Precision Frequency Metrology with Optical Lattice Clocks



William Bowden

Jesus College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Contents

1 Introduction **4**

2 Principle of Operation of an Optical Lattice Clock **13**

2.1 Laser cooling of strontium 13

2.2 Lamb-Dicke spectroscopy 16

2.3 Trapping in a magic-wavelength lattice 20

2.3.1 Atomic polarisability calculation 22

2.4 Motional sideband thermometry 24

3 Sr1: The First Generation Strontium Lattice Clock **29**

3.1 Experimental apparatus 30

3.1.1 Clock laser distribution and universal synthesiser 31

3.1.2 Transfer cavity 34

3.1.3 Atomic servo-control 35

3.2 Stability of Sr1 40

3.3 Uncertainty budget for Sr1 49

3.3.1 Blackbody radiation shift 51

3.3.1.1 Theoretical treatment of the BBR shift 51

3.3.1.2 Experimental treatment of the BBR shift 53

3.3.2 Lattice shifts 56

3.3.2.1 Higher-order lattice shifts 58

3.3.3 Density shift 60

3.3.4 Background gas collisional shift 64

3.3.5 Quadratic DC-Stark shift 65

3.3.5.1 Standard methods for DC-Stark shift evaluation . . . 66

3.3.5.2 Singlet Rydberg EIT spectroscopy for electrometry . 68

3.3.5.3 Rydberg theory: Single electron model and Stark maps 68

3.3.5.4 Experimental setup for Rydberg electrometry 69

3.3.5.5 EIT spectroscopy of Rydberg states . . . 73
3.3.6 Zeeman shifts . . . 78
3.3.7 Probe shift . . . 80
3.3.8 Doppler shift . . . 81
3.3.9 Servo offset . . . 84
3.4 Conclusion . . . 87

4 Sr2: The Second Generation Strontium Lattice Clock 89
4.1 Experimental apparatus . . . 90
4.1.1 Pyramid MOT . . . 92
4.1.2 In-vacuum cavities . . . 95
4.1.2.1 DC-Stark shift considerations . . . 98
4.1.2.2 Cavity aided non-demolition atom detection . . . 100
4.1.2.3 Flexure mount . . . 105
4.1.2.4 Laser stabilisation to in-vacuum cavities . . . 108
4.1.3 Zeeman slower . . . 113
4.1.4 MOT and compensation coils . . . 117
4.1.5 Laser systems . . . 119
4.1.5.1 461 nm laser system . . . 119
4.1.5.2 2.92 µm system for a triplet, magneto-optical trap . . . 120
4.1.5.3 497 nm repump laser . . . 122
4.1.6 Magneto-optical trap light delivery and imaging . . . 124
4.1.7 Cavity and clock light delivery . . . 127
4.1.7.1 Path length stabilisation for Sr2 . . . 129
4.1.7.2 Non-demolition detection laser system . . . 130
4.2 Validating the performance of Sr2 . . . 136
4.2.1 First stage cooling . . . 136
4.2.2 Magnetic trap lifetime . . . 138
4.2.3 Second-stage cooling and state preparation . . . 139
4.2.4 Cavity enhanced lattice . . . 145
4.2.4.1 Cavity-enhanced 2D lattice . . . 146
4.2.4.2 Cavity read out of the atomic state population . . . 148
4.2.5 Clock operation . . . 150
4.2.6 DC-Stark shift cancellation . . . 151
4.3 Conclusion . . . 153

5 Future Plans 156

A My contribution to the project **159**

B Design of the Vacuum system **162**

C List of Publications **166**

Bibliography **168**

List of Figures

1.1 Topology of an atomic clock 6
1.2 Working principle of a frequency comb 9

2.1 Term diagram for strontium showing the relevant transitions 14
2.2 Outline of the experimental cycle 17
2.3 Excitation spectrum of the clock transition for a 350 ms Rabi pulse . 21
2.4 Polarisability of the two clock states showing the position of the magic-wavelengths 24
2.5 Sideband spectrum used for thermometry 28

3.1 Rendering of the vacuum chamber showing the relevant design features 31
3.2 Diagram of the clock laser optical setup including distribution 32
3.3 Diagram of the transfer cavity setup 35
3.4 Measurement protocol used to derive the error used to lock the clock laser to the atomic transition 36
3.5 Illustration of how a frequency ramp leads to a steady-state offset . . 40
3.6 Scaling of the Allan deviations with averaging time for different noise processes 42
3.7 Allan deviation of interleaved clock comparison 45
3.8 Illustration of the Dick effect 47
3.9 Time trace of the estimated BBR shift of clock transition based on the temperature of the science chamber 56
3.10 Loss rate measurement of atoms held in the magic wavelength lattice used to estimate the frequency shift induced by background gas collisions 65
3.11 Characterising the DC-Stark Shift 67
3.12 Stark map of the 5s75d $^1D_2$ Rydberg state 69
3.13 Diagram of the experimental setup used for the Rydberg EIT spectroscopy 71

3.14 Time trace showing EIT signal on the probe beam induced by 461 nm probe beam . . . 73
3.15 Example of a Rydberg EIT spectrum . . . 74
3.16 Splitting of Rydberg manifolds induced by applied electric and magnetic fields . . . 77
3.17 Contour plots of χ^2 used to estimate fitting error . . . 78
3.18 Measurement of the quadratic Zeeman shift of the clock transition in ^{87}Sr . . . 80
3.19 Diagram of the optical setup used to stabilise the delivery path length of the clock light to the atoms . . . 83
3.20 Phase chirp of the probe AOM arising from RF induced heating . . . 84
3.21 Dependence of frequency chirp of the probe AOM on RF drive power 85
3.22 Analysis of the residual offset in the lock of the laser to the atomic transition based on the excitation fraction noise . . . 86

4.1 Photo of Sr2 and surrounding optics . . . 91
4.2 Diagram showing the working principle of a pyramid MOT . . . 93
4.3 Rendering and photo of the pyramid MOT structure and in-vacuum cavities . . . 94
4.4 Photo of the v-groove jig used to secure cavity mirror during gluing . 98
4.5 Configuration of the voltages applied to the in-vacuum electrodes used to generate the three orthogonal electric fields . . . 99
4.6 Outline of the non-demolition detection scheme . . . 104
4.7 Rendering of cavity flexure . . . 107
4.8 Alignment of the cavity mode overlap of the in-vacuum build up cavities 108
4.9 Diagram of the locking scheme used to stabilise the lattice laser to the in-vacuum cavities . . . 113
4.10 Rendering of the Zeeman slower and slowing beam prism periscope . 115
4.11 Magnetic field profile of the Zeeman slower field produced by the permanent magnets . . . 116
4.12 Residual magnetic fields of the Zeeman slower at the trap centre with and without compensation magnets . . . 117
4.13 Rendering of the MOT and compensation coils . . . 118
4.14 Frequency map of the 461 nm laser system . . . 120
4.15 Optical setup used to deliver the MOT and repump light to the chamber 126
4.16 Optical setup used to deliver the light to the Sr2 in-vacuum cavities . 128

4.17 Optical path length stabilisation scheme used in Sr2 130
4.18 Overview of non-demolition atom detection scheme 132
4.19 Diagram of the high speed Mach-Zehnder interferometric modulator 134
4.20 Relative detuning between the Sr2 cavity and the probe beam with and without feed-forward corrections from the lattice laser lock 135
4.21 Loading curve for Sr2 pyramid MOT 137
4.22 Measured trap lifetime of magnetically confined metastable strontium 138
4.23 Cooling sequence for the metastable triplet MOT 140
4.24 Time of flight measurement of triplet MOT 142
4.25 Triplet MOT lifetime 144
4.26 Heating rate and lifetime of atom held in the cavity enhanced lattice 146
4.27 Verifying overlap of cavities modes as needed for the 2D lattice 148
4.28 Measurement of the atom number using the in-vacuum cavity 149
4.29 Stability of the frequency difference between Sr1 & 2 when probed synchronously 151
4.30 Allan deviation of the Stark shift induced by the in-vacuum electrodes 152

B.1 Rendering of vacuum system 163

ASE amplified spontaneous emission

AOM acousto-optic modulator

BBR blackbody radiation

CF ConFlat

DAC digital-to-analogue converter

DAQ digital acquisition system

DDS direct digital synthesiser

ECDL extended cavity diode laser

EDM electrical discharge machining

EIT electromagnetically induced transparency

EOM electro-optic modulator

FSR free spectral range

LBO lithium triborate

LO local oscillator

LLI local Lorentz invariance

LPI local position invariance

MOT magneto-optical trap

MZI Mach-Zehnder interferometer

NDD non-demolition detection

NEG non-evaporable getter

NPL National Physical Laboratory

PDH Pound-Drever-Hall

PLS path length stabilisation

PMT photomultiplier tube

PPLN periodically-poled lithium niobate

QPN quantum projection noise

RGA residual gas analyzer

SHG second harmonic generation

TA tapered amplifier

TAI temps atomique international

UDP uniformly distributed probability

UHV ultra-high vacuum

ULE ultra-low expansion

VCO voltage controlled oscillator

Abstract

Strontium optical lattice clocks are proven to be excellent timekeepers. Their exceptional performance makes them a leading candidate for possible redefinition of the SI second and a valuable tool for searching for new physics beyond the standard model. This thesis reports on the development of strontium (Sr) optical lattice clocks at National Physical Laboratory (NPL). We will provide an overview of the principle of operation of such systems and various effects that limit their accuracy and stability. In particular, we will focus on our two strontium lattice clocks: Sr1, the work horse in our laboratory for several years now, and Sr2, the new addition to NPL's optical clock family. In regards to Sr1 we will provide an updated evaluation of the various sources of systematic uncertainty which limit its accuracy to 1×10^{-17}. In particular, we will outline a new method using precision spectroscopy of highly excited Rydberg states to measure residual electric fields which must be characterised as they can give rise to Stark shifts of the clock transition. Next we will outline the development of Sr2: NPL's next generation lattice clock, highlighting novel design features and cooling techniques such as: a compact pyramid MOT, in-vacuum buildup cavities used for the optical lattice and non-demolition detection of the atom number, and finally laser cooling using a metastable transition as 2.92 µm. Lastly, we will present preliminary results from Sr2 focusing on how we can leverage two lattice clocks to realize a single ultrastable frequency standard.

Chapter 1

Introduction

Fundamental to the scientific method – the process that structures our knowledge of the physical world – is the act of measurement. Scientists are constantly striving to develop new techniques or refine existing practices in the hopes such progress will enable new theories to be tested or old hypotheses dispelled. Through this continuous process of discovery, it is clear that our understanding of the natural world is underpinned by our ability to measure its properties.

Of all base units making up the International System of Units, scientists have succeeded in measuring the "second" with the highest level of accuracy. Starting with the first caesium referenced clock built by Louis Essen and Jack Parry at UK's NPL in 1955, atoms have underpinned our most accurate timekeepers. Much progress has been made since these first atomic clocks; currently the record-holder for relative accuracy is the strontium lattice clock developed at JILA which 'ticks' at a frequency characterized to 18 digits of accuracy [101] [1]. To put this level of accuracy into perspective, such a clock would run for 15 billion years before accumulating a timing error of one second. Atomic clocks are at the vanguard of a emerging class of qunatum technologies, which now includes accelerometers, gyroscopes and magnetometers, that

[1] Recently, a preprint has been published on Arxiv from the ytterbium lattice clock group at NIST demonstrating agreement between two clocks both with fractional uncertainties of 1.4×10^{-18} [90]

harness in the intrinsic properties of quantum systems for precision measurement. Through scientific and engineering advancements, our improved ability to control and manipulate these complex systems will allow us to examine our world at a previously unachievable level of detail.

A clock is composed of three parts, a local oscillator (LO), a counter, and a reference that work in unison to track the passage of time. As the name implies, the LO produces a periodic signal that serves as the ticking of the clock. It can take on many forms: ranging from an electronic circuit in a wrist watch to a rotating neutron star used for a pulsar clock. Regardless of the choice of LO, the variation in the oscillation period needs to be small if one hopes to realise a stable clock. The role of the counter is to keep track of the passage of time by counting the number of the LO ticks. Finally comes the frequency reference to which the local oscillator is 'disciplined' by means of a feedback loop to ensure it ticks at the correct frequency. Ideally the reference frequency should be insensitive to external perturbation making it immune to changes in its environmental conditions, e.g. temperature, stray electric and magnetic fields, etc.

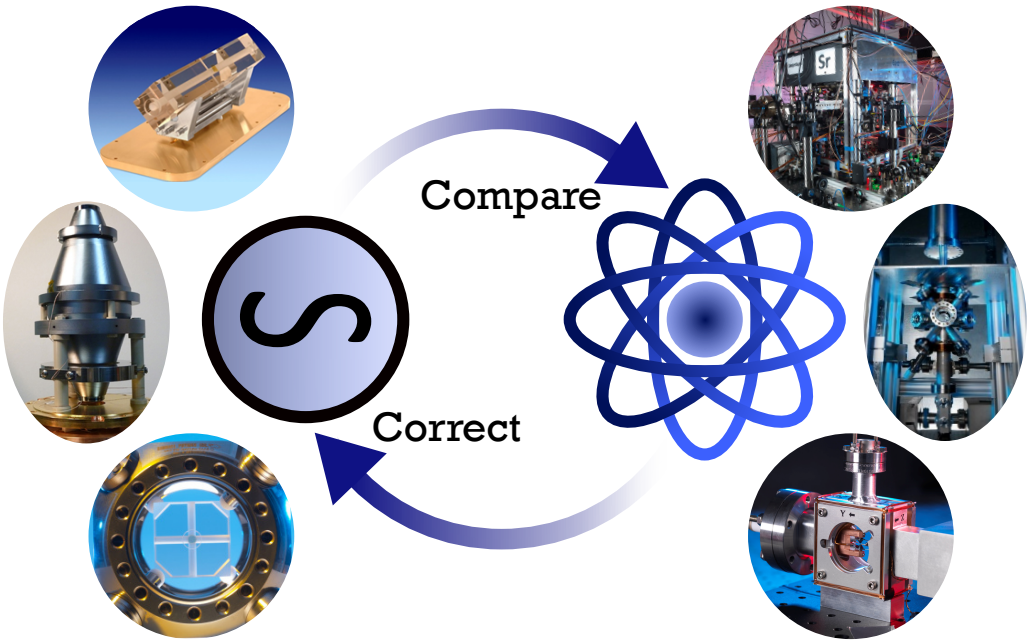


Figure 1.1: A few examples of local oscillators and atomic references that work together to realise stable and accurate frequency references. During operation, the LO is spectroscopically compared to the resonance frequency needed to drive a transition between two atomic levels within the reference atom. Based on the atomic response, any detuning of the LO from resonance is corrected in order to realise an accurate and stable frequency reference. Examples of LOs are shown on the left (from top to bottom): NPL's master oscillator is a 30 cm ULE cavity operating at 1064 nm; a cryogenic silicon cavity operating at 1550 nm which is held at a temperature of 124 K developed at the German metrology institute PTB – it holds the current record of short term stability at 4×10^{-17} [87]; NPL's novel cubic design that exploits the symmetry of the cavity by mounting it in such a way as to make the length insensitive to vibrations—making it well suited for transportable clocks and space applications [157]. On the right are several examples of atomic references, including (from top to bottom): NPL's lattice clock which is the focus of this thesis; NPL's primary standard caesium fountain; and finally another NPL's optical clocks based on a single ytterbium ion held in a endcap trap [104].

To understand the working principle of an atomic clock it is instructive to identify the parts the system corresponding to these three fundamental building blocks. For an atomic clock, the LO is a microwave source or laser which generates an electromagnetic field that is used to excite a transition in an atom. To stabilise the frequency of this field we rely on Planck's famous equation relating the frequency of a photon

to its energy:

$$E = hf \tag{1.1}$$

where h is Planck's constant. By performing precision spectroscopy, we can use this relation to find the resonance frequency of the LO which matches the energy spacing between the states in the atom. Therefore it is the energy levels of the atom which serve as the reference for the clock. The advantage of such a reference is that the transition frequency depend directly on the fundamental laws of physics. Furthermore, the indistinguishable nature of atoms ensures the reproducibility of the standard. An example of such a frequency reference is the hyperfine transition in caesium which serves as the current definition for the SI second, stated as:

> The second is the duration of 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium 133 atom.

After 60 years of progress since the first caesium beam clocks, state of the art cold atom fountains used to realise the SI second today report fractional frequency uncertainties at the 2×10^{-16} level [44, 74, 75]. However, in the past two decades a new class of clocks, operating using transitions in the optical domain, has emerged to surpass the performance of their microwave counterparts. The benefit of moving from the microwave to the optical domain is that the higher oscillation frequency makes for a intrinsically more stable and accurate clock. However, this transition is not without its challenges. For example, unlike microwave fields that can be counted using standard electronic circuitry, optical fields oscillate too fast for such detectors. This problem was solved with the realisation of the first broadband femtosecond frequency comb which provided a phase coherent connection between the RF and optical domains [58]. In the frequency domain, the spectrum of the pulsed femtosecond laser is

a collection of separated comb teeth that serve as a 'ruler' to which other frequencies can be compared. For the comb to work, the spacing between comb teeth ν_{rep} (set by the pulse repetition rate) and the start of ruler ν_{ceo} (set by phase shift of the optical field with respect to the carrier envelope from pulse to pulse) must be stabilised, see figure 1.2. Fortunately, these frequencies are designed to be in the RF domain and are easily stabilised using phase-locked loops. Once the ruler spacing and offset are fixed, the comb can be used to determine other optical frequencies by measuring a beat frequency between LO and the n^{th} comb mode using the following relation:

$$\nu_{\mathrm{LO}} = n \times \nu_{rep} + \nu_{ceo} + \nu_{\mathrm{beat}} \tag{1.2}$$

The absolute frequency of several optical transitions have now been measured, with the $5s^2\,{}^1S_0 \rightarrow 5s5p\,{}^3P_0$ in strontium being the most prevalent. As a divalent atom, strontium's electronic structure is separated into a series singlet (total spin S equal to zero) and triplet ($S = 1$) levels. Dipole induced transitions between singlet and triplet states are forbidden given they violate the selection rule $\Delta S = 0$ making the lowest lying triplet state remarkably long lived. In the case of strontium, this ultra-narrow transition allows us to reach remarkably high Q-factors for our clock. Furthermore, the $J = 0 \rightarrow J' = 0$ transition results in a reduced sensitivity to external perturbations such as stray electric and magnetic fields.

As the estimated uncertainty of optical standards surpass the current SI realisation of the second, we face the problem of how to characterise this new class of atomic clocks. Presently, the only way to verify the performance and reproducibility of such systems is through direct comparison against other optical clocks. This requires that links are established between different metrological institutes which are capable of transferring precision timing signals needed to compare optical clocks. This can be achieved using a microwave link distributed via satellite as was the case for the 2015 International Timescales with Optical Clocks (ITOC) campaign, but to achieve

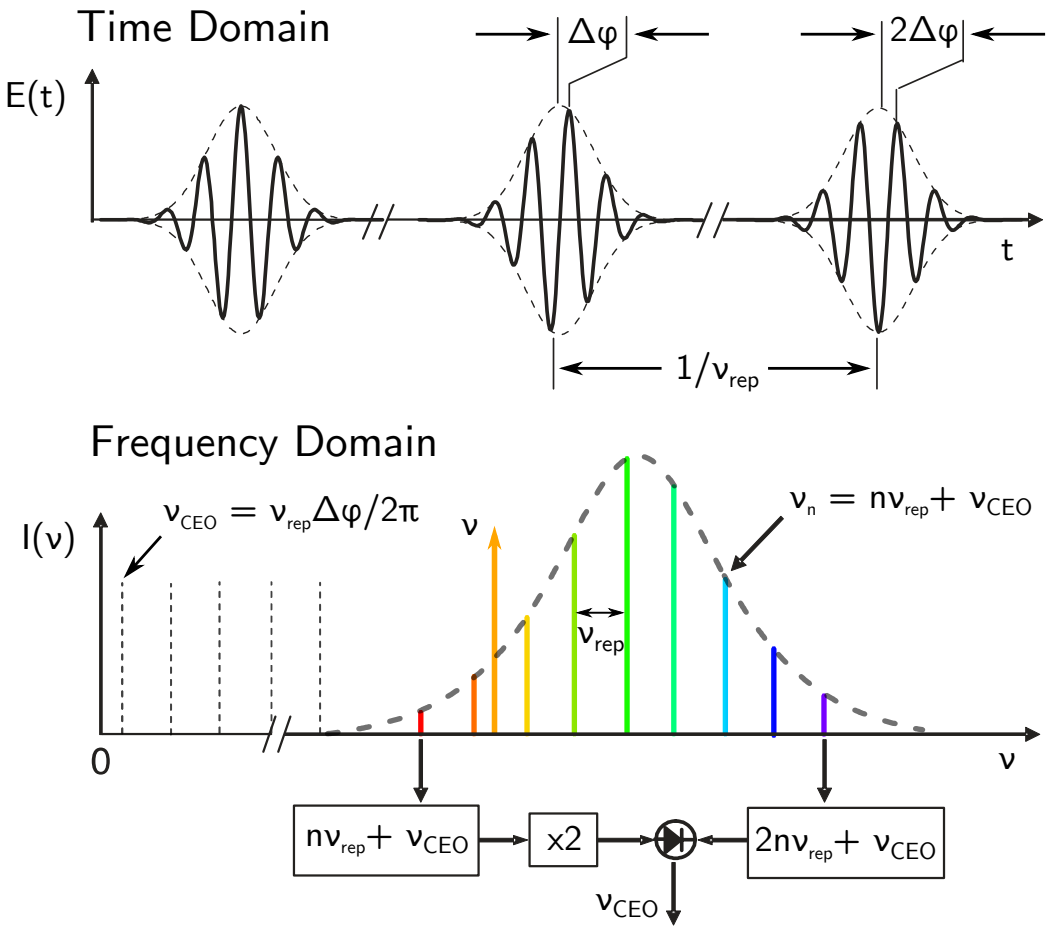


Figure 1.2: Diagram of the frequency comb showing the spectrum in both the time and frequency domain. The two relevant parameters that fully define the spectrum are the repartition rate ν_{rep} and position of the first comb tooth ν_{ceo}. If the comb spans a full octave, ν_{ceo} can be measured via optical heterodyne detection by frequency doubling the spectrum and then combining it with the original comb on a photodiode as shown in the figure. Once both ν_{rep} and ν_{ceo} are known, the comb can be used to measure other optical frequencies.

the best stability we must use optical signals distributed over a fibre network [114] as currently underway with the EMPIR project Optical frequency transfer – a European network (OFTEN).

Large scale, cross-continental comparisons of optical clocks are not only important for performance validation, but the unparalleled accuracy to which they are done also serve to test our fundamental theories of physics. In particular, two out of the three cornerstones of Einstein's equivalence principle – local Lorentz invariance (LLI) and local position invariance (LPI) – can be tested with atomic clocks.[2]. In a recent optical fibre comparison between the strontium lattice clocks at PTB (Germany), SYRTE (France) and NPL, LLI was tested by looking for daily variations in the frequency ratio between the clocks. In particular, the lack of any observed variation between the clocks constrains the Robertson–Mansouri–Sexl parameter β_{RMS} which accounts for the possibility of a preferred inertial reference frame in which the speed of light is isotropic. Accounting for sources of uncertainty and the instability of the clocks, the comparison data constrains $|\beta_{RMS}| < 1.1 \times 10^{-8}$, improving upon the previous best measurement by a factor of two [32].

Atomic clocks are also an excellent tool for testing LPI – where 'position' accounts for both invariance in both space and time. Depending on the type of transition, its resonance frequency will have different sensitivities to various fundamental constants such as the fine structure constant α or the mass ratio of the electron to the proton μ. Variation in the values of these constants, which is an example of an LPI violation, could be seen by looking for changes in the frequency ratio between clocks over time or space. Using optical clocks to hunt for possible variations of the fine structure α on time scales of hours to years provides complementary measurements to those made

[2]The third pillar is the weak equivalence principle (WEP) which states that inertial and gravitational mass are equivalent. Once again atom based sensors, in particular interferometry using ultracold atomic clouds provide some of the best constraints of possible violations of general relativity. The European space agency project STE-QUEST: Space-Time Explorer and QUantum Equivalence Principle Space Test which plans to launch a differential atom interferometer operating on the two rubidium isotopes hopes to test the weak equivalence principle at 1×10^{-15} level.

using the natural fission reactor located in the Oklo region of Gabon [30] (geological timescales) or astronomical detection using the absorption spectra of distant quasars [156] (Cosmological time scales). Taking the strontium clock transition frequency as an example, its sensitive to changes in α K_α, defined as:

$$K_\alpha = \frac{d}{d\alpha}\log \nu(\alpha), \tag{1.3}$$

is estimated to be a meager $K_\alpha = 0.006$ compared with the Yb^+ octopole transition which has $K_\alpha = -5.95$ [125]. However to detect changes in the value α, it is advantageous to choose two transitions which maximise the difference in K_α, making Sr and Yb^+ appealing candidates for such tests. So far, experiments looking for temporal variation of the fine structure constant using Yb^+ compared the the frequency ratio of its octupole to quadrupole ($K_\alpha = 1.03$) transitions. Both the results from references [42, 56], give upper bounds of $\dot{\alpha}/\alpha$ of $2 \times 10^{-17}\,\mathrm{yr}^{-1}$. An equally stringent limit for $\dot{\alpha}$ was measured by comparing the Hg^+ ($K_\alpha = -2.94$) and Al^+ ($K_\alpha = 0.008$) clocks over the course of year [121]. Because electronic transitions are insensitive to the electron-to-proton mass ratio μ, optical frequency ratios are not very effective at constraining $\dot{\mu}/\mu$; however $\dot{\mu}$ can be measured precisely by comparing optical clocks against microwave clocks. Once again looking at variations of the two clock transitions in Yb^+ against Cs fountains, references [42, 56] were able to set an upper bounds n $\dot{\mu}/\mu$ of $1 \times 10^{-16}\,\mathrm{yr}^{-1}$ level. Another approach for detecting changes in μ is to apply the precision spectroscopy techniques used for atoms to molecules. For rotational and vibrational transition, the energy scales as $E_{\mathrm{vib}} : E_{\mathrm{rot}} \propto \mu^{-1/2} : \mu^{-1}$. Furthermore the presence of hyperfine structure, Λ-doubling, hindered rotation, and other effects can further enhance their sensitivity to fundamental constants [26]. Another interesting approach would be to forego one of the clocks by replacing it with a cavity. The length of the cavity depends on the attractive electromagnetic force holding the bulk material together and therefore changes with α ($K_\alpha = 1$). This comparison could

be very precise over short timescales and is therefore of interest for theories predicting abrupt changes or fast oscillations of α. Such experiments could help to address many outstanding questions regarding the nature of dark matter or possible standard model extensions. For example, it has recently been proposed that ultralight scalar dark matter can couple to electromagnetic fields resulting in oscillations of the fine structure constant with an amplitude proportional to the dark matter density and frequency equal to the dark matter mass [7]. In an attempt to explain the dynamics of our nascent universe, it has been proposed that such dark matter particles should have masses in the range of 10^{-25} to 10×10^{-20} eV [53] corresponding to an oscillation period between 100 years to hours. Spectroscopy comparing the splitting of two near-degenerate levels for different isotopes dysprosium has so far resulted in a null result for oscillations of α in this mass range [154] and constrains possible linear variation to below $|\delta\alpha/\alpha| < 1 \times 10^{-16}$ yr^{-1} [68]. Another possible scalar field that is predicted to give rise to oscillation in α is the 'dilaton' which couples directly to matter. The dilaton appears in several theoretical frameworks such as the Kaluza–Klein extension and string theory which incorporate higher dimensions [150] in an attempt to unify fundamental forces. Extending on these ideas, a network of clocks could help map the spatial distribution of dark matter. In such an experiment one would look for fluctuations in the relative timing of remote clocks which would be attributed to variations in the local dark matter density [33, 155].

All of the applications mentioned so far require the comparison of accurate and stable clocks. The focus on this thesis is the development and characterisation of such frequency standards. Chapter 2 provides an overview of the working principle of strontium lattice clocks. Chapter 3 will focus of the characterization of Sr1, NPL's first generation lattice clock, including a full evaluation of all known sources of systematic uncertainty. Finally, chapter 4 focuses on applying the lessons learned from Sr1 to construct NPL's next generation lattice clock.

Chapter 2

Principle of Operation of an Optical Lattice Clock

2.1 Laser cooling of strontium

The internal level structure of strontium is divided into a series of singlet and triplet states. Figure 2.1 shows the internal level structure of strontium. For the clock, we use the fermionic isotope ^{87}Sr which has a nuclear spin of 9/2 leading to a complicated hyperfine structure not shown in figure 2.1.

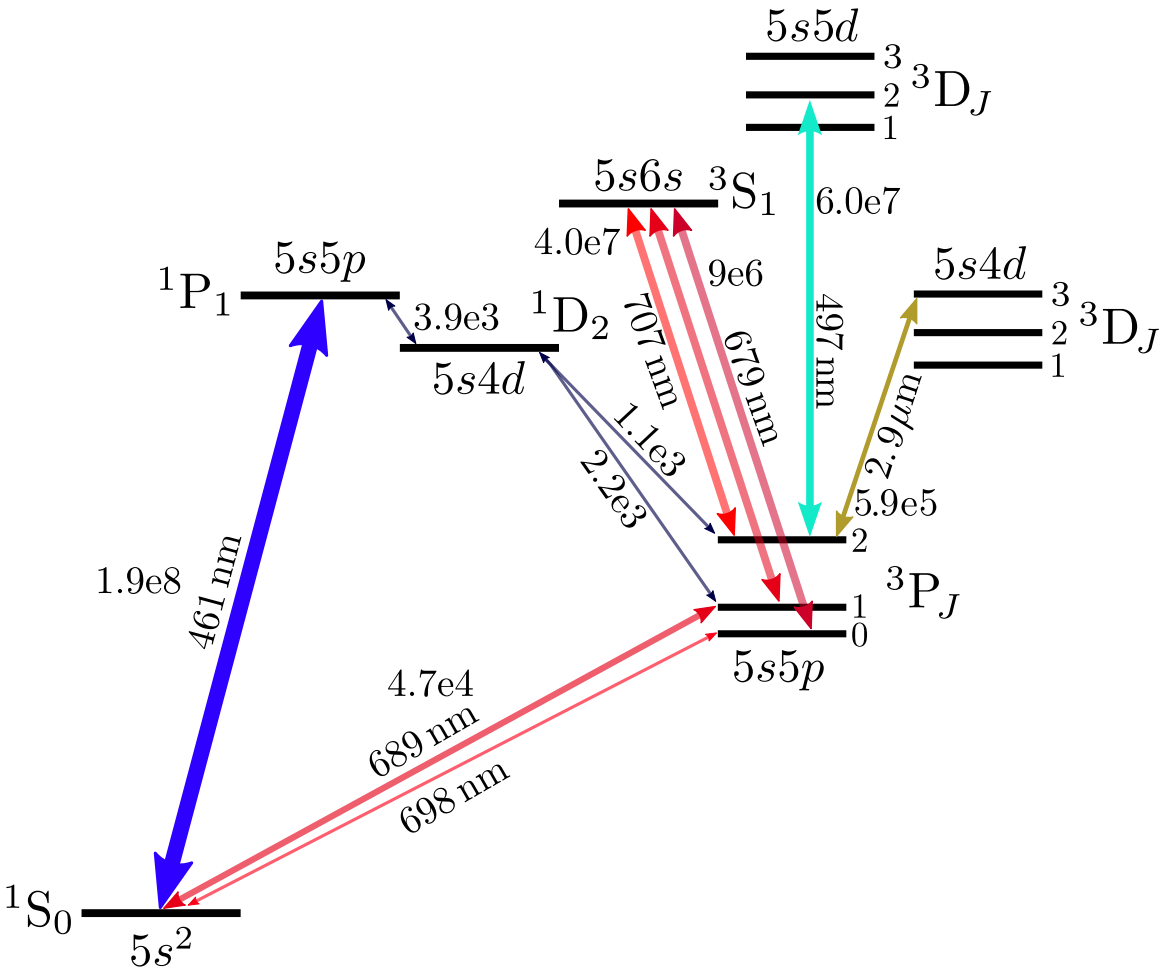


Figure 2.1: Term diagram for strontium showing the relevant atomic transitions. The ${}^1S_0 \rightarrow {}^1P_1$ and ${}^1S_0 \rightarrow {}^3P_1$ transitions at 461 nm and 689 nm are used for laser cooling. The $J = 0 \rightarrow J' = 0$ clock transition to which we lock our local oscillator is at 698 nm. During the 461 nm cooling cycle, atoms can get shelved into the ${}^3P_2$ state via the decay path through the ${}^1D_2$ state. Therefore, a repump laser at either 707 nm or 497 nm is needed to return atoms to the cooling transition. A repump laser at 679 nm is also needed to repump atoms from the upper clock state. Finally, an alternative second stage cooling option using a transition from the metastable ${}^3P_2$ and ${}^3D_3$ states resonant with 2.92 µm light is also highlighted. Included on the diagram are the decay rates of all relevant states in units of s^{-1}. Figure credit: Ian Hill

The techniques for laser cooling strontium and confining it in an optical lattice laser in order to perform precision spectroscopy are well established and the details can be found in literature with select examples being [163, 59, 134, 82, 15]. In summary, the experimental cycle is as follows. Starting with an effusive oven, strontium is heated

to produce a flux of atoms which are decelerated using a Zeeman slower [110] so that they can be trapped in the 'blue' magneto-optical trap (MOT). Both the MOT and Zeeman slower operate using 461 nm-light resonant with the $^1S_0 \rightarrow {}^1P_1$ transition. The transition is not fully closed as the $^1P_1$ state can decay to the $^1D_2$ state with a rate $\Gamma = 3.9 \times 10^3$ s^{-1} from which it populates the $^3P_2$ and $^3P_1$ states with branching ratios of 2/3 and 1/3, respectively. The $^3P_1$ state decays to the ground state, but the $^3P_2$ is metastable and must be repumped. We elect to excite atoms to the $5s4d\,{}^3D_2$ state using light at 497 nm [91] given its favorable branching ratio to the $^3P_1$ state, but there are several other possible transitions which can be used to return atoms to the cooling cycle [136, 96]. When efficiently repumped, the first stage MOT has a lifetime on the order of seconds.

The large scatter rate of the first stage cooling results in a Doppler limited temperature of a few millikelvin. Unfortunately this is too warm for efficient loading of the optical lattice given its trap depth of approximately 10 µK. Therefore, a second stage MOT operating on the $^1S_0 \rightarrow {}^3P_1$ transition is needed to further cool the atoms. Often referred to as the 'red' MOT given its wavelength is 689 nm, the small transition scattering rate of $2\pi \times 7$ kHz allows for a recoil limited temperature of 200 nK. The red MOT is operated in two steps. First a broadband stage during which the light is frequency modulated helps to ensure efficient transfer from the first stage cooling. Next, the modulation is switched off in order to reach the high phase space densities required for efficient loading into the optical lattice. Once in the lattice, a bias magnetic field is applied to separate the ten hyperfine levels via the Zeeman effect. The atoms are optically pumped to either of the $m_F = \pm 9/2$ stretched states by driving the $^1S_0$ $(F = 9/2) \rightarrow {}^3P_1$ $(F' = 9/2)$ transition using right– or left-handed circularly polarised light. Once the sample has been prepared in the correct state, the LO laser drives the clock transition. By alternating measurements of the two $m_F = \pm 9/2$ stretched states and averaging the result, the linear Zeeman shift is sup-

pressed. This entire cooling and spectroscopy cycle is outlined pictorially in figure 2.2 and is repeated indefinitely when operating the clock. Later in this chapter, I will focus on two specific aspects of this sequence: the optical lattice operating at the 'magic' wavelength and precision spectroscopy of tightly confined atoms.

2.2 Lamb-Dicke spectroscopy

Leaving the microwave domain to realise standards based on optical transitions was motivated by the observation that many perturbations which shift the clock states (Stark shift, Zeeman shift, collisions, etc.) result to smaller *fractional* uncertainties of the clock frequency. However, this is not the case for all systematics. Motional effects, such as the Doppler shift, scale linearly with frequency and can therefore significantly broaden and shift the clock transition. Even with the low temperature achieved from the second stage cooling, the residual motion of the atoms would prevent resolving Hz level spectroscopic features. For example at a 1 µK temperature, the atoms have a mean velocity on the order of $1\,\mathrm{cm}\cdot\mathrm{s}^{-1}$ and the resulting Doppler broadened FWHM σ_v of the transition is [17]:

$$\sigma_v = \sqrt{\frac{8k_bT\ln 2}{mc^2}}f_0 \cong 32\,\mathrm{kHz} \tag{2.1}$$

where k_b is the Boltzmann constant, m is the atomic mass, and c is the speed of light. Even at zero temperature, motional effects resulting from the momentum imparted by the absorbed photon during excitation can result in a significant shift of the clock frequency. This 'recoil' shift is given by:

$$\nu_{\mathbf{recoil}} = \frac{hk^2}{2m} \tag{2.2}$$

which for the strontium clock transition results in a shift of more than 4 kHz. If

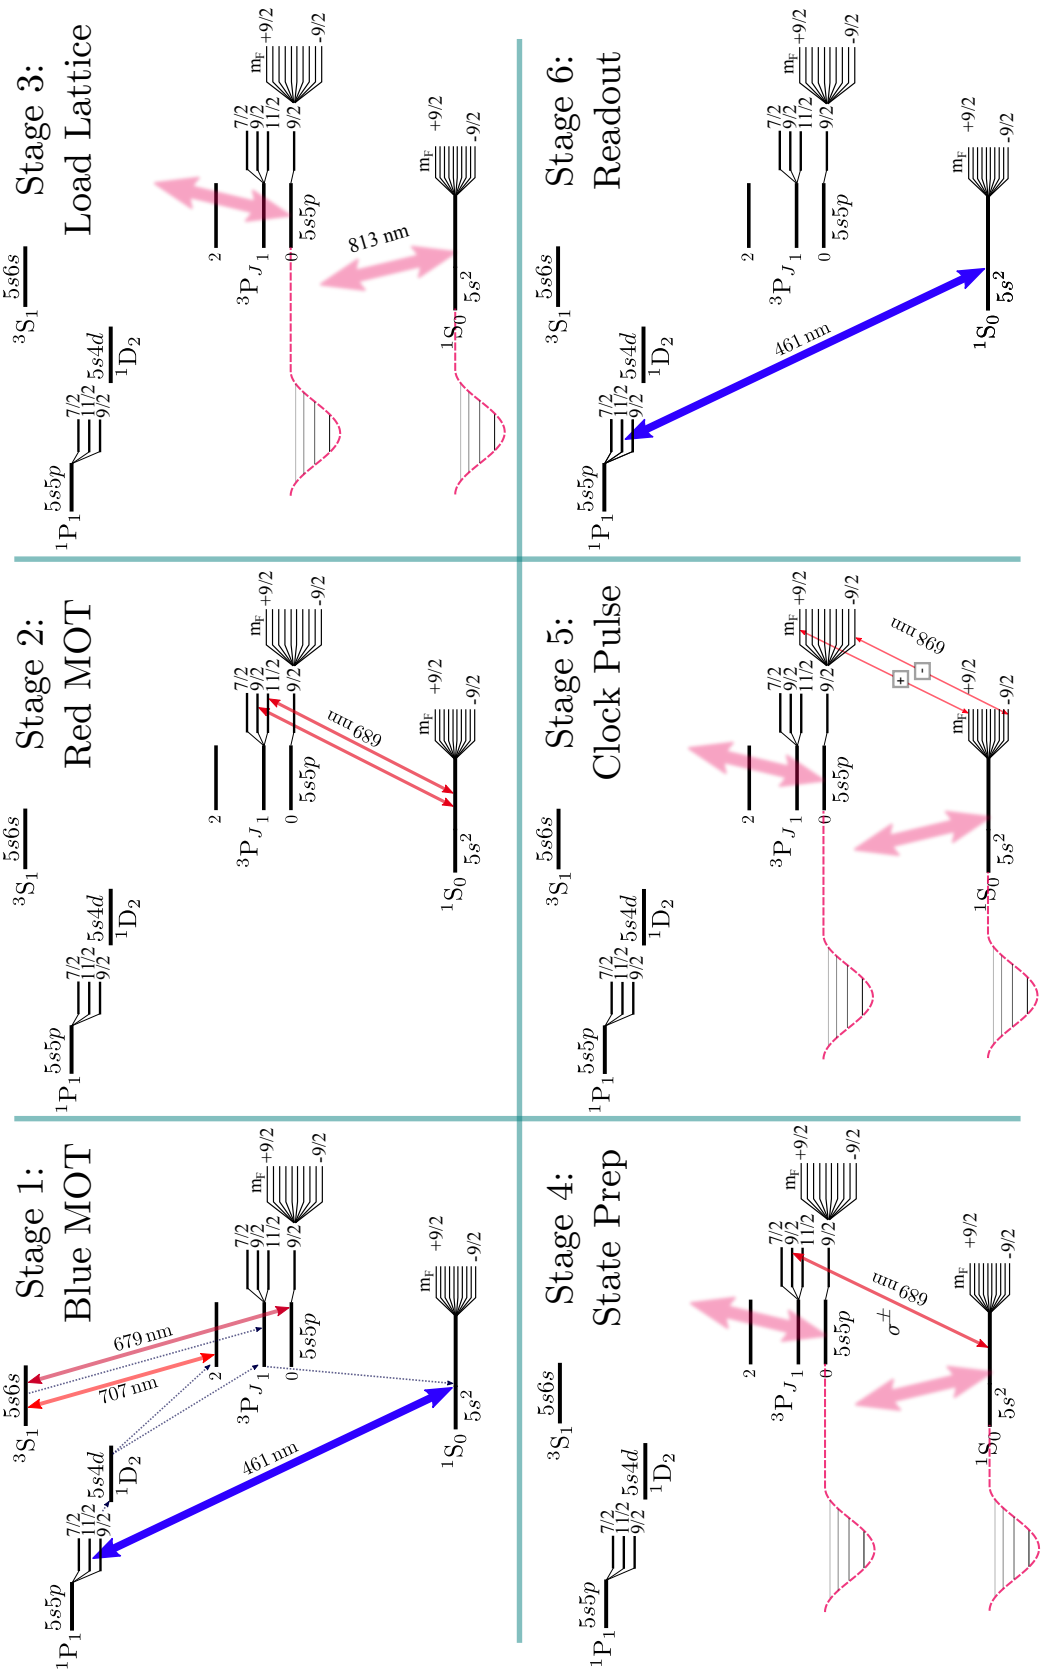


Figure 2.2: Outline of the experimental cycle showing the cooling stages, state preparation, and spectroscopy. Figure credit: Richard Hobson

unchecked, both of these motional effects – the Doppler broadening and recoil shift – would severely compromise the accuracy and stability of the clock. We suppress these effects, counter-intuitively, by *increasing* the motional energy of the atoms by tightly confining them in space. In such a regime, the dominant energy scale is set by the trapping potential rather than temperature or recoil energy. Specifically, when the trap frequency ν_z of the confining harmonic potential along the direction of the probe beam is larger than the recoil frequency we enter the so-called Lamb-Dicke regime [35] as indicated by the parameter η, defined as:

$$\eta = \sqrt{\frac{\nu_{\mathbf{recoil}}}{\nu_z}}, \tag{2.3}$$

being much less than one. Relating the trap frequency to the characteristic length $\Delta z \triangleq \sqrt{\frac{\hbar}{2m\omega_o}}$ of the harmonic oscillator, we see that reaching the Lamb-Dicke regime is equivalent to confining the atom to a region in space much smaller than the wavelength of the probing light. For optical clocks, where the wavelength is sub-micron, it is clear why reaching such a regime is challenging compared to microwave clocks where the wavelength of the probing field is a few centimeters.

To see how operating in the Lamb-Dicke regime helps to suppress motional effects, we must consider how the trapping potential modifies the excitation spectrum. In particular, the probe beam is capable of not only driving a transition between the ground and excited states of the atom (i.e. internal degrees of freedom), but can also excite the atoms to different motional states of the confining potential (external degrees of freedom). When the natural linewidth of the transition is much smaller than the motional trapping frequency ν_z, it is possible to spectroscopically resolve these 'motional' transitions which manifest themselves as distinct resonances separated from a central feature by integers of the trapping frequency. The central 'carrier' feature corresponds to a transition which leaves the motional degrees of freedom un-

changed, while the blue and red detuned 'sideband' resonances results from either increasing or decreasing by motional state by one quanta, respectively. For large motional frequencies these sideband features are well separated from the carrier transition. Furthermore, the sidebands are suppressed in amplitude as the Lamb-Dicke parameter is decreased. This can be seen by calculating the relative Rabi frequencies of the sideband transitions compared to the carrier Rabi frequency Ω_c which is given by [69]:

$$\begin{aligned} \frac{\Omega_{n\to n+1}}{\Omega_0} &= \eta\sqrt{n+1} \\ \frac{\Omega_{n\to n-1}}{\Omega_0} &= \eta\sqrt{n} \end{aligned} \tag{2.4}$$

Therefore, the main influence of the atom's motion on the clock frequency is through line-pulling of the carrier by asymmetric sidebands. This can be made negligible by increasing the trapping frequency to separate further the sidebands from the carrier while suppressing their relative strength. Aside from the asymmetric Rabi frequency, another important effect which changes the relative sideband amplitude is the relative atomic populations in the various motional states. The occupation probability of the various states is set by the ratio of the thermal energy to the trapping potential. In limit of $k_b T \ll \hbar\nu_z$, the atoms will predominantly reside in the lowest motional energy state. In the regime, the red sidebands are strongly suppressed as most of the atoms are incapable of decreasing their motional quantum number as they are already cooled to the ground state. We will see in section 2.4 we can use this effect to estimate the temperature of the sample.

2.3 Trapping in a magic-wavelength lattice

As seen in the previous section, it is necessary to confine the atoms in the Lamb-Dicke regime to avoid motional shifts arising from Doppler and recoil effects. This confinement is achieved by using a 1D optical lattice formed by two counter-propagating laser beams that trap the atoms via a position-dependent Stark shift. The resulting trapping potential $U(z, r, \omega)$ is set by the beam parameters (i.e. the waist ω_0, peak intensity I and k-vector k) and the atom's frequency dependent atomic polarisability $\alpha(\omega)$. It can be written as:

$$U(z, r, \omega) = \frac{1}{4} I_0 e^{-2r^2/\omega_0^2} \cos^2(kz)\alpha(\omega). \tag{2.5}$$

In the axial direction z along which the spectroscopy beam propagates, the atoms are tightly confined by the standing wave arising from the interference of the two beams. However, in the radial direction the trapping force is set by Gaussian profile of the beam and is much weaker. Approximating the trapping potential as harmonic in the axial and radial dimensions, we can deduce the trapping frequencies:

$$\begin{aligned} \nu_z &= \sqrt{\frac{2U_o k^2}{m}} \\ \nu_r &= \sqrt{\frac{4U_o}{m\omega_0^2}} \end{aligned} \tag{2.6}$$

For a typical operating trap depth of 12 µK, the induced Stark shift of the ground state is 250 kHz. Given that our targeted inaccuracy for the clock is less than 5 mHz, it would appear that we need to stabilise the laser intensity at the parts per billion level. Furthermore if the trapping potential is different for the ground and excited states, the resonance frequency to drive carrier transitions will be different depending on the initial state. It also has the practical drawback that the Stark shift is

inhomogeneous across the atomic sample resulting in a broadening of the transition linewidth. All these challenges would appear to render the lattice unsuitable for precision spectroscopy. However, with judicious choice of the trapping wavelength these problems can be mitigated. By choosing a 'magic'-wavelength for which the polarisabilities of the ground and excited states are identical, the differential Stark shift of the clock transition frequency is zero and independent of the lattice intensity. To find such a wavelength, we need to calculate the atomic polarisabilities of both states. Such a calculation will prove useful as the same differential polarisability will give rise to systematic frequency shifts induced by blackbody radiation (BBR) and residual DC electric fields.

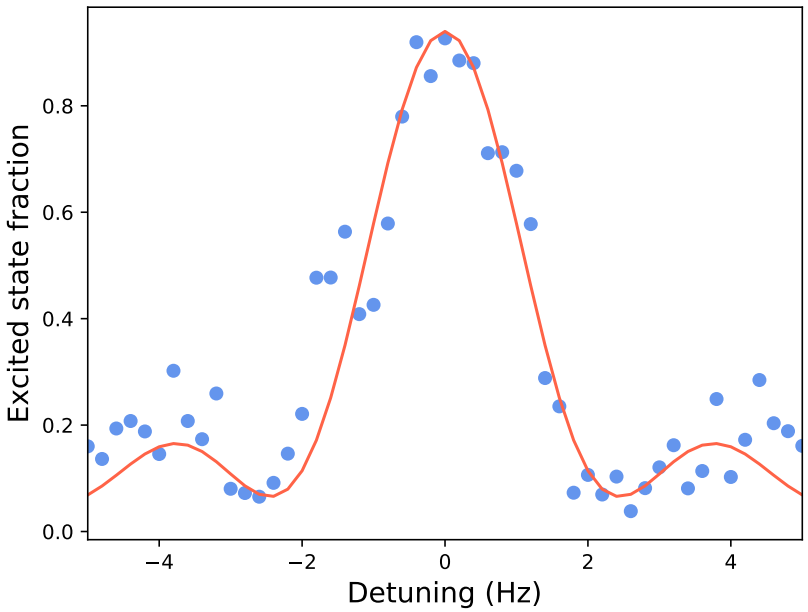


Figure 2.3: Excitation spectrum of the clock transition for a 350 ms Rabi pulse. The fit takes into account inhomogeneities in the Rabi frequency across sample by following the procedure outlined in reference [12]. The measured linewidth of 2.2 Hz combined with the resonance frequency of 429 THz results in a quality factor of 1.95×10^{14}.

