$ second-order Zeeman shift is also an order of magnitude smaller than that of many prominent ion clocks and several orders of magnitude smaller than that of the Cs clock transition [16].

2.2 Laser Cooling Sr

The source of the Sr atoms is an effusive oven that emits a hot atomic beam. To remove enough kinetic energy from this fast moving beam to trap the atoms, laser cooling is employed. Experiments that work with ultracold Sr typically use two laser cooling transitions: $^1S_0 \rightarrow ^1P_1$

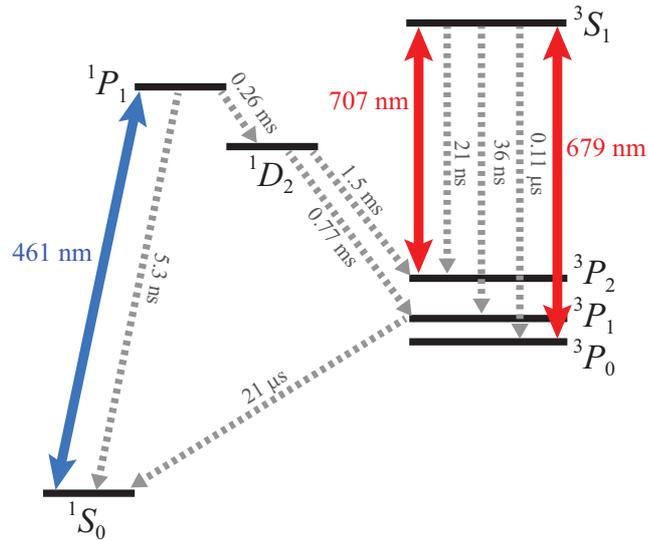


Figure 2.2: The repumping scheme. Lasers are represented as colored, solid arrows and decays are dashed arrows. The lifetimes of the decays are included next to the dashed arrows.

and $^1S_0 \rightarrow ^3P_1$. The $^1S_0 \rightarrow ^1P_1$ 461 nm transition is fully allowed by electric dipole selection rules and has a $\gamma = 30$ MHz linewidth. This transition is suitable for Zeeman slowing and 3D magneto-optical trap (MOT) cooling to a Doppler-limited temperature [81] of $h\gamma/2k_B = 0.72$ mK (where h is Planck's constant and k_B is Boltzmann's constant). The $^1S_0 \rightarrow ^1P_1$ transition is not completely closed since one in 4.9×10^4 decay events populates the 1D_2 state (Figure 2.2) [124]. Most of this population returns to the ground state through $^1D_2 \rightarrow ^3P_1 \rightarrow ^1S_0$; however, one in 1.4×10^5 decays out of 1P_1 results in shelving population in the metastable 3P_2 state via $^1P_1 \rightarrow ^1D_2 \rightarrow ^3P_2$.

To deal with this shelving, we use a 707 nm laser to drive the $^3P_2 \rightarrow ^3S_1$ transition. The 3S_1 state decays into the 3P manifold, so we also use a 679 nm laser to drive any population that decays into the long-lived 3P_0 clock state back into 3S_1 . Eventually all the shelved population returns to the ground state through $^3S_1 \rightarrow ^3P_1 \rightarrow ^1S_0$.

After the atoms have been cooled to their Doppler temperature on the $^1S_0 \rightarrow ^1P_1$ transition, we apply a second stage of 3D MOT cooling on the $^1S_0 \rightarrow ^3P_1$ intercombination line. This transition has wavelength of 689 nm and a linewidth of 7.5 kHz [32], which implies Doppler and recoil temperatures of 180 nK and 230 nK (respectively) [70]. In the ^{87}Sr intercombination line MOT,

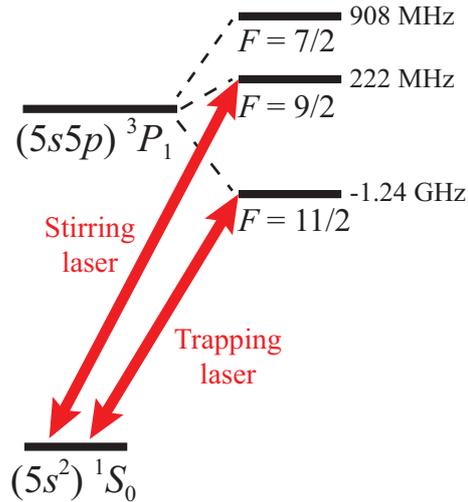


Figure 2.3: The three hyperfine states of 3P_1 . The hyperfine levels are labeled by the difference between the hyperfine transition frequency and the frequency of 3P_1 in absence of hyperfine structure. Using the language of Reference [86], the two lasers used in $^1S_0 \rightarrow ^3P_1$ cooling are the *trapping* and *stirring* lasers. These lasers are portrayed as red arrows.

complexities arise due to an interplay between the nuclear spin, the narrow linewidth, and the fact that the 1S_0 and 3P_1 levels have very dissimilar g -factors, leading to the lack of a trapping force for some of the atomic spin magnetic quantum states m_F [86]. This makes a traditional MOT based on a single $F \rightarrow F + 1$ hyperfine transition inefficient. A common solution to this problem is to use a *stirring laser* operating on the $|^1S_0, F = 9/2\rangle \rightarrow |^3P_1, F = 9/2\rangle$ transition to randomize m_F states along with the usual $F \rightarrow F + 1$ cooling laser driving the $|^1S_0, F = 9/2\rangle \rightarrow |^3P_1, F = 11/2\rangle$ transition [86]. After 500 ms of 461 nm MOT cooling and 260 ms of 689 nm MOT cooling, we typically obtain samples of ultracold ^{87}Sr at temperatures of a few μK .

2.3 The Magic Wavelength Optical Lattice

2.3.1 Optical Dipole Traps

Throughout both stages of laser cooling, an optical lattice is superimposed on the atom cloud trapped by the MOTs. The lattice potential collects the coldest atoms throughout the cooling process and retains them after the 689 nm MOT is shut off. Optical lattices utilize the ac Stark

energy shift E_{ac} to provide a conservative trapping potential. The ac Stark shift of quantum state n is given by [17]

$$E_{ac,n} = -\frac{1}{2}\alpha_n(\omega)|\vec{\mathcal{E}}_{ext}(\omega, t)|^2, \quad (2.1)$$

where $\alpha_n(\omega)$ is the atomic polarizability, $\vec{\mathcal{E}}_{ext}$ is an external electric field, and ω is the frequency of the external field. As Equation 2.1 shows, for ω chosen such that $\alpha_n > 0$, $E_{ac,n}$ is negative and it is energetically favorable for atoms to be positioned in the most intense regions of the electric field. Tightly focused laser beams can therefore confine atoms to small regions of space. Since lasers oscillate at frequencies of order 10^{15} Hz, which is much too fast to have a measurable effect on the trap dynamics, $|\vec{\mathcal{E}}_{ext}|^2$ in Equation 2.1 is replaced with its time average.

Lasers used for trapping are typically focused Gaussian beams [117],

$$\vec{\mathcal{E}}(r, z) = \vec{\mathcal{E}}_0 \frac{w_0}{w(z)} \exp \left[\frac{-r^2}{w^2(z)} - i(kz - \omega t) - ik \frac{r^2}{2R_c(z)} + i \tan^{-1} \left(\frac{z}{z_R} \right) \right], \quad (2.2)$$

where $|\vec{\mathcal{E}}_0|$ is the peak electric field amplitude, w_0 is the beam focal waist, $w(z) = w_0 \sqrt{1 + (z/z_R)^2}$, $R_c(z) = z [1 + (z_R/z)^2]$ is the beam's *radius of curvature*, $z_R = \pi w_0^2/\lambda$ is the *Rayleigh range*, $k = 2\pi/\lambda$, λ is the laser wavelength, and the unit vector $\vec{\mathcal{E}}_0/|\vec{\mathcal{E}}_0|$ is the laser polarization. This expression is written in cylindrical coordinates, with z the axial direction and $r = \sqrt{x^2 + y^2}$ the radial direction. The laser focus is located at $r = z = 0$.

The intensity of a harmonic complex electric field is given by [55]

$$I(r, z) = \frac{\epsilon_0 c}{2} (\vec{\mathcal{E}} \cdot \vec{\mathcal{E}}^*) = \frac{2P}{\pi w^2(z)} \exp \left(\frac{-2r^2}{w^2(z)} \right), \quad (2.3)$$

where the laser power $P = \int_0^{2\pi} \int_0^\infty r I(r, z) dr d\theta$, ϵ_0 is the permittivity of free space, c is the speed of light, and $*$ denotes complex conjugation. Expressing Equation 2.1 in terms of this laser intensity yields the potential for an *optical dipole trap*,

$$U_{ODT}(r, z) = -\frac{1}{2\epsilon_0 c} \alpha_n I(r, z) = -\frac{\alpha_n P}{\pi \epsilon_0 c w_0^2} \frac{\exp \left\{ \frac{-2r^2}{w_0^2 [1 + (z/z_R)^2]} \right\}}{[1 + (z/z_R)^2]}. \quad (2.4)$$

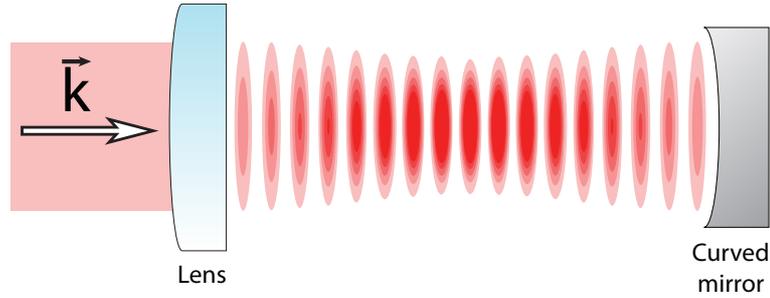


Figure 2.4: A retroreflected lattice. A laser traveling in the direction of \vec{k} is focused with a lens. A curved mirror retroreflects the light upon itself, generating periodic intensity fringes. These are used as an array of atom microtraps.

2.3.2 Retroreflected Lattices

Retroreflecting a trap laser upon itself generates a one-dimensional standing wave known as an *optical lattice*. A curved mirror is used to ensure that the beam foci of the incident and reflected beams overlap (Figure 2.4). This creates intensity fringes that are used as a periodic trap potential. Therefore, the electric field of the lattice is $\vec{\mathcal{E}}_{lat} = \vec{\mathcal{E}}(r, z) + \vec{\mathcal{E}}(r, -z)$, where $\vec{\mathcal{E}}(r, z)$ is the expression in Equation 2.2. This yields a trap potential U_{lat} of

$$U_{lat}(r, z) = -\frac{1}{4}\alpha_n \vec{\mathcal{E}}_{lat} \cdot \vec{\mathcal{E}}_{lat}^* = -\frac{4\alpha_n P}{\pi\epsilon_0 c w_0^2} \frac{\exp\left\{\frac{-2r^2}{w_0^2[1+(z/z_R)^2]}\right\}}{[1+(z/z_R)^2]} \cos^2\left(kz + \frac{kr^2}{2R_c(z)} - \tan^{-1}(z/z_R)\right). \quad (2.5)$$

Each intensity fringe provides axial confinement within a length of $\lambda/2$. It is common to make the approximations $\cos^2[kz + kr^2/2R_c(z) - \tan^{-1}(z/z_R)] \simeq \cos^2(kz)$ and $1 + (z/z_R)^2 \simeq 1$. For many experiments based on ultracold atoms in optical lattices (including Sr2 and our first generation Sr lattice clock, Sr1 [74]), these approximations are valid since, for typical beam parameters, the effects of the neglected terms are negligible over the dimensions of trapped atom samples.

Using these approximations, the lattice potential is

$$U_{lat}(r, z) = -U_0 e^{-2r^2/w_0^2} \cos^2(kz), \quad (2.6)$$

where $U_0 = 4\alpha_n P/\pi\epsilon_0 c w_0^2$ is the trap depth, defined as the difference between the asymptotic value

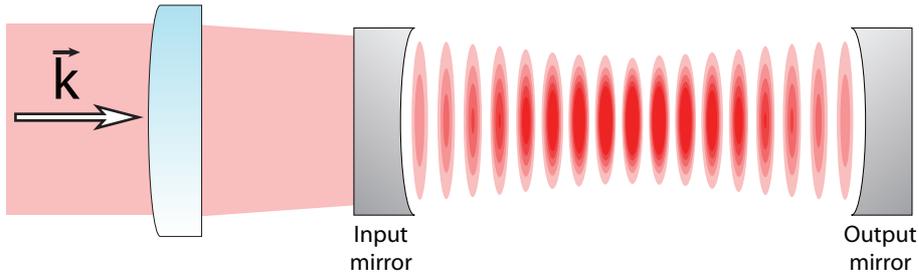


Figure 2.5: The cavity lattice. Incident light is coupled into the cavity with a lens, forming a standing wave in between the two mirrors.

of the potential and the potential minimum. Note that for the lattice, the trap depth is four times greater than in the optical dipole trap case due to constructive interference.

To trap large numbers of atoms, it is desirable to choose the laser power P such that U_0 is several times greater than $k_B T$, where T (again) is the sample temperature. In this regime, most of the sample is located near the potential minimum, where the trap is approximately harmonic. The *trap frequencies* are the radial and axial eigenfrequencies ω_r and ω_z at the minimum, which can be derived using normal mode analysis [41],

$$\omega_r^2 = \frac{1}{m} \left. \frac{\partial^2 U_{lat}}{\partial r^2} \right|_{r=0, z=0} = \frac{4U_0}{mw_0^2} \quad (2.7)$$

$$\omega_z^2 = \frac{1}{m} \left. \frac{\partial^2 U_{lat}}{\partial z^2} \right|_{r=0, z=0} = \frac{2k^2 U_0}{m}, \quad (2.8)$$

where m is the ^{87}Sr mass. Since w_0 is typically tens of microns, $\omega_z/\omega_r = kw_0/\sqrt{2} \gg 1$. The axial trap frequency is often expressed in terms of the lattice recoil energy $E_{rec} = \hbar^2 k^2 / 2m$ and the lattice recoil frequency $\omega_{rec} = E_{rec}/\hbar$,

$$\omega_z = 2\omega_{rec} \sqrt{\frac{U_0}{E_{rec}}}. \quad (2.9)$$

2.3.3 Cavity Lattices

The lattice employed by the Sr2 clock uses an optical cavity to enhance the lattice power. In this case, atoms are trapped in the fundamental transverse mode of a standing wave cavity

(Figure 2.5). The cavity is symmetric, meaning that the radii of curvature of the mirrors are equal, and atoms are trapped approximately in the center of the cavity. Assuming the cavity mirrors are separated by a distance L , we can derive the beam properties of the cavity lattice using the fact that, for a stable cavity, at $z = L/2$ the beam's radius of curvature $R_c(z)$ is equal to R (the radius of curvature of the mirror). $R_c(L/2) = R$ implies that the beam's Rayleigh range is

$$z_R = \frac{\pi w_0^2}{\lambda} = \frac{1}{2} \sqrt{L(2R - L)}, \quad (2.10)$$

which gives a cavity design equation for the beam waist w_0 . It is worth noting that Equation 2.10 is not valid for $L > 2R$ because in this case the cavity is unstable. It is straightforward to show with cavity stability analysis [117] that for a symmetric resonator, the criterion for a stable cavity is $0 < L < 2R$.

The cavity intensity inside and outside the cavity is given by [87]

$$I_{tr}/I_{in} = \frac{(1 - r_1)(1 - r_2)t}{(1 - t\sqrt{r_1 r_2})^2 + 4t\sqrt{r_1 r_2} \sin^2(kL)}, \quad (2.11)$$

$$I_{ref}/I_{in} = \frac{(\sqrt{r_1} - t\sqrt{r_2})^2 + 4t\sqrt{r_1 r_2} \sin^2(kL)}{(1 - t\sqrt{r_1 r_2})^2 + 4t\sqrt{r_1 r_2} \sin^2(kL)}, \quad (2.12)$$

$$I_{en}/I_{in} = \frac{(1 - r_1)t}{(1 - t\sqrt{r_1 r_2})^2 + 4t\sqrt{r_1 r_2} \sin^2(kL)}. \quad (2.13)$$

Here I_{in} is the incident laser intensity (taken to be constant in these equations), r_1 is the reflectivity of the input mirror, r_2 is the reflectivity of the output mirror, and t is the cavity loss (fraction of intensity lost after one pass through the cavity) from materials inside the resonator. The cavity is *resonant* when the round-trip phase accumulated by the electric field inside is equal to multiples of 2π , or equivalently kL is a multiple of π as in Figure 2.6. I_{tr} is the intensity transmitted through the cavity, and $I_{tr} = I_{in}$ when the cavity is resonant, $t = 1$ (no cavity loss), and $r_1 = r_2$. For $t < 1$, I_{tr} is always less than I_{in} . I_{ref} is the reflected intensity, which comes from the field reflected by the laser incident on the cavity as well as the field inside the cavity that leaks out of the input mirror. In the impedance matched case, the input mirror reflectivity is chosen such that $r_1 = r_2 t^2$. This

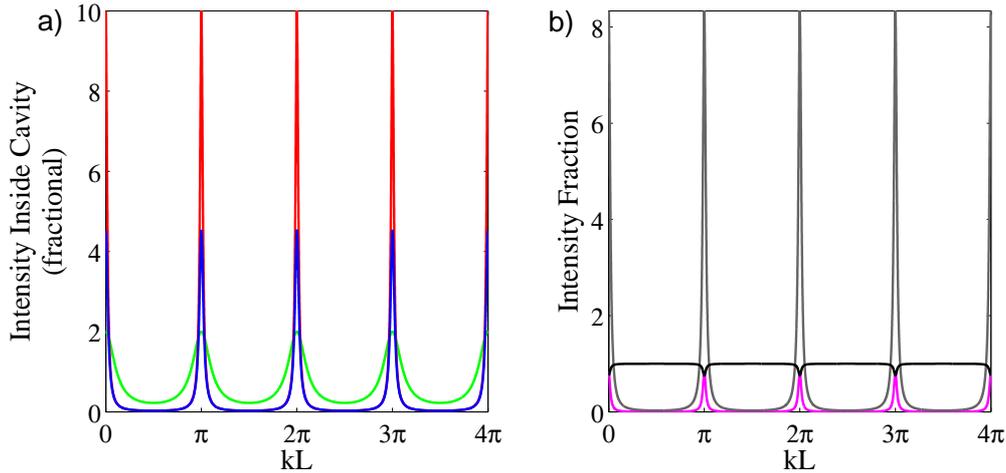


Figure 2.6: a) The intensity enhancement of an optical cavity. Plotted is I_{en}/I_{in} for the $r_1 = r_2 = r$ case. The green curve depicts the enhancement for $r = 0.5$ and no losses in the cavity ($t = 1$). For the red curve, $t = 1$ but r is increased to 0.9, which increases the enhancement and makes the intensity peaks narrower. In the blue curve, r remains at 0.9 but $t = 0.95$. Here the narrowness of the peaks remains but the intensity in the cavity is diminished. b) The cavity behavior for the $r = 0.9$ and $t = 0.99$ case. The gray curve is the one-way intensity I_{en}/I_{in} inside the cavity. The black curve is I_{ref}/I_{in} , showing the reflected intensity decrease when the cavity enhancement is maximal. The magenta curve is I_{tr}/I_{in} .

condition causes I_{ref} to vanish when the cavity is resonant, increasing the power inside the cavity.

Sr2 operates with mirrors from the same batch, so $r_1 = r_2$.

The enhanced intensity I_{en} is the motivation for using a cavity lattice. The standing wave inside the cavity is comprised of two counterpropagating electric fields, and the one-way intensity I_{en} is the intensity associated with one of those fields. When trapping atoms in this standing wave, both k and L are electronically controlled using Pound-Drever-Hall stabilization [10] to ensure that the cavity is resonant², which is necessary to achieve the maximum value of I_{en} (Figure 2.6).

The fact that light circulates several times in the cavity before leaking out means that the optical power stored in the cavity can be significantly larger than the incident power. Trapping atoms in a cavity-enhanced lattice allows us to leverage the extra power (compared to a retro-reflected lattice) for a larger beam waist w_0 while maintaining a good trap intensity and therefore a

² k is controlled by locking the laser to the cavity, and L is controlled by locking the cavity length with a piezoelectric actuator. See Section 3.4 for more technical details.

good value of the trap depth U_0 . This leads to a larger spatial extent of the trap potential, which proved to be critical for dealing with an important systematic frequency shift caused by atomic interactions (as I will explain in Section 5.5). According to Equation 2.13, the enhancement factor is

$$\frac{(1 - r_1)t}{(1 - t\sqrt{r_1 r_2})^2}. \quad (2.14)$$

To give a sense of the size of the enhancement, for $t = 0.99$ and $r_1 = r_2$, this factor can be as large as 25.

2.3.4 The Magic Wavelength

Although the ac Stark effect facilitates atom trapping, it also causes a frequency shift of the clock transition $\Delta\nu_{ac}$

$$\Delta\nu_{ac} = (E_{ac,3P_0} - E_{ac,1S_0})/h = -\frac{1}{2\epsilon_0 h c} \Delta\alpha(\omega) I(r, z), \quad (2.15)$$

where $\Delta\alpha$ is the *differential polarizability*

$$\Delta\alpha(\omega) = \alpha_{3P_0}(\omega) - \alpha_{1S_0}(\omega). \quad (2.16)$$

This shift is unwanted since the clock is based on the atom's bare frequency (Section 1.6). Therefore, we choose a lattice frequency ω at which $\alpha_{3P_0} = \alpha_{1S_0}$, causing the differential polarizability to vanish. This is called the *magic frequency*, and the corresponding wavelength is the *magic wavelength* [60, 57, 125].

The polarizability is given by [17]

$$\alpha_n(\omega) = 6\pi\epsilon_0 c^3 \sum_k \frac{A_{nk}}{\omega_{nk}^2 (\omega_{nk}^2 - \omega^2)}, \quad (2.17)$$

where the sum is performed over all hyperfine states except $k = n$. ω_{nk} is the frequency of a transition between hyperfine states n and k , and A_{nk} is the Einstein spontaneous decay coefficient

of this transition.

The polarizability is often expanded to make m_F dependence explicit,

$$\alpha_n = \alpha_{s,n} + \alpha_{v,n} \frac{m_F}{2F} \xi \hat{k} \cdot \hat{B} + \alpha_{t,n} \frac{3m_F^2 - F(F+1)}{2F(2F-1)} (3|\hat{\epsilon} \cdot \hat{B}|^2 - 1), \quad (2.18)$$

where $\alpha_{s,n}$, $\alpha_{v,n}$, and $\alpha_{t,n}$ are the *scalar*, *vector*, and *tensor* polarizabilities. Also, ξ is the *degree of ellipticity* of the lattice polarization, \hat{k} is the laser propagation direction, $\hat{\epsilon}$ is the direction of the lattice polarization (assumed to be nearly linear), and \hat{B} is the quantization axis of the atomic angular momentum. Using this form, $\Delta\nu_{ac}$ can be expressed as [121]

$$\Delta\nu_{ac} = - \left\{ \kappa_s + \kappa_v m_F \xi \hat{k} \cdot \hat{B} + \kappa_t [3m_F^2 - F(F+1)] (3|\hat{\epsilon} \cdot \hat{B}|^2 - 1) \right\} U_{lat}(r, z) \quad (2.19)$$

where the scalar, vector, and tensor shift coefficients κ_s , κ_v , and κ_t are proportional to their corresponding scalar, vector, and tensor polarizabilities. As explained in Section 2.6, the vector term is strongly suppressed in the Sr2 clock, so the desired wavelength is that at which the scalar and tensor terms cancel each other (Figure 2.7). The magic wavelength is approximately 813.4 nm [121].

2.4 Spectroscopy of Narrow Resonances in the Lamb-Dicke Regime

2.4.1 System Hamiltonian

To reference the clock oscillator to the Sr clock transition, this resonance will need to be interrogated. The fact that the atoms are trapped in an optical lattice means that the dipole approximation, which is often used for the spectroscopy of atoms, can no longer be made, leading to distinct qualitative differences compared to the free atom case.

To analyze this case, I will treat the electric field of the clock laser as $\vec{\mathcal{E}}_c = \vec{\mathcal{E}}_{c0} \cos(\vec{k}_c \cdot \vec{r} - \omega_c t)$, where $\vec{\mathcal{E}}_{c0}$ is a constant electric field amplitude, \vec{k}_c is the clock laser k -vector, \vec{r} is the position vector, and ω_c is the clock laser frequency. $\vec{\mathcal{E}}_c$ is linearly polarized because only π transitions

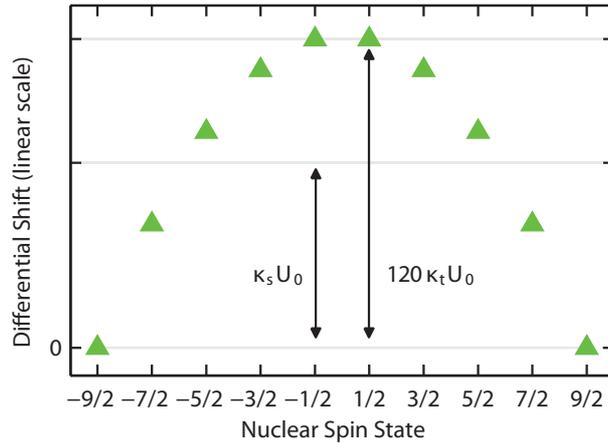


Figure 2.7: The differential ac Stark shift corresponding to the peak lattice intensity. Here $\hat{e} \cdot \hat{B} = 1$ (implying $\hat{k} \cdot \hat{B} = 0$). The scalar and $m_F = \pm 1/2$ tensor terms are included on the plot. The shift vanishes for the angular momentum stretched states $m_F = \pm 9/2$, which are used in clock operation.

(transitions that do not change m_F) are used for clock spectroscopy. This expression ignores the spatial properties of the clock laser beam, which is valid because this beam is typically made several times larger than the atom sample density to ensure uniform intensity across the sample.³ Also, the trapping potential will be treated as a harmonic oscillator. The interaction Hamiltonian H_{int} is $-\vec{d} \cdot \vec{\mathcal{E}}_c$, where \vec{d} is the dipole operator. Transforming to a frame rotating with the laser frequency and using the rotating wave approximation,

$$H_{int} = \frac{\hbar\Omega}{2} \left(e^{i\vec{k}_c \cdot \vec{r}} |e\rangle\langle g| + e^{-i\vec{k}_c \cdot \vec{r}} |g\rangle\langle e| \right). \quad (2.20)$$

Here $|g\rangle$ and $|e\rangle$ are the ground and excited electronic state vectors, $\Omega = -\langle e|\vec{d}|g\rangle \cdot \vec{\mathcal{E}}_{c0}/\hbar$ is the Rabi frequency, and $\vec{d} = e\vec{r}$ (where e is the electron charge).

As explained in Section 2.4.3, it is desirable to align the clock laser along the direction with the largest trap frequency, which is the lattice axial direction \hat{z} (also referred to as the “lattice axis”). In this case, $\vec{k}_c \cdot \vec{r} = k_c z$. Since the atoms are trapped in a harmonic potential, z should be regarded as an operator $z = \frac{z_0}{\sqrt{2}}(\hat{a} + \hat{a}^\dagger)$, where $z_0 = \sqrt{\hbar/m\omega_z}$ is the harmonic oscillator length, and \hat{a} and \hat{a}^\dagger are the annihilation and creation operators. These act on the harmonic oscillator

³ This is done to ensure that the clock laser excites the atoms uniformly across the sample density.

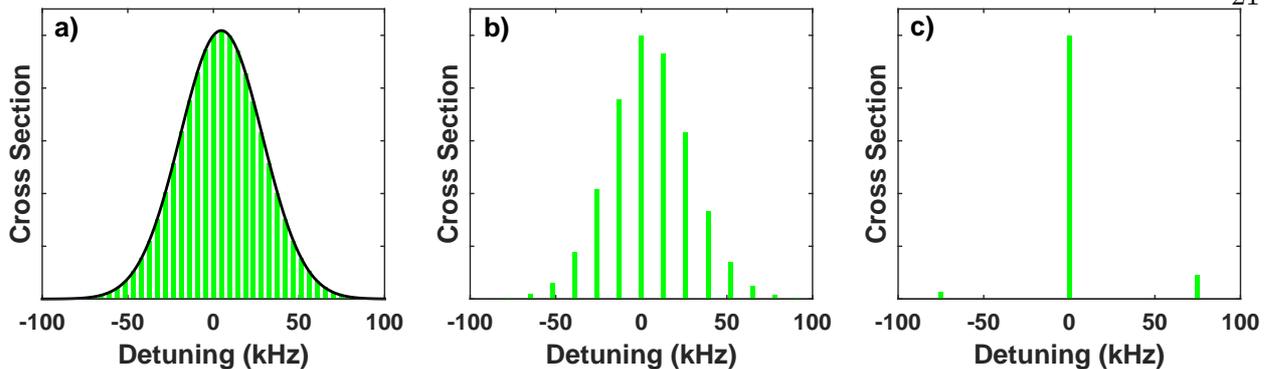


Figure 2.8: The normalized clock laser absorption cross section (from [123]) for weak excitation and $T = 3 \mu\text{K}$. The detuning is measured with respect to the free-atom resonance frequency in the atom's rest frame. a) The black line depicts free-atom motional shifts. This is a Doppler-broadened lineshape offset by the recoil shift $\nu_{c,rec}$. The green bars are the trapped-atom cross section for $\eta = 1$. Under this modest confinement, the effect of motional quantization begins to appear. b) For $\eta = 0.6$, motional transitions are well resolved, and each absorption event for nonzero detuning involves gaining or losing motional quanta. The frequency spacing between the motional transitions is ν_z , with the strongest transition now occurring at zero detuning instead of at $\nu_{c,rec}$. The trap frequency ν_z scales with $1/\eta^2$ (since $\eta = \sqrt{\nu_{c,rec}/\nu_z}$). c) For $\eta = 0.25$, sidebands are weak, and the zero detuning transition (or *carrier transition*), corresponding to no change of motional quanta, nearly dominates the cross section. Note that this carrier transition is not offset from zero detuning as in the free-atom case. Here the momentum from the clock laser photon is taken up by the lattice rather than the atom, resulting in recoil-shift-free spectroscopy [30].

state $|n_z\rangle$ of the lattice potential with the usual ladder operator behavior $\hat{a}|n_z\rangle = \sqrt{n_z}|n_z - 1\rangle$ and $\hat{a}^\dagger|n_z\rangle = \sqrt{n_z + 1}|n_z + 1\rangle$.

For many experiments, it is common to further simplify H_{int} with the *dipole approximation*, in which $e^{i\vec{k}_c \cdot \vec{r}} \simeq 1$. This is valid when $\vec{k}_c \cdot \vec{r}$ is much less than unity. For free atoms, the relevant length scale of the atoms' spatial extent is the Bohr radius a_0 , so $\vec{k}_c \cdot \vec{r} \sim k_c a_0 \ll 1$, and the dipole approximation holds. For atoms confined in the (approximately) harmonic potential of the lattice, the length scale is now the harmonic oscillator length z_0 , and $k_c z_0$ for an optical lattice is typically comparable to unity or greater. Therefore, the exponentials in Equation 2.20 must be kept.

The Hamiltonian of this system in the rotating frame (the same frame as Equation 2.20) is

$$H = -\hbar\Delta|e\rangle\langle e| + \hbar\omega_z \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) + \frac{\hbar\Omega}{2} \left(e^{i\eta(\hat{a} + \hat{a}^\dagger)} |e\rangle\langle g| + e^{-i\eta(\hat{a} + \hat{a}^\dagger)} |g\rangle\langle e| \right). \quad (2.21)$$

Here $\Delta = \omega - \omega_c$ is the detuning, $\eta = k_c z_0 / \sqrt{2} = \sqrt{\nu_{c,rec}/\nu_z}$ is the Dicke parameter, $\nu_z = \omega_z / 2\pi$,

$\nu_{c,rec} = h/2m\lambda_c^2 = 4.68$ kHz is the recoil frequency of the clock laser, and $\lambda_c = 698.4$ nm is the clock laser wavelength. This system is typically analyzed using the basis $|e, n_z\rangle = |e\rangle \otimes |n_z\rangle$ and $|g, n_z\rangle = |g\rangle \otimes |n_z\rangle$ [123]. In this case $|n_z\rangle$ is the z-directional motional quantum state. Although the trapped atoms also have motional excitations in the x and y directions, H in Equation 2.21 does not affect $|n_x\rangle$ or $|n_y\rangle$. Therefore, according to this model, the x- and y- directional excitations stay constant, so the trivial energy offset produced by these states is ignored. In this basis, an arbitrary quantum state $|\psi(t)\rangle$ can be expressed as $|\psi(t)\rangle = \sum_{n_z} [c_{g,n_z}(t)|g, n_z\rangle + c_{e,n_z}(t)|e, n_z\rangle]$, where the sum is taken from $n_z = 0$ to ∞ .

Spontaneous decay and laser decoherence are neglected in this treatment, implying that the spectroscopy of the clock states is fully coherent.⁴ Spontaneous decay can be neglected when the pulse time of the clock laser interrogating the atoms is much shorter than the spontaneous lifetime of the excited clock state. Also, laser decoherence can be neglected when the clock laser pulse time is much shorter than the coherence time of the laser. These approximations are valid in this work since the excited clock state is long lived and the clock laser is ultrastable.

2.4.2 Sideband Transitions

Analogous to the famous two-level Rabi solution, one can define a *sideband Rabi frequency* Ω_{eg} in terms of the matrix elements of the interaction Hamiltonian,

$$\frac{\hbar\Omega_{eg}}{2} = \langle e, n_e | H_{int} | g, n_g \rangle = \frac{\hbar\Omega}{2} \langle n_e | e^{-i\eta(\hat{a} + \hat{a}^\dagger)} | n_g \rangle, \quad (2.22)$$

where n_g and n_e are the motional state quantum numbers associated with the ground and excited electronic states. This expression has been evaluated [123] as

$$\Omega_{eg} = \Omega e^{-\eta^2/2} \sqrt{\frac{n_{<}!}{n_{>}!}} \eta^{|n_e - n_g|} L_{n_{<}}^{|n_e - n_g|}(\eta^2). \quad (2.23)$$

Here $n_{<}$ is the lesser of n_g and n_e , $n_{>}$ is the greater of the two, and $L_j^\alpha(x)$ is the j th generalized

⁴ Decay and decoherence could easily be treated with complex decay terms in the Hamiltonian of Equation 2.21 or with a Lindblad master equation and a Liouvillian superoperator to describe decay.

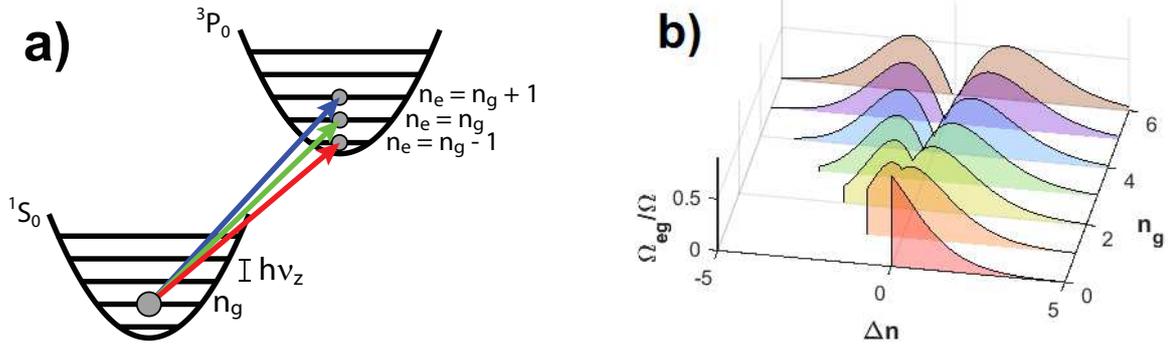


Figure 2.9: a) The carrier transition, for which $n_g = n_e$, is shown with a green arrow. It is also possible for atoms in the state $|g, n_g\rangle$ to be coupled to the excited electronic state when $n_e \neq n_g$, allowing for motional state addressing. The $n_g \rightarrow n_g - 1$ transition is the *red sideband*, and $n_g \rightarrow n_g + 1$ is the *blue sideband*. These resonances are separated in frequency from the carrier by ν_z . b) The sideband Rabi frequency for $\eta = 0.5$. Here $\Delta n = n_e - n_g$. Ω_{eg} is only defined at integer values of Δn , but the curves are plotted as continuous functions to make the behavior more apparent. At $n_g = 0$, atoms are initialized in the lowest possible motional state, so losing a motional excitation ($\Delta n < 0$) is impossible. Ω_{eg} begins to vanish for $\Delta n = 0$ as n_g gets large.

Laguerre polynomial. This expression describes the optical coupling for the $|g, n_g\rangle \rightarrow |e, n_e\rangle$ transition, which is nonzero even for *sideband transitions* in which the motional state changes when driving the electronic states (Figure 2.9). This is often interpreted as the Doppler shift when atomic motion is quantized [30] (Figure 2.8).

Solving the Schrödinger equation with the Hamiltonian from Equation 2.21 is considerably simpler in the *resolved sideband regime*. In this limit, sideband transitions are very well resolved from one another, making it possible to drive one sideband transition without exciting others.⁵

This is a good approximation for spectroscopy of the clock transition, which is very spectrally narrow. The system Hamiltonian in the resolved sideband approximation can be written in a 2×2 subspace

$$H/\hbar = \begin{pmatrix} -\Delta + \omega_z(n_e - n_g) & \frac{1}{2}\Omega_{eg} \\ \frac{1}{2}\Omega_{ge} & 0 \end{pmatrix}. \quad (2.24)$$

Equation 2.24 has an identical form to the familiar Rabi Hamiltonian, which describes the coherent

⁵ The resolved sideband regime occurs when $\gamma \ll \nu_z$, where γ is the full width at half maximum of the transition lineshape.

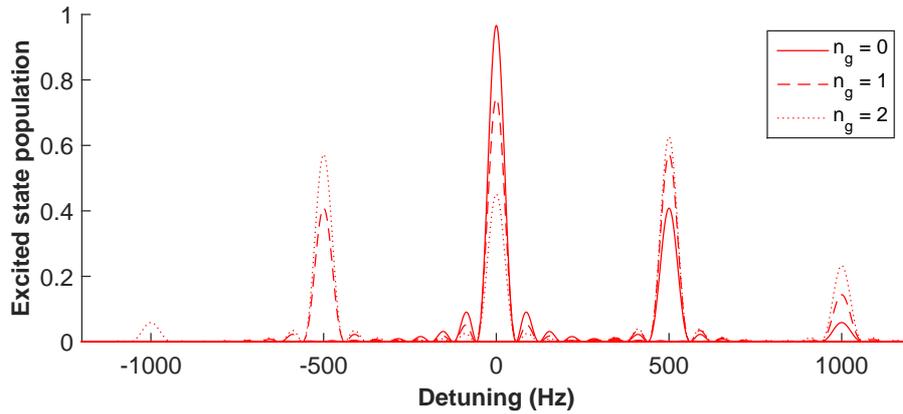


Figure 2.10: Sideband spectra for $\eta = 0.5$, $\nu_z = 500$ Hz, $\Omega = 100$ s⁻¹, and $t = \pi/\Omega$.

interrogation of a free atom at rest,

$$H_{Rabi}/\hbar = \begin{pmatrix} -\Delta & \frac{1}{2}\Omega \\ \frac{1}{2}\Omega^* & 0 \end{pmatrix}. \quad (2.25)$$

Solving the Schrödinger equation with H_{Rabi} for the excited state population fraction ρ_{ee} after a square pulse of length t ,

$$\rho_{ee} = \frac{|\Omega|^2}{|\Omega|^2 + \Delta^2} \sin^2 \left(\frac{t}{2} \sqrt{|\Omega|^2 + \Delta^2} \right). \quad (2.26)$$

Therefore, the excited state population fraction ρ_{n_e} implied by Equation 2.24 is

$$\rho_{n_e} = \frac{|\Omega_{eg}|^2}{|\Omega_{eg}|^2 + [\Delta - \omega_z(n_e - n_g)]^2} \sin^2 \left(\frac{t}{2} \sqrt{|\Omega_{eg}|^2 + [\Delta - \omega_z(n_e - n_g)]^2} \right). \quad (2.27)$$

Equation 2.27 describes transitions occurring for detunings $\Delta = \omega_z(n_e - n_g)$ (Figure 2.10).

2.4.3 The Lamb-Dicke Regime

The transition for which the motional state does not change ($n_g = n_e$) is a *carrier transition*. Carrier spectroscopy of lattice-trapped atoms has the advantage (compared to untrapped atoms) of being free of motion-related frequency shifts [30]. This is desirable since the clock is based on the unshifted bare frequency. However, at a sample temperature of a few μK , the Sr atoms may not

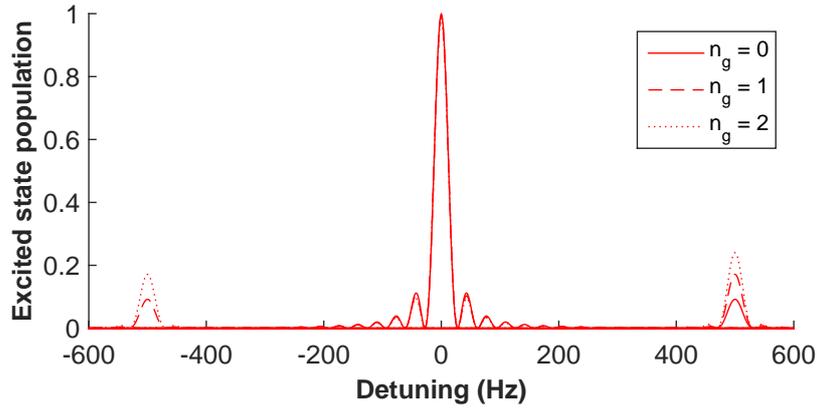


Figure 2.11: Sideband spectra for $\eta = 0.2$, with ν_z , Ω , and t the same as in Figure 2.10. The carrier transition is robust for different values of n_g . For $\eta \ll 1$, the sidebands would not be noticeable on this plot.

be entirely confined to the motional ground state of the lattice. As Figure 2.10 shows, the contrast of carrier transitions can degrade if one interrogates a gas of atoms that occupy different motional states.

To alleviate this problem, it is useful to work in the *Lamb-Dicke regime*. In this limit, the clock laser wavelength is much greater than the harmonic oscillator length along the lattice axis ($\lambda_c \gg z_0$, or equivalently $\eta \ll 1$). Here spectroscopy of the carrier transition is like that of an atom at rest, which is useful for an atomic clock based on a bare transition. The lineshape contrast is maximal, and Doppler and recoil shifts are absent. For typical Sr2 lattice parameters, $\eta = 0.2$, which is sufficient to reap these benefits (Figure 2.11).⁶ These advantages are not afforded to clocks based on free atoms (like Cs fountain clocks). Also, the magic wavelength discussed in Section 2.3.4 alleviates the main drawback to introducing the trap, namely the systematic ac Stark shift from the trap itself.

⁶ Probing the clock transition along the lattice radial direction, for which $\nu_r = \omega_r/2\pi = 120$ Hz, would yield $\eta = 6$. This is the reason for aligning the clock laser to the z direction.

2.4.4 Anharmonicity and Sideband Broadening

Since the Sr2 clock operates in the resolved sideband regime, measuring the frequency difference between the carrier and the red or blue sidebands (Figure 2.11) yields a value for ν_z . In fact, in the harmonic approximation, the frequency difference between the carrier and the red or blue sidebands is exactly ν_z . Since measurements of ν_z via the sideband frequencies will be important to the lattice Stark shift systematic measurement (Section 5.6), it is important to consider how these frequencies are related to ν_z in the true potential, which is anharmonic.

The motional part of the Hamiltonian in Equation 2.21 was approximated as harmonic, $H_{motion} = \hbar\omega_z(\hat{a}^\dagger\hat{a} + 1/2)$. Ignoring motion along the lattice radial direction, the full H_{motion} is

$$H_{motion} = \frac{p^2}{2m} - U_0 \cos^2(kz), \quad (2.28)$$

where $p = \frac{ip_0}{\sqrt{2}}(\hat{a}^\dagger - \hat{a})$, $z = \frac{z_0}{\sqrt{2}}(\hat{a} + \hat{a}^\dagger)$ as before, and $p_0 = \sqrt{m\hbar\omega_z}$. Expanding the potential in kz_0 ,

$$H_{motion} = \hbar\omega_z \left(\hat{a}^\dagger\hat{a} + \frac{1}{2} \right) - \frac{1}{12}\hbar\omega_z\eta_L^2(\hat{a} + \hat{a}^\dagger)^4 + \frac{1}{90}\hbar\omega_z\eta_L^4(\hat{a} + \hat{a}^\dagger)^6 + \dots, \quad (2.29)$$

and $\eta_L = kz_0/\sqrt{2} = \sqrt{\nu_{rec}/\nu_z}$ is the lattice Dicke parameter. The trivial potential offset $-U_0$ has been neglected. The blue sideband frequency ν_{bsb} (Figure 2.9) is therefore

$$\begin{aligned} \nu_{bsb} &= \langle n_z + 1 | H_{motion} | n_z + 1 \rangle / \hbar - \langle n_z | H_{motion} | n_z \rangle / \hbar \\ &\simeq \nu_z - (1 + n_z)\nu_{rec}. \end{aligned} \quad (2.30)$$

Terms of order η_L^4 and higher have been discarded since η_L is typically about 0.2. As Equation 2.30 shows, the blue sideband frequency is largest for atoms in the $n_g = 0$ state.

The effect of radial motion can be analyzed by expanding H_{motion} with the radial dependence included [12],

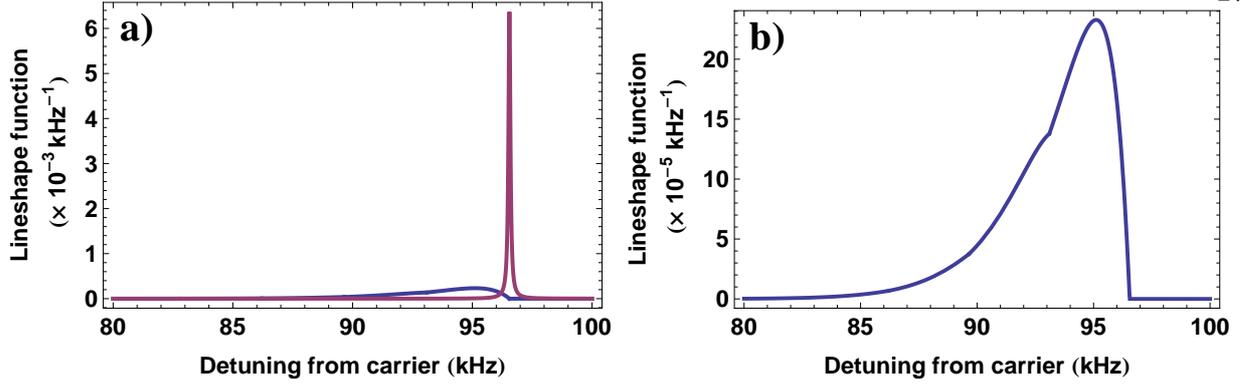


Figure 2.12: a) The unbrodened (purple) and broadened (blue) lineshape functions describing the blue sideband. These lineshape functions are area normalized. The carrier is located at zero detuning. b) The broadened lineshape function alone, showing broadening toward the carrier.

$$\begin{aligned}
 H_{motion} &= \frac{p^2}{2m} - U_0 e^{-2r^2/w_0^2} \cos^2(kz) \\
 &\simeq \frac{p^2}{2m} - U_0 \left(1 - k^2 z^2 + \frac{1}{3} k^4 z^4 - \frac{2r^2}{w_0^2} + \frac{2r^4}{w_0^4} + \frac{2k^2 z^2 r^2}{w_0^2} \right). \quad (2.31)
 \end{aligned}$$

Considering now the radial motional excitations n_x and n_y ,

$$\begin{aligned}
 \nu_{bsb} &= \langle n_x, n_y, n_z + 1 | H_{motion} | n_x, n_y, n_z + 1 \rangle / h - \langle n_x, n_y, n_z | H_{motion} | n_x, n_y, n_z \rangle / h \\
 &\simeq \nu_z - (1 + n_z) \nu_{rec} - \nu_{rec} \frac{\nu_r}{\nu_z} (n_x + n_y + 1). \quad (2.32)
 \end{aligned}$$

Since n_x and n_y do not change during spectroscopy, the r^2 and r^4 terms in Equation 2.31 do not contribute to ν_{bsb} ; however, the $r^2 z^2$ term that couples the radial and axial states results in the dependence on radial excitations. To obtain the correct sideband lineshape, we must thermally average over these states [12],

$$\mathcal{L}_{blue}(\delta) = \frac{2}{\pi \Gamma Z} \sum_{n_x, n_y, n_z} \frac{e^{-n_x h \nu_r / k_B T_r} e^{-n_y h \nu_r / k_B T_r} e^{-n_z h \nu_r / k_B T_z}}{1 + \frac{4}{\Gamma^2} (\delta - \nu_{bsb})^2}. \quad (2.33)$$

Here \mathcal{L}_{blue} is the lineshape function, Γ is the power broadened linewidth of a zero-temperature lineshape (modeled here as Lorentzian), T_r is the temperature in the radial direction of the lattice, T_z

is the axial temperature, and $Z = \prod_{x,y,z} \sum_i \exp(-n_i h\nu_i / k_B T_i)$ is the partition function. Although there are enough radial motional states in the potential to approximate the n_x and n_y sums as running from zero to ∞ , the sum over n_z should be taken from zero to the maximum number of axial excitations in the potential, which in this model is $\lfloor U_0 / h\nu_z \rfloor$ (where $\lfloor \dots \rfloor$ denotes the floor function).

This thermal averaging broadens the sideband asymmetrically toward the carrier (Figure 2.12). Physically, this broadening can be viewed in terms of atoms experiencing quantized motion in the axial direction ($h\nu_z > k_B T_z$) and classical motion ($h\nu_r \ll k_B T_r$) in the radial direction. Atoms rolling back and forth in the radial direction will experience different axial trap frequencies since the trap intensity they feel will vary depending on their radial position. The sharp sideband edge at large detunings from the carrier corresponds to when atoms are in the most intense region of the lattice laser ($r = 0$) and in the lowest axial motional state (Equation 2.30), resulting in the largest axial trap frequency. Therefore, in light of Equation 2.30, the frequency difference ν_{blue} between the blue sideband edge and the carrier is approximately

$$\nu_{blue} \simeq \nu_z - \nu_{rec}. \quad (2.34)$$

2.4.5 The Lattice in the Presence of Gravity

As I will explain in Section 5.6, the lattice ac Stark shift is studied as a function of the potential depth experienced by the atoms. This depth can be conveniently determined through $\nu_z = \nu_{blue} + \nu_{rec}$ (Equation 2.34), which is measured with sideband scans of the clock transition (Figure 3.19). According to Equation 2.9, $U_0 = (\nu_z / 2\nu_{rec})^2 E_{rec}$; however, this expression is only valid when gravity is ignored.

Gravity will change the sample's equilibrium position, moving the atoms out of the most intense region of the lattice. Since the position of the atoms is a function of U_0 , the lattice laser intensity at the location of the ultracold sample will not scale linearly with the laser power. Fortunately, the measured axial trap frequency can always determine the intensity at the location

of the atoms even in the presence of gravity. To demonstrate this, I analyze this situation with classical potential theory [41]. The lattice potential including gravity is

$$U_{grav} = -U_0 e^{-2x^2/w_0^2} \cos^2(kz) + mg(x \cos \theta + z \sin \theta), \quad (2.35)$$

where θ is the angle the lattice axial direction subtends with respect to the horizontal ($\theta = 16^\circ$ for Sr2). Here I take x , y and z to be the same coordinates as in Equation 2.6: z is the lattice axis, x makes an angle θ with respect to gravity, and y is taken to be perpendicular to both gravity and the lattice axis.⁷

The equilibrium coordinates x_{eq} and z_{eq} of U_{grav} are given by

$$U_x(x_{eq}, z_{eq}) = \left. \frac{\partial U_{grav}}{\partial x} \right|_{\substack{x=x_{eq} \\ z=z_{eq}}} = 0 \quad (2.36)$$

$$U_z(x_{eq}, z_{eq}) = \left. \frac{\partial U_{grav}}{\partial z} \right|_{\substack{x=x_{eq} \\ z=z_{eq}}} = 0. \quad (2.37)$$

When the two axes are coupled, the angular trap frequencies are the square roots of the eigenvalues of

$$\frac{1}{m} \begin{pmatrix} U_{xx}(x_{eq}, z_{eq}) & U_{xz}(x_{eq}, z_{eq}) \\ U_{xz}(x_{eq}, z_{eq}) & U_{zz}(x_{eq}, z_{eq}) \end{pmatrix}, \quad (2.38)$$

where U_{xx} , U_{xz} , and U_{zz} are second derivatives [41]. The equilibrium position of the atoms is affected enough by gravity to significantly diminish the potential depth at the location of the atoms (Figures 2.13a and 2.13b). However, as a function of ν_z , U_{lat} (Equation 2.6, which is the optical part of U_{grav}) evaluated at x_{eq} and z_{eq} agrees with the simple expression $(\nu_z/2\nu_{rec})^2 E_{rec}$ at the $0.1 E_{rec}$ level (Figure 2.13c).

According to Equation 2.9, in the absence of gravity $(\nu_z/2\nu_{rec})^2 E_{rec}$ is equal to $U_0 = 4\alpha_{1S_0} P / \pi \epsilon_0 c w_0^2$ (Equation 2.6), where P is the total trap power. As the above analysis shows,

⁷ y is ignored in this analysis because its contribution is trivial. The equilibrium position along y is always $y = 0$.

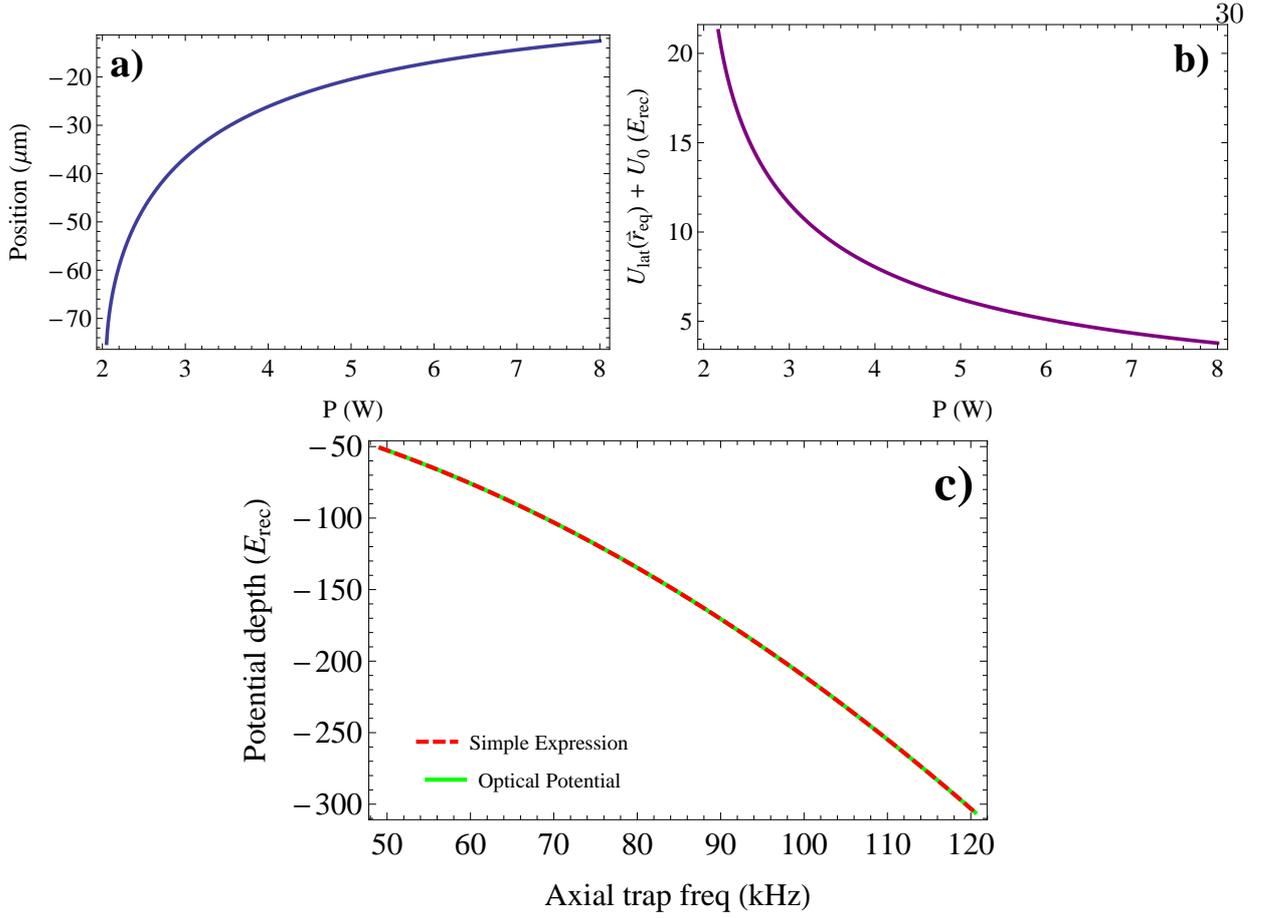


Figure 2.13: a) The lattice x direction as a function of the one-way lattice power. The z direction is only affected at the 0.1 nm level for the Sr2 experimental parameters $w_0 = 160 \mu\text{m}$ and $\alpha_{1S_0} = 286 \text{ a.u.}$ [105]. $x = z = 0$ is the laser focus where the atoms would be located without the influence of gravity. 8 W is the maximum lattice power our system can achieve. b) The difference between the potential depth at the equilibrium position of the atoms and the potential depth at the bottom of the trap (U_0) in units of E_{rec} as a function of one-way lattice power. The change is significant even for large lattice powers. c) The potential depth at the equilibrium position and the expression $-(\nu_z/2\nu_{rec})^2 E_{rec}$ as a function of ν_z .

when gravity is considered, $(\nu_z/2\nu_{rec})^2 E_{rec} = 2 \alpha_{1S_0} I_{eq}/\epsilon_0 c$, where I_{eq} is the laser intensity experienced by the atoms at their equilibrium position. Therefore, the potential depth at the location of the atoms is related to the blue sideband edge frequency ν_{blue} as

$$U_{atom} = \left(\frac{\nu_{blue} + \nu_{rec}}{2\nu_{rec}} \right)^2 E_{rec}. \quad (2.39)$$

2.5 Digital Servo to Atomic Spin Stretched States

2.5.1 Spin Polarization

The concerns about spectroscopy in a lattice raised in the previous section are addressed by working with a small value of the Dicke parameter η (Section 2.4.1) and operating in the resolved sideband regime. I now turn my attention to referencing the clock laser to the Sr clock transition. The free running clock laser suffers from frequency drift that would preclude its use as a frequency standard, which requires an accurate, stable frequency. Locking this laser to the clock transition removes this drift and ensures that the laser frequency matches that of the atoms trapped in the lattice. With the aim of providing an overview of how this lock works, it is important to first discuss magnetic effects on the clock transition due to interactions between external magnetic fields and the atom's spin.

Since $F = 9/2$ for both clock states, each state has the same 10 values of m_F . As explained in Section 2.1, the nuclear part of the clock transition g -factor determines its first-order Zeeman shift. The differential first-order Zeeman shift of the clock transition is

$$\Delta\nu_{Z1} = -\frac{\mu_B \delta g}{h} \langle F = 9/2, m_F | \vec{F} \cdot \vec{B} | F = 9/2, m_F \rangle = -\frac{\mu_B \delta g}{h} m_F B, \quad (2.40)$$

where μ_B is the Bohr magneton, δg is the clock transition differential g -factor, \vec{B} is an external dc magnetic field, and a π transition has been assumed. The g -factor has been measured [16] as $\mu_B \delta g/h = -108.4(4)$ Hz/G.

