A Transportable Ytterbium Optical Lattice Clock

by

R. J. Fasano

B.S., University of Maine, 2015

M.S., University of Colorado, 2018

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
2021

Committee Members:
Scott Diddams, Chair
Andrew Ludlow
R. Steven Nerem
Elizabeth Donley
Tara Fortier

Fasano, R. J. (Ph.D., Physics)

A Transportable Ytterbium Optical Lattice Clock

Thesis directed by Dr. Scott Diddams

In the last decade, optical clocks have overtaken conventional microwave frequency standards, achieving performance levels of 10^{-18} fractional frequency uncertainty and 10^{-16} fractional instability at one second of averaging time. A roadmap has been defined for redefinition of the S.I. second in terms of an optical frequency, but the required demonstrations of agreement between multiple optical lattice clocks have not yet been satisfied, owing to the large geographical distance between various international metrological institutes. We report on the development of a ^{171}Yb transportable optical lattice clock which can facilitate indirect frequency comparisons between distant clocks, targeting state-of-the-art performance in a rack-mounted form factor with low size, weight, and power. We also investigate several key systematic effects which are especially relevant to a transportable clock, outlining a path to 10^{-18} uncertainty.

Dedication

To all the friends and family who supported me along the way.

Acknowledgements

This project would not be possible without the contributions of countless friends and colleagues at NIST, and I can only name a few here. First, I thank Andrew Ludlow for giving me an opportunity to jump into such a fascinating field and involving himself closely in my development as a scientist. Chris Oates played an invaluable role in guiding me into the field of timekeeping and urging me to explore my ideas.

I thank the many people who have worked on the main Yb clock experiments; without their accumulated expertise, the ambitious task of building a transportable system would be insurmountable. I enjoyed working directly with Will McGrew, Nathan Hinkley, Marco Schioppo, Gianmaria Milani, Daniele Nicolodi, and Xiaogang Zhang through long clock measurements in my early years of grad school. Although my path diverged towards transportable applications, I have also greatly enjoyed working with new additions to the main clock team: Youssef Hassan, Jacob Siegel, Chun-Chia Chen, and Tanner Grogan.

Though our portable clock team was small, it was filled with formidable members who left resounding impacts on the experiment. Rich Fox has provided endless electronics wizardry and an adept ability to navigate bureaucracy. Yun-Jhih Chen brought new perspectives and a contagious happiness that bolstered us through tight deadlines. And Wesley Brand has contributed immense technical expertise, a willingness to get his hands dirty, and steadfast friendship.

Outside of our Yb team, I am grateful to Frank Quinlan and Megan Kelleher for aiding with the development of optical cavity systems, and to Craig Nelson, Archita Hati, and Marco Pomponio for providing frequency stabilization electronics and a transportable frequency comb.

I am also grateful to my family for the loving support along the many steps of this journey, and for instilling a scientific curiosity in me at a young age. This curiosity was supported by many teachers along the way: Mr. Choate and Mrs. Haskell in high school and Professor Michael Wittmann in college to name just a few. I thank Professor Neil Comins at UMaine for giving me early research experience in computational astrophysics, which has remained surprisingly useful to this day, and to Professor Murray Holland at CU for introducing me to the world of cold-atom physics during a summer internship.

And lastly, but perhaps most importantly, I am grateful to Alexa, who was there for every step of grad school, from late-night problem sets to long writing sessions. None of it would be possible with you.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction

1.1 Measurements through time

The system of units used by a society is a reflection of both its culture and its scientific progress. For pre-industrial societies, units of measurement were derived from the natural world: the duration of a day, the weight of a cereal seed, or the length of one's foot. While crude by modern standards, these units were sufficient for the daily tasks of agrarian life. As civilization grew, standard systems of weights and measures emerged, largely driven by trade and commerce. The age of industrialization drove a new demand towards standardization, motivated in part by the advent of interchangeable parts, and culminating in the development of physical standards representing the meter and the kilogram at the end of the 18th century in France. Around the same time, the invention of the marine chronometer allowed seafaring navigators to accurately determine their longitude, an extremely difficult task with the pendulum clocks of the time. This was the era of humanity's ascendancy from nature, as the peasant class traded fields and pastures for cities and factories, and as new timekeeping devices built from gears and springs superseded archaic measures derived from Earth's rotation or gravity.

This new hegemony was upended by the era of modern physics, ushered in by the twin discoveries of quantum mechanics and relativity. The notion that time is inextricably connected to mass and length through the universality of the speed of light re-anchored our system of units in the physical universe. In 1983, 78 years after Einstein's discovery of relativity, this concept was embedded into the SI system through the definition of the speed of light as *exactly* 299,792,458 m/s. An important consequence was the redefinition of the meter to

> the length of the path travelled by light in vacuum during a time interval of 1/299,792,458 of a second,

which was required for consistency with the 1967 definition of the second as

> the time duration of 9 192 631 770 periods of the radiation corresponding to the transition between the two hyperfine levels of the fundamental unperturbed ground-state of the caesium-133 atom.

Thus, our system of units came full circle, beginning with materials and cyclical processes from the natural world, transitioning temporarily to man-made artifacts, and returning at last to quantities as fundamental and universal as the speed of light and the "ticking" of atoms.

In 2019, the International Committee of Weights and Measures (CIPM) followed in the footsteps of the 1983 decision, abandoning the physical mass standard (the last unit based on a physical artifact) and instead realizing the standard electromagnetically through a device called the Kibble balance. At the same time, the entire SI system was redefined in terms of seven exact quantities:

- The Planck constant, $h = 6.62607015 \times 10^{-34}\ \mathrm{J \cdot S}$
- The elementary electric charge, $e = 1.602176634 \times 10^{-19}\ \mathrm{C}$
- The Boltzmann constant, $k_B = 1.380649 \times 10^{-23}\ \mathrm{J \cdot K^{-1}}$
- The Avogadro constant, $N_A = 6.02214076 \times 10^{23}\ \mathrm{mol^{-1}}$
- The speed of light, $c = 299792458\ \mathrm{m \cdot s^{-1}}$
- The ground state hyperfine transition frequency of the ^{133}Cs atom, $\Delta\nu_{\mathrm{Cs}} = 9192631770\ \mathrm{Hz}$
- The luminous efficacy of monochromatic radiation of frequency 540 THz, $K_{\mathrm{cd}} = 683\ \mathrm{lm \cdot W^{-1}}$

The promotion of an atomic transition frequency to equal footing with fundamental quantities like the Planck constant attests to the profound insight of William Thomson, who recognized in 1897 (following a conversation with James Clerk Maxwell) that atoms are “absolutely alike in every physical property” and “probably remain the same so long as the particle itself exists” [97]. This is not to say that atomic clocks are infallible, but simply that they offer a realization of the second with much greater accuracy and consistency than any other physical standard. For example, state-of-the-art cesium fountain clocks have demonstrated agreement at the level of several parts in 10^{16}, whereas an ensemble of replicas of International Prototype of the Kilogram (which served as the official mass standard until 2019) varied by greater than 2 parts in 10^{8} over the last century [33]. A deliberate advantage of the new SI system is that all but one of the seven base units (the exception being the mol) is directly traceable to the definition of the second [25].

At the time of this writing, we straddle another great step forwards - towards the redefinition of the SI second in terms of an optical transition frequency. Whereas microwave clocks have demonstrated agreement at the 10^{-16} level after weeks of averaging, optical lattice clocks can reach this agreement in seconds. Additionally, optical lattice clocks have demonstrated agreement below the 10^{-18} level after several days of averaging. Other clock architectures employing optical transitions have also demonstrated exceptionally high performance. Clocks based on trapped ions have achieved uncertainties below 10^{-18} [15]; though their instability is typically handicapped by quantum projection noise due to low atom number, order-of-magnitude improvements can be attained by bolstering interrogation time with the aid of an optical lattice clock [51]. Another novel architecture based on optical tweezer arrays has recently demonstrated instability comparable to state of the art optical lattice clocks [77, 113]. Spurred by promising results from optical clocks, the CIPM laid out a roadmap towards redefinition in 2016 [87], requiring the following milestones to be met:

- Three different optical clocks (either in different laboratories or of different atomic species) estimating fractional frequency uncertainty at the 10^{-18} level

- Three independent optical clock comparisons at the $< 5 \times 10^{-18}$ level
- Frequency ratios between at least five optical frequency standards agreeing at the $< 5 \times 10^{-18}$ level, each ratio measured at least twice by independent laboratories
- Three independent measurements versus three independent Cs primary clocks at the 3×10^{-16} level
- Regular contribution to International Atomic Time (TAI)

These five requirements can be directly associated with several key performance metrics for clocks, which are introduced in the following sections.

1.1.1 Uncertainty

The first requirement corresponds to the clock *uncertainty*, which is typically estimated following a rigorous evaluation of various systematic effects which can alter the clock frequency. We recently carried out a full evaluation of our both of our laboratory clocks, Yb1 and Yb2 [69], culminating in a fractional frequency uncertainty of 1.4×10^{-18} (Table 1.1). This evaluation relied on substantial previous efforts to investigate three of our most relevant systematic effects.

The first effect is the ac Stark shift of the lattice laser, which can be mitigated by operating at a *magic wavelength* where the shifts of the ground and excited states of the clock transition are equal [49]. By studying higher-order effects such as M1/E2 polarizability, hyperpolarizability, lattice anharmonicity, and shift averaging due to thermal motion of lattice-trapped atoms, we constrained the lattice Stark uncertainty at the 10^{-18} level [16]. The second effect is the dc Stark shift due to stray charges on the enclosure or static external fields, which we mitigate with an in-vacuum Faraday cage and constrain the shift below the 10^{-20} level with *in situ* measurements [7]. Lastly, the blackbody radiation (BBR) shift due to room-temperature radiation bathing the atoms, which is typically at the order of 10^{-15} of the clock frequency, can be compensated with frequency corrections based on temperature measurements of the vacuum enclosure [6].

Table 1.1: Uncertainty budget for our laboratory clocks, containing all known systematic effects which are relevant at the 10^{-18} level [69]. All numbers are in units of $10^{-18}\nu_{\text{clock}}$.

	Yb1		**Yb2**	
Effect	**Shift**	**Uncertainty**	**Shift**	**Uncertainty**
Background gas collisions	-5.5	0.5	-3.6	0.3
Spin polarization	0	0.3	0	0.1
Cold collisions	-0.21	0.07	-0.02	0.01
Doppler	0	0.02	0	0.01
BBR	-2,361.2	0.9	-2,371.7	1
Lattice light (model)	0	0.3	0	0.3
Travelling wave contamination	0	0.1	0	0.01
Lattice light (experimental)	-1.5	0.8	-1.5	0.8
Second-order Zeeman	-118.1	0.2	-117.9	0.1
DC Stark	0	0.07	0	0.04
Probe Stark	0.02	0.01	0.02	0.01
Line pulling	0	0.1	0	0.1
Tunnelling	0	0.001	0	0.001
Servo error	0.03	0.05	0.03	0.05
Optical frequency synthesis	0	0.1	0	0.1
Total	-2,486.5	1.4	-2,494.7	1.4

1.1.2 Reproducibility

The second and third requirements correspond to the *reproducibility*, or the extent to which separate realizations of optical frequency standards agree. Reproducibility is sometimes assessed during an uncertainty characterization; for example, in [69], we carried out 10 blinded comparisons between Yb1 and Yb2, measuring a frequency difference of $-7 \pm (5)_{\text{stat}} \pm (8)_{\text{sys}} \times 10^{-19}$ after accounting for all systematics. The statistical uncertainty is assessed through the Allan deviation, which describes the fractional differential frequency variation as a function of averaging time (Figure 1.1). The systematic uncertainty in the clock comparison is smaller than our single-clock uncertainty of 1.4×10^{-18} due to the common-mode cancellation of several effects.

Reproducibility characterizations are absolutely essential for lending credibility to optical standards. As we noted in [69],

> After four early comparisons, not included in the dataset presented here, it was discovered that a faulty wire had removed the grounding connection to the conductive windows of Yb-2, compromising Faraday shielding and leading to a mid-10^{-18} DC

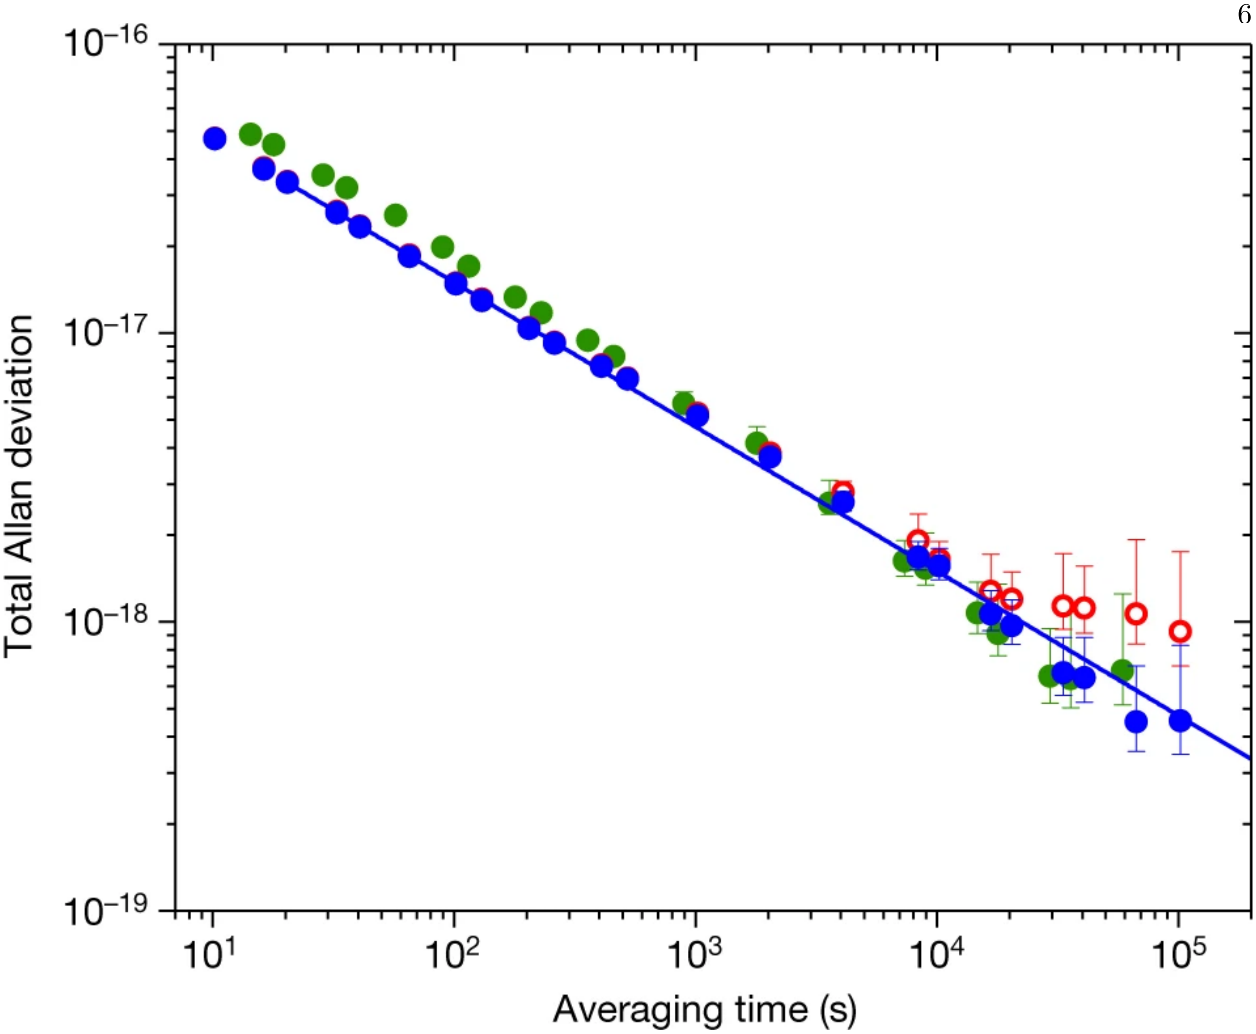

Figure 1.1: Allan deviation for the Yb1-Yb2 comparison, representing an upper bound on the single-clock instability for several different experimental configurations. Blue points: synchronized Rabi spectroscopy. The blue line is a fit to a white noise asymptote of $1.5 \times 10^{-16}/\sqrt{\tau}$. Green points: unsynchronized Ramsey spectroscopy. Red points: same as blue data, but not including corrections for blackbody radiation shifts assessed through real-time temperature measurements. Reproduced from [69].

> Stark shift between the systems. This experience underscores the indispensability of experimentally investigating reproducibility for substantiating an uncertainty budget.

While this example of a grounding error is a simple one, it illustrates a more general principle that uncertainty evaluations in isolation have their limits, and frequency measurements demonstrating reproducibility are critical.

Frequency ratio measurements between different atomic species and types of clock allow even

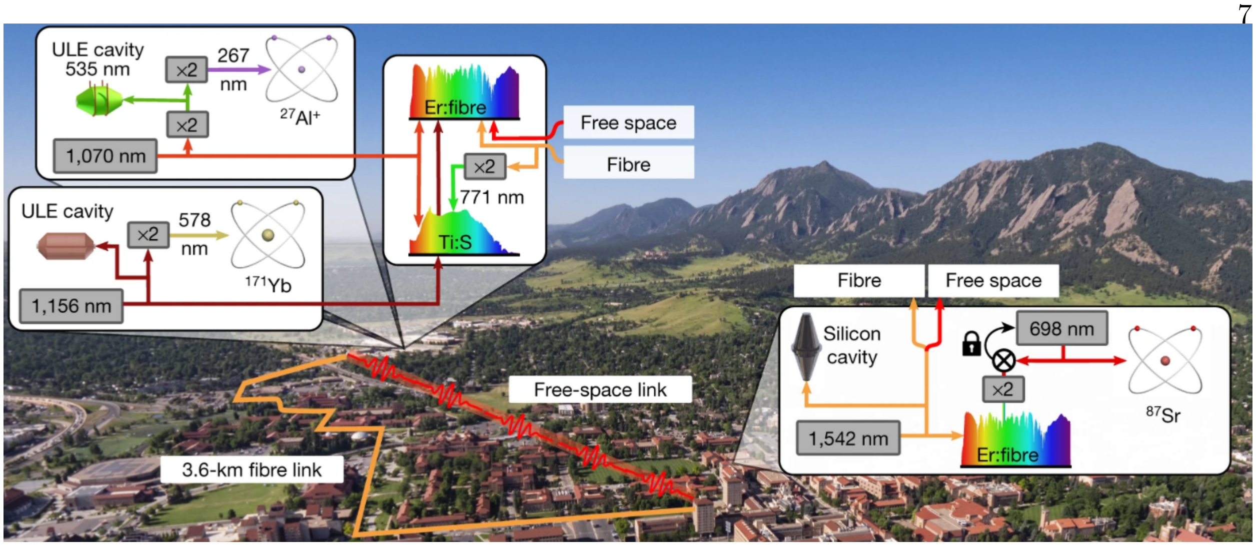

Figure 1.2: Boulder Atomic Clock Optical Network (BACON) utilized for frequency ratio measurements between ^{171}Yb, ^{27}Al+, and ^{87}Sr clocks. Reproduced from [22].

greater confidence in the reproducibility. Two Yb standards in the same lab may agree despite the omission of some unknown systematic effect which acts on each clock identically. In contrast, when comparing a Yb optical lattice clock to an Al$^+$ ion clock, the largest systematic effects can be very different: ion clocks are susceptible to large shifts due to micromotion of the ion within the trap, whereas optical lattices can cause large ac Stark shifts of the clock transition. Demonstrating agreement between two very different frequency standards can therefore inspire more confidence than agreement of nominally identical systems.

Seeking such validation for our clocks, we compared our Yb clocks against an ^{27}Al$^+$ ion clock at NIST and a ^{87}Sr optical lattice clock at JILA, both of which have been independently evaluated at the 10^{-18} uncertainty level [15, 14]; these three clocks currently possess the lowest uncertainty of any in the world. The comparisons utilized an existing optical fiber network in Boulder, CO, as well as a novel two-way time and frequency transfer (TWTFT) technique over a 1.5 km free-space link [22], shown in Figure 1.2. Over a series of comparisons between November 2017 and June 2018, we measured frequency ratios between the three atomic species with uncertainties of 5.9×10^{-18}, 8.0×10^{-18}, and 6.8×10^{-18} for Al$^+$/Yb, Al$^+$/Sr, and Yb/Sr respectively.

1.1.3 Inaccuracy

The fourth requirement corresponds to what is sometimes called *inaccuracy*, not to be confused with uncertainty. Inaccuracy pertains to the agreement between optical standards and the current SI definition of a second; it is absolutely imperative that the new definition be entirely consistent with the old one, regardless of what devices are used to realize it. Independent measurements between optical and Cs primary clocks at the level of uncertainty of the microwave standards ensure this continuity. We recently compared our Yb clocks against an international ensemble of primary and secondary frequency standards (PSFS) using a two-way satellite link [70]. The signal chain for these measurements is shown in Figure 1.3. Our 1156 nm clock laser, stabilized to the $^1S_0 \rightarrow {}^3P_0$ clock transition, served as the reference foran octave-spanning, self-referenced Ti:S frequency comb and downconverted to a repetition rate near 1 GHz. This microwave signal was mixed with a 1 GHz signal derived from a hydrogen maser, and the resulting frequency difference ($\approx$300 kHz) was transferred to a maser time scale at NIST which serves as a flywheel oscillator during intermittent periods of optical clock operation. This local hybrid microwave-optical timescale, which is additionally described in [112], was transferred to the BIPM using the hybrid Two-Way Satellite Time and Frequency Transfer/GPS Precise Point Positioning (TWGPPP) frequency transfer protocol [47], allowing comparison between our local timescale and PSFS.

Over an eight-month campaign consisting of 79 measurements, each with an average duration of 4.9 hours, we measured the absolute clock transition frequency to be 518,295,836,590,863.71(11) Hz. The uncertainty of 2.1×10^{-16} was dominated by measurement dead time and PSFS type B uncertainty, and is sufficient to satisfy the fourth roadmap requirement. For an additional consistency check, we used these results to compute a “loop misclosure” between Yb, Sr, and Cs standards, comparing ratios of absolute frequency measurements of Yb and Sr against Cs with a direct Yb/Sr optical frequency ratio measurement. We found a misclosure of $(0.8 \pm 2.4) \times 10^{-16}$, consistent with zero at a level limited by the microwave uncertainty.

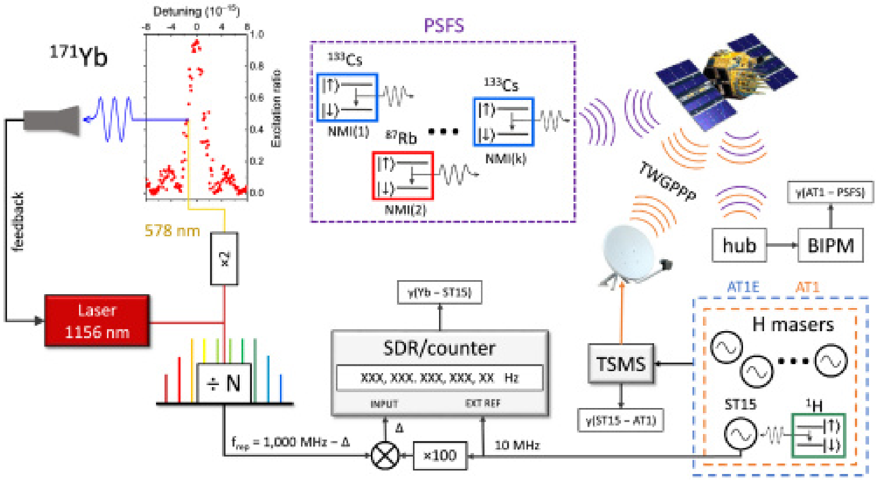

Figure 1.3: Frequency transfer chain from our Yb optical lattice clock to an international ensemble of primary and secondary frequency standards (PSFS) via satellite transfer. Reproduced from [70].

1.1.4 Maturity

The final requirement, regular contribution to International Atomic Time, is simply a statement of maturity. While optical clocks have existed for over a decade, they have mostly been engaged in iterative cycles of self-calibration and fundamental research rather than operating as a *clock* in the colloquial sense: a device which is used to tell time (as opposed to more esoteric pursuits like, for example, relativistic geodesy or searches for dark matter). Beyond absolute performance metrics like uncertainty and instability, clock operation requires a high uptime and robustness. It is in this arena, supporting development of an optical timescale [112] as well as our satellite comparisons against PSFS [70], that I first developed the skills and controls systems which would later be applied to operate our transportable clock with a high degree of robustness. These systems comprised a) a monitoring system which detects out-of-bounds experimental parameters and flags affected data to be rejected during post-processing and b) an algorithm to automatically relock our lattice laser to its reference cavity. Implementation of these systems allowed an average uptime of 75% measurement campaign in [70]. Later efforts supporting the transportable clock project extended the auto-relocking to all five lasers (Section 2.10) and incorporated automated alignment using high-bandwidth MEMS mirrors to compensate for beam pointing drift at critical points (Section 2.11.4).

Using the hybrid microwave-optical timescale and satellite transfer system described in the previous section, we have begun formally contributing to International Atomic Time (TAI), from which Universal Coordinated Time (UTC) is derived. Beyond a mere proof-of-concept, the high stability of our laboratory clocks allows us to contribute to global timekeeping with extremely competitive uncertainty despite only operating in this capacity for around 30 minutes per week on average.

1.1.5 Towards redefinition: outlook and remaining challenges

A cursory reading of the previous sections may give the impression that we are well poised to satisfy the requirements for the redefinition by 2026, when the CIPM will reconvene to decide whether a redefinition is justified yet. Indeed, during my doctoral research, our group and collaborators contributed towards satisfying each of the five requirements: demonstrating uncertainty in the 10^{-18} decade for both single clocks and clock comparisons [69], measuring frequency ratios against two other optical clocks [22], comparing with primary Cs standards [70], and contributing to International Atomic Time [112]. However, it cannot be denied that Boulder, Colorado is uniquely privileged to host state-of-the-art frequency standards across multiple atomic species. Many diverse high-performance optical clocks exist internationally: the Yb octopole clock at PTB in Braunschweig, Germany [44], a cryogenic Sr lattice clock at RIKEN in Tokyo [102], and a Sr lattice clock at NPL in Teddington, England [42] to name just a few. However, comparisons between geographically-distant institutes cannot be realized easily without compromising the uncertainty of these systems; satellite time and frequency transfer incurs an uncertainty at levels comparable to microwave clocks, easily spoiling the performance advantage of optical clocks, while optical fiber networks allow low frequency transfer uncertainty but with typically prohibitively high setup costs. It is unlikely that the third requirement, optical frequency ratios between at least five standards, will be satisfied using these techniques alone due to the sheer logistical challenges.

A more realistic approach is to construct transportable clocks which can serve as intermediate frequency standards enabling indirect comparisons between optical clocks around the world, as well as finding usefulness in numerous other applications. The focus of this dissertation is the development, construction, and early evaluation of a transportable optical lattice clock based on ^{171}Yb at NIST. While constructing an optical lattice clock alone is no easy feat, designing towards transportability poses great additional challenges. In rising to these challenges, we have made diverse innovations which will benefit the entire field of optical timekeeping, including experimentation with compact epoxy-sealed ultra-high vacuum chambers supporting rigorous environmental control of systematic effects, characterization and elimination of ac Stark shifts arising from the broadband background spectra in amplified diode lasers, and advanced control systems capable of autonomous clock operation and fault recovery.

In the remainder of this chapter, I will provide a brief overview of optical lattice clocks[1] which will segue into an analysis of the challenges associated with building transportable systems and an overview of the current state of the field of transportable optical clocks. Lastly, I will discuss several interesting applications that our research will enable, from conventional timekeeping to probing of the fundamental physics of our universe.

1.2 Optical lattice clocks

Before understanding the challenges involved in transportable optical lattice clock development, we should understand what systems such a clock consists of. At the core of any clock is a periodic oscillator with a stable frequency - in our case, the $^1S_0 \rightarrow {}^3P_0$ transition of neutral ytterbium. This transition has a frequency $\nu_0 \approx 518$ THz (578 nm), nearly 60,000 times higher than the Cs microwave transition and about 10^{10} times the frequency of a typical quartz oscillator. This transition is called *doubly-forbidden* as it violates two electric dipole selection rules: first, it links singlet and triplet spin states; and second, both states have an angular momentum quantum number $J = 0$. As a consequence of the extremely weak transition strength between the ground

---

[1] See Chapter 2 for a more complete and technical discussion

and excited states[2] , the excited state has a lifetime exceeding 20 s, or equivalently, a linewidth $\Delta\nu < 10$ mHz. The natural quality factor of the transition is therefore $Q = \nu_0/\Delta\nu \approx 5 \times 10^{16}$, offering an exceptionally pure frequency reference. For comparison, the hyperfine clock transition in ^{133}Cs has been measured with Hz level linewidths [5], conferring a quality factor on the order of 10^10, and room-temperature quartz oscillators can reach quality factors on the order of 10^6 [31].

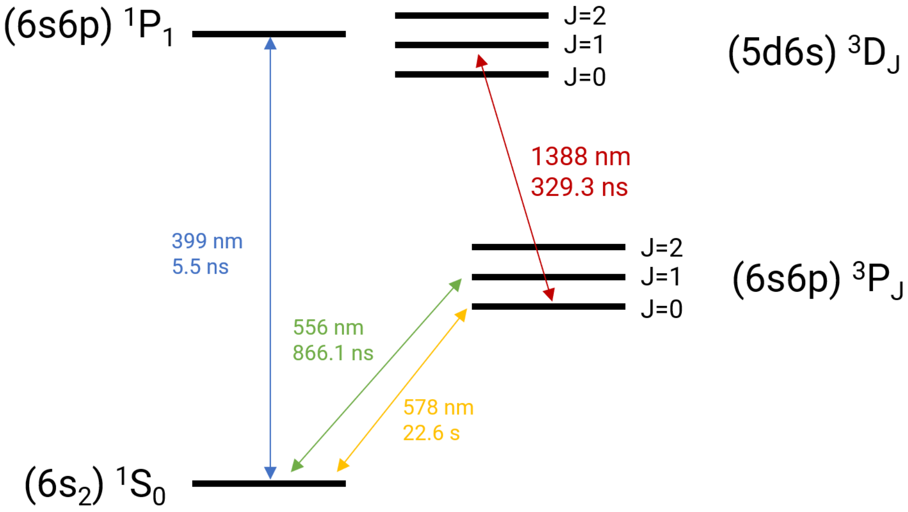

Figure 1.4: Atomic transition wavelengths and lifetimes used for a Yb optical lattice clock.

In order to confine the atoms during spectroscopy, a three-stage trapping process is employed, benefiting from several convenient, laser-accessible transitions which make Yb an appealing atom for timekeeping. First, a laser tuned near the strong, dipole-allowed $^1S_0 \rightarrow {}^1P_1$ transition at 399 nm is used to capture atoms from a thermal atomic beam and cool them to a Doppler limit of around 1 mK in a magneto-optical trap (MOT). Next, another laser at 556 nm is used for a second-stage MOT on the singly-forbidden $^1S_0 \rightarrow {}^3P_1$ transition, cooling the atoms to around 10 μK. Despite these low temperatures, Doppler broadening of the clock transition due to thermal motion would still result in a linewidth of about 30 kHz, greatly degrading the intrinsic quality factor. To hold the atoms nearly motionless during spectroscopy, the second-stage MOT is loaded into

---

[2] The transition strength is nonzero only due to internal hyperfine coupling of the nuclear spin in fermionic isotopes [85]. In bosonic isotopes, the transition strength can be tuned somewhat arbitrarily with the application of an external magnetic field.

an optical lattice operating in the Lamb-Dicke regime. In this regime, atomic motion is quantized in a harmonic potential with energy levels which are spaced much larger than the photon recoil energy, and quenched sideband cooling can be used to de-excite atoms to the motional ground state using the clock laser at 578 nm. In principle, the optical lattice can be tuned to any wavelength where both the ground and excited state polarizabilities are sufficiently high; however, it is common practice to tune to the so-called *magic wavelength* near 759 nm where the clock transition frequency is unperturbed by the intense lattice laser.

To exploit the ultra-high quality factor of the $^1S_0 \rightarrow {}^3P_0$ transition, we steer a laser, often referred to as a *local oscillator*, to track the clock transition. The laser is referenced to an optical cavity to assure coherence over typical spectroscopy durations from hundreds of ms to several seconds. In the simplest scheme, the local oscillator interacts with the cold atomic sample through *Rabi spectroscopy*, excitation of the cold atomic sample with a square (constant intensity) pulse. The excitation efficiency is assessed with a three-pulse sequence using the 399 nm laser to separately measure the ground- and excited-state populations as well as a background level, allowing a normalized excitation probability to be computed. This scheme utilizes an additional laser at 1388 nm to transfer the excited-state atoms to the $^3D_1$ state, where they decay back to the ground state quickly through the $^3P_1$ state. After repumping the $^3P_0$ population back to $^1S_0$ through this pathway, a subsequent measurement of fluorescence on the $^1S_0 \rightarrow {}^1P_1$ transition is used to measure the excited state population.

In order to realize a clock from this process, a *servo cycle* is implemented to determine frequency corrections for the clock laser after successive experimental cycles. At its simplest, a servo cycle consists of interrogation of the clock transition at a pair of equal and opposite detunings, with a frequency correction proportional to the difference between the excitation probabilities on the red- and blue-detuned sides of the transition.

In total, five lasers are required: the clock laser at 578 nm, the MOT lasers at 399 and 556 nm, the lattice laser at 759 nm, and the repump laser at 1388 nm. Nearly all of the equipment involved in optical lattice clocks pertains to the generation, manipulation, detection, and delivery of light from these five sources. Other necessary components include vacuum chambers for the atoms and optical cavities, electromagnetic coils to generate static magnetic fields or field gradients, and data acquisition, analysis, and control systems.

1.2.1 Transportable optical lattice clocks

Broadly speaking, there are two substantial areas of difficulty associated with the development of transportable optical lattice clocks. The first is *miniaturization*: reducing entire rooms full of electronics, optics, lasers, and vacuum systems to a compact package which can be easily transported is nontrivial[3] . Our two non-transportable clocks, Yb1 and Yb2, along with their optical cavities and laser systems, occupy approximately 50 m^2 of lab space in total. The clock vacuum chambers and surrounding optics each fill a volume of approximately 1 m^3. As of the writing of this dissertation, two transportable optical lattice clocks have been fully developed and deployed at labs outside of NIST. The first, developed by PTB [104], is built into a car trailer with an overall volume of 14.5 m^3, of which the physics package occupies about 0.6 m^3. The second, built by RIKEN [80], takes up 0.92 m^3 and weighs 370 kg. Our transportable clock, built into three server racks, targets a volume of approximately 1 m^3 and a maximum weight of 500 kg, though several miniaturization improvements are planned for a second generation as discussed in Chapter 5.

The second challenge is *robustness*. In a lab environment, we take for granted well-controlled temperatures, low levels of acoustic noise, and vibration damping granted by foot-thick, pneumatically floated optical tables; none of these are guaranteed when operating outside the lab. Indeed, PTB's deployment to a mountain road tunnel [35] was burdened by thermal hot spots inside the clock trailer, overheating of the lattice laser, and vibrations a factor of ten larger than typical laboratory conditions. The RIKEN group reported fewer challenges, perhaps owing to a more comfortable deployment inside a tower in Tokyo, though accelerations of up to $4g$ were experienced during transportation.

---

[3] Aside from the size and weight concerns, power consumption can be a significant concern for clocks operating in the field or in space. Even if the clock is to be deployed in a metrological institute or similar facility, it may be necessary to provide constant power during transportation to preserve the vacuum and temperature levels of the optical cavities and/or clock package.

In our transportable clock, substantial design efforts have been made in anticipation of robustness challenges. The physics package features a custom-patterned breadboard allowing direct mounting of optical mounts, and all optical path lengths are kept as short as possible to minimize beam misalignments. Laser modules requiring free-space optics were also designed with this philosophy, while several other laser modules are routed through optical fiber all the way from the laser to the physics package. In several critical locations, such as before optical fibers, MEMS mirrors are steered by optimization algorithms to maintain alignment in adverse conditions. For the clock laser, a rigidly-mounted optical cavity was designed with ultra-low acceleration sensitivity ($< 3.4 \times 10^{-11}/g$), along with several layers of thermal shields providing high isolation from the environment and extensive finite-element modeling to analyze the effects of thermal expansion of the cavity spacer and supports. Lastly, an extensive suite of control software and hardware has been implemented, supporting real-time monitoring of critical variables and automatic recovery from error conditions such as cavity unlocks or beam misalignments. As one early example, during time-of-flight temperature characterizations for the second stage MOT (Section 2.3), we operated the experiment remotely and autonomously for 27 hours with no human intervention, exploring temperature dependence on a three-dimensional grid of detuning, magnetic field gradient, and laser intensity. While operation outside of a lab is certain to be more difficult, the experimental uptime of the transportable system is already substantially better than the laboratory clocks, a testament to both the controls systems and the engineered robustness of the various laser systems.

1.3 Applications

Free from the constraints of a laboratory, we anticipate that optical clocks will find numerous additional applications beyond the scope of metrology. This section highlights just a few of these applications. It is noteworthy to mention that some of these applications such as clock networks for dark matter detection do not strictly require transportable clocks. However, our engineering efforts towards robustness during field deployment will also bolster the operation of non-transportable clocks for clock networks and optical timescales, as well as conventional laboratory clocks. Additionally, the modular optical systems described in Chapter 2 enable rapid construction of new clock systems with minimal alignment, an invaluable advantage towards the realization of optical timescales consisting of ensembles of many individual clocks.

1.3.1 Relativistic geodesy

The field of geodesy concerns itself with the measurement of the gravity potential[4] of Earth or other celestial bodies and the precise determination of the *geoid*, an equipotential surface close to sea level. The geoid serves as a common baseline for gravitational measurements, which are conventionally achieved through spirit levelling, a laborious point-to-point procedure which can accumulate sizeable errors over long distances. Even worse, the geoid is partially referenced with tide gauges to oceanographic sea level, which itself changes by at the level of millimeters per year due to ice mass fluctuations [99] and crustal motion [74].

More recently, optical clocks have reached uncertainties where *chronometric levelling* can match the accuracy of spirit levelling combined with gravity measurements. This technique entails the measurement of a relativistic frequency shift

$$\frac{\Delta\nu}{\nu_0} = \frac{\Delta U}{c^2}, \tag{1.1}$$

where ν_0 is the unperturbed frequency, c is the speed of light, and ΔU is the difference in geopotential relative to a reference clock. Near the surface of the Earth, this fractional frequency shift is of order 10^{-16}/m; an interesting consequence is that the denizens of Boulder, CO age faster than their friends at sea level by about 5 microseconds per year!

An important consequence for metrologists is that our meticulously prepared uncertainty budgets can be easily dwarfed by the geopotential shifts from clocks at different heights. Thus, it is important that clocks operate at heights which have been carefully measured relative to the geoid, a service which is provided domestically by the National Geodetic Survey. These considerations

---

[4] Geodecists often differentiate between *gravitational* fields arising from the mass of the Earth and *gravity* fields, which include centrifugal acceleration due to the rotation of the Earth. Here, I will use the term *geopotential* to include both effects.

are especially critical for the establishment of international timescales such as International Atomic Time (TAI): the frequency of each clock in the ensemble must be corrected to account for its height relative to the geoid.

Although the static geoid can be modeled globally with an uncertainty of several parts in 10^{18}, dynamic effects such as tidal variation or other time-varying mass redistribution (e.g. melting of polar ice caps) are less well-studied [75]. As the next generations of optical clocks forge ahead into the 10^{-19} decade, it is possible that the overall accuracy of optical timescales will be solely limited by geophysical phenomena. It has been proposed that the next generation of clocks could avoid these effects by operating in space [17]. Alternately, these challenges could be mitigated by the realization of a global network of optical clocks providing real-time geopotential measurements without the need for time-intensive point-to-point levelling. The first demonstration of a geodetic measurement with a Sr optical lattice clock by PTB [35] was compared with the terrestrial gravity database, showing consistency at the level of uncertainty of the clock measurements. It has been proposed that optical clock networks could be used to complement satellite-based gravitational measurements, offering higher temporal and spatial resolution [91] than current satellite missions such as GRACE [100] and GRACE-FO [54].

Beyond geodesy for practical purposes, optical clocks can also be used to test violations of general relativity by modifying equation (1.1) to

$$\frac{\Delta \nu}{\nu_0} = (1+\alpha)\frac{\Delta U}{c^2}, \tag{1.2}$$

where α designates a contribution from new physics beyond general relativity [98]. In a differential frequency measurement between the base and top of the 450 m Tokyo Skytree, the RIKEN transportable Sr optical lattice clock constrained the GR violation at the level of $\alpha = (1.4 \pm 9.1) \times 10^{-5}$, the best constraint to date from a terrestrial measurement.

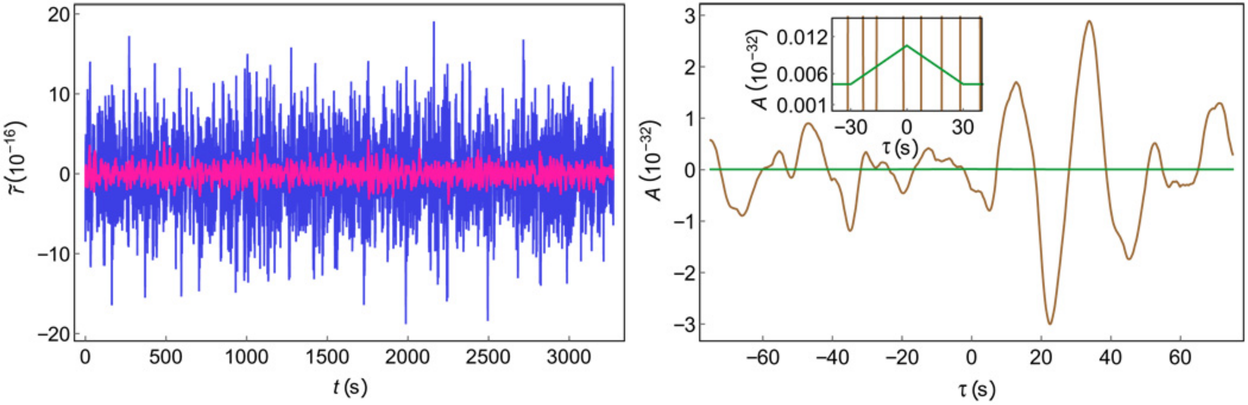

Figure 1.5: Clock comparisons for dark matter detection. Left: clock-cavity frequency differences for optical lattice clocks at NIST (pink) and LNE-SYRTE (blue). Right: cross-correlation of the NIST/SYRTE data (brown) and a fit to the expected cross-correlation function assuming a 30 s square perturbation. Reproduced from [107].

1.3.2 Dark matter detection

One candidate for non-particle dark matter is *topological defects* formed by stable field configurations of very light fields [24]. Such defects could yield transient clock frequency fluctuations while passing through the Earth, mediated by yet-unknown couplings to the Standard Model which could cause variations in the effective fine-structure constant. The defect size and energy scale can be assessed by cross-correlating frequency data from a network of clocks. In 2018, we published new constraints on dark matter coupling by forming a global network of optical atomic clocks at NIST, LNE-SYRTE in France, KL FAMO in Poland, and NICT in Japan [107]. Our efforts were sufficient to improve constraints on transient fine-structure constant variations by two orders of magnitude for topological defects comparable in size to the Earth.

However, although the detection technique did not require real-time synchronization of networked clocks, our constraints were limited by short operational overlap of the four metrological institutes due to time zone differences and the limitations of human operators. Additionally, our search was limited only to domain-wall topological defects with transverse dimensions much larger than the geographical span of our network. By realizing a denser network of clocks with higher uptime, these constraints could be significantly improved.

1.3.3 Very-long-baseline interferometry

Very-long-baseline interferometry (VLBI) is a technique in which signals from an ensemble of radio telescopes are combined, forming a synthetic aperture with a size of the entire ensemble. Similar to our dark matter detection network in Section 1.3.2, VLBI does not require physical synchronization of telescopes; rather, data are timestamped using a stable clock and combined in post-processing. The clock stability determines the maximum signal frequency which can be coherently recorded through the relationship

$$\omega\tau\sigma_y(\tau) << \delta\phi, \tag{1.3}$$

where ω is the angular frequency of the signal, τ is the integration time, $\sigma_y(\tau)$ is the Allan deviation, and $\delta\phi$ is the maximum acceptable phase excursion (typically chosen to be 1 rad). Hydrogen masers are typically used as VLBI frequency references, which are stable enough to coherently record wavelengths down to 0.87 mm [1]. The dominant decoherence mechanism is typically tropospheric turbulence, making investment in higher-performing clocks generally unwarranted. However, with recent efforts demonstrating reduced atmospheric decoherence, the stability of the frequency reference may soon limit VLBI coherence times. As proposed in [20],

> one possibility is to install at the antennas new-generation atomic clocks, e.g., the optical clocks... This could be feasible when commercial, unmanned optical clocks become available with the same reliability as the present H-masers.

Improved clocks are especially relevant for space-based VLBI systems [38], where atmospheric turbulence is not a factor. Thus, the development of transportable optical lattice clocks and the subsequent tech transfer to industry partners will enable a new generation of radio astronomy.

1.3.4 Satellite navigation

A global navigation satellite system (GNSS) consists of a constellation of satellites providing global coverage, each carrying an onboard clock which is periodically synchronized with a ground-based master clock. Precise position ranging is done by comparing transit times of signals simultaneously emitted from four or more satellites[5] . Positioning error can arise from uncertainty in the satellite ephemeris (knowledge of the position and trajectory), the instabilities of the various clocks, or effects associated with atmospheric turbulence or instrumentation error. To isolate the position and timing uncertainties, the signal-in-space ranging error (SISRE) is commonly quoted as a key performance indicator; present GNSS exhibit SISRE at the level of 1-10 m [72]. With current ephemeris uncertainty on the order of 0.2 m [73], a clock stability on the order of a few parts in $10^{-13}/\sqrt{\tau}$ with a synchronization period of 24 hours is sufficient to make the clock uncertainty negligible; thus, it has been argued that clock performance beyond this level is unnecessary [45].

However, it has also been shown that higher clock stability could enable improved orbital determination [43], reducing ephemeris uncertainty. In the planned Kepler architecture, inter-satellite synchronization and ranging is facilitated by optical links utilizing optical cavities to synchronize a constellation of 24 medium Earth orbit (MEO) satellites across three orbital planes far beyond the capabilities of modern GNSS [32]. Iodine-based optical frequency references onboard six low Earth orbit (LEO) satellites will maintain a system-wide timescale with a fractional instability of 10^{-15} between 0.1 and 10^{6} s. These LEO satellites also provide precise positioning of the MEO satellites through atmosphere-free laser ranging. The Kepler system is expected to achieve system-wide synchronization at the 1 ps level and SISRE at the few cm level. Although it is currently unclear whether these improvements will support better terrestrial GNSS in spite of limits to performance from atmospheric turbulence, optical lattice clocks could support better yet frequency instability by at least an order of magnitude at timescales exceeding 1 s.

Aside from GNSS applications, space-based optical lattice clocks could be used for gravitational wave detection [52], tests of local Lorentz invariance [37], and improved gravitational redshift measurements [89].

---

[5] The simplest ranging scheme requires only three satellites to triangulate the 3D coordinates of the receiver; a fourth satellite is added to account for synchronization error between the receiver and the constellation.

1.4 Dissertation outline

With the motivation and applications of transportable optical lattice clocks clear, Chapter 2 describes the development and implementation of such a system at NIST, up to and including narrow-line spectroscopy on the $^1S_0 \rightarrow {}^3P_0$ transition and closed-loop clock operation. Chapter 3 highlights several key systematic effects that are especially relevant or challenging for transportable clocks, as well as a broader discussion of other contributions to clock uncertainty. Chapter 4 presents a novel stretched-tetrahedral cavity design for stabilization of the clock laser, which has been optimized via finite-element modeling to allow ultra-low acceleration sensitivity in a rigidly-mounted design. Finally, Chapter 5 will offer a tantalizing glimpse into the next steps of our journey, including field deployment and overhauls of the atom trapping apparatus using a Yb dispenser and grating MOT.

