

A Compact Cold-Atom Microwave Clock Prototype



Martin Knapp

Department of Physics

University of Oxford

Experimental work undertaken at
the National Physical Laboratory



Submitted for the degree of Doctor of Philosophy

If I have seen further it is by standing on the shoulders of giants

—*Sir Issac Newton*

Acknowledgements

This DPhil has taught me many things, but the main one that stands out is how remarkably little I can do by myself, which is why the previous quote could not be more appropriate to describe my time studying for this DPhil. Therefore, I would like to thank a few *giants* in particular who have helped me over the past four years.

Firstly, my supervisors Patrick Gill and Chris Foot, I would like to thank you for your support, guidance, wisdom, and reassurance. To Mohsin Haji, I feel like none of this would have been possible had you not offered me a job at NPL all those years ago! Thank you for seeing potential in me and always helping me to develop professionally.

I have been very fortunate to work alongside many helpful colleagues at NPL, and to name them all here would perhaps require another thesis in itself! Nevertheless, I would like to thank a few individuals: Hugh Klein and Guilong Huang—for teaching me how to effectively work in the lab and troubleshoot experimental problems; Sean Mulholland—for your guidance and help with all my questions about the DPhil process; Richard Hobson and William Bowden—for looking out for me and keeping me focused in my first year; Krzysztof Szymaniec and Rich Hendricks—for answering all my cold atom questions (of which there were a lot!) and helping troubleshoot issues in the lab (I don't think I would have made anywhere near as much progress without your input); Sam Walby—for the fellow DPhil student camaraderie and useful discussions about cavities; Yuri Ovchinnikov—for lending lab equipment; Peter Tsoulos, Gary Hockley, Pravin Patel, and Conor Robinson—for helping debug my electronics/software problems. Finally, to the rest of the compact clocks team, both past and present, thank you for keeping our daily lunchtimes and occasional tea breaks fun and light-hearted.

Of course, no acknowledgements section would be complete without thanking my family—I would not be who I am today were it not for you.

Finally, to Chloe, thank you for showing me that there's more to life than just atomic clocks and cold atoms. Your never ending support and companionship are something I treasure dearly. I know post-DPhil life marks the start of the next chapter for us, and I can't wait.

Abstract

The transfer of quantum technologies from research laboratories into practical, compact devices brings with it the promise of enhancing a variety of real-world applications. This thesis details the design, construction, and characterisation of a compact cold-atom microwave clock prototype. The prototype consists of laser, control electronics, and physics package subsystems. As part of the physics package, a new design for a microwave cavity is discussed, which should lend itself well to a compact atomic clock. The prototype is intended as a benchtop experiment, where the procedure of interrogating a cloud of caesium atoms during free fall can be explored. Both open- and closed-loop spectroscopy of the clock transition have been performed. The linewidth of the observed Rabi spectrum, which is limited by a short pulse length chosen for diagnostic purposes, is 159(1) Hz. The corresponding short-term fractional frequency stability was measured to be $1.9 \times 10^{-9} \tau^{-1/2}$ out to 1000 s, which is still limited by experimental systematics rather than a fundamental limit of the atomic clock architecture. Through further optimisation and implementation of a Ramsey interrogation sequence with a longer drop time, this frequency stability can be improved to the level of low parts in $10^{-12} \tau^{-1/2}$. Such a clock would be highly suited to applications such as GNSS holdover and radio astronomy. The prototype developed here is a significant step towards the next generation of compact microwave clocks that rely on cold atoms at temperatures of tens of microkelvin rather than thermal atoms.

Publications and Conference Proceedings

1. M. Knapp, S. Walby, M. Haji, C. Foot and P. Gill, “Progress Towards a Compact Cold-Atom Microwave Clock,” 2023 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS), Toyama, Japan, 2023, pp. 1-3, doi: <https://doi.org/10.1109/EFTF/IFCS57587.2023.10272192> - Presented by M.Knapp
2. S. Walby; M. Knapp; J. Whale; A. Wilson; R. Hendricks; C.J. Foot; K. Szymaniec, “Normalised Detection of Clock States by Cold Atom Recapture Method,” 2022 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS), Paris, France, 2022, pp. 1-3, doi: <https://doi.org/10.1109/EFTF/IFCS54560.2022.9850717> - Presented by S.Walby

Conference Posters

1. M. Knapp, M. Haji, C. Foot and P. Gill, “Using Laser Cooled Atoms to Improve Atomic Clock Performance” Quantum Enhanced Sensing Symposium 2023, 13th-15th March 2023, Copenhagen, Denmark. Poster - Presented by M.Knapp
2. M. Knapp, S. Walby, M. Haji, C. Foot and P. Gill, “Progress Towards a Compact Cold-Atom Clock” 7th Post Graduate Institute Conference 2023, 7th-8th June 2023, University of Strathclyde, Glasgow. Talk (3rd Place Prize) - Presented by M.Knapp

Contents

Abstract	ii
List of Publications and Conferences	iii
Contents	iv
1 Introduction	1
1.1 A Brief History of Timekeeping	3
1.2 How Does an Atomic Clock Work?	6
1.3 What Makes a Good Atomic Clock?	8
1.3.1 Allan Variance	9
1.3.2 Short-term Stability	11
1.4 Why Do I Need an Atomic Clock?	14
1.4.1 Global Navigation Satellite Systems (GNSS) Holdover	14
1.4.2 Very Long Baseline Interferometry (VLBI)	15
1.5 Existing Compact Atomic Clocks	17
1.5.1 Vapour Cell Clocks	17
1.5.2 Hydrogen MASER	18
1.5.3 Literature Clocks	20
1.5.4 Clock Comparison	21
1.6 DPhil Goals	24
1.7 Experimental Overview	25
2 Background Theory	27
2.1 Atomic Structure	27
2.2 Atom-Light Interaction	30
2.3 Microwave Spectroscopy	31

2.3.1	Rabi Probability	31
2.3.2	Ramsey Probability	35
2.4	Laser Cooling	37
2.4.1	Radiation Force	38
2.4.2	Slowing an Atomic Beam	40
2.4.3	Optical Molasses	41
2.4.4	Doppler Cooling Limit	43
2.4.5	Sub-Doppler Cooling	43
2.4.6	The Magneto-Optical Trap (MOT)	45
2.5	Summary	48
3	Laser and Optoelectronics System	49
3.1	Laser Requirements	50
3.2	Cooling Light	51
3.3	Repumping Light	54
3.4	Frequency Stabilisation	56
3.5	Compact Laser System	59
4	Physics Package	61
4.1	Vacuum System	61
4.2	Trap Optics	64
4.3	MOT Coils	66
4.4	Compensation and Quantisation Coils	67
4.5	Caesium Source	70
4.6	Microwave Horn	70
4.7	Detection Optics	72
4.8	Compact Physics Package	74
5	Microwave Cavity	75
5.1	Cavity Theory	76
5.2	Cavity-Related Frequency Shifts	78
5.2.1	Cavity Phase Shifts	78
5.2.2	Cavity Pulling	79
5.2.3	Microwave Leakage	80

5.3	Cavity Design	80
5.4	Cavity Seals & Joins	84
5.5	Cavity Simulations	85
5.6	Cavity Summary	89
6	Electronics and Experimental Control	91
6.1	Hardware Overview	91
6.1.1	Microcontroller	94
6.1.2	ADC/DAC Board	95
6.1.3	AOM Control Board	96
6.1.4	Voltage Switch Board	97
6.1.5	Voltage & Current Supplies	99
6.1.6	Local Oscillator & Microwave Synthesis	100
6.2	Software Overview	103
6.2.1	Clock Locking	104
6.2.2	Clock Sequence	106
6.3	Compact Electronics System	109
7	Experimental Results	111
7.1	Cold-Atom Characterisation	111
7.1.1	Atom Number	111
7.1.2	Temperature of the Atom Cloud	113
7.2	Detection Characterisation	118
7.3	Detection Signal-to-Noise	121
7.4	Microwave Spectroscopy	125
7.5	Frequency Stability Analysis	130
7.6	Experimental Results Summary	134
8	Future Work and Conclusions	135
8.1	Future Work	135
8.2	Conclusions	139
	Bibliography	143
	Appendices	155

A	Timeline of the Project: My Contributions	155
B	List of Abbreviations	157

Chapter 1

Introduction

Throughout the years, frequency metrology has had a significant impact in the field of science and technology. Today, the best atomic clocks are capable of keeping time to an accuracy equivalent to losing one second over 13.8 billion years, which is approximately the age of the universe! [1]. However, these types of atomic clocks are the subject of research laboratories. They require racks of equipment, optical tables full of numerous lasers, and a team of scientists to maintain their operation. Ultimately, these clocks are not very portable, and their applications are limited primarily to tests of fundamental physics [2, 3, 4] and defining our nations' time scales (i.e. UTC).

Fortunately, another class of atomic clocks exist that are far more compact. Their small size and portability make them applicable for numerous applications outside of a laboratory environment, such as global navigation satellite systems (GNSS) holdover, telecommunication synchronisation, energy grid synchronisation, operational time scales, and radio astronomy. However, there is a trade-off; portability and compactness often come at the cost of compromised frequency accuracy and stability.

Most existing compact atomic clocks operate using thermal atoms confined in a vapour cell. Inside the vapour cell, there is typically a buffer gas, which serves two purposes. First, it increases the potential interrogation time and thus reduces collisions with the cell walls, a mechanism which leads to a relaxation and broadening of the resonance line. Second, it quenches the fluorescence radiation emitted through alkali atom collisions, which would otherwise cause relaxation in the atomic ensemble. However, the presence of the buffer gas leads to a temperature-dependent frequency shift that can ultimately limit long-term stability [5, 6]. To mitigate this, the atoms can be laser cooled. With laser cooling, atomic temperatures of a few tens of μK can be reached. At these temperatures, the kinetic velocity is a few orders of magnitude slower,

and a buffer gas is no longer needed. Thus, the interrogation time can be drastically increased; this leads to an improvement in short-term stability. Furthermore, the long-term stability is no longer limited by the presence of a buffer gas. The performance gains from laser cooling are attractive. However, laser cooling is typically associated with larger laboratory-scale clocks due to the additional equipment and lasers required. As technology progresses, so too does the demand for smaller atomic clocks with higher accuracy and stability performance. Hence, in recent years there have been several research efforts to develop compact laser-cooled atom microwave clocks [7, 8, 9].

This thesis reports on the apparatus that I have built and tested for the development of a compact cold-atom microwave clock. The experiment is a prototype for a clock architecture in which a cloud of laser-cooled caesium (Cs) atoms is trapped and then dropped through an interrogation region. This type of architecture has previously been explored using rubidium (Rb) [10, 11], but it is yet to be studied using Cs. The procedure of dropping the atoms and interrogating during free fall is perhaps one of the simplest experimental sequences for a clock, and the architecture lends itself well to miniaturisation. Most existing compact cold-atom systems employ a brief release and recapture interrogation method, which sees the atoms dropped only a short distance (a few mm). This limits the achievable interrogation time and potential short-term stability. If the atoms are dropped over a larger distance, say ~ 200 mm, this is enough to provide at least an order of magnitude improvement in the short-term stability without compromising too much on the overall size of the system. In the absence of all technical noise, one would expect the clock size to be limited by the distance the atoms are dropped. The target performance of this clock technology is a short-term stability¹ in the region of low 10^{-12} to mid $10^{-13} \tau^{-1/2}$.

Thesis Structure

Chapter 1 starts the thesis by introducing a few key concepts related to atomic clocks, focusing particularly on the working principle and the types of metrics used to quantify clock performance. With these topics in mind, a literature review is presented, discussing existing compact atomic clocks, both commercial systems and notable lab-based clocks from the literature. A small section then compares these systems to identify advantages and disadvantages before summarising with the approach taken in this work. Chapter 2 reviews some of the fundamental atomic physics

¹See section 1.3.1 for more information on how stability is quoted and its associated unit.

involved with the experiment, of which there are two main aspects: first, the interaction of atoms with microwave radiation, which governs the excitation of the clock transition; second, the interaction of atoms with optical radiation and how this is used for laser cooling and trapping. Chapters 3, 4 and 6 go into more detail about the three key subsystems used in the experimental setup, namely the laser system, the physics package, and the electronic control system. Each subsystem was built during this DPhil from the beginning; therefore, there is some discussion on design, construction, and operating performance. The subsystem chapters are summarised by considering aspects that might be modified to make the subsystem more compact. Chapter 5 discusses the design of a new microwave cavity developed in collaboration with another NPL project. Chapter 7 discusses the main experimental results obtained. This is broadly separated into the characterisation of the cold-atom cloud and measurements of the clock performance. Concluding remarks and an overview of potential ideas for future work are included in Chapter 8.

1.1 A Brief History of Timekeeping

In an age where we can effortlessly retrieve the time from our smartphones or wristwatches, it is easy to take for granted the intricacy and precision of modern timekeeping. But how did humans tell the time before clocks were invented? One of the earliest recorded methods for keeping time was derived from the oscillatory motion of celestial bodies. It takes one day for the Earth to rotate on its polar axis, one month for the moon to orbit around the Earth, and one year for the Earth to orbit around the Sun. Many civilisations measured this passage of time using sundials or shadow clocks. The path traced by the Sun's shadow could be segmented with markers to divide the length of the day. One of the oldest known sundials was discovered in Egypt's Valley of the Kings, dating back to approximately 1500 BC [12]. The main drawback with these clocks is that they need sunshine to work, i.e. they cannot tell the time at night or on an overcast day. For these reasons, humans had to develop artificial timekeeping devices.

Following sundials, the next generation of clocks operated by measuring the constant flow of a substance to denote an elapsed amount of time. In the time leading up to the late Middle Ages (c. 15th century), popular options were sand, water, and candle clocks. However, much progress was made when systems with a natural resonant frequency were used. If the resonant frequency ν_0 is known, then its inverse defines an increment of time $T = 1/\nu_0$. By measuring the number of cycles N of an oscillating system, the elapsed time can then be calculated by

multiplying the cycle number by the time increment. At the turn of the 17th century, Galileo Galilei began studying the motion of pendulums. He discovered that the pendulum exhibited a natural swinging frequency that depended only on the length of the pendulum and not the initial displacement [13]. In 1656, Dutch physicist Christian Huygens used this principle to develop the first pendulum clock [14]. Through a clever arrangement of mechanical gears, the swinging motion of the pendulum would drive the movement of hands on a clock face to tell the time. In addition, some of the energy from the pendulum's swing was fed back to the pendulum to keep it swinging indefinitely. This process is known as escapement. In the following years, variations in the designs of mechanical gears allowed improvements to be made to the stability of these timekeeping devices. However, perhaps one of the most notable improvements was made by John Harrison in 1761. Pendulum clocks were strongly susceptible to environmental effects, such as weather and temperature fluctuations. This posed a problem for seafarers who needed a method for calculating their longitude during lengthy voyages. Pendulum clocks simply were not suited to the harsh environments at sea, not to mention that the swaying motions of the ship would severely affect the pendulum. At the time, this was such a big issue that the UK government offered £20,000² to anyone who could solve it. John Harrison invented the marine chronometer, which was based on a spring-and-balance escapement to solve this very problem [15].

Mechanical clocks continued to improve until the 1930s when they were superseded by quartz crystal oscillators. The first quartz oscillators operated at frequencies of 100 kHz [16], thousands of times faster than their mechanical counterparts. They operate using the piezoelectric effect, in which a small electrical signal causes the quartz to resonate at its natural frequency. These oscillators were far less susceptible to environmental effects and, thus, could keep time for much longer. In 1956, the SI second was based on astronomical time, the fraction of an ephemeris year (1/31 556 925.9747). The apparent consensus was that quartz oscillators could keep better time than the Earth's rotation around the Sun. As our ability to measure the unit of time improved, it became clear that our definition of the second based on astronomical time was starting to weaken.

The next revolution in timekeeping came with the developments in quantum mechanics. There were many breakthroughs in quantum theory throughout the 20th century. However, two key pieces of work led to the invention of the atomic clock. The first was the Nobel Prize-winning work of Isador Rabi [17], who laid the foundations for spectroscopy of molecular and

²Today this would be worth approximately £2,000,000.

atomic beams. The second was a similar spectroscopic technique developed by Norman Ramsey [18] that allowed even finer precision over Rabi's method. The first operational atomic clock was built at the National Physical Laboratory (NPL) in 1955 by Louis Essen and Jack Parry³.

The 'swinging pendulum' of the atomic clock was now related to the 9.192 GHz transition frequency between the ground state splitting of a Cs atom. Essen and Parry's clock used a beam of Cs atoms, whose resonant frequency was measured using the Ramsey spectroscopy technique. In the years that followed, national laboratories around the world started to develop their own Cs beam frequency standards. Eventually, in 1967, at the 13th General Conference on Weights and Measures, a new definition of the SI second was adopted [19]:

“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.”

This was a major revolution for the field of timekeeping since the second was now defined by a fundamental constant of nature rather than a physical artefact. Atoms are inherently indistinguishable, which means that in theory, if two Cs clocks were built, they should 'tick' at the same rate to first order, regardless of where or how they were made. The best Cs clocks today are atomic fountain clocks, an idea first demonstrated in 1989 [20] but fully developed at SYRTE in Paris [21]. These clocks use lasers to cool Cs atoms to a fraction of a degree above absolute zero. The atoms are then launched vertically upwards by approximately a metre until they eventually fall back down under gravity. For a fraction of their trajectory, the hyperfine transition frequency of the atom is measured. The whole process takes ~ 500 ms, but by repeating this, we can measure the Cs hyperfine transition frequency with a relative uncertainty of $\sim 1 \times 10^{-16}$. Today, there are multiple Cs fountain clocks in several National Measurement Institutes (NMIs) around the world, all contributing to Coordinated Universal Time (UTC).

While microwave clocks still have their place in today's society, they are no longer the most accurate clocks we have. There are now clocks whose resonant frequency is based on an optical transition rather than a microwave transition. These are aptly named optical atomic clocks. The optical transitions operate in the THz regime, several tens of thousands of times higher in frequency than their microwave counterparts. This faster rate of 'ticking' allows for superior fractional frequency stability, approaching the 10^{-18} level. Furthermore, the magnitude of frequency shifting phenomena in optical clocks is comparable with that seen in microwave systems,

³Photo from figure 1.1 courtesy of the NPL photo library.




Figure 1.1: Jack Parry (left) and Louis Essen (right) with the first Cs atomic beam clock at NPL in 1955.

but the much higher clock frequency also results in superior fractional frequency accuracy. Optical clocks are complex instruments, and currently only a handful exist. This poses an issue regarding clock comparison to verify the accuracy and reproducibility of their measurements. If these difficulties can be overcome, one might imagine a redefinition of the second based on an optical transition [22].

1.2 How Does an Atomic Clock Work?

A clock is often comprised of three main components: an oscillator, a counter, and a reference. The oscillator generates a ‘ticking’ signal whose output is then counted by the counter to denote an elapsed time. If N is the number of cycles and t_{cycle} is the period of one oscillation from the oscillator, then $t_{\text{elapsed}} = Nt_{\text{cycle}}$. By itself, the oscillator’s output is typically not very stable. To ensure that the oscillator’s ‘ticking’ frequency does not drift over time, a well-defined stable reference is used to discipline the oscillator frequency and keep it fixed to the reference. Once disciplined, the oscillator output should theoretically be a stable ‘ticking’ signal.

Clocks typically fall into one of two categories: active or passive. In an active clock, the oscillator and the reference are merged. The counter directly counts the oscillations of the reference, eliminating the need for a local oscillator. This architecture is used in some clocks, but most systems today are passive. In a passive clock, the oscillator and the reference are separated, and the output of the oscillator must be steered to match the reference.

We can apply these descriptions to the classical example of a grandfather clock (an active system). The oscillator is the swinging motion of the pendulum whose reference frequency is defined by its length and gravitational acceleration. An elaborate gear mechanism known as the

escapement ensures sustained swinging and counts the oscillations of the pendulum to display the time on a clock face. However, a user must periodically adjust the clock to an external reference to provide the correct time. Now, we compare this with an atomic clock (normally a passive system). The oscillator takes the form of electromagnetic radiation, whose oscillation can be counted by an electronic circuit. The reference is the well-defined transition frequency between two discrete atomic energy levels, and an electronic feedback loop is implemented to keep the oscillator fixed to the reference.

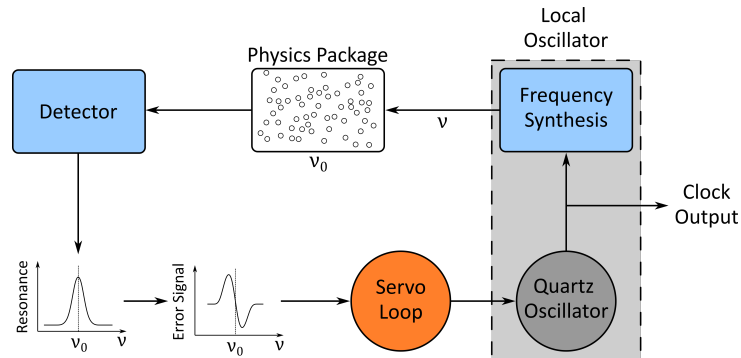


Figure 1.2: Schematic of a passive microwave atomic clock. In this system, a conventional design for the local oscillator is shown where a quartz crystal (~ 10 MHz) is used to reference a frequency synthesiser (~ 9.192 GHz for Cs).

Figure 1.2 shows how a passive microwave atomic clock can be experimentally implemented⁴. At the heart of all microwave clocks is a local oscillator (LO), which generates both the clock output and frequency output necessary for interrogating the atomic reference. In most commercial and laboratory clocks, the LO is based on a quartz oscillator with either a 5 to 10 MHz output. The quartz oscillator is a reference for a frequency synthesiser whose output is a microwave signal that matches the atomic frequency ν_0 . The microwave signal is then directed to the ‘physics package’ housing the atomic reference. This can be an ensemble of neutral atoms, a single trapped ion, or even an ensemble of trapped ions. The response from the interrogated reference is monitored by a detector. The resonance peak observed at the detector can be differentiated to derive an asymmetric error signal, which is used as the input to a servo loop. The output of the frequency synthesiser can be tuned directly or via the quartz oscillator. The servo loop acts on the quartz oscillator to keep the synthesised microwaves fixed to the atomic reference frequency ν_0 . This sequence runs continuously, and the output of the clock is formed from a pickoff of the quartz oscillator. The output now provides a stable 5 to 10 MHz signal, which is referenced to the atomic ensemble in the physics package.

The terms ‘atomic frequency standard’ and ‘atomic clock’ are often used interchangeably.

⁴The layout for an optical clock is slightly different and not discussed here.

For clarity, an atomic frequency standard is a device whose output provides a stable (usually) 10 MHz oscillating signal. An atomic clock consists of a continuously operating atomic frequency standard with a stable output. However, it also includes a counter to determine an elapsed time from a given initialisation time. It does this by counting the oscillations and providing a one pulse-per-second output (1 PPS)⁵. Both devices provide a frequency output, but an atomic clock also provides a time output. It seems over time that the phrase ‘atomic clock’ has become commonplace, most likely because it chimes more with a non-scientific audience. Throughout this thesis, the term ‘clock’ or ‘atomic clock’ is used.

1.3 What Makes a Good Atomic Clock?

There are two key metrics used to determine the performance of a clock: accuracy and stability. To explain both concepts, it is helpful to think of a marksman shooting a target, as seen in figure 1.3. The centre of the target represents the true value of the atomic transition frequency⁶, and each shot (black dot) represents consecutive cycles of the clock attempting to measure this frequency. Accuracy is a measure of proximity to the true value, while stability represents the spread of the measurements over time. Consider the third example, where the output is stable but not accurate. This would be analogous to a clock output with a constant shift from the true value. If the shift is known, it could be corrected by applying an offset. Another example shows a clock output that is accurate but not stable. Such a clock would be considered inaccurate on shorter timescales. However, after long periods of averaging, its mean output would be accurate. The ideal output is accurate and stable, as demonstrated in the first example.

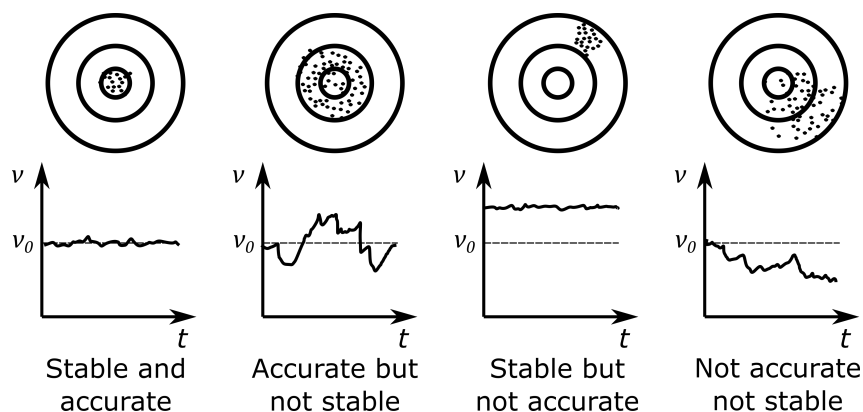


Figure 1.3: Schematic demonstrating the analogy between a marksman shooting a target and measurements of the atomic frequency (adapted from [23]). Accuracy represents how close a measurement is to the true value (target centre). Stability represents the spread of measurements.

⁵The key component of the atomic clock is the counter otherwise a pendulum clock with no clock face might be considered a 1 PPS frequency standard.

⁶The true value here is defined as the frequency in the absence of any shifts or perturbations.

As with all measurements, there are associated uncertainties with measuring the atomic transition frequency. Both systematic and statistical uncertainties will be present, and their quadrature sum denotes the total uncertainty of the measurement⁷. The systematic effects shift the output frequency away from the true value, i.e. they affect the clock accuracy. Accuracy is often quoted as fractional or relative accuracy to compare atomic clocks that use different atomic species and clock transition frequencies. There are many effects that can shift the clock frequency and these will depend on the clock in question. Common examples include magnetic or electric fields that can perturb the energy levels associated with the clock transition. If these shifts are measured, they can be corrected by simply applying an offset at the output. However, the evaluation of these shifts will not be perfect and any offset will be subject to its own uncertainty. What cannot be accounted for, on the other hand, is any random fluctuations in the frequency output. For example, an unstable magnetic field will cause instability in the clock as the energy levels can shift unpredictably. These random fluctuations lead to statistical uncertainties and ultimately affect clock stability.

1.3.1 Allan Variance

In an ‘ideal’ world, the continuous measurement of the reference in an atomic clock would indeed yield an unchanging value. But alas, we do not live in the ‘ideal’ world and must contend with noise. Noise can take on many forms, but ultimately it causes random fluctuations in the measurement process, which can be quantified using statistical uncertainties. Normally, these statistical uncertainties can be minimised by collecting more samples and then taking an average. This is readily seen in the equation for the conventional standard variance⁸, which is inversely proportional to the sample population. The standard variance is based on the calculation of deviations from a fixed average. When measuring frequency stability, this is an issue because the frequency in question is often not stable in time. Imagine trying to take repeated length measurements of a pencil, but the length of the pencil fluctuated over time. How would the average be defined in that scenario? For this reason, the conventional standard variance is unsuitable for frequency stability analysis. Fortunately, there exists the Allan variance [26], which is instead given by:

⁷There are now generally agreed procedures for quoting measurements and their uncertainties in the Guide to the Expression of Uncertainty in Measurement (GUM) [24].

⁸Standard variance, $\sigma^2 = \frac{1}{N-1} \sum_{i=1}^N (y_i - \bar{y})^2$ [25].

$$\text{AVAR}(\tau) = \sigma_y^2(\tau) = \frac{1}{2(M - 2n + 1)} \sum_{i=1}^{M-2n+1} (\bar{y}_{i+n} - \bar{y}_i)^2, \quad (1.1)$$

where σ_y^2 is the Allan variance (AVAR). M represents the number of fractional frequency measurements, evenly spaced in time by τ_0 . The average fractional frequency \bar{y}_i is calculated for different averaging times $\tau = n\tau_0$. Compared to the conventional standard variance, the Allan variance is calculated by summing the differences between averaged blocks of measurements. Figure 1.4 attempts to show this calculation schematically.




Figure 1.4: Schematic representation of how the conventional Allan variance (AVAR) is calculated. It starts with a series of evenly spaced fractional frequency measurements y_i . The average fractional frequency \bar{y}_i is calculated in blocks of $n\tau_0$. In this example $n = 4$. The averaged value from each block is plotted underneath and used in the summation from equation 1.1. This process is repeated for different averaging times τ , by varying n , and the resulting AVAR is plotted against τ .

Equation 1.1 is actually the formula to calculate what is known as the overlapping Allan variance (OAVAR). This is one of a few variations of the conventional Allan variance. In the OAVAR calculation, the averaged blocks of data overlap to be more efficient with the available measurements. The result is a smaller uncertainty on the frequency stability estimate. The quantity often quoted for clocks is the Allan deviation (ADEV), which is the square root of AVAR. ADEV is a function of the averaging time τ , and by varying n , the ADEV can be evaluated for different values of τ . The result is plotted in a sigma-tau plot similar to the one seen in figure 1.5.

The sigma-tau plot informs us of the expected fractional frequency stability for a given amount of averaging time τ . Its characteristic shape is modelled by the combination of several power laws, which are represented by straight lines on a log-log scale. The origin of these power laws can be understood if one considers the power spectral density $S_y(f)$ of the fractional




Figure 1.5: Example ADEV plot along with five power law dependencies.

frequency fluctuations. There are several different noise contributions to the fractional frequency fluctuations, and each noise source exhibits a power law relationship of the form $S_y(f) \propto f^\alpha$. Essentially, each noise source contributes differently for different frequencies, and this also means they contribute to the ADEV at different averaging times ($\sigma_y(\tau) \propto \tau^\mu$). These power law dependencies and their responsible noise sources are summarised in table 1.1. It is important to state that most clock sigma-tau plots will not exhibit all of these power laws. In fact, the ideal scenario is one in which the ADEV is said to continue ‘averaging-down’ with a $\tau^{-1/2}$ dependence. For a truly optimised clock in the absence of all technical noise, the stability is limited by white noise processes arising from the detection of the clock signal. Only after meticulous optimisation and stabilisation of various environmental parameters can this limit be reached.

Noise Source	α	μ
White Phase Modulation & Flicker Phase Modulation	-2	-1
White Frequency Modulation	-1	-0.5
Flicker Frequency Modulation	0	0
Random Walk Frequency Modulation	1	0.5
Linear Frequency Drift	2	1

Table 1.1: Power law dependencies for different noise sources. α represents the exponent of the power law for power spectral density. μ represents the exponent of the power law for averaging time τ . The term modulation is commonly used to denote how the noise modifies the signal.

1.3.2 Short-term Stability

The short-term stability of an atomic clock is often described by the ADEV performance from 1 s out to approximately 1 day [27]. Provided the clock averages down with a $\tau^{-1/2}$ dependence, simply knowing the ADEV at 1 s is enough to infer the expected performance. This means the short-term stability can be thought of as the ‘starting point’ on the sigma-tau plot. Thus,

a lower initial short-term stability is desirable to ensure that the clock can average down to smaller ADEV values⁹. On shorter timescales, the stability is largely dictated by the quality of the clock signal used for locking, and it can be estimated if a few parameters are known. Below is an example of a typical clock signal that might be used for locking a quartz oscillator to an atomic resonance signal.

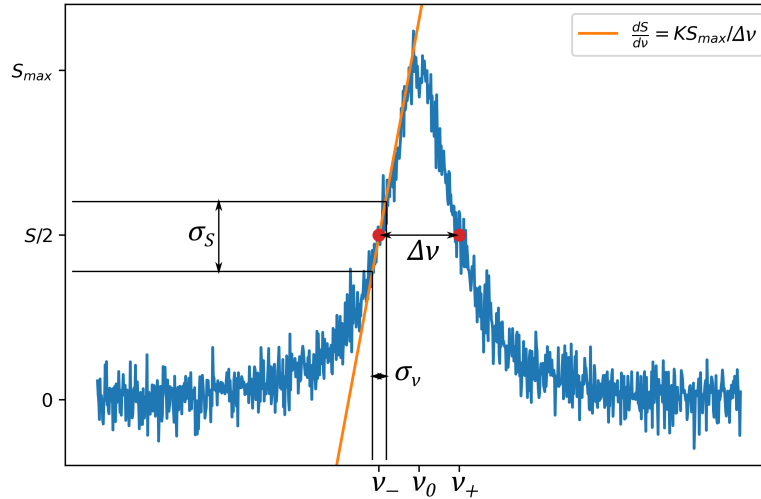


Figure 1.6: Example of a noisy atomic feature that might be used for frequency locking (adapted from [28]). The shape (and ultimately gradient at the lock point) of this feature will depend on the interrogation method.

Frequency locking is usually performed by modulating either side of the peak at the lock points $\nu_{-,+}$, which are typically separated by the linewidth of the feature $\Delta\nu$. Deviation from the centre frequency ν_0 is calculated from the difference between the upper and lower lock points. Noise on the signal σ_S translates to noise σ_ν on the stabilised frequency. The noise is related through $\sigma_S/\sigma_\nu = dS/d\nu|_{\nu_{-,+}}$, and the gradient (in general) at the lock point is given by $dS/d\nu = KS_{\max}/\Delta\nu$. The constant K is typically of order unity and depends on the line shape of the clock signal¹⁰. In the absence of all technical noise, σ_S at the lock point is governed by either photon shot noise or quantum projection noise depending on the clock. Both noise sources are indistinguishable from white frequency modulation, and thus follow the τ averaging law of $\sigma_S(\tau) = \sigma_S(\tau = 1\text{ s}) \tau^{-1/2}$. The predicted fractional short-term stability $\sigma_y = \sigma_\nu/\nu_0$ can then be written as [28]:

$$\sigma_y(\tau) = \frac{1}{K} \frac{1}{Q} \frac{1}{\text{SNR}} \sqrt{\frac{T_c}{\tau}}, \quad (1.2)$$

⁹Semantically, the term ‘lower’ stability sounds slightly counterintuitive; what this normally means is lower instability or smaller Allan deviation.

¹⁰ $K_{\text{Rabi}} \sim 1.516$, $K_{\text{Ramsey}} = \pi/2$ [28].

where $Q = \nu_0/\Delta\nu$ is the quality factor (Q-factor) of the atomic feature, SNR is signal-to-noise ratio defined as $\text{SNR} = S_{\text{max}}/\sigma_S$ ($\tau = 1$ s). The term T_c is also introduced, which is the experimental cycle time, i.e. the time it takes to measure one side of the peak at a lock point. This is required because the atoms are only interrogated for a fraction of T_c .

When designing new atomic clock systems, low short-term instability is the goal, and there are a few aspects that one can consider to achieve this. First, increase the Q-factor by increasing ν_0 or reducing the linewidth $\Delta\nu$. The advantage of optical clocks can be readily seen here, where the clock transition is several hundred THz, nearly five orders of magnitude higher than the microwave region. The linewidth of the feature can be reduced by increasing the interrogation time, provided that the linewidth of the clock signal is Fourier limited by the pulse time. Furthermore, increasing the SNR can also improve the $\tau^{-1/2}$ performance. This is achieved by eliminating all sources of technical noise until the clock is limited by photon shot noise or quantum projection noise (QPN). The photon-shot noise is negligible if large numbers of photons are detected from the atomic ensemble. QPN however, arises from operating on either side of the atomic feature where the transition probability due to clock state excitation is 50%. This means that the ensemble is driven into a superposition of the two clock states, and upon detection, decay to either state is permitted. The resulting state is no longer deterministic but now a random process characterised by probability. When detecting large numbers of atoms, fluctuations in this random process leads to noise in the detected signal known as QPN [29]. Once these limits are reached, the stability can only be improved by interrogating more atoms. Finally, reducing the cycle time T_c will also help improve the short-term stability. These are the three key areas that experimentalists will look towards when designing new systems.

Finally, it is worth noting that equation 1.2 only predicts the short-term stability and the expected performance provided the clock were to continue averaging down indefinitely. In reality, there will most likely be an experimental systematic that prevents the ADEV from averaging down and ultimately determines a limit in the stability. Thus, more thought must be given to consider what might potentially limit the clock in the mid- to long-term stability regions. This includes potential sources of random walk or long-term frequency drift. Random walk can be overcome by stabilising all possible environmental conditions that may influence the clock states. Frequency shifting phenomena can be predicted provided a suitable theoretical model can be constructed. If this is not possible, then a net frequency shift can be measured and calibrated out¹¹ once the clock is operational.

¹¹Provided the net shift is constant in time.

1.4 Why Do I Need an Atomic Clock?

In the previous sections we have seen how an atomic clock works, the next logical question would be “why bother?”. Since their inception, atomic clocks have played a crucial role in the development and innovation of various technologies. In the first instance, laboratory-scale clocks are used to realise the SI second and help to define UTC. Their extraordinary accuracy also makes them well-suited for tests of fundamental physics [2, 3, 4]. The drive to push atomic clock research to ever-increasing accuracy has led to enhanced sensitivity to environmental conditions. Suddenly, an atomic clock is no longer just a clock, but is now a device for sensing environmental changes. This has opened the door to a new field of research called quantum sensing¹² [30]. By making some minor adjustments to the clock’s experimental setup, it can be utilised for sensing magnetic fields, gravity, electric fields, and more.

Outside of the laboratory environment, many of the technologies and services we use on a daily basis would not function without the precise timing and synchronisation provided by atomic clocks. Services such as time scale dissemination, energy grid synchronisation, financial trading, telecommunications, and navigation systems all use atomic clocks in some way. The impact of atomic timekeeping on science and innovation has been significant. To list all applications in detail here would be too lengthy and is beyond the scope of this thesis, although there are a number of sources [28, 31, 32, 33] where a more in-depth discussion can be found. In particular, there are two main applications highlighted in this section which use compact atomic clocks and would also benefit from a higher-performance clock. The compact cold-atom clock developed in this research would be highly suited to these applications.

1.4.1 Global Navigation Satellite Systems (GNSS) Holdover

Global navigation satellite systems (GNSS) is the umbrella term given to multiple satellite constellations used for precise navigation and timing purposes. Within GNSS, there is the US GPS constellation, China’s BeiDou constellation, Europe’s Galileo constellation, and the Russian GLONASS constellation. Anyone with a GNSS receiver can find their position on Earth through a process known as trilateration. The time taken for a signal to be transmitted from a satellite to a receiver can be used to calculate the distance between the two. If three distances from three separate satellites are known, then a user can determine their position on Earth conclusively. However, to avoid the user needing to have their own atomic clock, a fourth satellite is required

¹²Quantum sensing is a broader field, there are sensing experiments that do not rely on atomic clocks.

to correct for any timing error in the user’s clock. This process would not be possible without compact clocks onboard the satellites to maintain precise timing and synchronisation. The main source of error in GNSS positioning currently comes from the satellite ephemerides, a term used to describe the positional data of the satellite, and having a more accurate or stable clock does not necessarily help improve the accuracy with which position can be determined. However, it does increase the resiliency during periods when the satellite cannot receive updates from the GNSS control centres. In this scenario, the onboard clock can act as a flywheel to continue providing GNSS to users who need it.

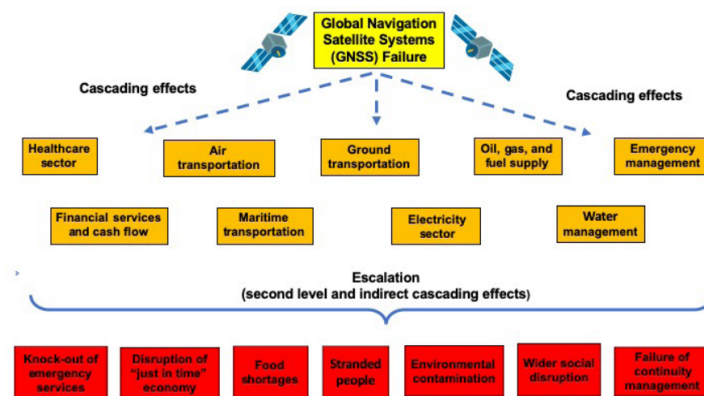


Figure 1.7: Schematic highlighting all of the services that depend on signals from GNSS [34].

In our modern society, GNSS has been dubbed an ‘invisible’ utility [34]. This is because a large proportion of our critical national infrastructure is heavily dependent on GNSS signals (see figure 1.7). If a sustained outage were to occur, accidental or malicious, the results would be catastrophic. Perhaps most worrisome is the concept of GNSS ‘spoofing’, whereby signals are not lost but intentionally altered to provide false positional data for receivers. A 2023 London Economics report [35] estimated the impact on the UK economy to be £7.6 billion for a seven-day disruption. These types of larger-scale disruption can be mitigated by placing compact clocks across strategic locations on Earth, which can serve as holdover references until GNSS is reinstated. Thus, there is a clear desire for clocks with higher accuracy and stability performance, which can provide extended holdover and ultimately increase the security and resiliency of the GNSS system. The clock developed in this work would be an ideal candidate for a ground-based GNSS holdover clock.

1.4.2 Very Long Baseline Interferometry (VLBI)

Another key area where compact atomic clocks are used is in radio astronomy. Similar to how we observe visible light from a star through an optical telescope, radio astronomy observes

radio waves from a star using a radio dish. The primary difference is that the wavelength of radio waves is approximately a million times longer than that of visible light. This means that the angular resolution of a radio telescope is approximately a million times worse than that of an optical telescope. With all telescopes, the minimum angular resolution is ultimately limited by diffraction, which is proportional to the ratio of the signal wavelength and its aperture. Achieving a desirable resolution from a single radio telescope requires an impractically large dish. However, the resolution can be increased by using two radio telescopes separated by a large distance. The effect is one large telescope with an aperture equal to the separation of the two smaller telescopes. But now the incoming signals can arrive at each telescope at different times. To determine the resulting fringe pattern, the incoming signals must be post-processed and cross-correlated. After this, an image of the astronomical object can be obtained. However, to perform the appropriate cross-correlation, the incoming signals must be timestamped. This can be done using a clock situated at each telescope or using one master clock that disseminates timing signals to all telescopes in the array. This type of interferometry is known as VLBI and is the operational foundation of radio telescope arrays. Such an array can be seen in figure 1.8 [36].