The clock transition Rabi frequency is also affected by the m_F state. Using the Wigner-Eckart theorem [17, 102] to make the m_F dependence explicit,

$$\begin{aligned} \Omega_{m_F} &= -\langle F = 9/2, m_F | \vec{d} | F = 9/2, m_F \rangle \cdot \vec{\mathcal{E}}_{c0} / \hbar \\ &= -\langle F = 9/2, m_F | F = 9/2, m_F, 1, 0 \rangle \langle F = 9/2 | \vec{d} | F = 9/2 \rangle \cdot \vec{\mathcal{E}}_{c0} / \hbar \\ &= -\frac{2m_F}{3\sqrt{11}} \Omega_{red}, \end{aligned} \quad (2.41)$$

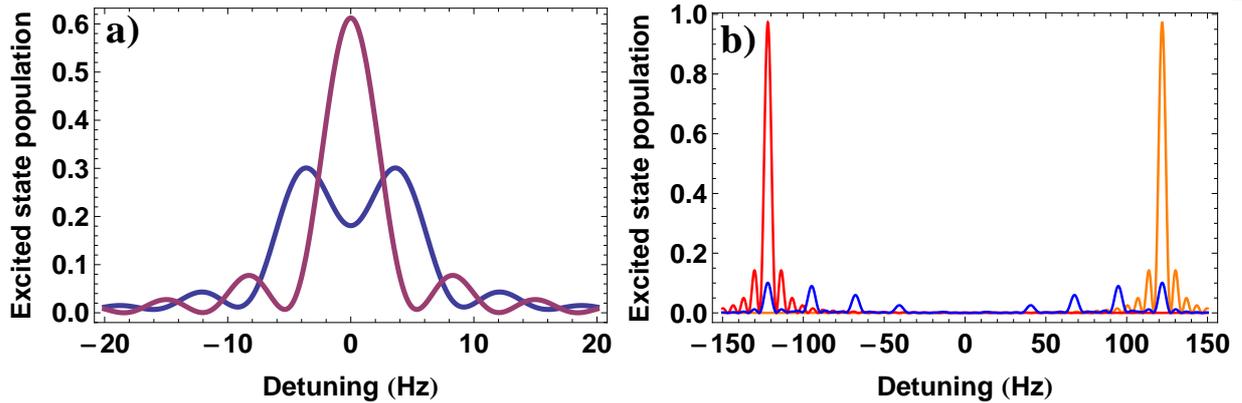


Figure 2.14: a) $\rho_{mixture}$ plotted $t = 160$ ms for $B = 0$ (purple) and $B = 10$ mG (blue). The latter is typical of stray laboratory magnetic fields. At zero field, the maximum achievable contrast is 0.61 (rather than unity, like in the case of the two-state population fraction ρ_{ee}). This is because the Rabi frequency Ω_{m_F} is different for each m_F state, each of which comprise 10% of the total population. At small fields, the lineshape broadens and loses contrast since the states are slightly split but unresolved. b) $\rho_{mixture}$ for $B = 250$ mG (blue) and spin-polarized lines (red and orange). At $B = 250$ mG, the ten π transitions are well resolved from each other. Polarizing the entire population into one of the two angular momentum stretched states, shown in red and orange, results in the Rabi lineshape of a two-state system.

where Ω_{m_F} is the Rabi frequency for the specific m_F state, $\langle F = 9/2, m_F | F = 9/2, m_F, 1, 0 \rangle$ is a *Clebsch-Gordan coefficient*, $\vec{\mathcal{E}}_{c0}$ is the clock laser electric field amplitude vector, $\langle F = 9/2 | \vec{d} | F = 9/2 \rangle$ is the *reduced dipole matrix element*, and Ω_{red} is the *reduced Rabi frequency* (Rabi frequency corresponding to the reduced matrix element).

Since state randomization during MOT cooling (Section 2.2) prepares the atoms in an incoherent mixture of m_F states with equal population in each state [86], the excited state population fraction after this cooling is

$$\rho_{mixture} = \frac{1}{2F+1} \sum_{m_F=-F}^F \frac{\Omega_{m_F}^2}{\Omega_{m_F}^2 + \Delta_{m_F}^2} \sin^2 \left(\frac{t}{2} \sqrt{\Omega_{m_F}^2 + \Delta_{m_F}^2} \right). \quad (2.42)$$

In this expression, ρ_{ee} from Equation 2.26 is evaluated for the appropriate Rabi frequency Ω_{m_F} and detuning $\Delta_{m_F} = \Delta + \mu_B \delta g m_F B / \hbar$ and then summed over. Compared to the two-state population fraction ρ_{ee} , typical stray magnetic fields of a few mG result in line broadening and contrast degradation of $\rho_{mixture}$ (Figure 2.14). This is undesirable since broader lines and poorer

contrast tend to worsen clock stability.

To deal with this problem, the Sr clock operates with a bias magnetic field B of about 300 mG to fully resolve the ten π transitions (Figure 2.14). The bias field direction \hat{B} , clock laser polarization \hat{e}_c , and lattice polarization \hat{e} are parallel. To take advantage of the best possible contrast, the atoms are spin polarized. Spin polarization is performed with an additional laser that is resonant with the $|^1S_0, F = 9/2\rangle \rightarrow |^3P_1, F = 9/2\rangle$ transition. The laser has circular polarization (with switchable helicity) and propagates in the direction of the bias field, driving σ transitions⁸ until atoms are polarized into one of the angular momentum *stretched states* ($m_F = \pm F$). The resulting population fraction is ρ_{ee} (Figure 2.14).

2.5.2 Locking to the Clock Transition

To reference the clock laser to the Sr clock transition, the laser is locked to the resonance center using a feedback loop (Figure 2.15). The excited state population fraction is measured on one side of resonance, then the clock laser is stepped by the full width at half maximum (FWHM) of the lineshape, and then the population fraction on the other side of resonance is measured. The difference between these population fractions is the error signal ϵ_{clock} , which the servo loop nulls by adjusting the laser frequency, locking the laser to the resonance center. This lock will be referred to as an *atomic servo*.

The FWHM γ can be changed by adjusting the pulse duration of the clock laser. The relation between γ and the probe time t is

$$\gamma \simeq \frac{0.8}{t}, \quad (2.43)$$

which is obtained by numerically finding the detuning at which ρ_{ee} in Equation 2.26 is equal to 1/2, assuming a π -pulse.⁹ When the FWHM of the measured lineshape obeys this equation, it is often referred to as the “Fourier-limited linewidth” for Rabi spectroscopy.

⁸ Transitions that change m_F by one

⁹ A π -pulse occurs when the Rabi frequency of the transition times t is equal to π , yielding the maximum signal contrast.

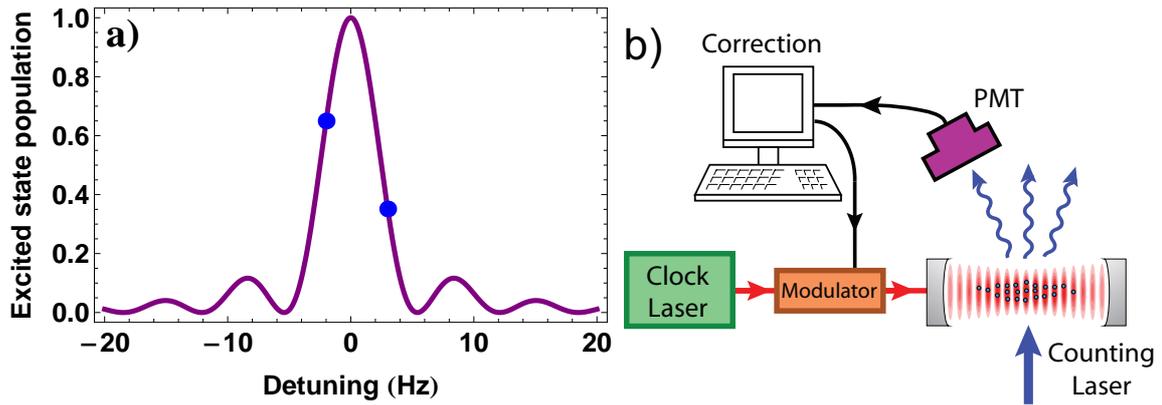


Figure 2.15: a) The clock transition is probed on the left side of resonance, driving the excited state population fraction to ρ_L (indicated by the left blue dot). The laser frequency is then stepped by the full width at half maximum of the transition, which is set using the clock laser probe time. The right side is probed, driving the population fraction to ρ_R (right blue dot). Measurements of these two population fractions yields an error signal $\epsilon_{clock} = \rho_L - \rho_R$. When the clock laser is detuned from the clock transition, $\epsilon_{clock} \neq 0$. The atomic servo adjusts the laser frequency until $\epsilon_{clock} = 0$, locking the laser to the resonance center. b) A simplified block diagram of the servo loop. The clock laser drives the atoms along the lattice axial direction. The counting laser fluoresces the ground state atoms, and some of the fluorescence is collected by the PMT. A computer stores the signal observed by the PMT, and once the population fraction has been measured on either side of resonance, the computer processes the correction needed to center the laser on resonance. The correction is made by varying an optical frequency modulator that changes the clock laser detuning in attempt to null ϵ_{clock} .

The atomic servo measures the population fractions ρ_L and ρ_R (Figure 2.15) destructively, meaning that each time either of these quantities are measured, the ultracold sample is lost. A new sample needs to be cooled and loaded into the lattice again, resulting in servo *dead time*, during which no information about the resonance center is added to the loop. For this reason, the atomic servo is digital. One cycle of preparing atoms, interrogating them with the clock laser, and probing the sample population takes about $T_c = 1$ s, where T_c is the *cycle time*. Corrections to the laser frequency f_{corr} are computed (Figure 2.15b) using a digital proportional-integral-differential (PID) filter,

$$f_{corr,n} = k_p \left[\epsilon_n + \frac{T_{ser}}{2T_i} (\epsilon_n + \epsilon_{n-1}) + \frac{T_d}{T_{ser}} (\epsilon_n - \epsilon_{n-1}) \right], \quad (2.44)$$

where ϵ_n is the n th value of $\epsilon_{clock} = \rho_L - \rho_R$ (Figure 2.15) for the data record; T_{ser} is the cycle

time¹⁰ of the servo; and k_p , T_i , and T_d are the proportional, integral, and differential gain constants, which are found empirically.

2.5.3 Digital Lock to Alternate Stretched States

So far, I have explained techniques for dealing with the broadening and reduced contrast that results from stray magnetic fields. Operating in a bias field ensures a nice Rabi lineshape rather than something broader. Spin polarization improves the signal-to-noise ratio of clock transition measurements. However, the bias field has the undesirable effect of frequency shifting the stretched states hundreds of Hz from their $B = 0$ values, which is an impediment for clocks based on a bare atomic resonance.

A *stretched-states servo* [112] is used to combat this problem. This technique involves four interrogations of the clock transition. Switching between the two stretched states, measurements are taken on either side of both resonances (Figure 2.16). The modulator that controls the laser frequency is an acousto-optic modulator (AOM), which shifts the frequency of the light passing through it by the frequency of a microwave drive injected into the device. For each measurement, the modulator drive frequency and the excited state population fraction are both recorded. With this information, the value of the modulator frequency needed to set the laser frequency on the resonance center is computed for both states using Equation 2.44. For each cycle of these four measurements, the ϵ_{clock} data is treated by the software as two independent atomic servo error signals (one for each stretched state).¹¹ This is accomplished by computing the correction for each stretched state only from data taken for that resonance.

The modulator frequency required to detune the clock laser to a given m_F state is

$$f_{m_F} = f_{Sr} + \frac{\mu_B \delta g}{h} m_F B, \quad (2.45)$$

where f_{Sr} is the modulator frequency that centers the laser on the zero-field resonance. With this

¹⁰ The servo cycle time depends on how it is run. For a single atomic servo running by itself, $T_{ser} = 2T_c$. For two servos updating in an alternate fashion, like in Section 2.5.3, $T_{ser} = 4T_c$.

¹¹ Note that the phrase “stretched-states servo” implies two atomic servos.

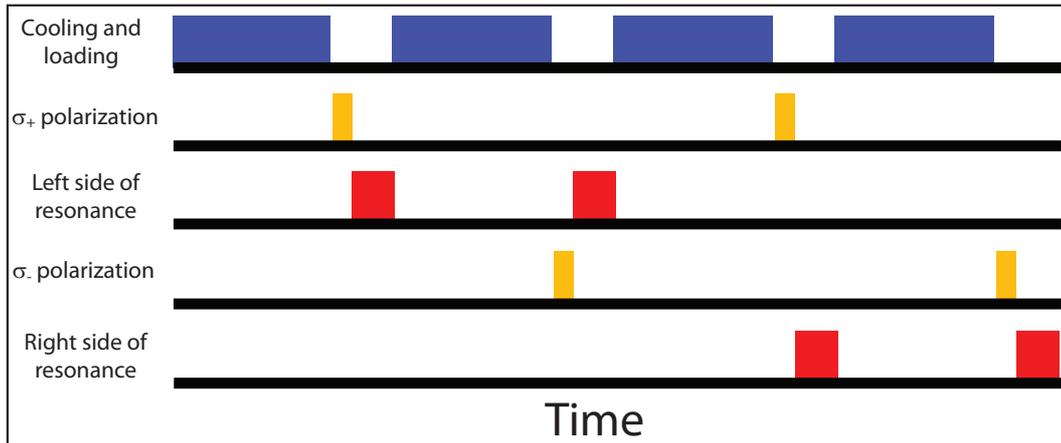


Figure 2.16: Timing diagram (not to scale) for the spin polarization and locking sequence. Atom cooling and lattice loading takes up most of one measurement cycle. In the first cycle, after the sample is prepared, atoms are spin polarized into the $m_F = F$ stretched state using σ_+ -polarized light. The clock transition is excited on the left side of resonance (Figure 2.15), and the excited state population fraction is recorded. In the next cycle, the helicity of the spin polarization laser is switched, polarizing atoms into the $m_F = -F$ stretched state. Again, the left side of resonance is excited, and the population fraction is recorded again. The third and fourth cycles are the same as the first two except that now the right side of resonance is measured.

stretched-state locking scheme, the data record of values for $f_{9/2}$ and $f_{-9/2}$ can be used to infer f_{Sr} as the *stretched state average frequency* $(f_{9/2} + f_{-9/2})/2$. This value of f_{Sr} can then be used to produce clock laser light oscillating at the zero-field resonance frequency.¹² These frequency data can also be used to infer the magnetic field through the stretched state frequency difference $f_{9/2} - f_{-9/2}$.

With the first-order Zeeman shift removed from the stretched state average frequency, this alternate stretched state scheme has done much to eliminate the problems associated with locking to an unpolarized resonance in the presence of stray magnetic fields. The one drawback is that the bias field creates a second-order Zeeman shift [16] that is not canceled upon averaging $f_{9/2}$ and $f_{-9/2}$. This shift is regarded as a systematic effect that must be evaluated to understand how much the clock laser frequency differs from the bare Sr frequency. This is a worthy trade off since the second-order Zeeman shift is one of the more manageable systematic effects in the JILA Sr clock.

¹² This can be accomplished by splitting off some of the clock laser light before it enters the modulator in Figure 2.15b and then sending this light through a second modulator, which has a frequency that tracks the computed value for f_{Sr} .

2.6 The Magic Wavelength Optical Lattice Revisited

With an understanding of the spin states involved in clock operation, the need for a bias field, and the bias field alignment, it is worthwhile to revisit the discussion of the magic wavelength optical lattice. Recalling Equation 2.19,

$$\Delta\nu_{ac} = - \left\{ \kappa_s + \kappa_v m_F \xi \hat{k} \cdot \hat{B} + \kappa_t [3m_F^2 - F(F+1)] (3|\hat{e} \cdot \hat{B}|^2 - 1) \right\} U_{lat}(r, z),$$

it is clear that the vector shift term is triply suppressed. Since \hat{e} and \hat{B} are aligned parallel to one another (Section 2.5.1), the value of $\hat{k} \cdot \hat{B}$ is minuscule, arising out of small imperfections in the alignment of \hat{e} and \hat{B} . The lattice is also linearly polarized, so ξ is also very small but nonzero due to imperfections in the polarization purity. Also, since the stretched-states servo cancels shifts linear in m_F , it diminishes the vector term as well. This strong suppression ensures that the vector shift does not affect the Sr clock.

Furthermore, since $m_F = \pm F$ for clock operation, the tensor term appears less ambiguous than before (Section 2.3.4) when I claimed that the magic wavelength causes the scalar and tensor terms to cancel (without elaborating on how this works given that the tensor term is different for each m_F state). The lattice Stark shift of the stretched-states servo is

$$\Delta\nu_{ac} = - [\kappa_s + 2\kappa_t F(2F-1)] U_{lat}(r, z). \quad (2.46)$$

Therefore, the desired magic wavelength is one for which the scalar and tensor shifts cancel when $m_F = \pm F$ and $\hat{e} \cdot \hat{B} = 1$.

2.7 Introduction to Stability and Systematic Effects

This chapter describes an ultracold gas of lattice-trapped Sr used as a frequency reference. The technique for referencing a laser to the clock transition has also been described. To characterize the performance of this system, the stability and systematic uncertainty must be determined (Section 1.6).

Clock stability is limited by the clock laser’s spectral noise. Because the servo loop described in Section 2.5.3 is digital, it has an inherent problem with aliasing, which degrades servo stability [28]. This problem is known as the “Dick Effect.”

Another important limitation to clock stability is the standard quantum limit, which, in the context of clocks, is called the “quantum projection noise limit” [54]. Since the clock is based on locking a laser to the half-maximum points of a Rabi lineshape (Figure 2.15), the quantum state of each atom is projected into either $|g\rangle$ or $|e\rangle$ with equal probability when the excited state population fraction is measured. Therefore, quantum projection noise (or QPN) is given by coin-toss (binomial) statistics. If several coins are tossed at once, the percentage that land on heads will fluctuate about 50% each time another toss is made. If the number of coins is increased, the percentage that lands on heads will fluctuate less. Similarly, more atoms leads to less QPN, which scales like $1/\sqrt{N}$, where N is the number of atoms.

As explained in Section 1.6, systematic uncertainty arises because the Sr lattice clock is based on the bare clock transition. Perturbations that shift this transition must be well understood and their uncertainties must be quantified so that the clock oscillator frequency can be appropriately corrected to the bare atom frequency. The uncertainty of the combined corrections is the total systematic uncertainty, which is the uncertainty in a clock’s ability to provide the bare atom frequency.

Most of the major systematic shifts that need to be considered for the Sr clock are electromagnetic perturbations. Despite operating the lattice at or near the magic wavelength, there can be a residual ac Stark shift from the lattice. Also, as explained in Section 2.5.3, the bias field creates a second-order Zeeman shift. A residual first-order Zeeman shift might also exist even though the stretched-states servo removes most of this effect. Patch charges or faulty electronics can cause dc electric fields that create clock transition Stark shifts. The clock laser itself causes an ac Stark shift. The ambient room-temperature heat in our lab is an electromagnetic field that causes a Stark shift. Furthermore, the clock transition is perturbed by atom-atom interactions.

There have been two evaluations of these systematic effects, as well as several other more

minor shifts, in the JILA Sr2 clock. However, before I begin a discussion of these measurements, I will describe the JILA Sr2 apparatus in the next chapter.

Chapter 3

The Sr2 Apparatus

3.1 The Sr2 Vacuum Chamber

3.1.1 Sr Oven

The different stages of the Sr2 clock measurement cycle were discussed in the previous chapter. To summarize, first a hot atomic beam from the Sr oven is slowed (with a Zeeman slower) and then trapped using a three-dimensional MOT. The first stage of MOT cooling operates on a dipole allowed transition at 461 nm, reducing the temperature of the atoms to about 1 mK. Next, the 461 nm MOT is shut off and a second-stage of three-dimensional MOT cooling begins. The second MOT operates on a singly forbidden (intercombination) transition at 689 nm, cooling the atoms to a few μK . These cold atoms are loaded into a one-dimensional magic wavelength optical lattice, and a clock laser interrogates the 1.5 mHz linewidth clock transition. After interrogation by the clock laser, the change in the Sr excited clock state population is measured. This measurement is processed by a feedback loop that corrects the clock laser frequency, locking the laser to the clock transition. This section describes the systems used to realize these measurement stages.

The hot atomic beam comes from a Sr oven that is based on a custom-made 2.75" vacuum nipple with a bend (Figure 3.1). The nipple contains solid Sr and a microchannel array to collimate the atomic beam. The oven is heated with two pairs of clamshell heaters. Homemade refractory ceramic fiber spacers keep the clamshells in place. Several thermocouples are attached to the bent nipple for temperature monitoring. For insulation, the oven is wrapped with a ceramic fiber

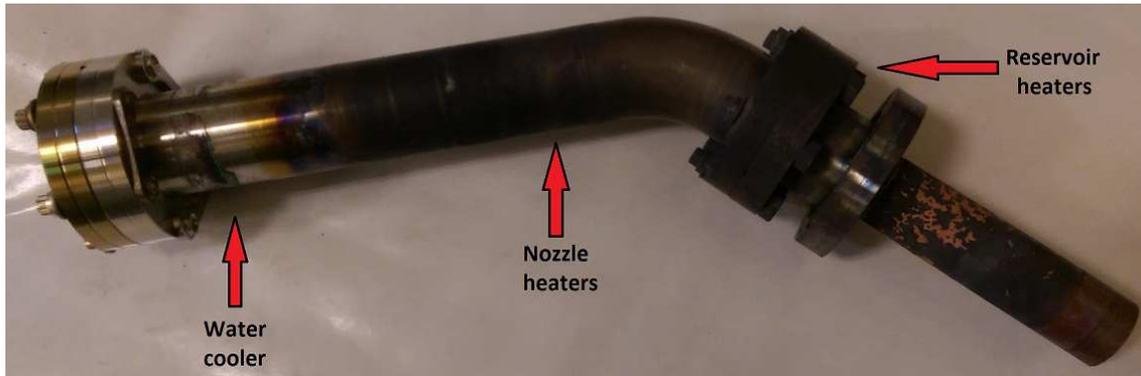


Figure 3.1: The Sr oven after many years of use. The clamshell heaters and water cooling block are not shown in this picture. One clamshell pair is wrapped around the conflat flanges closest to the bend (labeled “reservoir”). The strontium is contained near these flanges. Another clamshell pair is used to heat the “nozzle” region. The water cooling block is used where the oven joins with the rest of the chamber.

blanket, then a ceramic braid, and then aluminum foil. A water cooling block prevents heat from conducting away from the oven to other parts of the chamber.

The heaters are driven with ac current provided by variacs that output 60 ac volts. For each heater pair, both clamshells are wired in series, and their current can be gated with a solid state relay. Temperature controllers from Omega Engineering sense on two of the thermocouples and actuate on the solid state relays. Typically, the Sr reservoir is heated to 575 °C, and the nozzle (Figure 3.1) is heated to 625 °C.

3.1.2 Main Chamber

The Sr2 main chamber is an 8 inch extended spherical octagon from Kimball Physics. The atomic beam enters through one of the 16 1.33 inch conflats that makes a 16° angle with gravity. A pneumatic gate valve separates the main chamber region from the oven region (Figure 3.2). This valve is closed when the Sr2 clock is not running, protecting the main chamber region in case a leak were to occur in the oven region. Opposite the atomic beam is a sapphire viewport that is heated to 150 °C to prevent Sr deposition.

The remaining 14 1.33 inch conflats, the two 8 inch conflats, and the eight 2.75 inch conflats

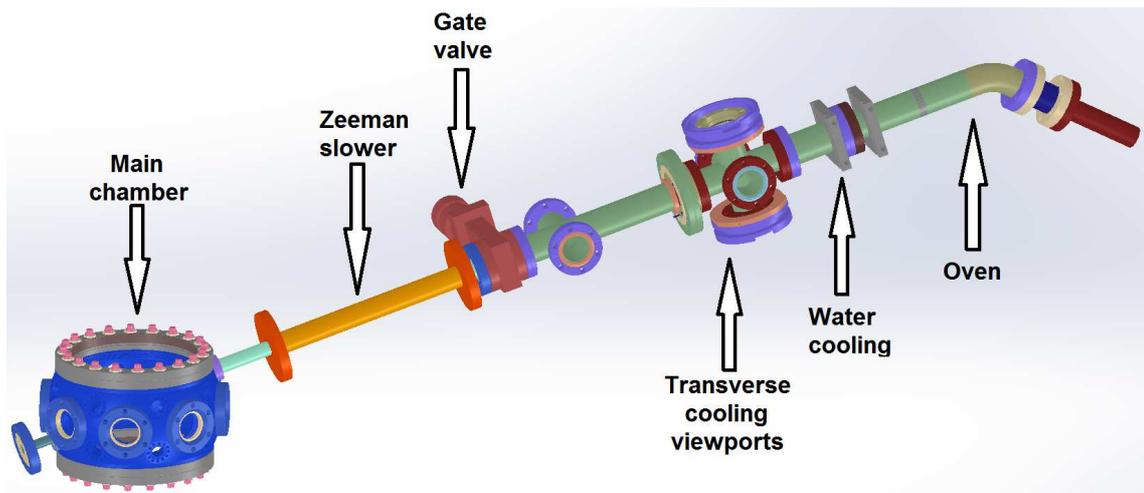


Figure 3.2: A CAD of the Sr2 vacuum chamber. A number of components—including pumps, coils, and valves—have been omitted to show the fundamentals of the chamber. The Zeeman slower is shown here as a simple nipple, around which the slower coils are wrapped in the actual machine. The Zeeman slower is joined with the main chamber using a flexible bellows. The side of the gate valve containing the main chamber is referred to in this work as the “main chamber region,” and the side of this valve with the oven is called the “oven region.” The oven region has two ion pumps as well as an all-metal valve for turbo and roughing pump access. The main chamber also has an all-metal valve for turbo and roughing pumps.

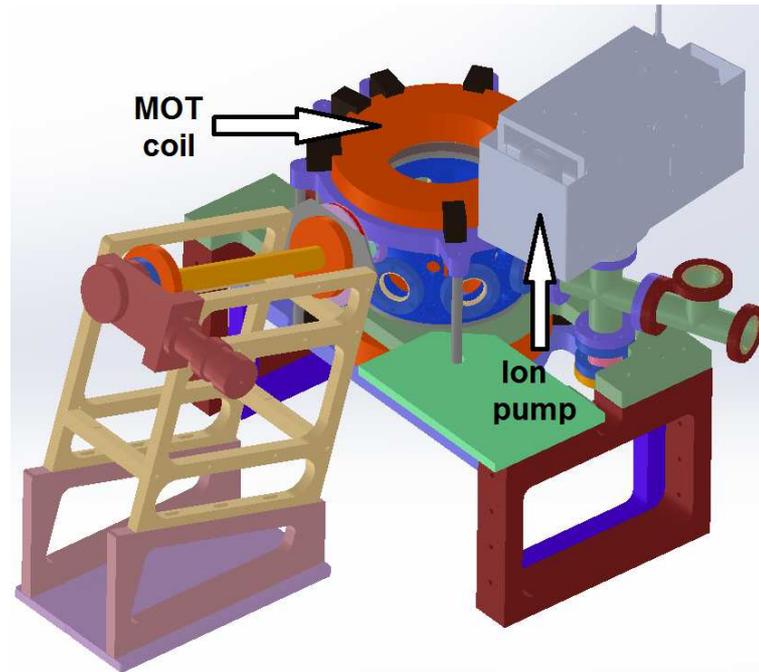


Figure 3.3: The main chamber region. The support structures and one ion pump are depicted. Opposite the four-way cross (on which the ion pump is attached) is a nude ion gauge. The top MOT coil can be seen as well. The two ion pumps in the oven region are not shown here.

are fitted with multiband-antireflection-coated fused silica viewports. The coating was performed by Research Electro-Optics and provides reflectivity at the 0.1% level for 461 nm, 689 nm, 698 nm, and 813 nm.

The main chamber has a 70 l/s ion pump (Figure 3.3). This configuration results in a pressure at the mid 10^{-10} torr level. At this pressure, the number of atoms in the lattice exponentially decays due to collisions with the chamber background gas with a time constant of a few seconds. This vacuum lifetime is sufficient for an atomic clock.

3.2 Cooling on the $^1S_0 \rightarrow ^1P_1$ Transition

3.2.1 461 nm Laser System

The first stage of cooling is based on the 461 nm $|^1S_0, F = 9/2\rangle \rightarrow |^1P_1, F = 11/2\rangle$ transition. The three hyperfine states of the 1P_1 level are unresolved [61], so cooling on this transition is

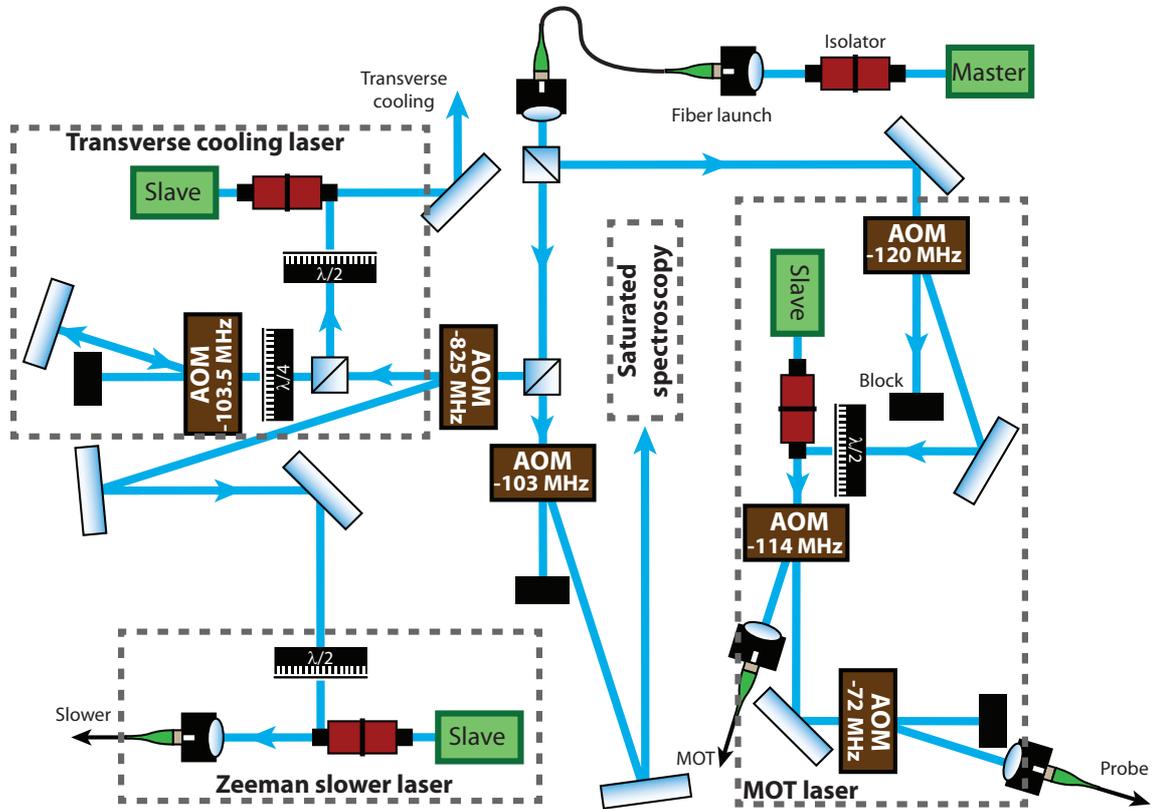


Figure 3.4: The 461 nm laser system. The master laser is a New Focus ECDL that is frequency stabilized using sub-Doppler saturation spectroscopy of a Sr discharge cell. A fiber is used to clean up the mode of the master laser. The AOMs here are labeled with the frequency shift they generate. All AOMs in this figure are aligned to provide light in their -1st order. The AOM before the saturated spectroscopy lock can be adjusted to a -155 MHz shift for ^{88}Sr trapping. The AOMs provide the detunings needed for the different cooling beams. The three slaves are injection locked by the master laser through the rejected light port out of the laser isolators. Light for the Zeeman slower and MOT is delivered to the Sr2 chamber by fibers. The slave laser for transverse cooling is sent to the chamber through free space.

affected by the entire 1P_1 hyperfine manifold. Sr2 had two generations of the laser system used for this cooling stage. In the first generation system, 922 nm light was frequency doubled to 461 nm using home-built linear doubling cavities based on periodically poled KTP. The 922 nm laser originated from an Eagleyard diode that seeded an Eagleyard tapered amplifier. This system suffered appreciable power degradation after a few hours of operation and was difficult to work with because of this. As we later learned, linear doubling cavities based on periodically poled KTP suffer from thermally induced dephasing, which compromises frequency doubling efficiency [68].

Fortunately, over the past few years, single frequency diode lasers at 461 nm became available. After encouragement from our research team, New Focus began selling external cavity diode lasers (ECDLs) at 461 nm, and Nichia also made an engineering sample of single frequency laser diodes at this color. Using this new technology, a second generation 461 nm laser system was constructed (Figure 3.4). This system uses a New Focus 461 nm ECDL referenced to the Sr $^1S_0 \rightarrow ^1P_1$ transition frequency with sub-Doppler saturation spectroscopy (Figure 3.5). Light from the New Focus ECDL injection locks three Nichia slave diodes, which are dedicated to the Zeeman slower, 461 nm MOT, and transverse cooling beams.

The Nichia diodes put out 100 mW of power. The New Focus laser originally emitted about 50 mW of power, but this number degraded to 35 mW over the course of a year and remained at this value ever since. Nevertheless, 35 mW is sufficient for our application.

3.2.2 Zeeman Slower and Transverse Cooling

70 mW of light from the Zeeman slower slave laser is delivered by a fiber to the Sr2 main chamber. Zeeman slowing is switched on and off with a mechanical shutter installed in the path of this slave laser (Figure 3.4) before the fiber that delivers it to the atoms. The slower is wound in *spin flip* or *zero crossing* configuration. It also has a compensation coil to cancel the slower field at the center of the main chamber where the atoms are trapped. After fiber delivery near the main chamber, the Zeeman slower laser passes through a $\lambda/4$ waveplate and enters the chamber through the heated sapphire window.

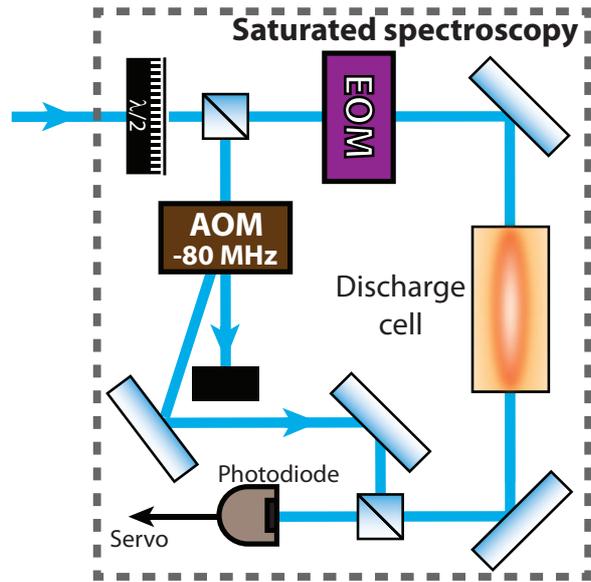


Figure 3.5: Saturated spectroscopy for referencing the New Focus master laser to the Sr $^1S_0 \rightarrow ^1P_1$ transition [87]. The photodiode output is mixed with the 40 MHz EOM frequency, and the output of this is sent to a lock-in detector referenced to the frequency at which the AOM intensity is chopped (90 kHz). The New Focus master laser external cavity piezoelectric transducer is locked to the resulting dispersive feature.



Figure 3.6: The Sr2 Zeeman slower. The right half of the slower has a zero crossing region. Coils to the left of the zero crossing have current flowing in the opposite direction from the coils on the right. The zero crossing configuration reduces the amount of wire needed for the slower, which reduces the resistive heating of the coils.

The output of the transverse cooling laser is split in two and sent into viewports mounted on six-way cross (Figure 3.2). One beam is periscoped up and sent through the horizontal direction of the cross. This light is retroreflected through the atomic beam to better counter balance the force in the horizontal direction. A similar retroreflected beam is present in the (nearly) vertical direction. When these transverse cooling lasers are well aligned, they boost the number of atoms trapped in the lattice by a factor of two.

3.2.3 Three-Dimensional MOT and Repump Lasers

A 3D MOT operating on the 461 nm transition is set up at the center of the main chamber. The slave laser designated for MOT power is delivered to the main chamber with an evanescent wave fiber-based one-to-three splitter from Evanescent Optics. The three output fibers are collimated to 2.5 cm diameter beams. The two horizontal MOT beams and one vertical beam are sent into the main chamber through two-color waveplates ($3\lambda/4$ for 461 nm and $\lambda/4$ for 689 nm), creating circular polarization. Each beam is retroreflected, double passing another two-color waveplate to provide the usual three pairs of σ_+ - and σ_- -polarized beams for the MOT [81].

The MOT light passes through an AOM before being fiber coupled to the main chamber (Figure 3.4). This AOM is useful for fast intensity switching and servoing. The counting laser, used for atom number counting (Section 2.5.2), is taken from the zeroth order of this MOT AOM. The counting laser has its own AOM to bring its detuning closer to zero.

The MOT coils are an anti-Helmholtz pair mounted on the top and bottom of our vacuum chamber (Figure 3.3). The coils are made of square-shaped copper tubing¹ that has an electrically insulating coating. Chilled water is flowed through the tubing to keep the coils from overheating, which would cause the electrical insulation to melt and short coil turns. The coils are driven by a 60 A power supply, and the coil current is controlled with a servo that actuates on a MOSFET bank. A Hall probe senses the coil current and feeds back through a PID filter to the gates of the MOSFETs. The MOSFETs are mounted on a cooling block to prevent overheating. The coils are

¹ By “square shaped” I mean that the cross section of the tubing is square shaped.

designed to provide a gradient at the center of the chamber of 60 G/cm for 50 A of current.

As explained in Section 2.2, repump lasers are needed to deal with shelving in the 3P_2 state. To target the many hyperfine states in the repumping process, the repump lasers are ECDLs that are frequency modulated by sweeping both the laser current and piezoelectric transducer (or “piezo”) voltage, spectrally broadening these lasers. These lasers are left free running; however, frequency drift is not an issue since drift over several hours of operation is insignificant compared to the large spectral widths of the modulated repump lasers.

With the Zeeman slower, transverse cooling, 461 nm MOT cooling, and repumping combined, this first stage of trapping and cooling takes 500 ms and results in roughly ten million atoms trapped at a temperature of about 1 mK.

3.3 Cooling on the $^1S_0 \rightarrow ^3P_1$ Transition

3.3.1 689 nm Laser System

As explained in Section 2.2, ultracold Sr is prepared using two stages of MOT cooling. The natural linewidth of the second cooling transition, $^1S_0 \rightarrow ^3P_1$, is 7.5 kHz. This is significantly narrower than the linewidth of ECDLs, which are typically of order 100 kHz. To ensure that the lasers used for the second stage of cooling provide sufficient power within the linewidth of the transition, the 689 nm laser system makes use of cavity stabilization.

This system is based on home-built ECDLs with diodes from Hamamatsu. These diodes were anti-reflection coated using the diode coating facilities at JILA. A master ECDL is stabilized to a cavity, resulting in a 300 Hz laser linewidth. To remove the effect of slow cavity drift, the master is frequency referenced to the $^1S_0 \rightarrow ^3P_1$ transition using saturated spectroscopy and a low-bandwidth servo loop (Figure 3.7). The trapping laser, operating on the $|^1S_0, F = 9/2\rangle \rightarrow |^3P_1, F = 11/2\rangle$ transition, and the stirring laser, operating on the $|^1S_0, F = 9/2\rangle \rightarrow |^3P_1, F = 9/2\rangle$ transition (Section 2.2), are slaved to the master with phase locks.

Each phase lock is based on one optical and one RF beat. The slave lasers are optically

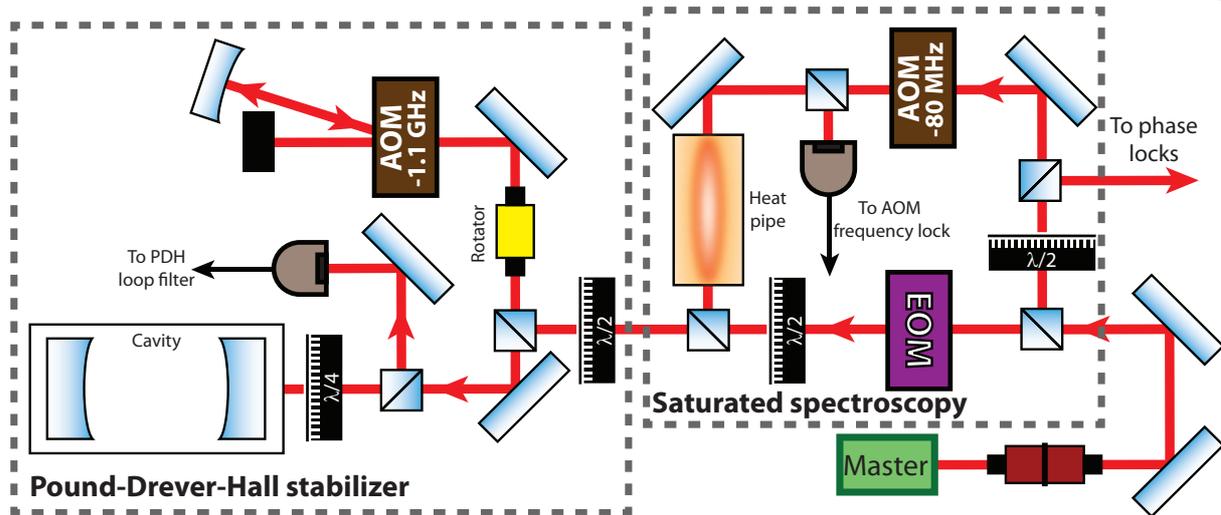


Figure 3.7: A home-built master EDCL is double passed through a 1.1 GHz AOM and stabilized to a cavity using the Pound-Drever-Hall technique [10]. The output of the Pound-Drever-Hall photodiode is mixed with the EOM frequency to achieve an error signal, to which the master laser current and external cavity piezo are locked. The cavity is under vacuum, temperature controlled, and has passive vibration damping. The stabilized laser linewidth is 300 Hz. The saturation spectroscopy setup is similar to Figure 3.5 aside from the fact that here the laser is locked to a heat pipe with Sr inside rather than a discharge cell. The photodiode output is mixed with the EOM frequency and then measured with a lock-in detector referenced to the AOM chopping frequency. A low-bandwidth loop filter takes the lock-in output as an error signal and feeds back onto the 1.1 GHz AOM frequency.

beat with the master, and the output of this beat on a photodiode is beat again with an RF frequency (Figure 3.8). The output of this second beat provides an error signal to which the current and external cavity piezos of the slave lasers are locked.² In this case, the phase lock mixers are intentionally saturated, providing a steeper and more linear slope of the error signal about zero. The frequency of each slave laser can be controlled using the synthesizer frequencies.

3.3.2 The 689 nm MOT

After the first stage of MOT cooling at 461 nm, the sample temperature is a few mK. At this temperature, the Doppler width of the $^1S_0 \rightarrow ^3P_1$ transition is roughly 1 MHz. To address all of the atoms, the slave lasers are spectrally broadened by frequency modulating the phase lock syn-

² This technique is sometimes referred to as an “RF offset lock.”

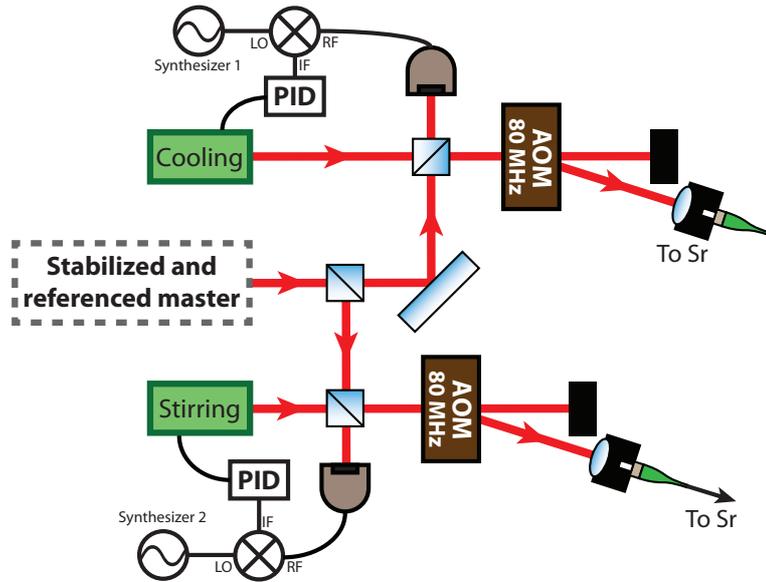


Figure 3.8: Two home-built ECDLs are phase locked to the stabilized and referenced master laser. The lock is performed by beating each laser with the master on a photodetector and then mixing the detector output with a synthesizer frequency. For both lasers, the stabilized output is sent to the Sr main chamber using fibers. AOMs are included for intensity stabilization and fast switching of each beam.

thesizers. Each synthesizer has an external frequency modulation (FM) port and a programmable frequency deviation. A triggerable two-channel arbitrary waveform generator controls the FM ports of the synthesizers.

The electric field amplitude of the frequency modulated laser \mathcal{E}_{mod} , assuming no phase noise, is [45]

$$\mathcal{E}_{mod} = \mathcal{E}_{0,mod} \sum_{n=-\infty}^{\infty} J_n(\beta) e^{-i(\omega+n\omega_{mod})t}, \quad (3.1)$$

with the auto correlation R_{mod}

$$R_{mod}(\tau) = \mathcal{E}_{0,mod}^2 \sum_{n=-\infty}^{\infty} J_n^2(\beta) e^{i(\omega+n\omega_{mod})\tau}. \quad (3.2)$$

Here $\beta = f_{dev}/f_{mod}$ is the modulation index, $\omega_{dev} = 2\pi f_{dev}$ is the frequency deviation, $\omega_{mod} = 2\pi f_{mod}$ is the modulation frequency, and J_n is a Bessel function of the first kind. For MOT cooling, we offset the broadened laser by the frequency deviation (Figure 3.10). Taking this into account,

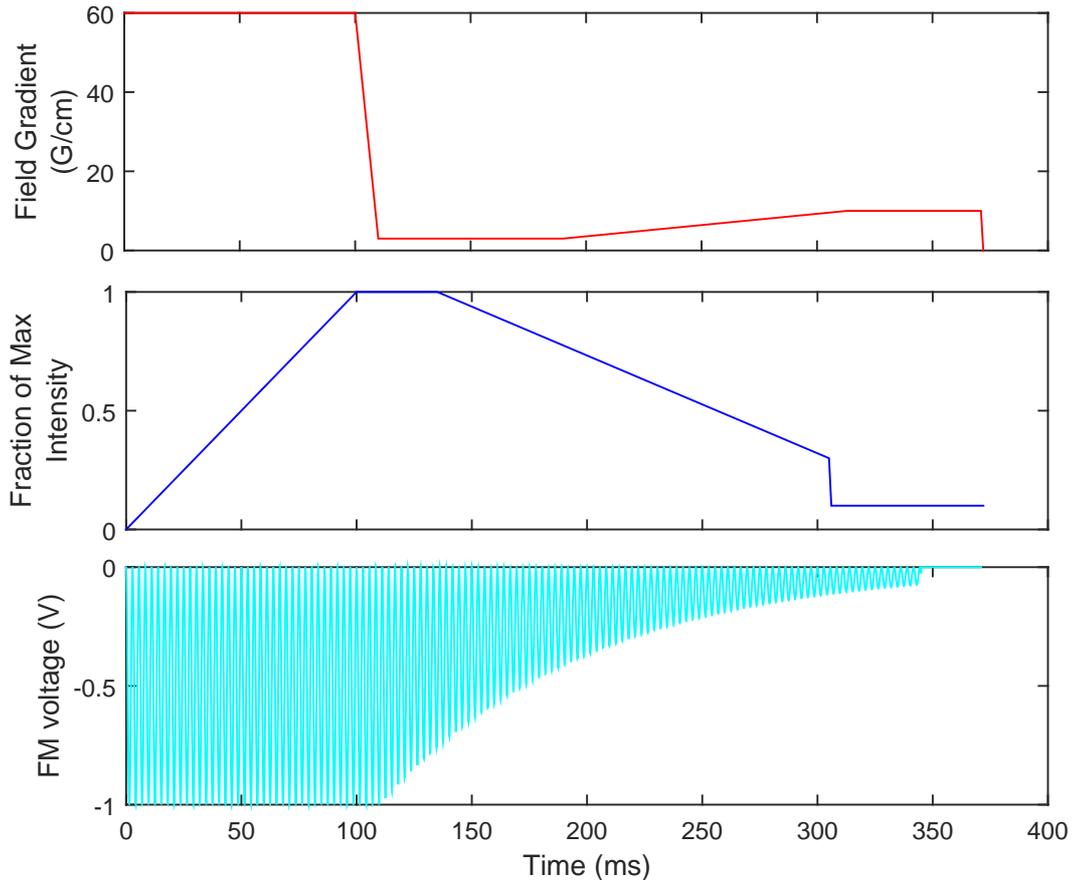


Figure 3.9: Timing diagram for intercombination line MOT cooling. The top curve (in red) is the gradient from the anti-Helmholz field. A current servo (Section 3.2.3) controls the gradient. The gradient ramps down from its 461 nm MOT value of 60 G/cm to 3 G/cm. After this ramp completes, the 461 nm MOT light is extinguished. Broadband cooling on the intercombination line, where the laser is significantly broadened from the FM, remains on. As the atoms cool, the frequency deviation of the FM decreases to cool more efficiently, and the gradient ramps up to compress the atoms. Once the gradient ramps up to 10 G/cm, the coil current is held for “single frequency cooling” (beginning at the 305 ms mark). In this stage, the laser is modulated with a small deviation until the 345 ms mark, at which point the modulation is shut off. The middle curve (blue) is the intensity of the trapping and stirring lasers as a fraction of their maximum value. These intensities are servo controlled using the AOMs shown in Figure 3.8. The control voltage for these intensity servos originates from a National Instruments DAC card. Intensity is ramped down as the frequency deviation shrinks. This is to compensate for the fact that as the FM deviation decreases and the sample cools, there is more power in the spectral components resonant with the atoms. The intensity is held constant during single frequency cooling. The bottom curve (cyan) is the arbitrary waveform generator voltage sent to the FM ports of the frequency synthesizers used in the trapping and stirring laser phase locks. The input voltage range of the FM ports is -1 V to 1 V, where 1 V corresponds to the programmed frequency deviation (3 MHz in this case). The oscillation frequency in this plot is 100 times smaller than the experimental value for the purpose of visualization.

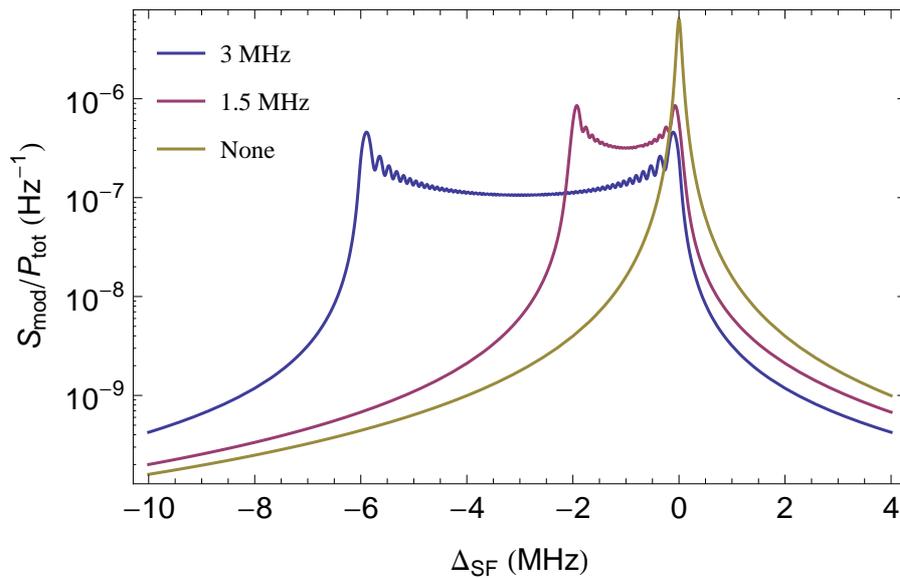


Figure 3.10: Equation 3.3 plotted for different values of f_{dev} (shown in the legend). This is the power spectrum that results from the FM external control depicted in Figure 3.9. Here $f_{mod} = 36$ kHz, as in the experiment. For this plot, Equation 3.3 is convolved with a Lorentzian to account for the fact that the laser has phase noise. The FM deviation starts at 3 MHz, collapsing until the modulation is shut off for the latter part of single frequency cooling. The center of the spectrum is shifted by the frequency deviation to ensure that the laser is always red detuned. At $\Delta_{SF} = 0$, the laser is detuned by -200 kHz from resonance.

the power spectral density S_{mod} (computed with the Wiener-Khinchin theorem) is

$$\frac{S_{mod}(\Delta_{SF})}{P_{tot}} = \sum_{n=-\infty}^{\infty} J_n^2(\beta) \delta(\Delta_{SF} - n\omega_{mod} - \omega_{dev}), \quad (3.3)$$

where Δ_{SF} is the difference between the laser detuning from resonance and the detuning during the single frequency cooling stage, P_{tot} is the total laser power, and $\delta(\dots)$ is the Dirac delta function. This spectral density is plotted in Figure 3.10 for typical experimental parameters. For this second stage of MOT cooling, the frequency deviation is dynamic, decreasing to allow for a laser linewidth that tracks the width of the atoms' velocity distribution as they are cooling and the velocity distribution narrows. Both the trapping and stirring lasers have this functionality.

After the trapping and stirring lasers are delivered to the optical table near the Sr2 main chamber, they are combined and divided into three beams using beamsplitters. These 1 cm diameter beams are aligned along the 461 nm MOT beams, and the two-color waveplates provide the necessary circular polarization. The intensities of these lasers are ramped on during the last 100 ms of 461 nm MOT cooling (Figure 3.9).

To ensure that the 689 nm MOT forms at the location of the optical lattice, Helmholtz coils are used to adjust the 689 nm MOT center (Figure 3.11). Two pairs of large, square coils are placed near the main chamber to provide two dimensions of horizontal fields, and a vertical bias field is created by coils that are mounted to the chamber above (below) the top (bottom) MOT coil. The coil currents are stabilized using servos that sense with Vishay Precision resistors and actuate with MOSFETs (Figure 3.17). The servo locks the voltage across the sense resistors to a voltage based on a stable LM399H source.

3.4 One-Dimensional Optical Lattice

3.4.1 The Cavity Lattice

Throughout both stages of laser cooling, the cavity lattice is left on. The cavity is comprised of two mirrors, both with a radius of curvature of $R = 20$ cm. The mirrors are held in place using

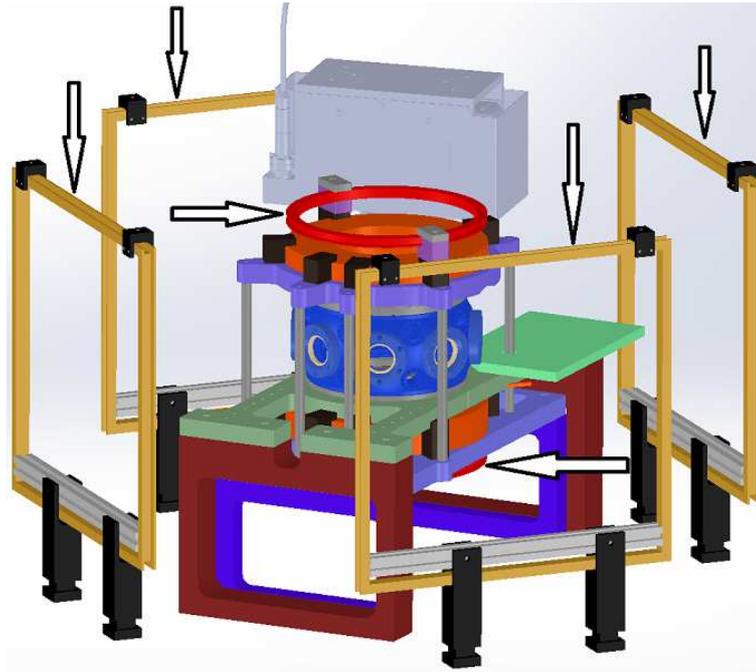


Figure 3.11: The three pairs of Helmholtz bias coils. These coils are indicated with arrows. For the remainder of this thesis, the square coils that form bias fields in the horizontal plane are called the “ xy bias coils” or “ xy coils,” the circular coils that make a vertical bias field are the “ z bias coils” or “ z coils,” and all three together will be referred to as the “ xyz bias coils” or “ xyz coils.”

custom mounts made by the JILA machine shop. The mounts clamp to two 1.33 inch conflat viewports on the main chamber. The mounts have adjustable tilt plates for cavity alignment (Figures 3.12 and 3.13). The lattice laser is locked to the cavity using Pound-Drever-Hall stabilization.

The mirrors are custom coated by Advanced Thin Films for reflection at 813 nm and transmission at 698 nm, allowing the clock laser to be aligned along the lattice axis (the need for which is explained in Section 2.4.3). The mirror surfaces are separated by about 27 cm, resulting in a $160 \mu\text{m}$ beam waist at the focus and a 9 cm Rayleigh range. The cavity output mirror (Figure 2.5) is adhered to a ring-shaped piezo stack. Scanning this piezo and measuring the cavity transmission on a photodetector, the transmission fringe observed on an oscilloscope is similar to that in Figure 2.6. Cavities are often characterized by their *finesse*, which is the transmission peak separation divided by the FWHM of the transmission fringes. Inferred from the cavity fringes, the Sr2 cavity lattice finesse is 120.

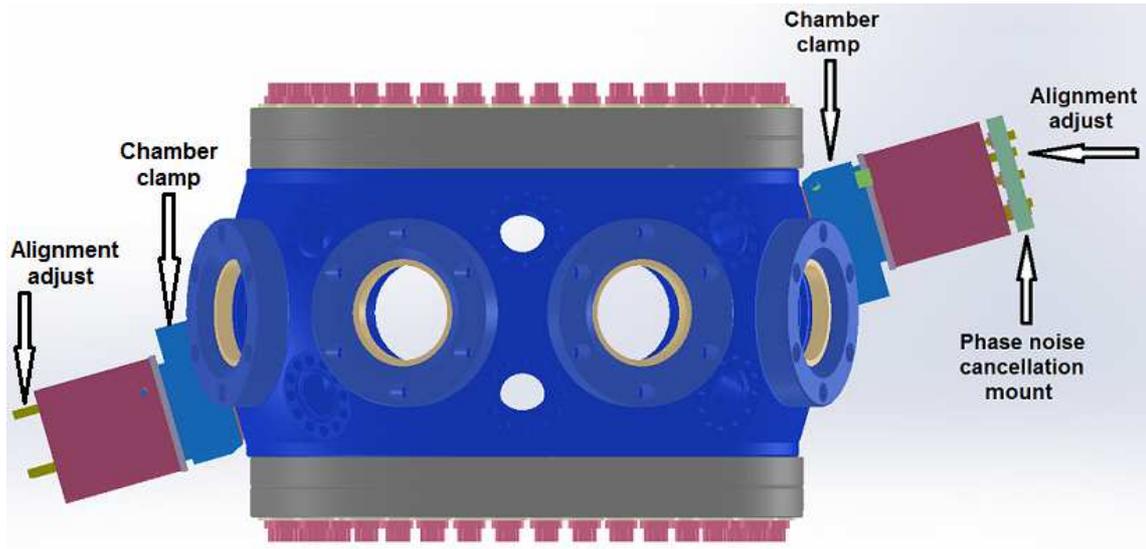


Figure 3.12: The two cavity lattice mirror mounts. These mounts were home-built in the JILA machine shop. The bodies of the mounts are anodized aluminum. The mounts clamp onto opposing 1.33 inch conflat viewports. Screws adjust the mirror position on a spring loaded tilt plate inside the mount body. On one of the mounts, there is an adjustable tilt plate for mounting the fiber phase noise cancellation mirror (see Section 3.6.2).

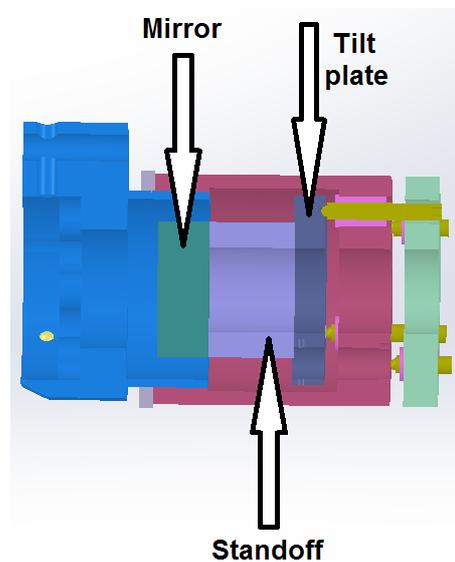


Figure 3.13: A cutaway of one lattice mount. The tilt plate inside the mount body changes the cavity alignment. The standoff is present because the other cavity mirror mount has a piezo stack actuator with similar dimensions, ensuring that the mirrors are equidistant from the center. Aside from tilt screws, ball bearings, and springs, the mounts are made out of anodized aluminum.