Chapter 2

The NIST transportable Yb optical lattice clock

2.1 Overview

Design of the transportable Yb optical lattice clock began in 2017. The initial design steps involved selection of lasers for all wavelengths used, considerations on the overall structure of the clock, and determination of a size, weight, and power (SWaP) budget. We targeted a volume of about 1 m^3, a maximum weight of 500 kg, and a maximum power of 1 kW. To enable convenient modularity and simple integration of various electronic and optical systems, we chose to build the clock in three separate server racks: a control rack, laser rack, and physics rack (Figure 2.1).

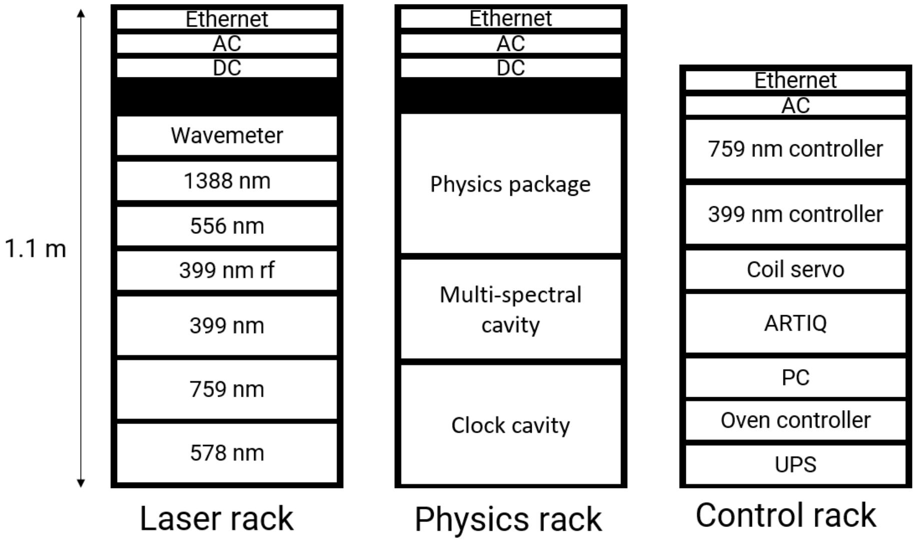

Figure 2.1: Rack layout for the transportable clock. Several modules on the back side of the control rack are not shown, including ion pump controllers, ion gauge readers, an optical shutter controller, and a direct digital synthesizer.

The control rack (21U height) includes a PC, ARTIQ real-time control system, data acquisition electronics, and laser controllers, all of which produce large heat loads that should be isolated from the optical systems. Also included are power supplies for the effusive atom source, electromagnetic coils, and ion pumps, as well as a 2U uninterruptible power supply (UPS). The UPS serves as a central hub for AC power, connected to 1U power strips in each rack, allowing critical systems like ion pumps to stay on during power outages or transportation.

The laser rack (24U height) contains the optical systems used to manipulate and deliver light at each of the five wavelengths to the atoms. Five laser modules are included: 399 nm, 556/1112 nm, 578/1156 nm, 759 nm, and 1388 nm; a total of 13 U of rack space is required. An additional 2U of space is required for electronics supporting the 399 nm laser module.

The physics rack (24U height) contains a multi-spectral cavity for stabilization of the 759 nm, 1112 nm, and 1388 nm light (5U); a dedicated, ultra-stable cavity for stabilization of the 1156 nm light (6U); and the clock physics package (7U).

Each rack also contains a rack-mounted passive Ethernet switch for networking of components like laser controllers, DAQ boards, and frequency servos. Additionally, a 1U switching power supply in each rack provides DC power to individual regulator boards in various rack drawers, each of which can cleanly supply different voltages such as ±5 V, ±12 V, and +15 V.

2.2 First-stage MOT

Atoms are sourced from an effusive oven operating at typical temperatures of 450-500 C, which sublimates solid Yb and collimates it into a thermal beam using a patterned nozzle. The atomic beam passes through a differential pumping tube, used to isolate the high-pressure region surrounding the effusive cell, and passes into a lower-pressure *science chamber* where trapping and spectroscopy takes place. In this region, a pair of electromagnetic coils in anti-Helmholtz configuration generates a quadrupole magnetic field within the chamber, well-approximated in the

MOT region by a linearly-varying field

$$\mathbf{B}(\mathbf{r}) = \frac{48a^2d}{(4a^2+d^2)^{5/2}}\mu_0 NI(\mathbf{x}+\mathbf{y}-2\mathbf{z}) \tag{2.1}$$

where a is the coil radius, d is the coil-to-coil spacing, N is the number of windings, I is the current, and μ_0 is the permeability of free space. This results in a spatially-varying Zeeman splitting of the ${}^1P_1$ state, $\Delta E(z) = g_J\mu_B m_J|\mathbf{B}(\mathbf{r})|$, where g_J is the Landé g-factor, μ_B is the Bohr magneton, and $m_J \in \{-1, 0, 1\}$ is the quantum number labeling a Zeeman substate of the $J = 1$ ${}^1P_1$ state. The atoms are trapped in a magneto-optical trap on the ${}^1S_0 \to {}^1P_1$ transition using three pairs of counter-propagating beams with wavenumbers $\mathbf{k}_i$, each with an effective detuning from the transition given by

$$\Delta = \Delta_0 - \mathbf{k}\cdot\mathbf{v} - g_J\mu_B m_J|\mathbf{B}(\mathbf{r})| \tag{2.2}$$

where Δ_0 is the detuning in the lab frame and zero-field, determined by the laser frequency. By solution of the optical Bloch equations, it can be shown that the scattering rate from the m_J transition is given by

$$R_i^{(m_J)}(\mathbf{r},\mathbf{v}) = \frac{\Gamma}{2}\frac{\beta_i}{1+\beta_T+\frac{4(\Delta-\mathbf{k}_i\cdot\mathbf{v}-g_J\mu_B m_J|\mathbf{B}(\mathbf{r})|)^2}{\Gamma^2}}, \tag{2.3}$$

where $\Gamma = 2\pi\times 29$ MHz is the linewidth of the ${}^1S_0 \to {}^1P_1$ transition, $\beta_i = I_i/I_{\text{sat}}$ is the saturation parameter of the beam relative to the saturation intensity of 60 mW/cm^2, and $\beta_T = \sum_i \beta_i$ is the total atomic saturation. The total force exerted by the radiation pressure is

$$\mathbf{F}(\mathbf{r},\mathbf{v}) = \sum_i\sum_{m_J}\hbar\mathbf{k}_i R_i^{(m_J)}\eta_i^{(m_J)} \tag{2.4}$$

The transition amplitudes into the three Zeeman substates are $\eta_i^{(0)} = (1-\xi_i^2)/2$ and $\eta_i^{(\pm1)} = (1\mp s\xi_i)^2/4$,where $\xi_i = \hat{\mathbf{k}_i}\cdot\hat{\mathbf{B}}$. The polarization handedness is $s=-1$ along the $z-$axis and $s=1$ along the x and y axes.

Some useful intuition can be gained by studying the forces along the $z-$axis, neglecting the other four beams except for their contribution to the total saturation. Along this axis, $\xi_i = 0$, and therefore the beams drive only the $\sigma_\pm$ transitions with $\Delta m_J = \pm1$. Equation (2.4) can be expanded around the phase-space origin in the small parameters

$$\epsilon_r = \frac{m_F \mu_F |\mathbf{B}(r)|}{\hbar\Delta}, \qquad \epsilon_v = \frac{\mathbf{k}\cdot\mathbf{v}}{\Delta} \tag{2.5}$$

to yield

$$F_z(z, v_z) \approx -\kappa z - \gamma v_z \tag{2.6}$$

which is equivalent to a damped harmonic oscillator with a damping force

$$\gamma = -8\hbar k^2 \frac{\beta}{(1+\beta_T+\frac{4\Delta^2}{\Gamma^2})^2}\frac{\Delta}{\Gamma} \tag{2.7}$$

and a restoring force

$$\kappa = -8\hbar k \frac{\Delta}{\Gamma}\frac{\beta}{(1+\beta_T+4\Delta^2/\Gamma^2)^2}\frac{\mu_F}{\hbar}\frac{\partial B}{\partial z} \tag{2.8}$$

The cooling provided by the two beams is opposed by a heating effect due to absorption [71]. The atom undergoes a random walk in velocity-space described by

$$\langle v(t)^2 \rangle = 2Dt = v_r^2 N \tag{2.9}$$

where $v_r = \hbar k/m$ is the recoil velocity and $N \approx R(0)t$ is the number of spontaneous emissions in a time t (limited by the number of absorptions in the same time). We find a diffusion constant

$$D = \frac{v_r^2 N}{2t} \approx \frac{v_r^2 R(0)}{2} = \frac{1}{2}\left(\frac{\hbar k}{m}\right)^2 \frac{\Gamma}{2}\frac{\beta}{1+\beta_T+\frac{4\Delta^2}{\Gamma^2}} \tag{2.10}$$

which determines the steady-state temperature,

$$k_B T = \frac{m^2 D}{\gamma} = -\frac{\hbar\Gamma^2}{32\Delta}\left(1+\beta_T+\frac{4\Delta^2}{\Gamma^2}\right) \tag{2.11}$$

The temperature is minimized for a detuning $\Delta = \sqrt{1+\beta_T}\Gamma/2$, which also minimizes the damping force in equation (2.7). At intensities far below the saturation intensity, this reduces to the well-known Doppler limit,

$$k_B T = \frac{\hbar\Gamma}{2} \tag{2.12}$$

The first-stage cooling transition at 399 nm has a Doppler-limited temperature of 700 uK. While this is too warm to load directly into the optical lattice, at these temperatures, the atoms can be efficiently transferred into a second-stage MOT for further cooling.

In order to improve the number of atoms which can be loaded into the magneto-optical trap, a far-red-detuned ($\Delta = -98$ MHz) slowing beam with 3 mm diameter and several mW of power is directed against the thermal atomic beam. Without a Zeeman slower, we obtain sufficiently high loading rate to trap over one million atoms over a typical loading time of 500 ms (Figure 2.2).

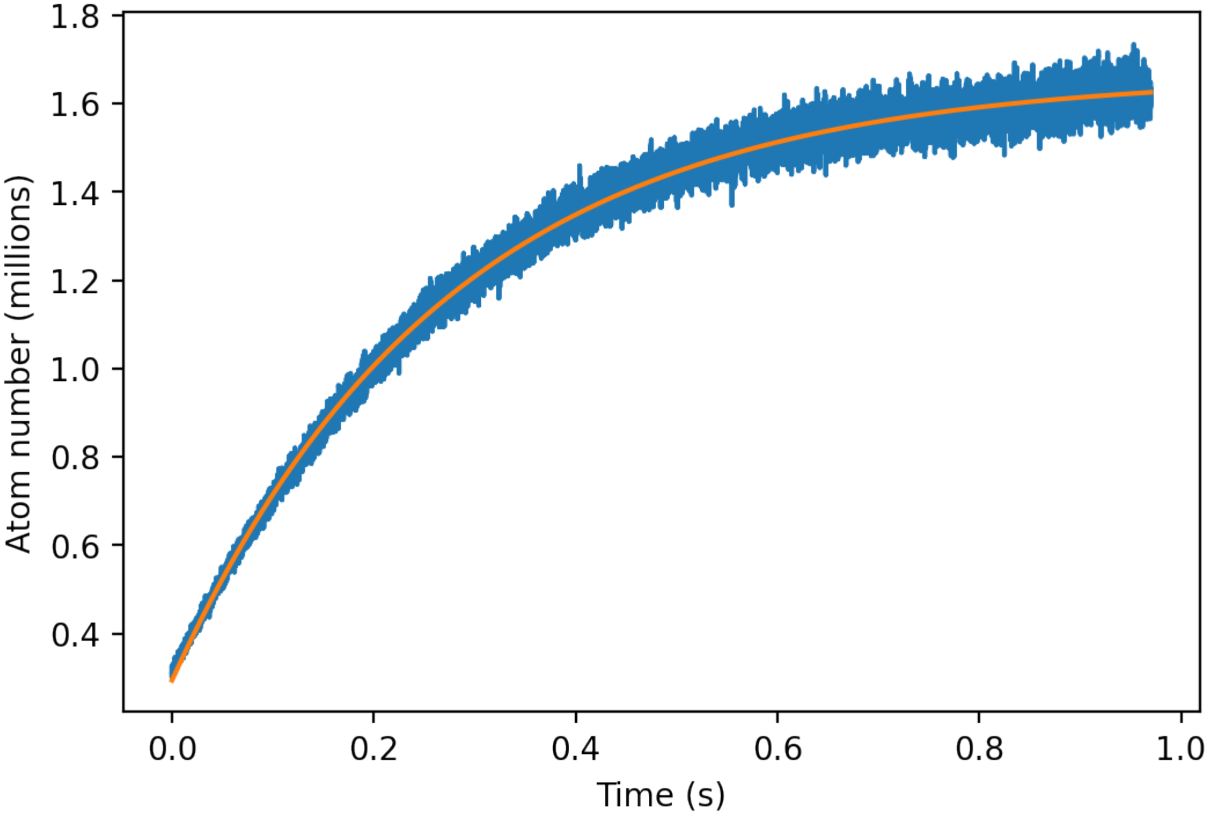

Figure 2.2: Loading curve for the first-stage MOT measured with atomic fluorescence from the cooling beams. The $y-$axis is scaled by comparison to a probe beam fluorescence with a well-known calibration of atoms per volt of PMT signal. The orange line is a fit of the form $N(t) = R\tau + (N_0 - R\tau)\exp(-t/\tau)$, yielding fit parameters $N_0 = 2.9 \times 10^5$, $R = 6.1 \times 10^6$ s^{-1}, and $\tau = 270$ ms. Typical loading time is 500 ms, over which >1 million atoms can be captured.

2.2.1 Atom number estimation

The atom number is assessed with a near-resonant probe beam which is combined with the slowing beam using a polarizing beamsplitter and introduced along the atomic beam axis. The probe light is pulsed on to heat atoms out of the trap over a typical timescale of 400 us and the fluorescence is measured using a photomultiplier tube (PMT). The total rate of fluorescence radiated from the N atoms is given by NR, where

$$R = \frac{\Gamma}{2} \frac{\beta}{1 + \beta + 4\Delta^2/\Gamma^2} \tag{2.13}$$

is the scattering rate in terms of the linewidth $\Gamma = 2\pi \times 29$ MHz, detuning $\Delta = -2\pi \times 3.25$ MHz, and saturation parameter $\beta = I/I_{\text{sat}} \approx 1$. The saturation fraction is determined by recording the fluorescent signal at varying probe intensity. Of this total fluorescence radiated into 4π steradians, our imaging setup captures only a fraction $f = T\Omega/4\pi \approx 1\%$, where Ω is the solid angle of the collection lens and T is the transmission of the vacuum viewport and collection lens. This captured photon flux can be converted to an effective optical power incident on the detector by multiplying the total scattering rate, collection efficiency, and photon energy $E = hc/\lambda$. Finally, this can be converted to a voltage output by the PMT by multiplying by the product of the radiant sensitivity $\mathcal{R}$ and the transimpedance gain G. Collectively, this yields a coefficient

$$\frac{N}{V} = \frac{1}{RE} \frac{4\pi}{\Omega T} \frac{1}{\mathcal{R}G} \tag{2.14}$$

where the first term gives the optical power, the second term gives the collection efficiency, and the third term gives the power-voltage conversion.

2.2.2 Light generation and delivery

The 399 nm light for slowing, trapping, and probing the atoms is provided by an optical layout in a 3U rack drawer, shown in Figure 2.3, and delivered to the physics package with three separate optical fibers. An additional 2U rack drawer is used for electronics for amplification, intensity stabilization, and laser stabilization. The light is generated using a external cavity diode laser (Toptica DL Pro HP) which outputs up to approximately 100 mW. A small amount of light is siphoned into a spectrometer setup which uses modulation transfer spectroscopy [109] on a Yb vapor lamp to stabilize the light to the ${}^1S_0 \rightarrow {}^1P_1$ transition (Fig. 2.3c).

The remaining light is split into three branches. First, an acousto-optic modulator (AOM) is used to send the first-order diffracted beam towards the probe fiber when an rf drive signal

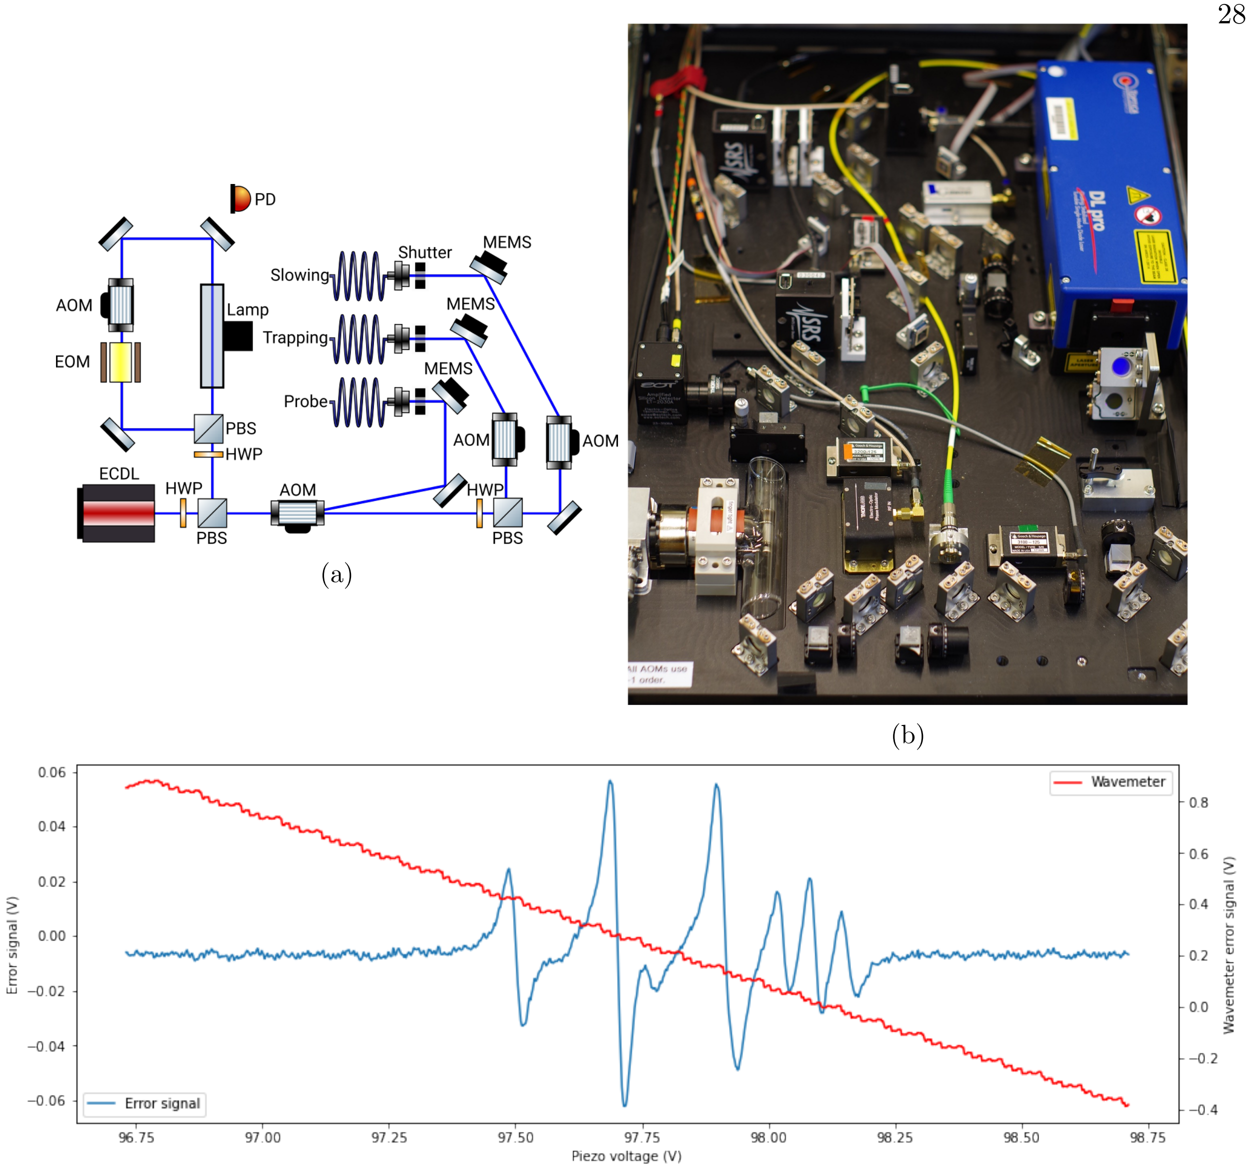

(a)

(b)

(c)

Figure 2.3: a) Schematic of 399 nm module. Beam paths are simplified for presentation and not all beam routing mirrors are shown. Abbreviations are as follows. HWP: half-wave plate, PBS: polarizing beamsplitter, AOM: acousto-optic modulator, EOM: electro-optic modulator, PD: photodiode, MEMS: fast-actuating mirror mount for automated alignment. b) Photo of 399 nm module. c) Modulation transfer spectroscopy allows the 399 nm light to be stabilized to a narrow, Doppler-free peak for various Yb isotopes. Blue line: demodulated error signal showing different isotopic peaks which can be tuned to by varying the ECDL piezo voltage. Red line: wavemeter PID error signal, proportional to the detuning relative to a target frequency. This signal can be monitored to detect and remove mode hops near the lock point by tuning the laser current.

is applied. When the rf drive is off, the zero-order beam is split into two paths for slowing and cooling light using a half-wave plate and a polarizing beamsplitter. Each of these two paths has an AOM to allow the frequency of the light to be tuned to the appropriate value. The three AOMs are also used to stabilize the intensity of the light delivered to the atoms via feedback to the rf amplitude, controlled by voltage-controlled attenuators. Although the light to each fiber path can be extinguished by switching off the rf power, we also use mechanical shutters (Stanford Research Systems SR475) to guarantee high levels of optical extinction to avoid the potential for shifts during spectroscopy. Without the shutters, we observed that just several nW of near-resonant probe power leaking through the AOM was sufficient to decrease the lifetime of the final stage of green cooling by 60%, limiting the transfer efficiency into the optical lattice.

2.3 Second-stage MOT

The atoms are transferred into a narrow-line MOT operating on the singly-forbidden $^1S_0 \rightarrow {}^3P_1$ transition at 556 nm. The green MOT beams are switched on during the blue MOT loading and for an additional 30 ms after the blue beams are extinguished to allow the atoms to reach thermal equilibrium. During this initial transfer stage, the green beams are detuned by 3-5 MHz and the magnetic field gradient is decreased to 10 G/cm; the total saturation parameter is $\beta_T \sim 140$. With these parameters, the minimum temperature is approximately 98 uK (equation 2.11). In order to load into the optical lattice, which has a typical trap depth of around 30 uK, the transfer stage is followed by a 40 ms compression stage. During the compression stage, the magnetic field gradient is linearly ramped up to 24 G/cm, the detuning to -0.5 MHz, and the intensity to $\beta_T \sim 3$. Afterwards, the atoms undergo a final stage of cooling at the end points of the ramp for an additional 5 ms; despite the short duration, this final stage reduces the atomic temperature from around 100 uK to a final temperature close to 20 uK, suitable for loading into the optical lattice.

2.3.1 Temperature measurement

After the compression and cooling stages, the MOT size is assessed by probing the atoms with the blue MOT beams and capturing the fluorescence on a CCD camera[1] . The final detuning and gradient were initially chosen by minimizing the final MOT size; later, they were slightly modified to maximize lattice loading efficiency, which depends on both the MOT size and temperature.

The atomic temperature can be assessed through time-of-flight imaging. At time $t = 0$, the atoms are released from the trap by extinguishing the trapping beams and switching off the magnetic field. The MOT is imaged at various delay times as the cloud ballistically expands and falls under gravity. A 2D Gaussian fit is used to extract the center position $(x(t), y(t))$ and rms width $\sigma(t)$ of the MOT. The camera can be spatially calibrated by fitting the vertical position of the MOT to a quadratic trajectory

$$y(t) = y_i + kv_i t + \frac{1}{2}gkt^2, \tag{2.15}$$

where $g = 9.8$ m/s^2 and k is a scale factor with units of pixels per meter[2] . The initial velocity v_i is generally nonzero and as large as cm/s due to motional relaxation of the MOT after the cooling/compression ramp; it depends on the ramp parameters such as the final intensity or detuning, as well as the ramp duration.

The scale factor coefficient k is used to convert the measured width in pixels to SI units, and the expansion is fit to the form[3]

$$\sigma^2(t) = \sigma^2(0) + \frac{k_B T}{m} t^2, \tag{2.16}$$

where σ is the rms width (standard deviation) of the Gaussian fit. As the fit is two-dimensional, we compute separate temperatures for the radial and axial directions. A typical result is shown in

---

[1] Using the retroreflected MOT beams instead of the probe beam results in less distortion of the MOT due to asymmetric radiation pressure. The MOT size was measured for various exposure times to determine the maximum pulse length which could be used without significantly heating the MOT, yielding a maximum exposure time of 200 us.

[2] This calibration was double-checked by imaging a ruler in the focal plane of the MOT; the two methods agree within the uncertainty established by the parabolic fit.

[3] The second term, which includes the temperature, should be multiplied by 1/4 or 2 if the $1/e^2$ or $1/e$ width, respectively, is used instead, as is sometimes done in the literature.

Fig. 2.4. The overall time-of-flight measurement process was performed while varying the intensity and detuning of the green MOT beams in order to minimize the temperature before attempting to load the cold atoms into the lattice.

2.3.2 Light generation and delivery

The 556 nm light for the second-stage MOT, as well as optical pumping to a single nuclear spin state, is provided by an all-fiber optical path in a 2U rack module (Figure 2.5). The light is first generated by a ytterbium-doped fiber laser at 1112 nm (NKT Adjustik) which is amplified with a separate fiber amplifier and doubled to 556 nm in a PPLN waveguide. A small amount of light at 1112 nm is picked off for frequency monitoring and stabilization; this light is split into one path to a wavemeter and another through an EOM to generate modulated sidebands for offset locking to the multi-spectral cavity. For the main optical path, after the frequency doubler, an AOM is used to tune light to the appropriate frequency for either cooling or optical pumping, as well as providing intensity stabilization via feedback to the rf amplitude. A 1x2 TTL-controlled switch routes light to either the cooling or optical pumping fiber, while a 1x1 switch directly before the 1x2 switch provides additional optical extinction during the clock spectroscopy and state readout stages. Finally, the 556 nm light is delivered to the physics package through a fiber.

2.4 Lattice loading

After second-stage cooling, the atoms are loaded into an optical lattice formed by a laser beam focused down to a typical waist diameter of 40 um and retroreflected on itself. The lattice is tuned to the magic wavelength near 759 nm, where the differential ac Stark shifts on the ground and excited clock states are nulled. The potential depth of the lattice can be given in dimensionless form, in units of the photon recoil energy $E_r = h^2\nu_{lattice}^2/2mc^2$,

$$U = \frac{4P}{\epsilon_0 c\pi w_0^2 E_r}\alpha(\nu_{magic}) = \left(\frac{h\nu_z}{2E_r}\right)^2, \tag{2.17}$$

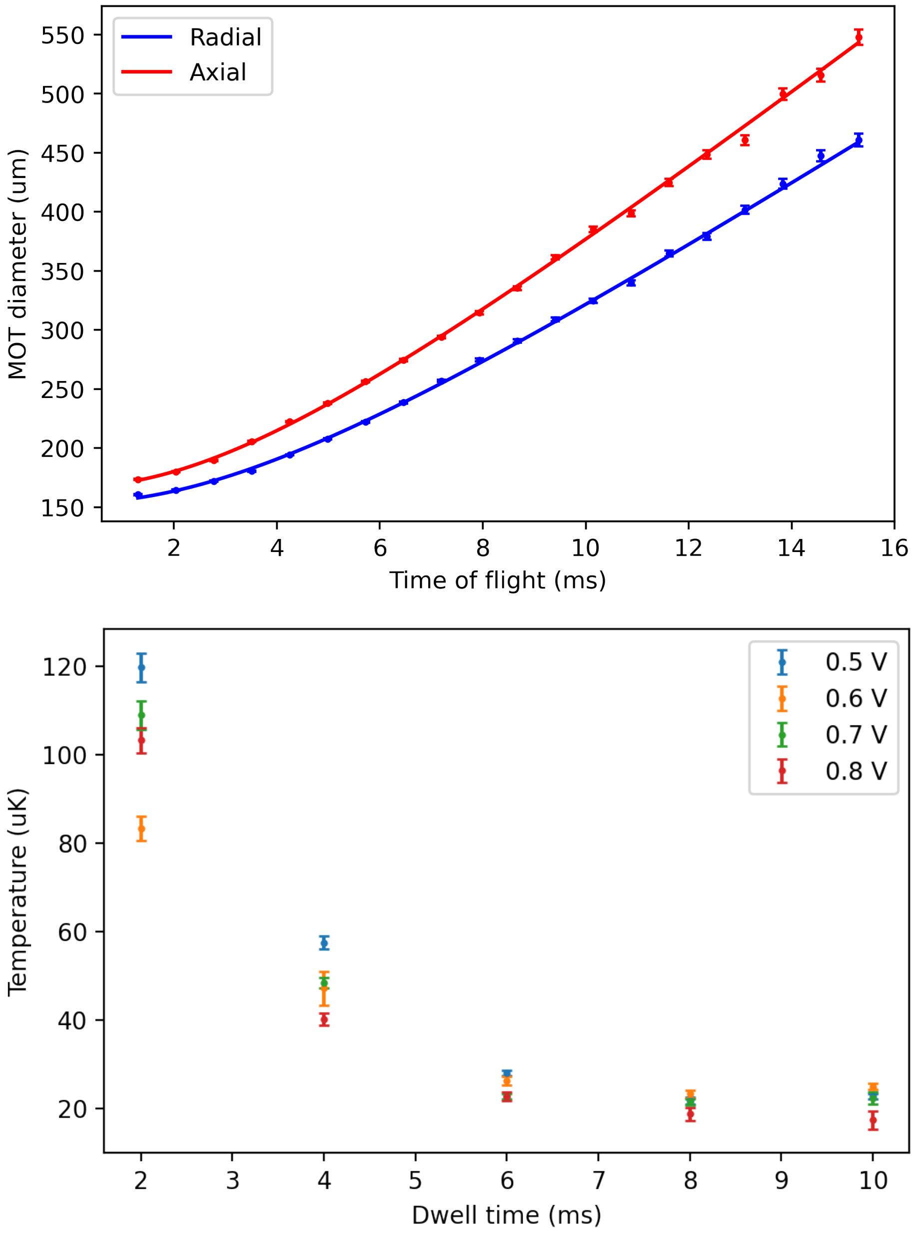

Figure 2.4: a) Ballistic expansion of the second-stage MOT as a function of the time-of-flight t. Each point represents the average of 10 experimental cycles, with error bars representing the standard error across the set of samples. The final ramp setpoints used in this cooling process were approximately -0.5 MHz detuning, 24 G/cm axial gradient, and 150 uW laser power with a beam $1/e^2$ diameter of 6 mm. The fit to (2.16) yields temperatures of 16.4(5) uK in the radial direction and 23.5(5) uK in the axial direction, with the uncertainty representing both statistical fluctuations in shot-to-shot MOT size and systematic uncertainty in camera scale factor. b) MOT temperature vs. length of the final post-ramp cooling stage. Voltage setpoints correspond to magnetic field gradients from 18-28 G/cm.

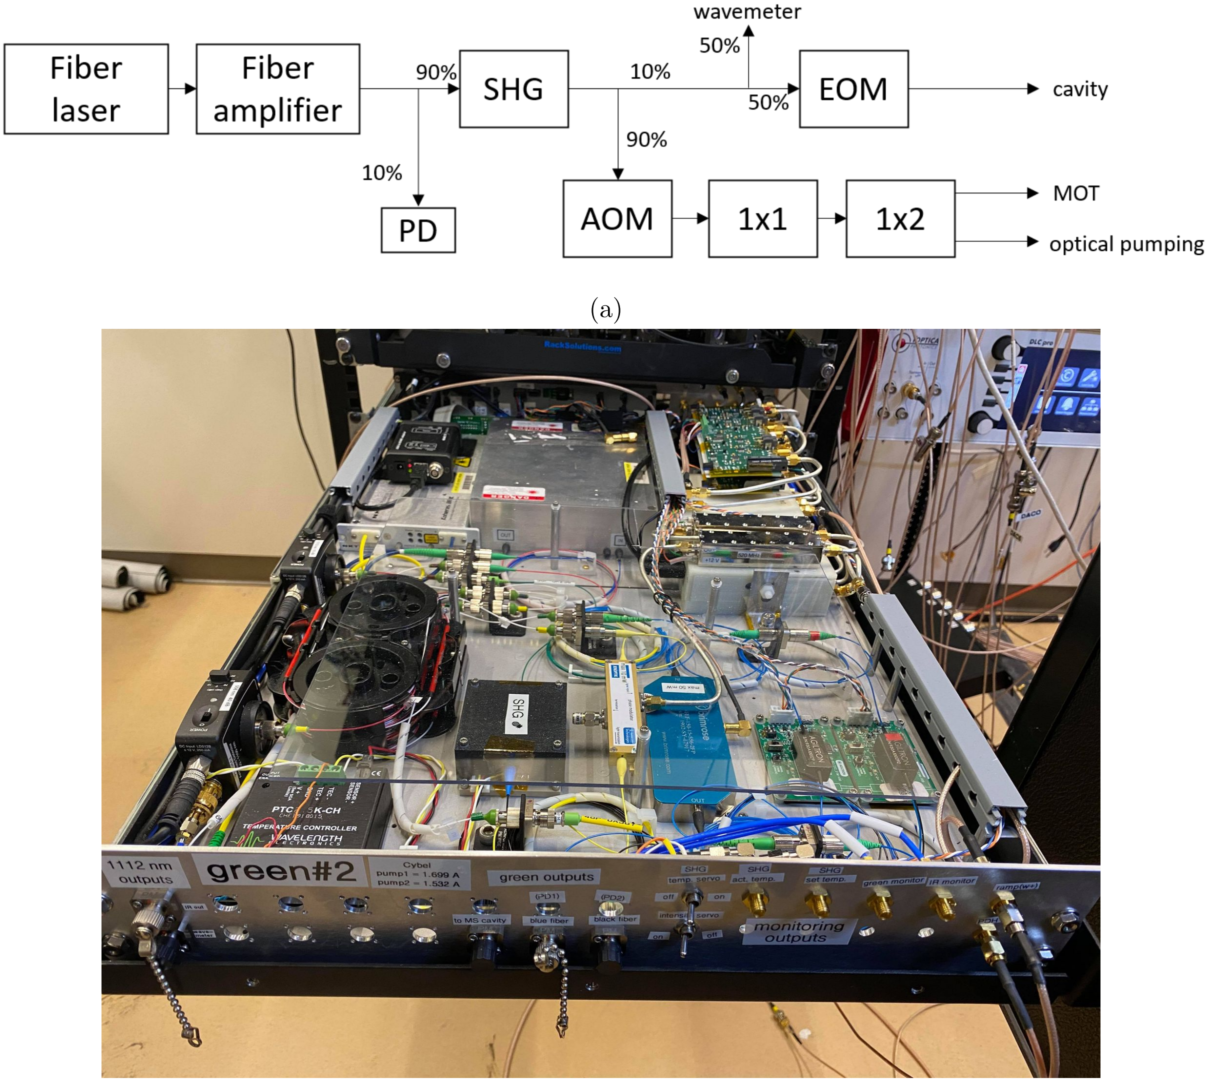

(a)

(b)

Figure 2.5: a) Schematic of 556 nm module. Abbreviations are as follows. SHG: second harmonic generation. PD: photodiode. AOM: acousto-optic modulator. EOM: electro-optic modulator. 1x1: 1x1 fiber switch. 1x2: 1x2 fiber switch. b) Photo of 556 nm module.

where P is the power of the incident lattice beam, w_0 is the beam waist radius at $1/e^2$ intensity, α is the E1 polarizability, c is the speed of light, and ϵ_0 is the permittivity of free space. The trap depth can either be computed in terms of the laser intensity and polarizability, or directly experimentally measured through determination of the trap frequency ν_z with motional sideband spectroscopy (Section 2.5.4). At maximum laser power, we can achieve trap depths of about 315

E_r, or an effective thermal depth of about 30 uK, corresponding to a longitudinal trap frequency $\nu_z = 70$ kHz.

The initial alignment of the lattice to the MOT can be time-consuming, since one must not only vary the input alignment but also keep the retroreflected beam overlapped with the incident beam. In order to expedite the alignment process, we deliver a resonant 556 nm beam through the same fiber as the lattice beam (an endlessly-single mode photonic crystal fiber, NKT LMA-PM-15). The beam is switched on during the final stage of cooling. When the input alignment of the lattice is near optimal alignment, the resonant green beam will extinguish the MOT. The input alignment is optimized by iteratively decreasing the power of the resonant beam and re-aligning to minimize the remaining atoms. Then, the retroreflector is aligned by coupling retroreflected lattice light back through the fiber and monitoring the back-coupled intensity using a pellicle beamsplitter and a photodiode in the 759 nm module.

The transfer efficiency into the lattice depends on the size and temperature of the MOT in comparison to the width and potential depth of the lattice beam. After lattice loading was achieved, the dependence of lattice-trapped atom number on the parameters of the final stage of the green MOT was studied (Figures 2.6 and 2.7). In general, we found that the simple practice of decreasing intensity by a factor of 50-100 and then choosing a gradient and detuning which minimizes MOT size is sufficient to achieve nearly optimal lattice loading. The transfer efficiency from the final stage of the green MOT into the lattice is approximately 5%, and up to 15,000 atoms can be loaded into the lattice at a typical trap depth of 370 E_r[4] .

2.4.1 Light generation and delivery

Light for the optical lattice is generated using an amplified diode laser (Toptica TA Pro) in a 3U rack drawer (Figure 2.8). To avoid large light shifts due to amplified spontaneous emission, the light is filtered with a reflective volume Bragg grating (VBG)[5] . The output mode of the

---

[4] To reduce lattice Stark shifts, we plan to ramp the trap depth down to 50-100 E_r between initial lattice loading and spectroscopy while cooling the atoms on the ${}^1S_0 \rightarrow {}^3P_1$ transition to reduce atom loss.

[5] See Chapter 3 for an in-depth discussion and uncertainty evaluation of this effect.

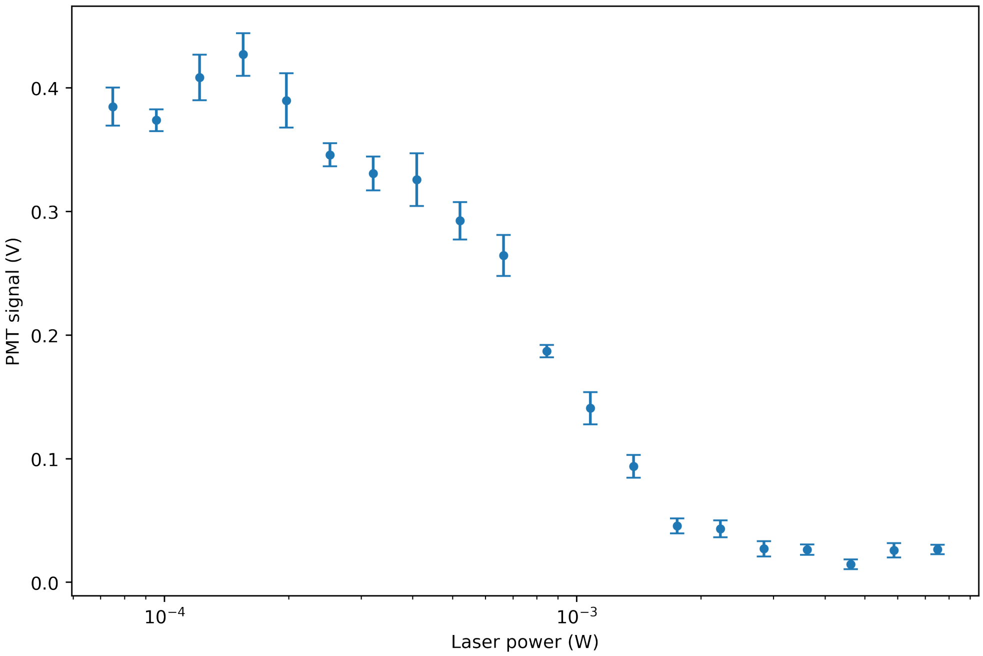

(a)

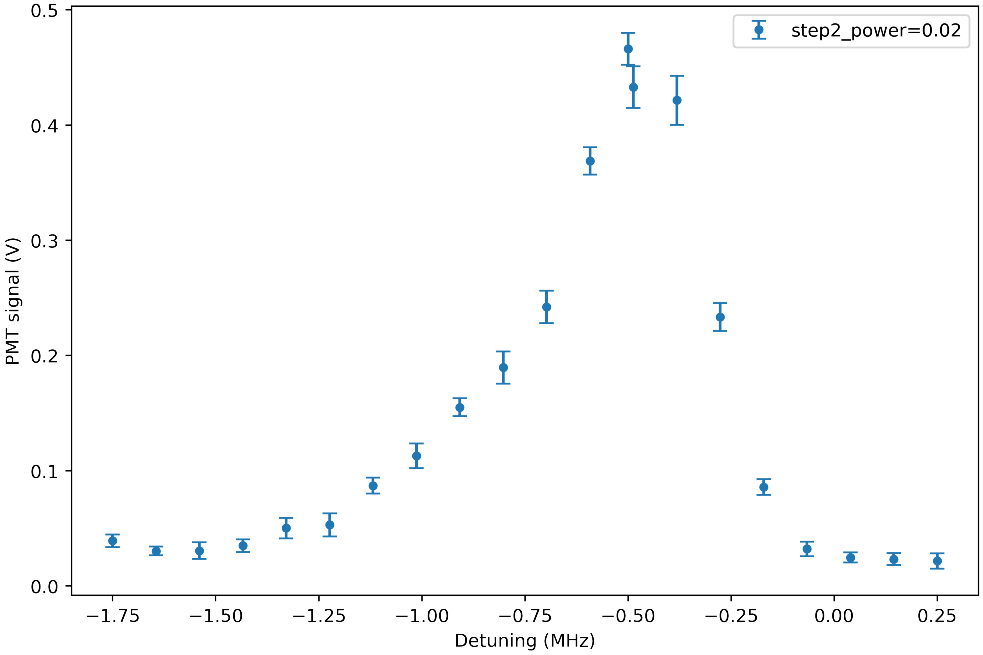

(b)

Figure 2.6: Study of lattice loading efficiency through the probe signal (which is proportional to atom number) versus a) green MOT final laser power and b) green MOT final detuning. The probe pulse lasts 3 ms and is switched on 49.1 ms after the final stage of the green MOT is extinguished.

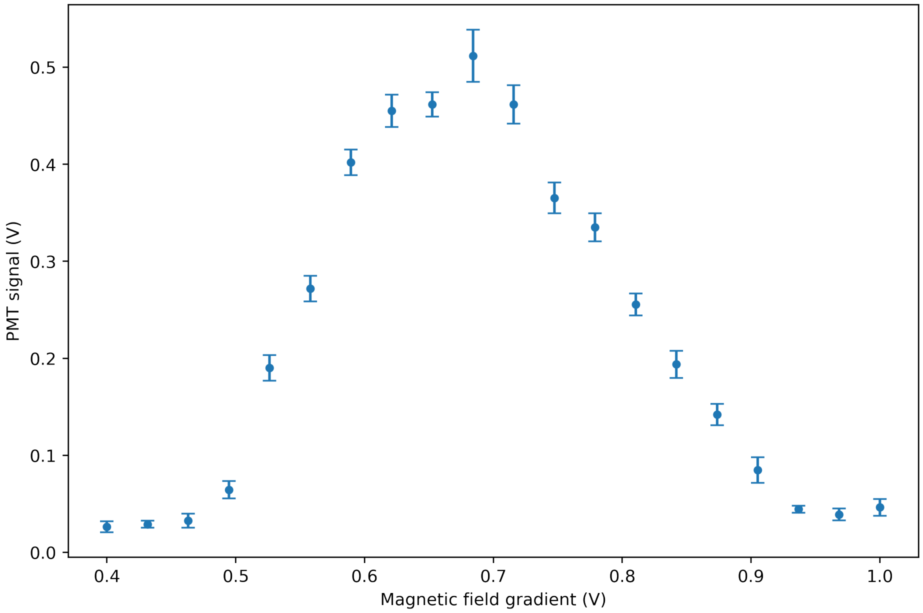

(a)

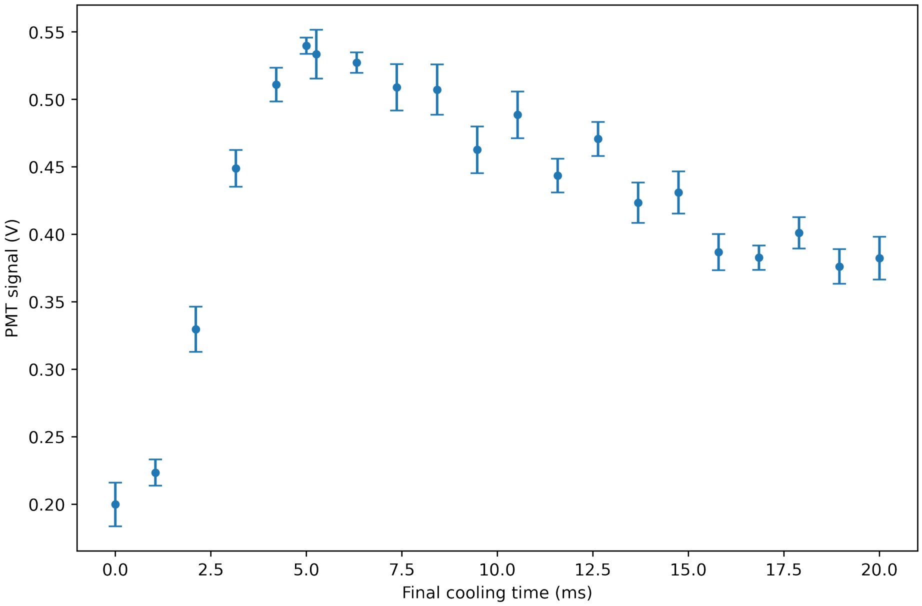

(b)

Figure 2.7: Study of lattice loading efficiency through the probe signal (which is proportional to atom number) versus a) green MOT final magnetic field gradient and b) green MOT final stage duration. The x-axis in a) is the setpoint of our current control servo, with a relation to the gradient of approximately 35 G/cm per V. The non-monotonic behavior in b) occurs due to cooling at short timescales (up to 6 ms) competing with atom loss from the trap over longer timescales.

tapered amplifier varies with propagation distance, and the VBG is placed at a distance from the amplifier where the reflected power is maximized, approximately 73 cm and 85 % diffraction efficiency. The grating is mounted in a temperature-stabilized housing to minimize thermal lensing and temperature-induced drifts of the center frequency from the 2 W of incident laser power. A MEMS mirror before the grating can be used to automatically tune the angle of the incident beam to maximize reflection at a target wavelength (typically the magic wavelength), while a second MEMS mirror after the grating can compensate for this misalignment to preserve coupling into a large-mode-area photonic crystal fiber (NKT LMA-PM-15). A pellicle beamsplitter is placed immediately before the fiber in order to monitor back-coupled light from the lattice retroreflector in the physics package with a photodiode. Additionally, there is a separate beam path to deliver 556 nm light via optical fiber and overlap it with 759 nm light to co-propagate in the photonic crystal fiber as an aid for initial lattice alignment.

