A Cold-Atom Fountain Clock from a Diffractive Optic

Ben Lewis

A thesis submitted for the degree of
Doctor of Philosophy

Department of Physics
University of Strathclyde
September 2022

Abstract

This thesis reports the design, assembly and characterisation of a new cold-atom fountain clock. Diffractive optics have previously been used to produce compact atom trapping systems. I build upon that research to launch the trapped atoms upwards, as in existing fountain clocks. Those clocks launch their atoms up to a metre high, passing through a microwave cavity. I also embed the diffractive optic within a microwave cavity allowing a small, integrated apparatus. This allows an interrogation time of 100 ms, intermediate between the 10 ms to 20 ms achievable without a launch, and the 500 ms observed in a full-scale fountain.

Coherent population trapping was used to test the new atomic fountain technique and observe its initial performance in a clock system, where it improved the clock by allowing a longer time-of-flight, with a best single-shot stability of 2.9×10^{-11} The diffractive optic was then mounted inside a microwave cavity, allowing in-situ atom trapping with only a single aperture for the trapping beam. Another clock sequence was realised here, using the microwave cavity to directly excite atomic transitions. The microwave excitation doubled the detection SNR, going from 42 to 95 with potential for further increases. The future combination of the fountain and in-cavity trapping techniques will produce a miniature version of the microwave fountain clocks which are central to modern frequency metrology. An analysis has been done of the potential performance such a clock could achieve, to highlight the most critical areas of design and guide future experiments.

Declaration

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

Signed: Ben Lewis Date: 21 Dec 2022

Acknowledgements

This thesis is only possible due to the help I have received from others throughout my PhD. My supervisors, Griff and Erling, have been a constant source of ideas, motivation, skills and knowledge. Your research group is a pleasure to work in, your doors are never closed, and you entertained my hare-brained ideas long enough to tell me if (and why) they were stupid.

The group at EQOP is a good one. Welcoming, collegial, fun. Most of what I learned over the last four years has come from your collective hands: be it optics, electronics, coding, writing or theory. My lab and office mates have kept me learning and entertained the whole four years. To those who joined me on bike rides, thank you, and I can only apologise for the number of hills!

My thanks also to my friends in the judo club, and elsewhere, who remind me that there is more to life than physics.

Sophie, thanks for bringing me on your adventures, and going on mine. Here's to the next one!

To my family, you've made me into the person I am today. You've been a safe haven when times were hard, especially during the pandemic. Adam, you've shown me that maybe PhD's aren't that hard after all. Mum and Dad, you've always supported me, backed whatever decision I made. Thank you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Timekeeping is a fundamental part of society. From the earliest calendars to the latest atomic clocks, measuring the progress of time has allowed us to understand the world around us. As technology has advanced, improved clocks have enabled both scientific research and technological applications. The earliest devices, water clocks and sand-glasses, could be used to time specific periods and determine rates such as the speed of a ship or the pulse rate of a doctor's patient.[1,2] Later, mechanical clocks were developed for navigation. The variation of local solar time against a reference clock allowed the calculation of latitude, and global position could be determined within 10s of km.[3] In the modern age, accurate clocks impact an enormous number of fields in ways both trivial and profound. Satellite navigation systems rely on a clock in each of the satellites,[4] general relativity has been probed by measuring the effect of time-dilation on clocks,[5,6] and almost the entire SI measurement system depends upon the accuracy of clocks that define the second.[7] The availability of a common global timescale with sub-second accuracy has become an accepted part of everyday life, enabling scheduling and communication across the world.

The most accurate clocks, by many orders of magnitude, are atomic clocks.[8] The resonant frequencies between atomic energy levels are universal among atoms of the same type, defined by the fundamental physics of the atomic structure. Because of this, clocks using the same atoms should always have *exactly* the same frequency, with no changes over time. In practice, external influences can perturb those 'perfect' frequencies, and I will discuss that in this thesis, but the potential accuracy of atomic references remains far superior to any mechanical oscillator.

When looking at the technological applications of clocks, it is clear that the usability of the clock is of vital importance. A clock that can be easily transported, powered and operated will be useful in many more situations than one which requires a laboratory full of equipment and staff to operate.

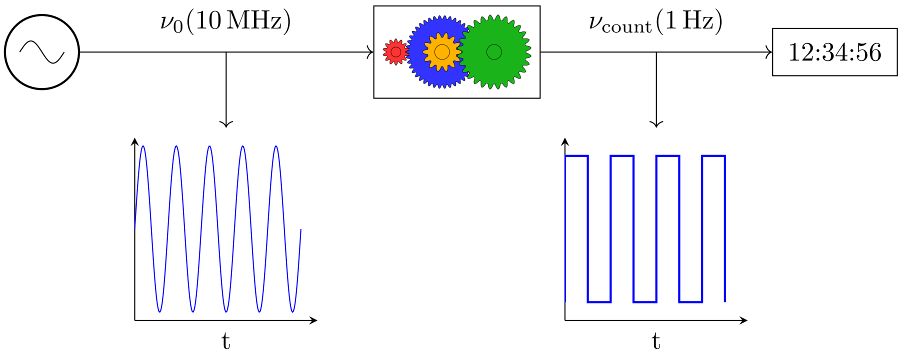

Figure 1.1: A general clock architecture consisting of a local oscillator, (electronic) clockwork and a counter. The shown waveforms and frequencies are examples.

The easier a clock is to use, the more applications it will find (for a given level of performance). There is now a great deal of focus on reducing the size, weight and power (SWAP) of existing high-performance clock technologies.

During my PhD I have worked to produce a clock which uses the techniques of a microwave cold-atom fountain clock but realises them in a more compact apparatus.

1.1 The Structure of a Clock

A clock consists of two main components: an oscillator and a counter. The oscillator, referred to as the local oscillator (LO), determines the clock frequency while the counter keeps track of how many cycles have occurred so the elapsed time is known (see Fig. 1.1). The clock performance will be determined by the stability and accuracy of the LO. There may also be some ‘clockwork’ that divides (or multiplies) the frequency of the LO to a frequency that is suitable for the counter. This frequency is often a standard of 1 Hz or 10 MHz as these are useful increments to count in common applications. For example a 32 kHz quartz crystal may be divided down to 1 Hz to drive a wristwatch display which only needs to update once per second. In optical clocks, this clockwork is of particular importance because the LO is at an optical frequency and the clockwork must transfer this frequency into the electrical domain whilst maintaing its frequency stability.[9]

In atomic clocks, the LO is enhanced by measuring its frequency, ν_{LO}, relative to the frequency of an atomic transition, ν_0. Any change or drift of the LO will be seen in this comparison and can be corrected. Fig. 1.2 shows this structure. The idea is that the atomic transition is far more stable than the LO in the long-term because it is determined by a comparatively simple physical system. The LO is more stable in the short-term and so a best-of-both scenario is achieved (the atomic system alone may not even have a well-defined short-term performance if it is passively

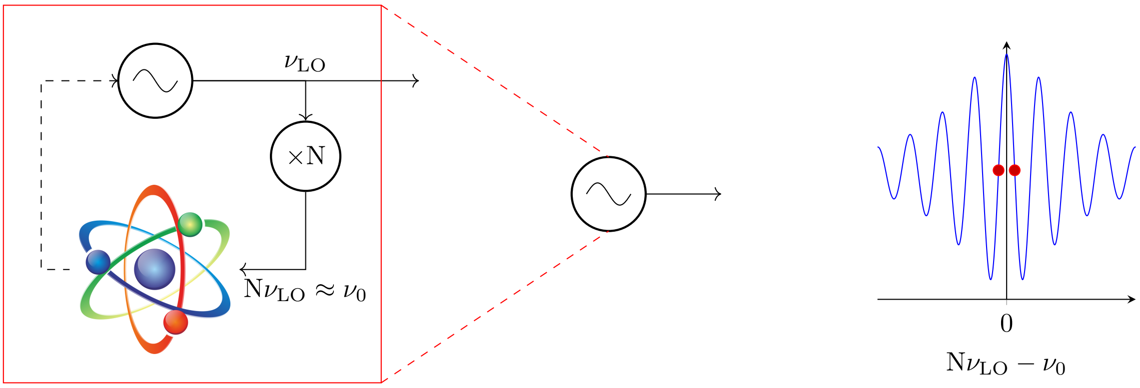

Figure 1.2: An atomic clock, where the local oscillator is locked to an atomic frequency standard (left). An example atomic signal that might be used for the lock (right).

interrogated, as opposed to an active oscillator). An unperturbed atomic transition should always have exactly the same frequency, in contrast to the classical oscillator which will drift over time (aging) even if unperturbed.

In general, the LO may be at a standardised frequency (e.g. 10 MHz) which will be multiplied (or divided) by a constant factor N (not necessarily an integer) so that it is close to the transition frequency, $N\nu_{\mathrm{LO}} = \nu_0 + \delta\nu$. The atoms are used to produce a signal that depends on $\delta\nu$ in some manner. Measurement of the signal allows calculation of $\delta\nu$, and is used to drive a feedback loop which maintains $\delta\nu = 0$, so the LO is locked to the atomic transition.

The principle challenge of building an atomic clock is to compare the LO and atomic frequencies with maximal precision, and without introducing systematic errors. This is often done by using the LO (after frequency multiplication) to drive the atomic transition directly. The resulting atomic state then provides a discriminator that is sensitive to $\delta\nu$.

1.2 Current clocks

Any proposed new clock should improve on the existing technologies in some way. It is important to understand the performance of the current generation of clocks and address their limitations. In general, clocks have two kinds of metric which can be improved; the performance of the clock in terms of stability (at various timescales) and the usability of the clock in terms of SWAP, cost, reliability, or ease of use. As might be expected, there is usually a tension between these different metrics. State of the art clock performance requires large, complex setups, whereas clocks which are compact or less expensive have much more limited performance. The performance of some next-generation clocks which have been demonstrated is illustrated in Fig. 1.3.

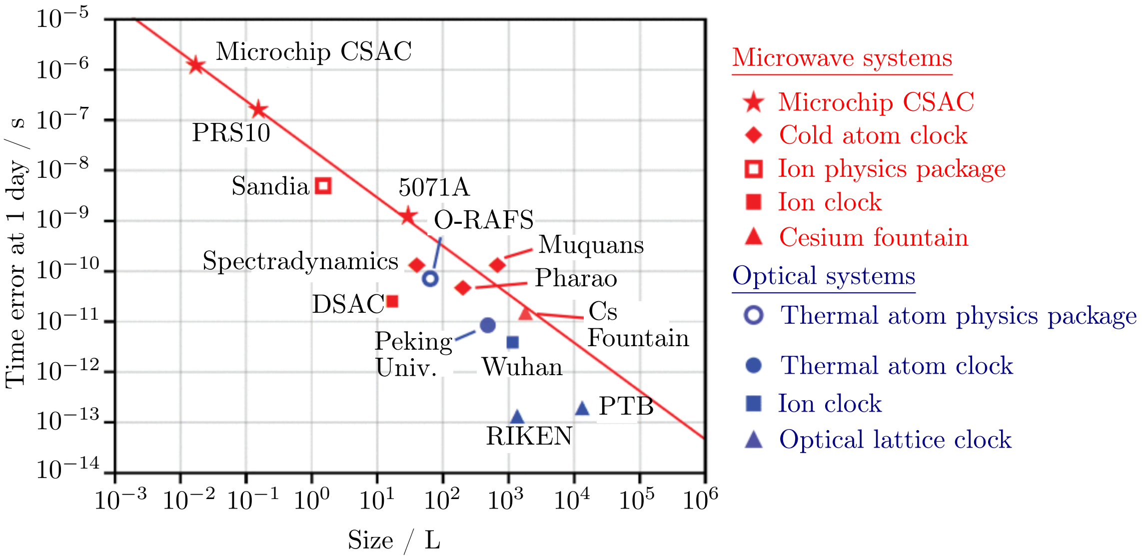

Figure 1.3: A comparison of existing clock performance, witha focus on developing next-generation clocks. Graph modified from [10], ©2021 IEEE. Y-axis is in time-deviation at 1 day, with 10^{-10} s equivalent to a fractional error of 2×10^{-15}.

Currently, the most accurate clocks in the world are *optical* clocks based upon cold atoms or trapped ions. These clocks can have an accuracy around 1 in 10^{18}, or 1 s in the age of the universe.[8,11] Optical clocks use a local oscillator (see Chapter 3) that oscillates at optical frequencies. In other words, they 'tick' extremely fast, at hundreds of THz. This, combined with long interrogation times of over 1 s, allows a very large number of clock cycles to be 'counted' in each measurement that is made, which is key to their impressive performance. The high frequency also has a number of downsides. It is hard to produce a highly stable oscillator at this frequency,[12,13] and it is also difficult to transfer the stability of the optical oscillations into the electrical domain so they can be used in electronic clock applications.[14,15] Further, each of these clocks requires a selection of lasers at multiple widely spaced frequencies, and so the optical systems become rather complex. These issues mean that it is currently very difficult to produce a compact, transportable optical clock. Transportable clocks do exist, and substantial progress is being made, but even the smallest require hundreds of litres of volume in addition to racks of electronic and optical equipment.[5,16–18] The difference between trapped-ion and cold atom optical clocks is worth examining. In some respects these are similar, as similar optical transitions and interrogation techniques are used for the clock itself. However, there are great differences in technology and performance. Ion clocks usually have only a single trapped ion which is trapped using an electromagnetic trap such as a Paul trap. The ions charge allows easier trapping, but only a single ion is used because they interact strongly with each other and this perturbs the clock. Because the trapping is electrical, not optical, it requires less optical access and can produce more compact physics packages. Cold atom clocks use *neutral* atoms, which must be trapped using optical forces in an optical lattice. Many more atoms can be trapped using this technique, in the range 1×10^5, which gives a much higher SNR due to the averaging power (see Section 3.2.4) and a better stability at short times (below 10^{-16} at 1 s for cold atom clocks but only 10^{-15} for trapped ion clocks.) The accuracy (and long-term stability) of both technologies is very comparable at a few 10^{-18}, because similar kinds of shift occur in each, with the major exception of shifts due to the trapping forces in each case. Even transportable clocks have achieved 10^{-17} accuracy.[17] An atomic clock will generally have improving instability over time, τ, as $1/\sqrt{\tau}$ so an optical lattice clock with 5×10^{-17} performance at 1 s would reach it's full accuracy specification in around 45 minutes, but a trapped ion clock might take 2 weeks to reach the same level.

In contrast to optical clocks, microwave atomic clocks use a local oscillator in the GHz range which is much more accessible using standard electronic components. These clocks usually work with neutral alkali atoms, such as Cs or Rb, and have somewhat simpler optical requirements because these atoms have relatively simple electronic structures.[19] Cs and Rb clocks are very similar because they have similar level structures, but there are some key differences. Cs has traditionally been used for the highest performing clocks, and the SI second is defined in terms of the ^{133}Cs frequency at around 9.912 GHz, which is higher than the ^{87}Rb frequency of 6.834 GHz. However it has been found that Rb does have some advantages, mainly that it has a much smaller clock shift due to collisions with other cold atoms, as discussed in Chapter 9. Rb also has fewer magnetic sublevels which can improve signal in some types of clock. In general, the species can be treated similarly. The best microwave clocks are the Cs fountain clocks which currently realise the SI second. These clocks have a short-term stability of around 10^{-13} at 1 s and an accuracy of a few 10^{-16}, and it is unlikely that any significant improvements will be made to this as the technology is now quite mature.[20–23] The field of compact microwave clocks is, likewise, much more developed than for optical clocks and has a number of commercially available clock systems.

Microwave atomic clocks can be broadly split into two groups, those that use a thermal gas of atoms in a vapour cell containing a substantial pressure of other gasses and those that use a collection of laser cooled atoms in a high or ultra-high vacuum system. The vapour cell systems usually have increased environmental sensitivity because the atoms are physically coupled to the temperature and pressure of their environment.[24,25] Vapour cell systems can be highly miniaturized and the chip-scale atomic clock (CSAC) is an example of this, having a volume of only 17 mL.[26] The cold-atom systems generally have much better long-term performance, with very little long-term drift, because the atoms are much more isolated from their environment. Additionally, cold temperatures correspond to smaller thermal velocities and low collision rates allowing extended interrogation times: 10s of ms or even 500 ms in a fountain clock instead of a few ms at most in thermal systems. This comes at the expense of requiring a vacuum and laser-cooling system. Cold-atom systems can still be miniaturised, though not to the same extent, and there are two commercial systems by Spectradynamics and MuQuans with volumes of 42 L and 700 L respectively.[27,28]

Finally, it is also worth considering those clocks not based on atomic transitions at all. Quartz oscillators are easy to use and have been extensively commercialised. Quartz is by far the most common kind of clock, and they can have good performance over short periods of time ($< 10\,\mathrm{s}$ to $100\,\mathrm{s}$),[29] but they are not at all accurate when compared to atomic clocks, having long-term drift of around 1 part in 10^8 per year for high-quality ovenized oscillators.[29–31] However, GPS signals contain timing information that is derived from high-performance atomic clocks, and this signal can be used to correct the drift of a quartz oscillator.[32] These GPS-disciplined oscillators (GPSDO) can provide a good compromise solution, with reasonable performance in both the short-term due to the quartz oscillator and the very-long term as the GPS time is locked within a few ns of International Atomic Time (TAI), which is referenced to a large number of atomic clocks.[33] This gives a stability around 10^{-12} at 1 s and a theoretical long-term performance limited by the accuracy of TAI, in the 10^{-16} region. However, this is a quite misleading comparison to atomic clocks. A GPSDO usually has fairly flat performance as τ increases from 1 s to about 1 day, before turning sharply downwards as $1/\tau$, whereas atomic clocks average purely as $1/\sqrt{\tau}$. This means that at many timescales of interest, from seconds to months, an atomic clock performs substantially better than a GPSDO. These systems are also completely reliant on a stable, high-quality GPS signal from a fixed antenna. This is significant when considering systems which must be robust to loss of GPS or applications which are mobile, underground, underwater, or in space.

We have identified that there is potential for improvements to compact cold-atom microwave clocks, by trapping atoms within the microwave cavity and therefore reducing the required size of apparatus. Further, use of grating trap technologies developed at Strathclyde (see Chapter 2) could reduce the amount and complexity of optics required for the clock. This thesis covers initial investigation and proof-of-principle for a clock using these techniques. The eventual performance targets for the clock are a few 10^{-13} stability at 1 s, 10^{-14} long-term accuracy, and $<10^2$ L system volume. The physics package should be $<10^1$ L.

1.3 Thesis Structure

This thesis consists of two parts.

Part I covers the theoretical background required for this thesis, including some novel contributions I have made towards clock stability statistics. Chapter 2 describes the optical interactions relevant to this thesis and their uses in trapping and cooling atoms. Chapter 3 describes the general structure of clocks and atomic clocks. The methods for measuring and describing clock performance are covered. Chapter 4 describes the Theo1 stability statistic and an improved algorithm that I developed for its calculation. Chapter 5 describes the coherent population trapping (CPT) mechanism and its use in atomic clocks.

Part II covers the design and construction of a cold atom microwave cavity clock. Chapter 6 describes the design process for the clock which I constructed for this thesis. Chapter 7 shows a new technique for launching a GMOT, and gives the initial characterisation measurements of that technique. The launch process is used as part of a CPT clock in Chapter 8, and the enhanced performance of the clock is shown. Chapter 9 gives estimations of the performance which could be achieved by an optimised version of the clock built for this thesis.

Chapter 10 concludes the thesis with a summary of the work that was completed, and an outlook on the future work which could build upon it.

1.4 Publications arising from this work

- **B. Lewis**, R. Elvin, A. Arnold, E. Riis, P. F. Griffin, "A grating-chip atomic fountain", Applied Physics Letters 121, 164001 (2022)
- **B. Lewis**, “A fast algorithm for calculation of Thêo1”, IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 67, 2187–2190 (2020)
- R. Elvin, G. W. Hoth, M. Wright, **B. Lewis**, A. Bregazzi, B. Keliehor, A. S. Arnold, P. F. Griffin, and E. Riis, “Progress of a compact microwave clock based on atoms cooled with a diffractive optic”, in Optical, opto-atomic, and entanglement-enhanced precision metrology ii, Vol. 11296 (International Society for Optics and Photonics, 2020),pp. 30–35.
- R. Elvin, M. W. Wright, **B. Lewis**, B. L. Keliehor, A. Bregazzi, J. P. McGilligan, A. S. Arnold, P. F. Griffin, and E. Riis, “Towards a compact, optically interrogated, cold-atom microwave clock”, Advanced Optical Technologies 9, 297–303 (2020)
- R. Elvin, G. W. Hoth, M. Wright, **B. Lewis**, J. P. McGilligan, A. S. Arnold, P. F. Griffin, and E. Riis, “Cold-atom clock based on a diffractive optic”, Opt. Express 27, 38359–38366 (2019)
- G. W. Hoth, R. Elvin, M. W. Wright, **B. Lewis**, A. S. Arnold, P. F. Griffin, E. Riis, F. Gruet, C. Affolderbach, and G. Mileti, “Impact of laser frequency noise in coherent population trapping with cold atoms”, in 2019 Joint Conference of EFTF and IFC (2019)

Part I

Theoretical Background

Chapter 2

Cold Atom Trapping

This chapter will introduce the conceptual background that is necessary for understanding cold-atom trapping. Much of this information has been explained and derived elsewhere, so this chapter serves as a summary and a pointer to other resources. The physics of atom-light interactions will be covered, followed by a discussion on how optical forces can be used to trap and cool atoms. Finally, more recent developments in compact cold atom traps will be discussed, covering the grating magneto-optical traps which are used in our group.

2.1 Atom-Light Interactions

In order for an atom to interact with an electromagnetic field, that electromagnetic field must perturb the atom, such that there is a change to the system Hamiltonian. In all the interactions of relevance to this thesis, the interaction will be through either an electric or magnetic dipole interaction. The details of such interactions can be found in many sources, such as [34].

The simplest example of such an interaction is a monochromatic electric field interacting with an atomic system of two states, $|1\rangle$ and $|2\rangle$, represented in column vector form as

$$|1\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} , \quad |2\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} . \tag{2.1}$$

For optical frequencies, this corresponds to a laser beam that is incident on the atom. In this case the Hamiltonian will have an interaction term

$$H_I = \hbar \begin{pmatrix} 0 & \Omega \cos(\omega t + \phi) \\ \Omega \cos(\omega t + \phi) & 0 \end{pmatrix} , \tag{2.2}$$

where $\Omega = X_{1,2}|E_0|/\hbar$ is the Rabi frequency which depends on $X_{1,2}$, the (polarisation-dependent) electric dipole matrix element between $|1\rangle$ and $|2\rangle$, and $|E_0|$, the electric field amplitude. Details of the dipole matrix element can be found in [34]. Note that ω, Ω and Δ (used later) refer to angular frequencies, while f, ν refer to ordinary frequencies.

The total system Hamiltonian will be formed from the sum of the interaction term and the unperturbed Hamiltonian, H_0, which contains the energies of the bare atomic states,

$$H_0 = \begin{pmatrix} E_1 & 0 \\ 0 & E_2 \end{pmatrix} , \tag{2.3}$$

giving us

$$H = H_0 + H_I = \hbar \begin{pmatrix} E_1/\hbar & \Omega\cos(\omega t + \phi) \\ \Omega^* \cos(\omega t + \phi) & E_2/\hbar \end{pmatrix} , \tag{2.4}$$

where $E_{1,2}$ are the energies of $|1\rangle, |2\rangle$ respectively. We may use the rotating wave approximation[34] (neglecting the most rapidly oscillating terms) and make a basis transformation to the rotating frame

$$e^{i(E_1/\hbar - \Delta/2)} |1\rangle \to |1\rangle , \qquad e^{i(E_2/\hbar + \Delta/2)t} |2\rangle \to |2\rangle \ , \tag{2.5}$$

with $\Delta = \omega - (E_2 - E_1)/\hbar$, the detuning of the laser from the atomic transition frequency. The effective Hamiltonian becomes

$$H = \frac{\hbar}{2} \begin{pmatrix} \Delta & \Omega \\ \Omega^* & -\Delta \end{pmatrix} , \tag{2.6}$$

where the phase, ϕ has been absorbed into the complex phase of Ω and Ω^*.

To see how an initial state will evolve with this Hamiltonian, we must solve the Schrödinger equation,

$$-i\hbar|\dot{\psi}\rangle = H\,|\psi\rangle \ . \tag{2.7}$$

If the system state is expressed as $|\psi(t)\rangle = c_1(t)\,|1\rangle + c_2(t)\,|2\rangle$, then the Schrödinger equation becomes:

$$i\dot{c_1}(t) = \Delta c_1(t)/2 + \Omega c_2(t)/2 \ , \tag{2.8}$$

$$i\dot{c_2}(t) = -\Delta c_2(t)/2 + \Omega^* c_1(t)/2 \ , \tag{2.9}$$

which has solutions of the form

$$c_1(t) = A\sin(\Omega' t/2) + B\cos(\Omega' t/2) \ , \tag{2.10}$$

$$c_2(t) = C\sin(\Omega' t/2) + D\cos(\Omega' t/2) \ , \tag{2.11}$$

where $\Omega'^2 = |\Omega|^2 + \Delta^2$ is the effective Rabi frequency. The solution for an initial state of $|1\rangle$ $(c_1(0) = 1, c_2(0) = 0)$ is

$$|\psi(t)\rangle = \left(\cos(\Omega' t/2) - i\frac{\Delta}{\Omega'}\sin(\Omega' t/2)\right)|1\rangle + \left(-i\frac{\Omega^*}{\Omega'}\sin(\Omega' t/2)\right)|2\rangle \ , \tag{2.12}$$

so the population of $|2\rangle$ is

$$P_2 = |c_2(t)|^2 = \frac{|\Omega|^2}{\Omega'^2} \sin^2(\Omega' t/2) \ , \tag{2.13}$$

which oscillates at the effective Rabi frequency, Ω'.

This Schrödinger description of light-atom interactions is sufficient for reversible processes. However, in real atomic systems, energy levels that are excited may be unstable and can decay by spontaneous emission of a photon. This is an irreversible process and cannot be modelled in the Schrödinger picture. The decay process occurs because the atom-light system is coupled to, and can emit photons into, the environment. This coupling process is modelled using the Lindblad formalism. In this formalism, the randomness of the decay process causes the pure quantum states to become a probabilistic mixture of states described by density matrices. The Schrödinger equation is extended to become the Gorini–Kossakowski–Sudarshan–Lindblad equation (also known as the Lindblad master equation) which operates on density matrices.[35] This equation contains an operator which models the spontaneous decay processes.

Phenomenologically, in the absence of excitation an excited state will decay over time, $P_2 \propto e^{-\Gamma t}$ where Γ is the decay rate determined by the details of the interaction between the two states. When the atom *is* excited, it can be shown that the system will evolve according to the 'optical Bloch equations'.[36] The transient behaviour is a little complex but a steady state will eventually be reached where the exciting field is balanced by spontaneous decay. In this case the population of the excited state is

$$P_2 = \frac{1}{2} \frac{s}{1 + s + \left(\frac{2\Delta}{\Gamma}\right)^2} \ , \tag{2.14}$$

where the saturation parameter, s, is

$$s = \frac{I}{I_{\text{sat}}} = \frac{2\Omega^2}{\Gamma^2} \ , \tag{2.15}$$

and I_{sat} is known as the saturation intensity. Notably, the steady-state excited population never exceeds 0.5. This is expected, because even in the Schrödinger case of no decay, with full oscillations between the two states, the average population in the excited state would only be 0.5.

In this steady state, photons are constantly absorbed from the laser beam and the re-emitted to maintain the steady state. Some photons are emitted as stimulated emission and will be emitted coherently back into the exciting laser beam. However some photons will be emitted by spontaneous decay, and these photons will be emitted in a random direction. These photons are also called scattered photons and will occur at a rate proportional to the excited state population, $\Gamma_{\text{sc}} = \Gamma P_2$. The spontaneously emitted photons cause a net force on the absorbing atom because the momentum of absorbed and emitted photons does not cancel. The net force is in the direction of the original laser beam and given by

$$F(\Delta, s) = \hbar k \Gamma_{\mathrm{sc}} = \hbar k \frac{\Gamma}{2} \frac{s}{1 + s + \left(\frac{2\Delta}{\Gamma}\right)^2} \ , \tag{2.16}$$

where k is the wavenumber of the beam.

2.2 Atom trapping

In order to trap cold atoms with optical forces, it is necessary for that force to depend upon the atoms' position and velocity. The velocity dependence allows a damping or cooling force, which reduces the atoms kinetic energy. Fortunately, suitable velocity dependence can be realised by proper engineering of Δ, the detuning of an incident laser beam. When an atom is moving with velocity $\vec{v}$ in a laser of wavevector $\vec{k}$ the apparent frequency of the beam (as observed by the atom) will be Doppler shifted by $\vec{k} \cdot \vec{v}$. This immediately produces a frequency-dependant force, $F(\Delta_0 - kv, s)$. If the beam is red-detuned for atoms at rest ($\Delta < 0$), then an atom moving towards the beam will be Doppler shifted closer to resonance, increasing the force and causing them to slow down faster. If two equal and opposing beams are present, then a dispersive force profile is produced, as in Fig. 2.1. The atom is more likely to absorb an atom from the beam which it is moving towards, so the motion slows over time. This system of opposed beams is referred to as an 'optical molasses' because it imposes a damping force analogous to that experienced by objects in a viscous liquid for sufficiently slow atoms.

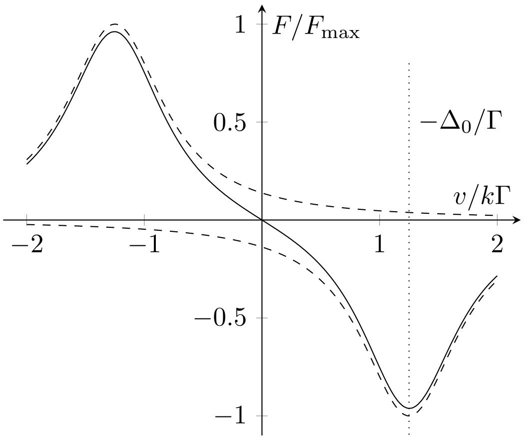

Figure 2.1: Net force applied to an atom in a 1D optical molasses as a function of velocity. The force from each beam (dashed) cannot cross zero, but the combination produces a symmetric restoring force. An initial detuning, Δ_0 is required to offset the peak force of each beam from zero, providing a strong gradient at equilibrium.

In order to extend this picture to 3D and also to understand magnetic effects, it is now necessary to consider the effects of field polarisation and the quantisation of angular momentum. If an atom is in a state with non-zero total angular momentum, F, then that angular momentum can be quantised along a particular direction, m_F with $|m_F| \leq F$. In the absence of a symmetry-breaking effect, the states with different m_F must be degenerate, but often there is such an effect. For example, if the states $|1\rangle$ and $|2\rangle$ have $F = 1$ then they are each subdivided into three states with $m_F = -1, 0, 1$. If there is a magnetic field, B, present along the quantisation direction, then these states will be separated in energy due to the Zeeman effect. The atoms magnetic dipole interacts with the magnetic field, adding a component to the atomic Hamiltonian, $V_{\mathrm{M}} \approx g_F \mu_B B m_F$ where g_F is a Landé g-factor and μ_B is the Bohr magneton.[34] The approximation is correct in the weak-field approximation where the magnetic interaction is much less than spin-orbit coupling. This splitting means that the atomic transition is split into multiple possible transitions between the upper and lower state manifolds. In principle, all these transitions could be driven, but only those which change m_F by $-1, 0, +1$ will have non-zero dipole matrix element. These types of transition are characterised as σ^-, π, σ^+ respectively, and are driven by different kinds of electric field polarisation. An electric field polarised along the quantisation axis will drive only π transitions, but σ transitions are driven if the field is left- or right-circular polarised around the quantisation axis. Any other polarisation or combination thereof can be decomposed into these three components, which are also called π and $\sigma^\pm$ polarisations. It is worth noting that I define these polarisations with respect to the chosen quantisation axis, not the direction of light propagation. There are opposing conventions that may be used elsewhere, so be cautious!

To understand how a trapping force can be generated, consider an $F = 0$ to $F = 1$ transition. The $F = 0$ ground-state only has the $m_F = 0$ sublevel, so the atoms are always available for transitions to all three $F = 1$ states, with $m_F = -1, 0, 1$. If a magnetic field gradient is present in the quantisation direction, z, so

$$\vec{B} \propto \begin{pmatrix} 0 \\ 0 \\ -2z \end{pmatrix}, \tag{2.17}$$

then the detuning of the $\sigma^\pm$ transitions will be $\Delta_\pm = \Delta_0 \pm \alpha z$ for some α. If we illuminate the atoms with opposing beams along the quantisation axis, then the amount of force from each beam depends on its polarisation. By giving the beams opposite polarisation, one beam can be shifted onto resonance while the other shifts away, see Fig. 2.2. This is directly equivalent to the optical molasses described previously, except the force is now dependant on position instead of velocity. In actuality the two effects occur in combination, to give both trapping and cooling forces in a magneto-optical trap (MOT).

The MOT can be extended to 3D quite easily by providing additional counter-propagating beams in each dimension. A pair of anti-Helmholtz coils can produce a quadrupole field which has zero

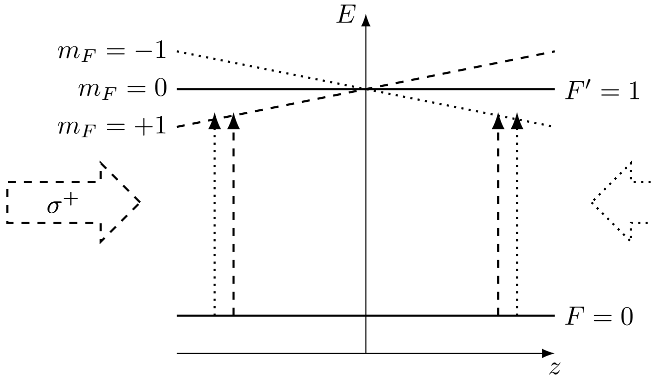

Figure 2.2: The energy levels in a MOT. The magnetic sublevels are Zeeman shifted by a magnetic gradient. The red-detuned σ^+ beam (dashed) comes on resonance as the atom moves left of centre, while the σ^- beam (dotted) comes on resonance if it moves right. The σ^+ and σ^- beams propagate to the left and right respectively, in order to generate a restoring force.

field centrally and gradient fields in all 3 dimensions,

$$\vec{B} \propto \begin{pmatrix} x \\ y \\ -2z \end{pmatrix} , \tag{2.18}$$

in the region near the origin. Strictly speaking the beams cannot be considered independent, as each beam changes the equilibrium state distribution which the other beams interact with. However, this simplification provides a good picture of the physics at work, and is much less complicated than trying to account for the fully coupled 3D behaviour. Additionally, real atoms may use transitions with other F numbers, giving a larger number of states and more complex pumping behaviour.