Aligning the cavity lattice makes temporary use of a laser resonant with the $|^1S_0, F = 9/2\rangle \rightarrow |^3P_1, F = 9/2\rangle$ transition. This laser is sent through the cavity (entering through the output mirror) and aligned to extinguish the atoms in the 689 nm MOT.³ A laser frequency that effectively extinguishes the atoms was found empirically and used for this purpose.⁴ The laser is also focused on the 689 nm MOT to ensure that the alignment procedure is sensitive. Comparing the positions of the lattice laser before the input mirror and the alignment laser after passing through the cavity, the lattice laser position is adjusted to be counterpropagating with the alignment laser. With this change in the lattice laser alignment, the cavity needs to be realigned by adjusting the cavity mirrors until the transmission indicates a single cavity mode. This procedure tends to steer the alignment laser off the atoms (since the output mirror surface often needs to be aligned during this process), so it needs to be iterated a few times until the cavity and 689 nm MOT are aligned. After a few iterations of this alignment procedure, the xyz coils (Figure 3.11) are varied to maximize the atom signal in the lattice.

The AR coating of the viewports has been observed to be birefringent. A cavity mode occurs for lattice laser polarization at a few degrees angle with respect to gravity, and another mode (resolved from the first) appears for lattice polarization nearly perpendicular to gravity. The cavity lattice laser polarization is rotated to align with the polarization mode perpendicular to gravity.

3.4.2 Tapered Amplifier Lattice Laser

Tapered amplifiers (TAs) can be problematic sources of lattice light [114]. Amplified spontaneous emission from the TA can cause a large ac Stark shift of the clock transition that varies in time. A better source of lattice light is a titanium sapphire (Ti:saph) laser, which has much greater spectral purity. However, a TA-based system is a more convenient source of 813 nm light since it requires little maintenance and start-up time. Therefore, when performing measurements that are insensitive to this amplified spontaneous emission shift, a TA system is preferable. For

³ The cavity mirrors transmit at this wavelength well enough for this procedure.

⁴ The detuning used here is -70 kHz from resonance. This could be a photoassociation resonance in ^{87}Sr .

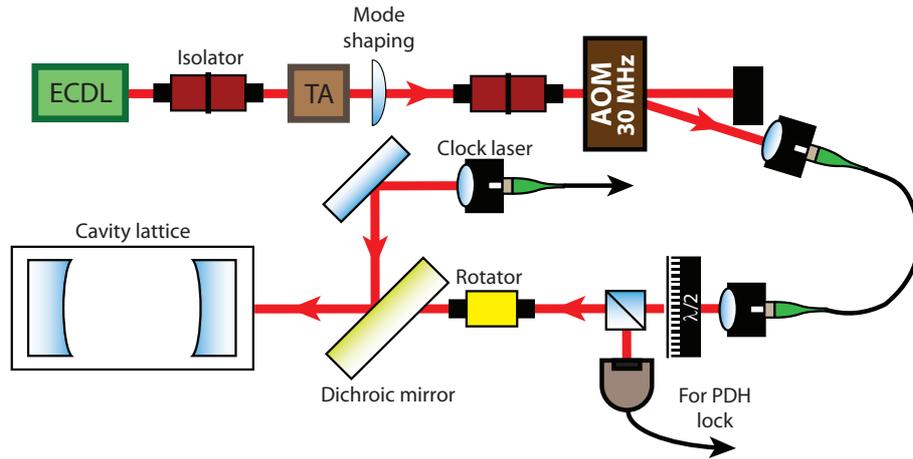


Figure 3.14: A free-running ECDL seeds a tapered amplifier (TA). Mode shaping optics (represented here as a single lens) collimate the output in the horizontal and vertical directions and then telescope it down to a smaller beam. An AOM is used for intensity control of the field in the cavity. This light is delivered via fiber from the laser source to the Sr2 main chamber. A Pound-Drever-Hall servo locks the laser to the cavity resonance.

measurements of the lattice ac Stark shift or for running the system as a clock, a Ti:saph-based lattice is necessary.

The TA for the lattice is seeded by an ECDL comprised of an Eagleyard diode laser in a home-built enclosure (Figure 3.14). The laser is centered on the magic wavelength every few hours (using a Wavemeter) and left free running. Laser frequency drift about the magic wavelength is not problematic since the drift rate is sufficiently small and this system is only used for measurements that are insensitive to lattice Stark shifts. Sidebands for a PDH lock to the lattice cavity are put on the laser by modulating the ECDL current. The laser diode current and the cavity lattice piezo stack are locked to the PDH error signal. The cavity transition is separated from the clock laser using a dichroic mirror and then measured on a photodetector. This detector is part of an intensity servo loop that actuates on the AOM just before the lattice light is fiber coupled (Figure 3.14). There is as much as 8 W of one-way power in the cavity.

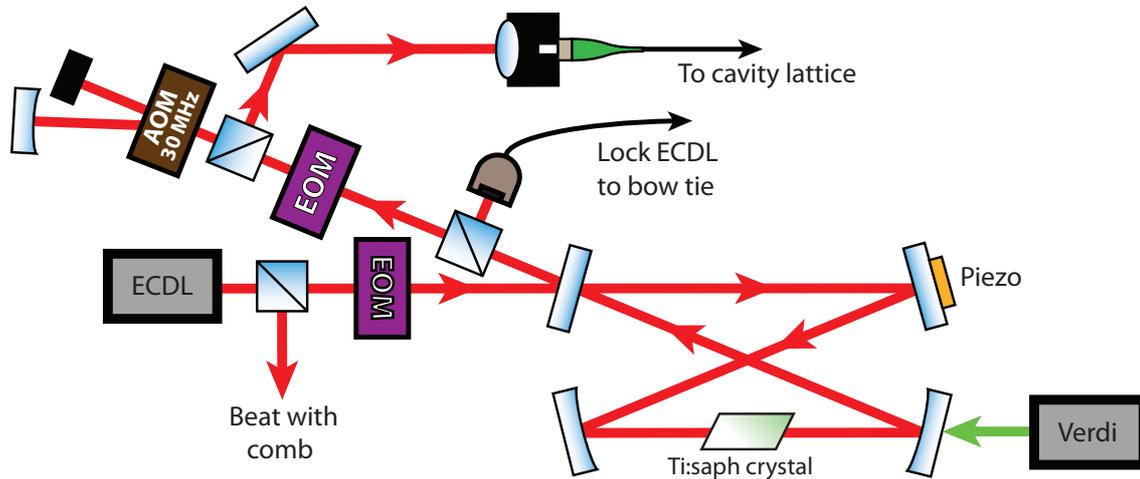


Figure 3.15: The home-built Ti:saph-based lattice laser. The frequency of this system is stabilized by frequency locking the Ti:saph to a Yb fiber comb (actuating on the ECDL external cavity piezo). A Ti:saph crystal is situated inside a bow tie cavity and pumped with an 18 W Coherent Verdi laser. The light reflected off of the bow tie cavity’s input mirror beats with the field leaking out of the cavity for a Pound-Drever-Hall error signal. This error signal is used to lock the ECDL current and bow tie cavity piezo for maximum cavity transmission. A double-passed AOM and the cavity lattice piezo lock this laser to the cavity lattice. An EOM before the double-passed AOM puts the sidebands on the light needed for a Pound-Drever-Hall lock to the cavity lattice.

3.4.3 Titanium Sapphire Lattice Laser

The lattice Ti:saph laser is a home-built system, injection locked [22] with an ECDL based on an Eagleyard diode. When the Ti:saph laser is needed, frequency drift of the lattice wavelength is unacceptable, so the diode is frequency locked (using the ECDL external cavity piezo) to a Yb fiber comb that is referenced to the NIST Boulder maser array. This stabilization procedure allows the absolute frequency⁵ of the lattice laser to be determined. In this case, a double-passed AOM provides the high-frequency lock to the cavity lattice. Although this double-passed AOM lock has a lower bandwidth (about 200 kHz, which is typical of AOM-based servos) than the current lock for the TA-based lattice laser, the Ti:saph laser also has a purer spectrum (on the order of a 1 kHz linewidth). Because of this, the AOM lock is sufficient for the Ti:saph laser. The maximum one-way lattice power using the Ti:saph laser is about 8 W. This system also uses an intensity servo that stabilizes the cavity transmission (actuating on the drive amplitude of the Ti:saph double-passed

⁵ The frequency in SI units

AOM). For clock operation, the intensity servo is set such that the one-way power is 2.3 W.

Ramping the MOT field from 60 G/cm to 3 G/cm (Figure 3.9) causes the MOT coils to shake, which shakes the cavity lattice as well. Although we checked to ensure that this shaking is well dampened by the time clock spectroscopy is performed, the 813 nm source can come unlocked from the cavity lattice during the ramp. For the TA lattice, the 2 MHz lock bandwidth is large enough for the laser to survive the shaking; however, the smaller-bandwidth AOM lock for the Ti:saph lattice occasionally unlocks when the lattice cavity shakes. To combat this, the ramp is made as gentle as can be, and the Ti:saph lattice also includes an auto-locking circuit.

3.5 Spin Polarization into the Angular Momentum Stretched States

3.5.1 Spin Polarization Bias Field Coils

With ^{87}Sr atoms cooled to ultracold temperatures and trapped in an optical lattice, spectroscopy of the clock transition is now possible. As explained in Section 2.5.1, clock spectroscopy is performed in a bias field that is aligned parallel to the lattice and clock laser polarizations. To this end, it is convenient to include another pair of Helmholtz coils mounted on the main chamber (Figure 3.16). These “spin polarization coils” (or just “polarization coils”) are also current stabilized using a sense resistor and MOSFET (Figure 3.17). Since the polarization coil field is not precisely in the direction of the lattice polarization (as defined by the cavity polarization modes), additional fields are provided by the xyz coils (Section 3.5.1) to align the bias field. For each of the xyz coils, the field used to align the magnetic bias to the lattice polarization and the field that centers the 689 nm MOT on the lattice are controlled with independent voltages, which are summed to create a set point for the xyz coil servos (Figure 3.17). The fields created by these voltages can be tuned and switched separately, acting as two independent coils.

The clock laser and lattice polarization are aligned by sending both these lasers through the same broadband-AR-coated polarizer. The bias field is aligned to these two polarizations using clock transition spectroscopy. Alignment using spectroscopy is possible because when the bias



Figure 3.16: A second pair of cavity mirror mounts are clamped to 1.33 inch viewports. These were constructed with forms around which 125-turn bias coils are wrapped. The mirrors are not installed in this work.

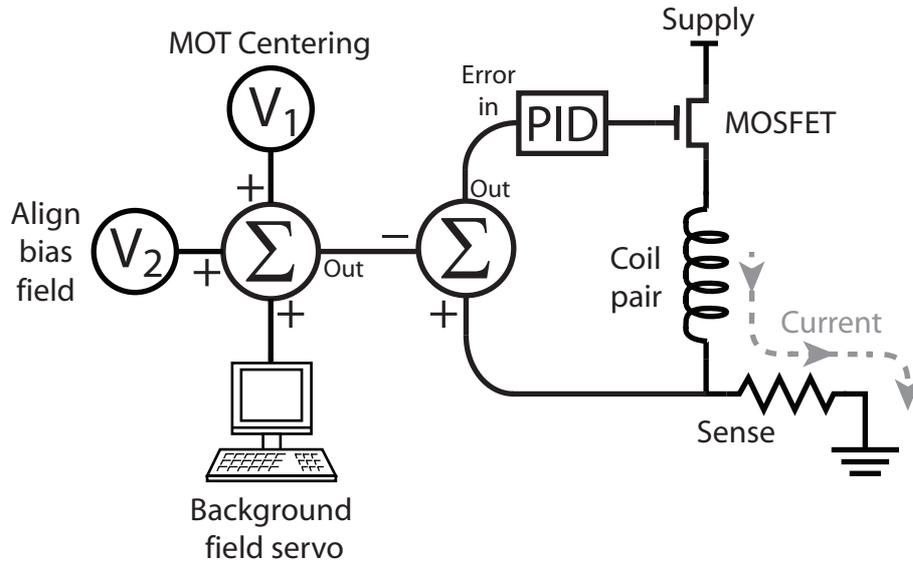


Figure 3.17: Bias coil current servo. The loop stabilizes the magnetic field and allows for field control. There is one of these servo loops for each pair of bias coils, allowing for independent control of the x , y , and z coils. V_1 and V_2 are tunable dc voltages based on LM399H voltage sources. These voltages are set by hand and do not typically need to be adjusted. They can also be tuned and switched independently. A digital-to-analog converter (DAC) computer card provides a third voltage for the background field servo, which will be discussed in Section 5.2.2.

field is not aligned with the clock laser polarization, scans about the clock resonance will show σ transitions that do not have the same $-108.4(4)m_F$ Hz/G magnetic field sensitivity as π transitions (Section 2.5.1). Therefore, the bias field is aligned by nulling the amplitude of the σ transitions using the xyz coil fields.

3.5.2 Spin Polarization

After the atoms are loaded into the lattice and the MOT fields are shut off, spin polarization is performed. The zeroth order of the stirring laser AOM (Figure 3.8) is sent through another AOM and then delivered by optical fiber to the main chamber. This acts as a spin polarization beam. The beam passes through a liquid crystal half waveplate and then a quarter waveplate for circular polarization. The liquid crystal waveplate allows for the resulting circular polarization to be switched between σ_+ and σ_- helicity using a TTL signal (motivated in Section 2.5.1). This beam is sent through clearance holes in the center of the polarization coil mounts (Figure 3.16).

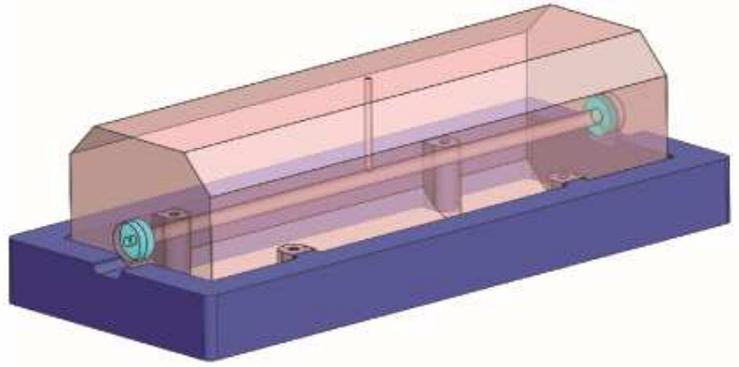


Figure 3.18: A CAD of the cavity used to stabilize the clock laser. Also pictured is the Zerodur cradle (in blue, false color). The cavity and cradle are separated by viton standoffs. The ultralow expansion glass spacer is 40 cm long.

During spin polarization, a small 50 mG field is applied to the polarization coils, and the fields from the xyz coils that align the bias field are shut off. The spin polarization beam is applied for 30 ms, driving σ transitions until the atoms are in one of the ground angular momentum stretched states (where the state is selected by the helicity of the liquid crystal waveplate). Like with the stirring laser, the $|^1S_0, F = 9/2\rangle \rightarrow |^3P_1, F = 9/2\rangle$ transition is used here. The field in this stage is kept small to ensure that the spin polarization direction is defined by the bias field, yet not so large as to shift some of the σ transitions far out of resonance with the spin polarization beam.

After this, the polarization coil current is increased and the xyz coil fields that align the bias field are turned on, resulting in a 300 mG field in the direction of the lattice polarization. The atoms are spin polarized along the polarization coil direction (Figure 3.16), yet the clock laser polarization is parallel to the lattice polarization; therefore, changing the bias field direction after spin polarization causes the atoms to become depolarized with respect to the clock laser polarization. This reduces the contrast of the π transitions used in the clock (Figure 2.14b). In practice, this contrast reduction is small (at the few percent level), so no attempt is made to fix it.

3.6 The Clock Laser

3.6.1 Cavity Stabilization

The 698 nm stable clock laser was constructed out of an ultralow expansion glass spacer that is 40 cm long [110, 89, 77, 8]. The mirrors were coated by Advanced Thin Films on fused silica substrates. The cavity is supported against gravity by viton hemispherical standoffs placed at four contact points. The standoffs were positioned to make the cavity insensitive to vibration noise [91].

The cavity and its cradle (Figure 3.18) are placed inside a heat shield and two-stage vacuum chamber. Active bipolar temperature control keeps the cavity temperature stable. The vacuum chamber resides on a Herzan active vibration cancellation stage. Also, to prevent acoustic noise, the vacuum chamber is surrounded by an acoustic shield.

The 689 nm light source is a home-built ECDL. First, this laser is current and piezo locked to a prestabilization cavity. The laser then double passes an AOM, which is used to lock the laser to the 40 cm cavity. A piezo on the prestabilization cavity also provides a lower bandwidth lock to the 40 cm cavity. Measurements of this laser's noise spectrum show that the fully stabilized clock laser has a 26 mHz linewidth [8].

3.6.2 Fiber Phase Noise Cancellation

To prevent the ECDL from unlocking due to laboratory bustle, the clock laser cavity stabilization setup is kept in a separate room that remains unoccupied while the Sr2 clock is running. The fully stabilized laser is sent via optical fiber to the Sr system. A fiber of the length required to transfer the clock laser between these rooms would cause broadening at the 1 kHz level, which would be devastating for the Sr2 system, if it were not for fiber phase noise cancellation [75]. Fiber phase noise cancellation reduces the decoherence from this fiber to levels that are currently negligible for Sr2 [36].

For fiber phase noise cancellation, the source laser passes through an AOM before getting launched into the fiber. A fraction of the fiber output is retroreflected back through both the fiber

and the AOM. The retroreflected electric field will acquire a phase factor $e^{2i(\omega_{AOM}t + \phi_{noise})}$, where ω_{AOM} is the frequency shift due to the AOM, ϕ_{noise} is a stochastic function that describes the decoherence (mostly due to the fiber) between the AOM and the retroreflecting mirror, and the factor of two comes from double passing both the fiber and the AOM. Using a beat between the retroreflected light and the source before the AOM, the retroreflected beam is phase locked to the source laser using the AOM frequency, removing the effects of the fiber.⁶

After the clock laser is delivered to the Sr lab using a phase-noise-canceled fiber, its power is not sufficient for both Sr1 and Sr2. Therefore, the light from the 40 cm cavity injection locks a diode, the output of which provides ample power for both systems. Light from the diode is divided equally between Sr1 and Sr2 and sent to these two experiments using fibers, which are again phase noise canceled. In this second stage of fiber noise cancellation, one can benefit from judiciously positioning the retroreflecting mirror that sends light back through the fiber.

Any decoherence encountered by the laser light as it travels from the phase noise cancellation AOM to the retroreflecting mirror is corrected by this cancellation technique. In other words, the retroreflecting mirror defines a plane at which the source laser coherence is reproduced (within the bandwidth of the phase noise cancellation lock). Therefore, the phase noise cancellation retroreflecting mirror is mounted to the same structure that holds the cavity lattice mirrors (Figure 3.12). In this case, the laser phase is adjusted to track any shaking of the lattice, removing the Doppler shift that would result from atoms moving with respect to the clock laser.

3.7 Digital Atomic Servo

3.7.1 Clock Laser Frequency Modulator and Drift Cancellation

Another AOM, labeled as the “modulator” in Figure 2.15b, is used to control the frequency of the clock laser for spectroscopy. The zeroth order of this AOM is retroreflected back through the fiber for phase noise cancellation, whereas the -1st order interacts with the atoms. The beam from

⁶ The phase lock actuates on a voltage-controlled oscillator (VCO) centered at twice the AOM frequency. The VCO frequency divided by two drives the AOM. The loop adjusts the VCO frequency to lock the beat to a stable RF reference at twice the desired AOM frequency.

the 698 nm injection-locked diode is too powerful for normal clock operation, so a neutral density filter with an optical depth of four is placed in the -1st order beam.

The frequency that drives this AOM comes from two synthesizers that are mixed together. One “scanning synthesizer” is responsive to the experimental control software that steps the clock laser frequency for scanning the clock transition or for the atomic servo. The other synthesizer has a linear frequency drift with a fixed drift rate of about -30 mHz/s. This drift rate is adjusted once or twice per day to cancel the drift of the 40 cm cavity. The cavity drift, which is attributed to material creep, is linear and can be canceled effectively. The scanning synthesizer, which has 1 μ Hz resolution,⁷ is connected to the Sr2 experimental control computer. The computer adjusts the clock laser frequency through a GPIB command that sets the frequency of the scanning synthesizer.

The modulator AOM that controls the clock transition frequency (Figure 2.15b) also has an intensity servo. Part of the -1st order is split off and sent to a photodiode. The photodiode output goes into a PI-I controller (one proportional-integral gain stage and one pure integrator), which actuates on the drive amplitude of the modulator AOM. Two versions of this servo have been used: one with 50 kHz of bandwidth and one with 250 kHz. The setpoint for this servo is provided by an adjustable dc voltage (based on a stable LM399H-based voltage source) that can be switched on and off using a 5 ns switch from Mini-Circuits. This allows for clean clock laser square pulses with an easily adjustable pulse duration and amplitude.

3.7.2 Camera Imaging and Photomultiplier Tube Detection

As explained in Section 2.5.2, the error signal for the atomic servo is the difference between excited state population fraction measurements. To measure the excited population fraction, three lasers are used. After driving the clock transition, the ground state is fluoresced with a resonant *counting laser* on the $^1S_0 \rightarrow ^1P_1$ transition (Figure 2.1). This strong transition scatters many photons, which are detected with a photomultiplier tube (PMT), yielding a ground-state population signal V_g that is proportional to the number of atoms in the ground state. The ground state

⁷ This is a Stanford Research Systems DS345.

population scatters so many counting laser photons that it is quickly heated out of the trap and lost to the experiment. The excited state population is not resonant with the counting laser, so it is unaffected. Next, the repump lasers (Figure 2.2) are applied, rapidly driving the excited state population to the clock transition ground state. Atoms are then fluoresced again on $^1S_0 \rightarrow ^1P_1$, and photons are counted again with the PMT, providing the excited state population signal V_e . The measured excited state population fraction is computed as⁸ $\rho_{meas} = V_e/(V_e + V_g)$.

When the counting laser is applied to measure the excited state population fraction, fluorescence is collected through two viewports. A 3 inch diameter lens collects fluorescence through the large 8 inch conflat viewport on the top of the main chamber, focusing this signal down onto the PMT with 2f imaging.⁹ The PMT output is sent to an oscilloscope computer card from GaGe Applied Technologies. Each time the counting laser is applied to the sample to measure the ground (excited) state population, V_g (V_e)—given by the time integral of the fluorescence signal measured by the PMT and the oscilloscope card—is added to the data record.

When V_g is measured, fluorescence is also collected by a 2 inch diameter lens outside one of the 2.75 inch conflat viewports. This light is imaged on an Andor CCD camera with a wedged window anti-reflection coated for 461 nm light. This camera is typically used for signal monitoring, but it has also been used for absorption imaging to calibrate the fluorescence counts in terms of atom number.

3.7.3 Clock Transition Scans

The software for scanning and locking the clock laser to the clock transition was written in MATLAB using the environment's GUI building functionality. To begin the digital servo, first the clock transition is found by scanning over a 200 kHz window centered about the last known scanning synthesizer frequency corresponding to the clock transition resonance. This broad scan is performed at zero bias field and with the neutral density filter removed from the clock laser

⁸ One detail not covered in this treatment is that once the atoms are all headed out of the lattice, a third measurement is taken with the counting laser. This yields the background counts, which are subtracted off of the excited and ground state population measurements (Section A.4).

⁹ When the lens is twice the focal length from the source

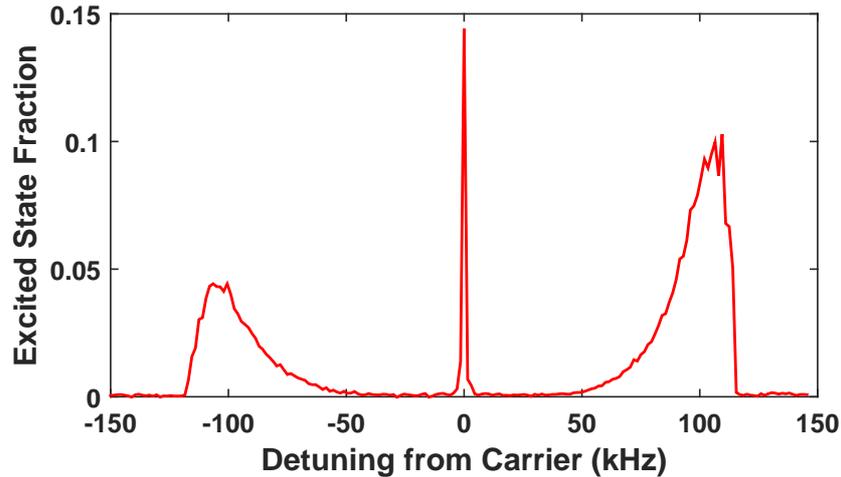


Figure 3.19: A typical sideband scan for 1.5 kHz frequency steps, a 160 ms pulse, and the neutral density filter (Section 3.7.1) removed from the clock laser. At this clock laser power, the clock transition is saturated and dephases, achieving a final excited state fraction of about 0.5. This value is not portrayed in this data set due to the coarseness of the scan. The carrier transition occurs at zero detuning, the blue sideband at positive detuning, and the red sideband at negative detuning. The sidebands are asymmetrically broadened due to the effect described in Section 2.4.4.

beam (Section 3.7.1) to power broaden the transition. This power broadening helps to locate the resonance quickly using coarse frequency steps of 1.5 kHz. A scan over this range and at this large power drives sideband transitions.

Data from this scan (Figure 3.19) resembles the calculation plotted in Figure 2.12. These data can be used to infer the trap frequency. As explained in Section 2.4.4, the frequency difference between the center of the carrier transition and the blue sideband edge, ν_{blue} , determines the lattice axial trap frequency ν_z as $\nu_{blue} = \nu_z - \nu_{rec}$.

The temperature in the lattice axial direction, T_z , can also be approximated using this scan. For the case of a 1D harmonic oscillator potential, the entire population in the lattice can make a blue sideband transition whereas all population that is not in the motional ground state can make a red sideband transition. Statistical mechanics gives the population in motional state n (where, again, $n = 0$ is the motional ground state) of $[1 - \exp(-h\nu_z/k_B T_z)] \exp(-nh\nu_z/k_B T_z)$. Therefore, the population participating in a blue sideband transition divided by that participating in a red sideband transition is $\exp(h\nu_z/k_B T_z)$. To a rough approximation, this ratio is equal to the

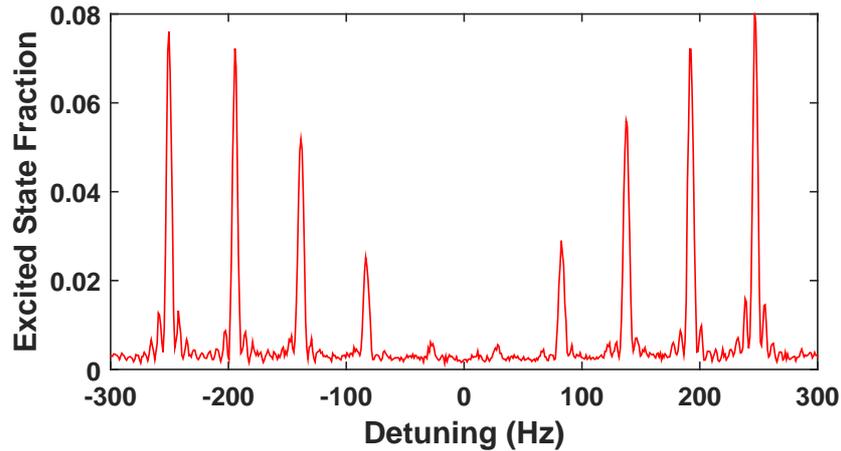


Figure 3.20: With the bias field applied, all 10 m_F states are resolved. Here the detuning is measured with respect to the zero field resonance. These data are mathematically described by the formalism presented in Section 2.5.1.

ratio of the maximum excited state fraction achieved for the blue and red sidebands. Therefore, $\rho_{max,blue}/\rho_{max,red} \simeq \exp(h\nu_z/k_B T_z)$. This approximation typically agrees within 15% of a more complex model that accounts for radial motion and trap anharmonicity [12].

With the clock transition located, the intensity is reduced with a neutral density filter and the clock laser intensity servo, and the bias field is applied. Without spin polarization, a scan of the carrier transition shows ten resonances corresponding to the different m_F states (Figure 3.20). This scan provides the clock laser AOM frequency corresponding to the two stretched states, and with this information the stretched-states servo (Section 2.5.3) can be initialized.

3.7.4 Digital Servo Software

The software requires the user to input the initial modulator frequencies corresponding to the stretched state resonances (obtained from clock transition scans) and the linewidth of the feature it is locking to. To lock to a lineshape, the software steps the scanning synthesizer half of the input linewidth on one side of the resonance center and then measures the excited state fraction. The scanning synthesizer is then stepped on the other side of the resonance by the input linewidth, and the excited state fraction is measured again. The first iteration of the computed error signal ϵ_{clock}

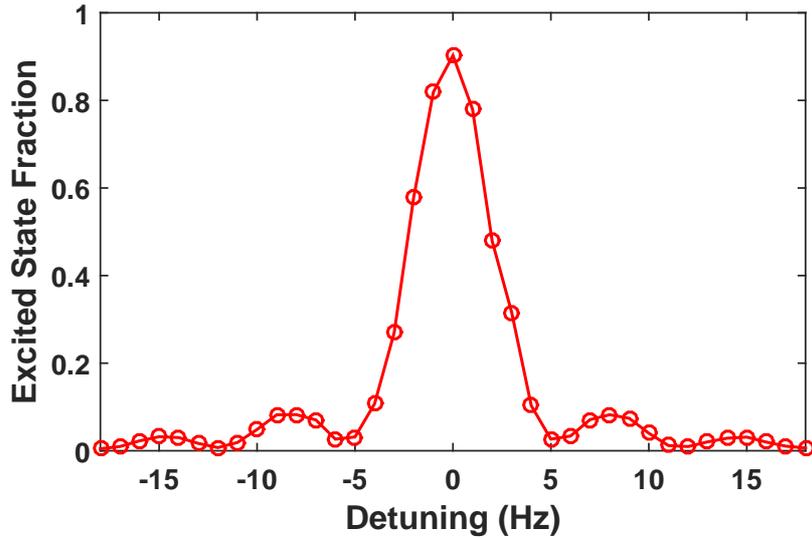


Figure 3.21: A scan of one stretched state in a bias field after the atoms have been spin polarized. Here the pulse duration is 160 ms. The clock laser intensity is adjusted for maximum contrast.

will be large in magnitude if the initial modulator frequencies are off from the resonant values. Laser frequency noise will also cause noise in ϵ_{clock} . A desktop computer records the fluorescence counts and the line center, which is calculated from ϵ_{clock} with Equation 2.44. The next iteration of the servo uses the last computed value of the line center, and with this information the measurement cycle begins again. The stretched-states servo locking procedure is performed for both stretched states using two independent servo loops, with the locking sequence described in Figure 2.16.

The lineshapes used in the stretched-states servo are Rabi π -pulses, which occur when $\Omega_{mF}t = \pi$ (see Equation 2.42). Lineshapes very similar to those predicted by the Rabi solution are typically observed (Figure 3.21).

Chapter 4

Clock Stability

4.1 The Allan Deviation

4.1.1 Definition

As explained earlier, total systematic uncertainty and stability are the two figures of merit used to gauge clock performance (Section 1.6). Stability is a measure of the noise of a frequency standard, for it quantifies the uncertainty in the mean clock oscillator frequency due to random processes such as frequency or phase fluctuations. In essence, better stability is useful because it implies less measurement time. If an oscillator needs to be calibrated to a frequency standard, this calibration measurement will typically involve averaging down the noise in a beat between the oscillator and the standard. The noisier the standard, the longer the calibration will take to average. Also, stability affects total systematic uncertainty. Systematic effects are often studied by modulating a parameter of the apparatus (like the amplitude of an electromagnetic field that perturbs the clock states) and measuring the corresponding frequency shift of the clock transition. Frequency noise in such a measurement needs to be averaged down to obtain precise uncertainties of systematic biases. Better stability means that evaluating these uncertainties will take less data acquisition time, which is usually the main practical limitation of systematic measurements.

According to the IEEE, the recommended statistic [50] for quantifying stability is the Allan Deviation, $\sigma_y(\tau)$, where

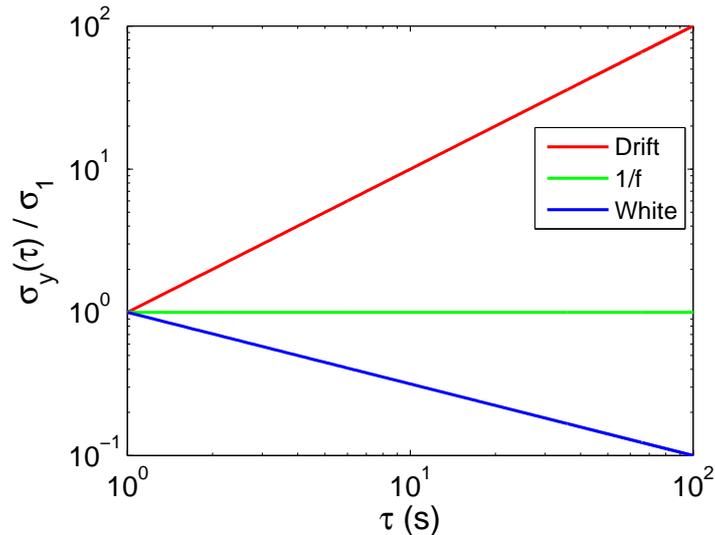


Figure 4.1: The three processes found in Sr2 that degrade stability. Each are scaled to σ_1 , which is their value at $\tau = 1$ s. For white noise, $\sigma_y(\tau)$ averages down in the same manner as the standard deviation of the mean, which decreases as the square root of the number of data points. There is no advantage to averaging pure $1/f$ noise since $\sigma_y(\tau)$ does not improve with τ . In the case of drift, stability gets worse with averaging time.

$$\sigma_y^2(\tau) = \frac{1}{2} E \left[\left(\frac{1}{\tau} \int_t^{t+\tau} y(t') dt' - \frac{1}{\tau} \int_{t-\tau}^t y(t') dt' \right)^2 \right]. \quad (4.1)$$

Here τ is the averaging time, $E[\dots]$ denotes the expectation value, and

$$y(t) = \frac{\nu_c(t) - \nu_{Sr}}{\nu_{Sr}}. \quad (4.2)$$

$\nu_c(t)$ is the cavity stabilized clock laser frequency, which is a stochastic as a function of time, and ν_{Sr} is the frequency to which the clock laser is referenced (that of the Sr atoms in this case).

It is straightforward to show (Section A.2) that

$$\sigma_y^2(\tau) = 2 \int_0^\infty S_y(f) \frac{\sin^4(\pi\tau f)}{(\pi\tau f)^2} df, \quad (4.3)$$

where $S_y(f)$ is the one-sided power spectral density of $y(t)$. The power spectrum of most atomic clocks can be described with a simple model $S_y(f) = \sum_{j=-2}^2 h_j f^j$, where h_j is constant [50, 100].

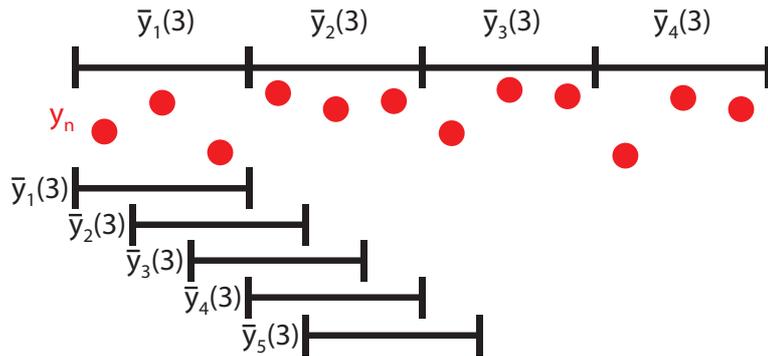


Figure 4.2: The Allan deviation estimator for overlapping and non-overlapping averages. The frequency data y_n (red points) is binned in two different ways. The binning scheme shown above the points is for the non-overlapping Allan deviation. The points $\bar{y}_n(m)$ are the average of the data within a given bin, and m is the bin size (which is 3 in this figure). The binning scheme below the points is for the overlapping Allan deviation. Again, the points $\bar{y}_n(m)$ are the average of the data within a given bin; however in this case the bins overlap.

Typically only three noise processes govern the behavior of the Sr2 clock: white noise, flicker frequency noise (or “1/f noise”), and linear drift.

For white noise, $S_y(f) = h_0$, and using Equation 4.3, $\sigma_y(\tau) = \sqrt{h_0/2\tau}$. For 1/f noise, $S_y(f) = h_{-1}/f$, and $\sigma_y(\tau) = 2\ln(2)h_{-1}$. In the case of drift at a rate α , $y(t) = \alpha t$, and Equation 4.1 yields $\sigma_y(\tau) = \alpha\tau/\sqrt{2}$. Therefore, these three processes can be distinguished by the behavior of the Allan deviation as a function of averaging time (Figure 4.1).

4.1.2 Estimators

Suppose $y(t)$ is sampled with a cycle time $2T_c$ and a total of N_y points.¹ Let $\{y_1, y_2, \dots, y_{N_y}\}$ be equal to $\{y(2T_c), y(4T_c), \dots, y(2N_yT_c)\}$. The Allan deviation is often estimated as

$$\sigma_y^2(2mT_c) \simeq \frac{1}{2(N_y - 1)} \sum_{n=1}^{N_y-1} [\bar{y}_{n+1}(m) - \bar{y}_n(m)]^2. \quad (4.4)$$

Here $\tau = 2mT_c$ is the averaging time and m is an integer. The way this estimator is defined, m runs from a minimum of 2 to a maximum of $N_y/2$. \bar{y}_n is the average of data binned in a particular

¹ Recall from Section 2.5.2 that the cycle time T_c is the time it takes to cool atoms, trap them in the lattice, and interrogate the clock transition once. It takes two interrogations of the clock transition to get one resonance center frequency, so $2T_c$ is needed to measure one point. If $y(t)$ were the stretched state average frequency, $4T_c$ would be needed to measure one point.

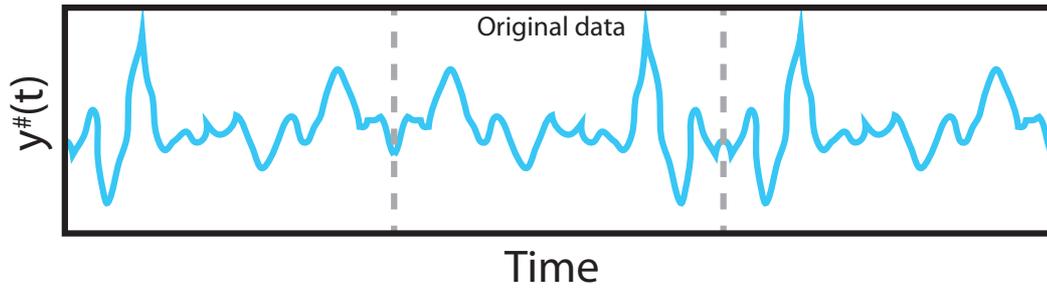


Figure 4.3: Frequency points used to calculate the total deviation. A set of points $y^\#$ is created by reflecting the original data set twice.

way. Overlapping and non-overlapping binning schemes are typical (Figure 4.2), but overlapping bins are much more utilized in the clock community [101]. When overlapping bins are used, this estimator is typically referred to as the “overlapping Allan deviation.” $N_{\bar{y}}$ is the total number of points after binning and averaging, and m is the bin size. Stability is typically analyzed by calculating the Allan deviation estimator for different values of m to observe how noise diminishes as a function of averaging time. The Allan deviation is usually plotted as a function of $2mT_c$. In essence, the Allan deviation estimator quantifies fluctuations in the difference between averaged subsets of data as a function of averaging time.

Another useful estimator is the *total Allan deviation* (or *total deviation*), $\sigma_{total}(\tau)$ [43]. This is calculated using a set of points constructed by reflecting the frequency data twice (Figure 4.3). Reflecting generates the points $\{y_{-N_y}^\#, y_{-N_y+1}^\#, \dots, y_{N_y}^\#\}$, which are defined as $y_n^\# = y_n$ for $1 \leq n \leq N_y$. For $1 \leq j \leq N_y - 1$, $y_{1-j}^\# = y_j$ and $y_{N_y+j}^\# = y_{N_y+1-j}$. Using this definition,

$$\sigma_{total}^2(2mT_c) = \frac{1}{N_y - 1} \sum_{n=2}^{N_y} \left[\bar{y}_n^\#(m) - \bar{y}_{n-m}^\#(m) \right]^2, \quad (4.5)$$

where $\bar{y}_n^\#(m) = \frac{1}{m} \sum_{k=0}^{m-1} y_{n+k}^\#$. This estimator is useful because it has less error at longer averaging times [43].

4.2 Effects that Limit Stability

4.2.1 Quantum Projection Noise

For atoms in uncorrelated states, the fundamental limit to clock precision is quantum projection noise (QPN). QPN is the result of the binomial statistics (Section 2.7) associated with projecting a coherent superposition of electronic states onto one of the levels of the clock transition. For a clock laser pulse chosen to drive the excited state population fraction ρ_{ee} to a value p (typically 1/2 during clock operation), the fluctuation in the excited state atom number σ_N is $\sqrt{Np(1-p)}$. This is simply the standard deviation of a binomial distribution, which is approximately Gaussian for $N \gtrsim 50$ and which averages down like white noise.

The fluctuation in the measured excited state population fraction σ_ρ due to QPN affects the clock signal. This is given by $\sigma_\rho = \Delta N/N = \sqrt{p(1-p)/N}$. The atomic servo cannot distinguish between these fluctuations and clock laser noise, so it will adjust the laser frequency in attempt to correct for the QPN. This negatively affects the stability of the locked clock laser.

When the clock laser is locked to a value of the excited state fraction, the slope s of the Rabi lineshape at the lock point is the factor that determines the effect of the excited state fraction fluctuation on the atomic servo.² Differentiating ρ_{ee} (Equation 2.26) with respect to detuning $\delta = \Delta/2\pi$ and evaluating this derivative at the detuning for which $\rho_{ee} = 1/2$ (since p is typically selected to be this value), s can be obtained numerically. Since QPN is a white noise process, the noise amplitude averages down with the square root of the number of points averaged. For an averaging time of τ and a cycle time of T_c , the number of points averaged is $\tau/2T_c$, and the clock stability due to QPN is

$$\sigma_{QPN}(\tau) = \frac{\sqrt{2}\sigma_\rho}{2s\nu_{Sr}} \sqrt{\frac{2T_c}{\tau}} = \frac{1}{\pi Q} \sqrt{\frac{T_c}{N\tau}}, \quad (4.6)$$

where $Q = \nu_{Sr}/\gamma$, γ is given by Equation 2.43, and $p = 1/2$. Equation 4.6 is not exact for Rabi

² Although there are two values of the slope (one on each side of resonance), which are equal in magnitude and opposite in sign, for convenience I take s to be the positive slope.

spectroscopy,³ but it is good to 3%. The fact that Equation 4.6 has $\sqrt{2}\sigma_\rho$ in the numerator instead of just σ_ρ is because the atomic servo requires two measurements of the excited state fraction to infer the center frequency of a lineshape.

The factor of 2 appears in the denominator of Equation 4.6 because the atomic servo error signal is twice as sensitive to clock laser frequency fluctuations as a single measurement of the excited state fraction. To see this, suppose the atomic servo is locked to an excited state fraction of $1/2$, and then the detuning due to laser frequency drift fluctuates by a small amount δ_ϵ . The two values of the excited state fraction measured by the atomic servo (one on either side of the lineshape) can be expanded about their lock points as $\rho_{ee}(\pm\delta_0 + \delta_\epsilon) \simeq \frac{1}{2} \pm s\delta_\epsilon$, where δ_0 is the detuning corresponding to an excited state fraction of $1/2$. The clock error signal is therefore $\epsilon_{clock} = 2s\delta_\epsilon$, so the sensitivity of the atomic servo to detuning fluctuations is $2s$. Since the atomic servo cannot discern between detuning fluctuations and QPN, it responds to QPN assuming the same $2s$ sensitivity.

4.2.2 The Dick Effect

The main limitation to lattice clock stability is the Dick Effect [28]. The Allan deviation due to the Dick Effect, $\sigma_{Dick}(\tau)$, is given by [42]

$$\sigma_{Dick}^2(\tau) = \frac{1}{\tau} \sum_{k=1}^{\infty} S_{locked}(k/T_c), \quad (4.7)$$

where $S_{locked}(f)$ is the one-sided power spectral density of the clock laser after it is locked to the atoms.⁴ According to this equation, the Dick Effect is stability degradation caused by discrete values of the clock laser's noise spectrum. This is simply signal aliasing, whereby noise components of the locked clock laser with frequencies equal to harmonics of the sampling rate $1/T_c$ appear to the atomic servo as lower-frequency signals. This causes high-frequency noise that is beyond the atomic servo bandwidth to appear as low-frequency noise within the bandwidth of the servo.

³ It is exact for Ramsey spectroscopy with infinitesimal pulses

⁴ Like the spectral densities defined in previous sections, this is the spectral density of the clock laser frequency fluctuation about the clock transition resonance in fractional units.

Since the servo cannot distinguish between aliased noise and actual low-frequency noise that it is designed to remove, stability is negatively affected as the servo attempts to correct for the aliased frequencies. The Dick effect is simply the sum of these contributions.

The derivation of this expression makes use of a widely utilized relation for discrete time signals (Section A.3), which is worth quoting because it provides intuition about the Dick Effect. For a continuous signal $\phi(t)$ sampled every T_c ,

$$S_{dis}(f) = \sum_{k=0}^{\infty} S_{con} \left(f + \frac{k}{T_c} \right), \quad (4.8)$$

where $S_{con}(f)$ is the one-sided power spectral density of the continuous signal, $S_{dis}(f)$ is the spectral density of $\phi(nT_c)$ (the sampled signal $\phi(t)$), and n is an integer. The terms for which the summation index $k \neq 0$ arise from sampling the continuous signal and describe aliasing. Comparing this relation to Equation 4.7, it is clear that the Allan deviation due to the Dick effect is just the aliased part of the locked clock laser power spectrum evaluated for low frequencies.

Note in Equation 4.7 that the Allan deviation averages like $1/\sqrt{\tau}$, which is characteristic of white noise (Section 4.1.1). Dick's original paper [28] explained this in the following way. Stability degradation arises from the laser noise spectrum windowed about harmonics of $1/T_c$. About sufficiently small windows, a complicated noise spectrum will have the characteristics of white noise; therefore, the Dick Effect is the sum of white noise, which is itself white.

Before quoting the standard form of the Dick Effect, another signals relation is worth mentioning. Suppose for two continuous signals $\phi_1(t)$ and $\phi_2(t)$, with Fourier transforms $\Phi_1(f)$ and $\Phi_2(f)$, the transfer function $G(f)$ relates them as $\Phi_2(f) = G(f)\Phi_1(f)$. It is straightforward to show that $S_2(f) = |G(f)|^2 S_1(f)$ (Section A.1), where $S_1(f)$ and $S_2(f)$ are the power spectra of $\Phi_1(f)$ and $\Phi_2(f)$. When these signals are sampled at a rate $1/T_c$, Equation 4.8 implies that, for $f = 0$,⁵ the spectrum of the sampled signal is $\sum_{k=0}^{\infty} S_2(k/T_c) = \sum_{k=0}^{\infty} |G(k/T_c)|^2 S_1(k/T_c)$. Motivated by this expression, it is advantageous to seek a transfer function that relates the locked spectral density

⁵ Here f is taken to be zero to give the same form as Equation 4.7.

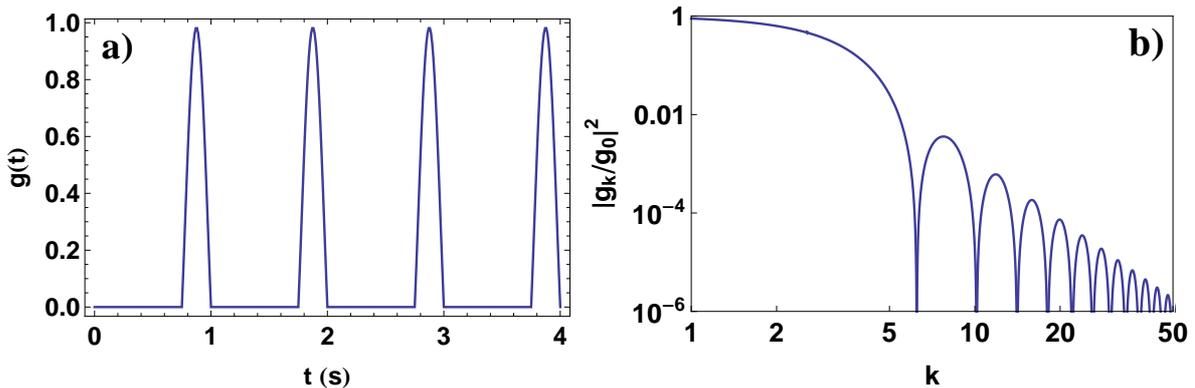


Figure 4.4: a) The sensitivity function for Rabi spectroscopy. Here $T_c = 1$ s and $t_p = 0.25$ s. At the end of the 1 s cycle, the clock laser pulse comes on and $g(t)$ is no longer zero. b) The squared magnitude of the normalized Fourier coefficients, $|g_k/g_0|^2$, for the same T_c and t_p .

to the spectrum of the free-running clock laser, which is easier to measure.⁶

With some work [42], it can be shown that the Dick Effect is

$$\sigma_{Dick}^2(\tau) = \frac{1}{\tau} \sum_{k=1}^{\infty} \left| \frac{g_k}{g_0} \right|^2 S_{clock}(k/T_c), \quad (4.9)$$

where $S_{clock}(f)$ is the one-sided power spectral density of the free-running clock laser, and

$$g_k = \frac{1}{T_c} \int_0^{T_c} g(t) e^{-2\pi i k t / T_c} dt \quad (4.10)$$

is the k th Fourier coefficient of the *sensitivity function* $g(t)$ [65]. The sensitivity function is the change in the excited state fraction due to an instantaneous phase step of the clock laser. g_0 in Equation 4.9 is simply g_k for $k = 0$. The sensitivity function depends on the type of spectroscopy used in the clock (i.e., Rabi, Ramsey), the measurement time, and the experimental cycle time.

For Rabi spectroscopy, when the clock laser pulse is on (at the end of a cycle),

$$g(t) = -\frac{\pi^2 \kappa}{(\pi^2 + \kappa^2)^{3/2}} \left[\sin\left(\sqrt{\pi^2 + \kappa^2}\right) + \sin\left(\frac{t - T_c}{t_p} \sqrt{\pi^2 + \kappa^2}\right) - \sin\left(\frac{t - T_c + t_p}{t_p} \sqrt{\pi^2 + \kappa^2}\right) \right]. \quad (4.11)$$

⁶ Here “free-running clock laser” means a laser that is cavity stabilized but not locked to the Sr atoms with the atomic servo.

In this expression, t_p is the pulse time. $g(t)$ is piecewise, and when the clock laser is off, $g(t) = 0$ (Figure 4.4). κ is the product of detuning and pulse time for which $\rho_{ee} = 1/2$ (where the atomic servo locks). Therefore, $\kappa = 2.509$, determined numerically by $\frac{1}{1+\kappa^2/\pi^2} \sin^2\left(\frac{1}{2}\sqrt{\pi^2 + \kappa^2}\right) = \frac{1}{2}$.

Equation 4.9 shows that the Dick Effect can be mitigated in two ways: reducing the sensitivity to laser noise, and reducing the laser noise amplitude. The sensitivity function can be reduced by decreasing the measurement dead time so that more of the experimental cycle is occupied by Rabi spectroscopy [120]. Also, better laser stabilization improves laser noise.

4.2.3 Long-Term Drift

Over timescales of hundreds of seconds and longer, long-term drift can dominate over QPN and the Dick Effect. The atomic servo removes frequency drift due to thermal variations in the length of the cavity used to stabilize the clock laser; however, the clock transition frequency itself can change. After all, the laser is locked to the perturbed clock transition, and laboratory perturbations are subject to drift.

One example of this drift is a time varying laboratory temperature. As I will explain in Section 5.8.1, the electric field associated with the ambient laboratory heat causes a blackbody radiation (BBR) Stark shift that has a strong effect on the clock transition. Even a 1 °C drift in laboratory temperature causes a measurable change in the BBR shift. The experimental signature of temperature drift would be an Allan deviation that averages down at first, then reverses direction and begins increasing linearly with averaging time (Figure 4.1) when drift has accumulated enough to become appreciable.

Long-term drift can be alleviated by stabilizing ambient temperatures and other sources of clock transition perturbations.

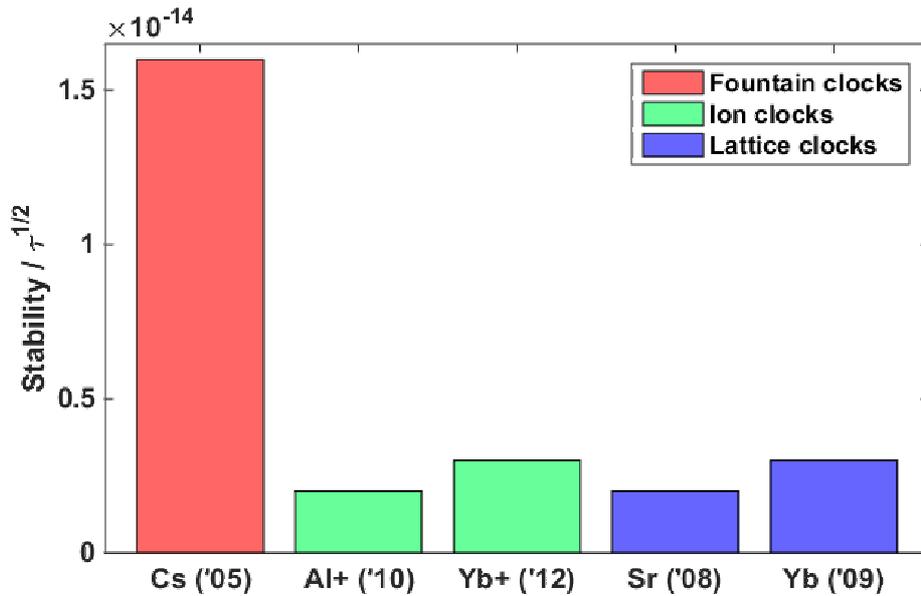


Figure 4.5: Fractional clock stability per square root of averaging time (in seconds) for several clocks. Stability quoted here is for a single clock rather than for a comparison. The record for best Cs clock stability, from BNM-SYRTE, is shown on the left [118]. The NIST Al⁺ ion clock [23] and PTB Yb⁺ ion clock [52] have comparable stability to that reported by lattice clocks prior to this work. The JILA Sr1 lattice clock [74] and the NIST Yb lattice clock [63] are also shown here.

4.3 Stability in Optical Clocks

4.3.1 The Stability of Lattice Clocks and Ion Clocks

With QPN and the Dick Effect explained, the stability of the two prominent types of optical clocks, ion clocks and lattice clocks (Section 1.5), can now be compared. Regarding QPN, ion clocks tend to use shorter clock transition wavelengths [23, 93, 52] than lattice clocks [63, 74], and this is beneficial for ion-based systems due to the factor of Q in σ_{QPN} (Equation 4.6). This is also why optical clocks are better than Cs, which is based on a microwave transition. Furthermore, ion clocks can operate with faster cycle times [23], which is a positive trait of these clocks due to a superior $\sqrt{T_c}$ factor in the σ_{QPN} expression. However, ion clocks only have one atom whereas N is of order 10^3 for lattice clocks. Since $\sigma_{QPN}(\tau) \propto 1/\sqrt{N}$, the larger atom number is a major advantage of lattice clocks. Due to larger atom numbers, the QPN of lattice clocks is typically at least an order of magnitude better than that of ion clocks.

The stability limitation of ion clocks tends to be QPN [54]. For lattice clocks, QPN is small and duty cycles are shorter (AKA more dead time), so the Dick Effect dominates. Prior to the work discussed in this thesis, the Dick-effect-limited stability of lattice clocks was about equal to reported ion clock stabilities (Figure 4.5). At the time, the laser causing the Dick Effect in the Sr1 clock was based on a 7 cm long cavity [72], which was similar to those used by other lattice clock teams.

4.3.2 Using Longer Cavities to Reduce Thermal Noise

When the Sr1 7 cm clock laser stabilization cavity was built, this new system made use of a clever cavity mounting scheme to remove ambient mechanical vibration noise [91]. Prior to these advances in cavity mounts, vibration noise was the dominant mechanism limiting cavity stability [126]. When vibration noise dominates, cavities should be made short since the amplitude of this noise scales with the size of the cavity. With vibration noise removed, the best stable lasers became limited by thermal fluctuations of cavity mirror coatings [92]. This *thermal noise* degrades cavity length stability, which negatively affects the frequency stability of clock lasers (since cavity stabilization is the stabilization of the laser frequency to the length of the cavity).

The 7 cm cavity in Sr1 was designed to be short to reduce residual vibration noise not canceled by the cavity mounting scheme. The tradeoff for using a short cavity is more thermal noise. The stability of a laser locked to a thermal-noise-limited cavity scales as $\delta L/L$, where L is the cavity length and δL is the fluctuation in L [90]. Although longer cavities could potentially suffer from worse residual vibration noise, if vibration noise cancellation could be made to work well for a longer cavity, this would yield better clock laser stabilization and a reduced Dick Effect.

This possibility motivated the construction of a longer 40 cm cavity at JILA to improve the stability of our lattice clock (Section 3.6.1) [77]. This cavity proved to be vibration insensitive and thermal noise limited at stability an order of magnitude better than that of the 7 cm cavity (Figure 4.6). Early characterizations of the clock laser based on the 40 cm cavity revealed a free-running

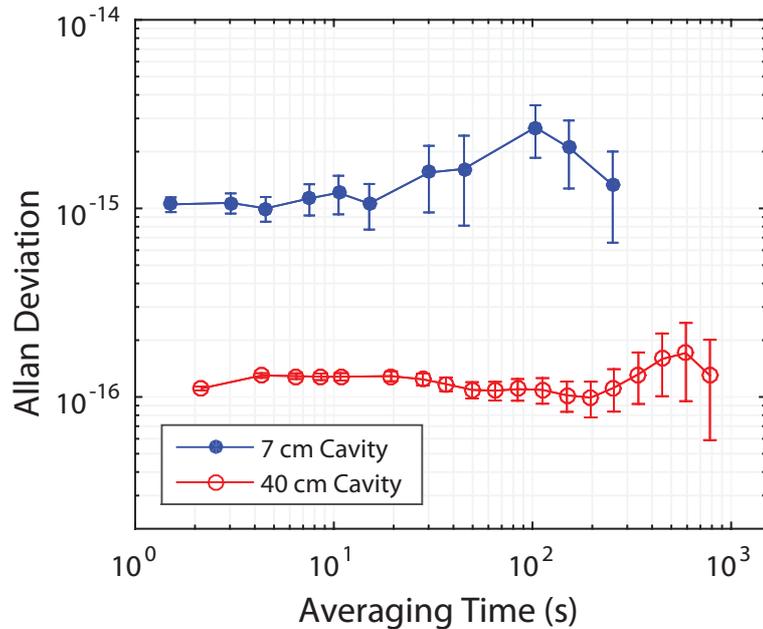


Figure 4.6: Overlapping Allan deviations of diode lasers stabilized to 7 cm and 40 cm cavities. Linear drift is removed to show the laser instabilities due to high frequency noise.

power spectrum⁷ with a flat Allan deviation, indicating a $1/f$ noise process [89]. This is consistent with a thermal-noise-limited laser, which also has a $1/f$ noise spectrum [92].⁸

For a laser with a $1/f$ noise spectrum, Equation 4.3 can be used to write $S_{clock}(f)$ in terms of the Allan deviation of the free-running laser, σ_{clock} ,

$$S_{clock}(f) = \frac{\sigma_{clock}^2}{(2 \ln 2)f}. \quad (4.12)$$

Using this expression, the Dick Effect is

$$\sigma_{Dick}^2(\tau) = \frac{\sigma_{clock}^2 T_c}{2 \ln 2 \tau} \sum_{k=1}^{\infty} \frac{1}{k} \left| \frac{g_k}{g_0} \right|^2. \quad (4.13)$$

Therefore, since the free-running laser stability with the 40 cm cavity is 10 times better than with the 7 cm cavity, and since the Sr1 stability with the 7 cm cavity was comparable to that of ion clocks, we expect this new cavity to make Sr2 clock stability an order of magnitude better than ion

⁷ The “free-running power spectrum” is the power spectrum of the clock laser stabilized to the cavity but not to the atoms.