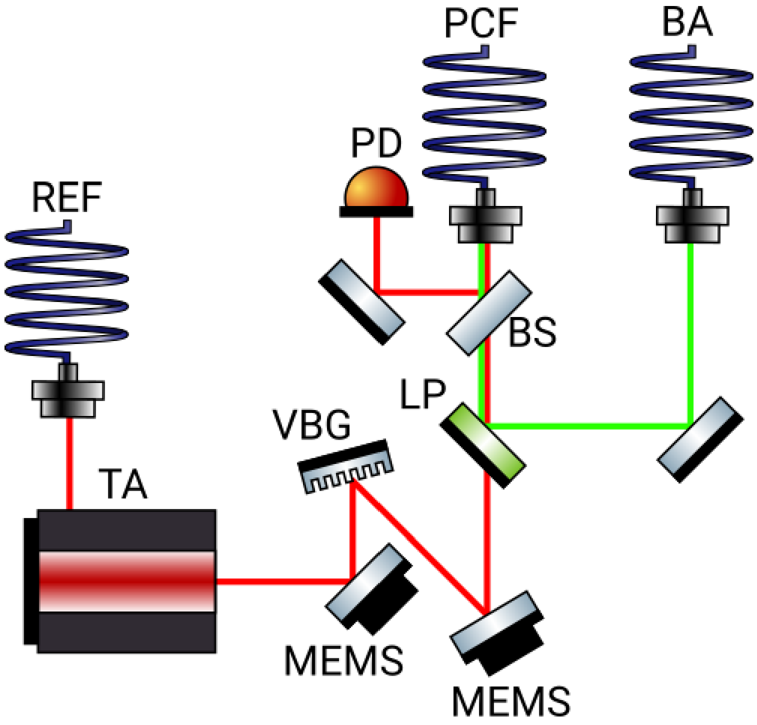

Figure 2.8: Schematic of 759 nm module. Beam paths are simplified for presentation and not all beam routing mirrors are shown. Abbreviations are as follows. LP: longpass dichroic beamsplitter, BS: pellicle beamsplitter, PD: photodiode, MEMS: fast-actuating mirror mount for automated alignment, VBG: volume Bragg grating, BA: fiber for resonant 556 nm blow-away beam for alignment aid, REF: reference fiber to wavemeter and multi-spectral cavity, PCF: photonic crystal fiber

2.5 Clock spectroscopy

2.5.1 Excitation and detection

The clock transition is driven using a constant-intensity clock pulse (Rabi spectroscopy) whose amplitude and duration are chosen to maximize the contrast of the lineshape. The lineshape is measured by determining a normalized excitation probability at varying detunings from the resonance. This probability is computed with an electron shelving technique consisting of three 399 nm probe pulses after the spectroscopy pulse. The first pulse heats the atoms in the ground state out of the lattice, resulting in a fluorescent signal which is integrated over the pulse duration to yield a signal V_g proportional to ground-state atom number. The second pulse measures the background signal V_{bg}, since there are nominally zero atoms left in the ground state. Before the third pulse, the repumping laser at 1388 nm is switched on for 5 ms to cycle the excited state atoms back to the ground state through the $^3P_0 \rightarrow {}^3D_1 \rightarrow {}^3P_1 \rightarrow {}^1S_0$ pathway. Then, the third pulse interacts with these atoms to yield a signal proportional to the excited state population, V_e. From these three signals, background-subtracted ground- and excited-state populations are computed as

$$V_g' = V_g - V_{bg} \tag{2.18}$$

$$V_e' = V_e - V_{bg} \tag{2.19}$$

Generally, this process can still result in small residual background signals which can compromise the calculation of the excitation ratio. We define population signals with a second round of background subtraction,

$$V_g'' = V_g' - V_{0g} \tag{2.20}$$

$$V_e'' = V_e' - V_{0e} \tag{2.21}$$

The two terms V_{0g} and V_{0e} are global offsets to the ground- and excited-state populations which are computed by averaging V_g' and V_e' over many experimental cycles with the lattice blocked. Finally, the normalized excitation probability is computed as

$$P = \frac{V_e''}{V_e'' + V_g''} \tag{2.22}$$

2.5.2 Transition lineshape

The normalized lineshape resulting from Rabi spectroscopy is described by

$$P(\Delta) = A\frac{\Omega^2}{\Omega^2 + \Delta^2} \sin^2\left(\frac{\sqrt{\Omega^2 + \Delta^2}T}{2}\right) \tag{2.23}$$

where Δ is the detuning from the clock transition, Ω is the Rabi frequency, T is the duration of the spectroscopy pulse, and A is the contrast. The spectroscopy time is typically chosen to drive a π

pulse for maximum contrast, $T = \pi/\Omega$. By numerically solving the transcendental equation (2.23), the Fourier-limited full-width at half-max is found to be approximately $0.87/T$.

The Rabi frequency can be determined by measuring the frequency of oscillations between the ground and excited state due to variation of the spectroscopy time while tuned on resonance with the transition (Fig. 2.10). This phenomenon, also known as Rabi flopping, is generally accompanied by a dephasing process due to incoherent interactions (e.g. glancing collisions with background gases, off-resonant photon scattering from the optical lattice, or excitation inhomogeneity due to the motional distribution of lattice-trapped atoms) which causes the ground and excited state populations to come to equilibrium with one another over long timescales.

Figure 2.9 shows several measurements of the ${}^1S_0 \rightarrow {}^3P_0$ lineshape at different Rabi frequencies. The narrowest feature achieved so far is a 5 Hz line with a spectroscopy time of 175 ms. Spectroscopy times of 1 s or longer will be possible with future vacuum improvements (Section 5.2).

2.5.3 Zeeman structure

As ^{171}Yb has a nuclear spin $F = 1/2$, the ground and excited states are each split into a manifold containing both spin states, $m_F \pm 1/2$. Thus, four transitions from the ground to excited state arise from nuclear spin effects. With zero compensation for the background magnetic field, we observe the two π transitions, which preserve nuclear spin, near ± 200 Hz; the spin-changing σ transitions are observed at $\pm$883 Hz. Based on known splitting sensitivities (200 Hz/G for π transitions and 949 Hz/G for σ transitions), we can estimate our magnetic field magnitude to be approximately 1 G. This is twice as large than the Earth's magnetic field at our location (0.52 G) and most likely originates from the ion pump magnets.

2.5.4 Sideband spectroscopy

Aside from the Zeeman structure, another class of transitions is resolvable: the lattice sideband transitions on which the longitudinal motional quantum number is raised or lowered,

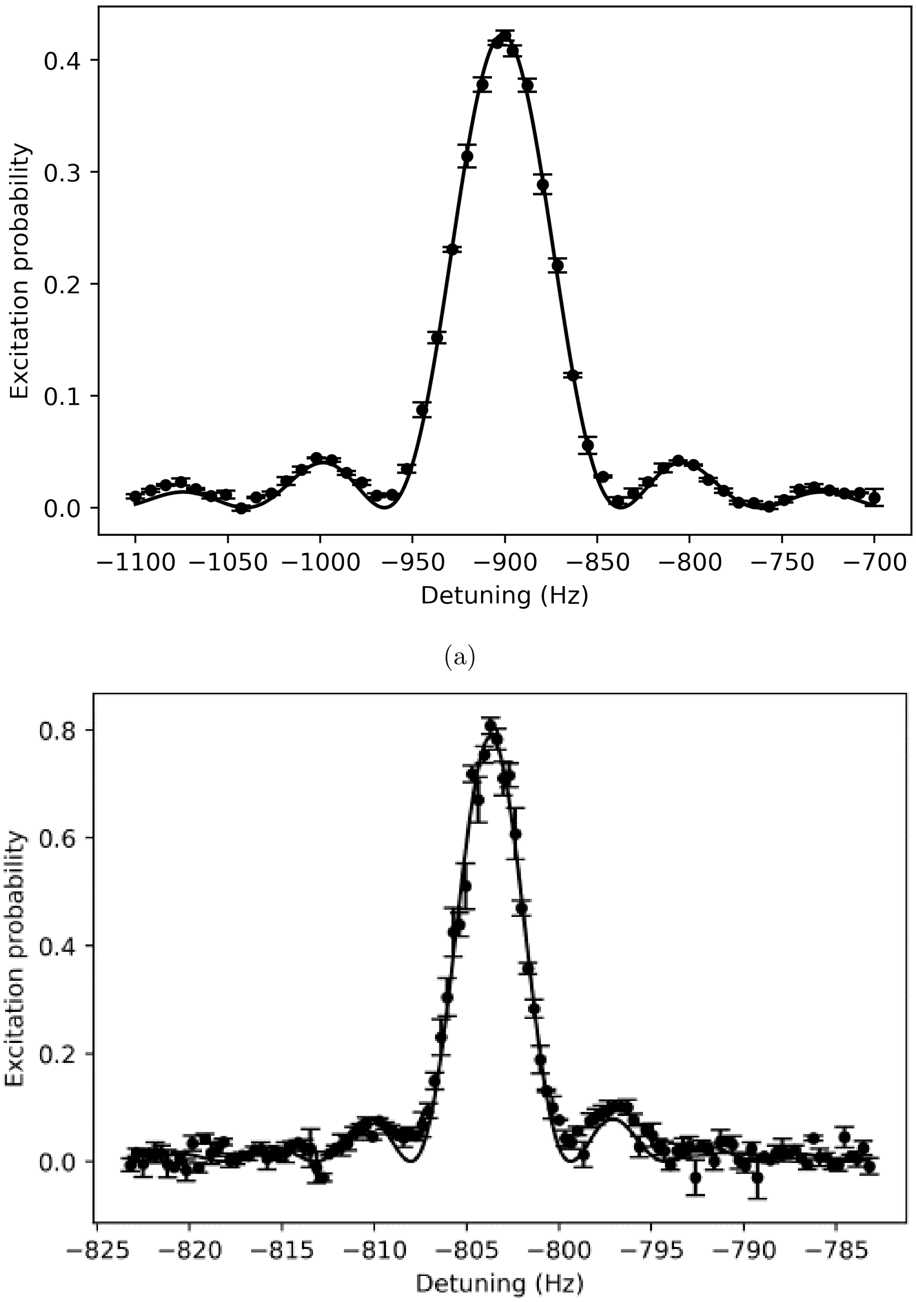

(a)

(b)

Figure 2.9: Clock transition spectroscopy. Each point (errorbar) represents the average (standard error) of three experimental cycles. a) 62 Hz line with 14 ms spectroscopy time fit to equation (2.23). b) 4 Hz line with 200 ms spectroscopy time. Contrast is improved due to optical pumping into a single nuclear spin state.

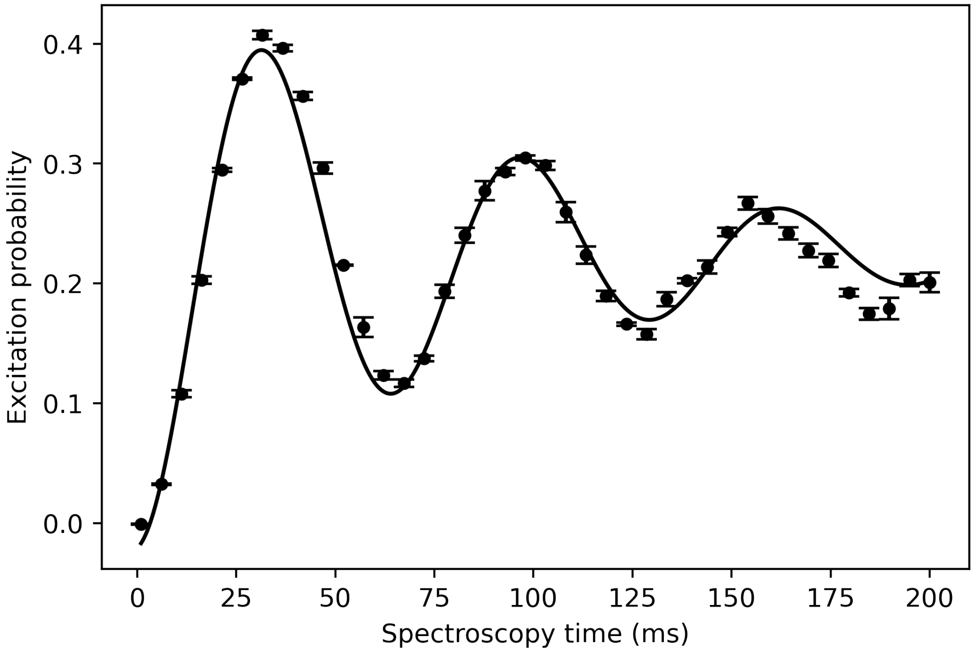

(a)

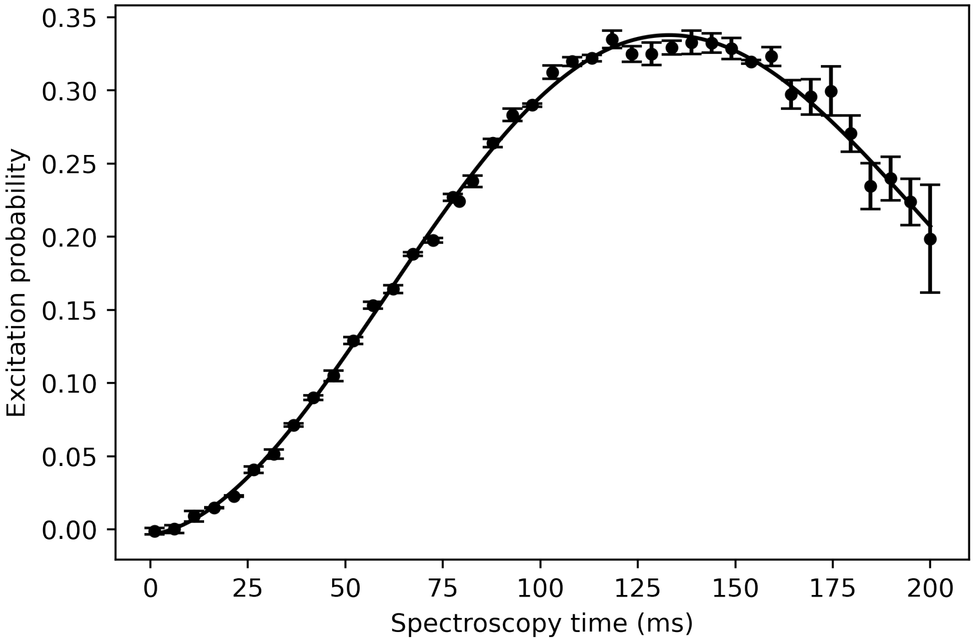

(b)

Figure 2.10: Rabi oscillations measured with a) 18.5 uW optical power and b) 1 uW optical power. The fits, shown by solid lines, yield Rabi frequencies of 96.3 Hz for a) and 22.9 Hz for b). Each point (errorbar) represents the average (standard error) of three experimental cycles.

$|^1S_0\rangle|n_z\rangle \rightarrow |^3P_0\rangle|n_z \pm 1\rangle$. The thermal distribution of atoms across various lattice motional states, as well as coupling between the longitudinal and transverse states, gives rise to broad sideband features, shown in Figure 2.11 and described mathematically by [9]. Sideband fits yield longitudinal and transverse temperatures of 9(2) and 21(2) uK respectively. The longitudinal temperature T_z can also be determined from the ratio of the areas of the red and blue sidebands, which satisfy

$$\frac{\int_0^{\nu_z} P_{\text{red}}(\Delta)d\Delta}{\int_{-\nu_z}^0 P_{\text{blue}}(\Delta)d\Delta} = 1 - \frac{\exp\left(-\frac{E_0}{k_BT_z}\right)}{\sum_{n_z=0}^{N_z}\exp\left(-\frac{E_{n_z}}{k_BT_z}\right)}. \tag{2.24}$$

Here, ν_z is the frequency of the sharp cutoff at the far-detuned side of either sideband. The energy associated with the longitudinal motional number n_z, as well as the transverse numbers n_x, n_y, is

$$E_{n_z}/h = \nu_z\left(n_z+\frac{1}{2}\right)+\nu_r(n_x+n_y+1)-\frac{1}{2}\nu_{\text{rec}}\left(n_z^2+n_z+\frac{1}{2}\right)-\nu_{\text{rec}}\frac{\nu_r}{\nu_z}(n_x+n_y+1)\left(n_z+\frac{1}{2}\right) \tag{2.25}$$

The first two terms are the eigenenergies of the familiar harmonic oscillator potential, while the last two include quartic corrections to the potential and couplings between the longitudinal and transverse states. Numerically solving (2.24) for the measured sidebands in Figure 2.11, we find a longitudinal temperature of 8.4 uK, in good agreement with the fit results.

2.5.5 Clock laser delivery

The 578 nm clock beam is introduced through the back surface of the lattice retroreflector mirror, which is coated for 90% reflectivity at 578 nm. Up to approximately 300 uW of clock light can be delivered to the atoms; for sideband spectroscopy, the lattice retroreflector was swapped to a mirror with nearly 100% transmission at 578 nm, allowing powers up to several mW to be used. The clock beam is overlapped with the lattice and focused down to a diameter of 100 um at the position of the lattice-trapped atoms. For these measurements, we borrowed laser light from a clock laser referenced to a high-finesse cavity used for our laboratory clocks Yb1 and Yb2. This was due to manufacturing delays in the laser module and reference cavity for our transportable clock; however, it was especially convenient to implement clock spectroscopy with a laser of demonstrably high performance.

Figure 2.11: Motional sideband spectroscopy. Taken with 150 ms spectroscopy time and averaging 3 experimental cycles per detuning. The red and blue curves are fits to the model [9], yielding longitudinal and transverse temperatures of 8(2) and 21(2) uK respectively.

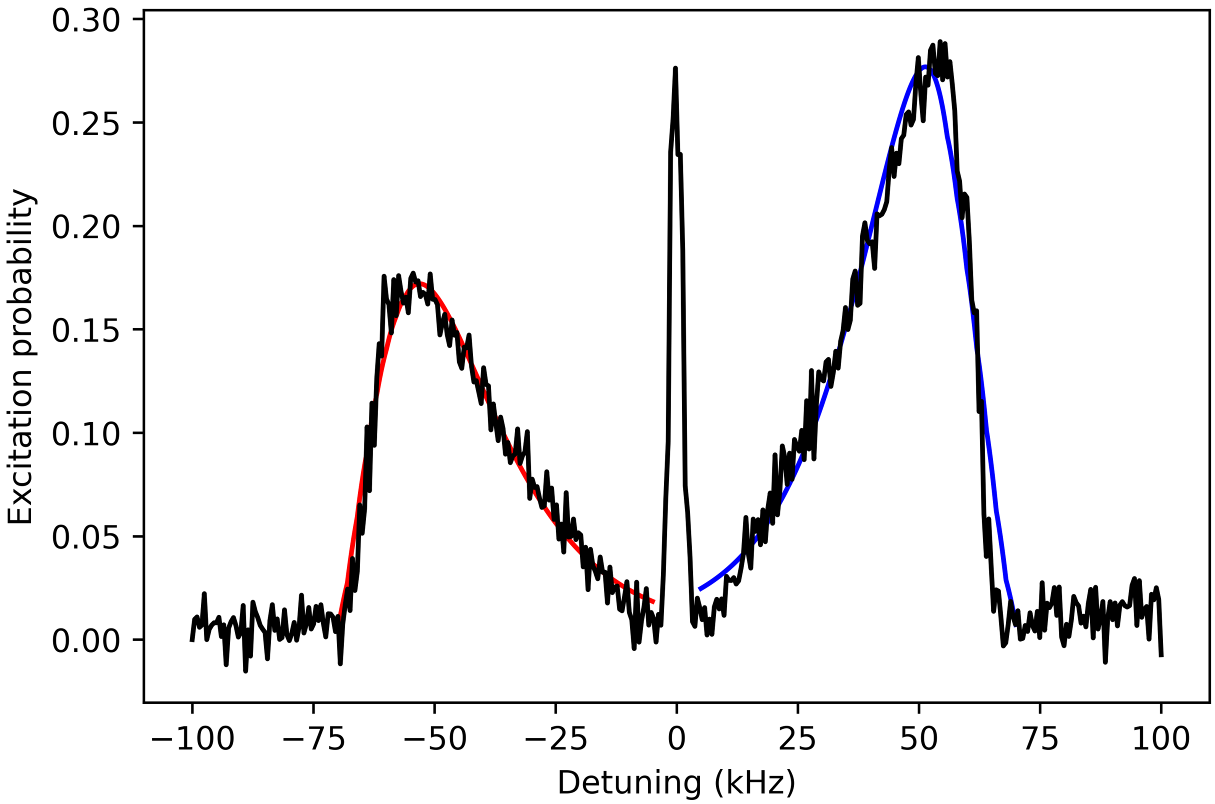

The light was delivered from a separate laboratory over a 20 m fiber. To compensate for phase noise introduced along this distance, we implemented phase-noise cancellation using a Michelson interferometer to compare the phase of outgoing light with light returned from a reflection from the flat fiber tip [66]. The frequency of the clock laser was controlled with a fiber AOM with a center frequency of 99 MHz, driven by a DDS (MOGLabs Agile RF Synthesizer) which is locked to a reference frequency near 198 MHz provided by our control hardware. The phase lock can track the reference for detunings up to around 10 kHz; to achieve larger detunings for sideband spectroscopy, we bypassed the phase noise cancellation and controlled the AOM directly with the reference DDS.

2.5.6 Clock servo operation

Once the clock transition can be resolved with high signal-to-noise, it is straightforward to operate in a servo mode tracking the atomic transition. Since the transition lineshape is symmetric, locking requires successive interrogations on either side of the line, resulting in an asymmetric error function. In our implementation, we actually simultaneously lock to two separate peaks, the $\sigma_\pm$ transitions in order to obtain first-order insensitivity to Zeeman effects. Before locking, we sweep across both peaks and fit Lorentzian lineshapes to determine three feedforward parameters: the *hop* Δ, which is equal to half the peak splitting; the *dev* δ, which is equal to half the linewidth; and the *offset* f_0, which is the center of the two peaks relative to the cavity frequency. Combining these three parameters with a given choice of signs for the hop and dev allows the laser to be tuned to the half-max position on the red or blue side of either peak. For example, tuning to $f = f_0 - \Delta + \delta$ selects the blue side of the σ_- peak, while $f = f_0 + \Delta - \delta$ selects the red side of the σ_+ peak.

During clock operation, a series of measurements are performed in order to update the offset f_0 and the hop Δ. Across four experimental cycles, the red and blue sides of each of the four peaks are probed and a frequency offset correction is computed, given by

$$\delta f_0^{(1)} = g_1[(P_1 - P_0) + (P_3 - P_2)], \tag{2.26}$$

where P_0 and P_1 are the excitation probabilities on the red and blue sides respectively of the σ_- peak and P_2 and P_3 are excitation probabilities on the red and blue sides of the σ_+ peak. The overall magnitude of the frequency correction is controlled by the gain g, which is typically on the order of unity. A second integrator is employed to provide slower, long-term corrections, computed by

$$\delta f_0^{(2)} \mapsto \delta f_0^{(2)} + g_2 \delta f_0^{(1)}, \tag{2.27}$$

The overall frequency update is

$$f_0 \mapsto f_0 + \delta f_0^{(1)} + \delta f_0^{(2)} \tag{2.28}$$

An independent lock is used to track the splitting of the $\sigma_\pm$ transitions, giving a real-time measurement of the local magnetic field strength at the atoms, by updating the hop according to

$$\Delta \mapsto \Delta + g_3[(P_1 - P_0) - (P_3 - P_2)], \tag{2.29}$$

After probing the red and blue sides of both peaks, we probe them in reverse order over the next four cycles in order to cancel out linear drift due to ULE creep of the reference cavity, which is typically on the order of 22 mHz/s uncompensated. We also employ feedforward drift compensation to null the cavity drift at the level of 1 mHz/s.

2.5.7 Clock comparison

After implementing the full clock servo system, we made a preliminary comparison against our laboratory clock Yb1. As a full uncertainty budget for the transportable system has not yet been completed, this measurement is primarily a proof-of-concept for testing clock operation. Additionally, several key systematic effects were not yet controlled over long timescales, such as lattice intensity, leading to a drift for time differences exceeding 10^3 s. The observed instability, approximately $10^{15}/\sqrt{\tau}$ (Figure 2.12), is worse than expected for the clock laser, which has supported clock instabilities at the $10^{-16}/\sqrt{\tau}$ level. This is primarily due to a short trap lifetime and a relatively low signal-to-noise ratio (≈ 50) in the normalized detection scheme, and we anticipate that the instability will be reduced by a factor of 7 or more with the implementation of optical pumping to a single nuclear spin state and future vacuum chamber upgrades (Section 5.2).

Figure 2.12: Clock comparison with Yb1. The dashed line is a fit to a white-noise asymptote with $\sigma_y(\tau) \approx 10^{-15}/\sqrt{\tau}$.

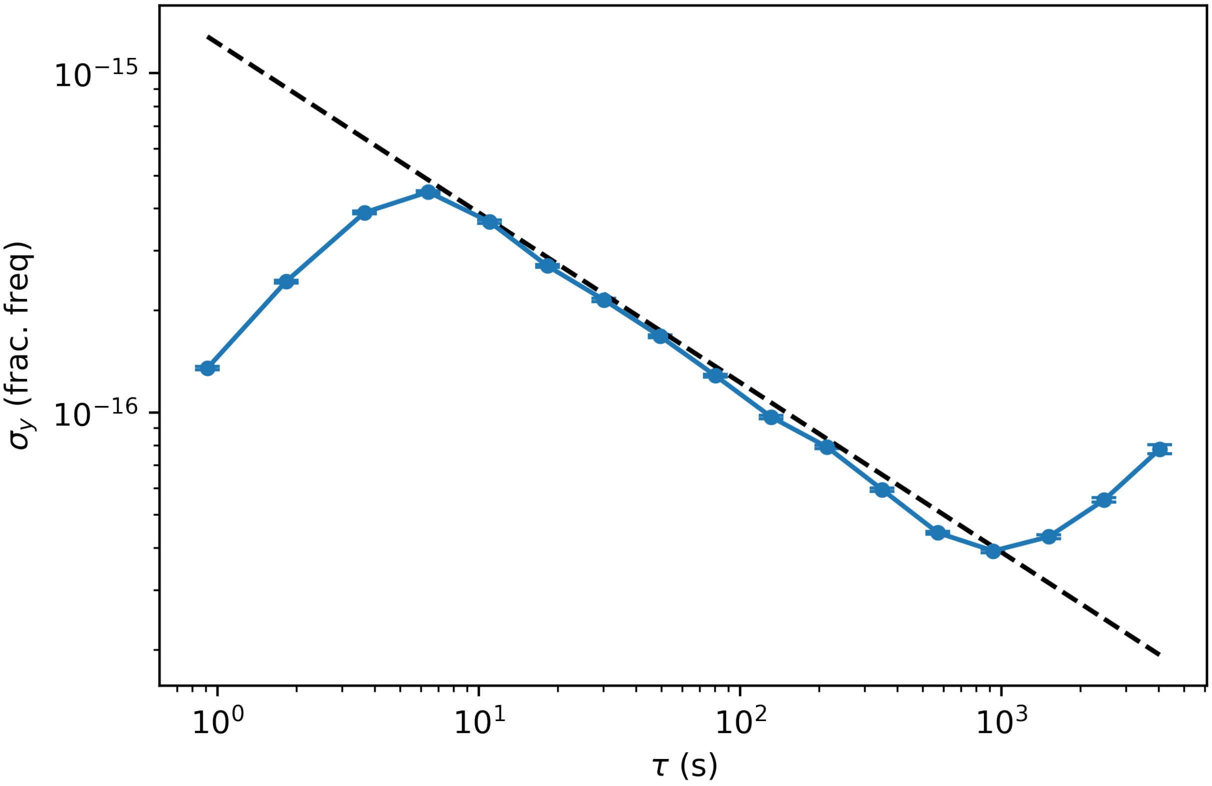

2.6 1388 nm

The 1388 nm module for repumping on the ${}^3P_0 \to {}^3D_1$ transition is shown in Figure 2.13. Light is generated with a diode laser and split into two branches. The first branch is sent to the multi-spectral cavity, passing through an EOM for generation and modulation of sidebands for locking. The second branch passes through an AOM for frequency tuning, then is delivered to the physics package through a 1x2 switch. The switch can be toggled via TTL to instead send light to a wavemeter, passing through a frequency-doubling module first to convert to a visible frequency. Like the 556 nm module, the 1388 nm module is entirely in-fiber all the way to the physics package.

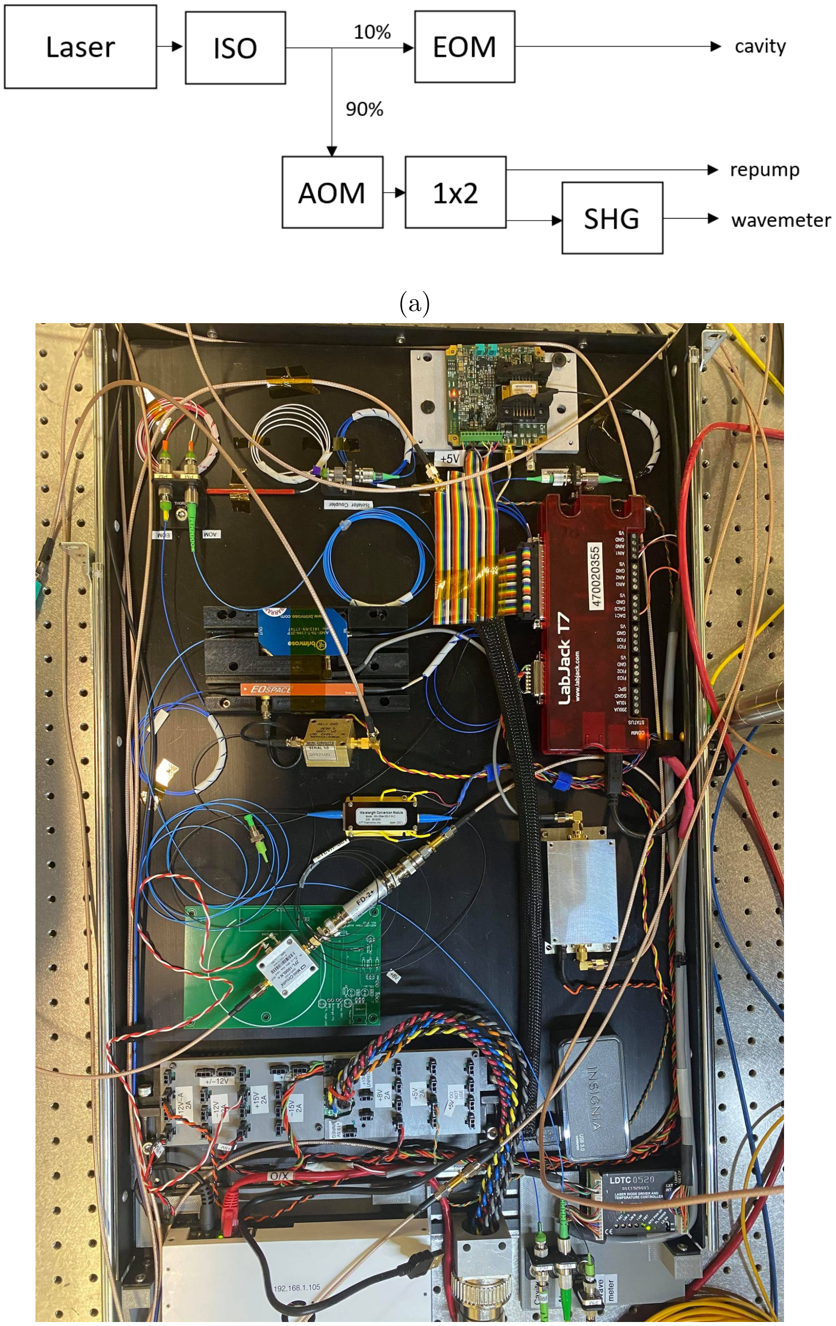

(a)

(b)

Figure 2.13: a) Schematic of 1388 nm module. Abbreviations are as follows. ISO: isolator. SHG: second harmonic generation. AOM: acousto-optic modulator. EOM: electro-optic modulator. 1x2: 1x2 fiber switch. b) Photo of 1388 nm module.

2.7 Vacuum chamber

The clock vacuum chamber is shown in Figure 2.14. The central chamber (called the *science chamber*), which serves as a Faraday cage and blackbody enclosure to constrain the dc Stark and blackbody radiation shifts (Chapter 3), is made of aluminum, which has several favorable properties compared to stainless steel. The density is approximately 64% lower, allowing the science chamber to weigh under 300g before bonding of the windows. The thermal conductivity is nearly a factor of 10 higher, reducing the magnitude of thermal gradients that could complicate blackbody radiation shift assessment. Outgassing rates are typically lower for aluminum, since the amount of hydrogen permeation is seven orders of magnitude lower than stainless steel. Lastly, aluminum is much cheaper and faster to manufacture, allowing the design to be rapidly iterated on during the early stages of vacuum testing.

Figure 2.14: Rendering of the clock vacuum chamber. The atom source and science chamber are differentially pumped by two ion pumps with integrated non-evaporable getters. After bakeout, the chamber is sealed with copper pinchoff tubes (not shown). The magnetic shield is hidden to expose the quadrupole coil heat sinks.

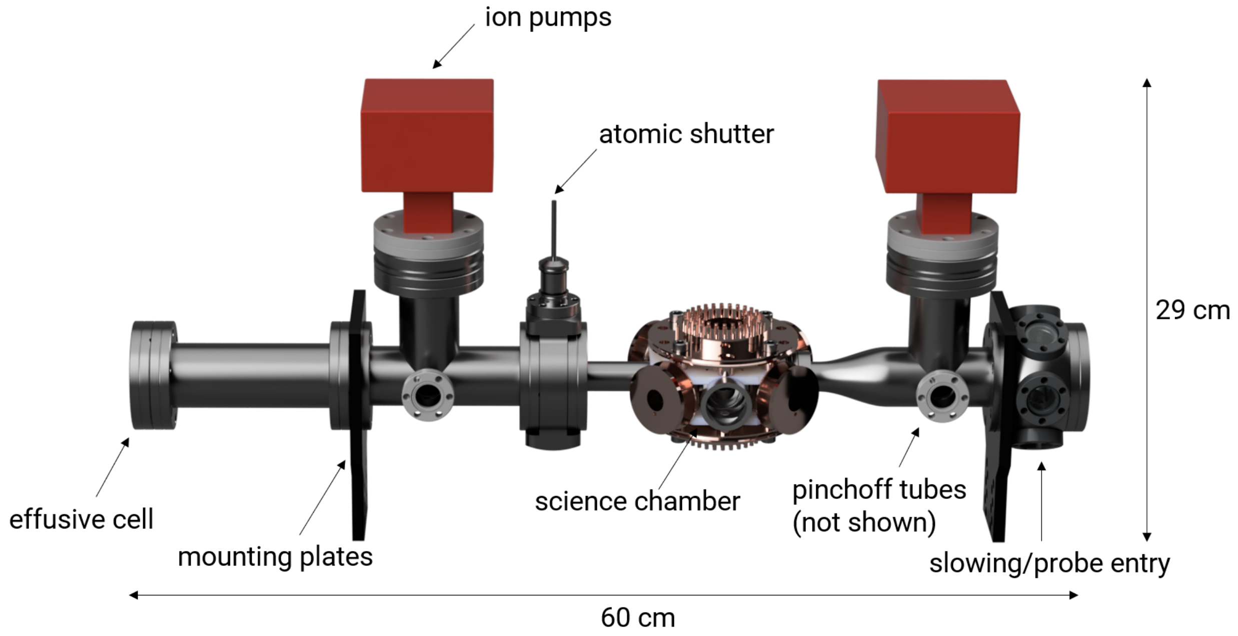

Rather than sealing the chamber viewports with conventional ConFlat vacuum hardware, we developed an epoxy-based sealing technique using a thermally conductive, UHV compatible epoxy (Epo-Tek H77). This allows windows to be mounted directly onto the chamber without additional space for bolt holes, which is vital for limiting window solid angle for BBR considerations while maintaining low overall size (Chapter 3). Additionally, this compact chamber size dramatically reduces the required current in the MOT coils, since the axial gradient scales strongly with the coil separation $2z_0$,

$$\frac{\partial B_z(z)}{\partial z} \propto \frac{R^2 z_0}{(R^2 + z_0^2)^{5/2}}, \tag{2.30}$$

where R is the coil radius. Whereas the quadrupole coils on our laboratory clocks require water cooling, the reduced current here is compatible with passive air cooling.

The atom source, a dual-filament effusive cell, occupies about 7 cm of in-vacuum length on the far left side of Figure 2.14. The vacuum tee adjacent to the effusive cell mounts an ion pump with an integrated non-evaporable getter (SAES NexTorr Z100), delivering a hydrogen pumping speed of 150 L/s (6 L/s from the ion pump). A compact ion gauge is also mounted on the tee. A narrow tube with 86 mm length and 4 mm diameter separates the low-pressure science chamber from the high-pressure oven region, limiting the conductance to under 1 L/s.

An atomic beam shutter based on a wobblestick and soft-shift solenoid is placed between the oven and the differential pumping tube to block the atomic beam for all stages of the experimental sequence other than MOT loading, preventing trap loss or clock frequency shifts from collisions with the hot atoms and blackbody shifts from the hot effusive oven. The shutter is controlled with an Arduino, which toggles between “open” and “closed” positions based on a TTL logic level. The physical position of the wobblestick in either position can drift over weeks of operation. To keep the shutter in the correct range, we use a magnetic sensor to read the positions in the two states and update position setpoints as needed.

To the right of the science chamber, a second vacuum tee mounts an identical ion pump and ion gauge. A hexagonal chamber at the end of the physics package holds an in-vacuum mirror which allows the slowing and probe beams to be introduced to the chamber and turned 90^o. The mirror mount is thermally contacted to a copper feedthrough allowing the mirror to be heated up to bake off any deposited ytterbium from the atomic beam; however, this has not yet been necessary over the first five months of operation.

2.7.1 Optics package

The vacuum chamber in the previous section is mounted rigidly in a 7U rack drawer which includes all of the optics for delivery of light to the atoms. A 1/2" thick, custom-patterned aluminum baseplate is used to mount the chamber and various modular optics breadboards. Each of the eight beams (399 nm MOT, 556 nm MOT, slowing, probe, optical pumping, lattice, clock, and repump) is introduced with an essentially identical fiber delivery module (aside from specific choices of optics coatings and fiber collimators). Each fiber delivery module contains a lens to collimate the light from the fiber to a given diameter, a half-wave plate and PBS cube to define the input polarization into the physics package, and a pellicle beamsplitter with a photodiode for intensity monitoring and stabilization. From the lower front fiber delivery module, the 399 nm and 556 nm MOT beams are overlapped with a shortpass dichroic mirror and split into three paths. The two horizontal paths are raised to the midplane of the chamber by a twisted periscope which contains a half-wave plate to preserve the linear polarization state. At the output of the periscope, the light passes through a quarter-wave plate to circularize the polarization, then enters the chamber. At the far side of the chamber, the light is retro-reflected by a mirror in a kinematic mount, passing through a quarter-wave plate twice to flip the handedness of the circular polarization. The vertical beam path is essentially the same, except no twisted periscope configuration is needed. The 399 nm probe and slowing beams are combined with a plate polarizing beamsplitter and routed to the far side of the chamber, where they are delivered to the atoms using a periscope and the in-vacuum mirror.

Above the MOT delivery module, the lattice and repump beams are overlapped using a longpass dichroic mirror, then focused to a waist at the atoms using an achromatic lens with $f =$200 mm focal length.

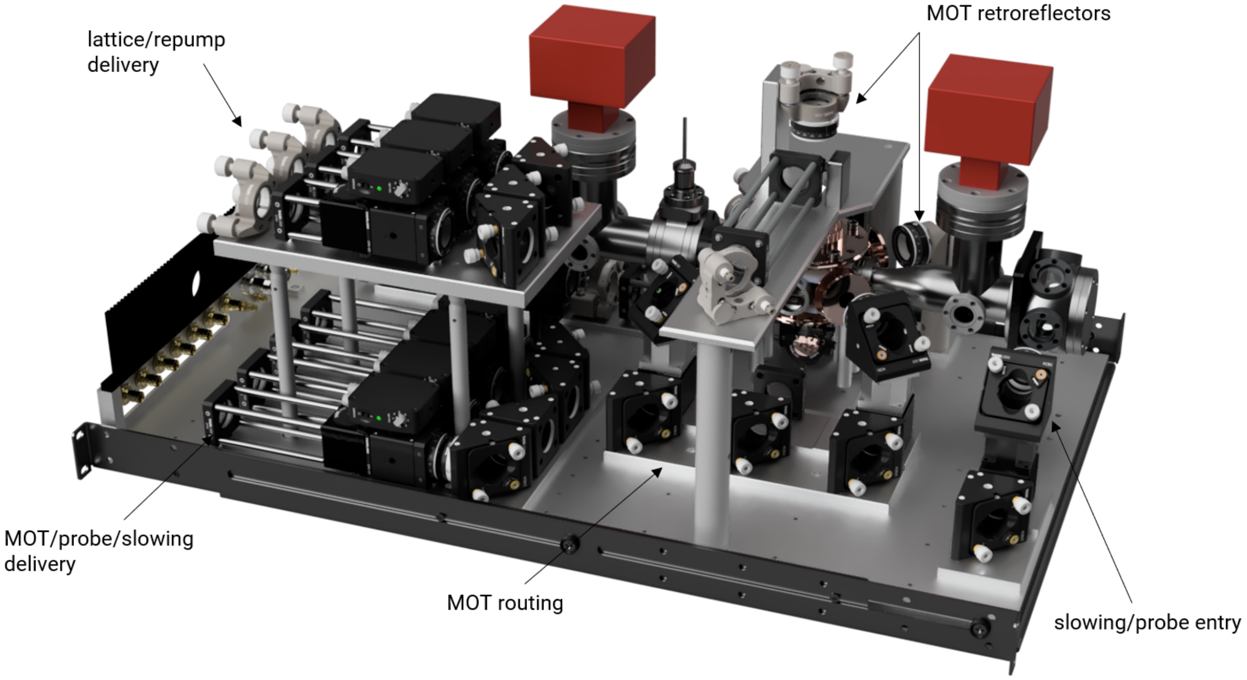

Figure 2.15: Rendering of the optics package. Presented for illustration, with some optics and other pieces hidden. The entire apparatus fits in 7U (31 cm) of rack space with a depth of 74 cm.

2.8 Magnetic coils

Four pairs of coils are mounted on a copper heat sink assembly surrounding the vacuum chamber. The largest coil pair, formed from 15 windings of 1x7 mm rectangular cross-section wire, produces the quadrupole field for magneto-optical trapping, typically with a maximum axial gradient of 35 G/cm and a maximum current of 24 A. Two identical pairs in the horizontal plane are used to null the ambient magnetic field, as well as providing a bias field to define a quantization axis for optical pumping and spectroscopy. A fourth pair in the vertical direction is used to null ambient vertical field components from the Earth and the ion pumps.

All four pairs of coils are driven with a custom four-channel current servo, occupying 2U of rack space, which allows fast switching of the coils using setpoints provided by DAC lines from the ARTIQ system (Section 2.11.1).

2.9 Magnetic shielding

As we plan to deploy the transportable clock to “metrologically-interesting” locations which could potentially possess local magnetic fields much stronger than the ambient fields in the lab, we designed a magnetic shield which fits around the chamber to screen the atoms from potentially large Zeeman shifts (Figure 2.16a). The shield is constructed from mu-metal, a high-permeability nickel-iron alloy. Finite-element modeling with COMSOL Multiphysics software [23] was used to compute the effectiveness of the screening along the three Cartesian axes. Before studying the actual shield geometry, the finite-element model was verified by simulating a spherical shell of variable permeability, whose screening effect can be computed analytically:

$$\frac{H_{in}}{H_0} = 9\left[5 + \frac{4a^3}{b^3} + 2(\mu + \frac{1}{\mu})\frac{b^3 - a^3}{b^3}\right]^{-1} \tag{2.31}$$

Here, H_0 is the ambient field, H_{in} is the field inside the shell, a and b are the inner and outer diameter of the shell respectively, and μ is the magnetic permeability of the shield. The finite-element model agreed with the analytical solution within 10% for relative permeabilities spanning six orders of magnitude, from 1 to 10^6; in general, the finite-element model underestimated the effectiveness of the screening effect, offering a conservative upper bound on the field interior to the shield.

Next, the geometry of the real shield was imported from a CAD model and the screening effect was simulated (Figure 2.16b). In particular, the degradation of the screening due to apertures required for laser beam access was investigated by calculation of the screening effect with and without the apertures; this led to a refined design that optimized the tradeoff of aperture size vs. screening factor. The shield attenuates external magnetic fields by at least a factor of 100 in all three directions.

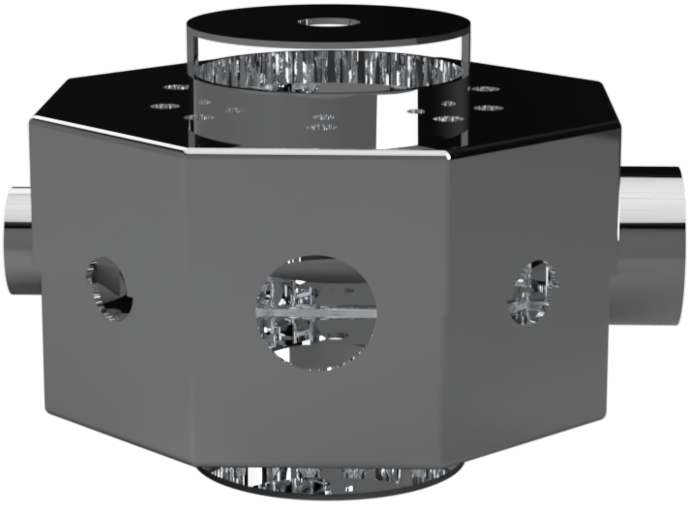

(a) Magnetic shield for the science chamber, constructed from two mu-metal halves which press-fit together. The shield has apertures for laser access, atomic beam entry/exit, and for airflow to the quadrupole coil heat sinks.

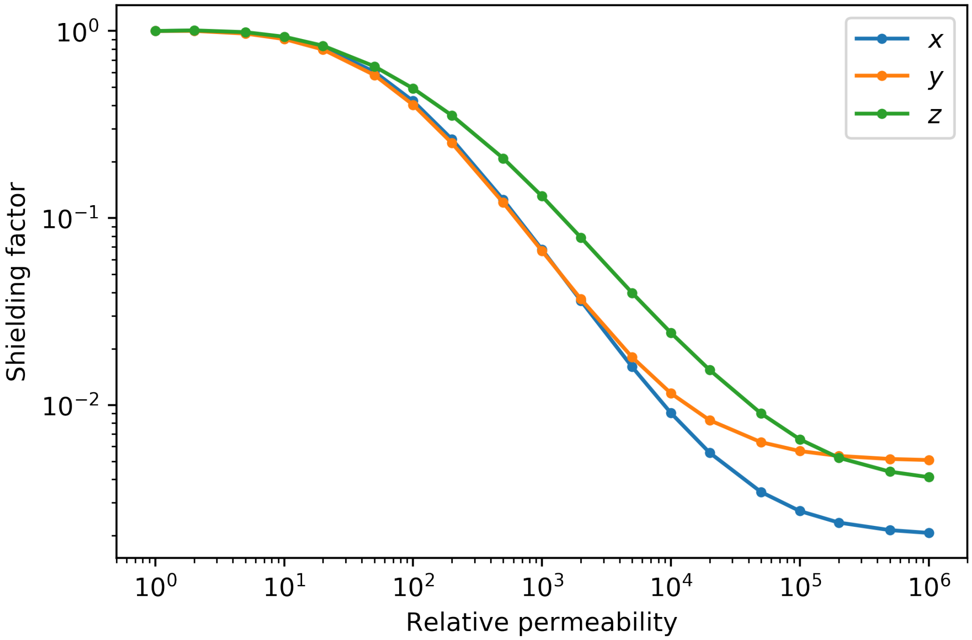

(b) Simulation of magnetic screening for the shield geometry. The shield is oriented with the vertical direction along z; the screening is worse along this axis due to cutouts required for airflow to the quadrupole coil heat sinks. The atomic beam axis is along y; the screening in this direction is slightly worse than along x due to the need to larger apertures to fit around the vacuum chamber.

Figure 2.16: Finite-element modeling of the magnetic shield.

2.10 Laser stabilization

In order to stabilize the light at 556 nm (stabilized at 1112 nm before frequency doubling), 759 nm, and 1388 nm, we used a 30 cm multi-spectral cavity (Stable Laser Systems) in a 5U rack module (Figure 2.17). The three colors are overlapped with dichroic beamsplitters and coupled into the cavity and the reflected and transmitted signals are monitored. A frequency lock is implemented with the electronic sideband technique [101] with an FPGA-based digital servo with closed-loop bandwidth up to 1 MHz [83]. The same digital servo is additionally used to stabilize the 399 nm laser to a vapor lamp using modulation transfer spectroscopy (Section 2.2).

