# Mini-Fountain Cold Atom Clock

Samuel Walby

Wadham College
University of Oxford

A thesis submitted for the degree of
*Doctor of Philosophy*
Trinity 2024

# Acknowledgements

I'm thankful to the many fantastic colleagues in the NPL caesium fountains group who've made the last four years of struggle and strife not just tolerable, but occasionally enjoyable. Special thanks must go to my industrial supervisor, Dr Krzysztof Szymaniec, who conceived of the mini-fountain project long before I ever got involved. His ability to persuade NPL into giving us enough money to get useful research done has been crucial, without which I'd probably still be banging on about what I did in my first year. His genuine enthusiasm to see me succeed has been essential for instilling confidence in me and my abilities. I must also thank Dr Rich Hendricks, who always has exactly the right solution to whatever problem I happen to be moaning about at any specific moment in time. I'm grateful also to my academic supervisor, Prof Chris Foot, for his many helpful suggestions and careful reading of this manuscript.

I'm reminded of a time where, after a discussion about what I want to do when I grow up, my Dad's colleagues "laughed at" how much money my pursuit of a PhD would cost him. Well, actually, they're all paying for it through their taxes, so the joke is on them!

Thanks also to my many skilled teammates at Oxford University Company of Archers who, time and time again and seemingly against all odds, have allowed me to continue my undefeated streak in student archery regional leagues: we brought back what future archers will refer to as our "glory days" and I can "retire" having surely set some kind of record.

Most importantly, I could not have achieved anything resembling, well, this, without the unconditional and unwavering support from my co-tenant, angry bunny rabbit, and scientifically-proven *second*-stubbornest husky in the world: 조희승. Fortunately, I am the stubbornest, so am a good match in the many (metaphorical) aroo-ing contests we often find ourselves in. Hopefully one day soon we have some real dogs of our own.

# Abstract

The use of atomic transitions as frequency references has lead to frequency becoming the most accurately-measurable quantity in modern science. This has allowed for the maintenance and dissemination of standard global timescales that continuously maintain consistency with each other to nanosecond precision, as well as advances in fundamental physics.

Current atomic fountain clocks reach fractional frequency instabilities of a few $10^{-16}$ after several of days of averaging. This is limited by the instability of several systematic effects that have been the subject of much scrutiny in order to further push the limits of current caesium, microwave-based technology. As this technology matures, it is natural to seek applications beyond National Metrology Institutes. However, caesium fountains are large, complex, and expensive pieces of equipment, that require a skilled team of people to build, operate, and maintain. The aim of this work is to design, build, and test a mini-fountain cold atom clock that achieves performance close to that of Primary Frequency Standards while reducing the size and complexity of the fountain physics package and optical control system.

This was achieved in several ways. A method of performing normalised clock state detection using only the MOT beams in the MOT chamber allows for the elimination of the large, traditional separated detection regions. Similarly, performing state selection in the MOT chamber with a waveguide, rather than a separate dedicated microwave cavity, allowed the centre of the MOT chamber to be only 14 cm below the centre of the Ramsey cavity. Cutting down on this space required careful simulation of the magnetic field, and a set of shields was designed to allow for a homogenous magnetic field throughout the interrogation region. A new Ramsey cavity design was integrated into the vacuum vessel itself, rather than being housed inside a separate, larger vacuum vessel. Finally, an all-in-fibre optical system for producing the cooling light was designed and implemented, which reduces the amount of free-space optical alignment to improve the robustness of the system.

The resulting rubidium fountain clock has a height reduced from 2.1 m to 0.8 m, and a footprint reduced from a 60 cm × 60 cm square to a 24 cm diameter cylinder, relative to NPL's Primary Frequency Standards, achieved a short term fractional frequency instability of $9 \times 10^{-13}/\sqrt{\tau/\mathrm{s}}$, and the long-term instability is likely below $10^{-15}$, although an evaluation of the systematic effects is not yet complete.

# Structure of the Thesis

Chapter 1 is a review of the literature on atomic fountain clocks. This technology was first proposed and tested in the 1990s, so there is a lot to discuss! One of the key aspects of fountain clocks, that the mini-fountain will replicate, is the ability to analyse all the systematic shifts to an accuracy of a few parts in $10^{16}$. Thus, particular attention is paid to the systematic shifts, the theory behind them, and ways in which they are measured.

Chapter 2 discusses the design considerations that went into the mini-fountain physics package and optical control system. While the desire to replace the the full-size detection zones with *something* was present from the inception of the mini-fountain project, the method of doing so, and the tests carried out to demonstrate the feasibility, are original contributions of the author. Similarly, the design of the Ramsey cavity, which separates the cavity dimensions from the method of vacuum sealing, is an original contribution of the author, with assistance from Martin Knapp. The design and simulation of the magnetic shields was carried out by the author, with CAD assistance from Peter Lovelock. The design for the optics was a joint contribution from the author and Krzysztof Szymaniec, and the work of constructing and testing the optics was performed by the author.

Chapter 3 discusses how the mini-fountain physics package was built. Simulations of the Ramsey cavity Q-factor were performed by Martin Knapp, and the rest of the cavity calculations and construction was performed by the author, with assistance from NPL's mechanical workshop in producing the parts, and from Krzysztof Szymaniec in sealing the MOT chamber.

Chapter 4 discusses how cold atoms were achieved in the mini-fountain, and how Ramsey fringes and frequency stability measurements were obtained. This work was performed entirely by the author. The software to control the fountain sequences was written by Rich Hendricks, with some small modifications by the author.

# Contents

# Chapter 1

# Literature Review

A property of atoms that makes them useful frequency references is that all atoms of a particular species are identical. If a frequency is defined relative to a specific atomic transition, then a *sufficiently-competent* experimenter can use this transition to realise the exact same frequency as any other experimenter, no matter where in the world they are. To that end, the SI second was defined in 1967 as “the duration of 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium-133 atom” [1]. Caesium was chosen because it is the stable alkali metal (all of which have relatively simple atomic structures) with the largest hyperfine splitting.

There are several well-known exploits of accurate time an frequency resulting from atomic clocks. In Global Navigation Satellite Systems (GNSS), precise timing is necessary to time the arrival on the ground of signals travelling at light speed from satellites millions of metres above with an error of only a couple of metres or so [2]. Telecoms networks need accurate time and frequency information to keep nodes sending and receiving data synchronised with each other, or else information is lost [3, 4]. In financial markets, a well-defined sequence of events is crucial for ensuring traceability and accountability [5]. Atomic clocks are also useful in fundamental research, where they have been used to precisely verify predictions of Einstein’s general relativity [6], and could yet be used to discern time-variations in what are currently thought to be fundamental constants of physics, such as the fine-structure constant [7, 8], which would violate general relativity and point the way to new fundamental physics, as well as hint that that our “fine tuned” universe may not stay this way for all eternity [9].

Atomic clocks come in a variety of shapes and sizes, and the performance level is generally proportional to the complexity of the clock. The smallest clocks are known as chip-scale atomic clocks (CSACs), where the entire system is contained in a package that measures several cubic centimetres. These use thermal atoms to achieve a fractional frequency accuracy of around $10^{-11}$, and are useful in telecoms, space, and aviation. The largest, most complex atomic clocks are optical clocks, such as strontium lattice clocks, where the optical table alone takes up several square meters. These, however, also achieve the highest accuracies, around $10^{-18}$, which are useful in fundamental physics research.

Also at the larger end of the scale lie atomic fountain clocks, which are currently used in National Metrology Institutes to realise the definition of the SI second. These are tall devices, with total heights of around 2 m, with modest optical benches to accommodate two separate lasers. These are the best-performing microwave clocks, routinely reaching accuracies and long-term stabilities of around $10^{-16}$.

One widely-used clock which is significantly smaller, but still large enough to typically only be practical in fixed installations, is the hydrogen maser, which achieve good short-term frequency stabilities of below $10^{-13}/\tau/\mathrm{s}$, but are not capable of long-term frequency accuracy. This is because the thermal hydrogen atoms interact with the teflon-coated walls of the storage bulb, which produces a long-term frequency drift which is unpredictable and not reproducible. A device with similar size, weight, and power, that takes a similar level of skill to operate, and is free from long-term drift, could potentially find application in areas where masers and other commercial clocks are commonplace, such as in local timekeeping applications for very-long baseline interferometry, GNSS monitoring, or telecoms networks.

This is where the mini-fountain fits into the current landscape of atomic clocks: it should be comparable in size, complexity, and cost to a maser but have the distinct advantage of having a long-term frequency accuracy better than $10^{-15}$, so it could be used as a holdover clock that would only drift by a nanosecond or so over an entire month.

There are other cold-atom clocks aiming fill this niche. For example, MuClock is a commercially-available device in a similar footprint to a maser that uses laser cooling to generate cold rubidium atoms, which are dropped into a microwave cavity. While the interrogation time is shorter than in a fountain, this also allows for a high repetition rate, such that the resulting short-term stability is around $4 \times 10^{-13}/\sqrt{\tau/\mathrm{s}}$

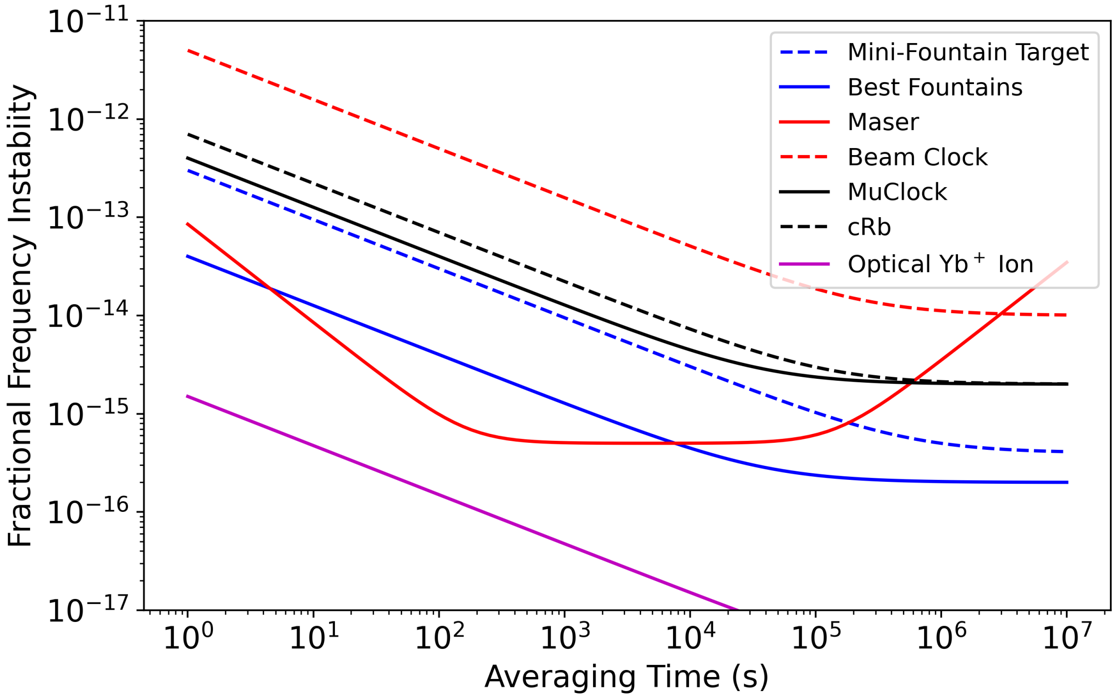

Figure 1.1: Fractional frequency instability vs averaging time for several different commercially-available clocks [10, 11, 12, 13], and with an optical Yb ion clock for comparison [14].

[12], and the flicker-floor is reached at around $2 \times 10^{-15}$ (solid black line in figure 1.1). Similarly, the Spectra Dynamics cRb clock is even smaller and achieves similar frequency stability [13] (dashed black line in figure 1.1) .

The advantage of the mini-fountain is that the flicker floor, after weeks to months of averaging, should be in the $10^{-16}$ range, which is currently only achieved by full-sized fountain clocks (and exceeded by optical clocks). Additionally, the fountain geometry has had much effort devoted to the study of the systematic effects, and thus is capable of being used as an absolute frequency standard.

## 1.1 Fountain Development History

Before atomic timekeeping, the second was defined in terms of Earth's orbit around the Sun. While not as universal as a definition based on atomic transitions, it was something that every human could easily conceptualise and keep track of. Timescales were primarily maintained and disseminated by astronomical observatories such as the Royal Greenwich Observatory [15] until 1972 [16].

The basis of an accurate measurement of time is an accurate measurement of frequency: if a frequency is well known, then keeping track of time is simply a case of counting oscillations. This section will discuss the key advances that lead to the development of atomic microwave frequency standards.

### 1.1.1 Ramsey's Separated Oscillating Fields

In the 1930s, beams of atoms and molecules traversing magnetic fields was a method of measuring the magnetic moments of nuclei. While it was known that atoms with different orientations of their magnetic moment relative to a magnetic field gradient would be deflected differently, the sizes and signs of nuclear magnetic moments was not yet known, as this would have required that the deflected atoms be probed with a rotating magnetic field [17] to excite transitions between magnetic sub-levels [18]. An equivalent experiment, using a spatially-varying field [19], was used initially to determine the magnetic moments of the proton and deuteron, later for measuring the nuclear moments of atoms in simple molecules, and was extended to use an oscillating magnetic field to excite magnetic dipole transitions between different $m_F$ states of the same hyperfine level [20]. For a spin-1/2 system, the transition probability for an atom or molecule in an oscillating field $B_1$, oriented perpendicular to a static field $B_0$ is [20]

$$P_{\text{Rabi}} = \frac{\Theta^2}{\left(1 - \frac{\omega}{\omega_0}\right)^2 + \Theta^2} \sin^2\left(\frac{1}{2}\omega t\sqrt{\left(1 - \frac{\omega}{\omega_0}\right)^2 + \Theta^2}\right), \tag{1.1}$$

where $\Theta = B_1/2B_0 \ll 1$, $\omega$ is the angular frequency of the oscillating field, and $\omega_0 = g_F\mu_B\sqrt{B_0^2 + B_1^2}/h$ is the Larmor frequency of the atom in the resulting magnetic field. There are similar formulae for higher spins [21]. Close to resonance, where $\delta = \omega - \omega_0 \ll \omega_0$, and with appropriate field strengths, this formula reduces to

$$P_{\text{Rabi}} \approx \left|\frac{\Omega t}{2}\right|^2 \frac{\sin^2(\delta t/2)}{\delta t/2}, \tag{1.2}$$

where $\Omega$ is the Rabi frequency at which the atoms' state oscillates between the $m_F = \pm 1/2$ levels. $P_{\text{Rabi}}$ can be made equal to 1 with $\delta = 0$ and $\Omega = 2/t$.

Advances in higher-frequency radio-wave generation and detection allowed this technique of observing transitions between $m_F$ levels, which required frequencies in the 100s of megahertz range or lower, to be extended to observing hyperfine transitions between different $|F\rangle$ states, with frequencies of several gigahertz [22].

It was discovered by Ramsey in 1950 that a narrower resonance profile could be achieved if the oscillating field was split into two regions, with a larger region of static magnetic field between them [23]. This method has the practical advantage of only needing to maintain a uniform oscillating field across much smaller regions than was required by the original Rabi method.

Qualitatively, this works because after the first interaction, if the angular frequency of the oscillating field $\omega$ is close to the Larmor frequency in the constant region $\omega_0$, then an angle will accumulate between the atomic magnetic moment and the fixed, uniform field. If each region of oscillating field applies a $\pi/2$ phase change to the magnetic moments, then the result is equivalent to that of a $\pi$-pulse, except that it is more discriminatory to the difference between $\omega$ and $\omega_0$. If these frequencies differ slightly then the accumulated phase is such that the $\pi$-pulse is not completed. For an atom travelling for a time $\tau$ through each interaction region and for $T$ in the uniform field between interactions, then the overall transition probability after a time $2\tau + T$ can be calculated from the Schrödinger equation to be

$$
\begin{aligned}
P_{\text{Ramsey}} = 4 \left( \frac{\omega_0}{a} \frac{B_1}{B_0} \right)^2 \sin^2 \left( \frac{1}{2} a\tau \right) & \\
\times \left[ \cos \left( \frac{1}{2} \delta T \right) \cos \left( \frac{1}{2} a\tau \right) - \frac{\delta}{a} \sin \left( \frac{1}{2} \delta T \right) \sin \left( \frac{1}{2} a\tau \right) \right]^2 &, \quad (1.3)
\end{aligned}
$$

where $\omega_0$ is the hyperfine transition frequency, $\delta$ is the detuning of the oscillating field in the interaction regions from resonance, $B_0$ is the strength of the uniform field, $B_1$ is the strength of the oscillating field, and $a = \sqrt{\delta^2 + (\omega_0 B_1/B_0)^2}$. This equation describes the Ramsey fringes (shown in figure 1.2) that arise by scanning the applied microwave frequency across resonance. A more compact form of equation 1.3 valid for small detuning $\delta = \omega - \omega_0 \ll \omega_0$, is given in terms of the Rabi frequency $\Omega$ such that [24]

$$
P_{\text{Ramsey}} = \left| \frac{\Omega\tau}{2} \right|^2 \left( \frac{\sin(\delta\tau/2)}{\delta\tau/2} \right)^2 \cos^2 \left( \frac{\delta T}{2} \right). \quad (1.4)
$$

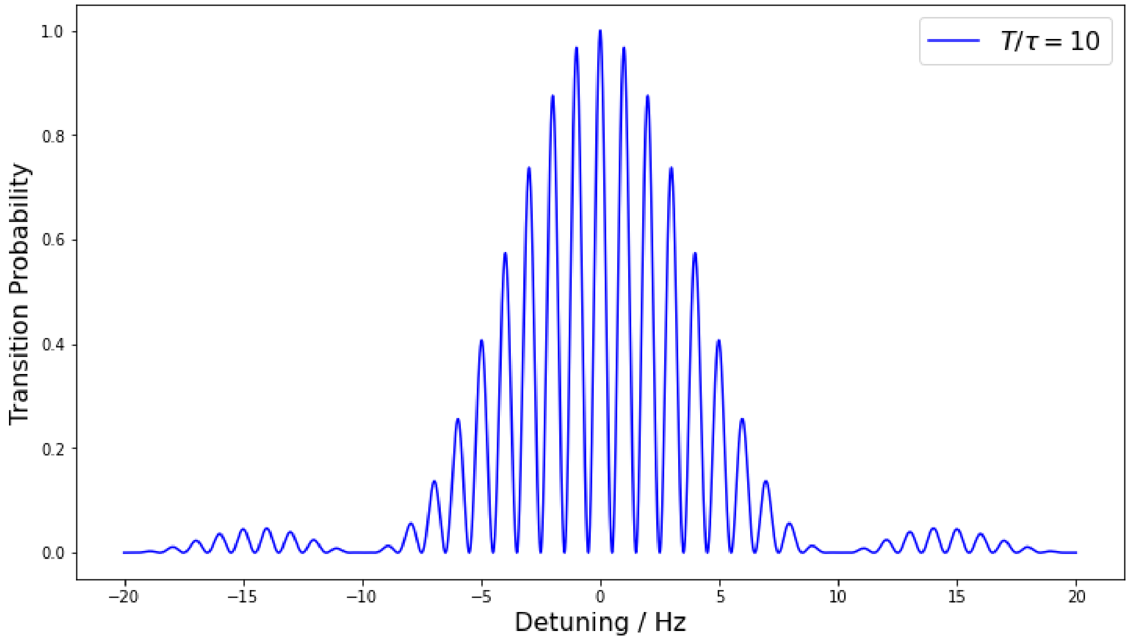

Figure 1.2: Theoretical example of transition probability vs frequency for a Ramsey-type interaction where $T/\tau = 10$.

If the intermediate time $T$ is longer than the pulse duration $\tau$, then the first minimum of $P_{\text{Ramsey}}$ occurs when

$$\cos(\delta T/2) = 0, \tag{1.5}$$

$$\delta = \pi/T, \tag{1.6}$$

such that the full-width-half-maximum (FWHM) of the central fringe is $\Delta\omega_{\text{Ramsey}} = \pi/T$, inversely proportional to the time between the interactions. A single pulse of length $2\tau$ would maximise the transition probability in the Rabi method, and would have a linewidth of $\Delta\omega_{\text{Rabi}} \sim 2\pi/\tau$. For equal interaction times, the Ramsey method thus gains a factor of 2 in linewidth, but, more crucially, allows easily for a much-extended interrogation time because it is easier to maintain a large region of uniform field than oscillating field: the Ramsey method can reduce the measured linewidth of an atomic transition by a factor of

$$\frac{\Delta\omega_{\text{Rabi}}}{\Delta\omega_{\text{Ramsey}}} = \frac{2T}{\tau}, \tag{1.7}$$

which can be as large as technology allows for.

### 1.1.2 Atomic Beam Clocks

Ramsey's method was put to use in applications aiming to standardise a new definition of time based on atomic transitions: the caesium beam clock [25]. The main attribute of the caesium beam clock that made it the most appropriate choice as a *primary* frequency standard was that the energy of the microwave transition used is relatively insensitive to the parameters of the experiment. The most significant frequency shift was the Zeeman shift due to the background magnetic field.

The Hamiltonian for the interaction between the nuclear angular momentum operator $\hat{\mathbf{I}}$ and the total electronic angular momentum operator $\hat{\mathbf{J}}$ for the ground state of alkali atoms ($J = 1/2$) is

$$\hat{H}_{\text{hfs}} = A_{\text{hfs}} \hat{\mathbf{I}} \cdot \hat{\mathbf{J}}. \tag{1.8}$$

The Hamiltonian for the interaction of atomic states $|F, m_F\rangle$ with an external field $\mathbf{B} = B_z \hat{\mathbf{z}}$ is

$$\hat{H}_B = \mu_B g_F B_z \hat{F}_z, \tag{1.9}$$

where the hyperfine Landé g-factor $g_F$ can be made out of the electronic g-factor $g_J$ and the nuclear-spin g-factor $g_I$ according to

$$g_F = g_J \frac{F(F+1) - I(I+1) + J(J+1)}{2F(F+1)} + g_I \frac{F(F+1) + I(I+1) - J(J+1)}{2F(F+1)}. \tag{1.10}$$

The result is the Breit-Rabi formula [26], which for an atom with total angular momentum $F = I \pm 1/2$ ($s$-states, there are similar formulae for non-zero $L$ [21]) and component $m_F$ along a field $B_z$ has two forms:

$$E_{|I\pm1/2,m_F\rangle} = -\frac{h\nu_0}{2(2I+1)} + g_I \mu_B m_F B_z + \frac{h\nu_0}{2}\left(\pm\sqrt{1 + \frac{4m_F}{2I+1}x + x^2}\right), \tag{1.11}$$

for the cases where $|m_F| \neq I + 1/2$, and

$$E_{|I\pm1/2,m_F\rangle} = -\frac{h\nu_0}{2(2I+1)} + g_I \mu_B m_F B_z + \frac{h\nu_0}{2}(1 \pm x), \tag{1.12}$$

for the cases where $m_F = \pm(I + 1/2)$, where $h\nu_0$ is the unperturbed hyperfine energy splitting between the $I + 1/2$ and $I - 1/2$ states, $\mu_B = 9.27 \times 10^{-24}\,\text{JT}^{-1}$ is the Bohr magneton, and $x = (g_J - g_I)\mu_B B_z/h\nu_0$ [22]. At small $B_z$ such that $x \ll 1$, the splittings of each level are close to linear and correspond to the Zeeman effect. At large fields the total angular momentum is not a good quantum number [27]. Caesium has a nuclear angular momentum $I = 7/2$. Transitions that are insensitive to magnetic-field strength, to first-order, can be used as "clock" transitions, such as the $6^2S_{1/2}\,|F = 3, m_F = 0\rangle \rightarrow |F = 4, m_F = 0\rangle$ transition. Neglecting $g_I$ as $g_I \ll g_J$, equation 1.11 reduces to

$$E_\pm = -\frac{h\nu_0}{16} \pm \frac{h\nu_0}{2}\sqrt{1 + x^2}, \tag{1.13}$$

thus, for small $x$, the actual frequency of the hyperfine clock transition has a small quadratic dependence,

$$\nu = \nu_0\left(1 + \frac{1}{2}x^2\right) \approx (\nu_0 + 4.27 \times 10^{10}\ B_z^2), \tag{1.14}$$

for $\nu_0$ in Hz and $B_z$ in Tesla. A field of $5\,\mu$T inside the intermediate region of the first atomic beam apparatus resulted in the measured mean central frequency of the Ramsey pattern deviating from the unperturbed frequency by 1 Hz [28]. The ease of relating the measured frequency back to the unperturbed frequency is the key aspect of the caesium beam clock that made it suitable to be adopted as a primary standard. Even the best hydrogen masers are unsuitable, as they have frequency shifts due to collisions between hydrogen atoms and the storage bulb walls that depend on the construction method and age of the teflon coating [29], and thus are not reproducible from device to device, and are time-varying.