⁸ A later, more detailed study revealed $1/f$ noise, white noise, and resonances in the power spectrum [8].

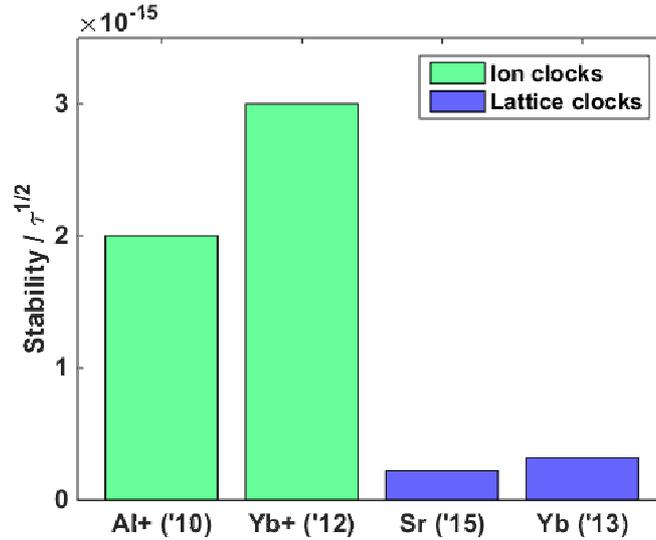


Figure 4.7: Best reported Sr and Yb stability results compared to the ion clock stabilities in Figure 4.5. Both the Sr and Yb advances were primarily due to better laser stabilization. The Yb result is reported in Reference [49], and the Sr2 result is discussed in Section 4.6.

clock stability.

As I will explain in the rest of this chapter, this order-of-magnitude improvement over ion clock stability is indeed what we have observed (Figure 4.7).

4.3.3 Self-Comparison

Self-comparison is a useful technique for characterizing clock stability. With self-comparison, a clock compares two independent atomic servos that are locked to the same stretched state and that update in an alternating fashion. If processed correctly, the center frequencies measured by the two servos yield the short-term stability of the clock.

Suppose that one atomic servo measures the center frequency ν_1 , and immediately after that the other servo measures ν_2 . Since these are scanning synthesizer frequencies, they contain the free-running clock laser drift (due to stabilization cavity length variations) that is not removed by drift cancellation (Section 3.7.1).⁹ To remove this drift, the center frequencies are processed as a

⁹ The atomic servo records the scanning synthesizer frequency required to set the free-running clock laser (after drift cancellation) on the clock transition resonance. In the presence of uncanceled clock laser drift, the servo will adjust the scanning synthesizer frequency to compensate for this drift, adding drift to the record of scanning

difference $\nu_2 - \nu_1$. This works because the uncanceled laser drift is typically negligible over the $4T_c$ it takes to make the two center frequency measurements. This difference also removes long-term drift of systematic shifts, so it does not determine long-term stability. However, $(\nu_2 - \nu_1)/2$ determines the short-term stability of the clock.

$\nu_2 - \nu_1$ is divided by 2 to extrapolate to the noise performance of the stretched state average frequency. Since the stretched state average frequency is an average of two atomic servo measurements (Section 2.5.3), its noise (using the usual error propagation mathematics) is $\sqrt{(\sigma/2)^2 + (\sigma/2)^2} = \sigma/\sqrt{2}$, where σ is the noise of a single center frequency measurement. The noise of $\nu_2 - \nu_1$ is $\sqrt{2}\sigma$, so this difference must be divided by 2 to have the same frequency fluctuation as the stretched state average frequency. Therefore, for frequency data $\{\nu_1, \nu_2, \nu_3, \dots\}$, where the odd subscripts indicate data taken by one atomic servo and the even subscripts are taken by the other, the self-comparison data is $\{\frac{1}{2}(\nu_2 - \nu_1), \frac{1}{2}(\nu_4 - \nu_3), \dots\}$

4.4 Synchronous Interrogation

4.4.1 Two Sr Clocks with a Shared Clock Laser Source

To study the stability of a frequency standard at different timescales, it is useful to beat this standard with a similar system. Therefore, we measure the Sr2 clock stability through a series of comparisons with the Sr1 system [89]. These comparisons were set up using a shared clock laser source.

As explained in Section 3.6.2, light from the clock laser is divided between the two experiments into two independent clock laser beams. Each system has its own AOM to realize independent atomic servos (Figure 4.8a). As the frequencies of the each scanning synthesizer (one for the Sr2 AOM and one for Sr1) are adjusted to lock the two clock laser beams, these synthesizer frequencies are recorded.¹⁰ One advantage to this setup is that a comparison between Sr1 and Sr2¹¹ can be

synthesizer frequencies.

¹⁰ Both synthesizers share a stable RF reference.

¹¹ Here “comparison” simply means the difference between the frequencies of the two clock laser beams after they are locked to their respective Sr clock

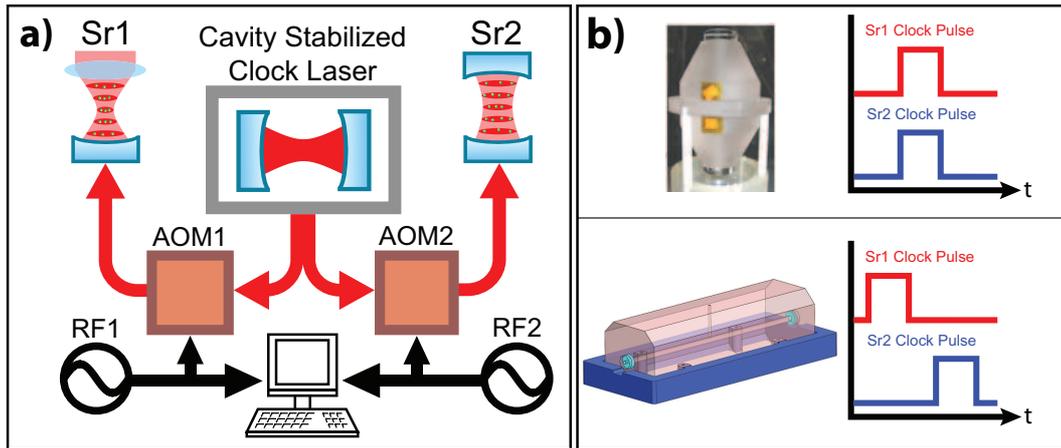


Figure 4.8: a) Comparison setup for the two JILA Sr clocks. b) Top panel: Stability with the 7 cm cavity (image on left) is studied using synchronous pulses (plot on right). Bottom panel: Asynchronous pulses (plot on right) are used to study stability with the 40 cm cavity (image on left). The Sr1 falling pulse edge and the Sr2 rising edge are separated by 10 ms for asynchronous interrogation. The synchronous and asynchronous measurements use 160 ms Rabi π -pulses for both systems.

performed in post processing. Since the two systems share a common clock laser source, taking the difference between the recorded values of the two scanning synthesizer frequencies provides the same information as beating the two clock laser beams (Figure 4.8a).

To rigorously study the stability with the new cavity (Figure 3.18), we use the setup shown in Figure 4.8a to perform comparisons between Sr2 and Sr1 with lasers stabilized to either the 7 cm or 40 cm cavities. Measurements with the 7 cm cavity made use of *synchronous interrogation* [9, 113]. In this case, the Rabi pulses of both Sr clocks were carefully synchronized such that they always occur at the same time (Figure 4.8b top). The PID filters of both systems are also set to have equivalent gain constants (Section 2.44), ensuring that both clocks respond to noise in the same way. Pulse synchronization is accomplished by triggering Sr1 off of the Sr2 clock.

Since the synchronous clock comparison occurs in post processing, data analysis must correctly match each Sr2 frequency measurement with the simultaneous Sr1 measurement; therefore, sorting errors in the data record could destroy the comparison. To ensure that the data record is robust against such errors, both systems used a shared counter (triggered by Sr2 to advance by 1

each measurement cycle) to stamp each measurement with a counter value. When comparing the Sr1 and Sr2 data, counter values can be checked to ensure that they agree.

With synchronous interrogation, the two clocks always sample the same laser noise. Therefore, the difference between the clock frequencies benefits from common-mode laser noise cancellation, which will suppress the Dick Effect. Also, studying the effect that synchronous interrogation has on stability is useful for quantifying the Dick Effect [89].

Synchronous interrogation is also useful for systematic evaluations. Systematic shifts can be evaluated using two synchronized clocks, where one clock modulates a parameter known to cause systematic bias and measures the resulting change in frequency with respect to the other clock. This would result in a systematic evaluation that is largely free of the Dick Effect [116]. However, a clock used for timekeeping is a stand alone system that has stability limited by the frequency noise of the clock oscillator. To quantify the stability of a system for use as a clock, synchronous comparisons will not work since they do not show the stability of independent systems. For this, we use *asynchronous interrogation*, which will be discussed in Section 4.5.1.

4.4.2 Excited State Fraction Correlation

Synchronous interrogation using the 7 cm cavity shows a clear effect on raw frequency data (Figure 4.9). This measurement scheme should also affect the excited state fraction measurements of these two systems [113]. When the two clocks are synchronized, clock laser noise will cause measurements of ρ_{ee} to fluctuate about 1/2, but synchronization means that these fluctuations will be correlated between the two clocks.

To understand the effect of synchronous interrogation on the excited state fractions, it is worthwhile to simulate what the corresponding measurements would look like. Treating laser fluctuations as a Gaussian white noise process, I simulate the excitation fraction data when both Sr1 and Sr2 are locked to set points of 1/2 (Figure 4.10). When the two atomic servos use pulses that are not synchronized, both laser noise and QPN (simulated for 1000 atoms) are uncorrelated. This is apparent because of the circle-shaped distribution of the simulated data in Figure 4.10a. Here an

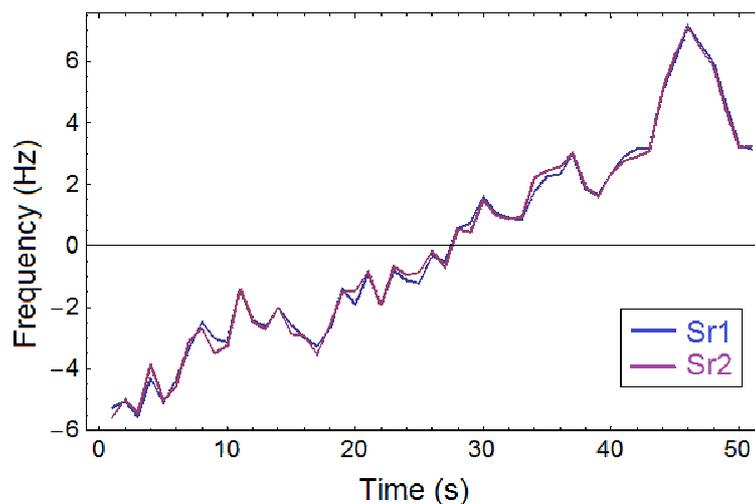


Figure 4.9: Raw data measured by synchronized clocks. The scanning synthesizer frequencies for both Sr1 and Sr2 during synchronized atomic servo operation are plotted with offsets between them removed. The Rabi π -pulses are 160 ms long. The y axis has been offset by the scanning synthesizer frequency at about 28 s. The linear laser drift removed by the atomic servos is apparent here. A striking correlation between the frequency data is apparent. For uncorrelated data, the average frequencies would agree but the instantaneous frequencies would not.

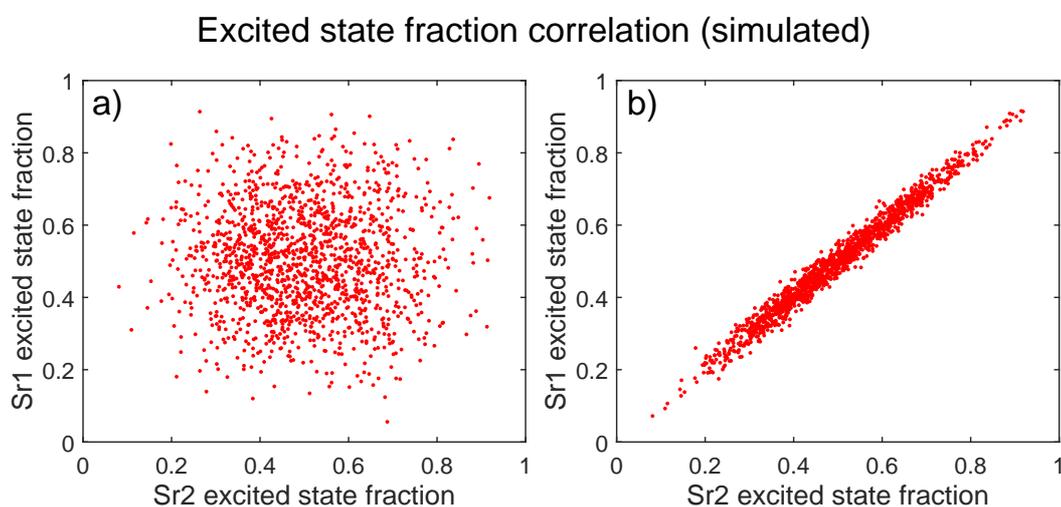


Figure 4.10: Simulated correlation measurement data. Appreciable laser noise is assumed. a) The two clocks are locked using pulses that are not synchronized. b) The clocks are locked with synchronized pulses.

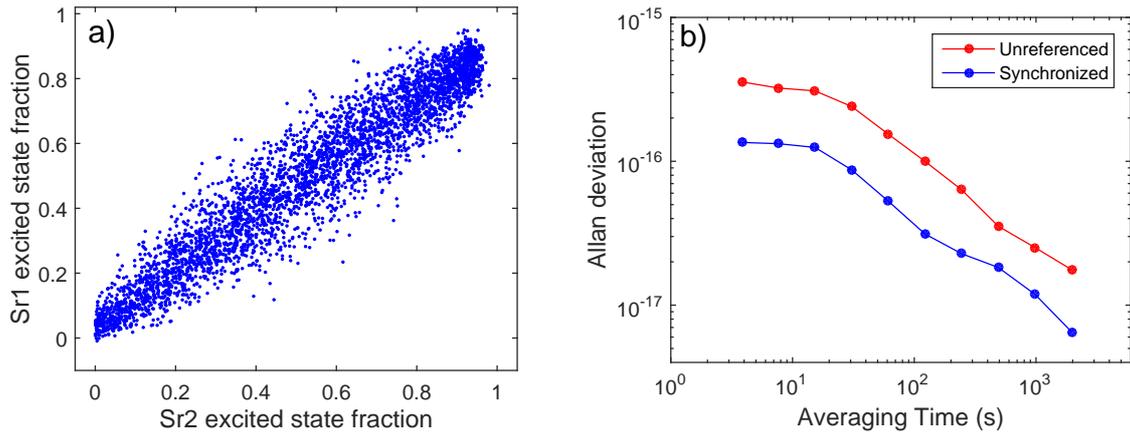


Figure 4.11: a) Correlation data using a clock laser stabilized to the 7 cm cavity. The width is due to uncorrelated noise, which can come from QPN, imperfect synchronization, or technical noise not common to both clocks. b) Overlapping Allan deviations of the Sr2 self-comparison data with and without the Sr1 reference. The 7 cm cavity is also used for this measurement.

excited state fraction measurement in one clock does not necessitate a value for the other system. When pulses are synchronized, the data has a diagonal stripe distribution (Figure 4.10b). Here laser noise causes the excited state fraction measurements to fluctuate, but the synchronization of the two clocks ensures that measured values are similar for both systems. The width of the horizontal stripe is given by quantum projection noise, which remains uncorrelated.

Using the noisier 7 cm cavity to test this situation, the 160 ms π -pulses from both clocks are synchronized and the same clock laser source is used. A correlation is observed in the data, showing that synchronous interrogation has succeeded in making much of the laser noise common mode between the two clocks (Figure 4.11a) [89].

The effect of synchronous interrogation on clock stability can also be seen with Allan deviations. This is done by running Sr2 and Sr1 in a synchronized fashion while Sr2 takes self-comparison data. Self-comparison data fits to $1.2 \times 10^{-15} / \sqrt{\tau}$ (Figure 4.11b, unreferenced data), where τ is in seconds, and the first few points shown on the unreferenced Allan deviation of Figure 4.11b have been ignored in the fit. If the Sr1 center frequency data is subtracted from the Sr2 measurements (where the difference is taken between points measured simultaneously), the boost in stability is

appreciable (Figure 4.11b, synchronized data). The stability now fits to $4.6 \times 10^{-16}/\sqrt{\tau}$, which is better than the unreferenced case by a factor of 2.6. This corresponds to a factor of 6.8 less averaging time.

The fact that the first few points of these Allan deviations are flatter is because for the first ~ 20 s, the Allan deviation is that of the free-running laser. After a few updates of the atomic servo, the lock to the atoms takes over and the stability begins averaging down. Because of this, Allan deviations limited by the Dick Effect or QPN are typically fit by excluding the points before the servo becomes effective and then performing regression to a $1/\sqrt{\tau}$ power law (and they will be fit this way from here on). Clock stability is usually quoted in terms of this fit.¹²

4.5 Asynchronous Interrogation

4.5.1 Using Asynchronous Interrogation to Measure Independent Clock Stability

As explained in Section 4.3.2, given the Allan deviation of the free-running clock laser based on the 40 cm cavity (Figure 4.6), we expect that upgrading to this cavity will provide a significant boost in stability (over the laser based on the 7 cm cavity). As the previous section shows, one can use synchronous interrogation to gauge the Dick Effect with this new laser. Also, after upgrading to the 40 cm cavity, it is important to measure the independent clock stability, which characterizes the Sr2 system as a frequency standard (Section 4.4.1). To do this, we use *asynchronous interrogation*, which ensures that the Sr1 and Sr2 π -pulses never overlap in time. With this interrogation scheme, the falling edge of the Sr1 pulse and the rising edge of the Sr2 pulse are separated by 10 ms (Figure 4.8b, bottom) so that the two clocks never sample the same laser noise.

It is important to note that two asynchronously interrogated systems are not 100% independent. Since the pulses have a fixed temporal relationship to one another, spectral components of the clock laser with frequencies that are integer multiples of $1/\Delta t_p$, where Δt_p is the separation between the pulse centers, will affect asynchronous interrogation differently than they would an

¹² For instance, “1 s stability” typically means the fit quoted at $\tau = 1$.

independent clock. To analyze this issue,¹³ it is useful to express the sensitivity function $g(t)$ (Section 4.2.2) in terms of the excited state fraction [106],

$$\rho_{ee}(t) = \frac{1}{2} + \frac{1}{2} \int g(t) \epsilon_\nu(t) dt, \quad (4.14)$$

where $\epsilon_\nu(t)$ is the stochastic frequency fluctuation of the laser about the clock transition frequency. The integral is taken over the clock cycle under consideration. This expression describes the behavior of ρ_{ee} when the clock laser is locked to the atoms. Here it is assumed that ρ_{ee} is locked to a value of 1/2. This expression is a first-order approximation to the effect of laser noise on the excited state fraction, treating $\epsilon_\nu(t)$ as a small-amplitude perturbation.

When ρ_{ee} deviates from 1/2, the atomic servo error signal $\epsilon_{clock} \simeq 2(\rho_{ee} - 1/2) = \int g(t) \epsilon_\nu(t) dt$ (using Equation 4.14). Since an asynchronous comparison is the difference between center frequencies computed by the Sr1 and Sr2 atomic servos—both of which operate with the same pulse durations, gain constants, and cycle times—it follows from Equations 2.44 and 4.14 that the noise in the comparison is given by the integral $\int [g_{Sr2}(t) - g_{Sr1}(t)] \epsilon_\nu(t) dt = \int [g(t + \Delta t_p) - g(t)] \epsilon_\nu(t) dt$. Here, $g_{Sr2}(t)$ and $g_{Sr1}(t)$ are the Sr2 and Sr1 sensitivity functions. The expression $g(t + \Delta t_p) - g(t)$ can be viewed as the sensitivity function of the asynchronous comparison.

Using the fact that $g(t)$ is periodic in T_c , the Fourier coefficients (Equation 4.10) of $g(t + \Delta t_p)$ are $g_k e^{2\pi i k \Delta t_p / T_c}$ (where g_k is the k th Fourier coefficient of $g(t)$). This leads to the Dick Effect for an asynchronous comparison σ_{asynch} of

$$\begin{aligned} \sigma_{asynch}^2(\tau) &= \frac{1}{\tau} \sum_{k=1}^{\infty} \left| e^{2\pi i k \Delta t_p / T_c} - 1 \right|^2 \left| \frac{g_k}{g_0} \right|^2 S_{clock}(k/T_c) \\ &= \frac{4}{\tau} \sum_{k=1}^{\infty} \sin^2\left(\frac{\pi k \Delta t_p}{T_c}\right) \left| \frac{g_k}{g_0} \right|^2 S_{clock}(k/T_c). \end{aligned} \quad (4.15)$$

Replacing the $\sin^2()$ factor with its average of 1/2 results in the Dick Effect for two identical but independent clocks, for which one would simply add two identical single-clock Allan deviations

¹³ My thanks to Michael J. Martin for working this out first.

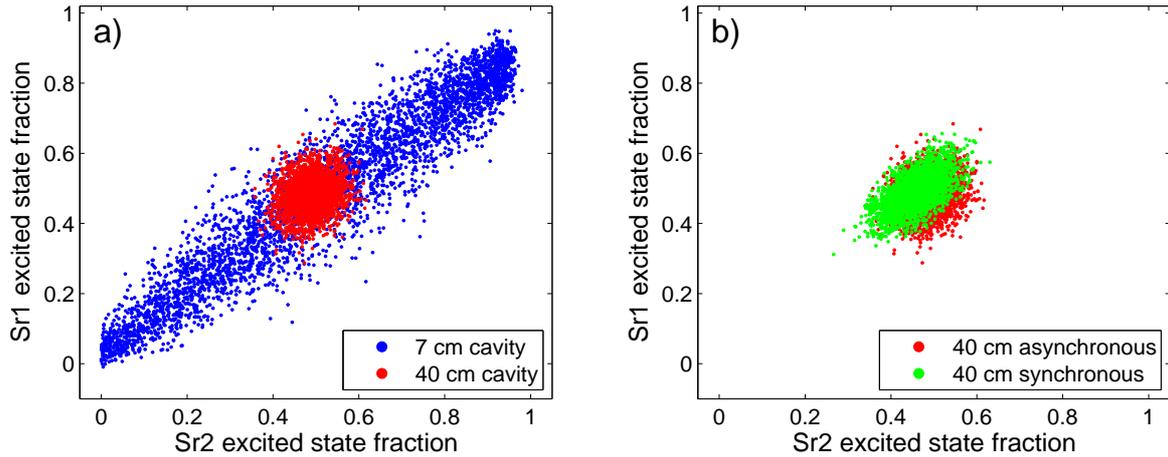


Figure 4.12: a) Synchronous data with the 7 cm cavity (from Figure 4.11a) and asynchronous data with the 40 cm cavity. The 40 cm cavity data shows a markedly narrower distribution, implying less laser noise. b) Synchronous and asynchronous data using the 40 cm cavity. Synchronous interrogation shows a bit of correlation.

in quadrature. Including this factor and using a good model for the clock laser power spectrum [8] changes the Dick Effect by 3% from its independent clock value. Therefore, asynchronous interrogation reproduces independent clock stability at the 97% level.

4.5.2 Correlations and Allan Deviation

Measuring excited state fraction correlations with the clock laser based on the 40 cm cavity, the distribution shows a dramatic difference. In Figure 4.12a, asynchronous data with the 40 cm cavity is compared with synchronous measurements using the 7 cm cavity [89]. The distribution is much narrower, implying less laser noise. Synchronous measurements with the 40 cm cavity show a bit of correlation (Figure 4.12b), indicating that the Dick Effect has not been completely suppressed below QPN [89].

Looking at Allan deviations of synchronous and asynchronous data taken with the 40 cm cavity, a significant stability improvement over the 7 cm cavity is observed (Figure 4.13). Self-comparison is also used to characterize the stability. Looking at self-comparison and synchronous data on the same plot (Figure 4.13), it is apparent that the Dick Effect plays a much smaller role

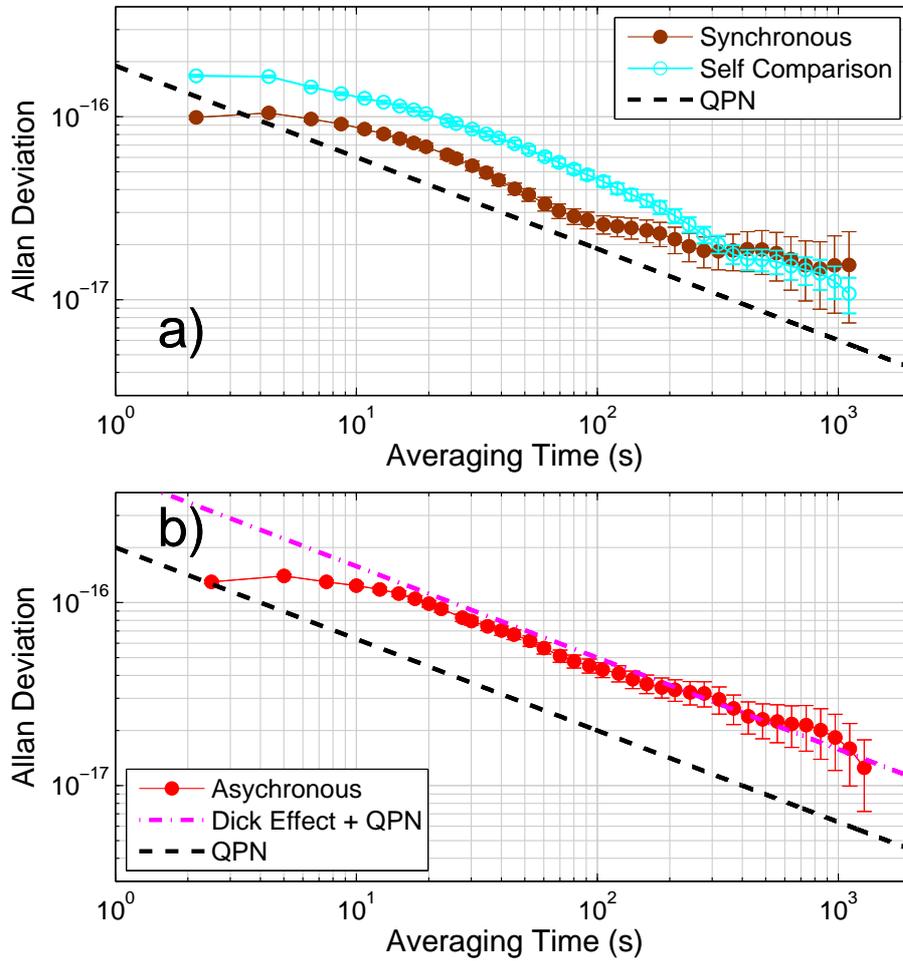


Figure 4.13: Overlapping Allan deviations with the 40 cm cavity. a) Synchronous and self-comparison data. b) Asynchronous data, which gives the stability of a comparison between two independent clocks. The Allan deviation fits to $4.4 \times 10^{-16}/\sqrt{\tau}$ for the asynchronous comparison or, extrapolating to the performance of a single clock, $3.1 \times 10^{-16}/\sqrt{\tau}$. Using the known clock laser power spectral density [8], the theoretical one-clock Dick Effect is $3.5 \times 10^{-16}/\sqrt{\tau}$.

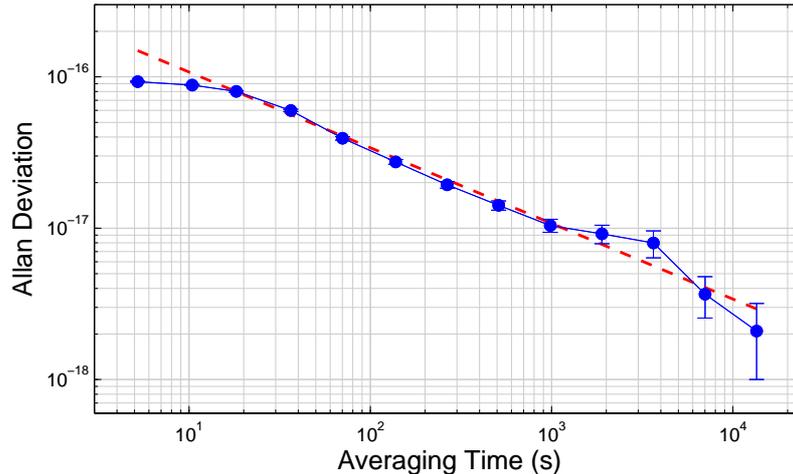


Figure 4.14: Overlapping Allan deviation for a 3.75 hour asynchronous comparison. 160 ms Rabi π -pulses were used again here.

in the stability [89].¹⁴ Synchronous interrogation now results in stability that is only a bit below the self-comparison (compared to the much larger effect shown in Figure 4.11b).

However, these measurements are only complementary to that shown in Figure 4.13b. This is an asynchronous comparison between Sr1 and Sr2, which yields the full independent clock stability. It is sensitive to long-term drift (that is not common to the two clocks), so it demonstrates the long-term stability as well. Drift, which would be evident by the Allan deviation increasing linearly with averaging time (Figure 4.1), is not present in the asynchronous data. This measurement achieves $3.1 \times 10^{-16}/\sqrt{\tau}$ one-clock stability (the stability of a comparison between two identical clocks divided by $\sqrt{2}$). When this stability was reported [89], it was the first instance of a lattice clock surpassing all ion clocks in an important measure of clock performance. The following year, the Yb lattice clock team at the Boulder campus of NIST also reported stability of $3.1 \times 10^{-16}/\sqrt{\tau}$ and averaged this down for 7 hours to 1.6×10^{-18} [49].

To study the Sr2 stability at longer averaging times, which ensures that drift is well controlled, Sr1 and Sr2 are compared asynchronously for 3.75 hours [13]. No drifts are apparent; however, the stability was a bit worse for this measurement, fitting to a value of $3.4 \times 10^{-16}/\sqrt{\tau}$ (Figure 4.14).

¹⁴ Recall that self-comparison data is sensitive to the Dick Effect whereas synchronous measurements are not.

4.6 High Stability Self-Comparison

So far, all stability measurements discussed used 160 ms π -pulses with $T_c \simeq 1$ s. To infer the best stability possible in Sr2, it is worthwhile to make the duty cycle larger, which reduces the Dick Effect (Section 4.2.2). However, the best possible stability of Sr2 should, in principle, be better than that of Sr1. An attempt to quantify the maximum Sr2 stability by comparing the two clocks would be difficult once Sr2 has significantly better performance.

The reason why Sr2 should have higher stability is because of its cavity lattice. As explained in Section 2.3.3, the cavity lattice is designed to have a large trap volume compared to the retro-reflected lattice in Sr1. The advantage of this is twofold. First, the large volume means that the cavity lattice has a much better spatial overlap with the 689 nm MOT, allowing Sr2 to trap more atoms than Sr1; therefore, Sr2 can obtain 2000 atoms with less MOT cooling time, resulting in less dead time and a bigger duty cycle. Second, the larger volume implies that if Sr2 and Sr1 operate with 2000 atoms, which is the number used for clock operation, then Sr2 will have lower density (by a factor of 30 for the actual trap parameters). This lower density allows Sr2 to use long pulses while still operating at 2000 atoms. This is useful because Sr1 has observed resonance contrast decay, broadening, and lineshape distortion of the clock transition due to atomic interactions [6, 78]. These effects (which compromise stability) increase with density, and for typical operating atom numbers, Sr1 was unable to achieve unbroadened lineshapes for pulse durations of 1 s and longer [78, 77].

Unfortunately, it is very inconvenient to upgrade the Sr1 lattice to a cavity-based trap because the vacuum chamber viewports on that system do not have anti-reflection coating at 813 nm. This coating on the Sr2 viewports (Section 3.1.2) is necessary to achieve good cavity enhancement, which is negatively affected by power loss through the viewports (Equation 2.14). Therefore, we must evaluate the Sr2 stability with self-comparison, which does not rely on another clock.

Since self-comparison is insensitive to long-term drift (Section 4.5.2), this technique alone cannot determine the full independent clock stability. However, if we use previous comparisons to

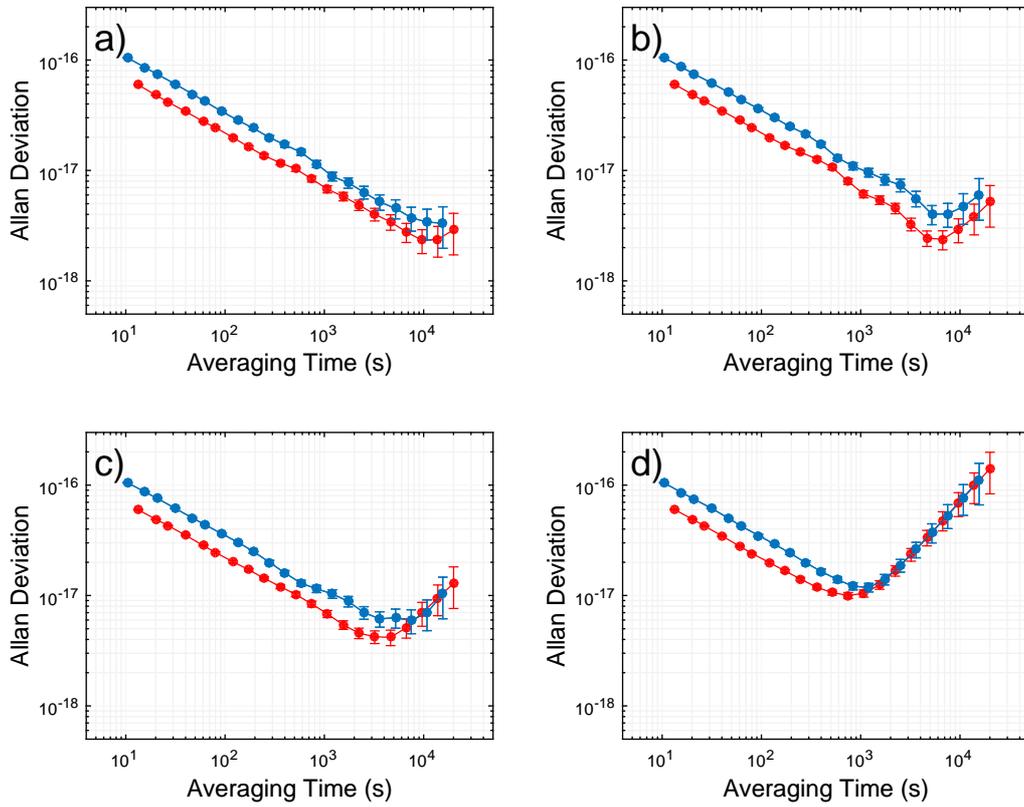


Figure 4.15: Simulations of stability data with drift. The noise is taken to be white, which is valid after the few 10s of seconds it takes for the atomic servo to kick in. The blue data has $3.4 \times 10^{-16}/\sqrt{\tau}$ stability, equivalent to that in Figure 4.14, and the red data fits to $2.2 \times 10^{-16}/\sqrt{\tau}$. The cycle times of both data sets match those used in our experiments. a) A drift rate of 1.4×10^{-18} over the measurement time. Drift is not easy to discern from either Allan deviation. The red data is turning up slightly, which could easily be mistaken for a different feature. b) 4.1×10^{-18} drift over the measurement time is apparent for both curves. c) 1.4×10^{-17} drift over the measurement time. d) 1.4×10^{-16} drift over the measurement time.

constrain the magnitude of long-term drift, self-comparison can determine the full stability at a level given by this constraint (Figure 4.15). Figure 4.14 shows that the total long-term drift over 3.75 hours of averaging is below 3×10^{-18} , or a $2.2 \times 10^{-22} \text{ s}^{-1}$ drift rate. Therefore, self-comparison stability is equal to the independent clock stability—so long as it is understood that the former should not be extrapolated below 3×10^{-18} or longer than 3.75 hours of averaging time (whichever comes first).

Using 1 s pulses and reducing MOT cooling time to achieve a 60% duty cycle¹⁵, the self-comparison (Figure 4.16) demonstrates one-clock stability of $2.2 \times 10^{-16} / \sqrt{\tau}$ [88]. This value, which is compared in Figure 4.7 with other prominent optical clocks, is the best stability per $\sqrt{\tau}$ achieved to date.

¹⁵ With two clocks operating at 50% duty cycle, one can realize a proposal by Dick, et al. [29], which was demonstrated in the microwave by Biedermann, et al. [5]. This involves using two clocks operating in an interleaved fashion. When the first system measures the clock transition, the second is accruing dead time for sample preparation. Then they switch roles, ensuring that the clock transition is always being measured and reducing dead time to zero. This would affect the Dick-Effect-limited stability dramatically.

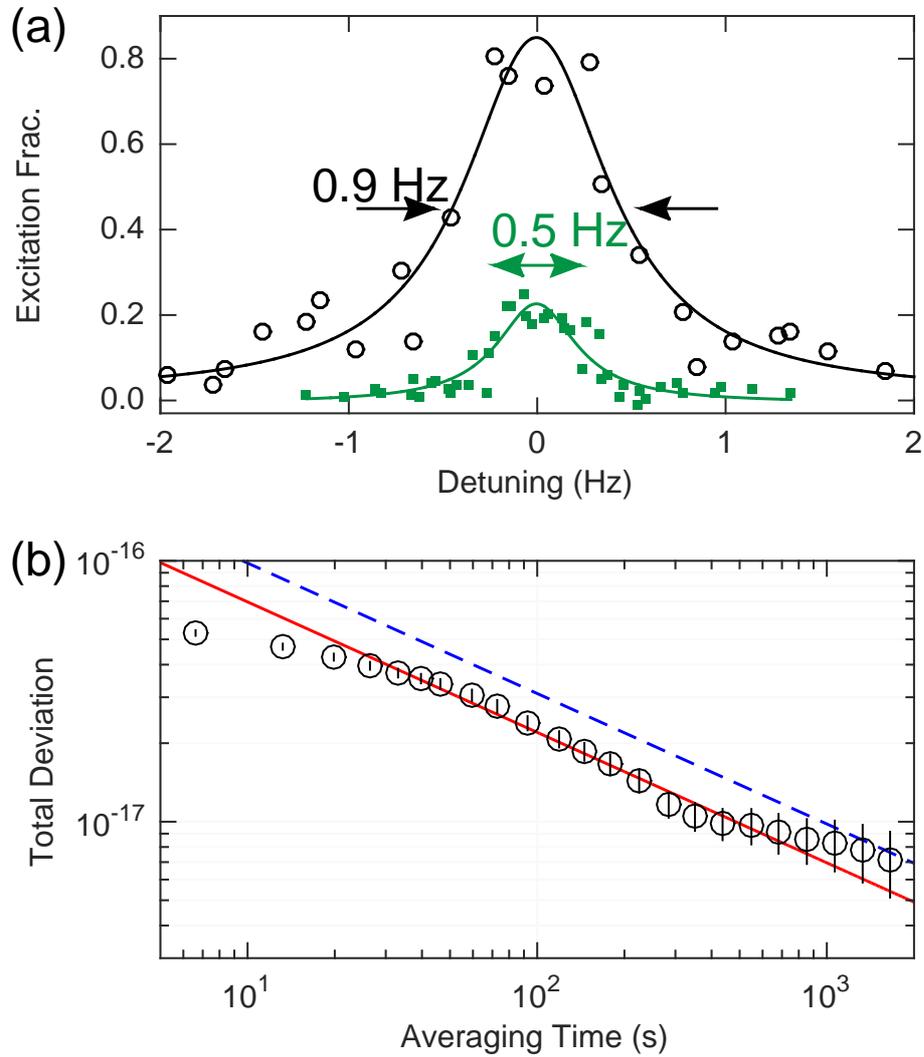


Figure 4.16: a) The effects of pulse duration on the clock transition. A typical line scan with a 1 s interrogation time is shown in open black circles. To explore the limit of coherence in Sr2, we scan the clock transition with a 4 s interrogation time and more atoms (solid green squares). Here the linewidth and contrast are affected by the Fourier width and atomic interactions [6, 78]. b) A new stability record (black circles) achieved by running with 1 s clock pulses and a 60% clock laser duty cycle for each preparation and measurement sequence. The stability fits to $2.2 \times 10^{-16}/\sqrt{\tau}$ (red line). The previous result is indicated with the blue dashed line.

Chapter 5

Systematic Uncertainty

5.1 Systematic Measurements

5.1.1 Introduction

Improved stability, demonstrated in Chapter 4, is critical for better systematic shift measurements. This is because systematic uncertainty is often measured by the clock itself (Section 4.1.1). A parameter known to cause a systematic shift (such as the lattice intensity or bias field) is modulated, and the frequency shift that results from changing this parameter is measured with a scheme similar to self-comparison. These measurements, which are typically limited by the stability of the clock, are extrapolated to operating conditions. Although most major systematic shifts are measured in this way, there is one important exception: the blackbody radiation (BBR) shift, for which a completely different approach must be taken.¹

Systematic uncertainty arises because in practice the clock laser is locked to the perturbed clock transition, but we want to extrapolate the locked clock laser frequency to the bare transition (Sections 1.6 and 2.7). With all systematic frequency shifts well known, real-time corrections could be applied to Sr2 so that, when running this system as a clock (with the frequency comb counter enabled), it ticks at a rate determined by the bare frequency.

¹ Other teams have measured the BBR shift by looking at the effect of temperature modulation on their clock transition frequency [2, 116]. This requires the ability to change the ambient temperature in a spatially uniform manner, and the Sr2 clock is not equipped for this kind of temperature control. This approach to the BBR shift is different than other systematic measurements performed in Sr2 (which are measured with a single clock) because it requires the clock that is under evaluation to be referenced to a second, highly stable frequency standard. A reference is needed because temperature cannot be modulated quickly, so drift of the clock laser cavity would ruin a BBR shift measurement with only one clock. Referencing to a second system removes this drift from the evaluation.

Sr2 has performed two systematic evaluations. Prior to these, the record for the best atomic clock² was 8.6×10^{-18} total systematic uncertainty [23]. The first Sr2 systematic evaluation broke this record with 6.4×10^{-18} total systematic uncertainty. The second evaluation, which began with the goal of realizing low- 10^{-18} -level performance, secured the standing of Sr2 as the best clock by achieving 2.1×10^{-18} total uncertainty.

Effect	Evaluation 1		Evaluation 2	
	Shift	Uncertainty	Shift	Uncertainty
1st-order Zeeman	-0.2	1.1	-0.2	0.2
2nd-order Zeeman	-144.5	1.2	-51.7	0.3
Probe Stark	0.8	1.3	0.0	0.0
Density	-4.7	0.6	-3.5	0.4
Lattice Stark	-461.5	3.7	-1.3	1.1
Dc Stark	-3.5	2.1	0.0	0.1
Static BBR	-4962.9	1.8	-4562.1	0.3
Dynamic BBR	-345.7	3.7	-305.3	1.4
AOM phase chirp	0.6	0.4	0.6	0.4
Servo offset	0.4	0.6	-0.5	0.4
Line pulling and tunneling	0	< 0.1	0	< 0.1
2nd-order Doppler	0	< 0.1	0	< 0.1
Background Gas Collisions	0	< 0.6	0.0	< 0.6
Total	-5921.2	6.4	-4924.0	2.1

Table 5.1: The uncertainty budget for both evaluations. The effects are listed in the order that they are discussed in this chapter. All uncertainties are quoted at the 1σ level.

5.1.2 Digital Lock-In Detection

The act of modulating a parameter to observe the change in the clock transition frequency is similar to lock-in detection. The parameter under consideration is modulated at a certain rate, and then the data record is demodulated at this rate in post processing, suppressing effects that do not change with the modulation. From here on this procedure will be referred to as “lock-in detection.”

An example of how this works for systematic evaluations is the following. A chosen parameter (such as the lattice intensity or bias field magnitude) switches between two states. Each state of

² The word “best” is often used to describe the clock with the lowest total systematic uncertainty despite that stability is also an important figure of merit for clock performance.

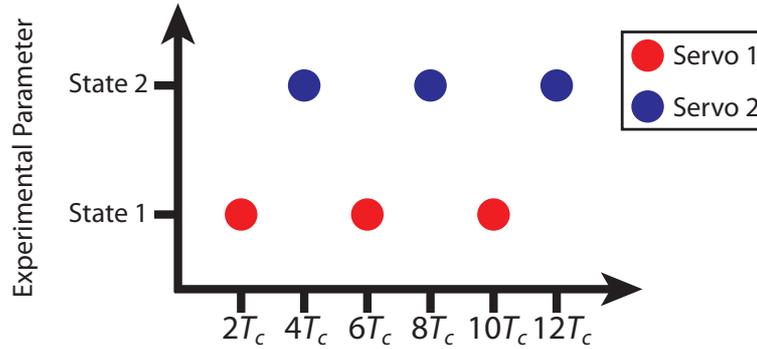


Figure 5.1: The lock-in measurement. A parameter is modulated by switching between two values (“State 1” and “State 2” in the Figure). If the cycle time is T_c , which involves one interrogation of the clock transition and an excited state fraction measurement, it requires $2T_c$ to measure the scanning synthesizer frequency corresponding to the resonance center (Section 2.5.2). When the parameter is in State 1, the clock transition is probed on either side of resonance, and the line center is determined by the digital PID filter (Equation 2.44). The parameter is then switched to State 2, shifting the clock transition resonance center, which is determined by the second servo. The red and blue points in the Figure depict the time and parameter value when the line center is determined.

this parameter corresponds to an independent atomic servo locked to the clock transition, and the lock-in alternates between these two servos (Figure 5.1). Aside from this modulation parameter, the atomic servos operate in the same conditions.

Suppose the scanning synthesizer center frequency data taken in the lock in measurement is $\{\nu_1, \nu_2, \dots, \nu_M\}$, where odd (even) subscripts mean that the data was taken with the lock-in parameter in State 1 (2), and M is always an even number. For each of these points, the lock-in parameter is known (either by measuring it each time a frequency point is acquired or by stabilizing the two lock-in states to be constant throughout the measurement). The lock-in data is $\Delta\nu_{lock-in} = \{\nu_2 - \nu_1, \nu_4 - \nu_3, \dots, \nu_M - \nu_{M-1}\}$. If clock laser drift is canceled well enough (Section 3.7.1) such that it is negligible over the $4T_c$ it takes to measure one lock-in point, then the lock-in data is equal to the clock transition frequency difference between the two states of the modulated experimental parameter.³

³ Many Sr lattice clock groups have performed systematic measurements while referenced to another clock; however, with good drift cancellation and point-string analysis (Section 5.1.5), the evaluations discussed in this thesis do not require a second clock as a reference.

5.1.3 Lever Arm

Extracting the information needed from lock-in data requires a model of how these data vary with the lock-in parameter. As an example, take the probe ac Stark shift. This is the systematic Stark shift caused by the clock laser. It is known that this shift is proportional to the clock laser power P . The proportionality constant K is related to atomic structure and experimental conditions such as the clock laser intensity distribution, so K is not known a priori. A lock-in point $\nu_2 - \nu_1$, taken when the power is modulated between P_1 and P_2 , is equal to $K(P_2 - P_1)$. Suppose P_0 is the clock laser value during normal clock operation. The systematic shift of the clock transition during normal operation is inferred from lock-in data as

$$\Delta\nu_{probe} = \frac{\langle\Delta\nu_{lock-in}\rangle P_0}{P_2 - P_1} = \frac{\langle\Delta\nu_{lock-in}\rangle}{\mathcal{L}}, \quad (5.1)$$

where $\langle\dots\rangle$ is the mean of the data set and, by analogy with simple machines, $\mathcal{L} = (P_2 - P_1)/P_0$ is known as the *lever arm*. Note that inferring $\Delta\nu_{probe}$ from $\langle\Delta\nu_{lock-in}\rangle$ using Equation 5.1 relies on the fact that the probe Stark shift is proportional to P . If this model were not accurate, Equation 5.1 would be an incorrect extrapolation of the data.

The uncertainty in $\Delta\nu_{probe}$, $\delta\nu_{probe}$, is

$$\delta\nu_{probe} = \frac{\sigma_{lock-in}}{\mathcal{L}}, \quad (5.2)$$

where $\sigma_{lock-in}$ is the statistical uncertainty in the mean of the lock-in data, $\langle\Delta\nu_{lock-in}\rangle$. $\sigma_{lock-in}$ is given by the stability of the clock (Section 4.6). As Equation 4.6 shows, the lever arm provides a means of measuring a systematic shift with precision much better than the clock stability. For many systematic measurements, $\mathcal{L} > 10$, allowing for a systematic shift to be measured at a given precision more than 100 times faster than if clock stability were relied upon alone.

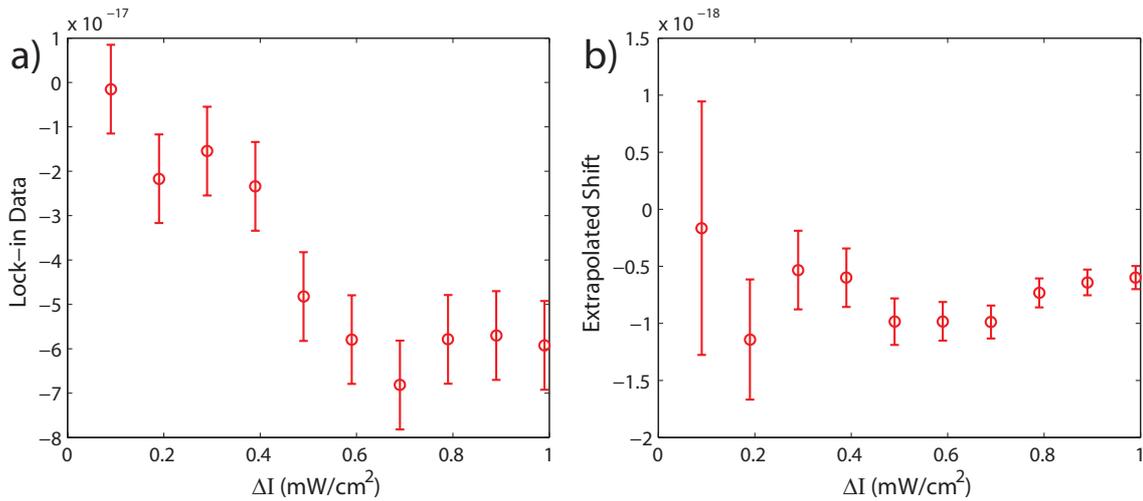


Figure 5.2: a) Simulated lock-in data. Here $\Delta I = I_2 - I_1$. The data scatter is such that χ_{red}^2 takes on the ideal value of 1 (despite looking ugly, as sparse $\chi_{red}^2 = 1$ data often does). b) Simulated shift extrapolated to the clock operation intensity. This is simply the lock-in data divided by the measurement lever arm.

5.1.4 Model Testing

Lock-in data can also be used to test the model of the frequency shift as a function of the lock-in parameter. To demonstrate this, I have simulated a probe Stark shift measurement (Figure 5.2). In this simulation, a clock laser intensity of $I_0 = 10 \mu\text{W}/\text{cm}^2$ is chosen as the clock operation value. The lock-in modulates between $I_1 = I_0$ and I_2 , and the frequency difference between these two conditions is averaged until it reaches 1×10^{-17} statistical uncertainty. This measurement is repeated for values of I_2 ranging from 0.1 to 1 mW/cm^2 .⁴ The mean of the lock-in data $\langle \Delta\nu_{lock-in} \rangle$ for each value of $I_2 - I_1$ is shown in Figure 5.2a. The scatter is simulated as Gaussian white noise with a reduced chi-square statistic $\chi_{red}^2 = 1$. The extrapolated shift, $\Delta\nu_{lock-in}/\mathcal{L}$, where $\mathcal{L} = (I_2 - I_1)/I_0$, is shown in Figure 5.2b. The simulation treats the shift as linear in intensity, and the mean and uncertainty of the mean of the extrapolated data is $(-6.9 \pm 0.5) \times 10^{-19}$. This agrees with the simulation's true value of -6.5×10^{-19} .

In the context of a real measurement, since extrapolating data to the clock operation condition

⁴ As I will show in Section 5.4, in the actual probe Stark shift evaluation, we study this effect as a function of clock laser pulse duration rather than intensity.

relies the shift being linear in intensity, it is tempting to worry if a next-order Stark shift term (quadratic in intensity) might be present in a subtle but significant way. An F -test for an additional term [4] can be used to see if the data requires a quadratic model. This test compares the chi-square statistic for different models to see if an extra term improves chi-square. In the simulated $\Delta\nu_{lock-in}/\mathcal{L}$ data, an F -test for an additional (quadratic) term results in the test statistic $F = 0.37$ (with 8 degrees of freedom). This corresponds to a test probability of 0.56.⁵ The F -test argues for another term if the test probability is lower than a critical value that is conventionally taken as 0.01 to 0.05 [4], so 0.56 is a strong indication that the data does not resolve quadratic behavior. This is expected since $\Delta\nu_{lock-in}$ values are simulated with a linear intensity dependence.

To study the effectiveness of the test, I add a small quadratic term (proportional to $I_2^2 - I_1^2$) to the simulated lock-in data. In this case, the extrapolated shift has the form $\Delta\nu_{lock-in}/\mathcal{L} = \alpha + \beta(I_1 + I_2)$, where α is the linear (in intensity) probe Stark shift at the clock operation condition and β is a positive constant. When $\Delta\nu_{lock-in}/\mathcal{L}$ is studied as a function of intensity, the mean of the extrapolated data will be offset from α by the mean of the $\beta(I_1 + I_2)$ term. Ideally, when β is small enough for the offset to be statistically unimportant, the F -test would conclude that the $\beta(I_1 + I_2)$ term is not needed to model the data. It would also be desirable for the F -test to conclude that $\beta(I_1 + I_2)$ is required when the offset is significant. For most values of β , these desired properties of the F -test hold true.

However, there is a small range of β values for which the mean is significantly affected but the F -test does not conclude that an additional term is needed. For example, β can be made as large as $5.8 \times 10^{-19} \text{ (W/cm}^2\text{)}^{-1}$, yet the F -test will show a model linear in intensity with a comfortable test probability of 0.15. For this β , the mean of $\Delta\nu_{lock-in}/\mathcal{L}$ is -2.3×10^{-19} , but β is still small enough that the simulation's true value of the clock operation probe Stark shift, $\alpha + \beta I_0 = -6.4 \times 10^{-18}$, is statistically consistent with α . Meanwhile for $\beta < 2.3 \times 10^{-19} \text{ (W/cm}^2\text{)}^{-1}$, the offset of the mean of the extrapolated data is within statistical uncertainty, and the F -test reflects that another term is not needed. Therefore, for this simulation, the linear shift model will introduce significant error

⁵ This probability is computed using the F -test probability distribution given in Reference [4].

for the range of $2.3 \times 10^{-19} \text{ (W/cm}^2\text{)}^{-1} < \beta < 5.8 \times 10^{-19} \text{ (W/cm}^2\text{)}^{-1}$, and the F -test will fail in this regime.⁶

Since this blind spot results from neglecting higher-order terms, it may seem as if the clear path forward is to always fit these terms even when the F -test argues that there is no statistical justification for doing so. However, for β in the blind spot range and below, fitting the simulated data to $a+b(I_1+I_2)$ (where a and b are fit parameters) increases the uncertainty in the extrapolated shift $a + bI_0$ to four times above that when linear intensity dependence is assumed—even when $\beta = 0$. Therefore, a wide range of probe Stark measurements that would be well described by a model linear in intensity are greatly affected by treating the quadratic term as unknown. Also, for β in the blind spot range, the value for the b fit coefficient often differs from the simulation’s true value by more than the fit uncertainty.

This analysis based on the probe Stark shift extends to the many systematic shifts that have higher-order terms that are not known precisely, such as shifts from the lattice that are nonlinear in intensity, the quadratic density shift, and the third-order Zeeman shift. Although the range of β values for which the statistical blind spot exists might seem insignificantly small, there has been a great deal of concern in the clock community over higher-order effects that have not been resolved but cause significant measurement bias [18, 121, 111]. If we choose to be concerned about higher-order terms that might be significant but will not show up in an F -test, it is often difficult to know how to proceed. Fits of higher-order terms will not be accurate, and they extract a big toll, destroying the total uncertainty of the clock over a shift that may be negligible. Furthermore, many of these shifts cannot be experimentally resolved even in the systems best equipped to do so.⁷ Being confronted with the possibility of effects that cannot be statistically tested for, fit, or measured, yet there is still concern about their size seems like a contrived scenario, but it

⁶ Note that this regime will vary with the parameters of the measurement set since the offset depends on the variable that is modulated.

⁷ For instance, an M1-E2 lattice light shift was predicted to be very large [111]. Before this prediction, it had not been considered in Sr clocks, and theory suggested that clock uncertainty budgets need to be amended to account for this effect. However, the weak dependence of this effect on trap intensity made it impossible to resolve in most lattice Stark shift measurements. The Sr clock at LNE-SYRTE could achieve unusually large lattice intensity, which should be conducive to observing an M1-E2 shift, yet they measured no shift within their precision [121]. They concluded the shift is much smaller than predicted.

does happen [111, 121]. In the face of this, and because even seemingly sound calculations of important shifts have turned out to be incorrect [94, 105, 111], we only accept higher-order terms in two situations: if our statistical tests demand them, or if another team has reported statistically resolved measurements of these effects that we have no cause to doubt. Although this approach technically leaves open a small possibility of uncorrected measurement bias (should we be unlucky enough for the coefficient of a higher-order effect to fall within the blind spot), to do otherwise would leave our experimental methods and conclusions highly sensitive to experimentally unsubstantiated concerns of higher-order effects.

5.1.5 Point-String Analysis

During a lock-in measurement, the scanning synthesizer data can be affected by drift of the 40 cm clock laser cavity. This is because if the rate of the drift cancellation oscillator (Section 3.7.1) is not optimized, the atomic servo removes the remaining drift by adjusting the scanning synthesizer frequency to compensate for it. We have observed that when this rate is far from optimized, the lock-in data (which is comprised of scanning synthesizer frequency measurements) is contaminated by this drift when it is analyzed as simple differences $\{\nu_2 - \nu_1, \nu_4 - \nu_3, \dots, \nu_M - \nu_{M-1}\}$. To guard against this problem, the residual drift is monitored by real-time plotting of the center frequency computed by one of the atomic servos to look for a linear slope. If a slope becomes apparent, the rate of the drift cancellation oscillator is reset to compensate. Also, the data is processed to remove any residual linear drift. This drift-insensitive method of analyzing lock-in measurement data is *point-string analysis* [109].

For residual linear drift at a rate r_{res} , the probe Stark data will be

$$\begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \nu_4 \\ \vdots \end{pmatrix} = \begin{pmatrix} \nu_{offset} + KP_1 \\ \nu_{offset} + KP_2 + 2T_c r_{res} \\ \nu_{offset} + KP_1 + 2 \times 2T_c r_{res} \\ \nu_{offset} + KP_2 + 3 \times 2T_c r_{res} \\ \vdots \end{pmatrix},$$

where ν_{offset} is an unmeasured technical offset to the scanning synthesizer frequency. If these data are used to create the vector

$$\Delta\nu_{3pt} = \begin{pmatrix} \nu_2 - \frac{1}{2}(\nu_1 + \nu_3) \\ \frac{1}{2}(\nu_2 + \nu_4) - \nu_3 \\ \nu_4 - \frac{1}{2}(\nu_3 + \nu_5) \\ \frac{1}{2}(\nu_4 + \nu_6) - \nu_5 \\ \vdots \\ \frac{1}{2}(\nu_{M-2} + \nu_M) - \nu_{M-1} \end{pmatrix},$$

then the drift will be canceled, making each entry in the vector equal to $K(P_2 - P_1)$ as desired.

Therefore,

$$\Delta\nu_{probe} = \frac{\langle \Delta\nu_{3pt} \rangle}{\mathcal{L}}, \quad (5.3)$$

and each entry in the vector $\Delta\nu_{3pt}$ is called a *three-point string*. This analysis is valid as long as the residual drift is sufficiently linear over the $6T_c$ it takes to measure three center frequencies. Control experiments with a known frequency offset added between the two states of the lock-in shows that three-point strings can effectively reproduce this added offset in the presence of uncanceled clock laser drift.⁸

The three-point string vector $\Delta\nu_{3pt}$ has $M-2$ elements whereas the vector of simple difference data $\{\nu_2 - \nu_1, \nu_4 - \nu_3, \dots, \nu_M - \nu_{M-1}\}$ has $M/2$ elements. Therefore, taking the uncertainty in the

⁸ For the control experiments, the rate of the drift cancellation synthesizer is intentionally set off its optimum.

mean of $\Delta\nu_{3pt}$ to be the uncertainty in the lock-in measurement will underestimate the error bar (since $\Delta\nu_{3pt}$ has more elements than the $M/2$ statistically independent simple difference points). The uncertainty in the mean of $\Delta\nu_{3pt}$ should be inflated to account for this issue.