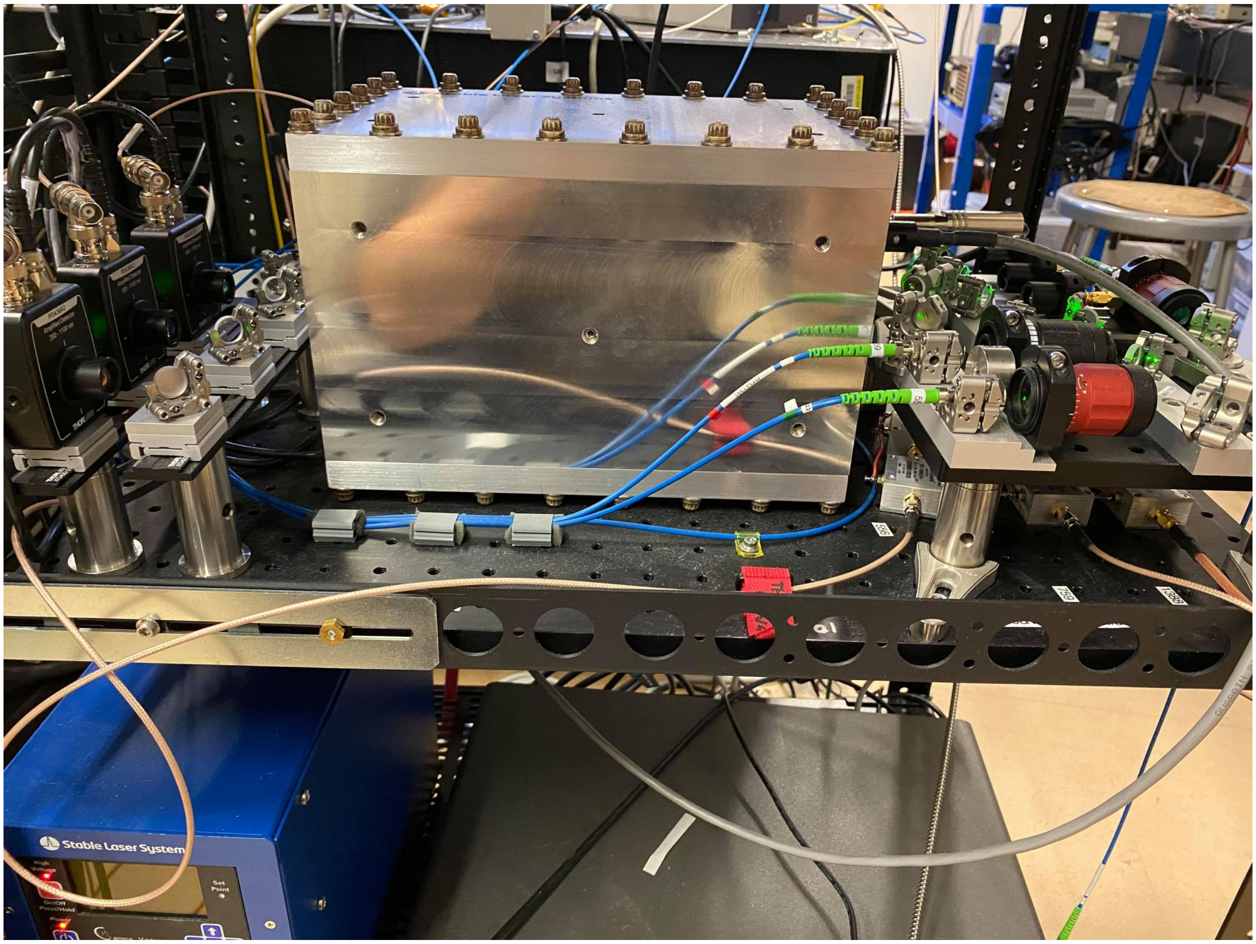

Figure 2.17: Multi-spectral cavity in a 5U rack module. The input fibers and photodiodes for reflection monitoring are on the right side, while transmission photodiodes are on the left.

The FPGA electronics are capable of generating the offset and sideband modulation signals up to 125 MHz, and by using a x8 frequency multiplier, over half of the 1.5 GHz free spectral range can be spanned, allowing the laser to be locked at any frequency. Additionally, the demodulation can be done digitally, not requiring any rf mixers or other components. Consequently, the total board footprint required for a complete single frequency lock is only about 270 cm^2 without any additional synthesizers or rf electronics.

The digital servo also implements a lock acquisition algorithm which can be used to recover from cavity unlocks. Either the dc level of the transmitted cavity signal or the rms amplitude of the reflected signal can be monitored, with the relocking algorithm triggered if the signal drops below a threshold. When this occurs, the integrator is disabled and the offset frequency is varied in a triangular wave of gradually increasing amplitude until the signal re-enters the threshold, at which point the integrator is re-enabled. It is possible that this procedure could reacquire lock to the wrong cavity mode[6] after a sufficiently violent unlock, so the frequency of each laser is monitored with a wavemeter, and a separate re-acquisition algorithm can be used to tune to the correct frequency prior to engaging the digital servo's relocking algorithm.

2.11 Experimental control and acquisition

Laboratory clocks typically require human operators to monitor many variables, such as laser frequencies and optical power levels, and to intervene when an issue is detected (in these examples, re-locking a cavity or re-aligning a beam into a fiber). After several years of experience operating our laboratory clocks, the anticipation of additional challenges outside of a well-controlled lab offered an exciting opportunity to re-imagine clock control architecture from the ground up. We have developed a sophisticated controls system allowing computer control and/or monitoring of many degrees of freedom and parameters which are expected to be relevant to clock operation.

Broadly speaking, the experiment relies on two classes of control variables. The first class consists of timing-critical variables that must be manipulated within each experimental cycle, like

---

[6] Or, for the 399 nm lock, to the wrong isotope

TTL signals used to open or close optical shutters. The second class consists of “set-and-forget” variables which are constant over many experimental cycles, such as the temperature setpoint for the second-harmonic-generation crystal in the 556 nm module. Separate controls systems have been built for these two classes of variables, as well as a middleware layer allowing the two systems to interact.

2.11.1 ARTIQ

The timing-critical variables are handled using ARTIQ [48], an open-source software and hardware platform whose development was largely funded by and developed in collaboration with the Ion Storage group at NIST. At the heart of the ARTIQ hardware ecosystem is an FPGA board (Kasli) which is responsible for guaranteeing precise control of a number of peripheral boards. In our experiment, we use peripheral boards for TTL signals, direct digital synthesis (DDS) of AOM drive frequencies, and digital-to-analog or analog-to-digital conversion.

In order to define experiments executed by ARTIQ, a simple but powerful programming language is offered, which is a subset of Python. We have further simplified experimental control by developing an application called Argent[7] , which allows timing sequences to be defined in a JSON representation and automatically converted to usable ARTIQ code. Argent offers a browser-based interface developed with JavaScript/React with which sequences can be created graphically. Any number of timesteps of specified durations can be accompanied by changes in TTL logic levels, output voltages, DDS frequencies and/or amplitudes. Input voltages can also be sampled at discrete streams spanning part or all of a timestep, and various reduction operations (mean, min, max, etc.) can be designated using the interface.

For more sophisticated experimental control than simple repetitive sequences, Argent implements a *variables* system which allows certain parameters to be updated after each cycle by querying the Argent server. This allows the experimental parameters to be controlled by external, asynchronous processes; this is generally accomplished using a Python client in a Jupyter notebook

---

[7] https://github.com/robertfasano/argent

environment. Applications of this system range from simple sweeps of one or more parameters to Bayesian optimization of experimental outcomes such as atom number. The variables system allows external control of timestep durations, DAC voltages, and DDS frequencies and amplitudes.

Output variables can also be defined to capture experimental results and send them to the server at the end of each experimental cycle. The server uses the SocketIO framework to emit signals which other external services can subscribe to, e.g. to write experimental results into a database.

While the Argent user interface allows great flexibility in sequence construction, there are some use cases which are more complicated than simple TTL/DAC/DDS/ADC events; for example, clock operation entails computation of a normalized excitation probability each cycle and determination of a frequency correction based on the results of multiple cycles. To accommodate use cases like this, Argent allows users to designate scripts containing any valid ARTIQ code to be selected and run at the end of each experimental cycle.

2.11.2 LabAPI

For non-timing-critical experimental parameters, we developed an object-oriented Python-based software framework called LabAPI[8] . LabAPI allows the definition of abstract objects representing experimental parameters which can be interfaced with physical devices using a simple syntax. LabAPI supports various types of parameters for different use cases:

(1) Float-valued parameters, e.g. coil current or laser frequency

(2) Bool-valued parameters, e.g. engage/disengage options for a frequency servo

(3) Multiple-choice parameters, e.g. a channel selector for a fiber-optic switch

(4) Read-only parameters representing experimental results or monitored variables, e.g. atom number or chamber pressure

---

[8] https://github.com/robertfasano/labAPI

The overarching goal of LabAPI is to provide a common syntax for controlling all devices involved in an experiment, and to leverage this uniformity to provide powerful features to the experimentalist. Although LabAPI can be run in an isolated kernel, it also includes a client-server framework allowing experiments to be decentralized across multiple processes or computers. When the server is started, it automatically detects defined experimental parameters and makes them accessible through a REST API, allowing the parameters to be controlled from external processes.

LabAPI also includes an optimization framework which allows experiments to be optimized using only a few lines of code. Many algorithms are supported, from simple grid searches (i.e. brute force optimization) to response surface modeling *a la* Bayesian optimization. The latter technique is a powerful tool for optimization of complex, potentially high-dimensional, experiments, and at its core it can be considered an automation of the scientific method. Starting with an initial population of $X \rightarrow y$ points, where X is a state vector representing the independent variables of the experiment and y is the outcome under optimization, a response surface is constructed which models the experimental landscape. Then, an intelligent sampling strategy is used to choose points to request from the experiment which are expected to most efficiently improve knowledge of the response surface [111]. This is done by numerically minimizing an *effective cost* on the response surface itself, defined by

$$c(X) = b\mu(X) - (1-b)\sigma(X) \tag{2.32}$$

Here, $\mu(X)$ and $\sigma(X)$ are the mean and uncertainty of the response surface at the point X, and b is a parameter controlling the *exploration-exploitation* tradeoff of the sampling strategy. If $b = 1$, then the effective cost is identical to the actual cost function that the response surface is modeling; this is referred to as *exploitation*, as the sampling strategy will exploit its knowledge of the cost function to find the optimum point. On the other hand, if $b = 0$ (*exploration*), then minimizing the effective cost will locate regions of high uncertainty; subsequently performing a physical measurement of this point is expected to yield large improvements in the model. An effective overall sampling strategy involves ramping from $b = 0$ to $b = 1$ over a prescribed number of experimental cycles. This way, the optimizer will start off by exploring uncertain areas to globally refine the model before meandering towards the optimum location.

Bayesian optimization can be an especially powerful tool for exploring high-dimensional cost landscapes with coupled parameters where human guidance is difficult, and has been successfully deployed in several other cold atom experiments [111, 4]. An illustrative example is the compression sequence during second-stage cooling at 556 nm (Figure 2.18). Prior to deliberate compression, we wanted to simply determine the optimal parameters which maximize atom number at the end of the ramp. Starting from zero knowledge of the parameter space aside from bounds for each parameter, the algorithm is able to optimize the six-dimensional cost function (start and stop points for magnetic field gradient, laser intensity, and laser detuning), converging to the global minimum in about 500 experimental cycles (≈5 minutes).

2.11.3 Monitoring

Most variables which are important for clock operation are measured each experimental cycle. Some of these measurements need to be performed with extremely precise timing; for example, the power of the probe beam is measured within a several-millisecond window while the probe light is delivered to the physics package. These stringent measurements are performed with ARTIQ and sent to the Argent server each cycle. The server emits a callback which triggers a LabAPI snapshot, measuring non-timing-critical variables such as laser frequencies. All of these measurements, both timing-precise and non-timing-precise, are written into an InfluxDB database for long-term persistence. A dashboard built with Grafana is used to monitor especially critical parameters like laser frequencies. These parameters have defined alert thresholds which are used to determine whether the clock is “locked” (operating properly) or “unlocked”; if an alert threshold is breached, Grafana issues an HTTP request to LabAPI which can trigger recovery functions. For example, the spectrometer lock for the 399 nm light can jump to the wrong isotope if perturbed. When this is detected, an optimization algorithm is used to tune the laser back to the correct isotope using a feedforward correction to the piezo voltage.

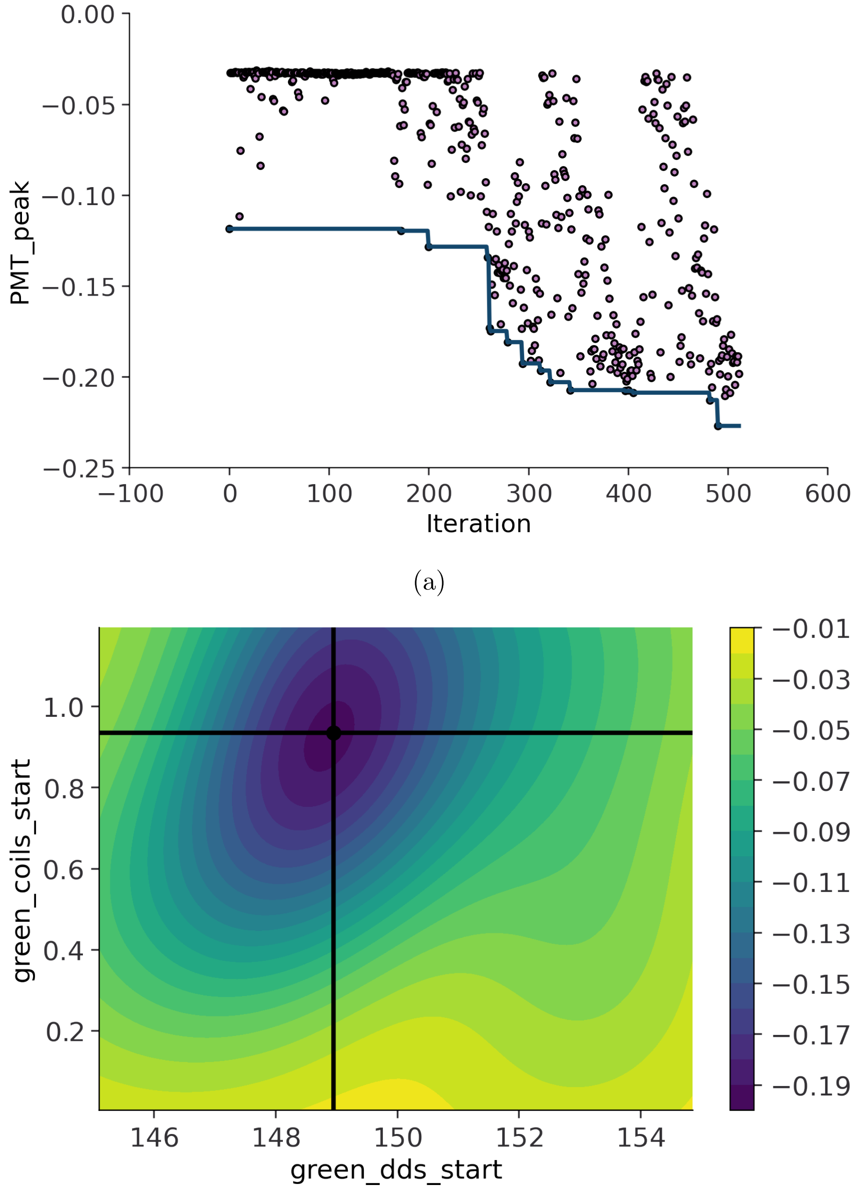

(a)

(b)

Figure 2.18: Six-dimensional optimization of start and stop setpoints for a linear ramp of magnetic field gradient and laser intensity and detuning during the compression stage of the 556 nm MOT. a) Convergence plot showing the improvement in the cost function (PMT peak voltage during the probe pulse, minus sign included by convention) as more experimental cycles are taken. The dark blue line shows the cumulative best performance. b) 2D slice through the optimum point of the 6D cost landscape. The x-axis is the starting AOM frequency (relative to detuning at around 154 MHz). The y-axis is the setpoint of our current servo used to define the MOT gradient, with a value of 1 being equal to approximately 35 G/cm axial gradient.

2.11.4 Auto-alignment

Beam pointing misalignment due to environmental perturbations can manifest in changes in the light delivered to the physics package via optical fiber, which can degrade the signal-to-noise of the clock stabilization scheme or require downtime to re-optimize the fiber coupling. To mitigate this effect, we installed MEMS mirrors (Mirrorcle Technologies) in critical beam paths. A closed-loop algorithm is used to tune the tip and tilt of the mirrors to maximize transmitted power measured by a photodetector in the physics package. Several different algorithms have been implemented successfully, including gradient descent [88], Nelder-Mead/simplex method [81], and Bayesian optimization [30]. Figure 2.19 shows an example optimization trajectory which uses the Nelder-Mead method to recover fiber coupling after intentional misalignment in just several seconds.

The MEMS mirrors are steered by the application of equal and opposite voltages to pairs of pins for the x and y directions. The manufacturer offers both analog and digital driver boards. The analog boards simply possess high-voltage amplifiers to amplify user-supplied voltages, followed by a Bessel filter to smooth out the control voltage before application to the mirror. The digital boards include all of the above, plus a digital-to-analog converter (DAC) before the amplifiers that can be programmed via the serial peripheral interface (SPI) protocol to set the applied voltages. We chose the digital variant to minimize the risk of unwanted mirror motion due to rf interference picked up by the control lines. We communicate with the driver board with a LabJack T7, which is also used for sampling of a photodiode voltage to determine the fiber coupling efficiency with each iteration of the optimization. A custom breakout PCB was designed to interface the LabJack T7, photodiode, and MEMS driver board with standard multi-pin and SMA connectors.

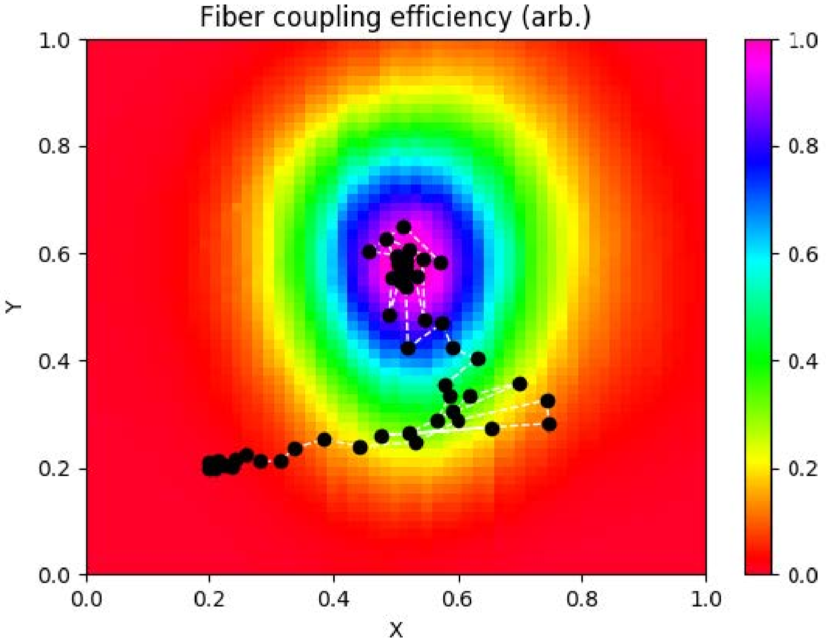

(a) Auto-align trajectory using a Nelder-Mead optimizer (solid points connected by dashed white line), overlaid on a measurement of the transmitted power vs. the coordinates X and Y obtained by grid search.

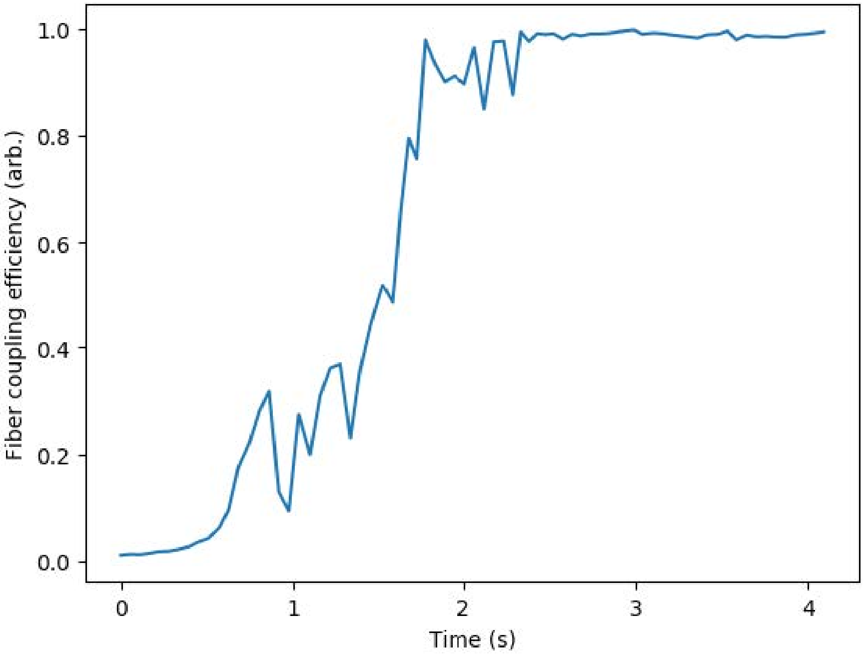

(b) Convergence plot of the optimization in a), demonstrating full recovery of fiber coupling in seconds.

Figure 2.19: Demonstration of automated alignment systems.

Chapter 3

Systematic uncertainties

In this chapter, we discuss the leading systematic effects in optical lattice clocks, with particular emphasis on effects which are exaggerated in transportable systems, as well as effects which constrain the design of transportable systems. An exhaustive uncertainty evaluation is beyond the scope of the present thesis; rather, we seek to establish a clear path to eventual 10^{-18} operation and show that no significant issues will arise from the transportable architecture. For a full discussion of relevant uncertainties at the 10^{-18} level in a Yb optical lattice clock, I refer the reader to our laboratory clock evaluation [69].

3.1 Lattice light shift

The optical lattice potential used to confine ultracold atoms during spectroscopy causes large ac Stark shifts of the ground and excited states of the clock transition. The potential is typically evaluated using a multipole expansion, with a lowest-order term from $E1$ polarizability given by

$$U_{E1} = -\frac{P}{\epsilon_0 c \pi w_0^2} \alpha_{(i)}(\nu_l) \tag{3.1}$$

Here, P is the laser power, w_0 is the beam radius at $1/e^2$ intensity, ϵ_0 is the permittivity of free space, and $\alpha_{(i)}(\nu_l)$ is the $E1$ polarizability of the electronic state $|i\rangle$ at a laser frequency ν_l. This potential results in a shift of the clock transition

$$\Delta\nu_{clock} = -\frac{P}{\epsilon_0 h c \pi w_0^2} \Delta\alpha(\nu_l) \tag{3.2}$$

where $\Delta\alpha(\nu) = \mathrm{Re}[\alpha_{(e)}(\nu) - \alpha_{(g)}(\nu)]$ is the differential polarizability between the two clock states. The E1 polarizabilities of the ground and excited states can be calculated using known atomic state parameters:

$$\alpha(\nu)_{(j)} = \frac{3\epsilon_0 c^3}{(2\pi)^3} \sum_i \frac{\Gamma_{ij} f_{ij}}{\nu_{ij}^4 - \nu^2\nu_{ij}^2 - 2\pi i \Gamma_{ij}\nu^3}, \tag{3.3}$$

where the sum is taken over all intermediate states coupled to the state of interest and j represents either the ground or excited state. Here, Γ_{ij}, f_{ij}, and ν_{ij} are the linewidth, branching ratio, and frequency of the transition from the ground state $|g\rangle$ or the excited state $|e\rangle$ to the intermediate state $|i\rangle$.

One of the greatest "tricks" which has enabled the modern supremacy of optical lattice clocks is operation at a so-called *magic wavelength* where the $E1$ polarizabilities of the ground and excited states are equal, and therefore the differential polarizability and the clock shift vanish. This technique was first demonstrated in 2003 [49], predicting suppression of the shift to as small as 5×10^{-18}, at which point higher-order terms in the multipole expansion become relevant. Later, it was discovered that these higher multipolar terms could be complicated by thermal atomic motion through the spatially-inhomogeneous lattice light field [96]. Additional shifts could arise from anharmonic corrections to the lattice potential or cross-dimensional couplings. Through the development of more sophisticated models combined with experimental light shift measurements, we characterized these effects at the $< 10^{-18}$ uncertainty level in our non-transportable clocks [16]. Motivated by empirical scaling of the atomic temperature with trap depth, the shift was parameterized with an effective model

$$\frac{\Delta\nu_{\text{clock}}}{\nu_{\text{clock}}} = -\alpha^*(\nu_l)U - \beta^* U^2, \tag{3.4}$$

where $\alpha^*(\nu)$ and β^* are the thermally-averaged polarizability and hyperpolarizability respectively and U is a dimensionless trap depth parameter:

$$U = \frac{P}{\epsilon_0 c \pi w_0^2 E_r}\alpha(\nu_{magic}) = \left(\frac{h\nu_z}{2E_r}\right)^2, \tag{3.5}$$

Here, $E_r = h^2\nu_l^2/2mc^2$ is the photon recoil energy of the lattice laser and ν_z is the longitudinal trap frequency, which can be determined with motional sideband spectroscopy. The conventional magic

Figure 3.1: a) Polarizability for the ground and excited states. b) Differential polarizability from the model in (a) (solid) and direct measurements (dashed).

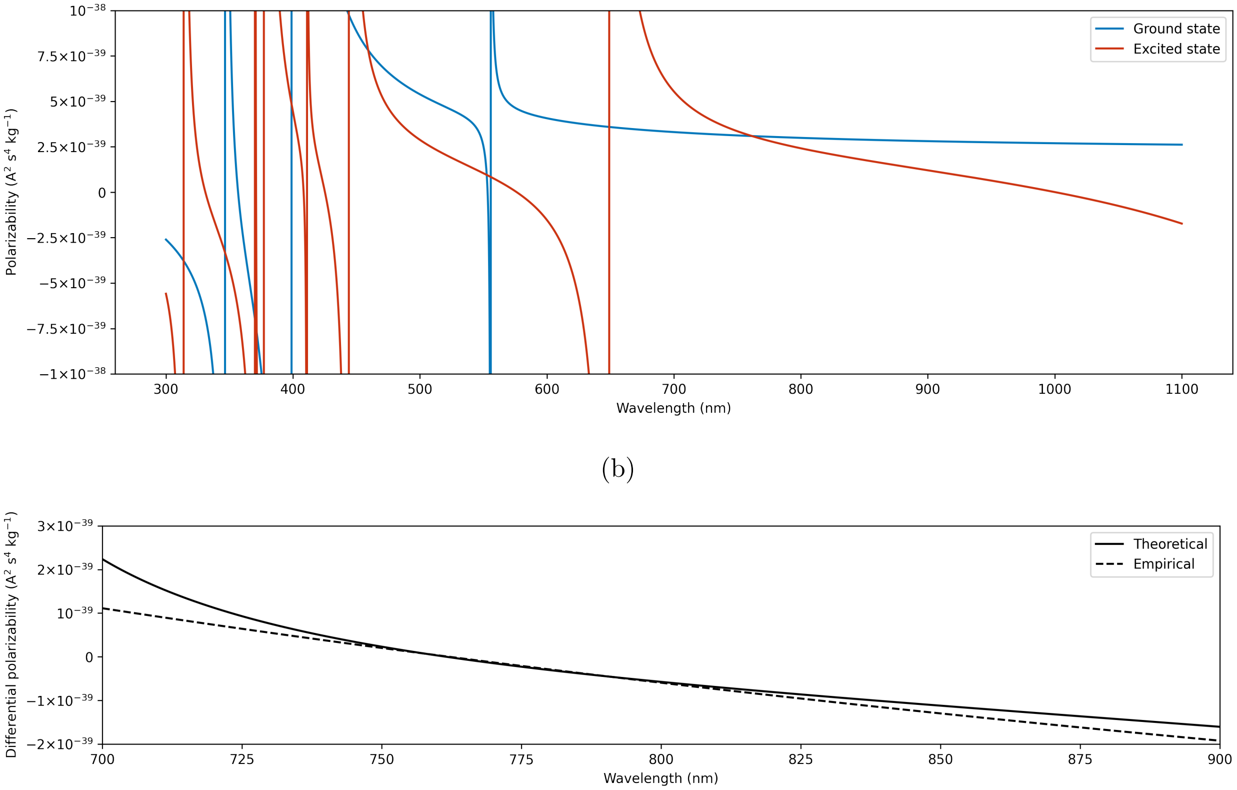

(a)

(b)

wavelength ν_{zero} can be identified as the point where $\alpha^*(\nu) = 0$. Using this effective model, we can also identify an *operational magic wavelength* ν_{opmagic}=394,798,267.7(5) MHz, at which the shift is insensitive to changes in trap depth due to partial cancellation between the linear and quadratic terms of (3.4), allowing the shift to be controlled within an experimental uncertainty of 8×10^{-19}.

With the operational magic wavelength carefully measured using our non-transportable clocks, it is fairly straightforward to reach similar lattice light shift uncertainty in a transportable clock using the same lattice laser by operating under similar conditions. This is especially true because of a useful feature of ^{171}Yb, which has nuclear spin $F = 1/2$ - the absence of tensor light shifts [2], which otherwise introduce geometric dependence and typically requiring separate characterization in each atomic apparatus. To achieve uncertainty in the 10^{-18} decade, the dominant consideration is simply operating at the magic wavelength ν_{zero} and similar trap conditions; therefore, all that is required is either an absolute frequency measurement of our lattice laser using a frequency comb, or a relative frequency measurement with a heterodyne beat against the lattice laser used for Yb1 and Yb2. However, in practice, one must be careful to also assess residual non-magic light shifts as described in the following section.

3.1.1 Background light shifts

The choice of an amplified diode laser to generate the optical lattice in the transportable clock, as opposed to a more conventional Ti:S laser, bears an important consequence: additional light shifts due to broadband amplified spontaneous emission, as shown in Fig. 3.3. These shifts have caused significant experimental challenges for clock operation [56, 55, 53, 10] due to their large magnitude - with typical amplifiers, background-induced light shifts at the Hz level (10^{-15} fraction of the clock frequency) or much higher are easily possible. There have been various attempts to mitigate ASE effects through amplifier optimization [115, 105], spatial filtering with optical fibers [13, 79], or spectral filtering using volume Bragg gratings [16, 53, 98], Fabry-Perot etalons [8], or absorption cells [26]. For an additional review of the origins and consequences of ASE, as well as an in-depth investigation of several strategies for ASE characterization, see [115]. In this section,

Figure 3.2: Lattice ac Stark shift characterization. a) Fractional clock shift vs. trap depth for different detunings from ν_{zero} ranging from -50 MHz (dark red) to 30 MHz (dark blue). b) Effective polarizability entering into equation (3.4) vs. detuning from ν_{zero}. Reproduced from [16].

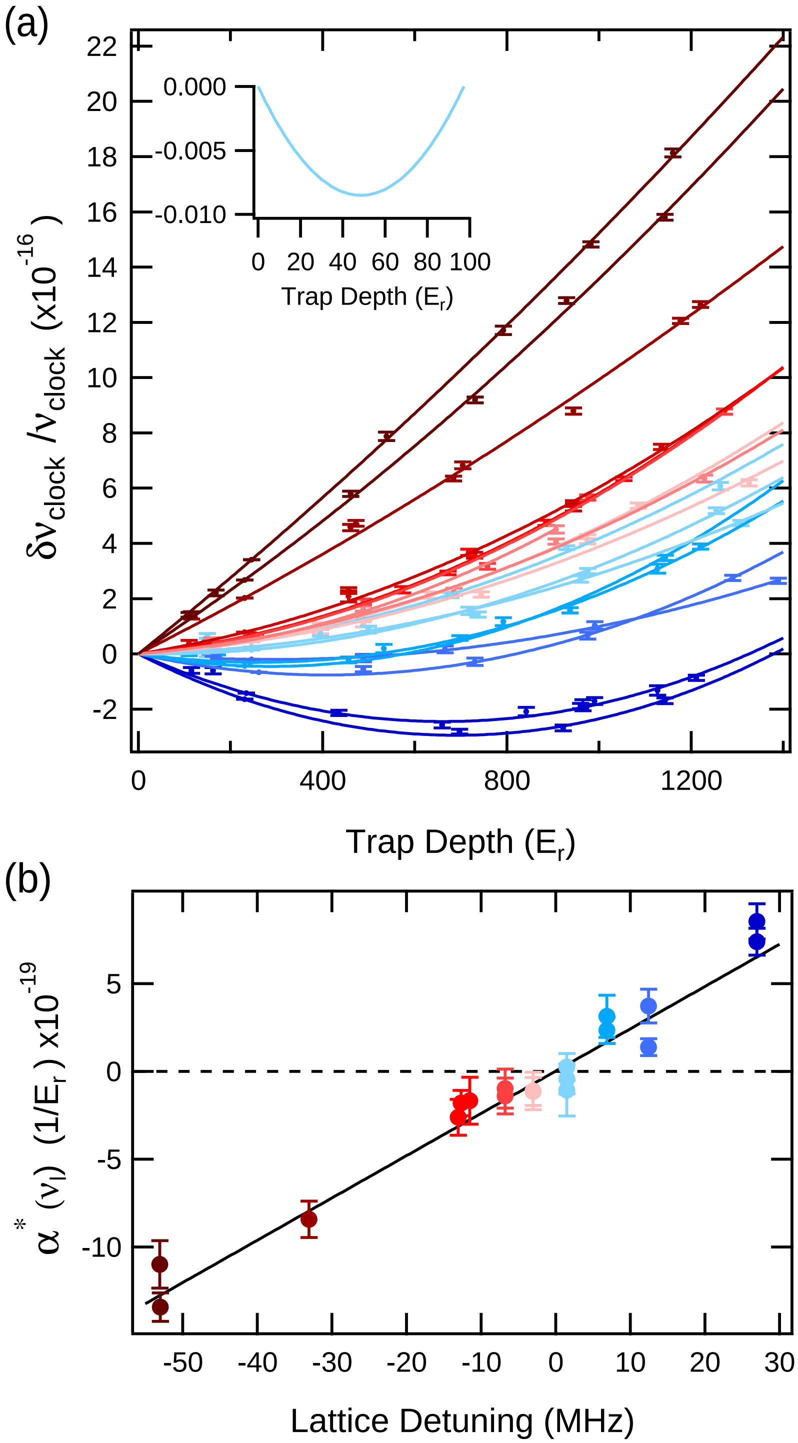

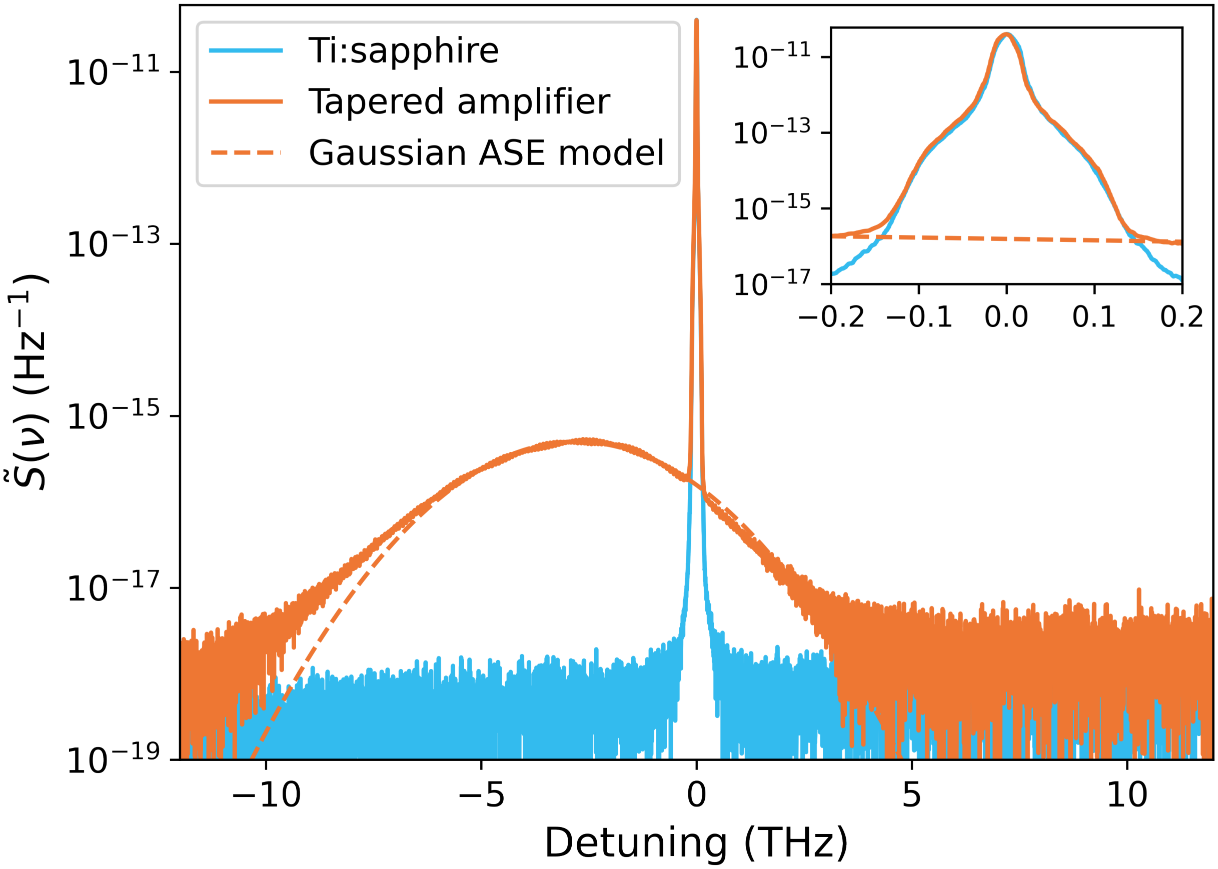

Figure 3.3: Optical spectrum analyzer measurements of background spectra of the Ti:S and TA systems as a function of detuning relative to the magic wavelength. Near the peak, the achievable resolution is limited by the optical rejection of the OSA, shown in the inset.

we present our own efforts for characterization and suppression, adapted from [28].

We assume that the laser spectrum can be described by two components: a spectrally-narrow "carrier" containing power P at a frequency ν_{trap} and a broadband background described by the power spectral density $S(\nu)$. By comparison with the ac Stark shift for a monochromatic laser spectrum given in (3.2), we can write the background light shift for a broadband spectrum:

$$\Delta\nu_{background} = -\int_0^\infty \frac{1}{\epsilon_0 hc\pi w_0^2}\Delta\alpha(\nu)S(\nu)W(\nu)d\nu, \tag{3.6}$$

Here, $W(\nu)$ is a weighting factor accounting for spatial dephasing of the ASE relative to the standing wave formed by the carrier. After reflecting from the lattice retroreflector, the ASE forms a continuum of standing waves that are generally incoherent with the trap. Therefore, the average

ASE intensity experienced by trapped atoms will be reduced. We estimate this effect by assuming that atoms are trapped in discrete positions $x_n = n\lambda_{magic}/2$ over a spatial distance d. The ASE intensity at a wavelength λ is therefore weighted by a factor

$$W(\lambda) = \frac{1}{N}\sum_{n=-N/2}^{N/2}\cos^2\left(\frac{\pi n\lambda_{magic}}{\lambda} + \phi(\lambda)\right) \tag{3.7}$$

where $N = 2d/\lambda_{magic}$ and

$$\phi(\lambda) = 2\pi\left(\frac{D}{\lambda} - \frac{D}{\lambda_{magic}}\right) \tag{3.8}$$

is the differential phase accumulated between the trap and ASE at $x = 0$. The round-trip distance between the center of the optical lattice and the retroreflector is given by D. The weighting factor converges to 1 at $\nu = \nu_{magic}$ and oscillates around 1/2 elsewhere with a period and decay rate depending on the size of the trap. As the divergence from 1/2 is only significant near the magic wavelength, where the differential polarizability is small, we will assume $W(\nu) = 1/2$ for all subsequent calculations.

It is convenient to instead express the clock shift in terms of dimensionless parameters separately describing the trap, the background spectrum, and the atomic response:

$$\frac{\Delta\nu_{clock}}{\nu_{clock}} = -\frac{1}{2}U\int_0^\infty \tilde{S}(\nu)\Delta\alpha'_{E1}(\nu)d\nu, \tag{3.9}$$

where

$$\Delta\alpha'_{E1}(\nu) = \frac{E_r}{h\nu_{clock}}\frac{\Delta\alpha(\nu)}{\alpha(\nu_{magic})} \tag{3.10}$$

and $\tilde{S}(\nu) = S(\nu)/P$ is the background power spectral density relative to the peak power.

3.1.2 ASE-induced light shifts

The ASE spectrum of the amplified diode laser in Fig. 3.3 is well fit by a Gaussian model,

$$\tilde{S}(\nu) = \tilde{S}_0\exp\left(-\frac{(\nu-\nu_{ASE})^2}{2\sigma_{ASE}^2}\right) \tag{3.11}$$

The parameters of this profile, and thus the resulting background light shift, can be tuned with the amplifier current and temperature. Except where specified otherwise, all measurements in this work are taken at an operational setpoint of 3210 mA and 20 °C, at which the spectrum is described by $\nu_{ASE} = \nu_{magic} - 2.8$ THz, $\tilde{S}_0 = 4.7 \times 10^{-16}$ Hz^{-1}, and $\sigma_{ASE} = 2.1$ THz.

Together with this empirically-observed Gaussian profile, we can obtain a closed-form estimate of the shift from (3.9) by making a linear approximation of the differential polarizability as

$$\Delta\alpha'_{E1} = \frac{d\Delta\alpha'_{E1}}{d\nu}(\nu - \nu_{magic}) \tag{3.12}$$

with $d\Delta\alpha'_{E1}/d\nu = 4.5(3) \times 10^{-26}$ Hz^{-1} [16]. This approximation agrees with (3.10) to better than 5% within the typical full-widths at half-max of ASE spectra observed in this work. Under these two approximations, the shift is

$$\frac{\Delta\nu_{clock}}{\nu_{clock}} \approx -\sqrt{\frac{\pi}{2}} U \frac{d\Delta\alpha'_{E1}}{d\nu} \tilde{S}_0 \sigma_{ASE} (\nu_{ASE} - \nu_{magic}) \tag{3.13}$$

With this expression, the fractional light shift can be estimated based on spectral measurements of the ASE profile. Notably, the shift scales as the product of the amplitude, width, and peak detuning of the ASE spectrum. For general non-Gaussian background spectra, the fractional shift can instead be evaluated numerically.

To validate our simple model given by equation (3.9), we undertook a series of shift measurements while varying the amplifier current and temperature in one of our non-transportable Yb optical lattice clocks (Yb1). These measurements interleaved (i) generating the optical lattice with the TA versus (ii) an approximately background-free Ti:S at the same laser frequency, the latter of which is filtered by a volume Bragg grating bandpass filter with a 26 GHz bandwidth. The lattice intensity was controlled with feedback to acousto-optic modulators for each laser using a common photodetector signal and control electronics. The TA was phase-locked to the Ti:S, which itself was referenced to a stable optical cavity. A liquid crystal waveplate and polarizer was used to select between the two lasers, and the inactive beam was further attenuated by switching off the rf drive from an acousto-optic modulator in that beam's path. The active beam was coupled into a single optical fiber and delivered to the atomic system, passing through a polarizing beam splitter to provide linearly polarized light for the trap. This configuration ensured that the optical lattices formed

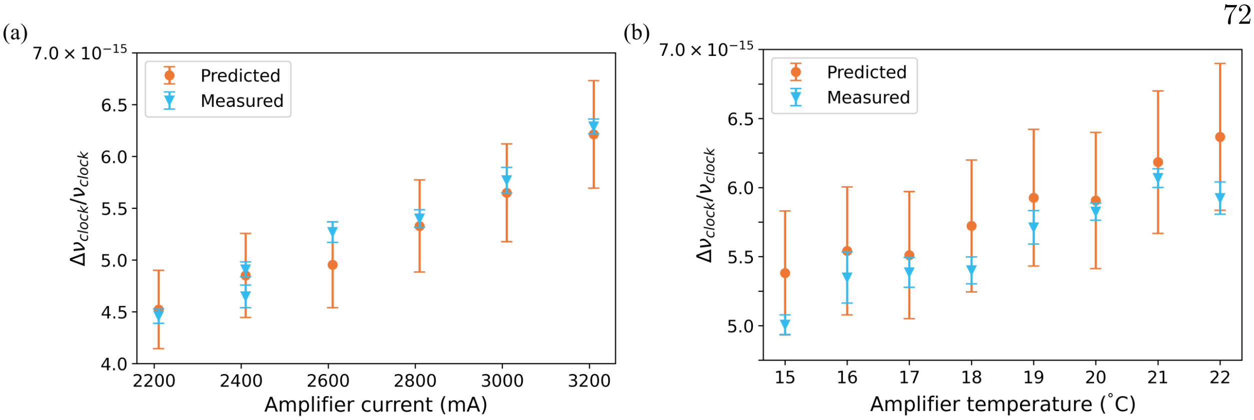

Figure 3.4: Comparison of measured light shifts with predictions based on spectral characterization while varying a) amplifier current and b) amplifier temperature. The error bars represent 1σ uncertainties given by the Allan deviation of the frequency difference between the TA and Ti:S. Uncertainties in predicted points are computed with standard propagation of error techniques based on uncertainties in the Gaussian model parameters, differential polarizability, and an assumed 5% uncertainty in trap depth.

by the two lasers had equal frequency, intensity, polarization, and spatial geometry. Furthermore, atomic cold collision shifts, which can substantially complicate lattice shift characterization [76], were suppressed in common mode due to the identical lattice conditions in the two cases.

Through motional sideband spectroscopy, the trap depth was assessed to be $U = 53$. The trap depth was also scaled by an experimentally-determined factor of 0.71 to account for thermal averaging of the laser intensity by the motional distribution of atoms within the lattice [16].

At each amplifier current and temperature, we additionally measured the ASE profile with an optical spectrum analyzer and extracted characteristic parameters with a Gaussian fit. These parameters are used to estimate the shift using (3.13) with an uncertainty given by standard propagation of error.

The predictions are in good agreement with the measured shifts (Fig. 3.4), demonstrating approximately linear dependence of the shift with amplifier current and temperature with slopes of 1.8×10^{-15}/A and 1.6×10^{-16}/K respectively. Therefore, to control the shift at the 10^{-18} level, the amplifier current and temperature would need to be stable at the 600 μA and 6 mK level.

Although this degree of control is practically realizable, previous efforts to characterize and control background light shifts have noted large time-varying effects [55, 10]. We observed substantial dependence of the shift on the alignment into the fiber used to deliver light to the atoms (NKT LMA-PM-15 large-mode-area photonic crystal fiber). Measurements of the ASE spectrum while intentionally degrading the alignment showed that significant spectral distortion could arise, which we attribute to a wavelength-dependent fiber coupling due to the spatial mismatch of the ASE and the fiber. This effect is discussed in greater detail in Section 3.1.8.

For the optical clock measurements in Fig. 3.4, the fiber pickoff for OSA measurements was located near the atoms, ensuring that any distortion induced by the photonic crystal fiber would be accounted for in the shift prediction. A multi-mode fiber was used to couple light to the OSA, and we verified that the observed laser spectrum was substantially less sensitive to misalignment than a single-mode alternative with 4.5 μm core diameter.

However, while these efforts make it possible to estimate the resulting background light shift at a given instant, degradation in fiber coupling at the few percent level modifies the shift by greater than 10^{-16}. In contrast, control of the amplifier-dependent shifts at the 10^{-16} level requires only modest efforts to control the temperature and current. Therefore, we find that achieving robust control of background light shifts is primarily limited by beam pointing stability and virtually impossible at the 10^{-18} level. Instead, these shifts must be suppressed, such as realized with spectral filtering below.