Although the simple picture of forces in Fig. 2.1 suggests cooling towards exactly zero velocity and position, in actuality random fluctuations of the scattering force (each photon is emitted in a random direction) cause a heating effect and the atoms become distributed in a thermal cloud around the equilibrium point. The minimum temperature reachable in this manner is the Doppler temperature, T_D, given by

$$k_B T_D = \frac{\hbar\Gamma}{2} . \tag{2.19}$$

This temperature can be understood qualitatively by considering the lifetime of the excited state, $1/\Gamma$, and then the ‘uncertainty relation’ between energy and time of $\delta E \delta t \geq \hbar/2$ sets the minimum energy scale that can be achieved in the system.[37] This qualitative explanation can be confirmed by detailed examination of the balance between Doppler cooling and the heating from spontaneous emission.[34] For ^{87}Rb (used in this thesis) the Doppler temperature is 146 µK.[19]

There are ways to cool below the Doppler temperature, the most common of which is polarisation-gradient cooling.[38] This cooling mechanism is achieved in the same optical setup as a MOT but without any magnetic field present. No trapping is present, only a velocity-damping force and so this technique is also a form of 'optical molasses'. On a small scale, the counter-propagating beams form an optical lattice, which alternates between different polarisations over wavelength scales. As atoms move across these 'polarisation gradients' they are optically pumped between different magnetic sublevels. This optical pumping can occur on a much slower timescale than $1/\Gamma$, and therefore allows a smaller energy uncertainty, and lower temperature. The optical pumping rate decreases as the optical intensity is decreased or the detuning is increased, giving $T \propto I/\Delta \propto \Omega^2/\Delta$. It can be made arbitrarily long, which seemingly gives arbitrarily small temperatures. The temperature is actually limited by the energy imparted by a single scattered photon. This recoil temperature, $T_R = \hbar^2 k^2/(mk_B)$, is around 0.37 µK for ^{87}Rb though temperatures of 2 µK to 10 µK are normally achieved.[39–41] Sub-recoil cooling methods do exist, though they will not be addressed in this thesis.[42,43]

Once the MOT technique developed, it quickly became the foundation of many experiments. A dense source of ultracold atoms provides access to a wide range of physics that would otherwise be inaccessible. Atoms released from a MOT are extremely well isolated from their environment and may remain so until acted on by an external force. Unfortunately, the released atoms will immediately start falling due to gravity and so the experiment time is limited by the size of the experimental system. One method to counteract this problem is to give the atoms an initial upwards velocity, causing the atoms to move upwards before falling back to their starting position. This technique is referred to as an atomic fountain.[44] Giving the atoms a velocity in another direction can be used to transfer atoms between different sections of an experiment. This initial velocity is usually imparted by using a frame transformation. If the entire atom cooling apparatus were moving, of course the atoms would inherit that initial velocity. However, this is quite impractical, so instead the *effects* of that motion are recreated. In an optical experiment, the dominant effect is the Doppler shifting of the laser beams. In the lab frame, each laser beam must be frequency shifted by an amount $\delta\omega = \vec{k} \cdot \vec{v}$ where $\vec{v}$ is the desired velocity. In the moving frame this frequency shift will be perfectly cancelled by the Doppler shift, restoring the initial cooling beam configuration. If a magnetic field gradient were present (e.g. in a MOT) this would complicate matters, as the gradient field would have to move with the desired frame, but the frame transformation is usually applied only during the sub-Doppler cooling, when no gradient is present. This technique is known as a 'moving molasses' and was first demonstrated in 1990.[45] Note that the atoms temperature is frame-independent because it refers to their velocity spread, not their absolute velocity.

2.3 Grating Magneto Optical Traps

Our group has been working on grating MOTs (GMOTs) for over a decade. Here I briefly review their operation, the motivation for their use, and the progress that has been made previously. I direct the reader to the previous theses that have been produced on this topic.[46–48]

One drawback of MOTs is that they may require a significant amount of optics to produce the six beams required (two per dimension). These optics take up space and require good optical access into the trapping region. Additionally, the MOT beams require careful alignment, and may need individual power stabilisation if long-term drifts are critical.[49] To facilitate more compact experiments, with fewer components, it would be beneficial to derive each of the MOT beams from a single source. This was first done using an inverted pyramid of mirrors to produce the six beams.[50] This technique was very effective, as the resulting beams are almost identical to the standard 6-beam MOT. However it has disadvantages; the pyramid construction may be difficult to integrate with more complex experimental setups, and the MOT forms within the pyramid structure giving very limited optical access to the MOT. Some work has been done to alleviate this second issue, using incomplete pyramids with space between the mirrors.[51] It was then shown that a MOT could also be formed using four beams in a tetrahedral configuration, generated by diffraction from etched grating patterns.[52,53] Because the grating is planar, it is often easier to integrate into a compact experimental setup.

A tetrahedral MOT does not have tetrahedral symmetry.[54] Only one beam is directed along the quadrupole field axis, and that beam has an opposite polarisation handedness to the other three. In a six-beam MOT, there are two counter-propagating beams along the axis, and four equally separated radial beams with opposite circular polarisation (relative to each beams propagation direction). A tetrahedral MOT maintains this distinction, but with only a single axial beam and three radial beams. The radial beams are angled below the radial plane to provide a balancing force to the single axial beam. It is harder to apply the 1D model of a MOT to this situation, because there is cross-coupling between all 3 dimensions, and a complex 3D optical lattice is formed. Fig. 2.3 illustrates the different MOT geometries.

The gratings we use are generally made from silicon, with a grating etched into their surface. The etching pattern may be produced by a number of lithographic techniques.[47] Electron-beam lithography has shown better performance than photolithography, perhaps due to increased resolution. The gratings in this thesis were manufactured using nano-imprint lithography. The gratings are then surface coated to improve their reflectivity. A number of coatings have been used such as gold and aluminium. For in-vacuum applications, gold may be unsuitable as it reacts with rubidium.

Previous work in our group has produced grating MOTs (GMOTs) which trap up to several 10^7 atoms and can cool atoms to 3 µK. In a GMOT, the overlap volume is considerably smaller than would be gained for a 6-beam MOT with the same beam size. This limits the number of atoms trapped. Further advances to the GMOT design may include 'chequerboard gratings' that use a single 2D grating to produce four diffracted trapping beams. Because the grating is not segmented into multiple regions, the trapping volume is greater (Fig. 2.3), increasing the trapped atom number.[40,55] However this technology is less well developed; previous work has struggled to balance the intensity of the diffracted orders and to remove the zero order reflection. Development is ongoing, but the linear grating MOTs are widely used within our group and their properties are much better understood, so they were used throughout this thesis.

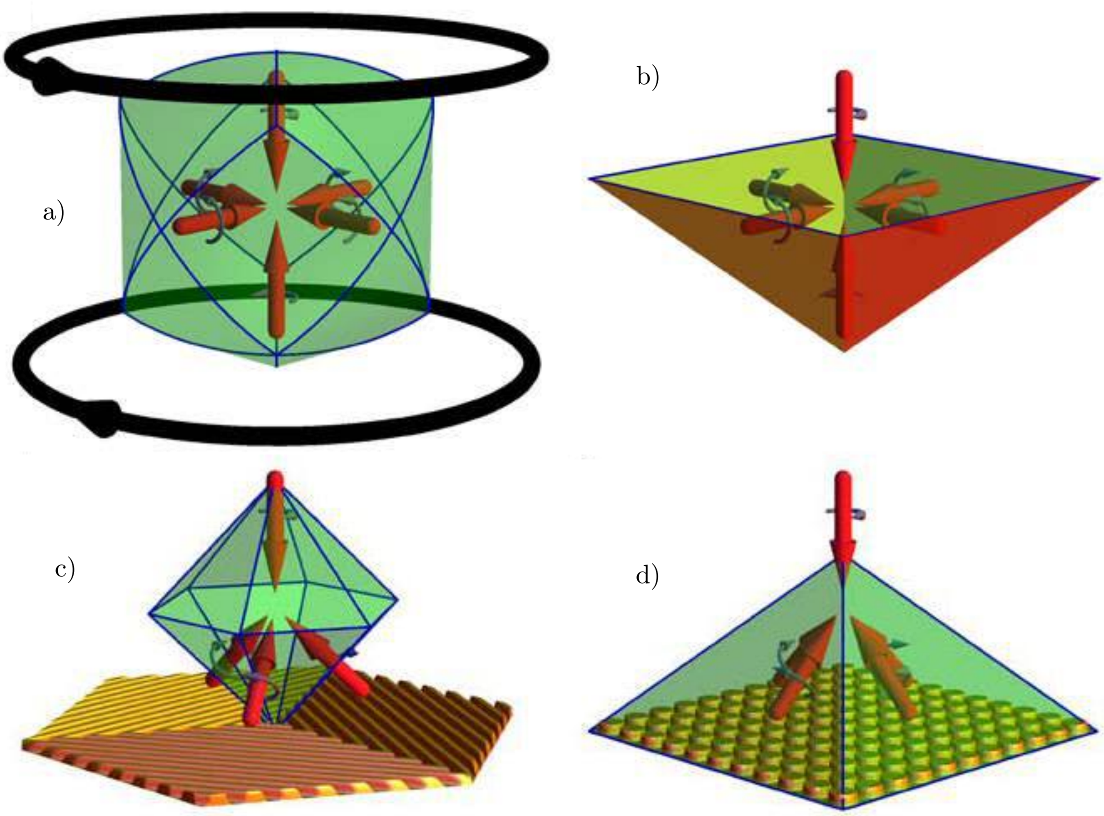

Figure 2.3: A comparison of different MOT geometries.[53] The beam k-vectors are shown by red arrows. Black tori (shown for (a) but implied for (b) - (d)) are magnetic quadrupole coils. Green zones indicate overlap volumes. a is (a) standard six-beam MOT, (b) is a pyramid MOT, (c) is a tetrahedral GMOT, and (d) is a chequerboard GMOT which increases the overlap volume for the same input beam. Reprinted by permission from Springer Nature Customer Service Centre GmbH: Nature Nanotechnology, A surface-patterned chip as a strong source of ultracold atoms for quantum technologies, C. C. Nshii et al. ©(2013)

Chapter 3

Atomic Clock Theory

This chapter will cover the background theory of atomic clock systems and their performance. Building on the clock architecture described in Chapter 1, I shall describe two methods by which atoms can be interrogated, so that a local oscillator may be locked to the atomic transition. Clock stability and its most common metric, the Allan variance, will be considered, with emphasis of the most common limiting factors for atomic clocks.

3.1 Clock Sequences

3.1.1 Rabi excitation

The task is to produce an atomic state which has maximal frequency sensitivity. The simplest solution is to excite the atomic transition with a simple pulse of field over some time τ. The atoms will undergo Rabi oscillations during that time, and the population in each state will change. Assuming the atoms are initially in $|1\rangle$, then equation (2.13) gives the transition probability of

$$P_{1\to2} = \frac{\Omega^2}{\Omega^2+\Delta^2}\sin^2\left(\sqrt{\Omega^2+\Delta^2}\frac{\tau}{2}\right) \ , \tag{3.1}$$

for a Rabi frequency of Ω. The detuning Δ is the angular frequency equivalent of $\delta\nu$ in Section 1.1 above, $\Delta = 2\pi\delta\nu$.

We see that the transition probability oscillates as Δ varies, shown graphically in Fig. 3.1. All the atoms will transfer if $\Delta = 0$ and the pulse area, $\Omega\tau = \pi$. This Rabi sequence is said to use a π pulse. The precision to which Δ can be measured is determined by the experimental signal to noise ratio (SNR) when measuring P_2 and the fringe width ($\Delta_{\mathrm{FWHM}} \approx \frac{1.6\pi}{\tau}$).

The effect of a Rabi pulse can also be shown on a Bloch sphere. In this visualisation, any pure

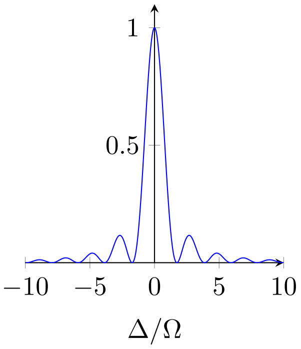

Figure 3.1: Rabi oscillations as a function of detuning from the atomic transition, Δ. The pulse area has been chosen such that $\Omega\tau = \pi$, giving full population transfer when $\Delta = 0$.

state of the two-level system is represented by a point on the unit sphere. The states $|1\rangle$ and $|2\rangle$ are represented by the poles:

$$|1\rangle \rightarrow \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \quad \text{and} \quad |2\rangle \rightarrow \begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix} , \tag{3.2}$$

and a general state is represented as

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|1\rangle + \sin\left(\frac{\theta}{2}\right)e^{i\phi}|2\rangle \rightarrow \begin{pmatrix} \sin(\theta)\cos\phi \\ \sin(\theta)\sin(\phi) \\ \cos(\theta) \end{pmatrix} = \vec{r} . \tag{3.3}$$

Converting the Schrödinger equation into the Bloch sphere representation, it can be shown that the state will evolve as

$$\dot{\vec{r}} = \vec{W} \times \vec{r} , \tag{3.4}$$

so $\vec{r}$ rotates around the new vector,

$$\vec{W} = \begin{pmatrix} \mathrm{Re}(\Omega) \\ \mathrm{Im}(\Omega) \\ \Delta \end{pmatrix} . \tag{3.5}$$

If $\Delta = 0$, then rotation is observed around a vector representing the magnitude and phase of Ω,

$$\vec{\Omega} = \begin{pmatrix} \mathrm{Re}(\Omega) \\ \mathrm{Im}(\Omega) \\ 0 \end{pmatrix} , \tag{3.6}$$

which can transfer perfectly from $|1\rangle$ to $|2\rangle$. As Δ is varied, $\vec{W}$ also changes and the initial $|1\rangle$ vector rotates more rapidly around a smaller circle, as shown in Fig. 3.2, and cannot complete a full transfer.

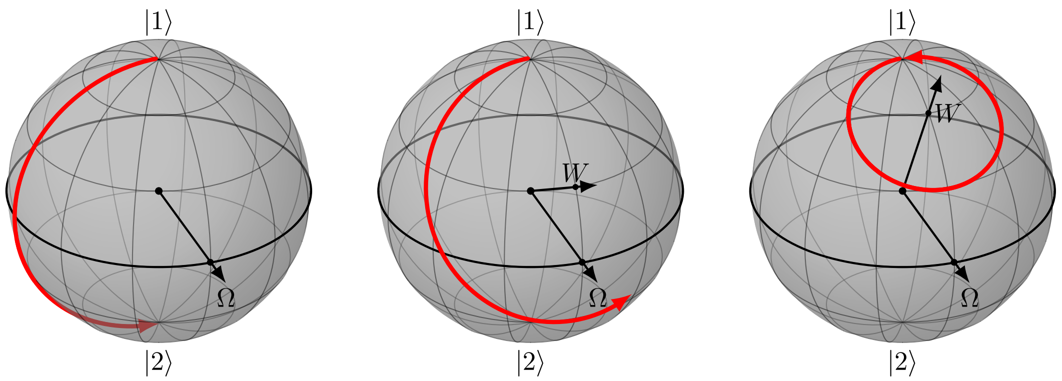

(a) $\Delta\tau = 0$ (b) $\Delta\tau = \pi\sqrt{3}/3$ (c) $\Delta\tau = \pi\sqrt{3}$

Figure 3.2: Bloch diagrams showing the effect of a Rabi pulse on an initial state $|1\rangle$, for different detunings, Δ. The initial state is rotated around $\vec{W}$.

3.1.2 Ramsey sequence

In many practical atomic clocks, a Ramsey sequence is used instead of the simple Rabi pulse. For this sequence, the π-pulse is split into two $\pi/2$ pulses separated by a 'Ramsey time', T. This has several benefits: the driving field need only be maintained for a small time, narrower fringes are achieved for the same total time. ($\Delta_{\mathrm{FWHM}} = \frac{\pi}{T}$, 60% narrower than a Rabi pulse with length T), and frequency shifts caused by effects during the pulses are 'diluted' by the longer Ramsey time. This can be extremely significant because during the Ramsey time the atoms are 'in the dark', not interacting with any electromagnetic field, and can ideally be made not to interact with any perturbing factors at all. In contrast the Rabi pulse requires a perturbation be applied for the entire sequence in order to realise the interrogation. At the time of this innovation, atomic beams were in use, where time and distance requirements are closely linked.[56] It is very difficult to engineer a uniform oscillating field over a region much larger than one field wavelength, which greatly limited the length of Rabi pulse which could be used. Ramsey's technique of separated fields allowed two separate field regions to be used, separated by large distances which correspond to longer times and narrower linewidths.

To model the Ramsey system two different Hamiltonians $H_{1,2}$ are required for during the pulses and between them respectively. These Hamiltonians, applying to dressed states in the rotating wave approximation as in equation (2.6), are

$$H_1 = \frac{\hbar}{2}\begin{pmatrix} \Delta_1 & \Omega \\ \Omega & \Delta_1 \end{pmatrix}, \quad H_2 = \frac{\hbar}{2}\begin{pmatrix} \Delta_2 & 0 \\ 0 & \Delta_2 \end{pmatrix}, \tag{3.7}$$

where $\Delta_{1,2}$ is the detuning of the LO from the atomic frequency during the two stages of the sequence. These detunings may be different if the pulses themselves cause a perturbation of the

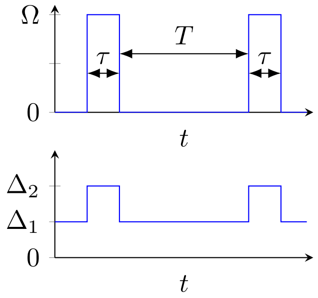

Figure 3.3: The Rabi frequency and detuning during a Ramsey sequence. Perturbations from the Rabi pulses cause the detuning to change as well as the Rabi frequency.

atomic energy levels, such as an AC Stark shift,[57] resulting in $\Delta_1 = \Delta_2 + \Delta_{\text{perturb}}$ as shown in Fig. 3.3. The meaning of Δ_2 requires some further explanation. The exciting field is not incident on the atoms between the pulses, so how can the detuning between this field and the atomic transition have a physical effect? The answer is that the field is assumed to still exist somewhere in the form of a local oscillator which ensures phase continuity between the two exciting pulses. Between the pulses, an additional phase difference can accrue between the local oscillator and the atomic phase. In essence, Δ_2 is important because it determines the phase difference between the atomic state and the exciting field at the start of the second pulse.

Solving the Schrödinger equation using H_1, H_2, H_1 for intervals of τ, T, τ respectively then gives the probability of transferring atoms from state $|1\rangle$ to state $|2\rangle$,

$$P_{1\to 2} = \left(\frac{\Omega}{\Omega'}\right)^2 \sin^2(\Omega'\tau)\left(\cos\left(\frac{\Delta_2 T}{2}\right) - \frac{\Delta_1}{\Omega'}\tan\left(\frac{\Omega'\tau}{2}\right)\sin\left(\frac{\Delta_2 T}{2}\right)\right)^2 , \tag{3.8}$$

where $\Omega'^2 = |\Omega|^2 + \Delta_1^2$.

In the case of perfect $\pi/2$ pulses, such that $\Omega'\tau = \pi/2$, and with $\Delta_1\tau \ll 1$, this can be simplified to

$$P_{1\to 2} \approx \frac{1}{2} + \frac{1}{2}\cos\left(\Delta_2 T + \Delta_1\frac{4\tau}{\pi}\right) \approx \frac{1}{2} + \frac{1}{2}\cos(\Delta_2 T) . \tag{3.9}$$

The effect of this Ramsey sequence can be visualised on a Bloch diagram (Fig. 3.4). The first $\pi/2$ pulse causes a rotation of the state vector from $|1\rangle$ to $(|1\rangle + |2\rangle)/\sqrt{2}$. During the Ramsey period, a phase difference of $\Delta_2 T$ is accrued between the LO and atomic state. The final $\pi/2$ pulse applies another $\pi/2$ rotation, but the accrued phase means this may now bring the state vector somewhere between $|1\rangle$ and $|2\rangle$.

A Ramsey scheme will produce fringes, similar to a Rabi scheme. However, the fringes are much closer to a true sinusoid, so many more fringes can be observed (within a Rabi envelope from the pulses). Additionally, a slightly narrower fringe width ($\delta\nu_{\text{FWHM}} \approx \frac{0.5}{T}$) is achieved for the same

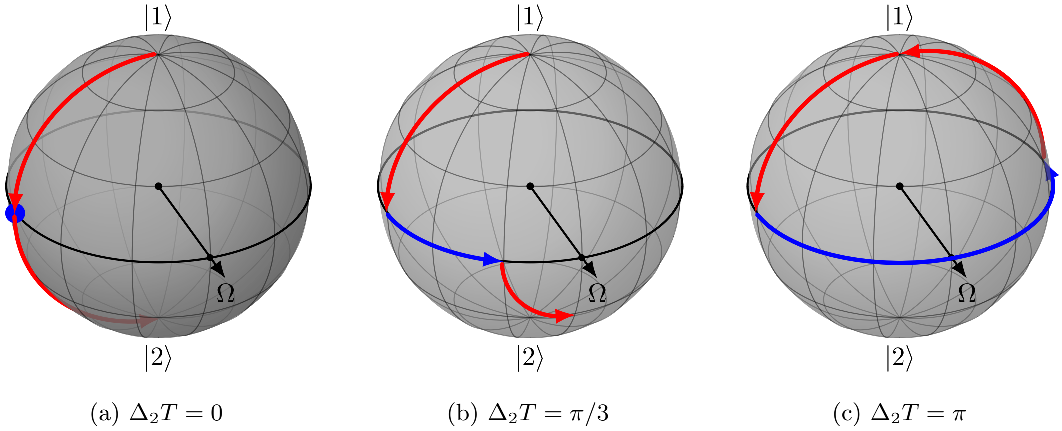

(a) $\Delta_2 T = 0$ (b) $\Delta_2 T = \pi/3$ (c) $\Delta_2 T = \pi$

Figure 3.4: Bloch diagrams showing the effect of a Ramsey sequence on an initial state $|1\rangle$, for different values of detuning, Δ_2. During the driving pulses (red), the state is rotated around $\vec{\Omega}$. During the Ramsey period (blue), the atoms may pick up an additional phase relative to the local oscillator. It is assumed that $\Delta_1 \ll \Omega$, so that $\vec{W} \approx \vec{\Omega}$.

sequence length, allowing more precise frequency measurements. More importantly, the reduced sensitivity to perturbations during the pulses (Δ_1) is advantageous for accuracy and long-term stability.

3.2 Clock stability

3.2.1 Clock comparisons

There is no single parameter that fully describes the performance of a clock. In order to evaluate a clock, it is necessary to compare it against other clocks. Assuming an ideal reference clock, we can monitor the difference between the clocks and see any deviations from the exact time. These deviations (noise) could have many different behaviours, which may or may not be important in a given application.

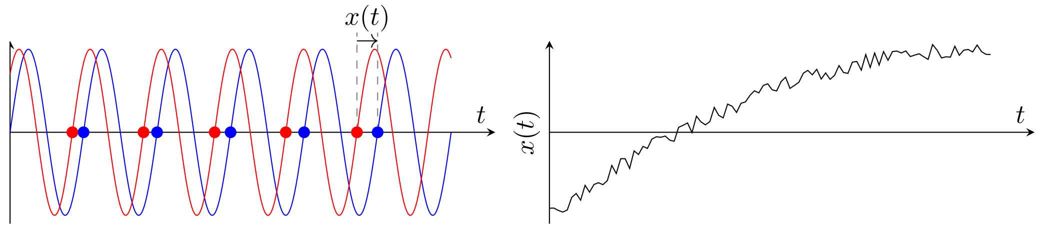

Figure 3.5: A clock comparison. The oscillator under test (blue) is compared to a reference oscillator (red). The time difference $x(t)$ can be measured and plotted over time.

For example, a constant frequency offset in a clock is often unimportant, as it can be calibrated out, except in the case of a primary standard clock where it would be extremely important. Some clocks may be highly stable over short timescales, maintaining their frequency with very low noise, but drift over the course of days or months. Other clocks have poor short-term performance, exhibiting high frequency jitter, but are extremely stable when averaged over long enough periods. To capture this timescale dependence, a range of statistics are used which describe performance as a function of either time or frequency. The noise may then be categorised, understood, and possibly mitigated according to the frequency/time-dependent behaviour.

The statistics describing clock performance are based on a general oscillator model, used to describe the LO behaviour in a well-defined manner. A clocks performance is usually entirely determined by its LO with the counter having a negligible contribution, so this is sufficient to describe the clock performance. The oscillator is modelled as a signal of a given frequency and amplitude

$$S(t) = (S_0 + \epsilon(t))\sin(2\pi\nu_0 t + \phi(t)) \ , \tag{3.10}$$

where S_0 and ν_0 are the nominal frequency and amplitude respectively, and $\epsilon(t)$ and $\phi(t)$ are time varying deviations from the expected behaviour. The 'ideal' time, t, is given by a reference clock. Normally the amplitude fluctuations, $\epsilon(t)$, are unimportant and the information of importance is contained in the phase fluctuations, $\phi(t)$. This is because the signal is usually processed in a manner that is relatively independent of the signal amplitude, e.g. using the zero-crossing times as reference times.

The phase fluctuations, $\phi(t)$, describes how far the clock has diverged from the ideal time, but it does so in a manner that depends on ν_0. Instead, the time-error function is used,

$$x(t) = \phi(t)/(2\pi\nu_0) \ , \tag{3.11}$$

which is the total difference of the 'clock time' from the reference time. Additionally, the oscillator has an instantaneous frequency,

$$\nu(t) = \nu_0 + \frac{1}{2\pi}\frac{d\phi}{dt} \ , \tag{3.12}$$

from which the fractional frequency error is calculated,

$$y(t) = \frac{\Delta\nu}{\nu_0} = \frac{\nu(t) - \nu_0}{\nu_0} = \frac{1}{2\pi\nu_0}\frac{d\phi}{dt} = \frac{dx}{dt} \ , \tag{3.13}$$

$$\implies x(t) = \int_0^t y(t')dt' + x(0) \ . \tag{3.14}$$

The functions $\phi(t), x(t), \nu(t), y(t)$ all describe the same behaviour, but have different interpretations and utilities.

In a real clock comparison, we will not have access to e.g. $x(t)$. Instead we will sample $x(t)$ over some finite period. This produces a series of N data points, $x_n(t_n)$. We usually assume that $t_n =$

$n\tau_0$ for integer n, i.e. that x is sampled uniformly with sampling period τ_0. Alternatively, we might have measured the clock frequency instead of phase to produce M fractional frequency samples, y_n. We can convert our data between these forms using the discretised versions of equations (3.13) and (3.14):

$$y_n = \frac{x_{n+1} - x_n}{\tau_0} \ , \tag{3.15}$$

$$x_n = x_0 + \tau_0 \sum_{i=0}^{n} y_i \ . \tag{3.16}$$

Note that there will be one more frequency sample than phase samples, $M = N + 1$

3.2.2 Allan variance

For many types of measurement, the most commonly used noise statistic is the standard variance, s^2. The standard variance quantifies the average amount of fluctuation around the 'true' value of the measurement.

$$s^2 = \left\langle (y(t) - y_0)^2 \right\rangle , \quad y_0 = \langle y(t) \rangle \ . \tag{3.17}$$

For clocks, such a measurement is usually inappropriate. Firstly, there is no notion of timescale involved with the standard variance. More importantly, for many types of noise the standard variance is not even well-defined. If a clock has any kind of long-term drift, or divergent (random-walk) noise, then the clock frequency will deviate further and further over time, leading to a divergent standard variance.

The Allan variance[58] is an attempt to rectify these problems, and is the most common statistic used to describe clock performance on a range of timescales. The Allan variance $\sigma_y(\tau)$ is a function of the timescale τ. To prevent divergence, the clock frequency fluctuations are taken relative to the clock frequency a finite period τ earlier, instead of the long-term average. One advantage of the Allan variance is that it has a relatively straightforward interpretation; the Allan variance is half the mean-square fluctuation of the average clock frequency between two consecutive periods of length τ. The half is introduced because the fluctuations of $y(t)$ and $y(t + \tau)$ both contribute independently to $\sigma_y(\tau)$.

Suppose that the (normalised) clock frequency, $y(t')$ is averaged over the two periods, $t < t' \leq t+\tau$ and $t + \tau < t' \leq t + 2\tau$ to give $\bar{y}_1$ and $\bar{y}_2$ respectively. The Allan variance is then given by

$\sigma_y^2(\tau) = \left\langle (\bar{y}_2 - \bar{y}_1)^2 \right\rangle /2$ where the $\langle\rangle$ indicates an average over all values of t.

$$\begin{aligned}
\sigma_y^2(\tau) &= \frac{1}{2}\left\langle (\bar{y}_2 - \bar{y}_1)^2 \right\rangle \\
&= \frac{1}{2}\left\langle \left(\frac{1}{\tau}\int_0^\tau y(t+\tau+t')dt' - \frac{1}{\tau}\int_0^\tau y(t+t')dt' \right)^2 \right\rangle \\
&= \frac{1}{2}\left\langle \left(\frac{1}{\tau}(x(t+2\tau) - x(t+\tau)) - \frac{1}{\tau}(x(t+\tau) - x(t)) \right)^2 \right\rangle \\
&= \frac{1}{2\tau^2}\left\langle (x(t+2\tau) - 2x(t+\tau) + x(\tau))^2 \right\rangle . \qquad (3.18)
\end{aligned}$$

The Allan variance can also be defined using a wavelet formalism.[59] In this case the oscillator frequency $y(t)$ is convolved with a wavelet $g(t)$ where

$$g(t) = \begin{cases} +1/\tau & ,\tau < t < 2\tau \\ -1/\tau & ,0 < t < \tau \\ 0 & ,\text{otherwise} \end{cases} , \qquad (3.19)$$

as is shown in Fig. 3.6(a). This convolution is equivalent to finding the value of $\bar{y}_2 - \bar{y}_1$ described earlier:

$$g(t) * y(t) = \bar{y}_2 - \bar{y}_1 \quad \Longrightarrow \quad \sigma_y^2(\tau) = \frac{1}{2}\left\langle (g(t') * y(t+t'))^2 \right\rangle . \qquad (3.20)$$

The Allan variance is then the mean-squared value of this convolution as the convolving waveform is shifted across the full width of the available $y(t)$. The wavelet formulation is useful as another picture for describing the statistic, and can extended to give other kinds of frequency statistic by using different wavelets $g(t)$.[60,61] The strength of the wavelet picture is its simplicity. While it may seem trivial here, the value may be clearer in Chapter 4 where it is used to describe the Theo1 statistic.

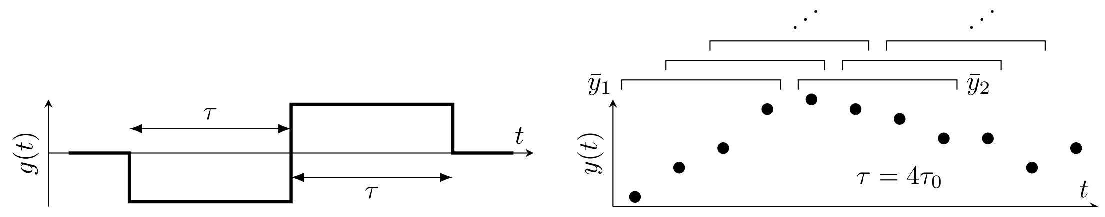

(a) (b)

Figure 3.6: The wavelet representation of the Allan variance (a). The discrete sums that form $\bar{y}_{1,2}$ for real calculations of the Allan variance (b). The regions of $\bar{y}_{1,2}$ are swept across the full dataset. A specific value of $\tau = 4\tau_0$ is shown, but the calculation would be repeated for different values of τ to give $\sigma_y(\tau)$.

When measuring a real clock, we will only be able to estimate the value of $\sigma_y(\tau)$ using our sampled measurements x_n or y_n. To do this we substitute $\tau = n\tau_0$ and replace the averages over t to averages over our entire set of samples.

$$\begin{aligned}\sigma_y^2(n\tau_0) &= \frac{1}{2n^2\tau_0^2}\left\langle (x_{i+2n} - 2x_{i+n} + x_i)^2 \right\rangle_i \\ &= \frac{1}{2n^2\tau_0^2(N-2n)} \sum_{i=1}^{N-2n} (x_{i+2n} - 2x_{i+n} + x_i)^2 \, , \end{aligned} \tag{3.21}$$

$$\sigma_y^2(n\tau_0) = \frac{1}{2n^2(M-2n+1)} \sum_{j=1}^{M-2n+1} \left(\sum_{i=j}^{j+n-1} y_{i+n} - y_i \right)^2 \, . \tag{3.22}$$

Fig. 3.6(b) shows how the Allan variance is estimated using this method.This method of estimating the Allan variance is sometimes referred to as the overlapping Allan variance, in contrast to earlier methods of which used only non-overlapped intervals in the estimation. That earlier method has been largely superseded by the overlapping Allan variance estimation.

Many types of noise follow a power-law relationship in which the noise amplitude is related to a power of the frequency or timescale, $S_y(f) \propto f^\alpha$, where $S_y(f)$ is the power spectral density (PSD) of the fractional frequency fluctuations, and α is the power-law exponent. Table 3.1 shows the most common types of power-law noise. One limitation of the Allan deviation is that it cannot distinguish between white phase noise and flicker phase noise. This is a limitation of the statistic, and is one reason that other statistics, such as the modified Allan deviation,[61] were developed. There may also be other non-ideal behaviours: linear frequency drifts or environmental dependence on temperature, pressure, magnetic field, or even gravitational environment. The Allan variance is normally plotted on a log-log plot, to summarise the clocks behaviour across many orders of magnitude (Fig. 3.7). The gradient of the plot will then reveal which types of noise are dominant at different timescales.