2.3.1 Atomic polarisability calculation

For atom-light interactions mediated by the dipole force, the polarisability is given by:

$$\alpha(\omega) = 2\sum_k \frac{\omega_{ik} |\langle\phi_i| \epsilon \cdot \hat{d} |\phi_k\rangle|^2}{\hbar(\omega_{ik}^2 - \omega^2)} \tag{2.7}$$

where $\{\phi_k\}$ is the set of all dipole allowed transitions from the state of interest ϕ_i which are responsible for the induced dipole moment. Although equation Equation 2.7 is correct, it is more helpful to rewrite it as $\alpha(\omega) = \epsilon_\mu \epsilon_\nu \alpha_{\mu\nu}(\omega)$ (repeated indices are summed over) where $\alpha_{\mu\nu}(\omega)$ is the Kramers–Heisenberg polarisability tensor given:

$$\alpha_{\mu\nu}(\omega) = 2\sum_k \frac{\omega_{ik} \langle\phi_i| \hat{d}_\mu |\phi_k\rangle \langle\phi_k| \hat{d}_\nu |\phi_i\rangle}{\hbar(\omega_{ik}^2 - \omega^2)}. \tag{2.8}$$

As a rank-two tensor, $\alpha_{\mu\nu}(\omega)$ can be decomposed into its three irreducible spherical tensor components: a scaler part given by its trace, an asymmetric vector component, and finally a traceless, symmetric tensors part. The functional form of these components depends on the angular momentum of states which are coupled by the light. For strontium, the precence of nucleur spin gives rise to hyperfine structure resulting in the relevant quantum numbers being the total angular momentum F and its projection m_F. As the full derivation is rather involved, I present only the final results [1]:

$$\alpha(\omega) = \alpha_s + i\alpha_v(\epsilon \times \epsilon^*) \cdot \hat{z}\frac{m_F}{F} + \epsilon_t \frac{3|\epsilon \cdot \hat{z}|^2 - 1}{2}\left[\frac{3m_F^2 - F(F+1)}{F(2F-1)}\right]. \tag{2.9}$$

Although the vector and tensor components are important as they introduce po-

[1] For the complete derivation I suggest Dan Steck's textxbook [135] which is available for free online. A more concise derivation relevant for strontium can be found in the thesis of Michael Martin [86].

tential systematic shifts of the clock frequency, it is the scalar component that dominates the trapping force leading to the greatest perturbation of the clock states. The scalar polarisability is given in terms of reduced matrix elements between the initial state denoted by (n, J) and all other states denoted by (n', J'):

$$\alpha_s(\omega) = \sum_{n'J'} \frac{2\omega_{nF\to n'F'} |\langle nF||\hat{d}||n'F'\rangle|^2}{3\hbar(\omega^2_{nF\to n'F'} - \omega^2)}. \tag{2.10}$$

By relating the reduced matrix element between levels to the experimentally measured Einstein A_T coefficients using the appropriate branching ratios, simple expressions for possibilities of the two clock states can be derived:

$$\begin{aligned} \alpha_{s,^3P_0} &= 2\pi\hbar\epsilon_0 c^3 \sum_{n'J'} \frac{2\omega_{nF\to n'F'} A_T(n', J')}{\omega^3_{n',J'}(\omega^2_{nF\to n'F'} - \omega^2)} \frac{2L'+1}{3} \\ \alpha_{s,^1S_0} &= 6\pi\hbar\epsilon_0 c^3 \sum_{n'J'} \frac{A_T(n', J')}{\omega^2_{n',J'}(\omega^2_{nF\to n'F'} - \omega^2)} \end{aligned} \tag{2.11}$$

Fortunately, all the relevant decay rates have been experimentally measured [165, 112, 127, 159]. To verify that all relevant transitions have been accounted for, one can use the Thomas-Kuhn sum rule which states the sum of all the oscillator strengths f_{ak} must sum to the number of valence electrons, which in this case is two. If we assume the missing levels are most likely high lying states close to the continuum, a crude approximation is to group them at continuum and assign to them the missing line strengths. For the $^1S_0$ state, the oscillator strengths sum to 1.98 – indicating that most, if not all, the relevant states are accounted for. However, for the upper clock state this is not the case as the sum is 1.57. Using the equation 2.11, we can calculate the polarisability of the two clock states, shown in Figure 2.4, in order to find the magic-wavelengths for which the polarisabilities are equal. There are several possible choices of such wavelengths such as 813 nm and 497 nm; both of which are 'red' detuned from the dominant transitions resulting in a positive polarisability

causing the atoms to be attracted to regions of high intensity. Alternatively, one could operate the lattice at the 'blue' detuned magic-wavelength at 390 nm. This is technically more challenging as the atoms are repulsed from the lattice and therefore one needs to implement a lattice in all three dimensions to trap the atoms [140].

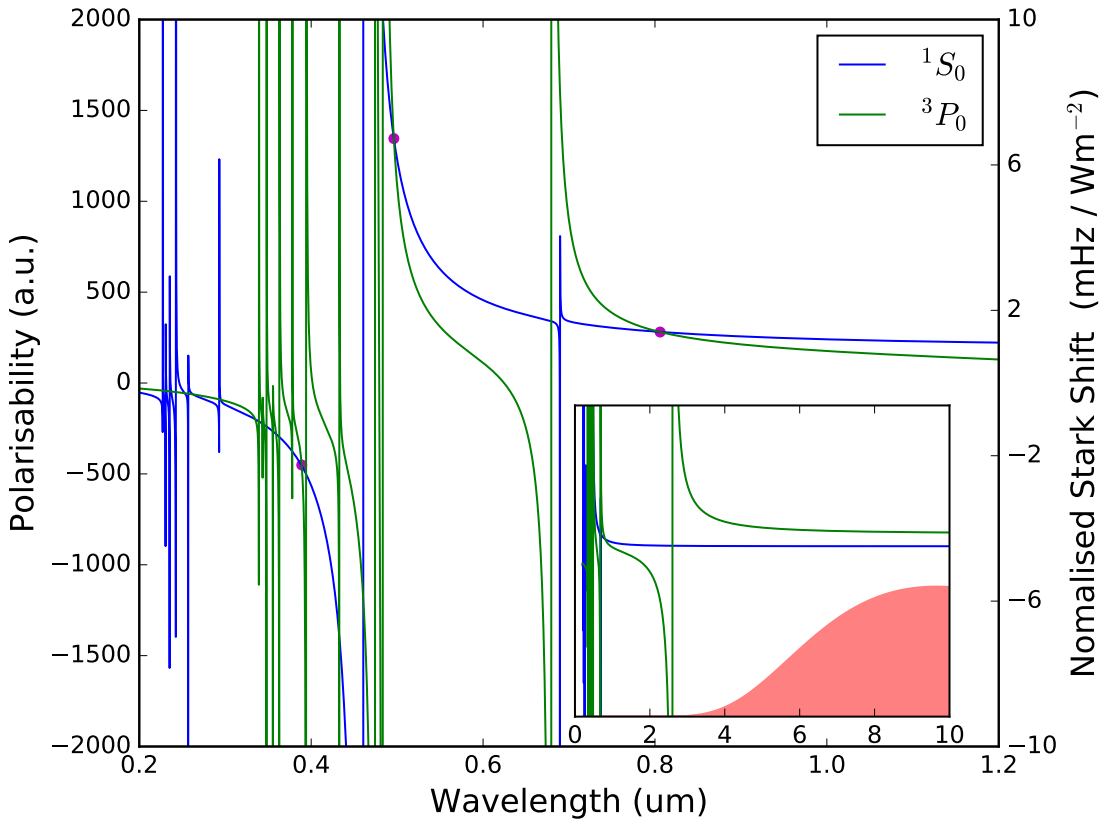


Figure 2.4: Polarisability of the two clock states expressed in atomic units (left) and normalized Stark shift (right). The magic-wavelengths occur where the polarisabilities cross such that the two states experience equal Stark shifts. Three examples at 390, 497 and 813 nm are marked with magenta dots. Inset shows the same polarisability extending to 10 µm along with the blackbody radiation spectrum (red, scaled to fit) at 300 K.

2.4 Motional sideband thermometry

The potential introduced by the one dimensional lattice in both the axial and radial direction is not perfectly harmonic. At the centre of the trap, the leading-order

anharmonic corrections to the potential are given by [103]:

$$V_{zz} = -\frac{(m\omega_z^2 z^2)^2}{12U_0} \tag{2.12}$$

$$V_{rr} = -\frac{(m\omega_r^2 r^2)^2}{8U_0} \tag{2.13}$$

$$V_{zr} = -\frac{m^2\omega_z^2\omega_r^2 z^2 r^2}{4U_0} \tag{2.14}$$

As a result, the resonance frequency to drive transitions between motional states will no longer be independent of the initial state as is the case with a perfect harmonic oscillator. This is useful for inferring the temperature of the sample because the motional excitation spectrum provides information regarding the probability distribution of atoms over the initial states. This section will provide a recipe for determining the temperature of the atomic sample using our standard spectroscopy toolkit. This will come in handy when characterizing the efficacy of laser cooling and will also be important for evaluating temperature dependent perturbation to the clock transition.

For the carrier transition– corresponding to the motional state remaining unchanged– any anharmonic corrections have no affect on the resonance frequency[2]. However, for the 1st-order 'sidebands' for which the axial motional state changes by ± 1, the distribution of atoms across several axial and radial motional states will distort the lineshape. Using perturbation theory, a change in resonance frequency Δ_{ij} between the motional state $|n_z, n_r\rangle \rightarrow |n_z', n_r'\rangle$ resulting from an anharmonic corrections V_{ij} is given by:

$$\Delta_{ij} = \langle n_{z'}, n_{r'}| V_{ij} |n_{z'}, n_{r'}\rangle - \langle n_z, n_r| V_{ij} |n_z, n_r\rangle \tag{2.15}$$

[2]This is only true when operating at the magic-wavelength such that the trapping potentials for the ground and excited state are identical. As the lattice is detuned from the magic-wavelength, motional effects begin to modify the carrier lineshape. Not surprisingly, this effect also provides information regarding the temperature [89].

For corrections listed in (2.12)—(2.13) for transitions between motional states $n_z' = n_z \pm 1$ (1st order sideband) and $n_r' = n_r$ [3], the resulting shift of the axial motional sideband frequency is given by:

$$\Delta_{zz}(n_z \pm 1) = \mp \frac{\hbar^2 \omega_z^2((2n_z \pm 1))}{8U_0} \tag{2.16}$$

$$\Delta_{rr} = 0 \tag{2.17}$$

$$\Delta_{zr}(n_r) = \frac{(n_r + 1)\hbar^2 \omega_z \omega_r}{4U_0} \tag{2.18}$$

The pre-factor $\frac{\hbar^2\omega_z^2}{8U_0}$ for Δ_{zz} is half the lattice recoil frequency and is equal to 1.8 kHz. Comparatively, the pre-factor for radial correction is scaled by the ratio of $\frac{v_r}{v_z} = \frac{\lambda}{\sqrt{2}\pi\omega_0} \approx .001$ where ω_0 is the waist of the lattice beam which for our trap is approximately 65 µm. To get a scale of the magnitude of these corrections, we calculate expectation value for motional state using the appropriate probability distributions:

$$\begin{aligned} p_z(n_z, T_z) &= \frac{1}{Z} e^{-\frac{h\nu_z n_z}{k_b T_z}} \\ P_r(n_r, T_r) &= \frac{n_r + 1}{Z} e^{-\frac{h\nu_r n_r}{k_b T_r}} \end{aligned} \tag{2.19}$$

where Z is the partition function and the extra factor of $(n_r + 1)$ for the radial function arises from the degeneracy of radial state (i.e. for a given n_r there are $n_r + 1$ combination of n_x and n_y that satisfy $n_x + n_y = n_r$). For a temperature of 1 µK and axial trap frequency of 65 kHz, the expectation values for the axial and radial motional states are $\langle n_z \rangle = 0.3$ while $\langle n_r \rangle = 62$.

To estimate the temperature from the sideband spectrum, the key realisation is that an excitation probability $P(\Delta)$ at a detuning Δ from the maximum carrier

[3] With the probe beam aligned along the lattice and perpendicular to any radial motion, it is not possible to change the radial motional state. This assumption is validated by the lack of any radial motional side bands in the excitation spectrum.

frequency is proportional to the number of atoms in all motional states $|n_r, n_z\rangle$ subject to the constraint $\Delta_{zz} + \Delta_{zr} = \Delta$ [12]. The number of atoms in a given motional state is proportional to the product of $p_z(n_z, T_z)$ and $p_r(n_r, T_r)$. Therefore, the excitation probability as a function of the detuning from the carrier can be expressed as:

$$\begin{aligned} P(\Delta) &\propto \sum_{nz} \sum_{nr} p_z(n_z, T_z), p_r(n_r, T_r) \delta(\Delta - \Delta_{zz}(n_z) - \Delta_{zr}(n_r) \\ &= \sum_{nz} p_z(n_z, T_z), p_r(\frac{4U_0(\Delta - \Delta_{zz}(n_z))}{\hbar^2 \omega_z \omega_r} - 1, T_r) \end{aligned} \tag{2.20}$$

where δ is the Dirac delta function and equation 2.18 was used to convert from Δ_{zr} to an equivalent radial state. This conversion requires that we treat the radial motional degree of freedom as a continuum rather than discrete which is valid given that $k_b T >> \hbar\omega_r$. In practice it was observed that including axial motional states beyond $n_z = 3$ made no difference to the final estimate for temperature. Figure 2.5 shows a fit to sideband spectrum using equation 2.20.

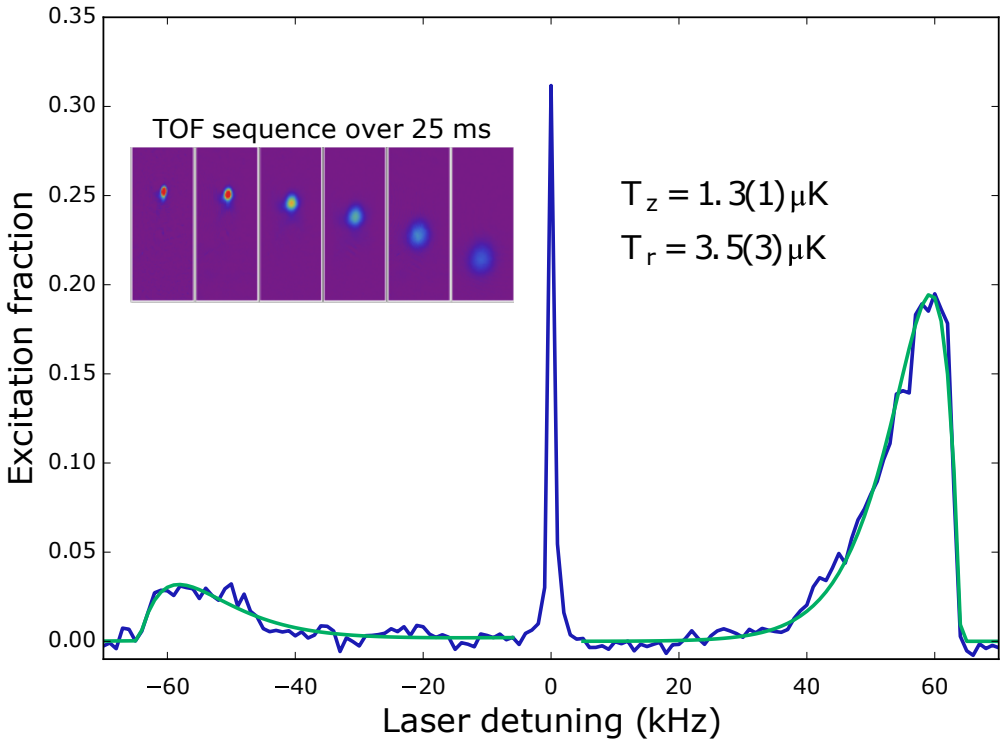


Figure 2.5: Spectrum showing the blue and red axial sidebands corresponding to changing the axial motional energy by $\hbar\omega_z$. Using equation 2.19, the relative amplitudes can be used to infer the axial temperature. The asymmetry of the sideband towards the carrier is a result of the radial temperature. Hotter atoms venture further from the centre of the beam where they encounter a more anharmonic trapping potential which reduces their motional frequency. The shape of sideband can therefore be used to estimate the average thermal energy in the radial direction. **Inset:** Time-of-flight images of atoms released from the red MOT, the rate of expansion gives an estimated temperature of 2 µK [51] which corroborates the temperature measurements using the sideband spectrum.

Chapter 3

Sr1: The First Generation Strontium Lattice Clock

In 1981, Hans Dehmelt outlined his plan for performing coherent optical spectroscopy in an atomic system in free space[31] – peculating that "an atomic line spectral resolution of about 1 part in 10^{18} may be realized in the not too far future." He envisioned using a RF Paul trap to bring a single ion into "a state of complete rest in free space," allowing it to be spectroscopically probed free of transit time broadening and other motional effects. We have followed a different route to realise Dehmelt's vision by using neutral atoms held in an optical lattice. Using a neutral system has the benefit of being able to trap and spectroscopically probe thousands to millions of atoms at once – helping to improve the stability of the clock.

However, such a route is not without its challenges. First, the dipole trapping forces used to confine the atoms in the lattice are minute compared to those applied to an ion in a Paul trap. As a result, our lattice struggles to capture atoms any hotter than 10 µK. Given that we evaporate strontium by heating to 500 °C to generate the flux of atom we intend to trap, it is clear we have to cool the atoms before we can trap them. To accomplish this, we employ the well-established techniques of laser cooling.

We cool using a standard two stage process. First we use a broad transition at 461 nm which efficiently captures the fast moving atoms and cools them to approximately a millikelvin. This is too hot for loading into the lattice requiring us to have a second stage of cooling using narrow transitions at 689 nm which results in a steady state temperature of approximately 1 µK. Once in the trap, we must ensure the atoms are not knocked out by collisions with other atoms, requiring that the cooling and trapping is carried out in an ultra-high vacuum environment. The first part of this chapter will outline these steps in more detail.

With atoms safely trapped and ready for spectroscopy, we may wonder, have we faithfully achieved Hans Dehmelt's original vision of an "atom in free space?" For example, what effect does the trap have on the atoms? As we saw in the previous chapter we can greatly suppress the lattice perturbation with judicious choice of wavelength, but how precisely can we find such a magic-wavelength? One quickly finds several more perturbations (residual background fields, interactions between atoms or background gas particles, Doppler shifts caused by atomic motion, etc.), all of which need characterisation before we can make any claims of accuracy. The last part of this chapter will detail this characterisation process culminating in an uncertainty budget for our lattice clock.

3.1 Experimental apparatus

As the fourth PhD student to work on NPL's strontium lattice clock much of the details regarding the design of the science package and laser system have been covered in the theses of my predecessors. In particular, information pertaining to the Zeeman slower, clock laser, repump lasers, and 461 nm-laser system needed for first stage cooling can be found in Ian Hill's thesis [49]. Additional details regarding the operation of Sr1 including the 689 nm-laser system, lattice laser, laser light distribution and the

science package (see figure 3.1) can be found in Richard Hobson's thesis [51]. Rather than duplicate their work here, I will focus on aspects of the experiments that have not yet been covered in detail or that provide context pertinent to understanding results presented later in the thesis.

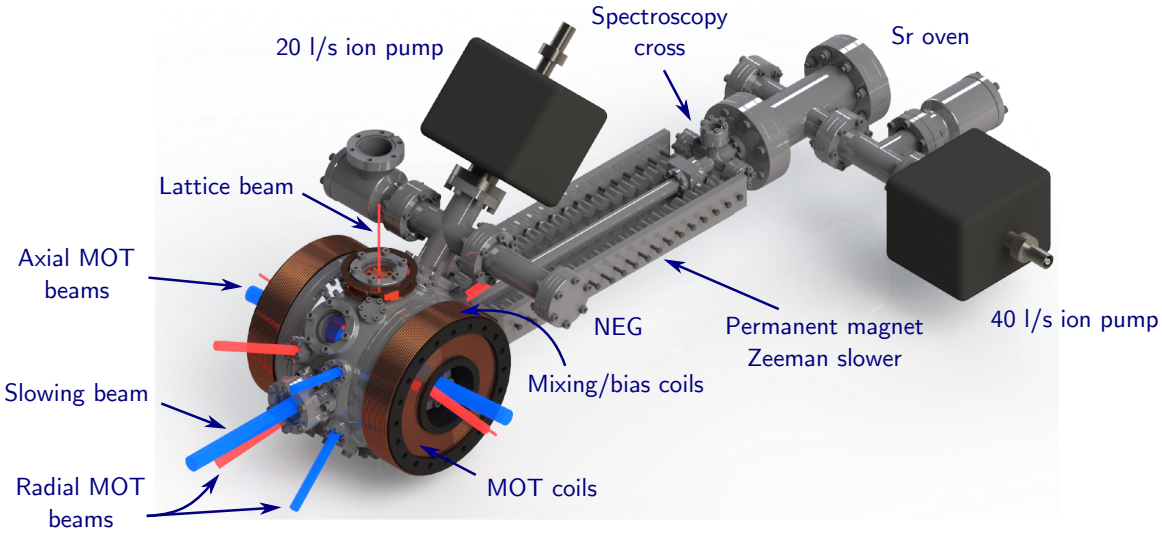


Figure 3.1: Computer rendering of the science package including the vacuum chamber, oven and Zeeman slower. Also included in the figure are the various laser beams needed for cooling, trapping, and spectroscopy.

3.1.1 Clock laser distribution and universal synthesiser

Inside the strontium lab, a 698 nm laser locked to a local cavity [80] spectroscopically probes the ultra-narrow clock transition. Along with being sent to the atoms, the light is also distributed to both frequency comb labs and a second tunable cavity which has a broadband high reflectivity coating extending from 630 nm to 1000 nm. By locking the tunable cavity to the clock laser and subsequently locking the 689 nm, 813 nm and 922 nm (sub-harmonic 461 nm) lasers to the cavity, frequency stability of the clock laser is transferred to these other systems.

An important consideration is that the clock cavity modes won't generally coincide with the atomic resonance or lie within the narrow bandpass filter used to filter the

beatnote with the frequency comb. Furthermore, the cavity length drifts over time requiring compensating corrections of the laser frequency. This frequency bridging and drift compensation is done with a series of acousto-optic modulators (AOMs) as shown in figure 3.2. Additional AOMs (AOM4, AOM6 and AOM7) at either end of the fibres that deliver light to the atoms and frequency combs are used to stabilise the optical path lengths [83]. By changing AOM3 frequency, the cavity drift is canceled. The frequency corrections applied to the cavity drift AOM are fed forward onto the transfer cavity AOM (AOM2) such that the length of the tunable cavity is kept fixed.

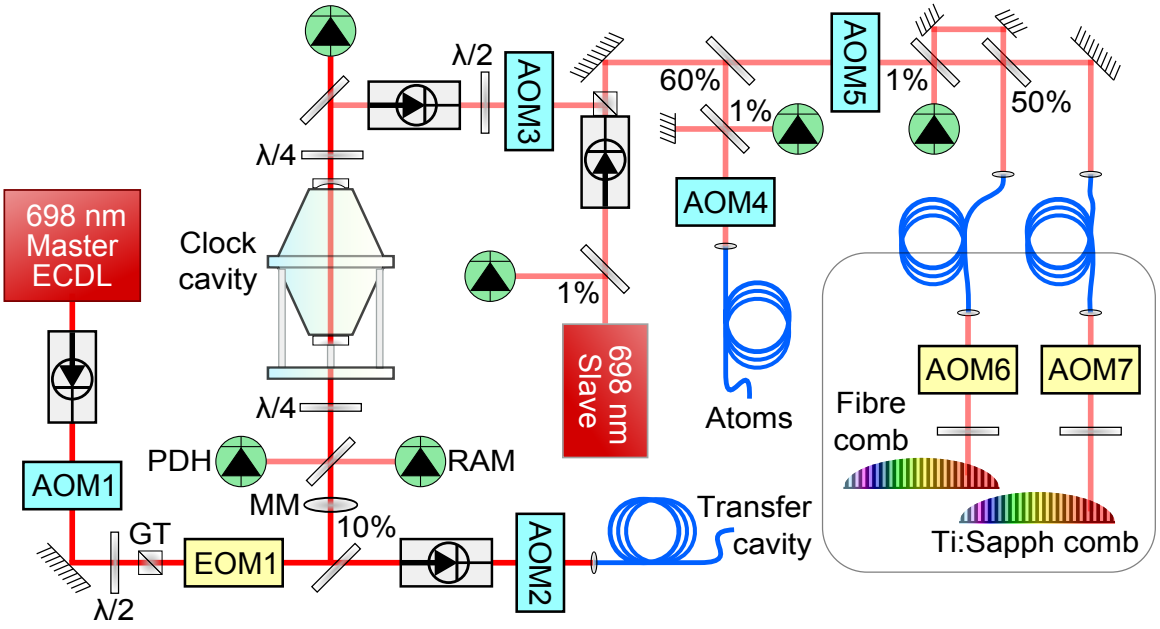


Figure 3.2: Diagram of the clock laser including its distribution to both frequency combs, transfer cavity, and the atomic system. These paths are phase stabilised via an interferometric comparison between the outgoing and reflected beams to ensure that the stability is not compromised. Figure credit: Ian Hill.

At one time, the cavity design exhibited near world-leading stability with a flicker frequency floor at low 10^{-15} level. However, as the technology improved it was surpassed in stability with much longer ULE cavities and, later, by cryogenic silicon resonators. Building new systems in an effort to keep up with this rapid progress is expensive; especially if NPL were to build separate cavities for each of its differ-

ent optical clocks. Instead, we have elected to build one master oscillator operating at 1064 nm and then transfer its stability to local cavities operating at the atomic transition wavelength using the frequency comb [142]. In this way, only a single high performance cavity is needed for all optical clocks. In practice the stability transfer is done as follows:

1. Detect the beatnotes between frequency comb and both the 698 nm (ν_A) and 1064 nm (ν_B) lasers.

2. Measure ν_{ceo} and subtract it from both ν_A and ν_B using a mixer.

3. Once ν_{ceo} has been removed, divide ν_A (or alternatively ν_B) by the mode number ratio m_B/m_A (m_A/m_B).

4. Subtract ν_B and $\nu_A \times m_B/m_A$, to generate a new RF frequency ν_C.

4. ν_C is proportional to the frequency difference between the 698 nm and 1064 nm lasers and is insensitive to change in ν_{ceo} and ν_{rep}.

5. The phase of ν_C stabilised to a reference signal by changing the frequency of the 698 nm laser.

In practice, to stabilise the synthesised beatnote ν_C between 698 nm and 1064 nm we phase lock it to a direct digital synthesiser (DDS) by actuating the frequency of the cavity drift AOM. Because the AOM is placed prior to the light being split into the paths going to the combs and science package, the light interrogating the atoms is free of any frequency instability caused by length fluctuations of our local cavity. The reference for the universal synthesiser is based on a newly constructed 48 cm ULE cavity similar to that presented in [45]. On medium to long timescales (1 s to 10 s) the cavity has a flicker floor at the low 10^{-16} level as measured via three-corner hat comparison with other stable lasers at NPL. For shorter timescales (less than 1 s)

more relevant for the stability of the atomic clock, the frequency noise increases—most likely resulting from vibrations – with work underway to improve such limitations.

3.1.2 Transfer cavity

To realise an operational lattice clock, one first has to undertake the laborious task of stabilising the frequencies of all the lasers needed for cooling, repumping and trapping. To accomplish this we employ a tunable length ULE cavity, the length of which is stabilised to the clock laser. Once referenced to the clock laser, the cavity serves as a stable reference to which other lasers can be locked using the standard Pound-Drever-Hall (PDH) approach as shown in figure 3.3. In such a configuration, the cavity is able to transfer the stability of the clock laser to the other lasers – at a fraction of the cost of alternative methods such as a multi-branch frequency comb.

The cavity's free-spectral range is 1.8 GHz which constrains the maximum offset between a cavity mode and the required laser frequency to be less than 900 MHz. To compensate the inevitable frequency mismatches we use either AOM or waveguide EOM placed between the laser and the cavities. The AOM placed between the cavity and clock laser is tuned such that when the 922 nm light is locked to cavity it is resonant with the principle cooling transition once doubled to 461 nm. Next the second stage cooling light at 689 nm is offset locked to the cavity using a waveguide EOM. By phase modulating the light at the offset frequency, sidebands are generated which can be locked to the cavity. Conveniently, the PDH modulation sidebands can be generated using the same EOM by phase modulating the sideband frequency. By tuning the offset frequency of the sideband locked to the cavity, the carrier frequency is changed [145]. The same scheme is used to lock the 813 nm laser to the transfer cavity. To operate our second generation lattice clock, Sr2, we also need to lock

[1]For the 636 nm and 689 nm lasers we use an EOM manufactured by Jenoptik while for 813 nm we elected to source the modulator from Photline (which has since joined iXblue).

636 nm to the transfer cavity. Once again, we elected to sideband offset lock the 636 nm laser to the cavity using a waveguide EOM [1]

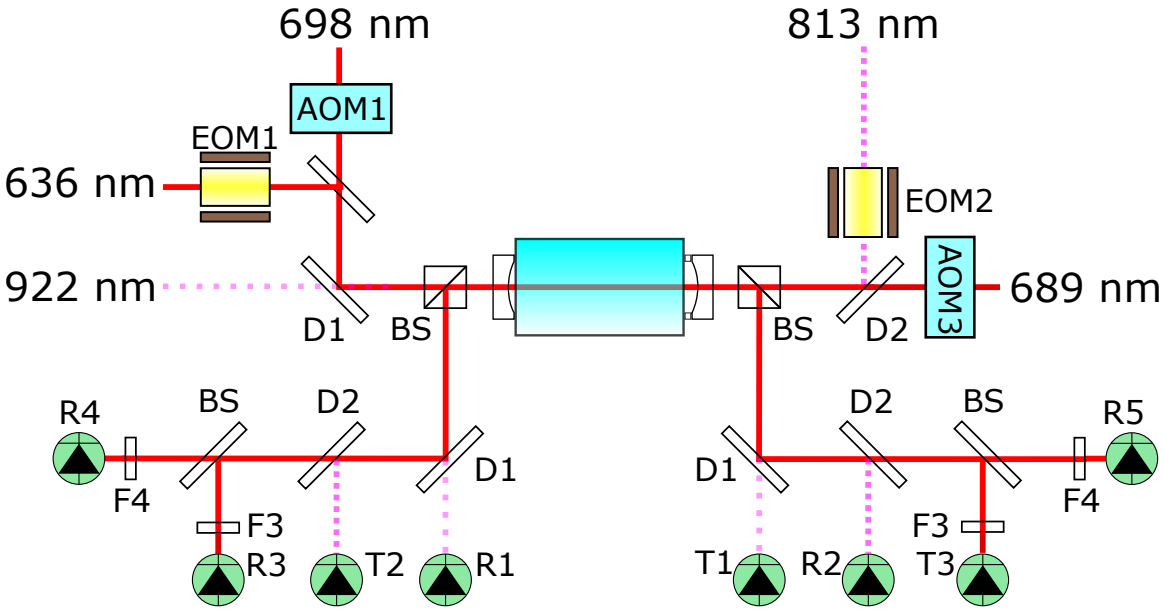


Figure 3.3: Optical setup of the transfer cavity system. By locking the cavity to the clock laser, its stability can be transferred to the other cooling and trapping lasers. Included in the diagram is a laser operating at 636 nm which is used for laser cooling in our second generation lattice clock which will be covered in chapter 4. Figure credit: Ian Hill.

3.1.3 Atomic servo-control

The role of the atomic servo-control is to steer the local oscillator towards the atomic resonance using an error signal derived by probing the atomic transition and measuring the excitation response. For Rabi interrogation, the error signal that steers the local oscillator is derived as follows:

1. The laser frequency is detuned below the atomic resonance by half the atomic linewidth Δ

2. The atomic sample is probed on the 'low' side of the resonance and the excitation fraction F_- is recorded

3. The laser frequency is detuned above the atomic resonance by half Δ

4. The excitation fraction F_+ on the 'high' side is measured

5. The error signal ϵ is the difference between the 'high' and 'low' side excitation fractions[2]

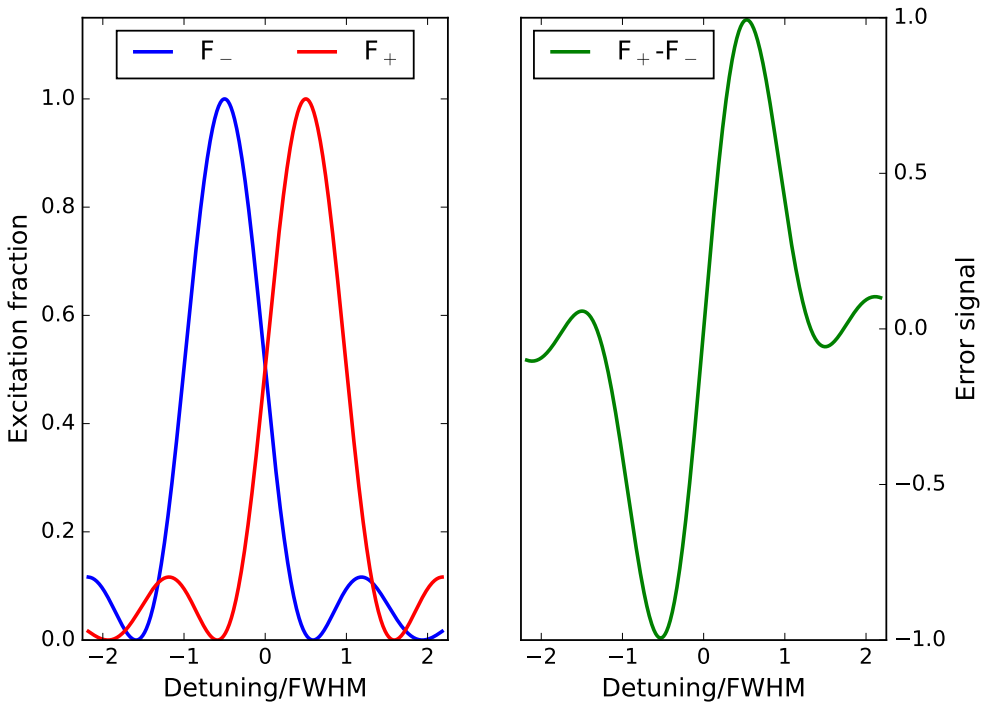


Figure 3.4: (left) Excitation line shapes for the two probe pulses taken at a detuning of $\pm\Delta/2$. (right) Error signal generated by taking the different in excitation fraction between the two probe pulses.

In practice, the excitation fraction is calculated by probing the line, then measuring the number of atoms remaining in the ground state N_g via fluorescence detection using resonant light at 461 nm. This detection probe quickly heats the atoms out of the trap leaving only atoms shelved in the excited state. These excited state atoms are then repumped to the ground state where the fluorescence detection is repeated to

[2]A similar procedure is followed when using Ramsey spectroscopy, except instead of stepping the detuning above or below resonance by half a linewidth, the phase of the last Ramsey pulse is stepped by $\pm\frac{\pi}{2}$.

measure the atom number in the excited state N_e. Finally, with no atoms remaining, a background image is taken for calibration of N_{bg}. Using these three measurements, the excitation fraction is calculated as:

$$F = \frac{N_e}{N_e + N_g - 2N_{bg}}. \tag{3.1}$$

For detection of the atomic fluorescence we use an EMCCD camera[3]. The entire sequence of cooling the atoms, loading them into the trap, and probing the clock transition takes approximately one second. We have to repeat this measurement twice – once for each side of the atomic line – to derive the error signal used to steer the clock.Returning to the atomic servo, it is helpful to rewrite the laser error signal in terms of the laser frequency ν relative to that atomic resonance ν_0 as:

$$\epsilon = F_+ - F_- \cong D(\nu - \nu_0) \tag{3.2}$$

where D is the linear discriminant which parameterises the sensitivity of the error signal to small laser detunings from resonance. By stepping the laser by half the atomic linewidth when measuring that excitation fraction, we maximise D giving us the greatest sensitivity to changes in laser frequency. The other way to increase the sensitivity is by decreasing the width Δ of the spectroscopic lineshape. Because we operate in a regime where the atomic lifetime far exceeds the interrogation time (<1 s), it's the finite probe time T_p that sets the width of the feature. For Rabi interrogation, the linewidth Δ is given by $0.8\mathrm{T_p^{-1}}$. Unfortunately we are limited to how long we can probe the transition by the coherence time of the laser. For exceedingly long pulses we observe reduced spectroscopic contrast C which degrades the error signal. For a

[3]Camera Model: Andor iXon Ultra. Alternatively we could have used a PMT, which would have been a far cheaper solution. However, the camera does provide information that a PMT does not, such a spatial excitation inhomogeneity which can be used to minimise magnetic field gradients or calibrate the vector light shift of the lattice [85]

Rabi pulse of length T_p and contrast C, the discriminator is given by:

$$D_{\text{Rabi}}(T_c, C) = 1.9 T_c C. \tag{3.3}$$

The corrections $\delta\nu$ applied to clock laser are equal to error signal scaled by a factor $\frac{0.8g}{T_p C}$ where g is the servo gain. Given our observed contrast of 0.9, we can calculate value for the gain such that the frequency is restored to line centre by each correction which is given by:

$$\frac{g}{T_p C} \times 1.9 T_p C = 1 \rightarrow g = 0.6. \tag{3.4}$$

In practice, this choice of gain will result in an under-damped servo loop. Empirically we find $g = 0.3$ is a good compromise between the servo quickly acquiring lock without exhibiting excessive oscillation in response to a disturbance. We also implement a second integrator which steers the cumulative error to zero which reduces our susceptibility to frequency offsets which can arise from the linear frequency drift of the laser. Hence, the total frequency correction for each clock cycle is given by:

$$\Delta\nu_n = \frac{g}{0.8 c T_p}\epsilon_n + \frac{g_2}{0.8 c T_p}\sum_{i=0}^{n-1}\epsilon_i. \tag{3.5}$$

Unfortunately, frequency offsets can still occur given that the feedback is not continuously applied. To illustrate why, consider a linear ramp of the local oscillator at a rate R and servo that applies corrections at a cycle interval T_c[4]. Neglecting any other frequency noise processes, in steady state the term $\propto \sum \epsilon_i$ in equation 3.5 must be constant. This implies the laser was on resonance when the frequency was measured such that the error signal was zero. For this to be true, the integral correction proportional to $\sum \epsilon_i$ made in the previous cycle must have compensated

[4]I assume the laser frequency is measured instantaneously at T_c and neglect any secondary effects due to averaging of the laser frequency over the probe pulse.

ramp offset accumulated during the current cycle. Therefore, for a positive ramp, the average laser frequency will be offset below the atomic resonance by $RT_c/2$, see figure 3.5. For our clock T_c is approximately 1 s and R can be as much as 50 mHz/s, which could result in a fractional frequency shift at the mid 10^{-17} level if gone unchecked. To combat this effect we apply a continuous ramp to the clock laser frequency. The ramp servos the double integral corrections ($\propto \sum \epsilon_i$) to zero. The gain of the ramp servo is chosen to have a time constant of approximately 100 s such that the servo properly tracks changes in the drift rate while not degrading the stability of the lock. The ramp is implemented with a DDS that has a maximum update rate of 100 Hz. To reduce further the sensitivity to linear drift, we probe the 'high' and 'low' sides of line following Thue–Morse ordering ('low-hi-hi-low-hi-low-low-hi') rather than directly alternating ('low-hi-low-hi') [2]. The servo corrections and frequency ramps are applied to the clock laser by actuating the frequency of the DDS used as the reference the phase lock to the universal synthesiser.

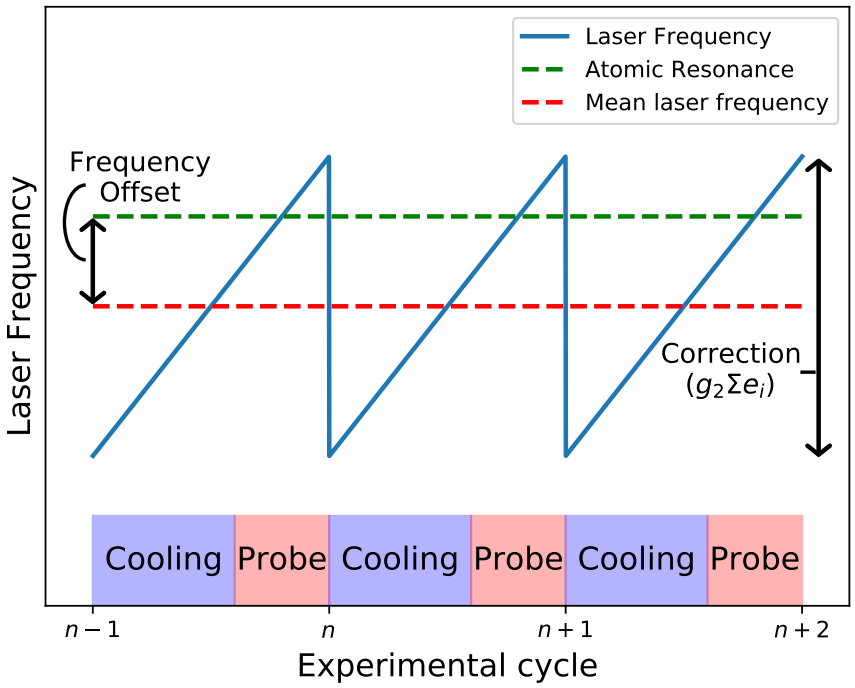


Figure 3.5: A linear ramp of the clock laser's frequency combined with the cooling deadtime in the experimental cycle can give rise to frequency errors. Each experimental cycle consists of a cooling stage and probing stage. In steady state, the mean laser detuning during the probe stage is zero. Therefore the error signal for that cycle is zero and integrated error ($\propto \sum \epsilon_i$) is constant. The applied correction at the end of the cycle is therefore $g_2 \sum \epsilon_i$ and precisely cancels the frequency error incurred by the laser during the cooling period such that the laser is once again resonance during the probe pulse in the following cycle. As result, the mean laser detuning averaged over the entire experimental cycle not zero – instead it is slightly below the atomic resonance for a positive linear drift of the laser frequency.

3.2 Stability of Sr1

At its foundation, the clock is a metrological apparatus that measures the resonant frequency of an atomic transition. As is the case with any physical measurement, we are interested in its accuracy and precision. The two figures of merit, although often confused, are fundamentally different:

Accuracy: The degree of closeness between a measured value and its *true*

value after attempting to correct for all sources of systematic uncertainty.

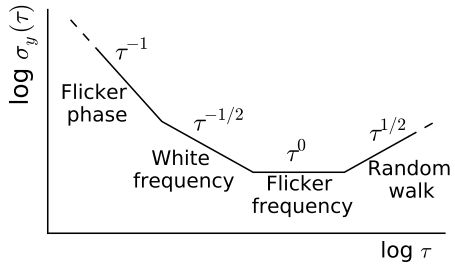
Precision: The variation in repeatedly measured values when all efforts are made to keep experimental conditions constant.

The accuracy of Sr1 – NPL's first generation lattice clock – will be covered in the next section, but first I will focus on precision. The precision of the clock is often limited by various noise sources which infiltrate the measurement. This noise can often be suppressed by performing several measurements and averaging the result. Therefore, is it useful to characterise how the precision of a measurement improves with the number of measurements averaged. This information is conveyed by the clock's stability; a more stable clock will need fewer measurements to reach the same precision. Because the cycle time needed to make a measurement differs between clocks, stability is typically given as a function of averaging time τ, rather than number of measurements n, in order to make meaningful comparisons between systems. The averaging time and the number of measurements are related by the cycle time T_c via $\tau = nT_c$. The stability of a signal $y(t)$ when averaged for a time τ is often expressed using the Allan deviation, defined as:

$$\sigma_y^2(\tau) \equiv \frac{1}{2\tau} \left\langle \left(\int_{t+\tau}^{t+2\tau} y(t')dt' - \int_t^{t+\tau} y(t')dt' \right)^2 \right\rangle \tag{3.6}$$

where $\langle - \rangle$ represents time averaging over the entire length of data. The physical interpretation of the Allan deviation and how it's a measure of stability is clear: if one bins the data in time into sections of length τ, the Allan deviation is the root-mean-square deviation in the average of adjacent bins. We will use Allan deviation extensively throughout this thesis as a way to estimate stability. Unfortunately, the Allan deviation does not always improve with averaging time; it could in fact get worse with time if parameters of the experiment are changing. The dependence of the

stability on averaging time is set by the characteristic noises processes perturbing the measurement. For example, if sequential measurements can be treated as statistically independent, then the Allan deviation is equivalent to the standard deviation of the mean and therefore would be expected to improve as $1/\sqrt{\tau}$. This illustrates the importance of understanding the noise correlations introduced between measurements when evaluating stability. Fortunately, we have a powerful tool for characterisation noise correlations: the Wiener-Khinchin theorem. The theorem states that the auto-correlation function of stationary noise processes and its corresponding power spectral density in the frequency domain are linked by Fourier transformation. Therefore, the stability of clock is determined solely by the power spectrum of the perturbing noise. It is useful to identify how the different types of noise processes (for example white, flicker or random walk) affect the Allan deviation. This information shown in figure 3.6.



Noise type	$L_y(f)$	$\sigma_y(\tau)$
White & flicker phase	f^2 & f	$1/\tau$
White frequency	h_0	$\sqrt{h_0/2\tau}$
Flicker frequency	h_1/f	$\sqrt{2h_1 \ln 2}$
Random walk	h_2/f^2	$\sqrt{\frac{2}{3}\pi^2 h_2 \tau}$

Figure 3.6: Allan deviations for various common noise types as characterised by their single sided power spectral density $L_y(f)$. White ($L_y(f) \propto f^2$) and flicker ($L_y(f) \propto f$) phase noise both lead to a scaling $1/\tau$ for $\sigma_y(\tau)$, but the constant of proportionality depends on the measurement bandwidth [123].

Achieving a stable frequency reference requires that the LO is precisely locked to the atomic transition. It is therefore paramount that the error signal used to discipline the LO is as free of noise as possible. In general, the noise that degrades the error signal can be divided into three categories: technical, laser-related and fundamental. For example, technical noise can arise from our fluorescence detection method used to measure the excitation fraction. Such detection methods will suffer from shot noise

which is inherent in any technique based on collecting a finite number of photons. Fortunately, high quantum efficiency of our camera (80%) and the moderate numerical aperture of our imaging system (NA = 0.1) allows us to gather enough signal such that the shot noise instability is not limiting. Beyond that, we also ensure that the laser intensity is sufficiently stable to ensure the background scatter in each of the three pictures used for deriving the normalised excitation is constant. A thorough discussion on the characterisation of sources of technical noise can be found in [4].

Once instability resulting from technical noise is reduced to a negligible level, one can start to address other sources of noise more fundamental to the operation of the clock. Specifically, the non-deterministic nature of projective quantum measurement – under which falls the fluorescence detection method we use to estimate the excitation probability F – gives rise to statistical fluctuations in possible outcomes. After spectroscopy, the fluorescence readout can be treated as N independent measurements used to estimate F, where N is the number of atoms. This holds assuming the system is not entangled such that measurement outcomes for each atom are uncorrelated. The expectation value follows a binomial distribution with a variance $(\Delta F)^2$ given by:

$$\sigma_F^2(\tau) = \frac{F(1-F)}{N} \cong \frac{1}{4N} \text{ for } F \cong 1/2. \tag{3.7}$$

This resulting variance in the measurement is often referred to in literature as quantum projection noise (QPN) and can be thought as the fundamental limit of instability one can achieve when using an ensemble of unentangled atoms. To see how QPN instability of the excitation fraction decreases with clock stability, we can compute its Allan variance $\sigma_F(\tau)^2$. We recognise that the QPN of subsequent measurements of the excitation fraction are uncorrelated resulting in the Allan variance improving in accordance with a white noise process as:

$$\sigma_F^2(\tau) = \frac{T_c}{\tau}(\Delta F)^2. \tag{3.8}$$

By scaling this instability by the discriminant of the error signal which converts between excitation probability and frequency corrections, we can calculate the QPN limited Allan variation of our frequency standard which, when expressed in fractional units, is given by:

$$\sigma_\nu^2(\tau) \approx \frac{T_c}{4\tau D^2 \nu_0^2 N} \tag{3.9}$$

or, rearranging to arrive at an expression more commonly used in literature:

$$\sigma_y(\tau) = \frac{\alpha}{Q\sqrt{N_{at}}}\sqrt{\frac{T_c}{\tau}} \tag{3.10}$$

where $Q := \nu_0/\Delta\nu$ is the quality factor, and $\alpha = 1/(2D\Delta\nu)$ is a factor of order one which depends on the excitation method used for interrogation (i.e. Rabi or Ramsey). The QPN instability for a clock cycle time of 1 s using 1000 atoms interrogated with Rabi pulse of length 300 ms is shown in figure 3.7.

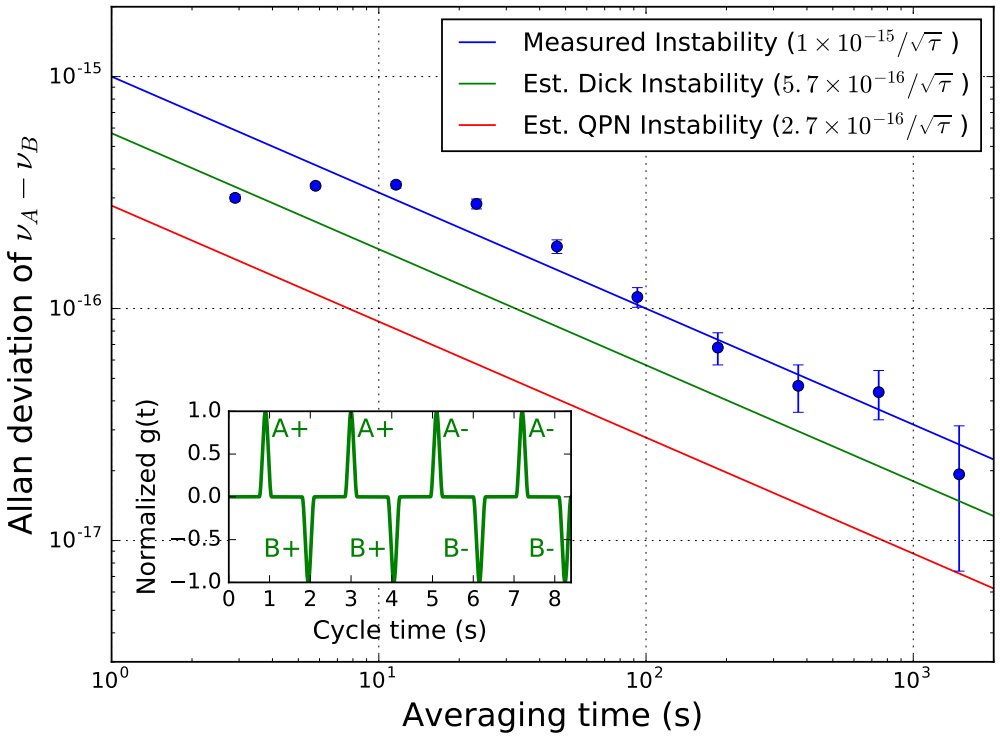


Figure 3.7: Allan deviation of the frequency difference for two clocks operated in an interleaved fashion. For this measurement the cycle dead time was 750 ms and the probe time was 300 ms. For comparison are the theoretical QPN instability for 1000 atoms and the estimated instability arising from the Dick effect for a laser with a flicker floor of 2×10^{-16}. **Inset:** Sensitivity function for interleaved operation used to estimate Dick instability of the frequency difference of the two clocks. Each servo consists of locks to the two Zeeman stretched states, denoted by ($\pm$), each of which requires two Rabi pulses to probe the 'high' and 'low' sides of the atomic line. The sensitivity function for clock B is inverted as we are computing the quantitative $\nu_A - \nu_B$.