3.1.3 Spectral filtering

Background light shifts can be reduced with optical bandpass filters centered at the magic wavelength. Various techniques have achieved passbands with <1 nm bandwidth, including complex prism geometries [3], atomic vapor filters [68], and various types of reflection gratings [67, 34] (potentially in conjunction with Fabry-Perot etalons [65]). Due to their relative simplicity and compact size, we have chosen a volume Bragg grating. Near normal incidence, the reflection transfer function is approximately [40]

$$R(\nu) = \frac{\sinh^2 \phi(\nu)}{\cosh^2 \phi(\nu) - \Gamma(\nu)^2} \tag{3.14}$$

where

$$\phi(\nu) = \frac{2\pi d\nu}{c} \sqrt{\frac{n_1^2}{4} - \left(\frac{\nu_0}{\nu} - 1\right)^2} \tag{3.15}$$

and

$$\Gamma(\nu) = \frac{2}{n_1}\left(\frac{\nu_0}{\nu} - 1\right) \tag{3.16}$$

Here, n_1 is the index modulation depth and ν_0 is the center frequency. Two distinct behaviors are evident in this transfer function. Near the center frequency ($\nu \approx \nu_0$), the transfer function is approximately constant and reaches a peak reflectivity

$$R(\nu_0) = \tanh^2 \frac{\pi d n_1 \nu_0}{c} \tag{3.17}$$

Far from the design frequency, the transfer function asymptotically varies as

$$R(\nu) \approx \frac{n_1^2}{4} \frac{\sin^2[\frac{\pi d}{c}(\nu - \nu_0)]}{(\frac{\nu_0}{\nu} - 1)^2}, \tag{3.18}$$

which is the product of a periodic modulation and a second-order filter rolloff. The transition point between the pass and rejection bands occurs where (3.15) changes from real to imaginary, at $\nu = 2\nu_0/(2 \pm n_1)$. The bandwidth can thus be approximated as $\Delta\nu_{BW} \approx \nu_0 n_1$. Representative values for gratings used in this work are $n_1 = 2.3 \times 10^{-5}$ and $d = 22$ mm, resulting in a theoretical bandwidth of 9.2 GHz and a peak reflectivity of 94%.

Volume Bragg gratings induce a wavelength-dependent distortion of the spatial beam profile [40]. To avoid complicated spatially-varying shifts in an optical lattice, we clean the beam profile with a single-mode optical fiber between the grating and the atoms. The actual transfer function is therefore the product of the spectral selectivity of the grating with the distortion-induced coupling losses of the fiber. This combined transfer function was experimentally characterized by measuring the reflected and subsequently fiber-coupled power from a Ti:S beam incident on the grating while varying the Ti:S frequency over several THz (Fig. 3.6). By adjusting the collimation of the incident

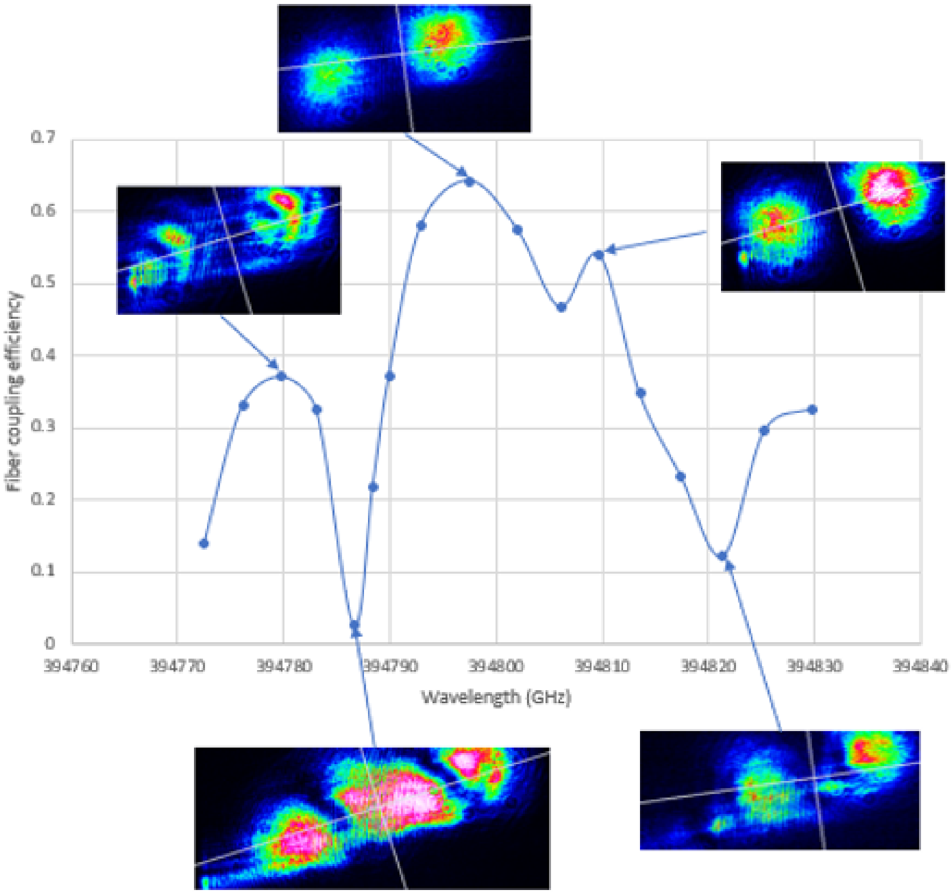

Figure 3.5: Observation of wavelength-dependent distortion of the diffracted beam, causing a change in fiber coupling efficiency independent of the grating diffraction efficiency. The beam images include ghost images to the left of the primary image due to reflections from an ND filter on the camera.

beam, a peak in-band reflectivity of 98% was achieved. The observed transfer function has a 20 dB/decade rolloff outside of an 11.5 GHz passband. The asymmetric ripples arise due to fiber coupling losses from beam distortion (Figure 3.5), while the characteristic modulation given by the idealized numerator of (3.18) is absent.

The light shift from the filtered background spectrum can be estimated by applying the filter transfer function (3.14) to the ASE profile (3.11) and numerically integrating the shift spectral density (3.9); we treat the experimentally-observed departures from the theoretical transfer function as an uncertainty in Section 3.1.4. For typical operating parameters of 3210 mA amplifier current

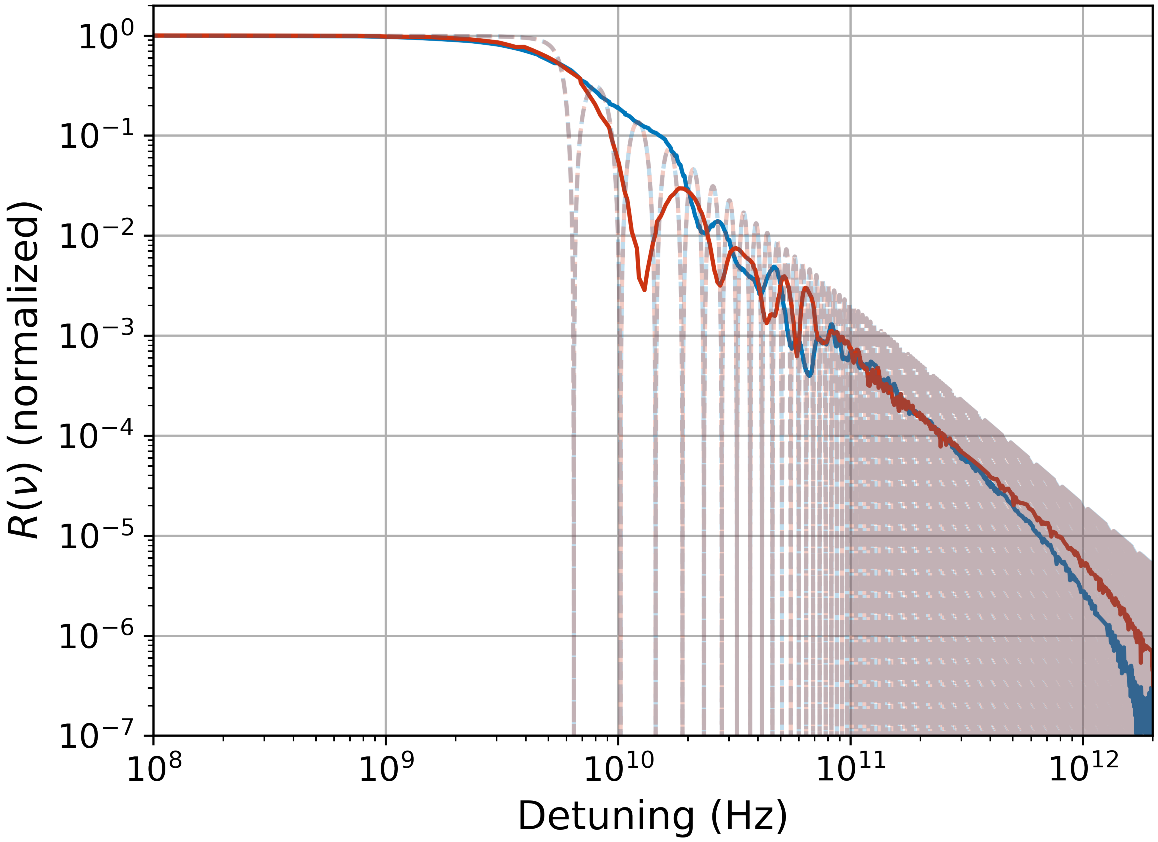

Figure 3.6: Volume Bragg grating transfer function, measured by varying the Ti:S frequency and measuring the reflected power. Positive detuning is shown in blue and negative in red. The theoretical profile is shown in light dashed lines.

and 20 °C temperature, the predicted shift is 8×10^{-21} for the experimental trap depth $U = 53$, including the thermal averaging factor of 0.71. We note that another recent work has estimated ASE light shifts below the 10^{-18} level using a VBG bandpass filter with 35 pm bandwidth in combination with a cavity with 6 MHz linewidth [80]. In either case, as discussed below, these calculations rely on a number of assumptions that must be verified to ensure reliable shift suppression at the predicted levels.

3.1.4 Background light shift uncertainty

While the theoretical framework outlined in Sections 3.1.2 and 3.1.3 predicts shifts that can be made negligible compared to the uncertainty budgets of current state-of-the-art clocks, it relies on several empirically-motivated assumptions and is potentially vulnerable to fluctuating systematic effects. In this section, we analyze these assumptions and effects to determine a conservative upper bound on the background light shift during normal operating conditions. All uncertainties are presented for the experimental trap depth $U = 53$ as described in Section 3.1.2.

3.1.5 Near-carrier spectrum

Within 100 GHz of the magic wavelength, our ability to resolve the spectrum is limited by the filter shape and optical rejection ratio of the optical spectrum analyzer. This limitation is referred to as the "close-in dynamic range," which is specified to be 60 dB at ±0.5 nm from a spectral peak for the OSA used in these measurements. This prevents direct observation of the background spectrum, potentially concealing strong spectral features associated with deviations from the Gaussian ASE profile or other lasing modes. Additionally, it could obscure any interaction between the coherent lasing carrier and the ASE background, which would otherwise challenge the assumption that they can be treated separately. These phenomena are especially worrying due to their proximity to the VBG filter passband, as any shifts in this region may persist even after filtering.

To highlight the need for caution here, we calculate the expected light shift from one side of the dynamic-range-limited OSA spectrum (red detuning) observed in Fig. 3.3 inset, rather than using the Gaussian model. For the case of the Ti:S laser, without the VBG filter, we find a shift estimate of 4×10^{-15}. After application of the filter, the remaining shift is still 6×10^{-16}. We emphasize that there is no physical motivation to expect such a completely asymmetric spectrum; nevertheless, it is the most rigorous claim we can make given the limitations of the OSA. Spectral characterizations using the OSA are therefore insufficient to confidently constrain the background light shifts below the 10^{-18} level, even for the spectrally pure case of the Ti:S laser.

We considered constructing several alternate measurement devices to improve the spectral characterization close to the carrier, including a filter Fabry-Perot cavity, an advanced grating spectrometer, or a combination of the two. Ultimately, we decided to make an optical heterodyne measurement between the TA and Ti:S. Compared to the other considerations, this choice was experimentally simple as it could be implemented almost entirely in fiber. The two beams were combined and directed onto a fiber-coupled detector and the beat signal was measured with an RF signal analyzer. In order to limit the contribution of near-DC components, measurements were made in the heterodyne frequency range of 250 MHz to 4 GHz, with the upper end of the range limited by the photodetector bandwidth. The Ti:S frequency was stepped in several GHz intervals to facilitate measurement out to ±100 GHz offset from the TA carrier. During the measurement, the Ti:S was locked to an internal etalon, while the TA was free-running; for each point in the scan, we combined multiple spectrum acquisitions with offset frequency axes to compensate for the effect of frequency noise. When relating an optical heterodyne measurement to the true laser spectrum, care must be taken to account for spurious spectral components corresponding to the complex-conjugate Fourier ambiguity (also known as *image* components). These can be identified by increasing the frequency of the laser under test (TA) relative to the local oscillator (Ti:S) and observing the change in received signal. Real components will appear at higher frequencies, while imaginary components will appear at lower frequencies. We developed a procedure to filter these spurious components during the acquisition. For each large step in the overall detuning scan, the relative frequency was varied by approximately ±100 MHz while several traces were recorded. When these traces were overlaid on a common frequency axis, the "real" signal overlapped, while the image components varied in apparent frequency, allowing them to be identified and statistically rejected. However, this procedure is not able to differentiate between positive and negative heterodyne frequencies for broadband features such as ASE. Therefore, the heterodyne measurement overestimates the ASE power by 6 dB, setting a conservative upper bound rather than a direct measurement.

A number of noise processes potentially contribute to the heterodyne measurement sensitivity. Thermal fluctuations (Johnson-Nyquist noise) yield a noise power $k_B T B$ for a measurement bandwidth B, resulting in a noise floor of -114 dBm/MHz at room temperature. Additionally, our electronic signal analyzer has a noise floor at -95 dBm/MHz. Another significant noise source was photoelectric shot noise. Consider a photodiode with responsivity $\mathcal{R}$, terminating resistance R, and an additional power gain of G. In terms of the optical powers of each beam, P_{TA} and $P_{Ti:S}$, the shot noise is $\delta P_e = 2e\mathcal{R}(P_{TA} + P_{Ti:S})RGB$, while the signal power is $P_e = 2\mathcal{R}^2 P_{TA} P_{Ti:S} RG$. The resulting signal-to-noise ratio is

$$\frac{P_e}{\delta P_e} = \frac{\mathcal{R}}{eB} \frac{P_{TA} P_{Ti:S}}{P_{TA} + P_{Ti:S}} \tag{3.19}$$

Since we measured the electronic power spectral density, rather than the total electronic power, the relevant bandwidth in this calculation is the resolution bandwidth of the signal analyzer, which was chosen to be several orders of magnitude smaller than the photodiode bandwidth. As long as shot-noise-limited performance is achieved in spite of other noise processes in the measurement system, the photodiode bandwidth only determines the spectral span accessible at a single setpoint of laser detuning; therefore, the product of the responsivity and the linearity threshold is the figure of merit for the photodiode. For these measurements, we used an InGaAs photodiode (Discovery Semiconductors xHLPD) with 0.25 A/W responsivity, 4 GHz bandwidth, and a linearity threshold of around 100 mW, achieving a shot noise floor of -94 dBm/MHz.

In Figure 3.7a, we combine OSA and heterodyne data to show the spectrum over eight decades of frequency range, from hundreds of kHz to tens of THz from the magic wavelength. We note that, although the heterodyne measurement is able to constrain background light much better than the OSA (by up to five orders of magnitude for detunings <100 GHz), shot noise still prevents direct observation of the ASE profile in this region as interpolated by the Gaussian fit [1] .

Using equation 3.9, we compute an upper bound on the shift spectral density, and determine the cumulative shift by integrating the shift spectral density outwards from the magic wavelength, shown in Fig. 3.7b. The VBG filter limits the region where the shift spectral density is large enough to cause appreciable shifts, causing convergence of the cumulative shift near the edge of the passband. Our upper bounds on the shift established by this composite spectrum are 6×10^{-15} without the filter and 1.4×10^{-19} filtered. In the filtered case, a notable limitation is the presence of spectral features detuned from the carrier at multiples of 3.5 GHz, corresponding to other

---

[1] Because the heterodyne measurements include ASE spectrum from both positive and negative heterodyne frequencies, essentially doubling the expected measurement of the ASE amplitude, comparison of the observed spectrum in this frequency region should be made relative to the interpolated Gaussian fit plus an additional 6 dB.

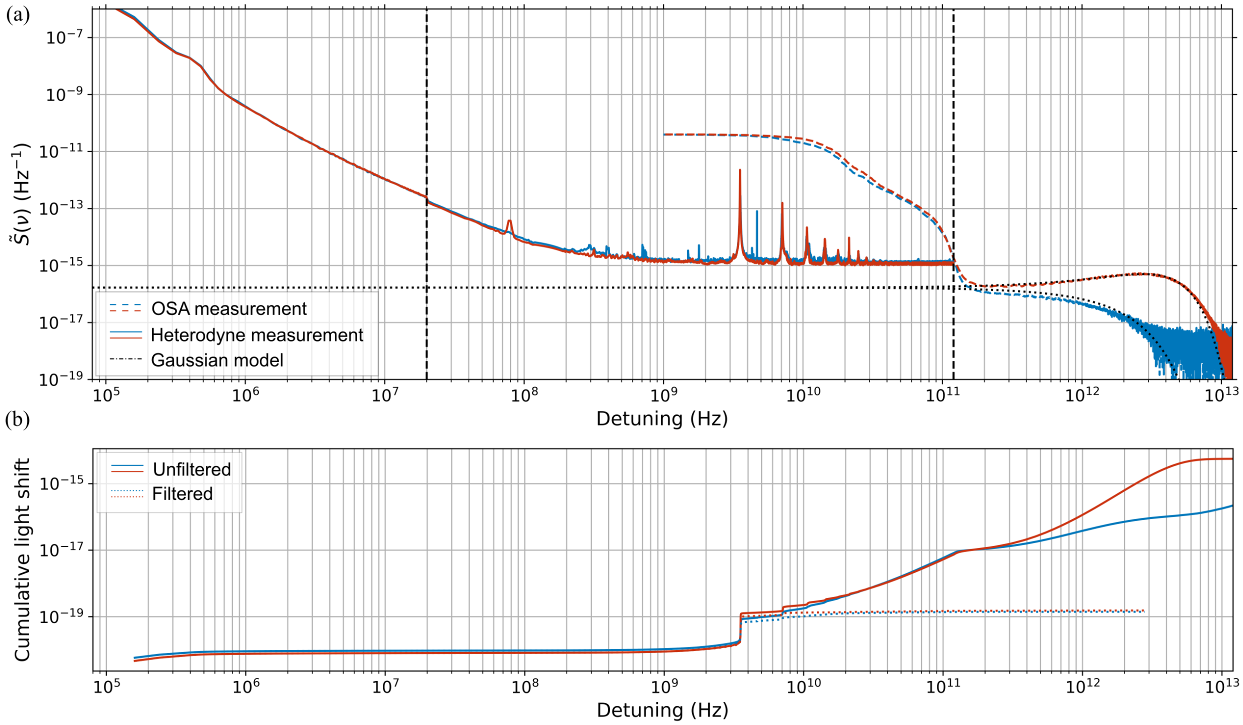

Figure 3.7: a) Composite spectral measurement of the tapered amplifier system, obtained from three different measurements: a ±20 MHz heterodyne measurement with 200 kHz resolution bandwidth (RBW), a ±100 GHz heterodyne measurement with 5 MHz RBW, and a ±10 nm optical spectrum analyzer measurement with 20 pm RBW. Black dashed lines mark the boundaries of the three measurements. The power spectral density normalized to the carrier power is plotted against the absolute value of the detuning from the magic wavelength carrier, with blue (red) showing spectral content at higher (lower) frequencies than the magic wavelength. The dash-dotted line shows a Gaussian fit to the ASE profile measured by the OSA. b) Cumulative shift upper bound, calculated by integrating the shift spectral density from the composite spectrum from the magic wavelength to a variable frequency. The filtered shift spectral density is calculated using the measured VBG transfer function in Fig. 3.6.

longitudinal modes of the seed laser. While there is partial cancellation of red- and blue-detuned features, we use the one-sided estimate as an upper bound; this results in a constraint at around the 10^{-19} level after filtering, which is easily sufficient for our targeted uncertainty budget, but may need to be confronted in future generations of optical lattice clocks. Similar features were observed on the Ti:S spectrum, resulting in one-sided shifts at the 10^{-20} level. The dominance of these features relative to the overall constraints highlights the importance of careful spectral characterization within the filter passband, which is beyond the limits of typical optical spectrum analyzers. Even without a separate laser for heterodyne characterization, monitoring the self-beating near-carrier spectrum on a high-bandwidth photodiode may be sufficient to detect multi-mode operation at the level required for 10^{-19} uncertainty.

3.1.6 Model extrapolation error

The range over which the ASE profile is observable is limited to approximately ±10 THz by the noise floor of the OSA (-80 dBm). The assumption of Gaussian ASE profiles is empirically supported, but extrapolation below the noise floor of the OSA is questionable, as deviations from the assumed Gaussian profile or the observed VBG transfer function could modify the shift, especially if significant spectral content overlaps with an electronic transition. To determine a conservative upper bound on the shift due to spectral content below the OSA noise floor, we assume a constant power spectral density at the noise floor and integrate from the magic wavelength to the edge of our mirror transmission window at 273 THz, again neglecting partial cancellation from the opposite detuning. We additionally assume that the VBG transfer function is constant at -70 dB within the range, neglecting increased attenuation below the noise floor of Fig. (3.6). This yields a constraint of -8×10^{-19} for the unfiltered spectrum and -2×10^{-20} with the filter applied. These constraints could be further tightened by employing a multipass Bragg filter [82] to more strongly attenuate light outside the passband, or by improving the signal-to-noise ratio of the characterizations of the laser spectrum or VBG transfer function.

We also observe that the Gaussian model falls off slightly faster than the observed ASE profile far from the peak (Fig. 3.3). By directly integrating the observed profile, we determine that this contributes an error of $< 5\%$ to the unfiltered shift estimate. This model error therefore contributes negligibly when the bandpass filter is employed.

3.1.7 Grating transfer function

For the purpose of simple numerical evaluation of the shift, the analytical bandpass transfer function (3.14) was used. However, realistic optical bandpass filters can diverge from the theoretical plane-wave transfer function (3.14) due to the effects of finite beams [40], non-monochromatic Gaussian beams with imperfect collimation [19] or potentially the non-Gaussian beam profiles characteristic of tapered amplifiers. We note that the filtered shift has low sensitivity to errors in the determination of the filter bandwidth - doubling the bandwidth would still produce a shift of only -3×10^{-20}. We take this value as an upper bound associated with the grating transfer function. Additionally, the experimentally-characterized reflectivity was used to compute the upper bound in Section 3.1.5 through numerical integration, accounting for deviations from the theoretical profile.

Because the center wavelength of a Bragg filter is tunable through the angle of incidence of light onto the grating, angular misalignment due to improper setup or beam pointing drift results in a wavelength offset of the passband. We estimate the angular bandwidth of our grating to be 200 μrad [40]. By calculating the shift for varying center frequency, we find a dependence of 8×10^{-21} per GHz. Therefore, detuning to the edge of the passband modifies the shift by 4×10^{-20}. We take this as a conservative uncertainty on the shift due to filter misalignment, noting that such a large detuning would result in a large power drop which would quickly be noticed and corrected. Additionally, our delivery module for 759 nm light includes a temperature-stabilized enclosure for the grating to mitigate filter drift due to thermal effects, and an auto-alignment system with MEMS mirrors is used to tune the angle of the incident beam to center the passband around the magic wavelength.

3.1.8 Fiber misalignment

During the unfiltered clock shift measurements in Section 3.1.2, it was observed that small changes in fiber alignment could change the measured shift substantially. To assess the magnitude of this effect, we systematically degraded the fiber alignment using a mirror mount before the fiber delivering light to the atoms as well as the pickoff for OSA spectral characterization. We use the relative decrease in transmitted power to parametrize the misalignment, with an additional sign factor corresponding to the direction of actuation of the mirror mount. For example, a misalignment of -40% indicates that the beam was deflected in the negative direction to reduce the transmitted power by 40% relative to the optimally-aligned case.

OSA spectra measured during this process show a clear trend in distortion and overall amplitude as a function of misalignment. Using these spectra, we compute the background light shift with and without the filter (Figure 3.9). Without filtering, the shift is susceptible to fiber misalignment effects at the level of 3×10^{-16} per percent of transmitted power loss. With filtering, this dependence is suppressed to 4×10^{-23} per percent. Therefore, beam pointing instability contributes negligibly to the uncertainty budget with an assumed worst-case degradation of 10%. However, exceptionally stable beam alignment would be required to constrain time-varying shifts below the 10^{-18} level if a filter were not employed.

3.1.9 Seed-amplifier coupling

The shift could also potentially change with increase in relative ASE level due to a degradation of seed-amplifier coupling. To characterize the sensitivity of this effect, we attenuated the seed laser with varying neutral density filters and measured the resulting optical spectrum with an OSA (Figure 3.8). With 30 dB of seed attenuation, the ASE level is increased only by a factor of approximately 60. Noting that seed-amplifier misalignment would result in a decrease in output power which would be noticed and quickly corrected, we very conservatively assume a maximum seed attenuation of 10 dB, producing a shift of -3×10^{-20}.

3.1.10 Total estimated shift uncertainty

Our uncertainty budget encompasses several constraints based on characterization limitations as well as uncertainty estimates for several time-varying systematic effects. To determine the total uncertainty, we model the former case as uniform distributions truncated at our upper bound and accordingly compute a standard deviation given by $\sigma = a/\sqrt{3}$, where $2a$ is the width of the uniform distribution. The latter effects are assumed to stem from Gaussian distributions with 2σ widths given by our worst-case estimates. We report all uncertainties at the 1σ level in Table 3.1. The total uncertainty on the magnitude of the filtered background light shift, computed as the quadrature

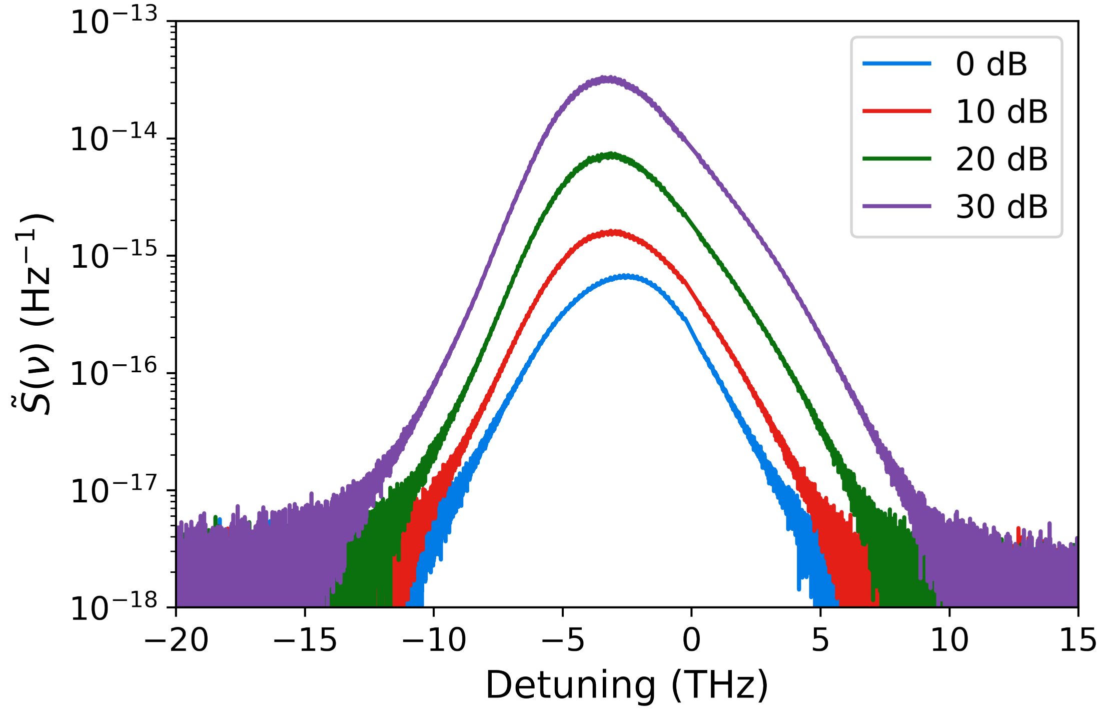

Figure 3.8: ASE spectra at varying seed attenuation relative to a nominal seed power of 35 mW. The spectra are processed to remove the carrier and interpolate the missing ASE spectrum.

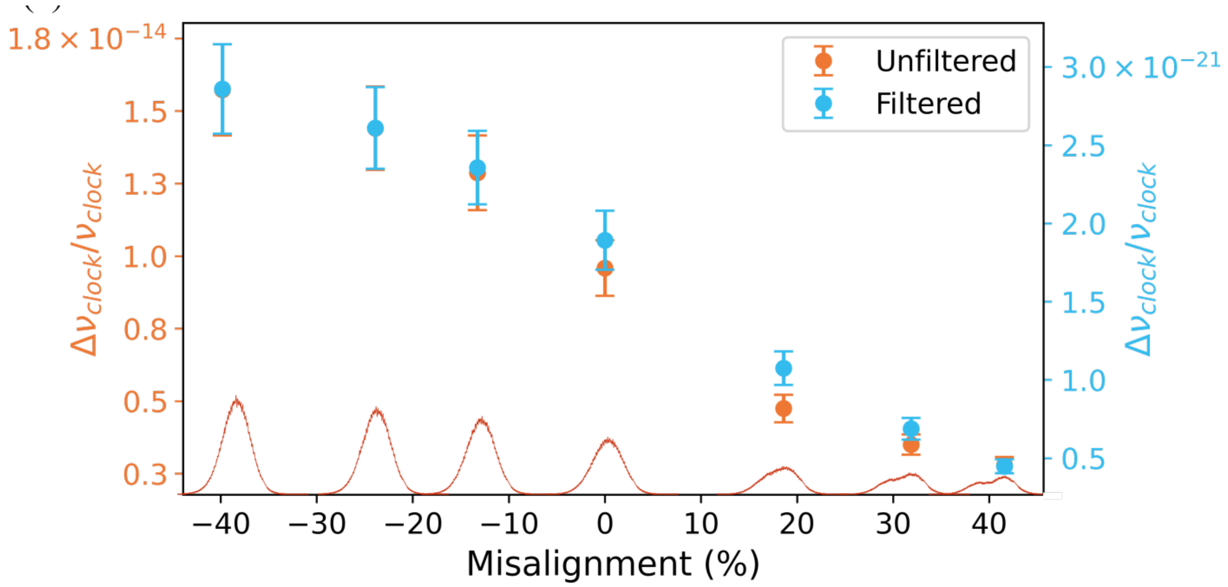

Figure 3.9: Estimates of the associated shifts with (red) and without (orange) spectral filtering as a function of misalignment. ASE spectra corresponding to each point are shown along the bottom.

Table 3.1: Background light shift uncertainty budget

Effect	**Uncertainty** ($\times 10^{-20}$)
Near-carrier spectrum	4
Model extrapolation error	0.6
Grating transfer function	0.8
Grating misalignment	2
Fiber misalignment	0.02
Seed-amplifier coupling	2
Total	5

sum of individual effects, is 5×10^{-20}, most significantly limited by our ability to constrain the amplitude of residual spectral content near the passband. We note that, due to the strong out-of-band rejection of the filter, even highly conservative assumptions on spectral content below the OSA noise floor at large detunings exceeding 10 THz do not contribute substantially.

3.2 Direct experimental validation

To corroborate our upper bounds on the background light shift, we measured the filtered shift using the interleaved atomic clock configuration from Section 3.1.2 with the addition of the volume Bragg grating filtering the TA (Figure 3.10a). This allows us to compare the background light shift from the filtered TA laser system relative to the intrinsically purer filtered Ti:S laser system. At an operational amplifier current of 3210 mA and a temperature of 20 °C, the frequency of the Yb clock transition was measured while switching between the filtered Ti:S and filtered TA for a number of trap depths between 50-120 E_r (Figure 3.10b). Each measurement was terminated at a statistical uncertainty near 1×10^{-17}. As before, the influence of cold collision shifts on this measurement were made negligible by operating at a constant, low atom number density. A weighted linear regression with zero intercept was used to compare data at different trap depths and scale to an operational trap depth of 50 E_r, at which we evaluate the shift to be $(0.1 \pm 1.0) \times 10^{-18}$. The reduced χ^2 of the fit is 1.02, indicating that the scatter of the various measurements are consistent with their stated uncertainties. Thus, under typical operating conditions, we find that the background light shift of the filtered TA system is consistent with zero at the 1×10^{-18} level.

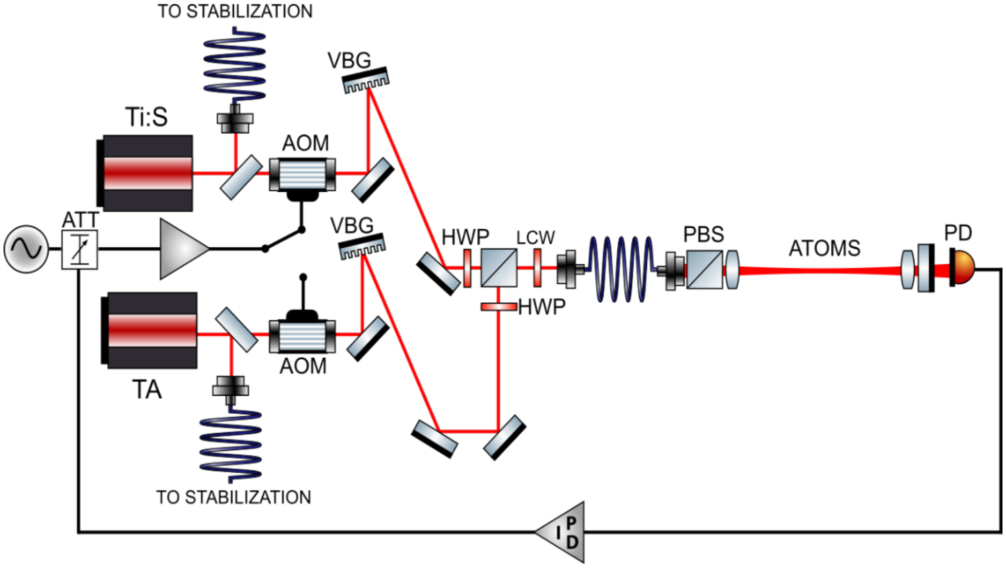

(a) Experimental setup for interleaved clock comparisons. During each cycle, the active beam is selected using a liquid crystal waveplate (LCW) and polarization beamsplitter (PBS). The intensity of the active beam, monitored with a photodetector (PD) behind the lattice retroreflector, is stabilized through the rf drive power of an acousto-optic modulator (AOM) using a variable attenuator (ATT); an rf switch selecting the active AOM further attenuates the inactive beam. Each laser is filtered by a volume Bragg grating (VBG); the Ti:S filter has 26 GHz bandwidth, while the TA filter has 11.5 GHz bandwidth. In the unfiltered measurements in Section 3.1.2, the volume Bragg grating in the TA beam path was replaced with a mirror.

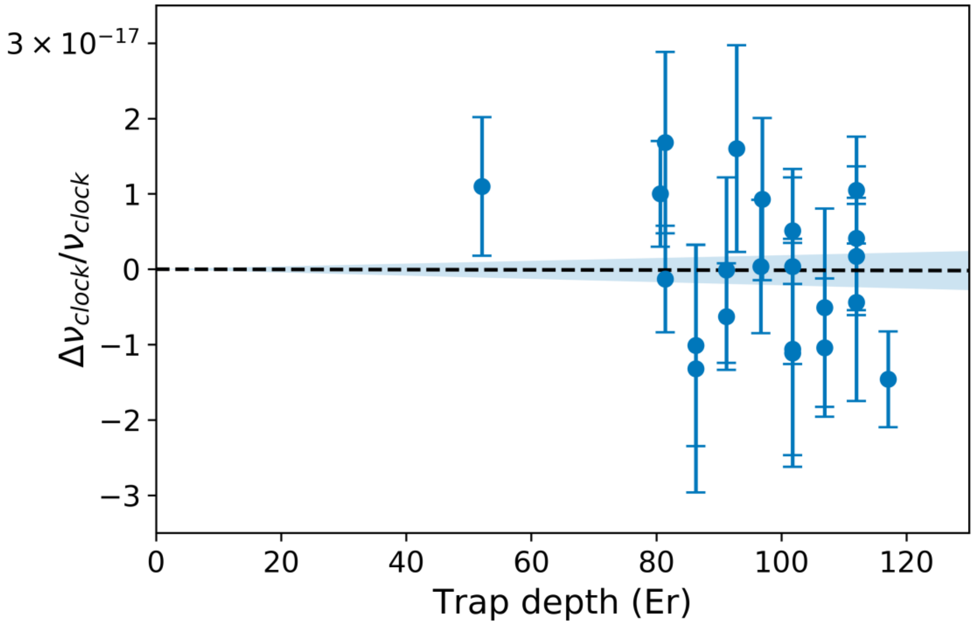

(b) Fractional clock shifts between the filtered TA and reference Ti:S configurations. The blue band shows the 1σ confidence region on the linear fit.

Figure 3.10: Background light shift experimental validation

3.3 Blackbody radiation shift

As opposed to shifts which can be tuned to zero with some uncertainty, like the lattice Stark shift in the previous section, the blackbody radiation shift exceeds 10^{-15} at room temperature. While the T^4 scaling can be exploited to suppress the shift to negligible levels at cryogenic temperatures [102], the added size, weight, and power is infeasible for transportable clocks, and therefore this shift must instead be characterized at better than the part-per-thousand level to achieve 10^{-18} uncertainty.

The shift can be written [84, 6]

$$\Delta\nu = -\frac{1}{2}\frac{\Delta\alpha(0)}{h}\langle E^2\rangle_T[1+\eta_{\mathrm{clock}}(T)] \tag{3.20}$$

where

$$\Delta\alpha(0) = h \times 3.62612(7) \times 10^{-6}\ \mathrm{Hz/(V/m)}^2 \tag{3.21}$$

is the differential static polarizability between the two clock states [94],

$$\langle E^2\rangle_T = [8.319430(15)\ \mathrm{V/cm}]^2(T/300\ \mathrm{K})^4 \tag{3.22}$$

is the mean-squared electric field of the thermal radiation, and

$$\eta_{\mathrm{clock}} \approx 0.01745(38)(T/300\ \mathrm{K})^2 + 0.000593(16)(T/300\ \mathrm{K})^4 \tag{3.23}$$

is the *dynamic correction* for the frequency dependence of the polarizability [6]. The dynamic correction accounts for approximately 2% of the total shift; thus, most of the meaningful variation of shift with temperature is contained in equation (3.22).

From equations (3.20-3.23), the fractional variation in the clock shift from a temperature variation δT around 300 K is

$$\frac{\delta\nu}{\delta T} = -17.1894(4)\ \mathrm{mHz/K} \tag{3.24}$$

Therefore, to achieve 10^{-18} uncertainty of the BBR shift, the effective temperature of the blackbody radiation bathing the atoms must be known to around 30 mK uncertainty.

Determination of the shift can be formulated as a boundary-value problem, wherein temperature measurements of the surfaces surrounding the atoms can be used to compute the effective temperature of radiation experienced by the atoms [11, 6, 14]. Typically, the temperature is logged during clock operation and the BBR shift is computed and subtracted in post-processing; however, in a live timescale, the shift could be subtracted in real-time. As we will see in the following section, the uncertainty can be constrained by surrounding the atoms with a blackbody surface with as high emissivity as possible. In our laboratory clocks, this is implemented through an in-vacuum copper shield with embedded platinum resistance temperature detectors (RTDs) [6]. The interior of the shield is coated with multiwall carbon nanotubes [58] to achieve a high emissivity $\varepsilon > 0.8$ from visible wavelengths to 20 um, spanning the dominant transitions which contribute to the BBR shift. The high thermal conductivity of the shield limits the magnitude of thermal gradients which can arise, simplifying the modeling of radiative exchange with the atoms. For the transportable clock, the aluminum vacuum chamber itself serves this role; in the following sections, I present finite-element modeling results which support a blackbody radiation uncertainty below the 10^{-18} level with this shield.

3.3.1 Effective temperature model

The effective temperature is given by an integral over the enclosure, or alternately as a summation over discrete surface elements,

$$\langle T^n \rangle = \frac{1}{4\pi} \oint T^n \left(\frac{d\Omega^{\mathrm{eff}}}{dA} \right) dA \to \frac{1}{4\pi} \sum_i T_i^n \left(\frac{\Omega^{\mathrm{eff}}_{\mathrm{i}}}{A_i} \right) A_i \tag{3.25}$$

The quantity Ω_i^{eff} is the *effective solid angle* of the surface element A_i, which is approximately equal to the product of the geometric solid angle and the emissivity (and therefore reduces to the geometric solid angle for a blackbody enclosure). This term serves as a weighting factor which determines how much each surface contributes to the overall effective temperature. The effective solid angles must obey the normalization condition

$$\sum_i \Omega_i = \sum_i \Omega_i^{\mathrm{eff}} = 4\pi, \tag{3.26}$$

which intuitively implies that the effective solid angle of a surface decreases as its emissivity decreases relative to the emissivity of other surfaces.

As a simple example, consider a sphere with one hemisphere a perfect blackbody ($\epsilon = 1$) and the other with $\epsilon = 0$. It can be shown (equation (3.29)) that the blackbody hemisphere contributes an entire 4π to the effective solid angle, since the other hemisphere does not radiate but simply reflects all incident radiation. Therefore, the certainty of the effective temperature can be improved by maximizing the emissivity of surfaces whose temperature can be measured accurately.

In practice, the effective solid angles of the different surface elements surrounding the atoms can be determined through a finite-element model of the vacuum chamber. A probe is placed at the center of the chamber and the probe temperature is monitored while varying the temperature of each surface. Varying the temperature of the i^{th} surface by $\Delta(T_i^4)$ causes the probe temperature to change by

$$\Delta(T^4) = \frac{1}{4\pi}\Delta(T_i^4)\Omega_i^{\text{eff}} \tag{3.27}$$

and therefore the geometric solid angle is determined to be

$$\frac{\Omega_i^{\text{eff}}}{4\pi} = \frac{\Delta(T^4)}{\Delta(T_i^4)} \tag{3.28}$$

This modeling was carried out using COMSOL Multiphysics [23] to determine the effective solid angle of each surface element surrounding the atoms. Prior to modeling the real chamber geometry, the model was validated on a simple toy model with an analytic solution: a spherical enclosure with emissivity ε_s fully surrounding the atoms except for an open aperture with solid angle Ω_A which exposes the atoms to the ambient temperature T_A. The analytical solution for the effective solid angle of the aperture is

$$\frac{\Omega_A^{\text{eff}}}{4\pi} = \frac{\Omega_A}{\Omega_A + (4\pi - \Omega_A)\varepsilon_s} \tag{3.29}$$

This simple model has two limiting cases which illuminate our strategy for controlling BBR uncertainties. In the limit of small chamber emissivity, $\varepsilon_s \to 0$, $\Omega_A^{\text{eff}} \to 4\pi$. In this limit, the chamber reflects radiation which enters through the aperture with perfect efficiency, so the atoms are maximally exposed to the ambient temperature. The other limit is a blackbody enclosure, $\varepsilon_s \to 1$, in which case $\Omega_A^{\text{eff}} \to \Omega_A$. Thus, we can understand the effective solid angle of a surface as its propensity to absorb or reflect the ambient radiation. To minimize the blackbody radiation uncertainty, our goal will be to surround the atoms to the greatest extent possible with perfect blackbody surfaces whose temperature can be measured precisely, and whose high emissivity prevents reflection of less-known ambient temperatures.

In practice, this entails minimizing the size of the entry and exit apertures for the atomic beam, whose effective temperature can be difficult to characterize. The clear aperture of the windows required for laser access should also be minimized, as their relatively low thermal conductivity complicates measurement of the temperature over the entire surface and allows thermal gradients due to laser heating. The interior chamber surface, which accounts for over 80% of the geometric solid angle, is coated with an optically black paint (AZ Technology MLS-SB-85-C) which has a reflectance at or below 4% between 250 nm and 2.5 um.

After model validation, the real science chamber was meshed with approximately 100,000 domain elements, with finer meshing on small regions such as the entry and exit apertures to guarantee numerical convergence (Figure 3.11). The effective solid angle of the windows and entry and exit apertures was calculated across a range of chamber and window emissivities (Figures 3.12, 3.13, and 3.14). To improve the accuracy of the results, the curves were scaled to match the analytic solid angles for the case with unit emissivities for all components, in which case the effective solid angles reduce to the geometric solid angles. These results have several important implications. First, similar to the toy model, the effective solid angles decrease with the emissivity of the chamber interior - thus, the black coating limits the uncertainty contributions of the windows and apertures. Second, the effective solid angles become independent of the window emissivity as the chamber emissivity tends towards 1. Therefore, careful reflectance measurements of the windows are not required. To compute the environmental BBR uncertainty associated with this geometry, we conservatively assume $\varepsilon_{chamber} > 0.8$ and $\varepsilon_{window} > 0.2$.

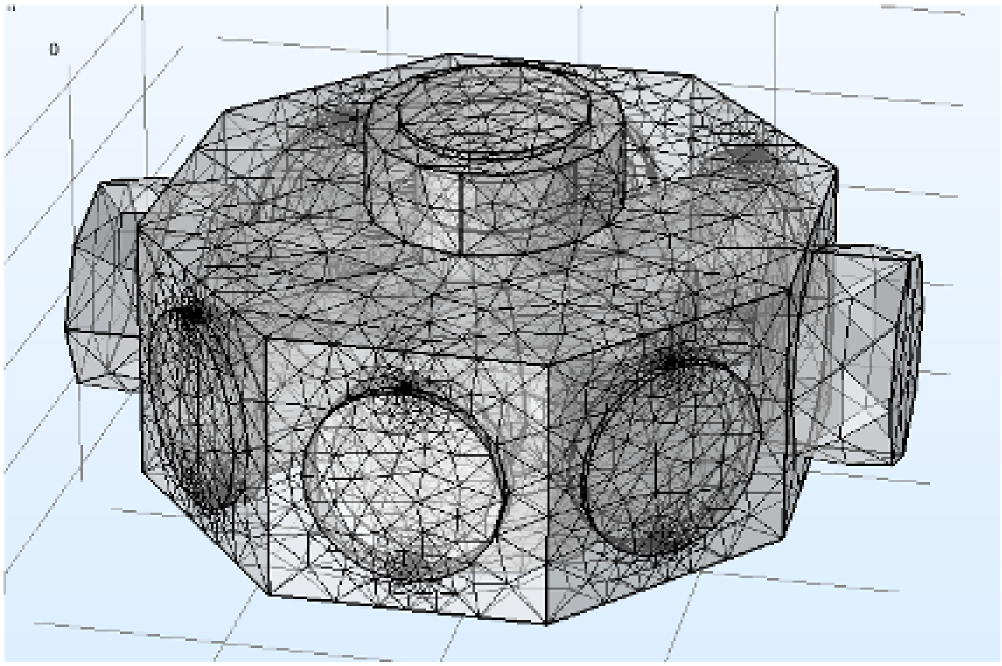

Figure 3.11: Meshed science chamber model with approximately 100,000 domain mesh elements used for calculation of effective solid angles.