Figure 1.8: *Very Large Array (VLA) Radio Telescope Array in New Mexico.*

VLBI measurements have allowed astronomers to precisely map out thousands of radio sources in the night sky. The motion of these sources across the sky is virtually undetectable due to the large cosmic distances. This makes them a perfect inertial reference frame. From this, astronomers have been able to measure Earth's position and orientation very accurately. Similarly, high-precision positional data between each telescope in the array can be measured and used to infer information about tectonic plate movements, which is an area of most interest to geologists.

There are typically two dominant sources of error in VLBI measurements: atmospheric distortion, which leads to path-length delays, and clock instabilities. The relative importance of each depends primarily on the observation frequency and quality of the weather conditions [37]. A higher-performance compact clock would ensure that errors arising from the atomic clock are always minimised. Additionally, for larger arrays where individual telescopes span multiple continents, the timing requirements for the coherent combination of signals become even more stringent. This is another aspect that could be mitigated by using higher-performance compact atomic clocks and is currently being reviewed for the Square-Kilometre Array (SKA) project [38].

1.5 Existing Compact Atomic Clocks

A variety of commercially available compact atomic clocks already exist, and this section intends to introduce a few of them. They all rely on different types of technology, and each has advantages and disadvantages, making them suitable for different applications. Current compact atomic clocks predominantly use microwave technology. Compared to optical clocks, the technology is more mature and much simpler. Although efforts are being made to develop compact optical clocks, the higher frequencies cannot be counted electronically in the same way that microwave frequencies can. Fortunately, the advent of frequency combs [39] has led to the ability to directly count optical signals, which can then be down-converted to the microwave region where regular electronics can take over. However, frequency combs currently take up a large footprint of an optical clock and are not ideal for compact systems. Nevertheless, advances have been made towards microcombs [40], which could serve as a comb source in a compact optical atomic clock. However, these systems are still the subject of research groups and are not available commercially. For this reason, microwave technologies are more attractive when considering compact and portable clocks. The clocks mentioned below are limited to those where the atomic reference uses neutral atoms rather than trapped ion systems.

1.5.1 Vapour Cell Clocks

The smallest and most portable atomic clocks are vapour cell clocks. In their simplest form, a laser is directed through a vapour cell and onto a photodetector. The vapour cell houses the atomic reference as well as a buffer gas. The buffer gas is comprised of an inert gas so that it does not react with the atomic reference. Its primary purpose is to slow the reference atom down and

prevent collisions with the walls of the vapour cell. Without the buffer gas, the kinetic velocity of the reference atoms is too high, and they would very quickly collide with a cell wall and lose coherence. Instead, the atoms can undergo many collisions with the buffer gas before either decohering or reaching a wall. This increases the achievable interrogation time and ultimately increases the Q-factor of the atomic resonance. Unfortunately, the buffer gas also limits these clocks in the long term [5, 6]. The interactions with buffer gas molecules and cell walls are complex processes; on longer timescales, they are very difficult to stabilise and maintain. There are typically two methods of clock state interrogation. The first is known as double-resonance optical pumping (DROP), where a resonant microwave cavity surrounds the vapour cell to excite the clock state transition [41]. The other is coherent population trapping (CPT), where the microwave source directly modulates the laser, thus forgoing the need for a bulky cavity [42]. The simplicity of the vapour cell architecture lends itself extremely well to miniaturisation. Their small size and affordability make them highly applicable to be deployed in applications outside of the lab. For example, these clocks are the main workhorse of telecommunications and satellite navigation systems [43].




Figure 1.9: (a) Example schematic of a CPT vapour cell clock. (b) Commercial chip-scale atomic clock (CSAC) - Microchip CSAC SA.45s [44].

1.5.2 Hydrogen MASER

The hydrogen maser clock was developed in the 1960s [45] and is still one of the most stable microwave clocks available today. Its operating principle is as follows: hydrogen atoms are first ejected from an RF discharge cell as a beam. The atomic beam is directed through a state-selector magnet, which deflects and removes those atoms in the state $F = 0$. The $F = 1$ state atoms are directed into a microwave cavity, which is resonant with the 1.420 GHz clock state transition. Atoms in the upper clock state $|F = 1, m_F = 0\rangle$ decay via stimulated emission down to the lower clock state $|F = 0, m_F = 0\rangle$. The stimulated emission is radiated into the cavity,

resulting in a small amplification of the microwave signal. When the atomic beam continually replenishes the atom number in the cavity, and the losses to the cavity walls are small, there is continuous oscillation, which can be detected by a coupling loop and external circuitry. This configuration is known as an active hydrogen maser because low cavity losses result in continuous oscillation. The active hydrogen maser is typically the size of a conventional washing machine appliance, which is largely determined by the size of the cavity. Moderately smaller masers exist using a dielectric-loaded cavity, although the cavity losses are too great in those cases, and continuous oscillation cannot be sustained. In those configurations, a second coupling loop is used to inject resonant microwaves, and the other is used to detect the microwave signal. This is known as the passive hydrogen maser.




Figure 1.10: (a) Schematic of an active hydrogen maser arrangement. (b) Commercial active hydrogen maser - Microchip Hydrogen Maser (MHM) [46].

The best-performing hydrogen masers can reach a fractional frequency stability in the low parts of 10^{-15} after 10 000 s of averaging [46]. However, this stability can only be achieved with precise stabilisation of environmental conditions. The long-term stability limitation is not fully understood, but there are a few elements of the maser design that could be responsible. The first is the bulk relaxation of the cavity material with temperature, which causes a detuning of the resonant frequency [47]. Other sources [48] suggest that the degradation of the Teflon coating on the walls of the storage cell could be the issue. Nowadays, an ensemble of hydrogen masers and other commercial atomic clocks form the basis for our nations' time scales. Periodic steering of the time scale is performed from measurements using atomic fountain clocks (which define the SI second). These time scales can then be disseminated as a reliable and trusted time source for industry and public use.

1.5.3 Literature Clocks

The majority of commercially available systems use thermal atoms as the atomic reference. The primary reason for this is that thermal atom systems are favoured for their relative simplicity and ability to be easily miniaturised. In the case of vapour cell clocks, the limitations come from the use of the buffer gas. One way to mitigate the effects of the buffer gas is to remove it entirely and use laser cooling techniques instead. Without the buffer gas, the thermal velocity of Cs for example is $\sim 194 \text{ m s}^{-1}$, which would rapidly cause a collision with a cell wall. Once laser cooled to $\sim 10 \text{ }\mu\text{K}$, the thermal velocity is reduced to $\sim 35 \text{ mm s}^{-1}$. In theory, this means that the atoms can be interrogated for longer before decohering. As shown in equation 1.2, a longer interrogation time is desirable to increase the Q-factor of the atomic signal, leading to improved short-term stability. The main drawback with laser cooling is that it often comes at the cost of increased experimental complexity and a larger overall size of the clock. Thus, in recent years there has been a push to develop compact cold-atom systems, which can hopefully provide superior performance over existing thermal clocks. With the exception of [49, 50], most cold-atom systems are not commercially available and are limited to the literature. There are many cold-atom clock arrangements from the literature [51, 52, 53, 54, 55], but they are not all discussed here. The few described below have some key elements that led to the clock arrangement chosen for this DPhil work.

Developments towards compact cold-atom sources have also greatly helped cold-atom clocks. One such example is the grating-MOT (g-MOT) [56]. The g-MOT is formed from a single incident laser beam on a diffraction grating. The diffracted beams generate the necessary arrangement for laser cooling and trapping, much like the conventional six-beam magneto-optical trap (MOT) approach. The g-MOT has been used to develop a compact cold-atom clock [8]. The clock sequence operates using a release and recapture method in which atoms are dropped and interrogated during a brief dark period. Fluorescence is used for state readout, but it also recaptures the atoms to recycle them for the next clock sequence, which ultimately keeps the overall cycle time small ($\sim 100 \text{ ms}$). The short-term stability reported is $2 \times 10^{-11} \tau^{-1/2}$, which at first glance is not dissimilar from some existing thermal systems. The stability is limited by the effective drop time and the effects of the lab temperature conditions.

Other systems include [11], where the atoms are cooled and trapped inside a loop-gap microwave cavity. The structure of the loop-gap cavity is inherently smaller than that of conventional cylindrical resonators and allows optical access to prepare the atoms directly inside the

cavity. Release and recapture is also used here with a Ramsey interrogation time of ~ 22 ms. Preparing the atoms directly inside the cavity allows them to be immediately interrogated once released from the trap, as opposed to waiting for the cloud to fall into an interrogation region. The short-term stability reported is $2.5 \times 10^{-11} \tau^{-1/2}$, which is ultimately limited by the phase noise of the LO.

The final system to note is a compact cold-atom beam clock [7], which leans more towards the vapour cell approach. A beam of cold-atoms is directed through a vapour cell and two CPT excitation regions separated by 4.6 cm. This leads to a 2.3 ms Ramsey dark time and a reported short-term stability of $3 \times 10^{-12} \tau^{-1/2}$. The results shown are preliminary, and the long-term performance is limited by specifics of the experimental implementation, including pressure drifts of the atom source and imperfections in the CPT optical implementation. The most notable feature is the single laser system used to generate the cooling, trapping, and CPT interrogation beams. This type of single-beam arrangement is very attractive for minimising the overall clock size.

1.5.4 Clock Comparison

In this section, the aforementioned clocks are compared. Table 1.2 shows the specifications of several commercially available clocks. These clocks represent only a handful of what is currently available today and were chosen to highlight the different sizes. A more in-depth review can be found in [57]. Similarly, several clocks from the literature can be seen in table 1.3. Given that most clocks from the literature are typically benchtop experiments and do not represent a truly miniaturised product, their sizes do not compare to commercial clocks. However, their ADEV can still be compared and is shown in figure 1.11.

Interestingly, some cold-atom systems show short-term ADEV performance similar to that of existing thermal systems. One possible reason for this is the use of release and recapture, which limits the interrogation time to $\lesssim 25$ ms. Additionally, the CPT interrogation method produces clock signals with only a few percent signal contrast, i.e. a poor SNR. Both factors are likely responsible for the slightly limited short-term stability performance.

Furthermore, the hydrogen maser is an example of a thermal atom clock which does demonstrate very good performance. In fact, its ADEV performance is in keeping with some of the newer commercially available cold-atom systems [49, 50]. However, its main drawback is its overall size and necessity for stringent control over environmental conditions. Any way to reduce the size or sensitivity to environmental fluctuations would be highly desirable.

Ref	Clock	Dimensions [cm]	Size [cm ³]	Weight [kg]	Power [W]	ADEV [$\tau^{-1/2}$]
[44]	SA45.s CSAC	$4.1 \times 3.5 \times 1.1$	16.4	0.035	0.12	3×10^{-10}
[58]	miniRAFS	$10.7 \times 6.75 \times 5.35$	390	0.45	-	1×10^{-11}
[59]	PRS10	$5.1 \times 7.6 \times 10.2$	393.3	0.6	14.4	2×10^{-11}
[60]	5071A	$13.4 \times 42.5 \times 52.4$	29841.8	30	50	1.2×10^{-11}
[50]	c-Rb*	$22.3 \times 38.1 \times 47.6$	40442	31	80	7×10^{-13}
[46]	MHM-H Maser	$107 \times 46 \times 76$	374072	216	75	1.5×10^{-13}
[49]	Mu-Clock*	$155 \times 55 \times 80$	682000	135	200	3×10^{-13}

Table 1.2: Main specifications of some commercially available clock arrangements, in ascending order of ADEV performance. Most commercial systems use thermal atoms, except those denoted by *, which use cold-atoms.

Ref	Clock	Interrogation	Probe Time [ms]	$\Delta\nu$ [Hz]	Cycle Time [ms]	ADEV [$\tau^{-1/2}$]
[55]	MACFS, Cs	Rabi, Cavity	10.3	200	100	2×10^{-11}
[8]	g-MOT, Rb87	Ramsey, CPT	10	50	100	2×10^{-11}
[7]	Vapour Cell Beam, Rb87	Ramsey, CPT	2.3	200	33.3	3×10^{-11}
[11]	Loop Gap Cavity, Rb87	Ramsey, Cavity	22	19.6	540	2.5×10^{-12}
[9]	Atom Chip, Rb87	Ramsey, two-photon	5000	0.1	16000	5.8×10^{-13}

Table 1.3: Main specifications of existing literature clock arrangements, in ascending order of ADEV performance. All of these systems use laser cooled atoms.




Figure 1.11: Various ADEV performances of commercially available (dashed line) and literature clocks (solid line). Squares indicate thermal systems. Circles indicate cold-atom systems.

In summary, rather than using release and recapture to interrogate the atoms, a dropping method would allow for a longer interrogation time and better short-term stability. Figure 1.12 shows the relationship between drop time and drop distance. One would expect that the size limitation of such a system would be linked directly to the distance the atoms are dropped. This allows for a conventional microwave cavity to be used as the interrogation method, avoiding additional lasers that would be needed in a CPT-style clock. Additionally, by implementing a low Q-factor cavity, the sensitivity to temperature fluctuations can potentially be reduced. This type of clock architecture would act as an intermediary between the vapour cell technology and the hydrogen maser technology. It could potentially be the size of a desktop PC but with a performance approaching that of a hydrogen maser at a much-reduced cost.




Figure 1.12: Atom drop time as a function of drop distance. This dependence ultimately dictates the limit for the interrogation time. Red area: represents the typical drop times used in a release and recapture interrogation architecture. Orange area: shows the potential increase in interrogation time if the atoms are dropped over a larger distance (i.e. the proposition of this work).

By selecting some modest estimates for certain aspects of the clock arrangement, we can use equation 1.2 to calculate the potential ADEV performance. The estimates are shown below in table 1.4 and assume that all technical noise has been eliminated and the stability is limited by white noise processes such as QPN. A QPN limited short-term stability is only achievable after meticulous optimisation of the clock. Therefore, a more reasonable target stability would be in the region of low 10^{-12} to mid $10^{-13} \tau^{-1/2}$ (also shown in figure 1.11). This is still competitive with some of the existing cold-atom systems.

K_{Ramsey}	$\pi/2$
T_{Ramsey}	150 ms
$\Delta\nu$	3.33 Hz
Q	2.8×10^9
N_{atoms}	10^6
SNR	1000
T_c	500 ms
σ_y	$1.6 \times 10^{-13} \tau^{-1/2}$

Table 1.4: Estimates of experimental parameters to determine potential ADEV performance for the compact cold-atom microwave clock.

1.6 DPhil Goals

The primary focus of this DPhil research was to develop a system in which the procedure of laser cooling and dropping the atoms could be explored and understood. This focus meant that achieving a final compact clock was not necessarily the target. However, thoughts and considerations were given as to how the system might potentially be made smaller. In fact, upon developing this experiment, characterisation studies were carried out, which should help guide design decisions for a future iteration of this system. Additionally, it was important to be realistic about what can be achieved in a reasonable amount of time for a DPhil, given that I would be the only person actively working on this. The goals of the DPhil can be summarised in a few key points:

1. The system should lend itself well to miniaturisation, i.e. there should be a clear route forward as to how the system could eventually be made smaller.
2. The system should demonstrate performance approaching that of a hydrogen maser, i.e. a short-term stability in the region of low 10^{-12} to mid $10^{-13} \tau^{-1/2}$, and continue to average down in the long term.
3. Developing the system should serve as a basis for understanding all aspects of an atomic clock arrangement.

To achieve these goals, it was necessary to design and build every aspect of the clock from the ground up. This included the laser system for cooling and trapping, a physics package for housing the atomic reference in a vacuum system, a microwave cavity for performing the clock state interrogation, electronic hardware for generating all the necessary signals, and control software for running the experimental sequence. Only once the entire system was built could the

characterisation and optimisation of the system begin. The remainder of this thesis details my progress towards constructing this compact cold-atom microwave clock.

1.7 Experimental Overview

This section provides a brief overview of how the clock arrangement works before delving into more detail in the following chapters. The experimental architecture can be seen in figure 1.13. Cs atoms are first laser-cooled and collected in a retroreflected magneto-optical trap (MOT). After the MOT, there is a brief stage of polarization-gradient cooling (PGC). This reduces the temperature of the atomic cloud to $\sim 11 \mu\text{K}$. The laser light is completely extinguished, with the repumping light being left on momentarily to optically pump all the atoms into the $F = 4$ clock state. The atoms then fall under gravity down a flight tube. During free fall, the atoms pass through an interrogation region where the $|F = 4, m_F = 0\rangle$ to $|F = 3, m_F = 0\rangle$ clock transition is excited with a microwave source applied at $\sim 9.192 \text{ GHz}$. The atoms continue to fall into a detection region where the same vertical beam used for trapping and cooling in the MOT is also used for state detection. A sequence of three laser pulses induces fluorescence from the cloud, and the normalised fraction of atoms which have undergone the clock transition is obtained. An oven-controlled crystal oscillator (OCXO) with 10 MHz output provides the reference for the 9.192 GHz microwave source. If the 10 MHz output drifts, so too will the 9.192 GHz output. This deviation away from the atomic resonance frequency is experienced by the atoms and is reflected in the normalised fraction measured during the detection. When the clock sequence is run continuously, the normalised fraction is used as the input to a servo loop, and feedback is directed to the OCXO to keep its output frequency locked to 10 MHz. A pickoff from the stabilised OCXO output can provide a stable 10 MHz source, which can then be used as a frequency reference.




Figure 1.13: Experimental overview of the compact cold-atom microwave clock prototype. The interrogation should ideally be performed using a microwave cavity, but there is also a microwave horn that can be used for shorter interrogation pulses. $F3, F4$ is shorthand for $F = 3, 4$.

Chapter 2

Background Theory

This chapter outlines the atomic physics governing the interaction of atoms and electromagnetic radiation. This is fundamental for two main aspects of the atomic clock. The first is the coupling between the atoms and the microwave field to excite the clock transition, and the second is the coupling between the atoms and the optical field for laser cooling and trapping. These descriptions of atom-light interaction and magnetic dipole transitions closely follow those in the literature [31, 33]. The laser cooling and magneto-optical trap theory can also be found in the previous references and other sources that provide a more in-depth review [61, 62].

2.1 Atomic Structure

Caesium (Cs) is an alkali metal with a single unpaired valence electron in its outer shell. For alkali metal atoms and alkali-like ions, the lower electronic configurations are filled to form closed shells. Therefore, the resultant total electronic angular momentum J of the atom comes solely from the single valence electron. This is illustrated in figure 2.1. The ground state of Cs has the following electronic configuration¹³:

$$1s^2 2s^2 2p^6 3s^2 3p^6 3d^{10} 4s^2 4p^6 4d^{10} 5s^2 5p^6 6s. \quad (2.1)$$

To be consistent with the general notation for atoms with LS-coupling, the ground state is written as the level $6S_{1/2}$, where the subscript now denotes the J value, and S represents the (total) orbital angular momentum. Relativistic effects lead to a splitting of the atomic energy levels called fine structure. The first excited state configuration, $6P$, has an orbital angular

¹³The letters s, p, and d represent the first three values $L=0, 1, 2$. Historically, this naming convention is associated with the names given by spectroscopists during some of the first observational experiments: sharp, principal, and diffuse.




Figure 2.1: Simplified classical model of an alkali atom. The closed shells formed by the lower ground states result in an effective atomic structure similar to the hydrogen atom. Various angular momenta are also labelled.

momentum $L = 1$. The spin-orbit interaction causes the electron spin angular momentum $S = 1/2$ to give either $J = 1/2$ or $J = 3/2$ for the total electronic angular momentum. This results in the first excited configuration to be split into $6P_{1/2}$ and $6P_{3/2}$. The transitions between these levels and the ground state correspond to wavelengths of 894.6 nm and 852.3 nm and are referred to as the D1 and D2 lines. Laser cooling and trapping are commonly performed on the D2 line.

On a smaller scale than the fine structure is the hyperfine structure, which comes from the interaction between the nuclear magnetic moment $\vec{\mu}_I$ (proportional to its spin I) and the magnetic field created by the electrons (proportional to their total electron angular momentum J). This interaction causes the two momenta to sum to give the total angular momentum $F = I + J$. Cs has a nuclear spin $I = 7/2$. The ground state $6S_{1/2}$ has $J = 1/2$, which splits into two hyperfine levels with $F = 3$ and $F = 4$. The splitting between these hyperfine levels is ~ 9.192 GHz; this is the clock transition. The upper fine-structure level $6P_{3/2}$, with $J = 3/2$, splits into four hyperfine levels with $F = 2, 3, 4, 5$. Both the fine structure and hyperfine structure for the D2 line can be seen in figure 2.2.

Each hyperfine level contains $2F + 1$ degenerate Zeeman states (also called m_F states). The degeneracy can be lifted when an external magnetic field is applied. This is necessary to resolve the $|F = 4, m_F = 0\rangle$ to $|F = 3, m_F = 0\rangle$ clock transition. The Zeeman shift is a crucial consideration for atomic clocks. Fortunately, some transitions are insensitive to magnetic fields to first order, and these are chosen as the clock transition. However, a second-order effect still exists, which can shift the clock states. Consequently, the accuracy and stability of the atomic clock depend on our knowledge of this field and our ability to control it.

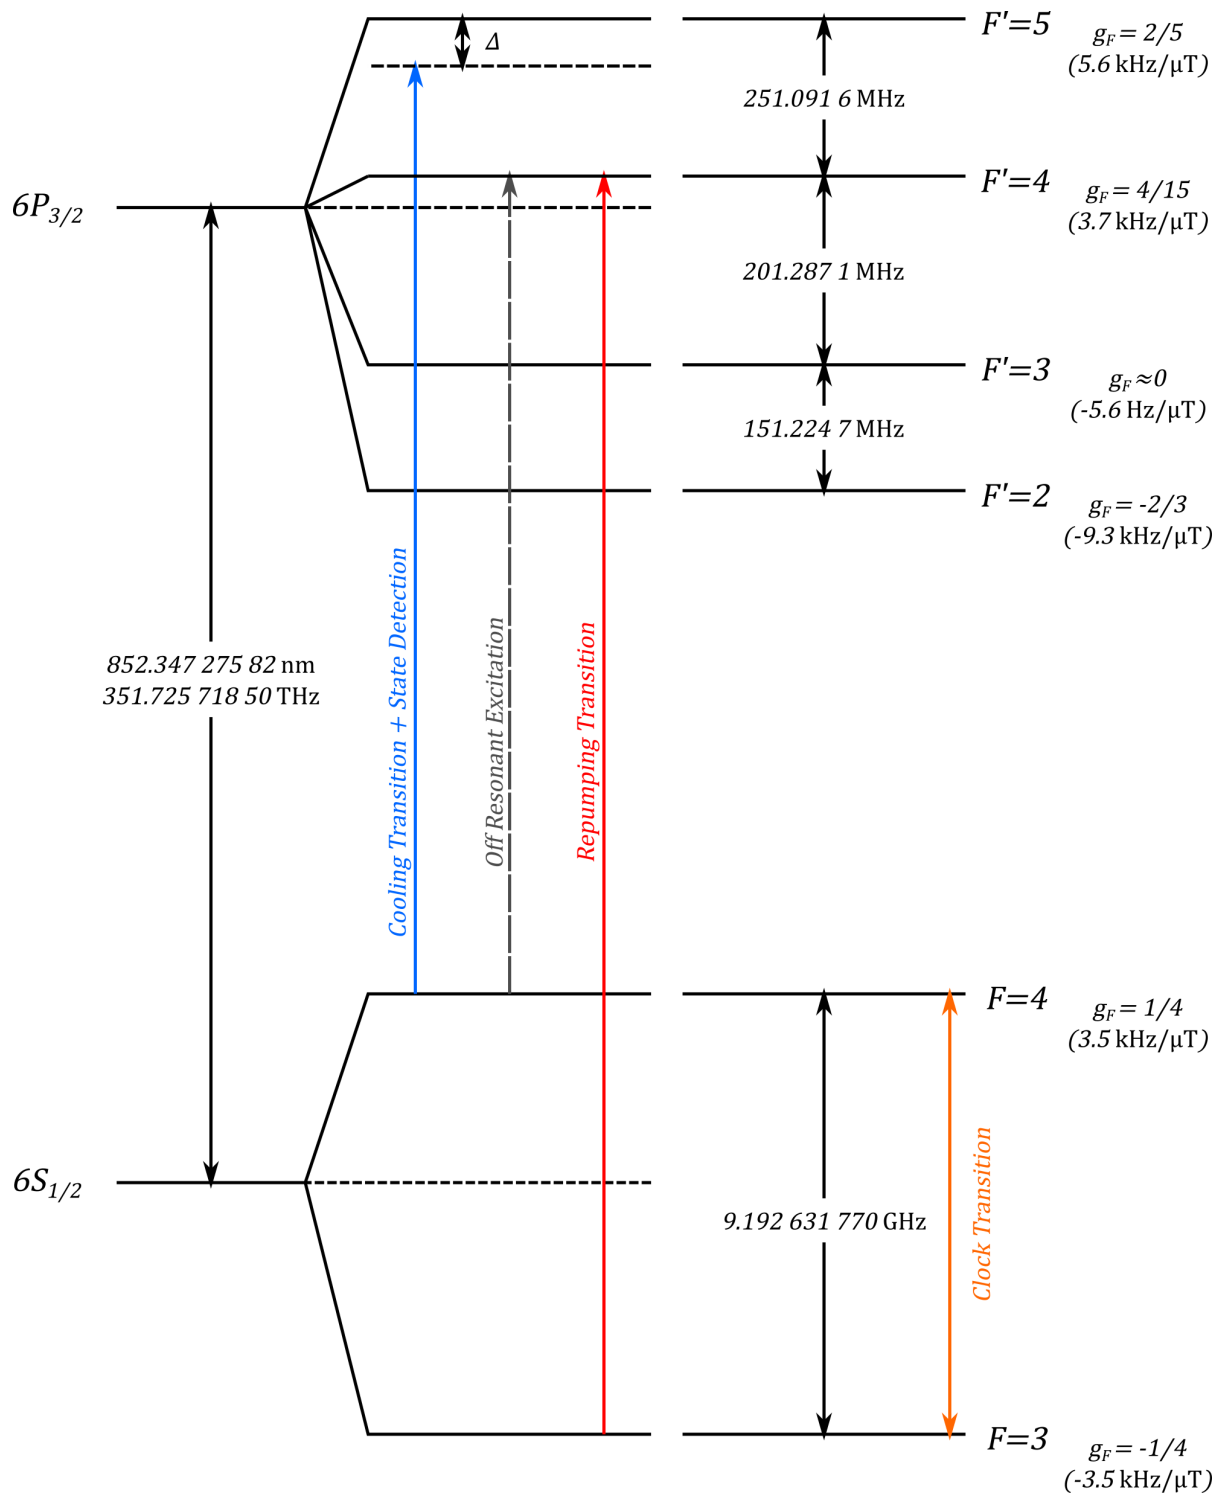


Figure 2.2: Cs D2 transition and hyperfine structure levels. Key transitions for the operation of the compact cold-atom clock are shown. For cooling and MOT operation, the detuning Δ varies from 2Γ to 12Γ . For state detection, $\Delta \sim 0.5\Gamma$. The excited levels are usually denoted as F' to differentiate them from those in the ground state configuration. The energy splittings shown are not to scale. Landé g -factors are given for each state along with the corresponding Zeeman shift (per μT).

2.2 Atom-Light Interaction

The operation of an atomic clock is based on the transition between two discrete energy levels within an atom. Therefore, knowledge of the conditions under which these transitions take place is essential to understand how an atomic clock operates. The Hamiltonian describes the energy of an atom, and it is instructive to separate it into a time-independent \hat{H}_0 and time-dependent component \hat{H}_{int} , which can be expressed as:

$$\hat{\mathcal{H}} = \hat{\mathcal{H}}_0 + \hat{\mathcal{H}}_{\text{int}} . \quad (2.2)$$

The time-independent term encompasses all of the static perturbations discussed in the previous section. The time-dependent term (often called the interaction Hamiltonian) encompasses the coupling between the atom and an external time-varying electromagnetic field. Such a field can be represented by the electric $\vec{E}(\vec{r}, t)$ and magnetic $\vec{B}(\vec{r}, t)$ field components, where \vec{r} is the vector between the nucleus and the electron. The total Hamiltonian becomes:

$$\hat{\mathcal{H}} = \hat{\mathcal{H}}_0 + \underbrace{e\vec{r} \cdot \vec{E}(\vec{r}, t)}_{\text{electric dipole}} - \underbrace{\vec{\mu} \cdot \vec{B}(\vec{r}, t)}_{\text{magnetic dipole}} + (\text{higher order terms}) . \quad (2.3)$$

The second term represents the electric dipole interaction, which is typically only relevant for optical transitions and is more important for laser cooling. The third term represents the interaction of the magnetic dipole moment $\vec{\mu}$ of the atom and the magnetic field of the electromagnetic radiation. The magnetic dipole transitions are most important for microwave atomic clocks and will be the focus for the remainder of the section. Higher-order terms such as the electric and magnetic quadrupole moments can occur, but these are much smaller and are not of interest here. The total atomic magnetic moment $\vec{\mu}$, is composed of three individual terms that determine the coupling to each momentum component of the atom. Hence:

$$\vec{\mu} = \vec{\mu}_L + \vec{\mu}_S + \vec{\mu}_I = -g_L \mu_B \vec{L} - g_S \mu_B \vec{S} + g_I \mu_N \vec{I}, \quad (2.4)$$

where μ_N is the nuclear magneton, and μ_B is the Bohr magneton. The quantities g_L , g_S , and g_I are respectively the nuclear, electronic orbital, and electron spin g-factors. The interaction Hamiltonian can be expressed using the total angular momentum of the atom \vec{F} (see figure 2.1):

$$\hat{\mathcal{H}}_{\text{int}} = g_F \mu_B \vec{F} \cdot \vec{B}(\vec{r}, t), \quad (2.5)$$

where g_F can be written in terms of g_J and g_I .

In the case of hydrogen-like systems and alkali atoms (with a single valence electron), the clock transition resides in the $S_{1/2}$ manifold where $L = 0$ and there is no orbital contribution. Therefore, the $\vec{\mu}_L$ term drops out of equation 2.4. Furthermore, when focusing on the transition between clock states, the system can be reduced to a theoretical treatment of a two-level atom. In this scenario, the nuclear moment $\vec{\mu}_I$ is negligible, which further simplifies the interaction Hamiltonian. In the next section, we shall see how this Hamiltonian can be used to derive the probability of excitation between the clock states when microwaves are applied to an atomic ensemble.

2.3 Microwave Spectroscopy

2.3.1 Rabi Probability

A microwave atomic clock uses a transition associated with the hyperfine splitting. Put simply, the clock operation consists of continuously measuring the fraction of atoms that undergo the clock transition after interacting with a microwave field. Our goal is to derive an equation that gives the probability of excitation during the microwave interaction. Such an equation will inform us of the expected line shape received at the detector after the atoms have been interrogated and detected. This was the same problem faced in early experiments of magnetic resonance [17].

For example, in Cs, the clock transition is often chosen to be the transition $|F = 4, m_F = 0\rangle$ to $|F = 3, m_F = 0\rangle$. A weak static magnetic field must be applied to the atoms to lift the Zeeman degeneracy. The upper clock state splits into nine Zeeman levels, while the lower splits into seven. Given that we only ever drive the clock transition, the problem can be simplified to a two-level system with the clock states labelled as $|e\rangle$ and $|g\rangle$. This also allows us to focus only on the 2×2 submatrix of the complete 16×16 density matrix:

$$\hat{\rho} = \begin{bmatrix} \rho_{e,e} & \rho_{e,g} \\ \rho_{g,e} & \rho_{g,g} \end{bmatrix}. \quad (2.6)$$

The states $|e\rangle$ and $|g\rangle$ have eigenvalues such that $E_e - E_g = \hbar\omega_0$, where ω_0 is the transition frequency. The time-independent Hamiltonian $\hat{\mathcal{H}}_0$ for the two states is then:

$$\hat{\mathcal{H}}_0 = \begin{bmatrix} \hbar\omega_0/2 & 0 \\ 0 & -\hbar\omega_0/2 \end{bmatrix}. \quad (2.7)$$

Here, the value $(E_e + E_g)/2$ is chosen as the origin for all energies. The microwave interaction will cause a coupling between the magnetic field component of the microwaves $B_z(t)$ and the magnetic dipole moment of the atom $\vec{\mu}$. In order to excite the clock transition, the B-field of the microwaves must be aligned parallel to the weak static field [33]. The interaction Hamiltonian takes the form of equation 2.5, although the nuclear moment term can be neglected in this simplified two-level system, and there is no electronic orbital momentum contribution. The off-diagonal matrix elements of the Hamiltonian become:

$$\langle e|\hat{\mathcal{H}}_{\text{int}}|g\rangle = \langle g|\hat{\mathcal{H}}_{\text{int}}|e\rangle^* = \frac{1}{2}\mu_B g_J B_z(t) \approx \mu_B B_z(t) = \hbar V_{e,g}, \quad (2.8)$$

the 1/2 factor cancels as $g_J \approx 2$, and $V_{e,g}$ is used to collect all the terms related to microwaves. Therefore, the total Hamiltonian can be written as:

$$\hat{\mathcal{H}} = \hbar \begin{bmatrix} \omega_0/2 & V_{e,g} \\ V_{g,e} & -\omega_0/2 \end{bmatrix}. \quad (2.9)$$

The elements of the sub-matrix (equation 2.6) can be found by solving the Liouville-von Neumann equation:

$$\frac{d\hat{\rho}}{dt} = -\frac{i}{\hbar}[\hat{\mathcal{H}}, \hat{\rho}]. \quad (2.10)$$

Substituting 2.6 and 2.9 into the Liouville-von Neumann equation yields two differential equations for the coherence $\rho_{e,g}$ and the fractional population difference $\rho_{e,e} - \rho_{g,g}$. Hence:

$$\frac{d}{dt}\rho_{e,g} = -i\omega_0\rho_{e,g} + iV_{e,g}(\rho_{e,e} - \rho_{g,g}), \quad (2.11)$$

$$\frac{d}{dt}(\rho_{e,e} - \rho_{g,g}) = 2i(V_{g,e}\rho_{e,g} - V_{e,g}\rho_{g,e}). \quad (2.12)$$

Equations 2.11 and 2.12 govern the dynamics of this two-level system and how it evolves with time during the microwave interaction. We work on the assumption that the microwave field has a sinusoidal time dependence of the form $B_z(t) = B \cos(\omega t + \phi)$. The quantity $V_{e,g}$ can then be written as:

$$V_{e,g} = b \cos(\omega t + \phi), \quad (2.13)$$

where b is the Rabi frequency in angular units and can be expressed as:

$$b = \mu_B B / \hbar. \quad (2.14)$$

It is beneficial to rewrite $V_{e,g}$ in terms $b_1 = b \cos(\phi)$ and $b_2 = -b \sin(\phi)$ as shown in equation 2.15. When making the rotating-wave approximation, the complex conjugate term (cc) is ignored¹⁴. Hence:

$$V_{e,g} = \frac{1}{2}(b_1 + ib_2)e^{-i\omega t} + \text{c.c.} \quad (2.15)$$

The state of the atoms depends on time t as well as time θ in the interaction region. Therefore, the total derivative can be written as:

$$\frac{d}{dt} = \frac{\partial}{\partial t} + \frac{\partial}{\partial \theta} \frac{d\theta}{dt}. \quad (2.16)$$

We can search for solutions to equations 2.11 and 2.12 using the forms:

$$\rho_{e,g} = \frac{1}{2}[a_1(\theta) + ia_2(\theta)]e^{-i\omega t}, \quad (2.17)$$

$$\rho_{e,e} - \rho_{g,g} = a_3(\theta). \quad (2.18)$$

Substituting these into equations 2.11 and 2.12 gives a set of three differential equations:

$$\frac{\partial}{\partial \theta} a_1(\theta) + \delta a_2(\theta) + b_2 a_3(\theta) = 0, \quad (2.19)$$

$$-\delta a_1(\theta) + \frac{\partial}{\partial \theta} a_2(\theta) - b_1 a_3(\theta) = 0, \quad (2.20)$$

$$-b_2 a_1(\theta) - b_1 a_2(\theta) + \frac{\partial}{\partial \theta} a_3(\theta) = 0. \quad (2.21)$$

These equations can be solved using matrix formalism and by taking the Laplace transformation. The solution takes the form:

¹⁴This is because the microwave frequency $\omega \sim \omega_0$. Terms in the Hamiltonian that oscillate with frequency $\omega + \omega_0 \sim 2\omega_0$ are neglected, while terms with frequencies $\omega - \omega_0$ are kept.

$$\begin{bmatrix} a_1(\theta) \\ a_2(\theta) \\ a_3(\theta) \end{bmatrix} = R(b_1, b_2, \delta, \theta) \begin{bmatrix} a_1(0) \\ a_2(0) \\ a_3(0) \end{bmatrix}, \quad (2.22)$$

where the 3×3 matrix $R(b_1, b_2, \delta, \theta) =$

$$\begin{bmatrix} C(\theta) + \frac{b_1^2}{\tilde{\Omega}^2} (1 - C(\theta)) & -\frac{\delta}{\tilde{\Omega}} S(\theta) + \frac{b_1 b_2}{\tilde{\Omega}^2} (1 - C(\theta)) & -\frac{b_1 \delta}{\tilde{\Omega}^2} (1 - C(\theta)) - \frac{b_2}{\tilde{\Omega}} S(\theta) \\ \frac{\delta}{\tilde{\Omega}} S(\theta) + \frac{b_1 b_2}{\tilde{\Omega}^2} (1 - C(\theta)) & C(\theta) + \frac{b_2^2}{\tilde{\Omega}^2} (1 - C(\theta)) & \frac{b_1}{\tilde{\Omega}} S(\theta) - \frac{b_2 \delta}{\tilde{\Omega}^2} (1 - C(\theta)) \\ -\frac{b_1 \delta}{\tilde{\Omega}^2} (1 - C(\theta)) + \frac{b_2}{\tilde{\Omega}} S(\theta) & -\frac{b_1}{\tilde{\Omega}} S(\theta) - \frac{b_2 \delta}{\tilde{\Omega}^2} (1 - C(\theta)) & 1 - \frac{b^2}{\tilde{\Omega}^2} (1 - C(\theta)) \end{bmatrix}, \quad (2.23)$$

and $C(\theta) = \cos \tilde{\Omega} \theta$, $S(\theta) = \sin \tilde{\Omega} \theta$. There are two new variables: δ the detuning from resonance, and $\tilde{\Omega}$ the effective Rabi frequency:

$$\delta = \omega - \omega_0, \quad (2.24)$$

$$\tilde{\Omega}^2 = b^2 + \delta^2. \quad (2.25)$$

If $\rho_{e,e}(0)$ and $\rho_{g,g}(0)$ represent the fractional population levels at the beginning of the interaction. Then they evolve as:

$$\rho_{e,e}(\theta) = \rho_{e,e}(0)[1 - P(\theta)] + \rho_{g,g}(0)P(\theta), \quad (2.26)$$

$$\rho_{g,g}(\theta) = \rho_{g,g}(0)[1 - P(\theta)] + \rho_{e,e}(0)P(\theta). \quad (2.27)$$

Using the above two equations, as well as 2.18, and solving for $P(\theta)$ we find:

$$P(\theta) = \frac{1}{2} \left(1 - \frac{a_3(\theta)}{a_3(0)} \right), \quad (2.28)$$

where the variable $a_3(\theta)$ is given in equation 2.23. It is often the case that the atoms start the interaction without coherence, i.e. $a_1(0) = a_2(0) = 0$. Thus, we finally arrive at the probability

of excitation¹⁵:

$$P(\tau) = \frac{b^2}{\tilde{\Omega}^2} \sin^2 \left(\frac{\tilde{\Omega}\tau}{2} \right), \quad (2.29)$$

where τ denotes the duration of the microwave interaction. The probability reaches its maximum when $b\tau = \pi$. This is known as the π -pulse condition and corresponds to a full population transfer. The pulse length τ and the microwave power (which is related to the field strength B and thus Rabi frequency b) are the two parameters that can be varied by the experimentalist. By appropriately tuning these two values, it is possible to drive any fractional or multiple integer number of π -pulses. In figure 2.3, the Rabi line shape is plotted as a function of Δ the microwave detuning¹⁶. Its full width at half maximum (FWHM), W , is given by $W_{\text{Rabi}} \approx 0.7987/\tau$. Thus, a longer probe time is beneficial to narrow the central feature used for locking¹⁷.




Figure 2.3: Rabi excitation probability as a function of microwave detuning. The line shape varies according to the length of the pulse applied, making the lobe features more pronounced at longer pulses.

2.3.2 Ramsey Probability

Instead of interrogating atoms with a long pulse, it is better for atomic clocks to use two pulses separated in time, as was demonstrated by Norman Ramsey in 1950 [18]. In each interaction region, there is a Rabi-style excitation driving a $\pi/2$ -pulse, and the interrogations are separated by a free evolution time T where they experience no microwaves. If the phase accumulated in the time T is an integral multiple of 2π then the resulting transfer is the same as if the two

¹⁵The Rabi excitation probability is typically quoted in terms of $\sin^2 x$. The following trigonometric identity $\sin^2 x = \frac{1 - \cos 2x}{2}$ is used to convert the $\cos x$ in R to match that seen in the probability.

¹⁶Where Δ is the detuning in units of Hz rather than angular frequency rads^{-1} , i.e. $\Delta = \delta/2\pi$.

¹⁷While maintaining the field amplitude required to satisfy $b\tau = \pi$.