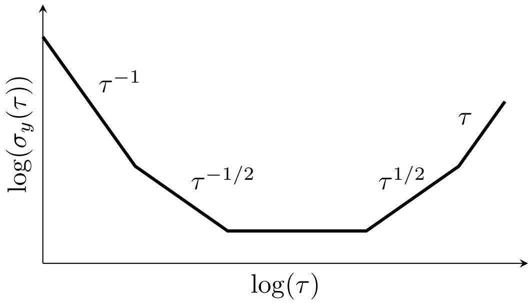

Figure 3.7: An idealised Allan deviation vs τ plot, showing different types of power-law noise. The type of power-law noise determines the gradient.

Table 3.1: Types of power-law noise. The PSD of frequency fluctuations, $S_y(f)$ is given for each noise type, alongside the Allan deviation and modified Allan deviation responses, $\sigma_y(\tau)$ and $\sigma_{y,\mathrm{mod}}(\tau)$.

Noise type	Abbreviation	$S_y(f)$	$\sigma_y(\tau)$	$\sigma_{y,\mathrm{mod}}(\tau)$
White Phase Noise†	WPM	$\propto f^2$	$\propto \tau^{-1}$	$\propto \tau^{-3/2}$
Flicker Phase Noise	FPM	$\propto f$	$\propto \tau^{-1}$	$\propto \tau^{-1}$
White Frequency Noise	WFM	$\propto 1$	$\propto \tau^{-1/2}$	$\propto \tau^{-1/2}$
Flicker Frequency Noise	FFM	$\propto f^{-1}$	$\propto 1$	$\propto 1$
Random Walk Frequency Noise	RWFM	$\propto f^{-2}$	$\propto \tau^{1/2}$	$\propto \tau^{1/2}$

† The Allan deviation cannot differentiate between WPM and FPM, and has a τ^{-1} response to both. Other statistics may be used to distinguish these noise types.

3.2.3 SNR and stability

Most atomic clocks will be dominated by white frequency noise at short times. This is because the atoms are interrogated periodically and do not maintain their phase between measurements. The noise in each measurement is generally uncorrelated so that each measurement is an independent frequency measurement with a random noise. Over time, more measurements are made and the frequency error decreases as $1/\sqrt{\tau}$ due to averaging.

We can derive the short-term stability by considering the transfer of measurement noise onto the clock output. In general, we will have some experimental signal, $S(\delta\nu)$, which depends on the detuning of our LO from the atomic resonance, $\delta\nu = \nu - \nu_0$. We use the observed signal to calculate the value of $\delta\nu$ and thereby apply a correction to the LO which drives $\delta\nu \to 0$. However we will also observe some noise on S, causing a random error in our calculation of $\delta\nu$ and a corresponding frequency error in our clock output because the LO correction is inaccurate.

We shall assume that the observed noise is white and normally distributed with standard deviation σ_S and that we can expand $S(\delta\nu)$ around the nominal detuning of $\delta\nu_0$:

$$S(\delta\nu) \approx S(\delta\nu_0) + (\delta\nu - \delta\nu_0) \left.\frac{dS}{d(\delta\nu)}\right|_{\delta\nu_0} . \tag{3.23}$$

We also assume that σ_S is sufficiently small, and $S(\delta\nu)$ sufficiently linear that this approximation causes negligible error. We can calculate that the observed error in S causes a normally distributed frequency error on the clock output with standard deviation of

$$\sigma_\nu = \frac{\sigma_\mathrm{S}}{\left| \left.\frac{dS}{d(\delta\nu)}\right|_{\delta\nu_0} \right|} . \tag{3.24}$$

In fractional frequency the clock will have a single-shot performance of $\sigma_y(\tau_0) = \sigma_\nu/\nu_0$, and as multiple frequency measurements are taken the performance will improve by a factor of $\sqrt{N_{\text{measure}}} = \sqrt{\tau/\tau_0}$ to give

$$\sigma_y(\tau) = \frac{\sigma_{\text{S}}}{\nu_0} \frac{1}{\left| \left. \frac{dS}{d(\delta\nu)} \right|_{\delta\nu_0} \right|} \sqrt{\frac{\tau_0}{\tau}} \ . \tag{3.25}$$

For a Ramsey clock, we know from equation (3.9) that we expect to see a sinusoidal signal, such that

$$S(\delta\nu) = A\cos(\Delta T) + B = A\cos(2\pi T\delta\nu) + B \ , \tag{3.26}$$

where parameters A and B are the signal amplitude and offset respectively, and depend on the experimental details. This gives

$$\frac{dS}{d(\delta\nu)} = -2\pi AT\sin(2\pi T\delta\nu) \ . \tag{3.27}$$

To maximise the gradient and therefore minimise clock error, Ramsey clocks usually operate on the side of the fringe where $\delta\nu_0 = \pm T/4$, giving

$$\left| \left. \frac{dS}{d(\delta\nu)} \right|_{\delta\nu_0} \right| = 2\pi AT \ . \tag{3.28}$$

Substituting this into equation (3.25) finally gives us the short-term stability,

$$\sigma_y(\tau) = \frac{2\sigma_{\text{S}}}{S_{\text{p-p}}} \frac{1}{2\pi\nu_0 T} \sqrt{\frac{\tau_0}{\tau}} = \frac{1}{\text{SNR}} \frac{1}{\pi Q} \sqrt{\frac{\tau_0}{\tau}} \ , \tag{3.29}$$

where $S_{\text{p-p}} = 2A$ is the peak-to-peak signal amplitude and $Q = 1/(2\nu_0 T)$ is the atomic quality factor. For other interrogation methods, such as a Rabi sequence, similar formulae will hold though they may differ by constant factors due to the exact lineshape observed. As we predicted, this is white frequency noise with $\sigma_y(\tau) \propto \tau^{-1/2}$.

3.2.4 Quantum projection noise

It will be possible to remove (or make negligible) many sources of noise such as detection or electronic noise, by suitable engineering of the clock system. However, whatever interrogation method is used, it must involve a measurement made upon the atomic system, a *quantum* system. The statistics of quantum measurements will then present a fundamental limit to the short-term stability.

Most atomic clocks can be modelled as a collection of N independent atoms, each acting as a 2-state quantum system. In the ideal case, all these atoms will be in the same quantum state, $|\psi\rangle$ at the end of the clock sequence. To detect a signal, the atoms must be projected onto some chosen basis set, $|1\rangle$ and $|2\rangle$. This projection is fundamentally random due to its quantum nature, and therefore there is an inherent noise in the measurement process, regardless of the physical measurement technique that is used. The probability of projecting each atom into $|2\rangle$ is $P_2 = |\langle 2|\psi\rangle|^2$. Because the atoms are independent, the total number of atoms which are projected into $|2\rangle$, N_2, will now be randomly drawn from a binomial distribution. This distribution will have a mean of $\langle N_2\rangle = NP_2$ and variance of $\sigma_{N_2}^2 = NP_2(1-P_2)$, and is usually well approximated by a normal distribution with the same values, because N is relatively large and P_2 close to 0.5.

In a Ramsey sequence, our signal will be $S = N_2 = NP_2$ with

$$P_2 = \frac{1}{2} + \frac{1}{2}\cos\left(2\pi T\delta\nu\right) \ , \tag{3.30}$$

from equation (3.9) and using $\Delta_2 = 2\pi\delta\nu$. From this we can calculate the observed noise:

$$\begin{aligned} \sigma_{\mathrm{S}}^2 = \sigma_{N_2}^2 &= N\left(\frac{1}{2} + \frac{1}{2}\cos\left(2\pi T\delta\nu\right)\right)\left(\frac{1}{2} - \frac{1}{2}\cos\left(2\pi T\delta\nu\right)\right) \\ &= \frac{N}{4}\sin^2(2\pi T\delta\nu) \ , \qquad (3.31) \\ \implies \sigma_{\mathrm{S}} &= \frac{\sqrt{N}}{2}\left|\sin(2\pi T\delta\nu)\right| \ . \qquad (3.32) \end{aligned}$$

We can also calculate the signal gradient,

$$\frac{dS}{d(\delta\nu)} = -N\pi T\sin(2\pi T\delta\nu) \ . \tag{3.33}$$

Finally, substituting into equation (3.25) gives the short-term stability of

$$\sigma_y(\tau) = \frac{1}{2\pi\nu_0 T}\frac{1}{\sqrt{N}}\sqrt{\frac{\tau_0}{\tau}} \ . \tag{3.34}$$

This is the quantum projection noise (QPN) limit for a Ramsey atomic clock, and shows that it is beneficial to have a large Ramsey time to ensure narrow clock fringes and a large number of atoms to increase SNR. It is interesting to note that the QPN limit does not require that measurements are made on the side of the fringe, where the fringe gradient is maximised. This is because the quantum noise is also maximised at this point. However, in practice other sources of noise (e.g. electronic noise) will add uniform noise at all points on the fringe, and so the side of the fringe is still a better frequency discriminant. It is also beneficial to operate around a symmetrical point of the fringe to minimise bias.

Chapter 4

Theo1

This chapter is based upon a published article.[62] That article was edited to fit the context of this thesis.

As discussed previously, in Section 3.2.2, there are many stability statistics used to describe the performance of a frequency source. One of the statistics used for this purpose is the 'theoretical variance #1' (Theo1 or Thêo1).[63] This chapter will firstly introduce the Theo1 statistic, with comparison to the Allan deviation, and then show an algorithm which improves the computational complexity of calculating Theo1 from $\mathcal{O}(N^3)$ to $\mathcal{O}(N^2)$, allowing faster calculation.

4.1 The Theo1 statistic

The Theo1 statistic is specialised to give more reliable estimates of stability at the longest averaging times, τ, achievable for a given dataset. Compared to the more commonly used Allan variance, Theo1 has increased confidence at long averaging times, and can be used to estimate stability up to 50% longer averaging times. Theo1 is also better able to identify which type of 'power-law' noise is present.[64,65] These properties have allowed Theo1 to be used for long-running experiments where datasets cannot easily be extended.[66] For a series of N time deviation points x_i, each separated by an interval τ_0, Theo1 can be defined[63] as

$$\text{Theo1}(\tau = 1.5k\tau_0, N) = \frac{T_k}{3(N-2k)(k\tau_0)^2} \ , \tag{4.1}$$

where $0 < k \leq (N-1)/2$, the averaging time is τ and

$$T_k = \sum_{i=0}^{N-2k-1} \sum_{\delta=0}^{k-1} \frac{1}{(k-\delta)} \left[(x_i - x_{i-\delta+k}) + (x_{i+2k} - x_{i+\delta+k}) \right]^2 \ . \tag{4.2}$$

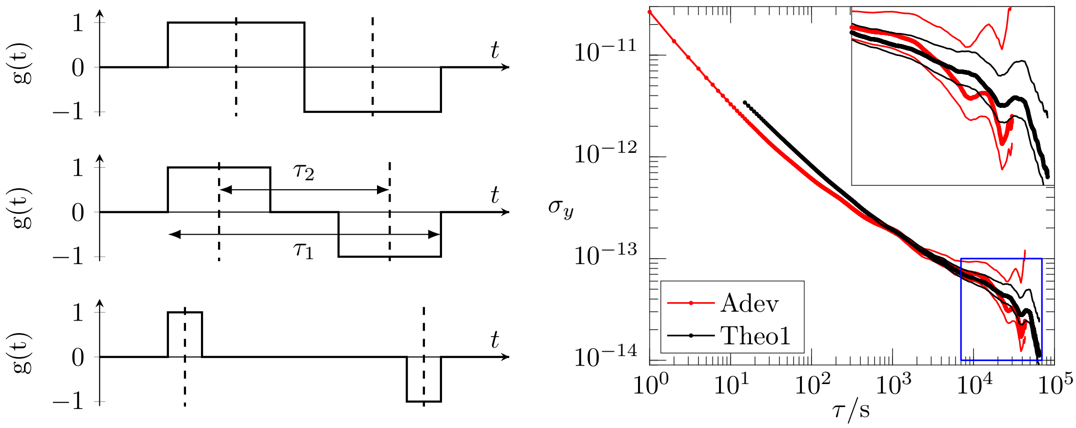

(a) (b)

Figure 4.1: A comparison of Theo1 and the Allan variance. (a) shows a selection of wavelets used in Theo1. The top wavelet is equivalent to that used for the Allan variance. For the other wavelets, a dead period is added in the centre, whilst keeping the same wavelet length, τ_1. This increases the effective sampling time, τ_2. (b) shows the Allan deviation and Theo1 (deviation) calculated for the same dataset,[67] with confidence intervals shown by thinner lines. Theo1 has significantly improved confidence at large τ (blue box, magnified in inset).

This definition can be broadly understood by considering the wavelet representation of stability variances, described earlier (Chapter 3). In the case of Theo1, a family of wavelets are used as illustrated in Fig. 4.1(a). The combination of these different wavelets requires the additional nested sum in equation (4.1). The use of wavelets with a larger effective averaging time, τ_2, allows Theo1 to provide results at a longer averaging period than the Allan deviation given the same dataset. Averaging across the larger family of wavelets provides additional averaging power to the statistic. An example of the difference between Allan deviation and Theo1 is shown in Fig. 4.1(b).

A naive implementation of this definition of Theo1 will have a complexity of $\mathcal{O}(N^3)$ because there are $\approx N/2$ values of k for which to calculate T_k, each taking $\mathcal{O}(N^2)$ operations due to the nested sum in equation (4.2) compared to equation (3.21). This can make computation prohibitively expensive for extremely large datasets or in applications requiring low latency, such as measuring the dynamic stability of an oscillator with a high data rate.[68]

In some cases it is not necessary to calculate Theo1 for every value of k, sometimes called an ‘all-τ’ calculation, and it may be sufficient to only use k equal to powers of two. However, the more sophisticated statistics TheoBr and TheoH,[69] which attempt to correct for bias in Theo1 relative to the Allan deviation, require the calculation of Theo1 for all k as a first step. There is a technique called ‘fast TheoBr’[70] which increases the speed of this calculation by averaging points within the initial dataset to reduce its size. However, for a fixed amount of averaging, the speed increase is only a constant factor and does not change the $\mathcal{O}(N^3)$ complexity. Similar statistics with $\mathcal{O}(N^3)$ complexity have been reduced to $\mathcal{O}(N^2)$ complexity by the use of an appropriate algorithm.[71]

4.2 Algorithm

One way to produce a faster algorithm for Theo1 is to find a recurrence relation between parts of the outer sum, which allows calculation of one term from the next without performing the full inner sum. This is made difficult by the term $1/(k-\delta)$ which forces a different coefficient before each terms as δ is incremented in the inner sum. However, the definition of T_k can be rearranged to move this awkward term outside the inner sum by swapping the order of the sums and using the substitution $v = k - \delta$, so that

$$T_k = \sum_{v=1}^{k} \frac{1}{v} A_{k,v} \ , \tag{4.3}$$

where $A_{k,v}$ is defined by

$$\begin{aligned} A_{k,v} = & \sum_{i=0}^{N-2k-1} (x_i - x_{i+v} + x_{i+2k} - x_{i+2k-v})^2 \qquad (4.4) \\ = & \sum_{i=0}^{N-2k-1} (x_i^2 + x_{i+v}^2 + x_{i+2k}^2 + x_{i+2k-v}^2 \\ & \quad + 2x_i x_{i+2k} + 2x_{i+v} x_{i+2k-v} - 2x_i x_{i+v} \\ & \quad - 2x_i x_{i+2k-v} - 2x_{i+v} x_{i+2k} - 2x_{i+2k} x_{i+2k-v}) \, . \qquad (4.5) \end{aligned}$$

Some of the expanded terms in equation (4.5) have similar forms, and can be expressed in terms of new summations $C^{(n)}$, defined as

$$C_j^{(1)} = \sum_{i=0}^{j} x_i^2 \ , \tag{4.6}$$

$$C_j^{(2)} = \sum_{i=0}^{N-j-1} x_i x_{i+j} \ , \tag{4.7}$$

$$C_{k,j}^{(3)} = \sum_{i=k}^{N-k-1} x_{i-j} x_{i+j} \ , \tag{4.8}$$

$$C_{k,j}^{(4)} = \sum_{i=0}^{N-2k-1} x_i x_{i+j} + x_{i+2k} x_{i+2k-j} \, . \tag{4.9}$$

It can then be shown by substitution that

$$\begin{aligned} A_{k,v} = & C_{N-2k-1}^{(1)} + (C_{N-2k-1+v}^{(1)} - C_{v-1}^{(1)}) + (C_{N-1}^{(1)} - C_{2k-1}^{(1)}) \\ & + (C_{N-v-1}^{(1)} - C_{2k-v-1}^{(1)}) + 2(C_{2k}^{(2)} + C_{k,k-v}^{(3)} - C_{k,v}^{(4)} - C_{k,2k-v}^{(4)}) \, . \end{aligned} \tag{4.10}$$

The calculation of T_k from the $C^{(n)}$ can be completed in $\mathcal{O}(N^2)$, so if the $C^{(n)}$ could all be calculated in $\mathcal{O}(N^2)$ then this would reduce the overall complexity of Theo1 to $\mathcal{O}(N^2)$. For $C^{(1,2)}$

the definition is already $\leq \mathcal{O}(N^2)$, but it can also be achieved for $C^{(3,4)}$ by using a recurrence relation between consecutive terms to avoid the full sum in equations (4.8) and (4.9):

$$\begin{aligned} C^{(3)}_{k,j} &= C^{(3)}_{k-1,j} - x_{k-1-j}x_{k-1+j} - x_{N-k-j}x_{N-k+j} \quad ,j < k \ , && (4.11) \\ C^{(4)}_{k,j} &= C^{(4)}_{k-1,j} - x_{2k-2-j}x_{2k-2} - x_{2k-1-j}x_{2k-1} \\ &\quad - x_{N-2k}x_{N-2k+j} - x_{N-2k+1}x_{N-2k+1-j} \quad ,j < 2k-1 \ . && (4.12) \end{aligned}$$

This allows almost all values of $C^{(3,4)}$ to be calculated in a recursive manner, the remaining values are

$$\begin{aligned} C^{(3)}_{k,k} &= C^{(2)}_{2k} \ , && (4.13) \\ C^{(4)}_{k,2k-1} &= 2C^{(2)}_{2k-1} - x_0 x_{2k-1} - x_{N-2k}x_{N-1} \ , && (4.14) \\ C^{(4)}_{k,2k} &= 2C^{(2)}_{2k} \ , && (4.15) \end{aligned}$$

and so $C^{(3,4)}$ can be calculated in the required $\mathcal{O}(N^2)$. Because the recurrence relations are for an incremented k value, the technique can only be used when calculating T_k for all values of k.

In order to calculate T_k it is sufficient to calculate

$$\begin{aligned} &C^{(1)}_{j}, \quad 0 \leq j \leq N \ , && (4.16) \\ &C^{(2)}_{j}, \quad j = 2k, 2k-1 \ , && (4.17) \\ &C^{(3)}_{k,j}, \quad 0 \leq j \leq k \ , && (4.18) \\ &C^{(4)}_{k,j}, \quad 0 \leq j \leq 2k \ , && (4.19) \end{aligned}$$

so only these values need to be held in memory. The recursion relations equations (4.11) and (4.12) can then be used to update $C^{(3,4)}_{k,j}$ to $C^{(3,4)}_{k+1,j}$ in place, allowing calculation of T_{k+1}. The memory requirement is only $\mathcal{O}(N)$. Specifically, it requires memory for $4N$ double precision values: N values for each of the input array x and $C^{(1,4)}$, and $N/2$ values for each of $C^{(3)}$ and the output array. The naive algorithm requires storage of $3N/2$ values so this is a significant increase but is still only 32 MB for a dataset with $N = 10^6$.

In order to calculate Theo1, the algorithm can proceed as follows:

1. Calculate $C^{(1)}$ using equation (4.6).
2. For each value of k from 0 to $\lfloor (N-1)/2 \rfloor$:
    (a) Calculate required values of $C^{(2)}$ using equation (4.7).
    (b) Add new values to $C^{(3,4)}$ using equations (4.13) to (4.15).
    (c) Update $C^{(3,4)}$ using equations (4.11) and (4.12).

(d) Calculate $A_{k,v}$ from the $C^{(n)}$ using equation (4.10).

(e) Calculate T_k from $A_{k,v}$ using equation (4.3).

An example implementation of the algorithm in C++ can be found in Appendix A.

4.3 Accuracy

Whilst the new algorithm for Theo1 is faster than the naive approach, it has more opportunities for floating point errors to accumulate. Equation (4.10) shows that terms of similar magnitude are subtracted from each other, allowing catastrophic cancellation to occur and leading to a loss of precision. The size of each term in equation (4.10) is $\leq \sum x^2$ and the size of the total is $\sim T_k$, so the fractional error might be expected to scale as $\sim \langle x^2 \rangle / T_k$. Theo1 is insensitive to any offset or linear change in x, so these components can be removed in order to reduce $\langle x^2 \rangle$ without ill-effect. This prevents a significant drop in precision that could be caused by a constant frequency or phase offset. This change alone is sufficient to prevent appreciable errors in most practical situations. However, in some cases with large datasets and where the long-term clock stability is dominated by frequency drift the errors could grow large enough to be significant.

In order to fully mitigate the floating point error it is necessary to use additional bits of precision. Whilst 128-bit floating point types are available in some environments, they are not generally supported in hardware and are therefore very slow to use. In contrast, many 64-bit CPUs have hardware support for multiplying two 64-bit integers into a 128-bit integer. By re-scaling the data and converting it to a 64-bit integer representation, this CPU instruction can be used to calculate Theo1 quickly without floating point errors. Terms $\propto x$ are stored as 64-bit integers, but terms $\propto x^2$ are stored as 128-bit integers. In C or C++ the GCC __int128 datatype may be used with multiplications implemented as

```
int64_t x1,x2;
__int128 result = (__int128) x1 * x2;
```

The use of a larger data types does cause a speed reduction of approximately 70%, possibly due to additional memory overhead. The conversion between datatypes is $\mathcal{O}(N)$ and takes negligible time in most cases.

The speed and fractional floating point error for the different methods are shown in Figs. 4.2 and 4.3 respectively. The fractional error was measured by comparing the value of Theo1 as calculated with the new algorithm to that calculated with the naive algorithm. Due to the slow speed of the naive algorithm, only points spaced at powers of two were compared, and the maximum of

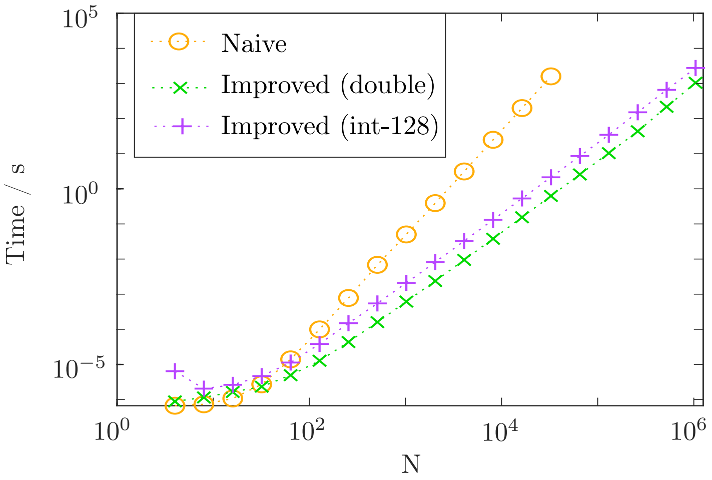

Figure 4.2: The time taken by different methods of calculating Theo1 on a desktop PC with an Intel i7-6700k CPU. The naive algorithm scales as $\mathcal{O}(N^3)$, but the improved algorithm scales as $\mathcal{O}(N^2)$. Using the int-128 datatype causes a reduction in speed, but the improved scaling remains. © 2020 IEEE

these errors was taken. Fig. 4.3 should be taken as indicative only as the details vary significantly depending on the noise type of the simulated data, although the 'int-128' method had negligible error in all cases tested. To exaggerate the errors seen, a white frequency noise with added linear frequency drift was simulated. The linear drift was chosen such that the frequency stability at the longest and shortest averaging factors was approximately equal. A linear frequency drift is particularly difficult for the simpler error reduction method (removing any linear component to x) to deal with, as the dominant x component is quadratic. Despite this, 1 to 2 orders of magnitude improvement was seen. For virtually all practical situations, this will be sufficient. However, if $N > 10^5$ and the dataset is know to be *significantly* dominated by low frequency noise, it may be worth spot checking some individual points against the naive algorithm. The int-128 algorithm could then be used if necessary.

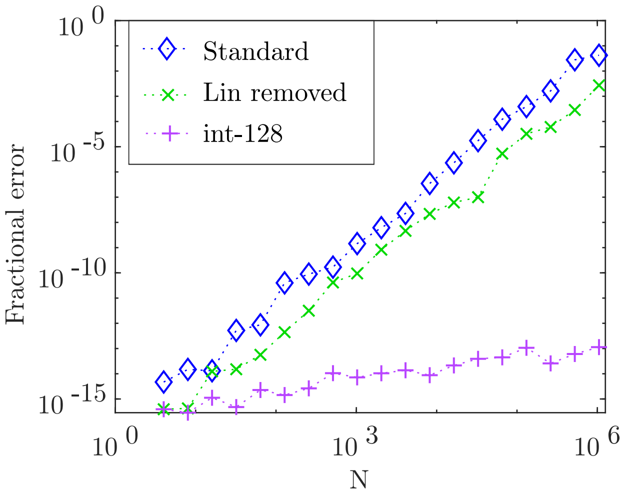

Figure 4.3: The maximum fractional error introduced by different methods of calculating Theo1, measured against the naive method. For the standard double precision method, errors grow rapidly with the dataset size. Removing any linear component helps reduce the shift at negligible computational cost. Using the int-128 datatype makes the errors negligible. Results depend strongly on the type of noise simulated, for this test white frequency noise with a linear frequency drift was used. © 2020 IEEE

Chapter 5

CPT clocks

Coherent population trapping is a feature of three-level quantum systems (or larger). When the system is excited by two separate fields, there may be a special state for which these two excitations interfere destructively, leaving no net excitation. This state is known as a ‘dark state’: because the atoms are not excited they cannot emit any fluorescence photons and appear dark. Because the destructive interference is phase-sensitive, this dark state is phase-coherent with the exciting fields and that phase-sensitivity can be used in an atomic clock.

5.1 Λ-system CPT

In Λ-system CPT, there are two stable ground states and an unstable excited state which couples to them both. In this system a dark state is formed as a coherent superposition of the two ground states.

The Λ-system can be modelled by extending our previous models of a two-level system. We now use two exciting fields on the $|1\rangle \rightarrow |3\rangle$ and $|2\rangle \rightarrow |3\rangle$ transistions with Rabi frequencies of $\Omega_{1,2}$ and detunings of $\Delta_{1,2}$.

$$\Delta_1 = \omega_1 - (E_3 - E_1)/\hbar \ , \tag{5.1}$$

$$\Delta_2 = \omega_2 - (E_3 - E_2)/\hbar \ , \tag{5.2}$$

However the CPT effect can also be viewed as a two-photon interaction between states $|1\rangle$ and $|2\rangle$, so it is useful to refer to the one and two-photon detunings of the system, $\Delta_{1\mathrm{ph}} = (\Delta_1 + \Delta_2)/2$ and $\Delta_{2\mathrm{ph}} = \Delta_1 - \Delta_2$, as illustrated in Fig. 5.1.

Using the results from Chapter 2, we can express the Hamiltonian of this system (in a dressed state picture) as

$$H = \frac{\hbar}{2}\begin{pmatrix} \Delta_{\mathrm{2ph}} & 0 & \Omega_1 \\ 0 & -\Delta_{\mathrm{2ph}} & \Omega_2 \\ \Omega_1^* & \Omega_2^* & -2\Delta_{\mathrm{1ph}} \end{pmatrix} , \tag{5.3}$$

and we can then observe that if the (dressed) state amplitudes are given by c_i, then the excited state amplitude will vary according to the Schrödinger equation as

$$i\dot{c}_3 = \Omega_1^* c_1 + \Omega_2^* c_2 - 2\Delta_{\mathrm{1ph}} c_3 \ . \tag{5.4}$$

A dark state requires zero excited population and therfore must have $c_3 = 0, \dot{c}_3 = 0$ which implies

$$\Omega_1^* c_1 + \Omega_2^* c_2 = 0 \ , \tag{5.5}$$

$$\implies |\psi\rangle = |\mathrm{Dark}\rangle = \frac{\Omega_2^*}{\Omega_{\mathrm{t}}}|1\rangle - \frac{\Omega_1^*}{\Omega_{\mathrm{t}}}|2\rangle \ , \tag{5.6}$$

where $\Omega_t = \sqrt{|\Omega_1|^2 + |\Omega_2|^2}$.

The dark state, $|\mathrm{Dark}\rangle$, is a superposition of $|1\rangle$ and $|2\rangle$. It is only an eigenstate in the case that $\Delta_{\mathrm{2ph}} = 0$ as then the dressed states $|1\rangle$ and $|2\rangle$ have the same energy, and $\dot{c}_{1,2} = 0$. If $\Delta_{\mathrm{2ph}} \neq 0$ then the dark state is unstable ($\dot{c}_{1,2} \neq 0$) and will evolve into a bright state. In this case, all of the system eigenstates will have some excited population.

If we extend our model to include spontaneous decay, then the atomic population will be pumped into the dark state provided that $\Delta_{\mathrm{2ph}} \ll \Omega_{1,2}, \frac{1}{\Gamma}$. The dark state is still slightly unstable for $\Delta_{\mathrm{2ph}} \neq 0$ and will evolve into a bright state in proportion to Δ_{2ph}. However, that bright state is then rapidly excited (proportional to $\Omega_{1,2}$) and decays (proportional to $\frac{1}{\Gamma}$). The decay has an even chance of entering the dark or bright states, so the steady state has the vast majority of atoms collected in the dark state.

It is worth noting that there is another effect in the CPT system that is not modelled above. We assumed that each laser couples only to it's own near-resonant transition, without affecting the

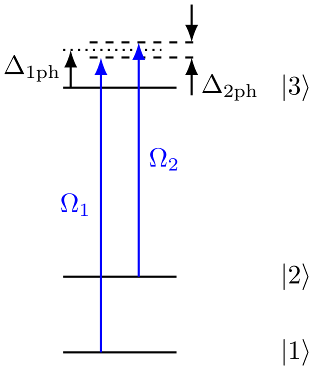

Figure 5.1: A Λ CPT system.

other transition in any way. This is inaccurate because both lasers can actually interact with both transitions, but they will be near-resonance for ‘their transition’ and detuned by $\approx \omega_1 - \omega_2$ for the other. This off-resonance excitation of another transition causes an AC Stark shift, effectively changing the energies, $E_{1,2,3}$. In a real atom, other transitions may also exist that will be off-resonantly excited causing additional shifts.

5.2 Ramsey CPT clock

In order to make a clock measurement, we perform a similar sequence to the Ramsey sequence discussed in Chapter 3. Two pulses of the CPT light fields are shone on the atoms with a Ramsey time in between where the atoms can evolve freely. This is an unusual Ramsey scheme as the two pulses are not performing $\pi/2$ pulses on the atoms. In fact the pulses are not any kind of unitary state operation, but instead optically pump the atoms. The first pulse is long enough to reach a steady state, where all atoms are in the dark state. The atoms do not need to be in any particular state before the first pulse; it also acts as a state preparation. During the Ramsey time, the exciting fields are removed, so we are left with the effective Hamiltonian of

$$H = \frac{\hbar}{2} \begin{pmatrix} \Delta_{2\text{ph}} & 0 & 0 \\ 0 & -\Delta_{2\text{ph}} & 0 \\ 0 & 0 & -2\Delta_{1\text{ph}} \end{pmatrix} , \tag{5.7}$$

and the $|1\rangle$, $|2\rangle$ states will pick up a relative phase during this time. To produce the most sensitive initial state, we would like $\Omega_1 = \Omega_2$ so that $|\text{Dark}\rangle = (|1\rangle - |2\rangle)/\sqrt{2}$. The evolution of the initial dark state throughout the Ramsey time will then be given by

$$\begin{aligned} |\psi\rangle &= \frac{1}{\sqrt{2}}(e^{i\Delta_{2\text{ph}}t}|1\rangle - e^{-i\Delta_{2\text{ph}}t}|2\rangle) , \\ &= \cos(\Delta_{2\text{ph}}t)\frac{1}{\sqrt{2}}(|1\rangle - |2\rangle) + i\sin(\Delta_{2\text{ph}}t)\frac{1}{\sqrt{2}}(|1\rangle + |2\rangle) \\ &= \cos(\Delta_{2\text{ph}}t)|\text{Dark}\rangle + i\sin(\Delta_{2\text{ph}}t)|\text{Bright}\rangle , \end{aligned} \tag{5.8}$$

so if we can detect the population of $|\text{Dark}\rangle$ we will see interference fringes. The second pulse of exciting fields will, in the same manner as the first pulse, optically pump the atoms back into the dark state. Atoms already in the dark state are unaffected, but atoms in the bright state will undergo Rabi oscillations to the excited state until they decay to the dark state. This decay process requires the spontaneous emission of at least one photon, which can be detected. The amount of spontaneous emission forms a signal proportional to the number of atoms in the bright state, and so we will have

$$S \propto |\langle\psi|\text{Bright}\rangle|^2 = \sin^2(\Delta_{2\text{ph}}t) . \tag{5.9}$$

These fringes are sensitive to $\Delta_{2\mathrm{ph}}$ and so can be used to lock the frequency difference between the exciting fields with $\Delta_{2\mathrm{ph}} = 0$. This frequency difference can be referenced to a local oscillator, forming the basis of a Ramsey-CPT clock.

5.3 Polarisation schemes in ^{87}Rb

While the level structure of Fig. 5.1 is sufficient to illustrate the mechanisms of CPT and Ramsey-CPT, real atoms have a more complex structure. In our clock we use atoms of ^{87}Rb, in which states $|1\rangle, |2\rangle$ correspond to the two stable ground states of the atom: $5^2\mathrm{S}_{1/2}$ states with $F = 1, 2$ respectively. State $|3\rangle$ can be any state of the D_1 or D_2 manifolds, though it must have $F' = 1$ or $F' = 2$ so that it couples to both ground states. All excited states decay with linewidth $\Gamma = 2\pi \times 6\,\mathrm{MHz}$. The neighbouring hyperfine levels will be detuned by $\sim 100\,\mathrm{MHz}$, and all others are THz away, so none of these should strongly affect the CPT-Ramsey mechanism though they may cause perturbations. However, each of the levels $|1\rangle, |2\rangle, |3\rangle$ are actually degenerate due to their non-zero total angular momentum, allowing for Zeeman splitting between different magnetic sub-levels. Fig. 5.2 shows the level structure we used in our experiment. A small magnetic bias field will be applied during experiments, so that a quantisation axis is defined and the magnetic degeneracy is broken. The magnetic sub-levels will be separated by less than the transition linewidth, so it is unavoidable that multiple sub-level transitions will be excited.