As an important aside, ideally to measure stability one would build two independent clocks and directly compare their relative frequency difference. Unfortunately with a single system we do not have this luxury. Instead, we estimate the instability by interleaving clock cycles which are used to steer two independent servos (clock A and B). The instability of the frequency difference $\nu_A - \nu_B$ between the servos is computed which, under the assumption the noise between measurements is uncorrelated, is twice the instability a single clock (one factor of $\sqrt{2}$ is gained because a single clock

will have half the cycle time and the other factor of $\sqrt{2}$ is because both clocks A and B contribute equal noise to the frequency difference). However, we will see in certain situations the assumption that the noise is uncorrelated does not hold and therefore for clarity I will always plot the Allan deviation of $\nu_A - \nu_B$ when estimating the clock stability.

As we can see from figure 3.7, the observed stability of the clock based on interleaved operation is worse than what is estimated by QPN noise. To explain the discrepancy, we need to take into account instability arising from laser frequency noise. One might assume that 'fast' laser noise (i.e. noise with a period much shorter than probe length) would have little effect on the excitation probability given it should average to zero over the interrogation time. Therefore, the job of the atomic servo would then be to correct slow changes of laser frequency with a long term stability set by how well it can resolve such excursions given technical – and quantum-projection noise. Unfortunately, the finite sampling of the laser frequency arising from 'dead time' in the cycle results in an aliasing effect causing noise at Fourier frequency matching the clock cycle to be down converted to slowly varying changes in the excitation probability as shown in figure 3.8. Misinterpreting this signal as arising from changes in the mean LO frequency, the atomic servo inadvertently degrades the stability of the laser through the feeding back of erroneous corrections.

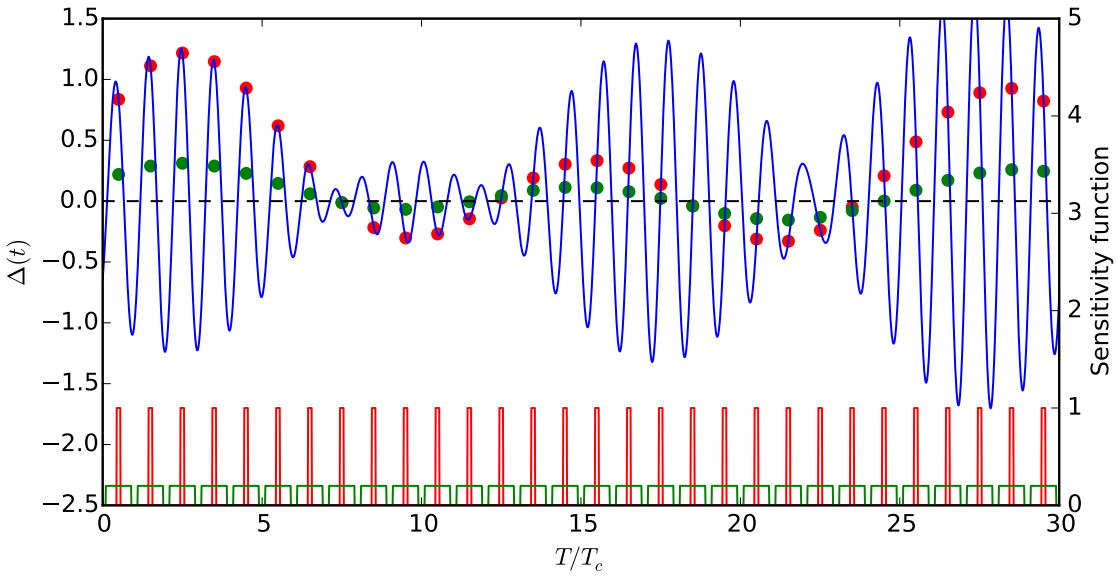


Figure 3.8: Aliasing of random local oscillator noise, focusing in particular on Fourier frequencies centered at $f = 1/T_c \pm T_c/10$ which cause instability on long time scales. Red sequence shows the very brief sampling of the local oscillator frequency which leads to large fluctuations in the lock offset (red points). In the green trace, the probe time is extended such that frequency noise is better averaged to zero (green points) leading to an improved stability. Figure credit: Richard Hobson, adapted from [160].

This effect was first explained by [34] and is well understood in the context of lattice clocks [116]. To see why this source of instability arises, we first define a sensitivity function $g(t)$ which relates fluctuations in the laser frequency $\Delta(t)$ over the length of the experimental cycle to a change in the final excitation fraction δF which is defined as:

$$\delta F = \frac{1}{2}\int g(t)\Delta(t)dt. \tag{3.11}$$

Given that the instability resulting from the Dick effect is dominated by noise at Fourier frequencies which synchronise with the clock cycle, it is constructive to Fourier decompose $g(t)$ using basis functions corresponding to harmonics of the clock cycle. The sensitivity to the mean detuning is given by $\tilde{g}_0 = \frac{1}{2}\int g(t)dt$ and provides the relevant information for steering the atomic servo. The other Fourier coefficients, equal

to $\tilde{g}_n = \frac{1}{2}\int g(t)\cos(\frac{2\pi n}{T_c})dt$, convert laser frequency noise at the various harmonics into fluctuations of the error signal. Given it is only a narrow band of frequency noise centered around each harmonic that causes the instability, we can approximate the noise spectrum as being effectively 'white' in nature. Therefore, we can think of the Dick effect resulting from infinitely many independent white noise processes whose amplitude is scaled by the relative sensitivity function for that frequency of noise. Under this approximation, the Allan deviation of the Dick effect induced instability is given by:

$$\sigma^2_{Dick}(\tau) = \frac{1}{\tau g_0^2}\sum_{m=1}^{\infty}|g_n|^2 S_{LO}(\frac{2\pi n}{T_c}) \tag{3.12}$$

where LO noise is given in terms of the single-sided spectral power density function $S_{LO}(f)$. For a more complete derivation of equation 3.12, see [51]. The estimated Dick instability for interleaved clock comparison when using a cavity with flicker noise limited floor at 2×10^{-16} is also shown in figure 3.7.

Fortunately, there are ways to mitigate the frequency instability arising from the Dick effect. If one is interested in comparing clocks that share the same local oscillator, by synchronising the spectroscopy pulses the Dick effect induced excitation noise will be identical for both systems and will therefore cancel. Such an approach improved the stability of the comparison of two strontium clocks by an order of magnitude [141], and has more recently been extended for comparing clocks of different atomic species using a common universal master oscillator transferred to the different clock transition wavelengths using a frequency comb [99]. Going beyond comparing clocks, to realise a frequency reference free of any Dick effect noise one has to eliminate the sensitivity to laser frequency noise. One approach is to engineer a flat sensitivity function by interleaving two clocks. When one system is cooling atoms and loading the lattice, the other is performing spectroscopy. If the time needed for cooling matches the probe time, there will be no dead time during which the clock frequency

is going unmeasured. Such a 'zero dead time' scheme has recently demonstrated for a ytterbium system [131] and is estimated to have a stability of $6 \times 10^{-17}/\sqrt{\tau}$ (however to verify such performance would require two more systems to realise an additional 'zero dead time' clock which has yet to be realised).

3.3 Uncertainty budget for Sr1

The motivation for realising a clock referenced to an atomic transition is the underlying assumption that all atoms of a given type are indistinguishable. Thus, all clocks using a particular element should 'tick' at the same rate. Unfortunately, this only holds true if the clocks are operating under the same environmental conditions as the interaction between the atom and its surroundings inevitably perturbs the resonance frequency. As a result, clocks that use the same species are no longer guaranteed to be identical and it becomes the job of horologists[5] to identify and characterise the possible sources of systematic frequency shifts. Such a process cannot be done perfectly and after completion one is left with a list of corrections with corresponding uncertainties[6]. As an example, the uncertainty budget for our ^{87}Sr clock is shown in table 3.1.

Determining a frequency correction is a two step process. First, the operating parameters and environmental conditions (e.g. the temperature, stray electromagnetic fields or probe laser intensity) need to be measured. Next, the operating conditions are converted to perturbations of the clock transition using the atomic sensitivities of the clock states to such parameters (i.e. their atomic polarisabilities or magnetic

[5] A maker of clocks or watches

[6] One typically works in fractional units where the shifts are normalised by the frequency of the clock. For the performance of the clock, it is the *fractional* uncertainty which is important – the timing error introduced by a 0.1 Hz error for microwave frequency standard operating at 9 GHz will be worse then a 0.1 Hz offset in optical frequency standard operating at 100 THz. In the case of the strontium $5s^2\ ^1\mathrm{S}_0 - 5s5p\ ^3\mathrm{P}_0$ clock transition, the scale factor to convert to fractional units is 2.33×10^{-15}, i.e. a 1 Hz uncertainty corresponds to a fractional frequency uncertainty of 2.33×10^{-15}.

Table 3.1: Systematic shifts and error budget for the ^{87}Sr lattice clock

Systematic effect	**Correction** ($\times 10^{-18}$)	**Uncertainty** ($\times 10^{-18}$)
Blackbody radiation	4856	6
1st-order (E1) lattice Stark	1.5	5
Higher-order lattice Stark	-1.6	4
Probe Stark	0	0.4
DC Stark	<0.1	<0.1
2nd-order Zeeman	287	2
Density	0	4
Background gas collisions	0.2	2
Servo offset	0	1
1st-order Doppler	0	<0.1
Total	5215	10

moments). Both steps introduce systematic uncertainties in the applied correction. In some cases when we have direct control of the environmental perturbation, we can vary it to map out its influence on the clock frequency in a similar way to how we interleaved clock instruments to evaluate stability. In particular, if the shift is a linear function of the perturbation and vanishes when it is switched off (e.g. the Stark shift induced by the lattice intensity), we can use extrapolation techniques to estimate its magnitude by the following recipe:

1. Perform interleaved measurements comparing the clock frequency at two different setpoints (P_1 and P_2) for the parameter of interest.

2. Measure the frequency difference $\delta\nu$ between the two setpoints.

3. Repeat the process several times to reduce the statistical uncertainty σ in the estimate of $\delta\nu$[7].

[7]The uncertainty as a function of averaging time is estimated by computing the Allan variance.

4. The shift $\delta\nu_0$ at any other setpoint P_0 can now be calculated using linear extrapolation as follows:

$$\delta\nu_0 = \delta\nu \frac{P_0}{P_2 - P_1} \tag{3.13}$$

5. The uncertainty σ_0 in the estimate for $\delta\nu_0$ is given by:

$$\sigma_0 = \sigma \frac{P_0}{P_2 - P_1} \tag{3.14}$$

It is therefore advantageous to maximise the lever arm, defined as $P_2 - P_1$, to minimise the uncertainty in the correction. We follow this procedure when evaluating several of the shifts in table 3.1.

3.3.1 Blackbody radiation shift

The leading source of uncertainty for many Sr lattice clocks results from the interaction between the atoms and the ambient BBR emitted from the surrounding environment [101, 38, 65]. For a strontium lattice clock operating at room temperature, the resulting shift induced by BBR is approximately 2 Hz making it by far the largest frequency correction applied to the clock. To characterise this frequency shift, we first need to estimate the radiative energy density resulting from BBR at the position of the atoms. Next, the differential polarisability of the clock states must be known in order to calculate the frequency shift induced by this radiation. Both of the parameters, the absolute temperature and the polarisability coefficients, also introduce uncertainty.

3.3.1.1 Theoretical treatment of the BBR shift

To calculate the shift of the clock transition induced by BBR, we return to the polarisability calculation shown in figure 2.4. We are interested in the region centred

around 10 µm that coincides with the peak of the BBR spectrum at room temperature. It is clear there are several transitions that affect the polarisability in this region with the most significant being the $5s5p\ ^3P_0 \rightarrow 5s4d\ ^3D_1$ transition at 2.6 µm. This illustrates why the BBR sensitivity of the Sr clock transition is high compared to other lattice clock species such as mercury (30× smaller fractional frequency shift), cadmium (18× less), magnesium (13× less), and ytterbium (2× less) which have principle atomic transitions at shorter wavelengths [95][8].

The resulting frequency shift is calculated by integrating the spectral energy density of the BBR field weighted by the differential polarisability $\Delta\alpha(\omega)$ of the clock states as:

$$\delta\nu_{BBR} = \frac{1}{2\epsilon_0 ch}\int \Delta\alpha(\omega) I(\omega) d\omega. \tag{3.15}$$

$I(\omega, T)$ is the electric field intensity per unit angular frequency emitted by a blackbody at a temperature T is given by Planck's equation:

$$I(\omega, T) = \frac{\hbar}{\pi^2 c^2} \frac{\omega^3}{e^{\frac{\hbar\omega}{k_B T}} - 1}. \tag{3.16}$$

At room temperature, the peak BBR wavelength is considerably longer than the resonance wavelengths of all the atomic transitions. As a result, a reasonable approximation for the BBR shift can be computed by performing a DC expansion of polarisability about $\omega = 0$ of the form $\Delta\alpha(\omega) = \Delta\alpha(0)(1 + \frac{1}{2}\frac{\Delta\alpha''(0)}{\Delta\alpha(0)}\omega^2 ...)$. Under this

[8]From a practical standpoint, the longer wavelengths required for cooling, trapping, and clock spectroscopy of Sr are convenient as they are readily accessible with diode lasers making for a more simple and reliable operation of the clock.

Table 3.2: Coefficients in eqn 3.17 for the BBR shift in atomic Sr [92, 126, 101]

	β_{st}	$\beta_{\mathrm{st}}\eta_1$	$\beta_{\mathrm{st}}\eta_2$	$\beta_{\mathrm{st}}\eta_3$	Total $\Delta\nu_{\mathrm{dyn}}$
Shift coefficient (Hz)	−2.13023	−0.1291	−0.01616	−0.00343	−0.1486
Uncertainty (Hz)	0.00006	0.0006	0.00007	0.00002	0.0007
Shift coefficient ($\times 10^{-18}$)	−4962.93	−300.7	−37.6	−7.97	−346.3
Uncertainty ($\times 10^{-18}$)	0.14	1.4	0.2	0.03	1.6

assumption, equation 3.15 becomes:

$$\begin{aligned}\delta\nu_{BBR} &\cong \frac{\Delta\alpha(0)}{2\epsilon_0 ch}\int I(\omega)d\omega + \frac{\Delta\alpha''(0)}{4\epsilon_0 ch}\int \omega^2 I(\omega)d\omega... \\ &\triangleq \beta_{\mathrm{st}}(\frac{T}{T_0})^4[1+\eta_1(\frac{T}{T_0})^2+\eta_2(\frac{T}{T_0})^4...]\end{aligned} \tag{3.17}$$

where T_0 is an arbitrary reference temperature taken to be 300 K. The majority of the BBR shift results from the 'static' polarisability. However, the transition at 2.6 µm contributes a non-negligible 'dynamic' correction at the low 10^{-16} level. The BBR coefficients for the static shift (β_{st}) and dynamic corrections ($\eta_1, \eta_2...$) are listed in table 3.2.

3.3.1.2 Experimental treatment of the BBR shift

With the theoretical framework describing the effect of BBR on the clock states understood, all that remains for characterising the BBR shift is determining the thermal energy density at the position of the atoms. If we hope to constrain the BBR shift for our lattice clock operating at room temperature to below 1×10^{-17}, it will require that the temperature uncertainty is less than 150 mK.

Many approaches are taken to tackle this problem such as placing thermometers inside the vacuum chamber to measure the temperature [101], operating the clock at cryogenic temperatures to reduce the frequency shift [148], or engineering temperature

controlled in-vacuum enclosures to provide a well-defined BBR environment [10]. We instead adopt the approach of placing external temperature sensors on the chamber and feeding the readings into a model to estimate the temperature at the position of the atoms. We use eleven Pt-100 temperature sensors attached to the outside of the science chamber that are calibrated to an absolute accuracy of $\pm 10\,\mathrm{mK}$ by the NPL temperature group.

We employ two models to estimate the BBR shift. The first does not use any *a priori* knowledge of the geometry of the vacuum chamber or the emissivities of the various surfaces. We assume only that the representative temperature of the BBR environment as seen by the atoms lies somewhere between the minimum and maximum temperatures measured by the sensors. Following the procedure described in [38], we assume that the probability density for the representative temperature is uniformly distributed between the two extrema. For such a distribution, the expected temperature is $\frac{T_{\mathrm{max}}+T_{\mathrm{min}}}{2}$ with an uncertainty of $\frac{T_{\mathrm{max}}-T_{\mathrm{min}}}{\sqrt{12}}$. The readings from the other nine sensors go unused in the model and only provide reassurance that one has not missed a hot or cold spot on the chamber. Assuming neither of the two temperature extrema go undetected, the maximum possible temperature error (in units of standard deviations σ) is $\sqrt{12}/2 = 1.7$. One can imagine extreme cases where this could occur, *e.g.* if one extrema has a very large emissivity compared to the rest of the chamber and therefore dictates the BBR environment. To ensure that such a situation does not arise we develop a second model that takes into account all temperature readings, the geometry of the chamber, and the emissivities of the various materials to estimate the BBR at the atoms' position. In this model, the average energy density $I_{\mathrm{avg}}(\omega, T)$ is computed by summing the radiation emitted by the various surfaces surrounding the atoms. The contribution from each surface is weighted by its emissivity ϵ and the solid angle A subtended with respect to the atoms' position [10]. The average

radiative energy density is given by:

$$I_{\text{avg}}(\omega, T) = \frac{\sum \epsilon_i A_i I(\omega, T_i)}{\sum \epsilon_i A_i} \tag{3.18}$$

where the sum is over every surface seen by the atoms and T_i is estimated using the temperature sensors. Care is taken to place sensors on every viewport as they have the largest emissivity and contribute most to the sum. The viewports are made from fused silica which is opaque for wavelengths above 4 µm. This strongly suppresses leakage of ambient BBR through the viewports as the radiation spectrum at room temperature falls off sharply below 10 µm. For both models, the uncertainty is reduced by decreasing the temperature inhomogeneity across the chamber. This fact is obvious for the uniformly distributed probability (UDP) model, but is not immediately self-evident for the solid angle approach. For the solid angle model, the improved temperature homogeneity reduces the sensitivity of the estimated BBR to variation in emissivity from surface to surface. In the limit of a perfectly uniform temperature, the radiation exiting from a surface is independent of emissivity as the emitted and reflected BBR are identical. Temperature gradients across the chambers primarily arise from a mismatch between the ∼15 °C water used to cool the MOT coils and the ambient room temperature at ∼21 °C. Under typical operation, this results in a maximum temperature difference across the chamber of between 0.7 to 1 K corresponding to a BBR correction uncertainty between 1×10^{-17} to 2×10^{-17} when applying the UDP model. We now heat the cooling water to match room temperature which results in the gradients reducing to 300 mK, limited now by the heating of the water as it propagates from the outer to the inner windings of the MOT coil.

To ensure the BBR correction was independent of the model used, we operated the clock continuously for several days and recorded the temperature estimates generated by both methods for comparison. Figure 3.9 shows the BBR corrections for both models over this time period. The discrepancy between the two models is less than

1×10^{-18}. The fractional uncertainty for UDP model is 6×10^{-18} while the uncertainty for the solid angle approach is 4×10^{-18}. We report the more conservative uncertainty for the BBR correction in table 3.1.

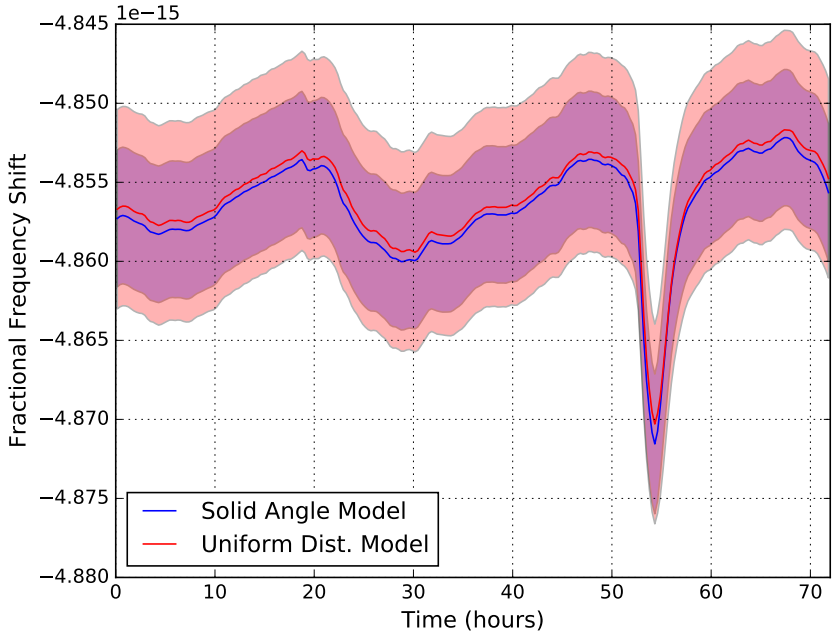


Figure 3.9: Comparison of the frequency corrections and corresponding uncertainty generated by the UDP and solid angle models. The disagreement is $< 1\times10^{-18}$ while the uncertainty is only marginally better for the solid angle approach. The temperature spike at 55 hour point resulted from a change in the processed cooling water temperature. The cooling water is a closed loop system shared by many experiments. When the heat load is changed, the servo controlling the process cooling water cannot responding instantly resulting in brief excursions of the water temperature.

3.3.2 Lattice shifts

To evaluate perturbations of the clock frequency introduced by the lattice, we use equation 3.13 given the linear scaling of the **E1**-Stark shift with intensity. We interleave clock measurements at two different trap depths, U_{ref} and U_L, and measure the differential frequency shift. Care is taken to ensure the density shift, which increases with trapping potential, is the same for both lattice intensity setpoints. This requires

us to scale the atom number with trap depth in accordance to [103]:

$$N_{\text{ref}} = N_L(\frac{U_L}{U_{\text{ref}}})^{3/2}. \tag{3.19}$$

Another important consideration when calculating the lattice shift is properly estimating the effect the atomic temperature has on trap depth. Intuitively, hotter atoms in higher motional states are more delocalised across the trap and thus experience a reduced time averaged lattice intensity. We therefore use an effective trap depth when extrapolating the lattice shift that takes into account thermal effects. In equilibrium, the equipartition theorem states that the thermal energy is divided evenly between the kinetic and potential energy degrees of freedom. Thus, an atom sitting in a 3D harmonic potential should possess a potential energy of $\frac{1}{2}kT$ in each direction. This energy comes at the expense of trap depth and is how we arrive at our effective trap depth given by:

$$U_{\text{eff}} = U_o - \frac{1}{2}(kT_z + 2kT_r) \tag{3.20}$$

where T_z and T_r are the axial and radial temperatures of the atomic sample. We are also careful to keep the lattice intensity the same when loading from the second stage MOT and only ramp to the final setpoint after the cooling beams are off to avoid loading hotter atoms when using a deeper trap.

For the lattice extrapolation, we operate at effective trap depths of $U_{\text{ref}} = 27E_r$ and $U_L = 128E_r$ where E_r is the recoil energy of the lattice given by $\frac{\hbar^2k^2}{2m}$. With the lattice tuned to the 'magic'-wavelength, the fractional frequency difference measured between interleaved atomic servos operating at the two lattice setpoints was $-1.2(11) \times 10^{-17}$. The uncertainty is purely statistical and was computed using the overlapping Allan deviation.

Therefore, under normal operating conditions corresponding to a trap depth of

$U_0 = 40E_r$, the linear extrapolated lattice shift is $-4.8(43) \times 10^{-18}$. The uncertainty is a combination of the statistical error in measurement and the uncertainty in the sample temperatures used for calculating the effective trap depth. However, this treatment of the shift neglects possible higher order perturbations that can cause the Stark shift to deviate from the assumed linear scaling with intensity. As the inaccuracy of the lattice induced fractional frequency shift pushes below 1×10^{-17}, higher order effects need careful characterisation [18]. Such a task proves challenging as the resulting shifts no longer scale linearly with trap depth and depend strongly on operating conditions.

3.3.2.1 Higher-order lattice shifts

In section 2.3, we evoked the dipole approximation to derive the trapping potential induced by the lattice and calculated its effect on the clock states using second order perturbation theory. Under such approximations, the lattice Stark shift can be canceled by operating the clock under a specific set of conditions (magic wavelength, lattice polarisation, interleaved interrogation $\pm m_F$ states, etc.). Furthermore, any perturbation resulting from deviations from these ideal conditions can be constrained using interpolation methods as the resulting shifts scale linearly with intensity. We justified these assumptions as the terms proportional $|E|^2$ dominate the interaction between the atoms and lattice.

Remaining within the E1 approximation for the trapping potential, the next non-zero perturbation term is the hyperpolarisability which scales as $|I|^2$ [9]. This results in a frequency shift of the clock transition given by:

$$\delta\nu_\gamma = \gamma(U_0/E_r)^2 \tag{3.21}$$

[9]The parity of the dipole operator results in terms proportional to odd powers of electric field to vanish.

where γ is the hyperpolarisability coefficient. Physically, this shift arises from two photon, dipole allowed, transitions from either of the two clock states driven by the lattice light. As a result, the states contributing most to this effect are those with resonant wavelengths near $\lambda_{\text{magic}}/2$. In the case of strontium, the dominant transitions are $|5s5p\,{}^3P_0\rangle \rightarrow |5s4f\,{}^3F_2\rangle$ at $\lambda/2 = 818.57/2$ nm and $|5s5p\,{}^3P_0\rangle \rightarrow |5s7p\,{}^3P_0\rangle$ at $\lambda/2 = 796.3/2$ nm. The hyperpolarisability varies with lattice polarisation as it modifies the coupling strength between the states. For ytterbium, mercury, and cadmium the presence of several transitions close to the two-photon resonance allows the differential polarisability to be completely suppressed by operating the lattice at a 'magic' ellipticity. Sadly, such a fortuitous effect does not exist from strontium. The hyperpolarisability coefficient for strontium has been measured for a linearly polarised lattice to be $\gamma = 0.46 \pm 0.18$ μHz [161].

We must also consider lattice shifts resulting from off resonant multipolar transitions beyond the dipole approximation. The first two are magnetic dipole (M1) and electric quadrupole (E2). As the M1 and E2 transitions are driven by $|B|^2$ and $|\partial_i E|^2$ respectively, one might assume their corresponding shifts would scale linearly with intensity. However, this is only true for the resulting Stark shift. The M1 and E2 transition also modify the trapping potential which shifts the energy of the motional states [139]. If the dipole potential is given by $-U_{E1}\cos^2(kz)$, the trapping potential resulting from M1 and E2 transitions will be proportional to $(U_{M1} + U_{E2})\sin^2(kz)$. Working within the harmonic trapping approximation, the axial trapping frequency for the ground (ω_g) and excited (ω_e) state become:

$$\omega_{e,g} = \sqrt{\frac{(U_{E1,(e,g)} + U_{M1,(e,g)} + U_{E2,(e,g)})k^2}{m}}. \tag{3.22}$$

The two trapping frequencies are now different as it is impossible to find a magic wavelength which makes U_{E1}, U_{M1}, and U_{E2} the same for both clock states. If we

operate the clock on the transition from $|g,n\rangle \rightarrow |e,n\rangle$, the change in resonance frequency $\Delta\nu_n$ scales with motional state n as:

$$\Delta\nu_n = (n + 1/2)(\omega_e - \omega_g). \tag{3.23}$$

Therefore, the combined **M1** and **E2** shift scales as $\sqrt{U_0}$ and can be re-written as:

$$\frac{\Delta\nu_n}{\nu} = (n + 1/2)\epsilon\sqrt{U_0}. \tag{3.24}$$

The value of ϵ has been experimentally measured to be $\epsilon = (0 \pm 7.1 \times 10^{-19})E_r^{-1/2}$ [161]. Therefore, under standard operating conditions we expect higher order lattice shifts to result in a fractional frequency shift of $1(4) \times 10^{-18}$. Returning to the **E1** component of the lattice shift, the contribution of higher order effects must be subtracted from total measured lattice shift to extract the linear component of the lattice Stark shift. Taking into account higher order lattice shifts, the estimated linear Stark shift at $U_0 = 40\,E_r$ is $-1(5) \times 10^{-18}$. Along with statistical uncertainty, the quoted error includes the contribution introduced by subtracting the high order lattice shift corrections and taking into account thermal effects when calculating the trap depth.

3.3.3 Density shift

The improved signal-to-noise ratio offered by large ensembles of cold atoms improves stability and is touted as a major advantage of lattice clocks. However at high densities with many atoms loaded into each lattice site, strong interactions can give rise to collisional frequency shifts. Scattering between atoms will impart a phase shift on the atomic wavefunction. The rate at which these collisions occur therefore sets the density dependent collisional frequency shift, i.e. the rate phase accumulates from these collisions.

Calculating the collisional phase shift involves expanding the incoming and scattered outgoing wavefunctions in terms of partial waves (s-wave, p-wave, d-wave...). Excluding s-wave collisions, the scattering potential associated with each partial wave contains a centrifugal barrier term that suppresses collisions at low temperatures[10]. As such, interactions between the cold atoms are dominated by s-wave collisions. For fermionic strontium, the Pauli exclusion principle dictates the two-particle wavefunction must be asymmetric under exchange, thus forbidding collisions of even parity (s-wave, d-wave...) between atoms in the same state [158]. Therefore, the collisional frequency shift for a sample of ultra-cold spin polarised fermionic strontium with every atom in the same internal state should be strongly suppressed. However, for both ^{87}Sr [137, 138] and ^{171}Yb [70] density shifts have been resolved, motivating a more detailed investigation. The main mechanisms that allow collisions to occur are as follows:

1) **Incomplete spin polarising**: The Pauli exclusion principle does not prevent collisions between atoms in different m_F states. To ensure a pure spin polarised sample, we implement a 'clear out' stage. This involves applying a π-pulse to shelve atoms in the correct sub-level into the upper clock state. Next, atoms in other m_F levels that remain in the ground state are removed from the lattice with a pulse of 689-nm light on resonance with the $^1\mathrm{S}_0 \rightarrow ^3\mathrm{P}_1$ transition. We then operate the clock in an inverted manor in which the clock pulse de-excites atoms to the ground state. The clear out stage also has the benefit that line pulling—the effect of the off resonance excitation spectrum of other m_F states on the error signal—is suppressed.

2) **Probe inhomogeneity**: Consider two atoms, both in the internal ground state $|gg\rangle$ and external motional states n_1 and n_2. The Pauli exclusions principle requires the two body wave function be asymmetric under exchange. Given the internal state is symmetric, the external state must be asymmetric and therefore

[10]For p-wave collisions, the centrifugal barriers is $\sim$75 µK, far greater than the atomic sample temperature of $\sim$2 µK [86]

immune to even order partial wave collisions. During spectroscopy the clock laser induces Rabi oscillation between the internal ground and excited states. If the Rabi frequency is inhomogeneous across the sample, the internal state of these two atoms will evolve differently. The internal state will no longer be purely symmetric resulting in the external state acquiring a symmetric component. Therefore, the inhomogeneous probe excitation opens a s-wave collisional channel which can give rise to a collisional frequency shift [117, 40]. This shift can be suppressed by reducing the temperature of the sample as one cause of the Rabi frequency inhomogeneity is the spread in Lamb-Dicke parameter across different motional states. Furthermore, ensuring the clock beam is aligned to the lattice further reduces the sensitivity to the radial motional state.

3) ***p*-wave collisions**: The picture of a centrifugal barrier preventing high order partial wave collisions at low temperatures is overly simplistic. Collisional shifts associated with p-wave interactions have been experimentally measured for Yb [71, 81] and Sr [167]. The detailed treatment that accurately reproduces the spectroscopic line-shapes requires a many-body treatment beyond the standard mean field approach [118]. Along with a frequency shift, the inelastic p-wave collisions lead to a loss of coherence resulting in broadening of spectral lines and loss of contrast. For moderate densities and short spectroscopy times used in our system, the decoherence resulting from inelastic p-wave collisions is negligible and our primary concern is the resulting frequency shift.

Regardless of the underlying mechanism which gives rise to the collisional frequency shift, its magnitude will be proportional to the rate of collisions and will therefore scale linearly with density. Given this scaling, we can characterise the resulting systematic shift via standard linear extrapolation techniques following the same procedure as with the lattice shift. To minimise the collisional shift and maximise the lever arm for the linear interpolation, we operate the clock with minimal atom

number. For as few as ∼200 atoms, the quantum projection noise will be less than the Dick instability and will therefore not compromise the performance of the clock. Using a lever arm of $3.2N_{\mathrm{ref}}$ where N_{ref} is the typical operational atom number, we measured a density shift of $-0.3 \pm 1.2 \times 10^{-17}$. Therefore, the extrapolated collisional frequency shift when operating the clock with a density N_{ref} is $-0.9 \pm 4 \times 10^{-18}$. Given the resulting frequency shift is consistent with zero, we include only the uncertainty in table 3.1 and apply no density correction.

At our clock's present level of systematic inaccuracy, the uncertainty caused by the density shift is not limiting. However, looking ahead to future systems with targeted inaccuracies below 1×10^{-18}, such density shifts become significant and must be addressed. A simple solution would be to use less atoms, however this strategy has the obvious drawback of compromising stability – especially for a zero-deadtime clock. Instead, we could elect to use a larger lattice waist (for example the system reported in [102] uses a 160 µm waist). However, to reach sufficient trap depths at such large mode waists requires an enhancement cavity. This in turn puts practical limits of the maximum achievable waists (i.e. long cavities are unstable, large radius of curvature mirrors need to very well aligned, and operating near an instability point is challenging). As a result, such a strategy is unlikely to reduce the density shift by more than an order of magnitude.

Another solution, paradoxically, is to increase interaction strengths by tightly confining atoms in a higher dimensional lattices. By using a two dimensional lattice to increase the interaction strength, the authors in reference [138] observed suppression of collisional shifts in agreement. As a complete explanation of such effects requires a full quantum treatment including mean-field effects, but semi-qualitatively the density shift arises from virtual exciting to non-fully symmetric states, the occupation of which is suppressed as their interaction energy increased. This idea has been extended further to three dimensional lattice [20]. In such a system, the interaction energies

between atoms trapped on the same lattice site perturbs the transition by hundreds of hertz– making it possible to complete eliminate collisional density shifts by only addressing lattice sites occupied by single atoms.

3.3.4 Background gas collisional shift

Collisional frequency shifts can also arise from scattering with background gas atoms present in the vacuum chamber. Given the large kinetic energy imparted to the trapped strontium atoms from such events, these collisions manifest themselves as a loss of atoms from the lattice. Therefore, the lattice lifetime (Γ_{vac}^{-1}) is a direct measure of the background gas collision rate. All that remains is to estimate the phase shift imparted to the atoms from such collisions. For that we follow the procedure outlined in [41] to arrive at the following expression for the background gas collisional shift:

$$\frac{\Delta\nu_{\text{backgnd}}}{\nu_0} \approx \frac{1}{\nu_0}\frac{\Gamma_{\text{vac}}}{\pi}\frac{\Delta C_6}{13.8 C_6} \tag{3.25}$$

where ΔC_6 is the difference in C_6 coefficient for the ground and excited state between strontium and the various background gas atoms. Using the theoretical C_6 coefficient for H_2 [94], the most probable vacuum contaminant resulting from the outgassing of the steel chamber, and the measured lattice lifetime of 8 s from figure 3.10, we calculate a background gas collisional shift of 2×10^{-18}. Fortunately, the C_6 coefficients between Sr and other atoms likely to be found in the vacuum (H, H_2, He, and Ar) do not differ by more than a factor of 2. Therefore, we assign an uncertainty of 2×10^{-18} to the background gas collisional shift to account for possible collisions with other atomic species. This uncertainty also encompasses the 10% uncertainty in the theoretical C_6 coefficient.

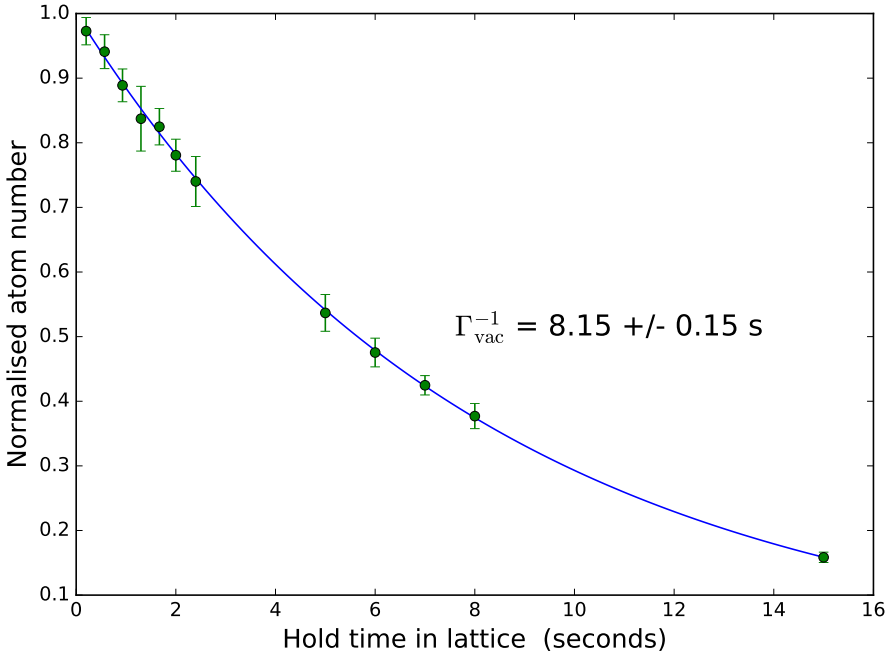


Figure 3.10: The lattice life is measured by loading atoms into the lattice and waiting a predetermined amount of time before imaging those that remain. This procedure is repeated for different hold times to map out the loss rate in the lattice. We observe no dependence of the trap lifetime on the lattice intensity indicating heating is negligible.

3.3.5 Quadratic DC-Stark shift

Uncharacterised DC electric fields can severely impact the accuracy of an atomic clock. For the clock transition in strontium, an electric field of $570\,\mathrm{Vm}^{-1}$ yields a DC Stark shift of of 2×10^{-15} [92]. Where dielectric surfaces are close to the atoms, shifts as large as $1{\times}10^{-13}$ have been observed [77]. Measurable DC Stark shifts have now been observed in strontium clocks built by JILA [103], which were attributed to residual charge on viewports, and PTB [38], caused by a differential voltage across the MOT coils. The DC stark shift of the clock transition depends on the static differential polarisability of the clock transition $\Delta\alpha(0)$ and the electric field $|\mathbf{E}|$. It is given by [92]:

$$d\nu_{DC}/\nu_0 = \frac{\Delta\alpha(0)|\mathbf{E}|^2}{2h\nu_0} = \frac{|\mathbf{E}|^2}{(1\mathrm{V/cm})^2)} \times 7.1705(2) \times 10^{-17}. \tag{3.26}$$

This differential polarisability is exactly what gives rise to the static BBR shift and has been precisely measured. While steps must be taken to reduce this residual electric field seen by the atoms, such as Faraday shielding [10], or UV discharge of dielectric surfaces [111], a characterisation of the remaining field is necessary.

3.3.5.1 Standard methods for DC-Stark shift evaluation

Characterisation of the DC-Stark shift is typically done by direct spectroscopy of the clock transition with an externally applied electric field as outlined in [88], see figure 3.11. First the frequency difference, $d\nu_+$, between interleaved measurements with and without an applied bias field $|\mathbf{E}_{bias}|$ is measured. The frequency shift will depend on magnitude background field $|\mathbf{E}_{bg}|$ and is given by:

$$d\nu_+ = \frac{\alpha(0)|\mathbf{E}_{bg} + \mathbf{E}_{bias}|^2}{2h}. \tag{3.27}$$

Repeating this measurement, but with the polarity of the applied field reversed will produce a shift given by:

$$d\nu_- = \frac{\alpha(0)|\mathbf{E}_{bg} - \mathbf{E}_{bias}|^2}{2h}. \tag{3.28}$$

With no residual electric field present ($\mathbf{E}_{bg} = 0$), the quadratic nature of the perturbation implies that $d\nu_+ - d\nu_-$ should be zero. If any frequency discrepancy arises between the two oppositely applied external field measurements, the culprit must be a residual electric field pointing along the same axis as the applied field. In the limit of the applied field being much larger than the background field, the frequency shift due to this residual electric field is given by:

$$d\nu_{bg} = \frac{(d\nu_+ - d\nu_-)^2}{8(d\nu_+ + d\nu_-)}. \tag{3.29}$$

Repeating this procedure for each orthogonal direction fully characterises any back-

ground electric field. The uncertainty in the resolved shift does not rely on any knowledge of the differential polarisability and depends only on the accuracy of the resolved shifts $d\nu_-$ and $d\nu_+$. Therefore, to accurately calibrate the DC Stark shift it is advantageous to use a large applied electric field. Adopting this approach to evaluate the DC Stark shift, we placed electrodes outside two of the axial viewports and one of the radial veiwports to apply electric fields at the atoms' position. Unfortunately, this proved challenging with our system as the metal vacuum chamber proves to be an excellent Faraday shield. Even with $\pm 10\,\mathrm{kV}$ applied to the plates, we failed to observe a DC Stark shift at the 10^{-16} level.

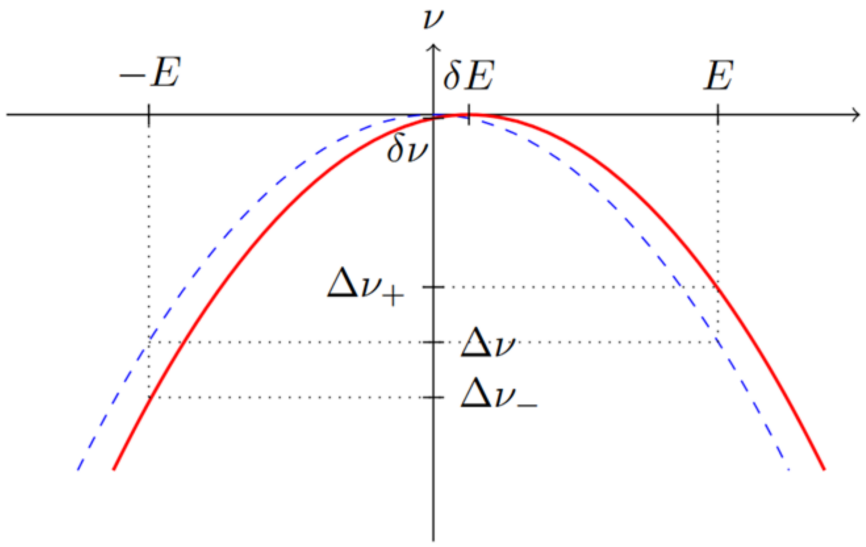


Figure 3.11: To detect a residual electric δE, one interleaves measurements of clock transition frequency with and without an applied electric field E to measure the induced stark shift $\Delta\nu_+$. The measurement is then repeated with field polarity reversed to measure $\Delta\nu_-$. The difference between the two measurements can be used to constrain the residual electric field along the axis of the applied field. Repeating this procedure twice more along the other two orthogonal directions will fully characterise DC-Stark shift arising from background electric fields. Figure adapted from reference [77].

3.3.5.2 Singlet Rydberg EIT spectroscopy for electrometry

Faced with the prospect of breaking vacuum to install electrodes inside that chamber, we began to look for alternate methods to characterise residual electric fields. One option is to use highly excited Rydberg states. With a polarisability eight orders of magnitude larger than that of the clock transition, these states are well suited for AC [5] and DC [109, 143, 37, 1] electrometry. Through discussion with Matthew Jones at Durham University, whose lab has done pioneering research on the Sr Rydberg series, we arrived at a spectroscopy scheme which allows us to characterise any residual electric field with a resolution better than $1\,\mathrm{Vcm}^{-1}$ electrometry using electromagnetically induced transparency (EIT) spectroscopy [14].

3.3.5.3 Rydberg theory: Single electron model and Stark maps

The theoretical Stark maps were calculated and experimentally verified by Matthew Jones' research group at Durham University. Even in divalent atoms such as strontium, where inter-electronic Coulomb interactions lead to perturbations of the Rydberg states [39], it can be shown that accurate wavefunctions [151, 152, 166], and Stark maps [93, 76, 50] can be obtained without recourse to the complex atomic structure calculations required for the clock states [126]. Figure 3.12(a) shows an example Stark map that is compared to experimental measurements taken in an atomic beam apparatus with a well-defined electrode geometry [46]. The data and the model are in quantitative agreement without any adjustable parameters. The predicted Stark map for the higher-lying state used to constrain the field in the lattice clock is shown in figure 3.12(b). In both cases, the shaded band indicates the theoretical uncertainty in the Stark map.

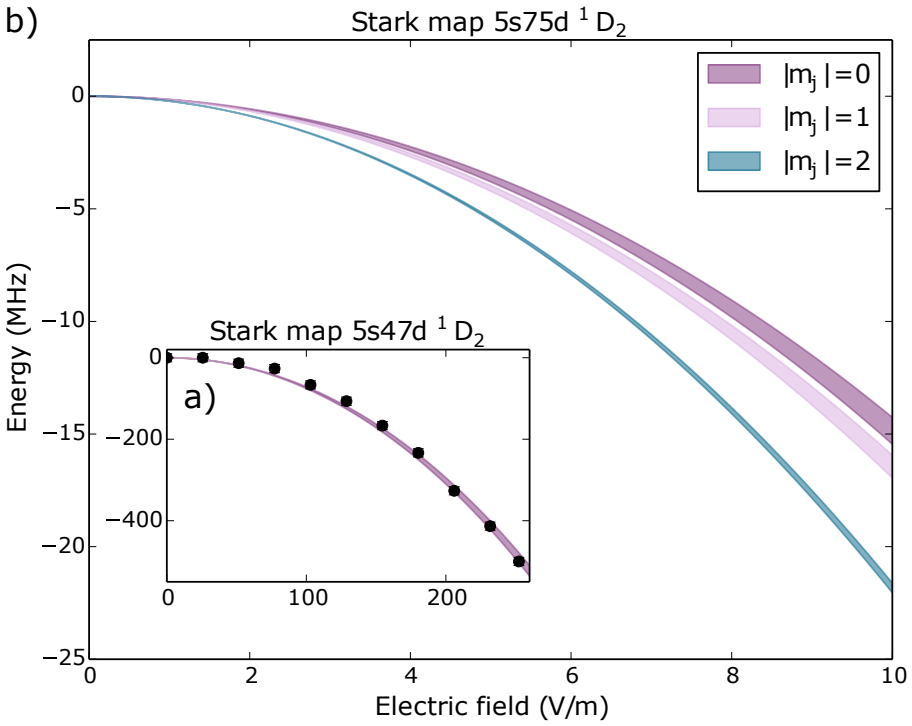


Figure 3.12: (a) Predicted Stark shift of the $5s47d$ ^{1}D$_2$ $m_J = 0$ state compared to experimental data (black dots) without any adjustable parameters (axes units are the same as main plot). (b) Predicted Stark shift of the three $|m_J|$ components of the $5s75d$ ^{1}D$_2$ state. The shaded area represents the uncertainty arising from the experimental determination of the zero-field energy (see text). Figure credit: Paul Huillery

3.3.5.4 Experimental setup for Rydberg electrometry

We follow the same methods used when operating the clock except we elect to use the bosonic ^{88}Sr isotope as it is easier to cool and naturally more abundant. Before implementing the electromagnetically induced transparency (EIT) probe pulse, the cloud of atoms is released from the red MOT for 5 ms, giving time for the magnetic field to settle to the desired bias value, and for the atoms to expand ballistically to a lower density value in order to improve the signal to noise level.

For the implementation of the EIT spectroscopy counter-propagating beams, one resonant with the $5s^2\ ^1\mathrm{S}_0 - 5s5p\ ^1\mathrm{P}_1$ transition at 461 nm and the other with tunable frequency centered at 413 nm, excite atoms to a chosen $5snd\ ^1\mathrm{D}_2$ Rydberg state, see

figure 3.13 for details. The resulting EIT signal is measured using 'lock in' detection of the 461 nm probe beam absorption via modulation of the 413 nm pump beam intensity by an optical chopper. A typical absorption measurement showing the modulated EIT signal induced by the pump beam is shown in the top of Fig.3.14. The probe beam is derived from a commercial frequency doubled diode laser system. Its power and waist, as defined by the $1/e^2$ radius of the intensity profile, are 800 fW and 120 μm, respectively. At this power and atomic number, the probe beam absorption is between $20-40\%$. A home built extended cavity diode laser (ECDL) provides 8 mW of pump light which is focused to an 80 μm waist at the atoms' position. The bottom inset in figure 3.14 shows typical spectra taken at zero magnetic field with and without the applied external electric field.