Rabi pulses are not separated, effectively giving a π -pulse. This method is known as Ramsey interrogation. The derivation of the excitation probability follows a similar (and slightly more involved) approach to that for a Rabi pulse, seen above, and gives the probability of transfer as:

$$P(\tau) = \frac{4b^2}{\tilde{\Omega}^2} \sin^2 \frac{1}{2} \tilde{\Omega} \tau \left[\cos \frac{1}{2} \tilde{\Omega} \tau \cos \frac{1}{2} (\delta T + \phi) - \frac{\delta}{\tilde{\Omega}} \sin \frac{1}{2} \tilde{\Omega} \tau \sin \frac{1}{2} (\delta T + \phi) \right]^2. \quad (2.30)$$

The Ramsey line shape is plotted in figure 2.4 as a function of the microwave detuning Δ . The FWHM, W , of the central Ramsey fringe is $W_{\text{Ramsey}} = 1/2T$ ¹⁸. Compared to a single pulse (the Rabi method), Ramsey interrogation offers a 0.63 reduction factor of the linewidth. This is advantageous for atomic clocks, as shown in equation 1.2, where the short-term stability benefits from a narrower atomic feature. Additionally, noise or disturbances during a long microwave pulse can cause a line-broadening, whereas the Ramsey signal remains narrow, although there can be a frequency shift from the total phase accumulated (average over the fluctuations in the long dark period). For these two reasons, Ramsey interrogation is favoured for atomic clocks and is used to great effect in an atomic fountain clock. Atoms are launched vertically upwards through a microwave cavity so that interrogation occurs once on the way up and again on the way down. The free evolution time, determined by the atoms' trajectory above the cavity, is $T \sim 500$ ms corresponding to a linewidth of ~ 1 Hz.




Figure 2.4: Ramsey excitation probability as a function of microwave detuning. $b\tau = \pi/2$ and $T/\tau = 10$. The inset shows a zoom on the central Ramsey fringe. In the immediate vicinity of the resonance, where $\delta \ll b$, $b \approx \tilde{\Omega}$ the probability can be approximated as $P(\tau) \approx \frac{1}{2} \sin^2 b\tau [1 + \cos \delta T]$.

¹⁸In units of Hz rather than angular units rad s^{-1} .

2.4 Laser Cooling

Since its advent in the 1970s, laser cooling has dramatically improved the development of atomic clocks. It has revolutionised the implementation of primary frequency standards and paved the way for the development of optical frequency standards. It has been shown in equation 1.2 that the stability of an atomic clock is directly proportional to the linewidth of the atomic feature used for locking. Therefore, it is desirable to reduce this linewidth as much as possible. One primary source of line-broadening stems from the motion of the atomic ensemble. At least in the case of an atomic gas, there will be a spread of atomic velocities governed by the Maxwell-Boltzmann distribution. This velocity spread leads to a first-order Doppler broadening of the atomic feature and a second-order Doppler shift. Methods to try to reduce the motion of the atomic ensemble have included using a buffer gas and storage cell to confine the atomic ensemble. However, these solutions have their limitations, as mentioned in earlier sections. As the development of laser technologies improved, it was proposed to use the atom-light interaction to reduce the thermal motion of atoms [63].

The laser cooling process relies on the scattering of photons from a laser beam. An incoming photon of energy $\hbar\omega$ carries momentum $\hbar k$. When the atom absorbs the photon, it gains energy $\hbar\omega$ by transitioning to an excited state. From the conservation of momentum, the atom receives a momentum kick $\hbar k$ away from the light source, thus slowing the atom down. When the atom decays back to the ground state, a photon is spontaneously emitted, carrying with it $\hbar\omega$ of energy. The direction of the spontaneously emitted photon is random, and thus, over many emission events, the average recoil effect is zero. Therefore, after multiple cycles of absorption and emission, the atom is eventually slowed down by the absorption process. We can estimate the size of this effect by considering the momentum change during recoil:

$$p_{\text{kick}} = \hbar k = Mv_{\text{recoil}}. \quad (2.31)$$

A Cs atom with velocity $v = 200 \text{ m s}^{-1}$ has momentum $Mv \sim 10^{-23} \text{ kg m s}^{-1}$. At first glance, this number pales in comparison to the typical momentum carried by a single photon $\hbar k \sim 10^{-29} \text{ kg m s}^{-1}$ (with a wavelength of 852 nm). The recoil velocity for Cs is 3.5 mm s^{-1} , which means for an initial velocity of 200 m s^{-1} , it requires $\sim 57\,000$ photons to stop one atom. The spontaneous lifetime of the excited state is tens of ns, which means that this many photons can be scattered in $\sim 1 \text{ ms}$, and the laser cooling process can be efficient with a laser beam of

modest power (a few mW as we shall see in the next section). The key is to ensure that many scattering events take place to slow the atom down.

2.4.1 Radiation Force

As mentioned in the previous section, the radiation force experienced by an atom originates from momentum transfer between the optical field and the atom. This is facilitated by the absorption of photons followed by spontaneous emission. The average force can be written as:

$$F_{\text{scatt}} = \hbar k \Gamma \rho_{2,2}, \quad (2.32)$$

where $\hbar k$ is the momentum transferred from a single photon, and $R_{\text{scatt}} = \Gamma \rho_{2,2}$ is the scattering rate for a single photon, which is equal to the excited state decay rate Γ , and the fractional population of the excited level $\rho_{2,2}$. To find the elements of $\hat{\rho}$, we use the following equation:

$$\frac{d\hat{\rho}}{dt} = -\frac{i}{\hbar} [\hat{\mathcal{H}}, \hat{\rho}] - \Gamma \begin{bmatrix} \rho_{2,2} & \frac{1}{2}\rho_{2,1} \\ \frac{1}{2}\rho_{1,2} & -\rho_{1,1} \end{bmatrix}. \quad (2.33)$$

Equation 2.33 is recognised as the Lindblad master equation and is a modified form of the Liouville-von Neumann equation (2.10) with an additional term Γ to account for the relaxation effects of the excited state, i.e. spontaneous emission. This extra term is important for the radiation force because the lifetime of the excited state limits the time-averaged rate of momentum transfer that results from a process of stimulated absorption followed by spontaneous emission. Before finding the solutions to equation 2.33, let us first make a few assumptions. We assume the atom is a two-level isolated system consisting of a ground state $|1\rangle$ and an excited state $|2\rangle$. The atom is at rest and exposed to a light field in the form of a single laser beam. The travelling wave of the light field along the z axis takes the form:

$$\vec{E}(z, t) = E_0 \cos(\omega t + kz) \vec{e}_\lambda. \quad (2.34)$$

The off-diagonal matrix elements of the total Hamiltonian become:

$$\mathcal{H}_{2,1} = \langle 2 | \hat{\mathcal{H}}_{\text{int}} | 1 \rangle = \langle 1 | \hat{\mathcal{H}}_{\text{int}} | 2 \rangle^* = \langle 2 | e \vec{r} \cdot \vec{E}(z, t) | 1 \rangle = \hbar \Omega \cos(\omega t + kz), \quad (2.35)$$

where the variable Ω is the optical Rabi frequency defined as:

$$\Omega = \left(\frac{E_0}{\hbar} \right) \langle 1 | e^{i\vec{r}} \cdot \vec{\epsilon}_\lambda | 2 \rangle. \quad (2.36)$$

With this definition in hand, the elements of $\hat{\rho}$ can be found using equation 2.33, the steady-state solutions take the form:

$$\rho_{2,1} = \rho_{1,2}^* = \frac{i\Omega}{2(\Gamma/2 - i\delta)(1 + s)}, \quad (2.37)$$

$$\rho_{2,2} = 1 - \rho_{1,1} = \frac{s_0/2}{1 + s_0 + 4(\delta/\Gamma)^2} = \frac{s}{2(1 + s)}, \quad (2.38)$$

where s is called the saturation parameter and is defined as:

$$s = \frac{|\Omega|^2/2}{\delta^2 + \Gamma^2/4} = \frac{s_0}{1 + 4(\delta/\Gamma)^2}, \quad (2.39)$$

and the on-resonance saturation parameter s_0 (i.e. s for $\delta = 0$) is given by:

$$s_0 = \frac{2|\Omega|^2}{\Gamma^2} = \frac{I}{I_{\text{sat}}}, \quad (2.40)$$

I_{sat} is the saturation intensity given by:

$$I_{\text{sat}} = \frac{\pi \hbar c \Gamma}{3 \lambda^3}. \quad (2.41)$$

Substitution into equation 2.32 gives the radiation force as:

$$F_{\text{scatt}} = \hbar k R_{\text{scatt}} = \hbar k \frac{\Gamma}{2} \frac{s_0}{1 + s_0 + 4(\delta/\Gamma)^2}. \quad (2.42)$$

For Cs, where $\Gamma = 2\pi \times 5.234$ MHz and $\lambda = 852.347$ nm, we find $I_{\text{sat}} = 1.1049$ mW cm⁻² [64]. From this, we can determine the ‘modest’ laser power required to bring an atom to rest more explicitly. If the atom must scatter $N_{\text{scatt}} \sim 57000$ photons to stop, and the excited state lifetime is $\tau = 1/\Gamma$, then the scattering rate required will be $R_{\text{scatt}} \approx N_{\text{scatt}}\Gamma$. By equating this estimate of R_{scatt} with equation 2.42, the intensity and thus power can be inferred. For a resonant laser beam, with beam radius $w_0 = 1$ cm, the required optical power to stop an atom is $P \sim 1.7$ mW, which is more than manageable from the output of a commercial laser diode.

In figure 2.5, the radiation force is plotted against the laser detuning δ . It reaches a maximum when $\delta = 0$. The force increases towards larger intensities, as shown by plotting for different

values of s_0 . As I tends to infinity, the force tends to a maximum value $F_{\max} = \hbar k \Gamma / 2$. This is to be expected as the population levels in both states tend towards $1/2$ as the transition saturates. Above saturation intensity, i.e. $s_0 > 1$, the profile experiences a power broadening. This is a direct consequence of the fact that absorption increases with increasing intensity, and this has a larger effect on atoms in the wings of the profile compared to the centre, where the atoms are already strongly saturated.




Figure 2.5: Scattering rate vs. detuning for various values of on-resonance saturation parameter s_0 . For $s_0 > 1$ a power broadening of the spectral profile starts to take place.

A more in-depth treatment as in [61, 62] reveals a second type of radiation force. The above derivation focuses on the scattering force, but there is also a so-called dipole force. It arises from the light shift of the ground and excited states, which depends on the strength of the optical field. These radiation forces are fundamentally different in nature. Spontaneous emission is irreversible, which makes the scattering force dissipative and ideal for cooling. On the other hand, the dipole force is conservative, which makes it suitable for trapping. The dipole force has much more significance in optical clocks, where neutral atoms are trapped in the potential wells of an optical lattice. The remaining parts of this section focus primarily on the scattering force and how it can be applied to various situations for laser cooling and trapping.

2.4.2 Slowing an Atomic Beam

The first experiments for laser cooling exploited the Doppler shift of moving atoms to make the scattering force velocity-dependent. This velocity-dependent force was then used for the deceleration of atoms. In its simplest form, laser cooling uses a single laser beam to cool an atomic beam effusing from an oven source. As shown in figure 2.6, the atoms leave the oven and

immediately encounter a laser beam shining in the opposite direction to their motion. Atoms travelling with velocity v will experience a Doppler-shifted laser frequency. If this shifts light to be on resonance, then the cooling process outlined at the start of this section will occur. The key here is to ensure that the laser frequency remains Doppler shifted into resonance as the atoms slow down. There are two notable methods that experimentalists employ to maintain resonance. The first is chirp cooling [65, 66, 67], where the laser frequency is swept (in time) to maintain resonance. The second uses the Zeeman effect [68, 69]. A Zeeman-slower consists of a tapered solenoid to generate a magnetic field that varies along the trajectory of the atoms in a way that consistently shifts the resonant frequency using the Zeeman effect.




Figure 2.6: Atoms effuse from an oven travelling with velocity v towards a counter-propagating laser beam (adapted from [62]). The process of photon absorption imparts a force F_{scatt} on the atoms opposing their motion. The direction of the spontaneously emitted photons is random, and thus, after many scattering events, their net force is zero.

2.4.3 Optical Molasses

The next step from cooling an atomic beam is to cool a gas, where the atoms are free to move in all three dimensions. Here, the concept of laser cooling is extended to include the effects of multiple laser beams all shining in different directions. For explanatory purposes, we consider just two counter-propagating beams in one dimension. An atom at rest will exhibit no Doppler shift, and thus, the net scattering is equal from each of the two beams. However, atoms moving along the axis of the beams will experience a force proportional to their velocity due to the Doppler shift. The key is that now both laser beams are slightly red-detuned from the atomic resonance. Therefore, the frequency of the beam opposing the atom's motion is shifted into resonance. The atom scatters more photons from this beam compared to the beam travelling with the atom's motion. The reverse will be true for atoms travelling in the opposite direction. This can be extrapolated to all dimensions using three intersecting orthogonal pairs of beams. Atoms entering the intersecting region experience a viscous damping in their velocity, hence the

name ‘optical molasses’.

The force experienced by the atoms in the optical molasses follows straightforwardly from equation 2.42. We start with the one-dimensional case of an atom moving in between two counter-propagating beams. We assume that the force on the atom is the sum of the forces from each beam $F_{\text{om}} = F_+ + F_-$, where:

$$F_{\pm} = \frac{\hbar k \Gamma}{2} \frac{s_0}{1 + s_0 + 4[\delta_{\pm}/\Gamma]^2}, \quad (2.43)$$

and $\delta_{\pm} = \delta \pm kv_z$ is the detuning for each laser beam containing the Doppler shift $\omega_D = -kv_z$.

By making a first-order expansion for low velocities, the sum can be written as:

$$F_{\text{lin}} = \frac{8\hbar k^2 \delta}{\Gamma} \frac{s_0}{[1 + s_0 + 4(\delta/\Gamma)^2]^2} v_z \equiv -\alpha v_z, \quad (2.44)$$

where the value α can be thought of as a damping coefficient. This force is plotted against velocity in figure 2.7, and we can see in the low-velocity limit that the force exhibits a linear dependence. It should be noted that the optical molasses is not a trap, and atoms that enter the molasses region will eventually diffuse out. The first experimental demonstration of optical molasses was in 1985 when three-dimensional laser cooling was performed on sodium atoms [70].




Figure 2.7: Force profile in an optical molasses as a function of velocity. The force from both counter-propagating beams is shown, as well as their sum, and the force in the low-velocity limit. For $s_0 = 1$, and $\delta = -\Gamma/2$.

2.4.4 Doppler Cooling Limit

The laser cooling process relies on the absorption and emission cycle of many photons. However, fluctuations in these two processes lead to a fundamental temperature limit. The randomly directed nature of the spontaneously emitted photons causes a random walk of the velocity in all directions. There is another random process that contributes to the heating, namely fluctuations in the number of scattered photons per unit time. In a time interval t , an atom will absorb a mean number of photons $N = R_{\text{scatt}}t = \Gamma\rho_{2,2}t$. Assuming Poissonian statistics, the variation in the absorbed number of photons is \sqrt{N} . This non-uniform nature leads to fluctuations in the force and causes a random walk in velocity along the direction of the laser beam. Both forms of random walk are analogous to Brownian motion (and contribute approximately equally), which leads to an effective heating of the atom. This limits the effective temperature that can be cooled to in the molasses. The limit is derived from the balance of heating with the laser cooling and the minimum temperature is found when setting $\delta = -\Gamma/2$, which gives the following:

$$T_{\text{D}} = \frac{\hbar\Gamma}{2k_{\text{B}}}. \quad (2.45)$$

This equation represents the lowest theoretically achievable temperature of the optical molasses¹⁹. For Cs, this limit is $T_{\text{D}} = 125.61 \mu\text{K}$ [64].

2.4.5 Sub-Doppler Cooling

However, early experiments with optical molasses yielded temperatures much lower than the Doppler limit [71]. Two important factors that led to sub-Doppler temperatures were missing from the initial theories of laser cooling. First, initial theories considered only two-level atoms, whereas alkali atoms have a more complex energy structure with multiple hyperfine levels, along with degenerate Zeeman levels. Second, an atom moving in between two counter-propagating laser beams will experience a spatially varying polarization along the direction of the laser beams. These two factors give rise to another phenomenon called the Sisyphus effect, or Sisyphus cooling, which is one type of polarization-gradient cooling (PGC). The explanation of sub-Doppler cooling mechanisms for multi-level atoms were described by two groups [72, 73] in 1989.

The Sisyphus effect is best explained by considering an atom moving in between two counter-propagating beams of orthogonal polarization. In the standing wave formed by the beams, the

¹⁹It is important to note that the intensity in each of the six beams is well below saturation where the absorption from each pair of beams is uncorrelated.

resulting polarization depends on the relative phase of the two beams and varies along the length of the intersecting region. The changing polarization causes a spatially varying light shift of the ground state hyperfine levels. This causes the atom to see hills and valleys in potential energy where the ground states shift up and down in energy, as shown in figure 2.8. A cooling mechanism occurs when the atom reaches the top of a hill and absorbs a photon, then decays back down to the bottom of the valley, thus losing energy. The transition rules between hyperfine sub-levels allow for multiple decay routes, but the transition to the state at the bottom of the valley has a higher probability for certain conditions. The period of oscillation for the polarization occurs over $\Delta z = \lambda/4$. Thus, the force is largest for those atoms that travel this distance during one optical pumping cycle. The idea of the atom climbing these hills and valleys, then losing energy at the top sending it to the bottom again is where the Sisyphus name comes from. This is analogous to the Greek myth in which Sisyphus was punished by having to push a great rock up a hill, only for it to fall to the bottom every time he neared the top.




Figure 2.8: Schematic of the Sisyphus cooling arrangement for an atom moving from left to right with ground state $J = 1/2$ and excited state $J' = 3/2$ (adapted from [62]). The standing wave formed by orthogonal linearly polarized counter-propagating beams results in a spatially varying polarization. The varying polarization modulates the light shift of the ground state energy levels creating potential valleys and hills. Sisyphus cooling occurs when an atom absorbs at the top of the potential hill and then decays back down to the potential valley.

The temperature limit for the Sisyphus cooling mechanism is determined when the energy lost in each Sisyphus cycle balances the recoil energy gained from spontaneous emission. The limit is known as the recoil temperature and is given by:

$$T_r = \frac{\hbar^2 k^2}{M k_B}. \quad (2.46)$$

In Cs, this temperature limit is $T_r = 198.34$ nK [64]. Note that the difference between the

Doppler and recoil limits are particularly large for Cs because it is the heaviest (stable) alkali metal.

In addition to Sisyphus cooling, there is another type of PGC mechanism called ‘motion-induced orientation’. This is more relevant for a standing wave formed from orthogonal circular polarization (σ^+ and σ^-). The resulting polarization is linear everywhere except it rotates through an angle 2π over one optical wavelength (analogous to a corkscrew). An atom moving through this polarization experiences a constant electric field magnitude but with changing direction. This redistributes the population of atoms across the various sublevels. This leads to an imbalance in the probability of photon absorption and a differential in the scattering force from the two laser beams. It is difficult to single out either one of these two sub-Doppler cooling mechanisms, since in a real three-dimensional optical molasses, both types of polarization-gradient will be present.

2.4.6 The Magneto-Optical Trap (MOT)

Until this point, we have only described methods to laser cool atoms. After some time, the atoms will eventually diffuse out of the laser cooled regions and heat up again. In this section, we explore how an optical molasses can be combined with an appropriate magnetic field to generate a radiation force that depends on position as well as velocity. This configuration is known as a magneto-optical trap (MOT) and is capable of confining an ensemble of cold-atoms tightly at the centre of an optical molasses.

Figure 2.9 illustrates how the MOT operates for the simple case of a $J = 0$ to $J' = 1$ transition. A magnetic quadrupole field is used to generate a gradient across the trap centre with $B = 0$ at $z = 0$. The gradient splits the upper $J' = 1$ level into a Zeeman triplet, where the splitting strength depends on the atoms’ position in the gradient. Counter-propagating beams that are red detuned from the $\Delta M_J = 0$ transition illuminate the atoms as in the optical molasses but now have circular polarization. Red-detuned light with right-hand circular polarization²⁰ is directed along the $+z$ direction, which drives σ^- transitions in the $z < 0$ region (where the magnetic field is anti-parallel to beam propagation). In the $z > 0$ region (where the magnetic field is parallel to beam propagation), the same beam instead drives σ^+ transitions, although with a reduced scattering rate due to the magnetic field gradient which increases the detuning. Similarly, an identically polarized counter-propagating beam travelling along the $-z$ direction will drive σ^- transitions in the $z > 0$ region and σ^+ transitions for $z < 0$. Thus, when both beams

²⁰Where the polarization is defined from the point of view of the source.




Figure 2.9: Magneto-optical trapping schematic for an atom with $J = 0$ to $J' = 1$ transition. The Zeeman splitting of the upper $J' = 1$ level depends on the atom's position z in the magnetic field. The atom is illuminated by two counter-propagating circularly polarized beams. The laser frequency for both beams is slightly red detuned from the $J = 0$ to $J' = 1$ transition.

are present, the force from the $+z$ propagating beam is greater than the counter-propagating beam in the region $z < 0$ and vice versa. The net force on the atom is directed towards the centre of the trap. This picture can be extended to all three axes, and the experimental arrangement of the MOT is shown in figure 2.10. Three pairs of counter-propagating beams form a molasses in their overlap region with appropriate circular polarization. A pair of anti-Helmholtz coils generate a quadrupole field²¹, with a field strength that is $B = 0$ at the centre and increases linearly in all directions (for small displacements around the zero point).

The forces that atoms experience in the MOT have a similar origin to those in the optical molasses. In one dimension, the forces from two counter-propagating beams sum to give the total force experienced by the atom, as described above, except now the detuning parameter δ_{\pm} contains the Zeeman shift $\omega_Z = \beta z$, and the Doppler shift term $\omega_D = -kv_z$ is still present. Hence:

$$\delta_{\pm} = \delta \mp kv_z \pm \beta z. \quad (2.47)$$

The force from each beam is the same as equation 2.43, except now δ_{\pm} has an extra term proportional to z . The constant β quantifies the Zeeman splitting and is represented as:

²¹If the polarity of the supply (and current) to the coils were reversed in figure 2.10, the magnetic field direction would reverse, thus requiring the axial and radial beams to switch their handedness of polarization. Either approach can be used to generate a MOT.




Figure 2.10: Experimental arrangement for the magneto-optical trap. Three pairs of counter-propagating laser beams form an optical molasses, and a pair of anti-Helmholtz coils generate a quadrupole field. The magnetic field around the centre of the trap is shown by the black arrows. Note the handedness of circular polarization across all six beams.

$$\beta z = \frac{\mu' dB}{\hbar dz} z, \quad (2.48)$$

where $\mu' = \mu_B (g_e M_e - g_g M_g)$ is the effective Bohr magneton, defined with $g_{e,g}$ the Landé g-factor, and $M_{e,g}$ the magnetic quantum number. The subscript denotes the excited and ground state levels. Summing the forces from each beam gives the total force on the atom in the low-velocity limit as:

$$F_{\text{MOT}} = -\alpha v_z - \kappa z. \quad (2.49)$$

The first term is the damping in the optical molasses. The extra positional-dependent term acts like a spring with constant $\kappa = \alpha\beta/k$. The positional-dependent nature of this force makes the MOT operate as a trap. Atoms that enter the optical molasses region experience the combination of a viscous damping force and a trapping force that pushes them to the centre. The first experimental demonstration of a MOT was performed using sodium atoms [74]. It was shown that $\sim 10^7$ atoms could be trapped with sub-millikelvin temperatures in a region less than 0.5 mm in diameter.

The practical implementation of the MOT is surprisingly robust. It does not depend on

precise alignment of the counter-propagating beams or a high purity of polarization²². The trap is relatively easy to construct given that the required magnetic fields can be produced by air-cooled coils, and the trap itself can be loaded directly from an alkali vapour. Although, MOTs loaded directly from a vapour tend to have considerably fewer atoms than those loaded from a slow atomic beam. Furthermore, advances in diode technology have allowed MOTs to be built for many atomic species. Therefore, nowadays, the MOT is the main workhorse for many cold-atom experiments.

2.5 Summary

The interaction of atoms with electromagnetic radiation is a key concept in the field of atomic timekeeping. It also plays a crucial role in certain aspects of the clock developed during this research. For instance, we use a retroreflected MOT to simultaneously cool and trap Cs atoms before subjecting them to a subsequent stage of PGC to reduce their temperature further. When released from the MOT, a Rabi pulse of microwaves excites the clock transition, and the fraction of atoms that have undergone the excitation is detected. In the next chapter, the optical layout required for laser cooling and trapping is discussed, along with design considerations for the development of a compact laser system.

²²This statement should be treated with caution, given that for practical applications we often want an optimised MOT.

Chapter 3

Laser and Optoelectronics System

This chapter details the laser system used for cooling, trapping and state detection in the proposed atomic clock. Before discussing the optical system constructed for this work, several laser requirements are outlined. After the physics package, the laser system tends to be the next limiting aspect of the overall size in most clock architectures. In this work, the laser system is comprised of conventional free space optics on an optical table, which was motivated primarily by simplicity and ease to set up. Hence, it would not be suitable in a truly miniaturised clock. The free space system serves to demonstrate what optical elements and laser systems are required, establishing a performance benchmark that needs to be achieved when the time comes to transition to a compact setup. This chapter concludes by discussing the improvements that could be made to miniaturise the free space system for a more compact solution.

The following sections show schematics of various free space optics for generating the light necessary for laser cooling and trapping. The legend below identifies different optical components [75].




Figure 3.1: Legend for optics used in various schematics. Additionally, for brevity, the many transitions discussed in the chapter follow the notation $F4 \rightarrow F'5$ corresponding to $|F = 4\rangle$ to $|F' = 5\rangle$.

3.1 Laser Requirements

For laser cooling and trapping caesium (Cs) in a magneto-optical trap (MOT), two key transitions are used (see figure 2.2). The first is the $F4 \rightarrow F'5$ hyperfine transition. The selection rules $\Delta F = 0, \pm 1$ dictate that the atoms can only decay back down to $F4$; this is known as a closed transition because there are no decay routes, and in theory, the cooling cycle can be performed many times. However, there is a non-zero probability for off-resonant excitation from $F4 \rightarrow F'4$. Once atoms occupy the $F'4$ state, the selection rules grant decay down to the lower hyperfine level $F3$. Once an atom is in the $F3$ state, it can no longer participate in the cooling cycle. Therefore, we must also drive the $F3 \rightarrow F'4$ repumping transition to bring atoms back into the cooling cycle.

The lasers themselves also have a few requirements. Some relate directly to the laser source, while others point to the need for auxiliary optics that must complement the laser source. A few of these key points are highlighted in this section. For laser cooling alkali metals, both a cooling and a repumping transition must be driven simultaneously. In conventional systems, this is achieved by using two lasers with two separate optical setups (as in this work). However, some approaches use a single laser (discussed further towards the end of the chapter). Furthermore, each laser must be able to address individual hyperfine transitions. To appropriately set the frequency of the laser and prevent frequency fluctuations, the lasers are frequency stabilised using a Doppler-free spectroscopy signal (the stabilisation procedure is described in more detail in later sections). This highlights the experimental complexity associated with a laser-cooled clock. Using two lasers and additional frequency stabilisation optics adds more footprint to the laser system, which increases the overall clock size.

Another important consideration is the power output from the laser source, which is directly related to the intensity. The optical intensity at the MOT should be at least of the order I_{sat} (for Cs, $I_{\text{sat}} = 1.1049 \text{ mW cm}^{-2}$, for a circularly polarized beam [64]). From equation 2.42, a higher optical intensity ensures that the transition is saturated and the force experienced by the atoms is maximised. This allows for the largest range of velocities to be captured and the fastest capture rate. Less time spent loading atoms into the MOT can reduce the overall cycle time of the experimental sequence. The intensity depends on the optical power P and the size of the MOT beams w_0 (where w_0 is the $1/e^2$ Gaussian radius of the beam [76]):

$$I = \frac{2P}{\pi w_0^2}. \quad (3.1)$$

This emphasises the need to select a laser source with sufficiently high output power to overcome various losses from the optical setup. Finally, the laser linewidth must also be considered. For cooling in the MOT, the linewidth must be narrow enough to probe only the red-detuned side of the cooling transition. If the frequency fluctuations are large enough to probe the blue-detuned side of the transition, then this will negate the effects of the laser cooling. For this reason, a laser with a linewidth less than the natural linewidth Γ of the transition is used, which for Cs is $\Gamma \sim 5.2$ MHz. The laser used for cooling is also used for state detection in the clock, which imposes a stability requirement since frequency fluctuations of the laser translate to noise in the detected number of atoms. For detection pulses on the order of a few 100 μ s in length, a laser linewidth of a few tens of kHz is usually satisfactory. The linewidth requirements for the repumping laser are less stringent, given that its sole purpose is to depopulate the $F3$ state. Most distributed feedback (DFB) and distributed Bragg reflector (DBR) laser chips have a modest linewidth of ~ 1 MHz, which is adequate. With frequency stabilisation, the linewidth can be further reduced to ensure that frequency noise on the timescale of the detection is sufficiently suppressed. With these key requirements outlined, attention is now drawn to the lasers and optical setups used in this experiment.

3.2 Cooling Light

The first transition required for laser cooling is $F4 \rightarrow F'5$. This is used as the cooling cycle in the MOT and also for clock state detection. Figures 3.2 and 3.3 show the optical setup used to generate the cooling light. The laser source is a Toptica DLPRO external-cavity diode laser (ECDL) laser, which can be used with a free space or fibre output. In this experiment, we have opted for the FibreDock to provide a clean Gaussian spatial profile of the beam when launching into the optical setup. The power output after the fibre is ~ 46 mW. The light is immediately split roughly 90:10 into two arms. Rather than using a beam splitter with a fixed split ratio, we use a $\lambda/2$ waveplate and a polarizing beam splitter (PBS) in combination. The split ratio is then controlled by adjusting the angle of the $\lambda/2$ waveplate. The 90% arm is directed to the trap, while the other 10% is used as a pickoff for frequency stabilisation. The pickoff is first sent through an acoustic-optical modulator (AOM) in a double-pass arrangement [77] and

then into a compact saturated absorption spectroscopy (SAS)²³²⁴ setup. The AOM arrangement allows the frequency to be shifted whilst maintaining the lock. The double-pass shifts the light by twice the applied RF frequency towards the blue side of the transition. To bring the laser back into resonance with the cooling transition, a second double-pass arrangement is used on the path of the laser beam sent to the trap, which red-shifts the light to counteract the first (lock) AOM. This combination of AOMs allows for frequency detuning to the red and blue side of the transition, equal to twice the difference between the two applied AOM RF frequencies. After the trap AOM, the cooling light is combined with the repumping light on a PBS. One arm contains only cooling light, while the other contains both cooling and repumping light. Both arms are coupled into polarization-maintaining (PM) fibres and directed towards the trap. The fibre containing only cooling light is split using a 50:50 PM fibre splitter to form the radial beams of the MOT. The other fibre, containing cooling and repumping light is used for the vertical beams in the MOT. The split ratio of the combination PBS is set such that each fibre output at the trap has an equal amount of cooling light. In a PM fibre, two orthogonal polarization modes can propagate. Any stress or strain on the fibre can cause crosstalk between the two modes and fluctuations of the output polarization. This crosstalk is quantified as the logarithmic ratio between the powers in each polarization mode, also known as the extinction ratio. PM fibre is specifically designed to reduce this crosstalk between the two modes, but care must be taken to align the input light with one of the axes of the PM fibre. Alignment is accomplished using $\lambda/2$ and $\lambda/4$ waveplates before the fibre collimator, which provides complete control over the input polarization. The alignment was optimised using the signal from a polarimeter²⁵, which monitors the output polarization state. The extinction ratio was measured to be > 25 dB in both cases. Attenuation and extinction of the cooling light are achieved by varying the RF power sent to the trap AOM. A mechanical shutter is in place after the trap AOM double-pass for complete optical extinction. After the entire optical setup, we were able to attain ~ 3 mW of cooling light in each arm of the MOT. For the MOT beam radius $w_0 = 6.1$ mm, we find that the intensity at the beam centre is $I = 4.6I_{\text{sat}}$.

²³The CoSy Doppler-free spectroscopy module from Toptica is a compact SAS with fibre input.

²⁴Although the CoSy is a form of saturated absorption spectroscopy, the signals it generates are not entirely described by the saturated absorption mechanism. There will also be hyperfine optical pumping to the lower ground state, which results in a narrow Doppler-free feature whose width is governed not only by the laser intensity but also the dimensions of the pump beam [78]

²⁵Thorlabs: PAX1000IR1/M.

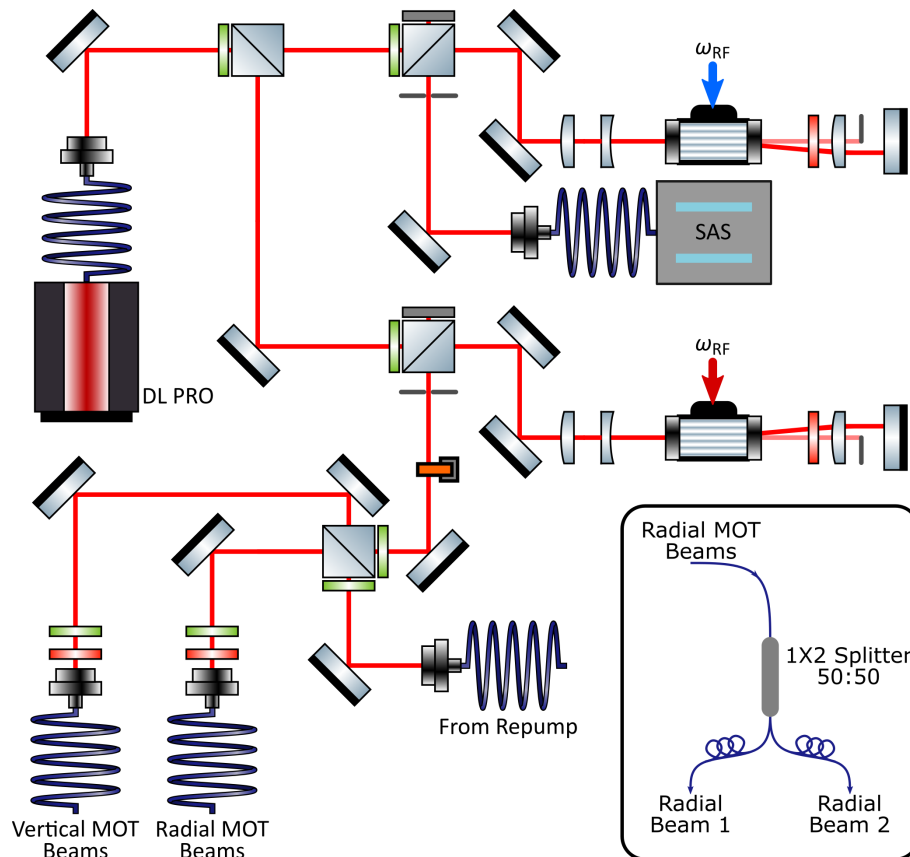


Figure 3.2: Optical schematic for generating the cooling light as well as the combination with the repumping light before fibre launching to the physics package. The radial MOT beams are split 50:50 using an in-line fibre splitter. The grey SAS box is the Toptica CoSy module.




Figure 3.3: Optical table layout of the cooling light optics in the laboratory.

3.3 Repumping Light

The second required transition is $F3 \rightarrow F'4$ to provide repumping light. Given that its primary purpose is to depopulate the $F3$ state, the repumping requirements are less stringent compared to the cooling light. Therefore, the optical setup is slightly more simplified. Figures 3.4 and 3.5 show the optical setup used to generate the repumping light. The laser source is a Thorlabs DBR (DBR852PN) butterfly laser with a pigtailed PM fibre output. The butterfly package contains a thermoelectric cooler (TEC) for temperature stabilisation and an integrated 30 dB optical isolator. The laser light is initially split 80:20 using a $\lambda/2$ and PBS arrangement (as described previously). The 20% arm of the PBS is sent to a homebuilt SAS setup for frequency stabilisation to the $F3 \rightarrow F'3$. The remaining 80% arm passes through an AOM, where the first-order diffracted mode blue-shifts the locked laser light to the required $F3 \rightarrow F'4$ repumping transition. Although the repumper does not require frequency shifting capability, the AOM and offset lock arrangement provide faster switching speeds ($\sim 1 \mu\text{s}$) via the AOM; the mechanical shutter is used only for complete optical extinction. By itself, the mechanical shutter switching speed is too slow ($\sim 1 \text{ ms}$ and has a time delay that depends on the positioning of the shutter) to perform the state detection sequence where the pulses are on the order of a few 100 μs . After the AOM, the laser light is coupled into a PM fibre and directed to the cooling optics where it is combined with the cooling light and ultimately sent to the trap. Before the fibre coupler, there is a single $\lambda/2$ waveplate for coarse alignment into one of the PM fibre axes. The quality of the MOT is not sensitive to the polarization of the repumping light since it only needs to give weak excitation out of the $F3$ state. As described in the previous section, the repumping beam is overlapped with the cooling light only in the vertical direction for the MOT. This was done to ensure that sufficient repumping power would be available for detection once the atoms had fallen to the lower chamber of the physics package. After the integrated pigtail fibre, the DBR provides $\sim 14 \text{ mW}$ of power output. By the time the repumping light reaches the trap, the power is reduced to $\sim 1.7 \text{ mW}$, which corresponds to $I = 2.9I_{\text{sat}}$.

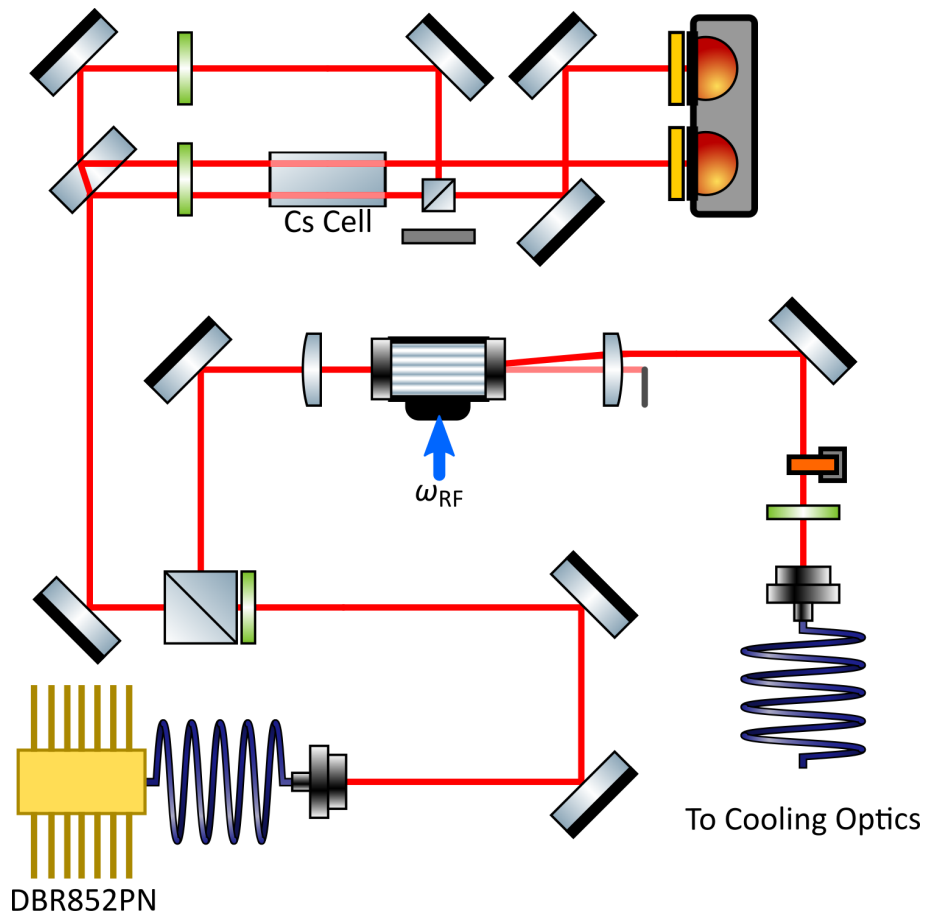


Figure 3.4: Optical schematic for generating the repumping light. Frequency stabilisation is achieved with a homebuilt SAS module. Band-pass filters are placed over the photodetectors to block background light.




Figure 3.5: Optical table layout of the repumper light optics in the laboratory.

3.4 Frequency Stabilisation

While the short-term frequency stability of the free-running laser diodes might be sufficient for some applications, most high-precision and frequency metrology applications require good laser frequency stability over long timescales. Therefore, it is necessary to stabilise the laser frequency to an appropriate stable frequency reference. Fortunately, for Cs, an atomic reference frequency can be easily obtained using a spectroscopy vapour cell. However, absorption spectroscopy with a single laser beam provides only Doppler-broadened features, which are typically too broad for precise frequency stabilisation. The most popular method to obtain high resolution is Doppler-free spectroscopy. In this arrangement, the individual hyperfine transitions can be resolved, which offers multiple locking options as well as narrower features for better stability. In most cold-atom experiments, a small pickoff of laser light is sent to an SAS setup for frequency stabilisation. A readily available atomic reference from a spectroscopy cell is appealing for neutral atom systems to keep things compact. In comparison, the wavelengths required for an atomic ion-based clock are not easily attainable from a spectroscopy cell (which does not contain ions). Therefore, ion-based clocks must resort to using reference cavities or wavemeters for stabilisation, which can increase the overall system size. In this experiment, the optical layouts for cooling and repumping have a small section dedicated to SAS. Once a Doppler-free spectroscopy signal is obtained from the SAS setup, it must be converted into a useful error signal for frequency stabilisation. This is achieved via frequency modulation (FM) locking. In FM locking, the laser frequency can be modulated by tuning either the current or piezo grating. As the laser scans over the spectral feature, the FM is converted to an amplitude modulation (AM). If the FM is smaller than the linewidth of the spectral feature, then the AM signal will be proportional to the first-order derivative of the spectral feature. A lock-in amplifier is used to demodulate the AM signal by mixing it with the original RF signal used for FM. This yields a DC signal, which can be used as an error signal. An example of both the sub-Doppler absorption spectroscopy and error signals can be seen in figures 3.6 and 3.7. The remainder of this section discusses how the cooling and repumping lasers are frequency stabilised using the above approach.