A Cs beam clock works by using an oven to produce a narrow beam of thermal caesium atoms, which has $k_B T \gg h\nu_0$, such that the $F = 3$ and $F = 4$ states are equally populated. In a high vacuum the atoms have a sufficient mean-free path to act as a beam, and Essen and Parry found that the pressure should be below $10^{-6}$ mbar to avoid degradation, due to collisions with background gas, to the detected beam signal [28]. These atoms are then state-selected by an inhomogeneous magnetic field, such that atoms in each hyperfine level are deflected in different directions. Atoms in the $F = 3$ level are directed towards the first arm of the U-shaped microwave cavity, where each arm is oscillating at a frequency $\nu_0$ in a $\mathrm{TE}_{011}$ mode, which minimises the Doppler effect by being very close to uniform in

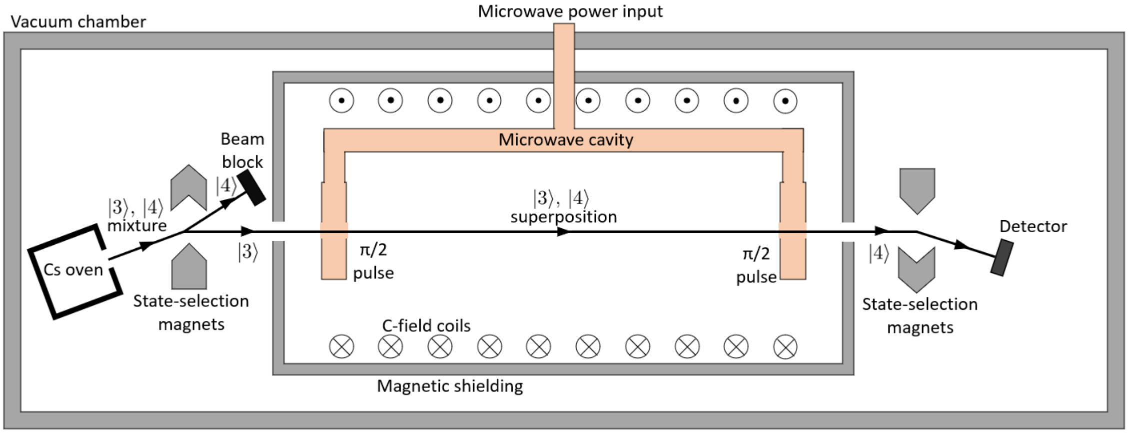

Figure 1.3: An illustration of a caesium beam clock experimental set-up.

phase and amplitude along the path of the beam, and across the diameter of the opening. This first arm drives a nominal $\pi/2$ pulse to put the atoms in a coherent superposition of $F = 3$ and $F = 4$ which evolves through the intermediate region in which there exists a weak, uniform field generated by coils within magnetic shields, which prevent unwanted external fields from perturbing the atoms. The second arm drives a second nominal $\pi/2$ pulse to complete the population transfer into $F = 4$, if the microwave field is on resonance. Ideally this is in-phase with the first, but in practice a small but non-negligible end-to-end phase shift is inevitable. A second set of state-selecting magnets deflects $F = 4$ atoms towards a hot wire which produces $\mathrm{Cs}^+$ ions that are detected on a metal plate by the current they produce. If the microwave frequency, which is sourced from some external oscillator, deviates from resonance, then there is a change in signal that is used as an error signal to adjust the oscillator frequency back towards resonance.

The first atomic beam clock achieved a Ramsey fringe width of 100 Hz, and a stability of 1 part in $10^8$ over an averaging time of 1 hour [28]. More complex hexapole magnets can be used to select not just a specific hyperfine level, but also focus only a particular velocity class of caesium atoms into the beam, such that there is a smaller spread of interaction times with the microwave field [30]. The need for strong magnetic field gradients so close to the interaction region can be removed entirely by replacing magnetic state selection with optical pumping. This uses optical transitions that excite, for example, the $F = 4 \rightarrow F' = 3$ transition, such that if the atom decays back to $F = 4$ it continues to absorb, but if it decays to $F = 3$ it stays there, such that atoms accumulate in the $F = 3$ level [31]. Another advantage of the optical pumping method is that virtually all the atoms in the beam end up in the desired hyperfine state, as opposed to only about half with magnetic state selection, and more complicated optical pumping schemes can preferentially populate the desired $m_F = 0$ Zeeman sub-level [32].

Today, caesium beam clocks are not state-of-the-art primary frequency standards; they have been surpassed in this role by caesium fountain clocks, which achieve an order of magnitude better stability and accuracy, and which themselves will one day be supplanted by optical clocks, which have currently been demonstrated to offer around two orders of magnitude better stability and accuracy.

Beam clocks are still in use in many settings, either as back-ups in national metrology institutes, or for their portability [33]. In the best caesium beam clocks, the widths of Ramsey fringes range from around 20 to 100 Hz [34][35], short-term stabilities are of the order of a few $10^{-13}/\sqrt{\tau/\mathrm{s}}$ [36], and long-term frequency accuracies are of the order of a few $10^{-15}$ [37]. This accuracy is sufficient to reveal other, more subtle systematic effects, for example the effect of time-dilation on the atoms as they pass through the apparatus, the change in the hyperfine splitting due to the atoms being immersed in blackbody radiation from the environment, and the effect of collisions between caesium atoms, amongst others [33]. Caesium beam clocks are available as off-the-shelf commercial products, typically achieving short-term stabilities of $\sim 10^{-12}/\sqrt{\tau/\mathrm{s}}$, long-term stabilities of $\sim 10^{-14}$, and accuracies $\sim 5 \times 10^{-13}$ [11].

### 1.1.3 Laser Cooling

As implied by equation 1.7, the linewidth of the Ramsey pattern can be narrowed by increasing the measurement time. With a caesium beam clock, this time is limited by the thermal velocities of the caesium atoms emitted from the oven, typically $\sim 200\,\mathrm{ms}^{-1}$ on average, compared to typical device lengths of around 1 m or less.

It was first suggested that laser light could be used to cool atoms in the 1970s [38], and in the 1980s, techniques were developed to cool, trap, and manipulate large numbers of neutral atoms.

One method is the optical molasses [39] [40]: three perpendicular pairs of counter-propagating laser beams are slightly red-detuned from resonance with an atomic transition, such that a moving atom will be Doppler-shifted into resonance with a beam opposed to its motion, and thus, over many cycles of absorption and spon-

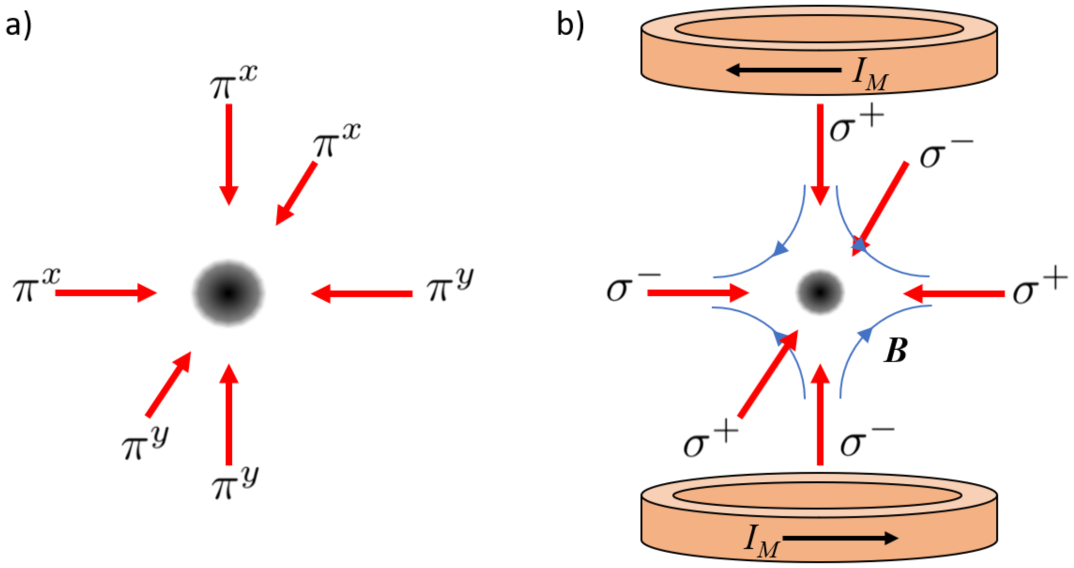

Figure 1.4: a) Simple schematic of an optical molasses, using perpendicular linear polarisations $\pi^x$ and $\pi^y$ for the cooling beams. b) Simple schematic of a MOT, indicating the required direction of circular polarisations $\sigma^+$ and $\sigma^-$ relative to the coil currents $I_M$, and the resulting $\boldsymbol{B}$ field near the trap centre.

taneous emission [41], experiences a drag force that slows it down. The photons absorbed during this process are re-emitted in random directions, giving the atom a series of random momentum kicks that limit the temperature that can be achieved.

The optical molasses technique can be extended by the use of anti-Helmholtz coils to generate a spherical quadrupole magnetic field, which has zero field strength at its centre but increases linearly away from the centre. The Zeeman effect of this quadrupole field causes atoms that move away from the trap centre to be shifted into resonance with the beam, with appropriate circular polarisation, that pushes the atom back towards the trap centre. Thus, the velocity-dependent force of the optical molasses is combined with a position-dependent, central potential to confine the atoms spatially too. This set-up is known as a magneto-optical trap (MOT). The optical molasses and MOT configurations are shown in figure 1.4. So long as multiple-scattering of photons in the atom cloud is negligible, the density of the cold atom cloud will be Gaussian in both cases [42].

### Near-Resonant Optical Scattering

Electric dipole transitions are characterised by the dipole matrix element, $\mathbf{D}_{eg}$, which comes from the spatial wavefunctions $\psi_g(\mathbf{r})$ and $\psi_e(\mathbf{r})$ of the ground and excited states $|g\rangle = |N, L, m\rangle$ and $|e\rangle = |N', L', m'\rangle$, such that

$$\mathbf{D}_{eg} = \langle g|\, e\hat{\mathbf{r}}\, |e\rangle \tag{1.15}$$

$$= \iiint \psi_g^*(\mathbf{r})\; e\mathbf{r}\; \psi_e(\mathbf{r})\; d^3\mathbf{r}. \tag{1.16}$$

Here, $|N, L, m\rangle$ are used to label the atomic states, rather than $|F, m_F\rangle$, because the principle quantum number $N$ and the orbital angular momentum $L$ are the quantum numbers relevant to the spatial wavefunctions. The total angular momentum operator is $\hat{J} = \hat{L} + \hat{S}$. For fine-structure transitions within the same principle level, the $\langle J|\, e\hat{\mathbf{r}}\, |J'\rangle$ matrix elements can be related back to the hyperfine matrix elements by

$$\begin{aligned}\langle F, m_F|\, e\hat{\mathbf{r}}\, |F', m_{F'}\rangle &= \langle J|\, e\hat{\mathbf{r}}\, |J'\rangle \\ &\times (-1)^{F'-1+m_F}\sqrt{(2F'+1)(2J+1)} \begin{Bmatrix} J & J' & 1 \\ F' & F & I \end{Bmatrix} \\ &\times \sum_{q=-1,0,1} \langle F, m_F|F', m_{F'}, 1, q\rangle\,, \end{aligned} \tag{1.17}$$

where the term in curly brackets is the Wigner 6j symbol [43] and the final summed-over terms are the Clebsch-Gordan coefficients.

The matrix element is related to the decay rate $\Gamma$ from the excited state to the ground state by [24]

$$\Gamma = \frac{4}{3}\frac{\omega_0^3\alpha}{e^2c^2}\frac{2J+1}{2J'+1}|\mathbf{D}_{eg}|^2, \tag{1.18}$$

where $\alpha = e^2/4\pi\epsilon_0\hbar c \approx 1/137$ is the fine-structure constant, $2J+1$ is the degeneracy of the ground state, and $2J'+1$ is the degeneracy of the excited state. Another important parameter is the Rabi frequency: the frequency at which a two-level system undergoes coherent population transfer,

$$\Omega_{e,g} = \frac{\mathbf{D}_{eg}\cdot\hat{\mathbf{E}}}{\hbar}, \tag{1.19}$$

where $\hat{\mathbf{E}}$ is the polarisation vector of the light's electric field. This is can be related to the saturation intensity, $I_{\text{sat}}$, by

$$\frac{2\Omega_{e,g}^2}{\Gamma^2} = \frac{I}{I_{\text{sat}}}. \tag{1.20}$$

The frequency and intensity-dependent scattering cross-section for an atom interacting with near-resonant light is then

$$\sigma_{\text{abs}}(\omega, I) = \sigma_0 \frac{\Gamma^2/4}{\delta^2 + (1 + I/I_{\text{sat}})\,\Gamma^2/4}, \tag{1.21}$$

where $\delta = \omega - \omega_0$ is the detuning and

$$\begin{aligned} \sigma_0 =& \sigma_{\text{abs}}(\delta = 0, I = 0) \\ =& \frac{\hbar\omega_0\Gamma}{2I_{\text{sat}}} \qquad (1.22) \\ =& 3|\hat{\mathbf{E}} \cdot \hat{\mathbf{D}}_{eg}|^2 \times \frac{\lambda_0^2}{2\pi} \times \frac{2J'+1}{2J+1}, \qquad (1.23) \end{aligned}$$

which depends on the angle between the polarisation of the electric field and the atom's dipole moment. Equation 1.21 can be re-written as a scattering rate

$$R_s = \frac{\Gamma}{2}\frac{I}{I_{\text{sat}}}\frac{1}{1 + 4(\delta/\Gamma)^2 + I/I_{\text{sat}}}, \tag{1.24}$$

which has a maximum value of $\Gamma/2$ for resonant light where $I \gg I_{\text{sat}}$.

### Doppler Theory

The radiation force exerted by a beam of intensity $I$, detuning $\delta$, and wavevector $\mathbf{k}$ on an atom of cross-section $\sigma_{\text{abs}}$ and resonant saturation intensity $I_{\text{sat}}$ is $\mathbf{F}_{\text{rad}} = \sigma_{\text{abs}} I/c\,\hat{\mathbf{k}}$. In one dimension, two counter-propagating beams result in a force, in the low-intensity limit, of [41]

$$F_{\text{rad}} = \hbar k \frac{I}{I_{\text{sat}}} \left( \frac{\Gamma^2}{(\delta + vk)^2 + \Gamma^2} - \frac{\Gamma^2}{(\delta - vk)^2 + \Gamma^2} \right), \tag{1.25}$$

on the atom, which can be expanded to show that, in the low-velocity limit, $F_{\text{rad}} \approx 4\hbar k^2 (I/I_{\text{sat}})\delta v/\Gamma^2 = -\alpha_D v$, which is a friction force with coefficient $\alpha_D$ for negative

$\delta$. A slightly more rigorous treatment [44] that keeps the next higher-order terms in the expansion of equation 1.25 results in

$$\alpha_D = -4\hbar k^2 \frac{I}{I_{\text{sat}}} \frac{2\delta/\Gamma}{(1 + (2\delta/\Gamma)^2)^2}. \tag{1.26}$$

The corresponding rate of energy loss is $R_{\text{cool}} = F_{\text{rad}} v = -\alpha_D v^2$. Spontaneous emission of a photon results in a momentum kick of $\Delta p = \hbar k$ to the atom in a random direction. Additionally, the scattering rate is not a constant, thus the force on an atom fluctuates. This means that $R_{\text{heat}} = \hbar^2 k^2 R_s / m$, where the two-beam, low-intensity, 1D scattering rate is

$$R_s = \frac{I}{I_{\text{sat}}} \frac{\Gamma}{1 + (2\delta/\Gamma)^2}. \tag{1.27}$$

The temperature can be identified with the mean thermal energy of the ensemble of atoms, and $R_{\text{heat}}$ balances $R_{\text{cool}}$ when [24]

$$k_B T = \frac{\hbar\Gamma}{4} \left( \frac{1 + (2\delta/\Gamma)^2}{-2\delta/\Gamma} \right), \tag{1.28}$$

such that a minimum temperature of $T_D = \hbar\Gamma/2k_B$, the so-called Doppler limit, may be reached when the laser detuning is $\delta = -\Gamma/2$.

It was, however, a pleasant surprise to see that temperatures of around $40\,\mu$K were achieved in an early optical molasses of sodium atoms [45]. The Doppler limit in this case should have been $240\,\mu$K. This occurred at $\delta = -1.5\,\Gamma$, which is also inconsistent with Doppler theory. It was soon discovered that the cooling mechanisms of atoms in a light field are more subtle than just the friction of the Doppler theory, and that the much colder temperatures can be explained by the effects of polarisation gradients on the atom-light interaction [46].

### Orthogonal Linear Polarisation

For counter-propagating, orthogonal linear polarisations $\pi^x\pi^y$, the resulting standing wave varies in ellipticity. This field is shown in figure 1.5. Each Zeeman sub-level is perturbed differently in regions of different beam polarisations: linear polarisation, being an equal supposition of each circular component, affects each sub-level equally; the Clebsch-Gordan coefficients for interaction with circularly-polarised light are different, such that each sub-level interacts more strongly with one handedness than the other sub-level [46]. The result is that each level's light shift oscillates with a spatial period of $\lambda/2$.

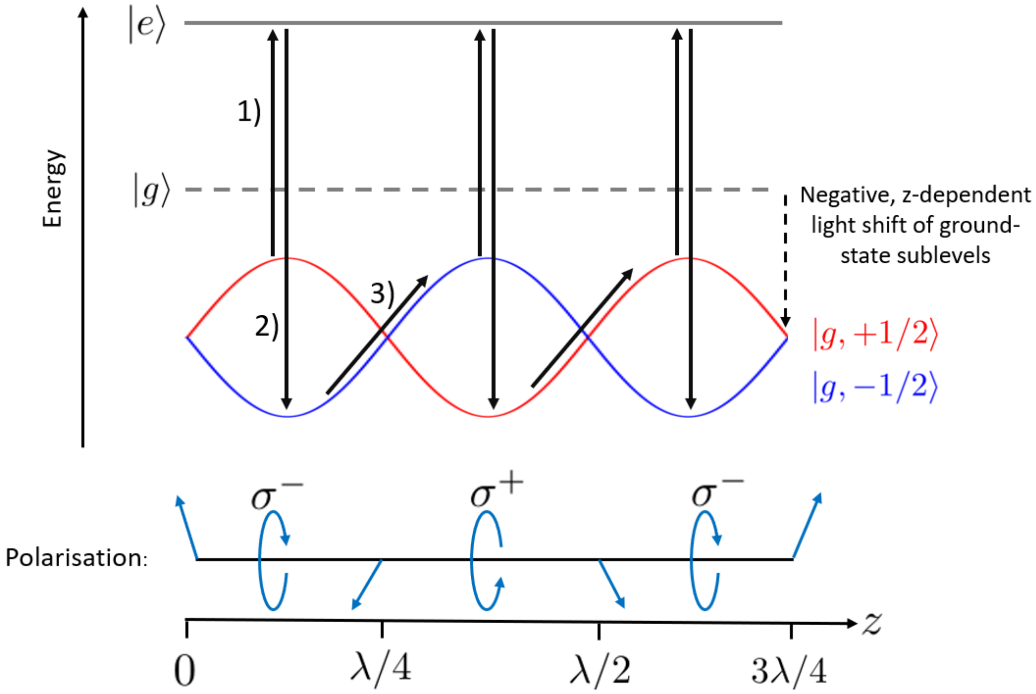

Figure 1.5: Schematic of Sisyphus cooling in counter-propagating beams of orthogonal linear polarisation. 1) Atom on top of potential hill absorbs a photon from one of the beams. 2) Atom emits photon to de-excite to potential valley. 3) Atomic velocity causes atom to climb potential hill, converting kinetic energy into potential energy.

In a region where the polarisation is $\sigma^-$, then atoms (assumed for simplicity to have no nuclear angular momentum, $J = 1/2$ in the ground state $|g\rangle$ and $J' = 3/2$ in the excited state $|e\rangle$) will be optically pumped by the light into the $m_J = -1/2$ sub-level. Similarly, regions of $\sigma^+$ polarisation will pump atoms into $m_J = +1/2$. This means that an atom will, on average, be more likely to move up a potential hill than down, and thus lose kinetic energy. This mechanism is known as Sisyphus cooling.

More detailed calculations [46] show that the Sisyphus friction coefficient, unlike the Doppler friction $\alpha_D$, does not depend on the laser intensity and increases with detuning. Thus, the equilibrium temperature is different for Sisyphus cooling than for Doppler cooling, as each emitted photon can have slightly different energy corresponding to decay to $m_J = +1/2$ or $m_J = -1/2$, and photons are absorbed at different rates from each beam, depending on where the atom is in the standing wave, which results in a fluctuating force on the atom. The result, when averaged over multiple wavelengths [46], is that, for large detuning $|\delta| \gg \Gamma$,

$$k_B T_S \approx \frac{\hbar}{16} \frac{I}{I_{\text{sat}}} \frac{\Gamma^2}{|\delta|}, \tag{1.29}$$

which, unlike the simple Doppler limit, is proportional to the beam intensity and inversely proportional to detuning. This fact will become very useful in section 1.1.4.

### Opposite Circular Polarisation

For counter-propagating beams of opposite circular polarisation, $\sigma^+\sigma^-$ (which, as they are counter-propagating, have the same handedness), the standing wave is of rotating linear polarisation [47], as shown in figure 1.6. Since the resulting polarisation is linear everywhere, each sub-level with the same $|m_J|$ has the same light shift everywhere. This configuration only provides an additional cooling mechanism for atoms with a ground state where $J \geq 1$ [46], because if $J = 1/2$ then both sub-levels have the same light-shift everywhere, there is no energetic advantage to either state.

An atom at rest may take its quantisation axis to be along the polarisation of the light field. Thus, the resulting linear polarisation excites $\pi$ transitions from $|g\rangle$ to $|e\rangle$, which, due to each transitions' different Clebsch-Gordan coefficients, causes de-excitations that put atoms preferentially in the $|g, 0\rangle$ sub-state, with equal but smaller proportions in each $|g, \pm 1\rangle$ sub-state. This also means that, since the $\pi$ transition to $m_J = 0$ is more intense than the $\pi$ transitions to $m_J = \pm 1$, that the light-shifts are also different: $\Delta_0 = 4/3\, \Delta_{\pm 1}$.

An atom moving along the $z$-axis will see the polarisation vector rotate around the $z$-axis. In a frame that rotates with the electric field vector, the atom sees a fixed electric field and a fictitious magnetic field due to the frame transformation, which adds a term $\hat{V} = \hbar k \hat{J}_z$ to the Hamiltonian. This term produces a velocity-dependent, non-zero $\langle J_z \rangle$ (an "orientation"), which causes an imbalance of the population of the $m_J = \pm 1$ states, which causes preferential absorption from the beam that will even out the populations, providing a frictional force on the atoms. This cooling mechanism is known as motion-induced orientation cooling, and results in a temperature

$$k_B T_{\text{mio}} = \frac{\hbar \Gamma^2}{2|\delta|} \frac{I}{I_{\text{sat}}} \left( \frac{29}{300} + \frac{254}{75} \frac{\Gamma^2/4}{\delta^2 + (\Gamma^2/4)} \right), \tag{1.30}$$

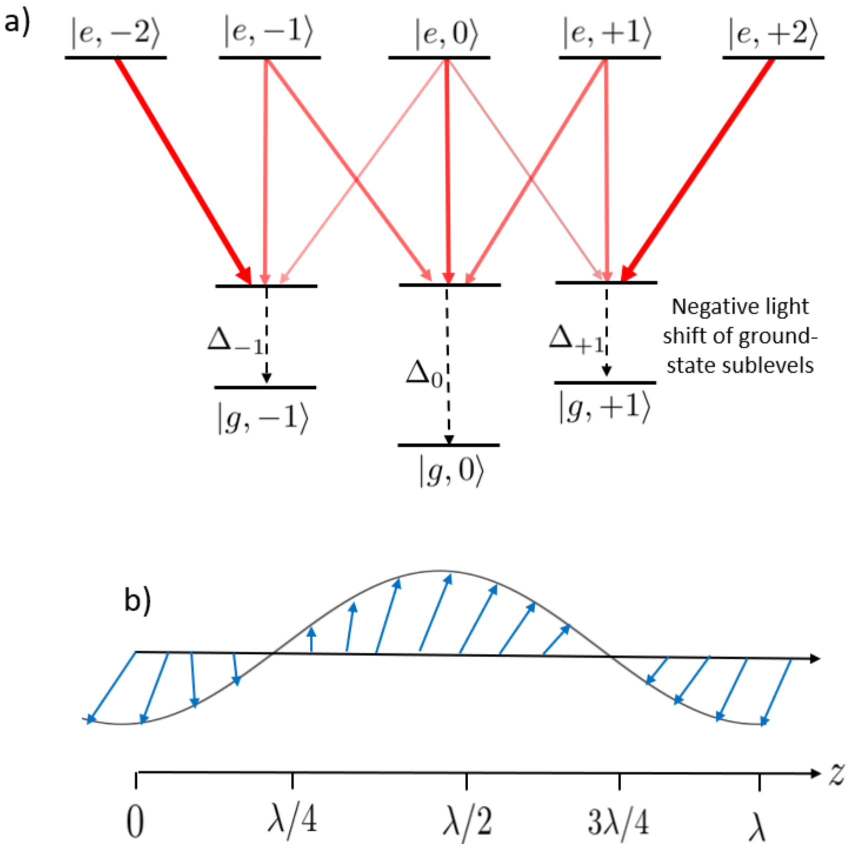

Figure 1.6: a) Energy-levels for a $J = 1 \rightarrow 2$ transition. The Clebsch-Gordan coefficients for each transition are indicated by the relative strength of each line. b) Counter-propagating $\sigma^+$ and $\sigma^-$ beams produce a helical standing wave that rotates with a period $\lambda$.

which, at large detunings $\delta \gg \Gamma$, reduces to

$$k_B T_{\text{mio}} \approx \frac{\hbar}{20} \frac{I}{I_{\text{sat}}} \frac{\Gamma^2}{|\delta|}, \tag{1.31}$$

which has the same form and similar magnitude as $T_S$. Both methods are limited to the recoil limit: an atom scattering a photon recoils with momentum $\hbar k$, such that the minimum temperature that can be achieved in, say, $^{133}$Cs using the 852 nm cooling transition, is 198 nK. In practice, temperatures this low are not easily achieved by these mechanisms, and the coldest atoms for use in Bose-Einstein condensation experiments are prepared using evaporative cooling [48].