In this more complicated system, it is possible to form many different CPT systems, sometimes simultaneously. The selected light polarisation will determine which systems are driven, and so

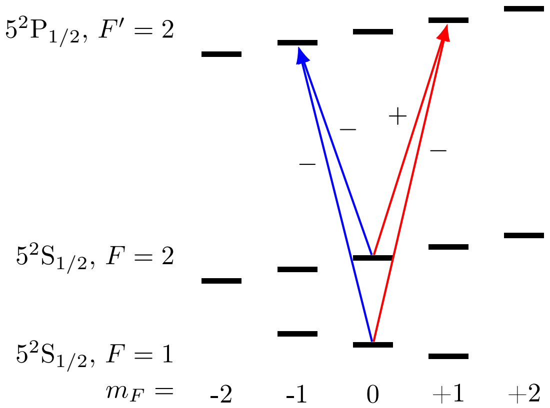

Figure 5.2: The relevant states for the CPT scheme used in our experiment. The desired dark state will use the shown transitions, which form two Λ systems. The red system uses only σ^+ transitions, and the blue system uses only σ^- transitions. The Clebsch-Gordan coefficient of the transitions are equal in magnitude but have varied signs, as annotated.

which dark state(s) are formed. In general, we would like to form a dark state from the $m_F = 0$ states as these are magnetically insensitive to first order (making the eventual clock system less magnetically sensitive). This means we cannot use π-polarised light, because we would have to use a 'forbidden' transition with $\Delta F = 0$ and $m_F = 0 \rightarrow 0$. We can use pure σ^+ or pure σ^- light but this will also pump the atoms into a stretched state ($F = 2, m_F = \pm 2$). A combination of different σ^+ and σ^- light is used to prevent this and concentrate atoms in the desired $m_F = 0$ dark state. This necessarily produces a 'double-Λ system where the dark state is excited by two distinct Λ systems, one coupling to an excited level with $m_F = +1$ via σ^+ light and the other to a level with $m_F = -1$ via σ^- light. The dark states for these two 'separate' systems must be identical in order to form a common dark state for the overall system. This is somewhat complicated because the signs of the transition Clebsch-Gordan coefficients (the relative weighting of the dipole matrix element that is dependent on the angular momentum projections) are different between the $\sigma^\pm$ systems.

The electric field at the atoms position can be described as a combination of frequency and polarisation components,

$$\overrightarrow{E} = E_1^+ e^{i\omega_1 t} \overrightarrow{\sigma^+} + E_2^+ e^{i\omega_2 t} \overrightarrow{\sigma^+} + E_1^- e^{i\omega_1 t} \overrightarrow{\sigma^-} + E_2^- e^{i\omega_2 t} \overrightarrow{\sigma^-} \ . \tag{5.10}$$

where $\overrightarrow{\sigma^\pm}$ are the vector representation of $\sigma^\pm$ polarisations. Equation (5.6) shows that the dark state for a Λ system depends only on the ratio Ω_1/Ω_2. For equal dark states in the two systems, we need this ratio to be consistent, so

$$\frac{\Omega_1^+}{\Omega_2^+} = \frac{\Omega_1^-}{\Omega_2^-} \ , \tag{5.11}$$

and this can be expressed in terms of the electric field using the Clebsch-Gordan coefficients as

$$\frac{E_1^+}{E_2^+} = -\frac{E_1^-}{E_2^-} \ . \tag{5.12}$$

Dark states between e.g. $|F = 1, m_F = +1\rangle$ and $|F = 2, m_F = +1\rangle$ will have a substantially different resonance frequency to our desired dark state because they have a first order magnetic shift of 1.4 MHz G^{-1}. (A bias field of 100 mG will separate the dark states by 140 kHz, clearly separated when the excitation has a 100 µs timescale as in Fig. 8.3.) These dark states will not contribute strongly as they will be significantly detuned when exciting the desired dark state. However dark states between e.g. $|F = 1, m_F = -1\rangle$ and $|F = 2, m_F = +1\rangle$ have only a small first order magnetic shift of magnitude 1.393 kHz G^{-1}, which is due to the non-zero nuclear g-factor for ^{87}Rb.[19,72,73] These dark states will have overlapping resonance with the desired dark state, causing interference between the three states and additional magnetic field dependence for the clock. However, if we use an excited state with $F' = 2$ then the $|F = 2, m_F = 1\rangle \rightarrow |F' = 2, m_F = 2\rangle$ and $|F = 1, m_F = -1\rangle \rightarrow |F' = 2, m_F = -2\rangle$ transitions prevent this from being a true dark state. We could use the $F' = 2$ state of either D_1 or D_2 manifolds. The D_1 line is simpler, with fewer levels and larger energy spacing between levels. An existing CPT clock was using the D_1 line, and we borrowed that laser for our experiment. Therefore we use the $5^2\mathrm{P}_{1/2}$, $F' = 2$ level as our excited level. If this technique were used long term, it would potentially be better to use the D_2 transition which could be generated from the same laser diode used for the MOT.

The combination of $\sigma^{\pm}$ light can be supplied in a number of ways and push-pull optical pumping (PPOP), Lin-$\|$-Lin, Lin-$\perp$-Lin and σ^+-σ^- are examples of schemes which have been used before. Lin-$\|$-Lin is a quite simple method experimentally, requiring only linear-polarised light at the two required frequencies.[74,75] The linear light contains both $\sigma^{\pm}$ components, and so will drive all the transitions needed. However, the Clebsch-Gordan sign change in not accounted for, and so there is no $m_F = 0$ dark state formed for this technique! Instead, this polarisation scheme only forms dark states between the $|F = 1, m_F = \pm 1\rangle$ and $|F = 2, m_F = \pm 1\rangle$, which are magnetically sensitive as described above.[48] The other polarisation schemes address this problem in different ways. Lin-$\perp$-Lin supplies the two beam frequencies in perpendicular linear polarisations, causing a sign difference when projected into the $\sigma^{\pm}$ components. This can be problematic because it requires separated beam paths for the two frequency components, giving the potential for additional noise in their frequency difference. PPOP uses a Michaelson interferometer to separate the $\sigma^{\pm}$ paths and adds path length to one arm until it accrues an additional π phase shift between the two frequencies.[76] The $\sigma^+ - \sigma^-$ polarisation scheme supplies only the σ^+ component in a single beam, but that beam is retroreflected through a quater-waveplate to provide the σ^- component.[77,78] The position of the retroreflecting mirror is adjusted until the extra path length produces a π phase shift between the two frequency components

We selected a $\sigma^+ - \sigma^-$ scheme for this experiment. The counterpropagating beams are beneficial because they come with an inherent Doppler-insensitivity (provided the intensity is balanced) and prevent the beam from exerting a net force on the atoms. Despite being a two-photon process, there can still be a residual Doppler effect because the two photons are at different frequencies (the Doppler effect is relative to the *microwave* wavelength). The Doppler-insensitivity can be achieved for other polarisation schemes by retroreflecting them as well, but then it is questionable why the extra complexity of that polarisation scheme is necessary. Additionally, counterpropagating components in the same polarisation state will interfere to form standing waves, making a detailed analysis of the system much more complicated, especially as atoms move across these waves due to finite velocity/temperature.

In any of these schemes it is advantageous to use counterpropagating beams which will mitigate any Doppler effect due to the atoms motion during the Ramsey time. Using counterpropagating beams also prevents the beams from exerting a net force upon the atoms, but can bring an additional complexity due to interference effects between the counterpropagating beams. The optical interference take place on a wavelength-scale and can usually be averaged over. However, for two-frequency beams there is also interference of the microwave signal formed by this frequency difference and this has a longer wavelength of 44 mm. The retroreflecting mirror must be placed the correct distance from the atoms to produce constructive microwave interference at the atoms position.

Part II

Experimental Work

Chapter 6

Cavity clock design

This chapter will detail the design of the microwave fountain clock that was produced during this PhD. The purpose and requirements of the project are described, and then an experimental sequence to achieve those requirements is specified. Details of the experimental apparatus required to fulfill these needs are given.

6.1 Experimental Aims

Cold-atom fountain clocks based on microwave cavities are the highest-stability microwave clocks that exist, as recognised by their realisation of the SI second. Our group aims to build a clock that uses the same fundamental design principles as these fountain clocks, but miniaturises the experimental apparatus as much as possible, targeting the few litre scale. In principle, most lasers, optical components and electronics are miniaturisable and can be integrated to a high degree, being a matter of suitable engineering.[14,79–81] Therefore, we first aim to reduce the volume of those components which are critical to the physics of the clock, specifically the components which directly interact with the atoms as they go through a clock sequence. For a microwave fountain clock, this is a variety of laser fields and the microwave cavity.

The atoms in a microwave fountain clock experience several experimental stages. These are trapping and cooling, launching, state selection, microwave interactions, and detection. In traditional fountain clocks these stages take place as the atoms move through an extended volume.[82] In order to miniaturise the clock, it is required that all of these stages take place within only a few mL. Microwave cavities usually have a size similar to the microwave wavelength ($\approx 44\,\text{mm}$ for ^{87}Rb).[83,84] Trapping a MOT of significant size typically requires laser beams of order 25 mm diameter. Achieving a MOT inside the microwave cavity will help minimise the total system volume.

The aim of the cavity clock experiment was to produce a cold-atom microwave clock, using a grating MOT inside a miniaturised microwave cavity produced by the University of Neuchâtel, Switzerland. All the clock interactions should occur within the cavity. The clock was to be a proof-of-principle, showing that this design was possible and could, in the future, be developed into a compact clock with state-of-the-art performance (a few 10^{-13} stability at 1 s and accuracy below 10^{-14}).

Because the experiment was to be a proof-of-principle, several design choices were made to prioritise experimental flexibility over performance. A large vacuum system would be used around the microwave cavity, with no magnetic shielding. Little attempt was made to achieve long-term stability. The verification of new techniques was the first priority, with short-term stability the secondary goal. Accordingly, efforts were made to allow a long Ramsey time, as this directly influences the achievable short-term stability. For this reason, a grating MOT launch method was proposed (Chapter 7) to extend the Ramsey time from around 10 ms to 100 ms and a design that would allow testing this method was chosen.

6.2 Experiment Overview

The experimental sequence is based on those of existing atomic fountain clocks, and consists of four stages:

1. MOT preparation - Trapping and cooling the atoms and (optionally) launching them
2. State selection - Preparing the atoms within a particular magnetic sub-level
3. MW pulses - Two Ramsey MW pulses separated by a Ramsey time
4. Detection - Detection of the population in each state

One of the major design constraints was the MW cavity supplied by the University of Neuchâtel, with whom we are collaborating. They already have the theoretical and technical expertise required to manufacture compact microwave cavities, which they have used for vapour-cell Rb clocks.[85–88] Their previous cavity designs have been made to accommodate a 25 mm vapour cell, so they are also able to accomodate the 25 mm laser beam required for our GMOTs. The MOT grating must be attached to the cavity base, opposite the large diameter entrace hole. These microwave cavities can be 3D-printed in aluminium which substantially decreases the cost and difficulty of manufacture compared to traditional machining. As aluminium is suitable for UHV, it is possible that future designs could integrate the cavity within a 3D-printed vacuum chamber.[89]

The microwave cavity has a loop-gap structure where the resonance frequency is set mainly by the

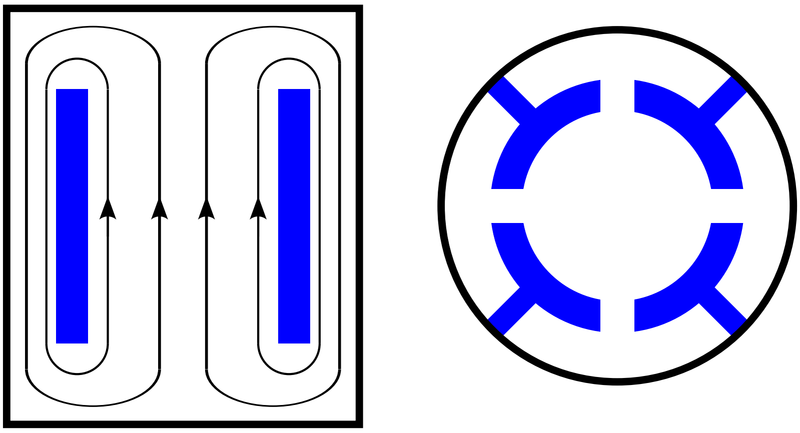

Figure 6.1: A schematic of the loop-gap cavity viewed from two axes. The electrodes (blue) are contained within a shield (black). Shield and electrodes are in electrical contact. The spacing between electrodes strongly affects the cavity resonance. The approximate magnetic field mode shape is shown by field lines. More details can be found in [85].

structure of electrodes that are present on the inner face of the cavity.[85,90] The cavity supports a TE_{011}-like mode that produces a magnetic field along the cylindrical axis of the cavity. The outer body of the cavity is not critical to the microwave resonance, and small holes can be created in this portion of the cavity. Beams can enter the cavity at these points, although they must fit between the cavity electrodes. Lasers for state selection and detection will require specific polarisation states, so they must use holes in the cavity side. An axial beam would be incident on the grating, leading to multiple beams with different polarisation states. The initial vision for the MOT region is shown in Fig. 6.2. I now consider each stage of the experiment in turn.

6.2.1 MOT preparation

This stage prepares the atoms into a known spatial distribution with restricted volume and velocity. The main purpose of this is to confine the atoms in phase-space such that the largest number of atoms may be detected after the Ramsey time has passed. For a small initial MOT after a sufficiently long Ramsey time the number of atoms which can be detected, N_{det}, using a beam of radius R, is given by

$$N_{\mathrm{det}} = N_0 \left(1 - \exp\left(\frac{-mR^2}{2k_B T t^2}\right)\right) \approx N_0 \frac{mR^2}{2k_B T t^2} \,, \tag{6.1}$$

where t is the atoms time of flight, N_0 is the initial number of trapped atoms, T is their temperature, m their mass and the initial MOT size has been neglected. The approximation applies when the MOT has expanded significantly beyond the size of the detecting beam (in the opposite case, we have simply $N_{\mathrm{det}} = N_0$). Clearly, it is advantageous to have a large initial number of atoms, at a very low temperature, and the ratio N_0/T is a reasonable metric for optimising the MOT.

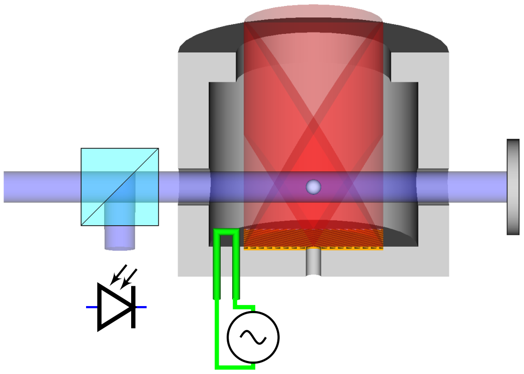

Figure 6.2: The cavity clock design. Trapping beams (red) enter from the top of the microwave cavity, striking a grating fixed to the cavity base. The cavity is excited by a microwave field (green) and has apertures to admit a horizontal beam (blue) for state preparation and detection.

This MOT preparation stage will be accomplished by means of a grating MOT, as introduced in Chapter 2, as that is the purpose of the experiment! The grating MOT will require a single input beam (unless launching, see Chapter 7) with two frequency components. The trapping and cooling is achieved by red-detuned excitation of the closed D_2 transition $F = 2 \rightarrow F' = 3$. This transition is nominally closed, but off-resonant excitation can excite atoms to the $F' = 2$ state, allowing decay back to $F = 1$. Therefore repumping on the $F = 1 \rightarrow F' = 2$ transition is required. An optical molasses is achieved by turning off the magnetic quadrupole field, allowing sub-Doppler cooling. An optimised molasses sequence with dynamic intensity and frequency control is required to achieve the lowest temperatures and maximise N_{det}.

6.2.2 State Selection

This stage prepares the atoms into a known internal state so that the MW pulses have the desired effect. The atoms must be prepared into either the $|F = 1, m_F = 0\rangle$ or $|F = 2, m_F = 0\rangle$ ground states as these are the states that will be excited by the Ramsey sequence (because they are first order magnetically insensitive). Several types of state selection are possible. The simplest is to pump all atoms into the $F = 1$ states by removing the repump beam during the last portion of the molasses. The remaining atoms will be spread across the magnetic sublevels but only those in $m_F = 0$ will contribute to the clock signal. The $m_F \neq 0$ states may cause a slight clock shift, as discussed in Chapter 9. This method can be improved by an additional π microwave pulse to transfer the $m_F = 0$ atoms to $|F = 4, m_F = 0\rangle$, and the other atoms can be pushed away using a resonant beam on any $F = 3$ transition. Finally, it is possible to optically pump directly into an $m_F = 0$ state, so that all of the atoms can be used in the clock procedure.[91] This method requires that the $m_F = 0$ state is 'dark' and not excited by the incident fields whilst every other state is

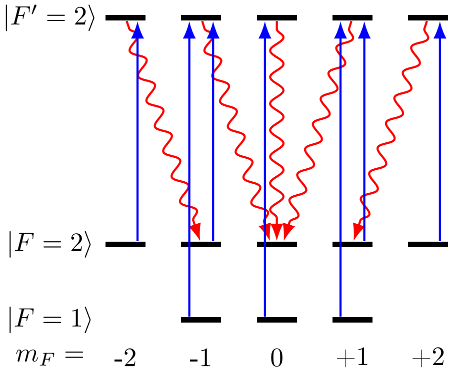

Figure 6.3: Optical pumping into $|F, m_F\rangle = |2, 0\rangle$. Not all decay pathways are shown, but $|2, 0\rangle$ is the only non-excited ground state.

excited. This can be achieved because of the selection rule that forbids the $\Delta F = 0, m_F = 0 \rightarrow 0$ transition.[92] This method was selected as no atoms are discarded, giving a larger total signal.

The selected pumping scheme is based on that in [93]. Pumping $F = 2 \rightarrow F' = 2$ and $F = 1 \rightarrow F' = 2$ with π-polarised light will pump all the atoms into $|F = 2, m_F = 0\rangle$ as shown in Fig. 6.3. The pumping can happen very rapidly because the process is limited only by the optical pumping timescale, $\frac{1}{\Gamma} \approx 25\,\mathrm{ns}$. A few photons must be scattered during the process which increases the atomic temperature by $1\,\mu\mathrm{K}$ to $3\,\mu\mathrm{K}$.[93] Because of the required π-polarised light, the pumping beam must be perpendicular to the bias magnetic field.

6.2.3 Microwave Pulses

The microwave pulses perform the actual Ramsey sequence which compares the LO frequency to the atoms. This is done with direct excitation of the microwave transition from $|F = 1, m_F = 0\rangle$ to $|F = 2, m_F = 0\rangle$.

The microwaves are applied with a microwave cavity as this reduces the ‘Doppler effect’ caused by phase gradients in the microwave field. The quality factor of the cavity will influence the MW phase uniformity, see Section 9.2.2. The MW transition is a π transition as $\Delta m_F = 0$, so the bias field must be aligned with the MW field, along the cavity axis.

6.2.4 Detection

In this stage it is necessary to detect the number of atoms in each of $|F = 1, m_F = 0\rangle$ and $|F = 2, m_F = 0\rangle$. Detecting the population in both states is necessary to calculate the fraction of atoms that were excited. If only one of these populations is measured then fluctuations in the total atom number are indistinguishable from fluctuations in the excited population, and contribute additional noise. Because all of the atoms should be in one of these two states, it is sufficient to detect the total population of each hyperfine level across all magnetic sublevels. Due to the close energy spacing it would be more difficult to measure the individual sublevel populations. The sublevel population can, and should, be measured during characterisation stages (see Fig. 8.7), but is not practical to measure during clock operation. The total population of $F = 2$ states can be measured simply by exciting the closed $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition. The amount of scattered light will be proportional to the number of atoms in $F = 2$ states, N_2. All the $|F = 1\rangle$ transitions are detuned by 6.8 GHz so will not be excited. Because the transition is closed, many photons can be scattered from each atom before the atom eventually decays to $|F = 1\rangle$ by off-resonant excitation to $|F' = 2\rangle$. Up to $8\frac{\Delta^2}{\Gamma^2} \approx 15000$ photons can be excited in this manner, with a reduction when strongly excited ($s > 1$) as this power-broadens the off-resonant transition. If a repump laser on the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition is added, the scattering will then be proportional to the total number of atom, $N_1 + N_2$ because atoms that were initially in $|F = 1\rangle$ are now also excited. The population of each state individually, and their ratio can then be calculated. Any polarisation can be used for the detection beams, as every ground state would be excited by every polarisation for these transitions.

The number of scattered photons can be measured using either fluorescence or absorption. Absorption allows all scattered photons to be detected as they are now absent from the beam, but is subject to intensity fluctuations and shot noise in that beam. The signal may be only a very small fraction of the beam power, and is around 1% maximum in this thesis! Fluorescence detects only a small fraction of the scattered photons due to the limited numerical aperture of the detection optics, but has (in principle) no background. The clock experiment will initially use absorption due to the limited optical access, though fluorescence could be advantageous in the future.

6.3 Optical systems

The optical system was designed to deliver all the optical fields required to drive transitions during an experimental sequence, as shown in Fig. 6.4(a). Optical fibres were used to take light from the optical system and deliver it to the main physics package. This decouples the two systems from each other and allows them to be separately aligned. Three polarisation-maintaining optical fibres were used: two carrying light for the MOT (as required by the launch procedure in Chapter 7), and one carrying light for the side beam. Both state detection and state preparation stages can be carried out using π-polarised light, so the side-beam fibre can carry all the light for these stages.

A nice feature of the optical transitions that have been selected for each stage, is that they are all D_2 transitions and can be driven by the same laser with some frequency shifting applied. All

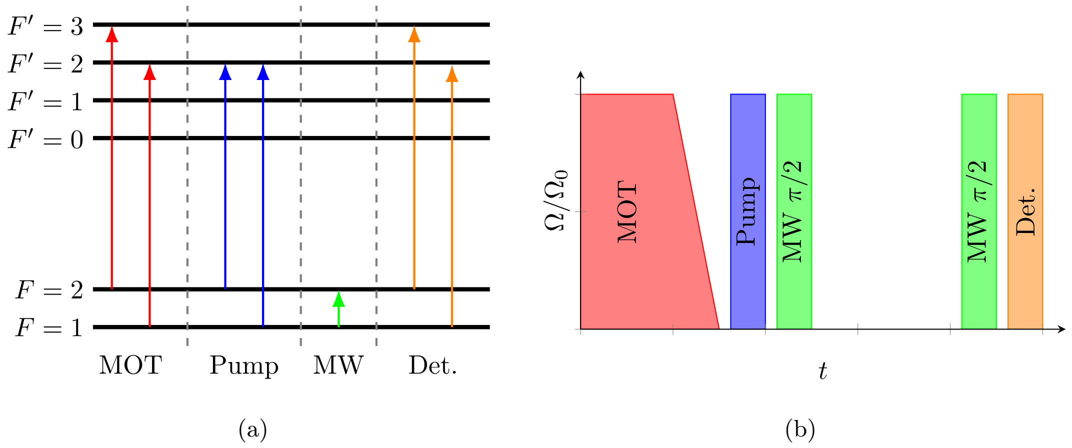

(a) (b)

Figure 6.4: (a) The transitions required for a full experiment sequence. (b) The corresponding experimental sequence. Indicative only, not quantitative.

the $F = 2$ transitions are within a 300 MHz span, easily covered by acousto-optical modulators (AOMs). The $F = 1$ transitions are only used to repump atoms, at times when an $F = 2$ transition is also active. This allows the use of an in-line electro-optic modulator (EOM) to generate the repump as a sideband of the $F = 2$ transition. A wide-bandwidth EOM is required to generate the repump for all stages of the experiment. The state-selection stage pumps on the $F = 2$ to $F' = 2$ transition, which is 6.835 GHz from the $F = 1$ to $F' = 2$ repump, but the MOT and detection stages pump the $F = 2$ to $F' = 3$ transition, which is 6.568 GHz from the $F = 1$ to $F' = 2$ repump. A single resonant EOM cannot usually cover both these frequencies, so a fibre-coupled EOM is used instead. Fibre-coupled EOMs are power-limited to around 10 mW output power,[94] so a tapered amplifier (TA) is placed after the EOM and the EOM is heated to 50 °C which allows higher power operation because photorefractive damage to the EOM can self-heal by ‘annealing’ at this temperature.[95,96] A block diagram of the optical system is shown in Fig. 6.5.

The seed laser is locked 300 MHz red-detuned from the crossover of the $F = 2 \rightarrow F' = 2, 3$ transitions using saturated absorption spectroscopy[97] and an AOM, which provides the frequency offset. This allows double-pass AOMs to produce the final laser frequencies. This allows dynamic frequency and intensity control of every beam, by changing the RF signal to the AOMs. Double-pass AOMs are used for two reasons; firstly, it allows greater frequency tuning range and secondly, the beam alignment is no longer sensitive to the AOM frequency.[98] The fibre-coupling is also much less dependent on the AOM frequency, although it is eventually limited by the bandwidth of the AOM. It may become necessary to shutter the laser beams, to achieve higher extinction and prevent light reaching the atoms during the Ramsey time. For this purpose space is left for a shutter in each beam path. The AOM frequencies are calculated as shown in Fig. 6.6. The state-selection

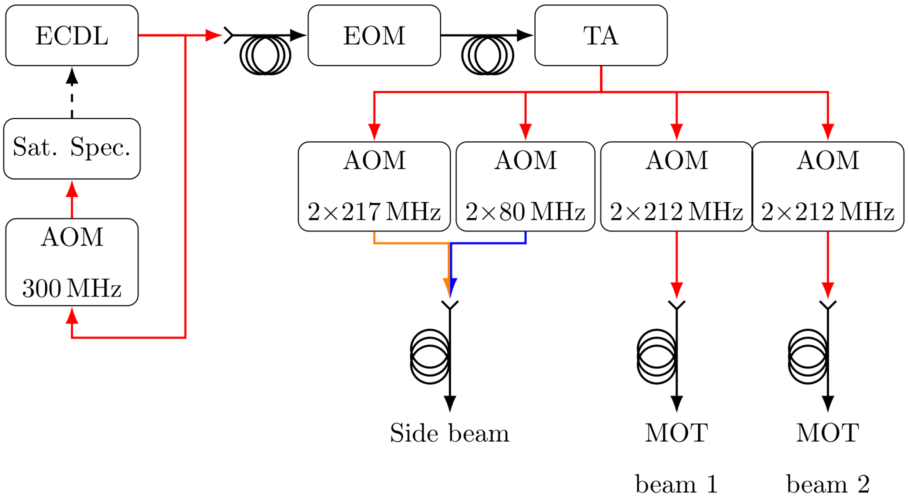

Figure 6.5: A block diagram of the optical system. The ECDL is locked by saturation spectroscopy to 300 MHz red-detuned from the crossover of the $F = 2 \rightarrow F' = 2, 3$ transitions. The EOM and TA add sidebands and amplify the light. Four double-pass AOMs (shown as 2×frequency) produce the final frequencies required.

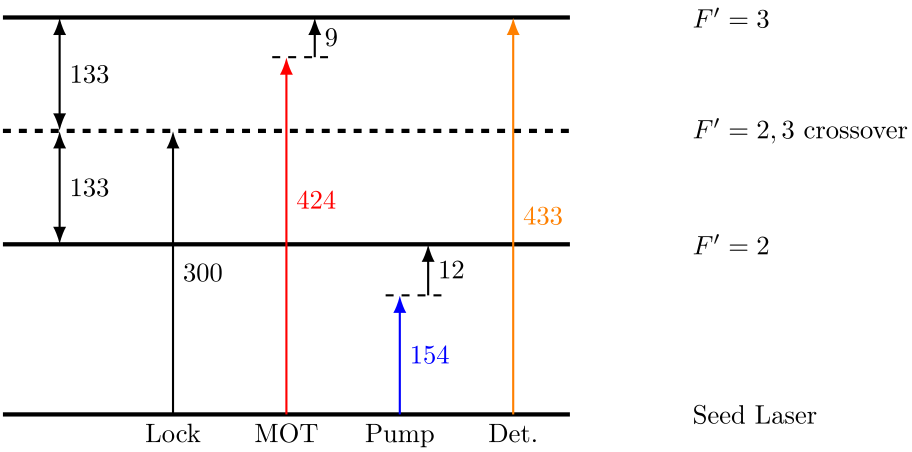

Figure 6.6: The AOM frequencies required for the clock. Offsetting the seed laser from the $F' = 2, 3$ crossover brings all the required $F = 1$ transitions within range of standard double-pass AOMs.

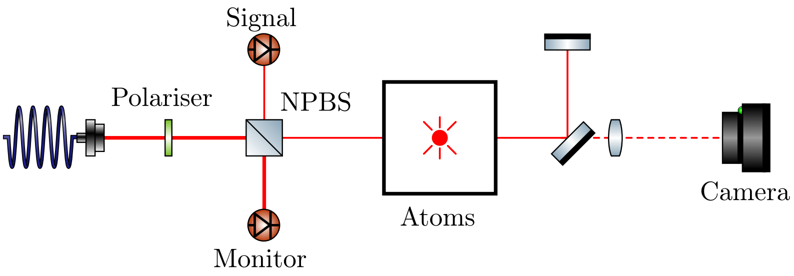

Figure 6.7: The optical system for the side-beam, simplified for clarity. A high extinction polariser in a rotating mount is used to align the beam polarisation with the quantisation axis. The non-polarising beam splitter(NPBS) has a 90:10 ratio of reflection:transmission. Removing one mirror allows the camera to image the atoms.

beam should be red-detuned approximately 2Γ for maximum pumping,[93] while the detection beam should be resonant to cause maximum fluorescence. The MOT beams are nominally detuned by 1.5Γ, but have significant tuning range to allow for molasses optimisation.

6.3.1 Beam-shaping optics

The beam-shaping for the MOT-beam will be discussed in Chapter 7. The beam-shaping for the side beam is done simply using a integrated fibre-connected collimation package. The beam then goes through a high-extinction polariser to ensure good polarisation-purity as required for the state selection stage. A 90:10 non-polarising beam splitter (NPBS) is used to divert 90% of the light onto a photodiode to monitor the beam intensity. The remaining light passes through the cavity, is retroreflected and then hits the same 90:10 NPBS so that the beam is directed to another photodiode, allowing a differential measurement of the absorption by the atoms. The side-beam optics are shown in Fig. 6.7. Removing one mirror allows the beam to strike a camera instead of being retroreflected, imaging the atoms so fluorescence or absorption images may be taken.

6.4 Physics package

The main physics package must hold the MW cavity in place, supply optical access for each of the required beams, allow control of the magnetic environment, and provide an ultra-high vacuum environment in which cold atom trapping is possible.

The main design decision for the physics package concerns the placement of the MOT gradient coils. These will be an anti-Helmholtz coil pair to supply the $\approx 15\,\mathrm{G\,cm^{-1}}$ field gradient required for a Rb MOT. These must switch off within 1 ms to 2 ms so that the sub-Doppler cooling can take place. The coils can be placed either inside or outside the vacuum chamber. Internal coils will be

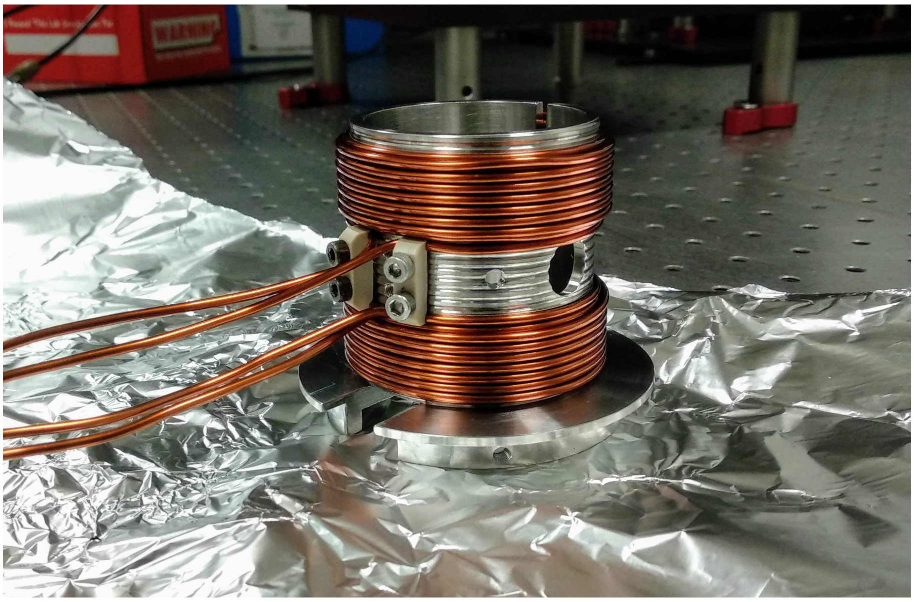

Figure 6.8: The coil former, which also acts as a mount for the MW cavity. Macor clips are used to hold the coil windings in place on the former, which has a spiral groove cut for the windings to fit into. The former has an internal diameter of 50 mm.

much smaller, and therefore require less power, and be easier to drive and switch. However it will be difficult to dissipate heat from in-vacuum coils and it will reduce the amount of space available in the vacuum system. External coils would require much more power (>100 W) and potentially water cooling. They would also create larger fields over a larger volume, potentially inducing eddy currents or magnetisation of other components.

Internal coils were selected to avoid the need for water cooling, and allow easier driving/switching. A coil former (Fig. 6.8) was designed which the coils would be wound onto, and which would also hold the MW cavity in place. The rest of the vacuum system could then be designed from mostly standard components, with optical and electrical feedthroughs as appropriate. A large $45\,\mathrm{L\,s^{-1}}$ ion-pump was specified with good conductance to the main chamber, because the outgassing properties of the 3D-printed cavity were not well known. It was later confirmed that the cavity did not outgas significantly. Fig. 6.9 shows the final design.