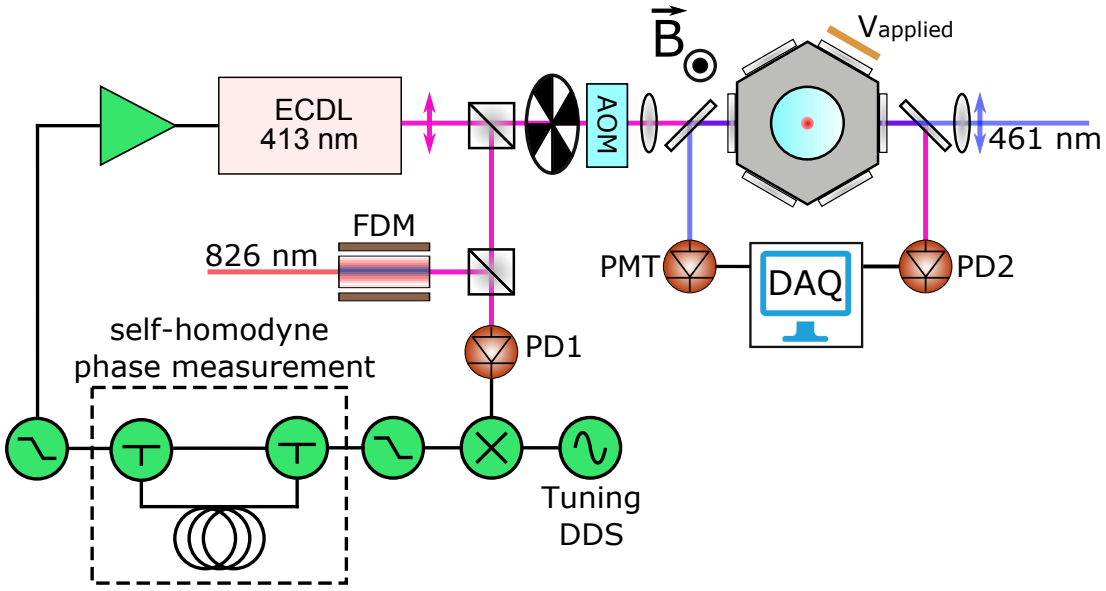


Figure 3.13: Diagram of the optical setup used for Rydberg EIT spectroscopy. The pump beam operating at 413 nm (arrows indicate polarization direction) is sent through an optical chopper and switching AOM before being focused onto the atoms. The atomic absorption of a counter-propagating 461 nm probe beam is detected using a PMT. The intensity of the chopped pump beam is monitored by a photodiode (PD2) and logged by a DAQ. In software, this signal is used to demodulate the PMT recording to extract the EIT signal. Also shown is the direction of magnetic field and the location of the electrode that is used to apply an external electric field. Also included is the A Ti:sapph laser that frequency doubled using a SHG module, see text for more detail. The pump beam is frequency stabilised to the Ti:sapph by sending the RF beat signal produced by photodiode (PD1) in a interferometer and locking the beat frequency such that the two paths interfere destructively. To tune the offset frequency, a DDS is mixed with the beat signal and lowpass filtered before being sent into the interferometer.

The long term frequency stability of the pump and probe beam is maintained to within 10 kHz by locking to a transfer cavity referenced to the 'clock' laser. In the case of the probe beam, the sub-harmonic at 922 nm is directly locked to the cavity. To stabilise the 413 nm pump laser frequency, a commercial Ti-sapphire laser, which is typically used to form the magic wavelength lattice at 813 nm, is first tuned to 826 nm and locked to the transfer cavity. This light is then frequency doubled by a lithium triborate (LBO) crystal [11] to produce 20 μW as needed to generate a beat-note with

[11] Purchased from Eksma Optics, Optolita UAB Mokslininku str. 11 LT-08412 Vilnius Lithuania

the pump beam. The beat-note signal is mixed with a DDS and the intermediate frequency is stabilised using a delay line offset lock scheme [133] via fast feedback to the diode current and slow feedback to the ECDL piezo. The DDS provides the necessary tunability needed for scanning the pump laser frequency.

In order to spectroscopically resolve the Zeeman sub-levels of the Rydberg state, an external magnetic field between 100-300 μT is applied orthogonal to the propagation direction of the pump and probe beams. In this low field regime, the Zeeman splitting of the intermediate $5s5p\ ^1\mathrm{P}_1$ state is negligible compared to its natural linewidth. The probe beam is linearly polarised orthogonal to the quantisation axis to enable a balanced access to all the m_J levels within the Rydberg manifold given the fixed polarisation of the pump light. Any background residual magnetic field is nulled using electron-shelving spectroscopy on the narrow-linewidth $5s^2\ ^1\mathrm{S}_0 - 5s5p\ ^3\mathrm{P}_1$ transition at 689 nm[3]. We resolve a Doppler-broadened linewidth of 40 kHz which constrains any residual field to below 2 μT.

To test the sensitivity of our method, an external plate electrode located directly opposite a radial DN40 viewport is used to apply a DC electric field in order to induce a Stark shift of the Rydberg states. Shielding from the metal vacuum chamber greatly attenuates the applied field at the atoms' position, meaning several kV potentials are needed to induce a substantial Stark shift. Such large potentials have the unfortunate effect of charging the dielectric ($\epsilon = 3.75$) fused-silica viewport resulting in an exponential decay of the applied electric field strength at the atoms' position as inferred from a reduction in the Stark shift with time. To ameliorate this effect, we interleave measurements with opposite field polarity to avoid charging any external surfaces.

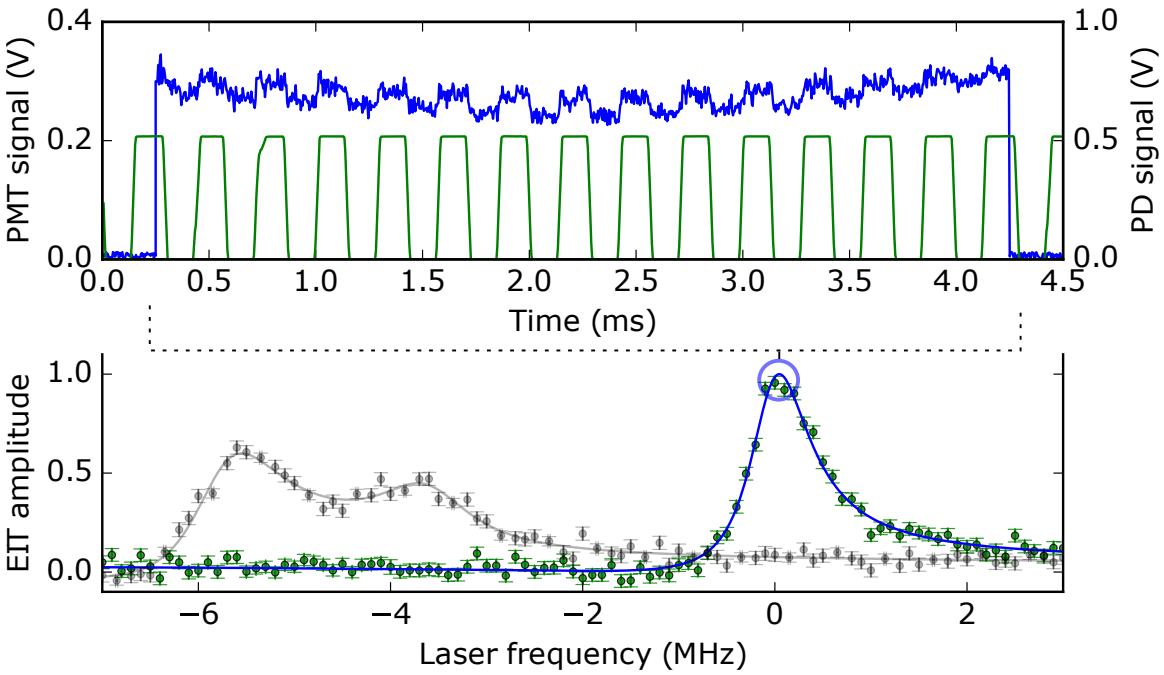


Figure 3.14: **Top:** Modulation of the pump beam intensity (green) and its corresponding modulated EIT signal induced on the probe absorption signal (blue). The pump beam is resonant with the $n = 53$ Rydberg state and both the external magnetic and electric fields set to zero. **Bottom:** An EIT spectrum is obtained by scanning the pump beam frequency and repeating such an absorption measurement. The spectra shown are with (grey) and without (green) an applied electric field.

3.3.5.5 EIT spectroscopy of Rydberg states

To better quantify the strength of the electric field, we must accurately measure the resulting Stark shift. This shift is also accompanied by a broadening resulting from the differential polarisability of the five m_F sub-levels. To resolve the line centres we employ a static magnetic field to separate better the levels via the Zeeman effect. Once clearly resolved, we fit the EIT signal T as a function of laser detuning Δ, using the following model:

$$T(\Delta) = \sum_{i=0}^{5} A_i \left(1 - \frac{(q_i \Gamma_i + \Delta - \Delta_{0,i})^2}{\Gamma_i^2 + (\Delta - \Delta_{0,i})^2}\right) \tag{3.30}$$

The model is composed of five Fano profiles each of which has four fitting parameters:

amplitude A, the linewidth Γ, the line centre Δ_0 and the Fano parameter q. Fig.3.15 shows a fit along with the normally distributed residuals.

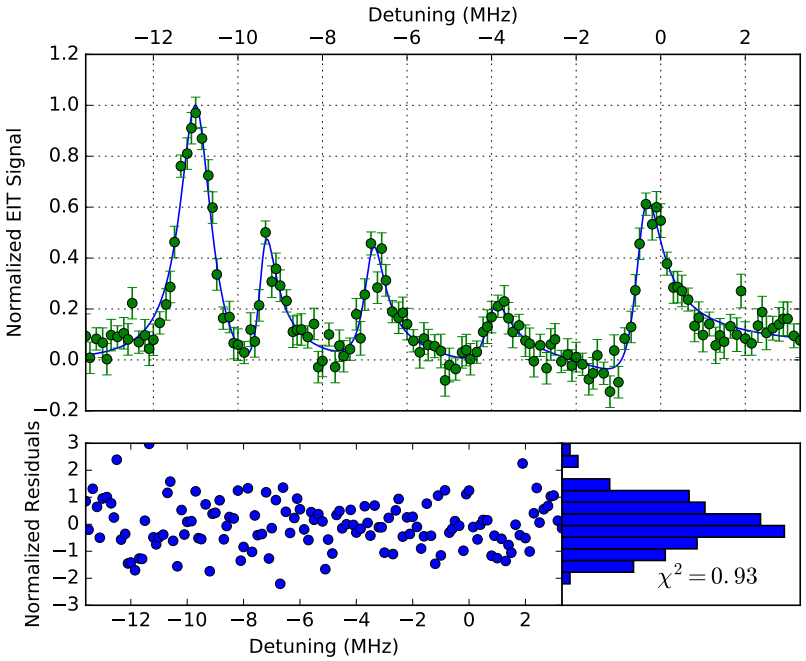


Figure 3.15: Example Rydberg spectrum taken at a magnetic field of 0.192 mT. The model shows good agreement with the observed data as indicated by the normally distributed residuals and the reduced χ^2 value.

As we do not have a method to directly measure the pump and probe frequencies, we have instead developed a method for measuring the applied electric fields that is based on the relative splitting of spectral lines, rather than the absolute detuning. In the absence of an electric field, the Zeeman sublevels split symmetrically with applied magnetic field $\boldsymbol{B}$. However, in the presence of an electric field $\boldsymbol{E}$, the Stark shift will result in an asymmetry in the spectrum, since the Stark shift depends on $|m_J|$. Example spectra with and without an applied electric field are shown in figure 3.16. In order to extract the line centres, five Fano profiles are fit to each spectrum corresponding to each of the five possible m_J transitions. The asymmetric effect of an applied electric field on the observed Zeeman shift of the m_J is clearly visible in figure 3.16.

To obtain the electric field from the spectroscopic data, the relative line positions are compared to a calculation of the combined Zeeman and Stark shift of each level. In the general case, the magnetic and electric field vectors are separated by an angle β, requiring transformation to a common basis. Choosing to work in the $|J, m\rangle$ basis defined by the magnetic field quantisation axis, the matrix elements of the Zeeman Hamiltonian are give by:

$$\langle J, m_1 | H_B | J, m_2 \rangle = -m_1 \mu B \delta_{m_1 m_2} \tag{3.31}$$

where μ, the magnitude of the magnetic dipole moment, is the Bohr magneton for a singlet state and δ is the Kronecker delta function. The Stark Hamiltonian H_E, with eigenenergies $\Delta_E(m_J, E)$ as shown in figure 3.12, is first commuted basis $|J, m'\rangle$ where m' is the angular momentum projection along the electric field direction. Next, the Hamiltonian is rotated by an angle β by applying the operator $R = e^{-i\beta J_y}$. The matrix elements of the Stark Hamiltonian in this rotated $|J, m_J\rangle$ basis are given by:

$$\begin{aligned}
&\langle J, m_1 | H_E | J, m_2 \rangle = \\
&\qquad \langle J, m_1 | R \sum_{m'} \Delta_E(m', E) | J, m' \rangle \langle J, m' | R^\dagger | J, m_2 \rangle \\
&\qquad = \sum_{m'} \langle J, m_1 | R^\dagger | J, m' \rangle \langle J, m' | R | J, m_2 \rangle \Delta_E(m', E) \\
&\qquad \triangleq \sum_{m'} d^J_{m_1, m'}(\beta) d^J_{m_2, m'}(\beta) \Delta_E(m', E)
\end{aligned} \tag{3.32}$$

where $d^J_{m_1,m'}(\beta)$ is the Wigner d-matrix for $J = 2$. In the final line, we used the identity $d^J_{m,m'}(\beta) = d^J_{m',m}(-\beta)$. Finally, the theoretical splitting is computed by diagonalising the Hamiltonian $H = H_\mathrm{E} + H_\mathrm{B}$. Using this approach, we fit the experimentally observed energy splitting by varying the electric field strength E and its angle β relative to the applied magnetic field in the model. As the Stark shift is

quadratic, our method only determines β modulo π. The only other fitting parameter is an overall two-photon detuning from the zero field resonance as we have no measure of the absolute frequency of the 413 nm laser.

Figure 3.16 shows the relative energy splitting for various magnetic fields with and without an applied electric field. An external electrode set to 2 kV, the maximum allowed by the high voltage supply, generated the applied electric field. From a fit to the Zeeman splitting, the electric field at the position of the atoms is estimated to be $5.75 \pm 0.11(\text{stat}) \pm 0.16(\text{sys})\text{Vm}^{-1}$. The fitting procedure also returned a value of $\beta = 0.47(1)\pi$ that is consistent with the axial magnetic field and radially applied electric field. An electric field of such magnitude would result in a fractional frequency shift of the clock transition equal to 2×10^{-19}. Given this is the largest field we can apply, it would have taken approximately a year of continuous operation to resolve this frequency shift given our fractional frequency instability, highlighting the utility of this method when applying large external fields is not possible.

Next the external field was switched off and the procedure was repeated. A fit to the resulting splitting revealed a residual electric field of $1.52^{+0.62(\text{stat})+0.05(\text{sys})}_{-.22\quad -.03}\text{Vm}^{-1}$ most likely due to patch potential on the surrounding chamber. The uncertainty for this electric field value is comprised of both statistical error resulting from the fitting procedure and systematic error arising the uncertainty in the Stark map for $5s75d\ ^1\text{D}_2$. The quoted statistical error corresponds to a 68% confidence interval as determined by the fitting procedure. A weak correlation observed between the uncertainty in β and the electric field is taken into account in this estimate [115]. The systematic error on the electric field value due to the uncertainty for the Rydberg polarisability was calculated by repeating the fitting procedure with revised Stark maps offset from the theoretically predicted value by $\pm\sigma$. Translating this electric field and corresponding uncertainty to the DC Stark shift of the $^1\text{S}_0$-$^3\text{P}_0$ clock transition results in a fractional frequency shift of $-1.6^{+0.4}_{-1.6} \times 10^{-20}$. The fractional uncertainty of the differential

polarisability of the clock states is negligible compared to that of the electric field and therefore has been ignored in the quoted uncertainty.

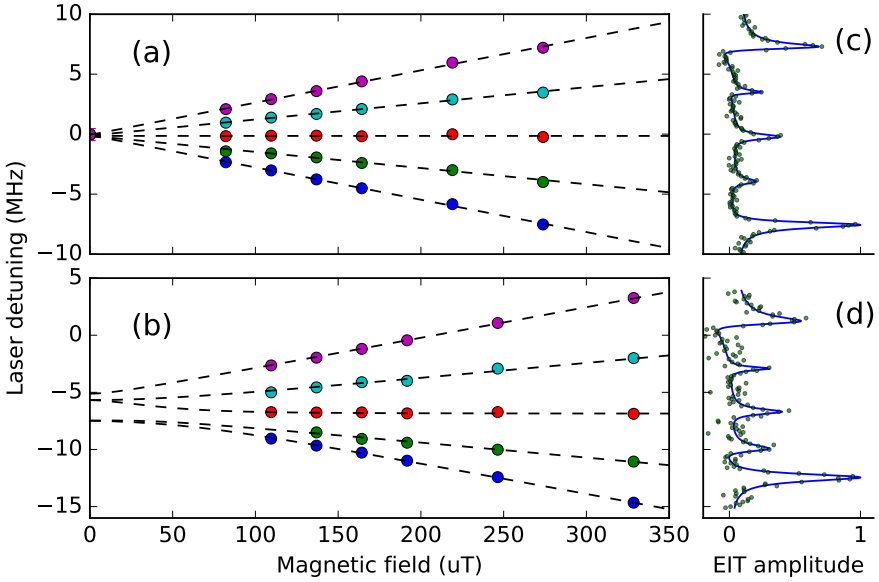


Figure 3.16: Inset *a* shows the Zeeman splitting without an applied electric field for the Rydberg state n = 75 (the error bars are smaller than the symbols). The line centres are extracted from the EIT spectra, an example of which is shown in inset *c* for the highest field case. Even without an applied field, a fit to the splitting reveals a slight asymmetry resulting from the tensor nature of the Stark shift consistent with a residual electric field of 1.52 Vm^{-1}. For the applied field case, shown in insets (*b*) and (*d*), the Stark shift is clearly visible and the fitting procedure returns an electric field of 5.75 Vm^{-1}

The error for the electric field and angle β are dominated by the uncertainty introduced by the fitting procedure. To characterise this uncertainty, we examine the sensitivity of the χ^2-value to changes in the fitting parameters. To do this we fix their values near those that minimise χ^2 and refit the data by only varying the two-photon detuning. By repeating this procedure for different electric field and β values, we produce a contour plot showing the change in χ^2 as function of these two fitting parameters. The region bounded by the contour corresponding to an increase in χ^2 of one sets the confidence region for the two fitting parameters. The extrema of this boundary defines the quoted $\pm 1\sigma$ uncertainties. Examples of such contour plots for the fits presented above are shown in figure 3.17.

This procedure also reveals correlations between the two fitting parameters that cannot be detected by the numerically computed covariance matrix. These correlations give rise to the asymmetric statistical errors for the measured DC-Stark shift. One can also see that χ^2 is symmetric about $\beta = \frac{\pi}{2}$. This results from the quadratic nature of the Stark shift. Therefore, our method returns the value of β modulo π. The numerical fitting package used for this work was the LMFIT: Non-Linear Least-Square Minimisation and Curve-Fitting developed for Python [100].

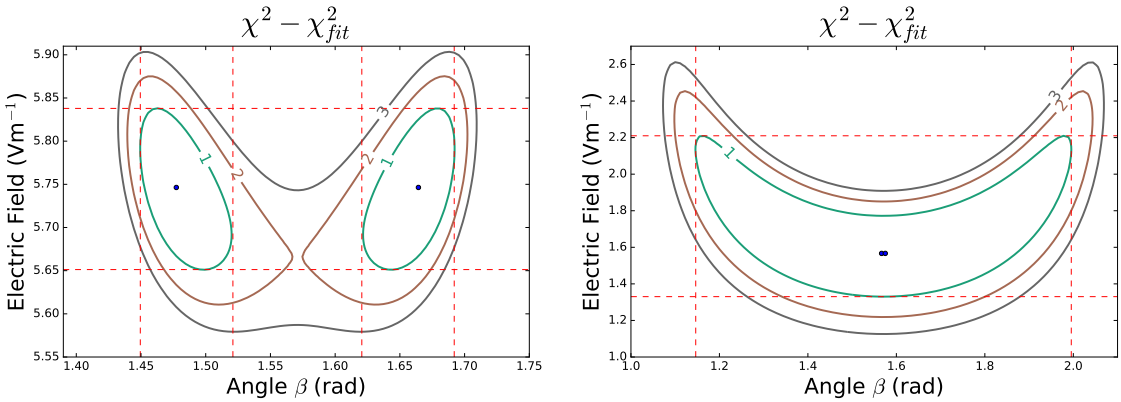


Figure 3.17: Contour plots showing the change in χ^2 resulting from varying the electric field strength and angle β in the fitting model. Contours are shown for the datasets with (left) and without (right) an externally applied electric field. From the contours we can determine the uncertainty for the fitting parameters along with correlation between their values. The blue dots indicate the values for E and β that minimise χ^2. The red dashed lines indicate the $\pm 1\sigma$ bounds for each fitting parameter based on the extrema of the contour corresponding to increase in χ^2 of one.

3.3.6 Zeeman shifts

For typical operation of the clock, a bias field of approximately 73 µT is applied using a Helmholtz coil pair to separate adjacent m_F sublevels. For the clock transition, the first-order linear Zeeman shift $\Delta\nu_B^{(1)}$ is given by:

$$\Delta\nu_B^{(1)} = \frac{\delta g \mu_B m_F B}{h} = \left(\frac{B}{1\,\mu\mathrm{T}}\right) \times 1.087(44) m_F \mathrm{Hz} \tag{3.33}$$

where δg, the differential g-factor between ground and upper state, has been experimentally measured [16]. For the applied field, the resulting separation between adjacent Zeeman sublevels is 70 Hz—easily resolvable given our spectroscopic resolution of <10 Hz. Two interleaved servos lock the clock laser frequency to the stretched $m_F = \pm 9/2$ states. The frequency standard is realised by averaging the two servos such that it is immune to a linear Zeeman shift. If we instead stabilised to a single stretched state, to reduce the fractional frequency inaccuracy of the linear Zeeman shift to below 10^{-17} would require an uncertainty in the applied field of less than 1 nT. The drawback of operating with an applied magnetic field is the resulting quadratic Zeeman shift given by:

$$\Delta\nu_B^{(2)}/\nu_0 = \left(\frac{B}{100\,\mu\mathrm{T}}\right)^2 \times -5.47(5) \times 10^{-16} \tag{3.34}$$

where the sensitivity coefficient is taken from [161]. Fortunately, the frequency splitting between the two stretched states provides a direct measure of the magnetic field strength and can be used to characterise the second order Zeeman shift[12]. Rewriting $\Delta\nu_B^{(2)}/\nu_0$ in terms of the stretched state splitting ν_{dif} gives:

$$\Delta\nu_B^{(2)}/\nu_0 = \left(\frac{\nu_{\mathrm{dif}}}{1\,\mathrm{kHz}}\right)^2 \times -5.75(4) \times 10^{-16} \tag{3.35}$$

where the sensitivity coefficient is the weighted average of several measurements as given in [101] and was experimentally verified using our clock as shown in figure 3.18.

[12]The assumption that stretched state splitting is directly proportional to the magnetic field is only true if the vector component of the lattice Stark shift is zero. The differential vector light shift is equal to $\delta\alpha_s^v m_F \xi \cos(\theta)$ where θ is the angle between magnetic field vector and k-vector of the lattice beam, $\delta\alpha_s^v$ is the vector shift coefficient measured to be 220(50) mHz/E_r and ξ is the ellipticity of the lattice polarisation. For our lattice, conservative upper bounds for the ξ and the deviation in θ from $\pi/2$ are 0.1 and 50 mrad, respectively. For our operational trap depth of 45 E_r, a conservative bound on the change in stretched state splitting due to the vector light shift is 0.5 Hz. This would introduce an error in the estimate of the quadratic Zeeman shift of $< 1 \times 10^{-18}$.

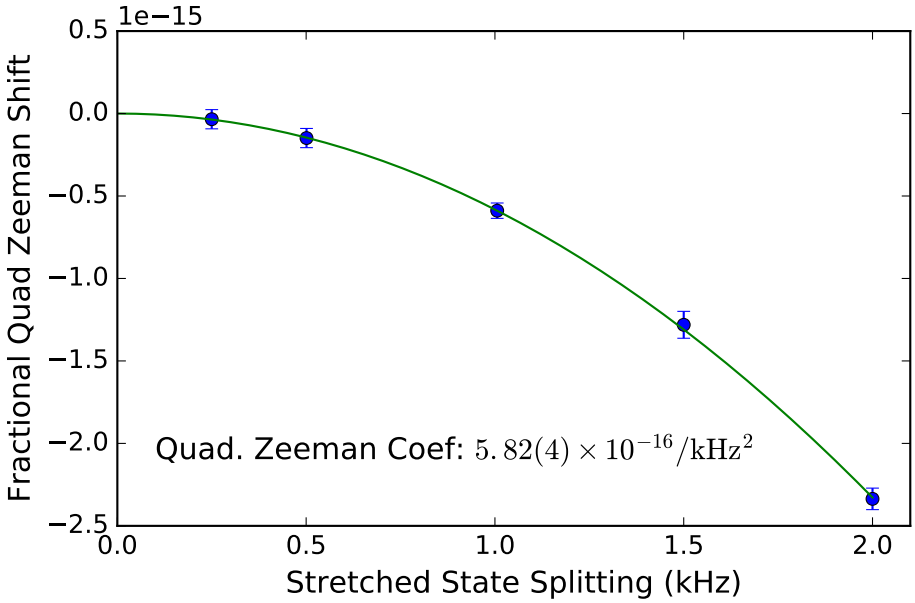


Figure 3.18: The quadratic Zeeman shift of the ^{87}Sr clock transition as a function of stretched state splitting. Data was taken by interleaved comparison to reference servo operating at a stretched state splitting of 1.006 kHz. The differential frequency shift data are then offset such that the fitted curve intersects the point (0,0) on the plot. The uncertainty in the quoted sensitivity coefficient is purely statistical and the influence of systematic effects remain to be evaluated.

The stretched state splitting also provides a measure of the long term instability of the magnetic field and shows it to be stable at the 20 nT level. For our operational stretched splitting of 706 Hz, the fractional quadratic Zeeman shift is $-287(2)\times10^{-18}$. The uncertainty is solely from the sensitivity coefficient as the contribution of the 20 nT magnetic field instability amounts to a change in the quadratic Zeeman shift at the 2×10^{-19} level.

3.3.7 Probe shift

Although the required intensity to drive the clock transition is small, the resulting probe induced Stark shift needs consideration. Given the theoretical lifetime (Γ^{-1}) of the upper clock state of 100 s, the necessary probe intensity I_π to drive a π-pulse

in a time T_p is:

$$I_\pi = \frac{4\pi^4\hbar c}{3\Gamma t_p^2 \lambda^3} = \frac{1.6\times10^{-5}\mathrm{W}}{\mathrm{cm}^2}\left(\frac{100\mathrm{ms}}{T_p}\right)^2 \tag{3.36}$$

where c is the speed of light and λ is the clock wavelength. The required intensity is marginally increased when taking into account the Clebsch-Gordan coefficient for the $m_F = 9/2 \rightarrow 9/2$ clock transition and the Lamb-Dicke parameter η_z associated with the trapping potential. These two effects reduce the Rabi frequency by a factor of 0.9045 and $1 - \eta_z^2 \approx 0.89$, respectively. Returning to figure 2.4, the differential polarisability at the clock transition is estimated to be $\Delta\alpha(698.4\,\mathrm{nm}) = 607.0$ a.u. Combining this with equation 3.36 and accounting for the effect of the Clebsch-Gordan coefficient and Lamd-Dicke parameter, the estimated probe shift as a function of pulse time is:

$$\Delta\nu_{\mathrm{probe}} = -1.65\times10^{-18}\times\left(\frac{100\,\mathrm{ms}}{T_p}\right)^2 \tag{3.37}$$

Therefore, the estimated probe shift for our operational pulse length of 200 ms is -4.1×10^{-19}. Resolving such a shift would require a prohibitively long period of operation given the stability of our clock. However other groups with more stable systems have resolved a probe related shift, for example [103] measured the Stark shift for a 1 s π-pulse to be $(-3.2\pm1.7)\times10^{-20}$ – in agreement with the derived value using equation 3.37. Given our inability directly measure it, the 4.1×10^{-19} estimated probe shift is quoted only as an uncertainty in table 3.1.

3.3.8 Doppler shift

The relative motion of the atoms with respect to the probe light compromises their ability to resolve the phase of the local oscillator. This motion can reduce the clock's stability and lead to Doppler shifts. The required velocity needed to induce a substan-

tial frequency shift is surprisingly small—just 3 nm/s induces a fractional frequency shift at the 1×10^{-17} level. As the atoms are confined in the anti-nodes of the optical lattice, keeping the atoms' position fixed amounts to ensuring the retroreflection mirror of the lattice is stationary. This is achieved passively by keeping sources of vibration to a minimum: all optical shutters use piezo cantilevers resulting in minimal mechanical noise when switching [13], the larger mechanical shutter for the camera is mechanically isolated from the optical table, and the MOT coil are rigidly mounted to the chamber.

Doppler shifts also arise if the probe beam is experiencing a phase chirp. This could result from a mirror moving, an optical fibre used to deliver the light expanding or contracting with temperature, or the optical path through the AOM changing resulting from thermal expansions as it heats with applied RF power. To compensate for this we actively stabilise the optical path of the probe using a path length stabilisation (PLS) servo that actuates the frequency of an AOM [83]. The optical setup is shown in figure 3.19. The stabilised optical path includes the AOM used to switch the probe light on and off during spectroscopy such that any phase chirp from heating is actively canceled following the same procedure described in [86]. The only uncompensated path over which the probe beam propagates is between the AOM and dichroic filter used to combined the clock light with the lattice.

To characterise the performance of the path length stabilisation, the probe beam was split and delivered along two separate fibres to the chamber. Independent PLS servos were set up on the two delivery paths. At the chamber, the two beams were overlapped and the resulting beatnote was counted to verify the frequency stability of the PLS servos. The resulting offset was consistent with zero at the $< 1 \times 10^{-18}$ level. However, this measurement was taken in steady state conditions without the probe AOM periodically switching on. Transient effects combined with the finite response time of the servo could lead to an uncompensated phase chirp. To evaluate possible

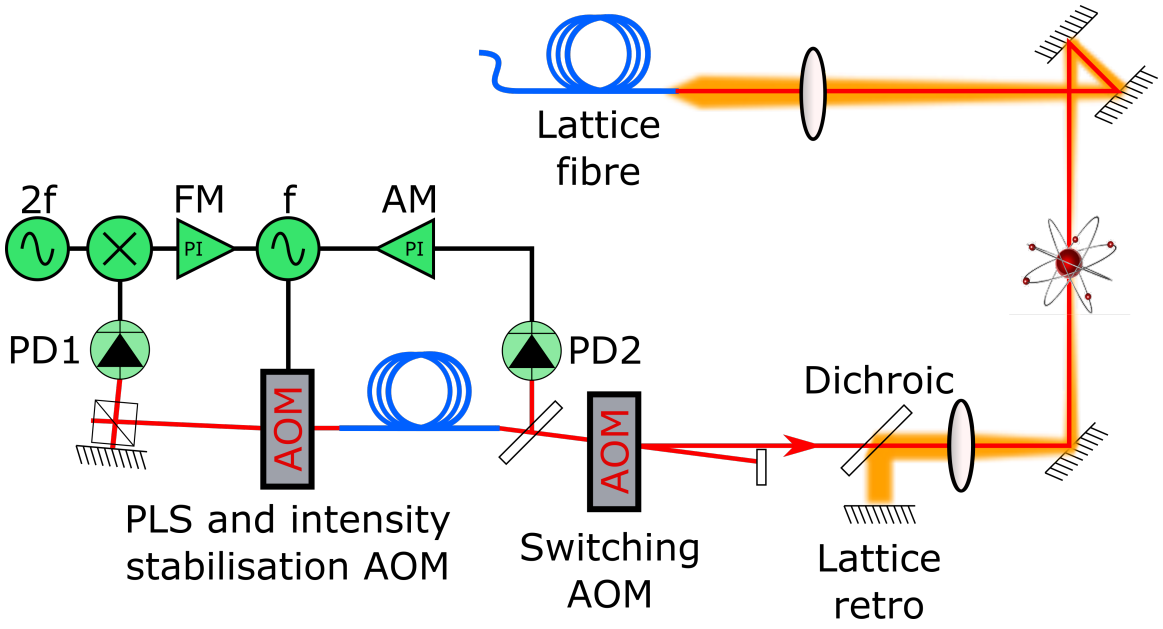


Figure 3.19: Diagram depicting the PLS setup used to deliver phase stable light to the science chamber. The light is shifted in frequency by an amount f using the first AOM before going through a 5 m fibre connecting the clock laser setup to the science chamber. Part of the light is sampled for intensity stabilisation before going into the switching AOM. The first-order diffraction is used to probe the atoms while the zero-order pass is retro-reflected back through the fibre. The reflected light is shifted again in frequency by the PLS AOM and then measured against the incoming light using a homodyne detection scheme. The resulting beat signal is phased locked to a stable RF reference at $2f$ via feedback to the AOM's drive frequency.

switching induced frequency offsets, we repeated the comparison while having the probe AOM switch as it would under normal clock operation. The only difference was the AOM RF power was set to 1 W to maximise the likelihood of observing thermally induced phase shifts. For normal clock operation with 60 db less power, such shifts will be considerably smaller. During the probe pulse, the beatnote between the two PLS servos is recorded in the time domain using a high speed oscilloscope (model: Tektronix MSO/DPO4000B) and transient drifts of the relative phase are extracted by performing piecewise fits of the resulting sinusoidal curve. We measured the relative phase for several pulses with and without the PLS servo engaged. With the PLS servo disengaged, the AOM induced phase chirp was measured to be 800 mrad/s based on

the average of 50 pulses, see figure 3.20. In fractional units, this phase chirp would correspond to an offset of 3×10^{-16}.

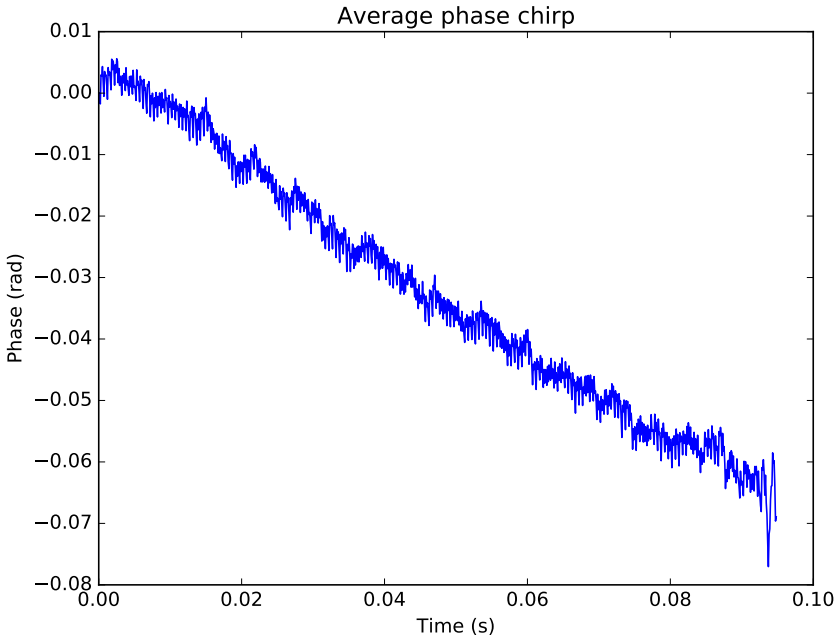


Figure 3.20: Estimated phase chirp during a 80 ms probe pulse with the PLS servo off. Based on 50 measurements, the estimated phase chirp is 800 mrad per second.

With the PLS engaged, we measured a suppression of the phase chirp by at least a factor of 5. The suppression could in fact be much better, but the signal-to-noise and technical limitations make it difficult to revolve phase chirps below 100 mrad per second using this method. Even with this pessimistic estimate for the servo suppression, the estimated phase chirp for our 200 ms pulse which uses 53 db less RF power is 3×10^{-22}. To verify the linear scaling of phase chirp with RF power, the procedure was repeated for several AOM drive powers, see figure 3.21.

3.3.9 Servo offset

Given the finite gain of servo used to lock the laser to the atomic resonance and the various noise processes perturbing the laser frequency, it is possible for there to be

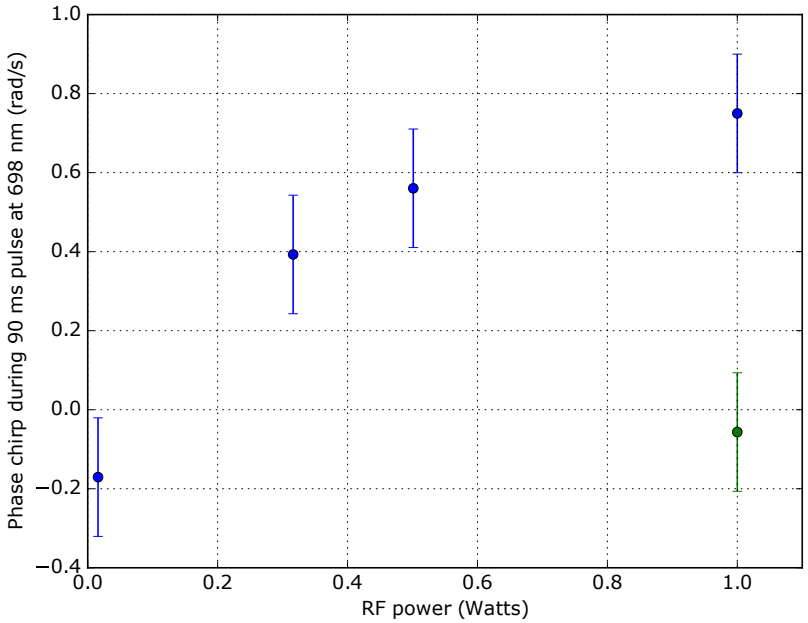


Figure 3.21: The measured phase chirp of the AOM as a function RF power with the PLS servo off (blue) and on (green). Each data point and corresponding error bar is a result of averaging the estimated phase chirp for 50 pulses.

an offset between the local oscillator and the atomic resonance. To suppress possible offsets, it is advantageous to maximise the sensitivity of the error signal ϵ to changes of the laser frequency ν.

To evaluate if both the atom servo and drift servo are functioning as intended, we save the error signals from every cycle and the residual offset in post processing. By averaging the error signal and scaling by the frequency discriminator, we derive the estimated frequency offset resulting servo error is given by:

$$\Delta\nu_{\text{servo}} = \frac{1}{1.8T_cN}\sum_{i=0}^{N-1}\epsilon_i \tag{3.38}$$

The servo error results from the finite gain of the control loop, uncompensated linear drift, or non-linear frequency excursions – all of which can be detected using this method. The Allan deviation of the servo offset is shown in figure 3.22. For measurement times greater than 3 h, the servo error is less than 1×10^{-18} and the mean

is consistent with zero within the statistical uncertainty. However, the assumption that frequency offset is proportional to the error signal is only true in the linear region of the discriminator. As the local oscillator detuning is increased, the slope of the error signal is reduced and scaling by the discriminator underestimates the frequency excursion. To ensure the laser frequency remains within the linear region of the error signal, we calculate the frequency detuning of every probe in post processing using the recorded excitation fraction and theoretical Rabi lineshape. This process in shown in figure 3.22 for a 200 ms probe pulse and clearly shows the local oscillator does not deviate outside the linear region of the discriminator.

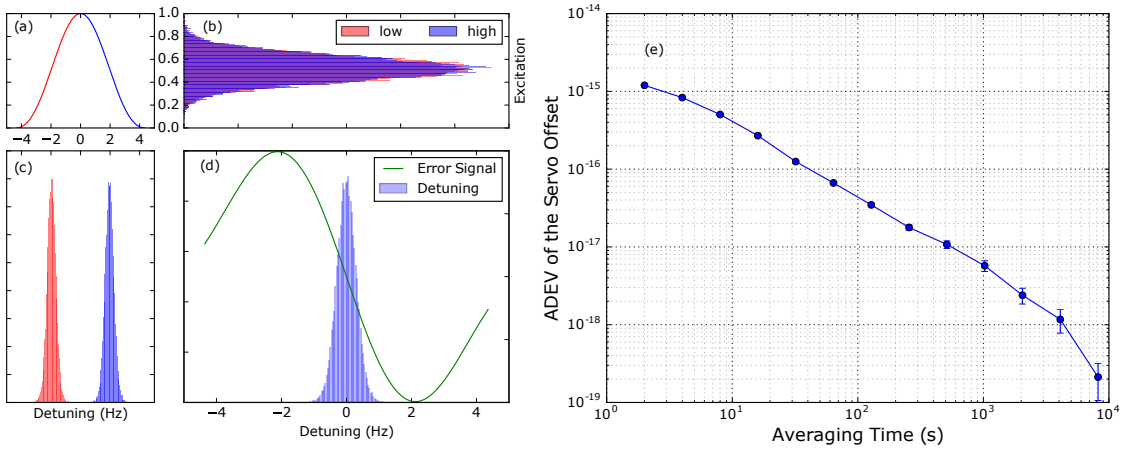


Figure 3.22: Sub-figure (a) shows the central Rabi lineshape for a 200 ms π-pulse used to convert the excitation fraction to laser detuning in post processing. The two measurements of the excitation fraction taken on the 'high' and 'low' side of the line by stepping the laser frequency by 2 Hz are shown in sub-figure (b). These distributions are converted to the laser detuning shown in (c) using the Rabi lineshape. The only assumption used in the reconstruction process is that the laser detuning does not exceed half the linewidth Δ. This is justified given the well separated distributions for the 'high' and 'low' detunings in (c). The frequency excursions of the clock laser are computed by subtracting the 2 Hz detuning steps from the 'high' and 'low' distributions shown in (c). The resulting distribution is shown in (d) and compared to the theoretical error signal to clearly show the clock laser frequency remains within the linear portion of the discriminator. Sub-figure (e) shows Allan deviation for the mean error signal in fractional frequency units scaled by the inverse of the frequency discriminant, see equation 3.38. The instability decreases faster than $\tau^{-1/2}$ as the shot-to-shot error signals are correlated.

3.4 Conclusion

This completes the characterisation of Sr1's stability and accuracy. In regards to our clock's accuracy, it is limited by BBR radiation resulting from the uncertainty of the chamber's temperature. Tackling this problem comes down to engineering a well-controlled blackbody radiation environment, but several solutions outlined in section 3.3.1 have already been demonstrated by other groups which we can adopt in the future. Therefore, we turn our attention to the next largest uncertainties which are associated with the dipole and higher-order Stark shifts induced by the lattice. Estimating the magnitude of these high-order lattice shifts has proved challenging to the strontium lattice clock community. There is discrepancy between groups over whether these shifts are [65, 38] or are not [101] relevant at our current levels of inaccuracy. Furthermore, recent experimental work measuring these higher-order effects [149] was inconsistent with theory [113]. Under such circumstances, the proven course of action is always take more data – motivating us to investigate these effects further. To expedite this process, it would be valuable to have a second clock to which we can compare our first system operating under different lattice conditions. Comparing both systems synchronously would eliminate the Dick effect and help to improve stability. Furthermore, if this new system had an integrated build up cavity for the lattice we could generate deeper traps where such non-linear effects become more pronounced.

The enhancement cavity can also benefit the clock's stability. As we saw, the instability arises due to the Dick effect caused by the deadtime in the experimental cycle needed for cooling and trapping. An enhancement cavity allows for a larger trap volume which helps loading efficiency, requiring less cooling time. The larger mode volume also helps to improve accuracy by reducing the density shift, which could be further suppressed by using higher dimensional lattices.

Another way the cavity can help to improve stability cavity is by using it to per-

form non-destructive detection of the number of atoms trapped in the lattice. The general principle would be to have a multi-wavelength enhancement cavity which is high finesse at both the magic-wavelength and the 461 nm-atomic resonance. The coupling of the atoms to the cavity will induce a frequency shift of the cavity's resonance proportional to the atom number. The following chapter will cover the design and preliminary characterisation of such a system.

Chapter 4

Sr2: The Second Generation Strontium Lattice Clock

Building on the lessons learned from Sr1, we undertook the task of realising a second strontium lattice clock. The main goals for this system were to improve upon the accuracy and stability of Sr1, while testing the feasibility of various design features that could simplify the experimental system and increase its robustness. This includes the use of a pyramid MOT and in-vacuum enhancement cavities for the lattice, along with investigating new techniques to improve the duty cycle of the clock in order to suppress the Dick effect. Such techniques include the possibility of replacing the 'red' MOT stage with with a triplet MOT operating on the $5s5p\ ^3\mathrm{P}_2 \rightarrow 5s4d\ ^3\mathrm{D}_3$ transition, and developing cavity-aided non-demolition detection (NDD) of the atom number trapped in the lattice. The following chapter will outline the design of Sr2 and present the initial results from this system.

[1]The details regarding the design of the vacuum system are given in appendix B.

4.1 Experimental apparatus

This section outlines the design of Sr2 focusing on the pyramid MOT, a compact permanent magnet Zeeman slower, and in-vacuum cavities– which are used for the magic-wavelength lattice and non-destructive detection.[1] I will highlight the key technical issues associated with building such a system and solutions we implemented to address these challenges. This section will also cover the various laser systems. In most cases we extended the existing infrastructure from Sr1, but for second stage cooling we built a new stable laser operating 2.9 µm. The main apparatus is shown in figure 4.1 along with the rendering showing how the pyramid MOT and in-vacuum cavities, along with the necessary lasers, are integrated into the setup.

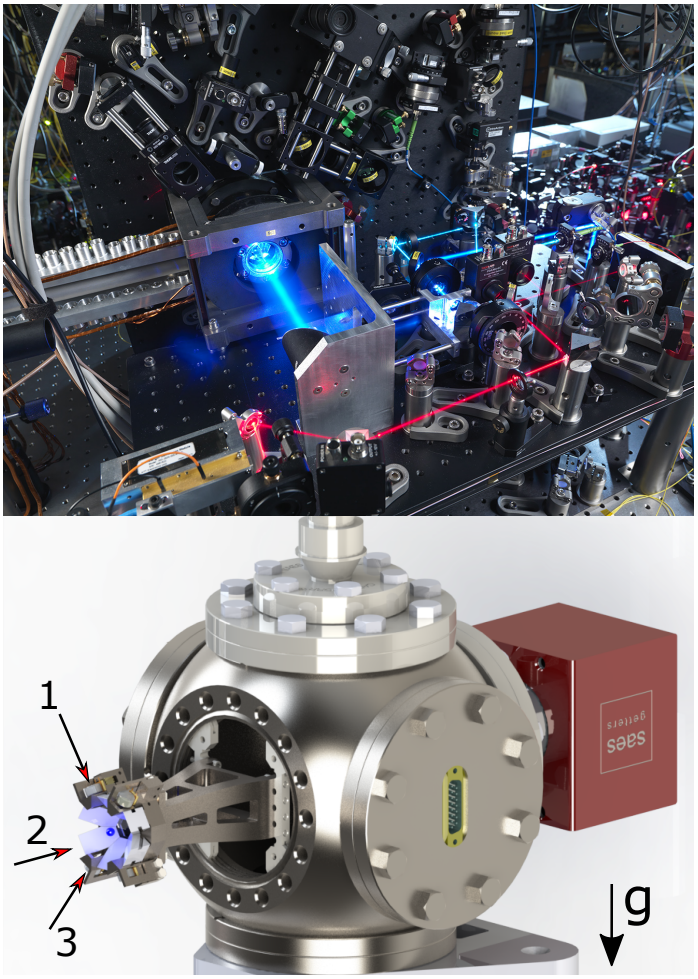


Figure 4.1: *Top:* Photo of the Sr2 vacuum system, coils and surrounding optics. The MOT and repump light delivery optics are located on the horizontal breadboard in front of the chamber (see figure 4.15 for details) while the optics for the in-vacuum cavities are on the vertical breadboard (see figure 4.16). *Bottom:* At the centre of the experiment is the pyramid MOT and the two in-vacuum enhancement cavities which are secured to the same titanium baseplate. Beams are sent into the system along three directions labelled 1, 2, and 3. Direction 1 is used to couple the magic-wavelength lattice beam into the enhancement cavity. This cavity axis is also has an anti-reflection coating for 698 nm as we probe the clock transition along this same direction. This cavity is also high-finesse at 461 nm and we use it for cavity aided non-destructive detection by coupling a weak near resonant probe beam into the cavity. Along direction 2, we send in the first and second stage cooling light for the MOT and the repump light. Finally along direction 3, we have the option to send in another magic-wavelength lattice beam into the second the enhancement cavity to form a 2D lattice.

4.1.1 Pyramid MOT

The pyramid MOT replaces the three retro-reflected beams needed for a typical MOT with a single beam incident on the pyramid retro-reflector [6, 22]. Besides requiring fewer mirrors and being simpler to align, there is no need for polarisation optics as the inversion of circular polarisation occurs naturally upon retroreflection. One challenge with the design is ensuring the beam is well reflected from the apex of the pyramid as this is directly inline with where the atoms are trapped for loading into the lattice. To have the mirrors meet at the apex without any gaps is challenging. We could have elected to have them precision cut as was demonstrated in [162], but we instead chose to use a prism fabricated from calcium fluoride (CaF_2). We also elected to use six radial trapping beams. The novel hexagonal symmetry is chosen partly to allow three radial axes of optical access (needed for the Zeeman slower and a 2D lattice), and partly to produce an extra pair of radial confinement beams to compensate for the $\sim 40\%$ lower MOT beam intensity incident on the prism mirrors compared with the intensity at center of the gaussian beam. Other designs have attempted to address the imbalance between the radial and axial directions by using either a conical reflector [67] or axicon [62] which effectively generate an innumerable number of concentric beam pairs in the radial plane. However, we were hesitant to implement such a design as no measure of temperature has been made in such a system. The focusing behavior of both the conical reflector and axicon lead to a large intensity in the centre of the trap which could lead to excessive heating. Figure 4.2 outlines the working princple of the MOT and its basic design.