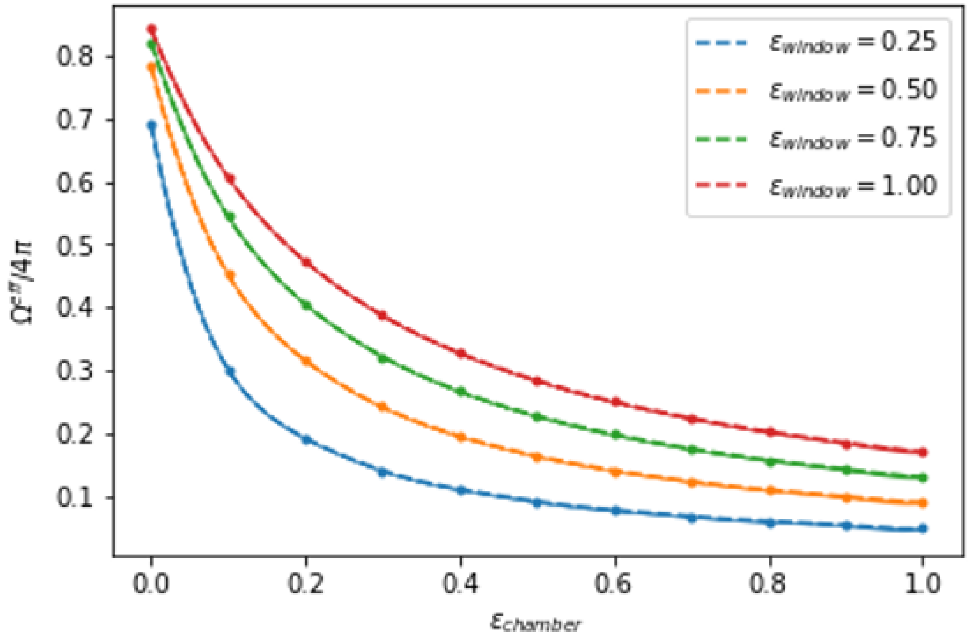

Figure 3.12: Calculated total effective solid angle for the eight chamber windows as a function of window and chamber emissivity. The shaded regions are 1σ uncertainty bounds from a Gaussian process fit to all four curves.

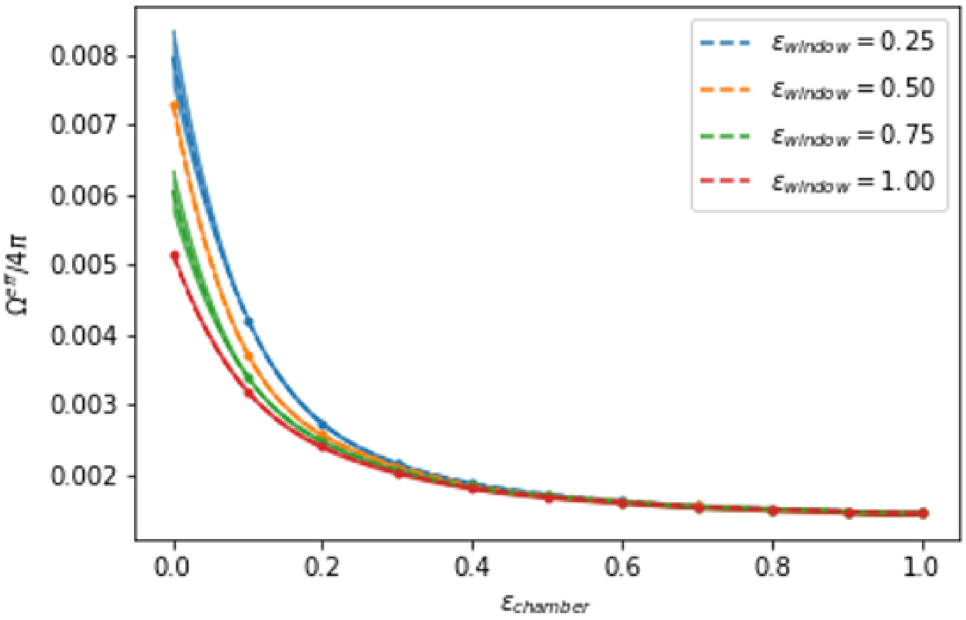

Figure 3.13: Calculated effective solid angle for the atomic entry aperture as a function of window and chamber emissivity. The shaded regions are 1σ uncertainty bounds from a Gaussian process fit to all four curves.

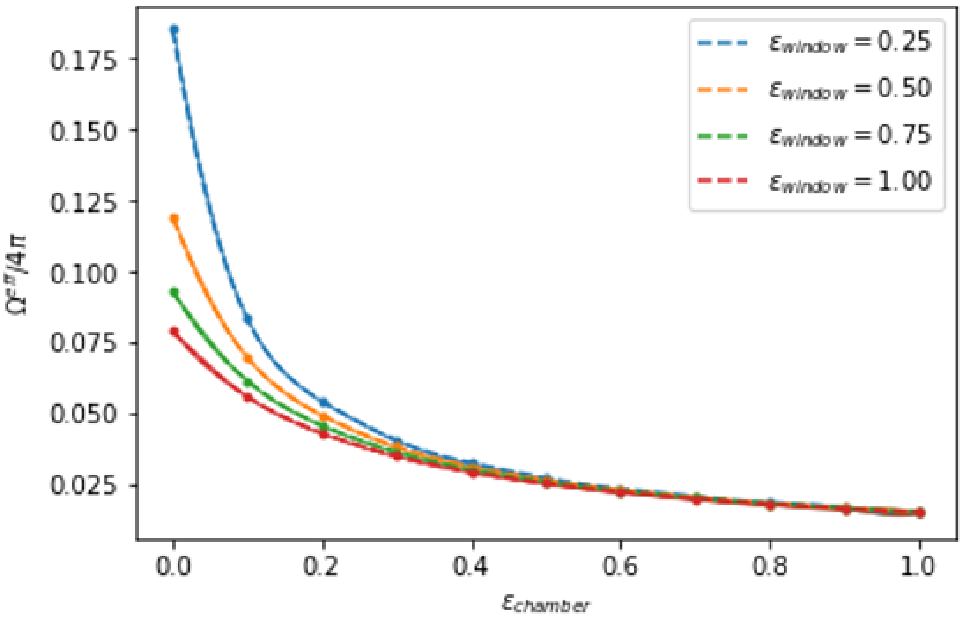

Figure 3.14: Calculated effective solid angle for the atomic exit aperture as a function of window and chamber emissivity. The shaded regions are 1σ uncertainty bounds from a Gaussian process fit to all four curves.

3.3.2 Propagation of uncertainty

Once nominal values and uncertainties are known for the solid angle and temperature of each surface element, the uncertainty of the effective temperature can be estimated with standard propagation-of-error techniques:

$$\delta T^4 = \sum_i \left[T_i^4 \delta f_i + f_i \delta T_i^4 + \mathcal{O}(\delta f_i \delta T_i^4) \right] \tag{3.30}$$

where we have used the shorthand $T^4 = \langle T^4 \rangle$ and $f_i = \Omega_i^{eff}/4\pi$ for neatness. The effective solid angle normalization constraint, $\sum_i f_i = 1$, also results in a second constraint

$$\sum_i \delta f_i = 0 \tag{3.31}$$

To enforce this constraint, we isolate element j in the sum and rewrite (3.30) as

$$\delta T^4 = \sum_{i \neq j} \left[T_i^4 \delta f_i + f_i \delta T_i^4 + \mathcal{O}(\delta f_i \delta T_i^4) \right] + T_j^4 \delta f_j + f_j \delta T_j^4 \tag{3.32}$$

Since the constraints enforce the relationships

$$f_j = 1 - \sum_{i \neq j} f_i \tag{3.33}$$

and

$$\delta f_j = -\sum_{i \neq j} \delta f_i, \tag{3.34}$$

we can further rewrite (3.32) as

$$\delta T^4 = \sum_{i \neq j} \left[T_i^4 \delta f_i + f_i \delta T_i^4 + \mathcal{O}(\delta f_i \delta T_i^4) \right] - T_j^4 \sum_{i \neq j} \delta f_i + (1 - \sum_{i \neq j} f_i) \delta T_j^4 \tag{3.35}$$

$$= \sum_{i \neq j} \left[(T_i^4 - T_j^4) \delta f_i + f_i (\delta T_i^4 - \delta T_j^4) + \mathcal{O}(\delta f_i \delta T_i^4) \right] + \delta T_j^4 \tag{3.36}$$

The first useful result is that the contribution from the uncertainty in effective solid angle δf_i vanishes for nominally equal temperatures, and therefore the effective solid angle needs to be known accurately only for elements with a large temperature difference from the chamber (e.g. the atomic shutter). Another useful result is that the contribution due to temperature uncertainty

δT_i^n is multiplied by the effective solid angle f_i, so careful temperature measurements need only be made for surfaces with large effective solid angles.

Despite the instructiveness of (3.35), substantial uncertainty is included in the higher-order differentials. For example, a contribution due to solid angle uncertainty which vanishes with nominally equal temperatures could be large if the temperature difference is nonzero. To properly account for all higher-order terms, the propagation of uncertainty is instead carried out using a Monte Carlo simulation, applying the following procedure for each of $N >> 1$ samples.

(1) Draw surface element temperatures from normal distributions characterizing the mean temperature and its uncertainty.

(2) Generate effective solid angles satisfying the normalization constraint, drawn uniformly from bounds established with the FE model.

(3) Compute the effective temperature and resulting shift.

For the first step, we distinguish between *measurement* and *operational* temperature uncertainties. The former quantifies the accuracy with which we can measure the temperature of each surface element, while the latter quantifies anticipated variation of the surface temperature during day-to-day operation. For example, the temperature of the chamber may be measurable to within 5 mK (the stated uncertainty of the resistance temperature detectors employed), but could vary by several K if the clock is operating in uncontrolled environmental conditions. For this reason, it is necessary to monitor the temperature during clock operation so that the time variation of the shift can be properly determined and compensated.

The results of the process are shown in Table 3.2. The total environmental uncertainty from the BBR shift is 7.2×10^{-19}.

Table 3.2: Environmental uncertainty budget for blackbody radiation shifts computed by Monte Carlo analysis. Ranges in effective solid angle are computed by assuming a chamber emissivity $\epsilon_c > 0.8$ and a window emissivity $\epsilon_w > 0.2$. The effective solid angle of the chamber is allowed to vary to satisfy the normalization constraint (3.26).

Element	Temperature (K)	Ω_{eff}	Environmental uncertainty ($\times 10^{-19}$)
Entry aperture	303(10)	0.00147-0.00153	5.0
Exit aperture	293(5)	0.013-0.016	2.3
Windows	293.0(1)	0.05-0.25	4.5
Chamber	293.000(5)	Variable	1.3

3.4 DC Stark shift

While lattice ac Stark shifts originate from a conceptually simple mechanism, the perturbation of the clock transition by a well-characterized lattice laser, they share the spotlight with a more subtle electromagnetic effect: dc Stark shifts due to stray charges on the vacuum chamber or elsewhere in the clock apparatus. Strategies to constrain this shift parallel our treatment of the BBR shift, essentially reducing this challenge to careful analysis of a boundary-value problem. Similar to our reliance on a thermally-conductive aluminum vacuum chamber to minimize temperature gradients, the low electrical resistivity of the black coating (10^5 ohms/sq.) reduces the extent to which patch charges can build up and surrounds the atoms with a Faraday shield to screen external fields from electronics or other sources. Furthermore, our analysis of effective solid angles through simulated perturbation of surface temperatures finds a direct experimental analogue here: by applying kV-level potentials to different combinations of vacuum windows, the size of the stray electric field can be constrained. We previously implemented this technique to constrain dc Stark shifts at the 10^{-20} level in our stationary clocks [7]; here I will describe the model from that publication before discussing its extension to the transportable system.

The six window electrodes are set to potentials V_i, resulting in a total electric field $\mathbf{E}_0 + \sum_i \mathbf{E}_i$, where $\mathbf{E}_0$ is the stray field and $\mathbf{E}_i$ is the field resulting from the potential V_i. This produces a clock shift

$$\delta\nu(\{V_i\}) = \delta\nu_0 + \sum_i a_i V_i + \sum_{ij} b_{ij} V_i V_j, \tag{3.37}$$

where $\delta\nu_0 = k\langle E_0^2\rangle$ is the shift arising from the stray field, $a_iV_i = 2k\langle\mathbf{E}_0\cdot\mathbf{E}_i\rangle$ accounts for combined effects of the stray and applied fields, and $b_{ij}V_iV_j = k\langle\mathbf{E}_i\cdot\mathbf{E}_j\rangle$ accounts for the applied fields alone. We restrict the voltage combinations such that pairs of opposite windows have either zero or equal and opposite potentials, which causes off-diagonal terms in b_{ij} to vanish. We measure shifts for different combinations of potentials, interleaved against a reference configuration with all windows grounded; thus, we are directly measuring the induced shift $\delta\nu(\{V_i\}) - \delta\nu_0$. A fit to equation (3.37) is used to extract the coefficients a_i and b_{ii} as well as the stray-field shift $\delta\nu_0$. Finally, a Monte Carlo analysis is used to derive a probability distribution for the stray-field shift.

While we have not yet carried out the experimental work to characterize dc Stark shifts in our transportable clock, it was necessary to anticipate this work during the design stage of the science chamber. The design strategies for characterizing BBR and dc Stark shifts are somewhat at odds with one another: reduction of window solid angles as viewed by the atoms also reduces the dc Stark shift for an applied potential. To ensure that sufficiently large shifts could be generated with realistic potentials, we created a finite-element model of the chamber geometry, simulated the electric fields at the center of the chamber vs. applied voltage, and computed the corresponding shift (Figure 3.15). In comparison with the chamber geometry of our stationary clocks, the shifts are over twice as large in the horizontal direction and about 40% larger in the vertical direction, guaranteeing a sufficiently large lever arm for precise constraint of the stray dc Stark shift.

To facilitate application of large voltages, the interiors of the windows are coated with electrically-conductive indium tin oxide (ITO) to serve as electrodes. Silver traces painted on the sides of the windows serve as electrical feedthroughs passing through the epoxy layer. The edge of the window is cut to a flat plane around the trace to avoid arcing between the trace and the vacuum chamber, and the interior window surface is insulated from the chamber with a silicone ring. Using this geometry, the application of kV-level voltages to measure or constrain the stray field is safe and straightforward.

However, we have also explored a welded titanium vacuum chamber (Chapter 5.2) for which no electrical feedthrough has been implemented. Instead, the ITO coating extends from the surface of the windows to their titanium weld rings, ensuring electrical continuity of the entire chamber interior. Lacking the ability to directly measure stray fields in this situation, we can instead attempt to constrain their magnitude through finite-element modeling. While the conductivity inside the science chamber is sufficiently high to ensure charge dissipation on timescales much shorter than the experimental cycle, persistent charge could accumulate on insulators further from the atoms, such as the slowing mirror. To estimate the maximum shift which could arise, we added a long tube to the finite-element model of the chamber and defined a circular insulating area at the end. The length and diameter of the tube are chosen such that the solid angle visible to the atoms is similar to the slowing mirror. With the rest of the chamber grounded, a large voltage is applied to the insulator. As the breakdown voltage of glass is approximately 10 MV/m, the maximum voltage which a 6 mm thick mirror can support before dielectric breakdown is about 60 kV. We find that shifts below 10^{-20} are produced even for applied voltages of 1 MV. While other groups have struggled to constrain dc Stark shifts due to in-vacuum optics [50], the geometry described here thus permits a high degree of confidence in shifts below 10^{-18} even without *in situ* measurement.

3.5 Background gas collisions

Collisions between lattice-trapped atoms and background gases in the vacuum chamber drives both trap loss for collisions with low impact factor and frequency shifts for more glancing collisions. This could be especially relevant to the transportable system due to the relatively high pressure level ($\approx$ 20 nTorr) of the current vacuum chamber. By measuring frequency shifts while intentionally provoking outgassing by heating a non-evaporable getter on Yb1, we determined that the collisional shift scales linearly with the trap loss rate with a coefficient of $-1.64(1) \times 10^{-17}$ s^{-1} [69]. In the transportable clock, the lattice $1/e$ lifetime was measured to be 200(24) ms, resulting in a shift of $(-8.2 \pm 1.2) \times 10^{-17}$. This shift will be reduced with an impending vacuum chamber upgrade. We note that the transportable system has been designed for rapid, plug-and-play interchangeability of vacuum systems, enabling much faster iterative development than conventional laboratory apparatuses.

Figure 3.15: Finite-element model for dc Stark evaluation. a) Meshed module of the chamber interior, formed by inverting a CAD model of the actual chamber. b) Fractional clock shift vs. voltage applied to horizontal electrodes. c) Fractional clock shift vs. voltage applied to vertical electrodes. In b) and c), triangular points represent direct clock shift measurements from [7].

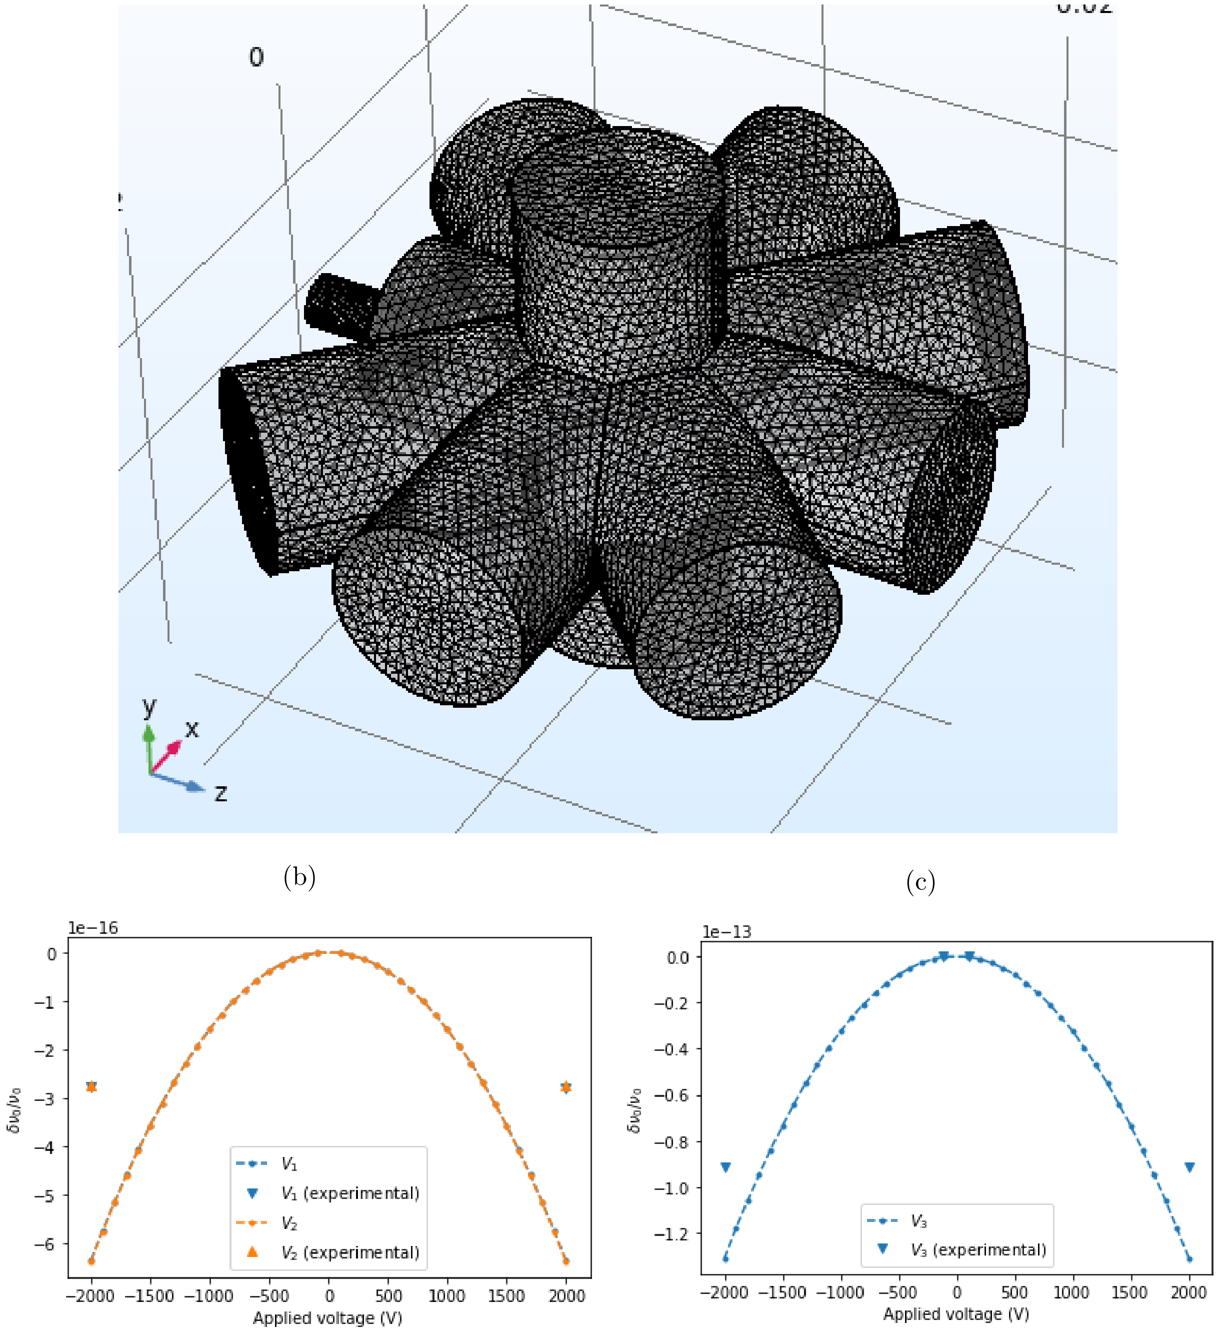
(a)

(b)

(c)

Chapter 4

Clock cavity

4.1 Background

Atomic clocks rely on high-finesse optical cavities to stabilize the frequency of the interrogating laser during and between spectroscopy cycles. The performance of the optical cavity (and of the clock laser) determines the ultimate stability of the clock in two ways. First, the maximum spectroscopy time, and therefore the quality factor Q with which the clock transition can be resolved, is limited by phase fluctuations of the interrogating laser. Second, laser noise can be aliased onto the clock frequency due to a phenomenon called the Dick effect [86]. Current state-of-the-art optical cavities achieve a noise floor at or below the 10^{-16} level, requiring careful design to mitigate various systematic effects. Cavities for transportable systems require even more careful attention to remain robust in spite of large vibrations or temperature changes. In addition, the requirement of rigid mounting for transportable systems imposes a new set of design challenges, as care must be taken not to deform the cavity through the mounting forces. In this chapter, I will first discuss potential sources of instability: Brownian noise, temperature fluctuations, acceleration sensitivity, and force sensitivity. Next, I will present a novel design for a rigidly-mounted cavity which simultaneously achieves a low thermal noise floor, low acceleration sensitivity, and low sensitivity to the holding forces.

Table 4.1: Typical parameters used to determine the Brownian noise instability (equation (4.1)) for our cavity. For these parameters, the thermal noise floor is 1.2×10^{-16}.

Parameter	Description	Value
T	Mirror temperature	303 K
σ	Poisson's ratio of mirror substrate	0.17
E	Young's modulus of mirror substrate	72.5 GPa
$\phi_{\text{substrate}}$	Mechanical loss angle of mirror substrate	1×10^{-6}
w_0	Laser $1/e$ field radius	513 um
L	Cavity length	10 cm
ϕ_{coating}	Mechanical loss angle of mirror coating	2.5×10^{-5}
d	Coating thickness	10 um

4.1.1 Thermal instability

In room-temperature cavities, performance is often limited by Brownian noise in the mirrors, contributing an instability

$$\sigma_{\text{therm}} = \sqrt{\ln 2 \frac{8k_B T}{\pi^{3/2}} \frac{1-\sigma^2}{E w_0 L^2} \left(\phi_{\text{substrate}} + \phi_{\text{coating}} \frac{2}{\sqrt{\pi}} \frac{1-2\sigma}{1-\sigma} \frac{d}{w_0} \right)} \tag{4.1}$$

The many parameters entering into this instability are described along with representative values for our cavity in Table 4.1. Benefiting from low-loss crystalline mirrors [21] with 10 m radius of curvature, our cavity targets a thermal noise floor of 1.2×10^{-16}.

Temperature fluctuations can also introduce instability over longer timescales through thermal expansion of the spacer. It is common to use ultra-low-expansion (ULE) glass, whose coefficient of thermal expansion (CTE) can be tuned to zero through the addition of titanium dioxide. The CTE zero-crossing is often chosen to be above room temperature so that the temperature can be maintained with heating only. For our cavity, we use ULE with a CTE zero-crossing at 40 ^{o}C, high enough to operate in warmer temperatures outside the lab. The nonzero CTE of the fused silica mirror substrates shifts the overall CTE of the spacer away from the ULE zero-crossing, which would reduce the zero-crossing temperature to 20 ^{o}C. To reduce this shift, we use ULE backing rings contacted to the backside of the mirrors [57], reducing the CTE shift to only -5 ^{o}C. The backing rings are 6.35 mm thick with a 25.4 mm outer diameter and 9 mm center bore.

4.1.2 Acceleration and force sensitivity

Another source of instability is the acceleration-induced deformation due to seismic noise or other sources of vibration [18]. Acceleration sensitivity can be mitigated by exploiting symmetries in the spacer design and/or by choosing the support points such that any deformation of the cavity results in equal deflection of the two mirrors and, consequently, no length change along the optical axis of the cavity[78].

Acceleration sensitivity is particularly challenging for transportable systems for two reasons. First, vibrations are expected to be much higher outside a well-controlled lab environment. Second, rigid constraints are required to prevent misalignment of the cavity during transportation. Care must be taken not to deform the cavity through over-determined constraints [39], as this could result in cavity length changes which are dependent on changes in the total constraint force due to acceleration or thermal expansion.

Several designs have been developed for rigidly-supported, force-insensitive cavities. A cavity held by stainless-steel flexure springs has demonstrated an acceleration sensitivity of $2(1)\times10^{-11}/g$, further suppressed below the $10^{-12}/g$ level with real-time feedforward corrections [59]. A similar concept was employed in the development of a transportable cavity constrained by flexible stainless steel stranded wires, achieving acceleration sensitivity from $0.7 - 12.3 \times 10^{-10}/g$ along the three Cartesian axes [39]. Both of these designs exploit a tradeoff in constraint force, choosing constraints which are strong enough to maintain optical alignment under acceleration, while weak enough to effectively decouple the cavity from the mounting structure.

An alternate approach was pioneered with a spherical cavity rigidly held at a “squeeze-insensitive angle” which grants an innate force insensitivity: changes in the applied support force produce nominally zero change in cavity length [60]. The spherical symmetry granted an acceleration sensitivity from $4.0(5) \times 10^{-11}/g$ to $3.1(1) \times 10^{-10}/g$ along the thre Cartesian axes. This philosophy was extended to a 5 cm cubic spacer by NPL [108], further analyzed in [110], and scaled up to 8.7 cm side length to achieve a lower thermal noise floor for a future space-borne cavity [92].

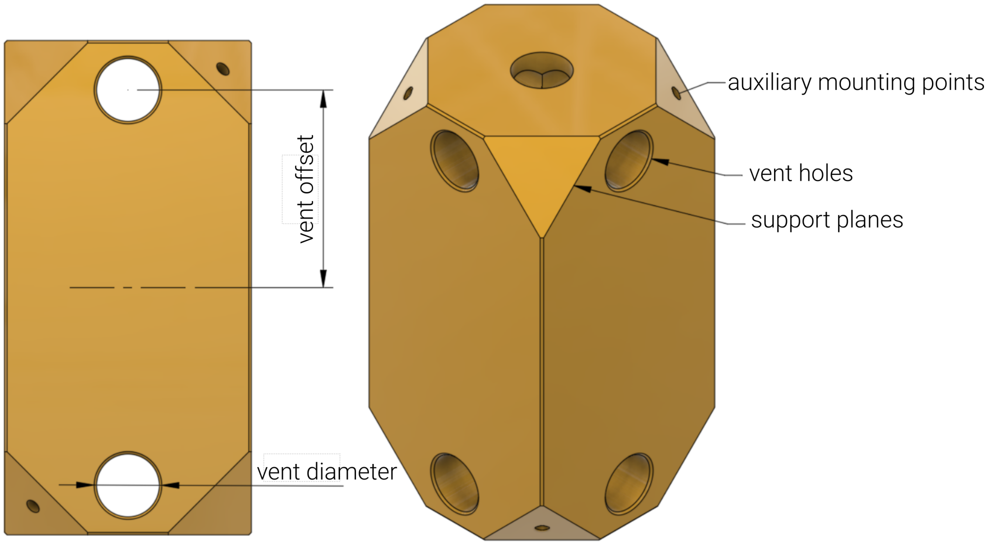

Figure 4.1: Stretched cubic cavity design.

In this design, the force sensitivity can be tuned to zero through the depth of the vertex cuts where the constraint forces are applied. An acceleration sensitivity of $2.5 \times 10^{-11}/g$ was demonstrated, thought to arise from mounting error relative to the nominal support points.

The inherent simplicity of force-insensitive cavities is appealing, as they allow higher mounting forces to be applied to stabilize the cavity against misalignment with a simple support geometry. However, a drawback of the spherical and cubic designs is the cubic scaling of the spacer volume and mass with the cavity length. Consequently, reaching a desired Brownian noise floor may require a prohibitively large spacer, which is especially inconvenient for transportable applications. We have developed an alternate design, where the cubic cavity is stretched along only the cavity axis, reducing the mass and volume by a factor of 4 relative to a cubic spacer with the same length. This results in a large force sensitivity which cannot be nulled by tuning the vertex depth alone. To account for this, we replace the centered venting holes with pairs of holes offset from the center; the offset and diameter of these holes offer two new parameters which can be tuned to zero the force sensitivity. In Section 4.2, we present a finite element analysis (FEA) model of the cavity spacer, demonstrating the ability to tune the force sensitivity to zero. In Section 4.3, we extend the FEA model to include deformation of the supports, investigating the physical validity of the model in Section 4.2. In Sections 4.4 and 4.5, we assess the ramifications of residual force sensitivity and acceleration sensitivity respectively due to fabrication errors. Through a careful and comprehensive analysis of the effects of various fabrication errors, we estimate the residual force sensitivity and acceleration sensitivity to be below $4 \times 10^{-11}/N$ and $5 \times 10^{-11}/g$ respectively. Additionally, the combination of a long spacer and low-loss crystalline mirrors allows thermal-noise-limited instability in the low 10^{-16} decade.

4.2 Cavity design and modeling

We assess the performance of the stretched cubic design using finite-element modeling with COMSOL Multiphysics [23]. The support forces and constraints are applied to circular boundaries on four of the vertex cuts. The diameter of these boundaries is determined through modeling of the elastic contact deformation of hemispherical supports and is typically around 1 mm (Section 4.4).

In the cubic design in the absence of acceleration, four equal support forces are sufficient to ensure static equilibrium. In the stretched design, the support forces do not point through the spacer center-of-mass, and therefore each exerts a nonzero torque. However, if the forces are balanced, these torques sum to zero by symmetry, resulting in static translational and rotational equilibrium. Nevertheless, in both designs, additional constraints are needed in the simulations to suppress rigid motion of the spacer [106]. To satisfy this requirement, we constrain the displacement of the support boundaries to be zero in directions perpendicular to the support force (parallel to the vertex plane), an assumption previously made in [110]. In Section 4.3, these constraints are relaxed, and the effects of support geometry and material properties on the force and acceleration sensitivity are explored.

Using this model, we calculate the on-axis length change of the cavity as a function of vertex depth, vent offset, and vent diameter for different lengths. We define the *force sensitivity* as the fractional change in cavity length per unit of force applied at each of the support points. The dimensions of the cavity in the two face directions orthogonal to the axis are fixed at 5 cm, and each support applies 50 N of force. Starting from the force-insensitive NPL design, stretching the cavity to an aspect ratio greater than 1 shifts the force sensitivity in the negative direction (i.e. the applied support forces will push the mirrors towards each other). A negative deformation can also be produced though the addition of mass near the ends of the spacer due to mirrors or ULE backing rings. These effects can be compensated by removing mass near the ends of the cavity by increasing the vertex cut depth, vent diameter, or vent offset. Except for aspect ratios very close to 1, tuning of the vertex depth alone is insufficient to achieve force insensitivity, requiring the split-bore design.

When the cavity is accelerated or subjected to a gravitational body load, we observe zero on-axis length change, though breaking the symmetry of the cavity by perturbing various features can produce a nonzero acceleration sensitivity (Section 4.5).

4.3 Constraint modeling

In the simple model of 4.2, the application of rigid tangential displacement constraints at the vertices implies that any acceleration or gravitational body load will be counteracted by the generation of tangential reaction forces (i.e. due to friction between the supports and the cavity). Consequently, the normal forces applied to the cavity are unperturbed by the acceleration. Realistically, the supports are typically an order of magnitude less rigid under acceleration than the spacer for common materials like Viton, PTFE, or Nylon; therefore, the assumption of rigid tangential constraints may not be valid. This implies that accelerations could result in reaction forces with normal components. If the force sensitivity is not exactly zero (as will generally be the case due to imperfect fabrication), then this can result in a nonzero acceleration sensitivity.

In general, the reaction forces generated in response to acceleration will be partitioned between the normal and tangential directions, with the relative amount of each dictated by the material properties and geometry of the supports. To gain some qualitative insight into this phenomenon, we study a toy model where a single cylindrical support is subjected to an effective load

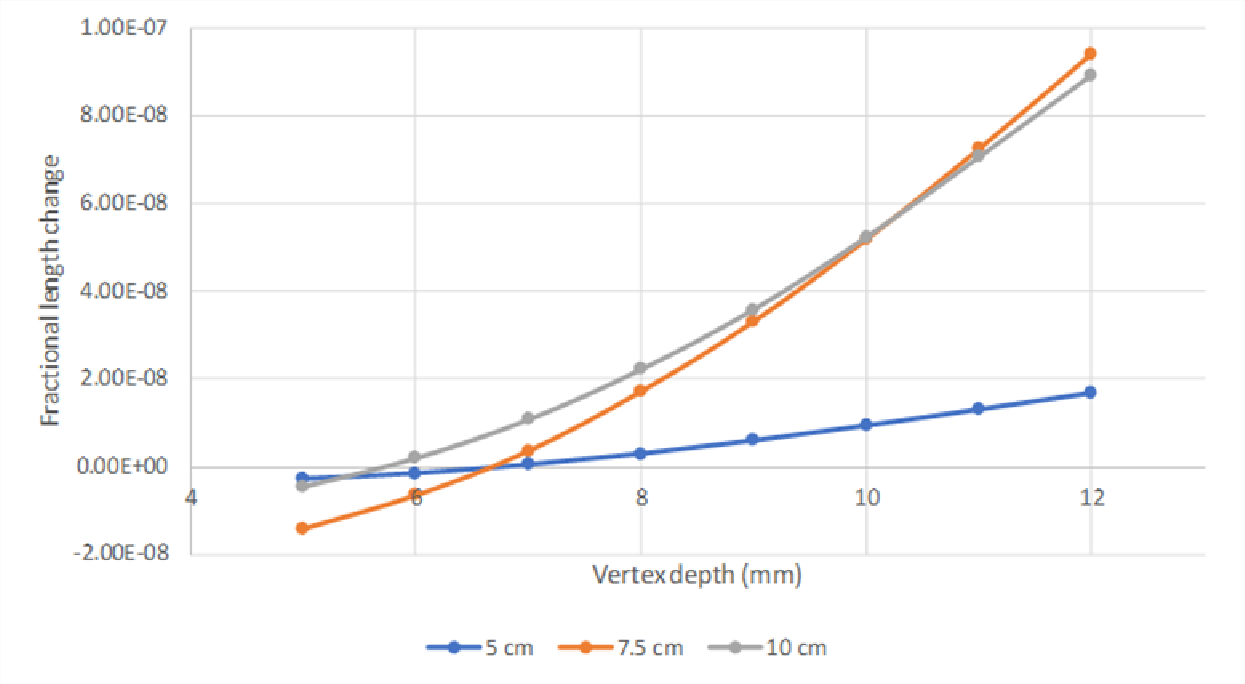

Figure 4.2: Deformation vs. vertex cut depth for three different spacer designs. Blue: the original NPL design with 5 cm length, 5.1 mm mirror diameter, 4 mm mirror thickness, and a single centered pair of vent holes. Gray: our design with 10 cm length, 12.7 mm mirror diameter, 6.35 mm mirror thickness, and two pairs of vent holes with 12.7 mm diameter and 40 mm offset. Orange: an intermediate case with 7.5 cm length, 12.7 mm mirror diameter, 6.35 mm mirror thickness, and two pairs of vent holes with 12.7 mm diameter and 24 mm offset. The length of the two non-optical axes is 5 cm in all three cases.

representing the gravitational force of the spacer, applied at an angle to its contact face (Figure 4.3). The backside of the support is fixed, and we study the displacement of the contact boundary as a function of the material properties and support geometry (Figures 4.4, 4.5). The ratio of normal to tangential displacement on the contact interface of the support can be used as a proxy for the strength of the constraints applied on the spacer. As the aspect ratio of the cylinder increases, the displacement becomes almost purely tangential. Therefore, the tangential constraints are weak, and the reaction forces on the other supports would be dominantly normal. Note that the opposite regime of strong tangential constraints cannot be physically achieved with this support geometry - the ratio of normal to tangential displacement peaks for a support with near unity aspect ratio, but is generally less than 1. Therefore, we find lacking support for the direct tangential constraints

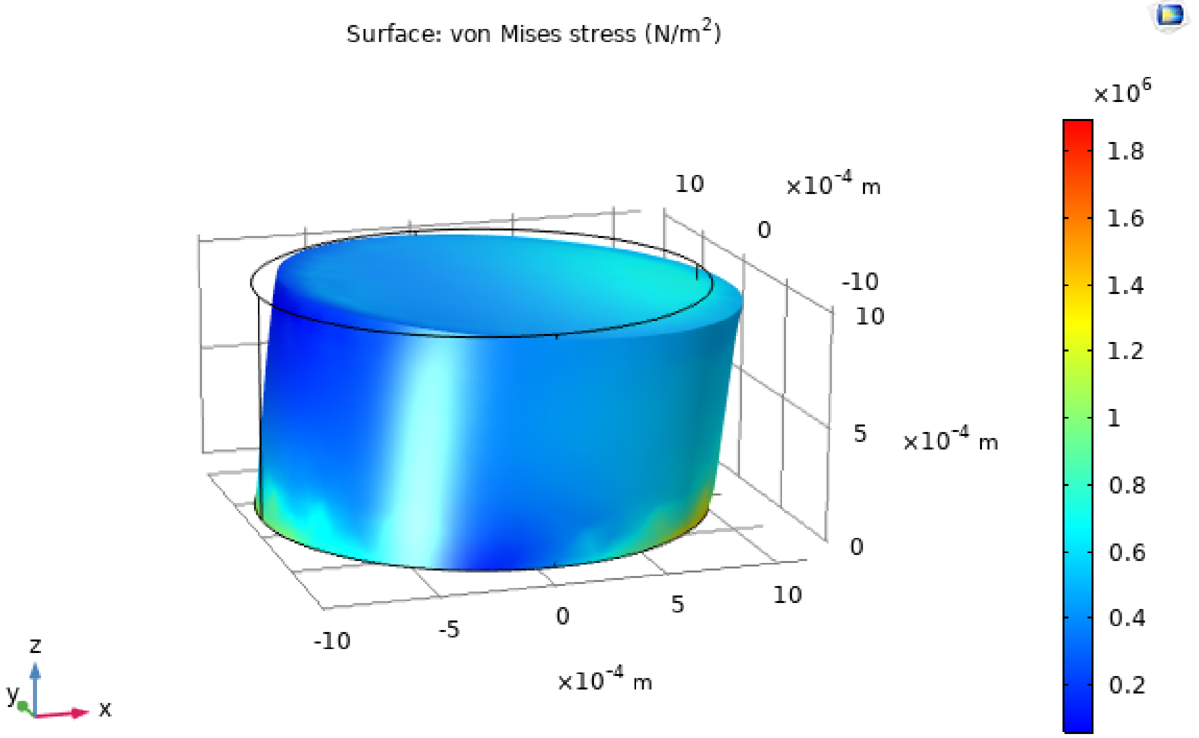

Figure 4.3: Toy model for qualitative simulation of support deformation and reaction force partitioning.

sometimes assumed in the literature and employed within the simple model of Section 4.2.

We corroborate this toy model by adding four cylindrical supports to the full spacer model. The compressional force of one support is varied, and we study the reaction forces of the other three supports, defining the *partitioning ratio* to be the ratio of the average normal reaction force to the applied force. A partitioning ratio of 1 means that the normal forces are self-balancing (normal constraint limit), while a ratio of 0 implies an imbalance between support forces (tangential constraint limit).

We observe that as the support length increases, the partitioning ratio generally increases as well, except for a local minimum at unity aspect ratio. This overall trend can be understood by comparison with the toy model (Figure 4.4) - as the support length increases, the tangential constraints grow weaker, and thus more normal reaction force is required to stabilize the cavity

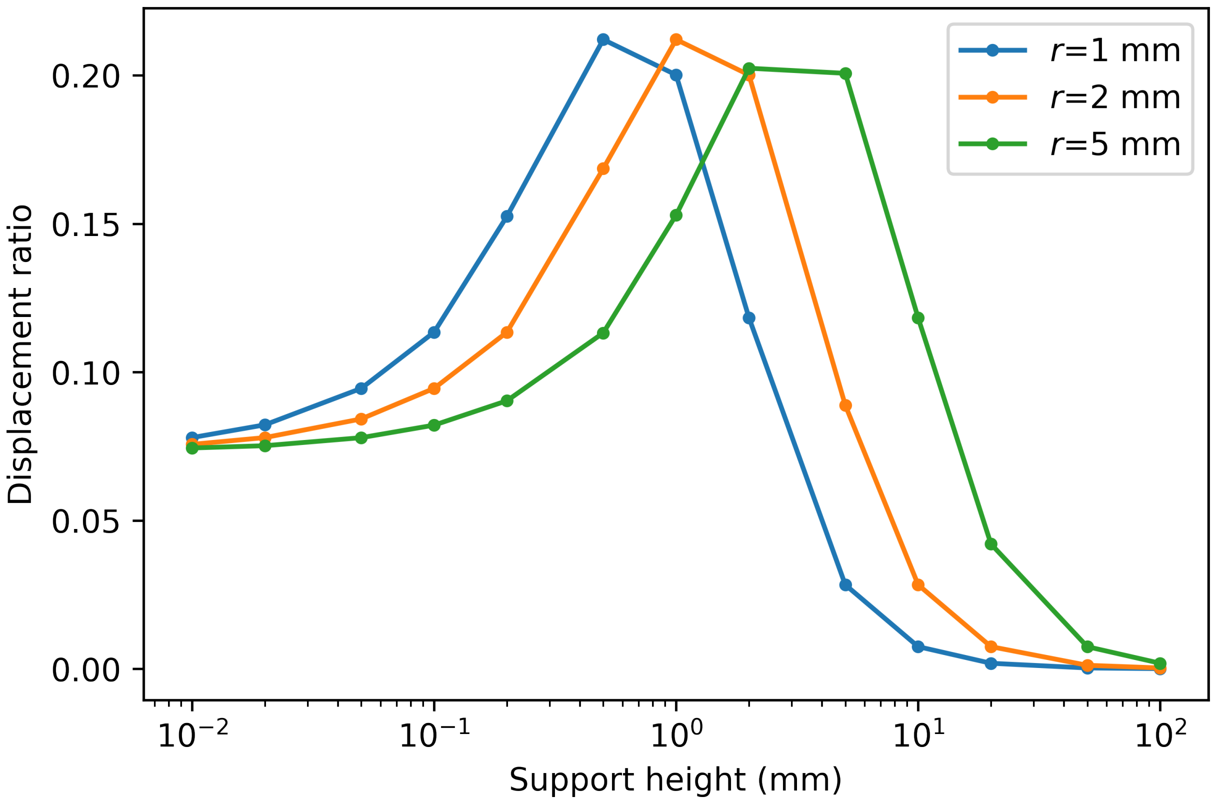

Figure 4.4: Toy model results for the normal:tangential displacement ratio as a function of support height and radius. The ratio is maximized for aspect ratios close to unity.

against the applied force.

Taken together, these two models give important insight to the physical constraints applied by the supports. In general, reaction forces generated to stabilize the cavity against an acceleration or an imbalanced support force will be partitioned between the normal and tangential components. This consideration is especially important in the study of acceleration sensitivities due to fabrication errors. A full acceleration sensitivity budget based on fabrication tolerances is presented in Section 4.5, but we examine two of the largest effects in closer detail here.

The first error is a lateral displacement of the spacer vent holes away from the optical axis of the cavity. We observe that this produces an acceleration sensitivity which increases with the force partitioning ratio, vanishing for direct tangential constraints (Fig 4.7a). The realistic case of hemispherical supports is approximately halfway between the two regimes.

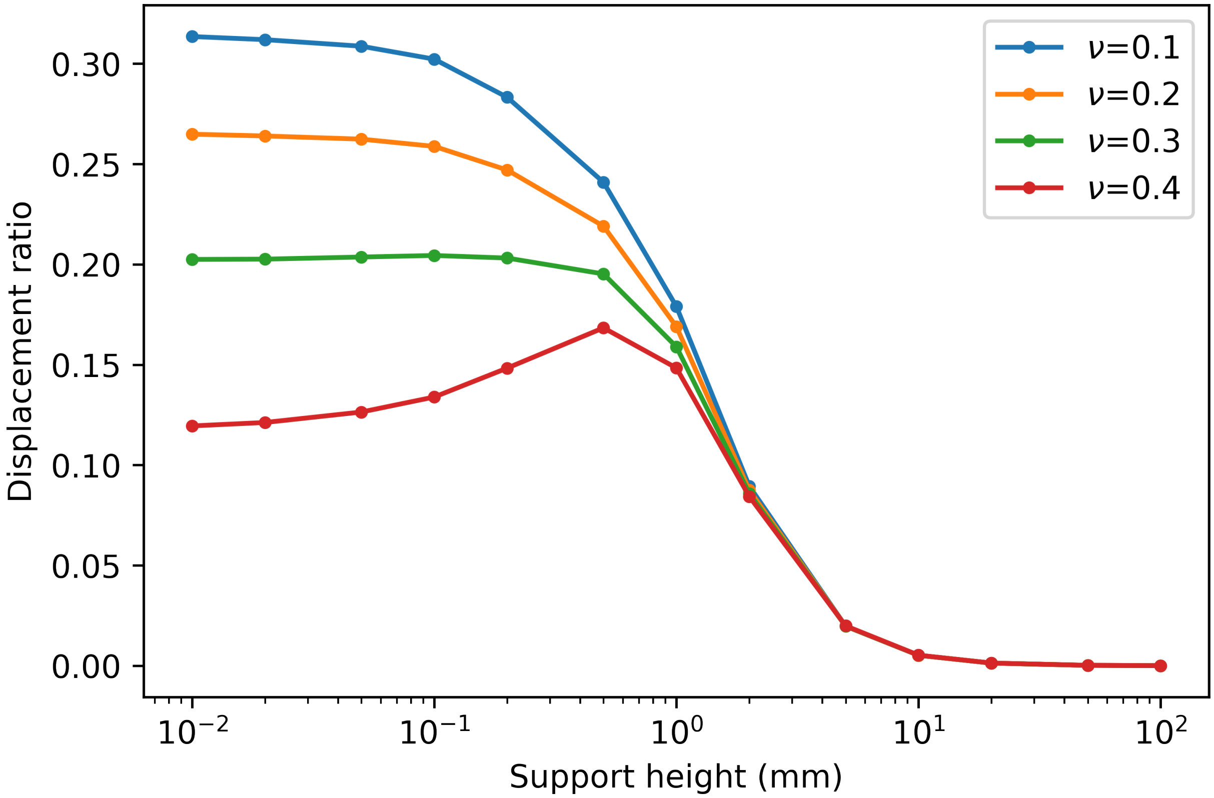

Figure 4.5: Toy model results for the normal:tangential displacement ratio as a function of support height and Poisson's ratio. For very short supports, higher Poisson's ratios result in relatively more tangential displacement through the Poisson effect.

The second error is an offset of the spacer relative to the nominal mounting points. Again, we observe that the acceleration sensitivity is zero for direct tangential constraints and generally increases with the partitioning ratio (Fig. 4.7b).

In summary, proper treatment of the constraints and reaction forces is crucial for estimating acceleration sensitivity due to fabrication errors.