### MOT Loading

The collected number of atoms $N$ in a MOT or molasses depends on the loading rate $R_L$, the trap lifetime due to thermal background gas collisions $\tau_L$, and also on the collision rate between cold atoms, which depends on the cold-collision cross-section $\beta$ and the trap volume $V$:

$$\frac{dN}{dt} = R_L - \frac{N}{\tau_L} - \beta\frac{N^2}{V}. \tag{1.32}$$

Neglecting cold-cold collisions (setting $\beta = 0$) and focusing only on the cold-thermal collisions, equation 1.32 can be solved, subject to the initial condition $N(0) = 0$, as

$$N(t) = R_L\tau_L(1 - e^{-t/\tau_L}), \tag{1.33}$$

which starts off rising linearly with a gradient of $R_L$, and eventually approaches a maximum of $R_L\tau_L$ after loading for several $\tau_L$. The lifetime can be estimated [42, 49] as

$$\tau_L = \frac{\sqrt{\pi m k_B T/8}}{\sigma_c P}, \tag{1.34}$$

where $P$ is the background gas pressure, $T$ is the background gas temperature, and $\sigma_c$ is the effective collision cross-section for loss from the trap, which depends on the species being trapped and the dominant background gas component. While there will inevitably be several different background gas species, such that the overall lifetime is contributed to by several different collision cross-sections, in practice ultra-high vacuum (UHV) systems that load a MOT have the background gas pressure typically at or below $10^{-9}$ Pa ($10^{-11}$ mbar), whereas the alkali metal vapour is typically around $10^{-7}$ Pa ($10^{-9}$ mbar), such that the dominant thermal background species is usually the thermal alkali atoms, for which the collision cross-section is particularly large [50]. The collisional cross-section is contributed to by the Van der Waals force for ground-state trapped atoms and by the much-larger resonant dipole-dipole force for excited-state atoms, particularly by glancing collisions in which only a small amount of energy is transferred from the fast atom to the slow atom. Experimental measurements result in a value for $\sigma_c$ somewhere between the ground and excited-state values [49], $\sigma_c = 20 \times 10^{-14}\,\mathrm{m}^2$ for Cs.

The loading rate has been estimated, using the kinetic theory of gasses, from the size of the trap, the temperature of the background vapour, and the frictional scattering force. The result is [42, 49]

$$R_L \approx n v_c^4 r^2 \left( \frac{m}{2k_B T} \right)^{3/2}, \tag{1.35}$$

where $n$ is the background density of thermal alkali atoms, the capture velocity

$$v_c = \sqrt{2ar}, \tag{1.36}$$

is the velocity of an atom that is slowed down to rest over a distance equal to the trap radius $r$ by half of the maximum scattering acceleration, which is $a = \hbar k \Gamma / 4m$. Combining equations 1.35 and 1.36 shows that $R_L \propto r^4$, such that the size of the trap, and thus, the radius of the trapping beams, is critical to the number of atoms loaded.

### Doppler-Free Spectroscopy

A key requirement of laser cooling experiments is to ensure the laser frequency stays resonant with the cooling transition. Short of being able to lase directly on an atomic transition of interest [51], this is done by locking a tunable laser to a reference cell. An example of a frequency-tunable laser is an extended-cavity diode laser (ECDL) [52, 53], where the cavity length, and hence the frequency, is controlled by a grating mounted on a piezoelectric translation stage. Another is a distributed Bragg reflector (DBR), where the grating is built into the structure of the diode [54, 55], and the frequency is controlled by the current to the diode.

However, when one scans a laser frequency across resonance and looks for absorption, all one would see is a broad, Gaussian lineshape which corresponds to the Doppler-broadening of the atomic transitions; the Maxwell distribution of atomic velocities at a temperature $T$ is

$$f_M(v) = \frac{1}{\sqrt{\pi} v_m} e^{-v^2/v_m^2}, \tag{1.37}$$

where the width of the Gaussian distribution is the most-probable speed $v_m = \sqrt{2k_B T/m}$, which for Cs at 300 K is 190 ms$^{-1}$. In an optical wave of 852 nm, this corresponds to a frequency width of 220 MHz, far larger than the 5.2 MHz natural

(Lorentzian) linewidth of each component of the Cs $D_2$ line.

Early experiments with He-Ne lasers used, in addition to the gas cell containing the He-Ne gain medium, a dilute cell of pure neon inside the optical cavity [56]. It was found that, when scanning the laser frequency by controlling the mirror separation, that a narrow peak (30 MHz across) in output power vs frequency was detected, which, due to power broadening by the intense beams inside the laser cavity, corresponds to about twice the homogenous linewidth. A similar experiment with the 3390 nm rotation-vibration line of methane resulted in a 150 kHz wide emission peak, dominated by homogenous broadening due to collisions between molecules [57]. This technique was extended to vapour cells external to the laser cavity [58], which, in principle, allows any suitable laser to be locked to any atomic transition within its emission range.

This method works according to the hole-burning model: an intense "pump" beam with $I \gtrsim I_{\text{sat}}$ and angular frequency $\omega$ interacts with atoms of velocity $v$ when $\omega_0 = \omega - kv$. This results in, for a two-level atom, steady-state populations in the lower level $N_1$ and upper level $N_2$ of

$$N_1(v,\omega) = f_M(v)\left(1 - \frac{1}{2}\frac{I}{I_{\text{sat}}}\frac{\Gamma^2/4}{(\omega-\omega_0-kv)^2 + (1+I/I_{\text{sat}})\Gamma^2/4}\right), \tag{1.38}$$

$$N_2(v,\omega) = f_M(v)\frac{1}{2}\frac{I}{I_{\text{sat}}}\frac{\Gamma^2/4}{(\omega-\omega_0-kv)^2 + (1+I/I_{\text{sat}})\Gamma^2/4}, \tag{1.39}$$

such that, at high $I$, there is equal population in the ground and excited states for the velocity class that the laser is resonant with. The absorption coefficient seen by a much weaker, counter-propagating "probe" beam of the same frequency is

$$\alpha(\omega) = \int_{-\infty}^{\infty} [N_1(v,\omega) - N_2(v,\omega)]\sigma(\omega - kv, I)dv, \tag{1.40}$$

which, for $N_1(v,\omega_0) = N_2(v,\omega_0)$, is zero. When the pump and probe are resonant with the transition, they are each seeking to interact with the atoms with $v = 0$, and the probe beam isn't absorbed. This apparent increase in transparency of the gas cell to the probe beam can be used as a frequency reference to lock the laser to, and has a linewidth

$$\delta\omega_{\text{SAS}} = \Gamma\sqrt{1 + \frac{I}{I_{\text{sat}}}}. \tag{1.41}$$

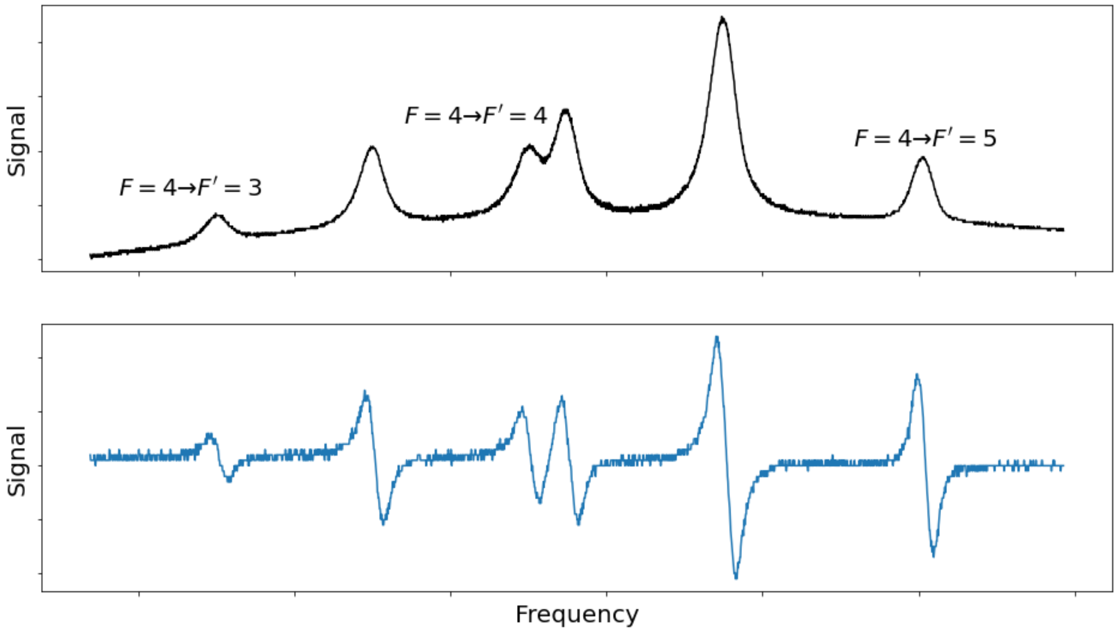

Figure 1.7: Top: saturated absorption spectrum for the $6\,^2S_{1/2}\,|F = 4\rangle \to 6\,^2P_{3/2}$ manifold of transitions in Cs. Bottom: the error signal fed to the laser controller is proportional to the gradient of the spectrum.

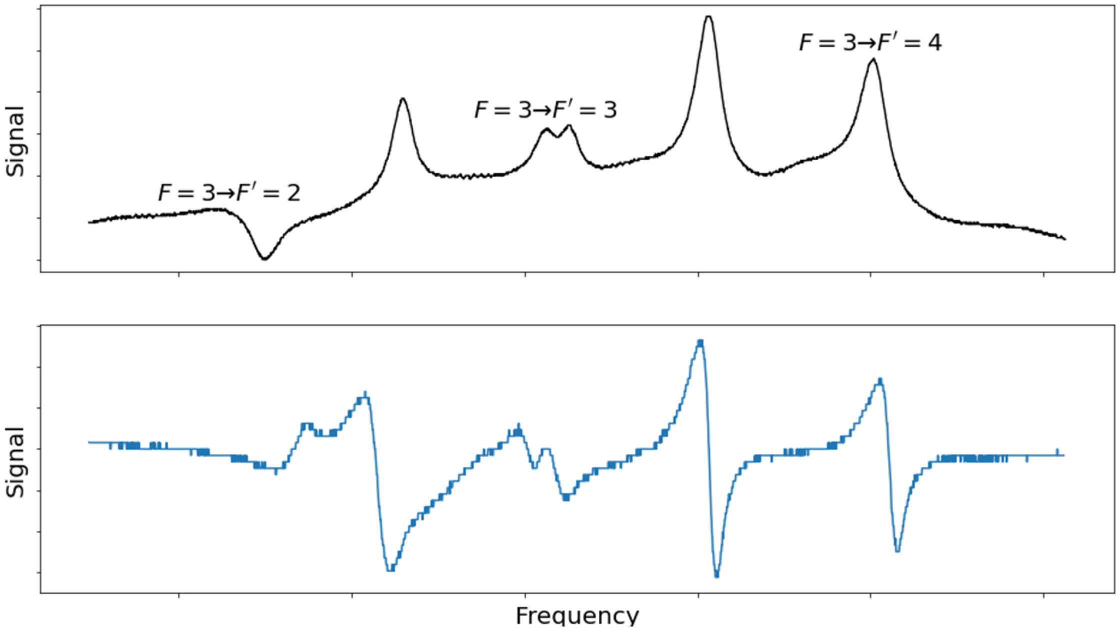

Figure 1.8: Top: saturated absorption spectrum for the $6\,^2S_{1/2}\,|F = 3\rangle \to 6\,^2P_{3/2}$ manifold of transitions in Cs. Bottom: the error signal fed to the laser controller is proportional to the gradient of the spectrum.

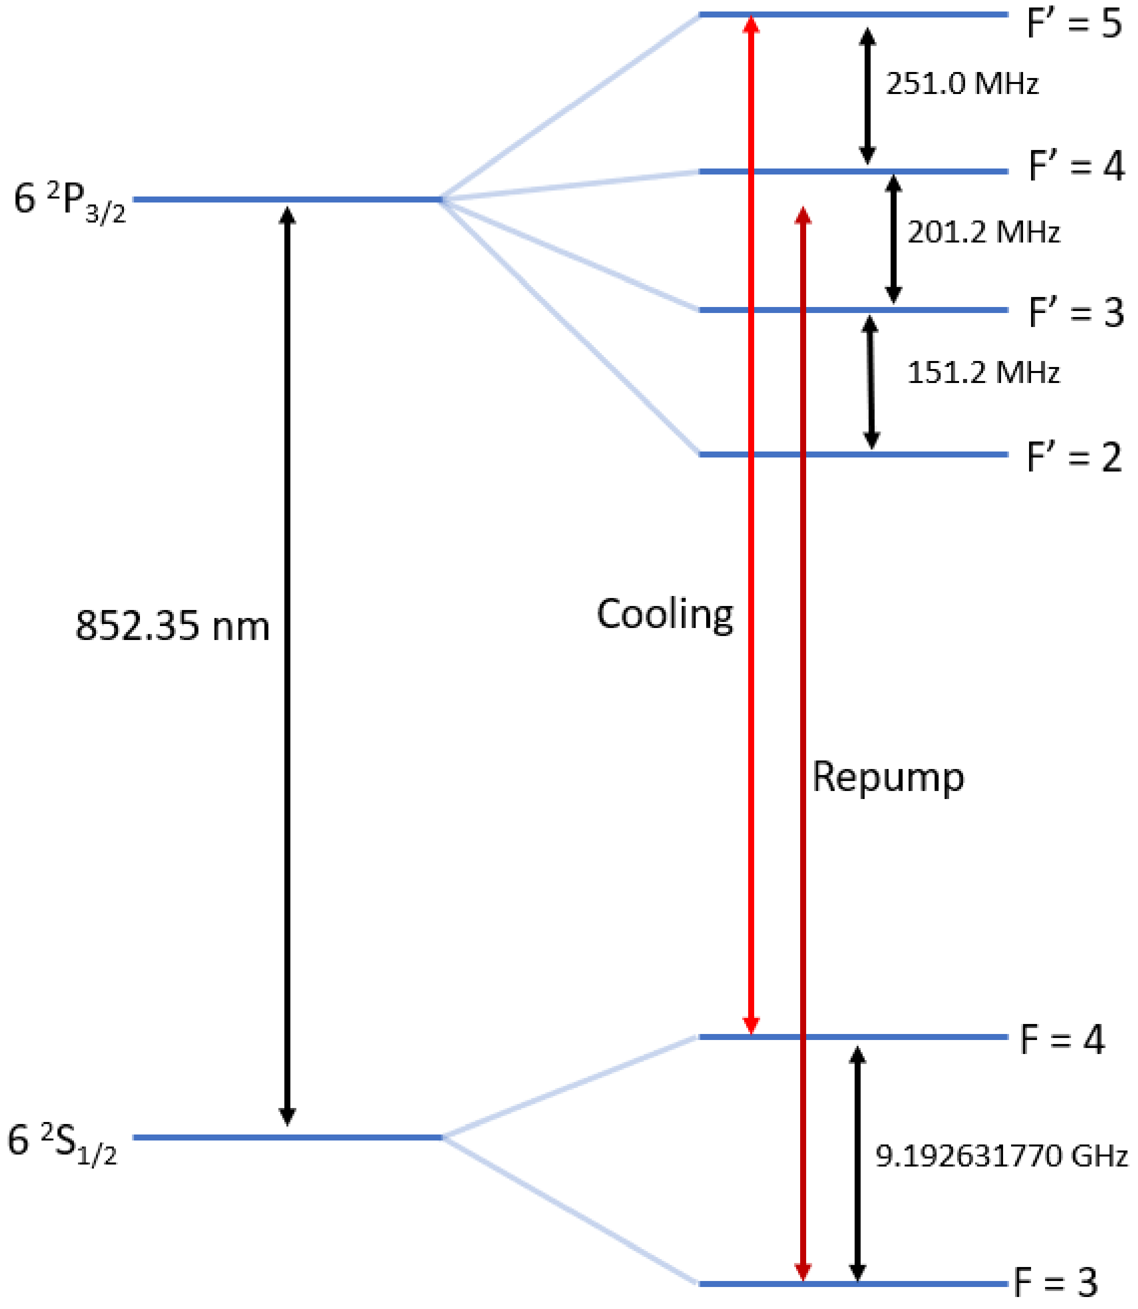

Figure 1.9: Energy level diagram of the $N = 6$ quantum number for $^{133}$Cs. The left shows the fine-structure slitting between the $6\,^2S_{1/2}$ and $6\,^2P_{3/2}$ states, and the right shows the further splitting into the hyperfine transitions.

Some example spectra are shown in figures 1.7 and 1.8, which correspond to the energy levels of figure 1.9.

In real saturated absorption experiments, a second probe beam, not overlapped with the pump, is often used to make a differential measurement that removes the Doppler-broadened background and common-mode laser intensity noise. Additionally, real multi-level atoms have several resonant frequencies corresponding to different transitions, such that for transition frequencies separated by less than the Doppler width, there will be an additional cross-over peak at a frequency exactly halfway between the transitions. This is because the pump and probe are counter-propagating, such that there is a velocity at which the positive Doppler shift to one beam is resonant with one transition and the negative Doppler shift to the other beam is resonant with the other transition.

A more detailed theory of saturated absorption [59], predicts some phenomena not described by the simple hole-burning model, including amplification of the probe by stimulated emission in some circumstances, such that the direction of the absorption feature changes sign, as in the $F = 3 \to F' = 2$ transition of figure 1.8.

In order to lock a laser frequency $\omega_L$ to a saturated absorption peak, the laser frequency must be modulated such that the peak height at $\omega_L - \delta\omega$ is compared to the peak height at $\omega_L + \delta\omega$. This can be done by either modulating the laser frequency directly, by, for example, modulating the current to a DBR laser, or by modulating the light by an electro-optic modulator (EOM) or acousto-optic modulator (AOM). The difference between the signal levels is used as an error signal (proportional to the gradient of the spectrum), which an electronic servo system uses to continuously keep the laser frequency on resonance.

### 1.1.4 Early Atomic Fountains

The first atomic fountain clock was suggested in 1953 as an upgrade to thermal beam clocks: atoms would be directed upwards out of an oven and pass through the same microwave waveguide twice, once on the way up and again on the way back down after reversing direction. This would be less susceptible to distributed cavity phase shifts than if the atoms passed through two separate cavities. However, in the initial attempts, there were not enough slow atoms emitted from an oven to detect a returning signal.

Laser cooling, which allows large numbers of atoms to be cooled to a few microkelvin, renewed interest in atomic fountain clocks [60]. The very first atomic fountain clock used sodium atoms cooled to $50\,\mu$K, performed Ramsey interrogation with two microwave pulses inside a waveguide, detected the returning atoms by ionisation, and obtained a Ramsey linewidth of 2.0 Hz [61]. Once it was demonstrated that laser cooling could produce very cold clouds of caesium atoms [62], the first caesium fountain improved upon this in several key areas [63]. Firstly, by launching using a moving-molasses technique: if the down-going beam is red-detuned, and the up-going beam blue-detuned, by an amount $\delta\nu$, then the atoms will be accelerated into a frame moving upwards at

$$v = \lambda\, \delta\nu, \tag{1.42}$$

in which they still see a molasses that provides cooling. The result was that the cloud of Cs atoms was cooled after launch down to $5.5\,\mu$K. This allowed the first caesium fountain to launch atoms through and above the microwave cavity, which prolongs the Ramsey interrogation time. Secondly, a $\mathrm{TE}_{102}$ microwave cavity, as opposed to a propagating waveguide, reduced the travelling-wave component of the microwave field, such that cavity phase shift errors were reduced. Additionally, a single beam for normalised optical detection, as opposed to hot-wire ionisation detection, improved the signal-to-noise ratio (SNR), which was limited by the quartz local oscillator from which the microwaves were synthesised. The short-term stability was $3 \times 10^{-12}/\sqrt{\tau/\mathrm{s}}$, where $\tau$ is the averaging time, which was comparable to the thermal beam clocks of the time. This was refined in a subsequent design that launched the atoms higher for a 500 ms Ramsey time (30 cm above the cavity), such that a 1.0 Hz FWHM linewidth was achieved [64]. The single-feed cavity was replaced by a symmetrically-fed $\mathrm{TE}_{011}$ cavity, which further reduces the distributed cavity phase shift. Separated optical detection regions for the atoms in the upper and lower hyperfine states, coupled with a more sophisticated quartz oscillator phase-locked to a hydrogen maser, improved the SNR, such that the frequency stability was $3 \times 10^{-13}/\sqrt{\tau/\mathrm{s}}$, an order-of-magnitude improvement on the previous design. Once the quartz local oscillator was replaced with a cryogenic sapphire oscillator (CSO) [65], short-term stabilities were limited by quantum noise inherent to the measurement process, at $4 \times 10^{-14}/\sqrt{\tau/\mathrm{s}}$ [66].

Today, atomic fountain clocks take much the same form as in figure 1.10c. However, virtually every atomic fountain clock in use at national metrology institutes has an additional state-selection cavity positioned just below the Ramsey cavity, rather than using a microwave pulse from a horn or antenna, as this minimises unwanted evolution of the atomic state due to stray microwave fields.

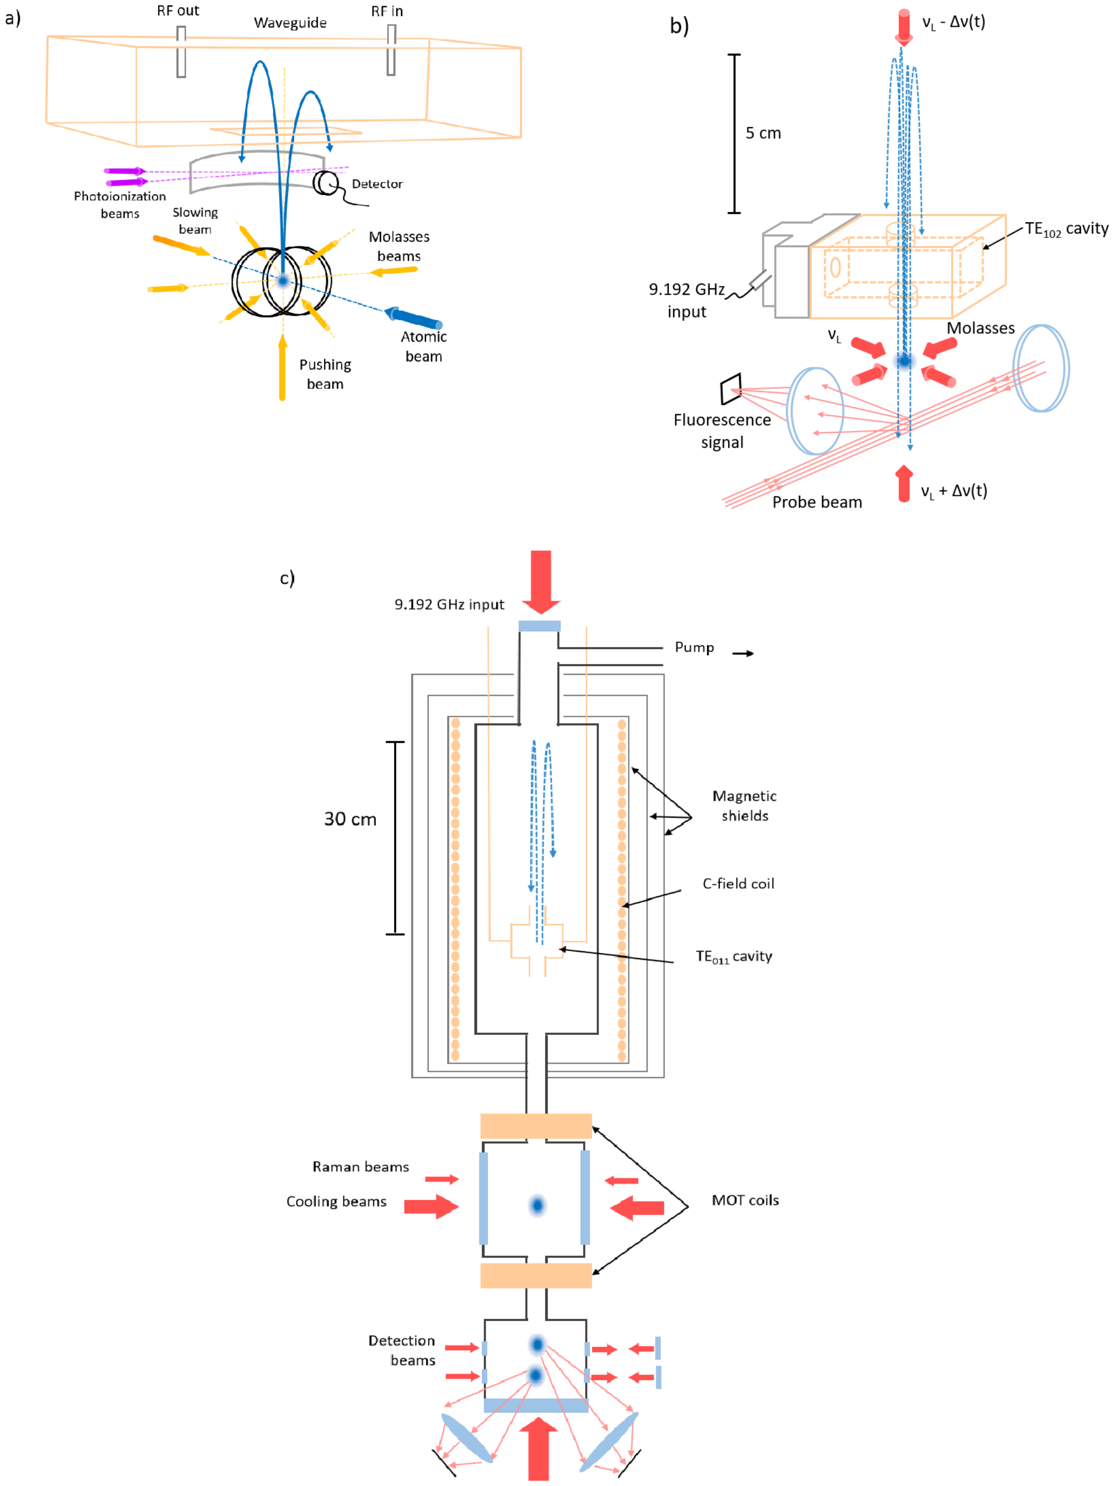

Figure 1.10: a) Initial Na fountain [61]. b) Initial Cs fountain [63], with a cavity and optical detection. c) Refined Cs fountain [64], with a symmetrically-fed cavity and separated optical detection regions.