Magnetic shielding was not included due to the additional complexity that would be required in a prototype system. However the gradient coils can be switched off hard using FETs to prevent any current flow. Shim coils (one Helmholtz pair per axis) were used for control of the DC magnetic field, cancelling the background field and applying a quantisation field axis. Non-magnetic components were specified throughout the vacuum system, and the (stainless steel) optical table was demagnetised before any components were placed. The vacuum system was baked at 100 °C

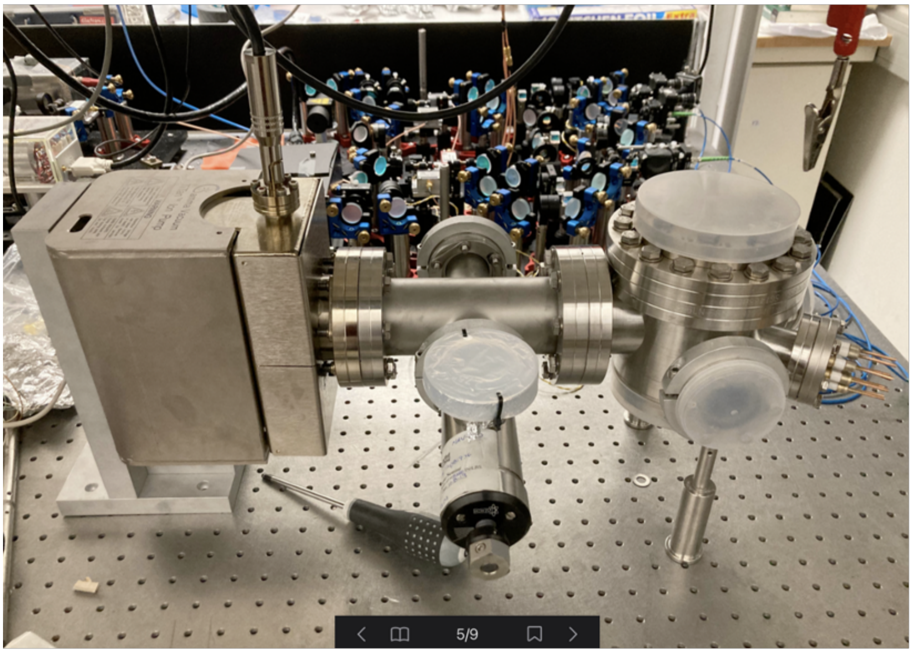

Figure 6.9: The vacuum system. The main physics chamber (right) is connected to the ion pump (left) by a long tube to reduce the influence of the magnetic field from the ion-pump magnets. The system is around 600 mm in length.

for > 1 day after pumping down to UHV levels, and the ultimate pressure of the system (according to the ion pump gauge) was around 1×10^{-10} mbar. The bake temperature was limited by the coaxial cable used to couple microwave signals into the cavity.

6.4.1 Microwave Cavity

The microwave cavity was supplied by Neuchâtel with one of our diffraction gratings pre-installed, see Fig. 6.10. The diffraction grating is a tri-sector grating with grating period of 1080 nm, giving a diffraction angle of 46° for light at 780 nm. The grating had a target etch depth of 195 nm and a 60:40 duty cycle of etched:non-etched surface. The grating was produced by nano-imprint lithography and coated with aluminium. It's important that the grating and coating are compatible with UHV and with Rb (gold coatings will react with Rb). The grating also had a 2 mm hole laser cut in it's centre, to allow launching (Chapter 7).

Neuchâtel specified that the cavity had initially had a quality factor, Q, of around 500, but this had decreased to around 185 when the grating was installed. This is an interesting phenomena that will likely require future investigation, as it suggests that the grating is significantly influencing the cavity mode.

On receipt of the cavity, it was tested to ensure that the resonance was correct (the cavity can

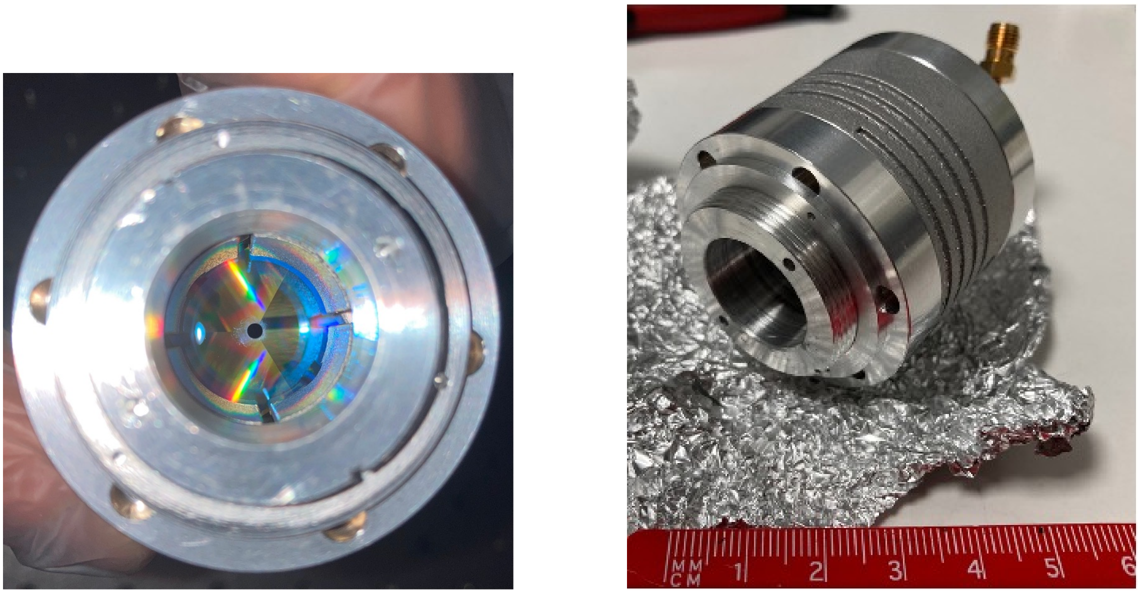

Figure 6.10: The microwave cavity with installed diffraction grating. Looking down the cavity axis (left) it is possible to see the grating with central hole and also the 4 cavity electrodes which strongly influence the cavity resonance.

be tuned to correct this), and to confirm the quality factor that was observed in Neuchâtel. The testing was done by measuring the reflection coefficient of the cavity as a function of frequency (the 'S_{11} parameter'), A coupler was used to separate the reflected power which was measured using an RF power detector. As the frequency was varied, a dip in reflected power was observed at the cavity resonance, see Fig. 6.11(a). Unfortunately, there was also a significant standing wave component caused by unwanted reflections in the RF path. This could probably be mitigated with more careful design and liberal application of attenuators, but it did not prove necessary. Examining the resonance lineshape in Fig. 6.11(b) gives the cavity resonance frequency and Q. Unfortunately the linewidth is not Lorentzian (or even symmetric) which means that exact fitting is difficult, but the lineshape does appear compatible with $Q = 185$ and the resonance frequency can still be pinpointed to 6835.8(2) MHz.

The cavity resonance frequency was also observed to change due to the physical changes it underwent when being installed in the experiment. When going from atmosphere to UHV, the frequency increased by 1.9(2) MHz, or 280 ppm, which is approximately the refractive index of air. When the cavity was baked at 85C-100C for 60 h and returned to lab temperature, a further increase of 3.3(2) MHz or 440 ppm. This was unexpected and the cause is unknown. It appears to have been a permanent change, perhaps due to changes in gasses adsorbed on the surface, or stress relaxation. Finally, a thermal sensitivity of $-130\,\mathrm{kHz\,K^{-1}}$ or $-19(1)$ ppm/K is seen. This is likely due to thermal expansion of the cavity, and should therefore be equal to the linear coefficient of expansion

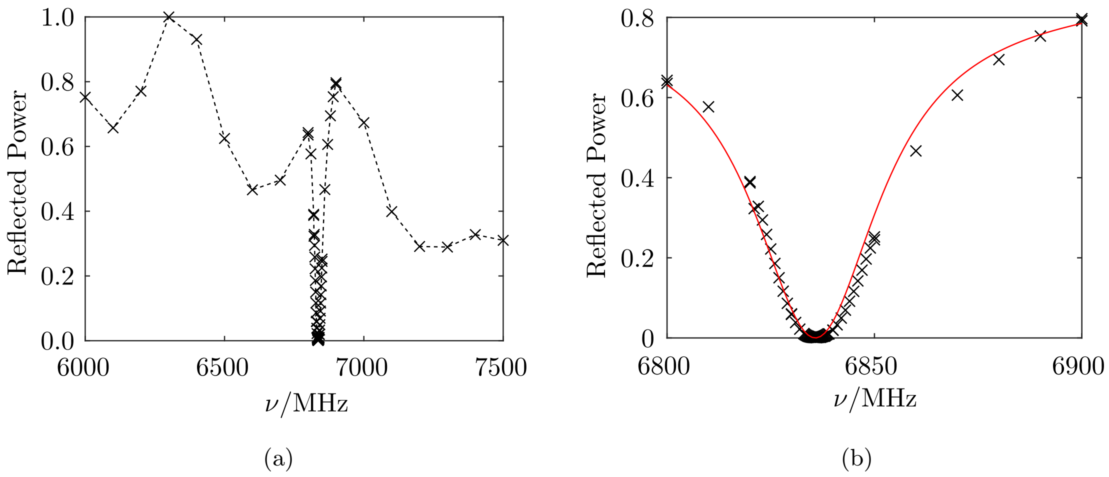

Figure 6.11: The reflected power from the cavity as a function of frequency. (a) shows that a standing wave is visible in the microwave system, but the resonance is still clearly visible. Dashed line is a guide to the eye. (b) shows the detailed lineshape. The red line is a Lorentzian with $Q = 185$ and a linear envelope applied which approximately accounts for the standing wave envelope.

for aluminium which is 23 ppm/K. The discrepancy is likely because the cavity temperature was improperly measured, or had not reached steady-state (the thermal time constant is 2 h to 3 h), or because the cavity is partially constrained by its mount.

Chapter 7

GMOT launch

This chapter will cover the grating-based launch mechanism that was developed in this PhD. The difficulties of launching with a GMOT are described, followed by a method to avoid these problems. That method was tested and the results are presented.

7.1 Grating launch mechanism

One limitation of cold-atom sensors where atoms are released from a MOT is that the interaction time is limited by the atoms falling out of the interaction region due to gravity. Launching cold atoms allows significantly longer interaction times as the atoms return to their original position after reaching a height $h = gt_{\text{launch}}^2/8$, where t_{launch} is the time for the atoms to return. In an atomic clock, launching allows a longer Ramsey time and therefore improves the short-term stability of the clock.[99,100] It is worth noting that the tradeoff between fountain height and clock performance is quite strong. A 10× increase in Ramsey time requires a 100× increase in fountain height. Sacrificing a small amount of performance due to a short Ramsey time may allow a substantial decrease in size.

As discussed in Chapter 2, the launch is typically achieved by manipulating beam frequencies during molasses, such that the atoms are cooled into a moving frame and a co-moving optical lattice is formed.[99] In a grating MOT, this is not normally possible. All of the beams have the same frequency because they are all generated by diffraction from one initial beam. The beams will actually form a highly stable optical lattice that is pinned to the grating, suggesting that it is not possible to use this technique for launching.

In order to achieve a launch with a grating, it is obvious that multiple frequencies must be present.

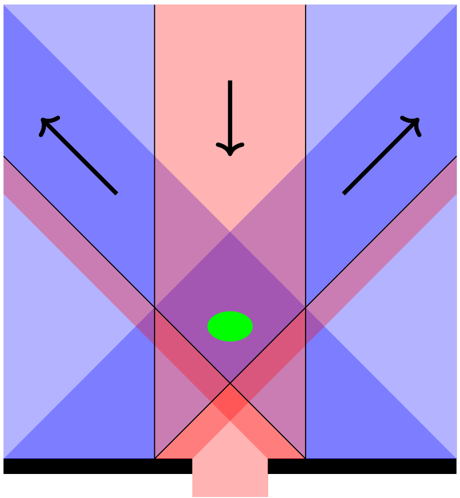

Figure 7.1: A modified GMOT for launching atoms, in cross-section. The inner (red) beam only strikes the atoms (green ellipse) from above but the outer (blue) beam only strikes the atoms from below. The grating has side lengths of 20 mm and the MOT forms ≈ 5 mm above its surface.

If all four tetrahedral beam frequencies were independent, then launching in an arbitrary direction would be possible. Two frequencies is the minimum that will be required to launch in a single direction, and the axial symmetry of a tetrahedral MOT suggests different frequencies for the incident beam and the diffracted beams. This produces a launch geometry similar to the (1,1,1) geometry that is used for some six-beam MOT launches.[49] Therefore, we will need to separate the parts of the input beam which strike the atoms from 'above' and 'below' (incident on the grating and diffracted from the grating, respectively). This can be done using the modification shown in Fig. 7.1. Two concentric beams are produced and strike the atoms from above and below respectively. A frequency detuning between the two beams will cause the optical lattice to move vertically and launch the atoms. It is necessary to have a hole in center of the grating, because any zeroth order reflection would be of the wrong frequency to strike the atoms from below, and could disturb the launch. The grating *should* be designed to prevent zeroth-order reflection by controlling the grating depth and duty cycle, but in real gratings design compromises must be made and there will always be some reflection, often at the level of a few percent.[40,47] The zeroth order would cause a small component of static optical lattice to be present, which is pinned to the grating. A static lattice would disrupt the transformation to a moving frame.

To launch the MOT with velocity $\vec{v}$ each beam must be frequency shifted by an amount $\delta\omega_i = \vec{k}\cdot\vec{v}$, where $\vec{k_i}$ is the wavevector of the i-th beam. The incident beam is directly opposed to $\vec{v}$, but the diffracted beams are angled away from $\vec{v}$ by the Bragg diffraction angle off the grating, $\theta = 40°$.

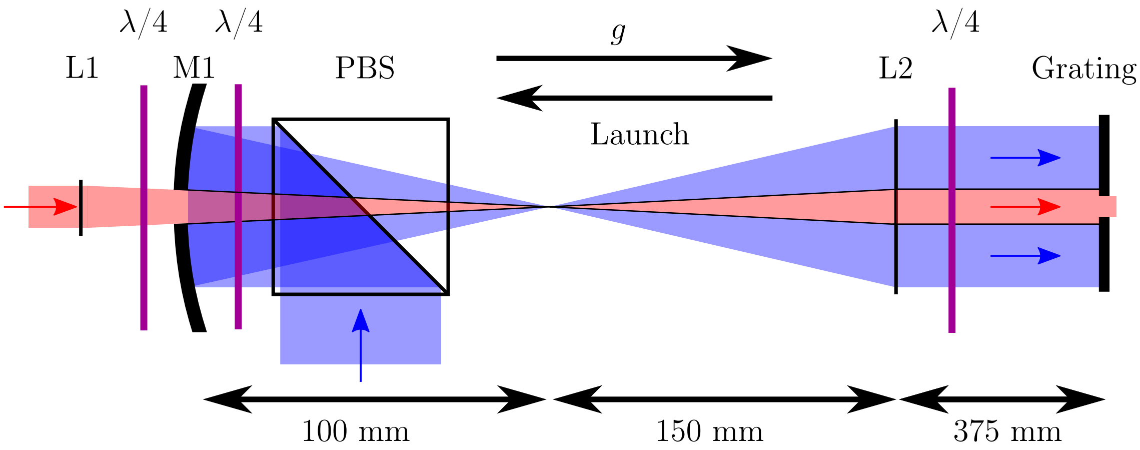

Figure 7.2: Optics for producing two concentric beams. Lens L2 forms an image of the mirror surface at the grating, greatly reducing diffraction effects. L1, L2, M1 have focal lengths of 200 mm, 150 mm and 100 mm respectively. The launch direction is opposite gravity, as indicated. Not to scale.

The frequency difference between the two beams will therefore be given by

$$\Delta f = \frac{1}{2\pi}\left(\delta\omega_1 - \delta\omega_2\right) = \frac{1}{2\pi}\left(kv\cos(\theta) - (-kv)\right) = \frac{v}{\lambda}(1+\cos(\theta)), \tag{7.1}$$

where $\lambda = 780\,\mathrm{nm}$ is the laser wavelength. A 100 ms flight time requires a launch velocity of $0.49\,\mathrm{m\,s^{-1}}$ and $\Delta f = 1.1\,\mathrm{MHz}$.

Producing the concentric beams does require some thought. Because the concentric beams must have the same polarisation, they cannot be combined using a simple polarising beam-splitter (PBS). Instead a physical aperture is used: a mirror with a hole in it. The central beam is fed through the hole, while the outer beam is reflected from the mirror face. As is shown in Fig. 7.2, the optical axes of both beams can be fully colinear, and the incident and reflected beams separated using a PBS. Using a sharp physical aperture has an unfortunate side-effect, the beams are strongly diffracted as they are clipped by the aperture. The diffraction would lead to spatial variations in intensity and mixing between the two beams. In order to mitigate the diffraction effects we use a re-imaging system; a lens is added which forms an image of the mirror (and its hole) at the grating surface. The optics system was designed to use a stock mirror with central hole (a Thorlabs CM254-100CH3-M02) as there are not many mirrors of this nature commerically available. This mirror is concave, but that is beneficial to forming the intermediate focus. The 3 mm diameter hole is magnified to form an 4.5 mm diameter inner beam. This diameter was chosen to keep the boundary between inner and outer beams a long way from the MOT for both the input and diffracted beams.

In order to form a good image, the inner beam should be clipped only at the mirror surface and

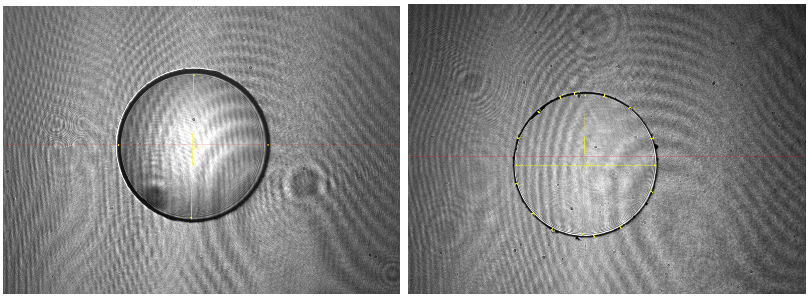

Figure 7.3: A comparison of the concentric beam profiles before (left) and after (right) a chamfer was added to the back of the mirror hole. The large-scale fringes are artefacts from the image sensor. The gap between beams is reduced from 150 µm to around 60 µm by using the chamfer. Beam profiles were taken at the grating position in Fig. 7.2, where a sharp image is formed.

not before. However, the beam is focusing through the mirror, and so is clipped more by the back surface of the mirror than the front. This leads to a gap between the inner and outer beams. To avoid this, a rotary tool was used to grind a chamfer into the back of the mirror hole. This reduced the gap between beams as shown in Fig. 7.3, but did not eliminate it as the chamfer could not extend all the way to the mirror surface without risk of damaging the mirror coating. Further, because the beam interacts with the atoms both before and after the grating surface, it is necessary that the beam should not have strong distortions in the $\approx 10\,\mathrm{mm}$ either side of the plane of the grating. Moving a camera through the focal plane allows imaging the beam as it propagates, which is shown in Fig. 7.4. A gap appears in the beam before it reaches the focal plane, but not after the focal plane. This corresponds to imaging a plane that is behind the mirror where the beams are combined, so it makes sense that this causes a larger distortion. Most importantly, no significant effects are seeen in the center of the beam, where the atoms are located. The MOT is confined to a small volume $\approx 1\,\mathrm{mm}$ in extent. The distortions in Fig. 7.4 should not affect the MOT launching, although they may have an effect on the MOT loading, which occurs throughout the full beam overlap volume.

7.2 Launch results

Using the optics of Fig. 7.2, the GMOT launch was characterised. This used the apparatus of Chapter 6 but without using the microwave cavity. This allowed improved optical access for imaging the MOT launch from a transverse axis.

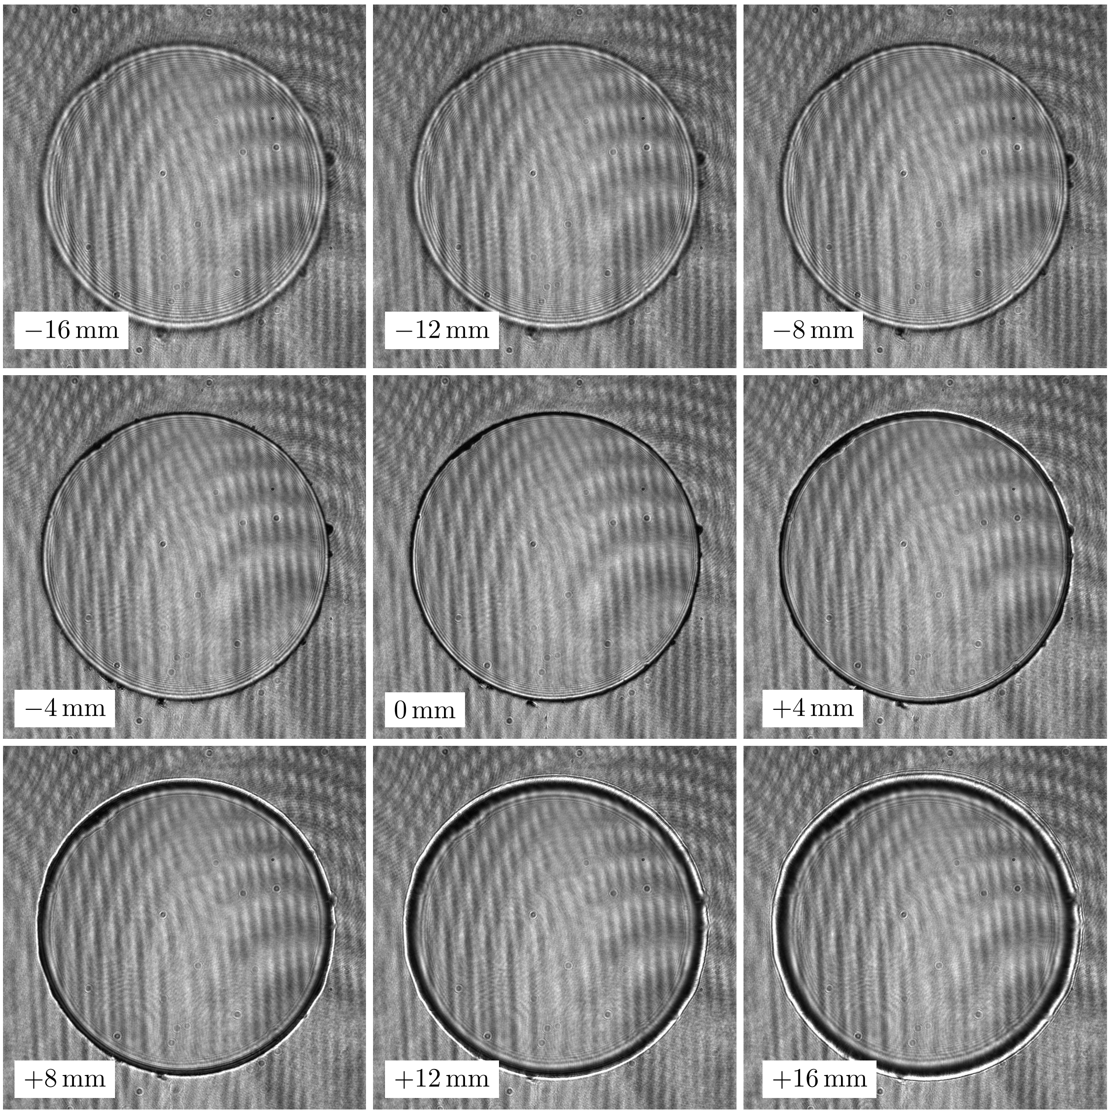

Figure 7.4: Beam profiles taken moving around the focal plane along the beam axis. Moving upstream of the focal plane (positive values) causes a gap in the beam to appear. Each image has sides of length 6.5 mm

7.2.1 Trapping atoms

Initially, it was difficult to align the concentric beams to the grating. However, a system was developed to ensure alignment. First, the inner beam and outer beam were aligned to each other, by ensuring that they were well aligned at their shared focus, and also in the far-field. Then the inner beam alone was aligned to the grating, using only optics after the beam combination. The relative alignment of the beams remains good so that both beams are aligned to the grating. Positional alignment was achieved using the transmitted beam through the grating hole. Angular alignment was achieved by observing the weak zeroth-order reflection from the grating around the hole. This was aligned to be coincident with the input beam, ensuring the beam was normal to the grating. At this point, varying the intensity balance between the two beams was sufficient to see a very weak MOT, which could then be optimised with further alignment.

Once the MOT was optimised, around 10^7 atoms could be trapped. A molasses sequence was then created which cooled the atoms to a sub-Doppler temperature. Throughout this process, both MOT beams were always kept at exactly the same frequency, mimicking a standard single-beam GMOT. To do this, an AOM driver based on a two-channel direct digital synthesis (DDS) was used to drive the AOMs in the two beam paths. This gives inherent phase-stability between the electrical signals driving the AOMs, with rapid and flexible control of the relative frequency and intensity of the two MOT beams. During the molasses both beams must be detuned from resonance by around 60 MHz in order to get sufficiently cold temperatures. To launch, the differential frequency between the two MOT beams must be dynamically controlled over a range of about 1.5 MHz with an accuracy around 50 kHz. Voltage-controlled oscillators do allow for tunability, but a relatively complex system would be needed to dynamically tune the frequency difference with the required accuracy. The DDS system is preferred for its simplicity. Additionally, the MOT is quite sensitive to the intensity balance between inner and outer beams, so closed-loop control of the intensity in each beam was implemented. Photodiodes were inserted into an unused area of each beam prior to combination, measuring the optical power. A homebuilt transimpedance amplifier and differencing circuit was used to produce an error signal proportional to the difference between desired and measured intensity. This error signal fed into the DDS controlling the AOM drive power, which has built in PID controllers.

The details of the molasses sequence changed over time due to optimisation and different experimental circumstances, but a representative example sequence is shown in Fig. 7.5. First, the magnetic-gradient field was turned off over 5 ms, then the optical intensity is reduced to 60%, so that a constant power can be maintained through the AOMs as the MOT detuning is linearly increased to 60 MHz over 2 ms. The final stage is to exponentially decrease the optical intensity

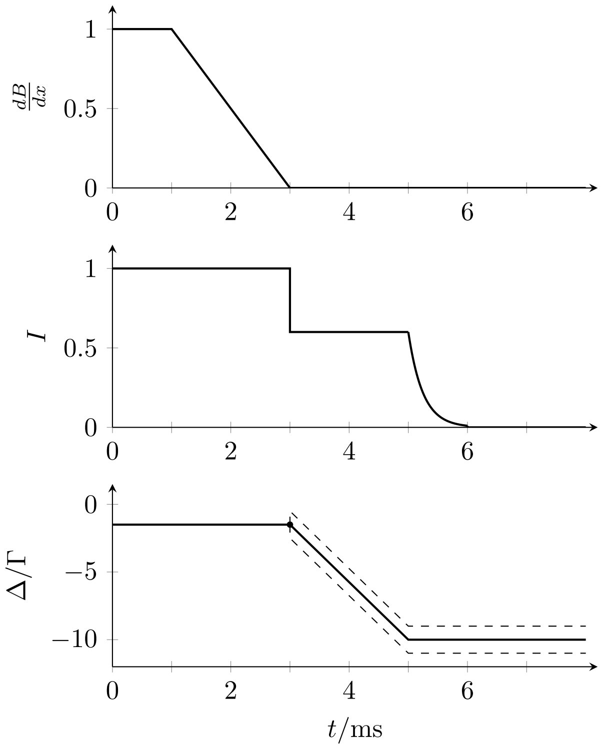

Figure 7.5: An example molasses sequence. The relative magnetic gradient (top), relative intensity (middle), and detuning of the MOT beams (bottom) all vary during the sequence. To launch the atoms, a frequency difference is introduced between the two MOT beams, as shown by the dashed lines. The frequency difference is exaggerated for clarity.

to 1% over 1 ms before extinguishing the MOT beams. A MOT temperature of 9 µK was measured in both the radial and axial directions. To measure the MOT temperature, time-of-flight fluorescence images were taken. When released from molasses the MOT starts to expand due to the thermal velocities of the atoms. The MOT size as a standard deviation, σ, after a time t will be approximately $\sqrt{x_0^2 + k_B T t^2/m}$ where x_0 is the initial MOT size and m is the mass of ^{87}Rb. By pulsing the MOT beams, the MOT can be made to fluoresce at different times after its release, giving images of the MOT for different values of t, and the temperature can be extracted from a plot of σ^2 against t^2, as shown in Fig. 7.6.

7.2.2 Launch

Once a working molasses sequence was found, a frequency difference was introduced between the two MOT beams to induce a launch. The frequency difference was constant throughout the entire molasses sequence starting at the time where the magnetic gradient coils were switched off, as in Fig. 7.5. A response was seen immediately as the atoms were launched upwards. The launch was then characterised with two main methods, both operating on a shared axis perpendicular to

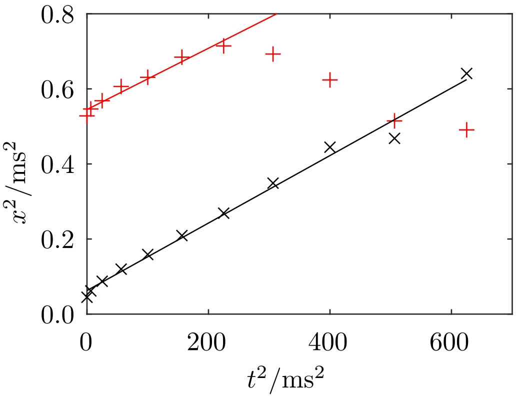

Figure 7.6: Atom cloud expansion after being released from molasses. The vertical (black ×) and horizontal (red +) size of the cloud as measured by standard deviation of a fitted Gaussian. As the atom cloud falls, its wings leave the overlap volume and fluoresce less leading to an apparent decrease in horizontal extent after 15 ms. Each point is from a single image of the atoms. The gradients of the lines of best fit correspond to temperatures of 9.4(4) µK and 8.5(9) µK in the vertical and horizontal directions respectively.

the launch direction. Firstly, fluorescence images were taken of the atoms, as they were excited using the main MOT beam. The second method was to use absorption techniques that had been developed over the years in the existing cold-atom CPT clock. The absorption from the side beam was used to measure the number of atoms that could be interrogated by the beam. The beam had a 7.5 mm $1/e^2$ diameter.

Turning the absorption and fluorescence signals into estimated atom numbers requires some calculation. The two techniques both operate with on-resonance beams, but use different intensity regimes designed to make the measurements less sensitive to the absolute power of the exciting beam. Fluorescence is done using a strong excitation ($s > 10$) to saturate the atomic transition. This means that each atom scatters photons at a rate of $\approx \Gamma/2$. If the camera system has a collecting solid angle of $d\Omega$ and an efficiency (counts/photon) of η then the detected number of counts will be

$$N_c = \frac{\Gamma}{2}\frac{d\Omega}{4\pi}\eta N_{\mathrm{at}} t_{\mathrm{exp}} \ , \tag{7.2}$$

after an exposure time t_{exp}.[19] In contrast, the absorption measurements operate in a low power regime with $s < 0.1$. This means that the transition remains unsaturated and each atom has a constant scattering cross-section of

$$\sigma_0 = \frac{\Gamma\hbar\omega}{2I_{\mathrm{sat}}} = (370\,\mathrm{nm})^2 \ . \tag{7.3}$$

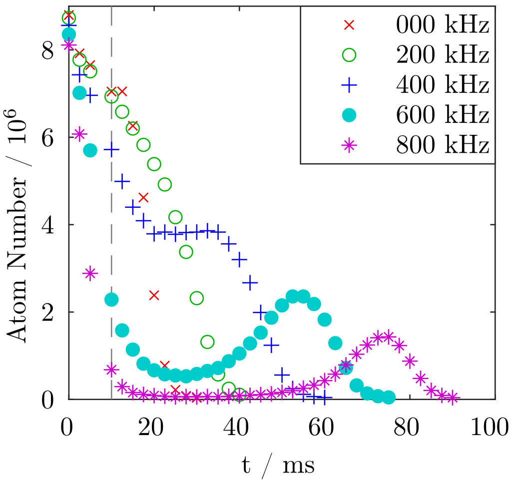

Figure 7.7: Detected atom number over time for a range of MOT launch beam frequency differences, Δf. The dashed line at 10 ms corresponds to the time at which images in Fig. 7.8 were taken.

The atoms then have an optical depth of $n\sigma_0$ where n is the area density of atoms (integrated along the imaging axis). By comparing images of the probe beam with and without the atoms present, the atomic density profile is calculated, and then integrated to give the total atom number.

After some optimisation of the moving molasses sequence, the results were quite promising, in that a measurable number of atoms could be seen at times up to 75 ms after the atoms were launched, see Fig. 7.7. However, there was significant atom loss relative to the initial atom number, much more than would be expected from the initially measured temperature of 9 µK. This temperature should lead to a MOT $1/e^2$ diameter of 9 mm after 75 ms and $> 40\%$ of the initial atom number should be in the observation window, not the observed $\approx 17\%$.

The time-of-flight images in Fig. 7.8, show that as Δf is increased, the MOT is launched faster, in good agreement with equation (7.1). However, the MOT also expands in the launch direction, indicating that the atoms are not being properly cooled in that direction. The reason for this reduced cooling is unknown, though I will give some hypotheses shortly. By integrating across the image in vertical and horizontal directions, the MOT shape along those axes can be examined, as in Fig. 7.9. This reveals that the cloud is no longer close to Gaussian in the launch direction, having a more flat-top profile. The cloud has an approximate FWHM of 1.15 mm in the launch direction after 10 ms, so an approximate temperature of 140 µK could be assigned.