The cooling transition is $F4 \rightarrow F'5$. However, for optimum cooling in the MOT and to perform PGC, it is necessary to red detune the laser frequency. The laser is stabilised to the $F4 \rightarrow F'5$ transition, and detuning is achieved using the dual AOM arrangement described in the previous section. One AOM shifts to the red, while the other shifts to the blue to provide complete control for detuning both sides of the resonance if needed. The SAS optics are provided

by a commercial off-the-shelf product by Toptica. The CoSy is a compact single-fibre input SAS module. The output from the CoSy can be seen in figure 3.6. It shows the three allowed hyperfine transitions $F4 \rightarrow F'3, 4, 5$, as well as three crossover resonances. The crossover resonances are a by-product of the counter-propagating pump-probe laser arrangement in the SAS setup and the multi-level nature of the atom. When the laser frequency is exactly midway between two transitions, there is a velocity class that experiences equal and opposite Doppler shifts, bringing both the pump and probe beam into resonance. The result is another peak exactly midway between two real transitions. This symmetry allows them to be distinguished from the other real transitions.




Figure 3.6: Sub-Doppler absorption spectroscopy signal (bottom trace) and corresponding error signal (upper trace) from the cooling setup. The x axis has been scaled relative to the $F4 \rightarrow F'3$ transition. All the transitions are identified above the sub-Doppler profile. $\rightarrow 4$ stands for the $F4 \rightarrow F'4$ transition. “CO 3,4” stands for the crossover peak between $F4 \rightarrow F'3$ and $F4 \rightarrow F'4$.

The cooling laser is driven by a commercial laser driver from Toptica (DLC Pro), shown in figure 3.8. This laser driver is designed to be used with Toptica’s ECDL series of laser sources and ensures the best noise performance. The FM is provided by the DLC Pro and applied to the laser’s current at a frequency of 20 kHz. A built-in digital lock-in amplifier is used to recover the error signal from a user-provided SAS signal. The ECDL has two frequency actuators: a piezo grating and the current source. Feedback is applied through two digital PID controllers, which are also built into the DLC Pro. The long-term performance of the lock was not measured, but it was observed that the laser would stay locked for several days without any issues. By this time, bulk relaxation of the fibre dock would cause misalignment and a drop in the output power.

This usually triggered the laser to become unlocked. Realignment of the fibre dock had to be performed to bring the laser back to its original working state.

Frequency stabilisation of the repumping laser was performed using a similar approach to the cooling laser. The desired transition for repumping corresponds to $F3 \rightarrow F'4$. However, fast switching speeds are required. Thus, the repumper is instead locked to the $F3 \rightarrow F'3$ transition, and an AOM is used to shift the frequency up by ~ 201 MHz to the desired repumping transition. Unlike the cooling optics, the SAS aspect for the repumper is homebuilt using conventional free space components. The output from the homebuilt SAS can be seen in figure 3.7. It shows the three allowed hyperfine transitions $F3 \rightarrow F'2, 3, 4$, and three extra crossover resonances.




Figure 3.7: Sub-Doppler absorption spectroscopy signal (bottom trace) and the corresponding error signal (upper trace) from the repumper setup. The x axis has been scaled relative to the $F3 \rightarrow F'2$ transition. All the transitions are identified above the sub-Doppler profile. $\rightarrow 3$ stands for the $F3 \rightarrow F'3$ transition. “CO 2,4” stands for the crossover peak between $F3 \rightarrow F'2$ and $F3 \rightarrow F'4$.

There are numerous benchtop drivers used to control the repumping laser, as shown in figure 3.9. Firstly, the laser is driven by a QubeCL [79] low-noise current source (RMS noise of 350 nA for 10 Hz to 1 MHz) and a Thorlabs temperature controller. Unlike the ECDL used for the cooling light, the repumping DBR only has one frequency actuator²⁶, which is its current. Fortunately, the laser driver has two modulation inputs, which can be used to apply a bias to the output current. These are used for both a slow ramp (~ 24 Hz) to sweep over the hyperfine features and a fast modulation (~ 500 kHz) for FM locking. A benchtop lock-in amplifier generates the

²⁶It is true the temperature could be used, but the response time is much slower. Additionally, an extra actuator was not needed.

dispersive error signal, which is fed back into a Vescent laser-servo unit (D2-125). The laser-servo unit has a built-in PI²D controller, which is used to feedback to the current of the DBR and keep its frequency stabilised to the $F3 \rightarrow F'3$ transition. It was observed that the laser would remain locked for weeks without issue. Unlike the ECDL, the DBR does not have an external grating, which mitigates the sensitivity to acoustic disturbances.




Figure 3.8: *Toplica's DLC Pro laser driver for the cooling laser. The CoSy control box is also pictured above the DLC Pro.*




Figure 3.9: *Repumping laser driving and frequency stabilisation electronics. A lock-in amplifier (not shown) is also used to generate the error signal.*

3.5 Compact Laser System

For a compact atomic clock, the optical setup described above is not entirely suitable. However, for this demonstrator system, the conventional free space approach is very useful for developing and understanding the clock system. This can lead to a subsequently more compact system after

some considered modifications.

First, an electro-optic modulator (EOM) can be used to generate the light for repumping. An EOM can be used as a phase modulator to imprint red and blue sidebands on the carrier frequency separated by twice the applied modulation. By modulating at ~ 9.192 GHz, the first-order blue sideband will be resonant with the repumping transition. EOMs were used as a method to perform chirp cooling in early laser cooling experiments [80]. When the frequency to the EOM is swept at the right frequency, the first-order sidebands are also swept to maintain resonance with the atoms. For experiments where an MOT is loaded directly from an alkali vapour, chirp cooling is not necessary, and thus, the EOM simply provides repumping capabilities. This has already been demonstrated in other cold-atom experiments [81, 82]. Using an EOM would eliminate the need for a second laser and an optical setup for the repumping light entirely.

Second, the optical setup needs to operate over long timescales without the need for user intervention or realignment. This is particularly crucial for a compact clock that might be subjected to large fluctuations in ambient temperatures or harsh vibrations. An all-fibre-based system would make the optics package more robust to any external perturbations and remove the need for any alignment procedures. Furthermore, compared to a free space setup, an all-fibre-based system would be easier to package into a small box. A truly miniaturised system would implement the entire optical setup and beam delivery on a photonic-integrated chip (PIC) [83]. However, this technology is still in its infancy and requires more research and development before it can be implemented in an atomic clock.

Finally, the ECDL source is also susceptible to acoustic and temperature perturbations. This would need to be replaced by a single diode source such as a DBR or DFB. It has been shown that an atomic fountain clock can be made far more robust without any degradation in performance when using a DBR laser [84]. If one considers switching to rubidium, where the laser cooling wavelength is 780 nm, it is possible to use a frequency-doubled telecom laser at 1560 nm. While this might seem more cumbersome, this technique is based on the very mature technologies developed at telecom wavelengths. Compact, robust fibre-based systems have already been demonstrated over the years, as they form the backbone of today's modern telecommunications infrastructure. This approach has already been used to great effect in other cold-atom sensors [85, 86] and even some commercial stable laser sources [87].

Chapter 4

Physics Package

At the heart of every atomic clock lies the ‘physics package’. This is the term given to the collection of components such as: the vacuum system housing the atomic reference; the clock state interrogation apparatus; magnetic field coils; and the detection system. In many atomic clocks, the physics package is often the most limiting component of the overall size. Thus, for compact clock developments, most research is focused here on attempting to miniaturise certain aspects to try and keep the size as small as possible. The physics package constructed here is comprised primarily of off-the-shelf components. Therefore, it is not an accurate representation of how a final clock of this variety might look. This approach was taken primarily for simplicity but also because the goal was to develop a system where the procedure of laser cooling and dropping the atoms could be explored and better understood. As a result, we have gained valuable insight, which will aid in the design of a potentially smaller second iteration. This is discussed further at the end of the chapter; until then, the remainder of the chapter details all the components of the physics package, most of which were designed and built from the ground up. Figure 1.13 provides a schematic overview of the physics package and its many components.

4.1 Vacuum System

Laser cooling and trapping require high vacuum pressures. At regular atmospheric pressures, the rate of collisions with background gas molecules is too high, which makes laser cooling impossible. This represents another major experimental complexity associated with a laser-cooled clock. The vacuum chamber constructed for this work can be separated into three primary sections: the magneto-optical trap (MOT) region, the clock-state interrogation region, and the detection region. The first section is the MOT chamber, where the atoms are cooled and trapped. This

is situated at the top of the vacuum system. The chamber is a ‘unitary spherical octagon’ from Kimball Physics and constructed from 316L stainless steel. There are eight DN16CF radial ports and two DN35CF axial ports. For optical access into the chamber, there are several fused-silica viewports. The viewport flange is made of 316LN stainless steel, into which the fused silica is brazed. To minimise optical losses, the viewports are also coated with an anti-reflection (AR) coating on both sides. All but one of the eight radial ports are used with an optical viewport. The remaining port allows for the connection of a small 90° elbow adapter where the caesium (Cs) dispensers are housed. Of the two DN35CF ports, the upper is used with an optical viewport for the vertical MOT beam. The lower is connected to a DN35CF-DN16CF reducer flange to connect to the interrogation flight tube. Figure 4.1 shows the assembly of the vacuum system prior to bakeout.




Figure 4.1: (a) Assembly of MOT chamber with viewports and vacuum elbow to house the Cs source. (b) Fully assembled vacuum system in an environmental chamber prior to bakeout.

This leads to the next area of the vacuum system, which is the clock-state interrogation region. Once released from the MOT, the atoms begin to free fall through the vacuum system, where they are interrogated during free fall. Ideally, the interrogation would be performed using a resonant microwave cavity. A microwave cavity was designed (more on this in chapter 5) but not implemented here due to timing constraints. Instead, the interrogation region consists of a DN16CF cross with optical viewports perpendicular to the drop path of the atoms. With this approach, a microwave horn antenna could be used to explore the dynamics of performing clock-state interrogation. The time it takes the atoms to pass the window is ~ 11 ms, which limits the potential interrogation time. In fact, this time will be slightly less since there will be significant spatial extent to the cloud at this point. However, this approach also offered the possibility of

experimenting with a Ramsey-style interrogation using two horn antennas. The atoms could be interrogated once in the MOT chamber and again when passing the DN16CF cross, separated by a dark period of ~ 148 ms. This type of interrogation region is not ideal and was not intended as the final design.

After interrogation, the atoms continue to fall into the detection chamber, which is housed in a DN35CF cube. Three AR-coated viewports provide optical access to monitor the fluorescence from the detection sequence and to retroreflect the vertical beam back up the flight tube. The two remaining ports are used for the pumping apparatus and a valve for connection to a bakeout station. The distance between the centres of the MOT and detection chambers is 200 mm, which corresponds to a free fall time of ≈ 200 ms.

The entire chamber is pumped with a SAES NEX Torr D100-5 pump, which is an integrated ion (5 L s^{-1}) and getter (100 Torr L) pumping system²⁷. This pump was chosen because of its relatively compact size and large pumping speeds. The measured pressure was $\lesssim 10^{-10}$ mbar, as inferred from the current draw of the ion pump. After assembly, the system was baked out for five days at a temperature of 150°C . This temperature was limited by one of the vacuum gauges used to monitor the pressure during the bakeout. The system was then disconnected from the bakeout station and valved off. The ion pump is used to maintain vacuum when the apparatus is on the optical table. Figure 4.2 shows both the CAD design of the physics package and the assembled system in operation.

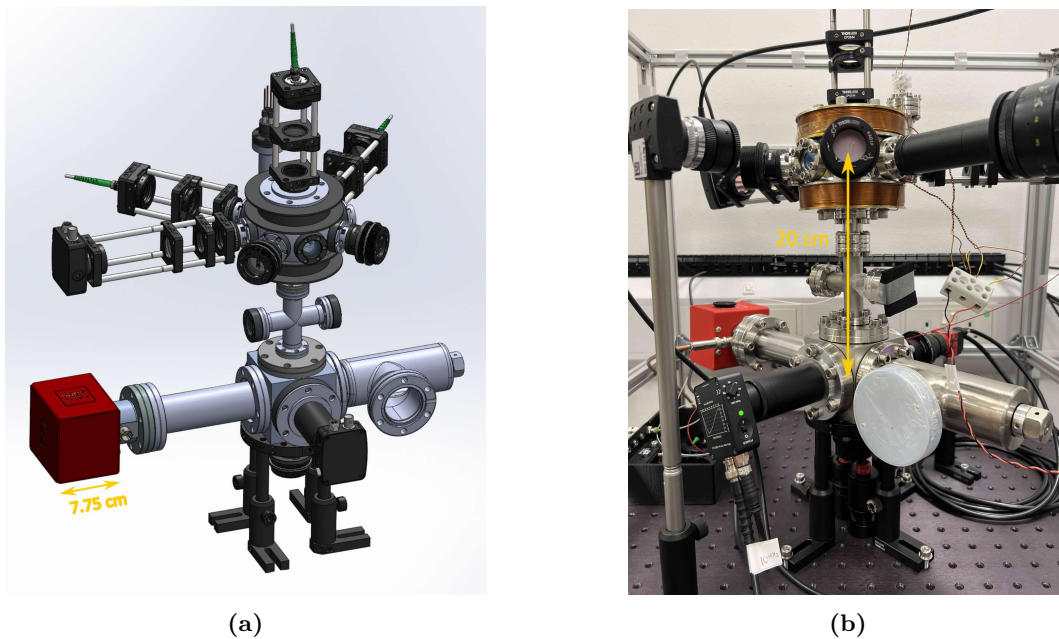


Figure 4.2: (a) Computer-aided drawing (CAD) of the vacuum chamber with trap and detection optics. (b) Vacuum chamber assembled and operational in the laboratory.

²⁷Pumping speeds (in litres per second and Torr litres) quoted are for O_2 .

4.2 Trap Optics

A retroreflected MOT was used with the aim of maintaining a compact form factor. The retroreflected MOT needs only three input beams rather than the conventional six. This simplifies the design of the optics required at the trap. The radial MOT beams traverse only the MOT chamber, whereas the vertical beam is directed down the entire flight tube and retroreflected back up. In this arrangement, the vertical MOT beams are also repurposed for state detection. A custom RMS²⁸ threaded adapter is attached to the viewport to allow optics to be mounted on the vacuum system. The trap optics are comprised of two components: a cooling-collimator and a retroreflector. The cooling-collimator is the input for the MOT light, and the retroreflector reflects the input light back on itself to form two counter-propagating beams.

The cooling-collimator consists of a terminated fibre end, which emits light diverging at an angle that depends on the numerical aperture (NA) of the fibre²⁹. The light passes through a $\lambda/4$ waveplate pre-aligned with the slow axis of the PM fibre to generate circularly polarized light. The still divergent beam is then collimated using a plano-convex lens, where the focal length dictates the beam radius. The resulting collimated radius is calculated using the following equations from Gaussian optics [76]:

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}, \quad (4.1)$$

$$z_R = \frac{\pi}{\lambda} w_0^2, \quad (4.2)$$

where w_0 is the initial radius of the beam and will be equal to half the mode field diameter (MFD) of the fibre. z_R is the Rayleigh range and depends on the wavelength λ and w_0 . Finally, z is the distance from the fibre, which will be equal to the focal length of the collimating lens. A lens with focal length $f = 59.8$ mm generates a beam radius of $w = 6.1$ mm. The focal length of the collimating lens was chosen such that the radius of the beam approximately maximised the available diameter of the viewport (~ 16 mm). The radius definition $w(z)$ matches that of equation 3.1 and contains $1/e^2$ of the total intensity.

Once collimated, the beam passes through the vacuum chamber and reaches the retroreflec-

²⁸A standard of thread used by the Royal Microscopical Society on almost all microscope objective lenses. Threading: 0.800" – 36.0. The RMS thread was chosen because it is close to the open diameter of the viewport and Thorlabs stock threaded components/adapters.

²⁹Fibre: PM780-HP, NA: 0.12, MFD: 5.3 μm @ 850 nm.

tor, which consists of another $\lambda/4$ waveplate and a mirror for retroreflecting. Precise angular alignment of the waveplate is not necessary given that the incident circular light is converted to linear polarization at an angle of 45° to the slow axis of the waveplate. Upon retroreflection, the linear polarization is unchanged, and the second pass through the $\lambda/4$ generates circular light except with the opposite hand to the incident beam. The reflecting mirror is also mounted in a kinematic mount for fine adjustment of the back-reflected beam. Figure 4.3 shows both the cooling collimator and retroreflector.




Figure 4.3: (a) Assembly of three cooling collimators. (b) CAD drawing of the retroreflector.

Good-quality optical molasses relies on the optimisation of a few experimental parameters: the magnetic field must be appropriately nulled (see section 4.4); the intensity in the beam pairs must be balanced; and the back alignment must be set correctly. Difficulties arise when trying to discern the difference between each of these and their effect on the cold-atoms. For example, an incorrectly nulled magnetic field could be overcompensated for by adjusting the intensity balance. To appropriately set these parameters, we used the following steps. The back alignment is first set using an iris placed in the cooling-collimator. Precise alignment ($< 5 \times 10^{-3}$ rad) is achieved when the retroreflecting mirror is adjusted, and the back-reflected beam passes through the iris even when it is entirely stopped down to a minimum aperture (1 mm). The magnetic field is nulled appropriately, as described in section 4.4. Only after the previous two parameters are set can the intensity imbalance be dealt with. Optical losses from cold-atom absorption and vacuum viewports can cause an intensity imbalance in the beams. To appropriately mitigate this, the distance between the terminated fibre end and the collimating lens can be adjusted to

generate a converging beam across the trap. This balances the intensity, but only at one point across the cloud. Thus, the balance across other parts of the beams is slightly compromised. The result will partially sacrifice the quality of the molasses, which could limit the potential number of trapped atoms or even the final cloud temperature. This is an unavoidable disadvantage of the retroreflected MOT compared to the conventional six-beam arrangement. The procedure for correcting the imbalance is as follows: the cloud is monitored with a camera; the MOT coils are turned on and off repeatedly while adjusting the lenses to ensure that the cloud does not shoot off in one direction. Once optimised, the atom cloud should expand uniformly in all directions and stay predominantly in the centre of the trap.

4.3 MOT Coils

The MOT coils provide the quadrupole field for trapping the Cs atoms. They are held in place by two circular aluminium formers, which fit tightly around the top and bottom viewports on the MOT chamber. In this design, the formers are bolted with the viewports onto the MOT chamber to form the vacuum seal. This means that they were attached to the system during the bakeout. Initially, this was chosen to keep things compact. However, in hindsight, having the ability to attach and detach the coils would have been a better option. Bakeout with the coils presented a minor issue. We had to ensure that the enamel coating on the wire used for the coil did not burn off at high temperatures. This would cause an electrical short between the coils and the chamber, which could compromise the current drawn through the coils. A small section of wire was tested in an oven at 180 °C for ~ 2.5 h. After the test, a continuity check was performed and there were no signs of an electrical short. However, upon bakeout of the assembled system, a short was found with one of the coils. This was overcome by driving the coils with a floating power supply (earth-isolated).

The two coils were wound using a custom coil winding machine at the NPL workshop. They have 378 and 383 turns respectively. They are driven in series in an anti-Helmholtz configuration to provide the necessary quadrupole field. To ensure that the gradient generated at the centre of the MOT was correct, the system was modelled in COMSOL. The results of the model are shown in figure 4.5, along with the B-field measured using a fluxgate magnetometer. A current of 500 mA produces a field gradient of $\sim 0.1 \text{ T m}^{-1}$ at the centre, which is the typical gradient used in an MOT [62].




Figure 4.4: (a) MOT coils fully wound prior to installation on the vacuum system. (b) Small coil of wire wound around a vacuum adapter, which was used for testing the enamel coating temperature rating.




Figure 4.5: COMSOL simulations for the magnetic field gradient generated along the axis of the coils for $I = 500$ mA. The green points represent the measured field.

4.4 Compensation and Quantisation Coils

The quality of the optical molasses is very sensitive to external magnetic fields. Stray magnetic fields will cause a Zeeman shift $\sim \mu_B B$ that perturbs the m_F sub-levels. If large enough, the shift can become comparable to the light shift U_0 and disrupt the Sisyphus cooling mechanism. For example, $T \sim U_0/k_B = \mu_B B/k_B = 2 \mu\text{K}$ implies that the field must be nulled to $\lesssim 3 \mu\text{T}$. This is an order of magnitude lower than the ambient magnetic field in the lab³⁰, which was measured to be $\sim 40 \mu\text{T}$ using a fluxgate magnetometer. To achieve low temperatures, the Earth's field must be appropriately compensated. This is performed via three pairs of 60×60 cm compensation coils that surround the experimental apparatus. The nulling process involves measuring the field without the vacuum chamber in the compensation coils and then adjusting the current for each

³⁰At the surface, the Earth's magnetic field typically ranges from $22 \mu\text{T}$ to $67 \mu\text{T}$ [88].

coil axes until the field is measured as $< 1 \mu\text{T}$. The vacuum chamber was then reintroduced. At $1 \mu\text{T}$, the molasses temperature is limited to $\sim 1 \mu\text{K}$. A low cloud temperature is critical to ensure that the spatial extent of the cloud remains small across the entire drop region. This nulling procedure is a suitable starting point, although microwave spectroscopy of the atoms can also provide a very good measure of the residual field (see section 7.4).




Figure 4.6: *Experimental apparatus with coils highlighted. Red: compensation coils. Blue: quantisation coils.*

The optical molasses region is not the only area that requires appropriate compensation. The 20 cm drop region traversed by the atoms must also be nulled. Interrogation of the clock transition requires a small and well-defined quantisation field (sometimes called C-field). This field is required to lift the Zeeman degeneracy and resolve the clock transition. This also leads to an unavoidable second-order Zeeman shift. Although, if the field is kept small, then so too is the shift. The quantisation field can be defined and controlled far more efficiently when all ambient fields are first removed. Given that the compensation coils are formed from a 60 cm^3 cube structure, the coil separation does not satisfy the ideal Helmholtz condition. This partially compromises the uniformity of the field across the centre of the coils. To understand the effect of this, the field profile was modelled using COMSOL (shown in figure 4.7). The centre of the compensation coil structure overlaps with the halfway point across the atoms' 20 cm drop path. During the time after being released from the trap to the detection point, the atoms

must not experience any vanishing magnetic field strength. This would cause a mixing of the Zeeman states, which would change the population levels and destroy any information after the microwave interaction. This leaves two options: optimise the compensation for the molasses region and deal with a $\sim 1.5 \mu\text{T}$ field during interrogation, or the reverse. We opted for the former to ensure a low cloud temperature and a higher number of atoms. A large field during the interrogation would lead to a large second-order Zeeman shift, which could be measured and corrected for. This balancing act is a main disadvantage of using the compensation coils.




Figure 4.7: *COMSOL model of the field profile generated from a pair of compensation coils with $I = 500 \text{ mA}$. The inset shows the field across the 20 cm drop region.*

Additionally, the coils only provide nulling of DC magnetic fields. In reality, there will also be AC fields, which, if not controlled, will affect the frequency stability of the clock. For this reason, magnetic shields would be a better solution, as they are capable of passively nulling both DC and AC fields. Not to mention, the compensation coils are currently one of the limiting factors in terms of the overall size of this experiment. The quantisation field is provided by a 55.5 cm square set of Helmholtz coils, which provide a uniform vertical field to align the atoms to the magnetic field component of the microwaves. The field is turned off during the cooling stage but immediately turned on prior to releasing the atoms from the trap. This is to ensure that efficient state preparation can be performed via optical pumping. The field remains on until the detection stage when it is eventually turned off to restart the experimental cycle. Even in a shielded system, the quantisation coil will still be needed inside the shields.

4.5 Caesium Source

Fortunately, the vapour pressure for Cs at room temperature³¹ is high enough such that the MOT can be loaded directly from a background vapour. This is very beneficial for compact systems, as it mitigates the need for first-stage cooling before loading into an MOT. In this experiment, the Cs source is provided by an alkali metal dispenser (AMD). To dispense the Cs, the AMD must be heated by passing a current through it. The dispensers are spot welded to a vacuum power feedthrough, as shown in figure 4.8. The feedthrough is mounted into a 90° vacuum elbow and attached to one of the eight radial ports on the MOT chamber. Control of the current can also provide control of the number of Cs atoms dispensed. However, it was found that too much current would degrade the strength of the fluorescence signal in the detection region. Oversaturation of the MOT chamber increases the number of collisions with background thermal atoms, which leads to an effective heating of the cold-atoms. The current was optimised relative to the detection signal strength shown in figure 4.8.




Figure 4.8: (a) Cs dispensers spot-welded onto a vacuum power feedthrough. (b) Strength of the detection signal plotted against dispenser current. Towards higher currents, the MOT chamber starts to saturate, and the signal is degraded.

4.6 Microwave Horn

In the absence of a microwave cavity, the interrogation region consists of a DN16CF cross, which provides optical access perpendicular to the free fall region. This was chosen so that a horn antenna could be used to direct microwaves through the viewport and still perform clock state interrogation. The horn antenna consists of a flared piece of waveguide, which acts to transition the microwaves from the waveguide structure into free space. This smooth transition

³¹10⁻⁶ mbar @ 20 °C [64].

reduces reflected waves and emphasises travelling waves. One of the main advantages of a horn antenna is that it provides a significant amount of directivity and gain. Figure 4.9 shows the field lines radiated from the horn antenna. The E-field is usually aligned parallel with the broad wall of the horn, whereas the B-field is aligned to the perpendicular axis. The travelling wave radiated by the horn antenna is initially a spherical wave but gradually approximates a plane wave in the far field. When the horn antenna is correctly orientated, the B-field can be made to align with the quantisation axis of the atoms. This alignment is crucial to drive the ‘field-insensitive’ transitions ($\Delta m_F = 0$) of the Cs atoms.

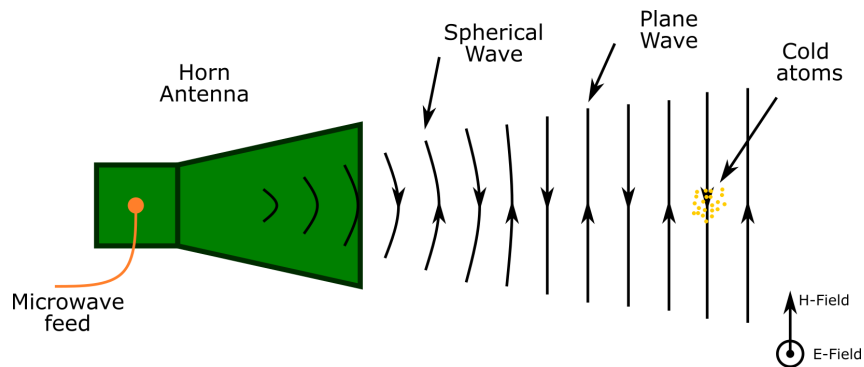


Figure 4.9: Schematic representation of the field lines emitted from the horn antenna.




Figure 4.10: Microwave horn antenna mounted next to a viewport on the MOT chamber.

The horn antenna is a good starting point but not an ideal long-term solution for clock-state interrogation. First, by its nature, the horn antenna radiates a travelling wave, which will lead to a first-order Doppler shift of the clock transition. The microwaves are directed perpendicular to gravity, which helps to minimise the shift, but any residual transverse velocity will be an issue. For a cloud temperature of $20\ \mu\text{K}$, the fractional frequency shift is of the order $\sim 10^{-10}$. This emphasises the need for optimised cooling and a low cloud temperature. Second, it won't be easy to guarantee the alignment of the B-field at the atoms. The inner walls of the chamber

are made of polished stainless steel, which will reflect the microwaves in various directions. Thus, it is likely that the atoms will experience a small B-field component perpendicular to the quantisation axis. This allows the neighbouring ‘field-sensitive’ transitions ($\Delta m_F = \pm 1$) to be driven, leading to another frequency shift of the clock transition [89]. The microwave horn antenna is not the intended final method for interrogation, although it can still be used to perform basic spectroscopy and establish the remaining experimental apparatus for performing closed-loop clock operation.

4.7 Detection Optics

After the microwave interrogation, fluorescence detection is used to identify the number of atoms that have undergone the clock transition. The same vertical laser beam used for cooling and trapping is used to induce fluorescence in the atoms. This fluorescence is collected by a lens and is focused onto a photodetector, where it can be monitored. Figure 4.11 shows the schematic of the lens arrangement used in this experiment.




Figure 4.11: Schematic of the detection collimator with retroreflector for doubling the collection efficiency.

Fluorescence from the atoms is emitted in all directions; however, we can only collect a small fraction of this emitted light. Therefore, it is important to quantify the collection efficiency of the detection system. This is calculated by comparing the solid angle of the collection optics (which is essentially a cone) to the solid angle of an entire sphere. The cone’s solid angle is found by integrating the surface element of a unit sphere in spherical coordinates³² which gives:

$$\Omega(\theta) = 2\pi(1 - \cos\theta), \quad (4.3)$$

where θ is the half-angle subtended by the cone. For the entire sphere $\theta = \pi$, thus the ratio

³² $\int_0^{2\pi} \int_0^\theta \sin\theta d\theta d\phi$.

becomes:

$$\frac{\Omega(\theta)}{\Omega(\pi)} = \frac{1 - \cos \theta}{2}, \quad (4.4)$$

where the angle θ is calculated from the distance and diameter of the first lens collection lens. In this experiment, there are two fluorescence detection systems: one for monitoring the MOT chamber and the other for monitoring the detection chamber. In both cases, the first collection lens is placed just outside the vacuum chamber to enable maximum collection efficiency. The detection efficiencies can be improved by up to a factor of two with the use of a retroreflector.

	$f_1 = f_2 = d$ [mm]	Lens Diameter [mm]	Collection Efficiency [%]
MOT	100	25.4	0.4
Detection	60	30	1.6

Table 4.1: Specifications of the two detection collimators used in the experiment (without retroreflector).

To obtain a strong detection signal, the entire cold-atom cloud is imaged onto the photodetector. This is achieved by using lenses with the same focal length to form a 1:1 imaging system; and photodetectors with a relatively large active area of 75.4 mm^2 . The detectors also feature built-in low-noise transimpedance amplifiers with variable gain. This allows the fluorescence signals to be amplified in the range of 0 to 70 dB. In this work, we operate with a gain of 70 dB, which reduces the bandwidth to 3 kHz (for a 50Ω load).




Figure 4.12: Detection collimator attached to the MOT chamber.

4.8 Compact Physics Package

The physics package in this work is a demonstrator, and more consideration is needed to make it fully miniaturised. However, in a second iteration, there are a few things that could be done differently to ensure a smaller footprint. Currently, the size-limiting aspect of this physics package is the compensation and quantisation coils. The next iteration of this system should incorporate magnetic shields and a solenoid to define the quantisation axis. This would be both smaller and better for minimising magnetic field instabilities.

After this, off-the-shelf vacuum components would need to be replaced with a custom chamber. The vacuum vessel would also have to incorporate the microwave cavity as well. The key is to minimise the time between releasing the atoms and starting the interrogation sequence. One potential solution that might help would be to use a compact cold-atom source such as a pyramid-MOT [81, 90, 91] or grating-MOT [92, 93, 94]. This would not only simplify the trap optics to a single-beam input but also mitigate the need for splitting the beam on the table. The compact cold-atom source could be placed upside down at the top of the flight tube, and the single-beam input could enter the chamber from beneath. This type of arrangement has been used to great effect to develop a compact cold-atom gravimeter [95].

The final key area for miniaturisation is the vacuum pumping apparatus. A compact alternative is to use a passively pumped vacuum system [96], which forgoes the need for a large ion pump and the unwanted static field generated by its magnets. However, one might question the lifetime of such a system and the effects when the passive getters become saturated. Most active-ion pumps incorporate a standard vacuum flange for universal mounting capabilities. The one used in this experiment has a DN35CF flange, but there are pumps that use a DN16CF flange, which is slightly smaller. Additionally, rather than using an angle valve for connection to a bakeout system, a copper pinch-off tube would be a better solution. In summary, these are a few potential ideas that, if implemented, would help to make the overall physics package much smaller.

Chapter 5

Microwave Cavity

The ideal solution for clock state interrogation would be to use a resonant microwave cavity. Over the past 30 years, research on atomic fountain clocks has seen a variety of microwave cavity designs [97, 98, 99]. The electromagnetic field produced inside the cavity forms a standing wave where the phase is constant in the region of the magnetic field anti-node (to first-order approximation). This is much more suitable for the clock compared to the microwave horn antenna. A microwave cavity was not implemented in this experiment due to time constraints. However, alongside my DPhil, there was another project ongoing at NPL to develop a mini-fountain microwave clock. Due to the significant overlap between my DPhil work and this project, a small collaboration was formed whereby I helped to design the microwave cavity for the mini-fountain system. A design in which many of the features would be transferable to my experiment or to other compact microwave clocks. Key members of this collaboration were Sam Walby (another DPhil student), their supervisor Krzysztof Szymaniec (NPL atomic fountain group) and Peter Lovelock (from the NPL design office). We also had many valuable discussions with both Rich Hendricks and Andrew Wilson. In the following sections, I describe some of the details and work conducted under this collaboration. The discussion is restricted mainly to the cavity design, as the remainder of the mini-fountain system is the primary focus of Sam Walby's DPhil research. My contribution was to identify an initial design and validate the feasibility of certain aspects using COMSOL modelling software.

For clarity, the cavity discussed below is for an atomic fountain clock. In this arrangement, the atoms are launched vertically upwards and pass through the cavity twice, once on the way up and again on the way down. Furthermore, the mini-fountain intends to use rubidium (Rb-87) rather than caesium (Cs). Therefore, certain dimensions that depend on the interrogating wavelength will be different from those elsewhere in the thesis. In fact, the ratio between the hyperfine

ground state splitting is $\nu_{\text{Cs}}/\nu_{\text{Rb87}} \sim 1.34$. Therefore, one could expect that a cavity for a Cs-based system would be $\sim 1/4$ smaller than that of Rb. This section is concluded by highlighting the applicability of this design to my system.

5.1 Cavity Theory

A microwave cavity supports many different resonant modes. Each mode corresponds to a particular arrangement of the electric and magnetic field lines inside the cavity. The modes occur at distinct frequencies, which for a cylindrical resonator are given by [100]:

$$f_{nml} = \frac{c}{2\pi\sqrt{\mu_r\epsilon_r}} \sqrt{\left(\frac{p'_{nm}}{a}\right)^2 + \left(\frac{l\pi}{d}\right)^2}, \quad (5.1)$$

where a and d correspond to the radius and length of the cavity respectively. The mode integers n , m , and l represent the number of variations for the standing wave pattern. The variable $p'_{n,m}$ is a constant that depends on the mode, the values for which can be found in lookup tables³³. For an atomic clock, where microwave excitation drives the ‘field-insensitive’ transition ($\Delta m_F = 0$), the magnetic field of the microwaves must align parallel to the quantisation axis of the atoms. This requirement somewhat limits the choice of mode that can be used in the cavity. In almost all fountain clocks, the mode of choice is TE_{011} ; this simulated field shape can be seen in figure 5.1.

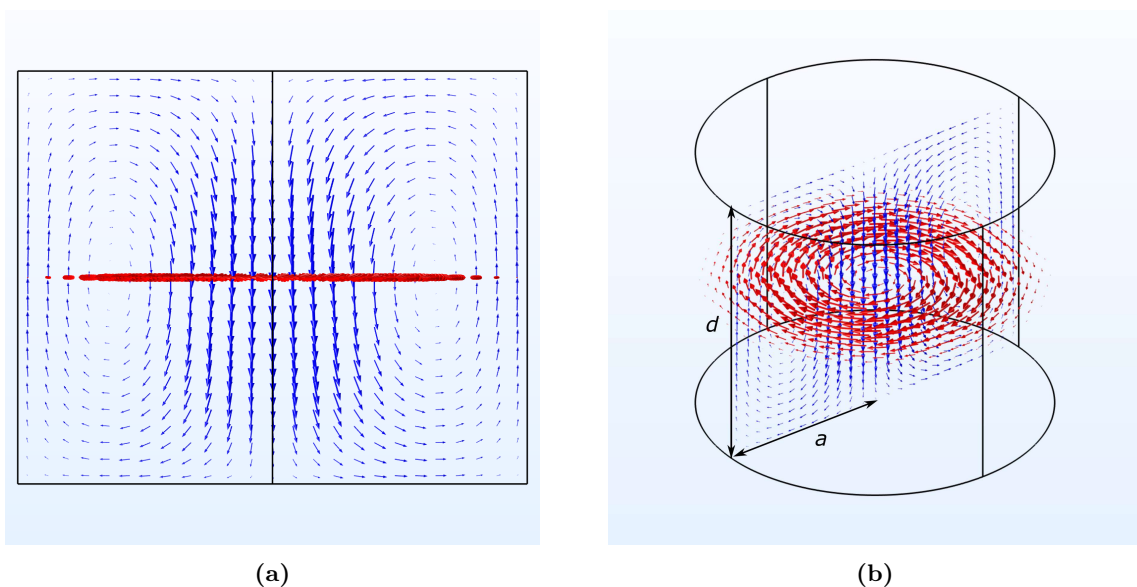


Figure 5.1: COMSOL model showing the electric field (red) and magnetic field (blue) of the TE_{011} mode. (a) 2D slice through the cavity showing the B-field. (b) 3D projection showing the looping E-field.

³³ $p'_{01} = 3.832$ [100].

This mode has a few properties that make it favourable for an atomic clock. First, across the centre of the cavity, the magnetic field lines are orientated vertically to ensure alignment with a quantisation field. Furthermore, the transverse components of the magnetic field are small, which helps to suppress 'field-sensitive' transitions. It is axially symmetric and exhibits a naturally high Q-factor. This allows holes to be introduced on the top and bottom for atom passage without too much degradation of the field inside. To appropriately excite the TE_{011} mode, it is common practice to attach two rectangular cavities on either side of the main cylindrical cavity, as shown in figure 5.3. Two small coupling holes about the midplane of the cylinder can then be used to feed microwaves between the rectangular and cylindrical cavities.




Figure 5.2

Figure 5.3: Schematic of a conventional cylindrical cavity fed by two symmetric rectangular cavities. The coupling holes about the midplane of the cylinder allow coupling between the two resonators.

A similar set of modes also exists for rectangular structures, given by equation 5.2. All of the variables have the same meaning as equation 5.1, except now a, b and d represent the three dimensions of the rectangular cavity:

$$f_{nml} = \frac{c}{2\pi\sqrt{\mu_r\epsilon_r}} \sqrt{\left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2 + \left(\frac{l\pi}{d}\right)^2}. \quad (5.2)$$

For efficient coupling between the rectangular cavities and the main cylindrical cavity, the resonators must be tuned to the same frequency, and the mode distribution across the coupling site should ideally be matched as closely as possible. For example, in an atomic fountain clock, where the mode of choice is the TE_{011} mode, and the magnetic field lines are orientated vertically,

the rectangular cavities should be excited with a mode in which the magnetic field lines are also orientated vertically across the coupling holes. A process commonly known as mode matching.

5.2 Cavity-Related Frequency Shifts

The ultimate accuracy and stability of an atomic fountain clock depend on several types of frequency-shifting phenomena. Of these shifts, some relate to the cavity and must be considered. They are discussed here.

5.2.1 Cavity Phase Shifts

Energy losses within the cavity are inevitable due to the surface resistivity of the cavity walls. The flow of this energy loss translates to a small travelling wave component that is superimposed on the standing wave. This leads to a spatial phase variation across the centre of the cavity. Atoms that experience a difference in phase between the two interactions will observe a frequency shift $\frac{\delta\nu}{\nu_0} = \frac{\delta\phi}{2\pi\nu_0 T}$. For example, a 10 μrad discrepancy translates to a fractional shift of 4.7×10^{-16} . The double-pass nature of the fountain arrangement somewhat helps to minimise this effect due to the velocity reversal³⁴. However, small misalignments in the launch trajectory and the residual transverse velocity of the atoms lead to a spread of trajectories. Therefore, the atoms do not sample the same part of the microwaves during both passes through the cavity. The resulting shift is known as the distributed cavity phase shift (DCP). To reduce phase gradients, high-Q cavities are used to minimise losses, as well as being fed by two opposing feeds that must be balanced in both phase and amplitude. The DCP shift is largely understood now with theoretical models [101, 102], and has been measured in most atomic fountains with an uncertainty of a few parts in 10^{-16} .

Separate from the DCP, there is also a dynamic cavity phase shift if the phase of the driving field changes between the two Ramsey pulses. The origin of this shift can arise from phase transients in the microwave switches used to generate the Rabi pulses. With appropriately designed switches, sub-microradian transients can be achieved [103], which leads to a frequency shift at the low 10^{-17} level.

³⁴However, for the clock arrangement in this thesis, the atoms are dropped rather than launched, meaning there is no velocity reversal. Hence, greater care would need to be taken to minimise the phase shifts between the two interaction regions.

5.2.2 Cavity Pulling

An additional area of concern is related to the tuning of the microwave cavity to the atomic resonance. A detuned cavity can lead to a frequency shift through an effect known as ‘cavity pulling’. A first-order shift arises due to the interference between the field supported inside the cavity and the field radiated by the atomic magnetic dipoles. The power emitted or absorbed by the atoms during interrogation will produce feedback to the field inside the cavity, resulting in a frequency shift [104]. A second-order shift occurs if the cavity is excited off-resonance [33]. When detuned, the atomic resonance is situated on the side of the cavity’s resonance profile, where the gradient is steepest. This effect is best visualised later in figure 5.6. The measured frequency experiences a shift due to the asymmetry of the microwave power experienced by the atoms when modulating either side of the atomic feature for locking.