State-of-the-art atomic fountain clocks [67, 68, 69] achieve short-term fractional frequency stabilities of a few $10^{-14}/\sqrt{\tau/\mathrm{s}}$, limited by quantum-projection noise, and long-term stabilities of a few $10^{-16}$, limited by environmental perturbations to the evaluation of the systematic shifts. The accuracies are also of a few parts in $10^{16}$, limited by the uncertainties in the systematic shifts.

### 1.1.5 Fountain Operation

Fountain clocks operate in a periodic, cyclic manner (with one notable exception [70]). This cycle can be broadly broken down into the following six stages: loading, launch, sub-Doppler cooling, state preparation, interrogation, and detection.

In the loading stage, a cloud of cold caesium atoms is prepared by laser cooling on the $6\,^2S_{1/2}\,|F=4\rangle \rightarrow 6\,^2P_{3/2}\,|F=5\rangle$ transition, with a small amount of repump light resonant with the $6\,^2S_{1/2}\,|F=3\rangle \rightarrow 6\,^2P_{3/2}\,|F'=4\rangle$, to keep atoms in the cooling cycle. Some fountains use only an optical molasses to load a relatively diffuse cloud from a background vapour [71, 72]. Others use a magneto-optical trap to load a denser cloud from a similar background vapour [73, 74, 75, 76]. Some use a combination of a MOT and a molasses. For example, retro-reflecting three perpendicular beams to form a MOT, but with a small hole in one of the mirrors, causes an imbalance in the radiation force in-line with the hole, forming a cold atom beam [77, 78]. This is called a low-velocity intense-source (LVIS). Similarly, additional coils can create a line of zero field aligned with the pinhole, to increase the flux of cold atoms [79, 80]. This is called a 2D MOT. This beam can be directed into a region of optical molasses where it is slowed and further cooled [81, 82]. Advantages of the multi-stage process are a larger number of loaded atoms (or conversely, a shorter loading time). A disadvantage is that they requires a larger, more-complex system. On the other hand, another key consideration is the density of the cloud. An advantage of the molasses-based methods is that the cloud density is lower, and more uniform, such that the collisional shift (section 1.2.4) is correspondingly smaller [82].

After loading, the cloud is launched. This is done by the moving-molasses technique, which minimises unwanted heating [63]. The frequency and intensity of each cooling beam is manipulated by acousto-optic modulators (AOMs) to achieve the appropriate detuning for the desired launch height and beam geometry. The cloud achieves the launch speed in only a few milliseconds. Some fountains allow the cloud to drop for a few milliseconds before beginning the launch procedure, to improve the overlap between the beams and the cloud [72, 82].

Once the cloud is moving, but still within the beam region, is is further cooled. The sub-Doppler cooling mechanisms achieve the lowest temperatures for large red detuning and low intensity, such that the AOMs are used to simultaneously increase the detuning by adjusting the input acoustic frequency, and reduce the intensity by reducing the input amplitude of the acoustic wave.

State preparation takes the launched cloud from being in an approximately equal mixture of all 9 $|4, m_F\rangle$ states, to the lower clock state, $|3, 0\rangle$. Firstly, some fountains increase the population of the $|4, 0\rangle$ state by optical pumping [83, 84]. A beam with vertical, linear polarisation (aligned with the vertical background magnetic field) resonant with the $6\ ^2S_{1/2}\,|F = 4\rangle \rightarrow 6\ ^2P_{3/2}\,|F' = 4\rangle$ transition can excite all the lower $|F = 4, m_F\rangle$ states into upper $|F' = 4, m_{F'}\rangle$ states, except for $|F = 4, m_F = 0\rangle$, as the $|F = 4, m_F = 0\rangle \rightarrow |F' = 4, m_{F'} = 0\rangle$ transition is forbidden by the selection rules. This means that once an atom is in this dark state, it stays there. It was found that around 20 scattered photons was sufficient to improve the clock-state population by a factor of around 7 [83, 84], with an associated increase in the cloud temperature that limited the improvement in returning signal after the ballistic flight to a factor of 4. The main step is to transfer the upper clock-state atoms to the lower clock state by a single microwave $\pi$-pulse (a pulse that rotates the atom's Bloch vector by $180°$) resonant with the clock transition. Afterwards, any atoms that remain in the upper level are then removed from the cloud by the radiation pressure of a short laser pulse resonant with the cooling transition, such that the only atoms continuing the ballistic flight are the lower clock-state atoms.

Interrogation is performed by the Ramsey method. On the way up through the Ramsey cavity, the cloud undergoes a nominal $\pi/2$ pulse (rotates the atom's sate vector by $90°$), such that it is put in a superposition of each clock state. The phase of this superposition evolves as the cloud reaches its apogee and returns again through the Ramsey cavity and undergoes another nominal $\pi/2$ pulse. If the microwave frequency was on-resonance, then the atomic population finishes entirely in $|4, 0\rangle$, however usually the microwave frequency is deliberately detuned such that the final result is an equal superposition of $|3, 0\rangle$ and $|4, 0\rangle$, which enables the detected signals to respond more sensitively to small changes in the microwave frequency.

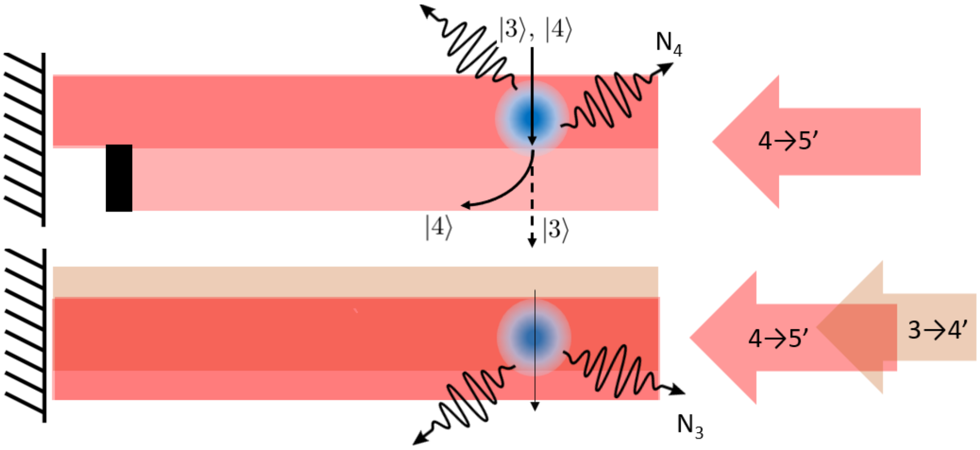

Figure 1.11: Illustration of the separated detection regions for fluorescence detection of $|F = 4\rangle$ (upper) and $|F = 3\rangle$ (lower) hyperfine states.

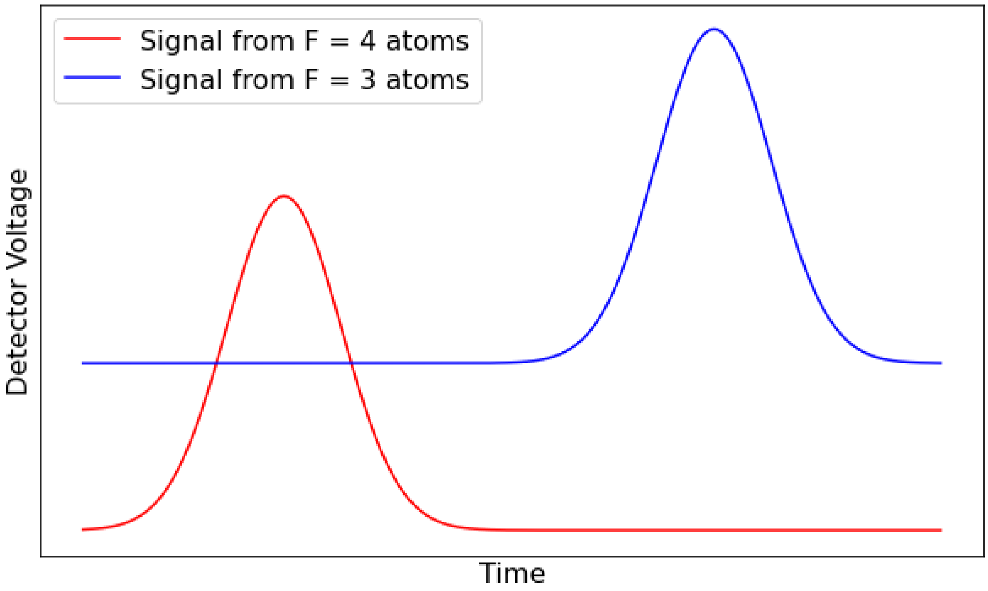

Figure 1.12: Illustration of the signals output by each detection photodiode. The different background level is because, in a fountain that uses a single-stage MOT trap, there is a higher density of background atoms in the lower ($F = 3$) detection region, because it is closer to the MOT region.

Detection is performed with two sets of laser beams, which are usually intensity-stabilised to reduce the effect of laser intensity noise [63]. The following scheme is illustrated in figure 1.11. The first beam is of the cooling light, $F = 4 \rightarrow F' = 5$, where atoms in the upper state fluoresce. This fluorescence is collected by a lens and focused onto a photodiode, with the area under the resulting voltage vs time plot (illustrated in figure 1.12) being proportional to the atom number in the upper state, $N_4$. This beam is partly retro-reflected such that, once a signal has been obtained after a few milliseconds, the atoms are accelerated until they are out of resonance with the light. The second beam is also a retro-reflected beam of the $F = 4 \to F' = 5$ cooling light, but with a small amount of $F = 3 \to F' = 4$ repump light overlapped, such that the atoms in the lower level are quickly repumped, and can be efficiently detected by fluorescence on the cooling transition to obtain $N_3$. The transition probability for the cycle is then

$$P = \frac{N_4}{N_3 + N_4}. \tag{1.43}$$

### 1.1.6 Fountain Short-Term Stability

The Allan deviation (sometimes known as the two-sample deviation, or just ADEV [85]) is a measure of the frequency instability of a clock. It is most commonly used to analyse the fractional frequency fluctuation $y(t)$, which is defined as

$$y(t) = \frac{\nu(t) - \nu_0}{\nu_0}, \tag{1.44}$$

where $\nu_0$ is the nominal frequency of the oscillator, and $\nu(t)$ is the instantaneous frequency at time $t$. In practice, instantaneous frequency is not an experimentally accessible quantity, so an average $\bar{y}(t)$ is taken over an interval $\tau$. The Allan deviation is defined as

$$\sigma_y(\tau) = \sqrt{\frac{1}{2(M-1)} \sum_{i=1}^{M-1} (\bar{y}_{i+1} - \bar{y}_i)^2}, \tag{1.45}$$

where $\bar{y}_{i+1}$ and $\bar{y}_i$ are adjacent fractional frequencies averaged over a time $\tau$, and $M$ is the number of fractional frequency values for each $\tau$. Put simply, it is half the average difference between two fractional frequency measurements made a time $\tau$ apart. An advantage of the ADEV over the standard deviation is that it does not diverge for non-stationary noise, such as linear frequency drift. There are variants of the ADEV, such as the overlapping-Allan deviation, which generates additional fractional frequency data by taking a moving average, such that the confidence at large $\tau$ appears higher. There is also the modified-Allan deviation, which better discriminates between white phase noise and flicker phase noise at short $\tau$ [86].

The frequency at which an atomic fountain probes the transition is alternated between $\nu_0 + \delta\nu$ and $\nu_0 - \delta\nu$, where $\delta\nu$ is the half-width-half-maximum (HWHM) of the Ramsey fringe, such that the returning atoms are detected in an equal superposition of the upper and lower clock states. The difference between successive observed transition probabilities is then used as an error signal to drive the oscillator frequency towards resonance. The Allan deviation for an atomic fountain can be expressed as [66]

$$\sigma_y(\tau) = \frac{2}{\pi Q_R}\sqrt{\frac{T_c}{\tau}}\left(\frac{1}{N_{\mathrm{at}}} + \frac{1}{N_{\mathrm{at}}\epsilon n_\gamma} + \frac{2\sigma^2_{\mathrm{det}}(N_{\mathrm{at}})}{N^2_{\mathrm{at}}} + \sigma_\gamma\right)^{1/2}, \tag{1.46}$$

where $T_c$ is the total fountain cycle duration, $\tau$ is the averaging time, and $Q_R$ is the resulting quality factor of the Ramsey interaction. The term in brackets is $1/\mathrm{SNR}^2$ where "signal" is the Ramsey fringe contrast and "noise" is the shot-to-shot ADEV of the detected transition probability, and has contributions from different sources. The first term $1/N_{\mathrm{at}}$ is quantum-projection noise [87]: for a number of atoms $N_{\mathrm{at}}$ in an equal superposition of the upper and lower clock states, the quantum fluctuations in the detected atom number are maximised and given by binomial statistics. The second term is the photon shot noise, which is similar to quantum-projection noise but much smaller, as the collection efficiency $\epsilon$ multiplied by the number of photons scattered per atom $n_\gamma$ is typically large. The third term in equation (1.46) is the uncorrelated noise in the detection system, where $\sigma_{\mathrm{det}}$ has contributions from electrical noise and background fluorescence (which are independent of cold-atom signal size) and laser intensity and frequency noise (which may be proportional to the cold-atom signal, and thus present an upper limit to the SNR regardless of how many atoms are detected).

The final term, $\sigma_\gamma$, is the local oscillator noise. Instability in the local oscillator frequency, due to ageing [88, 89] or environmental effects, will cause instability in the detected transition probability. However this is only part of the picture [90]. Because atomic fountains operate in a pulsed manner with period $T_c$, it can be shown [91, 92] that components of the local oscillator spectrum with frequencies that are multiples of $1/T_c$ are down-converted into frequency noise at the output of the locked oscillator [93]. Two ways to reduce the severity of this effect are to maximise the duty cycle, $T/T_c$ where $T$ is the Ramsey interrogation time, and to use a local oscillator with a low power spectral density at $f = 1/T_c$ and multiples thereof.

## 1.2 Systematic Shifts

In order to be called a *primary* frequency standard, the output frequency from the standard must be precisely related back to the internationally-agreed definition of the SI unit of frequency. For a caesium fountain clock, this means evaluating all the systematic effects to a level of uncertainty similar to the long-term instability of the output frequency, which is typically at the $10^{-16}$ level. Some effects produce large frequency shifts that are measured accurately, such that their uncertainty contribution becomes small; others are so small that their shift is about zero, but only verified to the limited precision of the measurement. To allow accurate comparisons between caesium fountain clocks, each must report all its systematic offsets and uncertainties [76, 71, 69].

### 1.2.1 Second-Order Zeeman Shift

The largest frequency shift is due to the background magnetic field. It is not desirable (or feasible) to maintain exactly zero field over the entire interaction region, because one needs to separate the clock transition, $|F=3, m_F=0\rangle \rightarrow |F=4, m_F=0\rangle$ from other $m_F \neq 0$ transitions. The background field (often called the "C-field") must be homogenous enough that the second-order Zeeman shift to the clock transition can be accurately evaluated. In the presence of a weak magnetic field, the frequency shift to the clock transition in caesium is

$$\frac{\Delta\nu_{Z2}}{\nu_0} = \frac{1}{2}\left(\frac{(g_J - g_I)\mu_B}{h\nu_0}\right)^2 \langle B^2\rangle \approx 4.63 \left\langle \frac{B}{[T]} \right\rangle^2, \tag{1.47}$$

where $\langle B^2\rangle$ is the average over the atomic trajectory between cavity passes. A typical mean C-field strength of 100 nT results in a fractional frequency shift of $4 \times 10^{-14}$. The most common way of measuring the C-field above the cavity is to launch the atoms to incrementally higher heights while tracking the position of the central Ramsey fringe of a transition with first-order magnetic sensitivity, usually $|3,1\rangle \rightarrow |4,1\rangle$ [73, 74], which, by equation 1.11, is detuned from the clock transition by

$$\Delta\nu_{1-1} = \frac{(g_J - g_I)\mu_B\langle B\rangle}{4h}. \tag{1.48}$$

The launch height must initially be only just to the cavity and increase slowly to unambiguously identify the central Ramsey fringe at each height. If the mean field in the cavity is sufficiently different to the mean field above the cavity, then the central Ramsey fringe may not even fall within the Rabi envelope at all [94]. One can also calculate the standard deviation, $\sigma_B$, of the C-field, as to calculate the field strength from equation 1.48 gives $\langle B \rangle^2$, not $\langle B^2 \rangle$, which are related by

$$\langle B^2 \rangle = \langle B \rangle^2 + \sigma_B^2, \tag{1.49}$$

such that the second-order Zeeman shift can be re-expressed in terms of the measurable quantities $\Delta\nu_{1-1}$ and $\sigma_B$ (at the full launch height) as

$$\frac{\Delta\nu_{Z2}}{\nu_0} = 8\left(\frac{\Delta\nu_{1-1}}{\nu_0}\right)^2 + 4.63\left(\frac{\sigma_B}{[T]}\right)^2, \tag{1.50}$$

such that the fractional uncertainty on the second-order Zeeman shift is taken as

$$\sigma_{Z2} = 16\frac{\Delta\nu_{1-1}}{\nu_0^2}\sigma_\nu, \tag{1.51}$$

where $\sigma_\nu$ is the uncertainty of the frequency of the central Ramsey fringe of the $|3,1\rangle \rightarrow |4,1\rangle$ transition, typically taken as 1 Hz. The detuning of the $|3,1\rangle \rightarrow |4,1\rangle$ transition in a C-field of 100 nT would be 700 Hz, such that the fractional frequency uncertainty on the clock transition due to the second order Zeeman shift is around $10^{-16}$ or less. The inhomogeneities of the C-field are typically around 1 nT, such that the shift associated with the second term in equation 1.50 is approximately $10^{-18}$, which is negligible.

Factors affecting the temporal stability of the C-field are the stability of the coil current, and any transient or accrued magnetisation of the high-permeability shields that surround the coil. It is common to periodically interrupt normal fountain operation to check the C-field.

There are other methods of evaluating the C-field strength. When varying the launch height, a separate antenna can direct microwaves into the vacuum above the cavity, as in figure 1.13b. A peak in transition probability vs frequency can identify the $|3,1\rangle \rightarrow |4,1\rangle$ transition, and thus the field strength, at the atoms' apogee [73, 95]. Another method involves using a coil to generate a lower-frequency (100s

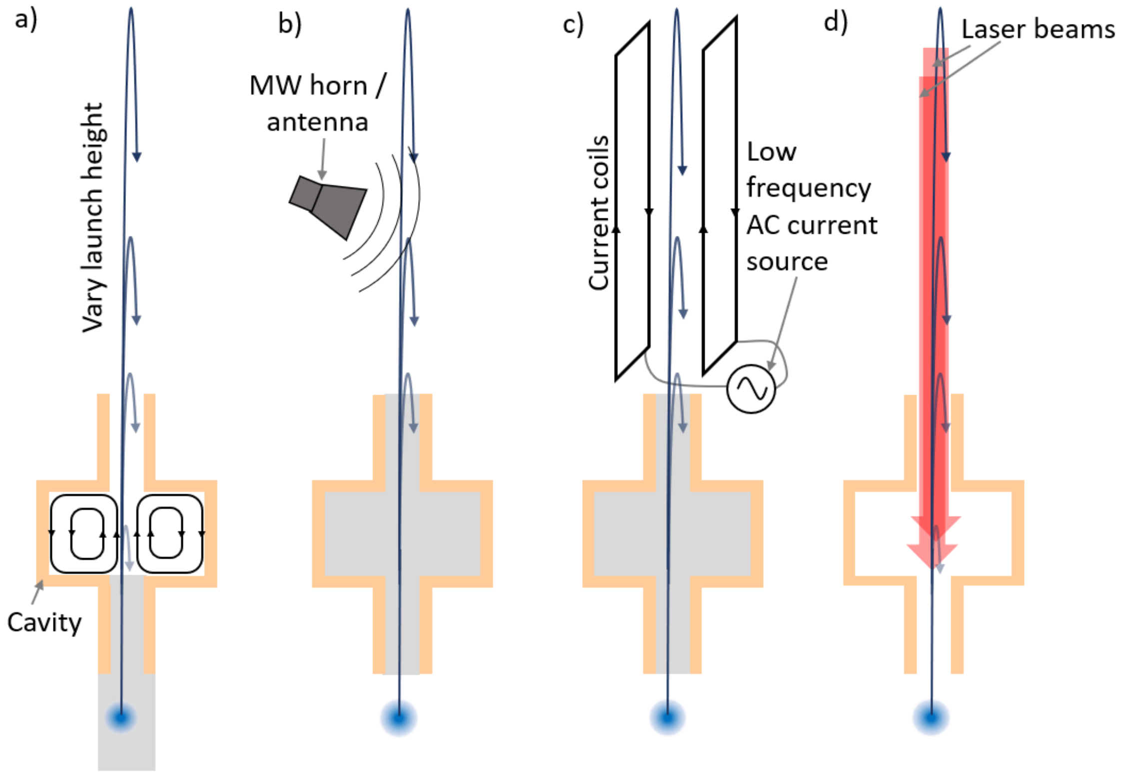

Figure 1.13: a) Standard C-field evaluation using the Ramsey cavity on a magnetically-sensitive transition. b) Microwave horn directed into interrogation region. c) Low-frequency coil to excite $\Delta m_F$ transitions. d) Two beams for stimulated Raman transitions on magnetically-sensitive transition. Greyed-out regions represent the inaccessible areas of each method.

of Hertz) oscillating magnetic field transverse to the C-field (figure 1.13c) [96], such that transitions between the lower clock state $m_F$ levels are excited at frequencies directly proportional to the field strength. In these two methods, there is no scope for mis-identifying any fringes, however these methods cannot map the field inside the cut-off waveguides immediately above and below the cavity, as the microwaves and radio waves can not propagate into these regions.

Another method is to employ two overlapping, co-propagating beams directed vertically along the atomic trajectory (figure 1.13d), with a frequency splitting that can be scanned around the $|3,1\rangle \rightarrow |4,1\rangle$ transition frequency [97]. A brief pulse of these beams will cause the atoms to undergo stimulated Raman transitions [98]. This method is unambiguous and can be applied to the entire flight region, but requires additional laser beams and optical access.

### 1.2.2 Black-Body Radiation and Light Shift

Atoms in the interrogation region are immersed in the blackbody radiation emitted from their surroundings, and the corresponding electric field perturbs the atomic energy levels. This large systematic shift requires accurate and precise evaluation.

Initially, the DC Stark shift for an external, static electric field was investigated [99, 100], which, for $S$ states, is a small quadratic dependence that was measurable with an external field applied to a beam clock [101]. It was subsequently realised that blackbody radiation produces a similar effect [102]. The energy density of the electromagnetic fields for temperature $T$ at a frequency $\nu$ are given by the Plank radiation law [103] as

$$u(\nu) = \frac{8\pi h\nu^3}{c^3}\frac{1}{e^{\hbar\omega/k_B T} - 1}. \tag{1.52}$$

While this energy is equally electric and magnetic, the effect of the magnetic field is negligible [102] [104]. Since infra-red blackbody radiation is much shorter in wavelength than microwaves, the shift to the clock transition is approximated by that of a static electric field with a strength that matches the root-mean-square of the radiation's field strength [105], which at 300 K is $\sqrt{\langle E^2\rangle} = 831.94\,\mathrm{V/m}$. Using third-order perturbation theory [104], the shift was calculated to be

$$\frac{\Delta\nu_{\mathrm{BBR}}}{\nu_0} = \beta\left(\frac{T}{T_0}\right)^4\left(1 + \epsilon\left(\frac{T}{T_0}\right)^2\right), \tag{1.53}$$

where $T_0$ is conventionally taken as 300 K and

$$\beta = -\frac{8}{7}\frac{\alpha_{10}}{h\nu_0}\langle E^2\rangle = -K_{\mathrm{Stark}}\langle E^2\rangle, \tag{1.54}$$

is due to the DC field which depends on the atom's scalar differential polarisability $\alpha_{10}$, and $\epsilon$ gives a small correction for the frequency spread of the radiation. In caesium, the calculations yielded $\beta = -1.69 \times 10^{-14}$ and $\epsilon = 0.014$, which was verified first directly by a Cs beam clock [106], later indirectly by an atomic fountain using an additional electric field (figure 1.14a) with an uncertainty of 0.2% [107], and later directly by surrounding the interaction region with a heated graphite tube (figure 1.14b) [108] [109]. The value of $\epsilon$ is only accurate to around 10%, which adds negligible additional uncertainty [110].