One possible cause of the reduced cooling is noise on the frequency difference of the two beams. Although a DDS system was used, the dynamic frequency changes were made via analog control. The launch velocity is directly proportional to the frequency difference, equation (7.1), so fluctua-

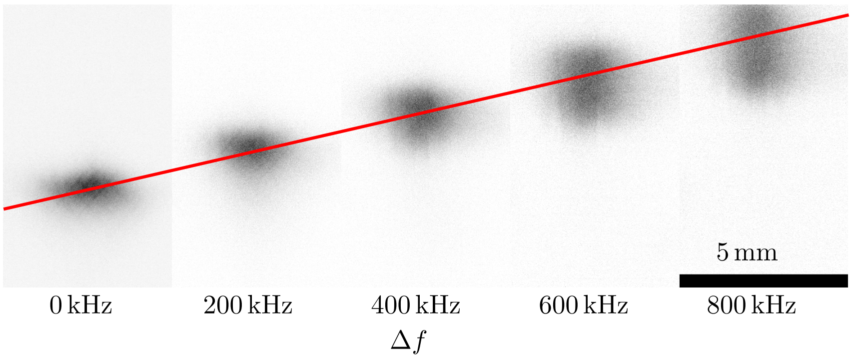

Figure 7.8: Time of flight images for a range of launch beam frequency differences, Δf from 0 to 800 kHz. All images were taken 10 ms after the launch ended. As Δf increases, the MOT is launched faster and it expands more in the launch direction. The red line represents the expected velocity as a function of Δf using equation (7.1). Gravity is downwards.

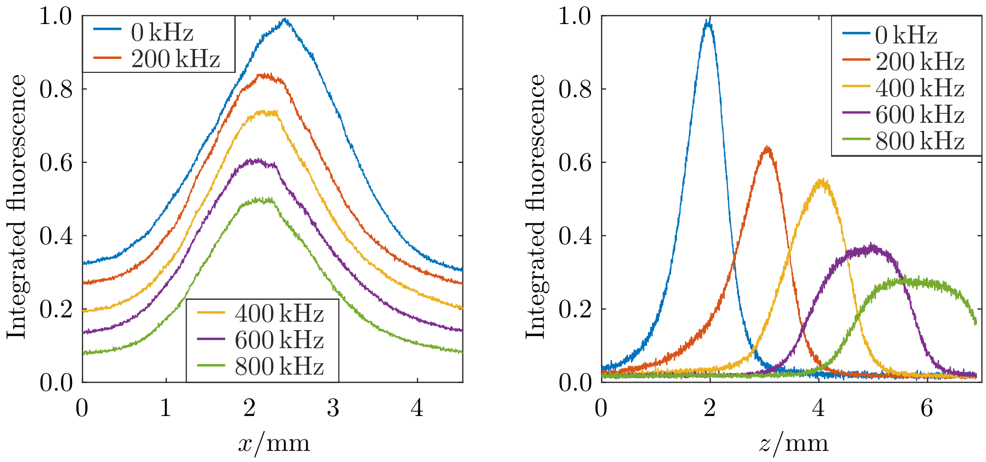

(a) (b)

Figure 7.9: The images from Fig. 7.8, integrated across the horizontal (a) and vertical (b) directions and then normalised. Curves of (a) are offset vertically for clarity. Dashed lines indicate the expected cloud position, assuming instantaneous acceleration to the final velocity when the frequency difference is imposed.

tions in the frequency difference could cause a spread of atomic velocities in the launch direction, equivalent to a higher temperature. However, it is not clear why this noise would only occur when doing a launch, and would get worse when launching faster. During an unlaunched molasses sequence the MOT beams are both detuned by 60 MHz. During a 75 ms launch they are detuned by 60.4 MHz and 59.6 MHz respectively, so it is surprising to have such large differences in behaviour. To check the differential frequency noise, the optical beatnote between the two frequencies was monitored on a photodiode during a mock molasses sequence. The result is shown in Fig. 7.10. During the frequency sweep there is an increased amount of noise, but during the final portion of molasses the frequency difference is highly stable at the exact value desired. The abrupt change in behaviour is because the frequencies are controlled by analog voltages during the frequency sweep but the final value, corresponding to a full scale signal, is fixed digitally. The velocity spread shown in the fluorescence images would correspond to frequency fluctuations of 200 kHz peak-peak, which is clearly not present.

Another possible cause is the introduction of a light of the wrong frequency to the moving molasses which prevents the frame-transformation to a moving frame from being exact. If light from the inner beam is incident on the MOT from below, or if light from the outer beam is incident from above, a static component of optical lattice would be introduced to what should be a pure moving lattice. This leakage could happen if the beam through the hole in the grating is partially reflected back to the MOT. I deliberately misaligned the input beam to an angle such that the reflected beam could not reach the MOT, but the reduced cooling was still observed, so such reflections were not the problem. Scatter or diffraction of the inner beam from the hole in the grating could still contribute. An image of the grating hole from its manufacturing report, Fig. 7.11, shows that there is some damage around the grating hole caused by the laser cutting, which could increase scatter.

Finally, the motion of the cloud during the launch could have an effect. However, the length of the molasses sequence did not have a strong effect on the amount of cooling that was seen. Molasses sequences as short as 2.5 ms total were used, during which the MOT moves < 1 mm, but no significant change was seen in the vertical spreading. Ultimately, the cause of this effect remains unknown.

In summary, the demonstration of a GMOT launch is a successful proof of principle. The launch allows the Ramsey time to be extended from 10 ms to 75 ms but also demonstrates anomalous heating, resulting in a loss of atoms and reducing the expected SNR. In the next chapter I shall use the GMOT launch to generate clock signals and examine the trade-off between Ramsey time and SNR.

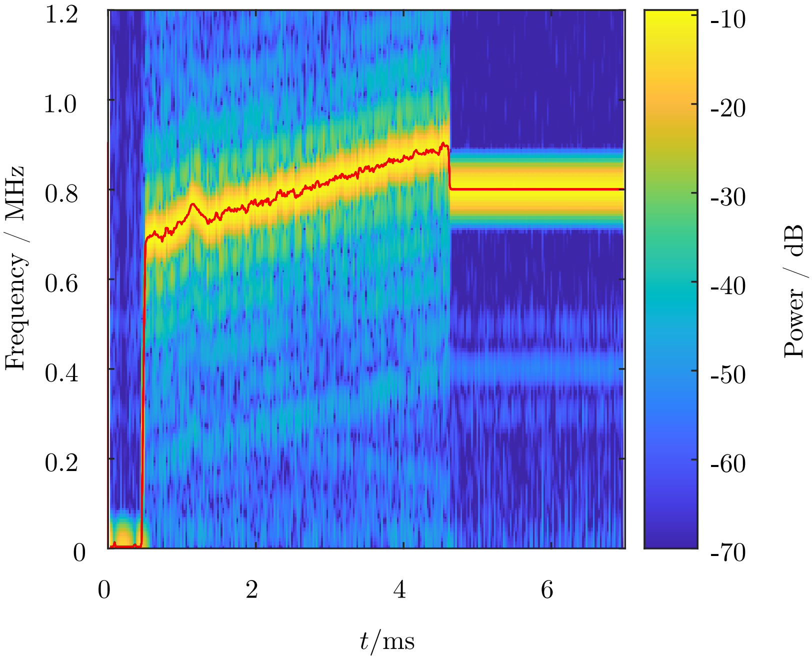

Figure 7.10: A spectrogram of the optical beatnote between the two MOT beams during a mock molasses sequence. The MOT beams are detuned 60 MHz over 4.2 ms and then remain at a fixed frequency. The optical power is kept high throughout to give a clear signal. The frequency width is due to FFT resolution. The red line indicates the estimated central frequency over time.

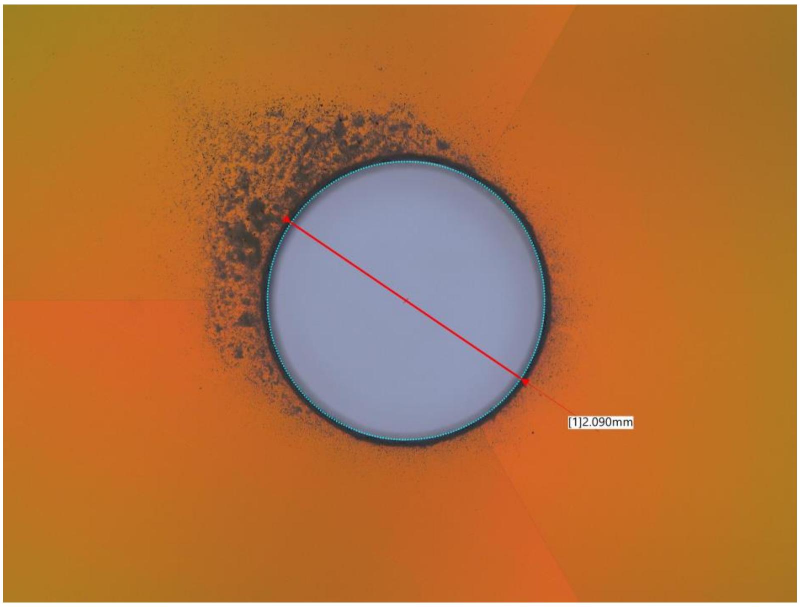

Figure 7.11: The grating hole, as imaged in the manufacturing report. Damage is observed from the laser cutting process.

Chapter 8

Clock Operation

This chapter details the initial tests of the GMOT launch method as part of a clock. Initial results were taken using a CPT-Ramsey interrogation, as described in Chapter 5, before switching to a MW cavity-based approach at a later time.

8.1 CPT results

Initial results were taken using a CPT-Ramsey technique. This is because we already had a working CPT-Ramsey clock in the group and wanted to test the GMOT launch using a known-working clock system. Additionally, these tests could be done without having to install the microwave cavity. The CPT light was carried from the CPT clock lab to the cavity-clock lab using a 20 m optical fibre. The same side-beam launch that would be used for state-selection and detection in the cavity clock, Chapter 6, was also appropriate for introducing the CPT light and measuring its absorption.

As explained in Chapter 5, we used a CPT system based on excitation of the D_1 line to the $F' = 2$ state. We used a σ^+-σ^- polarisation scheme, which required minor modification to the beam optics as shown in Fig. 8.1. This scheme has both frequency components in the same polarisation state, so they can both be launched into the same optical fibre without worry of polarisation crosstalk. Additionally, any polarisation drift due to the long fibre could be converted to intensity noise by a high-extinction polariser. Two CPT pulses are used, and the launch velocity is chosen such that the atoms will be returning through the CPT beam at the time of the second pulse.

In the CPT laser system the main laser (carrier frequency) is locked at a fixed frequency, resonant with the $F = 2 \rightarrow F' = 2$ transition. The second (sideband) frequency is added by an EOM at the

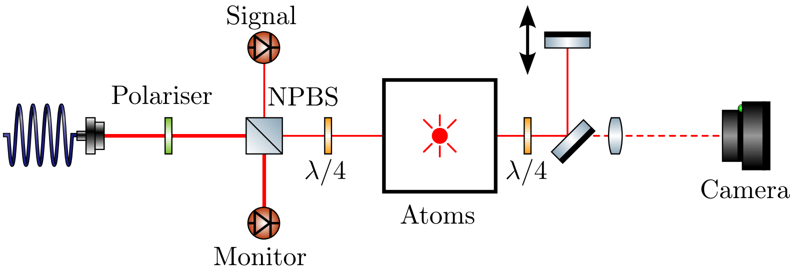

Figure 8.1: Two $\lambda/4$ waveplates are added to the side-beam launch, producing the circular polarisations required for σ^+-σ^- CPT. The retroreflecting mirror is placed on a translation stage so that the microwave phase of the forward and reverse beams can be matched at the position of the atoms. This is only sensitive to a few mm as the microwave wavelength is 44 mm.

hyperfine splitting frequency. When scanning frequencies, only the EOM frequency is changed, so it is Δ_1 which changes, and this causes changes to both $\Delta_{1\mathrm{ph}}$ and $\Delta_{2\mathrm{ph}}$. However $\Delta_{2\mathrm{ph}}$ determines most of the important behaviour.

The steady-state amount of absorption provides lots of information regarding the CPT resonance. As the EOM frequency is scanned over a large range an absorption feature is seen, Fig. 8.2(a). When the sideband is far from resonance, the atoms are quickly pumped into the lower hyperfine level and are not strongly excited. Significant absorption only occurs when *both* frequencies are near resonance. However, in the centre of the absorption feature there is a very sharp peak where the absorption no longer occurs. This is the formation of a dark state which does not interact with the light.

Closer examination of the central feature reveals that it has 3 separate peaks, see Fig. 8.2(b). The two side-peaks show the formation of magnetically sensitive dark states, as discussed in Chapter 5, with a small frequency offset. These peaks can be used to verify that the magnitude of the magnetic bias field is 220(10) mG. The central feature corresponds to our desired, magnetically insensitive dark state.

The CPT-Ramsey fringes that allow a clock measurement are taken from the absorption of the second pulse. When the atoms are pre-prepared in the dark state we should not see a pumping transient at the start of the pulse, although the photodiode will have a transient response, as in Fig. 8.3(a). However, if the atoms have evolved into the bright state that signal is modified by transient absorption as the atoms are pumped back to the dark state, shown in Fig. 8.3(b). The signal used for the clock divides the average transmitted signal during the first 100 µs of the pulse by that during the steady state (500 ms to 900 ms of the pulse). The 100 µs integration period was found to maximise SNR. A second stage of normalisation is also used, where the same quantity is

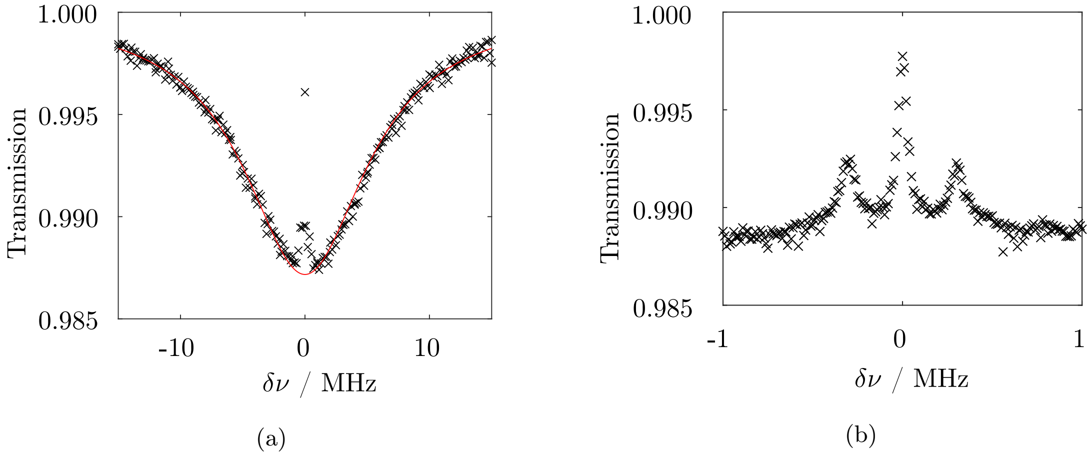

Figure 8.2: Absorption features seen when passing the CPT beam through the atom cloud. (a) The absorption line is seen and fits well to a Lorentzian with 6 MHz linewidth (red). (b) A detailed scan of the central CPT feature shows the formation of three different dark states, separated by the Zeeman splitting.

calculated for the reference photodiode which has not interacted with the atoms, and the signal is divided by that amount. This technique reduces the amount of noise by cancelling fluctuations in the laser intensity and was developed for the main CPT clock experiment.[48]

Exemplar Ramsey fringes are shown in Fig. 8.4. It is clear that longer Ramsey times (with faster launches) lead to a decrease in both fringe width and signal amplitude. The loss of signal amplitude is consistent with the reduction in atom number due to the launch, as measured by total absorption from the probe beam. The fringes are centred on an offset of $\approx 30\,\mathrm{Hz}$ which could be caused by a 228 mG magnetic field, consistent with the measured bias field.

For each of the Ramsey fringes of Fig. 8.4(a) we can calculate an SNR and fringe width. This allows calculation of the corresponding single-shot fractional instability using Section 3.2.3 and equation (3.29),

$$\sigma_{y,\mathrm{ss}} = \frac{2\sigma_\mathrm{S}}{S_\mathrm{p-p}} \times \frac{1}{2\pi\nu_0 T} \ , \tag{8.1}$$

with σ_S and $S_\mathrm{p-p}$ being the noise and peak-to-peak amplitude of the signal fringes respectively. The fringe width is determined by the Ramsey time T.

The fit residuals are broadly consistent with an electronic noise of $\sigma_\mathrm{S} = 0.0102(6)$, in the normalised units of Fig. 8.4, except for $T = 50\,\mathrm{ms}$ where the measured σ_S is 0.0127. This outlier was caused by an unknown transient effect which perturbed the experiment in the middle of this run, which is visible in the residual plot of Fig. 8.4(b). Using $\sigma_\mathrm{S} = 0.0102$ and the experimental fringe amplitudes we calculate the single-shot stability for each Ramsey time, as shown in Fig. 8.5(a).

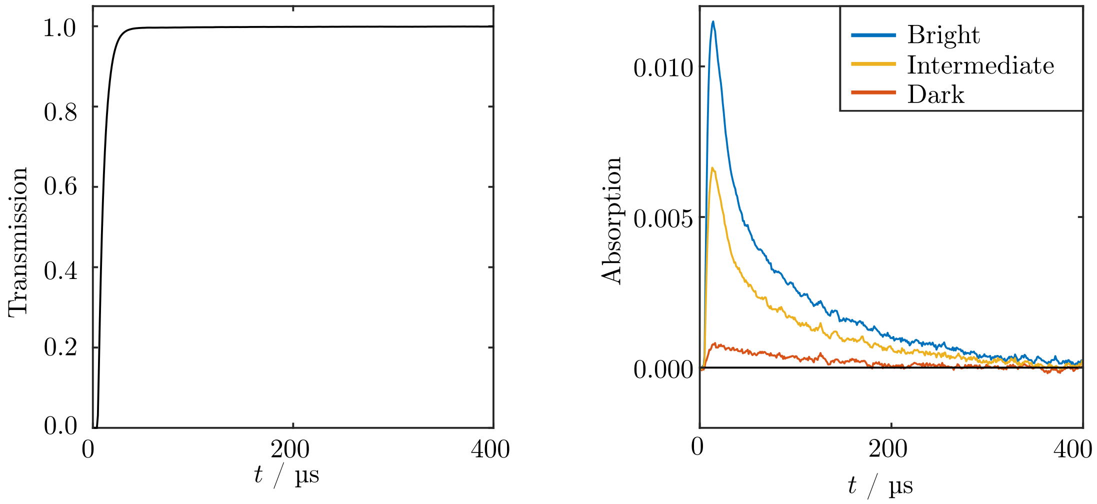

(a) (b)

Figure 8.3: (a) The transmission signal seen when a CPT pulse is recorded on the signal photodiode, with no atoms present. (b) When atoms are present, some additional absorption is seen corresponding to the number of atoms in the bright state. The atoms are pumped to the dark state, so only transient absorption is seen. The three curves shown are averages of 15 data runs each, and correspond to atoms that are fully in the bright state, in an intermediate state, and fully in the dark state.

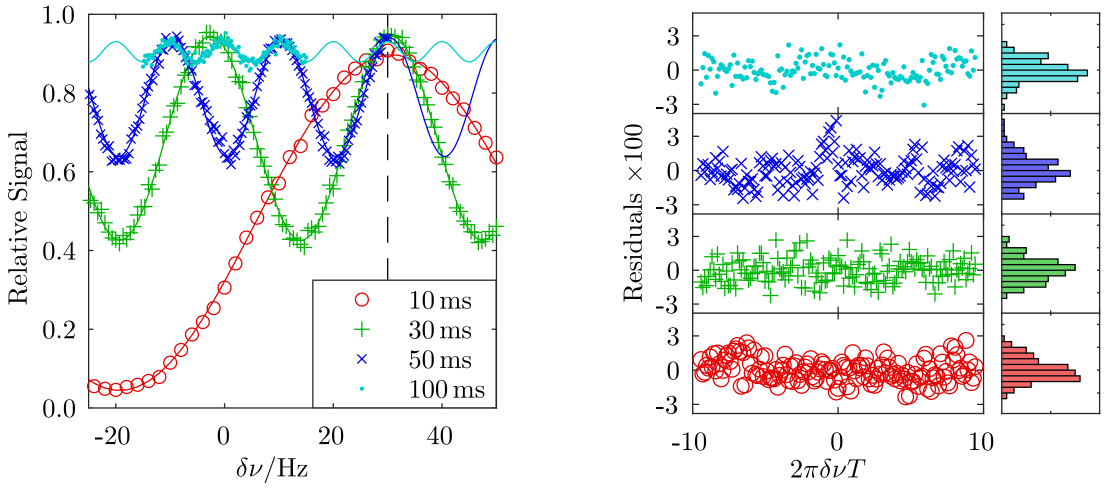

(a) (b)

Figure 8.4: (a) Clock fringes for different Ramsey times. Curves are best-fit sinusoidal fringes to the single-shot data points. The dashed line marks the centre fringe which is shifted 30 Hz from the accepted hyperfine transition frequency due to the second-order Zeeman effect. (b) Residuals from the fits in (a), with histograms.

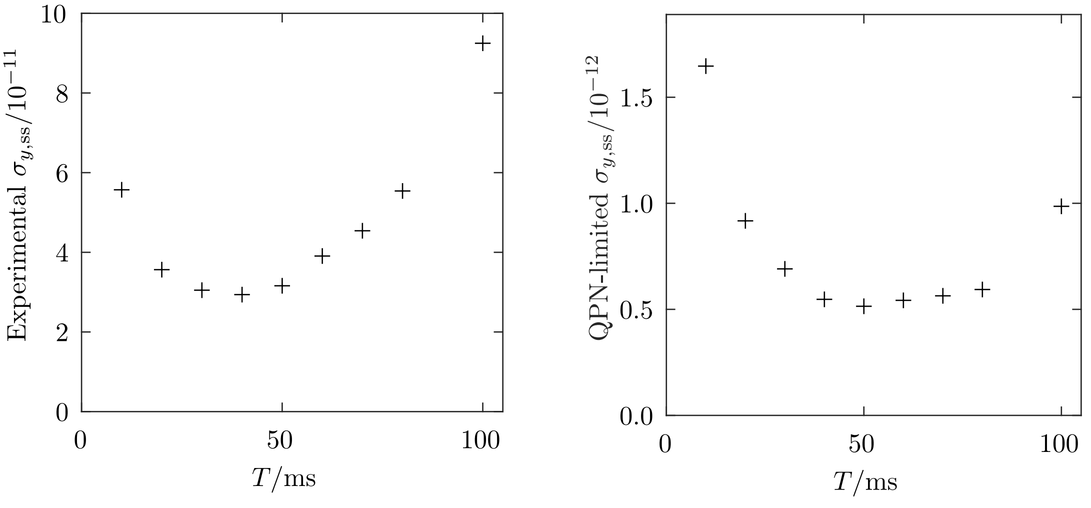

(a)

(b)

Figure 8.5: (a) The single-shot stability that was experimentally observed using CPT interrogation for different Ramsey times. The launch velocity was changed to allow each Ramsey time. (b) The estimated single-shot stability that would be observed if the measurements of (a) were QPN-limited.

The best stability of 2.9×10^{-11} was achieved at $T = 40\,\mathrm{ms}$.

A different clock design with a better detection method could achieve much lower noise and ultimately be limited by the atomic quantum projection noise (QPN).[101,102] In this case we would expect $\sigma_\mathrm{S} \propto \sqrt{N}$ for atom number N, and the impact of atom loss at longer Ramsey times would be mitigated by decreased noise. Fig. 8.5(b) shows this change in behaviour, using N proportional to the steady-state depth that is observed in the detection pulse. The calibration of absorption depth to atom number is quite uncertain (within a factor of 2), so the whole curve of Fig. 8.5(b) could move up or down but would retain its shape.

For QPN limited measurements, the best single-shot stability would be observed at $T = 50\,\mathrm{ms}$ but with only a 15% decrease in stability out to $T = 80\,\mathrm{ms}$. The fountain method would improve stability by a factor of up to 3.3 compared to the 10 ms result, which is a reasonable approximation of the performance achievable without a fountain. Additionally, good performance at longer Ramsey times allows the effects of any phase shift induced during the Ramsey pulses to be reduced, improving the accuracy and long-term stability.

8.2 MW cavity results

Once the GMOT launch had been characterised, the experiment moved towards the original design of a MW cavity clock. The cavity was installed, characterised and a MOT was trapped within.

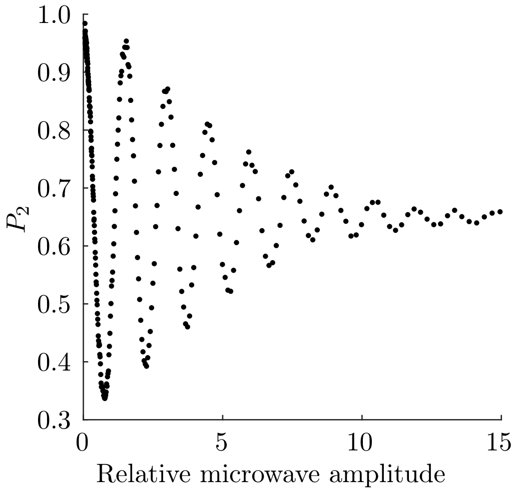

Figure 8.6: A 167 µs microwave pulse of varied amplitude is used to drive Rabi oscillations. A relative amplitude of 1 corresponds to 1 dBm microwave power.

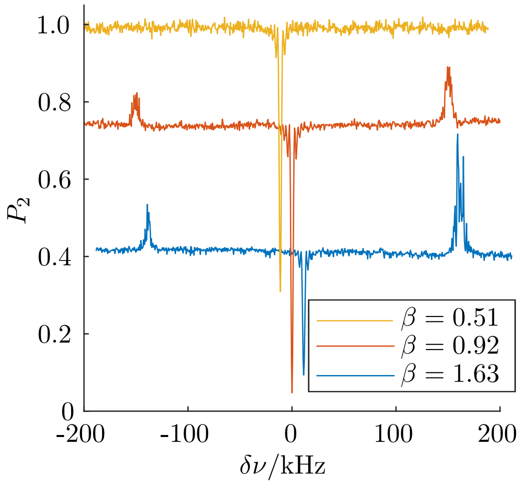

Figure 8.7: Optical pumping into a single state, $|F = 2, m_F = 0\rangle$. Increasing the EOM modulation depth, β, causes more efficient pumping. Curves horizontally offset for clarity.

At this point, the experiment was handed over to Alan Bregazzi (Strathclyde) and Etienne Batori (Neuchâtel) as my experimental time was complete. The results in the rest of this chapter were taken by them, while I wrote this thesis, and not by myself. The following data is shown to illustrate the results of my designing and building the clock experiment.

To validate the rest of the clock sequence and remove complexity, initial data was taken by dropping the atoms rather than launching them. Only a single MOT beam was required, but the Ramsey time was restricted to ≤ 20 ms. Once the clock has been validated and somewhat optimised, it is planned to re-introduce the GMOT launch for longer Ramsey times. Various measurements are taken using the fractional population, P_2 of the $F = 2$ states at the end of the chosen experimental sequence. The detection of this population is not fully characterised, and fractions below 0 or above 1 are sometimes seen, but relative changes in the detected fraction are reliable.

An on-resonance microwave pulse for 167 µs is used to excite atoms from $|F = 1, m_F = 0\rangle$ to $|F = 2, m_F = 0\rangle$. When varying the power of the pulse, Rabi oscillations are observed in P_2, as shown in Fig. 8.6. The decay of the oscillations is believed to be due to variations of the microwave power and orientation across the volume which the atoms occupy. Atoms within the cloud see different Rabi frequencies, causing them to gradually decohere. The amount of decoherence is higher than expected, corresponding to around 10% variations in microwave amplitude over the mm scale MOT. This data is very early and further investigation is required.

By changing the frequency of the MW pulse, the $m_F \neq 0$ states can also be excited. The distribution of atoms between states can then be observed, see Fig. 8.7, which is useful for evaluating the state-selection process. During the state-selection step, the EOM diverts power from the carrier frequency (exciting $F = 2 \rightarrow F' = 2$) to various frequency sidebands, one of which acts as a repump on the $F = 1 \rightarrow F' = 2$ transition. The amplitude of the EOM drive signal determines the modulation depth, β, of the EOM and how much power is diverted to each sideband. When the EOM is driven with a $\beta = 0.51$ the carrier and repump have 87% and 7% of the total output power respectively. There is insufficient repumping during the state selection and atoms will accumulate in the $F = 1$ states. P_2 has a low baseline, a negative peak when the $m_F = 0$ transition is excited, and positive peaks when the $m_F = \pm 1$ transitions are excited because atoms are excited out of the $F = 1, m_F = \pm 1$ states. When β is increased to 1.61 then the carrier and repump have 19% and 33% of the total power respectively and the $F = 1$ states are much more efficiently depopulated. P_2 has a much higher baseline, and only the central $m_F = 0$ peak is observed, showing that the atoms have been pumped into $|F = 2, m_F = 0\rangle$ as desired.

By adding a second MW pulse, a Ramsey sequence is formed and Ramsey fringes are observed in P_2, see Fig. 8.8. The fringes have the expected linewidth of $1/2T$ and are enclosed in an envelope determined by the pulse shape. For square pulses of length τ, as described in equation (3.8), the envelope has zeros where $\sin(\Omega'\tau/2) = 0$ which first occurs for $\delta\nu = \sqrt{15/16}/\tau$. The Ramsey fringes were taken using $\tau = 167\,\mu\text{s}$ and so the envelope zero occurs at $\pm 5.8\,\text{kHz}$. At $T = 10\,\text{ms}$ the fringes have an SNR of around 95 giving a single-shot stability of 2×10^{-11}. This is already equal to the best SNR we have ever observed in a CPT clock at Strathclyde,[103] highlighting the benefits of being able to scatter a much larger number of photons from each atom. The SNR is

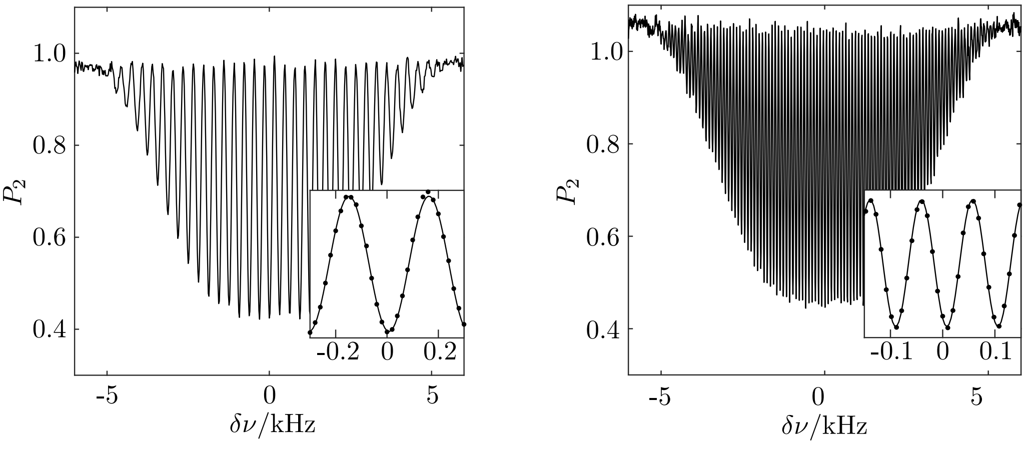

Figure 8.8: Ramsey fringes observed using the microwave cavity. Ramsey times of 3 ms and 10 ms were used (left and right respectively). Inset shows the central fringe

still limited by experimental noise, as the QPN and Dick effect are well below the experimental noise floor. The QPN would limit the SNR to over 1000, and the Dick effect limits the stability to $< 5 \times 10^{-12}/\sqrt{\tau}$. The Dick effect was calculated numerically from the noise specification of our reference oscillator, and is particularly bad because the duty cycle of the experiment is currently very poor. The system is running out of Rb and requires a long $\sim 3\,\mathrm{s}$ load time, while the Ramsey time is only $10\,\mathrm{ms}$.

The cavity clock experiment will continue, with detailed characterisation of the state-selection and detection systems, measurement of the clock short-term stability and re-integration of the fountain technique.

Chapter 9

Performance prospects

This chapter will consider the potential performance of the compact cavity-clock design. As the clock is at an early stage of engineering, this chapter aims to identify the most likely factors that will limit the clock, as these areas will require the most consideration in future designs. The chapter is largely based upon a study of clock accuracy budgets and potential clock shift mechanism which was undertaken during the lockdowns of the 2020 COVID-19 pandemic. During this time no experimental work was possible.

One nice feature of the cavity-clock design is that the experimental procedure is designed to be very similar, in principle, to those used for high-performance atomic fountain clocks such as the Cs fountain clocks. Therefore, much of the theoretical work considering clock performance limits has already been completed, and there are detailed accuracy budgets available.[20–23,104,105] These provide a good starting point to calculate what disadvantages are involved when moving to a compact system.

When considering the potential performance of the clock, a set of operational parameters should be selected. I have chosen the parameters shown in Table 9.1. These parameters are intended as an optimistic, but realistic, choice using the techniques described previously. The results in Chapters 7 and 8 suggest it is possible to detect this number of atoms at 100 ms. Larger holes may need to be added to the MW cavity to observe this number of atoms. Fortunately, due to the cavity geometry, the holes can easily be extended in a vertical direction (between the cavity electrodes) which is the MOTs longest dimension. Resolving the unknown issue which causes reduced cooling in the launch direction would make it easy to achieve this atom number. The pulse time, τ, is chosen to be relatively long as this decreases the effect of some shifts, see Sections 9.2.4 to 9.2.6.

Table 9.1: Key parameters which a future clock might achieve, used to predict the performance that could be realised.

Symbol	Parameter	Value
T	Ramsey time	95 ms
τ	MW $\pi/2$ pulse length	2.5 ms
T_c	Experiment cycle time	250 ms
N	Detected atom number	1×10^6

9.1 Short-term stability

For most atomic clocks, the short-term stability will consist of contributions from QPN (see Chapter 3), the Dick effect and experimental noise. As we are considering theoretical limits, we will assume that only the QPN and Dick effect limit the short-term stability. The real experiment will also have technical/electronic noise and photon shot noise. The photon shot noise limits the SNR to $\sqrt{N_{\text{photon}}}$, similar to the QPN which limits the SNR to $\sqrt{N_{\text{atom}}}$. We can say the QPN will dominate the photon noise provided that many photons are detected from each atom, which is the case for our detection techniques (unlike the CPT detection technique of Section 8.1). Additionally, detection schemes similar to ours have seen QPN-limited performance in practice, so this assumption is reasonable.[101,106]

The Dick effect is due to fluctuations of the LO frequency which occur during the experimental dead-time. These fluctuations cannot be detected by the atomic reference, and thus will never be corrected.[107] The Dick effect is reduced by having a low noise LO, and a high duty-cycle experiment. As we aim to have a compact clock, our LO must be based on a quartz crystal. Hydrogen masers and cryogenic sapphire oscillators both offer improved noise but are much larger and more expensive.[108,109] The Dick effect in Cs fountains using a quartz LO has been reduced to below $7 \times 10^{-14}/\sqrt{\tau}$, and to 0.9 to $1.7 \times 10^{-13}/\sqrt{\tau}$ in integrating sphere clocks.[110–113] Therefore a comparable performance could be expected for our clock.