Suppose a single center frequency measurement has uncertainty of σ . A single point in a measurement of $\Delta\nu_{probe}$, processed using a simple difference, is $(\nu_2 - \nu_1)/\mathcal{L}$. Using standard error propagation and assuming no uncertainty in \mathcal{L} , the uncertainty in this measurement is $\sqrt{2}\sigma/\mathcal{L}$, and the uncertainty in the mean is $2\sigma/\mathcal{L}\sqrt{M}$ (for large M). However, the first point of $\Delta\nu_{probe}$ processed as a three-point string is $[\nu_2 - \frac{1}{2}(\nu_1 + \nu_3)]/\mathcal{L}$. The single-point uncertainty and uncertainty in the mean are $\sqrt{3}\sigma/\sqrt{2}\mathcal{L}$ and $\sqrt{3}\sigma/\sqrt{2M}\mathcal{L}$ (again, for large M). Therefore,

$$\delta\nu_{probe} = \sqrt{\frac{8}{3}} \frac{\sigma_{3pt}}{\mathcal{L}}, \quad (5.4)$$

where σ_{3pt} is the statistical uncertainty in the mean of $\Delta\nu_{3pt}$.

This approach can be generalized to higher orders of drift [109]. Point strings of length n remove polynomial drift of order $n - 2$. If ν_k is the k th point of $\{\nu_1, \nu_2, \dots, \nu_M\}$, the n -point string is

$$\frac{1}{2^{n-2}} \sum_{m=1}^n (-1)^{k+m-1} \frac{(n-1)!}{(n-m)!(m-1)!} \nu_{k+m-1}. \quad (5.5)$$

This expression is valid from $k = 1$ to $k = M - n + 1$, which is the total number of point strings. When this formalism is used to remove drift, the uncertainty in the mean must be inflated by a factor

$$\frac{2^{n-1}}{\sqrt{\sum_{m=1}^n \left[\frac{(n-1)!}{(n-m)!(m-1)!} \right]^2}} \quad (5.6)$$

to extrapolate this uncertainty to that of statistically independent points.

5.2 Residual First-Order Zeeman Shift

5.2.1 Introduction

With an understanding of how many systematic measurements are performed and analyzed, I will now discuss the Sr2 systematic measurements. Before I discuss the systematic shifts that are measured with the lock-in technique, it is useful to first discuss the residual first-order Zeeman systematic shift. This is because a discussion of this shift provides a good context for introducing the background magnetic field servo, which is used for many other systematic measurements.

During normal clock operation, the stretched-states servo (Section 2.5.3) cancels most of the first-order Zeeman shift. However, if the background magnetic field (due to stray fields from lab equipment or Earth's field) drifts appreciably in between measurements of the two stretched-state center frequencies, there can be a residual first-order Zeeman shift.

Let $\{\nu_1, \nu_2, \nu_3, \dots\}$ be a vector of center frequency data taken by the stretched-states servo. Even (odd) subscripts correspond to data taken for the stretched state with a negative (positive) first-order Zeeman shift. Also, let the stretched state first-order Zeeman shift per Gauss be $\zeta = 9\mu_B\delta g/2h = 487.8 \text{ Hz G}^{-1}$ (Equation 2.40). For an external field of $B(t)$, which is comprised of the added bias field and possible a drifting background field, $\nu_1 = \nu_{Sr} - \zeta B(t_1)$ and $\nu_2 = \nu_{Sr} + \zeta B(t_2)$, where $t_2 - t_1 = 2T_c$ and ν_{Sr} is the zero-field clock transition frequency.⁹ When the external field is free of drift, the zero-field clock transition center frequency is inferred as $(\nu_1 + \nu_2)/2$ (Equation 2.5.3).

When the external field does drift, $(\nu_1 + \nu_2)/2 = \nu_{Sr} + \zeta [B(t_2) - B(t_1)]/2$. This stretched state average frequency, normally used to infer ν_{Sr} , now has a residual first-order Zeeman shift $\zeta [B(t_2) - B(t_1)]/2$. The shift of the entire data set is

$$\Delta\nu_{Z1} = \frac{1}{2}\zeta\langle B(t + 2T_c) - B(t) \rangle, \quad (5.7)$$

⁹ Given that the two stretched state resonance centers are not measured one after another—rather, the data used to determine these resonances is interleaved (Figure 2.16)—perhaps $t_2 - t_1$ should not be set equal to $2T_c$. However, to simplify the presentation and arguments in this Section, I consider the case of the resonance centers being measured one after another. All of the conclusions of this Section can still be applied to normal clock operation.

where $\langle \dots \rangle$ is an average over the time values for which the data was taken.

5.2.2 Background Field Servo

The Sr2 system can experience background magnetic field drifts at the level of 10 mG over a couple of hours. This is due to the magnet from an ion pump near the main chamber (Figure 3.3), which produces a field that changes as the laboratory temperature fluctuates.¹⁰ To remove this drift, we implement a background magnetic field servo [13].

Measurements for this servo take place with the spin polarization beam extinguished. The polarization coil current and V_2 (Figure 3.17) are also shut off.¹¹ In this state the clock transition is unpolarized and the only remaining field is the residual background field. For a typical background of about 10 mG, the unpolarized clock transition lineshape is broader and has worse contrast than in the zero-field case (Figure 2.14a).

The background field servo adds a field to cancel this background. The servo accomplishes this by adjusting the currents in the xyz coils to maximize the unpolarized clock transition contrast, which is biggest for zero background field. One iteration of this servo requires three measurements per xyz coil pair (nine total). The first three of these are taken holding the y and z coil currents fixed while the x coil current is switched between the values I_x , $I_x + \Delta I$, and $I_x - \Delta I$ (one for each measurement). Here I_x is the initial coil current and ΔI is a small current step.

For each of these three currents, the clock transition interrogates the center of the unpolarized line, and the excited state fraction is recorded. For a well-canceled background field, the excited state fraction would be the highest when the x coil current is I_x , and it would diminish equally for $I_x + \Delta I$ and $I_x - \Delta I$. If the background field drifted off its optimum, the currents $I_x + \Delta I$ and $I_x - \Delta I$ would yield different excited state fractions. Treating the on-resonance excited state fraction as a parabolic function of external field, a digital PI filter computes a new value for I_x that balances the excited state fractions measured for currents $I_x + \Delta I$ and $I_x - \Delta I$ (where ΔI , which

¹⁰ Temperature stabilization of the space around the vacuum chamber, which will be discussed in Section 5.8.8, helps with this problem.

¹¹ V_1 from Figure 3.17 is shut off after 689 nm MOT cooling, so it is off at this point as well.

is determined empirically to optimize the servo, remains fixed throughout servo operation). The x coil is then held while this same procedure is performed on the y and then z coils.

These servo measurements cannot occur while the stretched-states lock is running; therefore, normal clock operation involves switching between 81 cycles of the stretched-states lock and nine measurements for the background field servo.¹² The background field servo adjusts the coil currents using a DAC card. Each of the three outputs of the card are summed into a current stabilizer setpoint voltage for one of the xyz coils (Figure 3.17). After a few iterations of this servo, the background field remains well canceled while this servo is engaged.

In a typical day, before the background field servo is initiated for the first time, we find xyz coil servo voltages that roughly cancel the background field,¹³ allowing for the initial state of the servo to be nearly stabilized. However, to test the servo, we intentionally offset the xyz coils from their optima and observe the servo effectively correcting the coil currents to cancel the background field (Figure 5.3).

5.2.3 Shift Analysis

With the background field well canceled, the residual first-order Zeeman shift $\Delta\nu_{Z1}$ must be measured and assigned an uncertainty. This shift can be inferred with stretched-states servo data. As explained in Section 5.2.1, for a time varying magnetic field, the first two stretched-states servo points are $\nu_1 = \nu_{Sr} - \zeta B(t_1)$ and $\nu_2 = \nu_{Sr} + \zeta B(t_2)$. To analyze the shift, first I define a stretched state frequency difference vector as $\{s_1, s_2, s_3, \dots\} = \{\nu_2 - \nu_1, \nu_4 - \nu_3, \nu_6 - \nu_5, \dots\}$. Note that $s_2 - s_1 = \zeta [B(t_3) - B(t_1) + B(t_4) - B(t_2)]$.

Forming a new vector $S = \{s_2 - s_1, s_4 - s_3, s_6 - s_5, \dots\}$, the mean of S is

¹² The background field servo operates with the scanning synthesizer at the frequency corresponding to the zero-field clock transition resonance. This frequency is known from the stretched-states servo. Just before pausing the stretched-states servo to allow for the background field servo measurements, the position of the zero-field resonance center is computed by averaging the most recently acquired stretched state center frequencies.

¹³ This is done by setting the free running clock laser on the zero-field resonance and adjusting the coil currents by hand (without feedback) to coarsely maximize the excited state fraction.

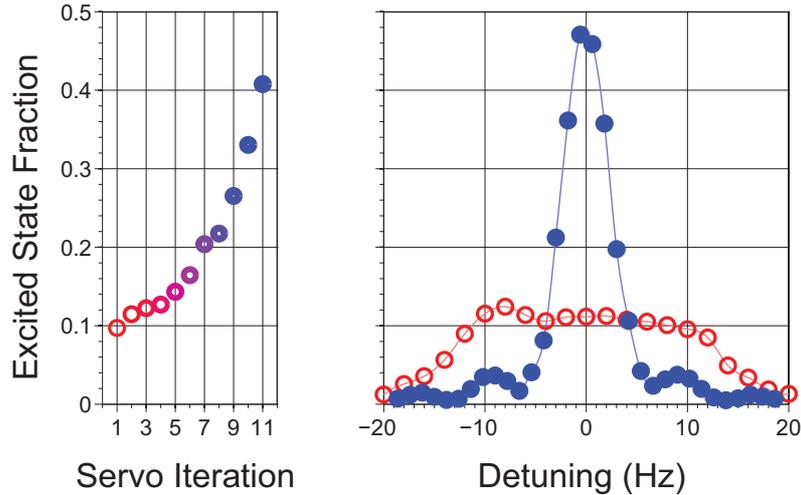


Figure 5.3: The zero-field-resonance excited state fraction as a function of servo iteration and detuning. To test the servo, the xyz coil currents are intentionally offset from their optima. This results in the broadened line of the right plot (red open circles). The servo is then turned on, and the excited state fraction gets larger as the servo runs (left plot). After 11 iterations, a scan of the zero-field clock transition shows the familiar Rabi lineshape (right plot, blue circles), indicating that the 10 m_F states are degenerate and the background field is well canceled.

$$\langle S \rangle = 2\zeta \langle B(t + 4T_c) - B(t) \rangle \quad (5.8)$$

$$\geq 4\zeta \langle B(t + 2T_c) - B(t) \rangle. \quad (5.9)$$

where this second expression is equal to $\langle S \rangle$ for drift linear in time and less than $\langle S \rangle$ for higher-order drift. Using Equation 5.7,

$$\Delta\nu_{Z1} \leq \frac{1}{8} \langle S \rangle. \quad (5.10)$$

This inequality puts an upper bound on $\Delta\nu_{Z1}$. Since $\Delta\nu_{Z1}$ is usually consistent with zero, an upper bound is sufficient. Applying this analysis to the full record of stretched-states servo data, $\Delta\nu_{Z1}$ is less than $(-1.6 \pm 2.0) \times 10^{-19}$ [88]. In practice, this value is quoted for $\Delta\nu_{Z1}$.

5.3 Second-Order Zeeman Shift

5.3.1 Introduction

As explained in Section 2.5.3, using a bias field for clock operation introduces a second-order Zeeman shift that cannot be removed with the stretched-states servo. This effect is predicted to shift the stretched states equally by $(-0.233 \text{ Hz G}^{-2})B^2$, where B is the bias field magnitude [16]. To measure this shift, we use a lock-in measurement with B as the modulation parameter. Like in Equations 5.1 and 5.2, the shift of the clock transition for normal clock operation is $\langle \Delta\nu_{lock-in} \rangle / \mathcal{L}$. The lever arm for the measurement is $\mathcal{L} = (B_2^2 - B_1^2) / B_0^2$, where B_1 and B_2 are the lock-in states and B_0 is the field during normal clock operation.

It is convenient to measure this shift as a function of the frequency difference between the stretched states (which is proportional to the bias field magnitude) for normal clock operation [121]. As explained in Section 2.5.1, the first-order Zeeman shift per unit field is $-108.4m_F \text{ Hz G}^{-1}$, which implies a stretched state frequency difference of 975.6 Hz G^{-1} . Using this number, the predicted value of the second-order Zeeman shift in units of fractional frequency per kHz of stretched state difference is $-5.71 \times 10^{-18} \text{ kHz}^{-2}$. Working in these units, we choose a clock operation bias field corresponding to a stretched state difference of $S_0 = 300 \text{ Hz}$. Therefore, the lever arm for the measurement can also be written as $\mathcal{L} = (S_2^2 - S_1^2) / S_0^2$, where S_1 and S_2 are the stretched state differences corresponding to B_1 and B_2 .

To measure the second-order Zeeman shift due to modulating the field, the lock-in data must be free of the first-order Zeeman effect. This is accomplished using four atomic servos (rather than two servos as discussed Section 5.1.2). For both lock-in states of the bias field, there is a stretched-states lock (comprised of two atomic servos) generating a stretched state average frequency (Section 2.5.3) that is not affected by the first-order Zeeman shift. The difference between the two stretched state average frequencies after one cycle through the four atomic servos is one lock-in data point.

Also, since a drifting background field would introduce unwanted variation in the second-order Zeeman shift measurement, the background field servo is enabled during this evaluation.

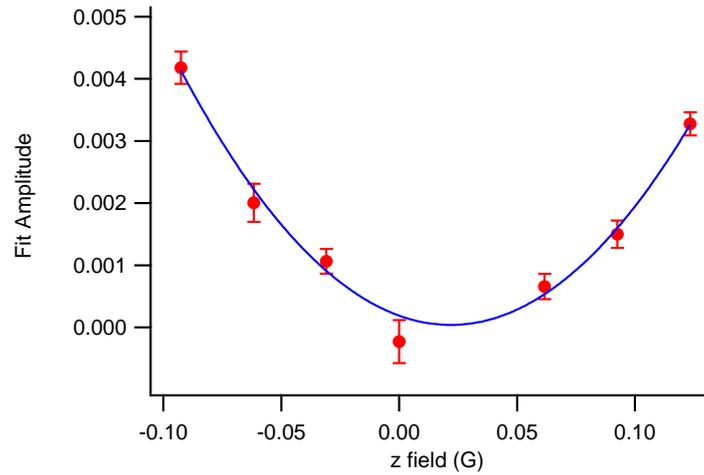


Figure 5.4: The data used to align the z coil field. The σ transition is scanned for different values of the z coil field. Measured lineshapes are fit, and the fit amplitudes are plotted here. These data are then fit to a parabola, and the parabola's offset is the desired field value.

5.3.2 The Lattice Tensor Shift

This measurement could be negatively affected by the lattice tensor shift if care is not taken.

As shown in Equation 2.19, the lattice tensor shift $\Delta\nu_t$ is

$$\Delta\nu_t = 36 \kappa_t (3|\hat{\epsilon} \cdot \hat{B}|^2 - 1) U_{lat}(r, z). \quad (5.11)$$

Due to the dependence on \hat{B} , if the bias field is modulated in a manner that changes the field direction, the measurement will be contaminated with a lattice tensor shift that differs between the two lock-in states. Since the overall bias field is comprised of fields from the polarization and xyz coils, lattice tensor shift contamination can be prevented by adjusting all four coil currents so that the bias field is well aligned with the lattice polarization $\hat{\epsilon}$ for both field states of the lock-in measurement. As explained in Section 3.5.1, \hat{B} and $\hat{\epsilon}$ are aligned by first aligning the clock laser polarization with $\hat{\epsilon}$ and then studying m_F -changing σ transitions (Figure 5.4). σ transitions become possible when \hat{B} is not aligned with the clock laser polarization.

For each state of the lock-in measurement, the polarization coil provides the bulk of the field

magnitude, and the xyz coils merely help to point the bias field in the correct direction.¹⁴ Two polarization coil currents are selected as lock-in parameter states, and then σ transition scans are used to adjust the xyz coils such that the bias field is aligned with the clock laser polarization. This procedure is repeated a few times to ensure that the bias field is well aligned.

5.3.3 First and Second Evaluations

The first evaluation of the second-order Zeeman shift is performed with lock-in parameters S_1 and S_2 ranging from 300 Hz to 1.2 kHz [13]. This evaluation measures a second-order Zeeman shift of $(-5.78 \pm 0.05) \times 10^{-16}$ kHz⁻². The second evaluation uses $S_1 = 291$ Hz and $S_2 = 1.66$ kHz, measuring a second-order Zeeman shift of $(-5.82 \pm 0.07) \times 10^{-16}$ kHz⁻². Extrapolating to the operating stretched state difference of $S_0 = 300$ Hz, the Sr2 second-order Zeeman shift implied by the second evaluation is $(-5.23 \pm 0.07) \times 10^{-17}$ (Figure 5.5a).

The shift per stretched state frequency difference squared is an atomic property, so evaluations of this by different research teams can be treated as independent measurements of the same quantity (Figure 5.5b). Performing a weighted average of several reported measurements [74, 13, 88, 121, 35], the shift and its root-reduced-chi-square-inflated uncertainty is $(-5.75 \pm 0.03) \times 10^{-16}$ kHz⁻². Extrapolating to the operating S_0 , the final Sr2 second-order Zeeman shift is -5.17×10^{-17} with an uncertainty of 2.5×10^{-19} [88].

5.4 Probe Ac Stark Shift

5.4.1 Introduction

As explained in Section 5.1.3, the probe Stark shift $\Delta\nu_{probe}$ is proportional to P , the clock laser power. This shift is measured using the lock-in technique. For state 1 (2) of the lock-in, a stretched-states lock uses clock laser π -pulses with a long (short) interrogation time, which means that a smaller (larger) Rabi frequency, and therefore less (more) laser power is needed to make a

¹⁴ Recall that the polarization coils point at an angle 16° with respect to the horizontal and the clock laser polarization is nearly horizontal

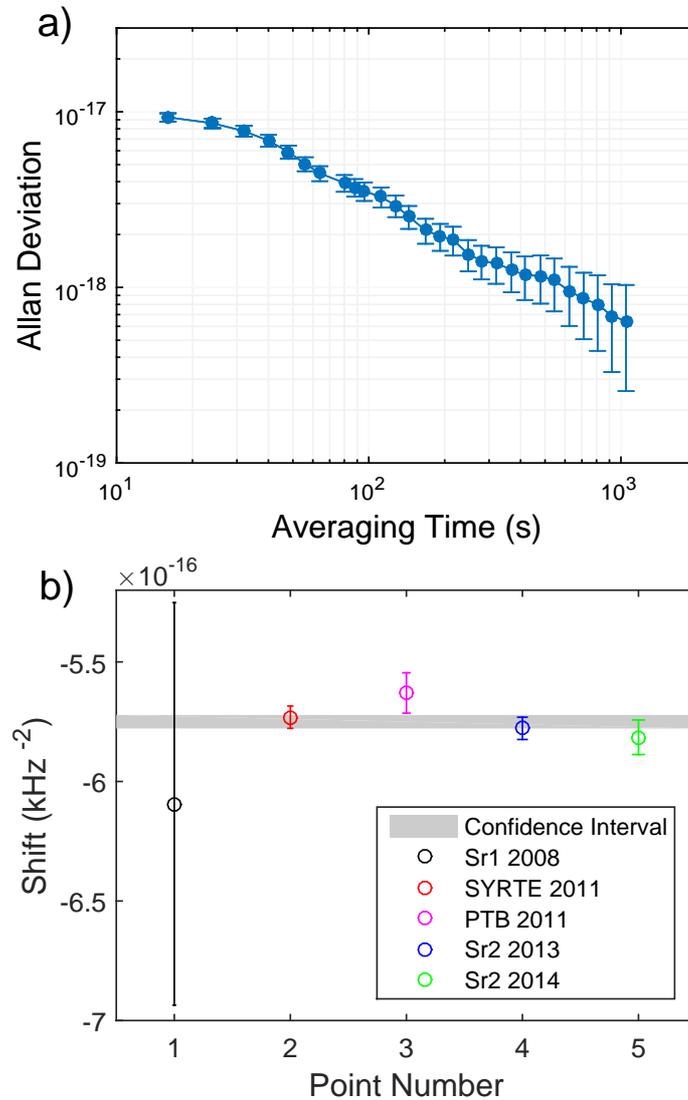


Figure 5.5: a) Overlapping Allan deviation of the second-order Zeeman shift extrapolated to the operating value of S_0 . This is the second Sr2 evaluation of this shift. b) Measurements of the fractional second-order Zeeman shift per stretched state frequency difference squared. These values have been reported by four strontium clocks: Sr1 [74], Sr2 [13, 88], the Sr clock at LNE-SYRTE (Paris) [121], and the PTB (Braunschweig, Germany) Sr clock [35]. The uncertainty of the weighted average is depicted as the gray band.

π -pulse. The difference between the two stretched state average frequencies is the lock-in data.

Two stretched-states servos are used in the lock-in measurement (rather than only two atomic servos) because the background field servo is run during this evaluation (Section 5.2.2). The background field needs to be canceled because the probe Stark shift has a tensor component that will vary with a drifting magnetic field. For the lock-in parameter values P_1 and P_2 , and an operating laser power P_0 , the lever arm for this measurement is $\mathcal{L} = (P_2 - P_1)/P_0$. It is easier to express \mathcal{L} in terms of the pulse duration used for each of these laser powers, t_0 , t_1 , and t_2 . Since π -pulses are used, the Rabi frequency for the measurement $\Omega = \pi/t$. Also, $\Omega \propto \mathcal{E}_{c0}$, where \mathcal{E}_{c0} is the clock laser electric field amplitude (Section 2.4.1). Since $\mathcal{E}_{c0} \propto \sqrt{P}$, $\Delta\nu_{probe} \propto 1/t^2$. Therefore, the lever arm is $\mathcal{L} = \left(\frac{t_0}{t_2}\right)^2 - \left(\frac{t_0}{t_1}\right)^2$.

5.4.2 First and Second Evaluations

The first evaluation of the probe Stark shift uses $t_2 = 50$ ms and $t_1 = 200$ ms [13]. In this case, the pulse duration for clock operation is $t_0 = 160$ ms, resulting in a lever arm of $\mathcal{L} = 9.6$. Modulating the clock laser interrogation time by this amount requires that the laser power be modulated by a factor of $(t_1/t_2)^2 = 16$ to maintain π -pulses. This is accomplished using the clock laser intensity servo (Section 3.7.1). Dividing the mean and standard deviation of the mean of the lock-in data by \mathcal{L} , the final probe Stark shift for the first evaluation is $(0.5 \pm 1.3) \times 10^{-18}$ [13].

This treatment of the measurement lever arm assumes that both pulses perfectly achieve a pulse area of π . The uncertainty in this assumption is assigned by studying the range of clock laser powers for which the measured lineshape still resembles a π -pulse. For power slightly larger or smaller than that corresponding to a π -pulse, the peak excited state fraction will be diminished. Studying the change in clock laser power required to diminish the excited state fraction, treating this range as uncertainty, and propagating it into the final processed signal, we calculate a negligible correction for this effect.

For the second evaluation, we would like the probe Stark uncertainty to be characterized at the mid- 10^{-19} level or better. This is because our measurement goal for BBR dynamic shift

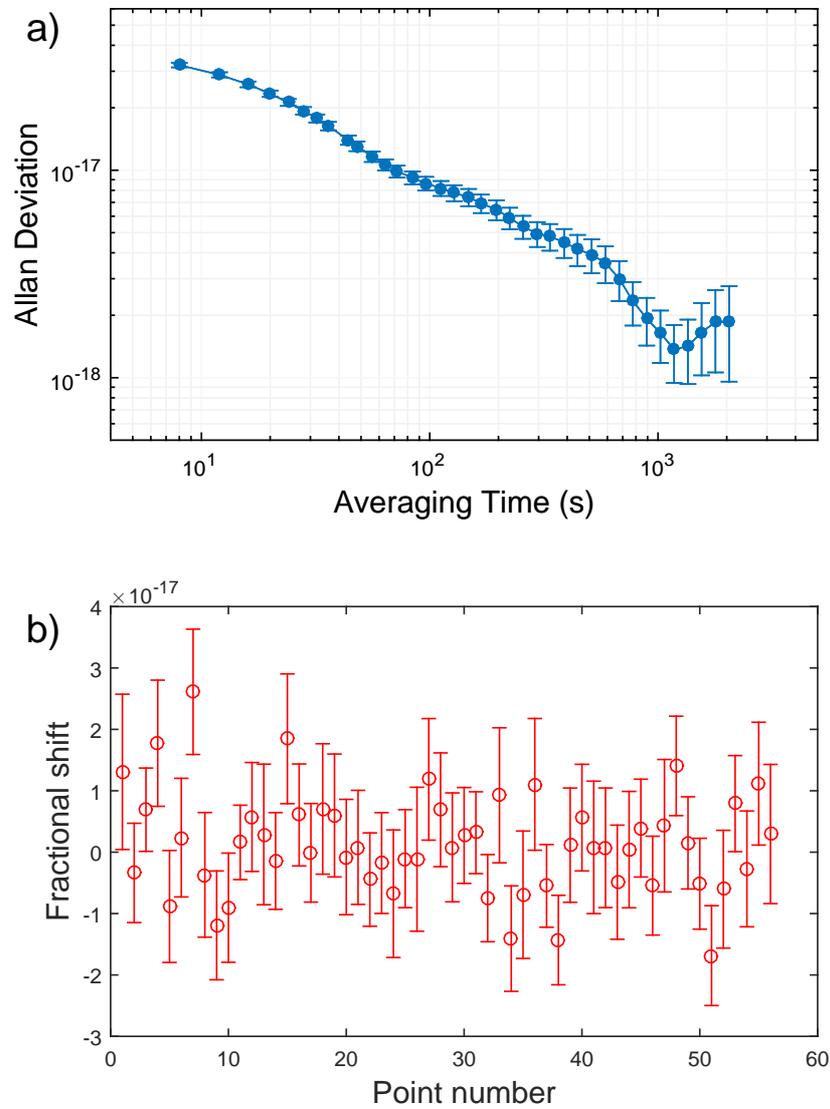


Figure 5.6: a) Overlapping Allan deviation of the first probe Stark shift evaluation. This is the measurement of the shift extrapolated to normal clock operation, meaning that it is the lock-in data divided by \mathcal{L} . b) Binned lock-in data. Each marker represents a mean of 25 points, and the error bars are the standard deviation of this mean. The reduced chi-square χ_{red}^2 is 0.92.

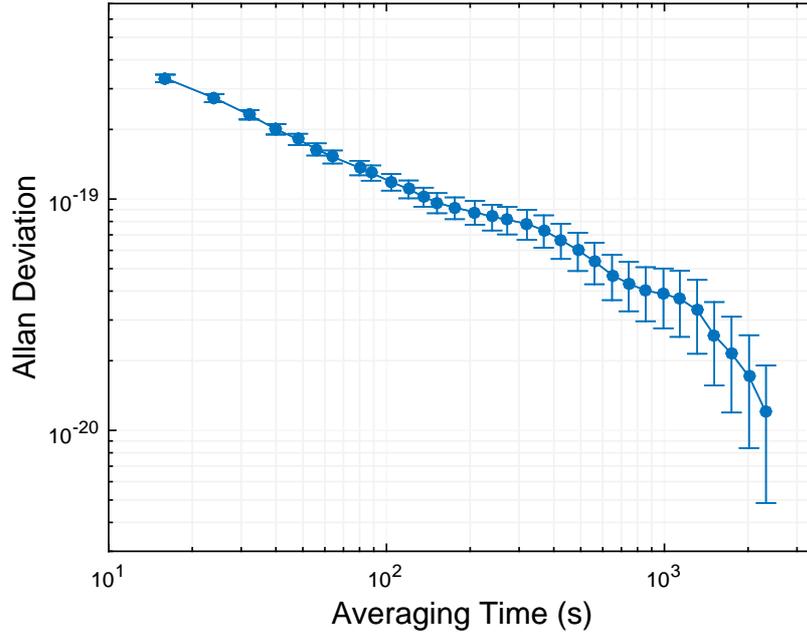


Figure 5.7: Overlapping Allan deviation of the second probe Stark shift evaluation. Here the measured shift extrapolated to clock operation conditions reaches the low- 10^{-20} -level in less than an hour of measurement time.

coefficient (discussed in Section 5.9) would yield low 10^{-18} level uncertainty, and we want to reduce as many other uncertainties as possible to negligible levels. To reduce the probe Stark uncertainty, it is advantageous to improve the measurement lever arm. This is accomplished with a neutral density filter attached to a stepper motor to move this filter in and out of the clock laser. This allows the clock laser intensity to be modulated by a factor of 83.5.

t_1 is fixed at 180 ms, and the clock laser intensity servo is adjusted for a π -pulse when the motorized neutral density filter is removed from the beam. For the same intensity servo value, t_2 is adjusted for a π -pulse when the filter is in the beam, resulting in $t_2 = 19.7$ ms. If $t_0 = 160$ ms as in the first evaluation, we would have already achieved a lever arm of 65.2. However, to bolster the Sr2 stability for our second evaluation [88], clock operation with 1 s pulses has now been demonstrated (Section 4.6). With $t_0 = 1$ s as the new pulse duration for normal clock operation, \mathcal{L} reaches a staggering value of 2546. With this amazing lever arm, the second probe Stark evaluation measures a shift of $(-3.2 \pm 1.7) \times 10^{-20}$.

5.5 Density Shift

5.5.1 p-Wave Density Shift

Although most major systematic shifts of the Sr2 clock transition come from electromagnetic perturbations, there is one important exception: the *density shift* [66]. This is a frequency shift due to interactions between particles in multiply occupied lattice sites. Since the ^{87}Sr lattice clock is based on spin-polarized ultracold fermions, the *s*-wave part of these interactions is suppressed unless there is a significant amount of excitation inhomogeneity [19, 99].

A substantial density shift was observed in the Sr1 clock, and early studies of this effect attributed the shift to inhomogeneous *s*-wave interactions [19]. Later, a study by the NIST lattice clock team identified *p*-wave interactions as the dominant partial wave interaction in ytterbium [64]. Further measurements in Sr1 revealed a *p*-wave contribution to the strontium density shift [7], meaning that *s*- and *p*-wave interactions most likely played a role in density shift measurements [109, 6]. More recent work in Sr1 showed improved excitation homogeneity, which suppressed *s*-wave interactions and left *p*-wave effects as the dominant density shift mechanism [78].

The *p*-wave density shift in the mean field approximation is given as [71]

$$\Delta\nu_p = \frac{\pi}{h} n(\vec{r}) E_T [v_{ee}\rho_{ee} - v_{gg}\rho_{gg} + (\rho_{gg} - \rho_{ee})v_{eg}], \quad (5.12)$$

where $n(\vec{r})$ is the particle density and E_T is the mean collision energy, which is proportional to $k_B T$. ρ_{gg} and ρ_{ee} are the ground and excited state population fractions, and v_{gg} , v_{eg} , and v_{ee} are the *p*-wave scattering volumes.

The density in a given lattice site is [110]

$$n(\vec{r}) = \frac{N_{site}}{\pi^{3/2} L_x L_y L_z} \exp\left(-\frac{x^2}{L_x^2} - \frac{y^2}{L_y^2} - \frac{z^2}{L_z^2}\right), \quad (5.13)$$

where N_{site} is the number of particles in the site. The lengths $L_{x,y,z}$ are given by

$$L_i = \sqrt{\frac{\hbar}{m\omega_i}} \sqrt{2\langle n_i \rangle + 1}. \quad (5.14)$$

Here $\omega_{x,y,z}$ is the trap frequency in the $\{x, y, z\}$ direction (Equations 2.7 and 2.8) and the occupancy factor $\langle n_i \rangle$ is

$$\langle n_i \rangle = \frac{1}{\exp\left(\frac{\hbar\omega_i}{k_B T}\right) - 1}. \quad (5.15)$$

In the limit $k_B T \gg \hbar\omega_i$, $L_i \rightarrow \sqrt{2k_B T/m\omega_i^2}$, which is the length scale of a thermal gas in a harmonic trap. For $k_B T \ll \hbar\omega_i$, L_i becomes the harmonic oscillator length $\sqrt{\hbar/m\omega_i}$. Since the radial trap frequency $\nu_r = 120$ Hz, L_x and L_y are both equal to $\sqrt{2k_B T/m\omega_r^2}$ to a very good approximation (for typical lattice temperatures of a few μ K).

It is interesting to look at how the density shift scales with the experimental parameters N , T , and U_0 . If one uses *sideband cooling* along the lattice axis, whereby the atoms are cooled to the motional ground state in this direction by driving a blue sideband transition, then $L_z \rightarrow \sqrt{\hbar/m\omega_z}$, and

$$\Delta\nu_{density} \propto NU_0^{5/4}. \quad (5.16)$$

This expression uses the fact that ω_r and ω_z are proportional to $\sqrt{U_0}$ (Equations 2.7 and 2.8). The measured density shift $\Delta\nu_{density}$ results from spatially averaging $\Delta\nu_p$. In this case, the p -wave density shift is temperature independent.

For $k_B T \gg \hbar\omega_z$,

$$\Delta\nu_{density} \propto \frac{NU_0^{3/2}}{\sqrt{T}}. \quad (5.17)$$

In Sr2, $k_B T$ is comparable to $\hbar\omega_z$,¹⁵ , and treating the density as thermal is an approximation that is good to 3%. Also, for consistency with Equations 2.7 and 2.8, this analysis has ignored

¹⁵ The biggest value of ν_z usually observed in Sr2 is 120 kHz. At this depth, the sample temperature is usually 7 μ K.

the influence of gravity. Gravity can be included by replacing U_0 in these scaling laws with U_{atom} (Section 2.4.5).

5.5.2 Density Shift in a Cavity Lattice

When operating a lattice clock, there is a tension between the density shift and QPN. A small atom number is good for reducing the density shift, yet a large N is preferred for minimizing QPN. Therefore, it is best to operate with N just large enough to make QPN a few times lower than the Dick Effect. For the 160 ms clock pulses used in the first Sr2 evaluation, the Dick Effect for typical experimental parameters is $3.6 \times 10^{-16}/\sqrt{\tau}$. Therefore, $N = 2000$ is a good choice since it results in QPN of $1.2 \times 10^{-16}/\sqrt{\tau}$. For the 1 s pulses used in the second evaluation, the Dick Effect is predicted to be $1.8 \times 10^{-16}/\sqrt{\tau}$, and QPN for $N = 2000$ is $2.1 \times 10^{-17}/\sqrt{\tau}$. Although the clock operation atom number could stand to be reduced by an order of magnitude for 1 s pulses, the Sr2 system is optimized for $N = 2000$,¹⁶ so it is convenient to use this value is used for both 160 ms and 1 s pulses. As I will show in Section 5.5.4, the density shift is manageable at $N = 2000$.

Although the density shift is large in Sr1 [74], it is significantly smaller in Sr2 due to the large-volume cavity lattice (Section 2.3.3). For both systems, the maximum lattice intensity is limited by the power available from the 813 nm light source. However, because of the intensity enhancement from the lattice cavity (Equation 2.14), the Sr2 system can achieve a good lattice depth for a trap beam waist of 160 μm (compared to the Sr1 waist of 40 μm). At this waist, Sr2 achieves a maximum axial trap frequency of $\nu_z = 120$ kHz (corresponding to $U_{atom} = 300 E_{rec}$ or 50 μK), whereas Sr1 measures $\nu_z = 80$ kHz (with $U_{atom} = 130 E_{rec}$ or 22 μK).

The two advantages of the cavity lattice that improve clock stability (Section 4.6) also improve the density shift. First, the large volume means that Sr2 can operate with a smaller density at $N = 2000$ than Sr1. For the same atom number, the Sr2 density is estimated to be 30 times less than that of Sr1 using a Poissonian lattice filling model [110].¹⁷ Second, the good spatial overlap

¹⁶ In particular, for the measured Sr2 detection noise of $\delta N = 11.4$ atoms, stability degradation due to detection noise is worse than it is for QPN when $N < 2 \delta N^2 = 260$ (Section A.4).

¹⁷ This model assumes that the lattice overlaps a 1/e 689 nm MOT cloud width of 160 μm , compared to 30 μm

with the 689 nm MOT means that Sr2 can trap an excess of atoms, which is useful for making density shift measurements with a large lever arm.

5.5.3 Large Lever Arm Measurement

Due to the ability to trap many more atoms than are needed, N is a good choice for a lock-in modulation parameter. Although Equation 5.17 suggests that U_{atom} might also be a good choice, it would result in a worse lever arm than if N were used. With a maximum U_{atom} of $300 E_{rec}$ and a minimum reliable U_{atom} of $71 E_{rec}$, the best achievable lever arm is 8.4. Meanwhile, the lattice can trap $N = 1 \times 10^5$. Extrapolating to an operating atom number of $N = 2000$, a lever arm of 50 is possible.

However, at some level, the density shift for Rabi spectroscopy becomes nonlinear in N [98]. To check the linearity of the density shift, this effect is measured as a function of ΔN , which is the difference in lock-in parameters [89]. For this measurement, we use the lock-in scheme described in Section 5.1.2, where two atomic servos (one for each value of the lock-in parameter) are locked to the same stretched state. Because the atom number fluctuates, the total PMT counts (Section 3.7.2), which are proportional to the atom number, are recorded for every excited state fraction measurement. Using the set of measurements of N , $\{N_1, N_2, N_3, \dots\}$, where N_k is controlled to be large when k is even and small when k is odd, the lever arm for this measurement is treated as a vector $\mathcal{L} = \{(N_2 - N_1)/N_0, (N_4 - N_3)/N_0, \dots\}$ (with $N_0 = 2000$). Density shift data is processed by dividing each lock-in point by the instantaneous lever arm, $\{(\nu_2 - \nu_1)N_0/(N_2 - N_1), (\nu_4 - \nu_3)N_0/(N_4 - N_3), \dots\}$. Here ν_k is the center frequency taken when the atom number is N_k .¹⁸

We study the density shift as a function of $\Delta N = N_{high} - N_{low}$, which is the difference between the large and small atom numbers used for each lock-in point. A plot of the difference between $\nu_{high} - \nu_{low}$ (the frequencies corresponding to N_{high} and N_{low}) divided by ΔN should appear constant as a function of ΔN if the density shift is linear, and this is the observed behavior

for the manner in which the Sr1 lattice is aligned.

¹⁸ Control of N is based on changing the 461 nm MOT laser intensity using the MOT light AOM (Section 3.2.3). Fluctuations in N of 15% about its average value are typical, and N can also drift downward when one of the Sr2 cooling laser systems needs to be reoptimized.

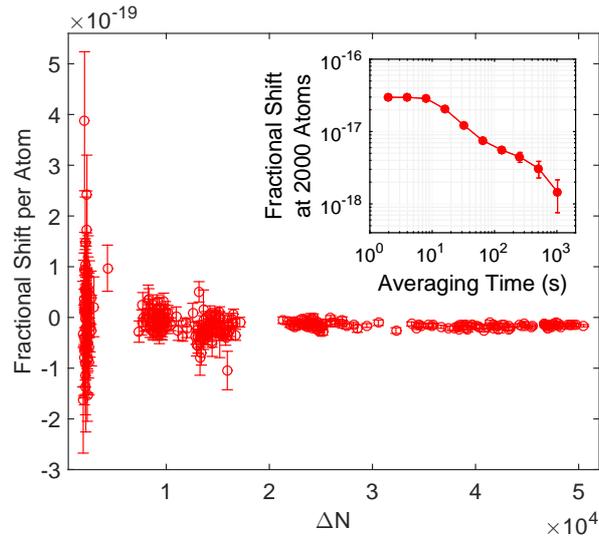


Figure 5.8: The measured Sr2 density shift as a function of ΔN (taken from Reference [89]). Each point on this plot represents an average over a bin of 30 measurements of $(\nu_{high} - \nu_{low}) / (N_{high} - N_{low})$. Our statistics show that these data extrapolated to $N = 2000$ do not support nonlinear terms within an uncertainty of 8.2×10^{-19} . This measurement was performed with $U_{atom} = 210 E_{rec}$, so if the data is extrapolated to the clock operation U_{atom} of $U_{clock} = 71 E_{rec}$, a nonlinear shift is not supported within a 1.6×10^{-19} uncertainty. Inset: A single 2000 s long density shift measurement with $\Delta N = 4.1 \times 10^4$ (average value). The shift per atom was measured and then scaled up to the clock operation atom number of $N_0 = 2000$.

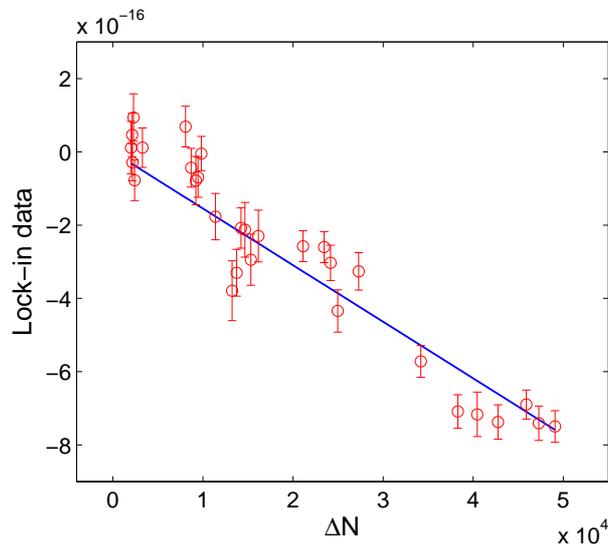


Figure 5.9: The linear trend is more apparent for the lock-in data. Here the data is binned more coarsely, using bin sizes of 600, for clarity. The blue line is a linear fit that assumes the shift goes to zero for $\Delta N = 0$.

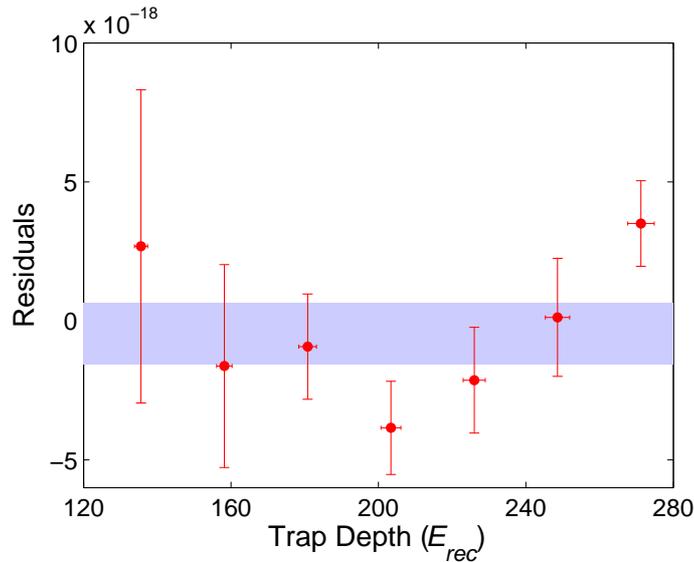


Figure 5.10: The density shift at $N = 2000$ studied as a function of U_{atom} . The residuals of a fit to the expected $U_{atom}^{3/2}$ power law (Equation 5.17) are plotted here to make the statistical behavior more apparent. The light blue band is the $\pm 1\sigma$ confidence interval. The fit coefficient is $(-3.9 \pm 0.3) \times 10^{-21} E_{rec}^{-3/2}$. The confidence interval has been inflated by the square root of $\chi_{red}^2 = 2.0$.

(Figure 5.8). An F -test¹⁹ shows that measurement does not require a quadratic model, arguing for a linear model within the measurement uncertainty of 1.6×10^{-19} . For large atom number modulation, the shift averages down to 1×10^{-18} uncertainty in about 20 minutes (Figure 5.8) [89].

The data in Figure 5.8 is overscattered, resulting in a reduced chi-square statistic of $\chi_{red}^2 = 1.7$. Although χ_{red}^2 is typically very close to unity for most of the Sr2 systematic evaluations, it is consistently around 2 for density shift measurements. An extra noise source could be present that has not been identified, such as variations in optical trapping conditions. To ensure that error bars are sufficiently inflated to account for possible unidentified sources of noise, all final uncertainties, including the 1.6×10^{-19} uncertainty from the data in Figure 5.8, are multiplied by $\sqrt{\chi_{red}^2}$ if $\chi_{red}^2 > 1$ [4].

We also study the density shift as a function of U_{atom} . For this measurement, we prefer to use the Ti:sapphire lattice laser instead of the TA system (Section 3.4), which was employed in the

¹⁹ With 305 degrees of freedom, the F -test results in a test statistic $F = 3.36$, corresponding to the probability 0.068.

probe Stark and second-order Zeeman shift evaluations. This is because any study of clock behavior determined by the lattice should use the light source that is part of normal clock operation. U_{atom} is determined using sideband spectroscopy to measure the trap frequency, which implies the potential depth accord to Equation 2.39. The data fits to a $U_{atom}^{3/2}$ model (Figure 5.10) with $\chi_{red}^2 = 2.0$.

5.5.4 Systematic Evaluation

Rather than use the results of the study of density shift linearity to assign a clock operation value to this effect, we opted to re-evaluate the density shift (Figure 5.11). This is because shortly after the linearity studies, changes to the Sr2 apparatus required the system to be re-optimized, potentially changing experimental parameters that could affect the density shift. For the systematic measurement, $U_{atom} = 119 E_{rec}$, and the lock-in modulates between 2400 and 12000 atoms (average values), leading to an average lever arm of 4.8. Extrapolating this to clock operation conditions quoted in the first evaluation of $N = 2000$ and $U_{atom} = 87 E_{rec}$, the final shift is $(-4.7 \pm 0.6) \times 10^{-18}$ [13].

This measurement was taken at values of U_{atom} lower than the best we could achieve (meaning that the maximum atom number, and therefore the lever arm, was also lower). Despite that the lever arm is low for this measurement, it ultimately proved to be a good determination of the Sr2 density shift due to the length of the data record. This measurement is also good enough for the second evaluation, which quoted new clock operation conditions of $N = 2000$ and $U_{atom} = 71 E_{rec}$, implying a final shift of $(-3.5 \pm 0.4) \times 10^{-18}$ [88].

5.6 Lattice Stark Shift

5.6.1 Lever Arm and Nonlinear Light Shifts

Using Equation 2.19, the lattice ac Stark shift for clock operation is

$$\Delta\nu_{ac} = - \left[\kappa_s(\lambda) + 36 \kappa_t(\lambda)(3|\hat{\epsilon} \cdot \hat{B}|^2 - 1) \right] U_{atom}, \quad (5.18)$$

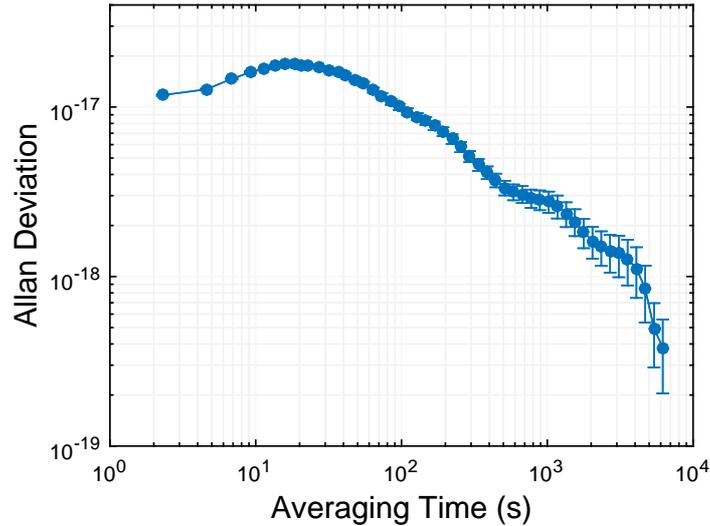


Figure 5.11: The density shift extrapolated to operating conditions of $N = 2000$ and $U_{atom} = 71 E_{rec}$.

where the vector term has been ignored (Section 2.6). As explained in Section 2.6, the magic wavelength for the optical lattice clock is the wavelength λ at which the scalar and tensor terms of $\Delta\nu_{ac}$ cancel one another for the $\hat{\epsilon} \cdot \hat{B} = 1$ case. Sr1 measured this wavelength, but at the time measurement precision was not high enough to resolve the tensor shift [74]. The Sr team at LNE-SYRTE achieved the precision needed to observe tensor shift effects, but they did not determine the scalar-tensor cancellation wavelength, reporting instead the wavelength at which $\kappa_s(\lambda) = 0$ [121]. Recently the Sr group at PTB measured this cancellation wavelength, obtaining a value of $c/368.554465(3)$ THz [33].

The two Sr2 lattice Stark shift evaluations are both performed with our Ti:saph lattice laser. The laser wavelength is stabilized to a frequency comb that is referenced to the NIST hydrogen maser array (Section 3.4.3), and the cavity lattice intensity is modulated for lock-in detection. Modulation is performed with the lattice cavity transmission intensity stabilizer (Section 3.4.2), which actuates on the drive amplitude of the Ti:saph double-passed AOM (Figure 3.15). U_{atom} is modulated by about $180 E_{rec}$, resulting in a small measurement lever arm of about 2.

The calculation of this lever arm assumes that the shift is linear in lattice intensity. However,

lattice hyperpolarizability, which is proportional to the square of the lattice intensity, has been considered since the first proposals of the Sr lattice clock [58]. Also, a frequency shift from the lattice laser coupling to M1-E2 transitions, which goes as the square root of the lattice intensity, has been discussed [111]. Attempts have been made to measure the size of these effects [121]. As I will discuss in Sections 5.6.3 and 5.6.4, we have very carefully considered the treatment of these shifts in the Sr2 systematic evaluations.

5.6.2 Technical Considerations

To perform a robust and rigorous evaluation of the lattice Stark shift, other systematic shifts that could couple to this measurement must be controlled. As with the probe Stark shift, the background field servo is used to prevent the tensor term (Equation 5.18) in the lattice shift from drifting. Also, we align the bias field to the lattice polarization to ensure that the tensor term is first-order insensitive to any residual magnetic field drift that may exist.

Another potential problem is an uncontrolled density shift, which can mimic a lattice Stark effect due to its $U_{atom}^{3/2}$ dependence (Equation 5.17). Therefore, as the trap intensity is modulated for lattice Stark measurements, we also modulate the atom number to ensure that $NU_{atom}^{3/2}$ is the same for both lock-in states. In this case, the density shift is common mode to both states, and therefore it cancels in the lock-in data. To thoroughly enforce the constancy of $NU_{atom}^{3/2}$ in the presence of a fluctuating N , lock-in points for which the magnitude of $N_1U_1^{3/2} - N_2U_2^{3/2}$ is largest are removed in post processing until the data run average of $N_1U_1^{3/2} - N_2U_2^{3/2}$ is well below our measurement precision.²⁰ We calculate the maximum acceptable value of the magnitude of $N_1U_1^{3/2} - N_2U_2^{3/2}$ from the shift coefficients obtained from studying the density shift as a function of N and U_{atom} . We then enforce that the data run average of $N_1U_1^{3/2} - N_2U_2^{3/2}$ is more than 10 times smaller than this maximum value.²¹

²⁰ For one lock-in point $\nu_2 - \nu_1$, N_1 (N_2) and U_1 (U_2) are the atom number and observed U_{atom} for which ν_1 (ν_2) is measured.

²¹ Great care must always be taken when removing data points. It is important to stress that no data point is removed based on its frequency value; rather, measurements of the atom number and trap depth are used to compute $N_1U_1^{3/2} - N_2U_2^{3/2}$, and the value of this difference determines whether frequency points should be removed. Additionally, the cuts are performed in an automated fashion that is blind to the frequency values (looking only at

Lastly, as explained in Section 3.4.2, a Ti:saph laser is needed for lattice Stark measurements. This is because the amplified spontaneous emission from the TA-based lattice laser system can cause severe frequency shifts. To get a rough idea of the magnitude of the TA amplified spontaneous emission shift, we measure the change in the clock transition frequency caused by putting a narrow optical bandpass filter (centered on the magic wavelength) in the TA-based lattice laser, observing a 10^{-15} -level effect. Furthermore, comparisons between Sr1 and the Yb lattice clock at NIST suffered from 10^{-15} -level drift that has been attributed to the TA lattice laser Sr1 used at the time. This implies that the Sr1 amplified spontaneous emission spectrum was time varying. In other measurements that used both an optical filter and a prestabilization cavity to remove amplified spontaneous emission, Sr1 still observed significant time varying frequency shifts due to their TA lattice spectrum, implying that filtration is not enough to remove this effect. The Sr group at LNE-SYRTE observed similar issues with TA systems [114].

Fortunately this spectral drift will have no measurable effect on TA-based evaluations of the probe Stark shift, second-order Zeeman shift, and other shifts that are not caused by the lattice. Since a lock-in measurement is the difference between two data points taken within a few seconds of each other, the drift rate of the amplified spontaneous emission shift is slow enough that it is common mode between these two points. However, shifts caused by the lattice are measured over days or weeks, and the TA spectrum drift is significant on these timescales. A spectrally pure Ti:saph laser does not suffer from these problems.

5.6.3 First and Second Evaluations

The Ti:saph frequency of the first lattice Stark shift evaluation [13] is 368.5545965(3) THz. If U_1 and U_2 represent the lock-in parameter values of U_{atom} , $\Delta U = U_2 - U_1$ is varied from 34 to 184 E_{rec} (Figure 5.12a). For each value of ΔU , lock-in data is averaged to 1×10^{-17} statistical uncertainty. We also measured 87 E_{rec} to be the lowest value of U_{atom} for which $N = 2000$ could be reliably obtained, so 87 E_{rec} is a good choice for U_{clock} , the clock operation value of U_{atom} .²²

$N_1 U_1^{3/2} - N_2 U_2^{3/2}$.

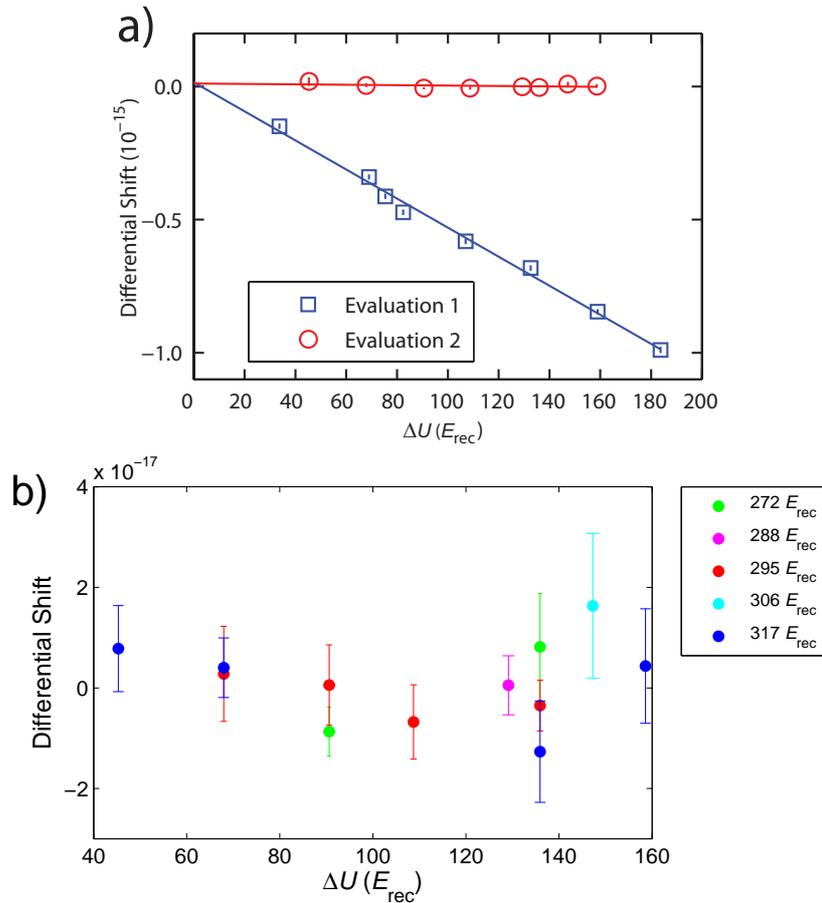


Figure 5.12: a) Lock-in data for the first evaluation (blue squares) and the second evaluation (red circles). Data from the first evaluation and other data (not shown) is used to compute the lattice frequency for the second evaluation. For the second evaluation, the fact that the shift remains zero for changes in the lattice intensity indicates that the magic wavelength has been identified. b) The lock-in data for the second evaluation with the values of U_2 specified (legend on right side).

To extrapolate this evaluation to the clock operation condition of $U_{clock} = 87 E_{rec}$, we must determine whether the data should be modeled as linear in U_{atom} or whether nonlinear terms like hyperpolarizability should be included (Sections 5.1.4 and 5.6.1).

The LNE-SYRTE reports of the hyperpolarizability and M1-E2 coefficients do not specify whether the density shift was under control during their measurements of these nonlinear effects (Section 5.6.1). Their determination of these coefficients involved modulating their lattice depth [119], and this could have been corrupted by an uncontrolled density shift. Also, measurements of

²² Using the analysis from Section 2.4.5, $U_{atom} = 87 E_{rec}$ corresponds to a one-way lattice power of 2.6 W.

nonlinear effects were performed by comparing two clocks with TA-based lattices [121]. Though the output of the TAs were filtered, one of the Sr systems used in these nonlinear shift measurements was later reported to still suffer from the effects of spectral drift [114].

Therefore, we perform an F -test on our data to see if hyperpolarizability and M1-E2 shifts are justified (Section 5.1.4). If hyperpolarizability were significant, the lock-in data $\Delta\nu_{lock-in}$ would need to be modeled as $k_1(U_2 - U_1) + k_2(U_2^2 - U_1^2)$, where k_1 and k_2 are constants. The F -test compares two fits of the extrapolated shift $\Delta\nu_{lock-in}U_{clock}/(U_2 - U_1)$: a constant fit (corresponding to a Stark shift linear in U_{atom}) and a fit to $k_1 + k_2x$ (linear shift plus hyperpolarizability), where $x = U_1 + U_2$. Comparing these two models, for 6 degrees of freedom the result is $F = 0.44$, corresponding to a test probability of 0.53. This is a comfortable margin above the 0.01 to 0.05 range below which a new term is required.

If an M1-E2 shift were significant, $\Delta\nu_{lock-in}$ should be fit to $k_1(U_2 - U_1) + k_2(\sqrt{U_2} - \sqrt{U_1})$. A second F -test compares a constant fit of $\Delta\nu_{lock-in}U_{clock}/(U_2 - U_1)$ to the fit $k_1 + k_2x$, where $x = 1/(\sqrt{U_1} + \sqrt{U_2})$. The result for 6 degrees of freedom is $F = 0.40$, resulting in the test probability 0.55. Therefore, the first evaluation does not support the addition of hyperpolarizability and M1-E2 terms.

Using reduced-chi-square-inflated least-squares fitting, the data from the first evaluation has a shift per E_{rec} of potential depth (at the atoms) of $(-5.31 \pm 0.05) \times 10^{-18} E_{rec}^{-1}$. For $U_{clock} = 87 E_{rec}$, the final value of the lattice Stark shift for the first evaluation is $(-4.62 \pm 0.04) \times 10^{-16}$.

The lattice Stark shift is known to also be linear in lattice frequency within tens of GHz of the magic wavelength [18]. Using this fact, data from the first evaluation and a handful of measurements taken at different lattice frequencies imply that the magic wavelength is 368.554485 THz. Therefore, our second evaluation [88] of the lattice Stark shift is performed with a lattice frequency of 368.5544849(1) THz. For these data, the lattice Stark shift has largely vanished (Figure 5.12a and 5.12b). We also achieve a lower value of $U_{clock} = 71 E_{rec}$ by operating the Sr oven reservoir at a hotter temperature of 600 °C.²³ Lowering the lattice depth means that the

²³ Using the analysis from Section 2.4.5, this corresponds to a one-way lattice power of 2.3 W.

trap will retain fewer atoms, so this U_{clock} is again determined as the lowest potential depth at which $N = 2000$ can be reliably obtained. Running the reservoir hotter increases the amount of trapped atoms overall (for any lattice depth), implying that $N = 2000$ can be achieved for a lower value of U_{atom} . We also find that the atom number is best when we load atoms into a deep lattice and then ramp down the lattice intensity to the final value rather than load in a shallow lattice.