A later calculation of $\beta$, using more up-to-date values of relevant transition fre-

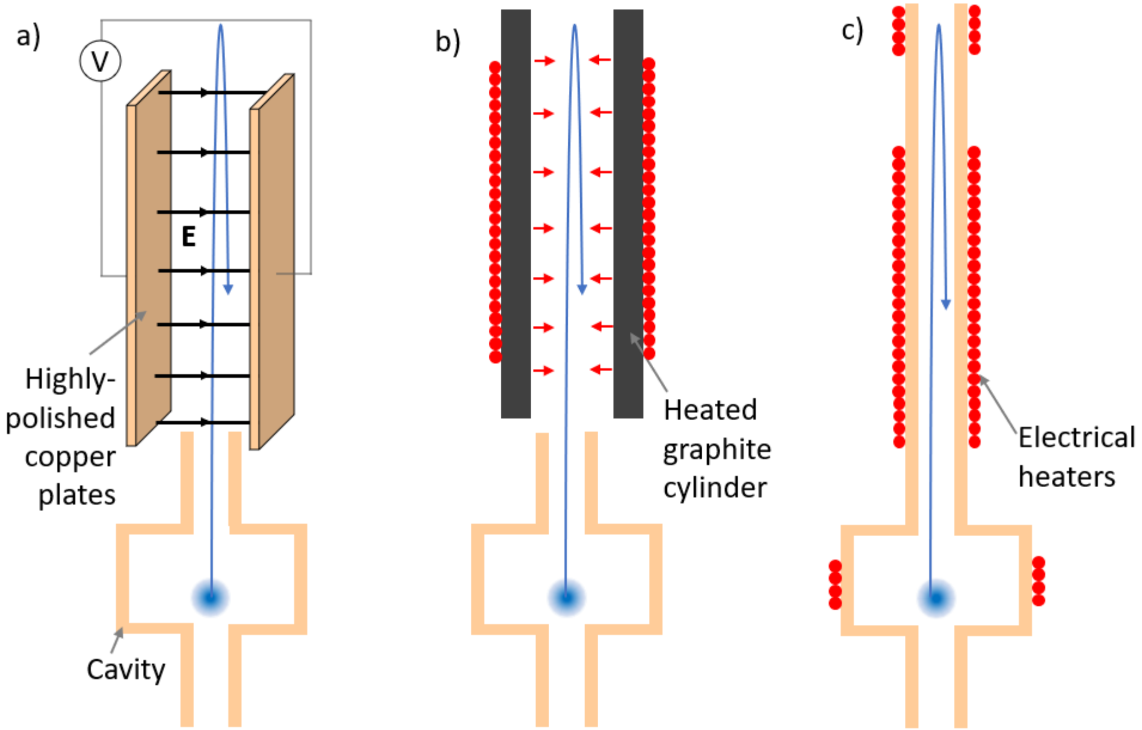

Figure 1.14: a) Indirect measurement of $K_{\text{Stark}}$ with a uniform electric field [107]. b) Direct measurement using a heated graphite cylinder [108][109]. c) Direct measurement by heating the cavity and flight tube [111].

quencies and dipole matrix elements, calculated a value that was lower than the previously confirmed value by 12%, with an estimated uncertainty of 5% [112]. This new value was seemingly corroborated by direct experimentation (figure 1.14c), by adjusting the temperature of the microwave cavity and flight tube [111]. This caused some consternation, because the frequency difference that arose by using each reported value of $\beta$ was $2 \times 10^{-15}$, far larger than the targeted systematic uncertainties.

This displeasing state of affairs continued until more attention was paid to the theory. Many-body relativistic calculations, using a complete set of intermediate states [113] [114] (rather than approximating with only lower-order terms) resulted in predictions for $\beta$ consistent with the original values, with the more-precise calculation resulting in $\beta = -1.710(6) \times 10^{-14}$ [113] becoming the currently-accepted value in caesium. Similarly, the value of $\epsilon$ is currently taken as $\epsilon = 0.013(1)$.

The resulting uncertainty in the blackbody correction can thus be expressed as

$$\sigma_{\mathrm{BBR}} \approx \sqrt{\left(\frac{T}{T_0}\right)^8 \sigma_\beta^2 + \beta^2 \left(\frac{T}{T_0}\right)^{12} \sigma_\epsilon^2 + \left(4\beta \frac{T^3}{T_0^4}\right)^2 \sigma_T^2}, \tag{1.55}$$

which means that, at $T = 300\,\mathrm{K}$, and $\sigma_T = 0.5\,\mathrm{K}$, the uncertainty is $\sigma_{\mathrm{BBR}} = 1.3 \times 10^{-16}$, around 1% of the size of the shift.

Characterising the BBR shift in a fountain clock can be carried out by attaching some well-calibrated thermocouples to the Ramsey cavity and flight tube, and plugging the resulting temperature into equation 1.53. Since these parts are typically made from high-conductivity copper, the measured temperatures typically agree with each other to around 0.2 K or better. However, it is not always the case that the entire solid angle the atoms see has the same temperature or emissivity. For example, some fountains used liquid nitrogen to cool the interrogation region down to around 80 K, as the uncertainty due to the blackbody radiation shift eventually became one of the largest, limiting uncertainties. Some room-temperature photons still made it through the opening in the cavity [115]. The actual temperature used in equation 1.53 is a time-average over the solid-angle weighted radiation temperature, which in the case of NIST's cryogenic fountain, raised the effective blackbody temperature by 0.3 K.

Higher frequency optical radiation from the cooling lasers also produces a light shift via the AC Stark shift. Many atomic fountain clocks use a pair of vertical cooling beams that are present through the entire interrogation. If these are not fully extinguished during interrogation, then a systematic shift occurs. For a cooling beam of intensity $I$ and detuning $\delta$ from resonance with a transition of saturation intensity $I_{\mathrm{sat}}$ and natural linewidth $\Gamma$, the light-shift is [24]

$$\frac{\Delta\nu_{\mathrm{ls}}}{\nu_0} = \frac{\Gamma^2}{8\delta} \frac{I}{I_{\mathrm{sat}}} \frac{1}{\nu_0}. \tag{1.56}$$

To minimise this effect, the intensity is reduced as far as possible with AOMs, the detuning is made as large as possible, and any residual light that may still pass through the AOM is blocked by a mechanical shutter. This extinction is confirmed by reducing the detuning of the laser light and looking for a change in the clock frequency, and is typically ruled-out to the $10^{-17}$ level or better.

### 1.2.3 Distributed Cavity Phase Shift

One of the motivations for fountains as upgrades over thermal beam clocks was eliminating the frequency shift as a result of the atoms passing through two separate cavities with slightly different phase profiles. However, even when passing twice through the same cavity, a frequency shift can occur because of the finite-temperature cloud expanding from a few millimetres on the upwards passage to over 1 cm on the downwards passage. Any transverse variations in the phase of the microwave field in the cavity causes an atom to see a microwave field of different phase from one passage to the next.

Phase shifts between cavity passes manifest as frequency shifts. For a relative phase shift of $\Delta\phi$ between cavity passes, the transition probability from equation 1.3 can be re-written [116] as

$$\begin{aligned} P_{\text{Ramsey}} &= 4\left(\frac{\omega_0}{a}\frac{B_1}{B_0}\right)^2 \sin^2\left(\frac{1}{2}a\tau\right) \\ &\times \left[\cos\left(\frac{1}{2}(\delta T - \Delta\phi)\right)\cos\left(\frac{1}{2}a\tau\right) - \frac{\delta}{a}\sin\left(\frac{1}{2}(\delta T - \Delta\phi)\right)\sin\left(\frac{1}{2}a\tau\right)\right]^2, \end{aligned} \tag{1.57}$$

and a simplified form valid close to resonance can be re-written as

$$P_{\text{Ramsey}} = \left|\frac{\Omega\tau}{2}\right|^2 \left(\frac{\sin(\delta\tau/2)}{\delta\tau/2}\right)^2 \cos^2\left(\frac{\delta T - \Delta\phi}{2}\right), \tag{1.58}$$

such that the transition probability is maximised when $\delta = \Delta\phi/T$, rather than $\delta = 0$. Thus, a phase shift presents itself as an apparent frequency shift of $\Delta\nu = \Delta\phi/2\pi T$. It is important to note that this shift is different to the second-order Zeeman shift and blackbody radiation shift: the energy levels of the atoms are not actually shifted by a change in cavity phase. Instead, the apparent frequency of the microwave radiation is displaced, which is indistinguishable from a change in the atom's resonant frequency.

It was known from experiments with thermal beam clocks that transverse phase gradients resulted in an apparent frequency shift [117]. These shifts were initially modelled in beam clocks with transmission-line theory and equivalent-circuit models of the microwave cavity assembly [118], and suggestions for reducing the effect included feeding each cavity from both sides to reduce the transverse running-wave part of the microwave field [119].

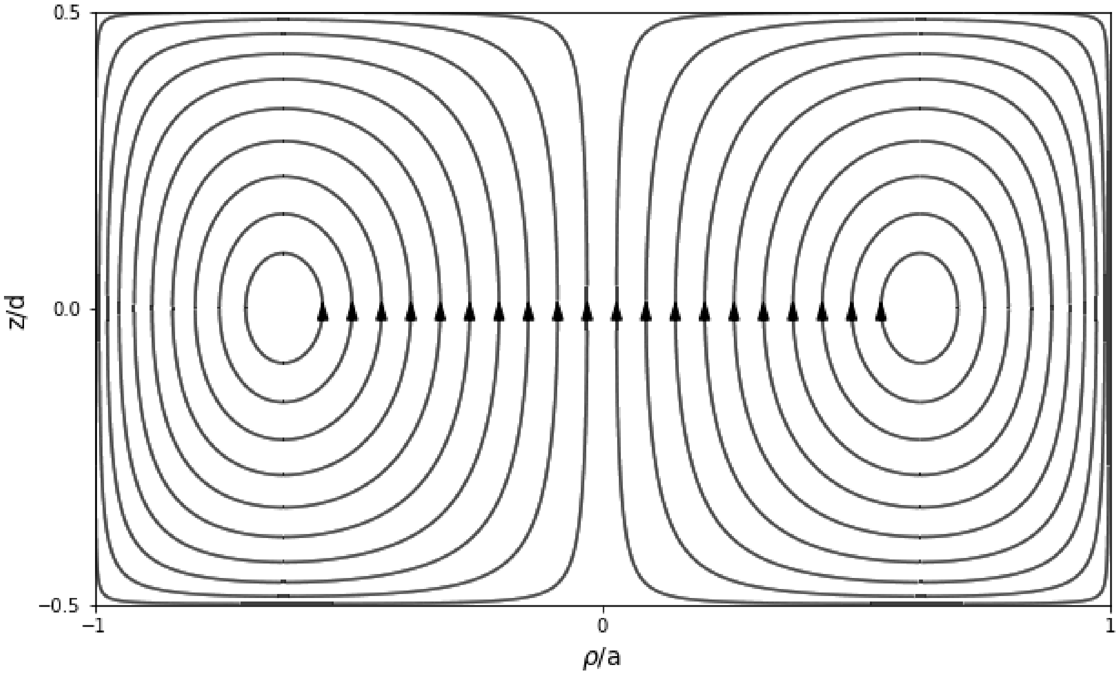

Figure 1.15: Magnetic field lines for a cylindrical $\mathrm{TE}_{011}$ cavity mode.

This analysis can be extended to an idealised cylindrical $\mathrm{TE}_{011}$ cavity [120]. The magnetic field of an ideal $\mathrm{TE}_{011}$ mode of height $d$ and radius $a$, with its centre at the origin, can be written [121] as

$$\begin{aligned} H_\rho &= H_0 \frac{\beta}{k_c} J_0'(k_c\rho)\sin(\beta z), && (1.59)\\ H_\phi &= 0, && (1.60)\\ H_z &= H_0 J_0(k_c\rho)\cos(\beta z), && (1.61) \end{aligned}$$

where the longitudinal wave number $\beta = \sqrt{k^2 - (p'_{01}/a)^2} = \pi/d$, the transverse wave number $k_c = p'_{01}/a$, $J_0$ is the Bessel function of the first kind, $J_0'$ its derivative, and $p'_{01} = 3.832$ is the first zero of $J_0'$. An $e^{i\omega t}$ term has been assumed and omitted. These magnetic field lines are shown in figure 1.15.

The unperturbed cavity mode is a standing wave, so the phase is uniform. However, the finite conductivity of walls leads to a modification of the field such that a small imaginary component is introduced, and the phase variations $\Phi(\rho, z)$ can be expressed as [120] [122]

$$\tan(\Phi(\rho, z)) = \frac{\mathrm{Im}(|H|)}{\mathrm{Re}(|H|)} \approx \frac{\mathrm{Im}(|H_z|)}{\mathrm{Re}(|H_z|)}. \tag{1.62}$$

Since the atoms should be passing parallel to the $z$-axis, the longitudinal phase variation cancels itself out on each passage, leaving only the transverse variation between passages to cause a frequency shift. This led to the conclusion that, in a high-Q copper cavity with surface wall resistance $R_s = 2.6 \times 10^{-7}\,\Omega$, the corresponding frequency shift would be kept below $10^{-16}$ if the phase shift is below $3\,\mu$rad, which was thought to be satisfied within about 4 mm of the cavity axis.

However, the analytical method can not be extended to real cavities with holes for feeding and atom passage. A two-dimensional numerical study using finite-element analysis [122] sought a primary field $\mathbf{H}_0$ as a solution for the full geometry with perfectly conducting walls, and a smaller component $\mathbf{h}$ which accounts for the wall losses. This was extended to include the flow of microwave power from the source in the $\rho$-$\phi$ plane [123], and it was found that the effect of the end-cap holes was to slow down the transverse phase variation across the opening, and that well-balanced symmetric feeds can eliminate transverse phase gradients [124], at the expense of increased longitudinal phase gradients.

A more-detailed 2D analysis of wall and end-cap losses decomposed the $z$-components of the secondary field, $h_z(\mathbf{r})$, into a Fourier series [125] such that

$$h_{z,m}(\mathbf{r}) \propto (1+i)\rho^m \cos(m\phi) \cos\left(\frac{\pi}{d} z\right). \tag{1.63}$$

The sum quickly converges as the contributions of higher $m$ become small, such that the most important terms are $m = 0$, 1, and 2. The $m = 0$ term corresponds to power fed from the mid-plane flowing towards the end-caps. The $m = 1$ term corresponds to power flowing across the cavity. Balanced, opposed feeds highly suppresses the contribution of the odd-$m$ terms. The $m = 2$ term corresponds to power flowing radially inwards at the positions of opposing feeds and radially outwards at 90° to the feeds, and can be minimised by detecting the atoms with a laser at 45° to the coupling holes [68]. 2D finite-element analysis with this revised analytical model [125] for a cavity with end-cap holes found that large phase shifts occur near the opening of these holes, as $\mathrm{Re}(H_z)$ is small but $\mathrm{Im}(h_z)$ is relatively large. Avoiding this effect was suggested by staggering the diameters of the waveguides-below-cut-off, such that the detected atoms do not pass through these regions.

When perfectly on-resonance, the rotation angle on the Bloch-sphere that an atom's state-vector turns through during a single cavity transit can be expressed as

$$\theta(\rho, z) = b\frac{\pi}{2}\int_{-d/2}^{z} H_z(\rho, z')dz', \tag{1.64}$$

where $b$ gives the amplitude of the field ($b = 1$ is the nominal optimum amplitude for $\pi/2$ pulses), and a full cavity transit corresponds to setting $z = +d/2$. Since the cavity field $H_z(\mathbf{r})$ is not uniform, it is common to introduce a sensitivity function $s(t)$ of the form [92]

$$s(t) = \begin{cases} 0 & \text{before cavity} \\ -\sin(\theta(\rho_1, z(t)))\sin(\theta(\rho_2, d/2)) & \text{during first pass} \\ -\sin(\theta(\rho_1, d/2))\sin(\theta(\rho_2, d/2)) & \text{above cavity} \\ -\sin(\theta(\rho_1, d/2))\sin(\theta(\rho_2, z(t))) & \text{during second pass} \\ 0 & \text{after cavity} \end{cases} \tag{1.65}$$

where an atom passes through the cavity at $\rho_1$ on the upwards passage and $\rho_2$ on the downwards passage. Changes in an atom's transition probability can thus be expressed as [92][126]

$$\delta P = \frac{1}{2}\int_{-\infty}^{\infty} s(t)\frac{d\Phi(\rho, z, t)}{dt}dt = -\frac{1}{2}\int_{-\infty}^{\infty} \Phi(\rho, z, t)\frac{ds(t)}{dt}dt, \tag{1.66}$$

Since the atom's vertical velocity $v_z$ is approximately constant over the time it spends traversing the cavity, it is straightforward to turn functions and integrals of $t$ into $z$, such that, for an atom that passes through the cavity at $\rho_1$ on the way up and $\rho_2$ on the way down, the change in transition probability due to each $m$ is [127][128], with $\Phi_m \approx -\operatorname{Im}(h_{z,m})/H_z$,

$$\delta P_m(\rho_1, \rho_2) = \frac{1}{2}(\sin(\theta(\rho_2, -d/2))\delta\Phi_m(\rho_1) - \sin(\theta(\rho_1, d/2))\delta\Phi_m(\rho_2)), \tag{1.67}$$

where

$$\delta\Phi_m(\rho) = b\frac{\pi}{2}\int_{-d/2}^{d/2} \cos(\theta(\rho, z))\operatorname{Im}(h_{z,m})dz. \tag{1.68}$$

The total change in probability is the sum over all $m$ terms,

$$\delta P(\rho_1, \rho_2) = \sum_m \delta P_m(\rho_1, \rho_2). \tag{1.69}$$

It is often convenient to measure the DCP effect in terms of $\delta P$ rather than a frequency shift because of its complicated power dependence. For even $b$, the atomic state vector is back where it started and no population is transferred between levels, such that the frequency shift is undefined [127, 129]. The change in transition probability is always well behaved.

This change in transition probability becomes a frequency shift due to the slope of the central Ramsey fringe. The mid-point of the fringe occurs when

$$\cos^2\left(\frac{\delta T}{2}\right) = \frac{1}{2}, \tag{1.70}$$

such that the detuning, in rad/s, is $\delta = \pi/2T$. The gradient of the Ramsey fringe is thus

$$\left.\frac{dP}{d\delta}\right|_{\delta} = \frac{T}{2} = \frac{Q_R}{4\nu_0}, \tag{1.71}$$

where $Q_R$ is the quality factor of the Ramsey fringe. Thus, the fractional frequency shift is

$$\frac{\Delta\nu_{\mathrm{DCP}}}{\nu_0} = \frac{2}{\pi Q_R}\delta P(\rho_1, \rho_2), \tag{1.72}$$

for a single atom passing on the way up at $\rho_1$ and on the way down at $\rho_2$. The resulting shift due to all atoms on all detected trajectories depends on the spatial distribution of the cloud during each pass, and is an average over all of the detected atoms.

While finite-element analysis can be a powerful tool in designing cavities and estimating their DCP errors, real cavities can differ from simulations due to factors like non-uniform oxidation of the copper surface, impurities lowering the conductivity, and mechanical imperfections. Quantifying the DCP with confidence requires in-situ frequency measurements.

Experimentally verifying a model for the DCP shift poses a challenge because it is difficult to isolate from other effects like microwave leakage [130]. However, since the dependence of DCP on microwave power is highly non-trivial, yet unique [127], then measurements of frequency vs microwave power at elevated microwave power may

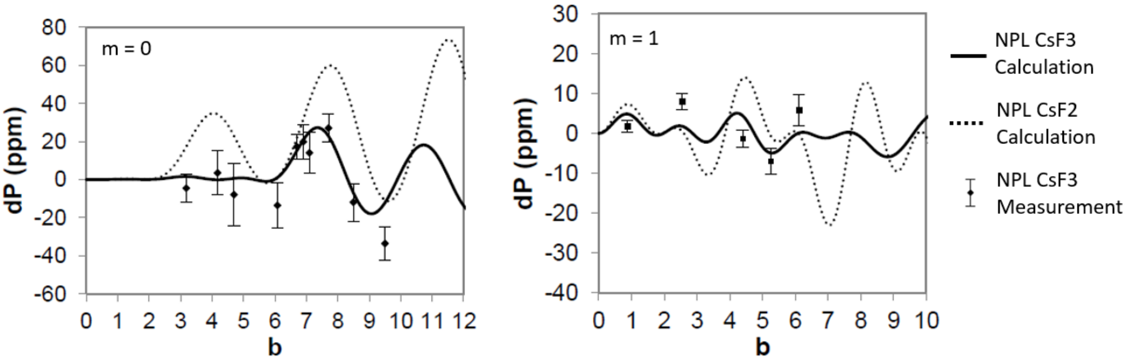

Figure 1.16: $m = 0$ (left) and $m = 1$ (right) DCP power dependence, comparing measurements of change in transition probability (in parts-per-million) vs microwave power (in multiples of $\pi/2$ pulses) to calculations [68].

constrain the shift at optimum power, which is usually in the $10^{-16}$ region [129]. An example of such a measurement is shown in figure 1.16. Since the shift also depends on the atom's position on each cavity pass, then accentuating the difference between the average $\rho_1$ and $\rho_2$ will also enhance the DCP effect, to up to the $10^{-15}$ level. This can be achieved by tilting the fountain, as gravity then gives a small transverse acceleration in the frame of the fountain. The $m = 1$ component can be further exaggerated by feeding the cavity only from one side, and comparing the frequency difference as a function of tilt [131]. If the fountain is mechanically well-aligned, then the frequency difference should vanish when the fountain is perfectly vertical. The DCP uncertainty is then the precision of the zero-crossing point multiplied by the frequency gradient in the vicinity of the crossing. These methods allow the DCP at optimal conditions to be measured with an accuracy of $10^{-16}$ or better [69, 132, 133], which is still large enough to make it the dominant systematic uncertainty in several modern fountain clocks [67, 68].

Reducing this uncertainty further can be achieved by obtaining more precise information about the positions at which the atoms cross the microwave cavity. By applying a tilt to the background C-field and observing the spectra of $\Delta m_F = \pm 1$ transitions, the position of the centre of the cloud could be measured to a precision of 0.1 mm [134]. By using horizontal coils to shift the position of the MOT in the initial cooling stage of the cycle, this precision could be improved to $50\,\mu$m, such that the DCP could be measured with an accuracy of $10^{-17}$ [135].

### 1.2.4 Collisional Shift

For caesium atoms cooled down to only a few microkelvin, the de Broglie wavelength of the atoms is comparable to the scale of the inter-atomic potentials. At a typical 1 $\mu$K temperature, a $^{133}$Cs atom has a wavelength of around $7000\,a_0$ (370 nm), where $a_0 = 5.29 \times 10^{-11}$ m is the Bohr radius. Similarly, the scattering length for atoms with anti-aligned and aligned electron spins are found to be $280\,a_0$ (14.8 nm) and $2400\,a_0$ (127 nm), respectively [136]. This results in a density-dependent, energy-dependent, state-dependent, and hence time-dependent shift to each atom's resonant frequency during the Ramsey interrogation [137]. This effect is often referred to as the spin-exchange shift, the density shift, or just the collisional shift.

In general, the wavefunction of an atom in a superposition of two states $|\alpha\rangle$ and $|\beta\rangle$ can be written as

$$|\psi\rangle = a\,|\alpha\rangle + b\,|\beta\rangle\,, \tag{1.73}$$

where $a$ and $b$ are the amplitudes of each state. The probabilities that, when measured, the atom will be found in state $|\alpha\rangle$ or $|\beta\rangle$ are $|a|^2$ and $|b|^2$, respectively. An alternative way to write this wavefunction is in an explicitly vector notation:

$$|\psi\rangle = \begin{pmatrix} a \\ b \end{pmatrix}, \tag{1.74}$$

where the first row implicitly represents the state $|\alpha\rangle$ and the second $|\beta\rangle$. Similarly, the Hermitian conjugate (transpose and complex-conjugate) of $|\psi\rangle$ can be written as

$$\langle\psi| = \begin{pmatrix} a^* & b^* \end{pmatrix}. \tag{1.75}$$

From these, the density matrix can be written as

$$\rho = |\psi\rangle\,\langle\psi| = \begin{pmatrix} |a|^2 & a^*b \\ b^*a & |b|^2 \end{pmatrix}. \tag{1.76}$$

The diagonal entries are the probabilities as before, and the off-diagonal entries are called coherences. It is the presence of these off-diagonal terms that distinguish a probabilistic mixture of $|\alpha\rangle$ and $|\beta\rangle$ from a coherent superposition.

The evolution equation for the density matrix $\rho$ under the influence of a Hamiltonian $\hat{H}$ is

$$i\hbar\frac{\partial\rho}{\partial t} = [\hat{H}, \rho]. \tag{1.77}$$

When this evolution is due to collisions, a frequency shift $\Delta\omega_c$ and line broadening $\Gamma_c$ occurs according to [137, 138]

$$\frac{\partial}{\partial t}\rho_{\alpha\beta} = -(\Gamma_c - i\Delta\omega_c)\rho_{\alpha\beta}, \tag{1.78}$$

where $\rho_{\alpha\beta}$ are the off-diagonal density matrix elements. Thus, determination of the collisional shift requires evaluation of equation 1.77 for a Hamiltonian that represents the interaction between cold Cs atoms in a superposition of clock states.