When plotting the clock frequency shift as a function of cavity detuning, both effects exhibit a dispersive shape with a minimum for zero detuning and a maximum shift when detuned by half the cavity linewidth. We can consider the size of both effects by approximating for small detunings. The slope of the first-order shift becomes [104]:

$$\left. \frac{d\Delta\nu_{\text{FO}}}{d\nu_{\text{cav}}} \right|_{\nu_{\text{cav}}=\nu_0} = -8\tau_{\text{eff}} \times \frac{\mu_0\mu_{\text{B}}^2 N_{\text{at}}}{2\pi^2\hbar V_{\text{mode}}} \times \frac{Q_{\text{cav}}^2}{Q_{\text{at}}}. \quad (5.3)$$

For the TE₀₁₁ mode in the case of a Rb cavity, for $N_{\text{at}} = 10^8$, $Q_{\text{cav}} = 10^4$, $Q_{\text{at}} = 6.8 \times 10^9$, $V_{\text{mode}} = 14.5 \text{ cm}^3$, and $\tau_{\text{eff}} = 10 \text{ ms}$. The fractional frequency shift $\left. \frac{1}{\nu_0} \frac{d\Delta\nu_{\text{FO}}}{d\nu_{\text{cav}}} \right|_{\nu_{\text{cav}}=\nu_0}$ is then $6.2 \times 10^{-17}/\text{kHz}$ of cavity detuning. In the case of an atomic fountain, where the accuracy reaches parts in 10^{16} , the tuning should ideally be better than 10 kHz. The slope of the second-order shift becomes [33]:

$$\left. \frac{d\Delta\nu_{\text{SO}}}{d\nu_{\text{cav}}} \right|_{\nu_{\text{cav}}=\nu_0} = \frac{8}{\pi^2} \times \frac{Q_{\text{cav}}^2}{Q_{\text{at}}^2} \times b\tau \cot b\tau. \quad (5.4)$$

If the exciting microwave power is set within 5% of the optimum condition ($b\tau = \pi/2$), and the same values as above are assumed. The fractional frequency shift $\left. \frac{1}{\nu_0} \frac{d\Delta\nu_{\text{SO}}}{d\nu_{\text{cav}}} \right|_{\nu_{\text{cav}}=\nu_0}$ is then $3 \times 10^{-20}/\text{kHz}$ of cavity detuning. The second-order shift is smaller than the first due to the narrow resonance of the atomic feature compared to the width of the cavity resonance.

5.2.3 Microwave Leakage

Finally, a frequency shift may also occur if the atoms experience stray microwaves outside the cavity region. There are a few areas where microwave leakage can occur: the frequency synthesis circuitry, improper feeding inputs, or even from the cavity itself. Given the uncontrolled nature of such fields, this shift can be difficult to evaluate and even disentangle from other shifts such as the DCP. In most well-designed atomic fountain systems, the effects of microwave leakage can be in the region of low parts in 10^{16} [105, 106].

5.3 Cavity Design

When developing a cavity for a compact system, there are a few factors that must be considered. One of which is the size of the cavity. Any effort to make the cavity system smaller would help reduce the overall size of the clock. However, this is limited by the physics of the cavity. The characteristic size is typically dictated by the length of the interrogating radiation, which is fixed at $\lambda_0 = 43.9$ mm (for 6.835 GHz in Rb-87). Furthermore, a compact clock must be able to operate outside of a laboratory environment, where temperature fluctuations can influence clock stability. For a cavity, this can be crucial since the dimensions define the resonant frequency. Additionally, a compact clock should be able to operate continuously for long time periods without the need for user intervention or re-tuning. With these factors in mind, we have incorporated a few key design elements. A computer-aided drawing (CAD) model³⁵ of the designed cavity can be seen in figure 5.4 and 5.5.

The microwave cavity is made from oxygen-free copper (OFC) and comprised of several different components, which all seal together to form part of the total vacuum vessel. Conventional cavity systems are usually attached to the underside of a vacuum flange, which is then placed inside a vacuum straight. There are two main benefits to the approach taken in this work; there is no need for additional vacuum feedthroughs for the microwave input to the cavity, and the overall radial footprint is reduced. The main Ramsey cavity is resonant with the TE_{011} mode and is formed from a cylindrical body (radius 30 mm and height 48.62 mm) and two endcaps. When fully assembled, the endcaps protrude into the main body with a slightly smaller diameter to form RF choke grooves (width 0.89 mm and depth 6.72 mm). The chokes are necessary to shift the unwanted degenerate TM_{111} mode [107]. Using COMSOL simulations, the shift was

³⁵Initial schematics were drawn by myself; subsequent revisions were drawn by Peter Lovelock to reduce the amount of material used.




Figure 5.4: CAD model of the cavity. (a) Vertical cross-section. (b) Axial cross-section, both RF windows can be seen as well as the feeding antenna into the ring waveguide. The zoomed inset shows the top right choke groove. The vertical seam is an artefact of the CAD software and does not represent two separate parts.

calculated to be ~ 340 MHz to lower frequencies.

There are holes in both endcaps to allow atom passage through the cavity. The dimensions of the holes extend above and below the cavity, acting as below cut-off waveguides to ensure sufficient isolation of the microwaves in the remainder of the flight tube. In addition to microwave isolation, the hole diameter also determines the final temperature that the cloud must reach after laser cooling. The extent of the cloud's ballistic expansion is determined by the final temperature when launched. The diameter of the holes was chosen to be 16 mm, a balance between providing sufficient microwave isolation and to ensure a high number of returning atoms. For a typical launch trajectory of 500 ms and a hole diameter of 16 mm, the cloud must be cooled to a temperature of $\lesssim 6 \mu\text{K}$. This is approximately three times higher than the temperatures seen in most atomic fountain systems [108]. Therefore, with sufficient optimisation, a similar performance would be expected here.

To feed microwaves into the main Ramsey cavity, there is a section of bent rectangular waveguide that forms a ring ‘hugging’ the main body. Two opposing holes about the midplane of the Ramsey cavity are used as microwave feeds from the ring waveguide. Compared to conventional fountain cavities, the feeds are also provided by two coupling holes but instead using two separate rectangular waveguides (see figure 5.3). Each waveguide has its own separate decoupled microwave input. Tuning of the phase and amplitude to each must be performed to balance the feeds and generate a standing wave in the cavity. In the ring waveguide arrangement, there is

only one microwave input that couples to the two feeds. This simplifies the design by eliminating the need to perform any balancing; however, there will be losses that compromise this somewhat leading to a DCP shift. Although, with precise manufacturing and geometry of the ring, this can be managed.




Figure 5.5: Exploded view of the cavity assembly with different parts labelled.

The TE_{011} mode is favoured for atomic fountain clocks due to its intrinsically high Q-factor ($\sim 25\,000$ to $30\,000$). A higher Q-factor results in smaller losses and better control over the DCP shift. However, a high Q-factor also leads to a high sensitivity to environmental fluctuations, e.g. the typical sensitivity coefficient for a conventional TE_{011} cavity made from OFC in a Rb fountain is $-114\text{ kHz}/^\circ\text{C}$ [109]. There are two methods to minimise the negative effects of temperature sensitivities. The first is to use a temperature stabilisation system to maintain resonance. The second is to intentionally operate the cavity with a lower Q, as demonstrated in figure 5.6. In the event of a temperature change, the cavity can become detuned, shifting the profile away from the atomic resonance. This leads to a change in the microwave power delivered to the cavity,

and thus, the atoms experience a non-optimal microwave pulse area. This is more manageable in the lower Q-factor cavity, as the gradient on the side of the slope is less steep.




Figure 5.6: Schematic showing a high-Q and low-Q cavity. Both cavities are initially tuned to provide the optimum microwave power P_{opt} to the atoms. When detuned, the atomic resonance is situated on the side of the profile and the atoms experience a different microwave power when the frequency is modulated for locking. The gradient of the profile dictates the discrepancy of the power delivered to the atoms. The low-Q cavity is less sensitive to potential detuning that may arise from environmental fluctuations.

In this design, the cavity is operated with a low loaded Q-factor (\sim few thousand) to reduce the environmental sensitivity and eliminate the need for temperature stabilisation of the cavity. This is much more suitable for a compact clock that would be expected to operate in environments where temperature stability cannot be guaranteed. The Q-factor of the Ramsey cavity is tuned by varying the size of the coupling holes from the ring waveguide.

The microwave cavity forms part of the total vacuum vessel, which means that the ring itself and the microwave feed into the ring need to be vacuum-tight. To simplify this, we decided to implement RF windows, which are placed over the feeding holes between the ring and the Ramsey cavity. The dielectric material of the windows acts as a physical barrier to form a vacuum seal, but they are also transparent to microwaves to allow feeding. The microwave input to the ring is provided by a single in-line loop antenna, which protrudes vertically from the roof of the ring waveguide. The antenna is formed from a piece of semi-rigid coaxial cable with the shield removed, leaving an exposed length of conductor protruding into the ring waveguide (see figure 5.5). The conductor makes a 90° bend and is shorted to the inner wall of the main cavity body. In this arrangement, the appropriate coupling to the TE_{106} mode is excited in the ring. We decided to use this type of feed to save space as the alternative would have been a right angle $\lambda/4$ antenna that enters the ring radially. A radial antenna was not a suitable option as this would intrude on the space already taken up by the quantisation coil that closely surrounds the microwave cavity.

5.4 Cavity Seals & Joins

Given that the cavity is comprised of several parts, there are a number of joins and seals that make the assembly partially complicated. Fortunately, there are a few design elements that help simplify this. An area of concern was the mating of surfaces that define the dimensions of the Ramsey cavity. In particular, the two endcaps that ultimately dictate the cavity length must also act as vacuum seals. Indium seals were considered, but there was uncertainty as to whether the unpredictable flow of pressed indium would allow for a flat mating between the endcap and main body surfaces. Discrepancies in the join could lead to detuning in the Ramsey cavity. We opted to use a simple bolted join for the lower endcap, which is then enclosed by another flange to form the vacuum seal. The enclosing flange is sealed using pressed indium to the main body of the cavity, i.e. the lower endcap is fully enclosed in the vacuum.

The upper endcap is electron-beam welded (EBW)³⁶ to the main body to form one single part. The most significant benefit is that the weld is localised only to the joining region, which results in low distortion of other bulk material. The final weld is very strong and suitable for maintaining UHV. EBW is performed after the initial tuning of the cavity is complete. Also, part of the upper endcap is a custom vacuum flange, which can be attached to the remaining flight tube using an indium seal. Fortunately, the radius of the cavity is defined by the machining of the main body and is thus not at risk of being compromised by an inadequate join or seal. The main body of the cavity is attached to a lower flange that connects to the rest of the vacuum vessel. This is a pressed indium seal for maintaining vacuum and encloses the lower endcap of the cavity as described previously. There are clearance holes in the body of the cavity, and the bolts thread into two semicircular titanium discs, which act like plate nuts to provide even pressure when tightening (a design feature added by Peter Lovelock, see figure 5.5).

Another key area is the attachment of the RF windows to the main cavity body. In this design, the windows also act as a vacuum seal to relax the requirements of the feed and seals of the ring waveguide. There was an interplay here between the window material and the type of join that could be used. An initial suggestion was to use metallized ceramic windows that could then be brazed to the cavity body. However, after consideration, it was decided that the material mismatch would cause issues in the brazing process. Instead, we investigated the possibility of using optical windows, as seen in [11]. A practical study was performed (by Sam

³⁶EBW is a fusion welding process in which a tightly focused high-energy electron beam is fired towards two parts to be mated. The kinetic energy of the beam is transferred to heat, which melts the two parts together to form the weld.



Figure 5.7: *Electron-beam weld on a test piece for the upper endcap. The location of the weld can be seen by the burred edge. Notably, there are no signs of the weld in the cross-section.*

Walby) to identify the best combination of window material and join type. The results suggested that a sapphire glass window and vacuum epoxy combination was the most appropriate option. Other joins included pressed indium and pressed C-solder³⁷. It was important to ensure that the difference in the coefficient of thermal expansion (CTE) between copper ($17.3 \times 10^{-6}/^{\circ}\text{C}$) [110] and the window material was minimised; otherwise, this could compromise the vacuum integrity during bakeout. Despite the factor of three difference in CTE for sapphire ($5 \times 10^{-6}/^{\circ}\text{C}$) [111], we still opted to use it, as it was the closest match of all materials tested.

The remaining joins are the cover for the ring waveguide and the feed into the ring (due to the RF windows, these seals are more straightforward); they include pressed indium to act as an electrical seal and prevent microwave leakage.

5.5 Cavity Simulations

COMSOL simulations were carried out to aid the decision-making process and ensure that the microwave field within the cavity would not be perturbed in any unwanted manner. The COMSOL simulations use either an eigenfrequency solver or a frequency domain solver. An eigenfrequency solver is useful for determining the resonant frequencies (or modes) for a given geometry, whereas the frequency domain solver studies the geometry response for a fixed user-defined frequency. Both were used across several simulations when designing the cavity.

The ring waveguide structure has already been investigated for a cavity designed at the Physikalisch-Technische Bundesanstalt (PTB); a large amount of theory is discussed in [112].

³⁷A particular type of solder capable of wetting to ceramics and glass.

We used COMSOL simulations to visualise the field distribution inside the ring waveguide. The primary goal of the ring waveguide is to feed microwaves into the Ramsey cavity, exciting the desired TE_{011} mode. For this to be achieved, there must be some overlap with the mode distribution across the feeding site. When the ring waveguide is excited with the TE_{106} mode (shown in figure 5.8), there is a symmetry at the coupling sites where the magnetic field lines are orientated vertically. This vertical orientation is ideal for coupling to the TE_{011} mode as the field lines should also be aligned vertically in the cavity. To correctly establish the TE_{106} mode in the ring waveguide, the height and length (determined by the average radius) of the ring must be tuned correctly. Note that the TE_{106} mode has the mode integer $m = 0$, which means that, in theory, the width of the ring should not have an effect on the resonant frequency. However, it does affect the Q-factor of the ring. As pointed out in [112], the length of the ring must be $3\lambda_g$ to achieve the necessary symmetry at the coupling sites. Here, λ_g is the guide wavelength, the distance between two equiphase planes in the guide. The final condition for TE_{106} excitation is the correct placement of the coupling antenna, which must be displaced by $\lambda_g/4$ around the ring from the symmetry plane. This coincides with a node of the electric field.

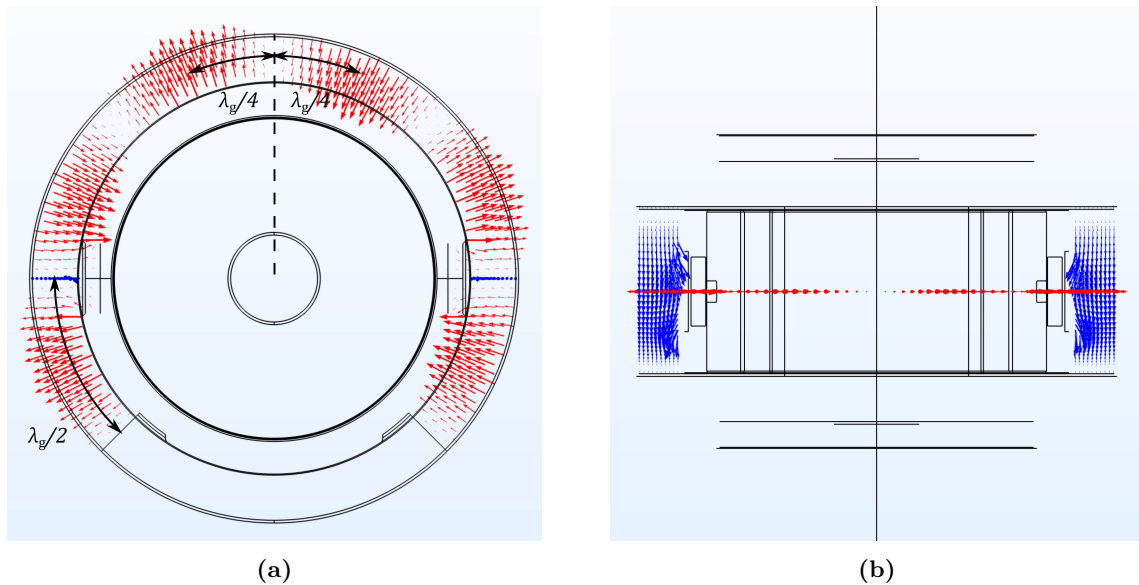


Figure 5.8: COMSOL model showing the electric field (red) and magnetic field (blue) of the TE_{106} mode inside the ring waveguide. (a) Cross-section showing the top-down view of the cavity and ring waveguide. (b) Cross-section showing the side view.

One of the largest unknowns with this design was the effect of the RF windows on the field distribution inside the ring and Ramsey cavity. It was suspected that this might cause a frequency shift of the resonances in both structures, but it was unclear how large this effect would be. Both the ring and the Ramsey cavity can be tuned after manufacture by removing cavity material from certain surfaces. This is something we envisaged having to do when first

receiving the cavity, as there might be some inaccuracies in the manufacturing process. However, if the shifts introduced by the RF windows were too large, then it might not be possible to appropriately tune the resonance back by removing material alone. We conducted a COMSOL study to understand whether this would be an issue. Initial simulations for both a simple ring and a cylindrical cavity were performed to identify a baseline resonance. Then, finer geometric details such as coupling holes and RF windows were incorporated to see how each modification would shift the frequency. Table 5.1 shows the effect of each geometric detail and the induced shift.

#	Geometry / Modification	Ring Waveguide		Ramsey Cavity	
		Frequency [MHz]	Shift [MHz]	Frequency [MHz]	Shift [MHz]
1	Simple Ring	6835.294	0	-	-
2	Ramsey Cavity	-	-	6830.358	0
3	2 + atom passage holes coupling holes + RF chokes	-	-	6829.208	-1.150
4	1 + 3	6833.684	-1.610	6829.208	0
5	4 + window cut-outs	6818.076	-15.608	6829.122	-0.086
6	5 + RF windows	6812.090	-5.986	6829.122	0
Total Shift			-23.204		-1.236

Table 5.1: Results from COMSOL simulations showing the effects of successively modifying the cavity structure to incorporate the RF windows and the frequency shifts induced. The first column labels the geometry, i.e. 1+3 means the geometry from 1 and 3 together. Lines 1 and 2 differ slightly from equations 5.1 and 5.2, this was never fully understood. Atom passage holes $\phi = 10$ mm, coupling holes $\phi = 4$ mm, RF windows $\phi = 10$ mm. Until line 6, everything is assumed to be vacuum-filled. Line 6 has the Ramsey cavity in vacuum and the ring in air.

The results show that the ring waveguide is most affected by the modifications, while the Ramsey cavity is less so. It would appear that the cut-out around the coupling hole for the window has the largest effect. To see if these shifts are manageable, we can compare them with the tuning sensitivities for each dimension shown in table 5.2. These can be found by differentiating equations 5.1 and 5.2. Shifts on the order of tens of MHz should be easy to compensate, given that typical machining tolerances go down to 0.1 mm. In conclusion, it was decided that the RF window approach would be manageable.

The next crucial aspect is the feeding antenna into the ring waveguide. Positioning the antenna is critical to ensure the correct coupling to the desired TE_{106} mode. Conventional feeds

Resonator	$\partial f/\partial a$ [MHz/mm]	$\partial f/\partial d$ [MHz/mm]
Ring waveguide	-113.5	-17.8
Ramsey cavity	-181.4	-28.6

Table 5.2: Tuning sensitivities with different dimensions in the ring and Ramsey cavity.

use a single $\lambda/4$ antenna that protrudes into the waveguide. In this design, we decided to use an in-line loop antenna to help reduce the radial footprint of the cavity whilst still achieving the desired coupling. Both types of antennae are illustrated in figure 5.9.




Figure 5.9: COMSOL frequency domain simulations highlighting the two main types of coaxial to waveguide transitions. In both cases, the same mode is excited with negligible change in coupling efficiency. The colour map shows the vertical component of the electric field. (a) Right-angle transition. (b) In-line transition.

The COMSOL simulation suggested that this approach would be feasible. However, to validate this, we modified an old Cs fountain cavity (with two separate rectangular feeding waveguides) to test the in-line loop antenna. A hole was drilled on the top of one of the waveguides (90° displaced from the original feed) and a loop antenna was inserted. By monitoring the S11 parameter (reflected power) on a network analyser, a coupling efficiency of ~ 18 dB into the waveguide was achieved. Compared to the unmodified waveguide, which had ~ 30 dB coupling. The coupled power was susceptible to the positioning of the loop antenna, which could not be achieved precisely by hand, hence the possible limit of ~ 18 dB. Furthermore, the hole drilled on the top of the waveguide was slightly displaced from the plane of the original input; this could also hinder the coupling. For the in-line loop antenna to work, it must be electrically

shorted to the inner wall. In this test, the antenna was only pressed (lightly touching at best) against the wall, and it is likely that the electrical connection was somewhat compromised. We decided to incorporate two design features for the real antenna. First, a small locating hole is drilled in the inner cavity wall to ensure a better electrical connection and to guide the antenna into position. Second, to make the antenna position more mechanically stable, the semi-rigid coax is extended above the cavity and secured alongside the flight tube. In this arrangement, disturbances from attaching and detaching the microwave input to the terminated end of the coax lead are minimised.





Figure 5.10: (a) Old fountain cavity system connected to the network analyser. (b) Modified waveguide with an in-line loop antenna. The position of the original antenna is shown for reference.


5.6 Cavity Summary

The cavity design outlined above is best suited to an atomic fountain clock arrangement where the atoms traverse the cavity twice along their trajectory to achieve Ramsey interrogation. This section discusses the applicability of the cavity (if it were it to replace the horn antenna) to the cold-atom microwave clock developed for this DPhil work. The time it takes for the atoms to fall through the cavity will set the limit on the achievable interrogation time. Since the atoms must first be prepared in the magneto-optical trap (MOT) and fall into the cavity, they will have accrued some initial velocity. The potential Rabi time τ_{Rabi} will be reduced for a larger initial displacement above the cavity. Figure 5.11 shows this dependence of potential interrogation time vs. initial displacement.

Currently, the distance from the centre of the MOT chamber to the point where the atoms enter the interaction region is ~ 39.1 mm, which limits the potential interrogation time to



(a)



(b)

Figure 5.11: (a) Potential interrogation time vs. initial displacement of the cloud above the cavity. The orange dashed line represents the minimum distance before the atoms can be interrogated in the current experimental arrangement. (b) The final manufactured body of the Ramsey cavity in the lab before characterisation measurements.

~ 44 ms. At first glance, this is only a marginal increase over systems that employ a much simpler and smaller ‘release and recapture’ interrogation. These predictions emphasise the importance of commencing the clock state interrogation as soon as the atoms are released from the MOT. This would involve amalgamating the MOT region and microwave cavity into one chamber [11]. If such a chamber could also be extended to incorporate a Ramsey interrogation scheme, this would be an attractive solution. One possible approach (as alluded to at the end of chapter 4), would be to place a pyramid-MOT upside down at the top of a cavity resonator. However, more investigation would be needed to ensure the pyramid structure does not disturb the field inside the cavity. In such an arrangement, the atoms could be probed immediately after they are released, fall for ~ 200 ms, and then be probed again before detection.

In summary, while the cavity designed above might not be directly applicable to this cold-atom clock prototype, it does have a few useful design features that could be transferred. First, the low Q-factor of the Ramsey cavity, which reduces sensitivity to temperature fluctuations would be desirable. Furthermore, both a single microwave feed and the use of RF windows would help to simplify the overall design. The cavity was manufactured and delivered while I was writing my thesis. Details of the physical implementation and characterisation process will be presented in Sam’s thesis³⁸.

³⁸In preparation at the time of publication.

Chapter 6

Electronics and Experimental Control

This chapter details the electronic hardware and software used to control the experiment. Compared to the physics package and laser system, the electronics typically do not take up much of the footprint for an atomic clock. Today's modern electronics have reached remarkable levels of compactness. This is readily seen in any smartphone, which has considerable functionality and enormous computing power, all in the size of something that fits in your pocket. This miniaturisation can be attributed to the continuous advancements driven by Moore's law, which states that the number of transistors in an integrated circuit (IC) doubles every two years. At first glance, the electronic requirements for an atomic clock are relatively simple compared to something like a smartphone. Thus, given appropriate design and construction, one would expect the overall footprint of the electronics could be made far smaller. Consequently, researchers typically focus their efforts on trying to miniaturise the physics package or laser system. Nevertheless, in this work we have taken the time to design some custom electronics and software to drive the experiment. This is also due to the fact that off-the-shelf solutions are few and far between. The many different aspects and components of this design are discussed in the following sections.

6.1 Hardware Overview

The development of electronics for an atomic clock can be approached in a few different ways. Conventional laboratory-based experiments use a desktop PC and control software such as LabVIEW (National Instruments) to interface with many pieces of benchtop equipment. This method is quick to set up and offers a user-friendly graphical interface. However, the more pe-

peripherals a LabVIEW-controlled experiment must communicate with, the more complex and cumbersome the system becomes. Therefore, this approach is best suited to static experiments and cannot be applied outside of a laboratory environment. For compact clock developments, an architecture that uses an embedded system is a much more suitable approach. An embedded system is a combination of hardware and software to perform a specific dedicated task. They are usually comprised of three main components:

1. Hardware – a small system-on-chip (SoC) such as a microprocessor that acts as a central processing unit (CPU) and interfaces with other peripherals such as analogue-to-digital converters (ADC) or digital-to-analogue converters (DAC).
2. Software – dedicated software for running a well-defined application.
3. Real-time operating system (RTOS) – an operating system that supervises the application software and defines the rules of execution (not always necessary in smaller systems).

A field-programmable gate array (FPGA) is a popular choice for the hardware aspect of an embedded system and has been used to great effect in many atomic physics experiments [113, 114, 115]. FPGAs are capable of parallel processing, which means that different parts of the chip can be programmed to perform dedicated tasks. These tasks can then run independently without having to share resources. The power of parallel processing is most apparent in applications where timing is critical. For example, an FPGA can appropriately respond to external events such as interrupts without compromising the timing of the main thread running the experimental sequence. Despite their advantages, they are not without drawbacks. A lot of training and technical expertise is required to program them; they tend to consume a lot of power and can be quite expensive. In contrast, a similar but slightly more straightforward approach is to use a microcontroller (MCU). Unlike an FPGA, most microcontrollers are not capable of parallel processing and typically can only operate a single thread at a time. However, they are cheaper, similar in size, consume less power, and are easier to program. For this reason, we have opted to use an MCU in this work. MCU-based systems have been demonstrated with: laser locking [116], laser spectroscopy [117], gravimeters [118], and even atomic clocks [119].

Figure 6.1 shows the architecture of the embedded system used to control this experiment. The system is managed and integrated using an MCU, which communicates with several other evaluation boards and custom printed circuit boards (PCB). A typical experimental sequence consists of laser cooling the atoms, microwave interrogation, and state detection. To perform such

a sequence, the electronics must be able to: turn the lasers on and off, switch magnetic field coils, sweep frequencies to modulate lasers, and read fluorescence signals from detectors. Management of all these tasks requires three types of signals: digital, analogue, and radio frequency (RF). Digital logic is handled by the MCU, whereas the other incorporated PCBs handle the analogue and RF aspects. Together, they make an embedded system capable of driving the atomic clock. Figure 6.2 shows the final integrated system.

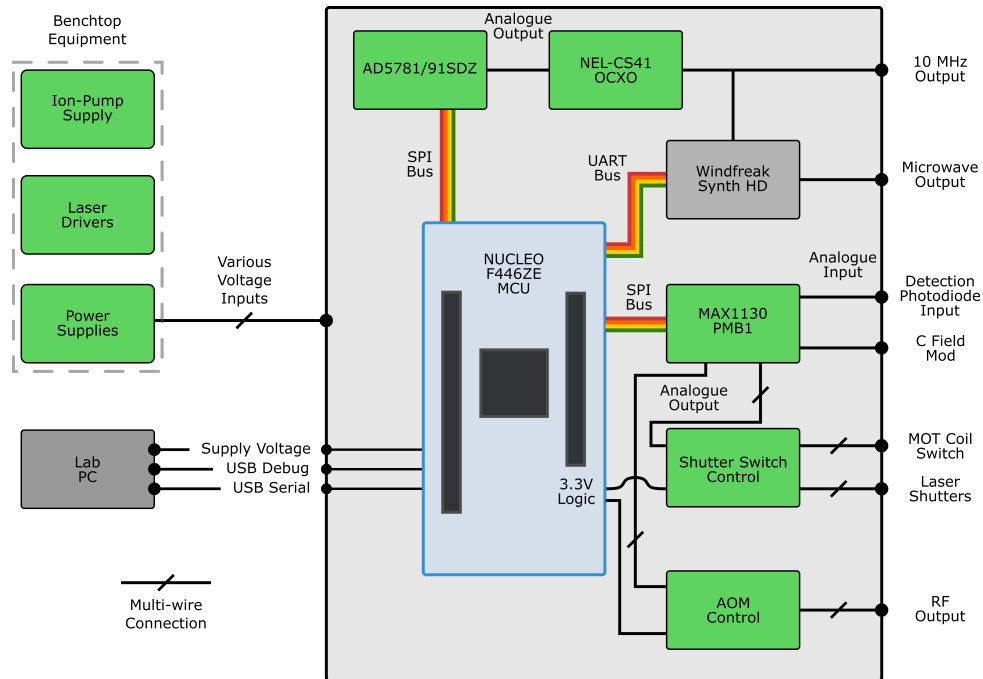




Figure 6.1: Architecture of the experimental electronics.



(a)



(b)

Figure 6.2: (a) Earlier state of the electronic hardware mounted in an enclosure. (b) Front panel of the electronic hardware where multiple input/output signals are routed.

The electronics for driving the lasers have already been discussed in section 3.4. This area of the experiment is driven using benchtop equipment, since there are readily available commercial units that provide adequately low noise. The remaining sections provide more details about each of the components within this architecture.

6.1.1 Microcontroller

At the heart of the electronics architecture is a NUCLEO-F446ZE MCU evaluation board. This MCU is responsible for housing the experimental code, interfacing with other parts of the system, and communicating with the lab PC. The NUCLEO-F446ZE eval-board was chosen due to its low cost and also because it is ‘Mbed enabled’, allowing it to be programmed easily using the Mbed Studio interactive development environment (IDE). Moreover, it has several specifications that are useful for the experiment. The onboard clock allows CPU speeds up to 180 MHz for executing large sections of code quickly. To interface with other boards and peripherals, there are up to 4 serial peripheral interfaces (SPIs) and 4 universal asynchronous receiver/transmitter (UART) communication buses available. Simple digital logic signals can be implemented using one of the 114 input-output (I/O) ports. The logic levels are only 3.3 V compatible, which is a minor limitation, but this was not an issue for this project as higher voltages were not needed. There is a large onboard flash memory of 512 kB, which stores the experimental code and various values for precomputed sweeps and ramps. The cortex M4 processor features a cycle count register, which is part of the data watchpoint and trace unit. This is used to monitor the length of time for code execution. Precise delays can be implemented by converting a delay time to a corresponding number of processor cycles and then waiting in a loop until the correct number of cycles has been completed. This is used to achieve accurate milli-, micro-, and nano-second delays [120]. Finally, the MCU connects to the lab PC via a single USB-micro connection to power the board, load the experimental code, and access the ST-LINK/V2-1 debugger tool onboard.




Figure 6.3: *NUCLEO-F446ZE evaluation board.*

6.1.2 ADC/DAC Board

Given that this experiment is in the development stage, where requirements can often change suddenly, we decided to use the MAX11300PMB1 evaluation board (or MAX board for brevity) to handle all the analogue output and input capabilities. This board offers 20 configurable mixed-signal ports supporting either a 12-bit analogue-to-digital converter (ADC), a 12-bit buffered digital-to-analogue converter (DAC), or general-purpose digital I/O. This board handles all the ADC and DAC functionality in the experiment that is not covered by the NUCLEO-F446ZE. Components such as voltage-controlled oscillators (VCOs), RF switches/attenuators, and detectors commonly operate within the range of 0 to 10 V. This was easily accommodated, given each port on the MAX board can support numerous voltage range options: 0 to 2.5 V, -5 to 5 V, 0 to 10 V, -10 to 0 V. To update the DACs or read ADC values, there is a 4-wire SPI interface that supports speeds up to 20 MHz. Although, realistically this is only as good as the sampling rate for the ADC and the update rate of the DAC, which were 200 kS/s and 2.7 kHz respectively³⁹.




Figure 6.4: MAX11300PMB1 evaluation board labelled with the significant channels. Red denotes controls for the repumping light, while blue denotes controls for the cooling light.

The MAX board requires only one 3.3 V voltage supply, although a minor issue was encountered relating to one of the large inductors mounted on the board. The proximity of the inductor to the MAX11300 IC caused EMI (electromagnetic interference), which translated to noise in the detection signal. This is discussed in more detail in section 7.3, but this resulted in a minor modification to the board to remove the inductor and another voltage regulator. This meant

³⁹The sampling rate for the ADC can be programmed to one of four values: 200 kS/s, 250 kS/s, 333 kS/s, or 400 kS/s. The DAC operates with an update rate of 40 μ s per port, but because the ports are cycled through sequentially, the resulting update rate will depend on the number of DAC-configured ports.

that the board had to be supplied by both 3.3V and 13V separately from a benchtop power supply.

6.1.3 AOM Control Board

As part of the laser and optoelectronics system, three acousto-optic modulators (AOM) were used. The AOMs provide a frequency-shifting capability and act as optical attenuators/switches. The AOMs are driven by a custom PCB that generates the necessary RF signals (shown in figure 6.5) without the need for a bulky benchtop RF source. The board consists of a voltage-controlled oscillator (VCO) with an output range of 70 to 120 MHz. Nominally, the output is 80 MHz, as this is the centre frequency specified by the AOM. After the VCO, the RF is directed through a voltage-controlled attenuator/switch, which can provide up to 47 dB of RF attenuation. To achieve a high diffraction efficiency, the AOM requires an RF driving power of ~ 30 to 33 dB. This property typically makes it unfavourable to use AOMs for compact atomic clocks, where power consumption is a key factor in the SWaP metric⁴⁰. This aspect needs more consideration if further miniaturisation were to be made. After the AOM control board, the RF signal is amplified by an in-line amplifier (Mini-Circuits ZHL-1-2W+) before being directed to the AOM. A single 12V source is required to power the AOM board, and this provides two channels to drive two AOMs. To sweep the laser frequency and intensity, two control voltages are provided from the MAX board. The tunability with these control voltages is shown in figure 6.6.




Figure 6.5: The control board for the AOMs with the other RF components that are packaged individually. The RF switch (ZASWA-2-50DRA+) is used before the AOM to switch the cooling light on and off quickly.

An additional in-line switch (Mini-Circuits ZASWA-2-50DRA+) was used on the path of the trap AOM for the cooling light. Initially, the cooling light was switched by the attenuator on the AOM control board. However, the functionality of the MAX board led to an uncontrollable variable time delay between the trigger signal for switching the light off and the actual turn-

⁴⁰Size, Weight and Power.

off time. This had a severe effect on the signal-to-noise ratio of the detection, as discussed in more detail in section 7.3. The additional switch is controlled by a digital logic signal from the MCU, which solves this timing issue. The switching speeds of the AOMs were measured by monitoring the response on a photodiode, which can be seen in figure 6.7. In the optical setup for the repumping light generation, the AOM requires a slightly different VCO frequency as the necessary shift is much larger (nominally 201 MHz). A minor modification was made to one of the AOM boards to incorporate a VCO with a higher output frequency range.




Figure 6.6: Cooling and repumping AOM characterisation. (a) VCO control voltage vs. output RF. The dashed line represents a linear fit to the data, whose coefficients can be seen in the legend. The units are MHz/V and MHz for the gradient and y -intercept respectively. (b) Diffraction efficiency of the AOM double pass systems.




Figure 6.7: AOM characterisation. (a) TFAS-1+ control voltage vs optical power after the double pass arrangement. Extinction measurement values below -60 dB are limited by the noise floor of the photodiode used in the test setup. (b) Switching speeds of the AOM as an optical shutter. The switching times are defined by the 90% - 10% thresholds.

6.1.4 Voltage Switch Board

In addition to using the AOMs to switch the intensity, we also use conventional mechanical shutters to ensure complete extinction. However, the digital logic signals provided by the MCU

do not provide enough current to drive the solenoid in the mechanical shutter directly. Therefore, the shutters are supplied using a different voltage (nominally 7 V) routed through a custom PCB (shown in figure 6.8) to provide the switching capability. The board has two output channels for both the cooling and repumping light, which can be toggled independently from digital logic signals provided by the MCU.




Figure 6.8: (a) Voltage switch board used for switching the mechanical shutters and the current to the MOT coils. (b) Mechanical laser shutter.

The switching speeds of the mechanical shutters were measured to be approximately two orders of magnitude slower than the AOMs (several 100 μs), which is comparable to the length of the detection pulses. Instead, we rely solely on the AOMs for shuttering capabilities. Although, if not completely extinguished, any remaining resonant light will induce an undesired frequency shift due to the AC Stark effect [121]. An evaluation of this shift was not done in this work, but it has been shown in several fountain clock arrangements that the effect can be suppressed to a low 10^{-16} level [106, 122]. Therefore, with sufficient isolation and appropriate detuning, a comparable performance would be expected here.




Figure 6.9: (a) Switching speeds of the mechanical shutter. The switching times are defined by the 90% - 10% thresholds. (b) Switching speed of the MOT coils as measured by a fluxgate magnetometer.

In addition to voltage sources, there is one current source to drive the C-field coils. This is provided by a benchtop laser diode driver (SRS model. 525). One of the ports on the MAX board is used as a DAC to control the external modulation input on the driver. The transfer function is 100 mA/V, and the voltage is switched between 0 V and 0.5 V to turn the C-field on and off.

6.1.6 Local Oscillator & Microwave Synthesis

One of the key components of any microwave atomic clock is the local oscillator (LO). The LO is a continuous voltage oscillator whose output frequency is periodically referenced to the atomic transition frequency. In the event of any drift, a correction is made to pull the LO output back to the atomic reference. Careful consideration must be given to the choice of LO used for the atomic clock. In a compact clock, the LO should ideally be small, have low power consumption, and be insensitive to environmental fluctuations. Furthermore, if the LO exhibits significant instabilities on timescales shorter than the experimental cycle time, this is likely to limit the final $\tau^{-1/2}$ performance via the Dick effect [124], which is caused by the frequency fluctuations of the LO that occur during the dead time of the clock cycle (when the atoms are not being measured). Clearly, these fluctuations cannot be corrected since the atoms never experience them. For an atomic fountain clock with a cycle time of 1 s, excess noise from a good quartz oscillator typically limits the frequency stability to $10^{-13} \tau^{-1/2}$ [125]. Therefore, larger laboratory-scale clocks like atomic fountains, use better LOs such as cryogenically cooled sapphire oscillators (CSO) [126], or better yet, optically stabilised microwave sources [127, 128]. However, neither of these solutions is suitable for compact developments due to their size and cost. In this system, where the expected quantum projection noise (QPN) limited short-term stability is in the region of low 10^{-12} to mid $10^{-13} \tau^{-1/2}$, this issue should not be a problem. However, it is important to note that the ongoing advances in compact microwave clock developments may soon reach the practical limit set by the availability of high-performance (10^{-13} at 1 s) quartz oscillators. Performance beyond this limit can only be achieved by improving quartz oscillator technologies or miniaturising some of the solutions used in larger laboratory-scale clocks.

The LO used in this experiment is an OCXO (NEL O-CS41-0T59IW-N-U-10-R). In an OCXO (oven-controlled crystal oscillator), the crystal and operating circuitry are housed within an oven designed to provide enhanced temperature stability. This places OCXOs at the upper end in terms of performance compared to other types of crystal oscillator arrangement (e.g. TCXOs, VCOs and XOs). The frequency stability of the free-running OCXO can be seen in figure 6.11.




Figure 6.11: (a) Tuning measurement of the OCXO. The coefficients of a linear fit around the zero crossing are shown. The inset shows a zoom of the zero crossing. (b) OADEV of the free-running OCXO. The error bars approximate the one-sigma confidence interval, $\sigma_y/\sqrt{N/M}$, where N is the number of frequency record samples that contribute to each bin and $M = \tau/\tau_0$ is the bin's τ_0 multiple.

To measure the stability, the 10 MHz output from the OCXO was sent to a phase noise analyser (Microchip 3120A) where it was compared against one of the UTC(NPL) hydrogen maser signals. The OCXO exhibits a short-term stability of $\sim 2.5 \times 10^{-13}$ at 1 s, but then starts to drift with longer averaging times. There is a noticeable ‘hump’ in the ADEV plot around 13 s whose origin is unknown. This could be due to environmental instabilities within the lab. There is an additional tuning input to the OCXO that can be used to fine-tune the 10 MHz output and to apply corrections from the clock’s servo loop. The tunability of the OCXO was characterised by measuring its output on a frequency counter (referenced to a maser signal), which can be seen in figure 6.11. A dedicated 18-bit DAC provided by an evaluation board (AD5781/91SDZ) is used to control the tuning voltage for the OCXO. Updating the DAC output is achieved via SPI communication from the NUCLEO-F446ZE MCU. Upon initialisation, the voltage is set to -0.118 V to correspond to a 10 MHz output.

The output of most LOs is in the RF domain (typically 10 MHz) and must be multiplied up to the GHz region to match the atomic transition frequency (9.192 GHz for caesium). This is achieved using a commercially available compact frequency synthesiser (Windfreak synth HD); a dual channel synthesiser that can provide up to 15 dBm output power over the frequency range 10 MHz to 15 GHz. Both aspects can be controlled via a UART communication bus between the synthesiser and the NUCLEO-F446ZE MCU. The OCXO acts as an external reference for the Windfreak, which now provides two methods for tuning the microwaves. The first is directly via the UART communication, which is used to modulate on either side of the atomic feature for locking. The second is indirectly by tuning the 10 MHz of the OCXO, which is used to

send corrections. Figure 6.13 shows the entire frequency chain, including the OCXO and the Windfreak synthesiser. A pickoff using an in-line power splitter is implemented to monitor a portion of the 10 MHz signal for frequency stability measurements.




Figure 6.12: (a) OCXO mounted on an evaluation board with SMA inputs/outputs. The evaluation board was designed as part of another NPL project. (b) Compact frequency synthesiser (Windfreak synth HD) used to generate the 9.192 GHz atomic transition frequency.