The six prisms that reflect the beams needed for radial confinement in the MOT measure 10 mm in width with a height of 12.7 mm. The prism mirrors are glued to a titanium baseplate using low out-gassing, ultra-high vacuum (UHV) compatible epoxy (EPO-TEK H77). The hexagonal hole, with an inradius of 4.3 mm, formed at the centre where the prisms meet is filled by the retro-reflecting CaF_2 prism. The apex

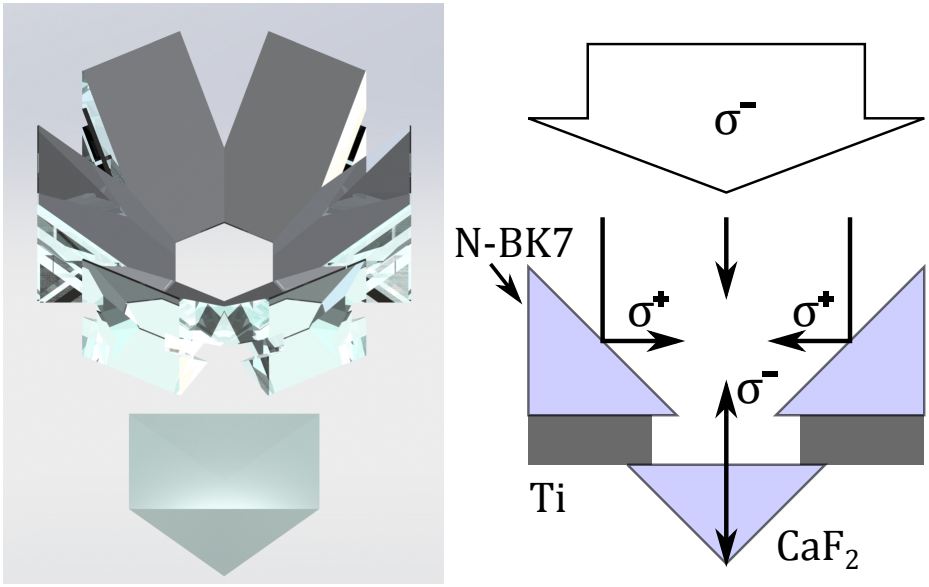


Figure 4.2: *Left:* Rendering of the in-vacuum optics used to realize a pyramid MOT. *Right:* Diagram showing a cross section of the pyramid MOT. A protected silver coating is applied to the hypotenuse of each N-BK7 prism; the CaF_2 prism is uncoated. The assembly measures a total of 33 mm across. The structure is illuminated by a circular polarised beam, upon 90° reflection, the polarisation handedness is mirrored ensuring the radial beams have the correct polarisation as needed for the MOT. The polarisation of the retroreflected beam is inverted twice such that it is restored to its original orientation.

of the prism is machined to a precise right-angle without beveling so that aberrations in the reflected beam are minimized. The CaF_2 prism is glued to stress-relieving flexure bars milled into the titanium baseplate, see figure 4.3. The flexure mounts mitigate thermal induced stress which is necessary to prevent cracking during UHV bakeout. CaF_2 was chosen primarily because of its transparency in the mid-infrared, making it compatible with realizing a second stage MOT operating on the $5s5p\ ^3P_2$ to $5s4d\ ^3D_3$ transition at 2.92 µm (see figure 2.1). However, all the pyramid MOT optics are also compatible with the more conventional route to microkelvin temperature with strontium, by cooling on the $5s^2\ ^1S_0$ to $5s5d\ ^3P_1$ transition at 689 nm [59, 97]. Another important feature of CaF_2 is the relatively low refractive index $n \approx 1.43$, which leads to near-critical total internal reflection at a 45 degree angle of incidence, thus minimizing polarization-dependent phase shifts upon retro-reflection of the MOT

beam [8]. Alternative optical materials with $n > 1.45$ would create a differential phase shift of more than 50 degrees between the s- and p-polarizations upon retro-reflection, severely undermining the MOT confinement force. In total, the entire retroreflector has an inradius of 33 mm and is illuminated with a beam whose waist ($1/e^2$ in intensity) measures 22 mm.

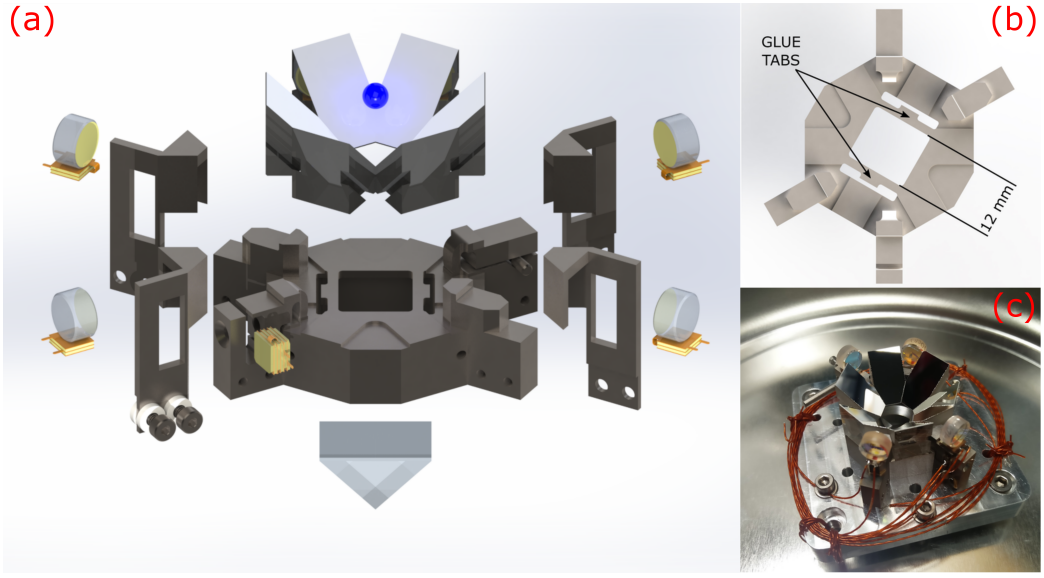


Figure 4.3: (a) Exploded rendering of the complete MOT and cavity system showing the various components: the hexagonal pyramid MOT and retro-reflector, the cavity mirrors and shear stack piezos onto which they are glued, the compensation electrodes, and finally the cavity flexure mount used to overlap the two 1D cavities (see text for more details). (b) The prism at the apex of the retro-reflector is glued to flexure bars that compensate for the thermal mismatch between the glass and titanium baseplate. Without compensation, the stress induced during thermal cycling the pyramid MOT structure cracked the prism. Furthermore, residual stress induces birefringence that could distort the polarisation of the trapping light. (c) Picture of the completed structure with the electrodes removed for clarity.

One important consideration is matching the thermal expansions between the titanium baseplate and the prisms to prevent the glass from cracking during baking. The radial prisms are are made from N-BK7 glass which has a thermal expansion coefficient of $7 \times 10^{-6}\,\mathrm{K}^{-1}$, well matched to that of titanium which is $8 \times 10^{-6}\,\mathrm{K}^{-1}$. CaF_2, on the other hand, has a thermal expansion coefficient of $19 \times 10^{-6}\,\mathrm{K}^{-1}$. Sure enough, when gluing a CaF_2 test piece to a titanium baseplate, which requires heating

the sample to 150 °C to cure the epoxy, the glass broke when it was cooled back to room temperature. To mitigate the thermal mismatch, the titanium MOT baseplate has two thin beams onto which the prism is glued that flex to relieve stress, see figure 4.3. The dimensions of the beams were chosen such that the deflection needed to account for the relative change in size when going from 150 °C to room temperature requires a maximum stress in the prism of ten percent of its elastic limit. The deflection of the beam Δx can be estimated using the following formula for a bar of length L, cross sectional moment of inertia I, and Young's Modulous E fixed on both ends and loaded in the centre with a force F:

$$\Delta x = \frac{FL^3}{48EI} \tag{4.1}$$

Using the above expression, the dimensions of the stress relief flexure bars where calculated to be 12 mm long and 1.6 mm wide – increasing to 2.6 mm at the point where the prism is glued.

4.1.2 In-vacuum cavities

For Sr2, we elected to integrate into the design in-vacuum build-up cavities for the lattice. There are several benefits of including build-up cavities, such as:

- The cavity ensures that the radial mode profile of the lattice beam is a near-ideal Gaussian which is important for ensuring the theoretical models of the atomic motion [12] and high-order lattice shifts [18, 60] match well with experiment.
- The running-wave component of the lattice beam is eliminated by the cavity, avoiding first-order contributions from electric quadrupole and magnetic dipole lattice shifts.
- The power enhancement factor, given by $\mathcal{F}/\pi$, enables the use of larger lattice

beam waists without compromising trap depth. Using larger trap volumes reduces the density shift for a fixed trap depth and atom number, and increases the number of atoms transferred from the MOT into the trap.

- The power enhancement reduces the required power incident on the cavity, helping to suppress the Stark shift from amplified spontaneous emission (ASE) incident on the atoms [11].

- Finally it could prove useful in the future for transportable lattice clocks as it eliminates the need for a high power lattice laser, reducing the cost and size of the system.

However, cavities also greatly increase experimental complexity and introduce several technical challenges which, if not properly addressed, could render the clock useless. First, the cavity mode must be locked to the lattice laser such that the power coupled into the cavity is stable. If not, fluctuations in the trap depth can quickly heat atoms. Therefore, the cavity must be designed to be mechanically stable and the laser must be free of significant frequency noise. These points will be covered in more detail in section 4.1.2.4. The other main concern is compromising the UHV environment by placing parts inside the chamber. Therefore, the material used for the cavity (e.g. the mirrors, piezos, wires and epoxies) must be vacuum compliant.

Aside from enhancing the lattice power, the high reflective coatings for cavities also include several other wavelengths; the reason for these additional wavelengths will be outlined in section 4.1.2.2. We elected to have different coatings for the two cavities inside the system, creatively named the 'red' and 'blue' cavities based on their reflectivity. The targeted design finesses and intended applications are outlined in section 4.1.

The titanium baseplate onto which the pyramid MOT mirrors are glued also serves as the spacer for the two cavities forming a compact, monolithic structure capable

Table 4.1: Finesse values of in-vacuum build up cavities

Wavelength (nm)	Blue Cavity	Red Cavity	Purpose
406	N/A	3000	Tune-out wavelength [a]
413	N/A	5000	Rydberg Generation
461	13,000	N/A	NDD ($^1S_0 \rightarrow {}^1P_1$)
497	1500	N/A	Magic-wavelength lattice
689	N/A	50,000	NDD ($^1S_0 \rightarrow {}^3P_1$)
698	AR coated	50,000	Super-radiance ($^1S_0 \rightarrow {}^3P_0$)
813	8000	8000	Magic-wavelength lattice

[a] The polarisability of the $^3P_0$ state goes to zero at 406.5 nm. Fortuitously, such a wavelength is commensurate with a magic wavelength lattice at 813 nm. It is therefore possible to modify the on-site lattice potential of only one of the states without shifting the potential minimum. A recent proposal [57] relies on a such a double lattice to realise a topological super-fluid with strontium.

of cooling and trapping the atoms. The cavity modes fit in the openings between the pyramid MOT mirrors, making them offset from the vertical by $\pm 30°$ when the structure is mounted in its final orientation. By not having the lattice horizontal, gravity helps suppress tunneling by lifting the degeneracy between adjacent sites [72]. The cavity mirrors have a diameter of 7.75 cm[2] and a radius of curvature of 10 cm. Given the cavity length of 3 cm, the radius of curvature of the mirrors result in a lattice waist of 100 µm at 813 nm. The cavity length is actuated using a 5 mm shear stack piezo onto which the mirrors are glued using EPO-TEK H77 epoxy. During gluing the mirrors are held in place by v-groove jig which ensures that they stay properly aligned, see figure 4.4. The stack was manufactured by Noliac and consists of two piezos which provide a total travel of 3 µm. The piezo stacks are glued to the baseplate using EPO-TEK H27D which is electrically conductive.

[2] Super polished substrates provided by REO and coated by FiveNine optics.

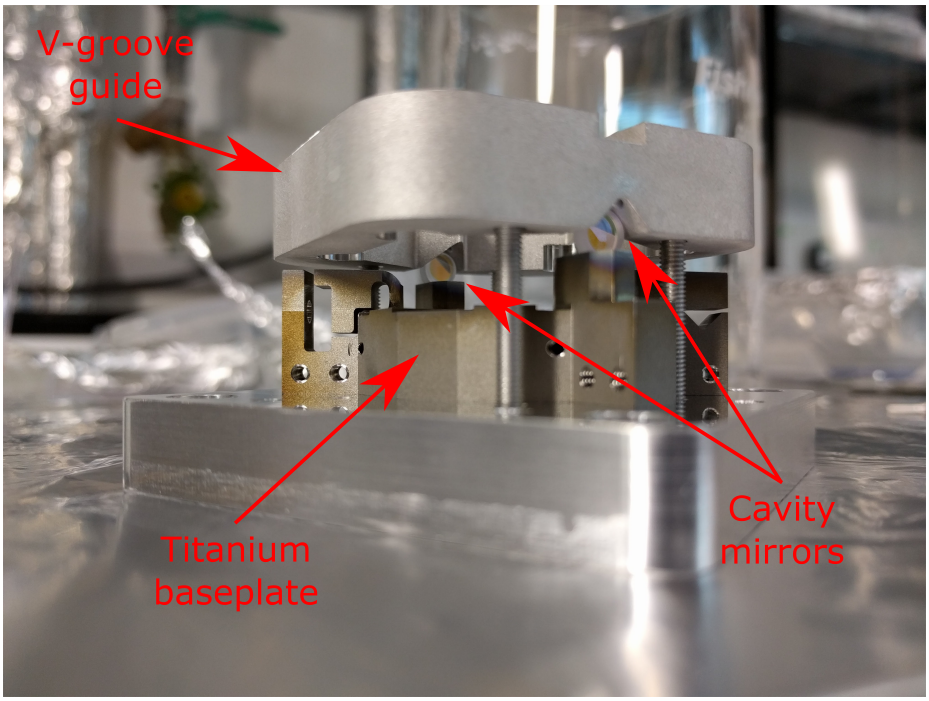


Figure 4.4: The mirrors are held in place during gluing using a v-groove jig. Alignment pins (not-shown) that fit into the jig are used to set the correct separation for the mirrors. The picture was taken during the testing of the cavity finesses. The final gluing of the cavity mirrors was done last, after the piezos and pyramid MOT prisms were secured.

4.1.2.1 DC-Stark shift considerations

One important consideration when placing the cavity and MOT mirrors inside the chamber is introducing large residual electric fields due to patch potentials on the substrates. This has been shown to lead to fractional shifts on the order of 1×10^{-13} [77]. A further annoyance is the large voltages applied to the cavity mirror piezos can also give rise to electric fields. Preliminary modeling using COMSOL indicated that an exposed $5\,\mathrm{mm}^2$ piezo placed under the cavity mirror with an applied potential $100\,\mathrm{V}$ would induce a Stark shift to the clock at the level of a $100\,\mathrm{Hz}$. To mitigate this problem, we use the dual stack piezo formed of two shear piezos orientated in opposite directions which sandwich an inner electrode. The top and bottom surfaces are connected to each other via a single beryllium-copper (BeCu) electrode. BeCu is

UHV compatible and is non-magnetic which prevents the switching of the MOT coils from perturbing the cavity. The outer two electrode plates are kept at ground while the control voltage is applied to the internal electrode such that is it well screened from the atoms. This prevents the electric fields generated by the piezos from Stark shifting the clock transition. According to the simulation, this simple design change of using the stack piezos attenuates the residual field by over 20 db.

To suppress further any residual fields from the cavity mirrors and piezos, metal electrodes are secured over mirrors as shown in figure 4.5. The electrodes are isolated from the metal baseplate by ceramic washers and can be set to arbitrary voltages allowing us to characterise residual electric fields. The electrodes also function as protective covers for the cavity mirrors preventing contamination from the atomic beam. Whether or not this is necessary in unclear as there is no direct line of sight between the mirrors and the oven, but was implemented as a precaution.

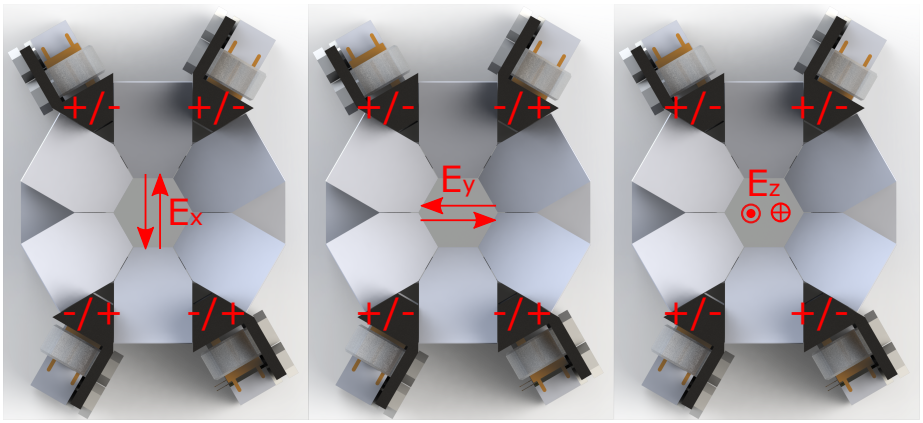


Figure 4.5: The in-vacuum electrodes are used to detect and compensate for the residual electric fields inside the chamber. They can apply fields in all three orthogonal directions by changing the configuration of the applied voltages. Electric fields in the plane of cavity (E_x and E_y) are set by applying opposite polarity voltages to neighboring electrode pairs, while an electric field in the perpendicular direction (E_z) is generated by applying an opposite polarity voltage to all electrodes relative to the chamber and cavity baseplate.

4.1.2.2 Cavity aided non-demolition atom detection

Along with the power enhancement of the lattice, the other advantage afforded by the in-vacuum cavity is the ability to perform NDD of the atom number. Techniques used to count the number of atoms coupled to a cavity mode based on the response of the resonator are well established [144]. These techniques have been applied to microwave clocks as means to generate entangled states which exhibit reduced quantum projection noise [24, 132, 52], with several groups looking to extend this work to optical clocks [153, 61].

Most NDD detection schemes rely on the interaction between the light with the atoms which imprints a phase shift on the beam as it propagates though the atomic medium. The phase shift measurement is similar to fluorescence detection as the signal is proportional to the number of atoms, but differs in the fact it does not depend on the power allowing for a much weaker probe beam [160, 78]. This makes the phase measurement advantageous as it reduces the atomic loss from the lattice resulting from scattering induced heating. This is often what is meant by non-demolition; the atoms remain in the lattice such that they can be probed again by the clock laser without having to be reloaded. This helps to improve the duty cycle of the clock and therefore suppresses frequency instability caused by the Dick effect as outlined in section 3.2. The cavity amplifies the induced phase shift by having the light pass through the atomic sample several times as it circulates withing the resonator. This increases the optical depth by a factor $\frac{2}{\pi}F$, where F is the cavity finesse [120]. Therefore, the cavity provides a convenient means of enhancing the atomic signal which can be detected by comparing the phase of the light incident on the cavity with the field leaking out.

The interaction strength between an atom and cavity is most often quantified by

the single-photon Rabi frequency g_0 given by:

$$g_0 = \frac{d}{\hbar}\sqrt{\frac{\hbar\omega}{2\epsilon_0 V_m}} \tag{4.2}$$

where d is the dipole matrix element transition which has a resonance frequency ω. The parameter V_m is the mode volume which is defined as:

$$V_m = \frac{\iiint_V |\mathbf{E}(\mathbf{r})|^2}{\max\left(|\mathbf{E}(\mathbf{r})|^2\right)}. \tag{4.3}$$

With a single photon inside the cavity resonant with the atomic transition, an atom placed in the resonator will undergo Rabi oscillations at a rate $2g_0$. Interestingly, g_0 does not depend on the mirror reflectivity. To put it another way, the finesse determines how much incident light was needed to couple a single photon into the mode – a detail entirely irrelevant from the perspective of the atom inside the cavity. However to understand the dynamics of the system, the field decay rate κ^{-1} along with the decay rate Γ of the excited atomic state become relevant, in particular vis-a-vis g_0. It is useful to parametrise the system by the atom-cavity cooperative C, given by:

$$C = \frac{4g_0^2}{\kappa\Gamma}, \tag{4.4}$$

which compares the coupling between the atom and the cavity to the coupling of the atom-cavity system to the surrounding environment arising from either loss of a photon or spontaneous emission. In the so-called strong coupling regime defined as $C \gg 1$, the atomic sample when coupled to the cavity will have a pronounced effect. In the few photon regime, a single atom can shift the cavity resonance by an amount far greater than its linewidth. It is this shift of the cavity resonance that provides us with the information regarding the atom number, motivating us to enhance our sensitivity by making C large. Unfortunately, this requires designing a cavity with a small mode volume which is somewhat incompatible with realising a state of the

art frequency reference. A smaller mode volume leads to fewer atoms being trapped (higher QPN) or a greater density shift for the same atom number. Instead, we can take a different route to increase the cavity phase shift by simply using more atoms. The single photon Rabi frequency increases with atom number according to $g = g_0\sqrt{N}$ giving rise to a new dimensionless constant, the collective cooperative defined as $C_N = NC$. In the limit of large atom-cavity detuning Δ, the collective cooperative between the atoms and cavity gives rise to a shift in the cavity resonance Δ_C given by [64]:

$$\Delta_C = \frac{g^2 N}{2\Delta}. \tag{4.5}$$

The other important consideration for NDD is that the atoms do not escape the trap during the probe time. The probe light can heat the sample either by scattering photons which, for n photons in the cavity mode, happens at a rate:

$$\Gamma_s = \frac{ng^2\Gamma}{\Delta^2}, \tag{4.6}$$

or by introducing a spatially varying AC-Stark shift ϕ_c given by:

$$\phi_c = \frac{2ng^2}{\Delta}. \tag{4.7}$$

which competes with the potential arising from the magic wavelength lattice. Table 4.2 shows the computed values for g_0, C, Δ_C, Γ_s and ϕ_c for the both the red cavity probing using 698 nm transition and the blue cavity probing using 461 nm.

There are two technical points that have yet to be addressed. First, how does one differentiate between changes of the cavity resonance arising from the presence of atoms versus other effects such as mechanical vibration or laser frequency noise? Secondly, how does the non-uniform coupling between the atoms and the cavity mode arising from the spatial mismatch between the trapping potential and the probing

Table 4.2: QED parameters for red and blue cavities

Parameter	Red Cavity	Blue Cavity
g_0: **Single photon Rabi frequency (kHz)**	10	540
C: **Single atom-cavity cooperative**	.36	0.05
Δ_c: **Cavity frequency shift (Hz)**	.02	77
ϕ_c: **Atom Stark shift (Hz)**	.1	288
Γ_c: **Scatter rate (s^{-1})**	$<<1$	27

light affect the signal? To address the second issue, it is possible to work instead with an effective atom number which is calculated by integrating over the atomic density profile weighted by the normalised mode volume [73], but this doesn't solve the problem that the spatial mismatch leads to a reduced cavity shift. Instead one might elect to use a trapping lattice which is commensurate with the probe lattice, for example as was implemented for rubidium by using a 1560 nm trapping lattice which was also frequency doubled to 780 nm for use as the probe [66]. Unfortunately, the additional requirement for our system that the lattice must be at the magic wavelength excludes such a solution. Instead we follow an approach similar to [79, 153] by probing simultaneously with two adjacent cavity modes such that the field at the centre of the trap is homogenous given the two standing waves are perfectly out of phase. In our scheme, which is shown pictorially in figure 4.6; the cavity length is tuned such that the two modes are placed symmetrically around the atomic resonance. In this scheme, two weak probe beams are resonant with cavity modes while stronger LOs placed symmetrically around the probe read out the imparted phase shift on the beam from the cavity. In the presence of atoms, the frequency spacing between these two modes increases given the dispersive nature of the atomic response. On the other hand, the signal is immune to technical sources of laser-cavity detuning noise given it is common to both modes and therefore cancels when

measuring the differential frequency shift. Technical details regarding how this scheme is implemented, specifically how the light used to probe the cavity is generated can be found in section 4.1.7.2.

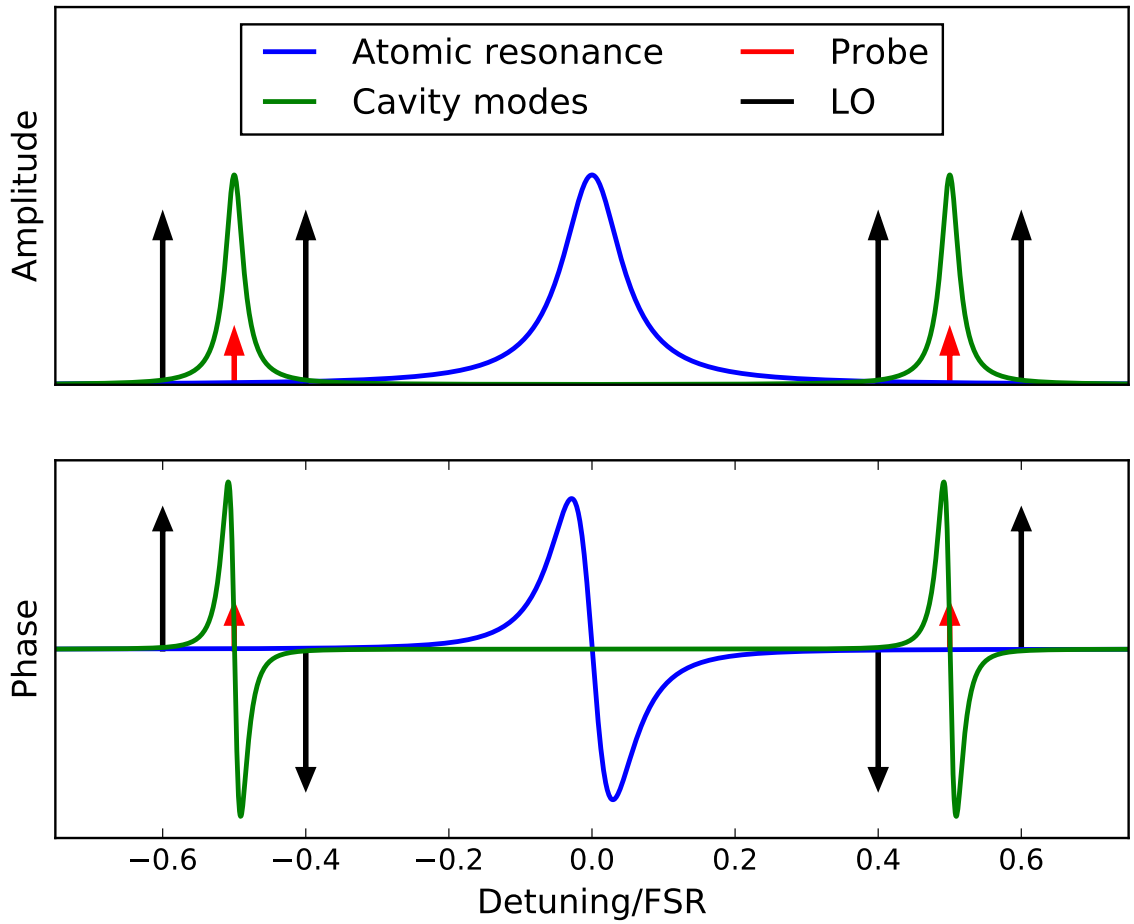


Figure 4.6: A simplified depiction of the NDD scheme considering only a single resonance. The scheme is based on a differential measurement of the detuning of two adjacent cavity modes centred around the atomic resonance. Two weak probe beams are resonant with the cavity modes and are used to produce a PDH error signal via interference with two nearby local oscillators. The phases are tuned such that the two error signals have opposite sensitivity such that the sum is insensitive to common mode detuning. However in the presence of atoms, the induced dispersion causes the cavity modes to shift in opposite directions to produce an error signal which is proportional to the atom number.

4.1.2.3 Flexure mount

The two in-vacuum, build-up cavities open the possibility to realise a 2D lattice at 813 nm. To form such a lattice the two cavity modes must overlap, requiring the relative separation between the modes be much less than the waist of the beams. The manufacturer's centration tolerances of 130 µm for the cavity mirrors eliminates trying to machine the cavity baseplate to the required levels of precision for the cavities to overlap. We therefore needed to include the ability to adjust their relative position after the mirrors are glued in place. To do this we integrated a flexure mount cut directly into the baseplate, which adjusts the tilt of the cavity mirror, see figure 4.7. The vertical tilt of the cavity mirror can can be adjusted by varying the diameter of the dowel pin which loads the flexure. A small change in cavity mirror tilt by an angle ϕ manifests itself in a translation of the beam in the centre of the cavity by an amount d equal to R$\phi/2$ where R is the radius of curvature of the cavity mirrors [23].

In regards to the design of the flexure, the main considerations are: (1) it applies sufficient clamping force to keep the dowel pin in place when loaded, (2) it has the necessary range of travel to compensate for the maximum displacement of the cavities, (3) the minimum change in cavity mode position is less than the cavity waist. The first two constraints are at odds as a stiffer mount will reduce the maximum displacement of the flexure prior to entering a regime of inelastic deformation. The clamping force and range of travel are determined by the dimensions of the flexure mount and the material properties of titanium, namely the Young's modulus and maximum tensile strength. Meeting the third design consideration is set by the minimum step size in dowel pin diameter and distance between the pin and pivot point of the flexure. We also considered trying to align the two cavities before gluing and then fixing them in place, similar to what was attempted in [130]. However we decided against this route given its high risk and technical challenges.

The flexure mount was designed using COMSOL which is easily interfaced with

Solidworks using the LiveLink toolkit. Once the general shape was decided, the specific dimensions were optimised using COMSOL's built-in routines. Finally, a prototype was built to verify the simulation in regards to clamping force and elastic deformation. The maximum cavity tilt was simulated to be 30 mrad which corresponds to clamping force of 3 N. Given the mirror radius of curvature, the maximum tilt angle corresponds to displacement of the cavity mode by ± 750 µm. In regards to the titanium dowel pins, machining tolerance sets the minimum incremental change in diameter to be 5 µm. Given that the dowel pin is approximately 12 mm from the flexure pivot point, this resolution corresponds to a change in cavity position of 20 µm.

One concern in our approach for overlapping the cavities is the mirrors could shift in position once installed in the chamber. To correct for minor changes in overlap, the opposite cavity mirror support has integrated flexure mount, but this time it is loaded with a Noliac shear stack piezo. The rationale for the shear piezo is it can actuate in both directions. In principle one could use a linear stack piezo and perform the initial alignment at half the maximum voltage. However in the event of the piezo failing or shorting, the cavities will no longer overlap. The shear stack provides 10 µm of travel and leads to an approximate displacement of cavity mode, given the lever arm, of 100 µm.

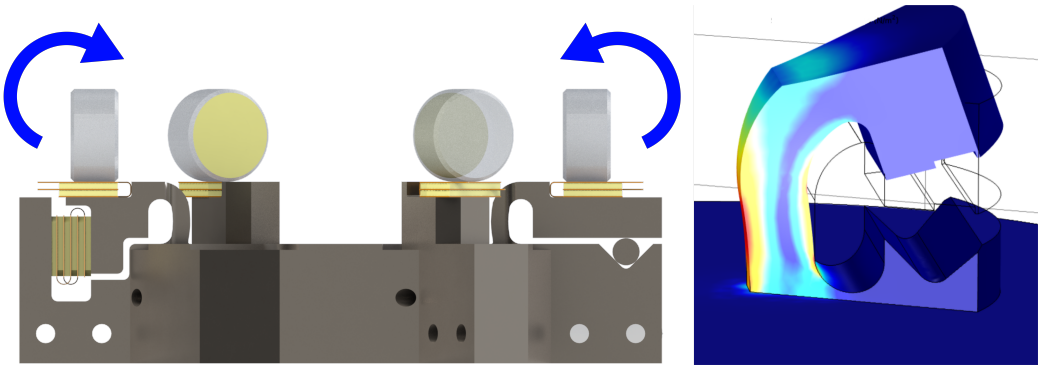


Figure 4.7: **Left:** Rendering of the cavity baseplate showing the integrated flexure mounts. The flexure on the right provides coarse adjustment by varying the diameter of the inserted dowel pin, while the flexure mount on the left uses a piezo for the fine tilt adjustment of the mirror which can be done while under vacuum. **Right:** Example of a COMSOL simulation of the flexure mount showing the induced stressed when deformed. The simulation was used to estimate the flexure mount's clamping force and elastic range of travel.

The flexures are machined in a two stage process. First the flexure shape is cut out of the blasplate using an electrical discharge machining (EDM) wire eroder. Next, the baseplate is removed and the flexure mounts are loaded using a dowel pin and metal block that matches the piezo dimensions. A final pass over the top of the baseplate leaves a flat level surface onto which the cavity mirrors are glued. The titanium dowel pins are machined using a lathe then ground to the required diameter. Given that it would require almost 100 dowel pins to cover the full range of travel with 5 µm resolution, we first measure the relative cavity mode displacement at two standard dowel pin diameters – 2.50 mm and 2.38 mm (imperial 3/32") – in order to pinpoint the required dowel pin diameter, see figure 4.8.

To measure the relative cavity mode, a 27G hyperdermic needle is attached to a three axis translation stage. The 813 nm light is coupled into both cavities and scanned over resonance while the needle is translated into the beam to extinguish the mode. Walking the needle position while monitoring the cavity transmission, the needle is used to map out the spatial position of the two cavity modes. Using this method, the relative offset between the cavities was measured with an accuracy

of 10 µm. Based on the measurements using the 2.5 mm and 2.38 mm dowel pins, the predicted diameter needed to overlap the cavities was 2.403 mm. To account for measurement uncertainty, we made pins ranging from from 2.385 to 2.415 mm.

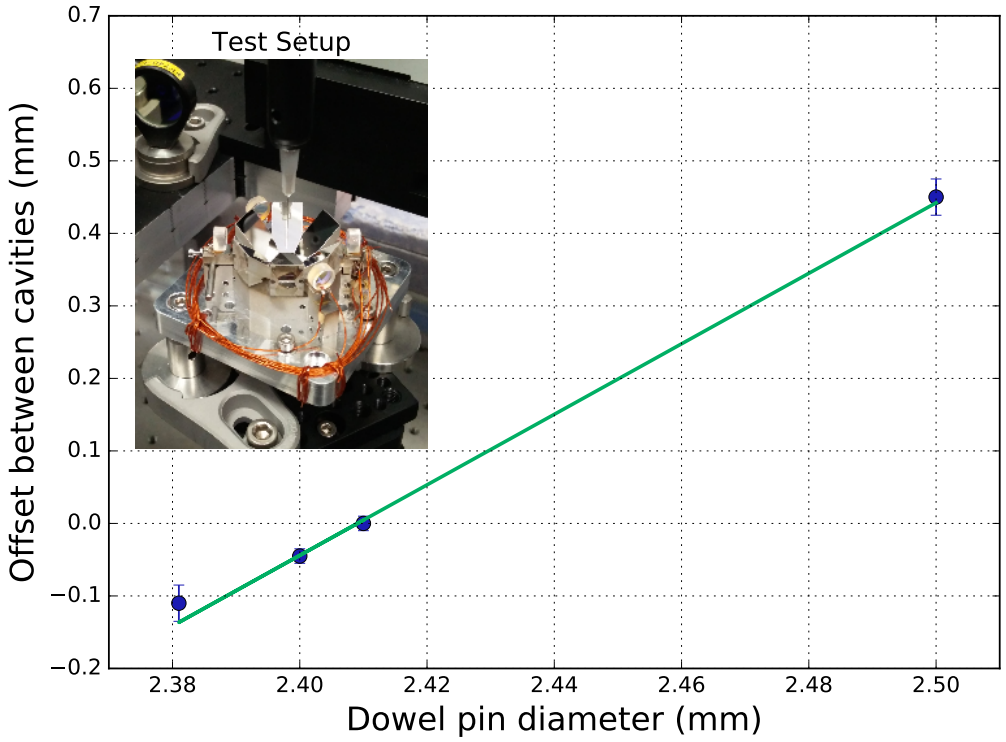


Figure 4.8: Plot showing the relative offset of the two cavity waists as a function of dowel pin diameter. A linear fit gives a slope of -11.7 which agrees well with that predicted by COMSOL. **Inset:** The jig used to measure relative cavity position consists of a hypodermic needle on a three-axis stage. The left flexure mount is loaded with a temporary dowel pin. To exchange dowel pins, a tapped hole in the top of the amount allows a M2 screw to be temporarily attached to serve as a handle.

4.1.2.4 Laser stabilisation to in-vacuum cavities

As outlined previously, there are many benefits to having a build-up cavity. However, it does increase the complexity of the experiment as the cavity length must be locked to the incident lattice light. In principle the shear stack piezo actuators should be

[3]In the future we may explorer replacing the analog loop-filter with a digital versions. This would allow us to easily implement notch filters to suppress these resonances and improve our bandwidth [124]

up for the job, but in practice mechanical resonances limit their response at high frequencies. For Sr2 the main cavity resonance is around 9 kHz which constrains the bandwidth of the piezo lock to about 2 kHz.[3] It is important that the cavity resonance tracks the lattice laser frequency as relative detunings between the two cause intensity and phase fluctuations of the intra-cavity field. This resulting modulation of the shape and depth of the trapping potential can resonantly excite the atoms into higher motional states through a process known as parametric heating. The heating rate scales with the square of the motional frequency, hence it is more of a concern for the axial direction compared to the radial as the trapping frequency ν_a is approximately 50 times larger than ν_r.

To estimate the required amplitude and phase stability of the lattice such that the lattice lifetime is not limited by parametric heating, I apply the results derived in [128]. The heating rates caused by fluctuations in the depth ($\epsilon(t)$) and position ($dz(t)$) of trapping potential – caused by changes in the intensity and phase of the circulating light, respectively – can be estimated using time dependent perturbation theory. Intensity and position fluctuations modify the harmonic oscillator Hamiltonian in accordance to:

$$H_{\epsilon,dz} = \frac{1}{2} m\nu_z^2 (1+\varepsilon)(z+dz)^2. \tag{4.8}$$

To calculate the heating rate, both the power spectral density $S_\varepsilon(f)$ of the peak trapping intensity I_0 (i.e. $I(t) = I_0(1+\varepsilon)$) and position fluctuation $S_z(f)$ must be estimated. Given the parity of the two perturbing Hamiltonians, to avoid heating caused by amplitude fluctuations we need to suppress noise around even harmonics of the trapping frequency; while for heating from position fluctuations, it is noise at odd harmonics of the trap frequency which leads to heating. Combining both effects,

the total heating rate Γ is given by:

$$\Gamma = \Gamma_I + \Gamma_z = \pi^2 \nu_z^2 S_\varepsilon(2\nu_z) + \pi^2 \nu_z^2 \frac{S_z(\nu_z)}{\langle z^2 \rangle} \tag{4.9}$$

where $\langle z^2 \rangle$ is the variance of the atomic wavefunction in the axial direction on a single lattice site. Changes in the position of each lattice site can arise from either vibration of the retro-reflecting mirror or fluctuations of the laser frequency. Mechanical vibrations of the mirror occur at acoustic frequency – much slower than the axial trapping frequencies – and are therefore not of a concern for heating. On the other hand, laser frequency noise can have significant spectral power at the trapping frequency giving rise to heating. The degree to which fluctuations in laser frequency changes the position of the lattice sites depends on the optical phase the light incurs between first passing through the atoms and when it returns after being retro-reflected. The longer the optical path between the atoms and retro-reflector, the greater affect a change in frequency will have on position. Specifically to convert between frequency fluctuations and changes in position one can simply scale the noise by a factor $\frac{2d}{\nu_L}$ where d is the distance between the atoms and retroreflector and ν_L is the lattice laser frequency. Rewriting equation 4.9 in terms of laser parameter and neglecting contributions of mirror fluctuations to motional heating, we arrive at:

$$\Gamma = \Gamma_I + \Gamma_z \approx \pi^2 \nu_z^2 \left(S_\varepsilon(2\nu_z) + \frac{4d^2 S_f(\nu_z)}{\nu_l^2 \langle z^2 \rangle} \right) \tag{4.10}$$

where $S_f(\nu_z)$ is the power spectral density of the laser frequency noise at the motional trapping frequency. Ideally, the heating rates arising from both frequency and intensity fluctuations should be less than $0.1\,\mathrm{s}^{-1}$. For intensity noise, a typical trapping frequency of 60 kHz gives an upper bound for the fractional intensity deviations at twice the trapping frequency of $\sqrt{S_\varepsilon(2\nu_z)} < 10^{-6}\ \mathrm{Hz}^{-\frac{1}{2}}$. For frequency fluctuations, the bound on laser frequency stability depends strongly on the distance between the

mirror and the atoms. For our cavity, this distance is approximately 15 mm which, given that for our trap $\sqrt{\langle z^2 \rangle} \approx 250\,\text{nm}$, results in a maximum frequency noise of $\sqrt{S_f(\nu_z)} < 3\,\text{kHz/Hz}^{\frac{1}{2}}$, corresponding to an integrated white noise linewidth of 23 MHz [36]– well above the specified free-running linewidth of our Ti:sapphire even before locking it to the transfer cavity.

An important consideration when using a build-up cavity for the lattice is that the power enhancement depends on the relative frequency detuning between the laser and cavity resonance. Therefore, the cavity can convert frequency noise of the laser or fluctuations in its length to variations of the intra-cavity intensity giving rise to heating. Previously we derived a limit for the maximum intensity noise of $\sqrt{S_\varepsilon(2\nu_z)} < 10^{-6}\ /\sqrt{\text{Hz}}$. Working under the assumption the laser is nominally on resonance such that the sensitivity of the intensity to laser frequency noise is quadratic, this constraint on intensity noise can be converted to a maximum relative laser to cavity frequency instability of $\sqrt{S_f(\nu_z)} < 700\,\text{Hz}/\sqrt{\text{Hz}}$ given the cavity linewidth is 700 kHz. This relative frequency noise would require that the laser has an integrated white-noise linewidth of 3 kHz relative to the cavity. Once again we have neglected instabilities of the cavity length as mechanical resonances in the support structure are well below the trapping frequencies relevant for parametric heating. However if the laser-cavity detuning caused by length fluctuations are large, it can enhance the sensitivity of the intra-cavity intensity to laser frequency noise given that the laser is no longer near resonance. The complexity introduced by the coupling between frequency noise and the intra-cavity intensity with a conversion factor that dynamically changes with cavity length fluctuations makes it difficult to estimate the performance specification to ensure a low heating rate.

Initially we simply tried to lock the cavity length to the laser via feedback to the in-vacuum piezos – sadly without much success. When looking at the transmitted intensity through the cavity, it exhibited intensity modulation at twice the mechanical

resonance frequency with a depth of 10% corresponding to an excursion of the cavity's frequency (length) at the 100 kHz (8 picometers) level. To improve the intensity stability, we locked the intensity of the light transmitted through the cavity. To implement this lock, we coupled light into the cavity through a waveguide, electro-optic modulator (EOM) which we used to apply auxiliary sidebands at 1 GHz which are off-resonant from any cavity modes. By changing the modulation depth, a variable amount of power is diverted from the carrier to control the intensity through the cavity. We also implemented a fast frequency lock of the laser to the cavity as shown in figure 4.9 using the EOM. We replicated the EOM electronic sideband offset lock [145], the same scheme used to stabilise the lattice to the transfer cavity, for both of the Sr2 in-vacuum cavities. Given the fast frequency response EOM, we can achieve lock bandwidths of 500 kHz via feedback to the frequency modulation port of the synthesiser used to generate the sideband. In order to prevent the lattice from drifting away from the magic wavelength, the Sr2 cavity piezos are used to servo sideband offset frequency to a set value. One added complication is the EOM to the transfer cavity is not in the optical path to the Sr2 cavity. This may not seem like an issue given that the EOM voltage is servoed to zero. However, for our NDD scheme it is crucial that the lattice light sent to Sr2 tracks the frequency noise of the transfer cavity even at high frequencies. Therefore, we feed-forward the fast frequency corrections applied to the transfer cavity EOM to an AOM before the Sr2 cavities.

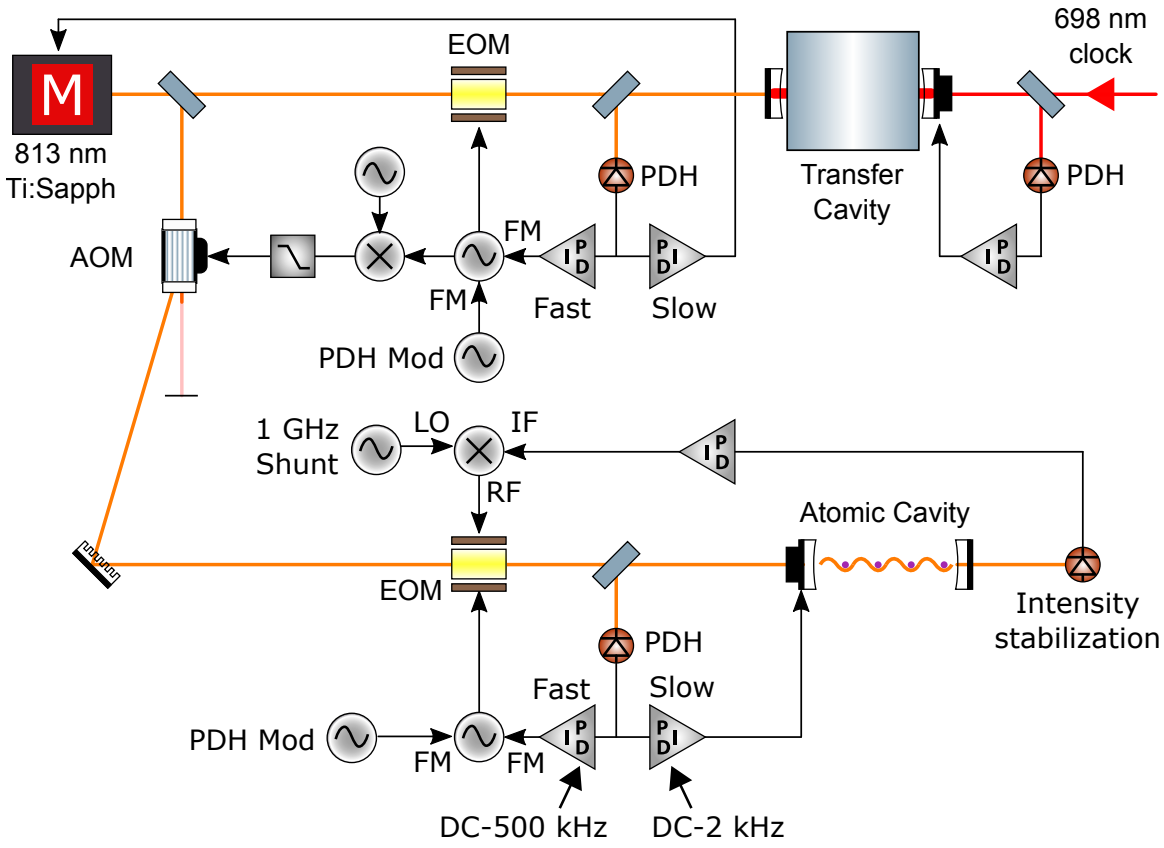


Figure 4.9: Schematic showing how the lattice laser is locked to the transfer cavity and the Sr2 in-vacuum cavities (diagram only shows one cavity). We implement a fast feedback lock to both the transfer cavity and Sr2 cavity using an EOMs operated in the electronic sideband configuration. For both locks we servo the DC voltage applied to the EOM to zero by feeding back to either the Ti:sapphire laser for the transfer cavity lock or the in-vacuum piezos for the Sr2 cavity lock. Further details regarding this scheme are given in the text. For the slow feedback to the Sr2 cavity piezos, the bandwidth is limited to 2 kHz as result of mechanical resonances in support structure. Figure credit to Alvise Vianello.

4.1.3 Zeeman slower

The thermal beam of atoms exiting the effusive source are too fast to be efficiently trapped by the MOT. The standard approach to slow the atoms to within the capture velocity of the MOT is to use a Zeeman slower. The ingenuity of the device is its use of a spatially varying magnetic field that Zeeman shifts the atomic levels to compensate for the changing Doppler shift as the atoms slow down. When the magnetic field is

tuned correctly, an atom will experience constant deceleration as it progresses down the slower. This deceleration is set by the magnetic field gradient and is typically chosen to be between 50-75% of the maximum deceleration a afforded by the finite photon scatter rate[4]. The maximum velocity class which can be decelerated by the slower is given by $\sqrt{2aL}$ where L is the slower length. For an oven temperature of 500 °C, 50% of atoms have an longitudinal exit velocity below 500 ms^{-1}. To slow an atom going at such a speed at a constant rate of $0.75a_{\mathrm{max}}$ would require a 23 cm long slower. Extending the slower beyond this value will lead to marginally decreasing improvements in the MOT loading rate.

One annoyance of implementing a Zeeman slower is that the atomic beam inevitably starts to coat the viewport through which the slowing beam is sent. For strontium this is particularly problematic as preventing the deposition of strontium (and the resulting decrease in the optical transmission) requires heating the viewport which compromises the temperature homogeneity of the chamber. To solve this problem, we couple in the slowing beam using a periscope formed from two silver coated prisms glued to a machined DN16 blank. The blank is sandwiched between the chamber and the DN16 viewport to protect the glass from being coated. The prisms will eventually be coated in strontium, but the reflectance should still be better than 50% [54]. The prism periscope is shown in figure 4.10.

For the slower we use a transverse magnetic field generated using permanent magnets as shown in figure 4.10. To estimate the required strength and position of the magnets, we first approximate them as point-like dipoles orientated to along the x-axis for which the magnetic field is given by:

[4]For the cycling transition at 461 nm this maximum deceleration is approximately 100,000 g's.

$$B_x = \frac{\mu_0 m}{4\pi}\left[\frac{2x^2 - y^2 - z^2}{r^5}\right] \tag{4.11}$$

$$B_y = \frac{\mu_0 m}{4\pi}\left[\frac{3xy}{r^5}\right] \tag{4.12}$$

$$B_z = \frac{\mu_0 m}{4\pi}\left[\frac{3xz}{r^5}\right], \tag{4.13}$$

where $r = (x^2 + y^2 + z^2)^{1/2}$, μ_0 is the vacuum permeability ($\mu_0 = 4\pi \times 10^{-7}$), and m is the magnetic moment. For a cylindrical magnet of radius r and height h made of a material with magnetic strength H, the magnetic moment is given by:

$$m = H\pi r^2 h. \tag{4.14}$$

Rather than the magnetic field strength, the strength of the permanent magnet is often specified by suppliers in terms of its residual inductance $B_r \triangleq \mu_0 H$. For neodymium N35 grade magnets used for the slower, B_r is typically on the order of 1.2 T at its surface.

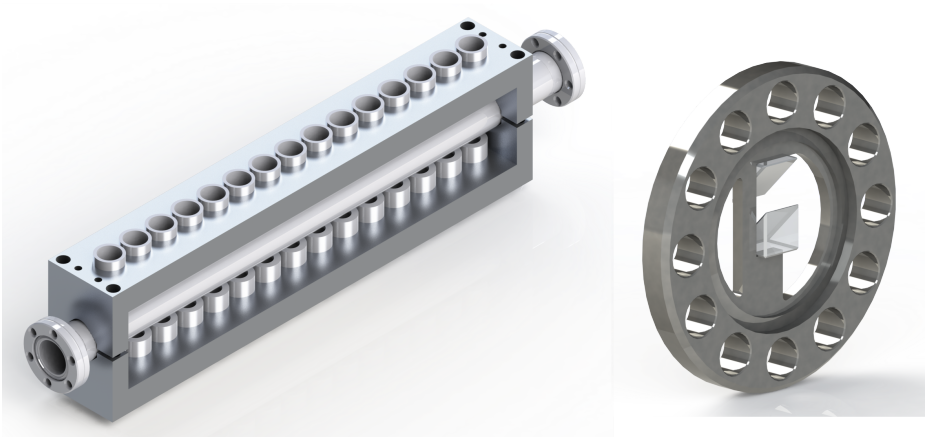


Figure 4.10: **Left:** Rendering of the permanent magnetic Zeeman slower. The strength of the magnet field is set by number of magnets placed inside the threaded casings and then fine tuned by setting their final positions. **Right:** Rendering of the prism periscope used to couple the slowing beam into the chamber while preventing the viewport from being coated with strontium.