In the low-energy case of cold atoms in a fountain clock, scattering can be analysed in terms of partial waves. The spatial wavefunction of a particle scattering off of a target can be expressed as a superposition of an incoming plane wave and series of scattered spherical waves,

$$\psi_k(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{z}} + \frac{e^{ikr}}{r}\sum_{l=0}^{\infty}(2l+1)f_l^{s,t}(k,r)P_l(\cos\theta), \tag{1.79}$$

where $P_l(\cos\theta)$ are the Legendre polynomials and the scattered wave amplitude $f_l^{s,t}(k,r)$ depends on the inter-atomic potential, which is different for the singlet (anti-aligned) and triplet (aligned) states [139]. Since the clock states both have $m_F = 0$, then, by the symmetry of the Legendre polynomials, the odd-$l$ terms do not contribute, and the $l = 0$ term dominates, known as s-wave scattering. To solve for $f_0^s(k,r)$ and $f_0^t(k,r)$ requires solving the Schroedinger equation, and the result [137, 140, 141, 142] is of the form

$$i\Delta\omega_c - \Gamma_c = n\sum_{\nu}\rho_{\nu\nu}(1+\delta_{\alpha\nu})(1+\delta_{\alpha\nu})\langle v(i\lambda'_\nu(E) - \sigma_\nu(E))\rangle, \tag{1.80}$$

$$\Delta\omega_c = n(\rho_{|3,0\rangle}\langle 2v\lambda'_{|3,0\rangle}(E)\rangle + \rho_{|4,0\rangle}\langle 2v\lambda'_{|4,0\rangle}(E)\rangle), \tag{1.81}$$

where the labels $\alpha = |3,0\rangle$, $\beta = |4,0\rangle$, such that the sum over $\nu$ means that $\nu = \alpha, \beta$, the total number density is $n$, $v$ is the relative velocity of colliding atoms, and $\lambda'_\nu(E)$ and $\sigma_\nu(E)$ are the energy and state-dependent scattering cross-sections. Since $v$ and $E$ are related by $E = \mu v^2/2$ (where $\mu$ is the reduced-mass of a two-atom system),

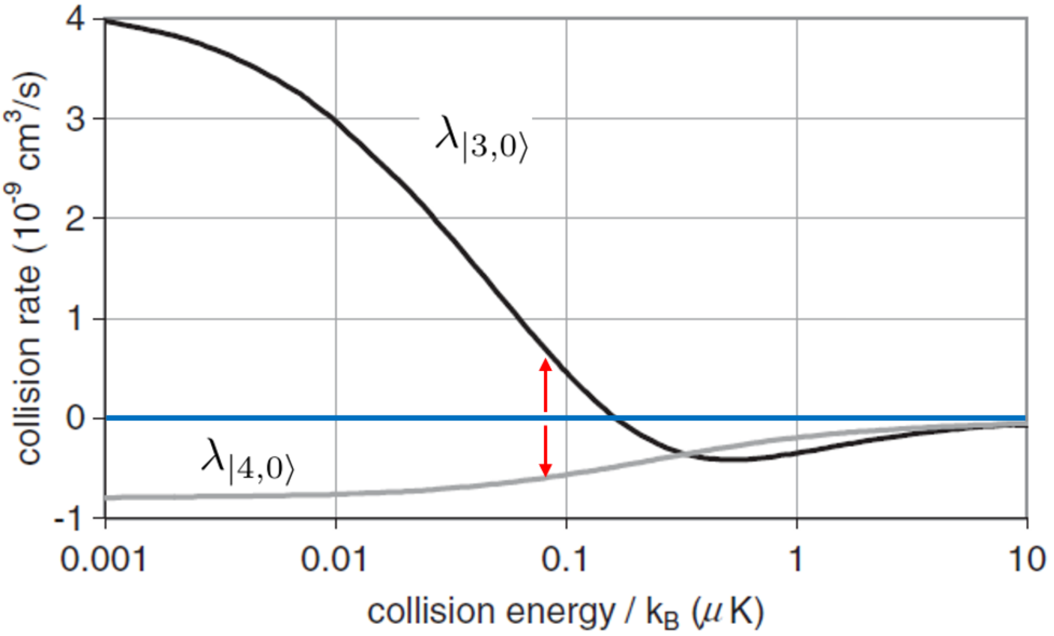

Figure 1.17: Theoretical collision rate vs collision energy for each clock state [145]. The sign of $\lambda_{|3,0\rangle}$ reverses, such that it has opposite sign but similar magnitude to $\lambda_{|4,0\rangle}$ at approximately $0.08\,\mu$K.

this can be re-written as

$$\Delta\omega_c = n\langle \rho_{|3,0\rangle}\lambda_{|3,0\rangle}(E) + \rho_{|4,0\rangle}\lambda_{|4,0\rangle}(E)\rangle. \tag{1.82}$$

It was initially assumed [137] and seemingly verified [143] that the collision energy was not important, because the changes in collision velocity were thought to cancel out the changes in collision cross-section at the $\sim\mu$K level. The size of the collisional shift was estimated to be approximately $\Delta\nu/\nu = -1\times10^{-22}\ n/\mathrm{cm}^3$ [137]. However, further analysis showed that this is only true for temperatures below 0.1 nK, and what the experiments were actually seeing was a *local* minima in clock-shift vs collision energy [141]. In fact, a reversal in the sign of the clock shift was predicted at $0.08\,\mu$K [141], due to reversal of the sign of $\lambda_{|3,0\rangle}$ [144], which leads to the possibility of operating a fountain with zero collisional shift [144, 145]. The energy dependence of the partial cross-sections $\lambda_{|3,0\rangle}$ and $\lambda_{|4,0\rangle}$ are shown in figure 1.17.

Since the collisional shift has a strong dependence on collision energy, it also depends on the initial cloud size and velocity distribution. Initial methods of evaluating the collisional shift focused on the density dependence: keeping the collision energies and state populations constant means that $\delta\omega_c \propto n$, such that measuring at two different densities allows the frequency shift to be extrapolated to zero density. Varying the temperature of the caesium source allows for a large variation in the number of atoms captured in the initial loading stage, but this is a slow method that takes around 1 hour to implement [96]. One quicker method of changing the atom number is to reduce the initial loading time [146], although this does not necessarily preserve the same relative density distribution. Similarly, the power in the state-selection cavity can be varied such that for the high density operation, the atoms experience a $\pi$ pulse, and for low density operation, the atoms experience a $\pi/2$ pulse, so only half of the population survives the pushing laser pulse. This method also does not guarantee an identical density distribution, because the field amplitude in the cavity is not quite uniform. The uncertainties on the extrapolated frequency shift ranged from $2 \times 10^{-16}$ [147] to $8 \times 10^{-16}$ [96, 146, 148] using these methods. The downside of these methods is that the fountain spends most of its time operating at low density, such that the signal-to-noise ratio, and thus the short-term stability, is degraded. The effective Allan deviation $\sigma_{\text{eff}}(\tau)$ of a fountain operating at two different densities such that $\kappa = N_H/N_L$ is the ratio of atom numbers is [145, 76]

$$\sigma_{\text{eff}}(\tau) = \frac{1}{\kappa - 1}\sqrt{\kappa^2\sigma_L(\tau)^2 + \sigma_H(\tau)^2}, \tag{1.83}$$

where $\sigma_H(\tau)$ is the ADEV at high-density and $\sigma_L(\tau)$ is the ADEV at low density. This means that, for $\kappa = 2$, then $\sigma_{\text{eff}}(\tau) = \sqrt{6}\sigma_H(\tau)$ if the fountain stability would otherwise be limited by quantum projection noise, and $\sigma_{\text{eff}}(\tau) = \sqrt{10}\sigma_H(\tau)$ if the fountain would otherwise be limited by local oscillator noise.

The reversal in sign of $\lambda_{|3,0\rangle}$ can be exploited to reduce the collisional shift. While this only happens at temperatures far below what even sub-Doppler cooling methods achieve, the instantaneous collision velocity between atoms can be much lower than the mean thermal velocity [149]. Once the cloud has expanded to much larger than its initial size, then the positions and velocities of the atoms become correlated: all the atoms in one particular vicinity are travelling in broadly the same direction and with similar speeds. The result is that part way through the ballistic flight, the value of $\lambda_{|3,0\rangle}$ changes sign, such that the total collisional shift, which is a time-average over $n$ and $\rho_{|3,0\rangle}\lambda_{|3,0\rangle} + \rho_{|4,0\rangle}\lambda_{|4,0\rangle}$, can be controlled by controlling the initial cloud size, initial temperature, and the ratio between $\rho_{|3,0\rangle}$ and $\rho_{|4,0\rangle}$ after the first Ramsey interaction. It was found that, for a cloud with initial radius of 1 mm and an initial temperature of $1\,\mu$K, then reducing the microwave power such that $\rho_{|4,0\rangle} = 0.4$ after the first Ramsey passage was sufficient to cancel out the collisional shift, with only a negligible decrease in fringe contrast [144]. However, this requires the initial cloud size and temperature to be stable to only a couple of percent or better to reach accuracy of $1 \times 10^{-16}$, such that extrapolation of the residual shift to zero is still necessary to ensure long-term frequency accuracy.

It was predicted early on that $^{87}$Rb would suffer a much smaller collisional shift than $^{133}$Cs [150]. Initial measurements with rubidium fountains found that the collisional shift was between 30 [151] and 50 [152] times smaller than in caesium. This is because the singlet and triplet scattering lengths are very similar to each other [153].

Another helpful feature of collisions between cold rubidium atoms is the relatively high temperature below which the frequency shift does not depend on collision energy. The temperature below which the Wigner-threshold regime is reached is approximately $100\,\mu$K, such that the exact collision energy in the single-microkelvin range does not affect the scattering cross-sections [154]. The collisional shift in $^{87}$Rb is approximately $\Delta\nu/\nu = -7 \times 10^{-24}\ n/\mathrm{cm}^3$ [152].

Collisions between cold atoms of the same species are not the only collisions of concern. Background atoms at room temperature have much more momentum than cold atoms, so when they collide with cold atoms, the cold atoms tend to be ejected from the cloud. However, glancing collisions in which little momentum is transferred may not eject a cold atom, but will result in a phase shift to its internal frequency due to the scattering event. There is a direct correlation between the loss of atoms from the cloud over the ballistic flight, and the frequency shift due to background gas collisions [155]. The result is that, for both background thermal Cs atoms and hydrogen gas (the otherwise-dominant species in UHV) the frequency shift is approximately [156]

$$\frac{\Delta\nu_{\mathrm{BGC}}}{\nu_0} = 2 \times 10^{-16} \frac{\Delta N}{N}, \tag{1.84}$$

where $\Delta N$ is the number of the atoms lost by collisions. Thus, for typical loss fractions of around 5%, the shift is estimated to be of the order $10^{-17}$ and thus negligible.

### 1.2.5 Cavity Pulling

An alkali atom has a magnetic dipole from its unpaired outer electron, which gives rise to interference between the magnetic field of the cavity mode and the AC magnetic field radiated by the atoms. This results in the (first-order) cavity pulling shift.

The total field in the cavity can be expressed as [157]

$$B(\omega, t) = B^{(0)}(\omega, t) + \frac{\mu_0 \omega^2}{V_m} \frac{\iiint \mathbf{m}(\mathbf{r}, \omega, t) \cdot \mathbf{H}_c(\mathbf{r}(t)) d^3\mathbf{r}}{\omega_c^2 - \omega^2 - i\omega\omega_c/Q_c}, \tag{1.85}$$

where $\mathbf{H}_c(\mathbf{r})$ is the normalised cavity mode as defined in equations 1.59, 1.60, and 1.61 by setting $H_0 = 1$, $\omega_c$ is the resonant angular frequency of the microwave cavity, $B^{(0)}(\omega, t)$ is the unperturbed field, $Q_c$ is the quality factor of the cavity mode and $V_m$ is the mode volume, which in general is

$$V_m = \iiint |\mathbf{H}_c(\mathbf{r})| d^3\mathbf{r}, \tag{1.86}$$

which for a $\mathrm{TE}_{011}$ mode in a cavity of radius $a$ and height $d$ is

$$V_m = \pi a^2 d \left(1 + \left(\frac{\pi a}{p'_{01} d}\right)^2\right) \frac{J_0^2(p'_{01})}{2}. \tag{1.87}$$

On the axis of the cavity, the field $\mathbf{H}_c(\mathbf{r}(t)) = \cos(\pi z(t)/d)\hat{\mathbf{z}}$. Assuming that, on the first cavity transit, that the cloud is small enough to treat as a single point, then each atom experiences this same magnetic field. The total magnetic moment $\mathbf{m}$ due to $N$ atoms can thus be written as $\mathbf{m} = N m_{1z}\hat{\mathbf{z}}$. The single-atom $z$-component of magnetic dipole moment of an atom undergoing a transition from a ground state $|g\rangle$ to an excited state $|e\rangle$ can be written, for a general state vector $|\psi(t)\rangle$, as [157]

$$m_{1z}(\omega, t) = 2\mu_B \langle\psi(t)|g\rangle \langle e|\psi(t)\rangle . \tag{1.88}$$

For a single pulse where the state vector, initially being in the ground state $|g\rangle$, rotates through an angle $\theta(t)$ around the Bloch-sphere, then

$$|\psi(t)\rangle = \cos\frac{\theta(t)}{2} |g\rangle - i \sin\frac{\theta(t)}{2} |e\rangle , \tag{1.89}$$

such that the magnetic dipole moment becomes

$$m_{1z}(\omega,t) = -i\mu_B \sin\theta(t). \tag{1.90}$$

This means that the modified cavity field is

$$B(\omega,t) = B^{(0)}(\omega,t)\left(1+\frac{\mu_0\omega^2}{B^{(0)}(\omega,t)V_m}\frac{iN\mu_B\sin\theta(t)H_z(\mathbf{r}(t))}{\omega_c^2-\omega^2-i\omega\omega_c/Q_c}\right), \tag{1.91}$$

where the real part of equation 1.91 represents a small change in field amplitude, and the imaginary part of equation 1.91 represents a phase shift, which results in a frequency shift. In this situation, the phase shift is

$$\begin{aligned}\Delta\Phi(t) =& \mathrm{Im}\left(\frac{\mu_0\omega^2}{B^{(0)}(\omega,t)V_m}\frac{iN\mu_B\sin\theta(t)H_z(\mathbf{r}(t))}{\omega_c^2-\omega^2-i\omega\omega_c/Q_c}\right) & (1.92)\\ =&\frac{N\mu_0\mu_B}{B^{(0)}(\omega,t)V_m}\frac{\omega^2(\omega_c^2-\omega^2)}{(\omega_c^2-\omega^2)^2+\omega^2\omega_c^2/Q_c^2}\sin\theta(t)H_z(\mathbf{r}(t)). & (1.93)\end{aligned}$$

For a $\pi/2$ pulse, the rotation angle $\theta(t)$ can be expressed as [154]

$$\theta(t) = -\frac{\mu_B B^{(0)}(\omega,t)}{\hbar}\tau_{\text{eff}} = -\pi/2, \tag{1.94}$$

where $\tau_{\text{eff}} = \int_0^\tau H_z(\mathbf{r}(t'))dt$. These can be re-arranged to show that

$$\Delta\Phi(t) = \frac{N\mu_0\mu_B^2}{\pi\hbar V_m}2\tau_{\text{eff}}\sin\theta(t)H_z(\mathbf{r}(t))\frac{\omega^2(\omega^2-\omega_c^2)}{(\omega_c^2-\omega^2)^2+\omega^2\omega_c^2/Q_c^2}. \tag{1.95}$$

This time-varying phase shift produces a change in transition probability according to the sensitivity function $s(t)$ (equation 1.65) and equation 1.66. For $0 \le t \le \tau$, $s(t) = \sin(\theta(t))$, and on the slope of the Ramsey fringe where $\omega = \omega_0 \pm \pi/2T$, one has

$$2\pi\Delta\nu = \frac{N\mu_0\mu_B^2}{\pi\hbar V_m}\frac{2\tau_{\text{eff}}}{T_{\text{eff}}}\frac{\omega^2(\omega^2-\omega_c^2)}{(\omega_c^2-\omega^2)^2+\omega^2\omega_c^2/Q_c^2}\int_0^\tau\frac{ds(t)}{dt}\sin\theta(t)H_z(\mathbf{r}(t))dt, \tag{1.96}$$

where $T_{\text{eff}} = \int_0^{T+2\tau} s(t)dt$. By using $\theta(t) = -\int_0^\tau \Omega H_z(\mathbf{r}(t'))dt'$ to simplify the integration, one has

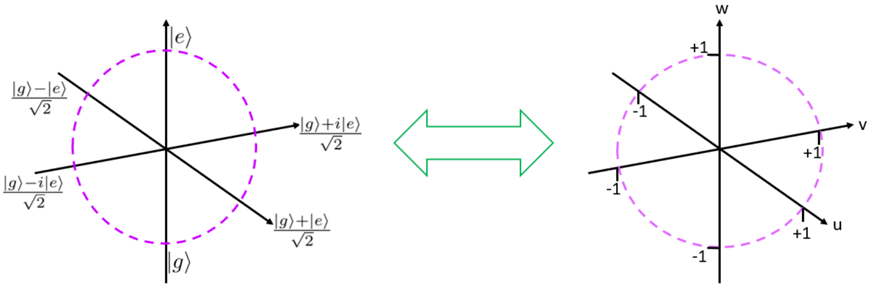

Figure 1.18: Two equivalent views of the Bloch sphere, in terms of atomic states (left) and coherences (right).

$$
\begin{aligned}
\Delta\nu =& \frac{N\mu_0\mu_B^2}{\pi^2\hbar V_m}\frac{\tau_{\text{eff}}}{T_{\text{eff}}}\frac{\omega^2(\omega^2-\omega_c^2)}{(\omega_c^2-\omega^2)^2+\omega^2\omega_c^2/Q_c^2}\frac{\pi}{4\tau_{\text{eff}}}\int_0^{\tau} H_z^2(\mathbf{r}(t))\sin(2\theta(t))dt && (1.97)\\
=& \frac{N\mu_0\mu_B^2}{\pi^2\hbar V_m}\frac{\tau_{\text{eff}}}{T_{\text{eff}}}\frac{\omega^2(\omega^2-\omega_c^2)}{(\omega_c^2-\omega^2)^2+\omega^2\omega_c^2/Q_c^2}\times 0.445, && (1.98)
\end{aligned}
$$

as the frequency shift, in Hz, due to the cavity pulling effect for a single cavity pass.

It is sufficient to consider only the first cavity passage, and neglect the second, for two reasons. Firstly, due to the thermal expansion of the cloud, there are far fewer atoms passing through the cavity aperture on the way down. Secondly, because the typical fountain operation is on the slope of the Ramsey fringe [154]. Coherent superpositions of $|g\rangle$ and $|e\rangle$ states can be represented as points on a unit-sphere. As shown in figure 1.18, there are two representations of the Bloch sphere: where the axis are represented by atomic state vectors, and where the axis are labelled by coherences $u$, $v$, and $w$ such that

$$
\begin{aligned}
u =& \rho_{ge} + \rho_{eg}, && (1.99)\\
v =& -i(\rho_{ge} - \rho_{eg}), && (1.100)\\
w =& \rho_{gg} - \rho ee, && (1.101)
\end{aligned}
$$

where $\rho_{ij}$ are the entries of the relevant density matrix, in the rotating frame of the light field. The evolution of the Bloch vector during the free-evolution time is $\delta T = \pi/2$, such that the second pulse, causing a $\pi/2$ rotation about the $u$-axis, has no effect because the atomic state vector is already aligned with the $u$-axis. This is illustrated by figure 1.19.

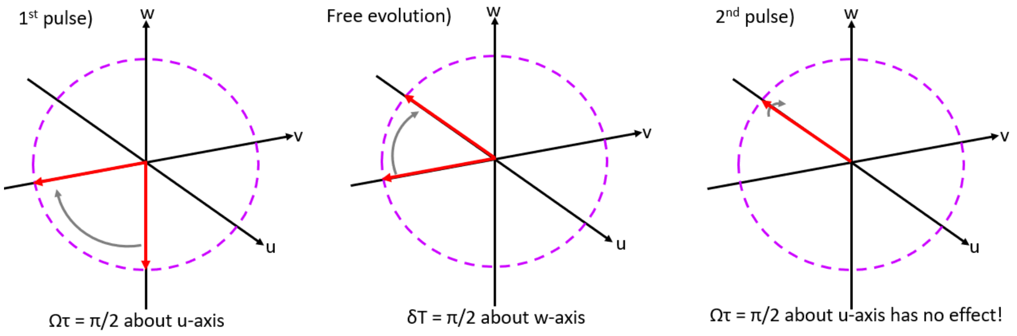

Figure 1.19: Ramsey pulse sequence represented on the Bloch sphere for $\delta = \pi/2T$.

Because the cavity pulling shift depends on the cavity's resonant frequency $\omega_c$, then thermal expansion of the cavity due to temperature changes affects the cavity pulling effect. The frequency shift to the cavity by a temperature change $\delta T$ is

$$\delta f_c = -\frac{\alpha \delta T}{1 + \alpha \delta T} f_0, \tag{1.102}$$

where, for copper, the thermal expansion coefficient is $\alpha = 17 \times 10^{-6}\,\mathrm{K}^{-1}$. At room temperature in a Cs fountain, this gives a frequency shift of $-156\,\mathrm{kHz/K}$. The cavity pulling shift vs cavity detuning is shown in figure 1.20.

The cavity pulling effect was first verified in a rubidium fountain. A higher atom number was used to exaggerate the effect, and the cavity's resonant frequency was varied by varying the temperature [157]. Because the effect depends on the atom number, it can be used to partially cancel out the collisional shift, which depends on the atom density [151]. Similarly, extrapolating to zero density is equivalent to extrapolating to zero atom number, such that cavity pulling is accounted for by accounting for the collisional shift.

Methods of controlling the uncertainty of the cavity pulling effect include using a low-Q microwave cavity to flatten the function [158], and controlling the resonant frequency of the cavity to within approximately 1 kHz of the atomic resonance by precise temperature control [157, 151]. Detuning the cavity away from the atomic resonance, and operating in a flatter part of the function, could reduce the cavity pulling effect, but at the expense of potentially worsening the DCP effect.

The second-order cavity pulling effect is a result of the mismatch of the cavity resonant frequency to the atomic transition frequency $\delta_c$, and the mismatch of the

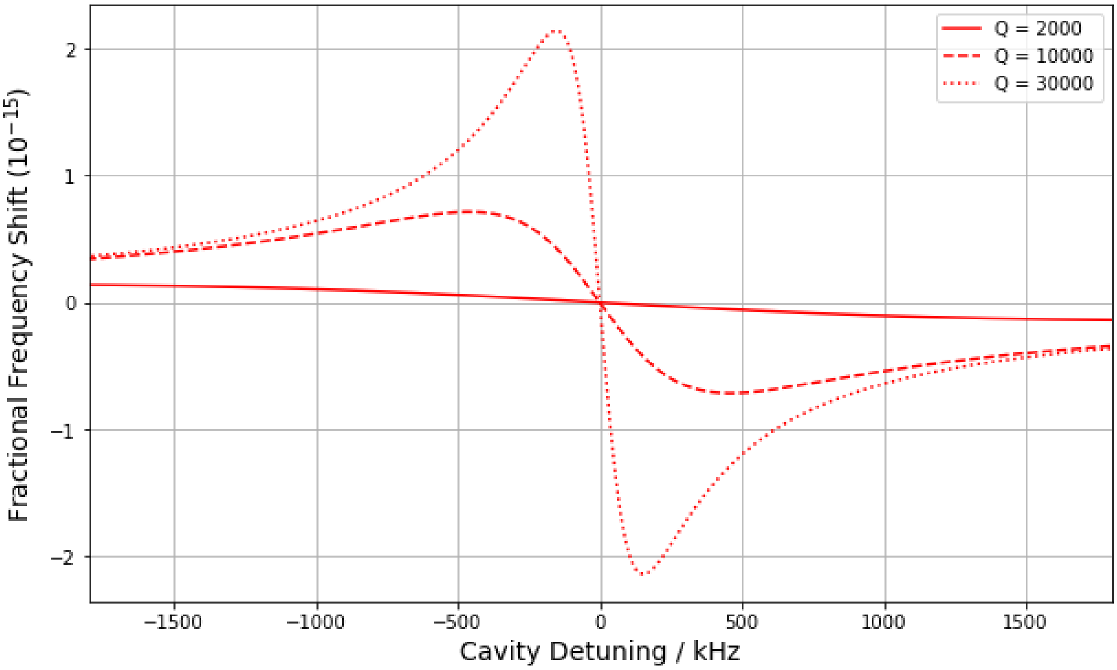

Figure 1.20: Cavity pulling vs cavity detuning for $N = 10^7$ atoms, $\tau_{\text{eff}} = 8.6\,\text{ms}$, $T_{\text{eff}} = 516\,\text{ms}$, and $V_m = 5.9 \times 10^{-6}\,\text{m}^3$ for temperatures of $\approx \mp 10\,\text{K}$.

microwave field pulse area $\Omega\tau$ from $\pi/2$, and is given by [159, 96]

$$\frac{\Delta\nu_{\text{cp2}}}{\nu_0} = \frac{8}{\pi^2}\left(\frac{Q_c}{Q_{\text{at}}}\right)^2\left(\frac{\delta_c}{\nu_0}\right)\Omega\tau\cot(\Omega\tau). \tag{1.103}$$

Typically this effect is also small compared to the fountain accuracy of $\sim 10^{-16}$, as $Q_{\text{at}} \gg Q_c$, $\nu_0 \gg \delta_c$, and $\Omega\tau$ can be set to $\pi/2$ to within a few percent.

### 1.2.6 Unintentional Excitation of Other Transitions

Since atoms have multiple Zeeman sub-levels, there exists the possibility of exciting transitions other than the clock transition. These can be broken down into three categories: Rabi pulling is the effect on the clock transition frequency of the wings of adjacent $\Delta F = \pm 1$, $\Delta m_F = 0$, $\pi$-transitions; Ramsey pulling is the effect of $\Delta F = \pm 1$, $\Delta m_F = \pm 1$, $\sigma$ transitions; and Majorana transitions, which occur when a sudden change in magnetic field orientation redistributes the atomic states into a superposition of $m_F$ states, with are detected with varying efficiencies. The most relevant transitions for Rabi and Ramsey pulling are shown in figure 1.21.