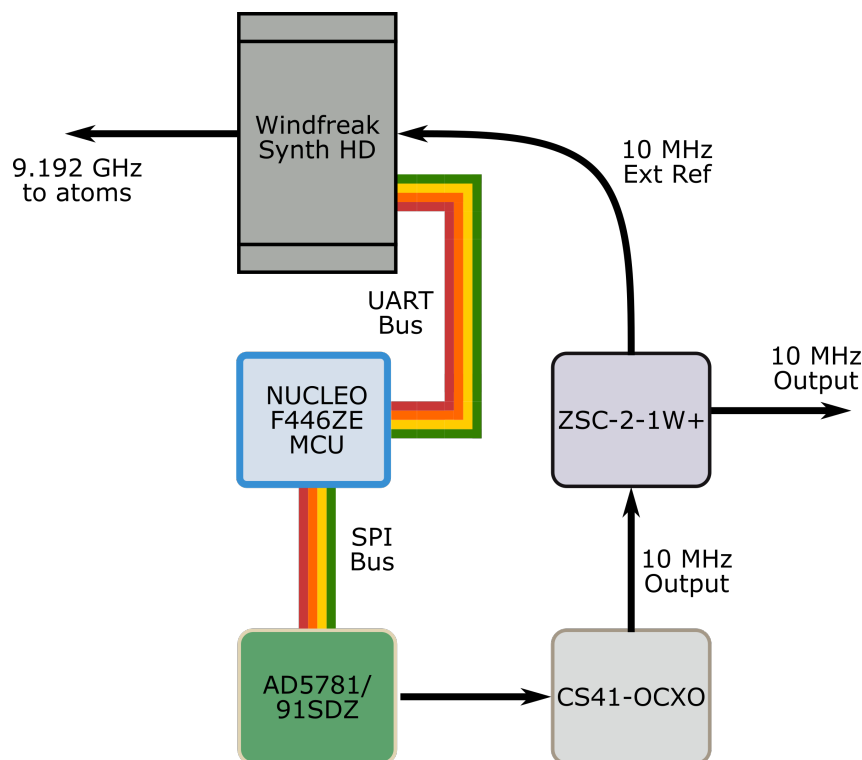


Figure 6.13: Schematic overview of the frequency chain used to generate the 9.192 GHz for interrogating the atoms.

6.2 Software Overview

The experimental sequence is a C++ program⁴¹, which is compiled and then saved to the MCU's onboard flash memory. The code is broken down into many different low-level functions, each containing a few lines of code to execute a specific task. Several functions are then grouped together to perform certain aspects of the experiment, such as: PGC cooling, microwave interrogation, or state detection. The program first initialises the MCU and all the peripheral boards. All outputs are set to their default values, which loads the MOT. Analogue ramps for various voltage sweeps are then precomputed and stored, ready to be called upon. This avoids having to compute values 'on-the-fly', which would otherwise waste valuable computational power. A single header file stores all the variables and constants that are repeatedly used throughout the experiment, which makes it easy to modify things if needed. The MCU then enters a loop where it idles waiting to receive a command to execute. The flow chart in figure 6.14 provides a visual representation of this step-by-step process.

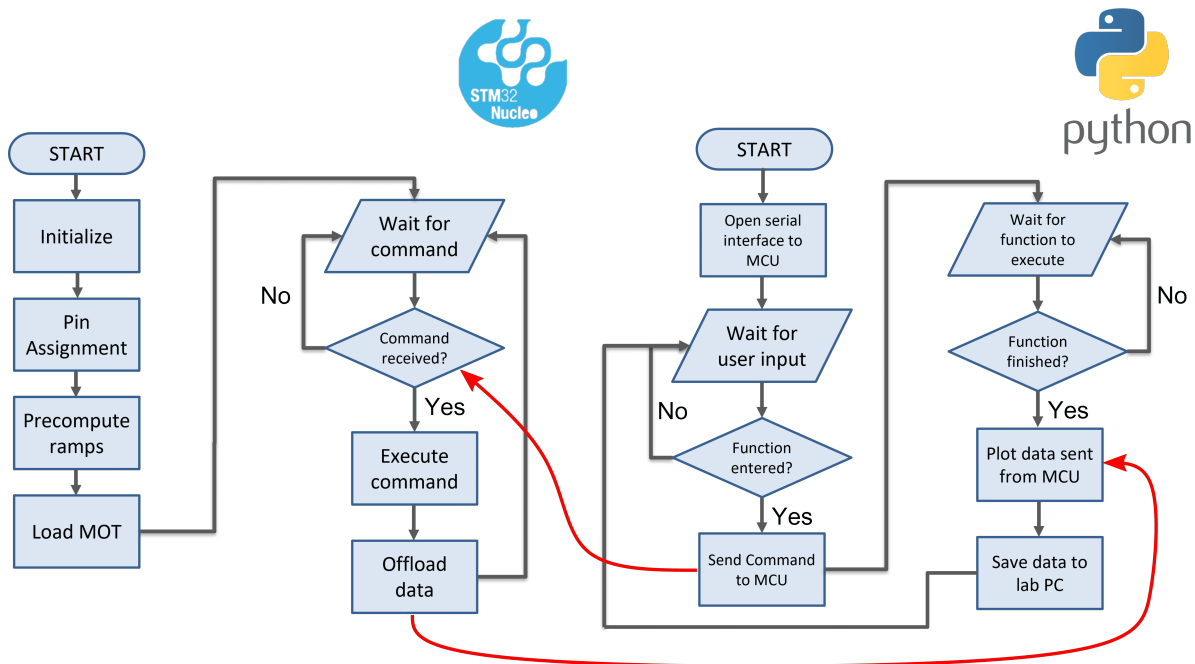


Figure 6.14: Flow charts of the software logic for both the NUCLEO-F446ZE MCU and the Python script running on the lab PC. Both scripts take turns idling whilst the other script is performing a task. Communication between the two is denoted by the red arrows. (Left) Experimental sequence loaded onto the MCU. (Right) Python script on the lab PC.

The MCU is connected to a lab computer via USB to supply 5V as a power source and to load the compiled code onto the board. Although most importantly, the USB connection facilitates a serial interface between the lab computer and the MCU. A small Python script acts

⁴¹The code is freely available to download [129]

as the front end for the experiment, where commands can be sent over the serial interface to the MCU dictating which functions to execute. The communication is two-way, meaning that the MCU can offload data to the lab PC, where it can be either plotted in real time or saved to the lab computer for later analysis. For efficient communication, the code on the MCU and in Python must be synchronised with each other. This is achieved by placing one of the programs in an idle state whilst the other one is active. For example, the Python script will be idle whilst waiting for the MCU to finish executing a function, or vice versa. The flow chart for the Python script is also shown in figure 6.14.

The MCU is programmed using Mbed OS, which has a built-in library of extensive drivers and functions. Furthermore, many developers already have prewritten drivers compatible with Mbed OS, which makes interfacing with peripheral boards quick and easy. However, it was found that some functions related to the SPI communications exhibited overheads in terms of execution time. This was most apparent when trying to change the value of the analogue outputs on the MAX board. This was overcome by writing a few custom functions to handle SPI communication, which increased read/write speeds by a factor of 10.

Given that the experiment is still in its development stage, most functions are used to perform specific measurements or for diagnostic purposes. This usually requires data to be offloaded to the lab PC, where it can be further analysed. However, these capabilities are not something that an end-user of a compact clock would necessarily want or need. The most critical function is the clock sequence, which continuously operates until suspended by user input, with almost no data offload required. It was observed that no aspect of the software or electronics impeded the experiment's ability to operate the clock sequence in a timely manner. The only possible time this could occur is during the calculation of the atomic fraction, which is measured to take 21 μ s. This delay is negligible compared to the few hundred ms required to load the MOT, which is the rate-limiting aspect of the experimental cycle time. Furthermore, the MCU is capable of averaging, integrating, and potentially filtering the ADC data 'on-the-fly'. This completely eliminates the necessity for a lab PC when the system is eventually operated in the field.

6.2.1 Clock Locking

In order to discipline the free-running OCXO to the atomic signal, periodic corrections must be sent to the OCXO's tuning input to keep its output fixed at 10 MHz. The general procedure for locking has already been discussed in section 1.3.2. The microwave frequency is alternated on either side of the atomic feature to obtain two measurements N_{high} and N_{low} . Their difference

$N_{\text{high}} - N_{\text{low}}$ is used to calculate an error signal, which determines the detuning of the average interrogating frequency to the true centre of the resonance. This is shown in figure 6.15, where the atomic feature takes the form of the characteristic Rabi profile (equation 2.29).




Figure 6.15: Example of an atomic feature obtained from a Rabi spectroscopy measurement. The error signal is calculated from $N_{\text{high}} - N_{\text{low}}$.

In the case of Rabi spectroscopy, the optimal lock points are $\pm 0.3805/\tau_{\text{Rabi}}$ (slightly displaced from the actual HWHM points). The error signal is used as an input to a digital PID that is implemented on the microcontroller. The PID system calculates corrections to send to the OCXO tuning input so that the error signal is always driven to zero (i.e. a fixed output of 10 MHz). Once both sides of the Rabi profile have been measured, a value for the error signal is obtained, and then e_k is calculated as the difference between the setpoint (which is 0) and the measured error signal. The PID consists of three components: a proportional term, an integral term, and a derivative term. Each of these is calculated using the following equations [130]:

$$w_k = K_P \cdot e_k, \quad (6.1)$$

$$p_k = (K_I \cdot T \cdot e_k) + p_{k-1}, \quad (6.2)$$

$$q_k = K_D \cdot \frac{1}{T} (e_k - e_{k-1}), \quad (6.3)$$

where the variables $K_{P,I,D}$ represent the gains for each term. The value T is the sampling time, i.e. how often corrections are sent, which is dictated by the experimental cycle time. The subscript k is used to denote different shots in the measurement cycle, and some terms must be stored in

memory to calculate the integrator and derivative terms. The total correction u_k is calculated as the sum of all three terms, which is then applied to the DAC controlling the OCXO tuning input.

$$u_k = w_k + p_k + q_k. \quad (6.4)$$

The gains and the sample time are predefined as constants before the code is compiled. Although, optimisation of the individual gains must be performed to achieve a good quality lock. It is common practice to impose a limit on the output of the PID to ensure that the corrections do not exceed the maximum voltage rating for the OCXO tuning input. The integral term is also limited to avoid integral wind-up (or saturation). This occurs when the integral term is allowed to grow beyond the physical limit of the system's actuator. Although, given the nature of the feedback here, this should not be a significant concern. Throughout the rest of the experimental work, we decided to implement only a PI controller to keep things simple. The derivative term is not always required for locking.

6.2.2 Clock Sequence

Most of the available functions that the MCU can execute are for diagnostic purposes or to perform certain measurements. By far the most important function is the one in which the experiment is operated as a clock. In this section, the three key steps of the clock sequence are outlined: state preparation, microwave interrogation, and state detection. Figure 6.16 shows the timing sequence of the key signals involved in the clock sequence.

The atoms are first laser cooled and trapped in the MOT, which is capable of trapping $\sim 10^8$ atoms in a loading time of 250 ms. Subsequently, both the cooling and repumping light are extinguished, along with switching off the MOT coils. A brief delay of 10 ms allows the residual magnetic fields from the MOT coils to die down. After this 10 ms, the cooling and repumping beams are turned on again to form an optical molasses. The atoms are further cooled in a PGC stage to a few tens of μK . The PGC stage involves rapidly increasing the detuning from 2.5Γ to 12Γ , and reducing the intensity from $4.6I_{\text{sat}}$ to $0.08I_{\text{sat}}$. These ramps are performed simultaneously over ~ 1.5 ms. After the PGC stage, there is another delay of 5 ms, while a small $\sim 5\mu\text{T}$ vertical C-field is turned on to define the quantisation axis and lift the Zeeman degeneracy. The cooling light is turned off, shortly followed by the repumper, which remains on briefly for 150 μs . Leaving the repumper on optically pumps the atoms to the $F4$ hyperfine

level. The C-field is necessary to achieve efficient optical pumping that spreads the atoms evenly across the nine Zeeman states.




Figure 6.16: Logic analyser trace of the signals sent to various peripherals in the electronic setup. 1. MOT loading ~ 250 ms. 2. PGC phase and optical pumping. 3. Microwave interrogation. 4. State detection and data offload. Captured using a Saleae Logic Pro 16 analyser with sampling frequencies: 125 megasamples per second (MS/s) – Digital, 3.125 MS/s – analogue. The logic of the cooling light RF switch is active high, i.e. 1 = off. For clock operation, the ~ 14 ms of data transfer is not necessary (MCU UART TX).

Once the cooling and repumping light are completely extinguished, the cloud begins to free fall. In this work, we only performed short drops of 5 ms so that the atoms remain within the MOT chamber. During this 5 ms of free fall, the microwaves from a horn antenna excite the clock transition. In an improved version of the clock, the total drop distance of 200 mm should be used, which gives a measurement time up to 200 ms and a reduced linewidth of the clock feature. Ideally, the interrogation would also be performed using a microwave cavity.

After interrogation, a sequence of three (plus a background) laser pulses is used to read out the number of atoms that have undergone the clock transition. The first pulse of cooling light induces fluorescence of the $F4$ atoms while also optically pumping them to $F3$. The second pulse of repumping light pumps everything back to $F4$, including those atoms initially in $F3$. A final pulse of cooling light induces fluorescence from $F4$, which now has all the atoms ($F4 + F3$). After the cold-atoms have sufficiently dispersed, another pulse of cooling light measures the background, which contains both scattered light and fluorescence from any thermal atoms. The sampling rate of the ADC and the SPI read speed are fixed, which means that the length of a detection pulse is controlled by specifying the number of samples to capture from the ADC. The time taken to read one sample is measured to be $\sim 3.42 \mu\text{s}$. An initial sample number of 64 was chosen, which is equivalent to a pulse length of $205.2 \mu\text{s}$. In the event that the atoms drop the entire distance, the detection is performed using only the vertical beams of the MOT. For short drop times, all the MOT beams are used.




Figure 6.17: Example of the detection pulses. The red area denotes N_4 . The orange area denotes N_{3+4} . The green area denotes the background pulse N_{BG} . The photodiode falloff after the second pulse is different from the first as the detection light is left on until the cold-atoms have dispersed from the detection area.

The fluorescence signal is recorded on an ADC and transferred to the MCU via an SPI bus. The MCU integrates the data (using a Riemann sum) to calculate the area under each detection pulse. The areas are used to calculate the normalised fraction of atoms that have undergone the clock transition N , given by:

$$N = \frac{N_4 - N_{\text{BG}}}{N_{3+4} - N_{\text{BG}}}. \quad (6.5)$$

As outlined in section 6.2.1, to lock the OCXO to the atomic signal, the microwave interrogation must be modulated either side of the atomic resonance. Therefore, the experimental sequence described above is run twice. One shot operates on the lower side of the resonance, and the other shot above the resonance. Two values of N are obtained, and $N_{\text{high}} - N_{\text{low}}$ forms the input for the digital PID, which is implemented by the MCU. The output of the PID is sent to the DAC controlling the tuning voltage of the 10 MHz OCXO. The 10 MHz OCXO is phase-locked to the 9.192 GHz synthesiser. If the output of the OCXO deviates from 10 MHz, so too does the 9.192 GHz. This is reflected in the fraction N and is ultimately corrected for by the PID feedback. This process is run continuously, and the result is a stable 10 MHz output from the OCXO referenced to the hyperfine ground state transition of caesium.

6.3 Compact Electronics System

The embedded system developed here has proven to be successful in running the experiment and implementing clock operation. The architecture could be applied to any number of quantum experiments or atomic sensors. Functionalities such as: sweeping voltages, toggling magnetic field coils, monitoring detectors, and modulating lasers are all common processes in many types of experiments. However, there are a few steps that could be performed differently to miniaturise further or improve this system.

The route towards miniaturisation starts with replacing benchtop equipment with dedicated evaluation boards. Nonetheless, evaluation boards are designed for universal applicability, which can still lead to large areas of redundant space or connections. The next step would be to populate a single PCB with all the individual ICs from each evaluation board. With careful design and optimised board layout, it would be possible to reduce the overall size of the electronics significantly. In this experiment, the size limiting factors are the laser drivers and power supplies, which are still provided by benchtop equipment. These should be the next area to be replaced with custom PCBs.

The current hardware arrangement is mounted in a partially enclosed box, which is ideal for easy access to probe various signals for debugging purposes. Although, without adequate shielding, the electronics will be more vulnerable to electromagnetic interference (EMI), which could induce unwanted noise. The next iteration should be fully enclosed to prevent this. Furthermore, the use of stripboard for routing power everywhere is acceptable for prototyping designs; however, the long tracks of copper can act as antennas, which will be susceptible to interference and noise. Both capacitive and inductive coupling can occur between neighbouring tracks if they are too close. For all these reasons, a future design should move to a single shielded PCB where these effects can be minimised.

Chapter 7

Experimental Results

This chapter summarises the main results obtained from the experiment and the route towards operating the system as an atomic clock. Most of the results are related to the optimisation of various experimental parameters and the challenges faced when the experiment was first set up. The first section discusses the characterisation of the cold-atom cloud and the measurements used to determine its properties. The key parameters are the total number of atoms trapped in the magneto-optical trap (MOT) and the cold-atom temperature. There is also a discussion of the implemented detection scheme and an analysis of its signal-to-noise ratio (SNR) performance.

When developing an atomic clock, there are generally two major experimental milestones. The first is demonstrating open-loop spectroscopy of the clock signal. The second is operating the clock by closing the servo loop and appropriately disciplining an oven-controlled crystal oscillator (OCXO) to the atomic signal. In this thesis, we were able to perform both, although various areas still need optimisation to improve performance, as discussed in the following sections.

7.1 Cold-Atom Characterisation

After the successful bakeout of the vacuum system and observing cold-atoms in the MOT, the next logical step is to characterise the cold-atom cloud. Measurements were taken to understand some of its basic properties, such as the atom number and temperature.

7.1.1 Atom Number

The number of atoms N in the MOT depends on the balance between the loading and loss rate, which is described by the following rate equation [131]:

$$\frac{dN}{dt} = \alpha P_{\text{Cs}} - (\beta P_{\text{Cs}} + \gamma)N, \quad (7.1)$$

where α is trapping cross-section of the MOT, P_{Cs} is the partial Cs pressure, β is the cross-section for collisions with background Cs atoms, and γ is the rate of loss due to collisions with non-Cs background atoms. The loading rate is characterised by αP_{Cs} and the losses by $(\beta P_{\text{Cs}} + \gamma)N$. The solution to equation 7.1 takes the form:

$$N(t) = N_{\text{eq}} \left(1 - e^{-t/\tau}\right), \quad (7.2)$$

where $N_{\text{eq}} = \frac{\alpha}{\beta} (1 - \gamma\tau)$ is the atom number at equilibrium and τ is the $1/e$ loading time for $t > 0$ (provided the lasers are turned on at $t = 0$). The absolute number of atoms in the MOT is deduced from fluorescence detection. The voltage V at the detector can be used to infer the atom number N by knowing that the voltage is derived from an optical power reaching the detector. The optical power can be estimated by knowing a few characteristics of the atom cloud and the fraction of light collected by the detector. This calculation is encapsulated in the following equation [132]:

$$N = \frac{V}{R_f \alpha} \frac{1}{R_{\text{scatt}}} \frac{1}{\hbar\omega} \frac{1}{(1-f)} \frac{1}{\Omega}. \quad (7.3)$$

The measured current at the photodetector is converted to a voltage V via a transimpedance-amplifier. A feedback resistor R_f across an op-amp is used to set the gain of the photodetector⁴², and α is the spectral responsivity of the detector, usually quoted in units of Amps/Watt. R_{scatt} is the scattering rate per atom, which has been discussed previously in equation 2.32. $\hbar\omega$ is the energy of a fluorescence photon. f is a dimensionless variable used to quantify the fraction of light lost in the collection optics and caesium (Cs) vapour. Ω is the collection efficiency of the detection optics, which can be found in table 4.1.

Both the loading rate and atom number were deduced by monitoring the fluorescence voltage whilst the MOT was loaded. Figure 7.1 shows the fluorescence voltage as a function of the load time. A larger atom number is advantageous for the clock primarily because it helps to reduce noise in the SNR. This shall be discussed in more depth in the later sections, but there will be several noise contributions within the SNR, all of which have different dependencies on the size of the signal (and thus atom number). To suppress these sources of noise, it is beneficial

⁴²For the PDA100A2, $R_f = 2.36 \text{ M}\Omega$ at 70 dB gain, $\alpha = 0.65 \text{ A/W}$.




Figure 7.1: Loading rate of the MOT. The fluorescence voltage is converted to atom number using equation 7.3. The orange line represents a fit to the data using equation 7.2. The uncertainties were determined from fitting only.

to operate with larger atom numbers. In addition, time must also be spent to load a new cloud for each shot of the clock cycle. Thus, a fast loading rate is desirable to keep the experimental cycle time as short as possible. In this experiment, the load time was chosen to be 250 ms, which translates to an atom number of $N = 4.8(0.9) \times 10^8$ (the uncertainty was deduced from general error propagation of equation 7.3 and estimates of the uncertainties in the dependent values. R_{scatt} is likely to contribute the most with an estimated uncertainty of 10 %).

7.1.2 Temperature of the Atom Cloud

In this experimental arrangement, the atoms undergo two stages of laser cooling. The first stage is cooling and trapping in the MOT to temperatures of a few hundred μK , which is comparable to the Doppler cooling limit (equation 2.45). To reduce the temperature further, the MOT coils are turned off and the atoms undergo polarization-gradient cooling (PGC), which can reduce the temperature down to only a few μK . Strictly speaking, the temperature is only well-defined for a system in thermal equilibrium, but it is often used to quantify the average kinetic energy of the atomic ensemble. Thus, the temperature of the atoms is inferred from the velocity distribution of the cloud. A low temperature is beneficial since it ensures the spatial extent of the cloud remains small while the atoms drop by ~ 200 mm. A significantly higher temperature would cause atoms to be lost, resulting in a weaker signal at the detection. The conventional approach to measuring cloud temperature is the time-of-flight method (TOF) [73, 133, 134]. In this method, a probe laser in the form of a sheet of light is placed beneath the cloud. When the cloud is released and falls through the probe, the fluorescence is recorded as a function of time and can be used to

infer the initial temperature. Unfortunately, the current experimental arrangement does not give sufficient optical access to perform TOF measurements. Instead, we use two other measurement processes to estimate the temperature of the cloud, both of which are described in this section.

Release and Recapture (Fluorescence)

The first approach involves release and recapture, in which the fluorescence of the atom cloud is measured both before and after a period of free expansion. The measurement sequence is as follows: the atoms are first cooled and trapped in the MOT, then the laser light and MOT coils are turned off so that the cloud undergoes ballistic expansion for some free evolution time, and finally, the laser light is turned on again to record fluorescence from the atoms. The initial temperature of the cloud can be inferred from the ratio of fluorescence between release and recapture. This is because atoms with a high enough initial velocity on release will leave the capture region during the ballistic expansion phase. Theoretical models describing the cloud expansion exist [135, 136, 137], which give an analytical expression for the expected signal at the detector (i.e. the number of atoms) as a function of time. We adopt a similar approach here but first make some assumptions.




Figure 7.2: Schematic of the release and recapture temperature measurement by monitoring the fluorescence. The initial size of the cloud at t_0 will increase in size after an expansion time $t = t_1 - t_0$. During this time, some of the atoms will leave the fluorescence capture region, which is set by an iris in front of the detector. Note the collection optics are not pictured.

It is assumed that: the cloud consists of non-interacting atoms with a Maxwell-Boltzmann velocity distribution, the initial spatial distribution of the cloud is Gaussian, the expansion is spherically symmetric, and the fluorescence from atoms travelling along the optical axis (z axis) is constant, thus reducing the problem to two dimensions. The probability density is written as [137]:

$$n(\vec{v}, \vec{r}) d^2\vec{v} d^2\vec{r} = \left(\frac{1}{v_0^2 \pi} \right) \exp \left(-\frac{(v_x^2 + v_y^2)}{v_0^2} \right) \left(\frac{1}{r_0^2 \pi} \right) \exp \left(-\frac{(r_x^2 + r_y^2)}{r_0^2} \right) d^2\vec{v} d^2\vec{r}, \quad (7.4)$$

where $v_0 = \sqrt{2k_B T/M}$, is the most probable velocity and r_0 is the most probable distance. Equation 7.4 is the product of the Maxwell-Boltzmann distribution (with v_x, v_y as the speeds along the x and y directions) and the initial ($t = 0$) Gaussian spatial profile of the cloud (r_x, r_y). The origin of the coordinates is chosen to be the initial position of the cloud, i.e. $x = y = 0$ is the centre of the MOT. A second coordinate system is needed to describe the position of the atoms in the cloud after some time t . To translate between these two coordinate systems, the following equations of motion are used:

$$v_x = (x - r_x)/t, \quad (7.5)$$

$$v_y = (y - r_y + \frac{gt^2}{2})/t. \quad (7.6)$$

After substitution, equation 7.4 is integrated over all area d^2r to obtain an expression for the probability density as a function of (x, y, t) . For the next step, it is logical to switch to cylindrical coordinates. Hence:

$$n(\rho, \phi, t) \rho d\rho d\phi = \left(\frac{1}{v_0^2 \pi} \right) \left(\frac{v_0^2}{r_0^2 + t^2 v_0^2} \right) \exp \left[-\frac{g^2 t^4 + 4gt^2 \rho \sin(\phi) + 4\rho^2}{4(r_0^2 + t^2 v_0^2)} \right] \rho d\rho d\phi. \quad (7.7)$$

The signal of interest in the experiment is the ratio of fluorescence between release and recapture, which can also be expressed as the probability density as a function of time only. Therefore, another integration is performed, except this time over the spatial extent of the imaged area by the detector, i.e. $\rho = 0 \rightarrow R$ and $\phi = 0 \rightarrow 2\pi$, where R is the radius of the imaged area set by an iris before the detector ($R = 2$ mm for this work). Equation 7.7 cannot be solved analytically. Instead, a Python script is used to solve the double integration numerically and also fit the result to the data to find the optimised values of T and r_0 .

The measurement procedure outlined at the start of this section was performed, and the results can be seen in figure 7.3. The temperature was measured to be 152(9) μ K, which relates to atoms released directly from the MOT without any further PGC. We attempted to measure

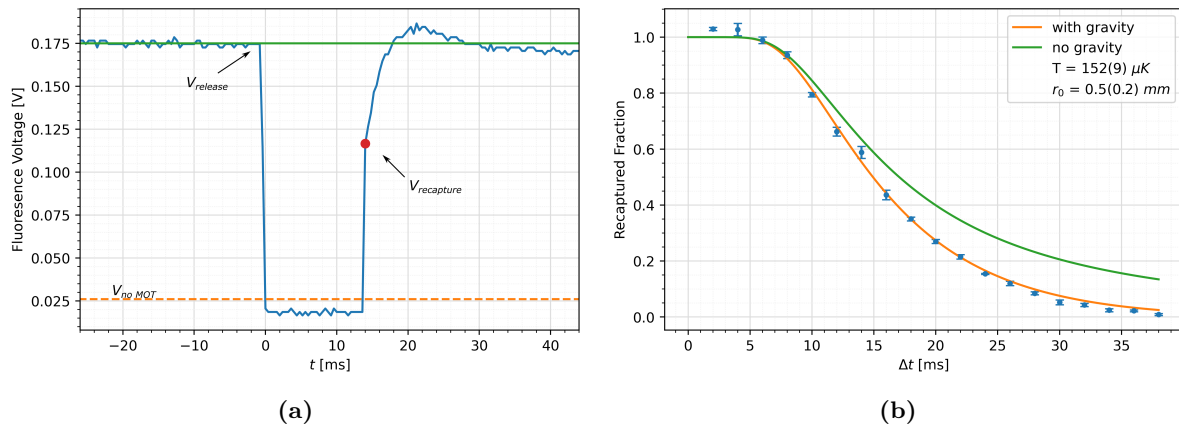


Figure 7.3: (a) Example plot of a release and recapture fluorescence signal ($\Delta t = 14$ ms) recorded at the detector with appropriate points labelled. The orange line indicates the background fluorescence with no MOT and must be subtracted from the release and recapture voltages. (b) The recaptured fraction was calculated for multiple dark times. The orange line represents the fit of the data to the integral of equation 7.7. A simpler model that does not take into account gravity is also shown for comparison. Each point represents five measurements of the fraction, and the error bars are the standard error on the mean [25]. The uncertainties were determined only from the fitting.

the temperature of the molasses after PGC using this method; however, there was an issue due to the ~ 10 ms delay between turning the MOT coils off and starting the PGC phase. This brief delay is necessary to allow the magnetic field sufficient time to decay. Unfortunately, 10 ms is enough time for the cloud to expand and for atoms to escape from the recapture region. This means that atoms have the opportunity to escape the recapture region on two occasions: once during the MOT coil decay and again during the dark period after PGC. This somewhat disrupts the ratio of release to recapture. For this reason, we rely on another method to measure the temperature of the molasses.

Release and Recapture (Spatial)

A similar release and recapture method involves spatially imaging the cloud with a camera. When the cloud is released and allowed to expand for some time Δt , the spatial distribution will evolve (irrespective of its initial shape) into a Gaussian of increased size. By imaging the cloud at different expansion times, the speed of expansion and, ultimately, the thermal velocity $v_{\text{th}} = \sqrt{2k_{\text{B}}T/M}$ can be deduced. The relationship between expansion speed and cloud size will follow:

$$\Delta_f^2 = \left(\frac{k_{\text{B}}T}{M} \Delta t^2 \right) 4 \ln 4 + \Delta_i^2, \quad (7.8)$$

where $\Delta_{i,f}$ represents the initial and final FWHM of the Gaussian profiles⁴³, k_{B} is the Boltzmann

⁴³The relationship between FWHM and variance σ is $\text{FWHM} = 2\sqrt{2 \ln 2} \sigma$.

constant, M is the mass of the atom, and T is the temperature of the cloud. For each expansion time, a Gaussian profile is fit to the image of the cloud and its FWHM Δ_f is determined. The dependence of FWHM vs expansion time is plotted in figure 7.4. The temperature of the cloud is extracted by fitting equation 7.8 to the data. The results indicate a cloud temperature of $11.3(0.3) \mu\text{K}$ in the x-direction and $12.7(0.4) \mu\text{K}$ in the y-direction. This is approximately an order of magnitude lower than the temperature of the MOT.




Figure 7.4: Example of a release and recapture (spatial) temperature measurement. Each point represents five measurements of the cloud's FWHM, and the error bars are the standard error on the mean [25]. The uncertainties quoted were determined by performing a weighted least-squares fit.




Figure 7.5: Example of one of the images taken by the camera, where $\Delta t = 20$ ms. A Gaussian profile is fit to the cloud in both x and y directions.

When fitting a Gaussian profile to the image captured by the camera, the FWHM is returned in units of pixels. Therefore, the imaging system must be calibrated to convert between pixels and mm. This is achieved by imaging a 1951 USAF Resolution Test Target (shown in figure 7.6); these targets are typically used to measure the resolution of an imaging system. They consist of a well-defined and spaced pattern of lines. The relative separation of the patterned lines can be used to determine the relative size of a pixel for any given magnification. A Python script was used to plot the image of the 1951 USAF Resolution Test Target, as well as the pixel values across a slice running through a set of patterned lines. The calibration is calculated using the known line separation and the distance between the troughs in the pixel value plot. In both the horizontal and vertical directions, the spacing was found to be 72 pixels/mm.

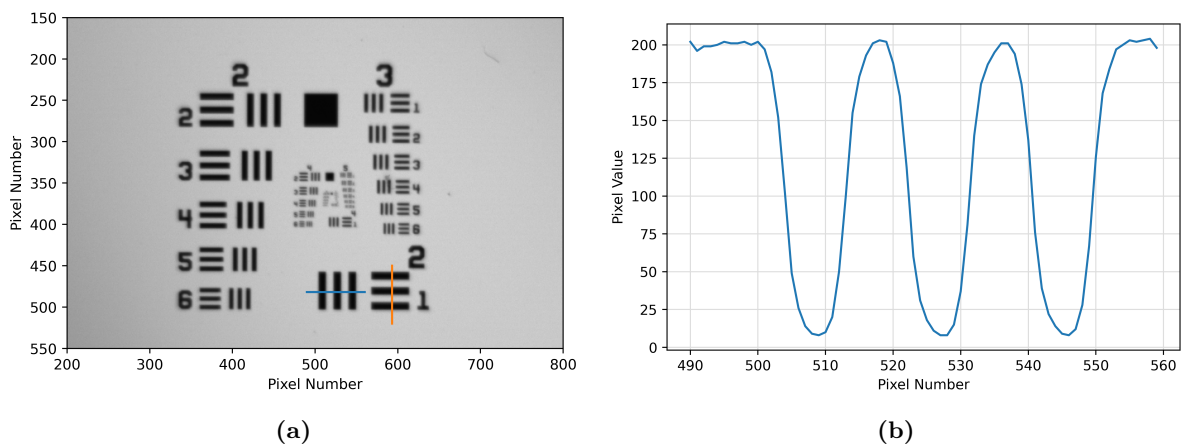


Figure 7.6: (a) Image of the 1951 USAF Resolution Test Target. The blue and orange lines indicate the slice taken through a set of line pairs. (b) Pixel value plotted against pixel number for the blue slice.

7.2 Detection Characterisation

The goal of the detection is to determine the number of atoms that have undergone the clock transition. From this number, it is possible to infer how closely tuned the synthesised microwaves are to the atomic resonance and ultimately send corrections if they start drifting away. Conventional detection sequences, like those seen in an atomic fountain, will detect atoms in both the $F3$ and $F4$ states separately. The setup consists of three spatially separated beams through which the atoms fall. The first beam is tuned to the $F4 \rightarrow F'5$ transition, and the time-integrated fluorescence signal can be used to infer N_4 . This first beam is a travelling wave that also pushes any remaining $F4$ atoms out of the atom cloud. After this, the atoms pass through a standing wave tuned to $F3 \rightarrow F'4$, which pumps the atoms into $F4$. Finally, the cloud is subjected to a third beam similar to the first, except here, the time-integrated signal is used to infer N_3 . In

some arrangements, it is possible to superpose the second and third beams together. The fraction $N = N_4 / (N_4 + N_3)$ calculates the normalised fraction of atoms that have undergone the clock transition and is applied to the servo system to determine the appropriate correction factor. The ratio N is independent of shot-to-shot fluctuations in atom number, which are usually on the order of a few percent.

In this experimental arrangement, we opted for a slightly different approach. The detection beams are instead temporally separated rather than spatially [138], which ultimately minimises the detection optics and maintains compactness. Figure 7.7 shows a schematic representation of the detection sequence.




Figure 7.7: Schematic representation of the compact normalised detection sequence with three laser pulses. 1. $F4$ readout. 2. Repump $F3$ to $F4$. 3. $F4 + F3$ readout. The solid (red) arrows indicate the resonant transition driven by the laser. The fluorescence recorded during the first and third pulses originates from those atoms which have been off-resonantly excited to $F'4$ (dashed grey arrow) and then spontaneously decayed to $F3$. Solid circles indicate the initial state of the atoms at the start of the pulse, whereas dashed circles indicate their final state after optical pumping.

In the operation of the clock, the microwaves are tuned on either side of the atomic feature (where the excitation fraction is 50%). Thus, the atoms leave the interrogation region in a coherent superposition of the two clock states $F3$ and $F4$. The vertical beam used for cooling and trapping in the MOT is repurposed for the detection sequence. A first pulse of cooling light ($F4 \rightarrow F'5$) reduces the vertical motion of the cloud and prevents it from falling further. During this time, the fluorescence is recorded. Again, the time-integrated signal is used to infer N_4 . This first pulse also optically pumps the atoms down to $F3$. Therefore, a quick second pulse of repumper light ($F3 \rightarrow F'4$) brings all the atoms (including those already in $F3$) back up to $F4$. A final pulse of cooling light then induces fluorescence from what is now $N_4 + N_3$. From this three-laser pulse sequence, the normalised fraction of atoms N can be obtained. A sufficient

amount of time after the cold-atoms have left the detection region, another pulse of cooling light is performed to determine the background fluorescence. The fluorescence signals in this detection scheme appear as an exponential decay because the atoms are optically pumped out of the $F4$ state. Figure 7.8 shows an example of the detection pulses.

The behaviour of the detection signals is first characterised without any microwave interaction; upon detection, the normalised fraction N should be equal to unity. However, this was not the case when the detection signals were first observed. The discrepancy in the height of the peaks (the time-integrated signal is the true measure) indicates that atoms are being lost between the two detection pulses. This is unsuitable for clock operation, as it diminishes any valuable information in the fraction N . To understand this, we tried varying a few experimental parameters related to the detection sequence. It was found that there was a dependency on the frequency detuning of the detection beam, as shown in figure 7.8.




Figure 7.8: (a) An example of the two detection peaks which exhibit a height discrepancy when the laser is tuned to resonance, i.e. $\delta = 0$. (b) The average fraction \bar{N} and the SNR (defined as \bar{N}/σ_N) in the detection as a function of detuning. This is for red detuning.

The results suggest that the height discrepancy is largest at smaller detuning. It is unclear what the underlying mechanism is for this. Towards larger detuning, the fraction imbalance is reversed. The ideal operating point is $\sim 0.5\Gamma$ since this is where the SNR is optimised. This is most likely due to the balance between a large fluorescence signal and operating at a point where any laser frequency fluctuations are minimised. This is an area that should be investigated in more detail in the future. The subject of SNR is crucial for the clock and is the primary focus of the next section.

7.3 Detection Signal-to-Noise

One of the key factors in determining the short-term stability of the clock is the signal-to-noise ratio (SNR). Before introducing microwaves to the system, the SNR of the detection system was first characterised. In this clock arrangement, there is no state selection prior to dropping the cloud. After the PGC stage, the repumper light is left on for 100 μs longer than the cooling light. This optically pumps all the atoms into the $F4$ state, where they will be evenly distributed across the nine Zeeman levels. Therefore, realistically, 1/9 of the atoms will be prepared in the $|F = 4, m_F = 0\rangle$ clock state. Furthermore, there is no clear-out of the other states with $m_F \neq 0$. This means that the valuable clock signal sits atop a background of fluorescence from the $m_F \neq 0$ atoms. When the microwave interaction is introduced, the maximum expected change in N will be $\sim 11\%$. However, when the clock is operated, the working fraction will be half of this, given that the locking operates on the side of the resonance signal. Without a more advanced state selection procedure, a low-noise detection system is crucial to resolve the useful clock signal.

Multiple experimental shots were taken to measure the SNR, each measuring the fraction N . The shots are performed without microwave interaction; thus, the average \bar{N} should be unity. The SNR is then defined as $\text{SNR} = \bar{N}/\sigma_N$, where σ_N is the standard deviation of all the measurements. To appropriately identify the dominant source of noise in the detection, the signal size (which is proportional to N_{at}) is varied and the dependence on SNR is monitored using equation 7.9. The total noise in the measured fraction N can be modelled as the quadrature sum of multiple different noise sources [138, 139], where each source has a distinct dependence on N_{at} , shown by the denominator in equation 7.9:

$$\text{SNR} = \frac{N_{\text{at}}}{\sqrt{\sigma_{\text{const}}^2 + \sigma_{\text{root}}^2 N_{\text{at}} + \sigma_{\text{prop}}^2 N_{\text{at}}^2}}. \quad (7.9)$$

Each noise source has a significant physical meaning. For example, σ_{const} encompasses noise in detection electronics without dependence on N_{at} . σ_{root} encompasses noise terms with $\sqrt{N_{\text{at}}}$ dependence, i.e. quantum projection noise (QPN) and photon-shot noise. The QPN will be negligible since there is no microwave excitation in these measurements. There will also be an additional sublevel partition noise arising from the optical pumping before the atoms are dropped [140]. The final term σ_{prop} encompasses all technical noise sources that are directly proportional to N_{at} , i.e. laser frequency or intensity noise.

The first analysis yielded an SNR of ~ 27 , shown in figure 7.9. The fit indicates that the

noise in the detection is dominated by some form of proportional noise. This means that as the cold-atom signal is increased, so too is the noise, which keeps the SNR constant. Usually, this would suggest a noise source that directly affects the cold-atom signal, such as laser frequency noise or laser intensity noise. However, during the experimental shots, it was observed that the analogue-to-digital converter (ADC) was not sampling the same portion of the decaying curve in the detection signals.




Figure 7.9: First SNR analysis with linear dependence, which indicates a σ_{prop} dominant noise term.

Further investigation showed that there was indeed an uncontrollable variable time delay between the trigger and the actual turn-on time for the cooling light. This is demonstrated in figure 7.10. At this stage of the experiment, shuttering capabilities were implemented using a DAC on the MAX board to control the voltage sent to an RF switch/attenuator. The MAX11300 integrated circuit (IC) operates by continuously cycling through the 20 available ports and updating the port value if new data is to be written (or read in the case of an ADC). It does this at a speed of $40 \mu\text{s}$ per port. This means that the time between sending a command to update the DAC value and the DAC value being updated can vary because it depends on where the MAX11300 IC is in the update cycle. The smallest value could be $40 \mu\text{s}$, provided that the port is next in the cycle. On the other hand, the update cycle could be $800 \mu\text{s}$ if the desired port was just updated⁴⁴. In the experimental code, there is a fixed time delay T between triggering the cooling light on and when the ADC starts to read samples for detection. This time delay was adhered to, but the uncontrollable delay t_v meant that the ADC would sample different parts of the detection curve across the two pulses. Sampling different parts of the curve leads to different

⁴⁴Provided all 20 ports are in use.

time-integrated signals, which ultimately affects the measured fraction N . From shot to shot, the delay varied randomly, severely hindering the SNR. This effect was present regardless of signal size, hence the lack of dependence on N_{at} in figure 7.9.