The QPN can be estimated using equation (3.34). With the parameters given earlier, this gives a stability of $1.1 \times 10^{-13}/\sqrt{\tau}$. A total short-term performance below $2 \times 10^{-13}/\sqrt{\tau}$ is achievable. This is highly competitive with other compact atomic clocks.[27,28,114,115]

9.2 Accuracy

The accuracy of an atomic clock is one of the most challenging performance metrics to calculate as a large number of factors could affect it. However the accuracy is also highly relevant for giving the ultimate performance of the clock, as it presents an upper bound for the long-term stability. To evaluate the potential accuracy of our clock, we can look at the error budgets for the existing full-scale Rb and Cs fountains, and see which terms would be most important in our clock. The error budgets for a number of these fountain clocks are summarised in Table 9.2. The largest clock shifts are usually the second-order Zeeman shift, the blackbody radiation shift (unless cooled by liquid nitrogen), and a relativistic shift due to gravitational potential. However these shifts are normally well characterised, and the shifts with most uncertainty are more varied, with contributions from the microwave spectrum, distributed cavity phase, collisional shifts and more.

I will now consider each type of shift in turn, and how it would affect our compact fountain clock. I have ordered the shifts by the likely uncertainty of the shift.

9.2.1 Second-order Zeeman shift

The $m_F = 0$ sublevels have a second-order magnetic sensitivity of $\delta\nu = \beta B^2$, with $\beta = 575\,\mathrm{Hz\,G^{-2}}$ for ^{87}Rb.[19] There will be a bias field, B_0 present to break the degeneracy of magnetic sublevels, so the shift can be approximated as $\delta f = 2\beta B_0 \delta B$, in the worst case where the magnetic field fluctuations, δB, are aligned with the bias field. Operating with a small magnetic bias field is therefore advantageous, but the field direction must be well aligned with the cavity axis to ensure the correct π transition is driven. This means that magnetic shielding is required to prevent any magnetic fluctuations from influencing the clock. The fountain clocks have typical bias fields of 1 mG to 4 mG. Because the current experiment is unshielded, it must operate at large bias fields $> 100\,\mathrm{mG}$ and this shift would be extremely limiting to long-term performance. If we used a bias field of 5 mG, then our clock sensitivity would be $8 \times 10^{-13}\,\mathrm{mG^{-1}}$. Field fluctuations at the 0.1% level would then give clock stability of 4×10^{-15}. This level of stability is reasonable for a DC magnetic field, and the large fountain clocks generally specify a magnetic shift uncertainty below 1×10^{-16} using multi-layer shields. However, one problem with our particular experiment is that a field gradient of $\approx 15\,\mathrm{G\,cm^{-1}}$ is required for the MOT, in the same region as our Ramsey sequence and only a few ms before that sequence starts. Switching the MOT coils off will cause eddy currents in nearby conductors which must decay by many $1/e$ lifetimes to reach the required field stability. The vacuum chamber, MW cavity and MOT coils would have to be jointly designed, giving only high resistance paths for eddy currents to ensure they would rapidly decay.[117] The MOT could also be held in the optical molasses before launching for a few extra ms or even tens of ms without

Table 9.2: Error budgets for existing fountain clocks

Effect	$\delta\nu/\nu \times 10^{-16}$											
	NPL-CsF2[20]		NIST-F2[21]		ITCs-F2[22]		PTB-CSF2[23]		SYRTE FO2-Rb[104]		NPL Rb[105]	
2nd-Order Zeeman	6160.0	0.8	2860.6	0.2	1070.4	0.8	998.5	0.6	3465.5	0.7	857.6	1.0
Blackbody shift	223.0	1.1	0.9	< 0.1	1.5	0.1	166.0	0.6	124.2	1.4	133.5	0.6
Resonant AC Stark shift	0.0	< 0.1	0.0	< 0.1			0.0	< 0.1				
Relativistic shift[1]	?[1]	0.5	1798.7	0.3	260.4	0.1	856.7	< 0.1	?[1]	< 0.1	12.6	0.3
Cold collisions[1]	100.0	0.4		< 0.2		2.2	3.2	6.2	8.0[2]	2.1[2]		
Background collisions	0.0	1.0	0.0	< 0.1		0.5	0.0	0.5	0.0	1.0	0.0	0.7
Rabi,Ramsey pulling	0.0	< 0.1	0.0	< 0.1			0.0	< 0.1	0	0.1		
Distributed cavity phase	0.0	1.1	0.0	0.4			0.0	1.5	0.4	1.0	0	1.3
Dynamic cavity phase	0.0	< 0.1				0.2						
Cavity pulling	0.0	0.2	0.2	0.2		0.3			8.0[2]	2.1[2]		
Microwave lensing/power	0.6	1.5	0.1	1.4[4]			0.0	4.0	0.7	0.7	0.9	0.5
Microwave leakage	0.5	0.6			1.2	1.5	0.0	1.0	0	0.5	0.0	1.0
MW spectral impurities	0.0	< 0.1	0.0	0.5			0.0	1.7				
DC Stark shift	0.0	< 0.1	0.1	< 0.1								
Bloch-Siegert	0	$\ll 0.1$	0	$\ll 0.1$								
Electronic/Software			0	< 0.1			0.0	1.0				
Total u_{B}		2.4		1.6		1.9		8.0		3.2		2.4

[1] The magnitude of some shifts is not reported, only the associated error.

[2] The FO2-Rb clock states only a combined error from cavity pulling and cold collisions.

[3] The NIST-F2 microwave lensing shift was disputed in [116] but this was not addressed by NIST. I have used the updated value, as was used in the ITCs-F2 evaluation.

significant detriment, allowing the transient magnetic field to dissipate.

9.2.2 Cavity phase shifts

Cavity phase shifts occur when the atoms experience a different microwave phase in the first and second pulses, due to their motion within the MW cavity. An ideal MW cavity would have equal MW phase at all points, but a real cavity will have small phase gradients. The phase shift will transfer directly into a frequency shift of $\frac{\delta\nu}{\nu_0} = \frac{\delta\phi}{2\pi\nu_0 T}$. A 1 mrad phase difference will cause a frequency shift of 2.3×10^{-13}. The cavity phase shift uncertainty is often significant in full-scale fountain clocks, despite their reduced sensitivity due to longer Ramsey times. In a large fountain, the atoms will pass through the entire cavity volume, and will enter regions with large phase variations. Our atoms should stay near the cavity centre where the phase is more uniform. However, our cavity has a much reduced quality factor, $Q = 185$. This leads to larger phase gradients because the forwards and reverse travelling waves are of unequal amplitude. Modelling the cavity as a one dimensional waveguide (a Fabry-Perot cavity) suggests that the reverse travelling wave has relative amplitude of $(1 - \frac{\pi}{2Q})$. Therefore the phase in the cavity varies with the displacement from the centre, x, as $\phi \approx \frac{\pi}{4Q}\tan(\frac{2\pi x}{\lambda})$. This leads to a frequency shift of 1.5×10^{-13} per 1 mm that the MOT has moved in between the first and second pulses. By changing the pulse timing we can apply the pulses when the atoms are at different heights, allowing us to map out the cavity phase and test this assumption. A higher Q-factor cavity would allow us to reduce the cavity phase shift, but this may or may not be possible. The current cavity is not optimised for high Q, but the Q decreased from ≈ 500 to 185 when the silicon diffraction grating was added. This is believed to be due to dielectric losses in the silicon, and may be hard to mitigate.

Separately, there can be dynamic cavity phase shifts, where the phase of the field driving the cavity changes between the first and second pulse. This can be caused by phase transients in the switch which produces the Rabi pulses.[118] Fortunately, there exist switches designed for other fountain clocks which exhibit extremely small phase transients below 2 µrad.[118–120]

9.2.3 Collisional shifts

There is a clock shift due to collisions between the atoms in the cloud. This takes the form of

$$\frac{\Delta\nu}{\nu_0} = a_1 n_1 + a_2 n_2 \approx a n \ , \tag{9.1}$$

for atoms that are sufficiently cold (below about 10 µK for Rb^{87}). $n_{1,2}$ is the density of atoms in each of the two hyperfine states. The approximation is valid when the fraction of atoms in each state is constant, e.g. during the Ramsey time when half the atoms are in each state. For $^{87}\mathrm{Rb}$ the value of a is approximately $5(1) \times 10^{-23}\,\mathrm{cm}^3$.[121,122] In Cs fountains this effect is > 30 times larger and is the focus of many design choices. We can expect a shift around 2.5×10^{-13} for 10^7 atoms in a cloud with $1/e^2$ radius of 1 mm. The total shift is averaged over the interrogation period and would be less than the initial shift, which occurs when the density is highest. Fortunately, the shift decreases strongly as the atom cloud size is increased. Expanding the MOT to $1/e^2$ radius of 3 mm before starting the interrogation results in an initial shift below 1×10^{-14}. This could be achieved by briefly turning off the cooling light to allow thermal expansion. Only atoms that are within the region covered by the inner coaxial beam can be properly launched.

There is also a shift due to collisions with background gas. These collisions are quite likely to eject atoms from the cloud due to the high speed of incidence by the background gas. However, 'glancing' collisions can cause small disturbances to the atomic coherence which then cause a clock shift. Increased background pressure will cause both a loss of atoms and a clock shift. According to [123], the fractional atom loss, ΔA, can be used to bound the clock shift:

$$\frac{\Delta\nu}{\nu_0} \leq \frac{\Delta A}{13.8\pi T\nu_0}[h\nu_0(E_1^{-1} + (E_1 + E_p)^{-1})] < \Delta A \frac{4\hbar}{13.8TE_1} \ , \tag{9.2}$$

where $E_{1,p}$ are the 'lowest resonant excitation energy' of Rb and the colliding species respectively. This is the D_1 energy of 1.59 eV for Rb. The shift is therefore bounded by $1.2 \times 10^{-15}\Delta A$.

9.2.4 Microwave leakage

In large fountain clocks, leakage of the microwaves outside of their cavities can cause shifts. For a compact fountain, the equivalent process is incomplete extinction of the microwave signal in-between Ramsey pulses. The remnant signal can drive the clock transition and alter the atomic phase. If the remnant microwave signal is in-phase with the original signal, it acts as a slightly increased Ramsey pulse area, and does not induce a clock shift.[124] However, a quadrature leakage component (out of phase by $\pi/2$) can cause shifts. Similarly, if the microwave leakage is symmetric about the centre of the Ramsey time, it does not induce a shift. The maximum phase shift which can be caused is $\delta\phi = \Omega_L t_L$ where Ω_L is the Rabi frequency of the (quadrature) leakage field, and t_L is the time for which the field is applied. This leads to an upper bound on the microwave leakage shift of $\frac{\delta\nu}{\nu_0} = \frac{X}{4\nu_0\tau}$ where X is the amplitude extinction ratio for the quadrature component of the leakage field (compared to during the microwave pulses).[125] An extinction ratio of −100 dB (1×10^{-5} in amplitude) leads to a maximum shift of 1.5×10^{-13}. It is likely that shifts would be below this level due to the symmetry constraints. Increasing τ would also decrease the shift and, in the limit, a Rabi sequence could be used instead of a Ramsey sequence so there could be no leakage between the pulses.

9.2.5 Unwanted magnetic transitions

In a MW fountain, it is important that only a single transition is excited among the ground state manifold, the $|F=1, m_F=0\rangle \rightarrow |F=2, m_F=0\rangle$ transition. If transitions occur between any other states it could lead to a frequency shift. Several different kinds of transition are possible. In general, to cause a clock shift an asymmetry is required between the $m_F > 0$ and $m_F < 0$ states.

Rabi and Ramsey pulling are caused by unwanted microwave transitions involving $m_F \neq 0$ states in the $F=1 \rightarrow F=2$ transition. Rabi pulling is due to π transitions. Atoms that are in the $m_F = \pm 1$ states will still interact with the microwave field, but their transition is detuned from resonance by the Zeeman splitting of $\nu_\mathrm{Z} = \beta_\mathrm{Z} B_0$ where $\beta_Z = 1.4\,\mathrm{MHz\,G^{-1}}$ for ^{87}Rb. For a bias field of $5\,\mathrm{mG}$, the Zeeman splitting is $\nu_\mathrm{Z} = 7\,\mathrm{kHz}$. According to equation (3.8), the off-resonant states will still be excited but the observed oscillations will have their amplitude decreased by the multiplying factor

$$\left(\frac{\Omega}{\Omega'}\right)^2 < \left(\frac{\Omega}{\Delta}\right)^2 = \left(\frac{\pi/2\tau}{2\pi\nu_\mathrm{Z}}\right)^2 = \frac{1}{\left(4\nu_\mathrm{Z}\tau\right)^2} \ , \tag{9.3}$$

so a large Zeeman splitting will help suppress the Ramsey pulling. We can bound the Rabi pulling shift by assuming that the oscillations due to the $m_F = \pm 1$ state are perfectly out of phase with the actual Ramsey fringes. This causes a maximum phase shift of

$$\delta\phi = \frac{f}{\left(4\nu_\mathrm{Z}\tau\right)^2} \ , \tag{9.4}$$

where f is the imbalance between $m_F = \pm 1$ states as a fraction of the detected atoms. The maximum frequency shift is

$$\left|\frac{\delta\nu}{\nu_0}\right| \leq \frac{1}{2\pi\nu_0 T}\frac{f}{\left(4\nu_\mathrm{Z}\tau\right)^2} \ . \tag{9.5}$$

This gives a shift below 2.5×10^{-15} for 5% imbalance.

Ramsey pulling involves the σ transitions between states with $\Delta m_F = \pm 1$. The effect is normally quite similar to the Rabi pulling effect, except that the Zeeman splitting is only half as big, and the transition strength is reduced because the σ transitions are only driven by the component of the field which is misaligned with the magnetic field axis. This tends to make the Ramsey pulling smaller than the Rabi pulling because the misalignment angle is small. However, in some situations the Ramsey pulling can be significantly enhanced. If a coherence is present between the atoms in the $F=1, m_F=+1$ state and those in the desired $F=1, m_F=0$ state, then a much larger clock shift may be observed[126]. This shift only decreases as $1/(\nu_\mathrm{Z}\tau)$ instead of $1/(\nu_\mathrm{Z}\tau)^2$, and is difficult to calculate (numerical modelling has previously been used). However, I think it is notable that a maximum shift of

$$\left|\frac{\delta\nu}{\nu_0}\right| \leq \frac{1}{2\pi\nu_0 T}\frac{f}{\left(4\nu_\mathrm{Z}\tau\right)} \ . \tag{9.6}$$

is approximately consistent with the maximum shifts simulated in [126], and I think this is a reasonable upper bound.

The initial coherence would be introduced by the state-selection process, so better design should prevent these coherences from forming in the first place. The coherences can be manipulated to vary in sign, allowing easier detection for verification of the shift suppression. Adding a small gradient field after state selection can cause any coherences present to oscillate at different frequencies due to varied Zeeman splitting, averaging to zero. However, this is difficult for our system as discussed in Section 9.2.1.

Majorana transitions ($\Delta F = 0, \Delta m_F \neq 0$) occur when the magnetic field changes direction rapidly (non-adiabatically) so that the m_F number is not conserved. This typically happens near a zero of the magnetic field. One fountain clock did observe a significant shift due to this effect in combination with unwanted σ component in their state-selection beam which pumped $m_F > 0$ states.[127] This shouldn't be a problem for our clock as the magnetic field is near-uniform after the state-selection. Any eddy currents must have settled by this point, or the 2nd order Zeeman shift will be a bigger problem than Majorana transitions!

9.2.6 Microwave spectrum impurities

If the microwave signal fed to the cavity is not a pure tone, then the additional frequency components may affect the atomic state. This is particularly a problem if the microwave signal is disturbed in a way synchronous to the experimental sequence. This problem has been analysed[128] and the resulting frequency shift is

$$\left|\frac{\Delta\nu}{\nu_0}\right| <\approx \frac{1}{4T\tau\nu_0\Delta}\frac{\Omega_1^2}{\Omega_0^2} \ , \tag{9.7}$$

where Δ is the modulation frequency and $\Omega_{0,1}$ are the Rabi frequencies of the main MW signal and its modulation sideband. For high modulation frequencies, the shifts are lower than this expression due to phase averaging within each microwave pulse. Even very small amounts of modulation may cause significant shifts: a $-100\,\mathrm{dBc}$ sideband can cause shifts up to 5×10^{-15}. For this reason, many primary fountains are run at a frequency such that 100 Hz noise has opposite phases on each experiment cycle, or have a randomised cycle length to prevent coherent modulation.[23]

There is also a second-order effect corresponding to sidebands which are incoherent with the experiment. The general form of the shift is quite complicated,[129,130] but if $\Omega_0\tau \approx \pi/2$ then the shift can be bounded by

$$\left|\frac{\delta\nu}{\nu_0}\right| \leq \frac{1}{2\pi\nu_0 T}\frac{\Omega_1^2}{\Omega_0^2}\alpha \ , \tag{9.8}$$

where

$$\alpha = \begin{cases} \frac{\pi}{2\Delta\tau}\left(\frac{\pi}{2}-1\right) & \frac{\Delta}{2\pi} \geq \frac{1}{8\tau} \\ \frac{2\Delta\tau}{\pi}\left(1+\frac{4\Delta\tau}{\pi}\right) & \frac{\Delta}{2\pi} < \frac{1}{8\tau} \end{cases} , \tag{9.9}$$

with the crossover, $1/8\tau = 50\,\mathrm{Hz}$. This limit was calculated using worst case assumptions in the formula provided by [130].

In general microwave spectrum impurities do not appear to be a large factor in most fountain clock accuracy budgets. Our errors will be somewhat worse because of a lower Ramsey time, but are unlikely to be a limiting factor provided that care is taken to avoid modulation coherent with the experimental sequence.

9.2.7 Microwave lensing

Microwave lensing occurs because the microwave field exerts a dipole force on the atomic cloud. The effect causes a slight focusing of one microwave state, and a defocusing of the other, depending on the frequency detuning. This leads to a difference in detection efficiencies, and an observed frequency shift. The shift is of order

$$\frac{\Delta\nu}{\nu_0} = \frac{\hbar\pi\nu_0}{mc^2} = 1.7 \times 10^{-16}, \tag{9.10}$$

so we can probably ignore it. The shift does seem to be difficult to evaluate precisely as it is a significant contributor in some fountain accuracy budgets.

9.2.8 Cavity pulling

A mismatch between the cavity resonance frequency and the atomic resonance causes frequency shifts. A first-order shift occurs due to interaction of the field emitted by the atoms with the cavity field causing a phase shift in the total field. A second-order shift occurs due to the change in field amplitude as the cavity is excited off-resonance. Both of these shifts have a dispersive shape, with no shift when on the cavity resonance, and maximum shift when detuned by half the cavity linewidth. The worst-case shift is taken, though this can be significantly reduced by accurately tuning the cavity resonance.

The first order cavity pulling shift is given by [131] as

$$\frac{\Delta\nu}{\nu_0} \leq \frac{\mu_0\mu_B^2 NQ}{2\pi^2\hbar V\nu_0}\frac{\tau}{T}S , \tag{9.11}$$

where V is the cavity mode volume and S is a numerical constant of order 1. This term should be negligible for us because the Q factor is low. With $N = 10^7$, $Q = 185$, $V = 12\,\mathrm{mL}$ the shift is below 3×10^{-17}. The background Rb vapour should not contribute because it has a randomised state, and the emitted field has no defined phase.

The second order cavity pulling shift is given by [22] as

$$\left|\frac{\Delta\nu}{\nu_0}\right| \leq \Gamma\frac{8}{\pi^2\nu_0}\left(\frac{Q}{Q_\mathrm{A}}\right)\Omega\tau\cot(\Omega\tau) \approx 2\pi\epsilon Q\left(\frac{1}{2\nu_0 T}\right)^2 , \tag{9.12}$$

where $\Gamma = \pi\nu_0/Q$ is the cavity resonance half-width, $Q_\mathrm{A} = 1/2\nu_0 T$ is the atomic line quality factor and ϵ is the fractional error in the microwave pulse area. $\epsilon = 0.05$, $Q = 150$ gives a shift of 3×10^{-17}.

In summary, both cavity pulling shifts should be negligible because of our low cavity quality factor. Even with a 100× improved quality factor,these shifts would be unlikely to dominate.

9.2.9 Blackbody shift

The blackbody radiation shift is an AC Stark shift from the interaction of the atoms with the surrounding thermal radiation. This is a relatively weak effect because the thermal radiation is not resonant with a transition. The shift is given by:

$$\frac{\Delta\nu}{\nu} = \beta\left(\frac{T}{T_0}\right)^4\left[1 + \epsilon \times \left(\frac{T}{T_0}\right)^2\right] , \tag{9.13}$$

with $T_0 = 300\,\mathrm{K}$. A breakdown of current experimental and theoretical values of β and ϵ is given in [132], which uses the best experimental value of $\beta = -1.2727(14) \times 10^{-14}$ and a theoretical value of $\epsilon = 0.011(1)$.[133] This gives a shift of $-1.2867(19) \times 10^{-14}$ at 300 K with a temperature coefficient of $1.72(2) \times 10^{-16}\,\mathrm{K}^{-1}$. These errors are negligible for us, provided minimal (10 K) temperature control measures are in place.

9.2.10 Negligible shifts

There are various other shifts which occur in fountain clocks which will be negligible in our clock.

Resonant AC Stark shifts can be removed entirely by the use of shutters on all laser beams.

Relativistic shifts have a magnitude of around 1.1×10^{-16} per metre of height above the geoid, which can easily be calculated to within a few metre accuracy.

The DC Stark shift due to static electric fields can be neglected because the microwave cavity has conductive surfaces (including the grating coating). Although patch fields may exist, these will be negligible. A 1 mm patch at a whole 1 V generates a field of $\approx 40\,\mathrm{V\,m^{-1}}$ at a distance of 5 mm, causing a clock shift[134] of only 3×10^{-17}.

The Bloch-Siegert shift is the AC Stark shift of the counter-rotating term from the rotating-wave approximation. This only occurs during the MW pulse, and gives a total shift of approximately

$$\frac{\delta\nu}{\nu_0} \approx \frac{1}{32\tau T\nu_0^2} = 3 \times 10^{-18}. \tag{9.14}$$

9.2.11 Summary

Most of the shifts in our clock are similar to those in full-scale fountain clocks. Many of the shifts scale as $1/T$ and so, in principle, our clock should be capable of reaching approximately one order of magnitude worse accuracy than a full-scale fountain, a few 10^{-15}. However, the nature of our experiment also has additional factors which would make this challenging:

- The low Q microwave cavity will cause large cavity phase shifts, at the 10^{-14} level.
- Switching the quadrupole coils will induce eddy currents, causing second-order Zeeman shifts.
- Low optical access necessitates a relatively small, dense atom cloud, allowing significant collisional shifts, despite Rb being much less sensitive to them than Cs.
- The atoms remain in the microwave cavity during the Ramsey time, so are more sensitive to microwave leakage due to incomplete extinction of the microwave field.

These effects are highlighted in Table 9.3, which summarises all the expected shifts. The worst shifts all have around 10^{-13} upper bounds. The shifts can all be significantly reduced by taking appropriate measures, but detailed evaluation is very hard precisely because the shifts are caused by the novel clock design and depend greatly on the exact implementation. In theory, each shift can be reduced to the 10^{-15} level, but the practicalities are very difficult. The distributed cavity phase shift is particularly nasty as it is probably influenced by the grating, which is integral to the new technique.

One interesting possibility is that a Rabi sequence could be used instead of a Ramsey sequence. Using a Rabi sequence changes a number of shifts, as shown in Table 9.3, but the main benefit is to drastically reduce the microwave leakage which would have occurred during the Ramsey time. Other potential mitigations for the worst shifts include

- Increasing the cavity Q or mapping the cavity phase gradients.
- Adding a waiting time between atom trapping and launching. This would allow magnetic transients to decay and the atom cloud to expand, reducing its density.
- High extinction, low phase transient microwave switches to prevent leakage

As might be expected, there may be compromises required between short-term and long-term performance. However, the cavity clock has competitive prospects in both areas.

Table 9.3: The scale of various clock shifts in a compact microwave fountain clock. Both a Ramsey and Rabi sequence are considered, as this effects some shifts.

Effect	$\delta\nu/\nu$	
	Ramsey	Rabi
2nd-Order Zeeman	$\approx 10^{-13}$	
Distributed cavity phase	$< 10^{-13}$	
Cold collisions	$< 10^{-13}$	
Microwave leakage	$< 10^{-13}$	$< 10^{-15}$
Rabi, Ramsey pulling	$< 10^{-14}$	$< 10^{-17}$
Cavity pulling	$< 10^{-16}$	$< 10^{-15}$
Dynamic cavity phase	$< 10^{-15}$	
Background collisions	$< 10^{-15}$	
MW spectral impurities	$< 10^{-15}$	
Microwave lensing/power	$< 10^{-15}$	
Blackbody shift	$< 10^{-15}$	
Relativistic shift	$< 10^{-15}$	
DC Stark shift	$< 10^{-17}$	
Bloch-Siegert	$< 10^{-18}$	
Resonant AC Stark shift	≈ 0	
Electronic/Software	≈ 0	

Chapter 10

Conclusion

Microwave cavity clocks have been a core part of the development of atomic frequency standards, from the first Cs beam clocks, to the Cs fountains that realise the SI second, to the space-borne clocks in GPS satellites. This thesis has followed my work developing a microwave cavity clock using an in-cavity GMOT to trap cold atoms.

I created an experimental sequence to fulfil the project goal of a clock based on cold-atoms trapped in a microwave cavity. The stages of atom trapping, cooling, state preparation, interrogation and detection were each considered and an apparatus was designed to achieve them all within the microwave cavity. By selection of compatible mechanisms at each stage, all the optical requirements could be satisfied by a single seed laser with appropriate frequency shifts and intensity control. The electrical, optical and vacuum systems were each commissioned to produce the working apparatus. A MOT was trapped within the system and atom trapping and cooling was verified.

The potential clock performance would be strongly improved by increasing the available interrogation time. To take advantage of this, methods for launching atoms from the grating MOT were investigated. It was realised that combining two concentric input beams on the grating would result in an overlapping beam configuration compatible with a moving molasses launch. Working backwards from this, a set of optics was developed that would create those beams with precise alignment and minimal diffraction. When applied in the experiment the new MOT beams did not compromise the atom trapping and cooling performance. Simply changing the relative frequency of the concentric beams allowed a launch to be realised. Experiments were conducted to optimise the launch and maximise the number of atoms seen at long interrogation times. An increased distribution of atomic velocities was observed in the launch direction as the launch speed increased, causing a reduction in the detectable atom number at long times. This effect was investigated but the mechanism was not identified.

To demonstrate the utility of this launch, an existing CPT clock experiment was used to measure clock signals under different launch conditions. The launch was shown to increase the clock performance substantially, with the optimal launch time being 40 ms resulting in 50% lower instability, but clock signals were clearly visible for launches up to 100 ms.

Having shown the viability of the launch technique using CPT, a microwave cavity was installed to complete the construction of the clock as designed. Preliminary data from using the cavity has shown all elements of the clock sequence are functional, with Rabi fringes, Ramsey fringes, state selection and detection. The evaluation of the experiment performance in this mode is ongoing, but the detection noise is already confirmed to be lower than using the CPT method.

To assess the potential performance of this new clock design, an investigation was undertaken into the probable sources of clock shifts. This highlighted that the new clock design is more susceptible to some kinds of shift (the second-order Zeeman effect and cavity phase shifts) compared to fountain-clock designs, and more resilient against others (cavity pulling). The clock design should be capable of competitive performance in both the short- and long-term, with potential short-term stability in the low $1 \times 10^{-13}\sqrt{\tau}$ and long-term accuracy in the 10^{-15} range. The accuracy in particular is quite difficult to evaluate at this early stage, and significant effort would probably be required to achieve this performance. The most problematic clock shift is likely to be caused by phase gradients in the microwave cavity, because the presence of the diffraction grating will disturb the cavity mode.

10.0.1 Future Work

This thesis is the first step in developing a new kind of clock. The results I have shown are promising, but further improvements to both the grating launch and clock system are clearly possible.

The CPT method which was used for preliminary work has been discarded due to its low SNR in comparison to using the microwave cavity, where only detection of the separate hyperfine levels is needed. However, a potential solution to this has been proposed which involves a coherent transformation of the CPT bright and dark states onto the two hyperfine levels using a Raman pulse, allowing clean detection. Work on this technique is ongoing for the existing CPT clock. A similar proposal would be to recreate a full Ramsey sequence using a Raman pulse instead of a microwave cavity. The advantage of these techniques is that it does not require the microwave cavity, and removes some shifts which are related to it. A supermirror or optical cavity could produce microwave phase gradients as good or better than any microwave cavity, and would not have the same eddy-current problems. However, the use of optical transitions will cause substantial AC Stark shifts instead, as is seen in other CPT clocks. This could be mitigated by operating far off-resonance (a THz or so), or using previously developed techniques for suppressing light shifts.

It is possible that the CPT detection method, which is currently greatly limited by SNR, could be improved by coherent transformation of the CPT bright and dark states onto the two hyperfine levels which can then be cleanly distinguished. However, it is likely that this would come with significant additional AC Stark shifts in comparison to the cavity clock due to the use of optical transitions in the CPT scheme. Research on this topic is being undertaken for the existing CPT clock.

The grating launch could be most improved significantly by identification and removal of the mechanism which is causing excess heating when launching at higher velocities. An immediate increase in the available atom number, maximum launch time, and clock stability would result. A different type of grating, using a chequerboard pattern, could potentially increase the trapped atom number several times as it offers a larger overlap volume from which to trap atoms.

The cavity clock system is currently undergoing rapid improvement and testing as it is still in an early development phase. The initial results are promising, but full performance and characterisation data are still being collected. The implementation of the GMOT launch within the cavity-clock system has yet to be attempted, and the combination of those two aspects of the project should bring significant improvements. The cavity clock system as a whole is still a prototype system, and many potential improvements are available. Implementation of magnetic shielding, as a minimum, would be required to meaningfully test the long-term stability that could be reached. The miniaturisation of the clock is clearly not complete as it takes up $2\,\mathrm{m}^2$ of optical table, but that was not the goal of this project. The microwave cavity itself, where all of the atomic physics occurs, has a volume of only 75 mL. The size of an eventual clock would likely be dominated by other components, and I believe is likely that magnetic shielding around the vacuum chamber, around the cavity would be the largest component. A total volume of a few litres seems a reasonable prospect.

Appendix A

Theo1 implementation

```
/*===============================================
* Calculates Theo1 of a given dataset.
* The dataset is assumed to be phase
* data, taken at unit time increments.
* Any linear component to the dataset is
* removed as a first step. Should be
* compiled with -O3 -ffast-math.
*
* X = input dataset
* N = number of elements in X
* T = Theo1 output
*
* Ben Lewis, University of Strathclyde
===============================================*/

/* Remove any linear part of the dataset*/
void removeLinear(double* X, int N){
    double midpoint = ((double) (N-1))/2;
    long double sum1 = 0;
    long double sum2 = 0;
    for(int i = 0; i < N; i++){
        sum1 += X[i];
        sum2 += X[i]*(i-midpoint);
```

```
    }
    double a = sum1/N;
    double b = sum2/N*12/((double)N*N-1);
    for(int i = 0; i < N; i++){
        X[i] -= a + b*(i-midpoint);
    }
}


/* The computational routine */
void Theo1(double *X, double *T, int N){
// Initialise arrays
int k_max = (N-1)/2;
double* C1 = new double[N];
double* C3 = new double[k_max+1];
double* C4 = new double[k_max*2];
// Preprocess by removing linear part
removeLinear(X,N);
// Calculate C1
double s=0;
for(int i = 0; i < N; i++){
  s += (X[i]*X[i]);
  C1[i] = s;
}
// Main loop
C3[0] = C1[N-1];
for(int k=1; k<=k_max; k++){
  //Calculate C2 values
  double C2_2k = 0;
  double C2_2k_1 = 0;
  for(int j = 0; j <= N-2*k-1; j++){
    C2_2k += (X[j]*X[j+2*k]);
    C2_2k_1 += (X[j]*X[j+2*k-1]);
  }
  C2_2k_1 += (X[N-2*k]*X[N-1]);
  // Update C3, C4 in place
```

```
    for(int v=0; v < k; v++){
        C3[v] -= (X[k-1-v]*X[k-1+v])
                 + (X[N-k+v]*X[N-k-v]);
    }
    for(int v = 1;v<=2*k-2;v++){
        C4[v-1] -= (X[2*k-1-v]*X[2*k-1])
                    + (X[2*k-2-v]*X[2*k-2])
                    + (X[N-2*k]*X[N-2*k+v])
                    + (X[N-2*k+1]*X[N-2*k+1+v]);
    }
    C3[k] = C2_2k;
    C4[2*k-2] = 2*C2_2k_1 - (X[0]*X[2*k-1])
                - (X[N-2*k]*X[N-1]);
    C4[2*k-1] = 2*C2_2k;
    //Calculate un-normalised T_k from C1-C4
    double T_k = 0;
    double A0 = C1[N-1] - C1[2*k-1]
                + C1[N-2*k-1] + 2*C2_2k;
    for(int v = 1;v<=k;v++){
        double A1 = A0 - C1[v-1]
                    + C1[N-1-v] - C1[2*k-v-1]
                    + C1[N-1-2*k+v];
        double A2 = C3[k-v] - C4[v-1]
                    - C4[2*k-v-1];
        T_k += (A1+2*A2)/v;
    }
    // Apply normalisation to get Theo1
    T[k-1] = (T_k/(3*(double)(N-2*k)*k*k));
}
//Release memory
delete[] C1;
delete[] C3;
delete[] C4;
}
```

Bibliography

1. Landels, J. G. Water-clocks and time measurement in classical antiquity. *Endeavour* **3,** 32–37. ISSN: 0160-9327. `https://www.sciencedirect.com/science/article/pii/0160932779900073` (2022) (Jan. 1, 1979).
2. Balmer, R. T. The Operation of Sand Clocks and Their Medieval Development. *Technology and Culture* **19,** 615–632. ISSN: 0040-165X. `https://www.jstor.org/stable/3103761` (2022) (1978).
3. Jespersen, J. & Fitz-Randolph, J. *From sundials to atomic clocks: Understanding time and frequency* ISBN: 0-16-050010-9 (National Institute of Standards and Technology, Dec. 1977).
4. Benedicto, J., Dinwiddy, S. E., Gatti, G., Lucas, R. & Lugert, M. *GALILEO: satellite system design and technology developments," European Space Agency* 2000.
5. Takamoto, M. *et al.* Test of general relativity by a pair of transportable optical lattice clocks. *Nature Photonics* **14,** 411–415. ISSN: 1749-4893. `https://www.nature.com/articles/s41566-020-0619-8` (2022) (July 2020).
6. Bothwell, T. *et al.* Resolving the gravitational redshift across a millimetre-scale atomic sample. *Nature* **602,** 420–424. ISSN: 1476-4687. `http://www.nature.com/articles/s41586-021-04349-7` (2022) (Feb. 2022).
7. *SI Brochure: The International System of Units (SI)* 9th edition (Bureau International des Poids et Mesures, 2019).
8. Bothwell, T. *et al.* JILA SrI optical lattice clock with uncertainty of 2.0×10^{-18}. *Metrologia* **56,** 065004. ISSN: 0026-1394. `https://doi.org/10.1088/1681-7575/ab4089` (2022) (Oct. 2019).
9. Zimmermann, M., Gohle, C., Holzwarth, R., Udem, T. & Hänsch, T. W. Optical clockwork with an offset-free difference-frequency comb: accuracy of sum- and difference-frequency generation. *Optics Letters* **29,** 310–312. ISSN: 1539-4794. `https://opg.optica.org/ol/abstract.cfm?uri=ol-29-3-310` (2022) (Feb. 1, 2004).