The final position of the magnets can be fine tuned by screwing the casing holding the magnets in and out of the sideplates. The final field is shown in figure 4.11 along with the theoretical profile as computed using the method outlined [48].

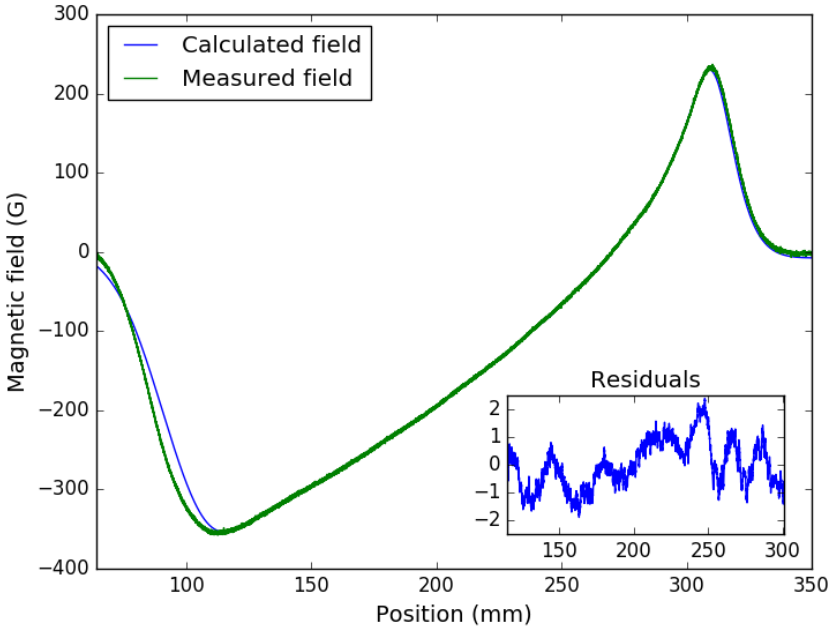


Figure 4.11: Comparison between the theoretically calculated slowing field profile and the one measured after fine tuning the position of the magnets. **Inset:** The residual error between the fields is below 1 G.

The final consideration for the Zeeman slower is compensating for residual magnetic fields at the position of the lattice. We considered using electromagnetic coils to produce the slowing field which could be switched off during spectroscopy. However, this adds an additional heat source leading to inhomogeneities in the BBR environment. Ideally we want a uniform temperature around the atoms and therefore elected to forgo this option. The three pairs of compensation coils which are aligned along the Cartesian directions (x,y,z) allow us to zero the magnetic field at the centre of the lattice and cancel the diagonal elements of the magnetic field gradient tensor (i.e. dB_i/dr_i). However, they cannot cancel gradients of the form dB_i/dr_j where $i \neq j$. Such gradients lead to a broadening of the clock transition given the spatial distribu-

tion of the atoms. To compensate the gradients generated by the slower, we make use of the fact that far from the slower its field resembles that of a dipole. We therefore compensate for the slower's field by using an additional magnet pair on the opposite side of the chamber which mirrors its field. Once properly tuned, the magnetic field gradients are all below $0.1\,\mathrm{mT\,m^{-1}}$, see figure 4.12.

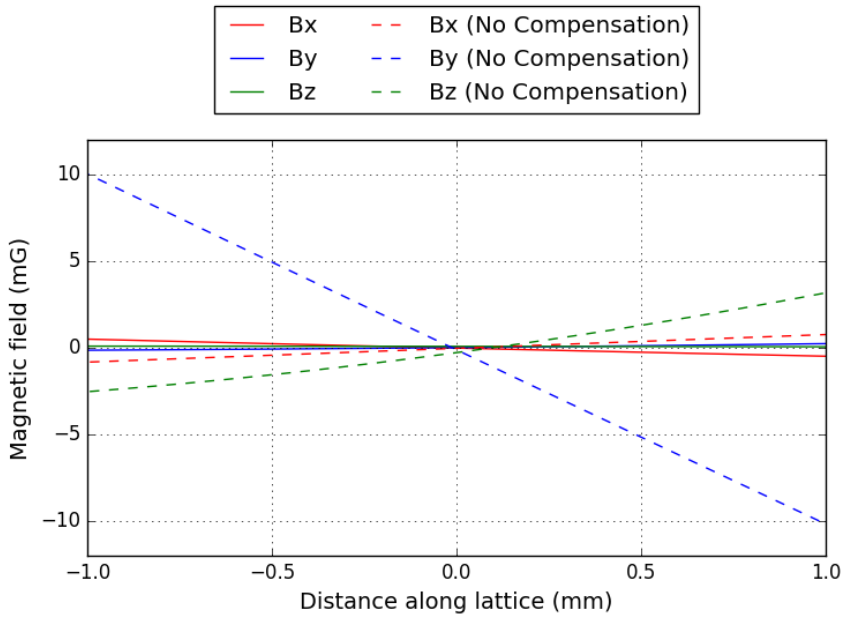


Figure 4.12: Variation of the magnetic field along the lattice formed by the 'blue' cavity. The field components Bx, By and Bz corresponds to the magnetic field components along the MOT axis, Zeeman slower axis, and vertical direction.

4.1.4 MOT and compensation coils

The first stage cooling of strontium requires large magnetic field gradients – in Sr1 we use $4\,\mathrm{mTcm^{-1}}$ along the coil axis which requires 70 A of current. To produce the field we use a pair of current carrying solenoids orientated in a anti-Helmholtz configuration. Ideally we would like to avoid the solution of using several windings to increase the gradients as the increased coil inductance will limit the switching time. Using a compact chamber helps to reduce the required power to produce the fields as

we can bring the coils close to the atoms. The coils are shown in figure 4.13, and are wound using hollow kapton-coated wire which allows them to be cooled by flowing water through their core. The coils are potted[5] in a 3D printed graphite former around which a secondary coil pair is wound in a Helmholtz configuration to provide a compensation field.

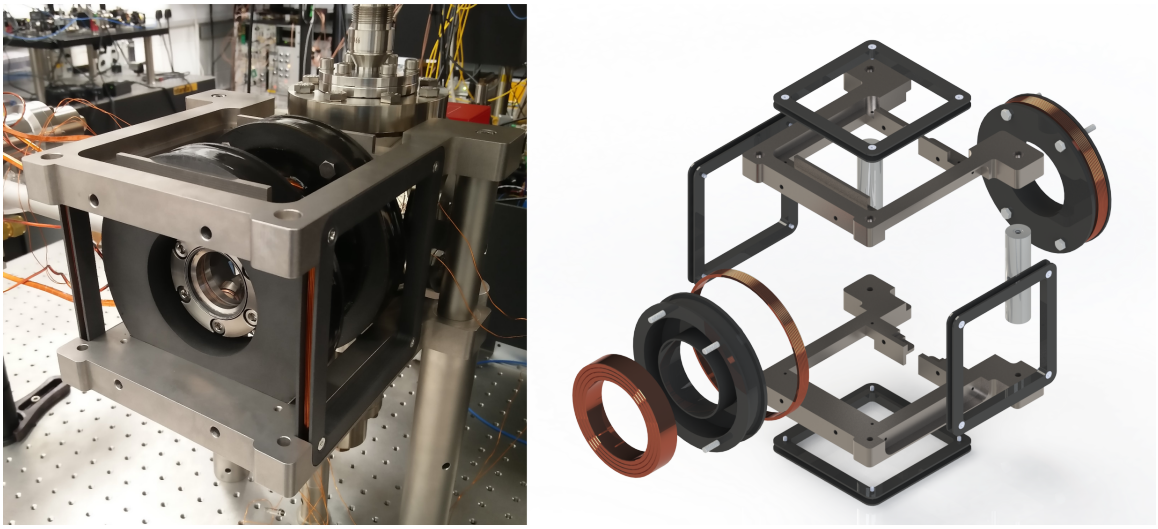


Figure 4.13: **Left:** Titanium coil mount holding the MOT and compensation coils. The anti-Helmholtz MOT coils are potted in a graphite former with the Helmholtz coils being wrapped around the exterior. The MOT coils provide axial magnetic field gradient of $0.07\,\mathrm{mTcm^{-1}A^{-1}}$ while the Helmholtz coils provide bias fields of $0.225\,\mathrm{mTA^{-1}}$. **Right:** Exploded rendering of the coil assembly showing the various parts.

One lesson taken from Sr1 was that directly attaching the coils to the chamber increases the temperature inhomogeneity of the thermal environment surrounding the atoms making the BBR induced frequency shift harder to characterise. For this reason the Sr2 MOT coils are held in place by a titanium mount such that they do not touch the chamber. The mount also is used to secure the other two pairs of compensation coils. To reduce further the temperature gradients caused by the coils, we cool the

[5]We use MG Chemicals 832HT-375ML Black High Temperature Epoxy Potting Compound as one coil cannot be removed from the chamber and must withstand a bakeout temperate of 130 °C

MOT coils using a dedicated chiller[6] such that we can set the water temperature to minimise gradients. The chiller provides a sufficient flow rate such that the water heats by less than a degree when the coils are running.

4.1.5 Laser systems

With the exception of the triplet MOT, the light needed for the operation of Sr2 is generated by extending the laser systems used for Sr1. This includes the cooling light at 461 nm, the repump beams at 497 nm and 679 nm, the optical lattice, and finally the clock laser. As previously mentioned, details regarding the Sr1 laser systems, with the exception of the 497 nm system, can be found in [49, 51].

4.1.5.1 461 nm laser system

To generate the 461 nm light needed for the MOT, Zeeman slower and probe beams an obsolete Toptica BoosTA previously operated at 813 nm was refurbished with a 922 nm tapered amplifier (TA) chip. This infrared light is then frequency doubled using periodically-poled lithium niobate (PPLN) waveguide SHG module, manufactured by NTT Electronics Corporation. Typical conversion efficiency achieved are approximately 150% /W, resulting in 100 mW of 461 nm light for 500 mW of 922 nm light coupled into the waveguide.

The TA is seeded using light from a homebuilt master ECDL [7] locked to the transfer cavity. It is the same design as our home-built clock laser, details of which can be found in [49]. A Toptica DL Pro – used to seed the commercial second harmonic generation system used for Sr1 – is offset frequency offset locked to the homebuilt ECDL. The frequency/phase discriminator needed for the offset lock is generated using a development board for the HMC984LP4E digital frequency/phase to voltage

[6] VascuTherm 5 manufactured by Thermotec.
[7] The ECDL uses a Toptica diode, model number: LD-0935-0100-AR-1

converter chipset. The offset locking scheme enables Sr1 and Sr2 to independently cool different strontium isotopes as needed for frequency ratio measurements between ^{87}Sr and ^{88}Sr. The various frequency detunings of all the relevant laser beams are shown in figure 4.14.

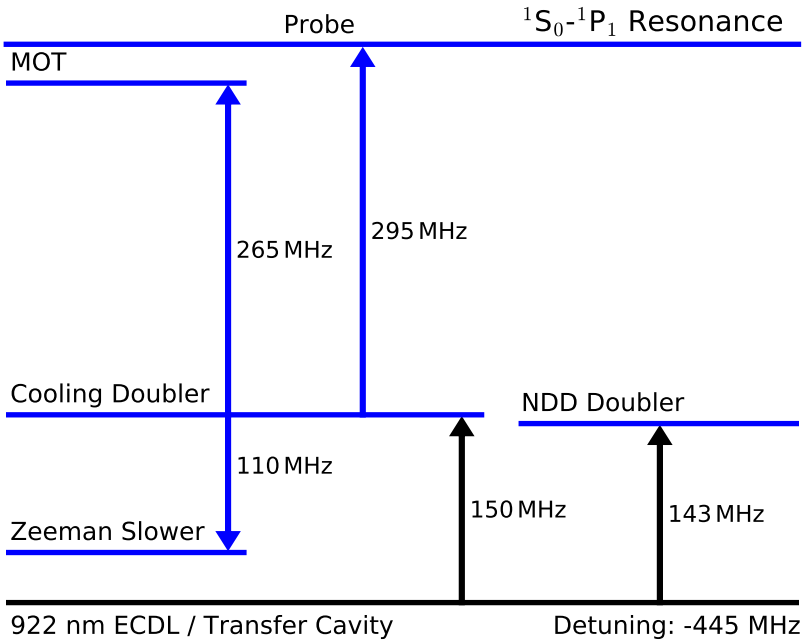


Figure 4.14: For Sr2 a new master ECDL was built and locked to the transfer cavity. The above diagram shows the various detunings of all the laser beams needed for cooling strontium and needed for the NDD scheme. All detunings are with respect to the $^1S_0$-$^1P_1$ transition. Labels in black indicate the laser is operating a 922 nm, while blue labels are used for the frequency doubled light at 461 nm. For clarity, frequency detunings in the infrared have been scaled to the visible by multiplying by two.

4.1.5.2 2.92 µm system for a triplet, magneto-optical trap

Using a metastable MOT to cool and trap alkali-earth earth atoms has previously been demonstrated with calcium [47, 43] and magnesium [84, 119] – but it has yet to be implemented in strontium. In the case of Ca and Mg, the triplet MOT or quenched cooling [29, 63] of the $^1S_0 \rightarrow {}^3P_1$ transition was a necessity as the inter-combination line is prohibitively narrow for direct laser cooling. However this is not the case for strontium – the inter-combination line at 689 is effective at reaching sub

microkelvin temperatures – motivating the reader to ask why one would pursue this option for cooling. There were several reasons why we chose to investigate using the triplet MOT: the hope was the MOT would produce colder samples more quickly compared to the conventional 689 nm cooling. Also, the long wavelength and low scatter rate, estimated to be approximately 100 kHz, results in a Doppler limited temperature of 2.4 µK. Moreover, the $J = 2 \rightarrow J = 3$ level structure means it is simpler to operate (we only need to address a single hyperfine transition) and allows for polarisation gradient cooling which opens up the possibility of pushing towards the recoil temperature of 12 nK. Another advantage of the triplet MOT over cooling on the inter-combination transition is it can be operated simultaneously with the first stage cooling. When the repump laser is turned off during 'blue' MOT, atoms will slowly accumulated in $^3P_2$ and can be captured in a the triplet MOT. This could reduce the time needed for laser cooling which is advantageous for realizing a stable lattice clock as the improved duty cycle helps to mitigate frequency instability arising from the Dick effect.

As a result of the low scatter rate and long wavelength, the intensity needed to saturate the transition is only 500 nW cm^{-1}. Producing enough power to saturate the transition is not difficult, rather the main challenge is generating the light in such a way to stabilise the frequency as required for laser cooling. We elected to make use of the transfer cavity which serves as a stable frequency reference for most of the cooling and trapping lasers. Unfortunately, the mirror coating does not extend into the mid-IR which prohibits directly locking 2.9 µm to the cavity. We instead generate the light using difference frequency generation in a PPLN non-linear waveguide module custom made for us by NTT Electronics Corporation. Given the freedom in choose the intermediate laser sources for frequency generation, we elected to leverage the high power afforded by the Ti:sapphire lattice laser to generate the 2.9 µm light. This required us to build an ECDL at 636 nm for frequency differencing with 813 nm. With

29 mW and 120 mW of transmitted light at 636 nm and 813 nm light, respectively, the device generates 800 µW of infrared light – implying a generation efficiency of 9.6% /W referenced to the transmitted power.

The hunt for high power diodes which lased reliably was arduous. Eventually we settled on the Oclaro HL63263DG after seeing other groups have success using it in a Littrow ECDL configuration [147]. Although better behaved than other diodes tested[8] , it is still temperamental – lasing only over a small current range with a mode hop free running range of about 1 GHz. However, it was sufficiently stable for laser cooling and routinely stays locked for over one day.

A fraction of the 636 nm light is picked off before being combined with the lattice and sent into waveguide EOM [9]. By phase-modulating the light, we can reference the ECDL frequency to the transfer cavity via a PDH lock to the first-order sideband via feedback to diode current. By changing the modulation frequency, we can control the carrier frequency relative to the cavity mode. This allows us to shift and modulate the carrier frequency during the experimental cycle which is needed for operation of the triplet MOT.

4.1.5.3 497 nm repump laser

For Sr2 we replaced 707 nm repump laser used to clear out the $^3P_2$ state with a laser operating at 497 nm which excites atoms to the $5s5d\ ^3D_2$. For ^{87}Sr, repumping is complicated by the presence of nuclear spin which splits the metastable state $^3P_2$ into five hyperfine states with a quantum number F ranging from 5/2 to 13/2. To account for this, we frequency modulate the laser to add first and second-order sidebands which simultaneously repump every hyperfine state. Experimentally, we find the optimal modulation to be 710 MHz with a absolute detuning referenced to the $F =$

[8]These included HL63142DG, HL63133DG (lases about as well as the Oclaro HL63263DG, but after two different diodes died we switched models), L638P150, and L638P200.

[9]Provided by Jenoptik, model number PM635.

$11/2 \rightarrow F' = 11/2$ of 38 MHz. In addition to the sideband modulation, we add an additional slower modulation at 15 MHz with deviation of approximately 50 MHz to account for residual frequency offsets from the hyperfine transition resonance. To see why this frequency modulation scheme is efficient we need to look at the relative frequency detunings of the hyperfine transitions as shown in table 4.3 [136].

	$F' = 5/2$	7/2	9/2	11/2	13/2
$F = 5/2$	-1584 (-35)	-1835			
7/2	-810 (29)	-1061	-1383		
9/2		-82 (48)	-404	-797 (42)	
11/2			766	373	-92 (38)
13/2				1719	1255 (-35)

Table 4.3: Relative detunings in MHz for hyperfine transitions between $^3P_2 \rightarrow ^3D_2$. The cells highlighted in green indicate the transitions we drive to repump atoms from the $^3P_2$ hyperfine levels. The values in brackets indicate the nearest detuning from the carrier or sideband (1^{st} or 2^{nd}) when the laser is frequency modulated at 710 MHz.

The choice of hyperfine transitions used for repumping was also dictated by the relative transition strengths $S_{FF'JJ'}$, set by the Clebsh-Gordan coefficients, which are given by:

$$S_{FF'JJ'} = (2F+1)(2J'+1)\begin{Bmatrix} J & J' & 1 \\ F' & F & I \end{Bmatrix}^2. \tag{4.15}$$

These values are tabulated in table 4.4.

	$F' = 5/2$	7/2	9/2	11/2	13/2
$F = 5/2$	0.48	0.39			
7/2	0.52	0.02	0.47		
9/2		0.59	0.06	0.39	
11/2			0.47	0.34	0.23
13/2				0.27	0.77

Table 4.4: Relative line strengths for the hyperfine transition as given by equation 4.15. The cells highlighted in green indicate the transitions we drive to repump atoms out of the $^3P_2$ hyperfine levels.

To generate the repump light we use a home-built ECDL operating at 994 nm which is frequency stabilised using a commercial wave-meter[10]. The diode outputs

[10]Model: High Finesse WS8

enough power to couple 130 mW into the waveguide EOM[11] which is used to apply the frequency modulation to the light. The light is then doubled using a waveguide enhanced PPLN SHG module[12] which, after fibre coupling and insertion losses, outputs 7 mW of 497 nm light.

4.1.6 Magneto-optical trap light delivery and imaging

Advantageous to the pyramid MOT design is it requires only a single beam—reducing the amount of required optics and optical access. However, electing to work with 2.9 µm proves challenging as it is incompatible with many of the materials used for standard optics. We therefore rely on parabolic mirrors and knife edge prisms for beam shaping and combining, respectively. Out of the waveguide difference module the 2.9 µm light, along with the 636 nm and 813 nm beams are collimated using a parabolic mirror. The three beams are then separated using a CaF_2 prism after which the 636 nm and 813 nm are picked off using a knife-edge mirror. Fortunately we do not have to collimate the mid-infrared light directly; it is sufficient to set the position of the parabolic mirror such that the 636 nm and 813 nm beams are well collimated. The 636 nm and 813 nm are separated using a dichroic filter and intensity stabilised using AOMs before the waveguide differencing crystal 9. The 813 nm is recycled for seeding the lattice build-up cavities. The lattice beam is coupled into a 50/50 fibre splitter with each branch having a waveguide EOM as needed for independently locking the two different cavities. Separately, the 461 nm light is combined with the repump beams using a dichroic splitter. The 2.9 µm and 461 nm beams are focused using a parabolic mirror and lens, respectively, and combined on a knife edge prism mirror. The co-propagating MOT and repump beams are collimated with a second

[11]The waveguide EOM is supplied by Photline.

[12]The second harmonic generation module is supplied by NTT Electronics Corporation. We purchased two devices; the first had a conversion efficiency of 145%/W while the second had an efficiency of 450%/W.

parabolic mirror before being sent to the chamber. See figure 4.15 for more details regarding the MOT light distribution system.

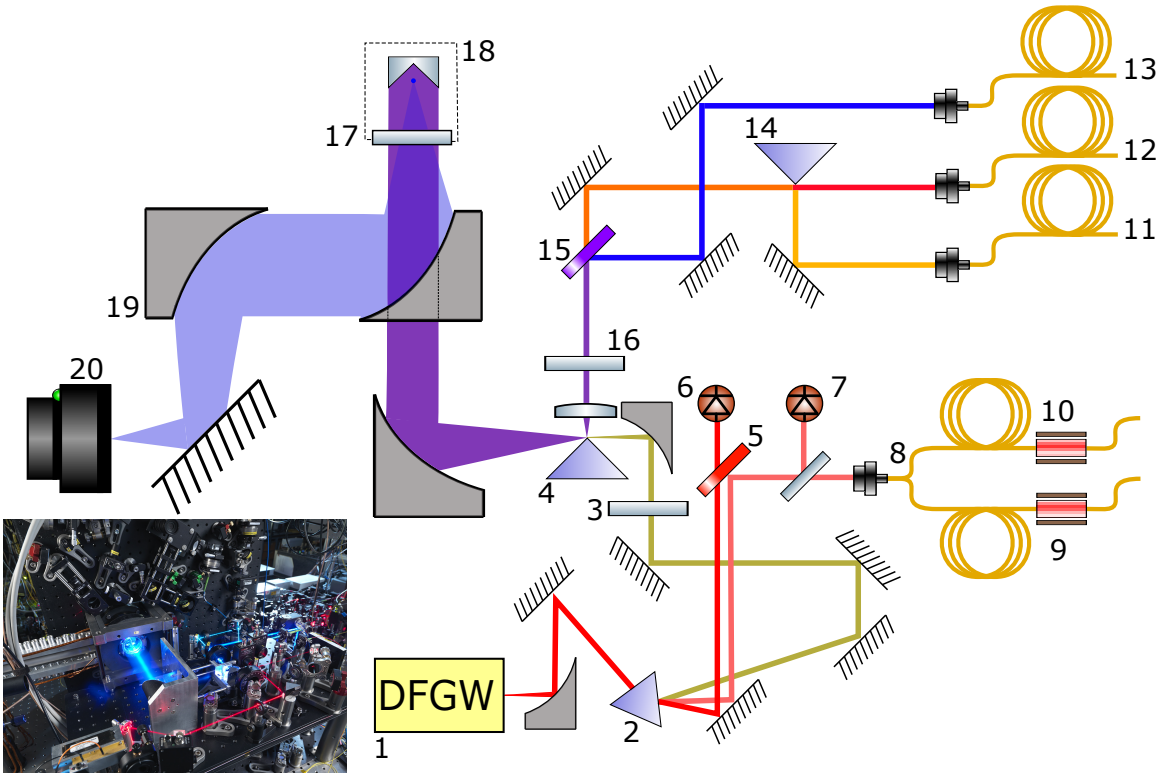


Figure 4.15: Diagram showing the optics needed for the MOT. The light for both the first and second stage cooling along with the repump light is combined on a breadboard and delivered to the chamber. There is an delivery fibre 689 nm which is not currently used. **Inset:** The optics setup and a legend for the diagram is shown below.

Label	Description
1	Difference frequency generation module
2	CaF_2 dispersion prism to separate 2.9 µm light
3, 16	Quarter-wave plates for 2.9 µm and 461 nm
4	Knife edge prism to combine 2.9 µm with other beams
5, 15	Long-pass dichroic filters
6, 7	Photodiodes for power stabilisation of 813 nm and 636 nm light
8	50/50 fibre splitter to deliver light to in-vacuum cavities
9, 10	EOM used to lock the 813 nm light to the two build up cavities
11, 12, 13	Fibre launch for repump, 689 nm, and 461 nm light
14	Knife edge prism to combine repump and 689 nm light
17	Low-OH fused silica viewport on vacuum chamber
18	In-vacuum pyramid MOT
19	Parabolic mirror pair used for imaging
20	PCO scientific CMOS camera

The other important consideration is imaging the atomic sample such that detection noise does not significantly compromise the stability of the clock. Given that much of optical access around the chamber is taken up by the MOT beams and cavity optics, finding room for a high numerical aperture imaging system proves challenging. We elected to collect the fluorescence with a three inch parabolic annulus mirror custom manufactured by Thorlabs. A 3.5 cm hole in the centre of the mirror allows passage of the MOT beams while the outer rim provides adequate collection of the fluorescence. The effective numeric aperture of the parabolic mirror imaging system is approximately 0.15—which is an improvement over Sr1. The collected fluorescence is imaged onto the camera[13] using a second parabolic mirror which provides a magnification of 1.5x. For wide-angle imaging, we also use a mirror and lens instead of the parabolic mirror setup smaller off-axis image distortion.

4.1.7 Cavity and clock light delivery

Several beams must be coupled into the two cavities; for the red cavity this includes the 689 nm and 813 nm light which enter through opposite sides, while for the blue cavity 813 nm, 461 nm and clock light all need coupling into the system. Included in the setup are photodiodes for the PDH signals, monitoring the lattice transmission, and finally for the NDD detection. Figure 4.16 shows the setup. For the NDD photodiode we use a home-built transimpedance amplifier based on low noise op-amp [14] and blue-enhance photodiode with a quantum efficiency of 91% at 461 nm [15]. The photodetector is shot-noise limited for power above 200 µW.

[13]We use a scientific CMOS camera manufactured by PCO, model number: pco.edge 4.2 LT

[14]part number: TI LMH6624 (We observed similar noise performance when testing the TI OPA657 which exhibits lower input current noise on account of its JFET-input stage, but suffers from larger voltage noise.

[15]part number: Hamamatsu S5973-02 (blue-enhanced)

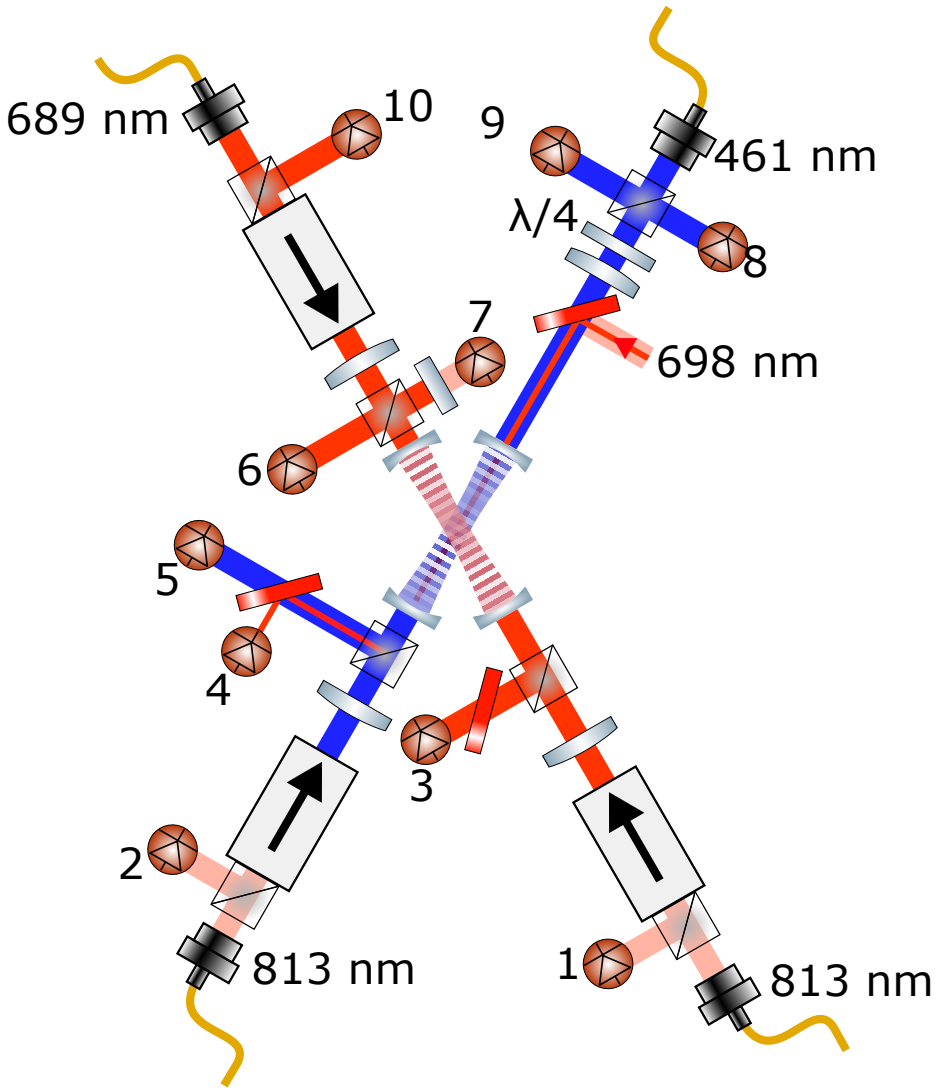


Figure 4.16: Diagram showing the optics for coupling in the lattice, NDD, and clock beams into the two cavities. Beam steering mirrors have been omitted for clarity. Below is a table indicating the function of all labeled photodiodes. For details regarding the delivery of 698 nm light and path length stabilisation, see section 4.1.7.1.

Photodiode	Purpose
1,2	PDH lock for lattices
3	Transmission of 689 nm beam
4	Transmission of 698 nm clock beam
5	Transmission of 461 nm beam
6	Intensity stabilisation of 689 nm delivery
7	Transmission of 813 nm beam
8	NDD signal of 461 nm probe beam
9	Intensity stabilisation of 461 nm delivery
10	PDH signal for 689 nm probe beam

4.1.7.1 Path length stabilisation for Sr2

As outlined in section 3.3.8, to realise a stable and accurate clock it is important to fix the optical path length from the cavity to the atoms. With Sr1, it was sufficient to place a retro-reflector at the end of the delivery fibre to reflect part of the light back to stabilise the path. However, there is still approximately one meter of uncompensated path between the path length stabilisation (PLS) reflector and the atoms. Ideally one wants to stabilise the optical path all the way to the lattice, but this is challenging as the small amount of probe intensity needed to drive the clock transition limits the signal used to stabilise the path.

One option is to use a strong, slightly off resonant probe beam (detuned by a few kilohertz) for path length stabilisation and then switch on a resonant sideband to drive the transition. The benefit to this scheme is the path length stabilisation is always engaged, but it has the drawback of introducing an AC-Stark shift caused by the PLS beam. We instead implement a new method which overcomes the problem of poor signal strength when using weak probe intensity without requiring an additional PLS beam. We invert the problem. Instead of stabilising the clock beam to the lattice, we stabilise the lattice to the clock beam. First the clock light is delivered to the Sr2 science chamber via a path length stabilised fibre using the same technique as for Sr1. Next we want to stabilise the path between retroreflector used to stabilise the clock delivery and the build-up cavity. To do this we take the 813 nm cavity transmission and split it using a beam splitter. Half the light then goes through an AOM before retroreflecting off the same mirror used for the clock delivery fibre PLS. The reflected lattice (now shifted in frequency by twice the AOM frequency) is compared to the light exiting cavity via homodyne detection of the resulting beat-note signal. The beat-note frequency is then phase-locked to a microwave reference traceable to temps atomique international (TAI). The phase-lock is implemented by actuating the RF reference used to clock the DDS driving the AOM. During clock

spectroscopy, a second RF frequency is applied to AOM to switch on the 698 nm probe beam. The PLS corrections are fed forward onto the clock light by clocking the 698 nm and 813 nm DDSs using the same reference. To ensure the corrections are fed forward with the correct amplitude, the ratio of the two DDS frequencies must equal the ratio of the two optical frequencies. This PLS is depicted in figure 4.17.

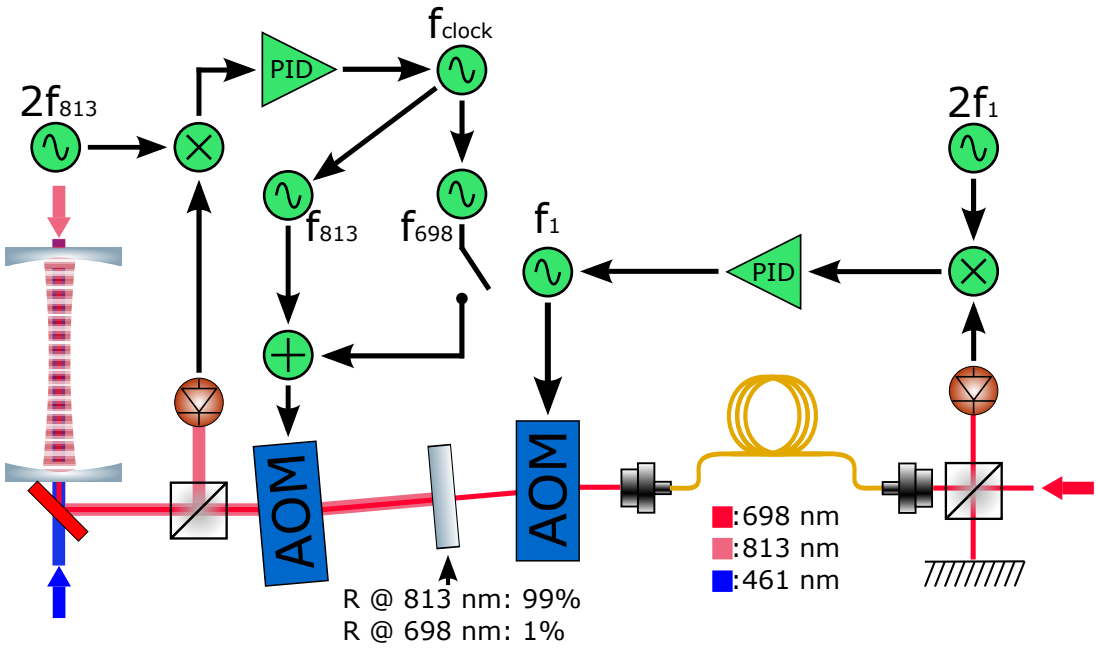


Figure 4.17: The optical path length stabilisation scheme used to deliver the clock light to the cavity. This is done in two steps. First, the path used to send light near the vacuum chamber is stabilised via phase-locking the homodyne beat signal between light picked off before the delivery fibre and retroreflected back through the fibre (shifted in frequency by $2f_1$) to a stable reference via feedback to the AOM. Next, the path between the retroreflector and the cavity is stabilised using same technique, except this time we use lattice light transmitted through the cavity to produce the homodyne beat signal. The beat-note is stabilised via feedback to the synthesiser clocking the DDS driving the AOM. The same synthesiser clocks a separate DDS which drives the same AOM and is used to switch on the clock probe beam during spectroscopy.

4.1.7.2 Non-demolition detection laser system

We elected to pursue a preliminary investigation of NDD using 461 nm transition. The choice was motivated by the broader cavity linewidth – relaxing the required

laser stability – and the larger atom induced cavity frequency shift compared with the 689 nm transistion. The topology of the measurement system used for NDD is as follows: both the lattice laser and the 922 nm master ECDL are locked to the transfer cavity (see section 3.1.2) to guarantee the long term frequency stability. Next, the Sr2 in-vacuum blue cavity is locked to the lattice light with a bandwidth exceeding 500 kHz. By locking both lasers to the transfer cavity, the passive long-term stability of the 461 nm NDD probe light relative to the Sr2 cavity is guaranteed. On short timescales, we compensate for length fluctuations of the Sr2 cavity by feeding forward the lattice lock corrections onto the NDD beam – details of which will be described later. The simplified diagram of scheme is shown in figure 4.18.

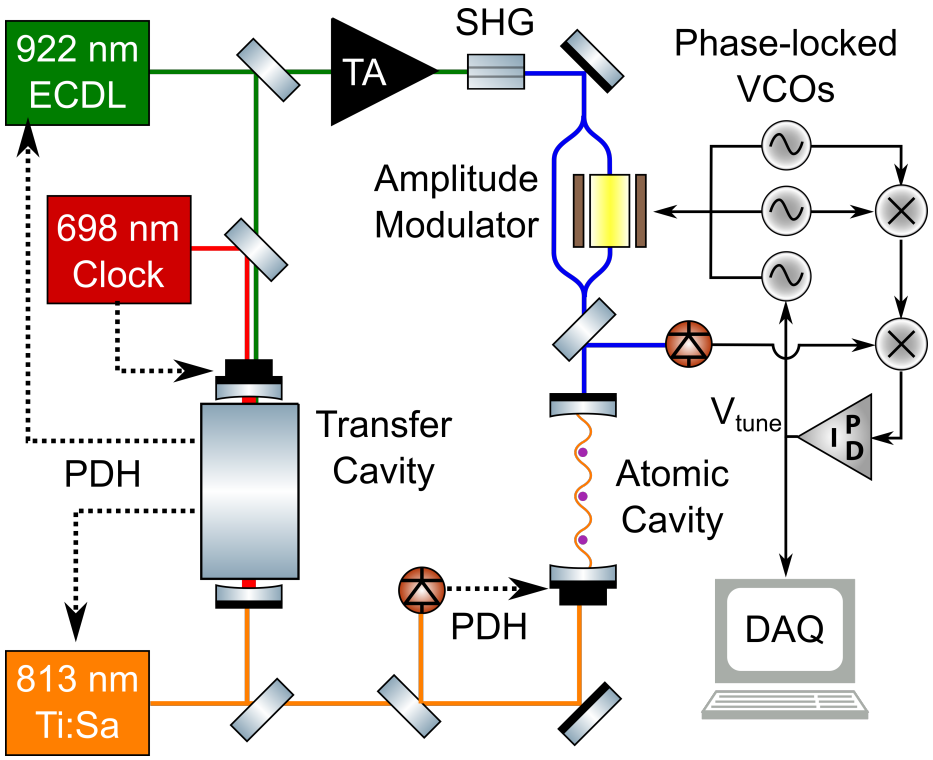


Figure 4.18: The NDD scheme operates as follows. Both the lattice light and 922 nm master ECDL are locked to the transfer cavity. The 922 nm ECDL seeds a TA whose output is frequency doubled using a SHG module. The probe light is then sent into the Mach-Zehnder interformetric modulator which is used to apply the sideband needed for generating the NDD signal. The probe light is sent to the cavity, which is locked to the lattice laser, and the reflected probe is detected using a photodiode. The signal from the photodiode is demodulated using the LO offset lock frequency to produce a NDD error signal. The error signal is used to lock the probe sidebands to the two cavity modes (see text for more detail). The voltage used to the lock the sidebands is proportional to the atom number and recorded using a DAQ. (Figure credit: Richard Hobson)

As outlined in section 4.1.2.2, we need six frequencies – two probes and four local oscillators – to produce the NDD error signal. To generate these frequencies, a single probe beam is sent into a tunable Mach-Zehnder interferometer (MZI) which we use as a high speed phase modulator to apply the necessary sidebands. Both MZI branches contain separate AOMs which can be used to tune continuously the relative phases of two paths. One branch also contains a waveguide EOM[16] allowing for fast

[16]The waveguide EOM is supplied by ADvR and exhibits a throughput loss of approximately 10 db while having 6 GHz modulation bandwidth.

modulation of the relative path lengths. Finally, a third reference beam derived by taking the zero-order out of the AOM in the non-EOM branch is used to stabilise the relative path lengths of both branches such that MZI output port going to the cavity exhibits complete destructive interference. Extinguishing this beam is critical because it will be resonant with the atomic transition and can lead to heating if it leaks into the cavity. To ensure complete destructive interference both the branches must contain equal power and have the correct relative phase. This is achieved as follows. A heterodyne signal between the output interferometer and the reference beam is measured using a balanced photodetector. The signal contains two frequencies: one at 75 MHz corresponding to the AOM drive frequency and a second at 105 MHz which is generated by applying a RF modulation signal to the EOM at 180 MHz. The beat signal at 110 MHz is phase-locked to a RF reference via feedback to the AOM in the MZI branch containing the EOM. This effectively locks the relative path-lengths between this branch and the reference beam. Next, the 75 MHz is demodulated to extract both quadratures of the signal (I and Q). The I and Q components are servoed to zero via actuating on the drive frequency and amplitude, respectively, of the AOM in the non-EOM branch. When all locks are engaged with the correct relative phases, we can suppress the intensity out of the MZI port by over 20 db. In one of the branches, we insert a tunable optical delay path to match the exact lengths of the two MZI branches. If the paths were not matched, the MZI amplitude modulation could only extinguish the light over a narrow frequency band and the broadband TA spontaneous emission would leak through.

Next the three modulation sidebands at $f_{\text{probe}} = f_{\text{FSR}}/4$ [17], $f_{\text{LO+}} = f_{\text{probe}} + f_{\text{PDH}}$, and $f_{\text{LO-}} = f_{\text{probe}} - f_{\text{PDH}}$ are applied to the EOM to generate the sidebands needed for the NDD measurement. These frequencies are produced by taking a master synthesiser operating at a frequency f_{probe} and then offset locking two voltage controlled

[17]Because MZI is locked to the transmission minimum, the response to changes in phase is quadratic requiring the modulation frequency be half the desired sideband frequency.

oscillators (VCOs) by $\pm f_{\mathrm{PDH}}$. The frequency f_{PDH} is also used to mix down the cavity reflection signal to generate the NDD error signal. As to not be limited by the maximum number of atoms we can load into the cavity, we use this signal to lock the probe sidebands to the match the free spectral range (FSR) of the cavity. We then read in this control voltage into the computer using a low noise DAQ[18]. If we did not implement this last lock, but instead read in the error signal, we would be constrained to how many atoms we could trap before the induced frequency shift was more than the cavity linewidth. Figure 4.19 shows a detailed description of the MZI setup.

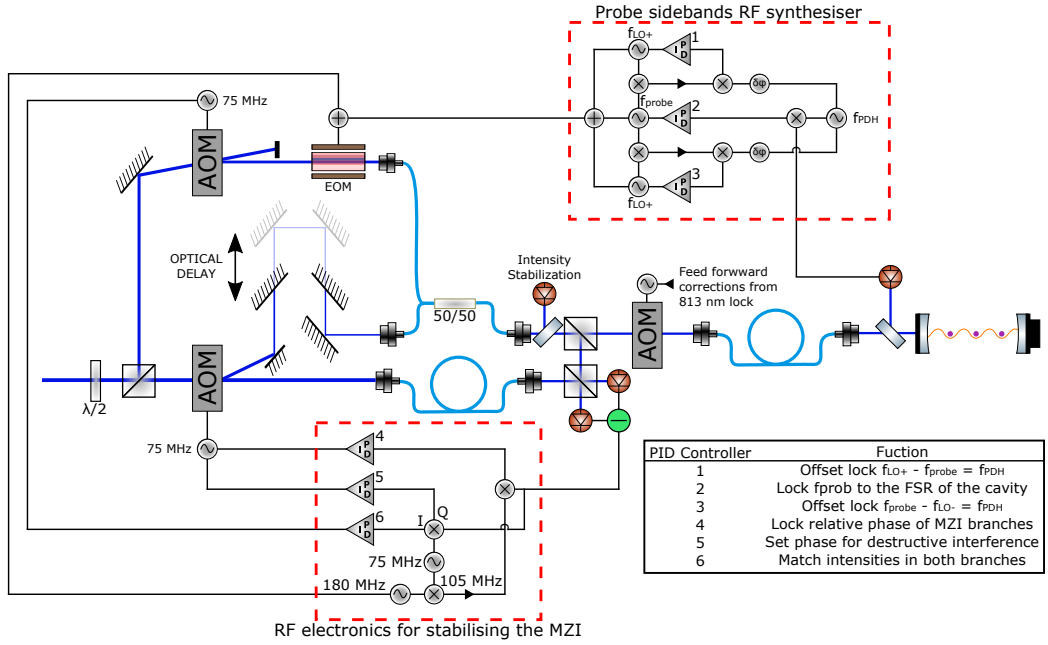


Figure 4.19: To generate the light needed for NDD, we construct a high speed amplitude modulator based on a Mach-Zehnder interferometer. By using AOMs in the two branches, the relative phases of the paths are locked such that the output port sent to the cavity exhibits complete destructive interference (see text for details). A waveguide EOM is then used to phase modulate the light to generate the sidebands as needed for the NDD scheme.

Although the 813 nm and NDD probe lasers are both locked to the transfer cavity, the lattice has the additional fast EOM lock to the Sr2 cavity. As a result, length

[18] We use a PicoScope 4262, it has two 16-bit digital-to-analogue converters (DACs) with a sample-rate 10 MS/s and a RMS noise of 8.5 µV (integrated over full measurement bandwidth).

fluctuations of the Sr2 cavity are not tracked by the probe beam. To solve this issue we feed-forward, with the appropriate scale factor, the corrections applied to the EOM used for lattice lock onto a switching AOM placed before the cavity that is used to turn on the probe light. This ensures that any mechanical vibrations of the cavity length are tracked by the probe beam. Figure 4.20 shows the effect feed-forward has on the relative frequency stability between the probe light and cavity. In principle, the zero-crossing of our error signal is immune to such fluctuations in the relative detuning, but if the probe is not resonant with the cavity it has the second-order effect of reducing the discriminant.

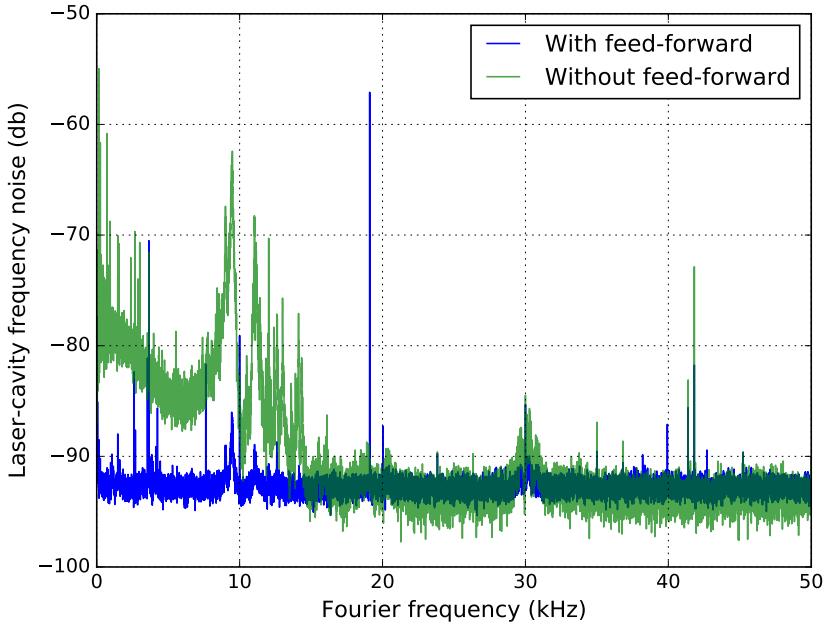


Figure 4.20: To decrease the relative frequency instability of the NDD probe beam with respect to the cavity, the corrections applied to the EOM used for locking of the lattice light to the in-vacuum cavity are fed-forward to the probe beam. Comparing the relative frequency noise with and without feed-forward corrections, this scheme is capable of suppressing the cavity noise, predominantly centreed around the 9 kHz cavity resonance, by almost 20 db. The drawback of such a scheme is it does not suppress any frequency noise on the NDD laser (the ECDL has piezo resonances at 2.6 kHz and 2.5 kHz which are clearly visible on both spectra). Also, the scheme transfers frequency noise from the lattice light onto the NDD beam, e.g. at 19 kHz which corresponds to the Ti:sapphire laser etalon dither used from locking its etalon.

4.2 Validating the performance of Sr2

4.2.1 First stage cooling

The two figure of merits for characterising the first stage cooling are the loading rate and MOT lifetime, the product of which gives the steady-state atom number – assuming density dependent loss effects are negligible. Although the lifetime of a repumped ^{87}Sr blue MOT is on the order of seconds, we typically operate the MOT for only a few hundred milliseconds in an effort to minimise the cooling time. Therefore, increasing the loading rate improves the clock stability by reducing the required loading time. The loading rate depends primarily on the effusive source and Zeeman slower which generate the slow flux of capturable atoms. We operate the slower with a 30 mW beam focused to a $1/e^2$ diameter of 3 mm, while for the MOT we use a 30 mW beam collimated to $1/e^2$ diameter of 44 mm. With an oven temperature of 470 °C, we measure a loading rate of $3(1) \times 10^6$ atoms·s^{-1} for ^{87}Sr as shown in 4.21. Using the repump scheme described in 4.1.5.3, we observe a MOT trapping lifetime of 2.3 seconds with 1 mW of power in both the 497 nm and 679 nm beams. We also measure the loading rate as a function of the MOT beam power. We observe that the loading rate initially scales linearly with power, but then begins to saturate above 25 mW. For all the atom number measurements, we calibrate the fluorescence signal by comparing it to the cavity NDD measurement, see section 4.2.4.2 for details.