Figure 7.10: Timing issue present in the detection system. A variable time delay t_v was introduced due to the operation of the MAX11300 IC and its update cycle. This delay meant the ADC sampled different parts of the detection curves (shown by the shaded areas). Notice how the first pulse samples more of the rising edge of the detection signal.

To solve this issue, two additional mini-circuits RF switches were implemented in-line immediately after the AOM board and before the RF amplifiers. Two switches are used to achieve maximum extinction. These are placed only in the path of the AOM with responsibility for switching the cooling light on and off. They are controlled by digital logic signals from the MCU. Thus, the delay between sending a switch command and the switch taking place is fixed (~ 20 ns). Cooling light switching is performed instead using these switches rather than the attenuator/switch on the AOM control board. This improved the SNR by an order of magnitude.

Another signal size versus SNR analysis was performed, and the results suggested a linear dependency (blue fit in figure 7.11). This indicated that the SNR was now dominated by a constant noise source, i.e. electronic noise in the detection system. This was easier to identify because there was a noticeable (few %) source of noise on the detection signals, an example of which can be seen in figure 7.12.

The ADC was identified as the origin of this noise source because it persisted even when the experimental sequence was run without cold-atoms and with the laser light blocked. Specifically, an inductor on the MAX board was likely causing electromagnetic interference (EMI) with the




Figure 7.11: SNR analysis. Blue: shows the SNR dependency after fixing the timing issue from the MAX board. Orange: shows the SNR after fixing the electronic noise from the MAX board. Green: SNR dependency after optimising the detuning of the detection beams.

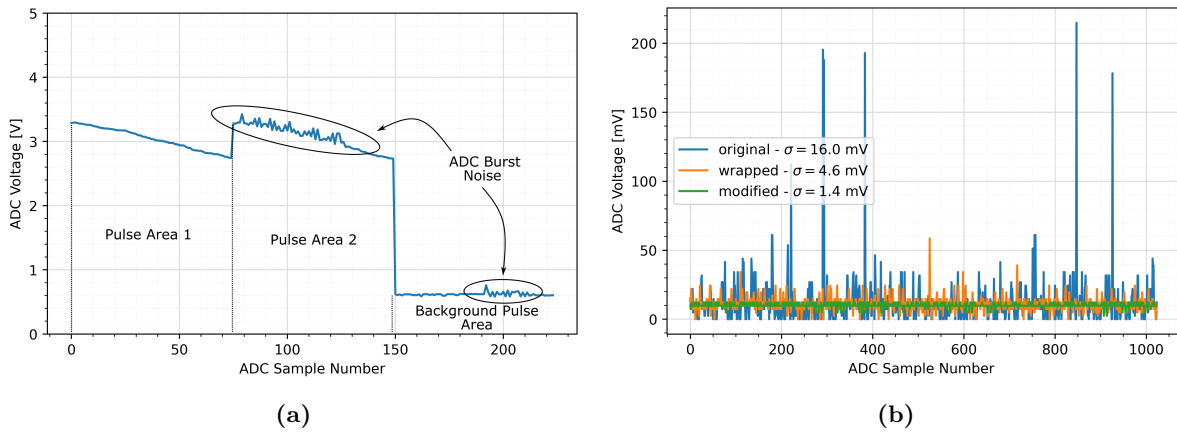


Figure 7.12: (a) Example of an experimental shot experiencing strong burst noise from the MAX board. (b) Examples of the voltage captured without cold-atoms or laser light present. Original modified board (blue), ADC circuitry wrapped with copper tape (orange), permanent modification with inductor removed and benchtop power supply (green). σ represents the standard deviation of all the ADC samples.

ADC due to its proximity. When the IC containing the ADC circuitry was covered (in grounded copper tape), the noise was reduced by a factor of 3.5. A permanent solution was created by removing the inductor and an onboard 13 V regulator, which meant that the board had to be supplied using a benchtop power source. The improvements in electronic noise with each of these modifications can be seen in figure 7.12, and the resulting change was an overall reduction in noise by a factor of 11.

After modifying the MAX board, the SNR analysis was performed again. The new SNR dependency can be seen in figure 7.11 as the orange line. It shows a clear improvement as the constant noise term has been suppressed towards the region of low signal size. An additional optimisation of the detection beam detuning (shown in figure 7.8) was able to push the SNR

to ~ 500 , which was eventually limited by a new source of proportional noise σ_{prop} . Further investigation is needed to identify the origin of this noise. Potential sources could be laser frequency or laser intensity noise. It was observed that a clear modulation could be superimposed on the detection signals if the modulation amplitude for frequency modulation (FM) locking of the cooling laser was increased. The modulation frequency is 20 kHz, which is comparable with the detection pulse lengths of a few hundred μs . To avoid this, it would be better to modulate the cooling laser at higher frequencies on the order of MHz, where the fast modulation would be averaged out and less noticeable. This step would have required the locking electronics to be switched out, which was decided against. Intensity fluctuations may arise from both the power and pointing instability of the detection laser. To mitigate this, the detection laser is operated with a high intensity (the same used for cooling and trapping), which has two effects: it reduces the impact of fractional fluctuations and saturates the transition to reduce the sensitivity to any fluctuations.

Rather than continuing to optimise the SNR, a decision was made to instead continue with the microwave interrogation as the SNR was high enough to resolve small fractional changes in N . It is worth noting that the SNR discussed above is related to the detection system. It is not the same SNR as defined in equation 1.2 that estimates the short-term stability of the clock. In this experiment, where there is no state selection, the average excited fraction of the clock signal will be ~ 0.11 . Therefore, the SNR of the clock signal will be $0.11 \times 500 = 55$. This emphasises the importance of either fully optimising the SNR until QPN-limited stability can be reached or implementing an advanced state selection to boost the fraction of excited atoms.

7.4 Microwave Spectroscopy

Ideally, the spectroscopy would be performed using a resonant microwave cavity. However, for reasons already discussed in chapter 5, a microwave cavity was not implemented into this system. Instead, a horn antenna was used, which is still suitable for performing basic microwave spectroscopy. Initially, the microwave horn was directed into one of the spare viewports on the MOT chamber. This allowed for spectroscopy with pulse times up to ~ 25 ms, after which the atoms started to fall out of the MOT chamber and the recapture region due to gravity.

The clock transition corresponds to the $|F = 4, m_F = 0\rangle$ to $|F = 3, m_F = 0\rangle$ transition. To appropriately resolve this transition, a magnetic field must be applied to the atomic ensemble, which lifts the Zeeman degeneracy. This gives rise to a whole host of other available transitions,

as shown in figure 7.15. The splitting of the levels, and thus the separation of these transitions is dependent on the magnetic field. Fortunately, the clock transition is insensitive to magnetic fields to first order. However, observation of the other transitions during spectroscopy can be used as a measure of any residual magnetic field. As described previously, the nulling fields were adjusted relative to a fluxgate magnetometer before the chamber was introduced into the coil system. If the field nulling were perfect, then ideally, all the Zeeman states would be collapsed, and a spectroscopy scan should show a single feature. This was not the case in our experimental setup; therefore, each axis of the nulling coils was optimised until all the peaks were combined as much as possible.

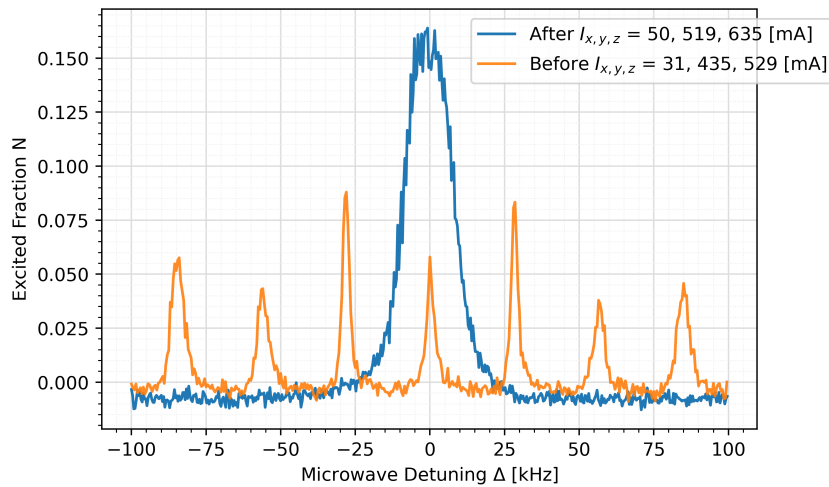


Figure 7.13: Microwave spectroscopy scan after optimising each axis of the magnetic field nulling coils. Also shown are the currents in each axis $I_{x,y,z}$ both before and after nulling.

Figure 7.13 shows the spectroscopy traces both before and after optimising the compensation coil currents. Nulling the field relative to the fluxgate magnetometer still shows some Zeeman structure with a splitting of $\sim 8.1 \mu\text{T}$ (determined from fitting). After optimisation, the peaks have collapsed into a single feature with a residual linewidth of $\sim 16.1 \text{ kHz}$ (determined from fitting). This might lead one to conclude that the residual field is inhomogeneous across the centre of the trap, i.e. the atoms experience different Zeeman splittings, and the net effect averages out across all the atoms in the trap, leading to a single broadened feature. However, if this were the case, the inhomogeneous broadening should also appear in the individual separated features. It is more likely that the field is still not fully minimised. The source of such residual fields could arise from the presence of magnetic material near the physics package. Upon carrying out these measurements, it was discovered that some of the materials used for the construction of the vacuum chamber exhibited magnetic properties. Ideally, magnetic materials should be avoided in all areas of the physics package for this very reason. Some of the bolts and other vacuum

components are made of stainless steel 304, which exhibits weak magnetic properties. This was an experimental oversight and was not realised until late in the remaining time available to take measurements. Addressing this would have required dismantling the vacuum system, which would have taken considerable time, but this may be carried out in the future. For cold-atom applications, where magnetic properties are usually a concern, stainless steel 316 or any of its variants⁴⁵ should be used, which has the lowest magnetic permeability of all the stainless-steel alloys. Note, however, that when 316 stainless steel is welded or stressed in other ways by machining, then its non-magnetic properties can be destroyed (even for 316 LN). If not stainless steel, then titanium or aluminium are other good options.




Figure 7.14: Transitions between Zeeman levels in the ground state of Cs. Blue arrows indicate π transitions. Green arrows indicate σ transitions. Each of the different transitions are labelled. Note that there are a total of 21 allowed transitions, although 6 of them occur at the same energy. The clock transition is labelled as 0.




Figure 7.15: Wide scan spectroscopy signal to show all the transitions. Each of the 15 transitions are labelled in correspondence with figure 7.14. The C-field current was set at 200 mA, which is relatively high equating to a splitting of 56 kHz (determined from fitting).

⁴⁵316L – ‘L’ for low carbon variant, 316LN – ‘N’ for nitrogen-enhanced variant.

When optimising the magnetic field, the microwave power was set to 15 dBm (the maximum power output from the Windfreak frequency synthesiser). This was arbitrarily chosen to observe the effects of the microwaves. However, for the operation of the clock (as was discussed in section 2.3.1), it is more efficient to drive a π -pulse, which excites a full population transfer. This is satisfied when $b\tau = \pi$, i.e. for a given pulse length τ , there will be a corresponding Rabi frequency b (and thus microwave power) to satisfy the π -pulse condition. Usually, the pulse length is set by the clock arrangement, which in this experiment would be dictated by the duration that the atoms are in free fall. To identify the correct microwave power, Rabi oscillations were performed [141, 142, 143]. The amplitude of the oscillations very quickly tends towards what should be a 50% excitation fraction (in our experiment this would be $1/9 \times 1/2 = 1/18$ given that we measure fluorescence from all nine Zeeman states). This indicates a strong dephasing mechanism, which could arise from inhomogeneous microwave power across the atoms or an alignment mismatch between the microwaves and the quantisation field.




Figure 7.16: Rabi oscillations with varying microwave power but fixed pulse length of $\tau_{\text{Rabi}} = 5$ ms.

The π -pulse condition is found by setting the microwave power to that which matches the first maximum of the oscillation profile. Once set correctly, the microwave scan is reduced to focus solely on the clock transition. The result can be seen in figure 7.17, along with the fitting of equation 2.29 to the data.

There are two measurements shown, one with an SNR of 250 and another with 500. The 250 SNR measurement was performed using the free-running OCXO as the reference for the frequency synthesiser (Windfreak). To set the microwave frequency, a command is written to the frequency synthesiser specifying the exact frequency to a resolution of 0.1 Hz. However,




Figure 7.17: Rabi signal for a pulse time of $\tau_{\text{Rabi}} = 5$ ms. There are two measurements shown, one with 250 SNR (greyed) and another with 500 SNR (solid colour). The 250 SNR profile is referenced using the OCXO, and the 500 SNR profile is referenced to a UTC(NPL) hydrogen maser signal. This explains the discrepancy between the frequency shift of the profiles. The fit to the data indicates a maximum excitation of 0.1.

the accuracy of the microwave output is determined by the accuracy of the 10 MHz reference. Measurements were performed (figure 6.11) to identify the exact tuning voltage corresponding to a 10 MHz output. However, the OCXO will have drifted since then, shifting the ideal tuning voltage away from the current setpoint. Therefore, the profile exhibits a $-18.4(0.5)$ Hz shift (determined from fitting). To validate this, the same measurement was performed again (now with 500 SNR) except using one of the 10 MHz UTC (NPL) hydrogen maser signals as a more accurate reference. The resulting shift (determined from fitting) is now $3.1(0.3)$ Hz, which is more aligned with what one might expect given that most of the frequency-shifting phenomena in Cs have a positive sign⁴⁶.

Another significant point to note about this result is that the fitted FWHM matches the Fourier limited width for a pulse length of $\tau_{\text{Rabi}} = 5$ ms. This indicates that there are no significant broadening effects that could degrade the potential short-term stability. This is to be expected given that the atoms are interrogated during free fall, during a period when they are isolated from any external perturbations. Ideally, by dropping the atoms down the entire flight tube, a longer pulse time (up to 200 ms) could be achieved, thus giving an even narrower feature. There is no reason to believe that there are narrower sources of broadening below 159 Hz that might limit a clock with a longer drop time.

The next step of the experiment is to perform closed-loop clock operation, whereby the above

⁴⁶The collisional shift, black body radiation shift, and second-order Doppler shift are the only systematics with a negative sign.

atomic feature is used to discipline the free-running OCXO to the atoms. Before doing this, the expected short-term stability (equation 1.2) can be estimated using some of the experimental parameters and characteristics of the atomic feature above.

K	1.516	$\pi/2$
τ (or T for Ramsey)	5 ms	150 ms
$\Delta\nu$	159 Hz	5 Hz
Q	5.75×10^7	2.8×10^9
SNR	$0.1 \times 500 = 50$	50
T_c	285 ms	430 ms
σ_y	$1.2 \times 10^{-10} \tau^{-1/2}$	$3 \times 10^{-12} \tau^{-1/2}$

Table 7.1: The expected short-term stability (equation 1.2) calculated using the current experimental parameters. Also shown is a stability estimate if the atoms were dropped the entire distance down the chamber using a 150 ms Ramsey interrogation.

7.5 Frequency Stability Analysis

The concept of frequency stability and Allan deviation (ADEV) analysis was discussed in section 1.3. In this section, those concepts are applied here to characterise the frequency stability of the experiment when operated as a clock. To characterise the stability of a clock, it must be compared against a second more accurate clock as the reference. Fortunately, at NPL, we have access to the hydrogen maser signals used to derive UTC(NPL). We use the hydrogen maser signal as a reference clock because it is likely that the stability of the experimental clock will be worse relative to the maser. To measure the frequency stability, the nominal 10 MHz output of the clock is compared to the maser's 10 MHz signal on a phase noise analyser (Microsemi 3120A). The 3120A measures the phase difference between the reference and device-under-test (DUT) at regular intervals, then uses software to calculate the ADEV. Section 1.3.1 showed how to calculate ADEV from samples of frequency data, but it can also be calculated using phase data.

Before measuring the frequency stability of the compact cold-atom clock, the noise floor of the 3120A was measured. This was done by using a power splitter to split the hydrogen maser signal in two and then using the maser signal for the reference as well as the DUT input. The result is shown in figure 7.19 and indicates a noise floor of $2.5 \times 10^{-14} \tau^{-1/2}$, which represents the limit of stabilities that can be measured with this setup. Also shown in figure 7.19 is the stability of the OCXO when integrated into the electronics but not locked to the atomic signal.

The stability of the OCXO has been severely degraded. The stability is now nearly three orders of magnitude worse than the free-running OCXO (before integrating it into the electronics). This indicates the current electronic setup is degrading the OCXO's performance. One

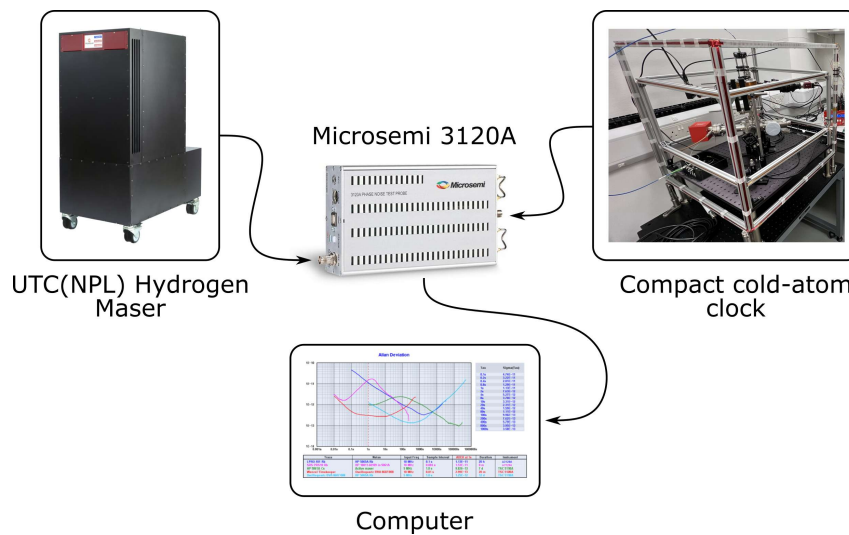


Figure 7.18: Schematic of the stability measurement. The 10 MHz from the compact cold-atom clock is compared against a hydrogen maser using a 3120A phase noise analyser.

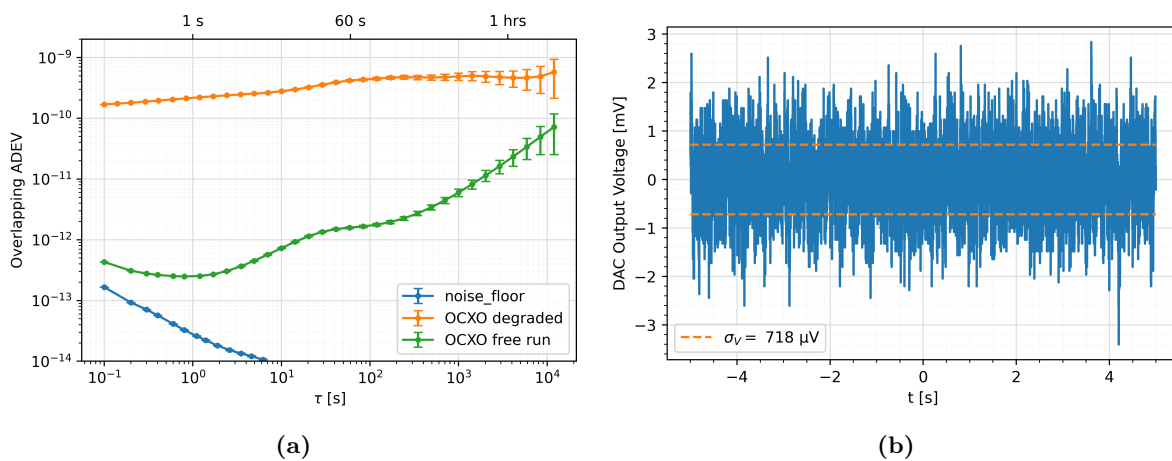


Figure 7.19: (a) ADEV measurements of the OCXO. Blue: maser vs. maser. Green: OCXO vs. maser. Orange: OCXO in electronics setup vs. maser. (b) Time domain trace of the output voltage from the DAC sent to the tuning port of the OCXO. The signal was measured on an oscilloscope when AC coupled, and the noise was estimated to be the one-sigma deviation.

possible source of this problem would be the tuning input to the OCXO. Tuning of the OCXO is performed using a DAC, which is necessary to provide corrections when locking to the atomic signal. However, any noise present on the tuning voltage will translate into noise on the OCXO's output. A crude measurement of the DAC's output noise was performed by monitoring the signal on an oscilloscope, as seen in figure 7.19. If we take the standard deviation as an estimate of the noise, this can be translated into an instability on the OCXO's output frequency via the tunability $-0.311(0.002)$ ppm/V from figure 6.11. This conversion translates to an instability of $2.2 \times 10^{-10} \tau^{-1/2}$, which closely matches the instability of the degraded OCXO measured at $\tau = 1$ s. This is only an estimate, as there will most likely be a whole spectrum of different noise components, each contributing to the stability of the OCXO on various timescales.

Despite this noise issue, we attempted to lock the OCXO to the cold-atom signal to see if the instability of the OCXO could be effectively reduced in the medium and long term. The clock's experimental sequence has already been described in section 6.2.2. For these measurements, Rabi spectroscopy with a pulse time of $\tau_{\text{Rabi}} = 5$ ms was used. The cycle time of the clock is $T_c = 285$ ms. This includes the time taken to laser-cool the atoms, trap them, release them from the trap, and perform the detection before restarting the cycle. After two cycles measuring either side of the atomic signal, a correction is ready to be sent to the OCXO. The time between corrections is often called the 'loop' time or 'attack' time. Typically, the ADEV performance on timescales smaller than the 'attack' time is determined by the free-running stability of the OCXO. This is because the OCXO acts as a flywheel oscillator between the atomic measurements. On timescales longer than the 'attack' time, the ADEV performance is dominated by the stability of the cold-atoms. At averaging times similar to the 'attack' time, there is a characteristic servo 'hump' in the ADEV as corrections are sent to the OCXO.




Figure 7.20: ADEV performance of the experimental setup when operated as a clock. Green: free-running OCXO degraded by electronics. Orange: first stability test ($K_{P,I} = 2$). Blue: second stability test with PI gains ($K_{P,I} = 0.5$). Dashed lines: fitting of power laws. Servo 'hump' occurs at $\tau = 0.570$ s.

In figure 7.20, two measurements of the clock's stability are shown in orange and blue. The orange line indicates the first measurement run without optimisation, and the results indicate that after averaging for ~ 700 s, the frequency instability of the OCXO is successfully reduced by the steering of the more stable atomic signal. Another measurement run was performed with reduced PI controller gains; the result is shown by the blue curve where the crossover point has shifted to ~ 30 s. After this, the ADEV then continues to improve (i.e. average down) with the characteristic $\tau^{-1/2}$ trajectory of $1.9 \times 10^{-9} \tau^{-1/2}$ (determined by fitting to the data). It then

reaches a floor of $\sim 6 \times 10^{-10}$ after a few 1000 s. On timescales longer than this, the ADEV begins to exhibit a random walk of $6.8 \times 10^{-13} \tau^{-1/2}$, which is most likely due to instabilities in the environment. However, further measurements would be necessary to confirm this.

A few conclusions can be drawn from these results. First, one might question why the $\tau^{-1/2}$ performance does not follow the $1.2 \times 10^{-10} \tau^{-1/2}$ behaviour that was estimated at the end of the previous section (see table 7.1). There are possibly two effects at play here. First, it seems that the act of applying corrections to the OCXO is degrading its short-term stability. If the corrections are not instantaneous but rather slowly varying over time, this will disturb the OCXO's short-term stability. The slew rate of the DAC used to tune the OCXO is specified to be on the order of a few μs , which should not be the issue. It is more likely that the tuning voltage is being low-pass filtered by the presence of capacitance between the DAC and the OCXO tuning port. This also potentially explains the sensitivity to the gains of the PI feedback, i.e., reduced gains result in smaller corrections, which will have less effect on short-term stability. In theory, even if the gains are not optimised and drive oscillations about the setpoint, these oscillations should eventually average out on longer timescales, and the $\tau^{-1/2}$ dependence should take over. Second, if the short-term stability of the OCXO has been degraded by the corrections from the feedback, this will impact the ability to operate at the lock points effectively. In fact, we can imagine the data points in figure 7.17 to have relatively large error bars in the x axis as well as the y axis (when referenced to the OCXO). This will significantly affect the SNR used to determine the estimated stability. Both of these effects are likely to be responsible for the degraded stability seen in figure 7.20.

Direct voltage tuning of the OCXO is a common feedback method, and its main drawback is the stringent requirement for a low-noise voltage source or an OCXO with a low tuning coefficient. Other feedback schemes allow the OCXO to remain free-running, thus eliminating the feedback noise issue [144]. This typically involves mixing a larger fixed microwave frequency with a smaller (few hundred MHz) frequency component; both are referenced to the OCXO. The instability in the OCXO is corrected by feeding back to the small frequency component to stay resonant with the atoms. A pick-off of the small frequency signal is then downconverted to generate the appropriate 10 MHz output signal. These schemes are slightly more complicated but can mitigate the need for an OCXO with a low tuning coefficient (which are generally more expensive).

If the issues with the stability of OCXO and feedback implementation can be resolved, the stability should follow the expected behaviour from table 7.1. It would then be beneficial to

either increase the probe time (which has been a primary aim of this project by dropping the atoms for a larger distance) or return to the issue of improving the SNR.

7.6 Experimental Results Summary

This chapter has presented the main results obtained from the experimental setup. Release and recapture measurements indicate an atom cloud temperature of $T = 11.3(0.3) \mu\text{K}$. The laser cooling is most likely limited by the presence of uncontrollable residual fields from magnetic material used in the construction of the vacuum chamber. The atom number estimated from the measured fluorescence for a load time of 250 ms is $N = 4.8(0.9) \times 10^8$.

We also characterised the SNR performance of a detection scheme in which the state of the atoms is measured using a three-laser pulse sequence. The shape, contrast, and SNR of the clock signal are critical aspects to estimate the potential stability performance of the clock. The SNR of the detection system was measured to be 500, which is limited by some form of proportional noise. More investigation is needed to identify the origin of this noise source, but it is suspected that either laser intensity or frequency noise is responsible.

We have performed open-loop spectroscopy of the clock transition using a short drop time of 5 ms. The characteristic Rabi line shape has been observed, and from this, the estimated stability was calculated as $1.2 \times 10^{-10} \tau^{-1/2}$. A short Rabi pulse time was initially chosen for diagnostic purposes, and this could be extended up to ~ 25 ms with the horn directed into the MOT chamber. For interrogation times longer than this, the atoms can be dropped down through the entire flight tube (~ 200 mm). For this, Ramsey interrogation would need to be implemented using two horn antennas to apply the microwaves once in the MOT chamber and again at the DN16 cross-section of the vacuum apparatus, which acts as an interrogation region. This would, theoretically, lead to two orders of magnitude increase in frequency stability, as evidenced by table 7.1.

Furthermore, we have demonstrated clock operation and measured the short-term stability to be $1.9 \times 10^{-9} \tau^{-1/2}$. This stability is limited by the current implementation of the corrections to the OCXO and its degraded short-term stability rather than the cold-atoms. Despite this, we have successfully disciplined an OCXO to the atomic reference and demonstrated that its stability can be improved after averaging for several tens of seconds. These results represent a first step towards a compact cold-atom microwave clock system and provide a good basis for further optimisation in future work, as described in the next chapter.

Chapter 8

Future Work and Conclusions

8.1 Future Work

The experimental results that were obtained within the time available for this project can be optimised further. In this section, I would like to discuss potential areas of improvement and also look ahead to what could be implemented in a potential second iteration of this system.

Optimisation of the noise issue with the tuning voltage to the oven-controlled crystal oscillator (OCXO) would be the first aspect to be revisited and improved. Currently, the electronics setup consists of multiple evaluation boards, which are all situated inside an open enclosure; this leads to susceptibilities with electromagnetic interference (EMI) and crosstalk between neighbouring boards. Noise improvements would include the use of a shielded enclosure (or several), appropriate filtering, and the design of specialised low-noise circuitry. Some of these issues were also raised at the end of section 6.3. This should bring the short-term stability closer to the behaviour expected from equation 1.2 and that predicted in table 7.1. Furthermore, in this experiment, a simple trial-and-error approach was used to set the gains of the PID feedback. Instead, it would have been more efficient to implement one of the more popular methods, such as Ziegler-Nichols tuning [145].

Following this, the next focus should be to optimise the signal-to-noise ratio (SNR) more completely. The ideal scenario would be to demonstrate quantum projection noise (QPN) limited short-term stability, whereby the dominant noise source arises from statistical fluctuations during the detection process. However, in practice, this can only be achieved after meticulous optimisation and stabilisation of various experimental parameters, which could include additional hardware. If this increases the overall footprint of the experiment, then this will not be ideal for the final compact clock. Nevertheless, a good understanding of the potential limit that

can be reached would be very insightful. If this limit can be demonstrated at shorter diagnostic times, such as $\tau_{\text{Rabi}} = 5 \text{ ms}$, the next step would be to increase the interrogation time. This would involve dropping the atoms down the entire flight tube ($\sim 200 \text{ mm}$).

These improvements should provide short-term stability at the level of low $10^{-12} \tau^{-1/2}$, as evidenced by the estimates in table 7.1, which assumes the current limited SNR of 500 and a longer interrogation.

One of the most significant limitations of this type of architecture is the delay between releasing the cloud from the magneto-optical trap (MOT) to the cloud entering the interrogation region. This delay wastes valuable interrogation time, especially since the atoms are moving slowest immediately after their release. An approach (alluded to at the end of section 5.6) to mitigate this would be to integrate a microwave cavity with a compact cold atom source such as a pyramid- or grating-MOT. However, the presence of the cold-atom source will most likely disrupt the field distribution inside the cavity. More consideration should be given to this idea, which could be aided by COMSOL simulations. An alternative approach would be to briefly stop the atoms in place when they enter the interrogation region by momentarily turning the trapping lasers back on to form a one-dimensional molasses. Consider the example of dropping the atoms halfway down the flight tube, freezing them in place, and then letting them continue to drop. The ratio between the time the atoms drop the first 10 cm and the time the atoms drop the last 10 cm is ~ 2.4 . This represents a significant increase in the potential interrogation time. However, stopping the atoms' vertical motion induces transverse heating; thus, there will be a trade-off between increasing the cloud size and losing atoms versus a longer interrogation time. The transverse root-mean-square (rms) velocity will increase by [146]:

$$v_{\text{rms}}^{\perp} = \sqrt{\frac{\alpha N}{3}} v_{\text{rec}}, \quad (8.1)$$

where $\alpha = 3/10$ to account for the dipole radiation pattern from spontaneous emission. N is the number of photons scattered to bring the cloud to a stop. $v_{\text{rec}} = \hbar k_L / M$ is the recoil velocity from each spontaneously emitted photon. According to this model, a caesium (Cs) atom falling for 10 cm will need to scatter ~ 400 photons before coming to a standstill. In doing so, the transverse rms velocity will increase by $\sim 6.3 v_{\text{rec}}$. If we assume an initial distribution of $10 v_{\text{rec}}$ (corresponding to $\sim 20 \mu\text{K}$), the final convolved transverse velocity is $\sim 12 v_{\text{rec}}$. Therefore, the displacement after falling the remaining 10 cm will be $12 v_{\text{rec}} (0.142 \text{ ms}) = 6 \text{ mm}$. This is significant but relatively small compared to the size of the entrance and exit apertures to the

interrogation region (~ 16 mm)⁴⁷. Thus, most of the atoms would be expected to reach the detection region. This estimate will be affected by any intensity imbalance in the beams or incorrectly nulled magnetic fields, which would increase the heating rate and final size of the cloud. This would be an interesting experiment to perform and could be tested in the existing experimental arrangement.

In this experiment, we did not implement any advanced state selection procedure to maintain experimental simplicity. Instead, all the atoms are pumped into the $F4$ state where they will be evenly distributed across the nine Zeeman levels. This means, at best, we can only ever hope for $1/9$ of the atoms to contribute to the useful clock signal. Two modifications could be implemented to improve the SNR. The first would be to perform a clear out of all other non-clock state atoms. Instead of preparing into $F4$, the atoms are pumped to $F3$ after cooling, and a π -pulse of microwaves is applied to transfer clock state atoms to the upper clock state in $F4$. A resonant laser pulse is then applied to clear out the remaining $F3$ atoms, effectively removing them from the cloud. After this, the clock sequence continues as normal. This would increase the maximum clock signal (i.e. fraction of excited atoms) to unity, resulting in an order of magnitude improvement for the SNR. However, the strength of the fluorescence signal also decreases by $8/9$ due to the lost atoms⁴⁸. The second approach would be to implement an all-optical state selection process whereby all the atoms are instead prepared in the clock state rather than cleared out [147]. For either of these schemes, one might consider transitioning to rubidium (Rb-87) instead of Cs, where the total electronic angular momentum numbers F are smaller. Thus resulting in a higher probability of atoms being prepared in the clock state. If the system developed in this work employed Rb rather than Cs, one would expect $1/5$ of the atoms to be prepared in the clock state. Figure 8.1 shows two examples of optical pumping schemes that prepare the clock state.

To implement either of these schemes and still maintain compactness, one can consider how this could be performed in the MOT chamber using the same light used for cooling and trapping. Perhaps locking to a different peak in the saturated-absorption spectroscopy (SAS) signal could allow for detuning to reach both cooling and optical pumping transitions. Another approach would be to implement a routine to unlock the laser, jump the lock point to another transition, and then relock the laser. Furthermore, polarization switching capability would have

⁴⁷16 mm is currently set by the size of the DN16 cross used in the interrogation region, although this might be made smaller in a second iteration.

⁴⁸The clock signal N will increase, but the fluorescence signals recorded at the photodetector will be weaker, which could run the risk of being comparable to the noise of the detection electronics.




Figure 8.1: Optical pumping schemes that could be used for implementing state selection in Rb-87. (a) Preparation of $|F = 2, m_F = 0\rangle$ clock state. Two transitions must be driven using π polarized light. Electric dipole selection rules dictate that the $|F = 2, m_F = 0\rangle \rightarrow |F' = 2, m_F = 0\rangle$ transition is forbidden. A similar scheme can also be implemented in Cs, the only difference being the larger F numbers. (b) Preparation of $|F = 1, m_F = 0\rangle$ clock state. Only a single transition is driven using σ^+ and σ^- polarized light, which is already fulfilled for the MOT arrangement in the experiment. The excited state $F'0$ has only one decay channel. Note this scheme is only present in Rb due to the presence of the $F'0$ state.

to be employed to toggle between circular light for cooling and linear light for optical pumping. Electronically controlled variable waveplates exist⁴⁹, but it would be necessary to characterise these optics and their suitability for use in the clock.

Finally, this research has focused only on improving short-term stability, i.e., reducing the statistical uncertainty in the measurement of the clock frequency. There will be systematic uncertainties that shift the output frequency of the clock leading to inaccuracies. A stable but inaccurate clock will exhibit a constant drift over time, which can often be calibrated out when measured against a more accurate clock. To evaluate the accuracy of this system, each individual shift must be measured along with its corresponding uncertainty. The combination of all shifts forms an error budget, and the resulting uncertainty determines the final accuracy of the system. To understand the prospective accuracy for this experiment, we can look at the published error budgets from atomic fountain clocks [105, 106] and see which systematics would be most relevant here. In fountain clocks, most of the systematics can be measured with a residual fractional inaccuracy of a few parts in 10^{-16} . The system developed here is much simpler, and a similar accuracy would not be expected. The closest competitor would likely be a commercial beam clock (e.g. Microchip 5071A) whose fractional inaccuracy is 5×10^{-13} (for the highest specification model). Accuracy limitations in commercial beam clocks are usually due to their inability to compensate for the second-order Doppler and second-order Zeeman shift.

⁴⁹Liquid crystal variable retarders.

Even though the atoms are dropped in this system, the velocities reached are still a few orders of magnitude slower than in a beam clock. When dropped the total distance, the final velocity reached by the cloud translates to a fractional second-order Doppler shift of $\sim 10^{-17}$, which will most likely be negligible against other shifts.

Currently, the experiment is highly sensitive to magnetic field instabilities in this unshielded arrangement. To address this, the operating C-field is set relatively high ($\sim 5 \mu\text{K}$). Thus, the fractional second-order Zeeman shift is 10^{-10} , which would certainly be a limiting factor to the long-term stability. This field could be reduced in a second iteration with appropriate multi-layer shielding. However, the issue then becomes how accurately the second-order shift can be measured. This is a rather delicate measurement dependent on the mean C-field strength, its inhomogeneity, and its temporal instability. In a fountain clock, this is performed by mapping out the field across the atoms' trajectory by launching the cloud to different heights above the cavity. In this experiment, mapping out the field would be complex because the distance the atoms are dropped cannot be varied. Therefore, the second-order Zeeman shift could potentially limit the fractional accuracy. By assuming some conservative values, we can estimate the expected field instability required to achieve, say 5×10^{-14} , an order magnitude better than the commercial beam clocks [60]. For a C-field of $0.5 \mu\text{T}$, the instability would have to be smaller than $0.01 \mu\text{T}$. This emphasises the need for suitable shielding and an appropriate method to accurately measure the field across the dropped path of the atoms.

8.2 Conclusions

This thesis has detailed my work to develop a compact cold atom microwave clock prototype. A relatively straightforward experimental clock sequence involving the release and interrogation of a cloud of cold atoms during free fall is proposed. This method has already been implemented in some existing clock arrangements, but most are restricted to short drop distances, which limits the achievable interrogation time. Here, we propose to drop the cloud over a larger distance, thus opening up the possibility of increased interrogation time and enhanced short-term stability performance. To demonstrate the feasibility of this approach, we have constructed an experimental arrangement where the procedure of laser cooling and dropping a cloud of cold atoms can be investigated. The experiment is comprised of three key areas: a laser system, a physics package, and electronics for experimental control. Each of these components was designed and constructed from the ground up.

A dual laser system comprised of conventional free space optics on an optical bench was assembled to perform laser cooling/trapping, state preparation, and state detection. An ultra-high-vacuum system forms the main body of the physics package and houses the Cs atomic reference. Off-the-shelf vacuum components were used to maintain flexibility and keep options open for potential modifications. Additionally, we have employed a retro-reflected MOT to generate cold atoms with the aim of keeping things compact. We have also designed and assembled an embedded electronics system that was successful in controlling the experiment, performing measurements, and running the entire clock sequence. The electronics are comprised primarily of evaluation boards, forgoing the need for bulky benchtop equipment. The electronics architecture developed here offers significant versatility and can be adapted for various other clock arrangements or experiments involving cold atoms.

In tandem with this DPhil research and in collaboration with another compact clock development at NPL, we have designed a microwave cavity for use in a mini-fountain clock. While the final designed cavity is not directly applicable to this work, it does have many design features that could be transferred to other compact microwave clocks such as this one. For example, a single feeding ring waveguide to simplify phase and amplitude balancing criteria, an intentionally low Q-factor to minimise sensitivity to temperature fluctuations, RF windows to seal the main cavity from the ring to simplify microwave inputs, and an assembled package that forms part of the total vacuum vessel. Throughout the design process, several simulations using COMSOL modelling software were performed, which not only supported a number of these design choices but also validated the feasibility of some new design concepts.

After constructing the experimental apparatus and successfully observing cold atoms in the MOT, the focus turned towards characterising the system. We have investigated a detection scheme whereby the atoms are probed using a three-laser pulse sequence to determine the normalised fraction of atoms excited during the clock state interrogation. This scheme uses three temporally separated beams rather than the conventional three spatially separated detection zones used in larger lab-scale clocks. This scheme is effective in reducing the overall clock size but is sensitive to the detuning of detection beams.

The SNR of the detection system was measured to be 500, which is suspected to be limited by laser frequency or intensity noise in the detection beams. Although, even if optimised, these results emphasise the necessity to perform a more advanced state selection process. Implementing this would enable up to an order of magnitude improvement to the short-term performance, and the added experimental complexity accompanying this would not significantly impact the overall

clock size.

Even though a cavity was not implemented in this experiment, microwave spectroscopy was still performed using a horn antenna. We have demonstrated Fourier-limited Rabi spectroscopy with a short pulse time of 5 ms. This length was chosen primarily for diagnostic purposes and was never intended as the final interrogation time. In addition to open-loop spectroscopy, we have also attempted to operate the experiment as a clock by disciplining a free-running OCXO using the Rabi signal. The short-term stability was measured to be $1.9 \times 10^{-9} \tau^{-1/2}$, which is an order of magnitude worse than the estimated behaviour given the open-loop data because the implementation of the OCXO feedback is not appropriately optimised.

This prototype is still limited by a few experimental systematics, which should be overcome with more optimisation. After optimisation, the stability can become competitive with some of the existing systems. However, the results here are encouraging as they highlight areas to improve, which can ultimately feed into the design of a second iteration with the potential to be significantly smaller and exhibit enhanced performance. Once the desired short-term stability can be demonstrated in this larger experimental setup, the attention can shift more towards the engineering challenges of miniaturisation. We envisage that the final size for a system like this could be approximately the size of a desktop computer ($\sim 50\,000 \text{ cm}^3$). Such a clock would be highly suited for applications such as GNSS holdover, radio astronomy, and operational timescales.

The work presented here represents a significant step towards such a system.

This page intentionally left blank.