Performing an F -test on these data to look for hyperpolarizability and M1-E2 effects (in the same manner as the first evaluation), the test yields $F = 0.17$ for hyperpolarizability and $F = 0.03$ the M1-E2 shift for 22 degrees of freedom.²⁴ These correspond to probabilities of 0.68 and 0.86, respectively. Therefore, we again model the data from the second evaluation as linear in U_{atom} . Extrapolating the data to $U_{clock} = 71 E_{rec}$, we obtain a final value of the lattice Stark shift of $(-1.3 \pm 1.1) \times 10^{-18}$. The final shifts from the first and second evaluations imply a magic wavelength of 368.5544845(3) THz, which is statistically consistent with the value used for the second evaluation. These evaluations also imply a lattice Stark shift per U_{atom} per detuning from the magic wavelength of $(-4.73 \pm 0.02) \times 10^{-17} \text{ GHz}^{-1} E_{rec}^{-1}$.

5.6.4 Hyperpolarizability

Although we have chosen not to include previously reported values of the LNE-SYRTE hyperpolarizability coefficient $(0.45 \pm 0.10) \mu\text{Hz } E_{rec}^{-2}$ [114], the uncertainty in the correction at $U_{clock} = 71 E_{rec}$ would be only 1.2×10^{-18} . It is also possible to process the data from the second evaluation (Figure 5.12) in a manner that is sensitive only to higher-order terms. In the presence of hyperpolarizability, the extrapolated shift will have the form $k_1 + k_2(U_1 + U_2)$; therefore, differences between extrapolated shift data points that have different values of $U_1 + U_2$ will depend only on hyperpolarizability (and other nonlinear terms). With this analysis, and treating hyperpolarizability as the most significant nonlinear shift, we infer a hyperpolarizability coefficient of $(0.3 \pm 0.3) \mu\text{Hz } E_{rec}^{-2}$. Although this is consistent with the LNE-SYRTE value, it is also consistent

²⁴ The degrees of freedom for this F -test is equal to the number of points minus 2. Here I quote 22 degrees of freedom, yet in Figure 5.12, the two plots of the second evaluation data have 8 and 12 points. This is because the F -test is performed on unbinned data, whereas the points in Figures 5.12a and 5.12b are binned.

with zero.

The magic wavelength of 368.554465(3) THz measured by the PTB Sr group [33] differs from our value by 19.5 MHz. One difference in their approach that can account for some of this discrepancy is that they assume the LNE-SYRTE hyperpolarizability coefficient and subtract it off of each of their lock-in measurements.

5.7 Dc Stark Shift

5.7.1 Introduction

A large systematic shift caused by dc electric fields was recently discovered by the LNE-SYRTE group [69]. Their in-vacuum lattice cavity mirror substrates had acquired patch charges that caused clock transition shifts at the 10^{-13} level, and they were able to remove these charges using UV treatment [69]. Motivated by this discovery, we looked for a dc Stark shift in Sr2, and an effect was indeed present [13]. The PTB group also measured a significant dc Stark shift that they attributed to an unwanted voltage across their MOT coils [33].

We measure the dc Stark shift using electrodes, with one placed above (below) our top (bottom) vacuum chamber viewports. Obtaining a value for the shift requires a measurement of the clock transition frequency for each of three conditions: the top and bottom electrodes held at voltages V and $-V$ (respectively), the top and bottom electrodes held at voltages $-V$ and V , and both electrodes grounded. The chamber metal is grounded, and V is usually around 90 volts. When the electrode field is applied, the clock transition frequency is

$$\nu_{\pm} = \nu_{Sr} - \frac{1}{2\hbar} \Delta\alpha(0) |\vec{\mathcal{E}}_{bg} \pm \vec{\mathcal{E}}_{elec}|^2, \quad (5.19)$$

where ν_{Sr} is the unshifted clock transition frequency, $\Delta\alpha(0)$ is the differential dc polarizability, $\vec{\mathcal{E}}_{bg}$ is the background electric field, and $\pm\vec{\mathcal{E}}_{elec}$ is the dc field from the electrodes. The \pm sign corresponds to the two directions of the electrode field. When the electrodes are grounded, the clock transition frequency is

$$\nu_{gnd} = \nu_{Sr} - \frac{1}{2h} \Delta\alpha(0) \mathcal{E}_{bg}^2. \quad (5.20)$$

These three frequency measurements can be combined to obtain the applied shift,

$$\nu_{app} = -\frac{1}{2h} \Delta\alpha(0) \mathcal{E}_{elec}^2 = \frac{1}{2}(\nu_+ + \nu_-) - \nu_{gnd}, \quad (5.21)$$

and the background shift,

$$\nu_{bg} = -\frac{1}{2h} \Delta\alpha(0) \mathcal{E}_{bg}^2 = \frac{(\nu_+ - \nu_-)^2}{16\nu_{app}}. \quad (5.22)$$

5.7.2 Electrodes and Lock-In Sequence

Two steel quadrant electrodes provide the dc electric field, one of which is mounted above the top viewport and the other below the bottom viewport (Figure 5.13a). Both electrodes are cut into four pieces, which can be hooked up in three different ways to create fields along three orthogonal axes (Figures 5.13c and 5.13d). Also, a hole is drilled in the center of both electrodes for MOT laser clearance. Finite element analysis of the chamber and electrodes shows an electric field that is reasonably uniform over the ultracold sample (Figure 5.13b).²⁵ The electrode voltage is provided by two low-noise commercial power supplies, one for the positive voltage and one for the negative. Reed relays allow us to choose which of these voltages are applied to which electrode using TTL logic.

The dc Stark shift sensitivity to the squared magnitude of this electric field is $-7.2 \times 10^{-17} \text{ V}^{-2} \text{ cm}^2$; therefore, a small applied field of only a few V/cm^2 (Figure 5.13b) will be sufficient to produce a large effect.²⁶ To ensure that this measurement is stable and repeatable, it is important to use the background field servo, which prevents the tensor component of the dc Stark shift from drifting. To facilitate the background field servo and the three clock transition mea-

²⁵ High field uniformity is not a requirement for this measurement. All that is needed is for the field to have the same magnitude when reversed and to be constant over the time it takes to measure one value of ν_+ and one value of ν_- .

²⁶ I calculated this sensitivity using the known value of the Sr clock transition's dc polarizability [82]

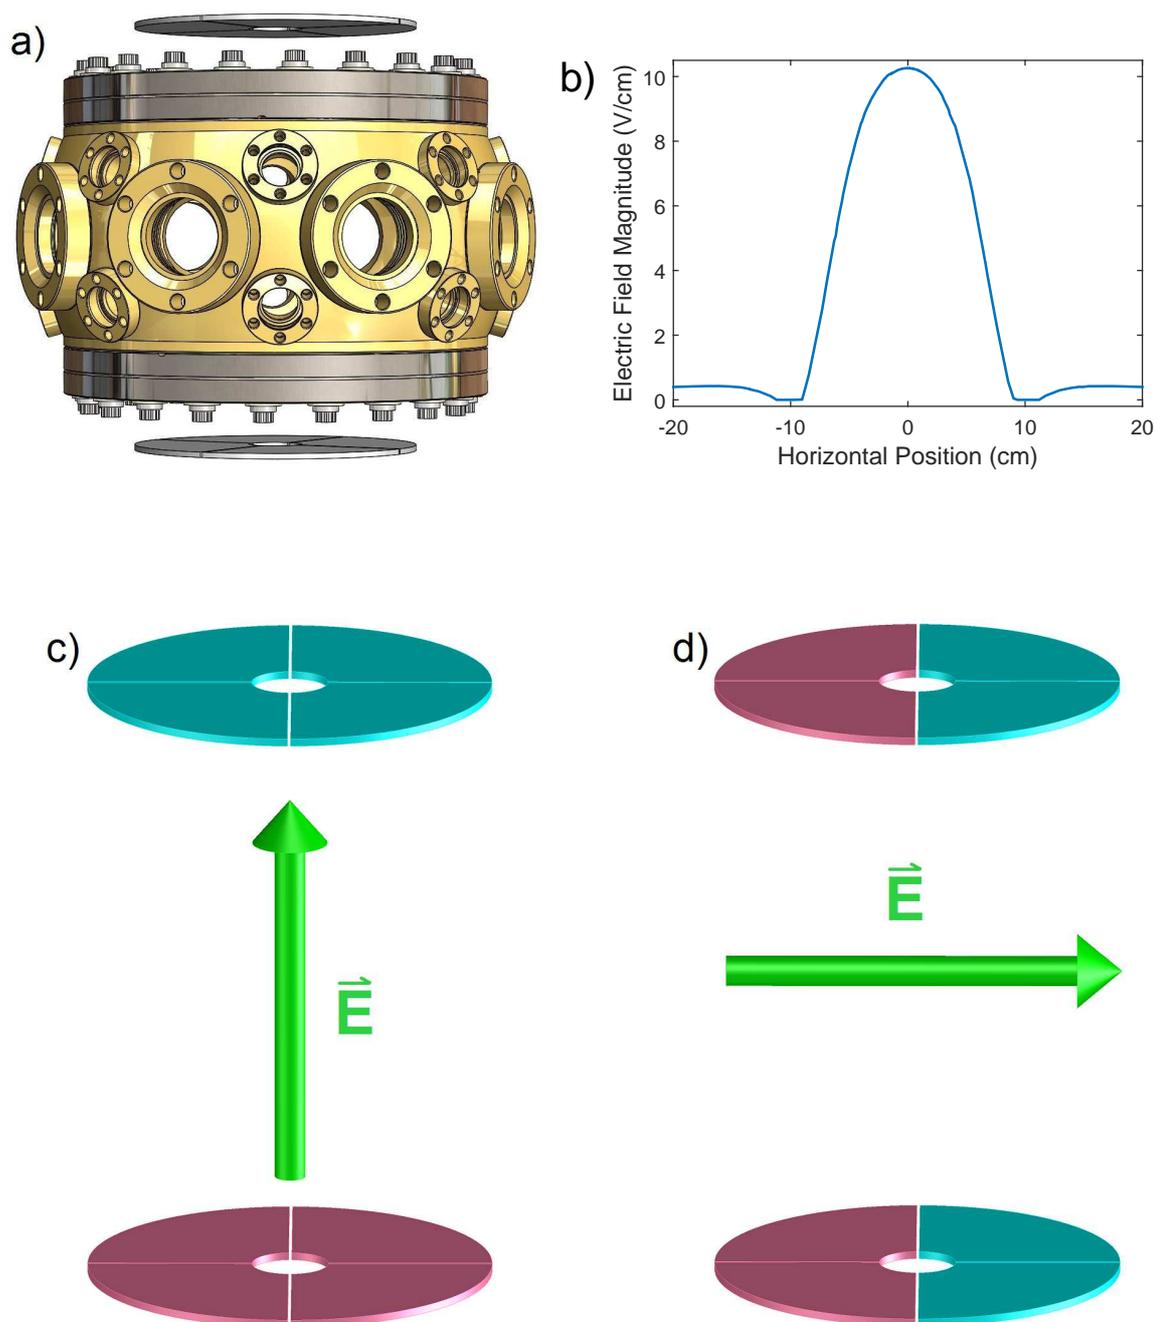


Figure 5.13: a) Electrodes are mounted above and below the vacuum chamber. They are separated by 17.5 cm. b) Finite element model of the electrode field. The potential difference between electrodes is 180 V. In this plot the field points in the vertical direction. Moving about the horizontal plane, with the vertical position fixed at the center of the chamber, the field is reasonably uniform over a 1 cm length scale. Along the vertical coordinate (not shown), the field is even more uniform. c) The electrodes have holes for MOT laser clearance and are cut into quadrants. The quadrants are colored cyan (pink) if they are held at a negative (positive) potential. If all four quadrants of the top (bottom) electrode are shorted together, the electric field is vertical. d) Holding halves of the electrodes at different potentials produces a horizontal field. The quadrants can also be arranged to create a field in the other horizontal dimension.

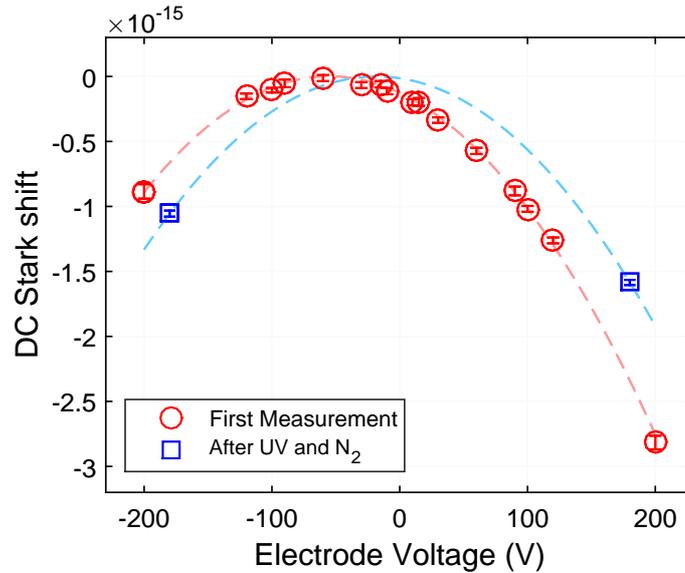


Figure 5.14: The first check for a dc Stark shift (red circles), and the shift after UV and nitrogen gas treatment (blue squares). Both data sets are from measurements along the vertical direction. The red data implies a background dc Stark shift of -1.3×10^{-16} . After removing the shift with UV light and nitrogen, the shift is -1.6×10^{-18} .

measurements needed to compute the background dc Stark shift (Equation 5.22), the lock-in sequence uses four atomic servos. For two of these servos, the electrodes are grounded and a stretched state lock is performed, yielding the zero-magnetic-field clock transition frequency that is needed for the background field servo. For the other two atomic servos, there is one servo for each direction of the electrode field, and they both operate on the $m_F = 9/2$ stretched state. After one cycle through the four servos, the three center frequencies measured for $m_F = 9/2$ are used to infer the background shift.²⁷ The frequency measured for $m_F = -9/2$ is used only for the background field servo and does not contribute to the dc Stark shift data.

5.7.3 Observation of Dc Stark Shift and Removal of Background Field

During the first check for a dc Stark shift in Sr2, we were surprised to find a large background shift in the vertical direction. Fitting dc Stark shift data measured as a function of the potential difference between the electrodes (Figure 5.14), a background shift of $(-1.27 \pm 0.05) \times 10^{-16}$ is

²⁷ These three frequencies are taken to be ν_+ , ν_- , and ν_{gnd} in Equation 5.22.

observed. This corresponds to a background field of $(1.33 \pm 0.03) \text{ V/cm}$.²⁸

Simple checks for background fields caused by our MOT coils did not turn up anything; therefore, we hypothesized that patch charges on the Sr2 vacuum viewports were causing the shift. In attempt to remove this shift, we treated the vacuum viewports with radiation from a UV lamp [69]. The UV treatments did have a positive effect on the shift but progress was very slow, and we projected from our measurements that it could take weeks or months to remove a -1.3×10^{-16} effect. After a couple of days of this procedure, the UV treatment was temporarily set aside so that we could install temperature sensors in the vacuum chamber (discussed in Section 5.8.5). This installation required that the vacuum chamber be vented to high-purity nitrogen. After venting, the shift was orders of magnitude smaller.²⁹

UV treatment followed by a nitrogen vent reduced the shift in the vertical direction to $(-1.6 \pm 1.0) \times 10^{-18}$ (Figure 5.14). We also looked for a horizontally oriented background dc Stark shift and obtained a value of $(-1.9 \pm 1.9) \times 10^{-18}$. The shift proved stable, so for the first evaluation we combine the vertical and horizontal measurements for a final dc Stark shift of $(-3.5 \pm 2.1) \times 10^{-18}$ [13].

The horizontal shift was persistently below measurement resolution even when the vertical shift was quite large. If a patch charge on a Sr2 viewport were responsible for the shift, this patch would have the field of an electric monopole [46]. For a monopolar field to cause a dc Stark shift in only the vertical direction, it would need to be located directly above or below the atoms (depending on the sign of the charge).³⁰

²⁸ Despite that the electrodes are separated by 17.5 cm, the fit tells us that the applied field behaves as if the plates are separated by $(41.4 \pm 0.3) \text{ cm}$. We did not attempt to track down the cause of this discrepancy since it does not affect our measurements.

²⁹ Later, cleaning the top Sr2 viewport created a 10^{-15} -level dc Stark shift. We wiped the outer viewport surface with material that should have removed static charge buildup, but the shift did not change. Therefore, we repeated the same procedure of UV treatment followed by venting the chamber to clean nitrogen. This removed the shift again.

³⁰ Reference [46] shows that laser light causes electron emission from anti-reflection coated viewports. The parts of the top and bottom viewports illuminated most brightly by the MOT lasers are directly above and below the atoms, so this is where positive patch charges resulting from electron emission would be located. Given the sign of the background field, the top viewport would need to be the bad actor. However, this explanation alone is insufficient to explain the Sr2 dc Stark shift because it is unclear why electron emission would not be significant for the five other viewports (all with the same coating) that MOT lasers are sent through.

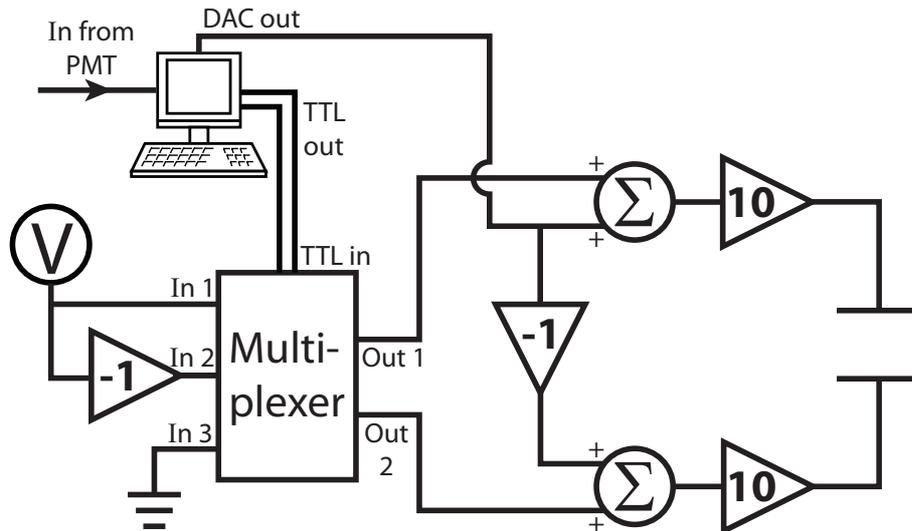


Figure 5.15: The control loop that cancels the background shift. While the stretched-states servo is running, two TTL lines from the computer set both multiplexer outputs to ground. In this case, the computer DAC card provides a voltage that, after amplification by a factor of 10, charges the electrodes (shown here as a capacitor). The gain = 10 stages are composite amplifiers based on high-voltage op-amps controlled for noise reduction using low-noise, lower-voltage op-amps. The DAC voltage is periodically adjusted to cancel the background dc Stark shift. Background shift cancellation utilizes a fixed dc voltage source depicted as a circled V. To measure ν_+ , the multiplexer sets $\text{Out } 1 = V_s$ and $\text{Out } 2 = -V_s$, where V_s is the voltage of the dc source (similarly, for ν_- , $\text{Out } 1 = -V_s$ and $\text{Out } 2 = V_s$). The potential difference of $\pm 20V_s$ across the electrodes is summed with the amplified DAC voltage used to cancel the background Stark shift. If the background shift is not well canceled by the DAC, $\nu_+ - \nu_-$ will not be zero. Excited state fraction measurements made by the Sr2 PMT are processed by the computer to create the error signal $\nu_+ - \nu_-$. A PI filter (with gain constants optimized empirically) uses this error signal to compute the voltage that the DAC must apply to balance ν_+ and ν_- and cancel the background shift. After one of these measurements, the multiplexer again grounds its outputs, and clock operation resumes with the new DAC voltage.

5.7.4 Active Control of the Dc Stark Shift

The dc Stark shift was measured again for the second evaluation. This time around, the horizontal shift remained unresolved, but the vertical shift seemed to be less stable. To remove the (possibly time varying) vertical shift in a dynamic way, we use active control similar to the background field servo. This “electrode servo” is based on measurements using an applied electric field (Figure 5.15). During clock operation, the stretched-states servo pauses after tens of cycles, and the electrode servo takes two center frequency measurements. With the quadrant electrodes

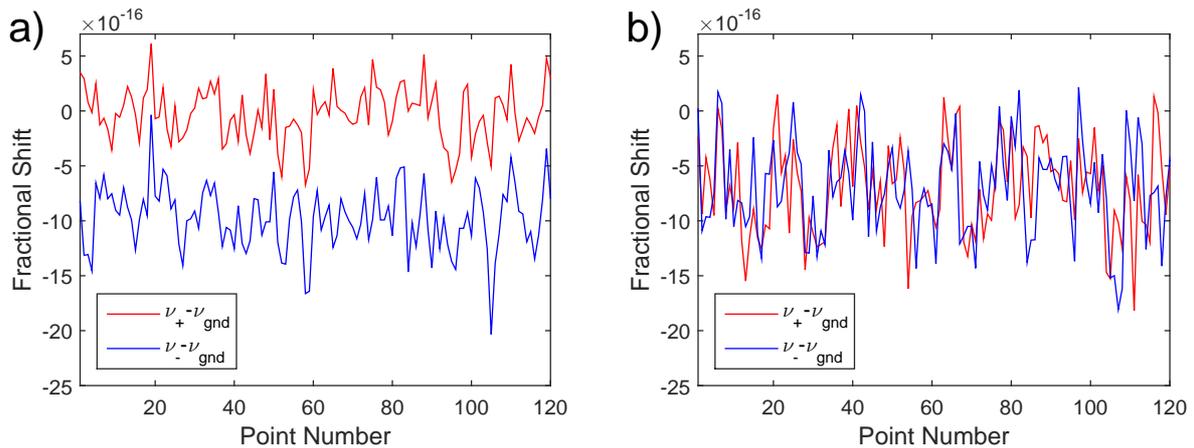


Figure 5.16: Dc Stark shift frequency data. A background shift means that $\nu_+ - \nu_{gnd}$ and $\nu_- - \nu_{gnd}$ will be different. a) Frequency data before the first UV treatment and N_2 vent were performed. It is only a coincidence that the background field nearly cancels the applied field in the + direction. b) Frequency data with the shift under active control.

configured for a vertical field, the servo measures ν_+ (ν_-), which is the center frequency when the field is pointing up (down). The difference between these center frequencies is $\nu_+ - \nu_- = 2 \Delta\alpha(0) \mathcal{E}_{bg} \mathcal{E}_{elec}$, which is proportional to the background field. This frequency difference is treated as an error signal that uses a digital proportional-integral (PI) filter to compute the voltage at which these two frequencies are equal.³¹ The electrodes are then held at this voltage, and the stretched-states servo resumes.

A key feature of this servo is the fact that the voltage that cancels the background field and the voltage that creates the applied shift (used to measure $\nu_+ - \nu_-$) come from independent sources (Figure 5.15). The DAC voltage that cancels the background field is always on, whereas the voltage that creates the applied shift is fixed and is only used when $\nu_+ - \nu_-$ is being measured. This means that $\nu_+ - \nu_-$ measures the residual field that is not canceled by the DAC voltage. In principle, this technique can be scaled up to servo a background field that has components in all three dimensions. There would need to be a DAC voltage and composite amplifier (Figure 5.15) for each quadrant of both electrodes.³²

³¹ The fact that the error signal $\nu_+ - \nu_-$ is also proportional to the magnitude of the applied electrode field \mathcal{E}_{elec} is not an issue for this application. \mathcal{E}_{elec} is fixed for this measurement, so it just acts like additional proportional gain.

³² We had these electronics ready, but in the end only vertical shift control proved important.

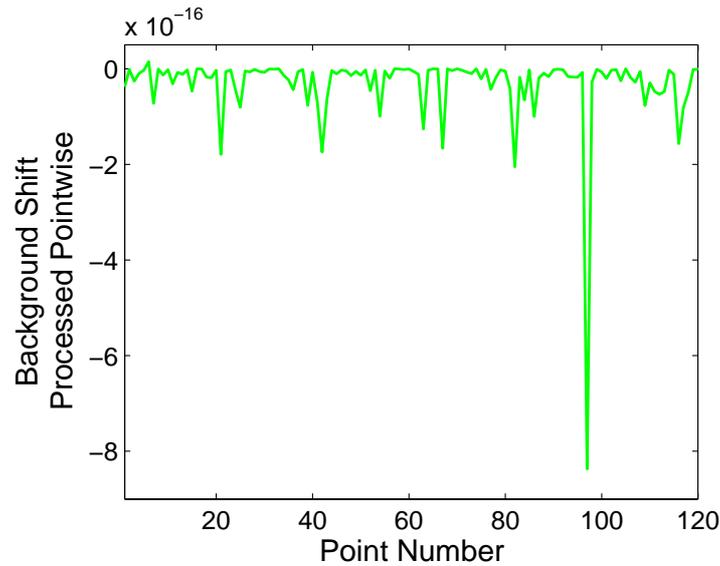


Figure 5.17: Data from Figure 5.16b processed point by point using Equation 5.22. Processing data in this way biases the mean toward negative values.

With the electrode servo operating, we observe a clear effect (Figure 5.16b). It is worth noting that a well-canceled shift means that the dc Stark effect can be evaluated rapidly. To see this, assume that each center frequency in Equation 5.22 has the same statistical uncertainty σ . Propagating σ to ν_{bg} , the (squared) uncertainty in the background shift is

$$\delta\nu_{bg}^2 = \frac{1}{2} \frac{\nu_{bg}}{\nu_{app}} \left(1 + 3 \frac{\nu_{bg}}{\nu_{app}} \right) \sigma^2. \quad (5.23)$$

In this expression, ν_{bg} is treated as the background field plus the cancellation field from the servo. Therefore, with the background field well canceled, the dc Stark shift averages down very rapidly. About 20 minutes of averaging results in a total shift of $(-0.1 \pm 1.1) \times 10^{-19}$ [88].

Care must be taken when processing measurements of ν_+ , ν_- , and ν_{gnd} to obtain values of ν_{bg} . Since the applied electric fields for these measurements are small, ν_{bg} computed for each point (using Equation 5.22) has unusual noise properties. The fact that the data scatters near or above zero causes erratic fluctuations when it is processed point by point, making the usual statistical analysis impossible (Figure 5.17). Simulating this measurement, the same fluctuations toward negative numbers are reproduced when the simulated data is processed as in Figure 5.17, resulting in a data

mean that is much more negative than the simulation's true value. However, the simulation does show that the true value can be recovered if ν_{bg} is computed as $\langle(\nu_+ - \nu_-)^2\rangle/8\langle\nu_+ + \nu_- - 2\nu_{gnd}\rangle$. Here $\langle(\nu_+ - \nu_-)^2\rangle$ is the mean of the numerator of Equation 5.22 processed point by point, and $8\langle\nu_+ + \nu_- - 2\nu_{gnd}\rangle$ is the mean of the point-by-point-processed denominator. The uncertainty in this ratio is the propagated standard deviation of the mean of the numerator and denominator, which is shown by the simulation to be a good assignment of uncertainty for the measurement. Given the conclusions of this simulation, all dc Stark background shift data sets are averaged as $\langle(\nu_+ - \nu_-)^2\rangle/8\langle\nu_+ + \nu_- - 2\nu_{gnd}\rangle$.

5.8 Static Blackbody Shift

5.8.1 Introduction to the Blackbody Radiation Shift

The blackbody radiation (BBR) shift causes the largest systematic uncertainty in many prominent lattice [114, 74] and ion [52, 76] clocks. Of all optical clocks currently in development, Sr unfortunately has the largest BBR shift [104]. The effect is essentially the ac Stark shift caused by the electric field from the ambient heat in the laboratory. Since the electric field associated with heat is distributed over a range of frequencies, Equation 2.15 can be converted to an integral form to express the BBR shift,

$$\Delta\nu_{BBR} = \frac{-1}{2\epsilon_0\hbar} \int_0^\infty \Delta\alpha(\omega) \mathcal{I}(\omega) d\omega, \quad (5.24)$$

where $\mathcal{I} = dI/d\omega$ is the electric field intensity per unit angular frequency. For a blackbody, \mathcal{I} is the Planck distribution

$$\mathcal{I}(\omega) = \frac{\hbar}{\pi^2 c^2} \frac{\omega^3}{e^{\hbar\omega/k_B T} - 1}. \quad (5.25)$$

As Equation 2.17 shows, the polarizability $\Delta\alpha(\omega)$ varies over frequency scales given by atomic transition energies. Since room-temperature BBR peaks at about a 10 μm wavelength, whereas the relevant atomic resonance wavelengths are orders of magnitude shorter, the atomic polarizability

is nearly constant over the room-temperature BBR spectrum. Therefore, Equation 5.24 is often approximated by performing a *dc expansion* $\Delta\alpha(\omega) = \Delta\alpha(0) + \frac{1}{2}\Delta\alpha''(0)\omega^2 + \dots$ on the differential polarizability, which results in

$$\Delta\nu_{BBR} \simeq -\frac{\Delta\alpha(0)}{2\epsilon_0 ch} \int_0^\infty \mathcal{I}(\omega)d\omega - \frac{\Delta\alpha''(0)}{2\epsilon_0 ch} \int_0^\infty \omega^2 \mathcal{I}(\omega)d\omega \quad (5.26)$$

$$= \nu_{stat} \left(\frac{T}{T_0}\right)^4 + \nu_{dyn} \left(\frac{T}{T_0}\right)^6. \quad (5.27)$$

The remaining terms in the polarizability expansion have been neglected since they result in corrections that are well below the current total uncertainty of the Sr clock. The first term in Equation 5.27 is called the *static shift*, which approximates the polarizability as its dc value, and the second term is the *dynamic shift*. The reference temperature $T_0 = 300$ K, and

$$\nu_{stat} = -\frac{\pi}{60c^3\epsilon_0} \left(\frac{k_B T_0}{\hbar}\right)^4 \Delta\alpha(0), \quad (5.28)$$

$$\nu_{dyn} = -\frac{4\pi^4}{21} \left(\frac{k_B T_0}{\hbar}\right)^6 \sum_k \left(\frac{A_{ek}}{\omega_{ek}^6} - \frac{A_{gk}}{\omega_{gk}^6}\right). \quad (5.29)$$

Here A_{ek} and ω_{ek} (Equation 2.17) are the Einstein coefficients and atomic resonance frequencies associated with coupling the excited clock state to the hyperfine state k , and A_{gk} and ω_{gk} correspond to the ground state coupling to k (where $|k\rangle \neq |g\rangle$ and $|k\rangle \neq |e\rangle$).

5.8.2 Sources of Uncertainty

As Equation 5.27 shows, physically there are two sources of uncertainty in the BBR shift: uncertainty in the atomic response to BBR, which is described by the shift coefficients ν_{stat} and ν_{dyn} , and uncertainty in measurements of the ambient temperature T at the location of the atoms. In 2008, uncertainty in the Sr1 clock from the shift coefficients was 7×10^{-17} [95, 74]. In 2012, the PTB Sr group performed an important measurement of the differential dc polarizability of the clock transition, resulting in $\nu_{stat} = (-4.9629 \pm 0.0001) \times 10^{-15}$ [82]. With the static coefficient uncertainty at 1×10^{-19} , it can be neglected for both Sr2 evaluations. The PTB group also calculated

the dynamic coefficient at 1.6% uncertainty [82]. Later, another team calculated ν_{dyn} with a bit lower uncertainty [105], resulting in $\nu_{dyn} = (-3.48 \pm 0.04) \times 10^{-16}$. Prior to these calculations, the dynamic coefficient had 7% uncertainty [105, 84].

Improving on the uncertainty in ν_{dyn} will not be necessary for our first systematic evaluation, which has the goal of breaking the previous clock record of 8.6×10^{-18} total uncertainty; however, the second evaluation needs a more accurate determination of ν_{dyn} to reach low- 10^{-18} total uncertainty. Although uncertainty in ν_{stat} is negligible, the static shift is the dominant contributor to the portion of the BBR shift uncertainty caused by δT (the uncertainty in T). Therefore, one can think of the static shift and the dynamic shift as two separate effects that must be tackled with different strategies. Reducing the uncertainty in the dynamic shift through a better determination of ν_{dyn} requires measurements of Sr atomic properties [105]. Meanwhile, reducing the uncertainty in the static shift due to δT requires improvements to the Sr2 apparatus.

For many years, different Sr teams measured T by attaching several temperature sensors to their clock vacuum chambers [74, 114, 33]. T was determined from this array of sensors as the average of the hottest and coldest measurements, and δT was taken as half of the difference between these two extreme values. With this scheme, δT was usually reported to be about 1 K, which results in a BBR shift uncertainty of 7×10^{-17} . To improve the Sr BBR shift uncertainty beyond this, a setup based on a cryogenic environment was proposed [83]. Cold environments, which have successfully mitigated the blackbody shift in ion clocks [103] and Cs fountains [48], would reduce the shift immensely since $\Delta\nu_{BBR}$ is a rapidly decreasing function of T .

Cold environments are unattractive to us because of their complexity and invasiveness. A cold environment would at least require the Sr2 main vacuum chamber to be vented to install the necessary hardware, meaning that much of the apparatus would need to be dismantled to bake the chamber and restore the vacuum. This feat could take the better part of a year. Also, it is possible that a good cryogenic system would demand an entirely new clock apparatus designed around it. These technical issues motivated us to come up with a simpler, room-temperature strategy for improving BBR shift uncertainty.

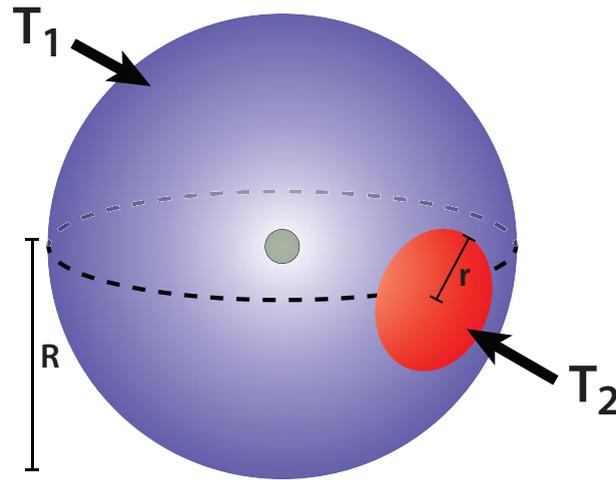


Figure 5.18: In this simple model, the chamber is treated as a sphere with a patch at temperature T_2 and the rest of the chamber at T_1 . The atoms are located at the center.

5.8.3 Radiation Thermometry

Determining the BBR shift with arrays of sensors mounted to a vacuum chamber has large uncertainty because it is difficult to characterize and model the chamber at the level required to accurately infer the ambient temperature where the atoms are trapped. The vacuum chamber provides a complex boundary that would need to be measured with a fine-mesh sensor array and well modeled (accounting for both the realistic geometry and material properties of the chamber) to predict the temperature at the center with low uncertainty. Instead, we opt to put a well-calibrated thermometer inside the vacuum chamber near the location of the atoms. This allows us to directly measure the radiative temperature at the chamber center.

To see why an in-vacuum radiation thermometer is a dramatic improvement over a chamber sensor array, it is constructive to consider a simple model. Treating the chamber as a spherical shell at two temperatures (Figure 5.18) with the atoms at the center, one can gain intuition about radiation thermometry even in a non-thermal (out of equilibrium) environment. Assuming steady state, the electric field intensity per unit frequency at the atoms is $\mathcal{I}(\omega) = \Omega_1 \mathcal{I}_1(\omega) + \Omega_2 \mathcal{I}_2(\omega)$. Here \mathcal{I}_1 and \mathcal{I}_2 are the intensities per frequency at the two boundary surfaces (Figure 5.18). Ω_1 and Ω_2 are the area fractions, which are the ratios of the areas of the surfaces at T_1 and T_2 (Figure

5.18) to the total surface area $4\pi R^2$. Using this $\mathcal{I}(\omega)$ in the BBR expression from Equation 5.26,

$$\Delta\nu_{BBR} = \frac{\nu_{stat}}{T_0^4}(\Omega_1 T_1^4 + \Omega_2 T_2^4) + \frac{\nu_{dyn}}{T_0^6}(\Omega_1 T_1^6 + \Omega_2 T_2^6). \quad (5.30)$$

First I will use this model to analyze the error inherent in using measurements of the chamber temperature to infer the temperature at the atoms. One typically calculates the atom temperature from measurements of the chamber surface temperatures as $T_{surf} = (T_1 + T_2)/2$, and the uncertainty in T_{surf} is computed as $\delta T_{surf} = (T_2 - T_1)/2$. Since other Sr teams typically report 1 K temperature variation across their chambers, I will take $\overline{\delta T_{surf}} = 1$ K. The BBR shift is then calculated by evaluating Equation 5.27 at $T = T_{surf}$. The uncertainty in this approach results from propagating δT_{surf} to $\Delta\nu_{BBR}$, and the error is the difference between the true shift (Equation 5.30) and the BBR shift calculated with T_{surf} .

Writing the true shift (Equation 5.30) in terms of T_{surf} and δT_{surf} and expanding it in $\delta T_{surf}/T_{surf}$ (since T_{surf} is roughly 300 K), the error ν_{err} between the calculated and true values is

$$\nu_{err} \simeq -(\Omega_1 - \Omega_2) \delta\nu_{surf}. \quad (5.31)$$

$\delta\nu_{surf}$ is the BBR shift uncertainty calculated by assuming that δT_{surf} is the measurement uncertainty,

$$\delta\nu_{surf} = \left| \frac{\partial}{\partial T} \Delta\nu_{BBR} \right|_{T=T_{surf}} \delta T_{surf}. \quad (5.32)$$

Since in reality, the spatial behavior of the chamber temperatures is not well known, Ω_1 and Ω_2 should be regarded as unknown quantities in Equation 5.31. Both Ω_1 and Ω_2 range from 0 to 1, so it is reasonable to bound ν_{err} within $\pm\delta\nu_{surf}$; therefore, $\delta\nu_{surf}$ is a good choice for the uncertainty. For $\delta T = 1$ K, $\delta\nu_{surf} = 7 \times 10^{-17}$.

Now suppose a thermometer is measuring the radiative temperature inside the chamber very near the atoms. The sensor absorbs thermal energy and provides a readout of its own temperature,

T_{sens} . Since the sensor equilibrates quickly, the heat energy it radiates per unit time can be treated as proportional to T_{sens}^4 (via the Stefan-Boltzmann law). Assuming the sensor to be an ideal blackbody, $T_{sens}^4 \propto \int_0^\infty \mathcal{I}(\omega)d\omega$, or

$$T_{sens} = (\Omega_1 T_1^4 + \Omega_2 T_2^4)^{1/4}. \quad (5.33)$$

Although the sensor exactly provides the static part of the true shift (Equation 5.30), which is also proportional to the fourth power of temperature, it does not provide the exact dynamic part. The error ν_{err} in the radiation thermometer case is the difference between the true shift and the BBR shift calculated with T_{sens} . To compare with the case of the chamber sensor array, it is useful to write the radiation thermometry BBR shift error in terms of T_{surf} and δT_{surf} . Expanding this,

$$\nu_{err} \simeq 6 [1 - (\Omega_1 - \Omega_2)^2] \nu_{dyn} \left(\frac{T_{surf}}{T_0} \right)^6 \left(\frac{\delta T_{surf}}{T_{surf}} \right)^2. \quad (5.34)$$

Assuming $\Omega_1 = \Omega_2$ (which maximizes the error) and taking $T_{surf} \simeq T_0$, $\nu_{err} = -2.3 \times 10^{-20}$, which is negligible for the total uncertainty goals in this work.

In summary, for a distribution of chamber temperatures that are within a degree of each other, the radiation thermometer provides the static shift and a good approximation to the dynamic shift. As I will explain later, a more complicated model of these processes (Section 5.21) also concludes that radiation thermometry error is negligible.

5.8.4 Kirchhoff's Radiation Law

Section 5.8.3 regarded the sensor as a device that treats T_{sens}^4 as proportional to the frequency-integrated spectral distribution of heat intensity incident on the sensor surface. For simplicity, the case of perfect sensor surface absorption and emission (an ideal blackbody) was considered; however, even sensors that are coated with highly absorptive paint are not ideal blackbodies. If imperfect absorption and emission significantly affect radiation thermometry in Sr2, we would need to perform

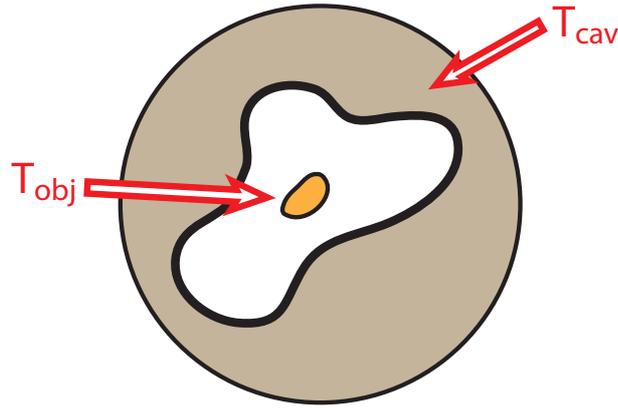


Figure 5.19: The cavity considered in the simple explanation of Kirchhoff’s radiation law given in the text. The cavity temperature is fixed, and the temperature of the object inside can vary. The temperature of the object can only change by heat emission or absorption.

the difficult task of characterizing the absorptive and emissive properties of the sensor surface.³³

Fortunately, for a well-equilibrated environment, Kirchhoff’s radiation law [85] means that imperfect absorption and emission are not concerns. To develop an understanding of Kirchhoff’s law, it is instructive to consider a simple case. Suppose a cavity is carved out inside a mass of thermally conductive material, which has a fixed temperature T_{cav} (Figure 5.19). Inside the cavity, there is another thermally conductive object at temperature T_{obj} , which can vary. The inner surface of the cavity and outer surface of the object inside the cavity have emissive and absorptive properties that are uniform over each surface; however, emission and absorption at these surfaces is not necessarily perfect.

The electromagnetic power incident on the object is $P_{in} = \epsilon_{cav}P_{cav} + \rho_{cav}P_{out}$. ρ_{cav} is the *reflectivity* of the cavity inner surface, which is defined as the reflected power divided by the incident power [85]. ϵ_{cav} is known as the *emissivity* of the cavity inner surface [85]. For a mass at temperature T , the emissivity is defined as the emitted power divided by the power of a perfect blackbody at temperature T . An emissivity of unity is considered perfect since heat radiation is maximal for a perfect blackbody. P_{cav} is the power of a perfect blackbody at temperature T_{cav} . P_{out} is the power directed away from the object in the cavity, due to either emission from that object or reflection

³³ These surface properties are difficult to characterize because they generally depend on wavelength, angle of incidence or emission, and temperature.

off of the object's surface. This quantity can be expressed as $P_{out} = \epsilon_{obj}P_{obj} + \rho_{obj}P_{in}$. Here ϵ_{obj} and ρ_{obj} are the emissivity and reflectivity of the object's surface, and P_{obj} is the electromagnetic power of a mass at temperature T_{obj} .

Solving for the incident and output power,

$$P_{in} = \frac{\epsilon_{cav}P_{cav} + \rho_{cav}\epsilon_{obj}P_{obj}}{1 - \rho_{cav}\rho_{obj}}, \quad (5.35)$$

$$P_{out} = \frac{\epsilon_{obj}P_{obj} + \rho_{obj}\epsilon_{cav}P_{cav}}{1 - \rho_{cav}\rho_{obj}}. \quad (5.36)$$

In steady state, $P_{in} = P_{out}$, implying that

$$\frac{\alpha_{obj}\epsilon_{cav}P_{cav} - \alpha_{cav}\epsilon_{obj}P_{obj}}{1 - \rho_{obj}\rho_{cav}} = 0. \quad (5.37)$$

Here I have made use of α_{obj} (α_{cav}), the *absorptivity* of the object's (cavity's) surface. The absorptivity is defined as the absorbed intensity divided by the incident intensity, which is unity for a perfect absorber. This implies that $\alpha_{obj} + \rho_{obj} = 1 = \alpha_{cav} + \rho_{cav}$. In steady state, the second law of thermodynamics requires that $T_{cav} = T_{obj}$, which implies that $P_{cav} = P_{obj}$ (since the blackbody power depends only on temperature). Under this condition, the expression in Equation 5.37 can only be zero when

$$\frac{\alpha_{cav}}{\epsilon_{cav}} = \frac{\alpha_{obj}}{\epsilon_{obj}}. \quad (5.38)$$

Since the absorptive properties of the two surfaces can vary independently with respect to one another (same for the emissive properties), Equation 5.38 can only hold if both ratios are equal to a constant. Given that the absorptivity and emissivity are maximal at unity, it follows that

$$\frac{\alpha_{cav}}{\epsilon_{cav}} = \frac{\alpha_{obj}}{\epsilon_{obj}} = 1. \quad (5.39)$$

This is Kirchhoff's radiation law.

Now suppose that the object inside the cavity in Figure 5.19 is a temperature sensor. In steady state, the power absorbed by the sensor is equal to the power emitted,

$$\alpha_{obj}P_{in} = \epsilon_{obj}P_{obj}. \quad (5.40)$$

Since $P_{obj} \propto T_{obj}^4$ and P_{in} is proportional to the integral of the incident intensity per unit frequency,

$$T_{obj}^4 \propto \frac{\alpha_{obj}}{\epsilon_{obj}} \int_0^\infty \mathcal{I}_{in} d\omega. \quad (5.41)$$

Note that this expression resembles a relation that Section 5.8.3 relied upon, which assumed the proportionality $T_{obj}^4 \propto \int_0^\infty \mathcal{I}_{in} d\omega$. As long as the proportionality factor in $T_{obj}^4 \propto \int_0^\infty \mathcal{I}_{in} d\omega$ is a constant, the analysis of Section 5.8.3 is valid. Equation 5.41 shows why Section 5.8.3 had to assume the sensor to be a perfect blackbody in order to use $T_{obj}^4 \propto \int_0^\infty \mathcal{I}_{in} d\omega$. Although the proportionality factor in Equation 5.41 is constant, α_{obj} and ϵ_{obj} are generally variable (as a function of temperature, for instance); however, for a perfect blackbody, $\alpha_{obj} = \epsilon_{obj} = 1$, and Equation 5.41 reduces to the desired relationship. As I have shown in this Section, the difficult-to-realize case of a perfect blackbody is not required for accurate radiation thermometry since Equation 5.41 also reduces to $T_{obj}^4 \propto \int_0^\infty \mathcal{I}_{in} d\omega$ when Kirchhoff's law holds.

This simple development of Kirchhoff's law is for equilibrium that is global to the entire sensor-cavity system. In practice, the Sr2 vacuum chamber is near global equilibrium but has a 1 K temperature variation across it. In this situation, Kirchhoff's law for global equilibrium holds approximately; however, the error in this approximate treatment might not be small enough for the high temperature accuracy needed for clock radiation thermometry. Fortunately *local thermodynamic equilibrium* ensures that the sensor surface properties do not effect the measurement [53]. If an object is in local thermodynamic equilibrium, each point on its surface absorbs and radiates like a body at a well-defined thermodynamic equilibrium temperature; however, the temperature across this object is not necessarily uniform. Local thermodynamic equilibrium exists in the presence of small temperature variations in time and across the body of the object under consideration.

The criteria for local equilibrium are well met by our system. Temperature variations in time are slow and small (shown later in Figure 5.27), and variations in space occur over length scales much larger than the sensor dimensions. In local equilibrium, Kirchhoff's law applies at each point on the sensor surface [53], meaning that imperfect absorption and emission do not affect temperature measurements.

5.8.5 First Evaluation: Silicon Diode Radiation Thermometers

For the first evaluation, the radiation thermometers are silicon bandgap temperature sensors from Lake Shore Cryotronics (part DT-471-SD) [13]. These diode sensors are useful for their high signal-to-noise ratio and low sensor self-heating properties. Two diode sensors, which have factory-specified uncertainties of 16 mK, are used for this work. The sensors are coated black with vacuum-compatible paint (VacKote 48816, Ball Corp) to improve their absorptivities and emissivities. Each sensor is held near the center of the chamber with a borosilicate glass tube (7 mm diameter, 1 mm glass thickness) that has a mounting tip at the end (Figure 5.20a). The glass tube is useful because it is a poor thermal conductor (conductivity of $\sim 1 \text{ W/mK}$) that is compatible with ultrahigh vacuum. Electrical feedthroughs allow us to read the sensor from outside the vacuum. One sensor is fixed at 2.5 cm from the center of the chamber to prevent it from being coated by the strontium beam. The other sensor can move between the center of the chamber or 2.5 cm away from the center using an edge-welded bellows (Figure 5.20b). This allows the movable sensor to measure the temperature at the center of the chamber where the atoms are located. Of course, the clock cannot be running with the movable sensor in this position, so the sensor is retracted by 2.5 cm during clock operation. The mounting structure for the fixed sensor is 8.1 cm from the sensor to the conflat base, and for the movable sensor this dimension is 14.4 cm. The fixed sensor is also referred to as the “short sensor,” and the movable sensor is also called the “long sensor.”

Prior to installation in the main chamber, the sensors were put in a test chamber. This test chamber was baked at 200°C (limited by the melting temperature of silver solder) and pumped down to the low 10^{-10} torr level to ensure that contaminants from the black coating and sensors

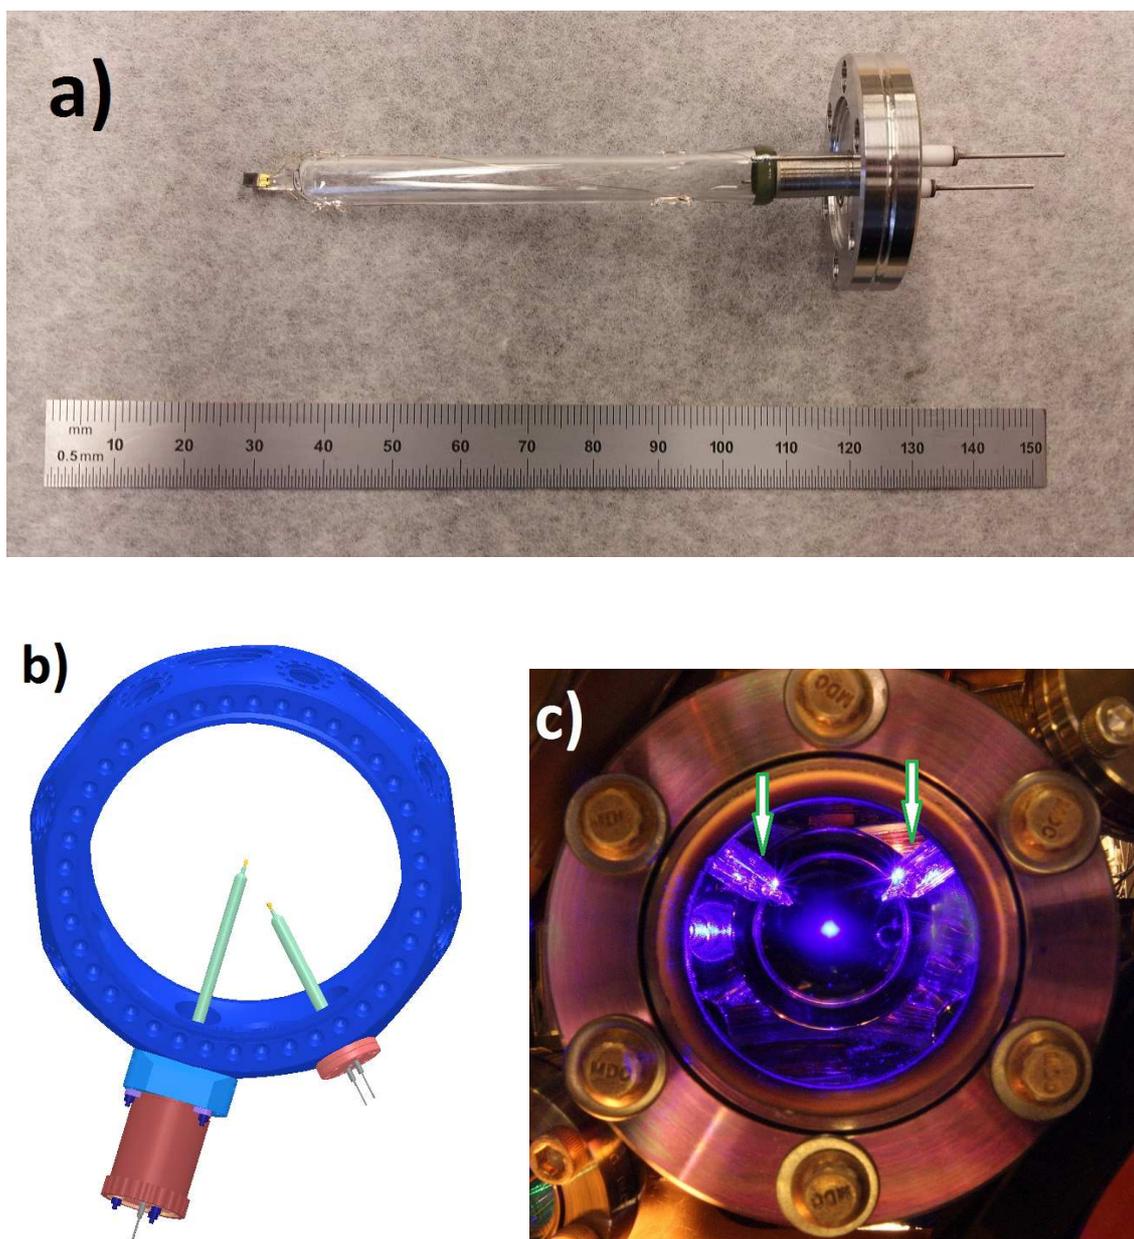


Figure 5.20: a) The shorter, fixed sensor. The black-coated sensor is at the tip of the glass tube. A glass-to-metal seal joins the glass tube to a shorter steel tube, which is welded to a 1.33" vacuum conflat. The conflat has a hole in the center (which cannot be seen in the photo) that allows access to the inside of the hollow glass and metal tubes. Since the insides of the tubes are at atmosphere, cartridge heaters for vacuum baking the sensors can be inserted in the tubes after the sensor mounts are installed in the main chamber. The leads of the sensor are connected to two electrical feedthroughs in the conflat for temperature measurements outside the vacuum chamber. The hook-up wire is phosphor bronze for low thermal conductivity, and this wire is joined with the sensor leads and electrical feedthroughs using silver solder. b) A CAD of both sensors in the chamber. The movable sensor is positioned near the center of the chamber. The mechanism used to move this sensor in and out is depicted at its base. c) The sensors installed in the vacuum chamber. The image is taken through one of the main chamber 2.75" viewports. The arrows point to the two glass tubes that hold the sensors. The glowing spot in the center of the chamber is the 461 nm MOT.

were removed. Surprisingly, pumping the sensors to this pressure took several weeks. At the time we attributed this to the black coating retaining water vapor. To transfer the sensors to the Sr2 system, the main chamber was vented to clean nitrogen at a pressure a few PSI greater than that of the Boulder atmosphere. With this positive main chamber pressure, removing two 1.33” conflat viewports from the main chamber to make space for the sensor does not significantly contaminate the chamber with water vapor. The sensors were transferred from the test chamber to the main chamber in air. Once installed, the nitrogen was shut off and the chamber was pumped out. The chamber did not pump down well at first until we added cartridge heaters (Figure 5.20a caption) inside the sensor glass tubes to bake the mounting structures.³⁴ Within a few days, the main chamber pressure had returned to normal. We refer to the technique of venting the vacuum chamber to clean nitrogen to provide positive pressure so that the chamber can be opened safely as a “nitrogen backflow.”

The diode is measured with electronics from Lake Shore. These electronics drive the diodes with a constant $10\mu\text{A}$ current and measure the resulting diode drop. Lake Shore provides the fit parameters (for a high-order fit function) used to calibrate the diode drop to a temperature standard. The Lake Shore circuit is queried by the Sr2 control software, recording temperature data each time the clock transition is measured. For clock operation, the BBR shift can be removed from each frequency measurement in post processing using this temperature record and Equation 5.27.

5.8.6 First Evaluation: Non-Thermal Heat Bias

There are several sources of systematic uncertainty that must be mitigated to ensure accurate radiation thermometry. The phenomenon with the greatest potential to bias temperature measurements is non-thermal heat.³⁵ The sensors are factory calibrated to a thermal heat distribution (realized inside a highly isothermal cavity), so non-thermal heat will shift the clock transition in

³⁴ These were baked at about 100°C .

³⁵ “Non-thermal heat” is heat that occurs in an out-of-equilibrium environment; therefore, it does not have a blackbody spectrum.

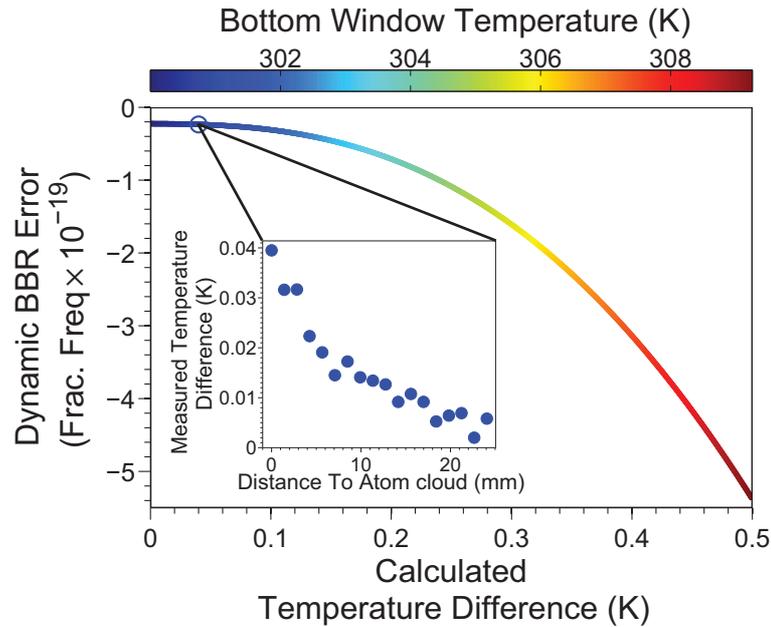


Figure 5.21: Simulating the non-thermal heat bias. The bottom 8” viewport temperature is raised above the rest of the chamber, which is fixed at 301 K. The non-thermal heat correction is then computed. The inset is a measurement of the chamber gradient using the movable thermometer.

a manner that the sensors will not predict.³⁶ Temperature gradients are the signature of non-thermal heat, so it is important to identify gradients and reduce temperature nonuniformity around the Sr2 vacuum chamber.

Section 5.8.3 argued that variations in temperature of 1 K across the vacuum chamber cause a negligible non-thermal heat bias. We also confirm this with a more elaborate analysis. We simulate the heat spectrum from the Sr2 chamber and viewports, assuming local thermodynamic equilibrium and allowing for temperature to vary across the chamber [13]. The emissivities of the chamber metal and viewports are included. To bound the non-thermal bias, we check an extreme case in which the temperature of the entire 8” bottom viewport is raised above that of the rest of the chamber. We find that for a 1 K variation across the chamber, the non-thermal bias is below

³⁶ The BBR shift is caused by the small part of a heat distribution that interacts with atomic transitions out of the clock states. Much of a room temperature blackbody distribution is too far detuned from any resonance involving one of the Sr clock states to have a significant effect on the BBR shift. On the other hand, the sensors measure the integrated heat spectrum. Using the integrated spectrum to predict the shift relies on the heat distribution being well known, which is the case in thermodynamic equilibrium. This is rarely the case for non-thermal heat; however, as I show in this Section and Section 5.8.3, the non-thermal heat bias in our system is well below the total systematic uncertainty of Sr2.

the 10^{-19} level (Figure 5.21). This conclusion holds even if the material emissivities are changed by 20%.³⁷ We also measure the gradient inside the chamber using the movable sensor, and we observe a 40 mK change over 2.5 cm (Figure 5.21 inset). Note that as long as imperfect sensor absorptivity and emissivity do not affect the measurement (Section 5.8.4), non-thermal heat only changes the dynamic BBR shift. This is because when we can ignore imperfections in heat transfer at the sensor surface, the sensor provides the total integrated heat, which is a direct measurement of the static shift (Equation 5.26) even in somewhat non-thermal conditions.