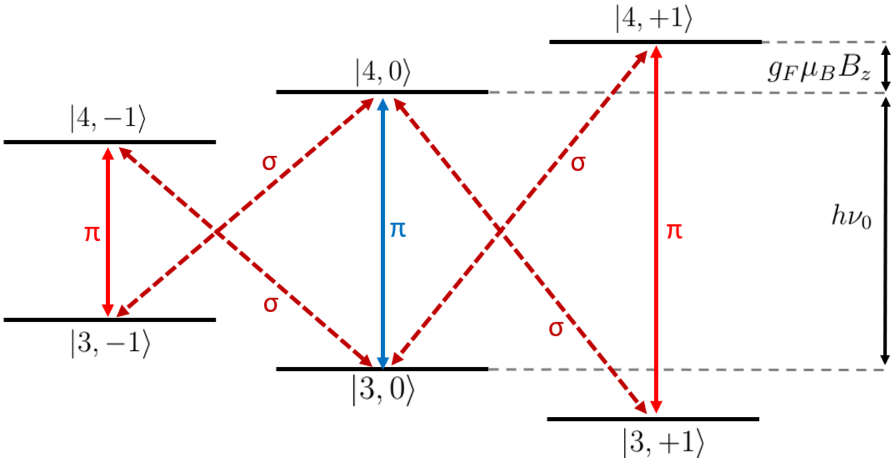

Figure 1.21: The $F = 3, 4$ and $m_F = -1, 0, 1$ states involved in Rabi pulling ($\pi$-transitions) and Ramsey pulling ($\sigma$-transitions). The levels are split linearly by an external weak field $B_z$.

### Rabi Pulling

The Rabi pulling effect comes from an imbalance of populations in the $|3, -1\rangle$ and $|3, +1\rangle$ states. Small populations could be off-resonantly transferred from $|4, \pm 1\rangle$ to $|3, \pm 1\rangle$ in the state-selection cavity, and survive the subsequent clearing laser pulse. Any asymmetry in these populations results in a sloping baseline of transition probability vs frequency, which alters the position of the central fringe.

This effect is distinct from the second-order Zeeman effect, blackbody shift, and collisional shift because it is not shifting the atomic energy levels, but nor is it shifting the phase of the microwave field, like for the DCP or cavity-pulling effects. Rather, it modifies the resulting transition probability as seen by the detection method, as it includes atoms that are not in, or did not start out in, the clock state.

Rabi pulling was an important systematic effect in caesium beam clocks where the population imbalance, velocity distributions, and lineshapes depended on the details of the state-selection magnets [160]. The Rabi pulling effect was around $10^{-12}$ for C-field strengths up to 8 $\mu$T, and generally reduced for stronger C-fields [160, 161]. However, increasing the C-field strength also increases the current stability required to keep the second-order Zeeman shift uncertainty low, such that the Rabi-pulling shift was a persistent problem. One technique for reducing this effect was to detect the third-harmonic of the sinusoidal frequency modulation, which is insensitive to the slope caused by adjacent transitions.

The situation is simpler in a fountain clock because all atoms have the same narrow velocity distribution. The population in each of the $|3, \pm 1\rangle$ states can be seen by the relative height of the Ramsey fringes of these detuned transitions, such that the total Rabi pulling shift can be calculated by [159]

$$\frac{\Delta\nu_{\mathrm{Rabi}}}{\nu_0} = \frac{1}{(2\pi)^6 \nu_0 T^2 \tau^2 (\Delta\nu_{1-1})^3} \frac{P_{+1} - P_{-1}}{P_0}, \tag{1.104}$$

where $T$ is the Ramsey time, $\tau$ is the cavity-passage time, $\Delta\nu_{1-1}$ is the detuning of the $|3, \pm 1\rangle \rightarrow |4, \pm 1\rangle$ transitions from the clock transition, and $P_{\pm 1,0}$ are the detected transition probabilities (by fluorescence collection) at the frequencies of the neighbouring $\pi$ transitions and clock transition, respectively. From equation 1.104, one can see that, for typical C-fields and interaction times, a 1% imbalance of the $|3, \pm 1\rangle$ populations, relative to the clock state population, results in a very negligible shift of the order of $10^{-21}$.

### Ramsey Pulling

The Ramsey pulling effect also originates in sub-level population imbalances: between the $|3, -1\rangle$ and $|3, +1\rangle$, or $|4, -1\rangle$ and $|4, +1\rangle$ states, which could be excited from the clock state by the small but unavoidable transverse magnetic field components of the $\mathrm{TE}_{011}$ Ramsey cavity mode, or by misalignment between the axial magnetic field components of the cavity mode and the background C-field. The clock state is coupled to these states by $\Delta m_F = \pm 1$, $\sigma$-transitions. This alters the populations of the clock states.

The result is a perturbation that has the shape of Ramsey fringes with a phase difference relative to the clock fringe that corresponds to the average phase difference between the clock state and other states [161], which may be up to $90^\circ$, and thus presents a slope across the central Ramsey fringe.

An analytical perturbative approach found that, while the general shape of Ramsey pulling vs C-field strength in a Cs beam clock was similar to that of Rabi pulling, the size was much bigger [161, 162]. A combined model for Rabi and Ramsey pulling in a fountain clock found that detuning the state-selection cavity, as well as inducing Zeeman coherences between the clock states and the magnetically-sensitive states, results in a relatively large shift around $10^{-15}$ [163]. In normal operation, the combined shift due to Rabi and Ramsey pulling was estimated to be around $10^{-18}$.

### Majorana Transitions

Majorana transitions are transitions that occur between different $m_F$ states of the same hyperfine level when the orientation of the external magnetic field changes much faster than the frequency of the Zeeman splitting, such that the atomic moment does not stay aligned with the field.

The resulting wavefunction can be described as a superposition of all $2F + 1$ available $m_F$ states with some probability amplitude that depends on the gyromagnetic ratio $g_F$ and on the magnetic field $\mathbf{H}(\mathbf{r}, t)$, but not on the initial $m_F$ state of the atoms [21, 164].

Majorana transitions can occur in a fountain clock if the sign of the vertical magnetic field component $H_z(\mathbf{r})$ changes sign, which may occur below the openings of the magnetically shielded region. Typically, the state-selection cavity is located inside the shielded region, which suppresses any additional Rabi and Ramsey pulling due to Majorana transitions. However, if Majorana transitions occur on the way down after the second Ramsey pass, and the resulting $m_F$ distribution is not symmetrical around $m_F = 0$, then uncancelled components of the resulting Ramsey pattern persist, and an apparent frequency error results due to different detection efficiencies of each $m_F$ state when using circular polarisation for detection [165]. This becomes particularly severe if the detection polarisation is not well-aligned with an axis of the optical fibre delivering it to the detection region [165]. However, Majorana transitions can be sufficiently suppressed by ensuring that the sign of $H_z(\mathbf{r})$ does not change throughout the atom's trajectory.

## 1.2.7 Microwave Leakage

Stray magnetic fields may come from the microwave electronics or cavities, which can affect the state of the atoms before, during, or after interrogation. The effective pulse areas are typical much smaller than the $\pi/2$-pulses undergone in the Ramsey cavity, but can still result in a noticeable frequency shift.

This effect was first recognised in thermal beam clocks [166, 167], where the averaging over the thermal velocity distribution meant that experimental observations of the effect were limited to microwave powers near optimum power. In a laser-cooled fountain, however, clear Ramsey fringes remain visible at many times the optimum microwave power.

One of the most likely places for microwave leakage to affect the atomic states is on the way down between the second Ramsey interaction and the optical detection. The way up is less important because state selection ensures a pure state enters the Ramsey cavity, and the space between the state-selection cavity and the Ramsey cavity is usually a waveguide-below-cut-off, where microwave fields at the clock frequency can't propagate. In this region, the change in observed transition probability between successive shots with angular frequency detuning $\pm\delta_m = \pm\pi/2T$ due to a leakage field is [168]

$$\delta P = 2\sin(\Omega\tau_{\text{eff}})\sin(\delta_m T)\int_0^{t_d}\Omega'_L(t)\cos(\delta_m t)dt, \tag{1.105}$$

where $\Omega$ and $\tau_{\text{eff}}$ are as defined in section 1.2.5, $T$ is the Ramsey time, $t_d$ is the time between the second cavity pass and detection, and $\Omega'_L$ is the component of the leakage field that is out of phase with the cavity field (a component in-phase causes a change in fringe contrast but not a frequency shift [169]). There are higher-order terms that have a quadratic dependence on $\Omega'_L$, but these are typically less important if $\Omega'_L t_d \ll \pi/2$. When discussing power-dependent effects, it is better to talk about $\delta P$ than $\Delta\nu$, because frequency shifts diverge at integer-$\pi$ pulses because the fringe contrast approaches zero. The microwave leakage model of equation 1.105 for the common case of $\Omega'_L \propto \Omega$ is shown in figure 1.22.

Microwave leakage can also occur during the Ramsey time, particularly if the vacuum vessel acts as a large microwave cavity to form a standing wave. There are certain conditions under which the shift due to microwave leakage in this region vanishes: if the leakage Rabi frequency $\Omega'_L(t)$ is symmetric about $T/2$, or if the pulse area inside the cavity is exactly $\Omega\tau_{\text{eff}} = (2n+1)\pi/2$. This effect can be accentuated if the pulse areas seen by an atom during each cavity transit are different, which is inevitable due to the transverse thermal expansion of the cloud during the flight, and the slight variation of the microwave fields strength over the cavity aperture [168]. However, this effect is generally much smaller than the effect of microwave leakage after the Ramsey interrogation.

Microwave leakage fields can be reduced by detuning [71] or switching off [69] the state-selection cavity when not in use, by having the atoms spend the entire Ramsey time in a waveguide-below-cut-off [132], and by taking precautions in the microwave synthesis electronics. The frequency bias due to microwave leakage is often assigned

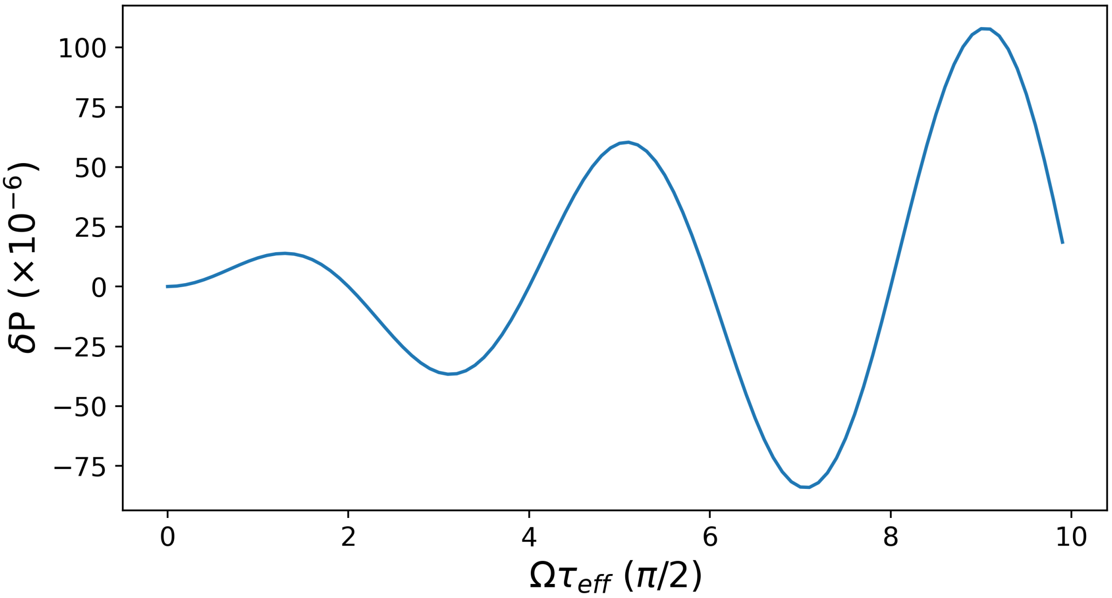

Figure 1.22: Change in transition probability $\delta P$ vs Ramsey cavity microwave power for $\Omega_L'$ 120 dB below the Ramsey cavity field amplitude $\Omega$, and $t_d = 0.2\,\mathrm{s}$.

a value of 0, with uncertainties around $1 \times 10^{-16}$ or so [76, 71, 69].

### 1.2.8 Microwave Spectrum Impurities

Spurious frequency components in the microwave field can affect the state of the atoms according to equation 1.3, such that, if there is an imbalance between the power of sidebands at $\pm\delta$, then the wings of this pair of Ramsey fringes overlap at the clock frequency and a frequency shift can occur, in a similar manner to Rabi and Ramsey pulling. Furthermore, a spectrum of frequency impurities mean that an atom sees many different microwave frequencies, which complicates the situation. This effect was first studied in depth for thermal beam clocks, and it was found that this shift depends on the sideband frequency [170, 171], the sideband power, the frequency modulation scheme, and the point in the microwave synthesis chain at which it is introduced [172].

The continuous nature of thermal beam clocks meant that the first-order contribution vanished, due the averaging over continuous arrival times of the atoms in the beam, such that one was left with only the smaller second-order contribution. For example, the sideband may have a fixed frequency relative to the carrier. In this case, a spur of detuning $\delta$ and Rabi frequency $\Omega_1 \ll \Omega$, results in a shift of [171, 172]

$$\Delta\omega_{\text{spur},2} = \frac{\tau}{T}\frac{\Omega_1}{\Omega}(y + z_1\cos(\delta T) + z_2\sin(\delta T)), \tag{1.106}$$

where $y$, $z_1$ and $z_2$ are functions of $\delta$ and $\Omega$, but not $\Omega_1$, according to

$$y = \frac{1}{\Omega\tau\sin(\Omega\tau)} \tag{1.107}$$

$$\times\left[\frac{\Omega}{\delta}\cos(\Omega\tau) + \frac{1}{1-(\delta/\Omega)^2}\left(\frac{\delta}{\Omega} - \frac{\Omega}{\delta}\cos(\Omega\tau)\cos(\delta\tau) - \sin(\Omega\tau)\sin(\delta\tau)\right)\right], \tag{1.108}$$

$$z_1 = \frac{\cos(\Omega\tau)}{\Omega\tau\sin^2(\Omega\tau)}\frac{1}{(1-(\delta/\Omega)^2)}\left((\cos(\Omega\tau)-1)\sin(\delta\tau) + \frac{\Omega}{\delta}(1-\cos(\delta\tau))\sin(\Omega\tau)\right), \tag{1.109}$$

$$z_2 = \frac{\cos(\Omega\tau)}{\Omega\tau\sin^2(\Omega\tau)}\frac{1}{(1-(\delta/\Omega)^2)}\left((1+\cos(\Omega\tau))(1-\cos(\delta\tau)) - \frac{\Omega}{\delta}\sin(\Omega\tau)\sin(\delta\tau)\right). \tag{1.110}$$

At optimal power ($\Omega\tau = \pi/2$) the $z_1$ and $z_2$ terms vanish, however in practice this is only achieved to within 1% or so. Figure 1.23 shows the shift vs the power of the spurious microwave frequency, It suggests that spurs closer to the carrier need to be suppressed by at least $-40$ dB, although it also suggests that there is not a simple relation between the shift at elevated microwave power ($\Omega\tau = 3\pi/2$, $5\pi/2$, ...) and the shift under optimal conditions.

This complexity is underscored by figure 1.24, which shows the shift vs spur frequency for different microwave excitation powers. In general, the shift is larger for frequency spurs with detuning $\delta \sim 2\pi/\tau$, such that spurs within the Rabi envelope of the Ramsey pattern should be kept below $-60$ dBc or so to ensure negligible frequency shift.

This effect is, however, only the second-order shift, which was relevant for thermal beam clocks. Since fountains operate in a cyclic manner, this opens up the possibility for shifts due to additional microwave frequencies that are synchronous with the fountain cycle. If some harmonic of $\delta \times T_c$ equals some integer $N\pi$, then the spur will be at the same phase in the next cycle. If the phase of the spur is different on the way up and the way down, then a non-zero frequency shift will not average out, according to

$$\frac{\Delta\nu_{\mathrm{spur},1}}{\nu_0} = \frac{1}{\pi\nu_0 T}\frac{\Omega_1}{\Omega}\sin\left(\frac{1}{2}\delta T\right)\cos\left(\frac{1}{2}\delta T + \phi\right). \tag{1.111}$$

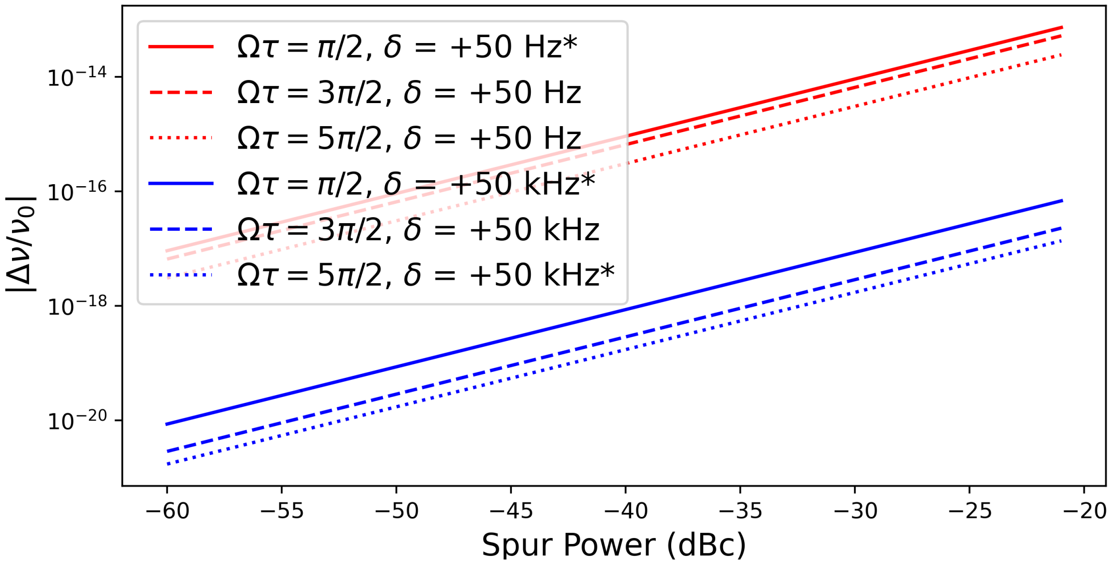

Figure 1.23: Plot of fractional frequency shift due to microwave spectrum impurities for varying spur power, at 50 Hz (red) and 50 kHz (blue). Asterisk denote where the shift is actually negative.

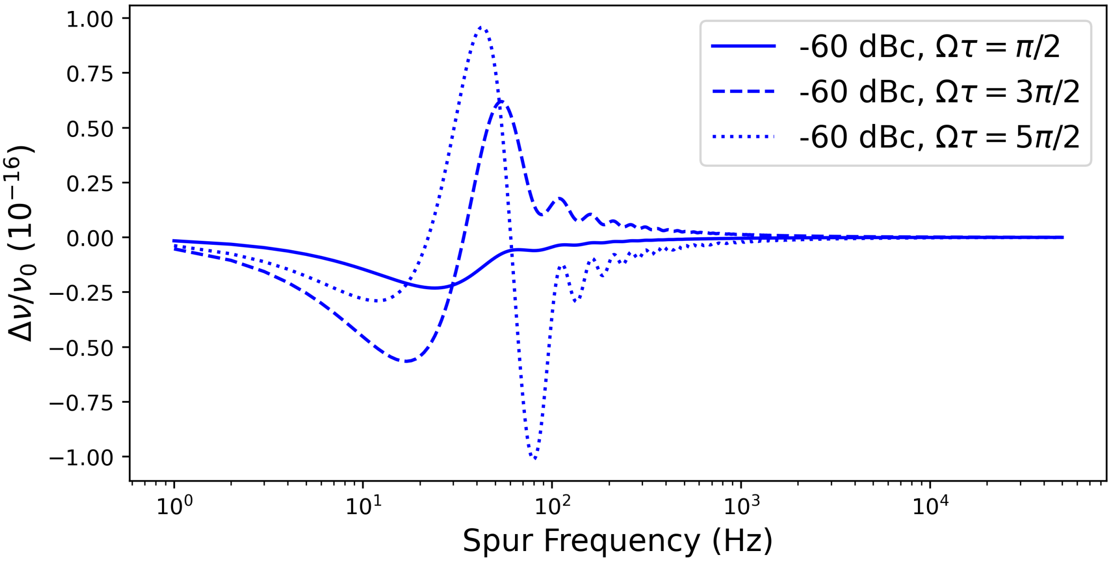

Figure 1.24: Frequency shift vs spur detuning, for different excitation powers and with a fixed relative spur power of −60 dB below the carrier.

For example, for a cycle time of 1.0 s, a Ramsey time of 0.5 s, and a 1.0 Hz spur at $-60$ dBc where the way up coincides with a maxima, and the way down with a minima, the frequency shift is maximised [173] and has size of $1.6 \times 10^{-14}$, which is very large. The cycle time $T_c$ can be adjusted to avoid synchronicity with large frequency spurs in the microwave field, such that the uncertainty associated with microwave spectrum impurities is typically below $10^{-16}$.

### 1.2.9 Microwave Lensing

With many of the larger systematic effects being evaluated with an uncertainty of only around $1 \times 10^{-16}$, some attention must be paid to the smaller, more subtle effects. One such effect is microwave lensing, which arises from the fact that the atomic wavefunctions are modified by the absorption of photons due to recoil inside the cavity. For an infinite plane wave, an atom recoils with momentum $\hbar k$, where, for Cs, the wavevector is $k = 192.66\,\mathrm{m}^{-1}$. This resulting energy $h\Delta\nu = \hbar^2 k^2/2m$ may lead one to naively assume a frequency shift of

$$\frac{\Delta\nu}{\nu_0} = \frac{\hbar k^2}{4\pi m \nu_0} = 1.5 \times 10^{-16}, \tag{1.112}$$

which is comparable to the quoted uncertainties of the most advanced fountain clocks. However, this is not the full story. The photons are not absorbed from an infinite plane wave, so the recoiling wavevector is not so straightforward. The upper and lower clock states end up recoiling slightly differently: they couple to different spatial wavepacket separations, which themselves correspond to different integer number $n$ of scatters. This is shown in figure 1.25. A frequency shift thus arises if there is a detuning-dependence on the phase of the different $|g, n\rangle$ and $|e, n\rangle$ components in the detected superposition of wavepackets [174].

There are various treatments of the microwave lensing shift that quantise the transverse motion of the atoms while keeping the vertical trajectory classical, either by numerically calculating the resulting shift from a single atomic wavepacket [174] or by taking many atoms from statistical distributions of position and momentum that depend on the atomic cloud's initial size and temperature distribution [175], or by making low-power approximations and finding an analytical solution [176]. In each case, the results are broadly the same: the opposite lensing of the lower and upper states in the superposition causes a partial cancellation of the recoil shift,

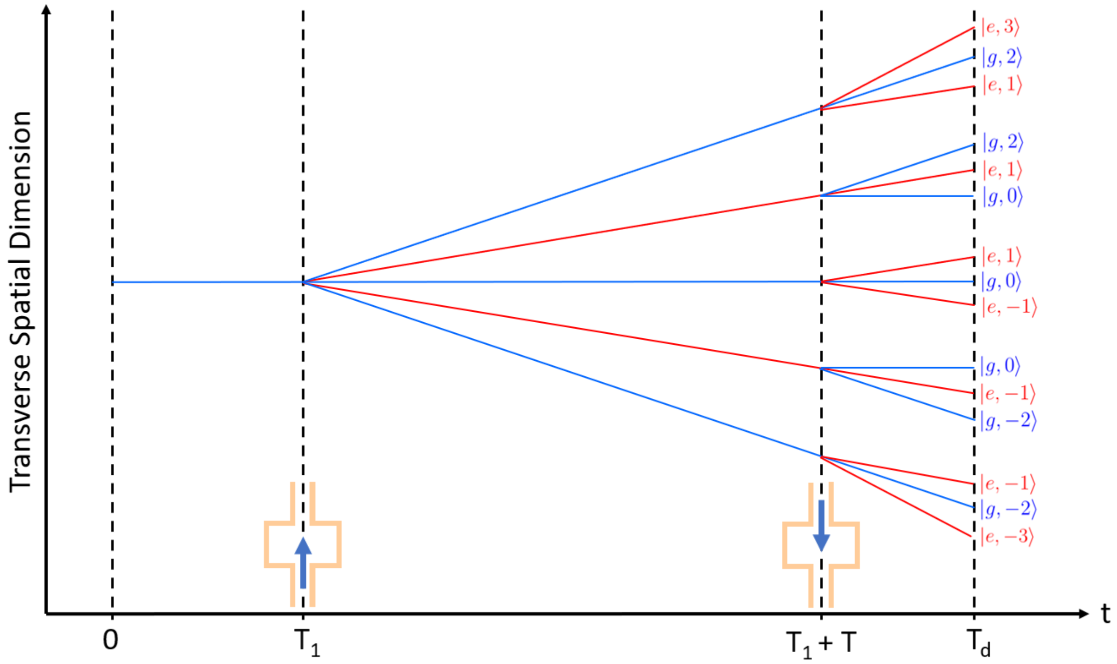

Figure 1.25: Coupling between internal states $|g\rangle$ and $|e\rangle$ to the atom's transverse velocity change, in multiples of the recoil momentum $\hbar k$, as a function of time. For clarity, two scatters are shown for the first interaction, and one for the second.

and the resulting shift is smaller than the simple plane-wave recoil shift of equation 1.112, in a way which depends strongly on the initial size and temperature of the cloud, as well as the times between the launch and the first Ramsey interaction, and between the second interaction and detection, and on the geometry of the detection system.

### 1.2.10 Gravitational Shift

The final systematic shift to consider is the gravitational shift. This is only noticeable when comparing clocks at different heights, as a sure-footed observer in an inertial frame measuring proper time need not concern themselves with the exact force of gravity in the local environment.