Bibliography

- [1] R. M. Godun, “Atomic clocks compared with astounding accuracy,” *Nature*, vol. 591, no. 7851, pp. 534–535, 2021. URL: <https://www.nature.com/articles/d41586-021-00738-0>
- [2] S. N. Lea, “Limits to time variation of fundamental constants from comparisons of atomic frequency standards,” *The European Physical Journal Special Topics*, vol. 163, no. 1, pp. 37–53, 2008. URL: <http://link.springer.com/10.1140/epjst/e2008-00808-6>
- [3] B. M. Roberts *et al.*, “Search for transient variations of the fine structure constant and dark matter using fiber-linked optical atomic clocks,” *New Journal of Physics*, vol. 22, no. 9, p. 093010, 2020. URL: <https://iopscience.iop.org/article/10.1088/1367-2630/abaace>
- [4] P. Wcisło *et al.*, “New bounds on dark matter coupling from a global network of optical atomic clocks,” *Science Advances*, vol. 4, no. 12, 2018. URL: <https://www.science.org/doi/10.1126/sciadv.aau4869>
- [5] O. Kozlova, S. Guérandel, and E. de Clercq, “Temperature and pressure shift of the Cs clock transition in the presence of buffer gases: Ne, N₂, Ar,” *Physical Review A*, vol. 83, no. 6, p. 062714, 2011. URL: <https://link.aps.org/doi/10.1103/PhysRevA.83.062714>
- [6] G. A. Pitz, D. E. Wertepny, and G. P. Perram, “Pressure broadening and shift of the cesium D_1 transition by the noble gases and N₂, H₂, HD, D₂, CH₄, C₂H₆, CF₄, and ³He,” *Physical Review A*, vol. 80, no. 6, p. 062718, 2009. URL: <https://link.aps.org/doi/10.1103/PhysRevA.80.062718>
- [7] J. D. Elgin *et al.*, “A cold-atom beam clock based on coherent population trapping,” *Applied Physics Letters*, vol. 115, no. 3, p. 033503, 2019. URL: <http://aip.scitation.org/doi/10.1063/1.5087119>
- [8] R. Elvin *et al.*, “Cold-atom clock based on a diffractive optic,” *Optics Express*, vol. 27, no. 26, p. 38359, 2019. URL: <http://arxiv.org/abs/1909.04361>
- [9] R. Szmuk *et al.*, “Stability of a trapped-atom clock on a chip,” *Physical Review A*, vol. 92, no. 1, p. 012106, 2015. URL: <https://link.aps.org/doi/10.1103/PhysRevA.92.012106>
- [10] F. G. Ascarrunz *et al.*, “A Portable Cold 87Rb Atomic Clock with Frequency Instability at One Day in the 10⁻¹⁵ Range,” in *2018 IEEE International Frequency Control Symposium (IFCS)*, IEEE, pp. 1–3, 2018. URL: <https://ieeexplore.ieee.org/document/8597585/>
- [11] S. Lee *et al.*, “A compact cold-atom clock based on a loop-gap cavity,” *Applied Physics Letters*, vol. 119, no. 6, pp. 1–5, 2021. URL: <https://pubs.aip.org/apl/article/119/6/064002/41912/A-compact-cold-atom-clock-based-on-a-loop-gap>
- [12] “University of Basel Kings’ Valley Project,” [Date Accessed: 2024-04-02]. URL: <https://aegyptologie.unibas.ch/forschung/projekte/university-of-basel-kings-valley-project/report-2013/>

- [13] P. Ariotti, "Galileo on the Isochrony of the Pendulum," *Isis*, vol. 59, no. 4, pp. 414–426, 1968. URL: <https://www.journals.uchicago.edu/doi/10.1086/350426>
- [14] J. G. Yoder, *Unrolling Time: Christiaan Huygens and the Mathematization of Nature*, Cambridge University Press, 1989.
- [15] D. Sobel, *Longitude: The True Story of a Lone Genius Who Solved the Greatest Scientific Problem of His Time*, Bloomsbury USA, 2007.
- [16] W. Cady, "The Piezo-Electric Resonator," *Proceedings of the IRE*, vol. 10, no. 2, pp. 83–114, 1922. URL: <http://ieeexplore.ieee.org/document/1666707/>
- [17] I. I. Rabi *et al.*, "The Molecular Beam Resonance Method for Measuring Nuclear Magnetic Moments. The Magnetic Moments of ${}^6\text{Li}$, ${}^7\text{Li}$ and ${}^{19}\text{F}$," *Physical Review*, vol. 55, no. 6, pp. 526–535, 1939. URL: <https://link.aps.org/doi/10.1103/PhysRev.55.526>
- [18] N. F. Ramsey, "A Molecular Beam Resonance Method with Separated Oscillating Fields," *Physical Review*, vol. 78, no. 6, pp. 695–699, 1950. URL: <https://link.aps.org/doi/10.1103/PhysRev.78.695>
- [19] "The 13th Conférence Générale des Poids et Mesures (CGPM)," *Proceedings of the 13th CGPM*, 1967. URL: <https://doi.org/10.59161%2Fcgpm1967res1e>
- [20] M. A. Kasevich *et al.*, "rf spectroscopy in an atomic fountain," *Physical Review Letters*, vol. 63, no. 6, pp. 612–615, 1989. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.63.612>
- [21] A. Clairon *et al.*, "A cesium fountain frequency standard: preliminary results," *IEEE Transactions on Instrumentation and Measurement*, vol. 44, no. 2, pp. 128–131, 1995. URL: <http://ieeexplore.ieee.org/document/377790/>
- [22] P. Gill, "Is the time right for a redefinition of the second by optical atomic clocks?" *Journal of Physics: Conference Series*, vol. 723, no. 1, p. 012053, 2016. URL: <https://iopscience.iop.org/article/10.1088/1742-6596/723/1/012053>
- [23] J. R. Vig, "Introduction to Quartz Frequency Standards," *Research and Development Technical Report SLCET-TR-92-1*, pp. 1–55, 1992.
- [24] BIPM and IEC and IFCC and ILAC and ISO and IUPAC and IUPAP and OIML, "Evaluation of measurement data — Guide to the expression of uncertainty in measurement, 2008." URL: <https://www.bipm.org/documents/20126/2071204/JCGM%5C.101%5C.2008%5C.E.pdf/325dcaad-c15a-407c-1105-8b7f322d651c>
- [25] I. Hughes and T. Hase, *Measurements and their Uncertainties a practical guide to modern error analysis*, Oxford University Press, 2010.
- [26] D. Allan, "Statistics of atomic frequency standards," *Proceedings of the IEEE*, vol. 54, no. 2, pp. 221–230, 1966. URL: <http://ieeexplore.ieee.org/document/1446564/>
- [27] Vectron International, "Oscillator Short Term Stability and Allan Deviation," [Date Accessed: 2024-04-02]. URL: <https://ww1.microchip.com/downloads/aemDocuments/documents/VOP/ApplicationNotes/ApplicationNotes/Oscillator+Short+Term+Stability+and+Allan+Deviation.pdf>
- [28] F. Riehle, *Frequency Standards: Basics and Applications*, Wiley, 2003. URL: <https://onlinelibrary.wiley.com/doi/book/10.1002/3527605991>

- [29] W. M. Itano *et al.*, “Quantum projection noise: Population fluctuations in two-level systems,” *Physical Review A*, vol. 47, no. 5, pp. 3554–3570, 1993. URL: <https://link.aps.org/doi/10.1103/PhysRevA.47.3554>
- [30] C. L. Degen, F. Reinhard, and P. Cappellaro, “Quantum sensing,” *Reviews of Modern Physics*, vol. 89, no. 3, p. 035002, 2017. URL: <http://link.aps.org/doi/10.1103/RevModPhys.89.035002>
- [31] J. Vanier and C. Tomescu, *The Quantum Physics of Atomic Frequency Standards: Recent Developments*, CRC Press, 2015. URL: <https://www.taylorfrancis.com/books/9781466576971>
- [32] L. Maleki and J. Prestage, “Applications of clocks and frequency standards: from the routine to tests of fundamental models,” *Metrologia*, vol. 42, no. 3, pp. S145–S153, 2005. URL: <https://iopscience.iop.org/article/10.1088/0026-1394/42/3/S15>
- [33] J. Vanier and C. Audoin, *The Quantum Physics of Atomic Frequency Standards*, Adam Hilger, 1989. URL: <https://www.taylorfrancis.com/books/9781420050851>
- [34] G. Pescaroli *et al.*, “Cascading Effects of Global Positioning and Navigation Cascading Effects of Global Positioning and Navigation Satellite Service Failures,” *UCL IRDR and Mullard Space Science Laboratory Special Report*,” 2019. URL: <https://doi.org/10.14324/000.rp.10076568>
- [35] London Economics, “The economic impact on the UK of a disruption to GNSS, 2023.” URL: <https://www.gov.uk/government/publications/report-the-economic-impact-on-the-uk-of-a-disruption-to-gnss>
- [36] “International Collaboration Expands NRAO’s NINE Program with New Hub,” [Date Accessed: 2024-04-02]. URL: <https://public.nrao.edu/news/2018-nine-mapp/>
- [37] M. Rioja *et al.*, “THE IMPACT OF FREQUENCY STANDARDS ON COHERENCE IN VLBI AT THE HIGHEST FREQUENCIES,” *The Astronomical Journal*, vol. 144, no. 4, p. 121, 2012. URL: <https://iopscience.iop.org/article/10.1088/0004-6256/144/4/121>
- [38] S. Schediwy, “A Clock for the Square Kilometre Array,” in *Proceedings of From Antikythera to the Square Kilometre Array: Lessons from the Ancients — PoS(Antikythera & SKA)*,” 2013. URL: <https://pos.sissa.it/170/031>
- [39] T. W. Hänsch, “Nobel Lecture: Passion for precision,” *Reviews of Modern Physics*, vol. 78, no. 4, pp. 1297–1309, 2006. URL: <https://link.aps.org/doi/10.1103/RevModPhys.78.1297>
- [40] Z. L. Newman *et al.*, “Architecture for the photonic integration of an optical atomic clock,” *Optica*, vol. 6, no. 5, p. 680, 2019. URL: <https://opg.optica.org/abstract.cfm?URI=optica-6-5-680>
- [41] J. Vanier and C. Mandache, “The passive optically pumped Rb frequency standard: the laser approach,” *Applied Physics B*, vol. 87, no. 4, pp. 565–593, 2007. URL: <http://link.springer.com/10.1007/s00340-007-2643-5>
- [42] J. Vanier, “Atomic clocks based on coherent population trapping: a review,” *Applied Physics B*, vol. 81, no. 4, pp. 421–442, 2005. URL: <http://link.springer.com/10.1007/s00340-005-1905-3>
- [43] G. M. Saxena and B. Ghosal, *Rubidium Atomic Clock*, WORLD SCIENTIFIC, 2020. URL: <https://www.worldscientific.com/worldscibooks/10.1142/11249>

- [44] Microchip Technology Inc., “SA.45s CSAC - Datasheet, 2023,” [Date Accessed: 2024-04-25]. URL: <https://www.microsemi.com/product-directory/clocks-frequency-references/3824-chip-scale-atomic-clock-csac#resources>
- [45] H. M. Goldenberg, D. Kleppner, and N. F. Ramsey, “Atomic Hydrogen Maser,” *Physical Review Letters*, vol. 5, no. 8, pp. 361–362, 1960. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.5.361>
- [46] Microchip Technology Inc., “MHM-2020 Active Hydrogen Maser - Datasheet, 2023,” [Date Accessed: 2024-04-25]. URL: <https://ww1.microchip.com/downloads/aemDocuments/documents/FTD/ProductDocuments/Brochures/MHM-2020-Active-Hydrogen-Maser-00003090.pdf>
- [47] N. Demidov *et al.*, “Studies of a short and long - term stability of an active hydrogen maser with stand alone cavity auto tuning,” in *2012 European Frequency and Time Forum*, IEEE, pp. 488–493, 2012. URL: <http://ieeexplore.ieee.org/document/6502430/>
- [48] D. Morris, “Special Hydrogen Maser Workshop,” in *Proceedings of the 22th Annual Precise Time and Time Interval Systems and Applications Meeting*, pp. 349–354, 1990. URL: <https://www.ion.org/publications/abstract.cfm?articleID=14479>
- [49] MuQuans, “MuClock - Datasheet, 2019,” [Date Accessed: 2024-04-25]. URL: https://www.muquans.com/wp-content/uploads/2019/03/muquans_muclock.pdf
- [50] Spectradynamics, “cRb-CLOCK - Datasheet, 2023,” [Date Accessed: 2024-04-25]. URL: <https://spectradynamics.com/product-sheets/cRb-Clock-2023.pdf>
- [51] C. Xi *et al.*, “Coherent Population Trapping-Ramsey Interference in Cold Atoms,” *Chinese Physics Letters*, vol. 27, no. 11, p. 113201, 2010. URL: <https://iopscience.iop.org/article/10.1088/0256-307X/27/11/113201>
- [52] F.-X. Esnault *et al.*, “High-stability compact atomic clock based on isotropic laser cooling,” *Physical Review A*, vol. 82, no. 3, p. 033436, 2010. URL: <https://link.aps.org/doi/10.1103/PhysRevA.82.033436>
- [53] M. Langlois *et al.*, “Compact Cold-Atom Clock for Onboard Timebase: Tests in Reduced Gravity,” *Physical Review Applied*, vol. 10, no. 6, p. 064007, 2018. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.10.064007>
- [54] P. Liu *et al.*, “Scheme for a compact cold-atom clock based on diffuse laser cooling in a cylindrical cavity,” *Physical Review A*, vol. 92, no. 6, p. 062101, 2015. URL: <https://link.aps.org/doi/10.1103/PhysRevA.92.062101>
- [55] D. R. Scherer *et al.*, “Progress on a Miniature Cold-Atom Frequency Standard,” *arXiv*, pp. 154–163, 2014. URL: <http://arxiv.org/abs/1411.5006>
- [56] M. Vangeleyn *et al.*, “Laser cooling with a single laser beam and a planar diffractor,” *Optics Letters*, vol. 35, no. 20, p. 3453, 2010. URL: <https://opg.optica.org/abstract.cfm?URI=ol-35-20-3453>
- [57] B. L. S. Marlow and D. R. Scherer, “A Review of Commercial and Emerging Atomic Frequency Standards,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 68, no. 6, pp. 2007–2022, 2021. URL: <https://ieeexplore.ieee.org/document/9316270/>
- [58] Safran, “MINIRAFS Rb Atomic Frequency Standard - Datasheet, 2023,” [Date Accessed: 2024-04-25]. URL: <https://safran-navigation-timing.com/wp-content/uploads/2023/03/miniRAFS-SAFRAN-Datasheet.pdf>

- [59] Stanford Research Systems, “PRS10 - Datasheet, 2023,” [Date Accessed: 2024-04-25]. URL: <https://www.thinksrs.com/downloads/pdfs/catalog/PRS10c.pdf>
- [60] Microchip Technology Inc., “5071A - Datasheet, 2023,” [Date Accessed: 2024-04-25]. URL: <https://www.microsemi.com/product-directory/cesium-frequency-references/4115-5071a-cesium-primary-frequency-standard#resources>
- [61] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*, Springer New York, 1999. URL: <http://link.springer.com/10.1007/978-1-4612-1470-0>
- [62] C. Foot, *Atomic Physics (Oxford Master Series in Atomic, Optical and Laser Physics)*, Oxford University Press, 2005.
- [63] T. Hänsch and A. Schawlow, “Cooling of gases by laser radiation,” *Optics Communications*, vol. 13, no. 1, pp. 68–69, 1975. URL: <https://linkinghub.elsevier.com/retrieve/pii/0030401875901595>
- [64] D. A. Steck, “Cesium D Line Data, 2023.” URL: <https://steck.us/alkalidata/>
- [65] W. Ertmer *et al.*, “Laser Manipulation of Atomic Beam Velocities: Demonstration of Stopped Atoms and Velocity Reversal,” *Physical Review Letters*, vol. 54, no. 10, pp. 996–999, 1985. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.54.996>
- [66] J. V. Prodan and W. D. Phillips, “Chirping the light—fantastic? Recent NBS atom cooling experiments,” *Progress in Quantum Electronics*, vol. 8, no. 3-4, pp. 231–235, 1984. URL: <https://linkinghub.elsevier.com/retrieve/pii/0079672784900193>
- [67] R. N. Watts and C. E. Wieman, “Manipulating atomic velocities using diode lasers,” *Optics Letters*, vol. 11, no. 5, p. 291, 1986. URL: <https://opg.optica.org/abstract.cfm?URI=ol-11-5-291>
- [68] V. S. Bagnato *et al.*, “Continuous Stopping and Trapping of Neutral Atoms,” *Physical Review Letters*, vol. 58, no. 21, pp. 2194–2197, 1987. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.58.2194>
- [69] W. D. Phillips and H. Metcalf, “Laser Deceleration of an Atomic Beam,” *Physical Review Letters*, vol. 48, no. 9, pp. 596–599, 1982. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.48.596>
- [70] S. Chu *et al.*, “Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure,” *Physical Review Letters*, vol. 55, no. 1, pp. 48–51, 1985. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.55.48>
- [71] P. D. Lett *et al.*, “Observation of Atoms Laser Cooled below the Doppler Limit,” *Physical Review Letters*, vol. 61, no. 2, pp. 169–172, 1988. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.61.169>
- [72] J. Dalibard and C. Cohen-Tannoudji, “Laser cooling below the Doppler limit by polarization gradients: simple theoretical models,” *Journal of the Optical Society of America B*, vol. 6, no. 11, p. 2023, 1989. URL: <https://opg.optica.org/abstract.cfm?URI=josab-6-11-2023>
- [73] D. S. Weiss *et al.*, “Optical molasses and multilevel atoms: experiment,” *Journal of the Optical Society of America B*, vol. 6, no. 11, p. 2072, 1989. URL: <https://opg.optica.org/abstract.cfm?URI=josab-6-11-2072>
- [74] E. L. Raab *et al.*, “Trapping of Neutral Sodium Atoms with Radiation Pressure,” *Physical Review Letters*, vol. 59, no. 23, pp. 2631–2634, 1987. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.59.2631>

- [75] “GW Optics Component Library,” [Date Accessed: 2024-04-02]. URL: <https://www.gwoptics.org/ComponentLibrary/>
- [76] D. S. Simon, *A Guided Tour of Light Beams*, IOP Publishing, 2020. URL: <https://iopscience.iop.org/book/978-0-7503-3467-9>
- [77] E. A. Donley *et al.*, “Double-pass acousto-optic modulator system,” *Review of Scientific Instruments*, vol. 76, no. 6, 2005. URL: <https://pubs.aip.org/rsi/article/76/6/063112/353058/Double-pass-acousto-optic-modulator-system>
- [78] D. A. Smith and I. G. Hughes, “The role of hyperfine pumping in multilevel systems exhibiting saturated absorption,” *American Journal of Physics*, vol. 72, no. 5, pp. 631–637, 2004. URL: <https://pubs.aip.org/ajp/article/72/5/631/1038946/The-role-of-hyperfine-pumping-in-multilevel>
- [79] “QubeCL,” [Date Accessed: 2024-04-02]. URL: <https://www.ppqsense.com/qubecl-2/>
- [80] W. Ertmer, “Laser cooling and storage of free atoms,” *Physica Scripta*, vol. 36, no. 2, pp. 306–311, 1987. URL: <https://iopscience.iop.org/article/10.1088/0031-8949/36/2/020>
- [81] K. I. Lee *et al.*, “Single-beam atom trap in a pyramidal and conical hollow mirror,” *Optics Letters*, vol. 21, no. 15, p. 1177, 1996. URL: <https://opg.optica.org/abstract.cfm?URI=ol-21-15-1177>
- [82] H. J. Lee *et al.*, “Raman Cooling of Atoms in an Optical Dipole Trap,” *Physical Review Letters*, vol. 76, no. 15, pp. 2658–2661, 1996. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.76.2658>
- [83] A. Isichenko *et al.*, “Photonic integrated beam delivery for a rubidium 3D magneto-optical trap,” *Nature Communications*, vol. 14, no. 1, p. 3080, 2023. URL: <https://www.nature.com/articles/s41467-023-38818-6>
- [84] S. Lee *et al.*, “Operating Atomic Fountain Clock Using Robust DBR Laser: Short-Term Stability Analysis,” *IEEE Transactions on Instrumentation and Measurement*, vol. 66, no. 6, pp. 1349–1354, 2017. URL: <http://ieeexplore.ieee.org/document/7803597/>
- [85] R. Caldani *et al.*, “A prototype industrial laser system for cold atom inertial sensing in space,” *The European Physical Journal D*, vol. 73, no. 12, p. 248, 2019. URL: <http://link.springer.com/10.1140/epjd/e2019-100360-2>
- [86] D. O. Sabulsky *et al.*, “A fibered laser system for the MIGA large scale atom interferometer,” *Scientific Reports*, vol. 10, no. 1, p. 3268, 2020. URL: <https://www.nature.com/articles/s41598-020-59971-8>
- [87] “ILS Series - Frequency-stabilized laser systems dedicated to atom cooling and trapping,” [Date Accessed: 2024-04-08]. URL: <https://www.ixblue.com/photonics-space/turn-key-frequency-stabilized-laser-systems/>
- [88] “The Earth’s Magnetic Field: An Overview,” [Date Accessed: 2024-04-02]. URL: http://www.geomag.bgs.ac.uk/education/earthmag.html#_Toc2075549
- [89] A. de Marchi, G. D. Rovera, and A. Premoli, “Pulling by Neighbouring Transitions and its Effects on the Performance of Caesium-Beam Frequency Standards,” *Metrologia*, vol. 20, no. 2, pp. 37–47, 1984. URL: <https://iopscience.iop.org/article/10.1088/0026-1394/20/2/002>
- [90] J. Arlt *et al.*, “A pyramidal magneto-optical trap as a source of slow atoms,” *Optics Communications*, vol. 157, no. 1-6, pp. 303–309, 1998. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0030401898004994>

- [91] W. Bowden *et al.*, “A pyramid MOT with integrated optical cavities as a cold atom platform for an optical lattice clock,” *Scientific Reports*, vol. 9, no. 1, p. 11704, 2019. URL: <https://www.nature.com/articles/s41598-019-48168-3>
- [92] J. Lee *et al.*, “Sub-Doppler cooling of neutral atoms in a grating magneto-optical trap,” *Journal of the Optical Society of America B*, vol. 30, no. 11, p. 2869, 2013. URL: <https://opg.optica.org/abstract.cfm?URI=josab-30-11-2869>
- [93] J. P. McGilligan *et al.*, “Diffraction-grating characterization for cold-atom experiments,” *Journal of the Optical Society of America B*, vol. 33, no. 6, p. 1271, 2016. URL: <https://opg.optica.org/abstract.cfm?URI=josab-33-6-1271>
- [94] C. C. Nshii *et al.*, “A surface-patterned chip as a strong source of ultracold atoms for quantum technologies,” *Nature Nanotechnology*, vol. 8, no. 5, pp. 321–324, 2013. URL: <https://www.nature.com/articles/nnano.2013.47>
- [95] Q. Bodart *et al.*, “A cold atom pyramidal gravimeter with a single laser beam,” *Applied Physics Letters*, vol. 96, no. 13, 2010. URL: <https://pubs.aip.org/apl/article/96/13/134101/892720/A-cold-atom-pyramidal-gravimeter-with-a-single>
- [96] B. J. Little *et al.*, “A passively pumped vacuum package sustaining cold atoms for more than 200 days,” *AVS Quantum Science*, vol. 3, no. 3, 2021. URL: <https://pubs.aip.org/aqs/article/3/3/035001/570595/A-passively-pumped-vacuum-package-sustaining-cold>
- [97] F. Zheng *et al.*, “Microwave-Vacuum Integrated Cavity with a Low Temperature Sensitivity for Cs Fountain Clocks,” in *2021 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS)*, IEEE, pp. 1–4, 2021. URL: <https://ieeexplore.ieee.org/document/9604254/>
- [98] S. Jefferts, R. Drullinger, and A. DeMarchi, “NIST cesium fountain microwave cavities,” in *Proceedings of the 1998 IEEE International Frequency Control Symposium (Cat. No.98CH36165)*, IEEE, pp. 6–8, 1998. URL: <http://ieeexplore.ieee.org/document/717869/>
- [99] K. Gibble, S. N. Lea, and K. Szymaniec, “A microwave cavity designed to minimize distributed cavity phase errors in a primary cesium frequency standard,” in *2012 Conference on Precision electromagnetic Measurements*, IEEE, pp. 700–701, 2012. URL: <http://ieeexplore.ieee.org/document/6251122/>
- [100] D. M. Pozar, *Microwave Engineering, 4th Edition*, John Wiley & Sons, Inc, 2012.
- [101] R. Li and K. Gibble, “Phase variations in microwave cavities for atomic clocks,” *Metrologia*, vol. 41, no. 6, pp. 376–386, 2004. URL: <https://iopscience.iop.org/article/10.1088/0026-1394/41/6/004>
- [102] R. Li and K. Gibble, “Evaluating and minimizing distributed cavity phase errors in atomic clocks,” *Metrologia*, vol. 47, no. 5, pp. 534–551, 2010. URL: <https://iopscience.iop.org/article/10.1088/0026-1394/47/5/004>
- [103] G. Santarelli *et al.*, “Switching atomic fountain clock microwave interrogation signal and high-resolution phase measurements,” *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 56, no. 7, pp. 1319–1326, 2009. URL: <http://ieeexplore.ieee.org/document/5116858/>
- [104] S. Bize *et al.*, “Cavity frequency pulling in cold atom fountains,” *IEEE Transactions on Instrumentation and Measurement*, vol. 50, no. 2, pp. 503–506, 2001. URL: <http://ieeexplore.ieee.org/document/918177/>

- [105] K. Szymaniec *et al.*, “Evaluation of the primary frequency standard NPL-CsF1,” *Metrologia*, vol. 42, no. 1, pp. 49–57, 2005. URL: <https://iopscience.iop.org/article/10.1088/0026-1394/42/1/007>
- [106] K. Szymaniec *et al.*, “First accuracy evaluation of the NPL-CsF2 primary frequency standard,” *Metrologia*, vol. 47, no. 4, pp. 363–376, 2010. URL: <https://iopscience.iop.org/article/10.1088/0026-1394/47/4/003>
- [107] N. Ashby, S. Romisch, and S. Jefferts, “Endcaps for TE₀₁ cavities in fountain frequency standards,” in *IEEE International Frequency Control Symposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum, 2003. Proceedings of the 2003*, IEEE, pp. 1076–1083, 2003. URL: <http://ieeexplore.ieee.org/document/1275240/>
- [108] R. Wynands and S. Weyers, “Atomic fountain clocks,” *Metrologia*, vol. 42, no. 3, 2005.
- [109] F. Zheng *et al.*, “A Low Temperature-Sensitive Ramsey Cavity for Rb Fountain Clocks,” *Lecture Notes in Electrical Engineering*, vol. 910 LNEE, no. CsnC, pp. 393–402, 2022. URL: https://link.springer.com/10.1007/978-981-19-2576-4_35
- [110] “C110 OFE Copper CTE - Holmedodsworth file assets,” [Date Accessed: 2024-04-02]. URL: https://www.holmedodsworth.com/assets/file_assets/
- [111] “Sapphire CTE - ThorLabs,” [Date Accessed: 2024-04-02]. URL: https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=6973&tabname=Sapphire
- [112] R. Schroder, U. Hubner, and D. Griebisch, “Design and realization of the microwave cavity in the PTB caesium atomic fountain clock CSF1,” *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 49, no. 3, pp. 383–392, 2002. URL: <http://ieeexplore.ieee.org/document/990959/>
- [113] P. Cheng *et al.*, “An electronic controller based on FPGA for the Ramsey-CPT atomic clock,” in *2017 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFC)*, IEEE, pp. 616–617, 2017. URL: <http://ieeexplore.ieee.org/document/8088975/>
- [114] T. Pruttivarasin and H. Katori, “Compact field programmable gate array-based pulse-sequencer and radio-frequency generator for experiments with trapped atoms,” *Review of Scientific Instruments*, vol. 86, no. 11, 2015. URL: <https://pubs.aip.org/rsi/article/86/11/115106/356481/Compact-field-programmable-gate-array-based-pulse>
- [115] A. Roy *et al.*, “An FPGA based all-in-one function generator, lock-in amplifier and auto-relockable PID system,” *Journal of Instrumentation*, vol. 14, no. 05, pp. P05 012–P05 012, 2019. URL: <https://iopscience.iop.org/article/10.1088/1748-0221/14/05/P05012>
- [116] K. Huang *et al.*, “Microcontroller-based locking in optics experiments,” *Review of Scientific Instruments*, vol. 85, no. 12, 2014. URL: <https://pubs.aip.org/rsi/article/85/12/123112/109108/Microcontroller-based-locking-in-optics>
- [117] E. E. Eyler, “Instrumentation for laser physics and spectroscopy using 32-bit microcontrollers with an Android tablet interface,” *Review of Scientific Instruments*, vol. 84, no. 10, 2013. URL: <https://pubs.aip.org/rsi/article/84/10/103101/360181/Instrumentation-for-laser-physics-and-spectroscopy>
- [118] B. S. Malek *et al.*, “Embedded control system for mobile atom interferometers,” *Review of Scientific Instruments*, vol. 90, no. 7, 2019. URL: <https://pubs.aip.org/rsi/article/90/7/073103/360512/Embedded-control-system-for-mobile-atom>

- [119] W. Deng *et al.*, “Embedded Control System for Atomic Clock,” in *Advances in Intelligent and Soft Computing*, vol. 135, pp. 61–67. URL: http://link.springer.com/10.1007/978-3-642-27708-5_9
- [120] Anthony VH, “Cycle counter and nanosecond delays on Cortex-M,” [Date Accessed: 2024-04-02]. URL: https://www.anthonylvh.com/2017/05/18/cortex_m-cycle_counter/
- [121] M. Kumagai *et al.*, “Evaluation of caesium atomic fountain NICT-CsF1,” *Metrologia*, vol. 45, no. 2, pp. 139–148, 2008. URL: <https://iopscience.iop.org/article/10.1088/0026-1394/45/2/003>
- [122] T. P. Heavner *et al.*, “First accuracy evaluation of NIST-F2,” *Metrologia*, vol. 51, no. 3, pp. 174–182, 2014. URL: <https://iopscience.iop.org/article/10.1088/0026-1394/51/3/174>
- [123] S. Aubin *et al.*, “Trapping Fermionic 40K and Bosonic 87Rb on a Chip,” *Journal of Low Temperature Physics*, vol. 140, no. 5-6, pp. 377–396, 2005. URL: <http://link.springer.com/10.1007/s10909-005-7322-5>
- [124] G. J. Dick, “LOCAL OSCILLATOR INDUCED INSTABILITIES IN TRAPPED ION FREQUENCY STANDARDS,” in *Proceedings of the 19th Annual Precise Time and Time Interval Systems and Applications Meeting*, pp. 133–147, 1987.
- [125] G. Santarelli *et al.*, “Frequency stability degradation of an oscillator slaved to a periodically interrogated atomic resonator,” *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 45, no. 4, pp. 887–894, 1998. URL: <http://ieeexplore.ieee.org/document/710548/>
- [126] A. Mann *et al.*, “A high stability atomic fountain clock using a cryogenic sapphire interrogation oscillator,” in *Proceedings of the 1998 IEEE International Frequency Control Symposium (Cat. No.98CH36165)*, IEEE, pp. 13–17, 1998. URL: <http://ieeexplore.ieee.org/document/717871/>
- [127] J. Millo *et al.*, “Ultralow noise microwave generation with fiber-based optical frequency comb and application to atomic fountain clock,” *Applied Physics Letters*, vol. 94, no. 14, pp. 8–11, 2009. URL: <https://pubs.aip.org/apl/article/94/14/141105/324458/Ultralow-noise-microwave-generation-with-fiber>
- [128] S. Weyers, B. Lipphardt, and H. Schnatz, “Reaching the quantum limit in a fountain clock using a microwave oscillator phase locked to an ultrastable laser,” *Physical Review A*, vol. 79, no. 3, p. 031803, 2009. URL: <https://link.aps.org/doi/10.1103/PhysRevA.79.031803>
- [129] M. Knapp, “GitHub repository - cold_atom,” [Date Accessed: 2024-04-02]. URL: https://github.com/martinknapp18/cold_atom
- [130] D. Ibrahim, *PID-based Practical Digital Control With Raspberry Pi and Arduino Uno*, elektor, 2022.
- [131] R. W. G. Moore *et al.*, “Measurement of vacuum pressure with a magneto-optical trap: A pressure-rise method,” *Review of Scientific Instruments*, vol. 86, no. 9, 2015. URL: <https://pubs.aip.org/rsi/article/86/9/093108/282746/Measurement-of-vacuum-pressure-with-a-magneto>
- [132] C. G. Townsend, “Laser Cooling and Trapping of Atoms A thesis submitted for the degree of Doctor of Philosophy Merton College , Oxford Michaelmas 1995,” Ph.D. dissertation,” 1995.

- [133] P. D. Lett *et al.*, “Optical molasses,” *Journal of the Optical Society of America B*, vol. 6, no. 11, p. 2084, 1989. URL: <https://opg.optica.org/abstract.cfm?URI=josab-6-11-2084>
- [134] C. D. Wallace *et al.*, “Measurements of temperature and spring constant in a magneto-optical trap,” *Journal of the Optical Society of America B*, vol. 11, no. 5, p. 703, 1994. URL: <https://opg.optica.org/abstract.cfm?URI=josab-11-5-703>
- [135] T. M. Brzozowski *et al.*, “Time-of-flight measurement of the temperature of cold atoms for short trap-probe beam distances,” *Journal of Optics B: Quantum and Semiclassical Optics*, vol. 4, no. 1, pp. 62–66, 2002. URL: <https://iopscience.iop.org/article/10.1088/1464-4266/4/1/310>
- [136] L. Russell *et al.*, “Measurements on release–recapture of cold 85Rb atoms using an optical nanofibre in a magneto-optical trap,” *Optics Communications*, vol. 309, pp. 313–317, 2013. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0030401813007189>
- [137] I. Yavin *et al.*, “A calculation of the time-of-flight distribution of trapped atoms,” *American Journal of Physics*, vol. 70, no. 2, pp. 149–152, 2002. URL: <https://pubs.aip.org/ajp/article/70/2/149/1055770/A-calculation-of-the-time-of-flight-distribution>
- [138] S. Walby *et al.*, “Normalised Detection of Clock States by Cold Atom Recapture Method,” in *2022 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS)*, IEEE, pp. 1–3, 2022. URL: <https://ieeexplore.ieee.org/document/9850717/>
- [139] R. J. Hendricks *et al.*, “Cs Fountain Clocks for Commercial Realizations—An Improved and Robust Design,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 66, no. 3, pp. 624–631, 2019. URL: <https://ieeexplore.ieee.org/document/8485663/>
- [140] S. Weyers *et al.*, “First performance results of PTB’s atomic caesium fountain and a study of contributions to its frequency instability,” *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 47, no. 2, pp. 432–437, 2000. URL: <http://ieeexplore.ieee.org/document/827431/>
- [141] S. Blatt *et al.*, “Rabi spectroscopy and excitation inhomogeneity in a one-dimensional optical lattice clock,” *Physical Review A*, vol. 80, no. 5, p. 052703, 2009. URL: <https://link.aps.org/doi/10.1103/PhysRevA.80.052703>
- [142] H. De Raedt *et al.*, “Quantum simulations and experiments on Rabi oscillations of spin qubits: Intrinsic vs extrinsic damping,” *Physical Review B*, vol. 85, no. 1, p. 014408, 2012. URL: <https://link.aps.org/doi/10.1103/PhysRevB.85.014408>
- [143] X. Ouyang *et al.*, “Distortion of Rabi oscillations in a compact cold-atom clock,” *Physical Review A*, vol. 103, no. 4, p. 043118, 2021. URL: <https://link.aps.org/doi/10.1103/PhysRevA.103.043118>
- [144] S. Mulholland *et al.*, “Laser-cooled ytterbium-ion microwave frequency standard,” *Applied Physics B*, vol. 125, no. 11, p. 198, 2019. URL: <http://link.springer.com/10.1007/s00340-019-7309-6>
- [145] J. G. Ziegler and N. B. Nichols, “Optimum Settings for Automatic Controllers,” *Journal of Fluids Engineering*, vol. 64, no. 8, pp. 759–765, 1942. URL: <https://asmedigitalcollection.asme.org/fluidsengineering/article/64/8/759/1155342/Optimum-Settings-for-Automatic-Controllers>

- [146] M. A. Joffe *et al.*, “Transverse cooling and deflection of an atomic beam inside a Zeeman slower,” *Journal of the Optical Society of America B*, vol. 10, no. 12, p. 2257, 1993. URL: <https://opg.optica.org/abstract.cfm?URI=josab-10-12-2257>
- [147] L. Han *et al.*, “Optical state selection process with optical pumping in a cesium atomic fountain clock,” *Chinese Physics B*, vol. 30, no. 8, p. 080602, 2021. URL: <https://iopscience.iop.org/article/10.1088/1674-1056/ac0698>

This page intentionally left blank.

Appendix A

Timeline of the Project: My Contributions

In my first year (2019/20), I spent a great deal of time investigating an idea that would hopefully provide enhanced performance to the already well-established vapour cell technology for compact atomic clocks. This was an idea proposed by my supervisors at NPL before I started the DPhil; however, I attempted to validate the concept with theory. After the first year, theoretical results suggested that this enhanced vapour cell approach would not prove feasible due to its increased complexity and minimal performance gains. As a result, I decided not to pursue this idea any further. After discussions with my supervisors and some of the research from my literature review, I opted to pursue the ‘dropping and interrogating’ architecture. This was motivated primarily by the fact that it was more straightforward to implement. I would be the sole person actively working full-time on this experiment; therefore, it was important to be realistic about what could be achieved within a reasonable amount of time for the DPhil. This is not to discredit the help or guidance I have received but rather to highlight that the onus was on myself to solve experimental problems and seek the advice of others.

My second year (2020/21) started amid the UK’s lockdown due to the COVID-19 pandemic. This prevented me from returning to the lab to begin experimental work. During the work-from-home period, I continued design work and investigated the potential idea of implementing a pyramid-MOT into the experiment. This is an idea that I think would be a good step towards miniaturisation but, unfortunately, was unable to implement here. When I eventually returned to the lab, there was nothing but an empty optical table. The first task was to set up the two laser systems for the cooling and repumper light. Both were designed and constructed entirely

by myself. After this, it was time to focus on the physics package and vacuum apparatus. The first attempt at cleaning and bakeout of the vacuum system produced pressures at the 10^{-10} mbar level, which was more than adequate to see cold atoms in the MOT. With the vacuum system built and the laser system ready, I could start to search for the MOT, which, to my delight, I saw the first signs of just before Christmas break in December 2021.

Halfway through my second year (2020/21), I had the idea to collaborate with the atomic fountain group at NPL. As part of the collaboration, I helped design a microwave cavity for a mini-fountain project. The idea was that the cavity would have several design elements that would also be applicable to my experiment, and I could potentially build a second cavity to implement into my experiment. However, the design process took longer than expected, and I eventually had to step away from lab work just as the manufactured cavity had been delivered. Another benefit of the collaboration was that it opened a discussion channel in which I could share and receive technical advice about my experiment. This is where the idea for the compact normalised detection scheme originated, as some work had already been done on this.

After successfully seeing cold atoms in the trap, the next step was characterising. Therefore, much of my third year (2021/22) was spent building the electronics and writing the software to control the experiment. The entire experimental control system was developed by myself from the ground up. This included choosing commercial evaluation boards, designing custom PCBs, populating custom PCBs once delivered, and integrating everything into the embedded system shown in section 6.1. The software was also written entirely by myself, an area in which I did not have much experience before starting the DPhil. Once up and running, I could execute specific experimental sequences to perform cloud temperature measurements and characterise the detection scheme.

In the final year (2022/23), I spent most of the time continuing to optimise and was able to obtain the first Rabi signal in April 2023. However, from these first signals, it was clear that more optimisation was needed (primarily the SNR). From discussions with others at NPL, I was able to improve this. I then focused on building the final aspects of the experiment (namely the OCXO) to implement clock operation. In June 2023, I obtained the first ADEV result, and I could only optimise a small amount before stopping lab work to spend time writing the thesis. The final year also included several conferences that I attended and presented at.

Appendix B

List of Abbreviations

1PPS	One pulse-per-second
ADC	Analogue-to-digital converter
ADEV	Allan deviation
AM	Amplitude modulation
AMD	Alkali metal dispenser
AOM	Acousto-optic modulator
AVAR	Allan variance
CGPM	Conférence Générale des Poids et Mesures
COMSOL	finite element analysis, solver, and simulation software
CSAC	Chip-scale atomic clock
Cs	Caesium
CSO	Cryogenically-cooled sapphire oscillator
CPT	Coherent population trapping
CPU	Central processing unit
CTE	Coefficient of thermal expansion
DAC	Digital-to-analogue converter
DBR	Distributed Bragg reflector
DC	Direct current
DCP	Distributed cavity phase
DFB	Distributed feedback
DPhil	Doctor of Philosophy
DROP	Double-resonance optical pumping

DUT	Device under test
EBW	Electron-beam weld
ECDL	External-cavity diode laser
EMI	Electromagnetic interference
EOM	Electro-optic modulator
FM	Frequency modulation
FPGA	Field programmable gate array
FWHM	Full width half maximum
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HWHM	Half width half maximum
IDE	Interactive development environment
LO	Local oscillator
MCU	Micro-controller unit
MOT	Magneto-optical trap
NA	Numerical aperture
NMI	National measurement institute
NPL	National Physical Laboratory
OCXO	Oven controlled crystal oscillator
OFC	Oxygen-free copper
OADEV	Overlapping Allan deviation
PCB	Printed circuit board
PGC	Polarization gradient cooling
PIC	Photonic integrated circuit
PID	Proportional integral derivative
PTB	Physikalisch-Technische Bundesanstalt
PM	Polarization maintaining
QPN	Quantum projection noise
Rb	Rubidium
RF	Radio frequency
RMS	Root mean square
SAS	Saturated absorption spectroscopy

SI	International System of Units
SKA	Square-kilometre Array
SNR	Signal-to-noise ratio
SoC	System-on-chip
SPI	Serial peripheral interface
SWaP	Size weight and power
SYRTE	Systèmes de Référence Temps-Espace
TX/RX	Transmit/Receive
UART	Universal asynchronous receiver / transmitter
UHV	Ultra-high vacuum
UTC	Coordinated Universal Time
VLBI	Very long baseline interferometry
VCO	Voltage-controlled oscillator