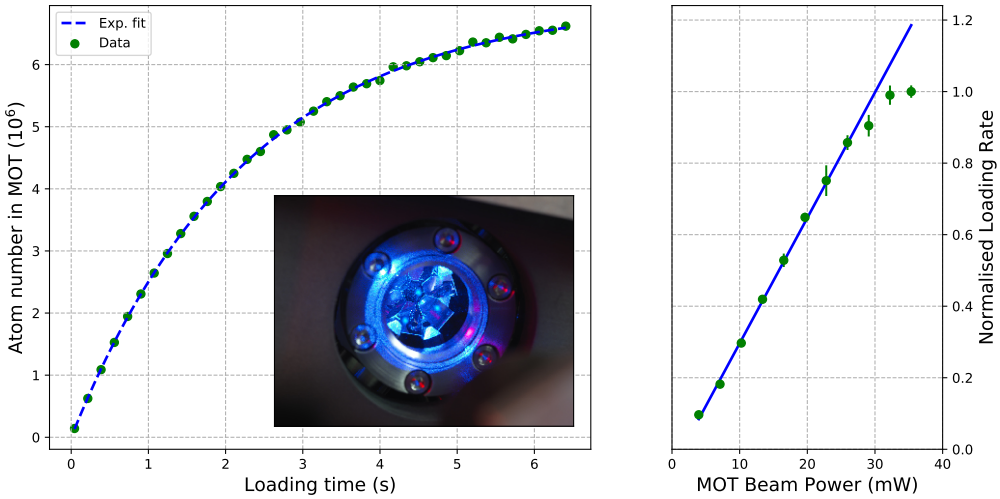


Figure 4.21: **Left:** Loading curve for ^{87}Sr MOT. From an exponential fit to the data, we measure a loading rate of $3(1) \times 10^6$ atoms·s^{-1} and a lifetime of 2.3 s. **Inset:** Picture of a ^{88}Sr blue MOT. **Right:** Normalized loading rate of the MOT as a function power in the beam.

As an alternative to using a Zeeman slower for generating the slow flux of cold atoms, one could elect to use a 2D MOT source as was reported in [107] which demonstrated a loading rate of 22×10^6 atoms/s for an oven temperature of 500 °C. However, one caveat when trying to compare the loading rate from various cold-atom sources is it depends on design parameters (e.g. the distance between the source and MOT or the oven aperture dimension) which are often not specified. Furthermore, the vapour pressure scales exponentially with temperature and even if a sensor is installed, it is hard to be certain the value it reads matches the temperature of the strontium. All of these factors makes it difficult to compare the loading rates presented for various designs.

4.2.2 Magnetic trap lifetime

With the repump light turned off during the first stage of cooling, atoms will be optically pumped into $^3P_0$ state. Atoms which happen to end up in either the $m_F = 1$ or 2 states will be magnetically trapped in the quadrupole field formed by the MOT coils. We can use these atoms to estimate the residual background gas pressure inside the chamber by measuring the collisional loss rate from the magnetic trap. The relationship between loss rate Γ and pressure has been measured empirically to be $\Gamma = 3.1 \times 10^8\,\mathrm{s}^{-1}\mathrm{mbar}^{-1}$ [98]. The loss rate as a function of hold time is shown in figure 4.22.

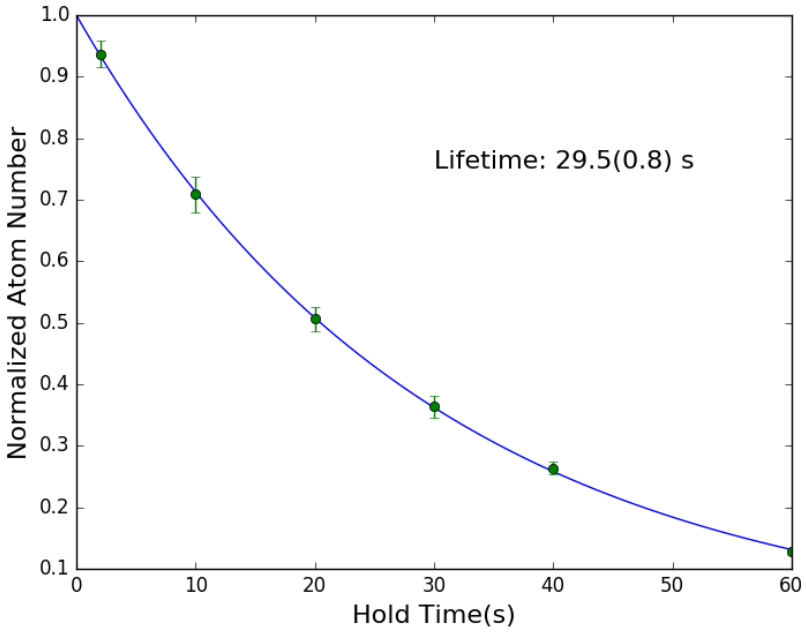


Figure 4.22: Magnetic lifetime as measured by holding the atoms in the magnetic trap for a variable delay. The remaining atoms are then repumped to ground state, recaptured in a blue MOT, and imaged.

The lattice lifetime was measured to be 30 s indicating a residual background pressure of 1×10^{-10} mbar. However, this loss rate coefficient was measured for much higher pressures and neglects other loss mechanisms which become non-negligible as the pressure pushes below 1×10^{-10} mbar. For example, black-body radiation can re-

pump atoms by driving transitions to the $5s4d\ ^3\mathrm{D}_3$ state which quench the lifetime of the $^3\mathrm{P}_0$ and induce losses. This effect was shown by JILA to limit the magnetic trap lifetime in their system to 41 s given they observed a dipole trap lifetime was 90 s [21]. This implies a BBR limited lifetime of 71 s, which is in reasonable agreement with the theoretical value of 145(50) [16]. Therefore, the estimated pressure of 1×10^{-10} mbar is most likely an upper bound and corresponds to an estimated background gas collisions induced fractional shift of the clock frequency of 5.3×10^{-19}.

4.2.3 Second-stage cooling and state preparation

For Sr2 we have replaced the conventional second-stage cooling at 689 nm with a MOT operating on a transition between the metastable $5s5p\ ^3\mathrm{P}_2$ triplet state and the $5s4d\ ^3\mathrm{D}_3$ state separated by 2.9 µm. Atoms are transferred to the triplet MOT immediately after the blue MOT, by switching off the 497 nm repump modulation sidebands to optically pump atoms into $^3\mathrm{P}_2$, $F = 13/2$ state. We keep the 497 nm light on to repump the $F = 11/2$ state which is not trapped in the MOT. This selective repumping enhances our atom number by approximately a factor for two. At the same time as atoms are transferred to the metastable state, the magnetic field gradient is stepped down from $4\,\mathrm{mT\,cm^{-1}}$ to $1.3\,\mathrm{mT\,cm^{-1}}$ simultaneously with the 2.9 µm light switching on. To compensate the Doppler broadening resulting from the thermal distribution of atom exiting the blue MOT, the 2.9 µm light is red detuned by 5 MHz and frequency modulated at 400 kHz with a depth of 1.5 MHz. Following this procedure, we can transfer approximately 50% of the atoms from the blue MOT into the metastable MOT. In the future, an additional laser at 688 nm or 487 nm resonant with the $5s5p\ ^3\mathrm{P}_1 \rightarrow 5s6s\ ^3\mathrm{S}_1$ transition could improve the atom number in the metastable MOT by transferring atoms from $^3\mathrm{P}_1$ to $^3\mathrm{P}_2$ [9].

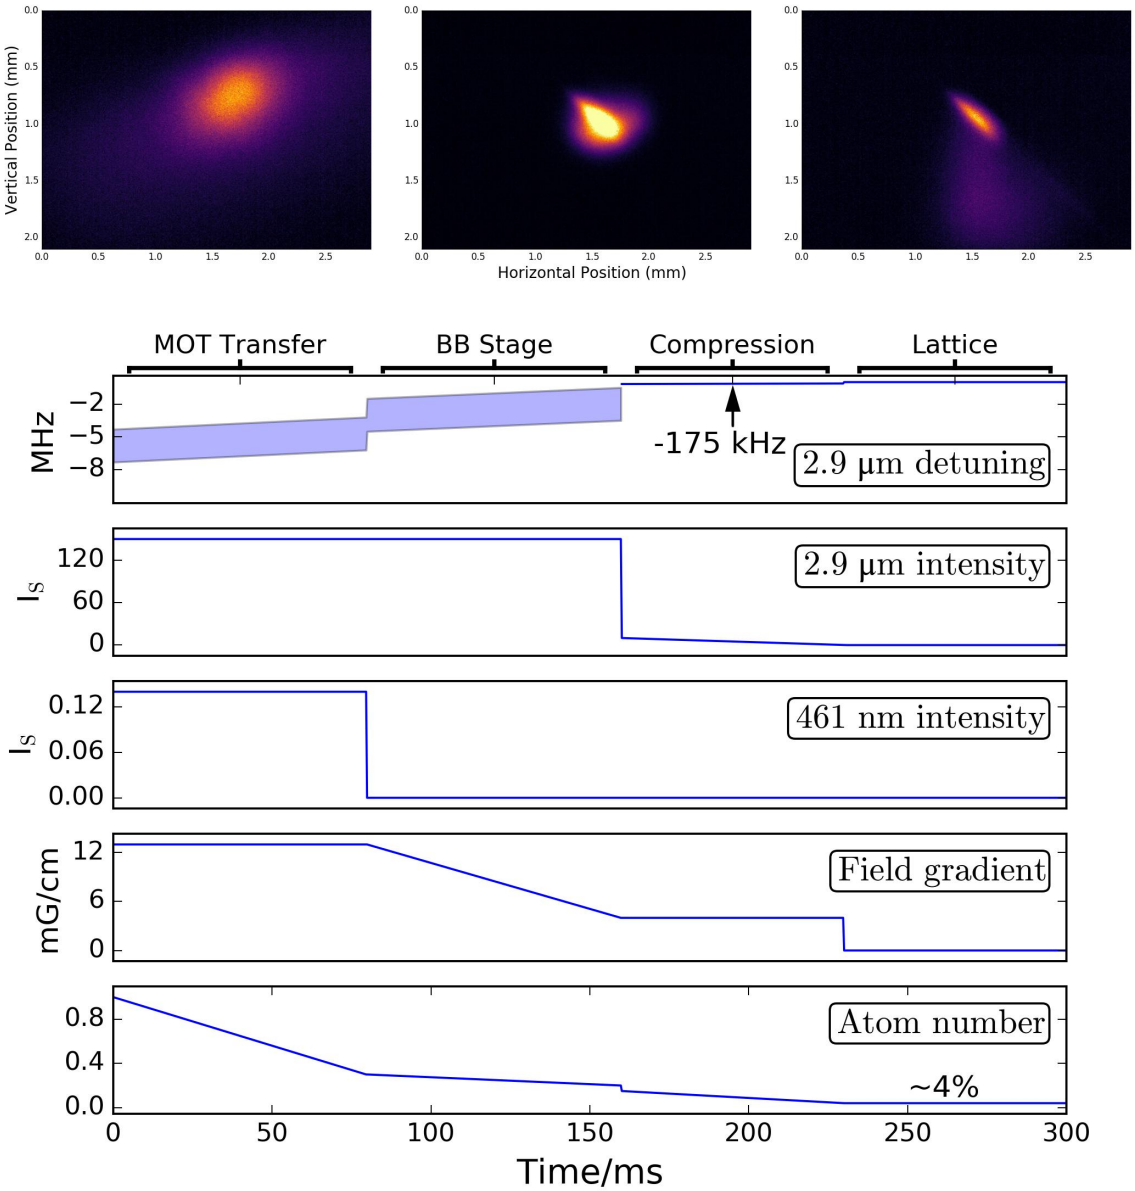


Figure 4.23: The triplet MOT sequence can be divided into three parts: the transfer stage, the broadband stage, and the narrowband lattice loading stage. Above we show the setpoints for various MOT parameters (e.g. the magnetic field gradient, frequency detuning and laser intensity), along with the relative atom number during these stages (see text for more details). Also included above are three fluorescence images showing atoms in the broadband stage, the narrow band stage, and the lattice with the MOT switched off.

Once the transfer is complete and the 461 nm light is switched off, a broadband

stage during which the modulation is left and field gradient is stepped to $0.4\,\mathrm{mT\,cm^{-1}}$. To cool and compress the atoms further, the laser detuning is reduced to $-3\,\mathrm{MHz}$. After 70 ms, the modulation is switched off and laser detuning is jumped to $-175\,\mathrm{kHz}$ and then slowly ramped towards resonance. To compensate for the increased spectral density resulting from switching off the modulation, the laser intensity is reduced by a factor of 100. Throughout this final cooling stage, atoms are loaded into the magic-wavelength lattice. In total, the metastable triplet MOT stage – including the transfer step – lasts 300 ms and results in 20,000 atoms being loaded into the lattice. One interesting possibility to improve the transfer of atoms from the metastable MOT into the lattice is the addition of a 497 nm 'drain' laser overlapped with the lattice to selectively repump atoms into the ground state as they enter the capture region. This has been successfully implemented as a means to load a dipole trap from a metastable calcium MOT [164] and could be a viable route to replenish strontium atoms held in a magic-wavelength lattice continuously. A similar scheme has been proposed as route to realise a continuous super-radiant laser [105, 106].

To estimate the atomic temperature, we load a triplet MOT with the lattice switched off. We then release the atoms causing the cloud to expand ballistically. After a fixed expansion time, we then repump the atoms and image the cloud. The rate of spatial expansion is directly related to mean thermal velocity and can therefore be used to estimate the thermodynamic temperature of the sample [19]. Assuming an initial Gaussian profile for the atomic density with a standard deviation σ_0, after a fixed expansion time t the new cloud size is given by:

$$\sigma(t) = \sqrt{\sigma_0^2 + \frac{k_b T}{m} t^2} \tag{4.16}$$

where T, m and k_b are the sample temperature, atomic mass and Boltzmann constant, respectively. The expansion of the triplet MOT is shown in figure 4.24 for time

of flight delays between 0 to 20 ms. Fitting the data using equation 4.16 gives an estimated temperature of 7 µK. This is sufficient for loading the lattice, but future investigation is needed to understand what factors are limiting the efficacy of cooling and preventing us from reaching the estimated Doppler limit of 2.2 µK.

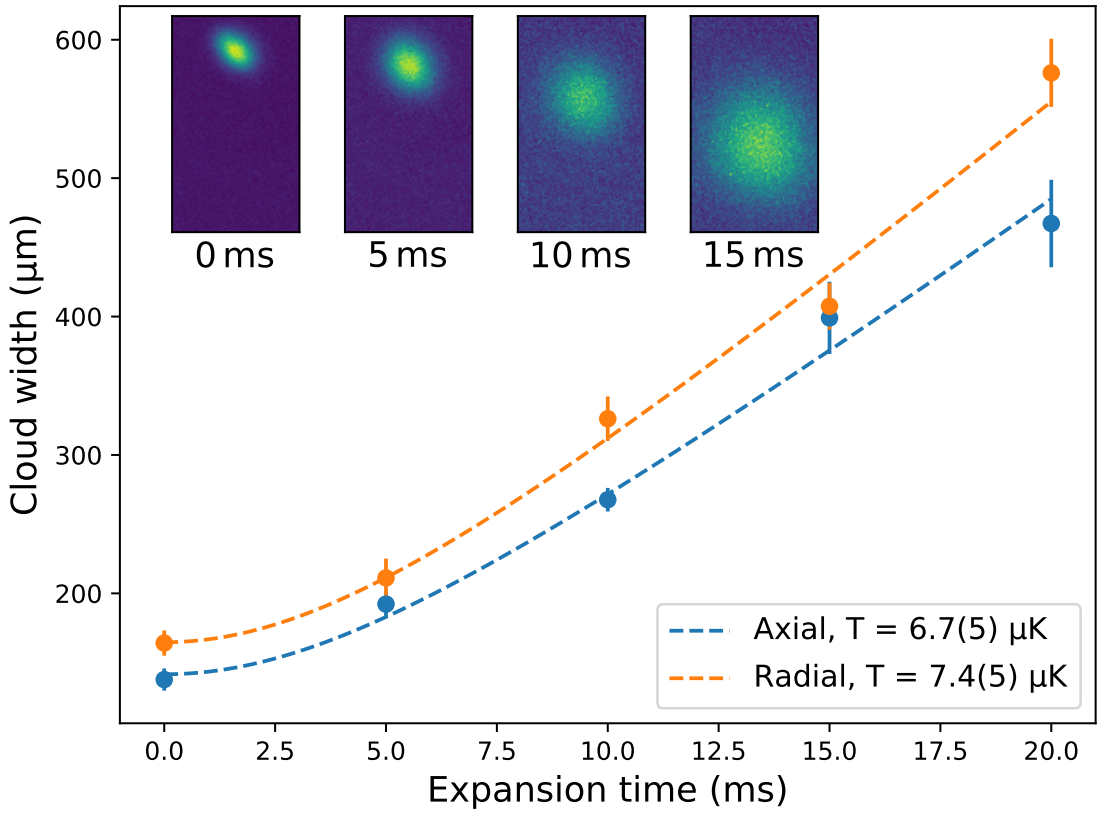


Figure 4.24: Time-of-flight measurement used to track the expansion rate of the atoms released from the triplet MOT. The expansion rate gives an estimated temperature of 7 µK.

We also measure the trapping lifetime of atoms held in the low intensity, narrow band MOT. We observe that the lifetime is strongly dependent on the trapping intensity indicating that single-body light assisted losses are the dominant decay mechanism, see figure 4.25. All the hold intensities are referenced to the setpoint used for normal MOT operation, I_0. For intensities below I_0, we see a decrease in the MOT lifetime. We observe similar behavior for the red MOT used for the second-

stage cooling in Sr1. At low intensities, the red MOT cannot compensate the force of gravity and the atoms literally fall out of the trap. We solved this issue in the first system with the addition of an upwards propagating beam that compensates the gravitational force. A similar effect could limit realizing a low intensity triplet MOT in Sr2. Another possible mechanism for the reduction in MOT lifetime with intensity is photoionization from the $^3D_3$ state. From the ground state the ionization energy is 218 nm [122]. Using transition wavelength for $^1S_0 \rightarrow {}^3P_2$ transition (671 nm [108]) and $^3P_2 \rightarrow {}^3D_3$ transition (2940 nm), the minimum wavelength needed for photoionization is calculated to be 361 nm – making such an explanation implausible. So far the reduced lifetime has not stopped us from loading the lattice, but it does warrant future study as the mechanism might also help to explain the hotter MOT temperature.

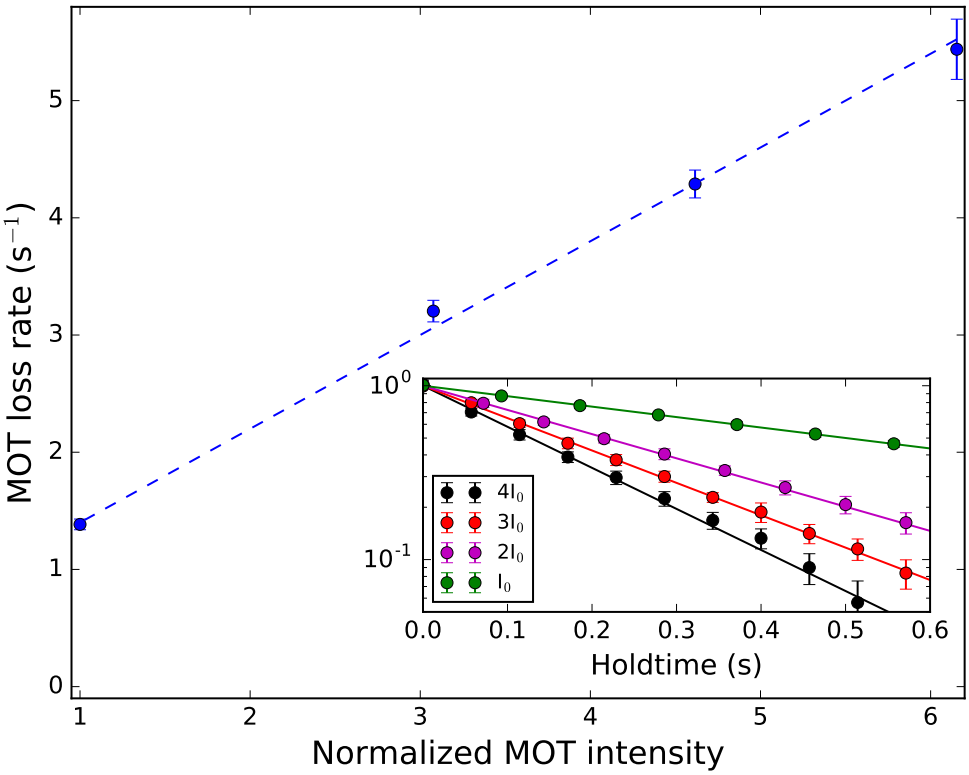


Figure 4.25: Lifetime of triplet MOT for various intensities. The decay rate scales linearly with intensity, indicating the lifetime is limited by a light assisted loss mechanism. **Inset:** Lifetime measurements of the triplet MOT for various intensities normalised to I_0, the intensity used for typical operation.

Finally with the goal to operate the system as a clock, the atoms need to be transferred to the $5s^2\ {}^1\mathrm{S}_0$ ground state. To eliminate the 1^{st}-order Zeeman shift of the clock frequency we average interleaved measurements using different Zeeman sub-levels with opposite linear magnetic field sensitivities. This requires that the sample is spin polarised, alternating between the two $m_F = \pm 9/2$ stretched states, before probing the clock transition. This is achieved by driving the $5s5p\ {}^3\mathrm{P}_2$ to $5s4d\ {}^3\mathrm{D}_3$ transition using the MOT beams in a bias magnetic field for 15 ms. By red detuning the beams by 1.4 MHz, atoms are preferentially pumped into the $m_F = +13/2$ stretched state. Next, we repump the atoms into the ground state using the

497 nm. Selection rules ensure the sample remains spin polarised during this process with approximately 80% of atoms ending up in the $m_F = +9/2$ $5s^2\ {}^1\mathrm{S}_0$ ground state. Unfortunately, this process is not capable of spin polarising the sample in the $m_F = -9/2$ Zeeman level as this would require a blue detuned beam which would heat the sample. To spin polarise the sample in the opposite stretched state, we first polarise the sample in the $m_F = +9/2$ and then rotate the magnetic field adiabatically by 180° in 40 ms. Next, we non-adiabatically flip the field back to the original direction– leaving the atoms in the $m_F = -9/2$ state.

4.2.4 Cavity enhanced lattice

From the triplet MOT, we directly load the atoms into the magic wavelength lattice. The enhancement provided by the cavity allows us to exceed the maximum trap depth achievable in Sr1 with only a few milliwatts of incident optical power. As outlined in section 4.1.2.4, the drawback of using the enhancement cavity is the trapping potential depends strongly on the detuning between the cavity resonances and the laser. Any instability of either the laser frequency or the cavity length is converted to amplitude modulation of the intra-cavity intensity which gives rise to parametric heating. We measure this heating rate by loading atoms into the lattice and holding them for a fixed duration before performing sideband spectroscopy to measure the atomic temperature. See section 2.4 for more details regarding this technique. With the slow feedback to the piezo, fast feedback using the EOM, and intensity stabilisation all engaged (see figure 4.9 for an explanation of these locks), we measure a negligible heating rate, see figure 4.26. To prove further that heating is negligible, we measure how long atoms remain trapped in the lattice. We observe that the decay in atom number is exponential with a lifetime comparable to that observed for the magnetic trap, indicating the dominant loss mechanism is from collisions with background gas particles. However if we disengage the fast feedback to the EOM, we observe that this

lock was critical in suppressing heating as the lattice lifetime drops to one second.

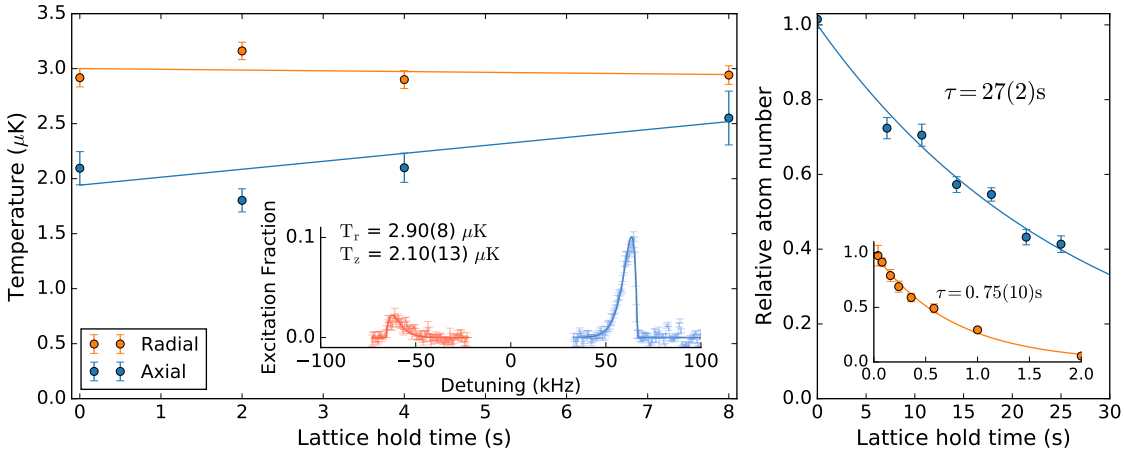


Figure 4.26: **Left:** The heating rate was measured by holding the atoms in the lattice for a fixed period of time followed by measuring the sideband profiles to extract the temperature. **Left, inset:** Example of a sideband scan. **Right:** Exponential decay of the lattice-trapped atom number. Based on the similar lifetime observed with magnetically trapped atoms in the $^3P_2$ state, see section 4.22, we can infer than lattice induced heating does not lead to a significant atom loss. **Right, inset:** Same lifetime measurement, but with only feedback to slow piezo engaged for locking the cavity to the lattice laser. Without the high-bandwidth lock of the laser to the atomic cavity, frequency noise induced modulation of the trap depth quickly heats the atoms.

4.2.4.1 Cavity-enhanced 2D lattice

We also verified that the relative position of the cavity modes went unchanged during the assembly and baking of the vacuum chamber to ensure we could operate a 2D lattice. To find the intersection point of the two cavity modes, atoms were translated along the blue cavity by changing the magnetic field position using the compensation coils. Once the crossing point was found, the magnetic filed was then stepped to move the atomic cloud in the plane perpendicular to the two cavity modes. By separately measuring the loading efficiency into each of the cavities as a function of the MOT position, we can locate the centre of each trap. Given that the fields generated by the MOT and compensation coils are known, it is easy to convert between magnetic field

and spatial position of the field zero. The result of this measurement is shown in figure 4.27 and indicated the offset between the two cavities is approximately 20 µm. We also attempted to change the cavity overlap using the shear-stack piezo. However we were unable to observe any measurable change in the relative cavity position when the piezo was activated. It is unclear why it was unsuccessful, but possible explanations are: the piezo failed, the glue interface was not strong enough to support the shear force, or the piezo did not induce a significant tilt of the mirrors. Fortunately the modes were sufficiently well overlapped that the final adjustment was not needed, but if we choose to include the piezo-based fine adjustment mechanism for future systems it will need redesigning.

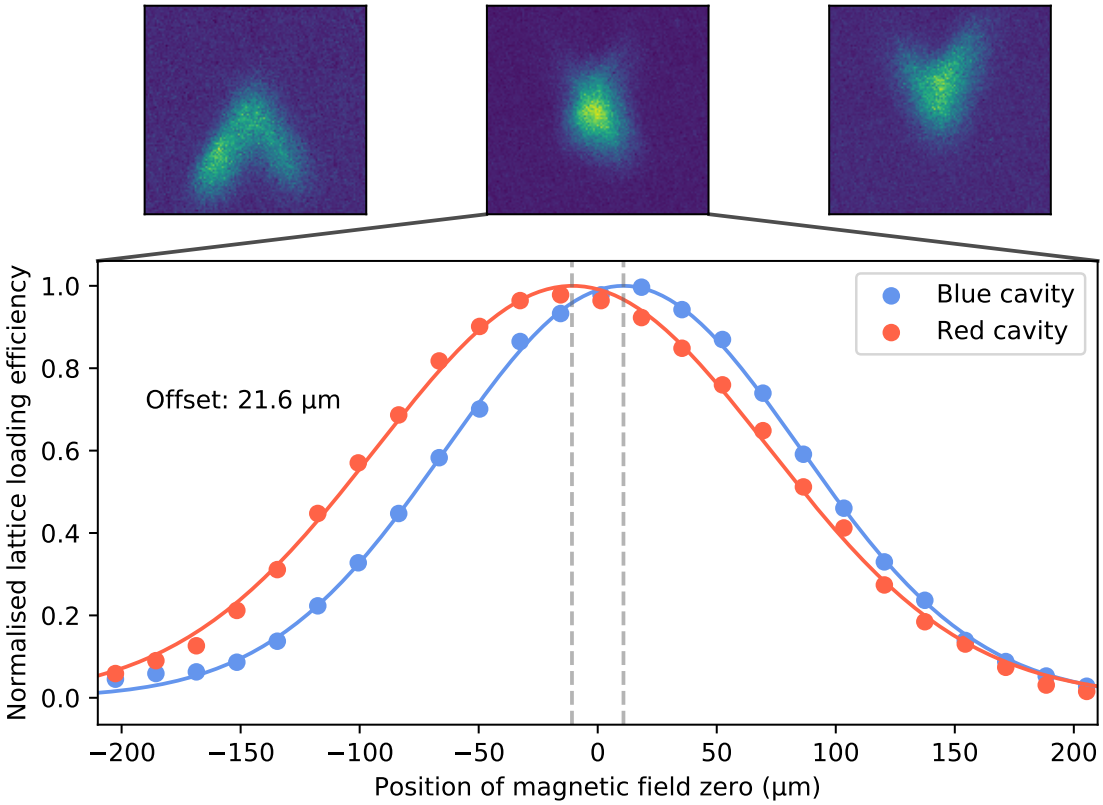


Figure 4.27: To find the crossing point of the two cavities, the position of the MOT is scanned along the blue cavity by moving the magnetic field as shown above in the three fluorescence images. Once the intersect is located, a similar process is repeated, with the exception that the MOT is now moved in plane perpendicular to both lattices, to measure the position of each trap centre. From Gaussian fits to the positionally dependent loading efficiency of both lattices, the relative offset between the two modes is estimated to be 20 µm.

4.2.4.2 Cavity read out of the atomic state population

To test the NDD scheme, we loaded atoms into the lattice and measured the splitting of the adjacent cavity modes as outlined in section 4.1.2.2. To measure the induced heating rate, we keep the probe light on until all the atoms are lost from the lattice as shown in figure 4.28. For this measurement, we did not implement the lock of the probe beams to the cavity mode splitting and instead just recorded the NDD error signal. This explains the initial shape of the decay curve which is slightly flat-

topped. For this measurement, we started with a sufficient number of atoms in the cavity to shift the modes by approximately half their FWHM. As atoms are lost, the induced cavity shift decreases and the error signal is restored to the linear regime. By fitting the tail end of the curve, we can extract the measurement induced loss rate which limits the trapping lifetime to 3.3 ms for a 200 µW probe beam. This is quite promising for realising non-demolition detection, as a probe pulse of less than 1 ms should allow for sufficient signal-to-noise without inducing appreciable atom loss. By looking at the residual of the fit, we can also extract the RMS detection noise which can be converted to units of atom number using the calibrated discriminant. This results in a sample-to-sample RMS noise of 50 atoms.

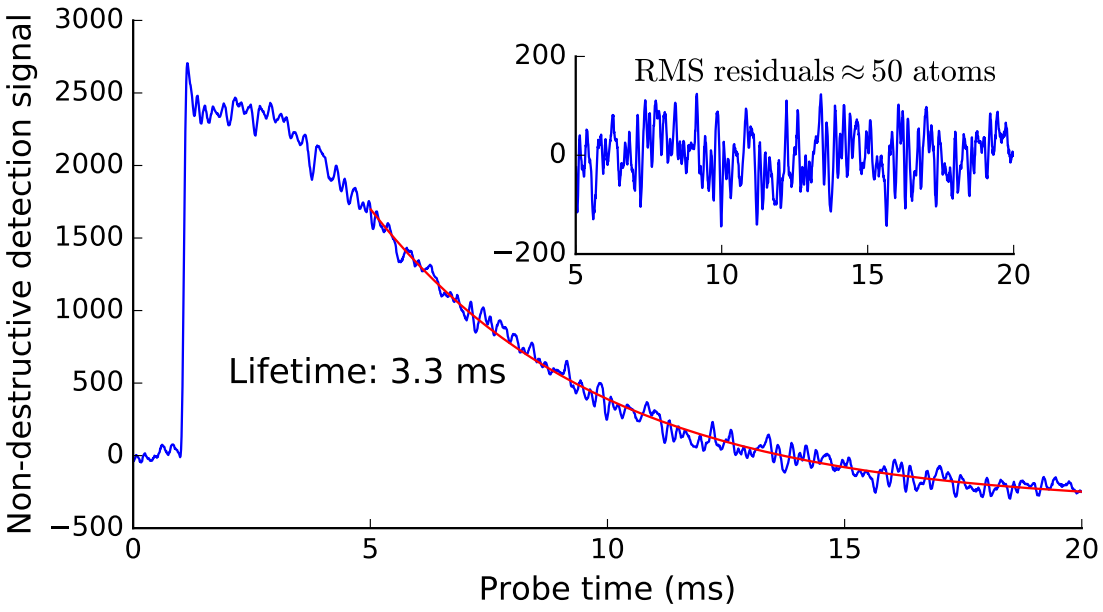


Figure 4.28: NDD of the atom number held in the cavity enhance lattice. Probe induced heating limits the lifetime of the trapped atoms to a few milliseconds. The RMS noise in the signal, inferred by fitting data to an exponential decay curve and measuring the residuals, is approximately 50 atoms.

4.2.5 Clock operation

With spin-polarised atoms loaded in the lattice, we can begin to operate the clock. As explained in section 3.2, with two operational systems there are tricks we can now play to improve stability. In general, these techniques involve operating both clocks simultaneously with coordinated spectroscopic pulses in order to suppress instabilities caused by the Dick effect. If the probe pulses are synchronised, the resulting excitation fraction noise caused by laser frequency fluctuations will be identical and therefore cancel when we calculate the frequency offset between the clocks. For certain systematic shifts which we evaluate by measuring the frequency offset between interleaved clock measurements under different experiment conditions (*e.g.* measuring the lattice shift between two different trap depths), the cancellation of Dick instability will speed up the characterisation process. Figure 4.29 shows the frequency instability of a comparison between Sr1 and Sr2 when the two systems are synchronously probed along with the resulting excitation fraction which exhibits clear correlations which result from common mode laser frequency noise.

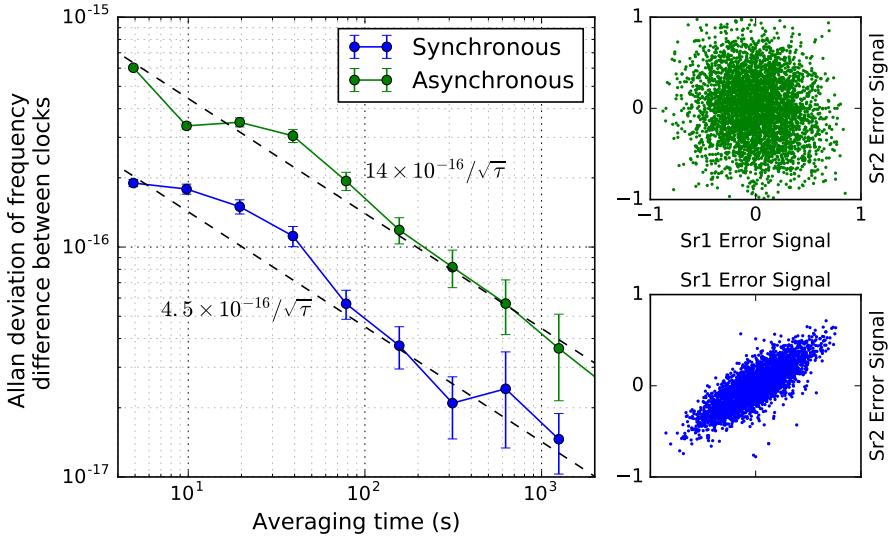


Figure 4.29: Allan deviation of the frequency difference between Sr1 and Sr2. By synchronising the probe pulses for both systems, instability arising from the Dick effect is common to both clocks and is therefore suppressed. For asynchronous interrogation, this common-mode rejection of the Dick noise does not occur which is evident by increased instability of the comparison. To illustrate further how synchronised comparison improves stability, the error signals for both clocks derived from the excitation fraction are plotted against each other. For synchronous interrogation, the two signals are strongly correlated.

4.2.6 DC-Stark shift cancellation

We make use of the improved relative frequency instability enabled by synchronous operation of both clocks to evaluate the DC-Stark shift in Sr2. To measure any residual electric fields that could arise from the in-vacuum piezos or parasitic charges that have built-up on dielectric surfaces, we apply voltages to the in-vacuum electrodes and follow the procedure outlined in section 3.3.5. To ensure we don't inadvertently charge up dielectric surfaces near the atoms when we apply voltages to the electrodes, we also apply a compensation voltage to the electrodes during the cooling sequence for the same duration, but with opposite polarity. An example of the DC-Stark measurement is shown in figure 4.30, which reveals a Stark shift of $-9(3) \times 10^{-19}$ arising from a residual electric field along the axis of the Zeeman slower. Repeating

this measurement along the vertical direction and MOT axis gave electric field values of $-2(1) \times 10^{-19}$ and $0(7) \times 10^{-19}$, respectively, giving a total DC stark shift of $12(8) \times 10^{-19}$. The uncertainty is purely statistical and could be improved by a longer averaging time.

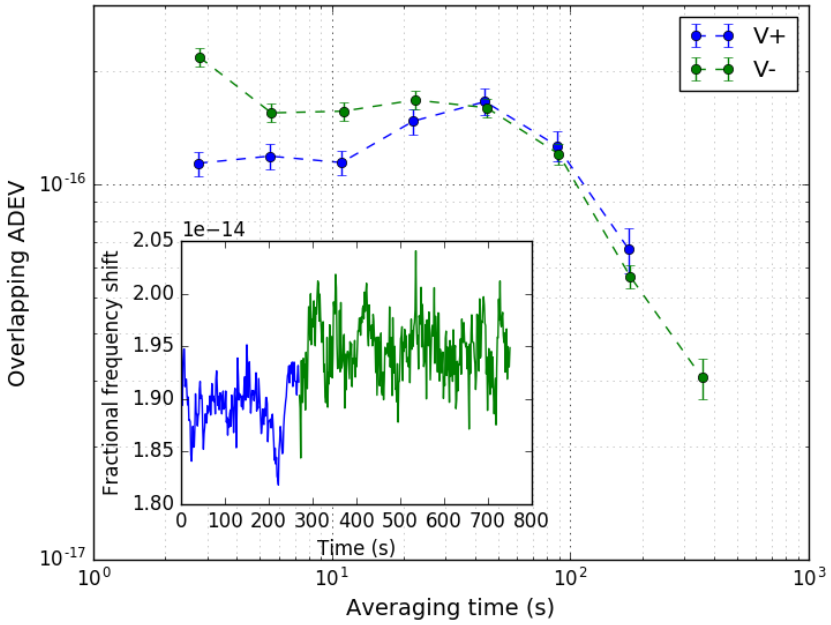


Figure 4.30: Allan deviation of the Stark shift induced by the electrodes for two different polarities. Both Stark shifts are measured with respect to Sr1 which is operated simultaneously with synchronised probe pulses such that we remove the frequency instability arising from the Dick effect. **Inset:** Time traces showing the measured Stark shifts. Any difference in the induced frequency shift between the two polarities must arise from a residual electric field along the direction of the applied field.

From the measurements, we can also extract a value for the electric field induced by the electrodes and compare it to values predicted by the COMSOL simulations. As shown in figure 4.5, for the two radial directions the simulation is in good agreement. However, for the vertical direction the measured field is much weaker than predicted. A possible explanation for the discrepancy is the COMSOL model only included the cavity baseplate whereas in reality the entire chamber is grounded which might lead to additional screening effects.

Table 4.5: Electric fields generated by the in-vacuum electrodes.

Direction	Simulation (V/m)	Measured (V/m)
Vertical	720	843
Along Zeeman Slower Axis	1245	1672
Along MOT Axis	2200	624

4.3 Conclusion

We now have two working strontium lattice clocks at our disposal. In regards to Sr2, we chose to pursue several innovative techniques in the hopes of improving its performance compared to Sr1. The two major design changes were the use of in-vacuum enhancement cavities and a pyramid MOT. The inclusion of an enhancement cavity will help us better characterise the lattice related Stark shifts that were a major source of systematic uncertainty for Sr1. Furthermore, the cavities also serve the secondary function of measuring the number of trapped atoms in the lattice non-destructively with reduced detection noise. Compared with the free space lattice used in Sr1, the enhancement cavity is technically more challenging to implement. To avoid excessive heating due to instability of the trapping potential, it is imperative that the laser-to-cavity detuning is stable– requiring the implementation of a multi-stage control loop acting on the both the laser frequency and cavity length. Furthermore, putting the enhancement cavity in the vacuum can compromise the UHV environment – which is also true for the pyramid MOT. However, we demonstrated that with careful choice of materials and sufficient pumping, these issues can be mitigated. In fact, we showed that Sr2 has a better vacuum than Sr1 by more than a factor of three. Another risk of placing dielectric surfaces close to the atoms is that they can cause parasitic electric fields that give rise to Stark shifts. However, we mitigated such effects by the installation of in-vacuum electrodes which allow us to cancel residual Stark shifts at the low 10^{-18} level.

It was clear why we took on these challenges and added experimental complexity for the build-up cavity, but for the pyramid MOT such risks could appear unjustified. The rationale was twofold. First, the pyramid MOT was essential for realising the metastable MOT as it would have been very challenging to implement it using a conventional six-beam geometry. The hope was that this new MOT would reduce the cooling time compared to the conventional second stage cooling without compromising temperature. Unfortunately, we have been unable to demonstrate such improvements yet, but this work is still ongoing. The second reason for developing the pyramid MOT was that it could be used as a platform for future clocks. Its compact size and low optical power requirements make it attractive for a transportable system. The compact size and reduced required optical access also makes the pyramid MOT easier to integrate into a cryostat– which appears to be an inevitability if one wishes to target 10^{-19}-level inaccuracy with strontium. Below I have summarised the various design changes made for Sr2 in a table that highlights their pros and cons.

Table 4.6: Tradeoff comparison of the technical changes implemented in Sr2.

Feature	Pro	Con
Pyramid MOT	• Needs less power ($< 30\,\mathrm{mW}$) • Requires less optics • Works a range of wavelengths (e.g. $2.9\,\mu\mathrm{m}$) • Compact form factor	• Requires UHV compatible materials and epoxies • Could give rise to perturbation (e.g. Stark shifts)
Metastable MOT	• Less stringent requirement for laser frequency stability • Easy to spin polarise using optical molasses • Possibility for sub-Doppler cooling • Could enable laser cooling without perturbing the clock	• Could only reach $7\,\mu\mathrm{K}$ • Could not reduce required cooling time • Working with mid-infrared light is challenging
In-vacuum cavity	• Large waist improves trapping efficiency and reduces density shift • Deeper trap helps to measure lattice shifts • No need for high-power laser (reduces cost and size)	• Can increase parametric heating • Requires control loop to lock cavity to laser • Could give rise to perturbations (e.g. Stark shifts)
Cavity NDD	• Reduced detection noise • Enables NDD • Could enable the generation of spin-squeezed states	• Needs better freq. stability than fluorescence detection • More complicated than standard imaging

Chapter 5

Future Plans

With two operational clocks now at our disposal, the next logical step is to see if they agree. Although this task is outside the scope of this thesis, work is currently underway in our lab towards this goal. This includes conducting a full systematic evaluation of Sr2 while addressing sources of noise limiting the stability of the clocks– relevant for both for synchronous and asynchronous comparison. In regards to the accuracy of Sr2, the first step will be to make use of the enhancement cavity to reach trap depths previously unachievable with Sr1 to thoroughly investigate lattice related perturbations–including higher order effects. This will not only be of use to our group, but the entire Sr lattice clock community. The time needed to perform this characterization can be reduced if we use Sr1 as a stable reference to which we synchronously compare Sr2. Previously the stability of such a synchronous comparison was limited by a combination of QPN, detection noise, and uncompensated optical paths between the local oscillator and the clocks. To address this issue we recently installed a lattice buildup cavity in Sr1 along with a PLS scheme similar to that used for Sr2. The cavity should improve the stability of the system by allowing for more atoms to be loaded into the lattice and eliminating the remaining uncompensated optical path. With these upgrades, I am optimistic we can demonstrate agreement

between the systems at the 10^{-17} level in less than 1000 s of averaging. The stability of such a comparison could further be improved by extending the probe length beyond the coherence time of the laser and looking at the correlation between the two clocks as was demonstrated for Al^+ [28].

The other obvious route we will pursue is to implement a zero-deadtime clock. If such an ultra-stable clock is realized, it is interesting to think about its possible uses; in particular if we could transfer its stability to other frequency references at NPL. If both the Yb^+ ion clock and our zero-deadtime system were to run simultaneously using the same master oscillator, we could increase probe time for the ion system by feeding forward our corrections during their clock pulse, similar to what has been proposed in [55]. This could improve the stability of ratio measurements between the two clocks helping to constrain time variations of fundamental constants. Extending our perspective beyond NPL, a major focus in the future will be clock comparisons via the cross continental optical fibre network which has recently been extended to include NPL. Comparisons have already taken place this year requiring considerable coordinated efforts across the various institutes. It is therefore imperative that our clocks are stable to maximize the investment of time and effort going forward.

Achieving zero-deadtime operation will no doubt prove challenging as it requires progress on two fronts: increasing the achievable probe time by improving stability of the local oscillator and decreasing deadtime in the experimental sequence arising from the cooling stages. In regards to the LO, work is underway to compensate for mechanical vibrations which are thought to be currently limiting the system using a scheme similar to [146]. To reduce deadtime in the sequence we plan to implement NDD to recycle atoms for several clock pulses. As shown in section 4.2.4.2, heating can be kept to a minimum if the probe time is kept to less than 1 ms which should provide sufficient signal. Looking to the future, if we wish to implement spin squeezing protocols we will need to reduce technical sources of noise plaguing the NDD scheme.

One likely culprit is broadband frequency noise introduced by the tapered amplifier used to seed the waveguide doubler. We observed that even with the probe sideband switched off, broadband noise appears on the signal whenever the local oscillators are near a cavity mode indicating spurious laser light is coupling into the cavity. To solve this issue we are currently constructing a filter cavity which will be placed immediately before the Mach-Zehnder modulator to remove ASE.

It is still an open question just how far we can push the stability and accuracy of optical lattice clocks. With new techniques such as those presented in this thesis and being developed elsewhere, I am confident there is still much room for progress. We are fortunate to be in Europe, as our location enables us to verify directly the performance of state of the art clocks using the optical fibre time transfer network. The ability to compare these systems at this level of accuracy and stability will allow us to test our fundamental theories of nature at an unprecedented precision.

Appendix A

My contribution to the project

I joined the lattice team, which at the time consisted of Ian Hill and Richard Hobson, in the fall of 2014. Prior to my arrival, they had a completed clock laser system, atoms trapped in a 'red' MOT, and were eagerly awaiting the arrival of our new Ti:sapph lattice laser. My first six months were spent divided between the NPL and Oxford; Mondays and Tuesdays I attended lectures while the rest of the week was spent working in the lab learning to run the experiment.

My first tasks mostly involved electronics and optics. I built a low noise current supply capable of outputting 20 A which is used to the drive the Helmholtz coils needed for spectroscopy (the requirement for high output current is mostly for operating a clock using the even isotope ^{88}Sr). We use the same design for driving the mixing and compensation coils for Sr2. I also built a general purpose bi-polar, opto-isolation board which comes in handy for monitoring signals, breaking ground-loops and protecting delicate instrumentation from voltage spikes. I also spent a month or so trying to build a lock-in servo for actively keeping a slave diode injected using a design similar to that presented in [27], but eventually settled on the alternate solution similar to that presented in [129] with the exception of using a grating instead of a cavity. In regards to optics, I rebuilt part of the 689 nm laser system such that the

ECDL light was no longer directly used for cooling, but rather a diode was seeded by the transmission of the cavity. This increased the available power for laser cooling while suppressing laser frequency noise induced by the servo-sidebands. Once the lattice laser arrived, I also built the optics and electronics for locking it to the transfer cavity.

Once the lattice laser arrived, the entire lattice team (which now included Marco Manchetti, a visiting PhD student from the University of Birmingham) worked tirelessly for several months completing the clock in preparation for the ITOC clock comparison in the summer of 2015. During this time my major focus was writing the clock servo software and few other general purpose imaging programs helpful for diagnostics and experimental sequence optimization. Working with Ian, we also designed and tested a new optical shutter based on a piezo cantilever. Such an actuator has the benefit of providing fast switching speeds, while emitting minimal vibration to the optical table which can disturb the lasers. Details regarding the design and performance can be found in [13].

After getting the clock operational, it quickly became obvious we needed some degree of automation if we hoped to run the system for several days without human intervention. I implemented a computer controlled auto-relocking and remote servicing of all the lasers which proved to be a fairly reliable solution; during ITOC we had nearly 80% uptime over three weeks despite having only locked to the clock transition a few weeks prior to the start of the campaign. In regards to evaluating the performance of Sr1, I wrote much of the initial data analysis software (sideband thermometry fitting, modelling inhomogeneous broadened Rabi spectra, and analysing BBR systematic shift). Around this time we also proposed and demonstrated modified Hyper-Ramsey to eliminate the probe related Stark shifts when using ^{88}Sr for a clock. The idea was conceived by Richard Hobson, but I contributed to much of the experimental implantation and data analysis.

In my third year, my two main undertakings were implementing Rydberg spectroscopy for the characterization of the DC-Shift in Sr1 with the help of Richard and Ian and working on the design of Sr2 which mostly included the science package and laser systems. During this time, the optical fibre link connecting NPL to SYRTE and PTB was finished enabling direct frequency comparison of lattice clocks across Europe. Aside from helping to run the clock during such comparison campaigns, I also wrote the data analysis scripts needed to process the raw link data into meaningful clock comparison frequency ratios. This analysis was independently done by every institute and is not included in the thesis as the results are still in preparation at this time, but will hopefully be finished shortly.

My final year was dedicated to the construction of Sr2. Completing the work needed to realise a functional clock (excluding the frequency comb and cavity) were mostly shared between Richard and myself. There are far too many jobs to list out, but it was very much a team effort that proved to be highly efficient – going from baking out the vacuum system in January to getting our first cavity enhanced non-destructive detection signals in June of the same year.