10. Marlow, B. L. S. & Scherer, D. R. A Review of Commercial and Emerging Atomic Frequency Standards. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **68,** 2007–2022. ISSN: 1525-8955 (June 2021).
11. Brewer, S. M. *et al.* $^{27}\mathrm{Al}^{+}$ Quantum-Logic Clock with a Systematic Uncertainty below 10^{-18}. *Physical Review Letters* **123,** 033201. `https://link.aps.org/doi/10.1103/PhysRevLett.123.033201` (2022) (July 15, 2019).
12. Hagemann, C. *et al.* Ultrastable laser with average fractional frequency drift rate below 5×10^{-19}/s. *Optics Letters* **39,** 5102–5105. ISSN: 1539-4794. `https://opg.optica.org/ol/abstract.cfm?uri=ol-39-17-5102` (2022) (Sept. 1, 2014).
13. Świerad, D. *et al.* Ultra-stable clock laser system development towards space applications. *Scientific Reports* **6,** 33973. ISSN: 2045-2322. `https://www.nature.com/articles/srep33973` (2022) (Sept. 26, 2016).
14. Newman, Z. L. *et al.* Architecture for the photonic integration of an optical atomic clock. *Optica* **6,** 680–685. ISSN: 2334-2536. `https://opg.optica.org/optica/abstract.cfm?uri=optica-6-5-680` (2022) (May 20, 2019).
15. Fortier, T. & Baumann, E. 20 years of developments in optical frequency comb technology and applications. *Communications Physics* **2,** 1–16. ISSN: 2399-3650. `https://www.nature.com/articles/s42005-019-0249-y` (2022) (Dec. 6, 2019).
16. Cao, J. *et al.* A compact, transportable single-ion optical clock with 7.8×10^{-17} systematic uncertainty. *Applied Physics B* **123,** 112. ISSN: 1432-0649. `https://doi.org/10.1007/s00340-017-6671-5` (2019) (Mar. 27, 2017).
17. Cao, J. *et al.* A compact, transportable optical clock with 1×10^{-17} uncertainty and its absolute frequency measurement. *Applied Physics Letters* **120,** 054003. ISSN: 0003-6951. `https://aip.scitation.org/doi/10.1063/5.0079432` (2022) (Jan. 31, 2022).
18. Stuhler, J. *et al.* Opticlock: Transportable and easy-to-operate optical single-ion clock. *Measurement: Sensors* **18,** 100264. ISSN: 2665-9174. `https://www.sciencedirect.com/science/article/pii/S2665917421002270` (2022) (Dec. 1, 2021).
19. Steck, D. A. *Rubidium 87 D Line Data* 2.1.5 (Jan. 13, 2015). `https://steck.us/alkalidata/`.
20. Li, R., Gibble, K. & Szymaniec, K. Improved accuracy of the NPL-CsF2 primary frequency standard: evaluation of distributed cavity phase and microwave lensing frequency shifts. *Metrologia* **48,** 283–289. ISSN: 0026-1394. `https://doi.org/10.1088%2F0026-1394%2F48%2F5%2F007` (2020) (Aug. 2011).
21. Heavner, T. P. *et al.* First accuracy evaluation of NIST-F2. *Metrologia* **51,** 174–182. ISSN: 0026-1394. `https://doi.org/10.1088%2F0026-1394%2F51%2F3%2F174` (2020) (May 2014).

22. Levi, F. *et al.* Accuracy evaluation of ITCsF2: a nitrogen cooled caesium fountain. *Metrologia* **51,** 270–284. ISSN: 0026-1394. `https://doi.org/10.1088%2F0026-1394%2F51%2F3%2F270` (2020) (May 2014).

23. Gerginov, V. *et al.* Uncertainty evaluation of the caesium fountain clock PTB-CSF2. *Metrologia* **47,** 65–79. ISSN: 0026-1394. `https://doi.org/10.1088%2F0026-1394%2F47%2F1%2F008` (2020) (Dec. 2009).

24. Huang, M., Klimcak, C. M. & Camparo, J. C. *Vapor-cell clock frequency and environmental pressure: Resonance-cell volume changes* in *2010 IEEE International Frequency Control Symposium* (June 2010), 208–211.

25. Moreno, W., Pellaton, M., Affolderbach, C. & Mileti, G. Barometric Effect in Vapor-Cell Atomic Clocks. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **65,** 1500–1503. ISSN: 1525-8955 (Aug. 2018).

26. *SA.45s CSAC datasheet* (Microsemi). `https://ww1.microchip.com/downloads/en/DeviceDoc/00002985.pdf` (2022).

27. *cRb-Clock Datasheet* (Spectradynamics Inc., 2019). `https://spectradynamics.com/product-sheets/cRb-Clock-2019.pdf` (2022).

28. *MuClock datasheet* (MuQuans). `https://www.muquans.com/wp-content/uploads/2019/03/muquans_muclock.pdf` (2022).

29. Salzenstein, P., Kuna, A., Sojdr, L. & Chauvin, J. Significant step in ultra-high stability quartz crystal oscillators. *Electronics Letters* **46,** 1433–1434. ISSN: 1350-911X. `https://digital-library-theiet-org.proxy.lib.strath.ac.uk/content/journals/10.1049/el.2010.1828` (2022) (Oct. 14, 2010).

30. *Wenzel BTULN datasheet* (Wenzel, 2022).

31. Time and Frequency from A to Z. *NIST.* `https://www.nist.gov/pml/time-and-frequency-division/popular-links/time-frequency-z` (2022) (May 12, 2010).

32. M. A. Lombardi, A. N. Novick & V. S. Zhang. *Characterizing the performance of GPS disciplined oscillators with respect to UTC(NIT)* in *Proceedings of the 2005 IEEE International Frequency Control Symposium and Exposition, 2005.* (Aug. 29, 2005), 677–684.

33. Guo, W. *et al.* Foundation and performance evaluation of real-time GNSS high-precision one-way timing system. *GPS Solutions* **23,** 23. ISSN: 1521-1886. `https://doi.org/10.1007/s10291-018-0811-1` (Jan. 2, 2019).

34. Foot, C. J. *Atomic Physics* 325 pp. ISBN: 978-0-19-850696-6 (Oxford University Press, Oxford, 2005).

35. Manzano, D. A short introduction to the Lindblad master equation. *AIP Advances* **10,** 025106. `https://aip-scitation-org.proxy.lib.strath.ac.uk/doi/10.1063/1.5115323` (2022) (Feb. 2020).
36. Loudon, R. *The Quantum Theory of Light* ISBN: 0-19-850177-3 (OUP, 2000).
37. Hilgevoord, J. The uncertainty principle for energy and time. *American Journal of Physics* **64,** 1451–1456. ISSN: 0002-9505. `https://aapt-scitation-org.proxy.lib.strath.ac.uk/doi/10.1119/1.18410` (2022) (Dec. 1996).
38. Dalibard, J. & Cohen-Tannoudji, C. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *JOSA B* **6,** 2023–2045. ISSN: 1520-8540. `https://www.osapublishing.org/josab/abstract.cfm?uri=josab-6-11-2023` (2019) (Nov. 1, 1989).
39. Bize, S. *et al.* High-accuracy measurement of the 87Rb ground-state hyperfine splitting in an atomic fountain. *Europhysics Letters* **45,** 558. ISSN: 0295-5075. `https://iopscience-iop-org.proxy.lib.strath.ac.uk/article/10.1209/epl/i1999-00203-9/meta` (2022) (Mar. 1, 1999).
40. McGilligan, J. P. *et al.* Grating chips for quantum technologies. *Scientific Reports* **7,** 384. ISSN: 2045-2322. `https://www.nature.com/articles/s41598-017-00254-0` (2018) (Mar. 24, 2017).
41. Wallace, C. D. *et al.* Measurements of temperature and spring constant in a magneto-optical trap. *JOSA B* **11,** 703–711. ISSN: 1520-8540. `https://opg.optica.org/josab/abstract.cfm?uri=josab-11-5-703` (2022) (May 1, 1994).
42. Kasevich, M. & Chu, S. Laser cooling below a photon recoil with three-level atoms. *Physical Review Letters* **69,** 1741–1744. `https://link.aps.org/doi/10.1103/PhysRevLett.69.1741` (2021) (Sept. 21, 1992).
43. Kerman, A. J., Vuletić, V., Chin, C. & Chu, S. Beyond Optical Molasses: 3D Raman Sideband Cooling of Atomic Cesium to High Phase-Space Density. *Physical Review Letters* **84,** 439–442. `https://link.aps.org/doi/10.1103/PhysRevLett.84.439` (2022) (Jan. 17, 2000).
44. Hall, J. L., Zhu, M. & Buch, P. Prospects for using laser-prepared atomic fountains for optical frequency standards applications. *JOSA B* **6,** 2194–2205. ISSN: 1520-8540. `https://www.osapublishing.org/josab/abstract.cfm?uri=josab-6-11-2194` (2020) (Nov. 1, 1989).
45. Riis, E., Weiss, D. S., Moler, K. A. & Chu, S. Atom funnel for the production of a slow, high-density atomic beam. *Physical Review Letters* **64,** 1658–1661. `https://link.aps.org/doi/10.1103/PhysRevLett.64.1658` (2022) (Apr. 2, 1990).

46. Vangeleyn, M. *Atom trapping in non-trivial geometris for micro-fabrication applications* PhD thesis (University of Strathclyde, 2011). `http://photonics.phys.strath.ac.uk/wp-content/uploads/2013/08/MVAngeleyn_Thesis_final.pdf` (2019).
47. McGilligan, J. P. *Micro-fabricated Diffractive Optics for Quantum Sensors and Atomic Clocks* PhD thesis (University of Strathclyde, Aug. 16, 2017). 255 pp.
48. Elvin, R. *Phase-sensitive optical spectroscopy of a laser-cooled, microwave atomic clock* PhD thesis (Thesis PhD – University of Strathclyde, 2020). `http://digitool.lib.strath.ac.uk/webclient/DeliveryManager?pid=33509&custom_att_2=simple_viewer` (2020).
49. Donley, E., Heavner, T., O'Brien, J., Jefferts, S. & Levi, F. *Laser cooling and launching performance in a (1,1,1)-geometry atomic fountain* in *Proceedings of the 2005 IEEE International Frequency Control Symposium and Exposition, 2005.* (Aug. 2005).
50. Lee, K. I., Kim, J. A., Noh, H. R. & Jhe, W. Single-beam atom trap in a pyramidal and conical hollow mirror. *Optics Letters* **21,** 1177–1179. ISSN: 1539-4794. `https://opg.optica.org/ol/abstract.cfm?uri=ol-21-15-1177` (2022) (Aug. 1, 1996).
51. Bowden, W. *et al.* A pyramid MOT with integrated optical cavities as a cold atom platform for an optical lattice clock. *Scientific Reports* **9,** 1–9. ISSN: 2045-2322. `https://www.nature.com/articles/s41598-019-48168-3` (2019) (Aug. 12, 2019).
52. Vangeleyn, M., Griffin, P. F., Riis, E. & Arnold, A. S. Laser cooling with a single laser beam and a planar diffractor. *Optics Letters* **35,** 3453–3455. ISSN: 1539-4794. `https://www.osapublishing.org/ol/abstract.cfm?uri=ol-35-20-3453` (2019) (Oct. 15, 2010).
53. Nshii, C. C. *et al.* A surface-patterned chip as a strong source of ultracold atoms for quantum technologies. *Nature Nanotechnology* **8,** 321–324. ISSN: 1748-3395. `https://www.nature.com/articles/nnano.2013.47` (2019) (May 2013).
54. Shimizu, F., Shimizu, K. & Takuma, H. Four-beam laser trap of neutral atoms. *Optics Letters* **16,** 339–341. ISSN: 1539-4794. `https://opg.optica.org/ol/abstract.cfm?uri=ol-16-5-339` (2022) (Mar. 1, 1991).
55. McGilligan, J. P., Griffin, P. F., Riis, E. & Arnold, A. S. Phase-space properties of magneto-optical traps utilising micro-fabricated gratings. *Optics Express* **23,** 8948–8959. ISSN: 1094-4087. `https://www.osapublishing.org/oe/abstract.cfm?uri=oe-23-7-8948` (2020) (Apr. 6, 2015).
56. Ramsey, N. F. A Molecular Beam Resonance Method with Separated Oscillating Fields. *Physical Review* **78,** 695–699. `https://link.aps.org/doi/10.1103/PhysRev.78.695` (2019) (June 15, 1950).
57. Mathur, B. S., Tang, H. & Happer, W. Light Shifts in the Alkali Atoms. *Physical Review* **171,** 11–19. `https://link.aps.org/doi/10.1103/PhysRev.171.11` (2022) (July 5, 1968).

58. Allan, D. Statistics of atomic frequency standards. *Proceedings of the IEEE* **54,** 221–230. ISSN: 1558-2256 (Feb. 1966).
59. Percival, D. B. & Walden, A. T. *Wavelet Methods for Time Series Analysis* (Cambridge University Press, 2000).
60. Vernotte, F., Lenczner, M., Bourgeois, P.-Y. & Rubiola, E. The Parabolic Variance (PVAR): A Wavelet Variance Based on the Least-Square Fit. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **63,** 611–623. ISSN: 1525-8955 (Apr. 2016).
61. Allan, D. & Barnes, J. *A Modified "Allan Variance" with Increased Oscillator Characterization Ability* in *Thirty Fifth Annual Frequency Control Symposium* (May 1981), 470–475.
62. Lewis, B. A Fast Algorithm for Calculation of Thêo1. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **67,** 2187–2190. ISSN: 1525-8955 (Oct. 2020).
63. Howe, D. & Peppler, T. *Very long-term frequency stability: estimation using a special-purpose statistic* in *IEEE International Frequency Control Sympposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum, 2003. Proceedings of the 2003* (IEEE, Tampa, Fl, USA, 2003), 233–238. ISBN: 978-0-7803-7688-5. `http://ieeexplore.ieee.org/document/1275095/` (2020).
64. Mcgee, J. A. & Howe, D. A. TheoH and Allan deviation as power-law noise estimators. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **54,** 448–452. ISSN: 0885-3010 (Feb. 2007).
65. Schwindt, P. D. D. *et al. Miniature trapped-ion frequency standard with 171Yb+* in *2015 Joint Conference of the IEEE International Frequency Control Symposium & the European Frequency and Time Forum* (Apr. 2015), 752–757.
66. Parker, T. E., Jefferts, S. R., Heavner, T. P. & Donley, E. A. Operation of the NIST-F1 caesium fountain primary frequency standard with a maser ensemble, including the impact of frequency transfer noise. *Metrologia* **42,** 423–430. ISSN: 0026-1394. `https://doi.org/10.1088/0026-1394/42/5/013` (2022) (Sept. 2005).
67. Neelam *et al. Precise Time Synchronization and Clock Comparison through a White Rabbit Network based optical fibre link data from CSIR-NPL , India* (Mar. 24, 2021). `https://data.mendeley.com/datasets/cbncxcynv6/1`.
68. Howe, D. A., Ashby, N., Lirette, D., Hati, A. & Nelson, C. *Limited live-time measurements of frequency spectra* in *2011 Joint Conference of the IEEE International Frequency Control and the European Frequency and Time Forum (FCS) Proceedings* (May 2011), 1–4.
69. Howe, D. A., McGee-Taylor, J. & Tassett, T. *ThêoH Bias-removal Method* in *2006 IEEE International Frequency Control Symposium and Exposition* (June 2006), 788–792.

70. Taylor, J. A. & Howe, D. A. Fast TheoBR: A method for long data set stability analysis. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **57,** 2091–2094. ISSN: 0885-3010 (Sept. 2010).

71. Li, M., Liao, C.-S. & Hwang, J.-K. Rapid Calculation of Time Deviation and Modified Allan Deviation for Characterizing Telecommunications Clock Stability. *International Journal of Computers and Applications* **30,** 92–96. ISSN: 1206-212X. https://doi.org/10.1080/1206212X.2008.11441886 (2022) (Jan. 1, 2008).

72. Esnault, F.-X. *et al.* Cold-atom double-Λ coherent population trapping clock. *Physical Review A* **88,** 042120. https://link.aps.org/doi/10.1103/PhysRevA.88.042120 (2019) (Oct. 31, 2013).

73. Breit, G. & Rabi, I. I. Measurement of Nuclear Spin. *Physical Review* **38,** 2082–2083. https://link.aps.org/doi/10.1103/PhysRev.38.2082.2 (2022) (Dec. 1, 1931).

74. Taichenachev, A. V., Yudin, V. I., Velichansky, V. L. & Zibrov, S. A. On the unique possibility of significantly increasing the contrast of dark resonances on the D1 line of 87Rb. *Journal of Experimental and Theoretical Physics Letters* **82,** 398–403. ISSN: 1090-6487. https://doi.org/10.1134/1.2142864 (2022) (Oct. 1, 2005).

75. Zibrov, S. A. *et al.* Coherent-population-trapping resonances with linearly polarized light for all-optical miniature atomic clocks. *Physical Review A* **81,** 013833. https://link.aps.org/doi/10.1103/PhysRevA.81.013833 (2022) (Jan. 28, 2010).

76. Jau, Y.-Y., Miron, E., Post, A. B., Kuzma, N. N. & Happer, W. Push-Pull Optical Pumping of Pure Superposition States. *Physical Review Letters* **93,** 160802. https://link.aps.org/doi/10.1103/PhysRevLett.93.160802 (2022) (Oct. 13, 2004).

77. Liu, X., Yudin, V. I., Taichenachev, A. V., Kitching, J. & Donley, E. A. High contrast dark resonances in a cold-atom clock probed with counterpropagating circularly polarized beams. *Applied Physics Letters* **111,** 224102. ISSN: 0003-6951. https://aip.scitation.org/doi/10.1063/1.5001179 (2022) (Nov. 27, 2017).

78. Elgin, J. D. *et al.* A cold-atom beam clock based on coherent population trapping. *Applied Physics Letters* **115,** 033503. ISSN: 0003-6951. https://aip.scitation.org/doi/10.1063/1.5087119 (2022) (July 15, 2019).

79. Hummon, M. T. *et al.* Photonic chip for laser stabilization to an atomic vapor with $10^{\text{\&\#x2212;11}}$ instability. *Optica* **5,** 443–449. ISSN: 2334-2536. https://opg.optica.org/optica/abstract.cfm?uri=optica-5-4-443 (2022) (Apr. 20, 2018).

80. Cheng, P. *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/IFCS)* (July 2017), 616–617.

81. Sosa, K., Oreggioni, J. & Failache, H. Miniaturized saturated absorption spectrometer. *Review of Scientific Instruments* **91,** 083101. ISSN: 0034-6748. `https://aip.scitation.org/doi/full/10.1063/1.5144484` (2022) (Aug. 2020).

82. Wynands, R. & Weyers, S. Atomic fountain clocks. *Metrologia* **42,** S64–S79. ISSN: 0026-1394. `https://doi.org/10.1088%2F0026-1394%2F42%2F3%2Fs08` (2019) (June 2005).

83. Jefferts, S., Drullinger, R. & DeMarchi, A. *NIST cesium fountain microwave cavities* in *Proceedings of the 1998 IEEE International Frequency Control Symposium (Cat. No.98CH36165)* (May 1998), 6–8.

84. Schroder, R., Hubner, U. & Griebsch, D. Design and realization of the microwave cavity in the PTB caesium atomic fountain clock CSF1. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **49,** 383–392. ISSN: 1525-8955 (Mar. 2002).

85. Stefanucci, C. *et al.* Compact microwave cavity for high performance rubidium frequency standards. *Review of Scientific Instruments* **83,** 104706. ISSN: 0034-6748. `https://aip.scitation.org/doi/10.1063/1.4759023` (2019) (Oct. 1, 2012).

86. Affolderbach, C. *et al.* Study of additive manufactured microwave cavities for pulsed optically pumped atomic clock applications. *Applied Physics Letters* **112,** 113502. ISSN: 0003-6951. `https://aip.scitation.org/doi/10.1063/1.5019444` (2019) (Mar. 12, 2018).

87. Pellaton, M. *et al.* 3D printed microwave cavity for atomic clock applications: proof of concept. *Electronics Letters* **54,** 691–693. ISSN: 0013-5194 (2018).

88. Almat, N., Gharavipour, M., Moreno, W., Affolderbach, C. & Mileti, G. *Long-Term Stability Analysis Towards* 10^{-14} *Level for a Highly Compact POP Rb Cell Atomic Clock* in *2019 Joint Conference of the IEEE International Frequency Control Symposium and European Frequency and Time Forum (EFTF/IFC)* (Apr. 2019), 1–2.

89. Cooper, N. *et al.* Additively manufactured ultra-high vacuum chamber for portable quantum technologies. *Additive Manufacturing* **40,** 101898. ISSN: 2214-8604. `https://www.sciencedirect.com/science/article/pii/S2214860421000634` (2022) (Apr. 1, 2021).

90. Mehdizadeh, M. & Ishii, T. Electromagnetic field analysis and calculation of the resonance characteristics of the loop-gap resonator. *IEEE Transactions on Microwave Theory and Techniques* **37,** 1113–1118. ISSN: 1557-9670 (July 1989).

91. Tremblay, P. & Jacques, C. Optical pumping with two finite linewidth lasers. *Physical Review A* **41,** 4989–4999. `https://link.aps.org/doi/10.1103/PhysRevA.41.4989` (2019) (May 1, 1990).

92. Avila, G. *et al.* State selection in a cesium beam by laser-diode optical pumping. *Physical Review A* **36,** 3719–3728. `https://link.aps.org/doi/10.1103/PhysRevA.36.3719` (2019) (Oct. 1, 1987).

93. Duan, Y.-X. *et al.* State Preparation in a Cold Atom Clock by Optical Pumping. *Chinese Physics Letters* **34,** 073201. ISSN: 0256-307X. `https://doi.org/10.1088%2F0256-307x%2F34%2F7%2F073201` (2019) (July 2017).
94. *NIR-MPX800-LN datasheet* (iXblue, 2022). `https://www.ixblue.com/wp-content/uploads/2022/02/NIR-MPX800-LN%20SERIES_6.pdf` (2022).
95. Betts, G., O'Donnell, F. & Ray, K. Effect of annealing on photorefractive damage in titanium-indiffused LiNbO/sub 3/ modulators. *IEEE Photonics Technology Letters* **6,** 211–213. ISSN: 1941-0174 (Feb. 1994).
96. H. Gourard (iXblue). *private communication* Jan. 13, 2020.
97. Preston, D. W. Doppler-free saturated absorption: Laser spectroscopy. *American Journal of Physics* **64,** 1432–1436. ISSN: 0002-9505. `https://aapt-scitation-org.proxy.lib.strath.ac.uk/doi/10.1119/1.18457` (2022) (Nov. 1996).
98. Donley, E. A., Heavner, T. P., Levi, F., Tataw, M. O. & Jefferts, S. R. Double-pass acousto-optic modulator system. *Review of Scientific Instruments* **76,** 063112. ISSN: 0034-6748. `https://aip.scitation.org/doi/figure/10.1063/1.1930095` (2022) (June 2005).
99. Clairon, A., Salomon, C., Guellati, S. & Phillips, W. D. Ramsey Resonance in a Zacharias Fountain. *Europhysics Letters (EPL)* **16,** 165–170. ISSN: 0295-5075. `https://doi.org/10.1209%2F0295-5075%2F16%2F2%2F008` (2019) (Sept. 1991).
100. Kasevich, M. A., Riis, E., Chu, S. & DeVoe, R. G. RF spectroscopy in an atomic fountain. *Physical Review Letters* **63,** 612–615. `https://link.aps.org/doi/10.1103/PhysRevLett.63.612` (2022) (Aug. 7, 1989).
101. Itano, W. M. *et al.* Quantum projection noise: Population fluctuations in two-level systems. *Physical Review A* **47,** 3554–3570. `https://link.aps.org/doi/10.1103/PhysRevA.47.3554` (2022) (May 1, 1993).
102. Santarelli, G. *et al.* Quantum Projection Noise in an Atomic Fountain: A High Stability Cesium Frequency Standard. *Physical Review Letters* **82,** 4619–4622. `https://link.aps.org/doi/10.1103/PhysRevLett.82.4619` (2019) (June 7, 1999).
103. Elvin, R. *et al.* Cold-atom clock based on a diffractive optic. *Optics Express* **27,** 38359–38366. ISSN: 1094-4087. `https://www.osapublishing.org/oe/abstract.cfm?uri=oe-27-26-38359` (2020) (Dec. 23, 2019).
104. Guena, J. *et al.* *New measurement of the rubidium hyperfine frequency using LNE-SYRTE fountain ensemble* in *2008 IEEE International Frequency Control Symposium* (May 2008), 366–370.
105. Ovchinnikov, Y. B. *Development of NPL Rb fountain frequency standard* in *2012 European Frequency and Time Forum* (Apr. 2012), 96–100.

106. Esnault, F. *et al. Reaching a few 10-13 τ-1/2 stability level with the compact cold atom clock HORACE* in *2008 IEEE International Frequency Control Symposium* (May 2008), 381–385.
107. Audoin, C., Santarelli, G., Makdissi, A. & Clairon, A. Properties of an oscillator slaved to a periodically interrogated atomic resonator. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **45,** 877–886. ISSN: 1525-8955 (July 1998).
108. Parker, T. *Hydrogen maser ensemble performance and characterization of frequency standards* in *Proceedings of the 1999 Joint Meeting of the European Frequency and Time Forum and the IEEE International Frequency Control Symposium (Cat. No.99CH36313)* **1** (Apr. 1999), 173–176 vol.1.
109. Mann, A., Sheng, C. & Luiten, A. Cryogenic sapphire oscillator with exceptionally high frequency stability. *IEEE Transactions on Instrumentation and Measurement* **50,** 519–521. ISSN: 1557-9662 (Apr. 2001).
110. Li, W., Du, Y., Li, H. & Lu, Z. Development of low phase noise microwave frequency synthesizers for reducing Dick effect of Cs fountain clocks. *AIP Advances* **8,** 095311. `https://aip-scitation-org.proxy.lib.strath.ac.uk/doi/10.1063/1.5042492` (2022) (Sept. 2018).
111. Rossetto, N. *et al. Dick effect and cavity pulling on HORACE compact cold atom clock* in *2011 Joint Conference of the IEEE International Frequency Control and the European Frequency and Time Forum (FCS) Proceedings* (2011), 1–4.
112. Esnault, F. X., Rossetto, N., Holleville, D., Delporte, J. & Dimarcq, N. HORACE: A compact cold atom clock for Galileo. *Advances in Space Research. Scientific applications of Galileo and other Global Navigation Satellite Systems - II* **47,** 854–858. ISSN: 0273-1177. `https://www.sciencedirect.com/science/article/pii/S0273117710008124` (2022) (Mar. 1, 2011).
113. Yu, M. *et al.* Low phase noise microwave frequency synthesizer for cold atom clock. *AIP Advances* **9,** 045223. `https://aip.scitation.org/doi/full/10.1063/1.5093165` (2022) (Apr. 2019).
114. Kang, S., Gharavipour, M., Affolderbach, C., Gruet, F. & Mileti, G. Demonstration of a high-performance pulsed optically pumped Rb clock based on a compact magnetron-type microwave cavity. *Journal of Applied Physics* **117,** 104510. ISSN: 0021-8979. `https://aip-scitation-org.proxy.lib.strath.ac.uk/doi/abs/10.1063/1.4914493` (2021) (Mar. 12, 2015).
115. Bandi, T. *et al.* Compact high-performance continuous-wave double-resonance rubidium standard with 1.4 $\times$ 10-13 τ-1/2 stability. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **61,** 1769–1778. ISSN: 1525-8955 (Nov. 2014).

116. Gibble, K. Comment on 'First accuracy evaluation of NIST-F2'. *Metrologia* **52,** 163–166. ISSN: 0026-1394. `https://doi.org/10.1088%2F0026-1394%2F52%2F1%2F163` (2020) (Jan. 2015).

117. Dedman, C. J., Baldwin, K. G. H. & Colla, M. Fast switching of magnetic fields in a magneto-optic trap. *Review of Scientific Instruments* **72,** 4055–4058. ISSN: 0034-6748. `https://aip-scitation-org.proxy.lib.strath.ac.uk/doi/abs/10.1063/1.1408935` (2022) (Nov. 2001).

118. Santarelli, G. *et al. Phase transient measurement at the micro radian level for atomic fountain clocks* in *Proceedings of the 20th European Frequency and Time Forum* (Mar. 2006), 166–172.

119. Santarelli, G. *et al.* Switching atomic fountain clock microwave interrogation signal and high-resolution phase measurements. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **56,** 1319–1326. ISSN: 1525-8955 (July 2009).

120. Liu, D. *et al. Study the phase transient of the RF Mach-Zehnder interferometric switch* in *2017 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS)* (July 2017), 292–293.

121. Kokkelmans, S. J. J. M. F., Verhaar, B. J., Gibble, K. & Heinzen, D. J. Predictions for laser-cooled Rb clocks. *Physical Review A* **56,** R4389–R4392. `https://link.aps.org/doi/10.1103/PhysRevA.56.R4389` (2020) (Dec. 1, 1997).

122. Sortais, Y. *et al.* Cold Collision Frequency Shifts in a ^{87}Rb Atomic Fountain. *Physical Review Letters* **85,** 3117–3120. `https://link.aps.org/doi/10.1103/PhysRevLett.85.3117` (2020) (Oct. 9, 2000).

123. Gibble, K. Scattering of Cold-Atom Coherences by Hot Atoms: Frequency Shifts from Background-Gas Collisions. *Physical Review Letters* **110,** 180802. `https://link.aps.org/doi/10.1103/PhysRevLett.110.180802` (2020) (May 3, 2013).

124. Shirley, J. H. *et al.* Microwave leakage-induced frequency shifts in the primary frequency Standards NIST-F1 and IEN-CSF1. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **53,** 2376–2385. ISSN: 1525-8955 (Dec. 2006).

125. Fang, F. *et al.* Accurate Evaluation of Microwave-Leakage-Induced Frequency Shifts in Fountain Clocks. *Chinese Physics Letters* **31,** 100601. ISSN: 0256-307X. `https://doi.org/10.1088%2F0256-307x%2F31%2F10%2F100601` (2020) (Oct. 2014).

126. Gerginov, V., Nemitz, N. & Weyers, S. Initial atomic coherences and Ramsey frequency pulling in fountain clocks. *Physical Review A* **90,** 033829. `https://link.aps.org/doi/10.1103/PhysRevA.90.033829` (2020) (Sept. 18, 2014).

127. Wynands, R., Schroder, R. & Weyers, S. Majorana Transitions in an Atomic Fountain Clock. *IEEE Transactions on Instrumentation and Measurement* **56,** 660–663. ISSN: 1557-9662 (Apr. 2007).
128. Shirley, J. H., Heavner, T. P. & Jefferts, S. R. First-Order Sideband Pulling in Atomic Frequency Standards. *IEEE Transactions on Instrumentation and Measurement* **58,** 1241–1246. ISSN: 1557-9662 (Apr. 2009).
129. Audoin, C., Jardino, M., Cutler, L. S. & Lacey, R. F. Frequency Offset Due to Spectral Impurities in Cesium-Beam Frequency Standards. *IEEE Transactions on Instrumentation and Measurement* **27,** 325–329. ISSN: 1557-9662 (Dec. 1978).
130. Levi, F., Shirley, J., Heavner, T., Yu, D.-H. & Jefferts, S. Power dependence of the frequency bias caused by spurious components in the microwave spectrum in atomic fountains. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **53,** 1584–1589. ISSN: 1525-8955 (Sept. 2006).
131. Bize, S., Sortais, Y., Mandache, C., Clairon, A. & Salomon, C. Cavity frequency pulling in cold atom fountains. *IEEE Transactions on Instrumentation and Measurement* **50,** 503–506. ISSN: 1557-9662 (Apr. 2001).
132. Guéna, J., Abgrall, M., Clairon, A. & Bize, S. Contributing to TAI with a secondary representation of the SI second. *Metrologia* **51,** 108. ISSN: 0026-1394. `https://iopscience-iop-org.proxy.lib.strath.ac.uk/article/10.1088/0026-1394/51/1/108/meta` (2020) (Jan. 29, 2014).
133. Angstmann, E. J., Dzuba, V. A. & Flambaum, V. V. Frequency shift of hyperfine transitions due to blackbody radiation. *Physical Review A* **74,** 023405. `https://link.aps.org/doi/10.1103/PhysRevA.74.023405` (2020) (Aug. 11, 2006).
134. Vanier, J. & Audoin, C. *The quantum physics of atomic frequency standards* (Hilger, Philadelphia, 1989).