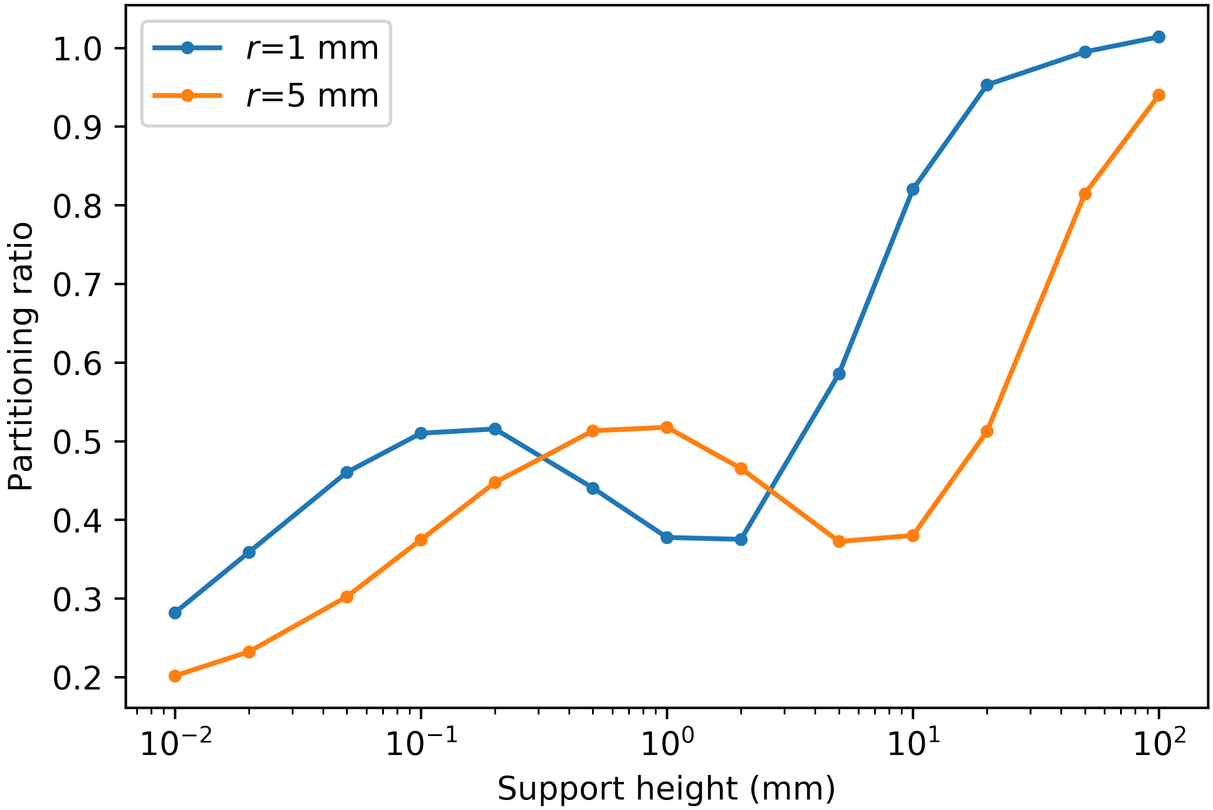

Figure 4.6: Partitioning ratio as a function of the radius and length of cylindrical supports. Flexible supports (long length, small radius) exhibit partitioning ratios near 1, indicating that the reaction forces are dominantly normal, or, equivalently, that the tangential constraints are weak; the opposite is true for more rigid supports with short length and large radius. The local minimum near unity aspect ratio is qualitatively similar to the single-support model results presented in Fig. 4.4.

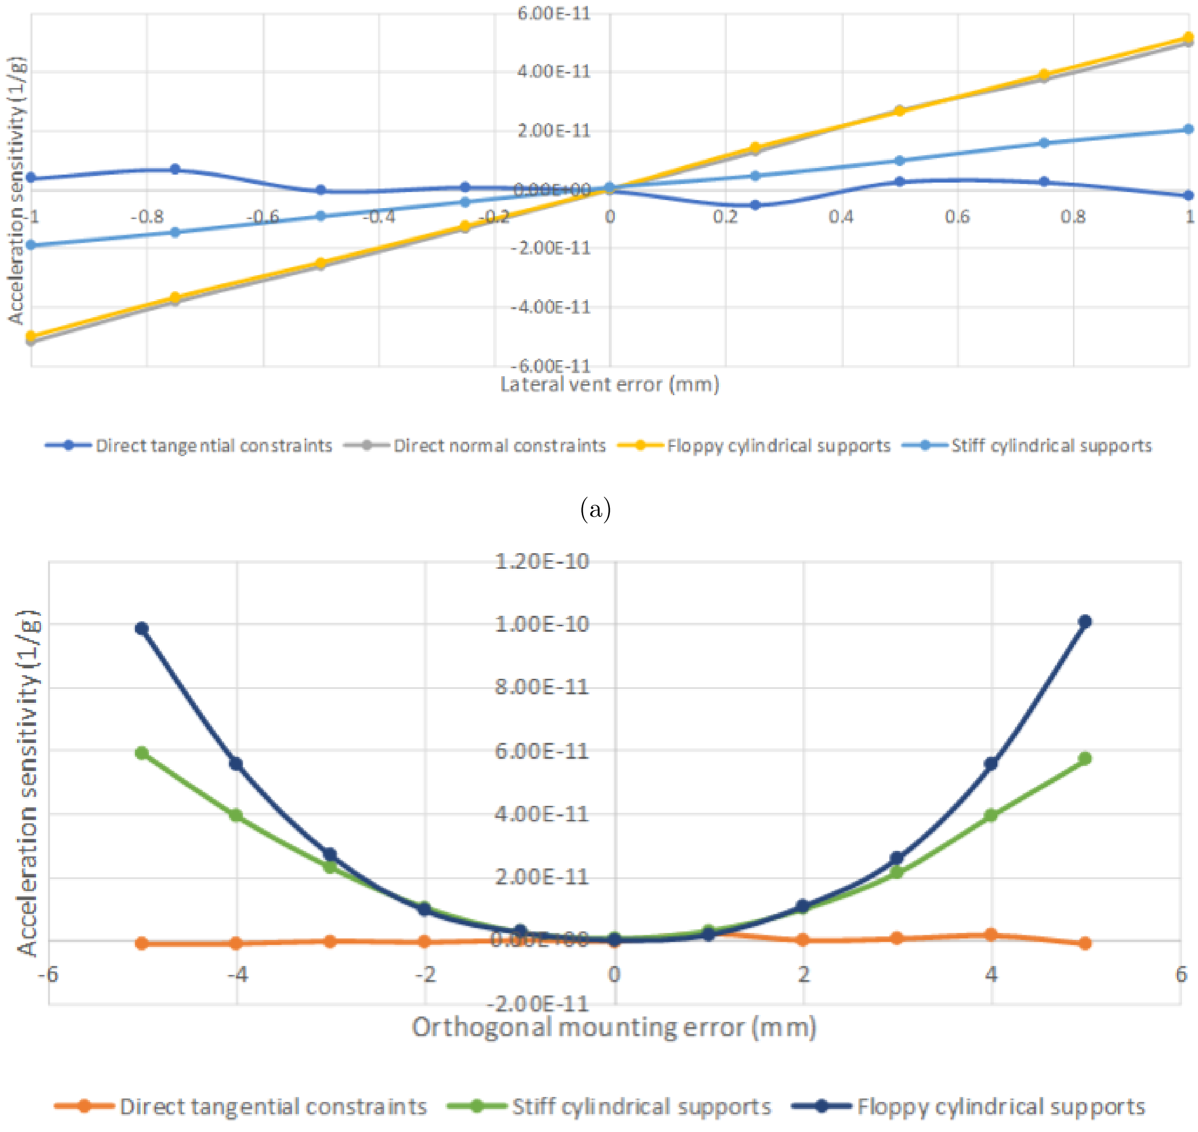

(a)

(b)

Figure 4.7: Acceleration sensitivity due to a) lateral positioning error in one spacer vent hole or b) spacer positioning error for several constraint cases. Direct tangential constraints: rigid displacement constraints applied directly to the support points. Direct normal constraints: lower (upper) support forces increased (decreased) by $mg\sqrt{3}/2$ to counteract the gravitational load. Stiff cylindrical supports: 100 GPa, 0.1 mm length, 8 mm radius. Floppy cylindrical supports: 0.1 GPa, 20 mm length, 8 mm radius. Hemispherical supports: 10 GPa, 8 mm radius, 0.5 mm contact radius.

4.4 Residual force sensitivity

Although the design of Section 4.2 grants perfect insensitivity of the cavity length to changes in holding force, the force sensitivity of the realized spacer will be nonzero due to fabrication errors. Consequently, thermal expansion due to temperature fluctuations in the support structure can cause a cavity length change via modulation of the applied force. With machining tolerances of 0.1 mm on the vertex depth, we predict the residual sensitivity will be below 4×10^{-11}/N.

An important consideration is the thermal stability required to reach a given frequency stability based on this residual sensitivity. We can gain analytical insight into this problem through the well-studied problem of the elastic deformation of two spheres in contact [29]. This simple model can be applied to our model by extending the radius of one sphere to infinity, representing the vertex plane. In this case, the deformed sphere contacts the spacer in a circular region with diameter

$$a = \left(\frac{3FR}{4E}\right)^{1/3} \tag{4.2}$$

where

$$\frac{1}{E} = \frac{1-\nu_{plane}^2}{E_{plane}} + \frac{1-\nu_{sphere}^2}{E_{sphere}} \approx \frac{1-\nu_{sphere}^2}{E_{sphere}} \tag{4.3}$$

and

$$\frac{1}{R} = \frac{1}{R_{sphere}} + \frac{1}{R_{plane}} \approx \frac{1}{R_{sphere}} \tag{4.4}$$

For a nylon sphere of radius $R = 0.75$ cm with $E = 3$ GPa and $\nu = 0.41$, the contact area is 0.5 mm^2.

The displacement of the backside of the hemispherical support due to the applied force F can be written

$$\Delta z = \left(\frac{9F^2(1-\nu^2)^2}{16RE^2}\right)^{1/3} \tag{4.5}$$

Under a temperature change ΔT, the sphere expands to a new radius $R' = R(1+\alpha\Delta T)$, where α is the coefficient of thermal expansion. The thermal expansion also causes an effective normal displacement equal to the difference between the perturbed and original radii:

$$\Delta z' = \Delta z + (R' - R) = \Delta z + R\alpha\Delta T \tag{4.6}$$

Under the replacements $R \to R'$ and $\Delta z \to \Delta z'$, equation (4.5) becomes

$$\Delta z + R\alpha\Delta T = \left(\frac{9F'^2(1-\nu^2)^2}{16R(1+\alpha\Delta T)E^2}\right)^{1/3} \tag{4.7}$$

Solving for the perturbed force, we find

$$F' = \frac{4}{3}\frac{E}{1-\nu^2}R^2\left(\Delta z/R + \alpha\Delta T\right)^{3/2}\left(1+\alpha\Delta T\right)^{1/2} \tag{4.8}$$

This result shows that in order to minimize force fluctuations due to thermal expansion, the supports should be small and made out of a low-CTE material with small Young's modulus and high Poisson's ratio. Thus, thermal expansion will be small and the effects will be absorbed into compression of the support rather than the spacer.

Together with the force sensitivity of the cavity, which can be estimated via FEA or measured directly, equation (4.8) can be used to estimate the frequency shift caused by a temperature change of the supports (neglecting expansion of the spacer or substrates). Figure 4.8 shows the shift vs. temperature change relative to the CTE zero-crossing of the ULE spacer for several different support materials, as well as the shift caused by thermal expansion of the spacer itself[1] . For small fluctuations, the shift is dominated by thermal expansion of the supports, while the expansion of the spacer will dominate for large temperature changes. At the 10^{-16} level, none of the three support materials contribute significantly to the total instability.

4.5 Fabrication errors

Although the stretched cubic cavity design is perfectly insensitive to acceleration, a sensitivity can arise due to fabrication errors. We estimate the severity of this effect by perturbing various design features and calculating the induced acceleration sensitivity both parallel and perpendicular to the optical axis. As noted in Section 4.3, proper modeling of the constraints and reaction forces is critical for estimating the effects of fabrication errors. For this analysis, we use cylindrical nylon supports with 0.5 mm contact radius and unity aspect ratio. This support geometry leads

---

[1] Despite operating at the zero-crossing of the spacer CTE, there is a second-order variation of cavity length with temperature.

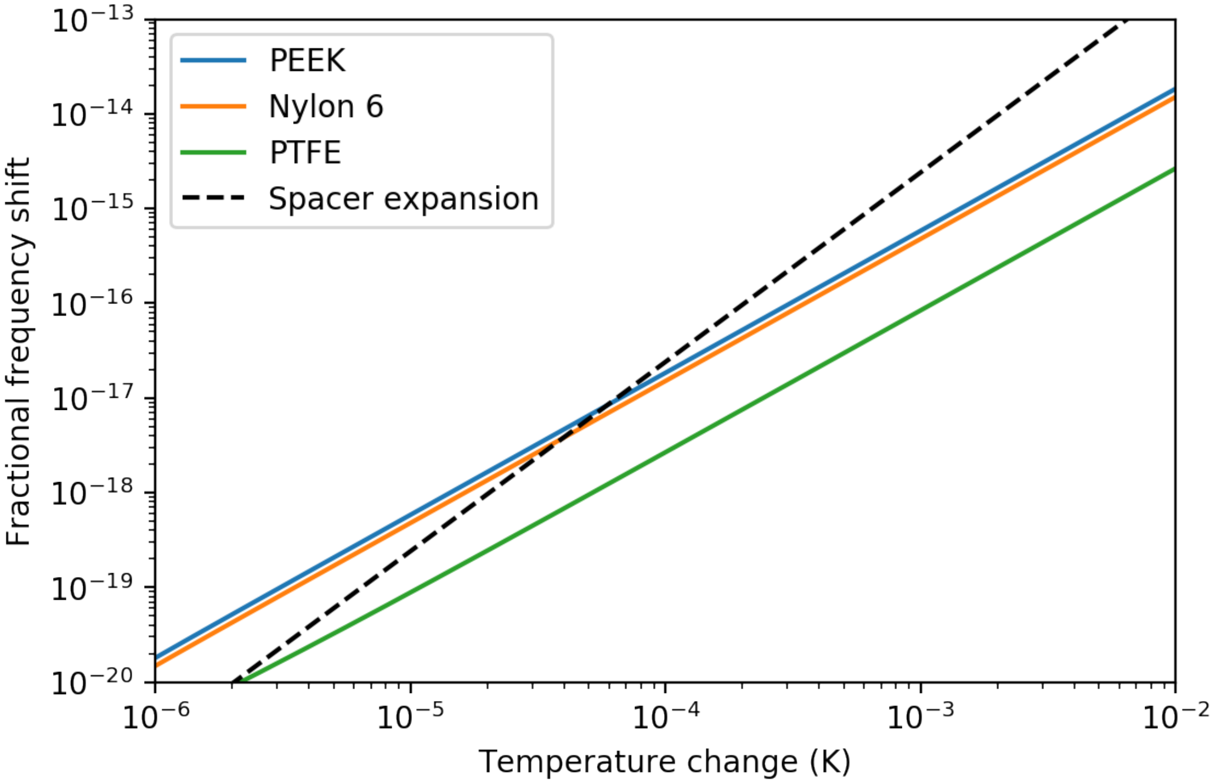

Figure 4.8: Frequency shift vs. temperature change for several support materials for a residual sensitivity of 4×10^{-11}/N and 50 N holding force. The expansion of the spacer assumes a linearly-varying spacer CTE with a coefficient of $2.4 \times 10^{-9}\ 1/K$ [57].

to conservative over-estimates of the acceleration sensitivities, while also being computationally simple, compared to hemispherical supports. A nominal 50 N holding force is applied to each support by prescribing an appropriate displacement of the back side of the cylinder. The back side of the cylinder is constrained against tangential motion, and the contact side has a no-slipping constraint with the cavity vertex plane.

4.5.1 Lateral acceleration

We first study accelerations perpendicular to the optical axis, which includes the gravitational force. We discuss the two largest contributions to the potential acceleration sensitivity of the realized cavity; all other fabrication errors yield acceleration sensitivities below the $4 \times 10^{-12}/g$ level for achievable machining tolerances. Our total acceleration sensitivity uncertainty for the lateral direction is $6.0 \times 10^{-11}/g$.

4.5.1.1 Vent position

Targeting a 0.1 mm tolerance, we perturb the position of a single vent hole. The acceleration sensitivity is dominated by errors in vent positioning perpendicular to the optical axis; positioning error parallel to the optical axis contributes negligibly. The vents perpendicular to the applied acceleration contribute an acceleration sensitivity of $2.8 \times 10^{-11}/g$, while the other two contribute $1.6 \times 10^{-11}/g$. We assume independent errors for all four vents, yielding a total acceleration sensitivity tolerance of $4.6 \times 10^{-11}/g$ for vent positioning.

4.5.1.2 Mounting

There are two separate effects to consider regarding the positioning of the cavity and the applied support forces: a) positioning error of the spacer relative to the support frame and b) positioning error of the supports within the vertex cut planes.

To accurately position the cavity before the supports make contact, we have added four small dowel holes at the auxiliary mounting points shown in Fig. 4.1. Based on the machining tolerances of the cavity spacer and the concentricity tolerances of the screws used for the auxiliary support forces, we estimate that the spacer can be positioned to within an overall accuracy of 0.15 mm. Positioning error is simulated by translating the support points on all for vertices. We assess that mounting error orthogonal to the cavity axis contributes an acceleration sensitivity tolerance of $2.0 \times 10^{-11}/g$, while mounting error in the parallel to the cavity axis contributes negligibly. Assuming independent errors for all three directions, we estimate the acceleration sensitivity tolerance to be $2.9 \times 10^{-11}/g$.

The second consideration is the error in the support forces themselves, independent of positioning error of the spacer relative to the support frame. We estimate that the support forces can be applied to within 0.15 mm of the center of the support planes, dominated by concentricity tolerances in the screws used to tighten the supports. This error results in an acceleration sensitivity of $1.3 \times 10^{-11}/g$ per support, and a total acceleration sensitivity of $2.6 \times 10^{-11}/g$ from the four supports.

Overall, the mounting effects contribute an acceleration sensitivity tolerance of $3.9 \times 10^{-11}/g$, slightly smaller than the uncertainty contributed by vent positioning error.

4.5.2 Axial acceleration

Now, we study accelerations perpendicular to the optical axis. Again, we find that only several effects are non-negligible, granting a total acceleration sensitivity uncertainty of $6.7 \times 10^{-11}/g$.

4.5.2.1 Vent diameter error

While error in the vent diameter contributed negligibly to the lateral acceleration sensitivity, it dominates the axial acceleration budget. Assuming a tolerance of 0.1 mm on vent diameter, we assess the acceleration sensitivity to be $2.9 \times 10^{-11}/g$ per vent and $5.8 \times 10^{-11}/g$ for independent errors in all four vents.

4.5.2.2 Mounting

We study the effects of spacer and support positioning error similarly to their treatment for lateral acceleration. We find now that mounting error along the cavity axis contributes $2.9 \times 10^{-11}/g$, while mounting error in the orthogonal directions contributes negligibly below the $10^{-12}/g$ level. Positioning error in a single support contributes $8.4 \times 10^{-12}/g$, with all four supports contributing $1.7 \times 10^{-11}/g$. Overall, mounting and support errors contribute a potential acceleration sensitivity of $3.4 \times 10^{-11}/g$.

4.6 Vacuum chamber

The vacuum chamber for the cavity is shown in Figure 4.9. The chamber is made primarily from aluminum, interfaced with stainless steel hardware with bimetal explosion-bonded flanges, and has outer dimensions of 27x35x22 cm and weighs approximately 8 kg. Two sides are removable

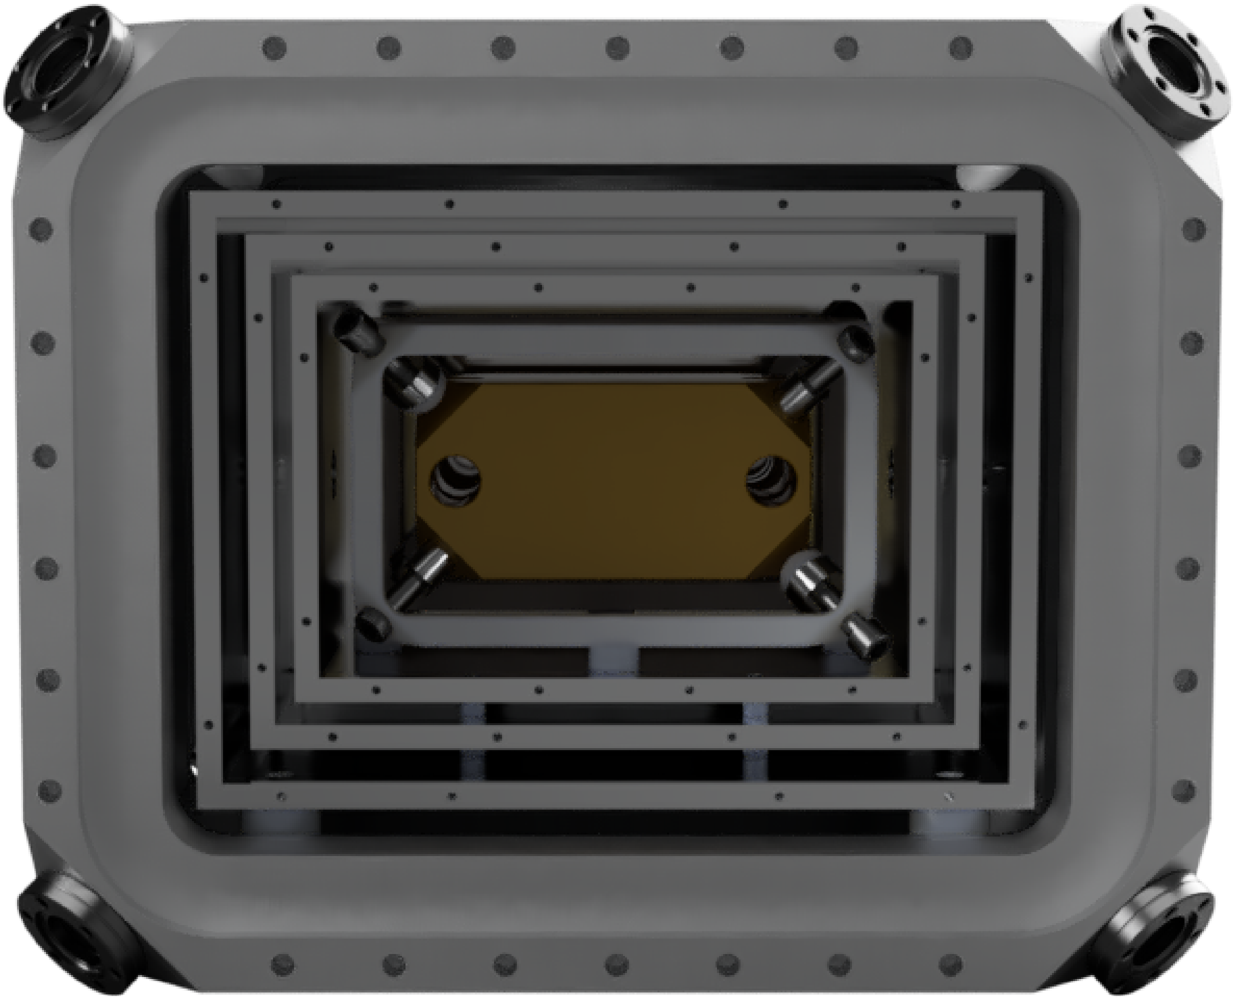

Figure 4.9: Cavity vacuum chamber rendering, opened to show the nested thermal shield and cavity mounting assemblies.

for insertion and access of the thermal shield assembly and cavity. The chamber is sealed with an elastometer gasket in the trench surrounding the open door, which is designed to support pressures below 10^{-8} Torr. The eight 1.33" ConFlat flanges are concentric with the cavity vertices, allowing *in vacuo* adjustment of the cavity mounting forces with magnetically-coupled rotation feedthroughs. This procedure requires the thermal shields to be removed, and will be done as part of the initial setup process during measurement of the acceleration sensitivity.

The thermal shield assembly consists of three separate shields, each consisting of six aluminum plates. The shields are connected to one another with polyetheretherketone (PEEK) spacers, the choice of material motivated by its low thermal conductivity of 0.25 W/(m K); four PEEK screws are used to fasten the innermost shield to the outer shield, passing through the middle shield and two layers of spacers. Vent holes on the shield assembly are staggered to provide vacuum conductance without direct line-of-sight through the entire assembly. Optical access to the cavity is granted by pairs of windows mounted in SM1 (1.035") threads on each shield. Alternating windows are angled by 5^o with respect to one another to avoid parasitic etalons.

4.7 Conclusion

We have altered previous cubic cavity designs by stretching the spacer along the optical axis, demonstrating that insensitivity to holding force and acceleration can be preserved. Our cavity, with 5 cm side length and 10 cm optical length, has a mass and volume a factor of 4 smaller than a cubic cavity of equivalent length. By analyzing the effects of fabrication errors, we estimate worst-case force and acceleration sensitivities of $4 \times 10^{-11}/\mathrm{N}$ and $5 \times 10^{-11}/g$ respectively. The spacer length and low-loss crystalline mirrors support a thermal-noise-limited instability near the 10^{-16} level. Manufacturing of the cavity and vacuum chamber is expected to finish in early 2022.

Chapter 5

Looking forward

While the work described in the previous chapters has culminated in a full demonstration of clock operation, there are ongoing efforts towards a full uncertainty characterization, the first field deployment, and upgrades to the clock apparatus, which will be described in this chapter. Efforts towards miniaturizing the atom trapping optics (Section 5.4) and atom source (Section 5.5) are complementary towards an eventual goal of improving clock stability with a zero-dead-time architecture requiring two atomic ensembles (Section 5.3).

5.1 Field deployment

Over the next year, various efforts will be undertaken to prepare for deployment, including:

- Stress testing with exaggerated temperatures and acoustic and vibrational noise
- Integration and testing of the clock laser and cavity
- Complete rack integration of all subsystems
- Full uncertainty evaluation in the low-10^{-18} decade
- Recovery of clock operation after short transportation

After these efforts, we will undertake a relativistic geodesy measurement as an interesting proof-of-concept use case for transportable optical lattice clocks. To obtain a large lever arm for differential geopotential measurements against a reference clock at NIST, we will take advantage of Colorado's mountainous topography. We have identified several peaks above 14,000 feet with facilities potentially amenable to clock operation: the US Army High Altitude Research Lab on Pikes Peak and the Meyer-Womble Observatory on Mount Evans. The former location offers line of sight to Table Mountain, where a relay node for free-space time and frequency transfer back to NIST has already been implemented [12, 27]. With the latter location, remote comparison with NIST could be implemented through a free-space link to a local facility which is connected by optical fiber to NIST.

In either location, the large gravitational redshift of approximately 3×10^{-13} offers a sixfold improvement in sensitivity to general relativity violations compared to the Tokyo Skytree measurement [98]. Even with a modest clock uncertainty of 10^{-17}, this large lever arm would permit a factor of three improvement in constraints on GR violations. Furthermore, the convenient access to indoor facilities above 14,000 feet, as well as the semi-permanent frequency transfer infrastructure, gives us a useful testbed for further improvement of constraints as the clock uncertainty improves.

5.2 Improving vacuum pressure

Our epoxy seals described in Section 2.7 support UHV pressures at the nTorr level or better after bakeouts. However, for some chambers, we have observed a long-term pressure rise, converging near 10^{-8} Torr after several months (Figure 5.1). This rise has been observed, with variations in rate, across three vacuum bakeouts and two separate chambers. We speculate that this could be caused by a slow diffusion of gases through the epoxy, resulting in a leak rate that increases over time before saturating. We are currently experimenting with a second science chamber in order to better understand this effect. We have also developed an alternate design for the science chamber, which is currently being manufactured, consisting of a titanium body with windows brazed to weld rings and welded on. The dimensions are nearly identical, such that it is a drop-in replacement for the aluminum epoxy chambers.

Figure 5.1: Chamber pressure vs. time since a bakeout, showing a gradual rise to the 10 nTorr level for our clock vacuum chamber (blue). A second test chamber (orange) exhibits a slower rise.

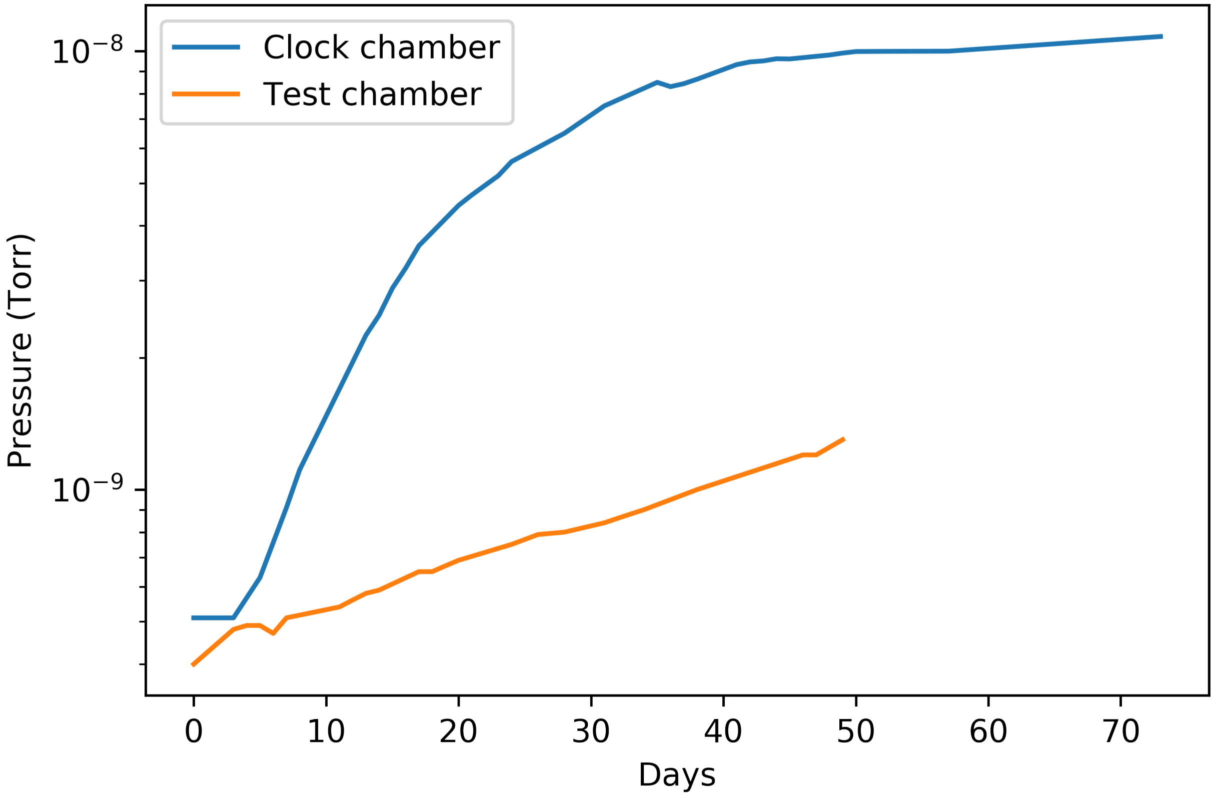

5.3 Dual-ensemble zero-dead-time clock

A significant, often dominant, source of instability in optical lattice clocks is the Dick effect. This phenomenon arises due to *dead time* in the clock cycle required for atom trapping and cooling, quantum state preparation, and state readout - all parts of the sequence other than interrogation of the clock transition. Due to this periodic interrogation with sub-unity duty cycle, laser frequency noise is aliased onto the clock servo at harmonics of the sampling frequency $1/T_c$, where T_c is the cycle time. Since this aliased noise is indistinguishable from noise within the detection bandwidth, it leads to spurious compensation by the clock servo which degrades the overall stability.

Elementary strategies to reduce the instability contributed by the Dick effect are a) reducing the short-term instability of the clock laser, and b) increasing the spectroscopy duty cycle. These are not mutually exclusive - efforts towards the first strategy will naturally lead to the second, as a more stable clock laser allows longer spectroscopy time without phase slips. Improvement of the clock laser has an added benefit that the higher quality factor allowed by increased spectroscopy time proportionally reduces the effects of all other sources of instability: quantum projection noise, technical detection noise, and photon shot noise.

An alternate strategy is to interleave interrogation between multiple atomic ensembles such that the clock laser is continuously exposed to the atomic wavefunction. We implemented this technique in 2016 [90], demonstrating that this *zero-dead-time* clock architecture allowed a clock instability of $6 \times 10^{-17}/\sqrt{\tau}$ compared to an anti-synchronized scheme with $1.4 \times 10^{-16}/\sqrt{\tau}$ (Figure 5.2). This stability improvement results in a six-fold reduction in averaging time required to reach a given stability. We also demonstrated an alternate mode of operation where one clock is operated with short measurement cycles to provide fast feedback to the cavity and extend the coherence time of the clock laser. This allowed Fourier-limited linewidths as low as 120 mHz (4 s spectroscopy time), limited primarily by decoherence due to lattice trap loss. While this configuration does not provide substantial suppression of the Dick effect, it can be applied to improve the quality factor of the transition, reducing the overall instability from all sources.

In [90], we opined that

> because of the relaxed requirements on OLO performance, this scheme can reduce the total complexity of state-of-the-art optical clocks, extending their usability outside the laboratory.

It seems paradoxical at first glance to consider adding a second physics package while attempting to miniaturize an optical lattice clock. However, high-performance cavities impose much more stringent demands on environmental isolation than the physics package. The impacts of environmental fluctuations on the physics package comprise slow misalignment due to beam pointing drift, as well as Doppler shifts due to vibration that can be effectively suppressed using phase-noise cancellation to the reference frame of the optical lattice; neither of these effects are likely to substantially degrade clock performance. In contrast, high-performance optical cavities require multiple layers

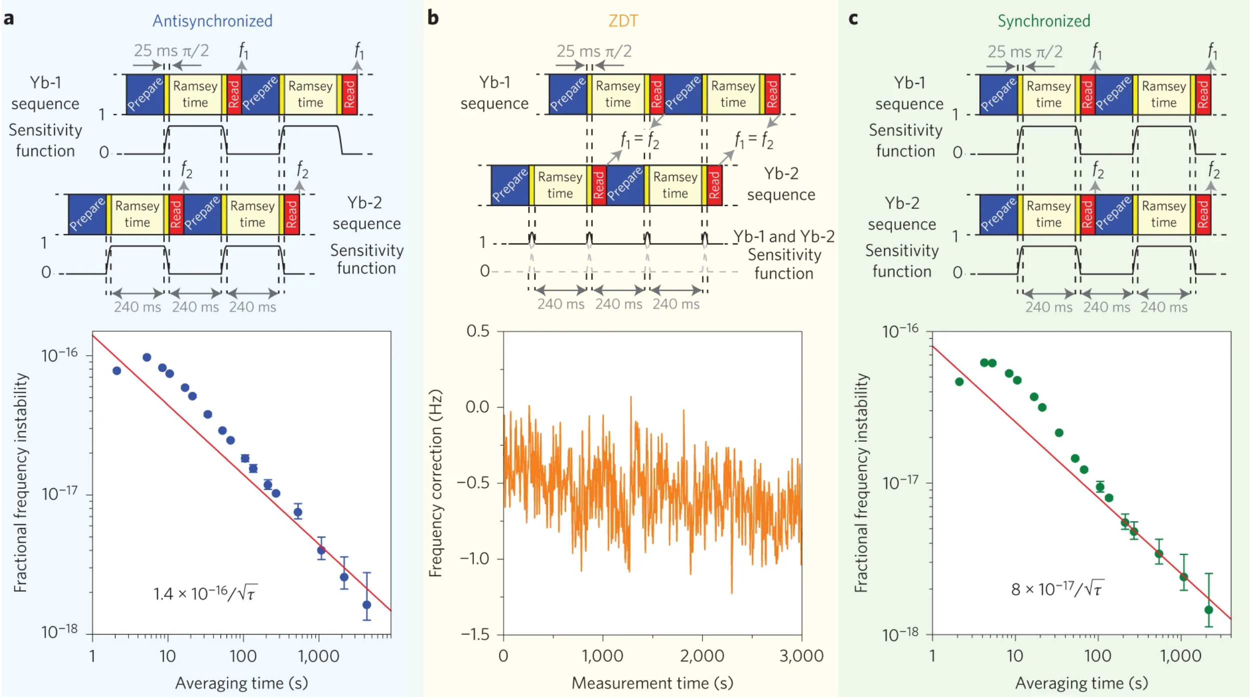

Figure 5.2: Zero-dead-time clock operation with Yb1 and Yb2. Reproduced from [90]. a) Antisynchronous interrogation, providing maximum sensitivity to the Dick effect, yielding a single-clock instability of $< 1.4 \times 10^{-16}/\sqrt{\tau}$. b) Zero-dead-time operation with a shared, phase-continuous frequency correction using both clocks. c) Synchronous interrogation, eliminating the Dick effect to estimate the performance improvements from the ZDT configuration without requiring four atomic ensembles. The true ZDT instability is $6 \times 10^{-17}/\sqrt{\tau}$, a factor of $\sqrt{2}$ smaller, since the ZDT clock accumulates updates twice as quickly.

of thermal shielding, a heavy vacuum chamber, acoustic isolation, and passive or active vibration isolation. Additionally, decreasing the cavity thermal noise floor into the mid-10^{-17} decade requires either a cavity length of tens of centimeters or cryogenic operation, neither of which is appealing in a transportable system. Either of these options would require a considerable engineering effort, whereas duplication of the physics package is made painless by the modular breadboard design which streamlines the setup time compared to conventional optical layouts involving pedestals and clamping forks. Moreover, two upgrades described in the following sections will substantially reduce the volume of the physics package, making a zero-dead-time architecture a natural choice for a transportable clock.

5.4 Grating magneto-optical trap

Despite being a laboratory workhorse for decades, the six-beam MOT is unappealing for compact systems due to the high volume of optics required for light delivery and polarization and intensity control, as well as the number of alignment degrees of freedom and free parameters which must be optimized. A first step towards reducing the complexity of a MOT is a tetrahedral configuration with four beams [95], the minimum required for three-dimensional confinement. This geometry can be realized with a single input beam and a triplet of diffraction gratings [103]. In this section, I will discuss the basic theory of grating MOTs, as well as experimental efforts towards realizing a grating MOT (gMOT) with ^{171}Yb.

5.4.1 Basic theory

To study this type of trap mathematically, we consider a uniform beam with saturation parameter β_0 traveling along the -z axis with wavevector $\mathbf{k}_0 = k(0,0,-1)$ and circular polarization $s = 1$. The beam diffracts from a triplet-patterned grating, producing three beams traveling upwards with wavevectors[1] $\mathbf{k}_i = k(\sin\alpha\cos\frac{2\pi i}{3}, \sin\alpha\sin\frac{2\pi i}{3}, \cos\alpha)$, saturation parameter β_i, and polarization $s = -1$. The diffraction angle α is measured from the $z-$axis and is equal to 70.5° in a pure tetrahedral configuration. The total force can be calculated with equation (2.4). We can repeat the series expansion around the phase-space origin, as we did for the six-beam magneto-optical trap. At lowest order, the transverse forces vanish by symmetry, leaving an axial component

$$F_z^{(0)} = -\hbar k \frac{\Gamma}{2} \frac{\beta_0 - 3\beta_i \cos\alpha}{1 + \beta_T + 4\Delta^2/\Gamma^2}, \tag{5.1}$$

which is balanced when $\beta_i = \beta_0/(3\cos\alpha)$. Given that the intensity is increased by a factor of $\cos\alpha$ due to beam compression upon diffraction, this condition is equivalent to a first-order diffraction efficiency $R_1 = 1/3$.

---

[1] This definition corresponds to a grating with a symmetry axis in the yz-plane.

At first order, we find a damped harmonic oscillator form

$$\mathbf{F}^{(1)}(\mathbf{r},\mathbf{v}) = -\kappa\mathbf{r} - \gamma\mathbf{v} \tag{5.2}$$

where κ and γ are 3x3 matrices with diagonal elements given by

$$\mathrm{diag}(\gamma) = (\gamma_r,\gamma_r,\gamma_z) = -2\hbar k^2 \frac{\beta_0}{(1+\beta_T+4\Delta^2/\Gamma^2)^2}\frac{\Delta}{\Gamma}\bigg(\sin\alpha\tan\alpha, \sin\alpha\tan\alpha, 2(1+\cos\alpha)\bigg) \tag{5.3}$$

$$\mathrm{diag}(\kappa) = (\kappa_r,\kappa_r,\kappa_z) = -2\hbar k \frac{\beta_0}{(1+\beta_T+4\Delta^2/\Gamma^2)^2}\frac{\mu_F B_1}{\hbar}\frac{\Delta}{\Gamma}\bigg(\sin\alpha\tan\alpha, \sin\alpha\tan\alpha, 4(1-\cos\alpha)\bigg) \tag{5.4}$$

Here, B_1 is the axial field gradient at the origin. It is interesting to note that the vertical trapping vanishes as $\alpha \to 0$, precisely because this situation is identical to a conventional 1D MOT with one beam having the wrong handedness.

5.4.2 Spatial intensity profile

In the previous section, we considered only the case of uniform beams and limited our attention to a region close to the phase-space origin; thus, the wavevectors alone suffice to describe the geometry of the acceleration profile generated by the superposition of the incident and diffracted beams. This simple analysis is suitable for analyzing the grating MOT through the lens of the damped harmonic oscillator model, offering considerable intuition to aid the choice of grating design parameters R_1 and α (Section 5.4.4. In order to assess the effectiveness of loading from an atomic beam, we will now extend the model to include the 3D intensity profile of the incident and diffracted beams.

We will start by relaxing the assumption of uniform beams, describing the incident beam through the wavevector $\mathbf{k}_0 = -k(0,0,1)$ and the intensity profile $I_0(x,y,z)$. We additionally generalized to a grating with N sectors. The N first-order diffracted beams have wavevectors

$$\mathbf{k}_n = k(\cos\phi_n\sin\alpha, \sin\phi_n\sin\alpha, \cos\alpha), \tag{5.5}$$

where α is the grating diffraction angle and

$$\phi_n = \pi\left(1 + \frac{2n+1}{N}\right) \bmod 2\pi \tag{5.6}$$

is the azimuthal angle determined by the sector. Note that this convention defines the sectors such that the beams originate in azimuthal angles in multiples of $2\pi/N$; the grating can be effectively rotated by adding an angular offset to ϕ_n.

Let's define our coordinate system such that the grating lies in the $xy-$plane at $z = 0$. In order to determine the intensity of one of the diffracted beams at a point $\mathbf{x} = (x, y, z)$, we can trace backwards from $\mathbf{x}$ along $\mathbf{k}_n$ until we reach the grating. This is described by the coordinate transformation

$$x_0 = x - z\frac{\mathbf{k}_n \cdot \hat{\mathbf{x}}}{\mathbf{k}_n \cdot \hat{\mathbf{z}}} = x - z\cos\phi_n \tan\alpha \tag{5.7}$$

$$y_0 = y - z\frac{\mathbf{k}_n \cdot \hat{\mathbf{y}}}{\mathbf{k}_n \cdot \hat{\mathbf{z}}} = y - z\sin\phi_n \tan\alpha \tag{5.8}$$

We can relate the intensity at $\mathbf{x}$ to the intensity of the incident beam on the grating as

$$I_n(x, y, z) = I_0(x_0, y_0, 0) \times R_1 \sec\alpha \times \Theta_n(\mathbf{x}) \tag{5.9}$$

where R_1 is the first-order diffraction efficiency and the $\sec\alpha$ term accounts for the beam compression. The $N-$sector mask is given by the Boolean $\Theta_n(\mathbf{x})$. To determine $\Theta_n(\mathbf{x})$, we define the sector from which the n^{th} beam originates to be bounded by two angles ϕ_n^+ and ϕ_n^-, such that

$$\phi_n^{\pm} = \left(\phi_n \pm \frac{\pi}{3}\right) \bmod 2\pi \tag{5.10}$$

We then determine the azimuthal angle of our test point $\mathbf{x}$ in a cylindrical coordinate system:

$$\phi_x = \mathrm{atan2}(y_0, x_0) + \pi \tag{5.11}$$

Note that we are using the *2-argument arctangent*, rather than the regular arctangent, and adding π to return an unambiguous angle in the $[0, 2\pi]$ range. We can then define the mask informally through

$$\Theta_n(\mathbf{x}) \overset{?}{=} \phi_n^- \le \phi_x \le \phi_n^+ \tag{5.12}$$

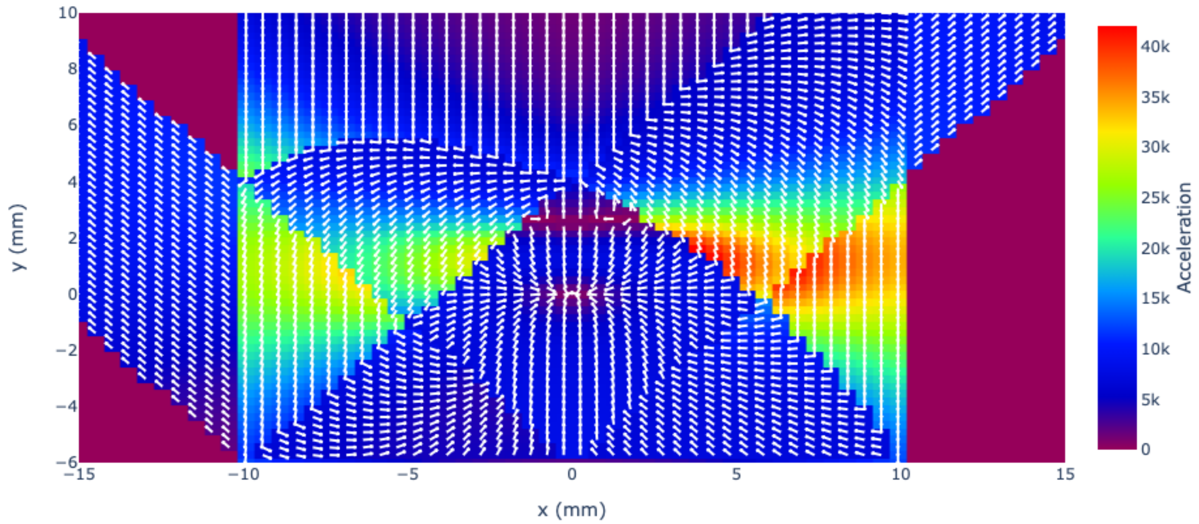

Figure 5.3: Acceleration profile for an atom at rest in a grating MOT.

This definition is "informal" in that it gets the idea right, but does not account for overflows of the $[0, 2\pi]$ angular range. The function is properly defined piecewise:

$$\Theta_n(\mathbf{x}) = \begin{cases} (\phi_n^- \leq \phi_x) \wedge (\phi_x \leq \phi_n^+) & \text{if } \phi_n^- \leq \phi_n^+ \\ (\phi_n^- \leq \phi_x) \vee (\phi_x \leq \phi_n^+) & \text{if } \phi_n^- > \phi_n^+ \end{cases}$$

Now, equation (5.9) can be used in combination with (2.4) to plot the trapping force in the vicinity of the grating (Figure 5.3). We model a 20 mm diameter grating, offset axially from the origin by 2 mm, with a diffraction angle of 45^o. The incident beam is tuned to $\Delta = -\Gamma/2$ and contains 30 mW of power with a 20 mm diameter, uniform intensity profile. The magnetic field is computed analytically for a pair of coils with 15 mm radius, 30 mm spacing, and 10 turns each; the gradient near the origin is 50 G/cm. The incident and diffracted beams overlap in a diamond-shaped region where a spatial restoring force exists. Outside of this region, the intensity profile is complicated by the existence of negative-order diffracted beams.

5.4.3 Beam loading

When loading from an atomic beam, there are two ways to orient the beam relative to the grating: axial loading, where the beam passes through the grating, and side loading, where the beam enters the trapping area parallel to the grating. The complicated intensity profile described in the previous section gives rise to several challenges which can inform the optimal loading direction.

The first challenge arises from the existence of even-ordered terms in the series expansion of equation (2.4), which vanish due to symmetry for a conventional six-beam MOT. These terms can introduce significant nonlinearities in the restoring force, causing anti-trapping behavior at a critical distance from the grating. An example of this phenomenon is shown in Figure 5.4a. The full analytical trapping force is computed using equation (2.4). The spatial profile is computed using the model in Section 5.4.2 for the same parameters as Figure 5.3. The range of currents from 5 to 50 A produces gradients from 8.3 to 83 G/cm at the center of the coils. For small magnetic fields, the restoring force is approximately linear around the origin. As the field increases, an unstable region moves closer to the origin, beyond which the force becomes anti-trapping. This phenomenon could degrade the trapping efficiency of axially-loaded gMOTs by rejecting atoms which pass the potential minimum and enter the unstable region.