Appendix B

Design of the Vacuum system

As explained in 3.3.4, we must maintain UHV inside the chamber to suppress frequency offsets arising from collisions with background gas particles. The exponential decay rate R associated with this collision-induced atom loss scales linearly with pressure. If we want a lattice lifetime exceeding 10 seconds, roughly that of Sr1, we need to target a pressure inside the chamber of below 3×10^{-10} mbar [98]. In order to engineer a vacuum system that meets this demand, it is important to estimate the main sources of outgassing. In our system, the major contributor to the background pressure is the strontium oven which out-gasses a significant amount of hydrogen at elevated temperatures. Once outgassing loads are known, the necessary pumping speed S can be calculated to achieve the target pressure. It is also important to take into account the conductance C from the position of the atoms to the pump as the geometry of the vacuum system can limit the efficacy of the pumping if the contaminants must travel a long way through a narrow passage before reaching the pump. In general, the effective pumping speed S_E is given by:

$$S_E^{-1} = S^{-1} + C^{-1}. \tag{B.1}$$

If the desired effective pumping speed to reach the target base pressure is prohibitively

large, either because of a lack of conductance or available pumps, one can isolate sources of outgassing from the atoms by adding differential pumping in combination with intermediate pumping stages. Neglecting other sources of outgassing in our system, by connecting the oven section via a small tube, the pressure in the main chamber is reduced by a factor C_{dp}/P_o where C_{dp} conductance of differential pumping tube and P_o is the pumping speed **before** the tube.[1]

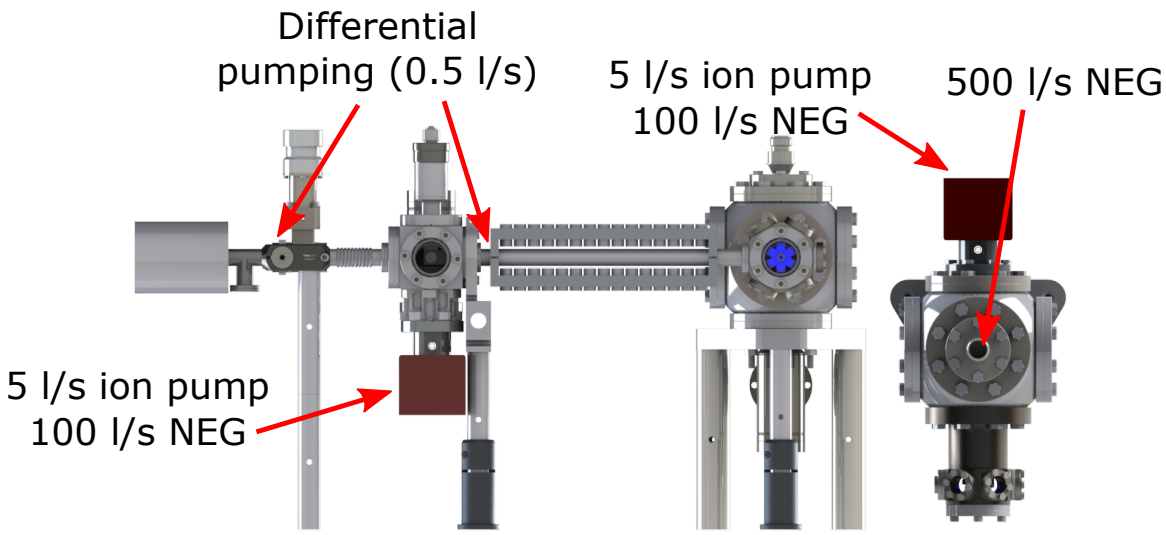


Figure B.1: Rendering of the vacuum chamber showing the source section, spectroscopy region, and pumping cube. Highlighted are location of the pumps and differential pumping tubes.

For our vacuum system, the spectroscopy section containing the pyramid MOT structure was custom made by Allectra GmbH from titanium as it has a much lower hydrogen outgassing rate compared to steel. Optical access is provided by one axial DN40 viewport used to deliver the MOT and repump beams along with six radial DN16 viewports used for coupling light in and out of the cavities. Opposite the DN40 viewport, the main chamber connects to a spherical cube (Kimball Physics MCF450-

[1]For a long tube of diameter D and length L, the conductance for hydrogen at room temperate is given by $44\,\mathrm{l}\cdot\mathrm{s}\cdot\mathrm{cm}^{-2}\ D^3/L$ [25], to calculate conductance for other species or temperatures one can make use of the fact the conductance is proportional to mean thermal velocity.

SphCube-E6) via a DN63 flange. The large flange size was chosen to maximise conductance. The cube contains a $500\,\mathrm{ls}^{-1}$ SAES CapaciTorr and $100\,\mathrm{ls}^{-1}$ combination NEXTorr combination non-evaporable getter (NEG) and ion pump. The main chamber is separated from the pumping cube by a 5 cm tube to accommodate one of the MOT coils which sits between the two sections. The estimated conductance from the centre of the pyramid MOT to the pump region is $200\,\mathrm{ls}^{-1}$. The pyramid MOT is attached to a titanium standoff which connects to the main chamber via 'groove grabbers' attached to the pumping cube as shown in figure B.1.

The other source of outgassing is from the effusive oven. We elected to use an externally heated stainless steel ConFlat (CF) nipple to generate the strontium vapour. The vapour exits the oven through a 2 mm diameter aperture. The opening is press fitted with microtubes measuring 7 mm in length with an inner (outer) diameter of 200 µm (350 µm). In regards to the outgassing of the stainless steel[2], which is intended to be operated at 500 °C, it is estimated to be $2 \times 10^{-8}\,\mathrm{mbar} \cdot \mathrm{l} \cdot \mathrm{cm}^{-2} \cdot \mathrm{s}^{-1}$. This corresponds to a total outgassing of $5 \times 10^{-7}\,\mathrm{mbar} \cdot \mathrm{l} \cdot \mathrm{s}^{-1}$ given the internal surface area of the oven. Without any differential pumping, the estimated equilibrium pressure inside the main chamber would be 2.5×10^{-9} mbar – which fails to meet our desired target pressure. We therefore need to reduce the pressure by a factor of 200 which is achieved by employing two stages of differential pumping. The oven region is pumped by a $10\,\mathrm{ls}^{-1}$ Gamma Vacuum ion pump and isolated from rest of the vacuum system by a $0.5\,\mathrm{ls}^{-1}$ hole made by drilling a 2 mm in a DN16d double-sided blank. Immediately after the differential pumping flange is a SAES NEXTorr pump providing $100\,\mathrm{ls}^{-1}$ of hydrogen pumping, followed by a second differential pumping tube[3] with a conductance of $0.2\,\mathrm{ls}^{-1}$. Therefore, the total reduction in the contribution of

[2]Initially we built and tested a titanium oven as it would have significantly less outgassing and would require only a single differential pumping stage, but after heating it to 450 °C and cooling it to room temperature the screws holding the CF flange became loose and the knife edge leaked. This possibly resulted from the fact titanium undergoes a phase transformation from an anatase to rutile structure around this temperature, but this is purely speculative.

[3]The tube measures 3.5 mm in diameter by 10 cm in length and is inset in the Zeeman slower tube.

the oven outgassing to the pressure in the main chamber is 10,000. It is important to highlight that although the NEXTorr pump and $0.2\,\mathrm{ls}^{-1}$ tube would provide sufficient isolation, the equilibrium pressure at the NEG would be 5×10^{-9} mbar – saturating it in approximately one year. Hence, the main function of the first differential pumping stage is to increase the longevity of the getter. The intermediate section also has a viewport used for diagnostic spectroscopy of the atomic beam and a rotatable shutter to block the atomic beam and possible blackbody radiation from the oven.

The other source of out-gassing needing consideration is the pyramid MOT structure comprising of the mirrors, piezos, wires and epoxies. After pre-baking the structure inside a dedicated UHV bake-out station, the steady state pressure was 4×10^{-10} mbar. Given that the pumping speed of the station of $20\,\mathrm{ls}^{-1}$ results in a pressure floor of 1×10^{-10} mbar when empty, a crude estimate for the out-gassing from the pyramid MOT and cavity structure is $6 \times 10^{-9}\,\mathrm{mbar} \cdot \mathrm{l} \cdot \mathrm{s}^{-1}$. The outgassing was dominated by hydrogen as indicated by an residual gas analyzer (RGA). Given the effective pumping speed in the main chamber is $150\,\mathrm{ls}^{-1}$, this out-gassing should result in a equilibrium pressure of 4×10^{-11} mbar. Hence, the combination of oven and pyramid MOT out-gassing should be 9×10^{-11} mbar resulting in a trap lifetime of 35 seconds.

Appendix C

List of Publications

1. S. Donnellan, I. R. Hill, W. Bowden, and R. Hobson. A scalable arbitrary waveform generator for atomic physics experiments based on field-programmable gate array technology. *Accepted to Review of Scientific Instruments*, 2019.

2. W. Bowden, R. Hobson, P. Huillery, P. Gill, M. P. A. Jones, and I. R. Hill. Rydberg electrometry for optical lattice clocks. *Physical Review A*, 96(2), 023419, 2017.

3. P. Delva et al. Test of Special Relativity Using a Fiber Network of Optical Clocks. *Phys. Rev. Lett.* 118, 221102, 2017.

4. W. Bowden, I. R. Hill, P. E. G. Baird, and P. Gill. A High-Performance, Low-Cost Laser Shutter using a Piezoelectric Cantilever Actuator. *Review of Scientific Instruments* 88, 016102, 2017.

5. I. R. Hill, R. Hobson, W. Bowden, E. M. Bridge, S. Donnellan, E. A. Curtis, and P. Gill. A low maintenance Sr optical lattice clock. *In Journal of Physics: Conference Series* Vol. 723, No. 1, p. 012019, 2017.

6. R. Hobson, W. Bowden, S. A. King, P. E. G Baird, I. R. Hill, and P. Gill. Modified hyper-Ramsey methods for the elimination of probe shifts in optical

clocks. *Physical Review A*, 93(1), 010501, 2016.

7. W. Bowden, W. Gunton, W. Semczuk, K. Dare, and K. W. Madison. An adaptable dual species effusive source and Zeeman slower design demonstrated with Rb and Li. *Review of Scientific Instruments*, 87(4), 043111, 2016.

8. M. Semczuk, W. Gunton, W., Bowden, and K. W. Madison. Anomalous Behavior of Dark States in Quantum Gases of ^{6}Li. *Phys. Rev. Lett.* 113(5), 055302, 2014.

9. Lecault, et al. High-throughput analysis of single hematopoietic stem cell proliferation in microfluidic cell culture arrays. *Nature methods*, 8(7), 581, 2011.

Bibliography

[1] R. P. Abel, C. Carr, U. Krohn, and C. S. Adams. Electrometry near a dielectric surface using rydberg electromagnetically induced transparency. *Phys. Rev. A*, 84:023408, Aug 2011.

[2] Marc Abrahams. How to pour the perfect cup of coffee. *The Guardian*, Jul 2010.

[3] T. Akatsuka, M. Takamoto, and H. Katori. Optical lattice clocks with non-interacting bosons and fermions. *Nature Physics*, 4(12):954–959, December 2008.

[4] Ali Al-Masoudi, Sören Dörscher, Sebastian Häfner, Uwe Sterr, and Christian Lisdat. Noise and instability of an optical lattice clock. *Phys. Rev. A*, 92:063814, Dec 2015.

[5] D. A. Anderson, S. A. Miller, G. Raithel, J. A. Gordon, M. L. Butler, and C. L. Holloway. Optical measurements of strong microwave fields with rydberg atoms in a vapor cell. *Phys. Rev. Applied*, 5:034003, Mar 2016.

[6] J.J. Arlt, O. Maragò, S. Webster, S. Hopkins, and C.J. Foot. A pyramidal magneto-optical trap as a source of slow atoms. *Optics Communications*, 157(1):303 – 309, 1998.

[7] Asimina Arvanitaki, Junwu Huang, and Ken Van Tilburg. Searching for dilaton

dark matter with atomic clocks. *Phys. Rev. D*, 91:015015, Jan 2015.

[8] RMA Azzam. Phase shifts that accompany total internal reflection at a dielectric–dielectric interface. *JOSA A*, 21(8):1559–1563, 2004.

[9] D. S. Barker, B. J. Reschovsky, N. C. Pisenti, and G. K. Campbell. Enhanced magnetic trap loading for atomic strontium. *Phys. Rev. A*, 92:043418, Oct 2015.

[10] K. Beloy, N. Hinkley, N. B. Phillips, J. A. Sherman, M. Schioppo, J. Lehman, A. Feldman, L. M. Hanssen, C. W. Oates, and A. D. Ludlow. Atomic clock with 1×10^{-18} room-temperature blackbody stark uncertainty. *Phys. Rev. Lett.*, 113:260801, Dec 2014.

[11] Slawomir Bilicki. Strontium optical lattice clocks : clock comparisons for timescales and fundamental physics applications, 2018.

[12] S. Blatt, J. W. Thomsen, G. K. Campbell, A. D. Ludlow, M. D. Swallows, M. J. Martin, M. M. Boyd, and J. Ye. Rabi spectroscopy and excitation inhomogeneity in a one-dimensional optical lattice clock. *Phys. Rev. A*, 80:052703, Nov 2009.

[13] W. Bowden, I. R. Hill, P. E. G. Baird, and P. Gill. Note: A high-performance, low-cost laser shutter using a piezoelectric cantilever actuator. *Review of Scientific Instruments*, 88(1):016102, 2017.

[14] W. Bowden, R. Hobson, P. Huillery, P. Gill, M. P. A. Jones, and I. R. Hill. Rydberg electrometry for optical lattice clocks. *Phys. Rev. A*, 96:023419, Aug 2017.

[15] Martin M Boyd. High precision spectroscopy of strontium in an optical lattice: Towards a new standard for frequency and time, 2007.

[16] Martin M. Boyd, Tanya Zelevinsky, Andrew D. Ludlow, Sebastian Blatt, Thomas Zanon-Willette, Seth M. Foreman, and Jun Ye. Nuclear spin effects in optical lattice clocks. *Phys. Rev. A*, 76:022510, Aug 2007.

[17] B. H. Bransden and C. J. Joachain. *Physics of atoms and molecules*. Harlow : Prentice Hall, 2nd ed edition, 2003. Previous ed.: 1983.

[18] R. C. Brown, N. B. Phillips, K. Beloy, W. F. McGrew, M. Schioppo, R. J. Fasano, G. Milani, X. Zhang, N. Hinkley, H. Leopardi, T. H. Yoon, D. Nicolodi, T. M. Fortier, and A. D. Ludlow. Hyperpolarizability and operational magic wavelength in an optical lattice clock. *Phys. Rev. Lett.*, 119:253001, Dec 2017.

[19] Tomasz M Brzozowski, Maria Maczynska, Michal Zawada, Jerzy Zachorowski, and Wojciech Gawlik. Time-of-flight measurement of the temperature of cold atoms for short trap-probe beam distances. *Journal of Optics B: Quantum and Semiclassical Optics*, 4(1):62, 2002.

[20] S. L. Campbell, R. B. Hutson, G. E. Marti, A. Goban, N. Darkwah Oppong, R. L. McNally, L. Sonderhouse, J. M. Robinson, W. Zhang, B. J. Bloom, and J. Ye. A fermi-degenerate three-dimensional optical lattice clock. *Science*, 358(6359):90–94, 2017.

[21] Sara Campbell. A fermi-degenerate three-dimensional optical lattice clock, 2017.

[22] Andrea Camposeo, A Piombini, F Cervelli, Francesco Tantussi, Francesco Fuso, and Ennio Arimondo. A cold cesium atomic beam produced out of a pyramidal funnel. *Optics Communications*, 200(1-6):231–239, 2001.

[23] Henning Carstens, Simon Holzberger, Jan Kaster, Johannes Weitenberg, Volodymyr Pervak, Alexander Apolonski, Ernst Fill, Ferenc Krausz, and

Ioachim Pupeza. Large-mode enhancement cavities. *Opt. Express*, 21(9):11606–11617, May 2013.

[24] Zilong Chen, Justin G. Bohnet, Joshua M. Weiner, Kevin C. Cox, and James K. Thompson. Cavity-aided nondemolition measurements for atom counting and spin squeezing. *Phys. Rev. A*, 89:043837, Apr 2014.

[25] P. Chiggiato. Vacuum technology for ion sources. 2014.

[26] Cheng Chin, V V Flambaum, and M G Kozlov. Ultracold molecules: new probes on the variation of fundamental constants. *New Journal of Physics*, 11(5):055048, 2009.

[27] Sheng-wey Chiow, Quan Long, Christoph Vo, Holger Müller, and Steven Chu. Extended-cavity diode lasers with tracked resonances. *Applied optics*, 46(33):7997–8001, 2007.

[28] C. W. Chou, D. B. Hume, M. J. Thorpe, D. J. Wineland, and T. Rosenband. Quantum coherence between two atoms beyond $q = 10^{15}$. *Phys. Rev. Lett.*, 106:160801, Apr 2011.

[29] E. Anne Curtis, Christopher W. Oates, and Leo Hollberg. Quenched narrow-line second- and third-stage laser cooling of ^{40}Ca. *J. Opt. Soc. Am. B*, 20(5):977–984, May 2003.

[30] Thibault Damour and Freeman Dyson. The oklo bound on the time variation of the fine-structure constant revisited. *Nuclear Physics B*, 480(1):37 – 54, 1996.

[31] H Dehmelt. Coherent spectroscopy on single atomic system at rest in free space ii. *Le Journal de Physique Colloques*, 42(C8):C8–299, 1981.

[32] P. Delva, J. Lodewyck, S. Bilicki, E. Bookjans, G. Vallet, R. Le Targat, P.-E. Pottie, C. Guerlin, F. Meynadier, C. Le Poncin-Lafitte, O. Lopez, A. Amy-

Klein, W.-K. Lee, N. Quintin, C. Lisdat, A. Al-Masoudi, S. Dörscher, C. Grebing, G. Grosche, A. Kuhl, S. Raupach, U. Sterr, I. R. Hill, R. Hobson, W. Bowden, J. Kronjäger, G. Marra, A. Rolland, F. N. Baynes, H. S. Margolis, and P. Gill. Test of special relativity using a fiber network of optical clocks. *Phys. Rev. Lett.*, 118:221102, Jun 2017.

[33] A. Derevianko and M. Pospelov. Hunting for topological dark matter with atomic clocks. *Nature Physics*, 10:933 EP –, Nov 2014.

[34] G John Dick. Local oscillator induced instabilities in trapped ion frequency standards. 1987.

[35] R. H. Dicke. Coherence in spontaneous radiation processes. *Phys. Rev.*, 93:99–110, Jan 1954.

[36] Gianni Di Domenico, Stéphane Schilt, and Pierre Thomann. Simple approach to the relation between laser frequency noise and laser line shape. *Appl. Opt.*, 49(25):4801–4807, Sep 2010.

[37] D. K. Doughty and J. E. Lawler. Spatially resolved electric field measurements in the cathode fall using optogalvanic detection of Rydberg atoms. *Applied Physics Letters*, 45(6):611–613, 1984.

[38] Stephan Falke, Nathan Lemke, Christian Grebing, Burghard Lipphardt, Stefan Weyers, Vladislav Gerginov, Nils Huntemann, Christian Hagemann, Ali Al-Masoudi, Sebastian Häfner, Stefan Vogt, Uwe Sterr, and Christian Lisdat. A strontium lattice clock with 3×10^{-17} inaccuracy and its frequency. *New Journal of Physics*, 16(7):073023, 2014.

[39] Thomas F. Gallagher. *Rydberg Atoms*. Cambridge University Press, Cambridge, 1994.

[40] Kurt Gibble. Decoherence and collisional frequency shifts of trapped bosons and fermions. *Phys. Rev. Lett.*, 103:113202, Sep 2009.

[41] Kurt Gibble. Scattering of cold-atom coherences by hot atoms: Frequency shifts from background-gas collisions. *Phys. Rev. Lett.*, 110:180802, May 2013.

[42] R. M. Godun, P. B. R. Nisbet-Jones, J. M. Jones, S. A. King, L. A. M. Johnson, H. S. Margolis, K. Szymaniec, S. N. Lea, K. Bongs, and P. Gill. Frequency ratio of two optical clock transitions in $^{171}Yb^{+}$ and constraints on the time variation of fundamental constants. *Phys. Rev. Lett.*, 113:210801, Nov 2014.

[43] Jan Grünert and Andreas Hemmerich. Sub-doppler magneto-optical trap for calcium. *Phys. Rev. A*, 65:041401, Mar 2002.

[44] J. Guena, P. Rosenbusch, P. Laurent, M. Abgrall, D. Rovera, G. Santarelli, M. E. Tobar, S. Bize, and A. Clairon. Demonstration of a dual alkali Rb/Cs fountain clock. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 57(3):647–653, March 2010.

[45] Sebastian Häfner, Stephan Falke, Christian Grebing, Stefan Vogt, Thomas Legero, Mikko Merimaa, Christian Lisdat, and Uwe Sterr. 8×10^{-17} fractional laser frequency instability with a long room-temperature cavity. *Opt. Lett.*, 40(9):2112–2115, May 2015.

[46] Ryan K Hanley, Alistair D Bounds, Paul Huillery, Niamh C Keegan, Riccardo Faoro, Elizabeth M Bridge, Kevin J Weatherill, and Matthew P A Jones. Probing interactions of thermal Sr rydberg atoms using simultaneous optical and ion detection. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 50(11):115002, 2017.

[47] Dirk P. Hansen, Janis R. Mohr, and Andreas Hemmerich. Magnetic trapping of metastable calcium atoms. *Phys. Rev. A*, 67:021401, Feb 2003.

[48] Ian R Hill, Yuri B Ovchinnikov, Elizabeth M Bridge, E Anne Curtis, and Patrick Gill. Zeeman slowers for strontium based on permanent magnets. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 47(7):075006, 2014.

[49] Ian Robert Hill. *Development of an apparatus for a strontium optical lattice optical frequency standard*. PhD thesis, Imperial College London, 2012.

[50] M. Hiller, S. Yoshida, J. Burgdörfer, S. Ye, X. Zhang, and F. B. Dunning. Characterizing high-n quasi-one-dimensional strontium rydberg atoms. *Phys. Rev. A*, 89:023426, Feb 2014.

[51] Richard Hobson. *An optical lattice clock with neutral strontium*. PhD thesis, Oxford University, 2016.

[52] Onur Hosten, Nils J. Engelsen, Rajiv Krishnakumar, and Mark A. Kasevich. Measurement noise 100 times lower than the quantum-projection limit using entangled atoms. *Nature*, 529:505 EP –, Jan 2016.

[53] Wayne Hu, Rennan Barkana, and Andrei Gruzinov. Fuzzy cold dark matter: The wave properties of ultralight particles. *Phys. Rev. Lett.*, 85:1158–1161, Aug 2000.

[54] J. Huckans, W. Dubosclard, E. Maréchal, O. Gorceix, B. Laburthe-Tolra, and M. Robert de Saint-Vincent. Note on the reflectance of mirrors exposed to a strontium beam, 2018.

[55] David B. Hume and David R. Leibrandt. Probing beyond the laser coherence time in optical clock comparisons. *Phys. Rev. A*, 93:032138, Mar 2016.

[56] N. Huntemann, B. Lipphardt, Chr. Tamm, V. Gerginov, S. Weyers, and E. Peik. Improved limit on a temporal variation of m_p/m_e from comparisons of Yb^+ and Cs atomic clocks. *Phys. Rev. Lett.*, 113:210802, Nov 2014.

[57] L. Isaev, A. Kaufman, G. Ortiz, and A. M. Rey. Topological superfluidity with repulsive alkaline-earth atoms in optical lattices, 2017.

[58] David J. Jones, Scott A. Diddams, Jinendra K. Ranka, Andrew Stentz, Robert S. Windeler, John L. Hall, and Steven T. Cundiff. Carrier-envelope phase control of femtosecond mode-locked lasers and direct optical frequency synthesis. *Science*, 288(5466):635–639, 2000.

[59] Hidetoshi Katori, Tetsuya Ido, Yoshitomo Isoya, and Makoto Kuwata-Gonokami. Magneto-optical trapping and cooling of strontium atoms down to the photon recoil temperature. *Phys. Rev. Lett.*, 82:1116–1119, Feb 1999.

[60] Hidetoshi Katori, V. D. Ovsiannikov, S. I. Marmo, and V. G. Palchikov. Strategies for reducing the light shift in atomic clocks. *Phys. Rev. A*, 91:052503, May 2015.

[61] Akio Kawasaki, Boris Braverman, Edwin Pedrozo-Peñafiel, Chi Shu, Simone Colombo, Zeyang Li, Özge Özel, Wenlan Chen, Leonardo Salvi, André Heinz, David Levonian, Daisuke Akamatsu, Yanhong Xiao, and Vladan Vuletić. Geometrically asymmetric optical cavity for strong atom-photon coupling. *Phys. Rev. A*, 99:013437, Jan 2019.

[62] J. A. Kim, K. I. Lee, H. R. Noh, W. Jhe, and M. Ohtsu. Atom trap in an axicon mirror. *Opt. Lett.*, 22(2):117–119, Jan 1997.

[63] Sebastian Kraft, Felix Vogt, Oliver Appel, Fritz Riehle, and Uwe Sterr. Bose-einstein condensation of alkaline earth atoms: ^{40}Ca. *Phys. Rev. Lett.*, 103:130401, Sep 2009.

[64] A. Kuzmich, N. P. Bigelow, and L. Mandel. Atomic quantum non-demolition measurements and squeezing. *EPL (Europhysics Letters)*, 42(5):481, 1998.

[65] R. Le Targat, L. Lorini, Y. Le Coq, M. Zawada, J. Guéna, M. Abgrall, M. Gurov, P. Rosenbusch, D. G. Rovera, B. Nagórny, R. Gartman, P. G. Westergaard, M. E. Tobar, M. Lours, G. Santarelli, A. Clairon, S. Bize, P. Laurent, P. Lemonde, and J. Lodewyck. Experimental realization of an optical second with strontium lattice clocks. *Nature Communications*, 4:2109 EP –, Jul 2013. Article.

[66] Jongmin Lee, Geert Vrijsen, Igor Teper, Onur Hosten, and Mark A. Kasevich. Many-atom cavity QED system with homogeneous atom-cavity coupling. *Opt. Lett.*, 39(13):4005–4008, Jul 2014.

[67] K. I. Lee, J. A. Kim, H. R. Noh, and W. Jhe. Single-beam atom trap in a pyramidal and conical hollow mirror. *Opt. Lett.*, 21(15):1177–1179, Aug 1996.

[68] N. Leefer, C. T. M. Weber, A. Cingöz, J. R. Torgerson, and D. Budker. New limits on variation of the fine-structure constant using atomic dysprosium. *Phys. Rev. Lett.*, 111:060801, Aug 2013.

[69] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Rev. Mod. Phys.*, 75:281–324, Mar 2003.

[70] N. D. Lemke, A. D. Ludlow, Z. W. Barber, T. M. Fortier, S. A. Diddams, Y. Jiang, S. R. Jefferts, T. P. Heavner, T. E. Parker, and C. W. Oates. Spin-1/2 optical lattice clock. *Phys. Rev. Lett.*, 103:063001, Aug 2009.

[71] N. D. Lemke, J. von Stecher, J. A. Sherman, A. M. Rey, C. W. Oates, and A. D. Ludlow. p-wave cold collisions in an optical lattice clock. *Phys. Rev. Lett.*, 107:103902, Aug 2011.

[72] Pierre Lemonde and Peter Wolf. Optical lattice clock with atoms confined in a shallow trap. *Phys. Rev. A*, 72:033409, Sep 2005.

[73] Ian D. Leroux, Monika H. Schleier-Smith, and Vladan Vuletić. Implementation of cavity squeezing of a collective atomic spin. *Phys. Rev. Lett.*, 104:073602, Feb 2010.

[74] Filippo Levi, Davide Calonico, Claudio E Calosso, Aldo Godone, Salvatore Micalizio, and Giovanni A Costanzo. Accuracy evaluation of ITCsF2: a nitrogen cooled caesium fountain. *Metrologia*, 51(3):270, 2014.

[75] Ruoxin Li, Kurt Gibble, and Krzysztof Szymaniec. Improved accuracy of the NPL-CsF2 primary frequency standard: evaluation of distributed cavity phase and microwave lensing frequency shifts. *Metrologia*, 48(5):283, 2011.

[76] G. Lochead, D. Boddy, D. P. Sadler, C. S. Adams, and M. P. A. Jones. Number-resolved imaging of excited-state atoms using a scanning autoionization microscope. *Phys. Rev. A*, 87:053409, May 2013.

[77] J. Lodewyck, M. Zawada, L. Lorini, M. Gurov, and P. Lemonde. Observation and cancellation of a perturbing dc stark shift in strontium optical lattice clocks. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 59(3):411–415, March 2012.

[78] Jérôme Lodewyck, Philip G. Westergaard, and Pierre Lemonde. Nondestructive measurement of the transition probability in a Sr optical lattice clock. *Phys. Rev. A*, 79:061401, Jun 2009.

[79] R. Long, A. K. Tuchman, and M. A. Kasevich. Multiple frequency modulation for low-light atom measurements in an optical cavity. *Opt. Lett.*, 32(17):2502–2504, Sep 2007.

[80] A. D. Ludlow, X. Huang, M. Notcutt, T. Zanon-Willette, S. M. Foreman, M. M. Boyd, S. Blatt, and J. Ye. Compact, thermal-noise-limited optical cavity for diode laser stabilization at 1×10^{-15}. *Opt. Lett.*, 32(6):641–643, Mar 2007.

[81] A. D. Ludlow, N. D. Lemke, J. A. Sherman, C. W. Oates, G. Quéméner, J. von Stecher, and A. M. Rey. Cold-collision-shift cancellation and inelastic scattering in a Yb optical lattice clock. *Phys. Rev. A*, 84:052724, Nov 2011.

[82] Andrew D Ludlow. *The strontium optical lattice clock: optical spectroscopy with sub-Hertz accuracy.* PhD thesis, University of Colorado at Boulder, 2008.

[83] Long-Sheng Ma, Peter Jungner, Jun Ye, and John L. Hall. Delivering the same optical frequency at two places: accuratecancellation of phase noise introduced by an optical fiber or other time-varyingpath. *Opt. Lett.*, 19(21):1777–1779, Nov 1994.

[84] N. Malossi, S. Damkjær, P. L. Hansen, L. B. Jacobsen, L. Kindt, S. Sauge, J. W. Thomsen, F. C. Cruz, M. Allegrini, and E. Arimondo. Two-photon cooling of magnesium atoms. *Phys. Rev. A*, 72:051403, Nov 2005.

[85] G. Edward Marti, Ross B. Hutson, Akihisa Goban, Sara L. Campbell, Nicola Poli, and Jun Ye. Imaging optical frequencies with 100 μHz precision and 1.1 μm resolution. *Phys. Rev. Lett.*, 120:103201, Mar 2018.

[86] Michael J. Martin. *Quantum Metrology and Many-Body Physics: Pushing the Frontier of the Optical Lattice Clock.* PhD thesis, University of Colorado, 2013.

[87] D. G. Matei, T. Legero, S. Häfner, C. Grebing, R. Weyrich, W. Zhang, L. Sonderhouse, J. M. Robinson, J. Ye, F. Riehle, and U. Sterr. 1.5 μm lasers with sub-10 mHz linewidth. *Phys. Rev. Lett.*, 118:263202, Jun 2017.

[88] Arthur Matveev, Christian G. Parthey, Axel Beyer, Nikolai Kolachevsky, Janis Alnis, Randolf Pohl, Thomas Udem, and Theodor W. Haensch. Systematic frequency shifts in spectroscopy of 1S-2S transition in atomic hydrogen. In *CLEO:2011 - Laser Applications to Photonic Applications*, page QTuJ6. Optical Society of America, 2011.

[89] M. McDonald, B. H. McGuyer, G. Z. Iwata, and T. Zelevinsky. Thermometry via light shifts in optical lattices. *Phys. Rev. Lett.*, 114:023001, Jan 2015.

[90] W. F. McGrew, X. Zhang, R. J. Fasano, S. A. Schäffer, K. Beloy, D. Nicolodi, R. C. Brown, N. Hinkley, G. Milani, M. Schioppo, T. H. Yoon, and A. D. Ludlow. Atomic clock performance beyond the geodetic limit, 2018.

[91] PG Mickelson, YN Martinez de Escobar, P Anzel, BJ DeSalvo, SB Nagel, AJ Traverso, M Yan, and TC Killian. Repumping and spectroscopy of laser-cooled Sr atoms using the (5s5p) $^3P_2$–(5s4d) $^3D_2$ transition. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 42(23):235001, 2009.

[92] Thomas Middelmann, Stephan Falke, Christian Lisdat, and Uwe Sterr. High accuracy correction of blackbody radiation shift in an optical lattice clock. *Phys. Rev. Lett.*, 109:263004, Dec 2012.

[93] J Millen, G Lochead, G R Corbett, R M Potvliege, and M P A Jones. Spectroscopy of a cold strontium rydberg gas. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 44(18):184001, 2011.

[94] J Mitroy, M S Safronova, and Charles W Clark. Theory and applications of atomic and ionic polarizabilities. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 43(20):202001, 2010.

[95] J. Mitroy and J.Y. Zhang. Dispersion and polarization interactions of the strontium atom. *Molecular Physics*, 108(15):1999–2006, 2010.

[96] P H Moriya, M O Araújo, F Todão, M Hemmerling, H Keßler, R F Shiozaki, R Celistrino Teixeira, and Ph W Courteille. Comparison between 403 nm and 497 nm repumping schemes for strontium magneto-optical traps, 2018.

[97] Takashi Mukaiyama, Hidetoshi Katori, Tetsuya Ido, Ying Li, and Makoto

Kuwata-Gonokami. Recoil-limited laser cooling of ^{87}Sr atoms near the fermi temperature. *Phys. Rev. Lett.*, 90:113002, Mar 2003.

[98] S. B. Nagel, C. E. Simien, S. Laha, P. Gupta, V. S. Ashoka, and T. C. Killian. Magnetic trapping of metastable $^3P_2$ atomic strontium. *Phys. Rev. A*, 67:011401, Jan 2003.

[99] Nils Nemitz, Takuya Ohkubo, Masao Takamoto, Ichiro Ushijima, Manoj Das, Noriaki Ohmae, and Hidetoshi Katori. Frequency ratio of Yb and Sr clocks with 5 $\times 10^{-17}$ uncertainty at 150 seconds averaging time. *Nature Photonics*, 10:258 EP –, Feb 2016.

[100] Matthew Newville, Till Stensitzki, Daniel B. Allen, and Antonino Ingargiola. LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python¶, September 2014.

[101] T. L. Nicholson, S. L. Campbell, R. B. Hutson, G. E. Marti, B. J. Bloom, R. L. McNally, W. Zhang, M. D. Barrett, M. S. Safronova, G. F. Strouse, W. L. Tew, and J. Ye. Systematic evaluation of an atomic clock. *Nature Communications*, 6:6896 EP –, Apr 2015. Article.

[102] T. L. Nicholson, M. J. Martin, J. R. Williams, B. J. Bloom, M. Bishof, M. D. Swallows, S. L. Campbell, and J. Ye. Comparison of two independent sr optical clocks with 1×10^{-17} stability at 10^3 **s**. *Phys. Rev. Lett.*, 109:230801, Dec 2012.

[103] Travis L Nicholson. *A new record in atomic clock performance.* PhD thesis, University of Colorado, 2015.

[104] P. B. R. Nisbet-Jones, S. A. King, J. M. Jones, R. M. Godun, C. F. A. Baynham, K. Bongs, M. Dolezal, P. Balling, and P. Gill. A single-ion trap with minimized ion-environment interactions. *Applied Physics B*, 122(3):57, 2016.

[105] Matthew A. Norcia, Julia R. K. Cline, Juan A. Muniz, John M. Robinson, Ross B. Hutson, Akihisa Goban, G. Edward Marti, Jun Ye, and James K. Thompson. Frequency measurements of superradiance from the strontium clock transition. *Phys. Rev. X*, 8:021036, May 2018.

[106] Matthew Andrew Norcia. New tools for precision measurement and quantum science with narrow linewidth optical transitions. 2018.

[107] Ingo Nosske, Luc Couturier, Fachao Hu, Canzhu Tan, Chang Qiao, Jan Blume, Y. H. Jiang, Peng Chen, and Matthias Weidemüller. Two-dimensional magneto-optical trap as a source for cold strontium atoms. *Phys. Rev. A*, 96:053415, Nov 2017.

[108] Oleksiy Onishchenko, Sergey Pyatchenkov, Alexander Urech, Chun-Chia Chen, Shayne Bennetts, Georgios A. Siviloglou, and Florian Schreck. The frequency of the ultranarrow $^1S_0 - {}^3P_2$ transition in ^{87}Sr. 2018.

[109] A. Osterwalder and F. Merkt. Using high rydberg states as electric field sensors. *Phys. Rev. Lett.*, 82:1831–1834, Mar 1999.

[110] William D Phillips and Harold Metcalf. Laser deceleration of an atomic beam. *Physical Review Letters*, 48(9):596, 1982.

[111] S. E. Pollack, M. D. Turner, S. Schlamminger, C. A. Hagedorn, and J. H. Gundlach. Charge management for gravitational-wave observatories using uv leds. *Phys. Rev. D*, 81:021101, Jan 2010.

[112] S. G. Porsev, Andrew D. Ludlow, Martin M. Boyd, and Jun Ye. Determination of Sr properties for a high-accuracy optical clock. *Phys. Rev. A*, 78:032508, Sep 2008.

[113] S. G. Porsev, M. S. Safronova, U. I. Safronova, and M. G. Kozlov. Multipolar

polarizabilities and hyperpolarizabilities in the sr optical lattice clock. *Phys. Rev. Lett.*, 120:063204, Feb 2018.

[114] K. Predehl, G. Grosche, S. M. F. Raupach, S. Droste, O. Terra, J. Alnis, Th. Legero, T. W. Hänsch, Th. Udem, R. Holzwarth, and H. Schnatz. A 920-kilometer optical fiber link for frequency metrology at the 19th decimal place. *Science*, 336(6080):441–444, 2012.

[115] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 1988.

[116] 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.

[117] A. M. Rey, A. V. Gorshkov, and C. Rubbo. Many-body treatment of the collisional frequency shift in fermionic atoms. *Phys. Rev. Lett.*, 103:260402, Dec 2009.

[118] A.M. Rey, A.V. Gorshkov, C.V. Kraus, M.J. Martin, M. Bishof, M.D. Swallows, X. Zhang, C. Benko, J. Ye, N.D. Lemke, and A.D. Ludlow. Probing many-body interactions in an optical lattice clock. *Annals of Physics*, 340(1):311 – 351, 2014.

[119] M. Riedmann, H. Kelkar, T. Wübbena, A. Pape, A. Kulosa, K. Zipfel, D. Fim, S. Rühmann, J. Friebe, W. Ertmer, and E. Rasel. Beating the density limit by continuously loading a dipole trap from millikelvin-hot magnesium atoms. *Phys. Rev. A*, 86:043416, Oct 2012.

[120] Daniele Romanini, Irène Ventrillard, Guillaume Méjean, Jérôme Morville, and Erik Kerstel. *Introduction to Cavity Enhanced Absorption Spectroscopy*, pages 1–60. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.

[121] T. Rosenband, D. B. Hume, P. O. Schmidt, C. W. Chou, A. Brusch, L. Lorini, W. H. Oskay, R. E. Drullinger, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, W. C. Swann, N. R. Newbury, W. M. Itano, D. J. Wineland, and J. C. Bergquist. Frequency ratio of Al^+ and Hg^+ single-ion optical clocks; metrology at the 17th decimal place. *Science*, 319(5871):1808–1812, 2008.

[122] J R Rubbmark and S A Borgström. Rydberg series in strontium found in absorption by selectively laser-excited atoms. *Physica Scripta*, 18(4):196–208, oct 1978.

[123] J. Rutman. Characterization of phase and frequency instabilities in precision frequency sources: Fifteen years of progress. *Proceedings of the IEEE*, 66(9):1048–1075, Sept 1978.

[124] Albert Ryou and Jonathan Simon. Active cancellation of acoustical resonances with an fpga fir filter. *Review of Scientific Instruments*, 88(1):013101, 2017.

[125] M. S. Safronova, D. Budker, D. DeMille, Derek F. Jackson Kimball, A. Derevianko, and Charles W. Clark. Search for new physics with atoms and molecules. *Rev. Mod. Phys.*, 90:025008, Jun 2018.

[126] M. S. Safronova, S. G. Porsev, U. I. Safronova, M. G. Kozlov, and Charles W. Clark. Blackbody-radiation shift in the Sr optical atomic clock. *Phys. Rev. A*, 87:012509, Jan 2013.

[127] J. E. Sansonetti and G. Nave. Wavelengths, transition probabilities, and energy levels for the spectrum of neutral strontium (SrI). *Journal of Physical and Chemical Reference Data*, 39(3):–, 2010.

[128] T. A. Savard, K. M. O'Hara, and J. E. Thomas. Laser-noise-induced heating in far-off resonance optical traps. *Phys. Rev. A*, 56:R1095–R1098, Aug 1997.

[129] Brendan Saxberg, Benjamin Plotkin-Swing, and Subhadeep Gupta. Active stabilization of a diode laser injection lock. *Review of Scientific Instruments*, 87(6):063109, 2016.

[130] S. Schiller, A. Görlitz, A. Nevsky, S. Alighanbari, S. Vasilyev, C. Abou-Jaoudeh, G. Mura, T. Franzen, U. Sterr, S. Falke, C. Lisdat, E. Rasel, A. Kulosa, S. Bize, J. Lodewyck, G. M. Tino, N. Poli, M. Schioppo, K. Bongs, Y. Singh, P. Gill, G. Barwood, Y. Ovchinnikov, J. Stuhler, W. Kaenders, C. Braxmaier, R. Holzwarth, A. Donati, S. Lecomte, D. Calonico, and F. Levi. The space optical clocks project: Development of high-performance transportable and breadboard optical clocks and advanced subsystems. *2012 European Frequency and Time Forum*, pages 412–418, April 2012.

[131] M. Schioppo, R. C. Brown, W. F. McGrew, N. Hinkley, R. J. Fasano, K. Beloy, T. H. Yoon, G. Milani, D. Nicolodi, J. A. Sherman, N. B. Phillips, C. W. Oates, and A. D. Ludlow. Ultrastable optical clock with two cold-atom ensembles. *Nature Photonics*, 11:48 EP –, Nov 2016.

[132] Monika H. Schleier-Smith, Ian D. Leroux, and Vladan Vuletić. States of an ensemble of two-level atoms with reduced quantum uncertainty. *Phys. Rev. Lett.*, 104:073604, Feb 2010.

[133] U. Schunemann, H. Engler, R. Grimm, M. Weidemuller, and M. Zielonkowski. Simple scheme for tunable frequency offset locking of two lasers. *Review of Scientific Instruments*, 70(1):242–243, 1999.

[134] F. Sorrentino, G. Ferrari, N. Poli, R. Drullinger, and G. M. Tino. Laser cooling and trapping of atomic strontium for ultracold atoms physics, high-precision

spectroscopy and quantum sensors. *Modern Physics Letters B*, 20(21):1287–1320, 2006.

[135] Daniel A Steck. Quantum and atom optics. *Oregon Center for Optics and Department of Physics, University of Oregon*, 47, 2007.

[136] Simon Stellmer and Florian Schreck. Reservoir spectroscopy of $5s5p\ ^3P_2$-$5snd\ ^3D_{1,2,3}$ transitions in strontium. *Phys. Rev. A*, 90:022512, Aug 2014.

[137] M. D. Swallows, G. K. Campbell, A. D. Ludlow, M. M. Boyd, J. W. Thomsen, M. J. Martin, S. Blatt, T. L. Nicholson, and J. Ye. Precision measurement of fermionic collisions using an ^{87}Sr optical lattice clock with 1×10^{-16} inaccuracy. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 57(3):574–582, March 2010.

[138] Matthew D. Swallows, Michael Bishof, Yige Lin, Sebastian Blatt, Michael J. Martin, Ana Maria Rey, and Jun Ye. Suppression of collisional shifts in a strongly interacting lattice clock. *Science*, 331(6020):1043–1046, 2011.

[139] A. V. Taichenachev, V. I. Yudin, V. D. Ovsiannikov, V. G. Pal'chikov, and C. W. Oates. Frequency shifts in an optical lattice clock due to magnetic-dipole and electric-quadrupole transitions. *Phys. Rev. Lett.*, 101:193601, Nov 2008.

[140] M. Takamoto, H. Katori, S. I. Marmo, V. D. Ovsiannikov, and V. G. Pal'chikov. Prospects for optical clocks with a blue-detuned lattice. *Phys. Rev. Lett.*, 102:063002, Feb 2009.

[141] Masao Takamoto, Tetsushi Takano, and Hidetoshi Katori. Frequency comparison of optical lattice clocks beyond the Dick limit. *Nature Photonics*, 5:288 EP –, Apr 2011.

[142] H.R. Telle, B. Lipphardt, and J. Stenger. Kerr-lens, mode-locked lasers as transfer oscillators for optical frequency measurements. *Applied Physics B*, 74(1):1–6, Jan 2002.

[143] T. Thiele, J. Deiglmayr, M. Stammeier, J.-A. Agner, H. Schmutz, F. Merkt, and A. Wallraff. Imaging electric fields in the vicinity of cryogenic surfaces using rydberg atoms. *Phys. Rev. A*, 92:063425, Dec 2015.

[144] R. J. Thompson, G. Rempe, and H. J. Kimble. Observation of normal-mode splitting for an atom in an optical cavity. *Phys. Rev. Lett.*, 68:1132–1135, Feb 1992.

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

[146] Michael J. Thorpe, David R. Leibrandt, Tara M. Fortier, and Till Rosenband. Measurement and real-time cancellation of vibration-induced phase noise in a cavity-stabilized laser. *Opt. Express*, 18(18):18744–18751, Aug 2010.

[147] William G. Tobias, Jason S. Rosenberg, Nicholas R. Hutzler, and Kang-Kuen Ni. A low-temperature external cavity diode laser for broad wavelength tuning. *Review of Scientific Instruments*, 87(11):113104, 2016.

[148] Ichiro Ushijima, Masao Takamoto, Manoj Das, Takuya Ohkubo, and Hidetoshi Katori. Cryogenic optical lattice clocks. *Nat Photon*, 9(3):185–189, Mar 2015. Letter.

[149] Ichiro Ushijima, Masao Takamoto, and Hidetoshi Katori. Operational magic intensity for sr optical lattice clocks. *Phys. Rev. Lett.*, 121:263202, Dec 2018.

[150] Jean-Philippe Uzan. The fundamental constants and their variation: observa-

tional and theoretical status. *Rev. Mod. Phys.*, 75:403–455, Apr 2003.

[151] CL Vaillant, MPA Jones, and RM Potvliege. Long-range rydberg–rydberg interactions in calcium, strontium and ytterbium. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(13):135004, 2012.

[152] CL Vaillant, MPA Jones, and RM Potvliege. Multichannel quantum defect theory of strontium bound rydberg states. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 47(15):155001, 2014.

[153] G Vallet, E Bookjans, U Eismann, S Bilicki, R Le Targat, and J Lodewyck. A noise-immune cavity-assisted non-destructive detection for an optical lattice clock in the quantum regime. *New Journal of Physics*, 19(8):083002, 2017.

[154] Ken Van Tilburg, Nathan Leefer, Lykourgos Bougas, and Dmitry Budker. Search for ultralight scalar dark matter with atomic spectroscopy. *Phys. Rev. Lett.*, 115:011802, Jun 2015.

[155] P. Wcislo, P. Morzynski, M. Bober, A. Cygan, D. Lisak, R. Ciurylo, and M. Zawada. Experimental constraint on dark matter detection with optical atomic clocks. *Nature Astronomy*, 1:0009 EP –, Dec 2016. Letter.

[156] J. K. Webb, M. T. Murphy, V. V. Flambaum, V. A. Dzuba, J. D. Barrow, C. W. Churchill, J. X. Prochaska, and A. M. Wolfe. Further evidence for cosmological evolution of the fine structure constant. *Phys. Rev. Lett.*, 87:091301, Aug 2001.

[157] Stephen Webster and Patrick Gill. Force-insensitive optical cavity. *Opt. Lett.*, 36(18):3572–3574, Sep 2011.

[158] John Weiner, Vanderlei S. Bagnato, Sergio Zilio, and Paul S. Julienne. Experiments and theory in cold and ultracold collisions. *Rev. Mod. Phys.*, 71:1–85, Jan 1999.

[159] H. G. C. Werij, Chris H. Greene, C. E. Theodosiou, and Alan Gallagher. Oscillator strengths and radiative branching ratios in atomic Sr. *Phys. Rev. A*, 46:1248–1260, Aug 1992.

[160] P. G. Westergaard, J. Lodewyck, and P. Lemonde. Minimizing the Dick effect in an optical lattice clock. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 57(3):623–628, March 2010.

[161] P. G. Westergaard, J. Lodewyck, L. Lorini, A. Lecallier, E. A. Burt, M. Zawada, J. Millo, and P. Lemonde. Lattice-induced frequency shifts in Sr optical lattice clocks at the 10^{-17} level. *Phys. Rev. Lett.*, 106:210801, May 2011.

[162] Xuejian Wu, Fei Zi, Jordan Dudley, Ryan J. Bilotta, Philip Canoza, and Holger Müller. Multiaxis atom interferometry with a single-diode laser and a pyramidal magneto-optical trap. *Optica*, 4(12):1545–1551, Dec 2017.

[163] Xinye Xu, Thomas H Loftus, John L Hall, Alan Gallagher, and Jun Ye. Cooling and trapping of atomic strontium. *JOSA B*, 20(5):968–976, 2003.

[164] C. Y. Yang, P. Halder, O. Appel, D. Hansen, and A. Hemmerich. Continuous loading of $^1S_0$ calcium atoms into an optical dipole trap. *Phys. Rev. A*, 76:033418, Sep 2007.

[165] Masami Yasuda, Tetsuo Kishimoto, Masao Takamoto, and Hidetoshi Katori. Photoassociation spectroscopy of ^{88}Sr: Reconstruction of the wave function near the last node. *Phys. Rev. A*, 73:011403, Jan 2006.

[166] S. Ye, X. Zhang, T. C. Killian, F. B. Dunning, M. Hiller, S. Yoshida, S. Nagele, and J. Burgdörfer. Production of very-high-n strontium rydberg atoms. *Phys. Rev. A*, 88:043430, Oct 2013.

[167] X. Zhang, M. Bishof, S. L. Bromley, C. V. Kraus, M. S. Safronova, P. Zoller,

A. M. Rey, and J. Ye. Spectroscopic observation of SU(N)-symmetric interactions in Sr orbital magnetism. *Science*, 345(6203):1467–1473, 2014.