There are two sources of heat that require special attention: the Zeeman slower window and the oven. The Zeeman slower window (Section 3.2.2) is heated to 150°C and has line of sight to the atoms. Also, although the oven region is long, the oven is hot and has line of sight to the atoms. We have confirmed that the clock can operate with the Zeeman slower window cooled to room temperature for several hours, so long as the window is heated when the experiment is off; therefore, the Zeeman slower window heater is not used during clock operation. Also, the oven region has a rotary feedthrough that is used to block the atomic beam. The feedthrough is motorized, which allows us to block the oven after atoms are cooled and ready for spectroscopy, preventing oven heat from reaching the atoms.

Measuring the distribution of temperatures around the vacuum chamber, we observe that 1 K is indeed the temperature variation. This means that non-thermal heat bias can comfortably be ignored.

5.8.7 First Evaluation: Immersion Error

Since we want the sensors to only measure the radiative temperature inside the chamber, temperature measurements will be biased by any heat that conducts from the chamber along the glass mounting tubes (Section 5.8.5) and into one of the sensors. This is known as *immersion error* [122]. Immersion error is mitigated by the fact that the glass tubes emit and absorb radiation,

³⁷ Since the metal emissivity is 4 orders of magnitude smaller than that of the viewports, the emissivities can actually change by a factor of 20 and the simulation will not be affected; however, 20% is a much likelier margin of error for the emissivity values.

which attenuates heat that conducts along the length of the tubes.

To quantify the immersion error, we heat the conflat base of the fixed (shorter) sensor about 8 K above room temperature. The response of the sensor is minor and occurs after 2 to 3 minutes. We also measure the radiative response of the sensor with a heat gun on its coolest setting. Pointing the gun such that emitted heat passes through both the top vacuum chamber viewport and the center of the chamber, the sensor responds within seconds. This means that the sensor equilibrates with the radiative environment orders of magnitude faster than with the vacuum chamber through heat conduction, further reducing immersion error.

To bound uncertainty due to immersion error, we use the data from heating the sensor conflat base. We extrapolate these data to the case of heating the flange by normal temperature variations of about 1 K. Treating this error entirely as uncertainty, 0.7 mK uncertainty is added to that of the sensor [13].

5.8.8 First Evaluation: Passive Temperature Control

The Sr2 chamber is surrounded by a black enclosure (Figure 5.22) [13]. This enclosure shields the atoms from radiation emitted by warm objects throughout the lab.³⁸ The laser light required for the experiment is delivered into the enclosure via optical fibers, meaning that there are no holes in the enclosure to admit laser beams. Without holes, there is no way for stray thermal radiation from outside of the enclosure to enter the chamber.

To ensure that the system is robust against the usual laboratory temperature changes, temperature control of the space around the vacuum chamber is also attractive. This prevents temporal and spatial variations in the laboratory temperature from coupling too strongly to the chamber. Since we only require the temperature to be stable at the 1 K level (out of few-K fluctuations), passive temperature stabilization will be sufficient. The black enclosure around the chamber is also insulating, providing the temperature control we need. For the first evaluation, measurements with

³⁸ The chamber viewports have emissivities of 0.8-0.9 for IR wavelengths corresponding to thermal radiation, so the absorptive properties of the viewports also help with stray heat.

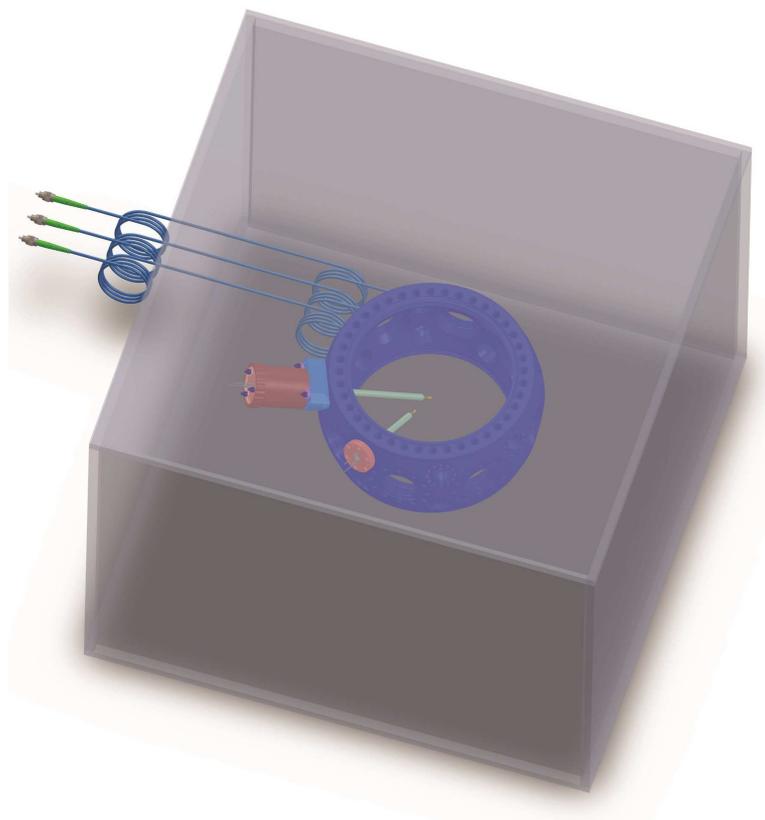


Figure 5.22: A cartoon of the black enclosure around the vacuum chamber. Light is delivered into this enclosure via optical fibers.

thermistors inside the enclosure show that the distribution of temperatures around the chamber is consistently around 1 K. Also, records from the diode sensors show that temporal fluctuations of the temperature inside the box are at the few hundred mK level even when the laboratory temperature changes by a few degrees K. This enclosure ensures that the temperature of the chamber does not vary in a markedly inhomogenous way, which further keeps the non-thermal error under control.

5.8.9 First Evaluation: Other Sources of Uncertainty

Low sensor self-heating is one major advantage to working with diode sensors compared to the widely used platinum resistance thermometers (PRTs). With PRTs, there is a competition between signal-to-noise and sensor self-heating. PRTs are resistors, so more current will provide

Effect	Correction (mK)	Uncertainty (mK)
Factory calibration	0	16
Immersion error	0	0.7
Sensor translation	40	20
Wire resistance	7.7	1.5
Lattice heating	-15	7.5
Total	32.7	26.7

Table 5.2: Uncertainty budget for the diode sensors. All uncertainties are quoted at the 1σ level. “Sensor translation” is the temperature difference between the center of the chamber and 2.5 cm from the center.

more signal but also more self-heating, which compromises accuracy. For example, a $100\ \Omega$ PRT with an excitation current of 1 mA will cause small-but-significant self-heating and produce a 100 mV signal; however, diode sensors are excited with only $10\ \mu\text{A}$ of current, which leads to negligible self-heating, yet the signal is a $\sim 0.5\ \text{V}$ diode drop.

The mounting structures for the sensors (Figure 5.20) use 2-wire (rather than 4-wire) measurements. This means that the resistance of the phosphor bronze leads will affect the measurement of the diode voltage drop used to infer the temperature. Estimating this effect based on the lengths of the leads and the diode calibration, the sensor measurement needs to have 7.7 mK added.³⁹ A generous uncertainty of 1.5 mK has been added to this correction [13].

As explained in Section 5.8.6, measurements with the movable sensor show that there is a 40 mK temperature difference between sensor readings at the center of the chamber (where the atoms are located) and 2.5 cm away. Since the sensors are 2.5 cm from the chamber center during clock operation, a 40 mK correction must be added to the final temperature reading. A large uncertainty of 50% is assigned to this correction [13].

Two other systematic biases were checked. When the lattice is turned on, the fixed sensor heats up by 15 mK. We treat this as a correction and add another generous 50% uncertainty. We also checked to see if the sensor was affected by any of the Sr2 magnetic fields and did not see a change.

³⁹ The sensor diode drop decreases with increasing temperature, so extra voltage will cause the sensor to read a temperature that is too cool.

Combining all systematic uncertainties in the thermometry measurements, we obtain a final temperature uncertainty of $\delta T = 26.7$ mK (Table 5.2) [13]. This corresponds to a static BBR shift uncertainty of 1.8×10^{-18} and a total BBR shift uncertainty of 4.1×10^{-18} . Since this work, two other lattice clock teams achieved BBR shift uncertainty below the 10^{-17} level. The Tokyo Sr group realized cryogenic lattice clocks, resulting in a 9×10^{-19} BBR shift uncertainty [116], and the NIST Boulder Yb group achieved BBR shift uncertainty of 1.0×10^{-18} using a uniform, well-characterized thermal environment [2].

5.8.10 Second Evaluation: Platinum Resistance Radiation Thermometers

To improve radiation thermometry for the second evaluation, we began a collaboration with the Sensor Science Division of NIST (Gaithersburg campus). Our collaborators had their misgivings about diode temperature sensors. In the 1980s, NIST tested diode thermometers for NASA and found these devices to have poor reproducibility after thermal cycling (which is an issue for us since we want to bake our sensors). The problem is that contact resistances at the junction between the silicon and the metal leads can vary significantly with thermal cycling. In the 2000s, Lake Shore published data showing more favorable thermal cycling properties [25], and new sensor models since then are likely to be even better. However, our collaborators did not feel that there was sufficient data to verify good diode sensor performance after several thermal shocks.

Our collaborators recommended platinum resistance thermometers. PRTs have been well studied for long-term stability and performance during thermal cycling; therefore, we opted for thin-film PRTs [14], which have good vacuum properties and are insensitive to calibration shifts caused by mechanical impacts. Motivated by a demonstration of a thin-film PRT that shifts only 1 mK after several significant thermal shocks [21], we chose 100Ω model C416 PRTs from Heraeus Sensor Technology.⁴⁰ The Heraeus PRTs are not factory specified at uncertainties better than what we achieved with diode sensors, but they can be calibrated at NIST at the 10 mK level or

⁴⁰ The demonstration in Reference [21] was for a different Heraeus model that is no longer sold. After contacting Heraeus, we were told that C416 is the replacement part for that used in Reference [21].

below.

5.8.11 Second Evaluation: Sensor Prequalification

Although Reference [21] demonstrated PRTs that shift by 1 mK after thermal cycling, this high performance was only observed in 20% of their sensors (all the same model). Motivated by these observations, we perform prequalification tests on our Heraeus sensors to identify those with the best temperature properties. These tests involve cycling the sensors between 0 °C and 200 °C. Those that only change at the 1 mK level after several cycles are identified as the most robust.

The lower cycling temperature of 0 °C is realized in an ice bath. The tests begin by measuring the initial sensor resistances at 0 °C. After the initial measurement, the sensors are heated in an oven at 200 °C for an hour, cooled to room temperature for about 20 minutes, and then returned to the ice bath for an hour. During this hour in the bath, the sensors' resistances are checked every 30 seconds, and we record these resistances when the sensors have clearly equilibrated. Each sensor is subject to 5 cycles of heating and cooling.

Since we are looking for 1 mK-level variations over several hours of measurement time, the temperature at which the sensors' resistances are checked must be stable at this level throughout the measurement (otherwise poor temperature stability could be mistaken for calibration shifts). This is why we need an ice bath, which can be stable to about 1 mK in the right conditions [47]. The bath is kept in a dewar. It is based on frozen distilled water broken up with a clean ice shaver. The shaved ice is made into dense slush with more distilled water. The bath must be carefully monitored, draining it or repacking it with ice as necessary, to prevent it from warming up due to the ice melting. The sensors are mounted inside test tubes, plunged in the bath, and measured with phosphor bronze leads that feed through the lid of the dewar. After 5 rounds of thermal cycling, one sensor maintained ± 0.5 mK reproducibility and one was repeatable to ± 2.0 mK. The ± 0.5 mK sensor is chosen for the fixed sensor mount and the ± 2.0 mK sensor is picked for the movable mount.

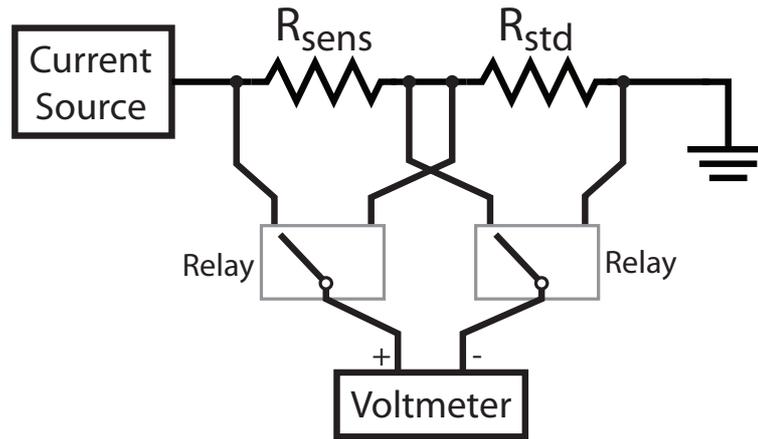


Figure 5.23: Electronics for measuring the PRT resistance. The sensor resistance R_{sens} is compared with a resistance standard R_{std} . A stable current drives the two resistors, and both voltage drops are recorded. Relays are used to switch between resistors. The measurement is repeated with the current reversed so that data can be processed in a manner that removes thermoelectric effects.

5.8.12 Second Evaluation: New Radiation Thermometry Hardware

To prevent lead resistance from affecting the thermometry, we upgrade to four-wire sensor measurements [88]. In this case, two pieces of phosphor bronze hookup wire are soldered to each sensor lead (one for sourcing/sinking current, one for measuring voltage). This requires new electronics (Figure 5.23) and new sensor mounting structures that each have four electrical feedthroughs. These electronics are based on comparing the PRT resistance to a $100\ \Omega$ commercial resistance standard (Fluke 742A-100).

The standard and the PRTs are wired in series, and a commercial low-noise current source (Keithley 6220) drives the sensors. The voltage drop across the PRTs and the standard is measured using an 8.5-digit voltmeter (Keithley 2002). The current direction is then reversed, and these voltage drops are measured again. For one PRT, the resistance inferred from these measurements is

$$R_{sens} = \frac{V_+ - V_-}{v_+ - v_-} R_{std}, \quad (5.42)$$

where V_+ (V_-) and v_+ (v_-) are the PRT and standard voltage drops measured for positive (neg-

ative) current. Also, R_{std} is the resistance of the standard (quoted at 1 ppm accuracy by the manufacturer). Recording the voltage drops with two currents (one in each direction) makes the measurement of R_{sens} insensitive to thermoelectric voltages, which do not change sign with the current. This method is also insensitive to any problems that would arise from the voltmeter having an out-of-date calibration, relying only on the resistance standard to provide accuracy.

5.8.13 Second Evaluation: Sensor Calibration at NIST

Since the sensors will ultimately be used in ultrahigh vacuum, their calibration must be performed under vacuum as well.⁴¹ We designed and built a vacuum chamber at JILA with the dimensions needed for the calibration facilities at NIST Gaithersburg (Figure 5.24). The sensors were installed (on their glass mounts) in the calibration chamber, which was then pumped down to a pressure similar to that of the Sr2 main vacuum. Like the diode sensors, the PRTs were painted black with VacKote 48816 to increase radiative coupling; however, the chamber pumped down in less than 24 hours this time, indicating that the difficulty obtaining good pressure with the silicon sensors was not due to the black paint.⁴² The chamber was then transported to NIST under vacuum. The safest and most cost effective option for transport was for me to hand carry the chamber on a commercial flight. We obtained an official Department of Commerce request for airport security to forgo x-ray scans of the chamber (which could damage the sensors) and perform chemical testing instead.

The NIST calibration is performed in a water bath [115]. The calibration chamber (Figure 5.24) is submerged in the bath, which is spatially uniform in temperature to 1 mK and has 1 mK temperature stability. Two standard PRTs, calibrated to the NIST ITS-90 temperature scale to 1 mK accuracy, are also submerged in the bath. This allows the Sr2 sensors to be calibrated with respect to the standard PRTs. PRT resistance as a function of temperature behaves according to

⁴¹ One reason for this is that sensor self-heating would be markedly different if the sensors were calibrated at atmospheric pressure. Convective heat transfer at the sensor-air interface mitigates self-heating, which will be worse under vacuum.

⁴² This may be because the Heraeus sensors are coated with a vacuum-compatible material, whereas the diodes have more structure that could trap air.

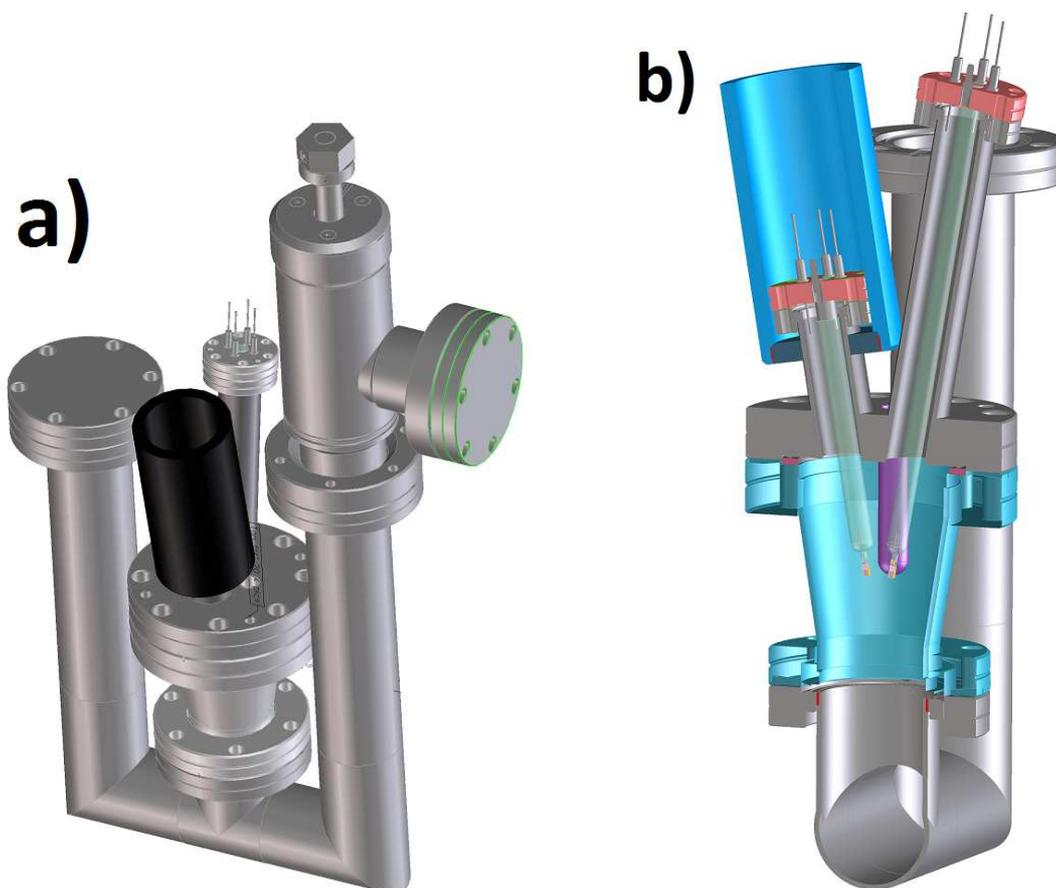


Figure 5.24: a) CAD of the chamber used for shipping. The chamber is pumped through the 2.75" conflat angle valve. The blanked-off arm of the chamber holds a nude ion gauge in the real setup. The 1.33" conflat is the base of the long sensor, and the sensor's four electrical feedthroughs can be seen here. Inside the black tube is the base of the shorter sensor. The tube, which is water tight, is present so that the chamber can be submerged (in the calibration bath) up to the long sensor base. b) A cross section of the chamber CAD (with false color for clarity). The 1.33" conflat sensor bases are in pink. The glass tubes that hold the sensors can be seen. The chamber is designed such that the two PRTs are situated close to each other (in the center of the cyan-colored reducer). The watertight tube is shown in cyan.

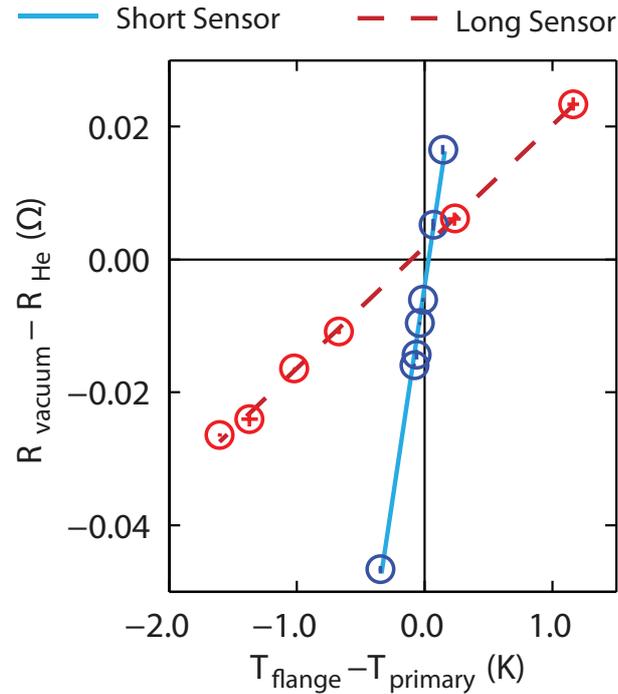


Figure 5.25: The calibration of the primary sensors. The primary sensors' resistances are measured with respect to the He calibration data as a function of the temperature difference across the mounting structures. The longer sensor is less sensitive to immersion error, as expected.

the Callendar van Dusen equation,

$$R = R_0(1 + AT + BT^2), \quad (5.43)$$

where T is in celsius (valid for $T \geq 0^\circ\text{C}$), and R_0 , A , and B are fit parameters for the calibration. For the Sr2 sensors, the manufacturer designed R_0 to be $100\ \Omega$. For platinum, $A = 3.9 \times 10^{-3}\ \text{C}^{-1}$, and B is on the order of $-1 \times 10^{-7}\ \text{C}^{-2}$.

Immersion error is a bigger problem for the NIST calibration than it was with the diode sensors in the Sr2 main chamber [88, 115]. This is because at NIST the bases of the sensors cannot be submerged in the bath (for the sake of electrical measurements), so they are coupled to the laboratory temperature. As the bath temperature is varied 5 to 10 $^\circ\text{C}$ above and below room temperature to calibrate the sensors, larger gradients than those present in the Sr2 main chamber will cause increased heat conduction through the mounting structures.

To fix this, a two-part strategy is employed [115].⁴³ First, a second pair of PRTs are calibrated, and each of these PRTs are affixed to the conflat flange at the base of one of the sensor mounting structures. With these “flange sensors,” the in-vacuum resistances of the “primary sensors” (AKA the Sr2 PRTs that will be used as radiation thermometers) are studied as a function of $T_{flange} - T_{primary}$, where T_{flange} is the temperature of the relevant mounting structure conflat as measured by one of the flange sensors [88]. Second, the primary sensors are calibrated in a manner that is free from immersion error. This is done using an exchange gas, which thermally couples the sensors and mounting structures more strongly to the walls of the chamber. High purity helium was selected as the exchange gas, and after venting the chamber to He, immersion error could no longer be detected.⁴⁴ Combining the data from both parts of this strategy, the immersion error for both sensor mounting structures can be understood (Figure 5.25).

The sensor resistances behave according to

$$R_{primary} = R_{He} + C_{IE}(T_{flange} - T_{primary}) + \Lambda \quad (5.44)$$

$$= R_0(1 + AT_{primary} + BT_{primary}^2) + C_{IE}(T_{flange} - T_{primary}) + \Lambda, \quad (5.45)$$

where the immersion error coefficient C_{IE} and offset Λ are free parameters in a linear fit of the data in Figure 5.25, and R_0 , A , and B come from the calibration of the primary sensors in the presence of the He exchange gas [88, 115]. To calculate the final temperature, this equation is solved for $T_{primary}$ as a function of the measured quantities $R_{primary}$ and T_{flange} . The flange sensors, which will need to be installed in the Sr2 system for this equation to hold, are useful because they provide real-time monitoring of the immersion error. It is worth noting the fit values of C_{IE} for the two sensors: $(0.1358 \pm 0.0075) \Omega K^{-1}$ for the short sensor and $(0.0183 \pm 0.0006) \Omega K^{-1}$ for the long

⁴³ My thanks goes out to Weston Tew for his excellent work on this problem and on the sensor calibration in general.

⁴⁴ Under vacuum, heat conduction along the glass tubes would cause the primary sensor readings to fluctuate with the lab temperature. Argon was tried as an exchange gas, but under Ar the sensors still displayed small fluctuations that tracked the laboratory temperature, indicating that immersion error was not suppressed. However, with a helium exchange gas, the primary sensors read steady values and only varied with the bath temperature. In this case, measurement precision was high enough that if few-degree laboratory temperature fluctuations could not be observed on the primary sensor with the short mounting structure, then immersion error is negligible (compared to other uncertainties) for the exchange gas measurements.

sensor. Despite that the long sensor mount is not even twice the length of the shorter mount, the immersion error for the shorter sensor is an order of magnitude larger. This is because immersion error decays exponentially with the tube length [122].

The final uncertainty of the sensors is 5.2 mK for the long sensor (Table 5.3) and 11 mK for the short sensor⁴⁵ [88]. Because propagated uncertainty in the solution to Equation 5.44 increases with the value of C_{IE} (Figure 5.27), the short sensor has larger total uncertainty than the long sensor; therefore, we use the longer sensor for determining the BBR shift. The short sensor will be used primarily to ensure that the two sensors agree after installation in the Sr2 chamber.

Two other phenomena were checked during the NIST calibration. First, the sensor self-heating is studied in vacuum and under He. This is done by measuring the sensor readout as a function of the driving current squared (which is proportional to the dissipated electrical power) and then extrapolating the trend to zero current. Self-heating data plotted as a function of squared current fits to a line, which can be used to extrapolate resistance measurements to zero self-heating [115]. For the data in Figure 5.25, all resistances are zero-heating extrapolated values. Second, the effect of vacuum pressure on the sensor readout is studied. In the presence of background gas, the sensor will not be sensitive to radiation alone since convection will be present. However, we observed that pressure effects are negligible below a vacuum pressure of 1×10^{-7} torr [115]. Since the Sr2 main chamber is at the 10^{-10} torr-level, pressure effects can be comfortably neglected.

5.8.14 Second Evaluation: Measurements in the Sr2 Chamber

After the NIST measurements, the calibration chamber with the sensors inside was hand carried (this time by Jun) back to JILA. We pumped the chamber down again, and the sensors were checked for agreement. They agreed within their combined uncertainties, indicating that transportation did not significantly affect calibration (Figure 5.26).⁴⁶ This is expected since thin-

⁴⁵ These uncertainties are computed for negligible gradients across the mounting structures. The total uncertainties of the sensors will be higher if significant gradients are present; however, significant gradients were not observed in the Sr2 main chamber.

⁴⁶ As our very knowledgeable collaborator Wes Tew (NIST Sensor Science Division) tells us, calibration shifts of thin-film PRTs are random. If multiple sensors of the same model are subject to the same calibration-shift-causing stress, the magnitudes and even the directions of their shifts will not be consistent. Therefore, agreement between

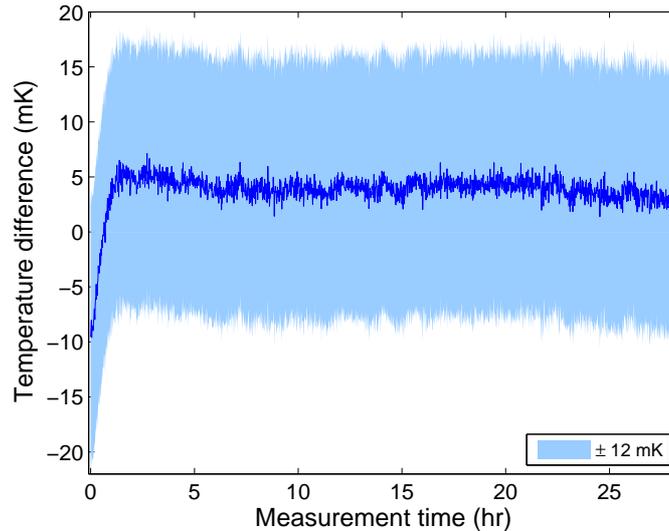


Figure 5.26: Difference between the two sensors after transport. This measurement is performed with the sensors in the calibration chamber. The data is represented in dark blue, and the light blue region is the combined total uncertainty of the two sensors assuming zero gradient. Flange sensor data has not been included. This data record is initialized right after shutting off the test chamber ion gauge, which is a heat source that warms the sensors relative to their conflat bases. With the sensors at markedly different temperatures than the bases, there will be heat conduction along the glass mounting tubes, and the sensors will disagree since their immersion error constants are so different. After an hour, the sensors have thermalized with the chamber, and sensor agreement is within the combined sensor uncertainty of 12 mK.

film PRTs are robust against calibration shifts caused by impacts, and for added security they were handled carefully throughout shipping and calibration.

The sensors are again installed in the Sr2 main chamber using a nitrogen blackflow technique (Section 5.8.5). For the second evaluation, greater care is taken to ensure that the environment inside the opaque chamber enclosure (Section 5.8.8) is uniform. In the first evaluation, one major heat source in the enclosure was the Andor CCD camera used to image the atoms. This is unnecessary for clock operation, so it is removed from the enclosure. Also, the 8 °C chilled water used to cool the MOT coils (Section 3.2.3) was a major heat sink in the first evaluation.⁴⁷ Replacing the two sensors indicates that calibration shifts did not occur at a significant level.

⁴⁷ Since water flows inside the MOT coils (Section 3.2.3), these coils are cooled very efficiently. Operating at full current (used in the blue MOT stage) for a few hours, the coils get hotter than room temperature. During clock operation the field is ramped, and the coils settle at a temperature lower than that of the laboratory. Therefore, before clock operation, we must turn the field ramps on and wait for an hour or two to let the inside of the enclosure thermalize.

the chilled water with 15 °C water balanced the heat from the coils much better, resulting in MOT coils very near room temperature for clock operation.

With better uniformity of the temperature inside of the enclosure around the vacuum chamber, we check sensor agreement to make sure that installation and baking at 100 °C did not affect the sensor calibration. This bake was similar to that of the diode sensors, in which cartridge heaters were used to bake the sensors while the main chamber was pumping down. Our checks confirm that sensor agreement is still within uncertainty (Figure 5.27a) and quite similar to the agreement observed in the calibration chamber after shipping (Figure 5.26) [88]. The temperatures measured by the two primary sensors show that the temperature inside the enclosure is stable to 140 mK (Figure 5.27b). Since the insides of the sensor mounting structures' glass tubes are exposed to atmosphere (Figure 5.20a caption), convection inside the tubes is suppressed by filling them with cotton insulation.

Although we use the flange sensors for thermometry measurements, we have found that they are unnecessary because gradients across the temperature mounting structures are so small (Figure 5.27c). The observed temperature difference between the long primary sensor and its flange sensor is (on average) -0.45 mK, which indicates a well-thermalized environment inside the enclosure. Even if the temperature difference across the long sensor mounting structure were as high as ± 1 °C, the increase in uncertainty (Figure 5.27d) would not significantly contribute to the final sensor accuracy; therefore, for the mounting structure temperature difference data shown in Figure 5.27c, the correction to the sensor uncertainty is below 1 μ K and can be comfortably ignored.

Lastly, the effect of translating the long sensor is studied. Since the translation mechanism is a hand crank that can heat the mounting structure, one must wait for an hour or so after the sensor is translated to allow the structure to thermalize with the ambient temperature. Therefore, the temperature in the enclosure will drift (Figure 5.27b) between data points measured at different positions of the long sensor. Fortunately, each of these points can be measured with respect to the short sensor temperature, which cancels ambient temperature drift. The temperature difference between the chamber center and 2.5 cm away (where the long sensor is positioned during clock

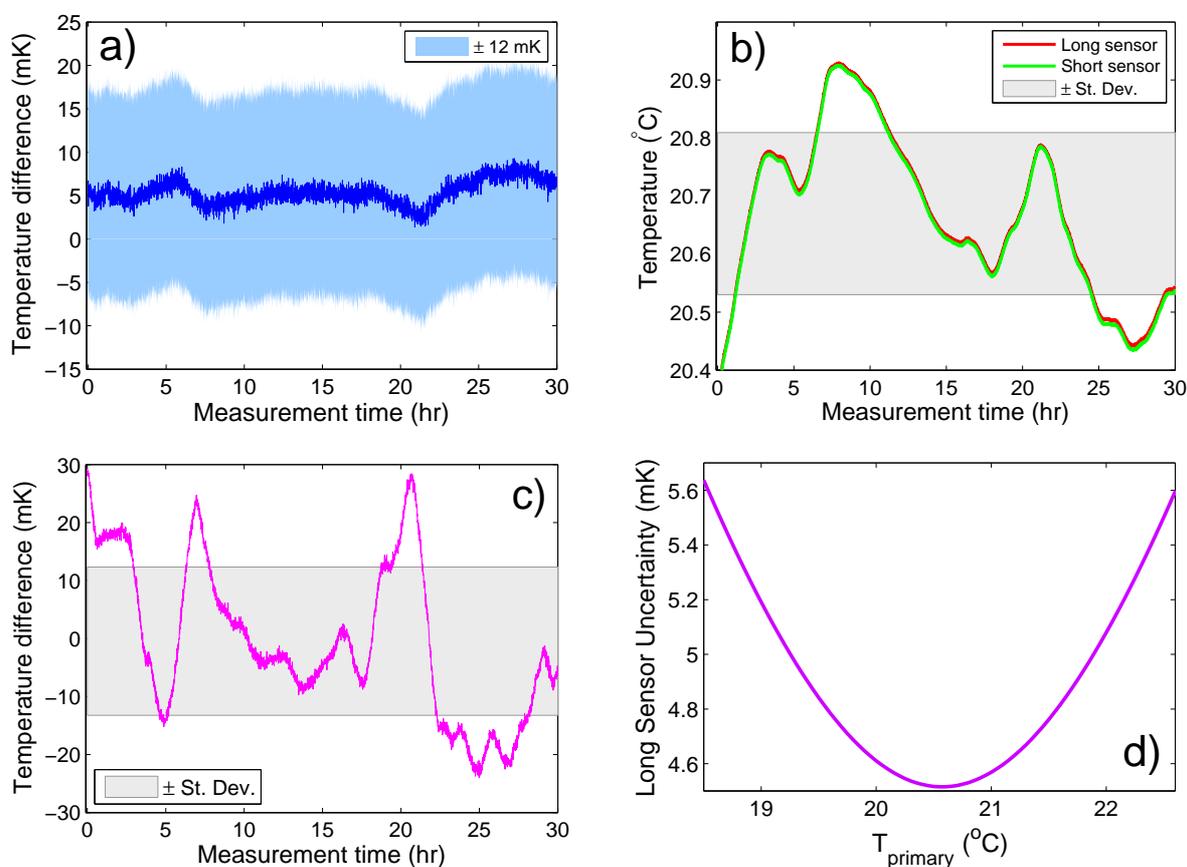


Figure 5.27: a) The difference between the long and short sensor after installation in the Sr2 main chamber. The data is in dark blue, and the light blue region is the combined uncertainty of the sensors. b) The temperatures measured by the sensors. The standard deviation of the data (shown as the grey band) reveals that the temperature fluctuates by 140 mK. The average temperature of this data set is 20.7 $^{\circ}\text{C}$, but in longer data records (not shown) it averages to 20.6 $^{\circ}\text{C}$. c) Measured temperature difference between the long primary sensor and its flange sensor. On average, the temperature difference is -0.45 mK, and the standard deviation of these data is 12.8 mK (shown as the gray band). d) Calculated total uncertainty of the long sensor, assuming $T_{\text{flange}} = 20.6^{\circ}\text{C}$. The total uncertainty is determined as the propagated uncertainty in the solution to Equation 5.44. This plot shows that temperature differences across the sensor mounting structures as large as $\pm 1^{\circ}\text{C}$ would not have a significant effect on the total clock uncertainty; therefore, the uncertainty accrued by the observed temperature difference of -0.45 ± 12.8 mK is orders of magnitude too small to be of any concern.

Effect	Uncertainty (mK)
Bath non-uniformity	1.0
Bath SPRT calibration	1.0
Bath temp. stability	1.0
Sensor self-heating	0.5
Electrical errors	0.07
Sensor translation	0.03
Thermal cycling	2.0
Calibration coefficients	4.5
Total	5.2

Table 5.3: Uncertainty budget for the primary sensor with the longer glass tube. All uncertainties are quoted at the 1σ level. Here “SPRT” is the standard PRT used in the calibration.

operation) is (1.45 ± 0.03) mK [88]. Again, this small gradient indicates that the chamber is well thermalized.

Combining all of the PRT sensor measurements at JILA and NIST, the total uncertainty of the long sensor is 5.2 mK (Table 5.3).⁴⁸ For the Sr2 operating temperature of 20.6°C , the static BBR shift is -4.5621×10^{-15} with an uncertainty of 3.48×10^{-19} [88].

5.9 Dynamic Blackbody Shift

5.9.1 A Scheme to Measure the Dynamic Coefficient

As explained in Section 5.8.2, dynamic BBR shift uncertainty is dominated by uncertainty in the coefficient ν_{dyn} . Furthermore, uncertainty in ν_{dyn} is dominated by insufficient accuracy in the Einstein A coefficients for transitions out of the clock states. The relationship between ν_{dyn} and the Einstein coefficients is given by Equation 5.29,

$$\nu_{dyn} = -\frac{4\pi^4}{21} \left(\frac{k_B T_0}{\hbar} \right)^6 \sum_k \left(\frac{A_{ek}}{\omega_{ek}^6} - \frac{A_{gk}}{\omega_{gk}^6} \right).$$

This expression is proportional to the curvature of $\Delta\alpha(\omega)$ (the differential polarizability of the clock

⁴⁸ Lattice light heating was not observed this time around. With the lattice at its operating power, no effect could be discerned from background temperature drift. This may be due to the reduction of scattered lattice light after realigning the lattice and cleaning up dust around the chamber. Scattered light must have been the cause of lattice light heating in the first evaluation because, at the position of the sensors, the Gaussian beam equation predicts that the trap light intensity is smaller than its maximum value by a factor of e^{48828} .

transition) at $\omega = 0$. Curvature in $\Delta\alpha(\omega)$ is caused by resonances, so ν_{dyn} is determined by the lowest-frequency transitions out of the clock states (as long as the corresponding resonance widths are substantial).

98.2% of the dynamic shift is caused by the ${}^3P_0 \rightarrow {}^3D_1$ transition at $2.603 \mu\text{m}$, and uncertainty in the Einstein A coefficient for this resonance dominates the uncertainty in ν_{dyn} [105]. In Yb, this same transition dominates ν_{dyn} (which is 6.6 times smaller than in Sr), and a scheme to measure the ${}^3P_0 \rightarrow {}^3D_1$ Einstein coefficient was demonstrated by the NIST Yb lattice clock team [3]. The NIST approach involves measuring the total decay rate of the 3D_1 state and then inferring the needed Einstein coefficient through the ${}^3D_1 \rightarrow {}^3P_0$ branching ratio. We use this same scheme to measure ν_{dyn} in Sr.

We did not need to measure this decay rate in our first evaluation. The ${}^3D_1 \rightarrow {}^3P_0$ Einstein coefficient had already been determined at the 1% level [82, 105], which results in 3.7×10^{-18} uncertainty in the dynamic BBR shift. This uncertainty is low enough for the first evaluation's goal of becoming the best clock, but for the second evaluation to reach low- 10^{-18} -level total clock uncertainty, the dynamic coefficient must be measured more accurately; therefore, we aim for a 0.5% measurement of ν_{dyn} . Since the ${}^3D_1 \rightarrow {}^3P_0$ branching ratio is known to 0.1%, and since a combined uncertainty of well below 0.1% arises from all terms in the ν_{dyn} expression aside from that due to the ${}^3D_1 \rightarrow {}^3P_0$ line, ν_{dyn} can be determined at our accuracy goal by measuring the 3D_1 decay rate to 0.5% [105].

The decay measurement begins by preparing atoms in the 1S_0 ground state, driving population into the 3P_0 state using a clock laser π -pulse, and then driving the ${}^3P_0 \rightarrow {}^3D_1$ transition with a $2.6 \mu\text{m}$ laser (Figure 5.28). The 3D_1 state decays into the 3P manifold. Of the 3P states, only 3P_1 has any measurable decay on experimental timescales. Photons emitted from ${}^3P_1 \rightarrow {}^1S_0$ decay are collected on a PMT, and the number of collected photons is recorded as a function of time.

Treating the five-level system depicted in Figure 5.28 with rate equations,

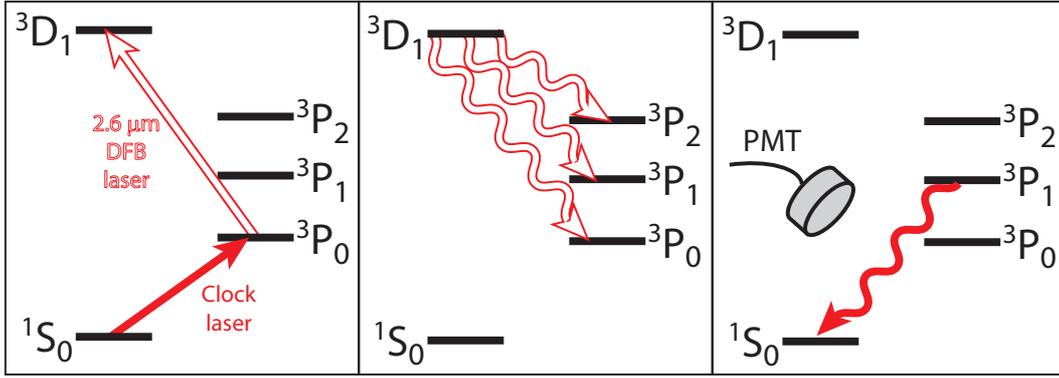


Figure 5.28: The decay measurement scheme. (Left) Population is prepared in the electronic ground state and then driven to 3D_1 using two laser pulses. (Center) The 3D_1 state decays into the 3P manifold. The branching ratios to 3P_2 , 3P_1 , and 3P_0 are 1.9%, 38.7%, and 59.5% (respectively). (Right) The population in the 3P_1 state then decays to the ground state. Photons emitted during this decay are collected with a photomultiplier tube.

$$\frac{dN_2}{dt} = -\frac{N_2}{\tau_{tot}} \quad (5.46)$$

$$\frac{dN_1}{dt} = \frac{N_2}{\tau_2} - \frac{N_1}{\tau_1}, \quad (5.47)$$

where N_1 (N_2) is the population in the 3P_1 (3D_1) state, τ_{tot} is the lifetime for 3D_1 to decay into the entire 3P manifold, τ_1 is the lifetime of the 3P_1 state, and τ_2 is the inverse of the decay rate for the ${}^3D_1 \rightarrow {}^3P_1$ pathway. Solving this,

$$N_1(t) = \frac{\tau_{tot}\tau_1}{(\tau_1 - \tau_{tot})\tau_2} N_0 \left(e^{-t/\tau_1} - e^{-t/\tau_{tot}} \right), \quad (5.48)$$

where the initial conditions are $N_1(0) = 0$ and $N_2(0) = N_0$ (assuming that the $2.6 \mu\text{m}$ laser pulse is infinitesimal). This is proportional to the number of photons emitted by the ${}^3P_1 \rightarrow {}^1S_0$ decay, which is proportional to the experimental signal (Figure 5.29).

5.9.2 Decay Measurement

One cycle of the 3D_1 decay measurement begins by laser cooling ${}^{87}\text{Sr}$ atoms and trapping them in the TA-based lattice. The atoms are spin polarized and driven into the 3P_0 state using a 3

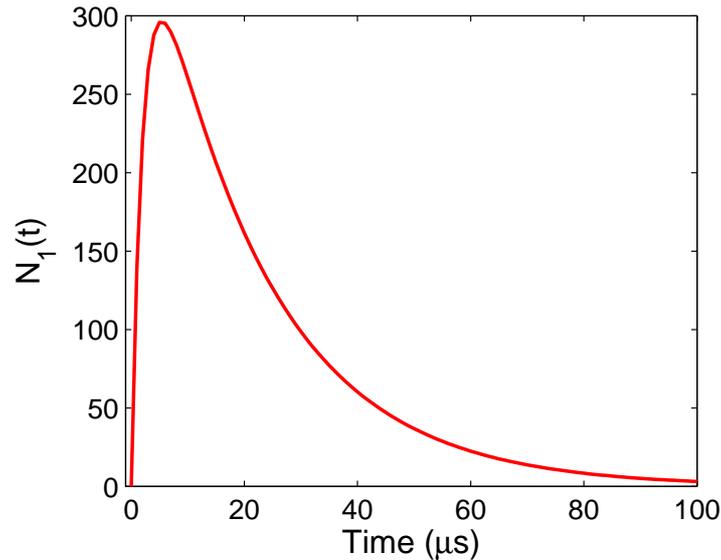


Figure 5.29: Equation 5.48 plotted as a function of time. Correct values have been used for the lifetimes $\tau_{tot} = 2.16 \mu\text{s}$, $\tau_1 = 21.28 \mu\text{s}$, and $\tau_2 = 5.58 \mu\text{s}$. An initial atom number of $N_0 = 1000$ has been assumed.

ms clock laser π -pulse. At this pulse time, the Fourier-limited clock transition linewidth is 270 Hz. This comparatively broad linewidth is useful because it makes the decay measurement less sensitive to clock laser drift, which is a greater concern (compared to other systematic measurements) since atomic servos are not used here. We also spin polarize the atoms because this allows us to drive population into the 3P_0 state with near unity efficiency.⁴⁹

With nearly all population in the 3P_0 state, the atoms are driven into the $|F = 11/2, ^3D_1\rangle$ state with a $2.6 \mu\text{m}$ DFB laser.⁵⁰ The DFB laser is from NanoPlus, and it is factory specified to have a 2603.1 nm center frequency, a 3 MHz linewidth, $\pm 1.5 \text{ nm}$ wavelength tunability, and 2 mW of output power. We observed output power closer to 4 mW. The frequency drift of the DFB laser is very low, and we can use it free running for an entire evening without having to adjust the frequency.⁵¹ The DFB laser is pulsed with a 90 MHz AOM.

⁴⁹ The peak excited state fraction for short π -pulses like these and a spin-polarized sample is typically 0.95. Recall that for unpolarized lines, the peak excited state fraction is 0.61 (Figure 2.14a caption).

⁵⁰ Prior to the decay measurement, we performed survey spectroscopy with an OPO laser and found the $|F = 9/2, ^3P_0\rangle \rightarrow |F = 11/2, ^3D_1\rangle$ line at 2603.128 nm. This is the Stark shifted resonance wavelength observed in the TA-based lattice.

⁵¹ The resonance frequency of the $|F = 9/2, ^3P_0\rangle \rightarrow |F = 11/2, ^3D_1\rangle$ transition does drift if the lattice intensity is

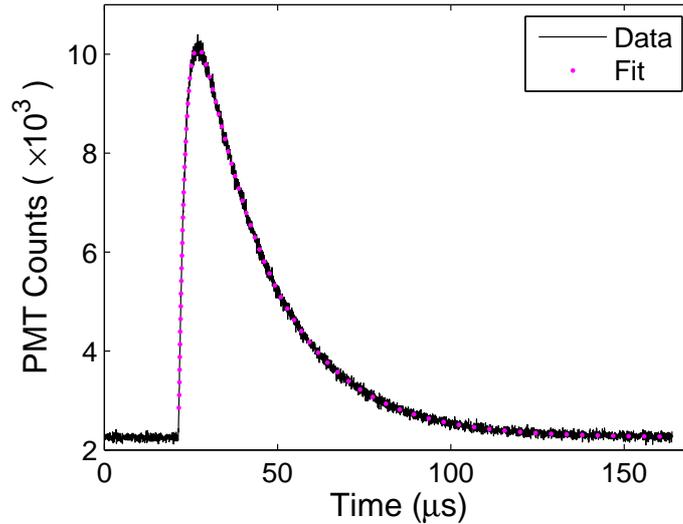


Figure 5.30: Decay data. Decay measurements are taken in 40 ns time bins. This figure is generated by summing all data for each time bin. The magenta dots are the fit of data taken after the pulse is off. The y-axis offset is due to dark counts on the PMT, and the x-axis offset is primarily a timing delay added to ensure that all the early data is captured.

$^3P_1 \rightarrow ^1S_0$ fluorescence resulting from the $^3D_1 \rightarrow ^3P_1 \rightarrow ^1S_0$ cascade is collected with a 3" diameter lens and focused onto a Hamamatsu PMT.⁵² To get rid of dark counts, the PMT is covered with two optical bandpass filters centered on 690 nm (the fluorescence is at 689.4 nm).⁵³ The output of the PMT is sent directly to an SR430 photon counter.

Since atoms decay back to the ground state, more decay measurements can be performed without having to prepare another ultracold sample; therefore, 50 decay measurements are made each time we cool and trap atoms. Decay back into the ground state depolarizes the atoms into other m_F states, so for each cooling and trapping cycle the spin polarization laser is applied for 5 ms in between decay measurements. Also, a bias field is turned on when the clock laser is pulsed, and then it is shut off before the DFB laser is applied. Keeping the field off during the $^3D_1 \rightarrow ^3P_1 \rightarrow ^1S_0$ decay prevents systematic bias due to Zeeman quantum beats [3].

The photon counter records the number of detected photons in time bins of 40 ns. Data is

not stabilized, resulting in a time varying ac Stark shift.

⁵² This is a second PMT in addition to the one used to measure the atom number and excited state fraction.

⁵³ We found that one optical filter was not enough to eliminate dark counts.

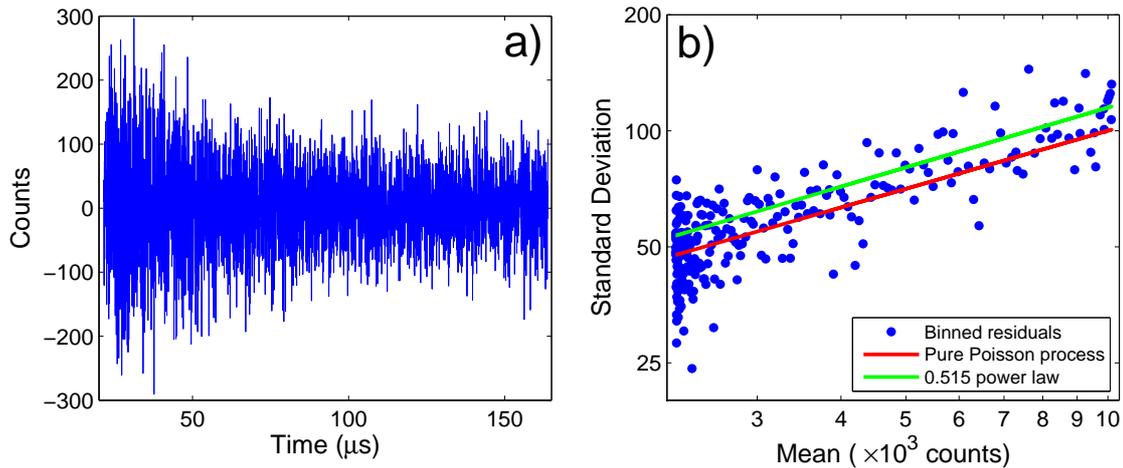


Figure 5.31: a) Fit residuals of the decay data. b) Comparing statistical moments of the residuals. For this plot, the residuals are binned (bin size of 15 points), and the mean and standard deviation is computed for each bin. Plotted on a log-log scale, a pure Poisson process would have a slope of 0.5. The trend is somewhere between a Poisson process and a power law with a 0.515 exponent.

stored in the SR430 local memory for speed and offloaded after several minutes. The decay data (Figure 5.30) was taken over one week in two shifts, one in the daytime and one overnight, for round-the-clock acquisition. The data set corresponds to about 8 million single decay measurements, and the trend in the data strongly resembles the rate equation solution shown in Figure 5.29 [88]. Motivated by the solution to the rate equations (Equation 5.48), the signal is fit with the function

$$y(t) = y_0 + A \left[e^{-(t-t_0)/\tau_{3P_1}} - e^{-(t-t_0)/\tau_{3D_1}} \right], \quad (5.49)$$

where y_0 , A , t_0 , τ_{3P_1} , and τ_{3D_1} are fit parameters. y_0 and t_0 are added due to the presence of technical offsets (Figure 5.30).

5.9.3 Fit Uncertainty

Binning performed by the photon counter introduces concerns about fit uncertainty. For each time bin, an offload of photon counter data provides the sum of all counts acquired since the previous offload. These sums have an averaging-like effect on the bins, smoothing statistical fluctuations. Therefore, in principle each bin should also have an error bar representing the statistical uncertainty

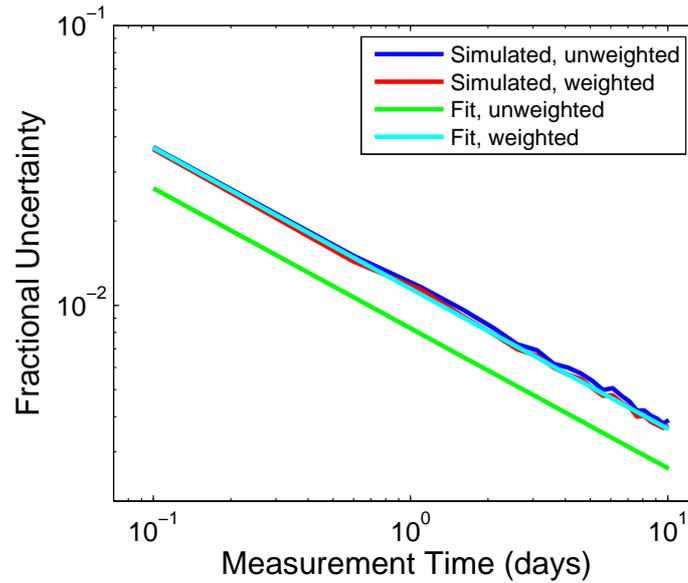


Figure 5.32: Simulated fit uncertainty in τ_{3D_1} . Here the experiment is simulated assuming Poissonian noise. For each simulation, a weighted and unweighted fit to Equation 5.49 are performed. The fit weights treat the statistical uncertainty of each simulated time bin as the square root of the bin mean. The blue and red lines represent the simulated uncertainties. To compute the simulated uncertainties, simulated data is generated and fit repeatedly, resulting in a distribution of fit values of τ_{3D_1} . For the blue (red) line, fit values are computed by an unweighted (weighted) fit. The simulated uncertainty is taken to be the standard deviation of this distribution, which describes the true scatter in the fit. In practice there is only one data set, so uncertainty in the fitted value of τ_{3D_1} must come from the uncertainty generated by the fit routine. If the fit routine provides the correct uncertainty, it will agree with the simulated uncertainty. The green and cyan lines are uncertainties generated by the fit routine, showing that an unweighted fit significantly disagrees with the simulated uncertainty whereas a weighted fit (cyan) gets the uncertainty correct.

in each sum, and a fit of the data to Equation 5.49 should be weighted by these error bars. The photon counter does not provide uncertainties, but it is possible to understand what they should be by studying the measurement noise process.

The noise is characterized through the fit residuals of the decay data (Figure 5.31a). Plotting the standard deviation of the residuals against the mean, we observe a trend consistent with a (nearly) Poisson process (Figure 5.31b). This implies that the measurement is dominated by photon counting noise, with no other noise sources significantly contributing. Since Poissonian noise has a standard deviation equal to the square root of its mean, the data fits should be weighted by the square root of the counts for each bin (Figure 5.32). Without weighting, the fit routine will

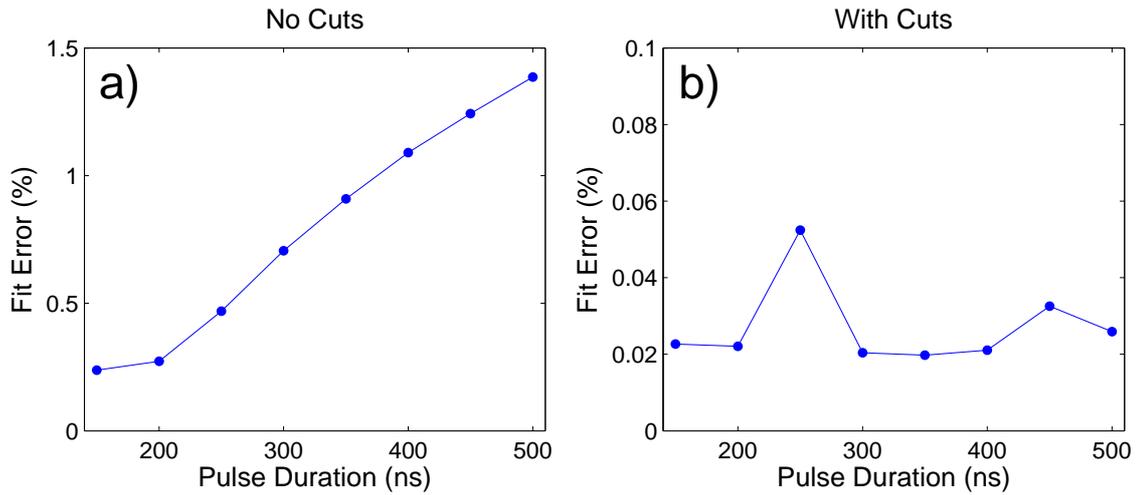


Figure 5.33: a) Theoretical fit bias due to finite pulses. This is based on numerically solving the Lindblad master equation for different pulse durations. The situation is treated as a four-state system of 1S_0 , 3P_0 , 3P_1 , and 3D_1 . 3P_2 is ignored because its branching ratio is small, and it does nothing except shelve population (including this state would merely increase the 3D_1 lifetime a bit, which would not affect the conclusion). A realistic Liouvillian is used in the calculation, assuming accurate values of state lifetimes and an estimate of the DFB laser linewidth based on factory specifications. The master equation solution provides numerical data for the system density matrix, which yields the 3P_1 population (proportional to the experimental signal) as a function of time. This population is fit to Equation 5.49, and the error is computed as $|\tau_{fit} - \tau_{true}|/\tau_{true}$, where τ_{fit} is the fitted value of τ_{3D_1} and τ_{true} is the true value used in the master equation. The error is nearly equal to the measurement accuracy goal of 0.5% for pulses of 250 ns. b) The same master equation data with points cut from the fit. Here the fit only considers data after the DFB pulse is off, which reduces error to a very comfortable level.

underestimate τ_{3D_1} uncertainty.

Fitting the data in this way, the measured 3D_1 lifetime is $\tau_{3D_1} = (2.18 \pm 0.01) \mu\text{s}$, which is slightly better than the 0.5% measurement goal [88]. The fit also yields a new best value for the 3P_1 lifetime (which is of interest to other research teams) of $\tau_{3P_1} = (21.28 \pm 0.03) \mu\text{s}$

5.9.4 Systematic Bias Due to the Finite DFB Laser Pulse

The function used to fit the data (Equation 5.49) is motivated by the solution to the rate equations (Equation 5.48). The initial conditions for the rate equations assume an instantaneous DFB laser pulse.⁵⁴ To check if the length of the pulse has any effect, I numerically solve the

⁵⁴ Equation 5.48 resulted from assuming the initial condition of all population in the 3D_1 state. This implies an instantaneous DFB laser pulse because for a finite pulse, population would leak into 3P_1 while the DFB pulse is on.

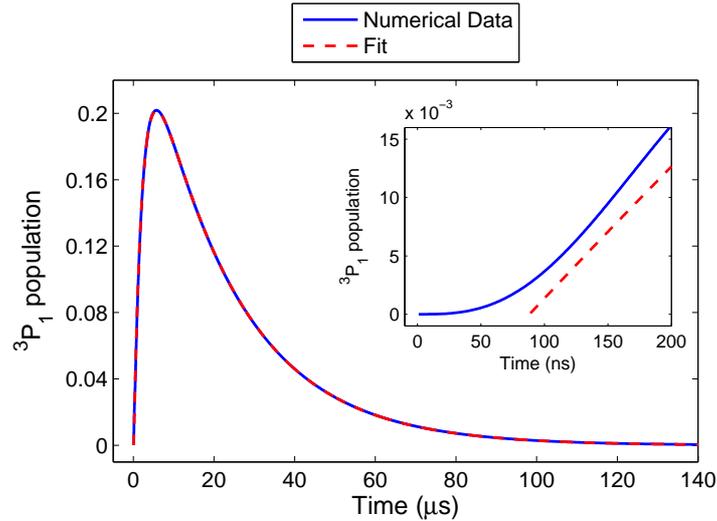


Figure 5.34: The master equation solution for a 500 ns pulse fit with Equation 5.49. Inset: Although the instantaneous pulse model appears to fit the data well, closer inspection at early times reveals disagreement. The non-instantaneous system response at time = 0 (due to dynamics while the pulse is on) causes the master equation solution to be qualitatively different than the fit function at short times. This causes a fit bias of nearly 1.5% (Figure 5.33), which is unacceptable for our accuracy goal of 0.5%.