Side loading poses its own challenges due to the requirement that the atom traverse a region of severely imbalanced radiation pressure from the negative-order and incident beams before entering the capture region. Figure 5.4b shows 20 simulated atomic trajectories within the potential of the same trap configuration of Figure 5.3. The atoms are initialized with a velocity v_x=10 m/s and enter at different distances from the grating. As well as a maximum capture velocity, side-loaded grating MOTs have a minimum capture velocity, below which atoms will be deflected before reaching the balanced region. The capture range of the trap configuration used here is approximately 9-25 m/s.

For the experiments described in Section 5.4.5, we used a side loading geometry, mounting the grating on a linear translation feedthrough to allow the grating position to be adjusted relative to the atomic beam.

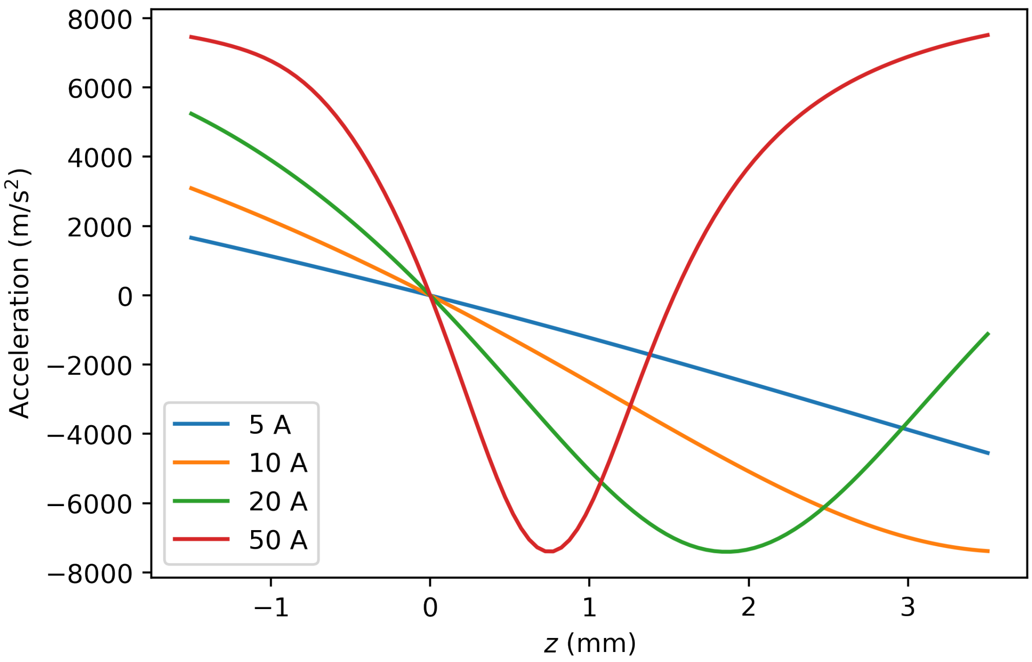

(a) Restoring force along the $z-$axis as a function of position and magnetic field.

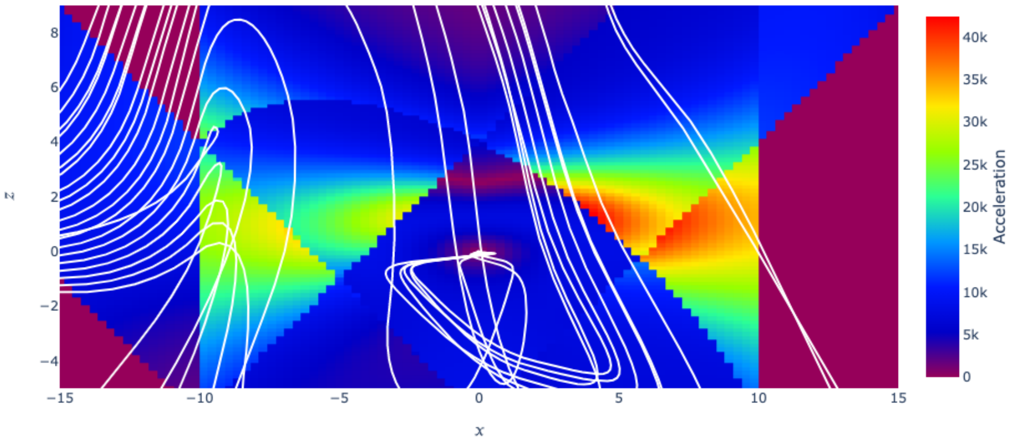

(b) Trajectories of atoms entering the grating MOT trapping potential at different distances from the grating.

Figure 5.4: Spatial effects in grating MOTs.

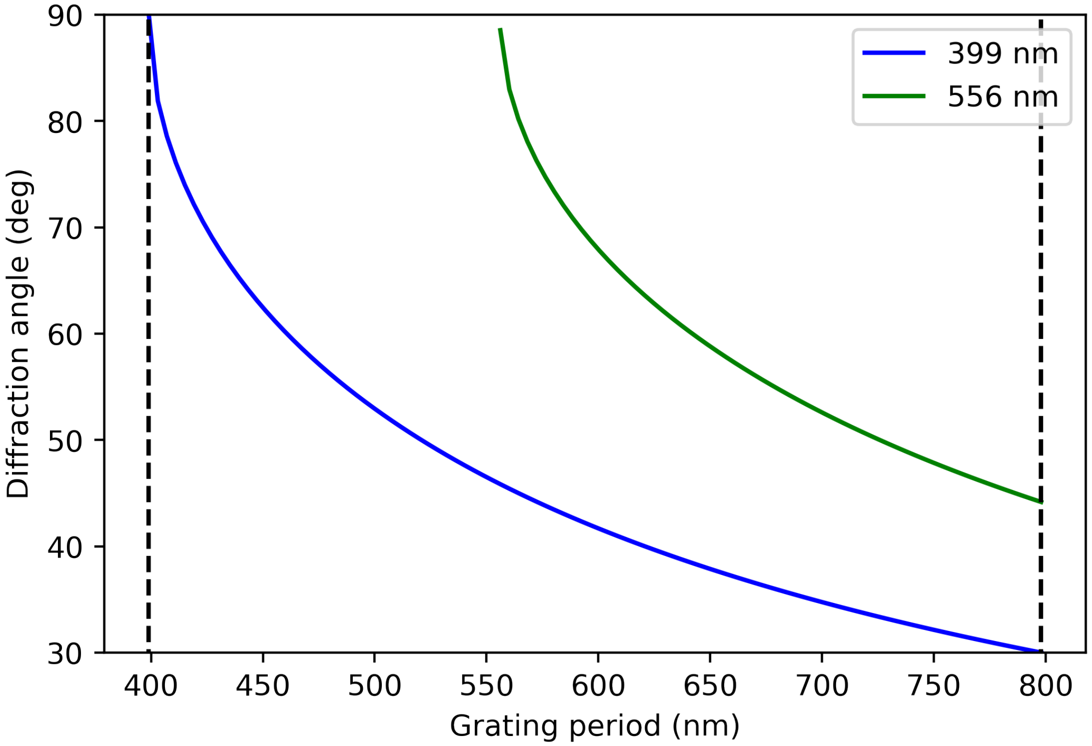

Figure 5.5: Diffraction angles for 399 nm and 556 nm light as a function of grating period.

5.4.4 Grating fabrication

We collaborated with gMOT pioneers at the University of Strathclyde to design, fabricate, and characterize several gratings for a two-color magneto-optical trap for ^{171}Yb. The first design criterion is the diffraction angle α, given by the Bragg condition

$$d \sin \alpha = m\lambda \tag{5.13}$$

where d is the grating period, m is the diffraction order, and λ is the wavelength (Figure 5.5). The viable range of grating periods is 556 nm $\leq d \leq$ 798 nm; the lower bound is necessary for first-order diffraction at both wavelengths, while the upper bound prevents second-order diffraction.

The second design consideration is the first-order diffraction efficiency, which should be 1/3 for a triplet grating to achieve force balance (equation (5.1)). While the grating period determines the angle, the grating duty cycle can be tuned freely to control the diffraction efficiency.

Lastly, the etch depth can be tuned to interferometrically reduce the zeroth-order diffraction, which can compromise the vertical force balance. This order can be perfectly eliminated for a single color by etching to $\lambda/4$ depth, or some tradeoff can be chosen to accommodate both colors simultaneously.

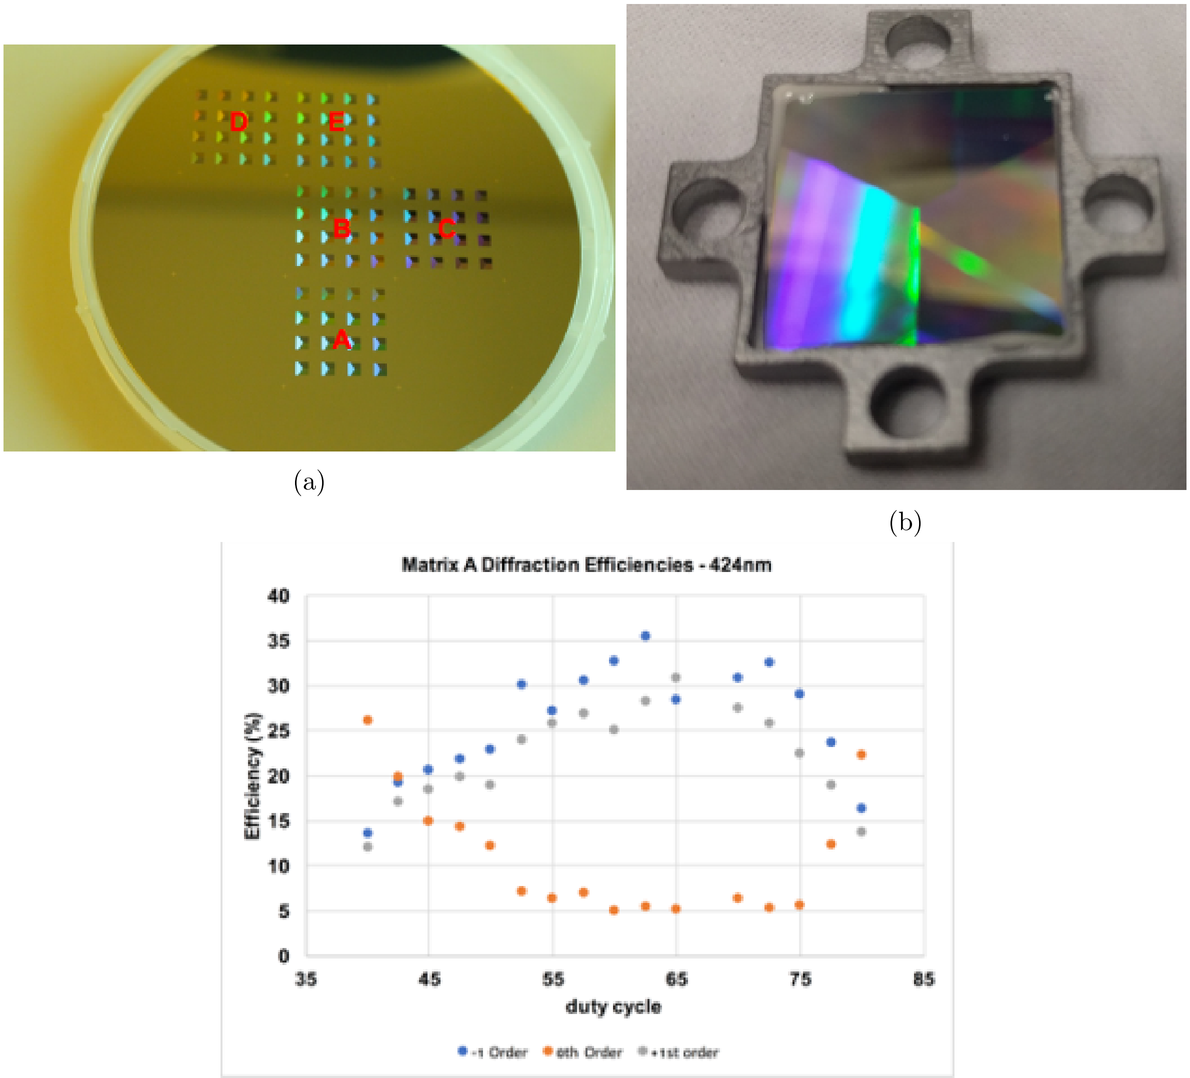

(a)

(b)

(c)

Figure 5.6: Grating fabrication and characterization. a) Dose test with 80 gratings etched on a single wafer to explore different design parameters. b) One of the final gratings fabricated for the experiment, mounted in a molybdenum adapter plate to minimize stress from thermal expansion during the vacuum bakeout. c) Characterization of diffraction efficiency vs. duty cycle at 424 nm.

5.4.5 Grating MOT experiment

We built a test apparatus to evaluate the effectiveness of magneto-optical trapping using the gratings designed in the previous section. Atoms were sourced from a compact effusive oven similar to the one currently used in the transportable clock. Trapping and slowing light was generated and delivered using a prototype of the 399 nm module in Section 2.2. To test the atom source and laser systems in isolation, we first realized a conventional six-beam magneto-optical trap, trapping 1.7×10^7 atoms with a loading rate of 2.8×10^7 s^{-1}.

To load atoms into the grating trap, we kept two of the counter-propagating beam pairs from the six-beam MOT in the plane parallel to the grating surface. These beams provided additional trapping force opposing the thermal beam, aiding with obtaining an initial trapping signal. After loading atoms into the assisted-gMOT, these transverse beams were extinguished to attempt to confine atoms entirely with the grating MOT. We observed loading with as little as 1 mW of transverse assistance, but the grating MOT exhibited a high loss rate which prevented effective trapping (Figure 5.7). We suspect that the transverse trapping abilities of the grating are insufficient to load from the strongly anisotropic atomic beam.

We have planned several improvements for a second round of grating MOT experimentation, but due to fabrication delays, this will not begin until 2022. Fabrication of new gratings with larger diffraction angles, providing more transverse trapping force, is underway. To mitigate the effects of beam non-uniformity, we have purchased beam-shaping optics to transform input Gaussian beams into top-hat profiles. We have also purchased a new injection-locked laser system providing nearly 400 mW of power at 399 nm (a factor of 4 improvement), which will enable the use of larger trapping beams.

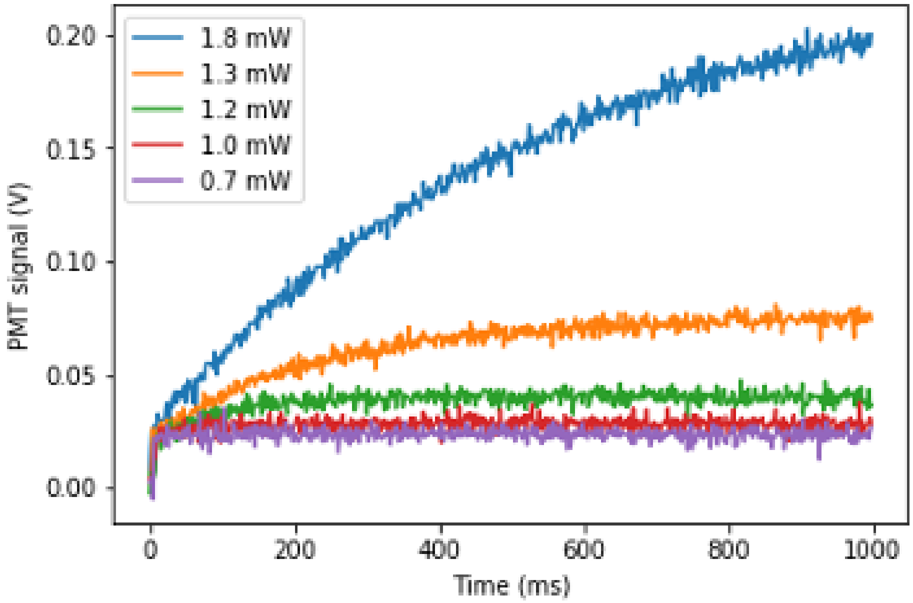

(a)

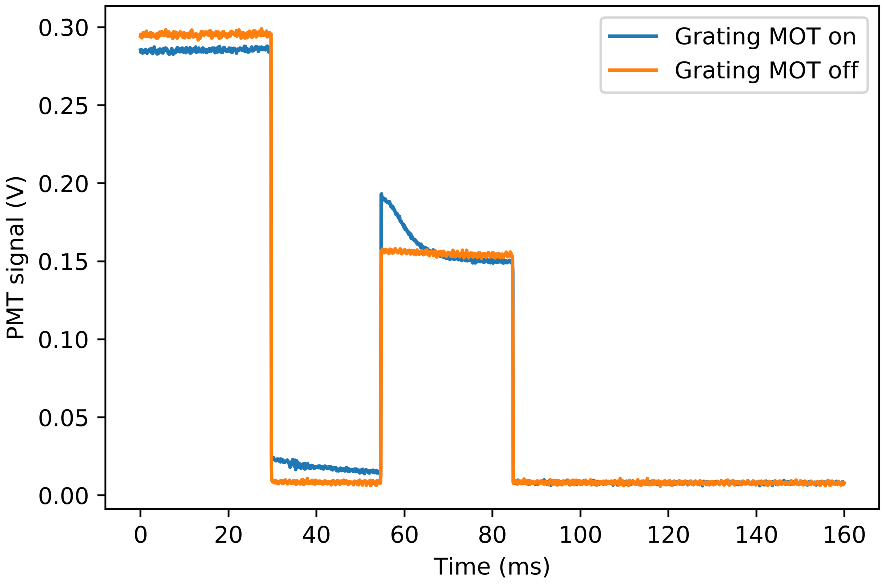

(b)

Figure 5.7: Grating MOT experimental results. a) Loading curves obtained by measurement of trapping beam fluorescence with a photomultiplier tube (PMT), plotted for different amounts of power in the transverse assistance beams. b) Transverse assistance beams are extinguished around 25 ms and the atoms are probed at around 55 ms to assess the capture ability of the gMOT.

5.5 Dispenser atom source

Upgrading from a six-beam MOT to a grating MOT would substantially reduce the volume of free-space optics surrounding the chamber. However, the vacuum chamber itself dominates the footprint of the physics package, with nearly half of the chamber occupied by the atom source and associated components: the effusive oven itself, the ion pump, and differential pumping region. This could be easily miniaturized by switching to a dispenser source. Since dispensers operate at a lower temperature than ovens, they can be positioned close to the science chamber without a need for differential pumping. Early experimentation by our group with a Yb dispenser has demonstrated successful second-stage trapping while maintaining extremely low pressure, enabling long trap lifetimes of several seconds; however, more work will be required to integrate the dispenser into our transportable chamber.

5.6 Conclusion

Transportable clocks will likely be required to facilitate international frequency comparisons required for redefinition of the optical second, as well as enabling numerous other applications such as relativistic geodesy, optical clock networks, and improved satellite navigation. In this dissertation, I have presented design and development efforts of the NIST Yb transportable optical lattice clock, up to and including spectroscopy of the clock transition and closed-loop clock operation. Over the next year, stress testing and complete rack integration will be carried out to prepare for field deployment. Finally, the clock will be deployed and operated in various metrologically-interesting locations, including a high-altitude measurement >4000 m above sea level and clock comparisons at other metrological institutes. State-of-the-art timekeeping outside the lab will enable many applications, including relativistic geodesy, searches for dark matter, next generations of very-long-baseline interferometry and global navigation satellite systems, and yet-unimagined new frontiers.

Bibliography

[1] Kazunori Akiyama, Antxon Alberdi, Walter Alef, Keiichi Asada, Rebecca Azulay, Anne-Kathrin Baczko, David Ball, Mislav Baloković, John Barrett, Dan Bintley, et al. First m87 Event Horizon Telescope results. ii. Array and instrumentation. The Astrophysical Journal Letters, 875(1):L2, 2019.

[2] James Roger Prior Angel and PGH Sandars. The hyperfine structure Stark effect i. Theory. Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences, 305(1480):125–138, 1968.

[3] LP Bakker, JM Freriks, and GMW Kroesen. A new ASE filter: the 20-fold prism monochromator. Measurement Science and Technology, 10(3):L25, 1999.

[4] Adam J Barker, Harry Style, Kathrin Luksch, Shinichi Sunami, David Garrick, Felix Hill, Christopher J Foot, and Elliot Bentine. Applying machine learning optimization methods to the production of a quantum gas. Machine Learning: Science and Technology, 1(1):015007, 2020.

[5] Andreas Bauch. Caesium atomic clocks: function, performance and applications. Measurement Science and Technology, 14(8):1159, 2003.

[6] K Beloy, N Hinkley, Nate B Phillips, Jeffrey A Sherman, Marco Schioppo, J Lehman, A Feldman, Leonard M Hanssen, Christopher W Oates, and Andrew D Ludlow. Atomic clock with 1×10^{-18} room-temperature blackbody Stark uncertainty. Physical review letters, 113(26):260801, 2014.

[7] Kyle Beloy, Xiaogang Zhang, William F McGrew, Nathan Hinkley, Tai Hyun Yoon, Daniele Nicolodi, Robert J Fasano, Stefan A Schaeffer, Roger C Brown, and Andrew D Ludlow. Faraday-shielded dc Stark-shift-free optical lattice clock. Physical Review Letters, 120(18):183201, 2018.

[8] Tom Bienaimé, Giovanni Barontini, Laure Mercier de Lépinay, Louis Bellando, Julien Chabé, and Robin Kaiser. Fast compression of a cold atomic cloud using a blue-detuned crossed dipole trap. Physical Review A, 86(5):053412, 2012.

[9] S Blatt, Jan Westenkær Thomsen, GK Campbell, AD Ludlow, MD Swallows, MJ Martin, MM Boyd, and Jun Ye. Rabi spectroscopy and excitation inhomogeneity in a one-dimensional optical lattice clock. Physical Review A, 80(5):052703, 2009.

[10] Benjamin Jacob Bloom. Building a Better Atomic Clock. PhD thesis, University of Colorado, 2014.

[11] BJ Bloom, TL Nicholson, JR Williams, SL Campbell, M Bishof, X Zhang, W Zhang, SL Bromley, and J Ye. An optical lattice clock with accuracy and stability at the 10^{-18} level. Nature, 506(7486):71–75, 2014.

[12] Martha I Bodine, Jean-Daniel Deschênes, Isaac H Khader, William C Swann, Holly Leopardi, Kyle Beloy, Tobias Bothwell, Samuel M Brewer, Sarah L Bromley, Jwo-Sy Chen, et al. Optical atomic clock comparison through turbulent air. Physical Review Research, 2(3):033395, 2020.

[13] V Bolpasi and W Von Klitzing. Double-pass tapered amplifier diode laser with an output power of 1 W for an injection power of only 200 μW. Review of Scientific Instruments, 81(11):113108, 2010.

[14] Tobias Bothwell, Dhruv Kedar, Eric Oelker, John M Robinson, Sarah L Bromley, Weston L Tew, Jun Ye, and Colin J Kennedy. Jila Sri optical lattice clock with uncertainty of 2.0×10^{-18}. Metrologia, 56(6):065004, 2019.

[15] Samuel M Brewer, J-S Chen, Aaron M Hankin, Ethan R Clements, Chin-wen Chou, David J Wineland, David B Hume, and David R Leibrandt. ^{27}Al+ quantum-logic clock with a systematic uncertainty below 10^-18. Physical Review Letters, 123(3):033201, 2019.

[16] Roger C Brown, Nate B Phillips, K Beloy, William F McGrew, Marco Schioppo, Robert J Fasano, Gianmaria Milani, Xiaogang Zhang, N Hinkley, H Leopardi, et al. Hyperpolarizability and operational magic wavelength in an optical lattice clock. Physical review letters, 119(25):253001, 2017.

[17] Luigi Cacciapuoti and Ch Salomon. Space clocks and fundamental tests: The aces experiment. The European Physical Journal Special Topics, 172(1):57–68, 2009.

[18] Lisheng Chen, John L Hall, Jun Ye, Tao Yang, Erjun Zang, and Tianchu Li. Vibration-induced elastic deformation of Fabry-Perot cavities. Physical Review A, 74(5):053801, 2006.

[19] Igor V Ciapurin, Leonid B Glebov, and Vadim I Smirnov. Modeling of Gaussian beam diffraction on volume Bragg gratings in PTR glass. In Practical Holography XIX: Materials and Applications, volume 5742, pages 183–194. International Society for Optics and Photonics, 2005.

[20] Cecilia Clivati, Roberto Aiello, Giuseppe Bianco, Claudio Bortolotti, Paolo De Natale, Valentina Di Sarno, Pasquale Maddaloni, Giuseppe Maccaferri, Alberto Mura, Monia Negusini, et al. Common-clock very long baseline interferometry using a coherent optical fiber link. Optica, 7(8):1031–1037, 2020.

[21] Garrett D Cole, Wei Zhang, Michael J Martin, Jun Ye, and Markus Aspelmeyer. Tenfold reduction of Brownian noise in high-reflectivity optical coatings. Nature Photonics, 7(8):644–650, 2013.

[22] Boulder Atomic Clock Optical Network BACON Collaboration. Frequency ratio measurements at 18-digit accuracy using an optical clock network. Nature, 591(7851):564–569, 2021.

[23] COMSOL, Inc. COMSOL Multiphysics Reference Manual, version 5.3.

[24] Andrei Derevianko and Maxim Pospelov. Hunting for topological dark matter with atomic clocks. Nature Physics, 10(12):933–936, 2014.

[25] Bureau International des Poids et Mesures. Si Brochure: The International System of Units. https://www.bipm.org/en/publications/si-brochure, 2019.

[26] R Dumke, M Volk, T Müther, FBJ Buchkremer, G Birkl, and Wolfgang Ertmer. Micro-optical realization of arrays of selectively addressable dipole traps: a scalable configuration for quantum computation with atomic qubits. Physical review letters, 89(9):097903, 2002.

[27] Jennifer L Ellis, Martha I Bodine, William C Swann, Sarah A Stevenson, Emily D Caldwell, Laura C Sinclair, Nathan R Newbury, and Jean-Daniel Deschênes. Scaling up Frequency-Comb-Based Optical Time Transfer to Long Terrestrial Distances. Physical Review Applied, 15(3):034002, 2021.

[28] RJ Fasano, YJ Chen, WF McGrew, WJ Brand, RW Fox, and AD Ludlow. Characterization and suppression of background light shifts in an optical lattice clock. Physical Review Applied, 15(4):044016, 2021.

[29] AC Fischer-Cripps. The Hertzian contact surface. Journal of materials science, 34(1):129–137, 1999.

[30] Peter I Frazier. A tutorial on Bayesian optimization. arXiv preprint arXiv:1807.02811, 2018.

[31] Serge Galliou, Maxim Goryachev, Philippe Abbe, Xavier Vacheret, Michael E Tobar, and Roger Bourquin. Quality factor measurements of various types of quartz crystal resonators operating near 4 k. IEEE transactions on ultrasonics, ferroelectrics, and frequency control, 63(7):975–980, 2015.

[32] Gabriele Giorgi, Tobias D Schmidt, Christian Trainotti, Ramon Mata-Calvo, Christian Fuchs, Mohammed Mainul Hoque, Jens Berdermann, Johann Furthner, Christoph Günther, Thilo Schuldt, et al. Advanced technologies for satellite navigation and geodesy. Advances in Space Research, 64(6):1256–1273, 2019.

[33] G Girard. The third periodic verification of national prototypes of the kilogram (1988-1992). Metrologia, 31(4):317, 1994.

[34] Alexei L Glebov, Oleksiy Mokhun, Alexandra Rapaport, Sébastien Vergnole, Vadim Smirnov, and Leonid B Glebov. Volume Bragg gratings as ultra-narrow and multiband optical filters. In Micro-Optics 2012, volume 8428, page 84280C. International Society for Optics and Photonics, 2012.

[35] Jacopo Grotti, Silvio Koller, Stefan Vogt, Sebastian Häfner, Uwe Sterr, Christian Lisdat, Heiner Denker, Christian Voigt, Ludger Timmen, Antoine Rolland, et al. Geodesy and metrology with a transportable optical clock. Nature Physics, 14(5):437–441, 2018.

[36] J Guéna, S Weyers, Michel Abgrall, C Grebing, V Gerginov, P Rosenbusch, Sebastien Bize, B Lipphardt, H Denker, N Quintin, et al. First international comparison of fountain primary frequency standards via a long distance optical fiber link. Metrologia, 54(3):348, 2017.

[37] Norman Gürlebeck, Lisa Wörner, Thilo Schuldt, Klaus Döringshoff, Konstantin Gaul, Domenico Gerardi, Arne Grenzebach, Nandan Jha, Evgeny Kovalchuk, Andreas Resch, et al. Boost: A satellite mission to test Lorentz invariance using high-performance optical frequency references. Physical Review D, 97(12):124051, 2018.

[38] Leonid I Gurvits. Space vlbi: from first ideas to operational missions. Advances in Space Research, 65(2):868–876, 2020.

[39] Sebastian Häfner, Sofia Herbers, Stefan Vogt, Christian Lisdat, and Uwe Sterr. Transportable interrogation laser system with an instability of mod σ y= 3× 10^{-16}. Optics Express, 28(11):16407–16416, 2020.

[40] Jonas E Hellstrom, Björn Jacobsson, Valdas Pasiskevicius, and Fredrik Laurell. Finite beams in reflective volume Bragg gratings: theory and experiments. IEEE Journal of Quantum Electronics, 44(1):81–89, 2007.

[41] Richard Hobson, William Bowden, Alvise Vianello, Ian R Hill, and Patrick Gill. Cavity-enhanced non-destructive detection of atoms for an optical lattice clock. Optics express, 27(26):37099–37110, 2019.

[42] Richard Hobson, William Bowden, Alvise Vianello, Alissa Silva, Charles FA Baynham, Helen S Margolis, Patrick EG Baird, Patrick Gill, and Ian R Hill. A strontium optical lattice clock with 1× 10^{-17} uncertainty and measurement of its absolute frequency. Metrologia, 57(6):065026, 2020.

[43] U Hugentobler, M Plattner, S Bedrich, M Heinze, V Klein, and D Voithenleitner. Optical clocks in future global navigation satellites. In Proceedings of the 2nd International Colloquium on Scientific and Fundamental Aspects of the Galileo Programme, 2009.

[44] N Huntemann, C Sanner, B Lipphardt, Chr Tamm, and E Peik. Single-ion atomic clock with 3× 10^{-18} systematic uncertainty. Physical Review Letters, 116(6):063001, 2016.

[45] Bernardo Jaduszliwer and James Camparo. Past, present and future of atomic clocks for gnss. GPS Solutions, 25(1):1–13, 2021.

[46] YY Jiang, AD Ludlow, Nathan D Lemke, Richard W Fox, Jeffrey A Sherman, L-S Ma, and Christopher W Oates. Making optical atomic clocks more stable with 10^{-16}-level laser stabilization. Nature Photonics, 5(3):158, 2011.

[47] Zhiheng Jiang and Gérard Petit. Combination of twstft and gnss for accurate utc time transfer. Metrologia, 46(3):305, 2009.

[48] Grzegorz Kasprowicz, Paweł Kulik, Michal Gaska, Tomasz Przywozki, Krzysztof Pozniak, Jakub Jarosinski, Joseph W Britton, Thomas Harty, Chris Balance, Weida Zhang, et al. Artiq and Sinara: Open Software and Hardware Stacks for Quantum Physics. In Quantum 2.0, pages QTu8B–14. Optical Society of America, 2020.

[49] Hidetoshi Katori, Masao Takamoto, VG Pal'Chikov, and VD Ovsiannikov. Ultrastable optical clock with neutral atoms in an engineered light shift trap. Physical Review Letters, 91(17):173005, 2003.

[50] Huidong Kim, Myoung-Sun Heo, Won-Kyu Lee, Chang Yong Park, Hyun-Gue Hong, Sang-Wook Hwang, and Dai-Hyuk Yu. Improved absolute frequency measurement of the ^{171}yb optical lattice clock at kriss relative to the si second. Japanese Journal of Applied Physics, 56(5):050302, 2017.

[51] May E Kim, William F McGrew, Nicholas V Nardelli, Ethan R Clements, Youssef S Hassan, Xiaogang Zhang, Jose L Valencia, Holly Leopardi, David B Hume, Tara M Fortier, et al. Optical coherence between atomic species at the second scale: improved clock comparisons via differential spectroscopy. arXiv preprint arXiv:2109.09540, 2021.

[52] Shimon Kolkowitz, Igor Pikovski, Nicholas Langellier, Mikhail D Lukin, Ronald L Walsworth, and Jun Ye. Gravitational wave detection with optical lattice atomic clocks. Physical Review D, 94(12):124043, 2016.

[53] SB Koller, J Grotti, A Al-Masoudi, S Dörscher, S Häfner, U Sterr, Ch Lisdat, et al. Transportable Optical Lattice Clock with 7$\times$ 10^{-17} Uncertainty. Physical Review Letters, 118(7):073601, 2017.

[54] Richard P Kornfeld, Bradford W Arnold, Michael A Gross, Neil T Dahya, William M Klipstein, Peter F Gath, and Srinivas Bettadpur. Grace-fo: the gravity recovery and climate experiment follow-on mission. Journal of Spacecraft and Rockets, 56(3):931–951, 2019.

[55] R Le Targat, Luca Lorini, Y Le Coq, M Zawada, J Guéna, M Abgrall, M Gurov, P Rosenbusch, DG Rovera, B Nagórny, et al. Experimental realization of an optical second with strontium lattice clocks. Nature Communications, 4:2109, 2013.

[56] Rodolphe Le Targat, Rafal Gartman, Luca Lorini, Bartłomiej Nagórny, Mikhail Gurov, Pierre Lemonde, Michal Zawada, and Jérôme Lodewyck. Comparison of two strontium optical lattice clocks in agreement at the 10^{-16} level. In 2012 IEEE International Frequency Control Symposium Proceedings, pages 1–4. IEEE, 2012.

[57] Thomas Legero, Thomas Kessler, and Uwe Sterr. Tuning the thermal expansion properties of optical reference cavities with fused silica mirrors. JOSA B, 27(5):914–919, 2010.

[58] John H Lehman, Mauricio Terrones, Elisabeth Mansfield, Katherine E Hurst, and Vincent Meunier. Evaluating the characteristics of multiwall carbon nanotubes. Carbon, 49(8):2581–2602, 2011.

[59] David R Leibrandt, James C Bergquist, and Till Rosenband. Cavity-stabilized laser with acceleration sensitivity below 10^{-12} g^{-1}. Physical Review A, 87(2):023829, 2013.

[60] David R Leibrandt, Michael J Thorpe, Mark Notcutt, Robert E Drullinger, Till Rosenband, and James C Bergquist. Spherical reference cavities for frequency stabilization of lasers in non-laboratory environments. Optics Express, 19(4):3471–3482, 2011.

[61] Nathan D Lemke. Optical lattice clock with spin-1/2 ytterbium atoms. University of Colorado at Boulder, 2012.

[62] Jun Liu, Tao Liu, Long Chen, Linbo Zhang, Guanjun Xu, Dongdong Jiao, and Shougang Zhang. A compact sub-hertz linewidth Fabry Perot cavity frequency stabilized laser for space application. Optics & Laser Technology, 136:106777, 2021.

[63] Jérôme Lodewyck, Philip G Westergaard, and Pierre Lemonde. Nondestructive measurement of the transition probability in a sr optical lattice clock. Physical Review A, 79(6):061401, 2009.

[64] Andrew D Ludlow, Martin M Boyd, Jun Ye, Ekkehard Peik, and Piet O Schmidt. Optical atomic clocks. Reviews of Modern Physics, 87(2):637, 2015.

[65] J Lumeau, LB Glebov, and V Smirnov. Tunable narrowband filter based on a combination of Fabry-Perot etalon and volume Bragg grating. Optics Letters, 31(16):2417–2419, 2006.

[66] Long-Sheng Ma, Peter Jungner, Jun Ye, and John L Hall. Accurate cancellation (to millihertz levels) of optical phase noise due to vibration or insertion phase in fiber transmitted light. In Laser Frequency Stabilization and Noise Reduction, volume 2378, pages 165–175. International Society for Optics and Photonics, 1995.

[67] Robert Magnusson and SS Wang. New principle for optical filters. Applied Physics Letters, 61(9):1022–1024, 1992.

[68] Jack B Marling, J Nilsen, LC West, and LL Wood. An ultrahigh-Q isotropically sensitive optical filter employing atomic resonance transitions. Journal of Applied Physics, 50(2):610–614, 1979.

[69] WF McGrew, X Zhang, RJ Fasano, SA Schäffer, K Beloy, D Nicolodi, RC Brown, N Hinkley, G Milani, M Schioppo, et al. Atomic clock performance enabling geodesy below the centimetre level. Nature, 564(7734):87–90, 2018.

[70] William F McGrew, Xiaogang Zhang, H Leopardi, RJ Fasano, Daniele Nicolodi, Kyle Beloy, Jian Yao, Jeffrey A Sherman, Stefan A Schaeffer, Joshua Savory, et al. Towards the optical second: verifying optical clocks at the SI limit. Optica, 6(4):448–454, 2019.

[71] Harold J Metcalf and Peter van der Straten. Laser cooling and trapping of atoms. JOSA B, 20(5):887–908, 2003.

[72] Oliver Montenbruck, Peter Steigenberger, and André Hauschild. Multi-gnss signal-in-space range error assessment–Methodology and results. Advances in Space Research, 61(12):3020–3038, 2018.

[73] Oliver Montenbruck, Peter Steigenberger, and André Hauschild. Comparing the 'big 4'- A User's View on gnss Performance. In 2020 IEEE/ION Position, Location and Navigation Symposium (PLANS), pages 407–418. IEEE, 2020.

[74] Nils-Axel Mörner. Eustasy and geoid changes. The Journal of Geology, 84(2):123–151, 1976.

[75] Jürgen Müller, Dominic Dirkx, Sergei M Kopeikin, Guillaume Lion, I Panet, G Petit, and PNAM Visser. High performance clocks and gravity field determination. Space Science Reviews, 214(1):1–31, 2018.

[76] Nils Nemitz, Asbjørn Arvad Jørgensen, Ryotatsu Yanagimoto, Filippo Bregolin, and Hidetoshi Katori. Modeling light shifts in optical lattice clocks. Physical Review A, 99(3):033424, 2019.

[77] Matthew A Norcia, Aaron W Young, William J Eckner, Eric Oelker, Jun Ye, and Adam M Kaufman. Seconds-scale coherence on an optical clock transition in a tweezer array. Science, 366(6461):93–97, 2019.

[78] Mark Notcutt, Long-Sheng Ma, Jun Ye, and John L Hall. Simple and compact 1-Hz laser system via an improved mounting configuration of a reference cavity. Optics Letters, 30(14):1815–1817, 2005.

[79] Robert A Nyman, Gaël Varoquaux, Brice Villier, Delphine Sacchet, Frédéric Moron, Yann Le Coq, Alain Aspect, and Philippe Bouyer. Tapered-amplified antireflection-coated laser diodes for potassium and rubidium atomic-physics experiments. Review of Scientific Instruments, 77(3):033105, 2006.

[80] Noriaki Ohmae, Masao Takamoto, Yosuke Takahashi, Motohide Kokubun, Kuniya Araki, Andrew Hinton, Ichiro Ushijima, Takashi Muramatsu, Tetsuo Furumiya, Yuya Sakai, et al. Transportable Strontium Optical Lattice Clocks Operated Outside Laboratory at the Level of 10^{-18} Uncertainty. Advanced Quantum Technologies, page 2100015, 2021.

[81] Donald M Olsson and Lloyd S Nelson. The Nelder-Mead simplex procedure for function minimization. Technometrics, 17(1):45–51, 1975.

[82] Daniel Ott, Marc SeGall, Ivan Divliansky, George Venus, and Leonid Glebov. High-contrast filtering by multipass diffraction between paired volume Bragg gratings. Applied Optics, 54(31):9065–9070, 2015.

[83] Marco Pomponio, Archita Hati, and Craig Nelson. Fpga-based Low-latency Digital Servo for Optical Physics Experiments. In 2020 Joint Conference of the IEEE International Frequency Control Symposium and International Symposium on Applications of Ferroelectrics (IFCS-ISAF), pages 1–2. IEEE, 2020.

[84] Sergey G Porsev and Andrei Derevianko. Multipolar theory of blackbody radiation shift of atomic energy levels and its implications for optical lattice clocks. Physical Review A, 74(2):020502, 2006.

[85] Sergey G Porsev, Andrei Derevianko, and EN Fortson. Possibility of an optical clock using the 6^1s_0ß$6^3p_0^o$ transition in ^{171,173}yb atoms held in an optical lattice. Physical Review A, 69(2):021403, 2004.

[86] Audrey Quessada, Richard P Kovacich, Irène Courtillot, André Clairon, Giorgio Santarelli, and Pierre Lemonde. The Dick effect for an optical frequency standard. Journal of Optics B: Quantum and Semiclassical Optics, 5(2):S150, 2003.

[87] Fritz Riehle, Patrick Gill, Felicitas Arias, and Lennart Robertsson. The cipm list of recommended frequency standard values: guidelines and procedures. Metrologia, 55(2):188, 2018.

[88] Sebastian Ruder. An overview of gradient descent optimization algorithms. arXiv preprint arXiv:1609.04747, 2016.

[89] S Schiller, GM Tino, P Lemonde, U Sterr, Ch Lisdat, A Görlitz, N Poli, A Nevsky, and Ch Salomon. The space optical clocks project. In International Conference on Space Optics—ICSO 2010, volume 10565, page 1056531. International Society for Optics and Photonics, 2017.

[90] Marco Schioppo, Roger C Brown, William F McGrew, Nathan Hinkley, Robert J Fasano, Kyle Beloy, TH Yoon, Gianmaria Milani, D Nicolodi, JA Sherman, et al. Ultrastable optical clock with two cold-atom ensembles. Nature Photonics, 11(1):48–52, 2017.

[91] Stefan Schröder, Simon Stellmer, and Jürgen Kusche. Potential and scientific requirements of optical clock networks for validating satellite-derived time-variable gravity data. Geophysical Journal International, 226(2):764–779, 2021.

[92] Thilo Schuldt, Klaus Döringshoff, Alexander Milke, Josep Sanjuan, Martin Gohlke, Evgeny V Kovalchuk, Norman Gürlebeck, Achim Peters, and Claus Braxmaier. High-performance optical frequency references for space. In Journal of Physics: Conference Series, volume 723, page 012047. IOP Publishing, 2016.

[93] Ramamurti Shankar. Principles of Quantum Mechanics. Springer Science & Business Media, 2012.

[94] Jeffrey A Sherman, Nathan D Lemke, Nathan Hinkley, Marco Pizzocaro, Richard W Fox, Andrew D Ludlow, and Christopher W Oates. High-accuracy measurement of atomic polarizability in an optical lattice clock. Physical Review Letters, 108(15):153002, 2012.

[95] Fujio Shimizu, Kazuko Shimizu, and Hiroshi Takuma. Four-beam laser trap of neutral atoms. Opt. Lett., 16(5):339–341, Mar 1991.

[96] AV Taichenachev, VI Yudin, VD Ovsiannikov, VG Pal'Chikov, and Christopher W Oates. Frequency shifts in an optical lattice clock due to magnetic-dipole and electric-quadrupole transitions. Physical Review Letters, 101(19):193601, 2008.

[97] Peter Guthrie Tait and William Kelvin Thomson. Elements of Natural Philosophy. Cambridge University Press, 1879.

[98] Masao Takamoto, Ichiro Ushijima, Noriaki Ohmae, Toshihiro Yahagi, Kensuke Kokado, Hisaaki Shinkai, and Hidetoshi Katori. Test of general relativity by a pair of transportable optical lattice clocks. Nature Photonics, pages 1–5, 2020.

[99] ME Tamisiea, JX Mitrovica, GA Milne, and JL Davis. Global geoid and sea level changes due to present-day ice mass fluctuations. Journal of Geophysical Research: Solid Earth, 106(B12):30849–30863, 2001.

[100] Byron D Tapley, S Bettadpur, Mo Watkins, and Ch Reigber. The gravity recovery and climate experiment: Mission overview and early results. Geophysical Research Letters, 31(9), 2004.

[101] James I Thorpe, K Numata, and J Livas. Laser frequency stabilization and control through offset sideband locking to optical cavities. Optics Express, 16(20):15980–15990, 2008.

[102] Ichiro Ushijima, Masao Takamoto, Manoj Das, Takuya Ohkubo, and Hidetoshi Katori. Cryogenic optical lattice clocks. Nature Photonics, 9(3):185–189, 2015.

[103] Matthieu Vangeleyn, Paul F Griffin, Erling Riis, and Aidan S Arnold. Single-laser, one beam, tetrahedral magneto-optical trap. Optics Express, 17(16):13601–13608, 2009.

[104] Stefan Vogt, Sebastian Häfner, Jacopo Grotti, Silvio Koller, Ali Al-Masoudi, Uwe Sterr, and Christian Lisdat. A transportable optical lattice clock. In Journal of Physics: Conference Series, volume 723, page 012020. IOP Publishing, 2016.

[105] D Voigt, EC Schilder, RJC Spreeuw, and HB Van Linden Van Den Heuvell. Characterization of a high-power tapered semiconductor amplifier system. Applied Physics B, 72(3):279–284, 2001.

[106] Kateryna Vyshenska. How to Provide Structural Stability in Thermal Expansion Simulations.

[107] P Wcisło, P Ablewski, K Beloy, S Bilicki, M Bober, R Brown, R Fasano, R Ciuryło, H Hachisu, T Ido, et al. New bounds on dark matter coupling from a global network of optical atomic clocks. Science Advances, 4(12):eaau4869, 2018.

[108] Stephen Webster and Patrick Gill. Force-insensitive optical cavity. Optics letters, 36(18):3572–3574, 2011.

[109] Wang Wen-Li, Ye Jie, Jiang Hai-Ling, Bi Zhi-Yi, Ma Long-Sheng, and Xu Xin-Ye. Frequency stabilization of a 399-nm laser by modulation transfer spectroscopy in an ytterbium hollow cathode lamp. Chinese Physics B, 20(1):013201, 2011.

[110] Eugen Wiens and Stephan Schiller. Simulation of force-insensitive optical cavities in cubic spacers. Applied Physics B, 124(7):1–17, 2018.

[111] Paul B Wigley, Patrick J Everitt, Anton van den Hengel, John W Bastian, Mahasen A Sooriyabandara, Gordon D McDonald, Kyle S Hardman, Ciaron D Quinlivan, P Manju, Carlos CN Kuhn, et al. Fast machine-learning online optimization of ultra-cold-atom experiments. Scientific Reports, 6(1):1–6, 2016.

[112] Jian Yao, Jeff A Sherman, Tara Fortier, Holly Leopardi, Thomas Parker, William McGrew, Xiaogang Zhang, Daniele Nicolodi, Robert Fasano, Stefan Schäffer, et al. Optical-clock-based time scale. Physical Review Applied, 12(4):044069, 2019.

[113] Aaron W Young, William J Eckner, William R Milner, Dhruv Kedar, Matthew A Norcia, Eric Oelker, Nathan Schine, Jun Ye, and Adam M Kaufman. Half-minute-scale atomic coherence and high relative stability in a tweezer clock. Nature, 588(7838):408–413, 2020.

[114] Guang-Yao Zhao, Juan Carlos Algaba, Sang Sung Lee, Taehyun Jung, Richard Dodson, Maria Rioja, Do-Young Byun, Jeffrey Hodgson, Sincheol Kang, Dae-Won Kim, et al. The power of simultaneous multi-frequency observations for mm-vlbi: beyond frequency phase transfer. The Astronomical Journal, 155(1):26, 2017.

[115] Chao Zhou, Sachin Barthwal, Wendong Zhang, Chuan He, Biao Tang, Lin Zhou, Jin Wang, and Ming-Sheng Zhan. Characterization and optimization of a tapered amplifier by its spectra through a long multi-pass rubidium absorption cell. Applied optics, 57(26):7427–7434, 2018.