Lindblad master equation for a variety of pulse durations (Figure 5.33a). This calculation uses realistic values for state lifetimes and the DFB laser linewidth. The numerical master equation data is fit using the function from Equation 5.49 to extract the 3D_1 lifetime, τ_{3D_1} . The error, which is the disagreement between the true lifetime value used in the master equation and the value obtained by the fit, is higher than the measurement accuracy goal of 0.5% for pulses longer than 250 ns (Figure 5.33a). The error is caused because the signal edge when the DFB laser is turned on gets smoothed (with respect to the instantaneous DFB pulse model) due to dynamics occurring during the DFB pulse (Figure 5.34).

Treating the decay data while the pulse is on would require a good model of the real DFB pulse shape and linewidth, and we would also need confirmation that these quantities do not vary in time. However, the system still admits an easy solution after the DFB pulse switches off.⁵⁵ The

⁵⁵ My thanks to Murray Barrett for being the first in our team to work out the analytic treatment of the population decay for finite pulses.

rate Equations 5.46 and 5.47 still hold in this case, but the initial conditions are now $N_1(0) = N_{01}$ and $N_2(0) = N_{02}$, where N_{01} (N_{02}) is the atom number in 3P_1 (3D_1) when the pulse is extinguished. The difference between this case and that used to derive Equation 5.48 is the number of atoms N_{01} that accumulate in 3P_1 when the DFB laser is on. Here $t = 0$ is taken to be the instant when the DFB pulse reaches zero intensity. The solution is

$$N_1(t) = \left(N_{01} + N_{02} \frac{\tau_{tot}\tau_1}{\tau_2(\tau_1 - \tau_{tot})} \right) e^{-t/\tau_1} - N_{02} \frac{\tau_{tot}\tau_1}{\tau_2(\tau_1 - \tau_{tot})} e^{-t/\tau_{tot}}. \quad (5.50)$$

Defining constants A and t_0 through the equations $Ae^{t_0/\tau_1} = N_{01} + N_{02} \frac{\tau_{tot}\tau_1}{\tau_2(\tau_1 - \tau_{tot})}$ and $Ae^{t_0/\tau_{tot}} = N_{02} \frac{\tau_{tot}\tau_1}{\tau_2(\tau_1 - \tau_{tot})}$, it is apparent that Equation 5.50 is faithfully modeled by the fit function that assumes an instantaneous pulse (Equation 5.49). Therefore, as long as we only fit data taken after the pulse switches off, the fit function will provide the correct 3D_1 lifetime. This conclusion is independently verified by fitting master equation data for times after the pulse is off (Figure 5.33b).

As the master equation analysis shows, two strategies can be used to reduce fit bias due to the finite pulse duration: keep the DFB pulses at 200 ns or shorter, and only fit data after the DFB pulse is off. To doubly suppress this bias, we take both approaches. We confirm the amount of data that must be cut from the fit by comparing (on an oscilloscope) the 200 ns DFB pulses on a high-speed photodetector to the electronic pulse used to trigger the photon counter.⁵⁶ Based on the master equation analysis, we assign a conservative upper bound of 0.1 ns on the systematic uncertainty in τ_{3D_1} due to the remaining finite pulse bias.

5.9.5 Fit Bias Due to Stray Electromagnetic Radiation

Stray electromagnetic radiation resonant with the $^3P_0 \rightarrow ^3D_1$ transition will affect decay dynamics and bias the τ_{3D_1} measurement. The two sources of stray radiation at a 2.6 μm wavelength are the heat in the Sr2 laboratory and stray beams from the DFB laser. The bias due to laboratory heat is extremely small, and it can be estimated using the Einstein approach to radiative processes

⁵⁶ The fit of the data in Figure 5.30 is plotted only for the range of data included in the fit. The fact that the fitted curve is absent right when the decay process begins is because early data is not included in the fit due to the issues described in this subsection.

[117]. Laboratory heat will cause unaccounted for absorption and stimulated emission on ${}^3P_0 \rightarrow {}^3D_1$, meaning that the 3D_1 decay rate will be

$$\frac{1}{\tau_{3D_1}} + B \int_0^\infty \rho(f)g(f)df, \quad (5.51)$$

for spontaneous decay and stimulated emission. Here B is the stimulated emission Einstein coefficient, $\rho(f)$ is the energy density of blackbody radiation from the lab, and $g(f)$ is the lineshape function of the transition. Since $g(f)$ is orders of magnitude narrower than $\rho(f)$, $g(f)$ can be treated as a delta function $\delta(f - f_0)$ in the integral, where f_0 is the ${}^3P_0 \rightarrow {}^3D_1$ resonance frequency. Using the relations between Einstein coefficients, the fractional systematic uncertainty in τ_{3D_1} is simply $1/(e^{hf_0/k_B T} - 1)$, which is the average number of blackbody photons at frequency f_0 . Although this calculation neglects state degeneracy (which modifies the answer by a factor of order unity), it also results in 10 ppb uncertainty, so this analysis could be off by a factor of 100,000 and the uncertainty would still be negligible. A similar treatment for absorption results in another very small correction. Therefore, we assign a conservative upper bound of 10 fs on the systematic uncertainty in τ_{3D_1} due to contamination from laboratory heat.

To study the uncertainty due to stray light, we set the AOM (used to pulse the DFB laser) in its off state and scan the DFB frequency across the ${}^3P_0 \rightarrow {}^3D_1$ line. We expose the atoms to residual light that may still be in the AOM diffracted order for three orders of magnitude more time than the experimental pulse duration. No excitation of the ${}^3P_0 \rightarrow {}^3D_1$ resonance is observed.⁵⁷

Treating the background signal fluctuation in this measurement as the upper bound on 3D_1 state excitation, another master equation calculation is used set a bound on the effect of residual light from the AOM diffracted order used in the decay data. The results show a negligible effect.

Although residual light in the AOM diffracted order is not significant, we do observe a stray light effect when the DFB laser frequency is such that the AOM zeroth order is resonant with the ${}^3P_0 \rightarrow {}^3D_1$ transition. This effect is observed for an artificially long exposure time of 200 ms, and it

⁵⁷ Excitation is measured by driving the clock transition, applying the DFB laser, and then measuring the clock transition excited state fraction in the usual way. Excitation into 3D_1 will diminish the excited state fraction.

appears to be coming from the AOM zeroth order beam (rather than from the DFB beam before it enters the AOM) because the effect disappears when we close a mechanical shutter installed after the AOM. After measuring the population we can excite for 200 ms of exposure to the stray AOM zeroth order light, the master equation is used to put an upper bound on the Rabi frequency from this light. From this information, the bias due to stray light in normal experimental conditions is inferred using the master equation.⁵⁸ Again, the effect is negligible. The small size of this effect is due to both the low intensity of the stray light and also because during the decay measurement, the DFB frequency is such that the AOM zeroth order is detuned by 90 MHz from the $^3P_0 \rightarrow ^3D_1$ resonance.

These studies of the effects of stray DFB light allow us to put an upper bound of 10 fs on fit bias due to stray DFB beams.

5.9.6 Density-Dependent Effects

During their measurement of the 3D_1 lifetime in Yb, the NIST Boulder lattice clock team observed a density dependence to fitted values of τ_{3D_1} [3]. Effects like superradiance and radiation trapping could cause the 3D_1 lifetime to vary with density. Uncertainty in extrapolating the measured lifetime to zero density was the dominant limitation to the precision of the NIST measurement [3].

The Sr2 experimental sequence used to measure decay data also records the atom counts (with the 461 nm counting laser and the PMT described in Section 3.7.2). Each cycle of the sequence involves preparing the atoms in the lattice, applying the clock and DFB lasers, measuring fluorescence from the $^3D_1 \rightarrow ^3P_1 \rightarrow ^1S_0$ cascade 50 times, and then fluorescing the $^1S_0 \rightarrow ^1P_1$ transition to record how many atoms remain after each set of 50 decay measurements. Every 30 cycles, an interrupt shuts off the clock laser and photon counter trigger so that the atom number can be recorded in a manner that is decoupled from the effects of the clock and DFB lasers. Treating

⁵⁸ To calculate this bias, I account for the fact that the stray light interacts with the atoms whenever the shutter is open during the decay measurement, which is on a tens of ms timescale. The calculation is performed by running the simulation with the stray light included and then fitting the result. The disagreement between the fitted and true values of 3D_1 is the fit bias.

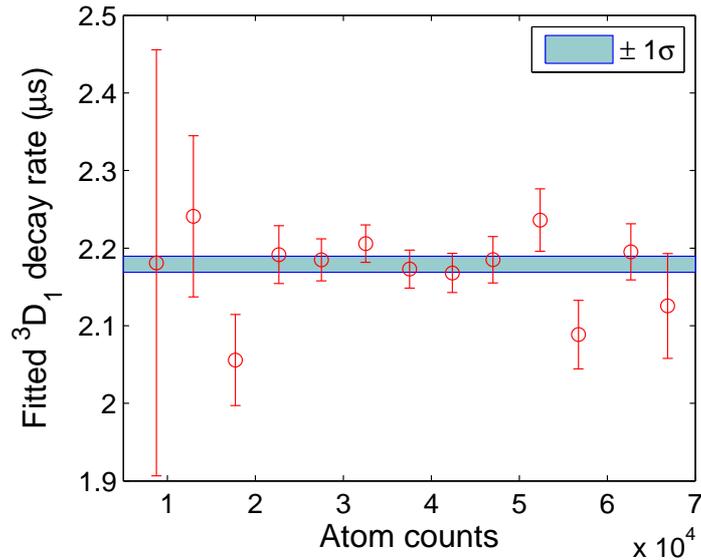


Figure 5.35: Fitted 3D_1 lifetime as a function of atom counts, which is proportional to the atom number. Here the data from Figure 5.30 is reanalyzed to look for density dependence. No trend is observed. The blue confidence interval represents the 1σ uncertainty in the weighted mean assuming the decay is constant in density. This uncertainty has been inflated by $\sqrt{\chi_{red}^2}$, where $\chi_{red}^2 = 1.12$ for these data.

this one-in-thirty measurement as the atom number for each of the previous 29 cycles, sets of decay data (as a function of time) are binned according to atom number, then summed and fitted as in Figure 5.30.

The fitted 3D_1 lifetime as a function of atom counts on the PMT does not show the same striking trend observed by the Yb group [3]; in fact, it does not seem to be affected by density at all [88]. To test the data for density dependence, a constant model is compared to a linear trend, which would be the first-order description of density-related effects. The F -test comparing these models results in $F = 0.045$ for 11 degrees of freedom. This corresponds to a test probability of 0.84, which concludes that the data does not resolve a density dependence; therefore, density effects are excluded from the uncertainty budget for the decay measurement.

Two more systematic uncertainties are considered. First, 0.4 ns timing uncertainty is assigned to the photon counting setup. Second, the decay measurement assumes that the 3D_1 lifetime does not depend on the large number of F and m_F states that are populated during this experiment.

Effect	Uncertainty (ns)
Fit uncertainty	10
Photon counter timing	0.4
Finite pulse duration	< 0.1
Hyperfine correction	< 0.1
Stray laser light	< 0.01
Laboratory heat contamination	< 0.01
Total	10

Table 5.4: Decay measurement uncertainty budget. All uncertainties are quoted at 1σ or as upper bounds on a 1σ uncertainty. Fitting uncertainty dominates this budget, with all other effects being negligible at the 1 ns level.

One typically sums the decay rate over all product hyperfine states, resulting in decays that do not depend on either the initial or final hyperfine quantum numbers [44]. This approximation breaks down when hyperfine quenching becomes important. The hyperfine quenching correction is calculated to be below 0.1 ns.⁵⁹

Combining all of these uncertainties, the final value of the 3D_1 lifetime is $\tau_{^3D_1} = (2.18 \pm 0.01) \mu\text{s}$ (Table 5.4) [88]. This results in a dynamic coefficient of $\nu_{dyn} = (-148.7 \pm 0.7) \text{ mHz}$, or -3.464×10^{-16} with uncertainty of 1.6×10^{-18} [88]. For the Sr2 operating temperature of 20.6°C , the dynamic BBR shift is -3.053×10^{-16} with uncertainty of 1.4×10^{-18} .

5.10 Minor Systematic Shifts

With all major systematic uncertainties now discussed, I will briefly mention a few minor uncertainties that were checked. First, phase transients can occur when the clock laser AOM (Section 3.7.1) is pulsed [26]. We study the AOM phase transients by looking at a beat of the -1st AOM order with the zeroth-order on a digital phase detector. The phase detector was found to have transients of its own that were measured, fit to a high-order polynomial, and removed from the final measurement. We calculate the effect of clock laser phase transients on the atoms using a model developed in Reference [34]. This results in systematic uncertainty due to an AOM phase chirp of $(6 \pm 4) \times 10^{-19}$. This value has been used in both evaluations.

⁵⁹ My appreciation goes out to our theory collaborator, Marianna Safronova, for performing this calculation.

Evaluations of several different optical clocks have considered a servo offset [63, 33, 52], which is steady state error in the atomic servo resulting in a frequency bias of the clock laser. This offset is computed as the mean of the entire history of atomic servo error signal data converted from excited state fraction units (Section 2.5.2) to frequency. This is done by dividing error signal data taken with a given clock laser pulse duration by 2 times the calculated slope of a Rabi lineshape (corresponding to the correct pulse time) at half the maximum excitation. The factor of 2 comes from the fact that the atomic servo error signal is the difference between excited state fractions measured on each side of a resonance, so the error signal change per unit frequency excursion is 2 times the Rabi lineshape slope. For the first evaluation, the servo offset is $(4 \pm 6) \times 10^{-19}$, and it is $(-5 \pm 4) \times 10^{-19}$ for the second evaluation.

A phenomenon called “line pulling” has the potential to cause small clock transition shifts. This occurs when the tails of off-resonant spectroscopic features make the clock transition resonance appear slightly off its true center. These features can be present due to imperfect spin polarization resulting in population in m_F states aside from $\pm 9/2$, residual clock laser ellipticity driving m_F -changing σ transitions, or clock transition sidebands that result from tunneling between lattice sites. Population in the $m_F = \pm 7/2$ states is most likely to cause off-resonant excitation, but some common-mode cancellation of this effect will be present in the stretched-states servo. This is because equal population in the $\pm 7/2$ states will pull the $\pm 9/2$ levels toward the clock transition bare frequency equally, and averaging the $\pm 9/2$ stretched state data will remove this effect. Calculating line pulling for these cases and bounding the calculations using the fact that, with our measurement precision, we do not observe any off-resonant features near the clock transition (where they would be the most harmful), we put a conservative upper bound on this effect of 1×10^{-19} . This is valid for both evaluations.

In the ideal case of a perfectly stable lattice, probing the Sr atoms with the clock laser aligned along the lattice axis will be free of a first-order Doppler shift. If mechanical noise causes the lattice to shake with respect to the clock laser, a first-order Doppler shift is possible. Fiber phase noise cancellation (Section 3.6.2) removes this problem, but a second-order Doppler shift (the second-

order term in a low-velocity expansion of the relativistic Doppler expression) can still exist. We estimate this effect to be at the 10^{-21} level. To buffer this calculation against errors, we give it a comfortable 1×10^{-19} upper bound for both evaluations.

The background gas in the Sr2 main chamber that limits the vacuum lifetime will collide with the trapped Sr sample and cause frequency shifts. For the Sr2 vacuum pressure, the background gas is mostly hydrogen. Using this fact, we employ the model of Reference [38] to put an upper bound of 6×10^{-19} on frequency shifts due to collisions between Sr and the background gas. This bound is used for both evaluations.

Finally, there is added uncertainty associated with treating the BBR static and dynamic shifts as two separate effects. When the static and dynamic uncertainties are added in quadrature, this ignores the fact that they are not statistically independent. With a bit of uncertainty propagation mathematics applied to Equation 5.27, it is straightforward to show that the static and dynamic shift uncertainties must be added in quadrature with a covariance-like term $\sqrt{48\nu_{stat}\nu_{dyn}} \left(\frac{T}{T_0}\right)^4 \frac{\delta T}{T_0}$, which is equal 8.2×10^{-19} for the first evaluation and 1.4×10^{-19} for the second. The value for the first evaluation is just barely negligible (a slightly larger value would affect the second digit of the total clock uncertainty), and for the second evaluation this correction can comfortably be ignored.

Chapter 6

Outlook and Concluding Remarks

I started as a first-year graduate student on the JILA Sr team shortly before Sr1 achieved 1.4×10^{-16} total uncertainty, which was the first time a lattice clock surpassed the total uncertainty of the Cs primary standard at NIST Boulder [74]. 1.4×10^{-16} total uncertainty [20] was a great accomplishment at the time, but despite our optimism about the future of lattice clocks, it is safe to say that it was unclear whether an optical lattice clock would ever overtake ion clocks. Three major issues needed to be addressed for lattice clocks to realize the best performance: clock stability, the density shift, and the BBR shift.

Although lattice clocks seemed promising because of their potential for good stability, this would require better laser stabilization. Unfortunately, for over a decade the best laser stabilization efforts had not managed to surpass laser linewidth records set in the late 90s [126]. Then cavity stabilization took a leap forward [89, 49], resulting in laser linewidths an order of magnitude narrower than the previous best [8]. Clock stability resulting from these laser improvements enabled the record total uncertainty we later achieved.

The density shift is not present in single-ion clocks. When Sr1 achieved 1.4×10^{-16} total uncertainty, a density shift was observed in Sr1 evaluations [74, 20]. This shift was a major contributor to Sr1 uncertainty budgets, yet its physical cause was confusing. It implied collisions between ultracold spin-polarized fermions, which naively seemed to be forbidden by the Pauli Exclusion Principle. After many experiments and theoretical studies aimed at understanding and mitigating this shift [19, 99, 37, 64, 73, 6, 89, 78], it is now a small contributor to the Sr2 clock uncertainty.

The BBR shift has limited many of the best optical clocks [114, 74, 52, 76]. Ion-based systems have an advantage on this front due to naturally smaller sensitivity of ion clocks to room-temperature BBR. When Sr1 first outperformed Cs, BBR shift uncertainty was at 1×10^{-16} . Given that the sensor arrays used to measure the BBR shift in many lattice clocks [74, 33, 114] were merely workarounds for systems not originally designed with this shift in mind, there was plenty of room for progress. However, beating the 3×10^{-18} uncertainty in the miniscule BBR shift of the Al^+ clock seemed a formidable task. When the time came for the lattice clock community to confront the BBR shift head on, the first three teams to make serious efforts on the BBR shift all succeeded using distinctly different techniques [13, 2, 116]. It is now clear that the BBR shift is not a limitation at the 1×10^{-18} level, and it will no doubt improve as the lattice clock community continues to work on this effect.

Despite the atomic clock performance records presented in this dissertation of 2.1×10^{-18} total uncertainty and $2.2 \times 10^{-16}/\sqrt{\tau}$ stability, lattice clocks still have tremendous potential. For the clock probe times, duty cycles, and atom numbers of 1×10^5 achievable in the Sr2 system, I calculate the QPN-limited stability to be $2 \times 10^{-18}/\sqrt{\tau}$. This would allow a clock to average to a given precision 10,000 times faster than Sr2 can currently achieve. Although the density shift would be a problem for large atom numbers in a 1D cavity lattice, multidimensional cavity lattices could fix this [109]. The Sr2 blackbody shift due to temperature uncertainty is already at the low 10^{-19} level, and the uncertainty of the full BBR shift could match this with more precise measurements of ν_{dyn} . The NIST Yb clock would also see improvement with a more precise measurement of ν_{dyn} [2], reducing their total BBR shift uncertainty to the 10^{-19} level. Even cryogenic environments have room for improvement [116] despite that advances in this technology would be more technically challenging. For example, colder cryogenic environments can be used to further reduce the BBR shift. Also, to confirm or reject previous determinations of higher-order lattice light shifts, it will be important for another Sr team to measure these effects.

Better clock lasers are also on the horizon. Novel cavity spacers based on silicon [59] and new mirrors with crystalline coatings [24] have been demonstrated to improve cavity stabilization.

Our team is in the process of combining this technology to build a new clock laser, which will surpass the previous best stable laser by a factor of 10 according to our simulations. Also, proof-of-principle tests of proposals for novel stable lasers that do not use cavity stabilization [80, 79] have been demonstrated [15]. If these proposed schemes become practical laboratory technology, clock stability could take another giant leap forward.

Timekeeping at the level realized by Sr2 would not lose a second in 15 billion years (greater than the age of the universe), yet there are good scientific reasons to push performance further. Sr2 can now resolve 2 cm height differences near the surface of Earth due to the gravitational redshift. With this sensitivity, optical comparisons of two remote clocks [96] would be sensitive to variations of Earth's surface due to tidal forces. Better clock precision combined with efforts to realize portable clocks (currently underway at NIST and PTB) could make these systems useful for relativistic geodesy. Clocks are also useful for studying fundamental physics like time variation of the fine structure constant or the electron-to-proton mass ratio [40, 51]. Although ion clocks are more sensitive to fundamental constant variation, with better precision lattice clocks may be in a position to make competitive measurements.

Sr clocks have also been used to put limits on the coupling of fundamental constants to gravity [11], which is predicted by some unification theories. Improvements to clock performance could serve to improve these limits. Furthermore, upgrading to the 40 cm cavity enabled Sr1 to spectroscopically observe many-body physics [78] and novel particle interactions [127, 108]. It is likely that many more interesting quantum effects would be uncovered with better measurement precision.

Theoretical proposals have argued for replacing the current system of global timekeeping with a worldwide network of optical clocks [62]. Although such a network is far from being realized, it exemplifies the potential of linking up precise optical clocks on an international scale. Finally, it was argued that atomic clocks are useful systems for dark matter searches [1, 27]. We may be entering a regime in which atomic clocks can answer some of the most pressing physical questions of our time.

In summary, the outlook for lattice clocks is bright, and I expect that the lattice clock community will continue with the same impressive pace it has maintained for the past 10 years. It has been a pleasure to participate in this wonderful field.

Bibliography

- [1] A. Arvanitaki, J. Huang, and K. Van Tilburg. Searching for dilaton dark matter with atomic clocks. Physical Review D, 91:015015, 2015.
- [2] K. Beloy, N. Hinkley, N.B. Phillips, J.A. Sherman, M. Schioppo, J. Lehman, A. Feldman, L.M. Hanssen, C.W. Oates, and A.D. Ludlow. Atomic clock with 1×10^{-18} room-temperature blackbody stark uncertainty. Physical Review Letters, 113:260801, 2014.
- [3] K. Beloy, J.A. Sherman, N.D. Lemke, N. Hinkley, C.W. Oates, and A.D. Ludlow. Determination of the $5d6s\ ^3D_1$ state lifetime and blackbody-radiation clock shift in Yb. Physical Review A, 86:051404, 2012.
- [4] P.R. Bevington and D.K. Robinson. Data Reduction and Error Analysis for the Physical Sciences. Mc Graw-Hill, 3rd edition, 2003.
- [5] G.W. Biedermann, K. Takase, X. Wu, L. Deslauriers, S. Roy, and M.A. Kasevich. Zero-dead-time operation of interleaved atomic clocks. Physical Review Letters, 111:170802, 2013.
- [6] M. Bishof, Y. Lin, M.D. Swallows, A.V. Gorshkov, J. Ye, and A.M. Rey. Resolved atomic interaction sidebands in an optical clock transition. Physical Review Letters, 106:250801, 2011.
- [7] M. Bishof, M.J. Martin, M.D. Swallows, C. Benko, Y. Lin, G. Quémener, A.M. Rey, and J. Ye. Inelastic collisions and density-dependent excitation suppression in a ^{87}Sr optical lattice clock. Physical Review A, 84:052716, 2011.
- [8] M. Bishof, X. Zhang, M.J. Martin, and J. Ye. Optical spectrum analyzer with quantum-limited noise floor. Physical Review Letters, 111:093604, 2013.
- [9] S. Bize, Y. Sortais, P. Lemonde, S. Zhang, P. Laurent, G. Santarelli, C. Salomon, and A. Clairon. Interrogation oscillator noise rejection in the comparison of atomic fountains. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 47:1253–1255, 2000.
- [10] E.D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. American Journal of Physics, 69:79–87, 2001.
- [11] S. Blatt, A.D. Ludlow, G.K. Campbell, J.W. Thomsen, T. Zelevinsky, M.M. Boyd, J. Ye, X. Baillard, M. Fouché, R. Le Targat, A. Brusch, P. Lemonde, M. Takamoto, F.-L. Hong, and H. Katori. New limits on coupling of fundamental constants to gravity using ^{87}Sr optical lattice clocks. Physical Review Letters, 100:140801, 2008.

- [12] S. Blatt, J.W. Thomsen, G.K. Campbell, A.D. Ludlow, M.D Swallows, M.J. Martin, M.M. Boyd, and J. Ye. Rabi spectroscopy and excitation inhomogeneity in a one-dimensional optical lattice clock. Physical Review A, 80:052703, 2009.
- [13] B.J. Bloom, T.L. Nicholson, J.R. Williams, S.L. Campbell, M. Bishof, X. Zhang, W. Zhang, S.L. Bromley, and J. Ye. An optical lattice clock with accuracy and stability at the 10^{-18} level. Nature, 506:71–75, 2014.
- [14] D. Boguhn and M. Koepke. Typical $R(T_{90})$ characteristics of platinum thin-film resistance thermometers in the temperature range from -50°C to $+660^{\circ}\text{C}$. International Journal of Thermophysics, 32:2379–2387, 2011.
- [15] J.G. Bohnet, Z. Chen, J.M. Weiner, D. Meiser, M.J. Holland, and J.K. Thompson. A steady-state superradiant laser with less than one intracavity photon. Nature, 484:78–81, 2012.
- [16] M.M. Boyd, T. Zelevinsky, A.D. Ludlow, S. Blatt, T. Zanon-Willette, S.M. Foreman, and J. Ye. Nuclear spin effects in optical lattice clocks. Physical Review A, 76:022510, 2007.
- [17] B.H. Bransden and C.J. Joachain. Physics of Atoms and Molecules. Prentice Hall, 2nd edition, 2003.
- [18] A. Bruschi, R. Le Targat, X. Baillard, M. Fouché, and P. Lemonde. Hyperpolarizability effects in a Sr optical lattice clock. Physical Review Letters, 96:103003, 2006.
- [19] G.K. Campbell, M.M. Boyd, J.W. Thomsen, M.J. Martin, S. Blatt, M.D. Swallows, T.L. Nicholson, T. Fortier, C.W. Oates, S.A. Diddams, N.D. Lemke, P. Naidon, P. Julienne, J. Ye, and A.D. Ludlow. Probing interactions between ultracold fermions. Science, 324:360–363, 2009.
- [20] G.K. Campbell, A.D. Ludlow, S. Blatt, J.W. Thomsen, M.J. Martin, M.H.G. de Miranda, T. Zelevinsky, M.M. Boyd, J. Ye, S.A. Diddams, T.P. Heavner, T.E. Parker, and S.R. Jefferts. The absolute frequency of the ^{87}Sr optical clock transition. Metrologia, 45:539–548, 2008.
- [21] R. Carvell, E. Usadi, N. Fox, R. Rusby, and S. Bruce. Innovative black body for on-board calibration. In Proceedings of the SPIE: Sensors, Systems, and Next-Generation Satellites, volume VII, pages 411–422, 2004.
- [22] S.-W. Chiow, S. Herrmann, H. Müller, and S. Chu. 6 W, 1 kHz linewidth, tunable continuous-wave near-infrared laser. Optics Express, 17:5246–5250, 2009.
- [23] C.W. Chou, D.B. Hume, J.C.J. Koelemeij, D.J. Wineland, and T. Rosenband. Frequency comparison of two high-accuracy Al^+ optical clocks. Physical Review Letters, 104:070802, 2010.
- [24] G.D. Cole, W. Zhang, M.J. Martin, J. Ye, and M. Aspelmeyer. Tenfold reduction of brownian noise in high-reflectivity optical coatings. Nature Photonics, 7:644–650, 2013.
- [25] S.S. Courts and P.R. Swinehart. Long-term stability of a cryogenic diode thermometer. In Proceedings of the 2002 Cryogenic Engineering Conference, volume 613, page 1636. American Institute of Physics, 2002.

- [26] C. Degenhardt, T. Nazarova, C. Lisdat, H. Stoehr, U. Sterr, and F. Riehle. Influence of chirped excitation pulses in an optical clock with ultracold calcium atoms. IEEE Transactions on Instrumentation and Measurement, 54:771–775, 2005.
- [27] A. Derevianko and M. Pospelov. Hunting for topological dark matter with atomic clocks. Nature Physics, 10:933–936, 2014.
- [28] G.J. Dick. Local oscillator induced instabilities in trapped ion frequency standards. In Proceedings of the 19th Precise Time and Time Interval Meeting, pages 133–147, 1987.
- [29] G.J. Dick, J.D. Prestage, C.A. Greenhall, and L. Maleki. Local oscillator induced degradation of medium-term stability in passive atomic frequency standards. In Proceedings of the 22th Precise Time and Time Interval Meeting, pages 487–508, 1990.
- [30] R.H. Dicke. The effect of collisions upon the doppler width of spectral lines. Physical Review, 89:472–473, 1953.
- [31] S.A. Diddams, J.C. Bergquist, S.R. Jefferts, and C.W. Oates. Standards of time and frequency and the outset of the 21st century. Science, 306:1318–1324, 2004.
- [32] R. Drozdowski, M. Ignaciuk, J. Kwela, and J. Heldt. Radiative lifetimes of the lower 3P_1 metastable states of Ca and Sr. Zeitschrift für Physik D, 41:125–131, 1997.
- [33] S. Falke, N. Lemke, C. Grebing, B. Lipphardt, S. Weyers, V. Gerginov, N. Huntemann, C. Hagemann, A. Al-Masoudi, S. Häfner, S. Vogt, U. Sterr, and C. Lisdat. A strontium lattice clock with 3×10^{-17} inaccuracy and its frequency. New Journal of Physics, 16:073023, 2014.
- [34] S. Falke, M. Misera, U. Sterr, and C. Lisdat. Delivering pulse and phase stable light to atoms of an optical clock. Applied Physics B, 107:301–311, 2012.
- [35] S. Falke, H. Schnatz, J.S.R. Vellore Winfred, T. Middelmann, S. Vogt, S. Weyers, B. Lipphardt, G. Grosche, F. Riehle, U. Sterr, and C. Lisdat. The ^{87}Sr optical frequency standard at PTB. Metrologia, 48:399–407, 2011.
- [36] S.M. Foreman, A.D. Ludlow, M.H.G. de Miranda, J.E. Stalnaker, S.A. Diddams, and J. Ye. Coherent optical phase transfer over a 32-km fiber with 1 s instability $< 10^{-17}$. Physical Review Letters, 99:153601, 2007.
- [37] K. Gibble. Decoherence and collisional frequency shifts of trapped bosons and fermions. Physical Review Letters, 103:113202, 2009.
- [38] K. Gibble. Scattering of cold-atom coherences by hot atoms: frequency shifts from background-gas collisions. Physical Review Letters, 110:180802, 2013.
- [39] P. Gill. When should we change the definition of the second? Philosophical Transactions of the Royal Society A, 369:4109–4130, 2011.
- [40] R.M. Godun, P.B.R. Nisbet-Jones, J.M. Jones, S.A. King, L.A.M. Johnson, H.S. Margolis, K. Szymaniec, S.N. Lea, K. Bongs, and P. Gill. Frequency ratio of two optical clock transitions in $^{171}\text{Yb}^+$ and constraints on the time variation of fundamental constants. Physical Review Letters, 113:210801, 2014.

- [41] H. Goldstein, C.P. Poole, and J.L. Safko. Classical Mechanics. Addison Wesley, 3rd edition, 2001.
- [42] C.A. Greenhall. A derivation of the long-term degradation of a pulsed atomic frequency standard from a control-loop model. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 45:895–898, 1998.
- [43] C.A. Greenhall, D.A. Howe, and D.B. Percival. Total variance, an estimator of long-term frequency stability. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 46:1183–1191, 1999.
- [44] G. Grynberg, A. Aspect, and C. Fabre. Introduction to Quantum Optics: From the Semi-Classical Approach to Quantized Light. Cambridge University Press, 2010.
- [45] J. L. Hall and M. Zhu. An introduction to phase-stable optical sources. In Proceedings of the International School of Physics “Enrico Fermi,” Course CXVIII, pages 671–702. North Holland Elsevier, Amsterdam, 1992.
- [46] M. Harlander, M. Brownnutt, W. Hänsel, and R. Blatt. Trapped-ion probing light-induced charging effects on dielectrics. New Journal of Physics, 12:093035, 2010.
- [47] A.H. Harvey, M.O. McLinden, and W.L. Tew. Thermodynamic analysis and experimental study of the effect of atmospheric pressure on the ice point. In Proceedings of the 9th International Temperature Symposium, 2013.
- [48] T.P. Heavner, E.A. Donley, F. Levi, G. Costanzo, T.E. Parker, J.H. Shirley, N. Ashby, S. Barlow, and S.R. Jefferts. First accuracy evaluation of NIST – F2. Metrologia, 51:174–182, 2014.
- [49] N. Hinkley, J.A. Sherman, N.B. Phillips, M. Schioppo, N.D. Lemke, K. Beloy, M. Pizzocaro, C.W. Oates, and A.D. Ludlow. An atomic clock with 10^{-18} instability. Science, 341:1215–1218, 2013.
- [50] D.A. Howe, D.W. Allan, and J.A. Barnes. Properties of signal sources and measurement methods. In Proceedings of the 35th Annual Symposium on Frequency Control, pages 1–47, 1981.
- [51] N. Huntemann, B. Lipphardt, C. Tamm, V. Gerginov, S. Weyers, and E. Peik. Improved limit on a temporal variation of m_p/m_e from comparisons of Yb^+ and Cs atomic clocks. Physical Review Letters, 113:210802, 2014.
- [52] N. Huntemann, M. Okhapkin, B. Lipphardt, S. Weyers, Chr. Tamm, and E. Peik. High-accuracy optical clock based on the octopole transition in $^{171}\text{Yb}^+$. Physical Review Letters, 108:090801, 2012.
- [53] F.P. Incopera, D.P. DeWitt, T.L. Bergman, and A.S. Lavine. Introduction to Heat Transfer. Wiley, 5th edition, 2006.
- [54] W.M. Itano, J.C. Bergquist, J.J. Bollinger, J.M. Gilligan, D.J. Heinzen, F.L. Moore, M.G. Raizen, and D.J. Wineland. Quantum projection noise: Population fluctuations in two-level systems. Physical Review A, 47:3554, 1993.
- [55] J.D. Jackson. Classical Electrodynamics. Wiley, 3rd edition, 1998.

- [56] J. Jespersen and J. Fitz-Randolph. From Sundials to Atomic Clocks: Understanding Time and Frequency. Dover, 1999.
- [57] H. Katori. Spectroscopy of strontium atoms in the Lamb-Dicke confinement. In Proceedings of the 6th Symposium on Frequency Standards and Metrology, pages 323–330. University of St Andrews, Fife, Scotland, 2002.
- [58] H. Katori, M. Takamoto, V.G. Pal’chikov, and V.D. Ovsiannikov. Ultrastable optical clock with neutral atoms in an engineered light shift trap. Physical Review Letters, 91:173005, 2003.
- [59] T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M.J. Martin, L. Chen, and J. Ye. A sub-40-mHz-linewidth laser based on a silicon single-crystal optical cavity. Nature Photonics, 6:687–692, 2012.
- [60] H.J. Kimble, C.J. Hood, T.W. Lynn, H. Mabuchi, D.W. Vernooy, and J. Ye. The quantum internet. In Proceedings of the 14th International Conference on Laser Spectroscopy, pages 80–89. World Scientific, Singapore, 1999.
- [61] H.-J. Kluge and H. Sauter. Level crossing experiments in the first excited 1P_1 states of the alkaline earths. Zeitschrift für Physik, 270:295–309, 1974.
- [62] P. Kómár, E.M. Kessler, M. Bishof, L. Jiang, A.S. Sørensen, J. Ye, and M.D. Lukin. A quantum network of clocks. Nature Physics, 10:582–587, 2014.
- [63] N.D. Lemke, A.D. Ludlow, Z.W. Barber, T.M. Fortier, S.A. Diddams, Y. Jiang, S.R. Jefferts, T.P. Heavner, T.E. Parker, and C.W. Oates. Spin-1/2 optical lattice clock. Physical Review Letters, 103:063001, 2009.
- [64] N.D. Lemke, J. von Stecher, J.A. Sherman, A.M. Rey, C.W. Oates, and A.D. Ludlow. p -wave cold collisions in an optical lattice clock. Physics Review Letters, 107:103902, 2011.
- [65] P. Lemonde, G. Santarelli, Ph. Laurent, F. Pereira dos Santos, A. Clairon, and C. Salomon. Properties of signal sources and measurement methods. In Proceedings of the 1998 IEEE International Frequency Control Symposium, pages 110–115, 1998.
- [66] P.J. Leo, P.S. Julianne, F.H. Mies, and C.J. Williams. Collisional frequency shifts in ^{133}Cs fountain clocks. Physical Review Letters, 86:3743–3746, 2001.
- [67] J. Levine. Introduction to time and frequency metrology. Review of Scientific Instruments, 70:2567–2596, 1999.
- [68] Z.M. Liao, S.A. Payne, J. Dawson, A. Drobshoff, C. Ebberts, and D. Pennington. Thermally induced dephasing in periodically poled KTP frequency-doubling crystals. Journal of the Optical Society of America B, 21:2191–2196, 2004.
- [69] J. Lodewyck, M. Zawada, L. Lorini, M. Gurov, and P. Lemonde. Observation and cancellation of a perturbing dc stark shift in strontium optical lattice clocks. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 59:411–415, 2012.
- [70] T.H. Loftus, T. Ido, M.M. Boyd, A.D. Ludlow, and J. Ye. Narrow line cooling and momentum space crystals. Physical Review A, 70:063413, 2004.

- [71] A.D. Ludlow. The Strontium Optical Lattice Clock: Optical Spectroscopy with Sub-Hertz Accuracy. PhD thesis, University of Colorado at Boulder, 2008.
- [72] A.D. Ludlow, X. Huang, M. Notcutt, T. Zanon-Willette, S.M. Foreman, M.M. Boyd, S. Blatt, and J. Ye. Compact, thermal-noise-limited optical cavity for diode laser stabilization at 1×10^{-15} . Optics Letters, 32:641–643, 2007.
- [73] A.D. Ludlow, N.D. Lemke, J.A. Sherman, C.W. Oates, G. Quémener, J. von Stecher, and A.M. Rey. Cold-collision-shift cancellation and inelastic scattering in a Yb optical lattice clock. Physical Review A, 84:052724, 2011.
- [74] A.D. Ludlow, T. Zelevinsky, G.K. Campbell, S. Blatt, M.M. Boyd, M.H.G. de Miranda, M.J. Martin, J.W. Thomsen, S.M. Foreman, J. Ye, T.M. Fortier, J.E. Stalnaker, S.A. Diddams, Y. Le Coq, Z.W. Barber, N. Poli, N.D. Lemke, K.M. Beck, and C.W. Oates. Sr lattice clock at 1×10^{-16} fractional uncertainty by remote optical evaluation with a Ca clock. Science, 319:1805–1808, 2008.
- [75] L.-S. Ma, P. Jungner, J. Ye, and J.L. Hall. Delivering the same optical frequency at two places: accurate cancellation of phase noise introduced by an optical fiber or other time-varying path. Optics Letters, 19:1777–1779, 1994.
- [76] A.A. Madej, P. Dubé, Z. Zhou, J.E. Bernard, and M. Gertszov. $^{88}\text{Sr}^+$ 445-THz single-ion reference at the 10^{-17} level via control and cancellation of systematic uncertainties and its measurement against the SI second. Physical Review Letters, 109:203002, 2012.
- [77] M.J. Martin. Quantum Metrology and Many-Body Physics: Pushing the Frontier of the Optical Lattice Clock. PhD thesis, University of Colorado at Boulder, 2013.
- [78] M.J. Martin, M. Bishof, M.D. Swallows, X. Zhang, C. Benko, J. von Stecher, A.V. Gorshkov, A.M. Rey, and J. Ye. A quantum many-body spin system in an optical lattice clock. Science, 341:632–637, 2013.
- [79] M.J. Martin, D. Meiser, J.W. Thomsen, J. Ye, and M.J. Holland. Extreme nonlinear response of ultranarrow optical transitions in cavity QED for laser stabilization. Physical Review A, 84:063813, 2011.
- [80] D. Meiser, J. Ye, D.R. Carlson, and M.J. Holland. Prospects for a millihertz-linewidth laser. Physical Review Letters, 102:163601, 2009.
- [81] H.J. Metcalf and P. van der Straten. Laser Cooling and Trapping. Springer, 2001.
- [82] T. Middelmann, S. Falke, C. Lisdat, and U. Sterr. High accuracy correction of blackbody radiation shift in an optical lattice clock. Physical Review Letters, 109:263004, 2012.
- [83] T. Middelmann, C. Lisdat, S. Falke, J.S.R.V. Winfred, F. Riehle, and U. Sterr. Tackling the blackbody shift in a strontium optical lattice clock. IEEE Transactions on Instrumentation and Measurement, 60:2550–2557, 2011.
- [84] D.A. Miller, L. You, J. Cooper, and A. Gallagher. Collisional energy transfer between excited-state strontium and noble-gas atoms. Physical Review A, 46:1303, 1992.
- [85] M. F. Modest. Radiative Heat Transfer. Academic Press, Elsevier Science, 2003.

- [86] T. Mukaiyama, H. Katori, T. Ido, Y. Li, and M. Kuwata-Gonokami. Recoil-limited laser cooling of ^{87}Sr atoms near the Fermi temperature. Physical Review Letters, 90:113002, 2003.
- [87] W. Nagourney. Quantum Electronics for Atomic Physics. Oxford, 1st edition, 2010.
- [88] T.L. Nicholson, S.L. Campbell, R.B. Hutson, G.E. Marti, B.J. Bloom, R.L. McNally, W. Zhang, M.D. Barrett, M.S. Safronova, G.F. Strouse, W.L. Tew, and J. Ye. Systematic evaluation of an atomic clock at 2×10^{-18} total uncertainty. Nature Communications, 6:6896, 2015.
- [89] T.L. Nicholson, M.J. Martin, J.R. Williams, B.J. Bloom, M. Bishof, M.D. Swallows, S.L. Campbell, and J. Ye. Comparison of two independent sr optical clocks with 1×10^{-17} stability at 10^3 s. Physical Review Letters, 109:230801, 2012.
- [90] M. Notcutt, L.-S. Ma, A.D. Ludlow, S.M. Foremand, J. Ye, and J.L. Hall. Contribution of thermal noise to frequency stability of rigid optical cavity via hertz-linewidth lasers. Physical Review A, 73:031804, 2006.
- [91] M. Notcutt, L.-S. Ma, J. Ye, and J.L. Hall. Simple and compact 1-Hz laser system via an improved mounting configuration of a reference cavity. Optics Letters, 30:1815–1817, 2005.
- [92] K. Numata, A. Kemery, and J. Camp. Thermal-noise limit in the frequency stabilization of lasers with rigid cavities. Physical Review Letters, 93:250602, 2004.
- [93] W.H. Oskay, S.A. Diddams, E.A. Donley, T.M. Fortier, T.P. Heavner, L. Hollberg, W.M. Itano, S.R. Jefferts, M.J. Delaney, K. Kim, F. Levi, T.E. Parker, and J.C. Bergquist. Single-atom optical clock with high accuracy. Physical Review Letters, 97:020801, 2006.
- [94] S.G. Porsev and A. Derevianko. Multipolar theory of blackbody radiation shift of atomic energy levels and its implications for optical lattice clocks. Physical Review A, 74:020502, 2006.
- [95] S.G. Porsev, A.D. Ludlow, M.M. Boyd, and J. Ye. Determination of Sr properties for a high-accuracy optical clock. Physical Review A, 78:032508, 2008.
- [96] K. Predehl, G. Grosche, S.M.F. Raupach, S. Droste, O. Terra, J. Alnis, T. Legero, T.W. Hänsch, T. Udem, R. Holzwarth, and H. Schnatz. A 920-kilometer optical fiber link for frequency metrology at the 19th decimal place. Science, 336:441–444, 2012.
- [97] F. Quinlan, T.M. Fortier, H. Jiang, A. Hati, C. Nelson, Y. Fu, J.C. Campbell, and S.A. Diddams. Exploiting shot noise correlations in the photodetection of ultrashort optical pulse trains. Nature Photonics, 7:290–293, 2013.
- [98] A.M. Rey. Private communication, 2012.
- [99] A.M. Rey, A.V. Gorshkov, and C. Rubbo. Many-body treatment of the collisional frequency shift in fermionic atoms. Physical Review Letters, 103:260402, 2009.
- [100] F. Riehle. Frequency Standards: Basics and Applications. Wiley-VCH, 2003.
- [101] W.J. Riley. Handbook of Frequency Stability Analysis. NIST Special Publication 1065, 2008.

- [102] M.E. Rose. Elementary Theory of Angular Momentum. Dover, 1995. Republished from Wiley, 1957.
- [103] T. Rosenband, D.B. Hume, P.O. Schmidt, C.W. Chou, A. Brusch, L. Lorini, W.H. Oskay, R.E. Drullinger, T.M. Fortier, J.E. Stalnaker, S.A. Diddams, W.C. Swann, N.R. Newbury, W.M. Itano, D.J. Wineland, and J.C. Bergquist. Frequency ratio of Al^+ and Hg^+ single-ion optical clocks; metrology at the 17th decimal place. Science, 319:1808–1812, 2008.
- [104] M.S. Safronova, M.G. Kozlov, and C.W. Clark. Blackbody radiation shifts in optical atomic clocks. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 59:439–447, 2012.
- [105] M.S. Safronova, S.G. Porsev, U.I. Safronova, M.G. Kozlov, and C.W. Clark. Blackbody-radiation shift in the Sr optical atomic clock. Physical Review A, 87:012509, 2013.
- [106] G. Santarelli, C. Audoin, A. Makdissi, P. Laurent, G.J. Dick, and A. Clairon. Frequency stability degradation of an oscillator slaved to a periodically interrogated atomic resonator. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 45:887–894, 1998.
- [107] R. Santra, K.V. Christ, and C.H. Greene. Properties of metastable alkaline-earth-metal atoms calculated using an accurate effective core potential. Physical Review A, 69:042510, 2004.
- [108] F. Scazza, C. Hofrichter, M. Höfer, P.C. De Groot, I. Bloch, and S. Fölling. Observation of two-orbital spin-exchange interactions with ultracold $\text{SU}(N)$ -symmetric fermions. Nature Physics, 10:779–784, 2014.
- [109] M.D. Swallows, M. Bishof, Y. Lin, S. Blatt, M.J. Martin, A.M. Rey, and J. Ye. Suppression of collisional shifts in a strongly interacting lattice clock. Science, 331:1043–1046, 2011. Supplementary Online Material.
- [110] M.D. Swallows, M.J. Martin, M. Bishof, C. Benko, Y. Lin, S. Blatt, A.M. Rey, and J. Ye. Operating a ^{87}Sr optical lattice clock with high precision and at high density. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 59:416–425, 2012.
- [111] A.V. Taichenachev, V.I. Yudin, V.D. Ovsiannikov, V.G. Pal’chikov, and C.W. Oates. Frequency shifts in an optical lattice clock due to magnetic-dipole and electric-quadrupole transitions. Physics Review Letters, 101:193601, 2008.
- [112] M. Takamoto, F.-L. Hong, R. Higashi, Y. Fujii, M. Imae, and H. Katori. Improved frequency measurement of a one-dimensional optical lattice clock with a spin-polarized fermionic ^{87}Sr isotope. Journal of the Physical Society of Japan, 75:104302, 2006.
- [113] M. Takamoto, T. Takano, and H. Katori. Frequency comparison of optical lattice clocks beyond the Dick limit. Nature Photonics, 5:288–292, 2011.
- [114] R. Le Targat, L. Lorini, Y. Le Coq, M. Zawada, J. Guéna, M. Abgrall, M. Gurov, P. Rosenbusch, D.G. Rovera, B. Nagórny, R. Gartman, P.G. Westergaard, M.E. Tobar, M. Lours, G. Santarelli, A. Clairon, S. Bize, P. Laurent, P. Lemonde, and J. Lodewyck. Experimental realization of an optical second with strontium lattice clocks. Nature Communications, 4:1–8, 2013.

- [115] W.L. Tew, T.L. Nicholson, and R.B. Hutson. Calibration of thin-film platinum sensors for use in the JILA Sr II clock. National Institute of Standards and Technology Internal Report, 2015. <http://dx.doi.org/10.6028/NIST.IR.8046>.
- [116] I. Ushijima, M. Takamoto, M. Das, T. Ohkubo, and H. Katori. Cryogenic optical lattice clocks. Nature Photonics, 9:185–189, 2015.
- [117] J.T. Verdeyen. Laser Electronics. Prentice Hall, 3rd edition, 1995.
- [118] C. Vian, P. Rosenbusch, H. Marion, S. Bize, L. Cacciapuoti, S. Zhang, M. Abgrall, D. Chambon, I. Maksimovic, P. Laurent, G. Santarelli, A. Clairon, A. Luiten, M. Tobar, and C. Salomon. BNM-SYRTE fountains: recent results. IEEE Transactions on Instrumentation and Measurement, 54:833–836, 2005.
- [119] P.G. Westergaard. Strontium Optical Lattice Clock: In Quest of the Ultimate Performance. PhD thesis, Observatoire de Paris, LNE-SYRTE, 2010.
- [120] P.G. Westergaard, J. Lodewyck, and P. Lemonde. Minimizing the Dick effect in an optical lattice clock. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 57:623–628, 2010.
- [121] P.G. Westergaard, J. Lodewyck, L. Lorini, A. Lecallier, E.A. Burt, M. Zawada, J. Millo, and P. Lemonde. Lattice-induced frequency shifts in Sr optical lattice clocks at the 10^{-17} level. Physical Review Letters, 106:210801, 2011.
- [122] D.R. White and C.L. Jongenelen. The immersion characteristics of industrial prts. International Journal of Thermophysics, 31:1695–1695, 2010.
- [123] D.J. Wineland and W.M. Itano. Laser cooling of atoms. Physical Review A, 20:1521–1540, 1979.
- [124] X. Xu, T.H. Loftus, J.L. Hall, A. Gallagher, and J. Ye. Cooling and trapping of atomic strontium. Journal of the Optical Society of America B, 20:968–976, 2003.
- [125] J. Ye, H.J. Kimble, and H. Katori. Quantum state engineering and precision metrology using state-insensitive light traps. Science, 320:1734–1738, 2008.
- [126] B.C. Young, F.C. Cruz, W.M. Itano, and J.C. Bergquist. Visible lasers with subhertz linewidths. Physical Review Letters, 82:3799–3802, 1999.
- [127] X. Zhang, M. Bishof, S.L. Bromley, C.V. Kraus, M.S. Safronova, P. Zoller, A.M. Rey, and J. Ye. Spectroscopic observation of SU(N)-symmetric interactions in Sr orbital magnetism. Science, 345:1467–1473, 2014.

Appendix A

Mathematical Relations Not Derived in the Main Text

A.1 Two Time Signals Related by a Transfer Function

Two time signals $x_1(t)$ and $x_2(t)$ have corresponding Fourier transforms $X_1(f)$ and $X_2(f)$, respectively. The frequency functions are related by a transfer function $H(f)$, such that $X_2(f) = H(f)X_1(f)$. Equivalently, according to the convolution theorem, $x_2(t) = h(t) * x_1(t)$, where $*$ denotes a convolution, and $h(t)$ is the inverse Fourier transform of $H(f)$. The cross correlation $R_{12}(\tau)$ between $X_1(f)$ and $X_2(f)$ is

$$\begin{aligned} R_{12}(\tau) &= \langle x_1(t)x_2(\tau+t) \rangle = \int_{-\infty}^{\infty} h(t') \langle x_1(t)x_1(t+\tau-t') \rangle dt' \\ &= \int_{-\infty}^{\infty} h(t') R_1(\tau-t') dt' = h(\tau) * R_1(\tau), \end{aligned} \quad (\text{A.1})$$

where $R_1(\tau)$ is the autocorrelation function of $x_1(t)$. Furthermore,

$$\begin{aligned} R_2(\tau) &= \langle x_2(t)x_2(t-\tau) \rangle = \int_{-\infty}^{\infty} h(t') \langle x_1(t-\tau-t')x_2(t) \rangle dt' \\ &= \int_{-\infty}^{\infty} h(t') R_{21}(-\tau-t') dt' = h(-\tau) * R_{21}(-\tau) = h(-\tau) * R_{12}(\tau), \end{aligned} \quad (\text{A.2})$$

where the cross correlation identity $R_{21}(-\tau) = R_{12}(\tau)$ was used. Putting equations A.1 and A.2 together, the relationship between the autocorrelation functions of $x_1(t)$ and $x_2(t)$ is

$$R_2(\tau) = h(-\tau) * h(\tau) * R_1(\tau). \quad (\text{A.3})$$

According to the Wiener-Khinchin theorem, the power spectral density (PSD) is the Fourier transform of the autocorrelation function,

$$\begin{aligned} S_2(f) &= \mathcal{F}[R_2(\tau)] = \mathcal{F}[h(-\tau)] \cdot \mathcal{F}[h(\tau)] \cdot \mathcal{F}[R_1(\tau)] \\ &= |H(f)|^2 S_1(f). \end{aligned} \quad (\text{A.4})$$

where \mathcal{F} denotes a Fourier transform. This is the desired relation.

A.2 Relating the Allan Deviation to the Power Spectral Density

The Allan deviation is given by

$$\sigma^2(\tau) = \frac{1}{2} E \left[\left(\frac{1}{\tau} \int_t^{t-\tau} y(t') dt' - \frac{1}{\tau} \int_{t-\tau}^t y(t') dt' \right)^2 \right], \quad (\text{A.5})$$

where $y(t)$ is the fractional frequency fluctuation of the local oscillator and $E[\dots]$ is an expectation value. The integral terms can be combined in impulse-response form,

$$\sigma^2(\tau) = \left[\left(\int_{-\infty}^{\infty} h_\tau(t') y(t' - t) dt' \right)^2 \right], \quad (\text{A.6})$$

where $h_\tau(t)$ is an asymmetric window filter:

$$h_\tau(t) = \begin{cases} -\frac{1}{\sqrt{2}\tau}, & -\tau < t < 0 \\ \frac{1}{\sqrt{2}\tau}, & 0 < t < \tau \end{cases}. \quad (\text{A.7})$$

The transfer function of this filter is:

$$\begin{aligned} H_\tau(f) &= \int_{-\infty}^{\infty} h_\tau(t) e^{-2\pi i f t} dt = \frac{1}{\sqrt{2}\tau} \int_0^\tau e^{-2\pi i f t} dt - \frac{1}{\sqrt{2}\tau} \int_{-\tau}^0 e^{-2\pi i f t} dt = \\ &= \sqrt{2} \frac{\sin^2(\pi f \tau)}{i \pi f \tau}. \end{aligned} \quad (\text{A.8})$$

Let $f(t) = \int_{-\infty}^{\infty} h_\tau(t') y(t' - t) dt'$,

$$\begin{aligned} \sigma^2(\tau) &= \langle f^2(t) \rangle = \langle f(t) f(t - \tau) \rangle |_{\tau \rightarrow 0} = R_f(0) \quad (\text{autocorrelation}) \\ &= \int_{-\infty}^{\infty} S_f(f) \left[e^{2\pi i f \tau} \right]_{\tau \rightarrow 0} df \\ &= \int_{-\infty}^{\infty} S_f(f) df = \int_{-\infty}^{\infty} |H_\tau(f)|^2 S_y(f) df = 2 \int_{-\infty}^{\infty} \frac{\sin^4(\pi f \tau)}{(\pi f \tau)^2} S_y(f) df, \end{aligned} \quad (\text{A.9})$$

where the relation developed in Appendix A.1 has been utilized. Here $S_y(f)$ is the PSD of $y(t)$.

Note that $S_y(f)$ is the two-sided PSD. It is convenient to define a one-sided PSD,

$$S^{1-sided}(f) = 2S^{2-sided}(f), \quad \text{for } 0 \leq f < \infty. \quad (\text{A.10})$$

This factor of 2 is chosen so that both spectral densities give the total power when integrated over all frequencies for which these functions are defined,

$$\begin{aligned} \text{Total power} &= \int_{-\infty}^{\infty} S^{2-sided}(f) df = 2 \int_0^{\infty} S^{2-sided}(f) df \quad (\text{since PSD is even}) \\ &= \int_0^{\infty} S^{1-sided}(f) df. \end{aligned} \quad (\text{A.11})$$

Therefore

$$\sigma^2(\tau) = 2 \int_0^{\infty} \frac{\sin^4(\pi f \tau)}{(\pi f \tau)^2} S_y^{1-sided}(f) df. \quad (\text{A.12})$$

A.3 Power Spectral Density of Discrete Time Signals

Suppose that $y(t)$ is a stationary time signal. This function is sampled with a period of T_c . For the n th cycle, the discrete time signal autocorrelation $R(nT_c)$ is defined as

$$R(nT_c) = \sum_{m=-\infty}^{\infty} y(mT_c) y[(m-n)T_c]. \quad (\text{A.13})$$

The PSD of the discrete signal $S_{dis}(f)$ is given by the discrete time Fourier transform of $R(nT_c)$, defined as

$$S_{dis}(f) = T_c \sum_{n=-\infty}^{\infty} R(nT_c) e^{-2\pi i f T_c n}. \quad (\text{A.14})$$

Let $R(\tau)$ be the autocorrelation of the continuous signal $y(t)$. Then let $S_{con}(f)$ be the PSD of $y(t)$, obtained via a continuous Fourier transform \mathcal{F} of $R(\tau)$. It follows that

$$\begin{aligned}
S_{dis}(f) &= T_c \sum_{n=-\infty}^{\infty} R(nT_c) e^{-2\pi i f T_c n} = T_c \sum_{n=-\infty}^{\infty} R(nT_c) \mathcal{F}[\delta(t - nT_c)] \\
&= \mathcal{F} \left[T_c \sum_{n=-\infty}^{\infty} R(nT_c) \delta(\tau - nT_c) \right] = \mathcal{F} \left[R(\tau) \sum_{k=-\infty}^{\infty} e^{-2\pi i k \tau / T_c} \right] \\
&= \sum_{k=-\infty}^{\infty} \mathcal{F} \left[R(\tau) e^{-2\pi i k \tau / T_c} \right] = \sum_{k=-\infty}^{\infty} S_{con} \left(f + \frac{k}{T_c} \right). \tag{A.15}
\end{aligned}$$

where $\delta(t)$ is the Dirac delta function. Here I have used the well-known relation between a sum over delta functions and a sum over complex exponentials. This is the desired relation between $S_{dis}(f)$ and $S_{con}(f)$.

A.4 Detection Noise

Suppose that the counts measured on the PMT after fluorescing the counting laser (Section 3.7.2) are C_e , C_g , and C_{bg} , which are the excited, ground, and background counts. Also, $V_e = C_e - C_{bg}$ and $V_g = C_g - C_{bg}$ (Section 3.7.2), and each are proportional to their corresponding atom numbers by a factor of β . The measured excited state fraction is

$$\rho_{meas} = \frac{C_e - C_{bg}}{C_g + C_e - 2C_{bg}}, \tag{A.16}$$

where the total atom number N is proportional to $C_g + C_e - 2C_{bg}$.

The detection noise is the number of counts measured when the atomic beam and lasers that cause background counts are blocked. The observed detection noise is $\delta N = 11.4$ atoms per measurement of C_e , C_g , and C_{bg} . Propagating this error into ρ_{ee} , the uncertainty due to detection noise is

$$\sqrt{\left(\frac{\partial \rho_{ee}}{\partial C_e} \beta \delta N \right)^2 + \left(\frac{\partial \rho_{ee}}{\partial C_g} \beta \delta N \right)^2 + \left(\frac{\partial \rho_{ee}}{\partial C_{bg}} \beta \delta N \right)^2} = \frac{\delta N}{\sqrt{2N}}. \tag{A.17}$$

This assumes that the clock is locked to $\rho_{meas} = 1/2$.

Just like in the case of QPN, this fluctuation in ρ_{meas} causes frequency instability in the clock laser. Using arguments similar to those in Section 4.2.1, the frequency stability due to detection noise is

$$\sigma_{det} = \frac{\delta N}{\sqrt{2} N s \nu_{Sr}} \sqrt{\frac{T_c}{\tau}}. \quad (\text{A.18})$$

Therefore, detection noise scales like $1/N$, QPN scales like $1/\sqrt{N}$, and the two stabilities cross one another when $N = 2 \delta N^2$. For atom numbers less than $2 \delta N^2$, detection noise dominates over QPN.