The fractional frequency difference between two clocks in (negative) gravitational potentials $W_1 < W_2$ is

$$\frac{\Delta\nu}{\nu_0} = \frac{W_1 - W_2}{c^2} \approx -\frac{g\Delta h}{c^2} \approx -1.06 \times 10^{-16} \frac{\Delta h}{\mathrm{m}}. \tag{1.113}$$

where $\Delta h = h_2 - h_1$ is the (positive) vertical displacement between the two clocks.

The SI second is defined on the geoid: an approximation of the average sea level at constant gravitational potential. Since most clocks are run by oxygen-breathing humans, they tend to be located at higher gravitational potential than the geoid, and the relevant correction to a primary frequency standard thus corresponds to the height between the clock and the geoid. The geoid is determined to different accuracies by both global and regional models, which use satellite tracking and altimetry, and surface gravity measurements [177]. Depending on the model used, the height of the geoid can vary by around 0.1 m, which results in a negligible frequency uncertainty of around $10^{-17}$ when compared to the uncertainty in the clock's height above the geoid, although this will be a more important limitation in the future, when the SI second may be redefined in terms of an optical transition with anticipated uncertainty of a few $10^{-19}$.

Atomic fountain clocks total about 2 m in height, so what is the correct height to use? The time-averaged height over the interrogation region, where the interrogation lasts for $T = 2v_c/g$ (where $v_c$ is the speed of the atoms through the cavity), such that

$$\langle z(t) \rangle = \int_0^T v_c t - \frac{1}{2} g t^2 \, dt = \frac{v_c^2}{3g} = \frac{2}{3} z_{\mathrm{max}}, \tag{1.114}$$

is the height above the centre of the cavity that corresponds to "the clock's" height above the geoid. The launch height is generally known to better than 1 cm, which adds negligible uncertainty.

A similar shift results from special-relativistic time dilation, as the atoms are moving through the radiation field in the cavity with a speed $v_c$, which nominally has a frequency $\nu_0$ only in the stationary lab frame. This is sometimes called the second-order Doppler shift (where the first-order is the DCP shift). This shift can be expressed as [178]

$$\frac{\Delta \nu_{\mathrm{td}}}{\nu_0} = -\frac{v_c^2}{2c^2}, \tag{1.115}$$

such that, for a typical Ramsey time of 0.5 s resulting from $v_c = 2.9\,\mathrm{ms}^{-1}$, the fractional special-relativistic shift is $-0.47 \times 10^{-16}$. As with the gravitational redshift, accurate knowledge of the launch height means that the uncertainty resulting from the time-dilation shift is negligible.

# Chapter 2

# Designing the Mini-Fountain

Atomic fountain clocks are large apparatus that have been developed over many years to achieve the best microwave frequency accuracy possible. In several areas, it should be possible to simplify the design to reduce the size of the device without severely degrading the clock's performance. In this chapter, the considerations behind the design choices of the mini-fountain are discussed and justified.

## 2.1 Compact Normalised Detection Method

Traditionally, detection of each clock state has been performed in a dedicated detection chamber, located between the interrogation region and the cooling chamber. This has the advantage of a lower background fluorescence due to thermal atoms in the detection signal, however it also adds extra bulk to the physics package. Additionally, there are two separate beams (and corresponding optical fibres), such that approximately $10^3$ photons are scattered and detected per atom, and the signal-to-noise ratio due to detection exceeds 2,000 or so. If the fountain is used with an ultra-stable microwave source such as a cryogenic-sapphire oscillator (CSO), and loads cold-atoms with a single-stage or multi-stage MOT, then the short-term stability reaches the quantum-projection noise limit of around $4 \times 10^{-14}/\sqrt{\tau/\mathrm{s}}$ [64, 68, 66].

With a quartz local oscillator, one expects a short-term stability limited by local oscillator noise to around $2 \times 10^{-13}/\sqrt{\tau/\mathrm{s}}$, corresponding to an SNR of around 500. A detection method that can achieve a very high SNR does not improve the fountain's short-term stability if other sources of noise are limiting it to a lower level. The aim of this section is to describe the development of a different, novel detection scheme that should achieve an SNR of at least 500, while not adding any additional bulk or complexity to the physics package by performing a normalised detection method using the MOT beams in the mini-fountain's MOT chamber.

### 2.1.1 Pulse Sequence

The traditional method separates the detection sequence in space, and uses beams that are not time-varying during the atom's transit through the detection region. Instead, a key innovation of the mini-fountain is to use a time-varying sequence of laser pulses in the MOT chamber.

This novel method is illustrated in figure 2.1: firstly, when the atoms re-enter the MOT beam region after interrogation, a pulse of the cooling light from all six MOT beams induces fluorescence from the upper-clock state, while atoms in the lower state continue to fall. This fluorescence is proportional to the population in the upper level. After a short time (around a millisecond), the cooling beams are turned off and a short pulse of repump light puts all the atoms into the upper level, including those that were initially in the lower clock state, such that another pulse of the cooling light induces fluorescence proportional to the total atom number. The beams are then turned off and the atoms allowed to free-fall, such that when the beams are turned back on, the background fluorescence is measured. This fluorescence can be collected by lenses and focused onto a photodiode, such that the computer controlling the experiment can process the data to extract the transition probability. An illustration of the voltage vs time graph for all atoms initially in the upper level, which in Cs is $F = 4$, is shown in figure 2.2. An exponential decay of the fluorescence is expected, as without the repump some Cs atoms will be off-resonantly excited

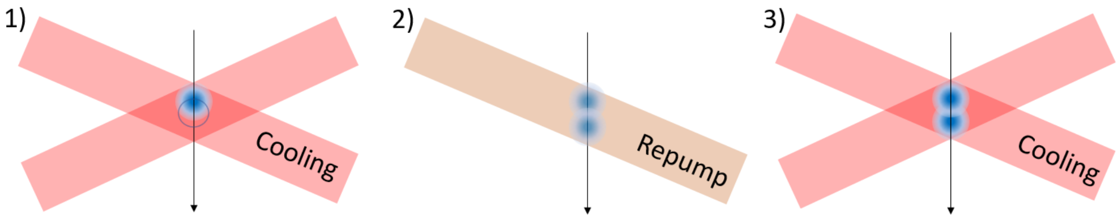

Figure 2.1: Illustration of the detection beam sequence. 1) Cooling light extracts fluorescence proportional to the upper level population. 2) Repump light transfers all population to the upper level. 3) Cooling light extracts fluorescence proportional to the total atom number.

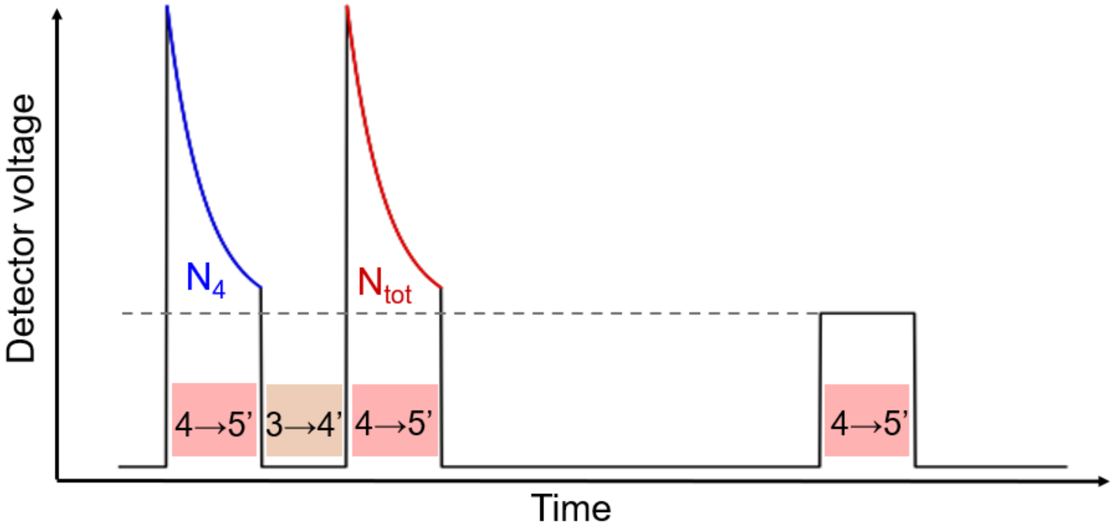

Figure 2.2: Voltage vs time seen by a photodiode for this compact detection method.

into the $F' = 4$ level, from which they can decay to the $F = 3$ level and no longer fluoresce. The scattering rate for an atom in a near-resonant light beam is [24]

$$R_s = \frac{\Gamma}{2}\frac{I}{I_{\text{sat}}}\frac{1}{1 + 4(\delta/\Gamma)^2 + I/I_{\text{sat}}}. \tag{2.1}$$

For $I \gg I_{\text{sat}}$, the scattering rate on the main cycling transition, $F = 4 \to F' = 5$ approaches $\Gamma/2$, and increasing $I$ does not increase the scattering rate. However, the nearest off-resonant transition in Cs is the $F = 4 \to F' = 4$ transition, at a detuning from the cycling transition of 251 MHz, such that $(\delta/\Gamma)^2 \gg I/I_{\text{sat}}$, such that the scattering rate for the off-resonant transition becomes

$$R_{\text{off}} = \frac{\Gamma^3}{8\delta^2}\frac{I}{I_{\text{sat}}}, \tag{2.2}$$

which is proportional to $I$ and does not saturate at moderately high intensity. Thus, the exponential decay rate of the fluorescence will be proportional to the intensity. It is anticipated that there will be some optimal intensity $I \gtrsim I_{\text{sat}}$ at which the area under the voltage-time curve is maximised.

Additionally, the scattering of photons heats the cloud and cause it to disperse, which reduces the signal extracted from the atoms later in the pulse sequence. To mitigate this, the detection beams are detuned a small amount to the red of resonance, such that a molasses is formed to cool the cloud.

Performing detection in the MOT chamber leads to a significant background fluorescence due to the thermal alkali atoms, from which the MOT was initially loaded. Only a small fraction of these atoms fluoresce, due to the broad Doppler detuning, however there are a large number of them. At a lab temperature of around 300 K, and an alkali vapour partial-pressure of $P = 10^{-9}$ mbar, the ideal gas law

$$P = nk_BT, \tag{2.3}$$

means that the number density $n \approx 2.4 \times 10^7\,\mathrm{cm}^{-3}$, which will be greater than the number density of returning cold atoms. Another contribution to the background fluorescence is stray light reflected by the metal walls of the vacuum chamber, which depends on the laser intensity but not its frequency.

### 2.1.2 Initial Tests with Demonstrator System

A pre-existing UHV caesium MOT system with an optical control system similar to that of a (1,1,1) geometry fountain clock was used to test this detection method. This system comprised of two extended-cavity diode lasers; one for generating the cooling light on the $6\,^2S_{1/2}\,|F=4\rangle \rightarrow 6\,^2P_{3/2}\,|F'=5\rangle$ transition, and another for generating the light for the repump $6\,^2S_{1/2}\,|F=3\rangle \rightarrow 6\,^2P_{3/2}\,|F'=4\rangle$ transition, each locked by saturated absorption spectroscopy. The spectroscopy set-up for this "demonstrator system" is shown in figure 2.3.

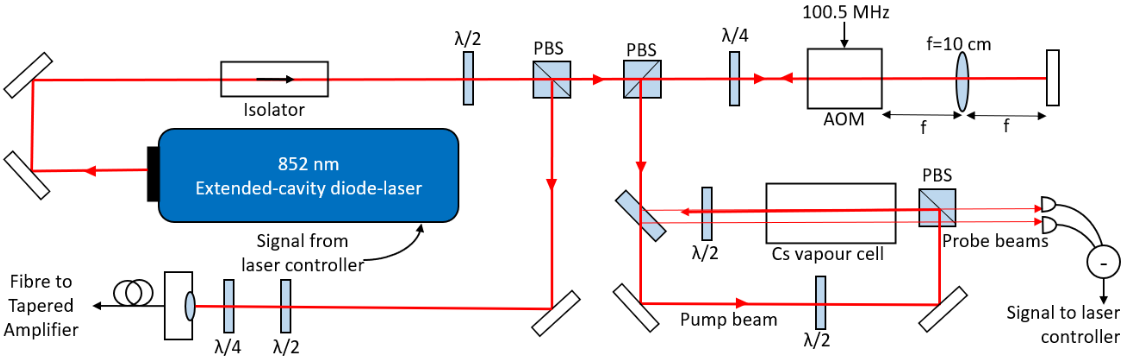

Figure 2.3: Laser and spectroscopy set-up for locking to the cooling transition. A similar setup, minus the double-pass AOM section, is used for the repump. $\lambda/2$ = half-wave plate, $\lambda/4$ = quarter-wave plate, PBS = polarising beam splitter, AOM = acousto-optic modulator.

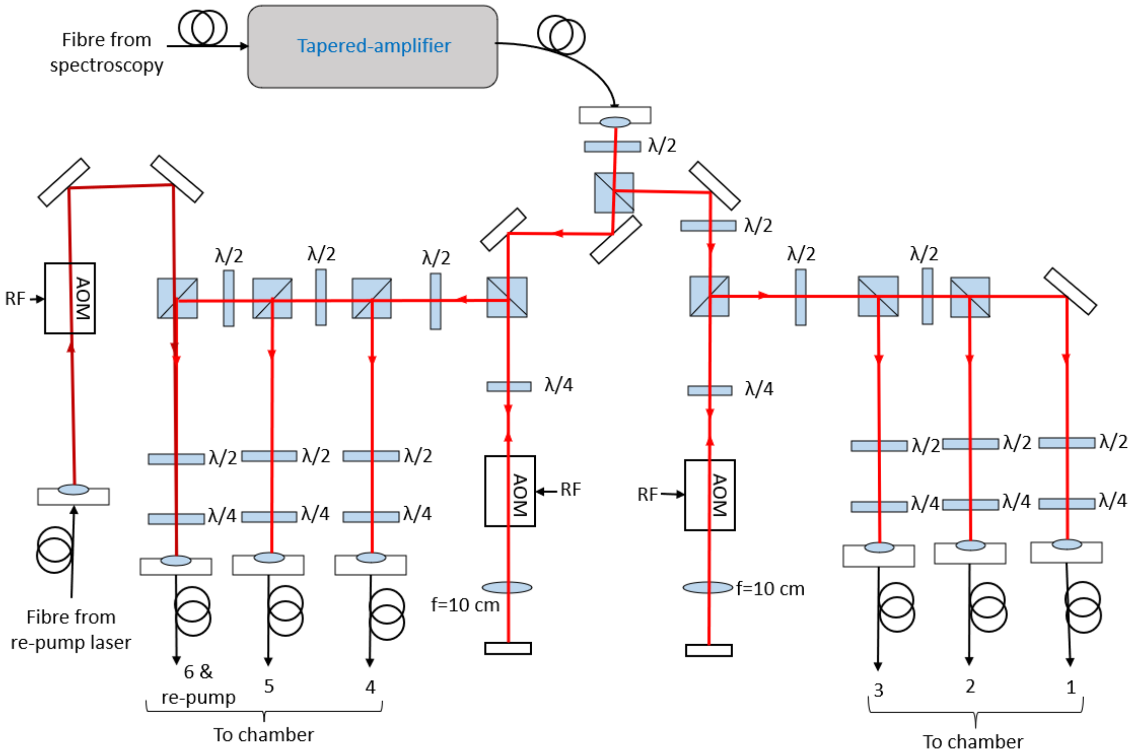

Figure 2.4: Splitting and frequency control for a (1,1,1) geometry atomic fountain clock. Laser light locked above resonance is amplified by the tapered amplifier, and split into two paths corresponding to the up and down-going MOT beam. Free-space AOMs bring the light back close to resonance, and each path is further split into three separate beams. The repump is locked below resonance and brought back to resonance by another AOM, and overlapped onto cooling beam 6. Each cooling beam has a final power of 10 mW, and the repump 2.5 mW.

Because the frequency-shifted light is locked onto resonance after a double-pass AOM, the cooling laser is 201 MHz below the transition frequency. It is brought back to (near) resonance by another pair of AOMs that control the frequency of the three up and down-going beams separately, as shown in figure 2.4. The geometry of this system is such that each MOT beam makes a $\pm 54^\circ$ angle to the vertical, so the moving-molasses technique requires that, for a launch velocity $v$, the three up-going beams be detuned by $\delta\nu = +\cos(54^\circ)v/\lambda$ and the three down-going beams be detuned by $\delta\nu = -\cos(54^\circ)v/\lambda$. The beams are delivered to the chamber by polarisation-maintaining optical fibres. They exit the fibres with linear polarisation, which is turned into the appropriate circular polarisation and collimated by 10 cm converging lenses.

The MOT uses a pair of anti-Helmholtz coils, each of 9.6 cm diameter, 741 turns, spaced 19 cm apart, to produce the spherical quadrupole magnetic field required for trapping. The optimal MOT coil current for maximising atom number in the trap was found to be 2.0 A. Having the spacing much larger than the coil radius is not ideal, but the shape of the (1,1,1) geometry MOT chamber does not easily allow for the ideal spacing. This is shown in figure 2.5.

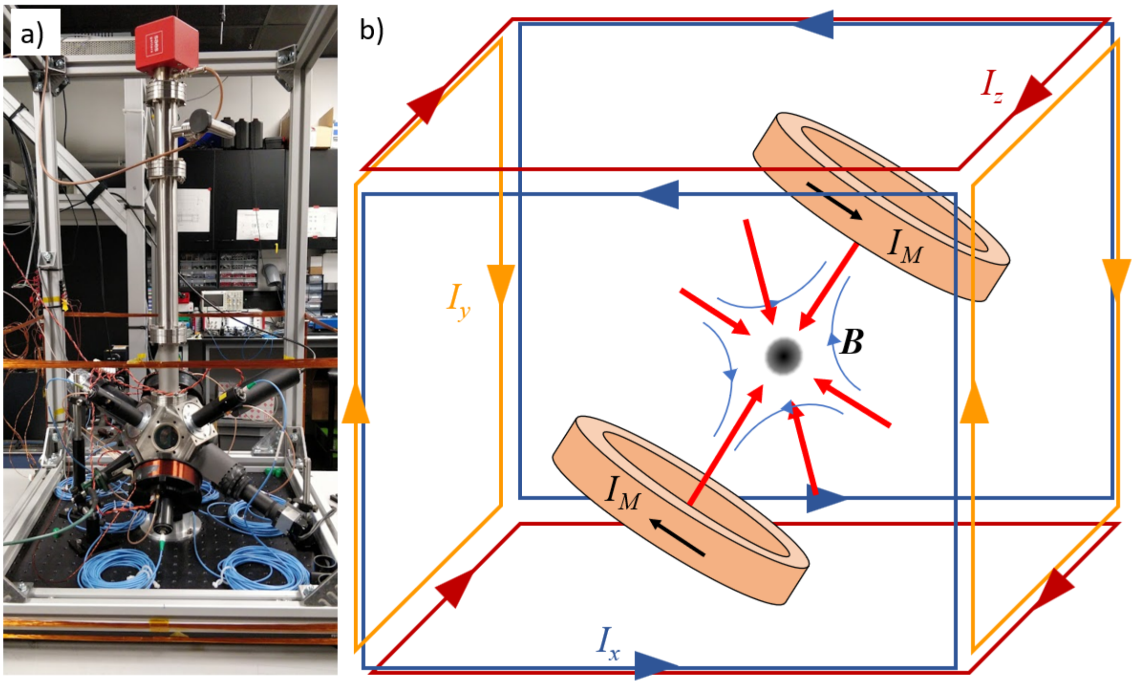

Figure 2.5: a) Photo of the (1,1,1) geometry demonstrator system, for initial tests of the compact detection method. b) Diagram of the MOT coils, MOT beams, resulting MOT field, and three pairs of rectangular Helmholtz coils, of spacing and diameter approximately 40 cm, rectangular compensation coils made from enamelled copper wire, wrapped around the support structure, each carrying a current of $I_x$, $I_y$ and $I_z$ in the two transverse directions and vertical direction, respectively.

Achieving the coldest atoms requires that many variables be well-controlled. The power in each counter-propagating beam must be balanced, as any an imbalance can drastically alter the motion of the atoms. This is controlled by manually rotating circular polarisers (quarter wave plates with a linear polariser mounted matched to an axis) inside the lens tubes of the top beams. Similarly, the alignment of each pair must also be accurate. This is controlled by kinematic fibre mounts in the ends of the lens tubes. The magnetic field compensation is adjusted by changing the current through each of three orthogonal pairs of coils, such that the magnetic field vector produced in the centre can cancel out any background field vector. The final way to affect the temperature of the cloud is the ramp-downs of intensity and detuning, as the intensity and detuning that result in optimal loading are drastically different from those that result in the coldest atom temperature.

As would be the case in a fountain clock, the first stage consists of 0.5 s of MOT loading, in which time approximately $10^8$ atoms are loaded. The laser beams are shut off for a few milliseconds, such that the eddy currents induced by turning off the MOT coils don't disturb the cold-atom cloud. This is followed by a further 30 ms of cooling in an optical molasses. Next, the detuning is increased and the intensities ramped down to implement sub-Doppler cooling.

The temperatures at each stage are measured with a CCD camera: by the end of the sub-Doppler cooling phase, the beams are turned fully-off and the cloud undergoes a period of free-expansion $t$, such that its radius $r(t)$ once the beams are restored can be used to measure its temperature according to

$$r(t)^2 = r_0^2 + v^2 t^2, \tag{2.4}$$

where $r_0$ is the cloud's initial size and

$$v = \sqrt{k_B T/m}, \tag{2.5}$$

is the cold-atom's mean thermal velocity. Fitting equation 2.4 to data of $r$ vs $t$ thus estimates $T$. This data for the MOT and sub-Doppler cooled molasses phases of the demonstrator system is shown in figure 2.6. The relatively hot final temperature of the molasses suggests that there is something slightly wrong with the set-up. Potential sources of imperfections could be laser frequency noise, instability in the beam intensity, or, most likely, uncancelled magnetic field gradients due to the very-magnetic screws attaching the Cs dispensers to the power feed-through flange.

Because the temperature was higher than ideal, only small launch heights resulted in detecting a large number of atoms. The detection pulse sequence was implemented for 100 ms launch times, such that for a 0.5 s MOT loading time, the returning detected signal under both peaks was $S_{\mathrm{det}} \approx 10\,\mathrm{mVs}$.

In the absence of state-selection and microwave interrogation, the cloud is launched and detected all in the upper level. If one were to hypothetically implement clock

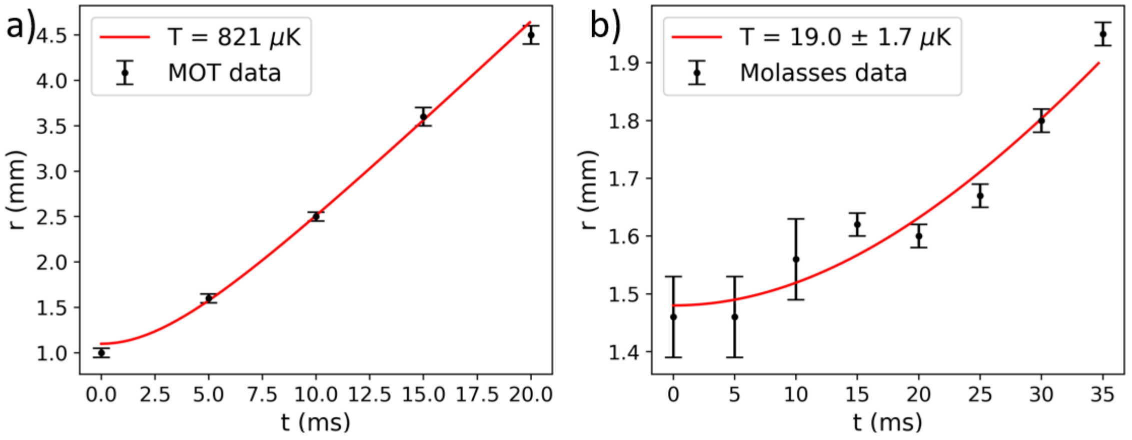

Figure 2.6: Radius vs free-expansion time for the MOT (left), with a temperature of 821 ± 16 $\mu$K, and the molasses (right) with a temperature of 19.0 ± 1.7 $\mu$K.

operation, equation 1.46 would reduce to

$$\sigma_y(\tau) = \frac{1}{\pi Q_R}\sqrt{\frac{T_c}{\tau}}\frac{1}{\mathrm{SNR}}, \tag{2.6}$$

where

$$\mathrm{SNR} = \frac{N_{\mathrm{at}}}{\sqrt{\sigma_{\mathrm{const}}^2 + N_{\mathrm{at}}^2\sigma_{\mathrm{prop}}^2}}, \tag{2.7}$$

where $\sigma_{\mathrm{const}}$ is constant noise that does not depend on the cold atom signal, such as electrical noise and background fluorescence noise, and $\sigma_{\mathrm{prop}}$ is noise that scales with the cold atom signal, such as laser noise. The SNR can be experimentally measured by using the Allan deviation: shot-to-shot, this is equivalent to the standard deviation if the data is stationary with only white noise, but stationarity is not guaranteed to be the case in practice. The signal-to-noise ratio is thus the noise, the Allan deviation of the observed fluorescence ratios at $\tau = T_c$, divided by the signal, the hypothetical clock transition probability (nominally 1.0).

The SNR for each set of detection beam parameters was measured by taking a run of 300 shots, and discarding the first 10, as thermal effects in the AOMs mean that the atom numbers are lower at the very beginning of a run. The optimal SNR was found to occur at intensities of 5.4 mW per beam, a detuning of $-0.8\,\Gamma$, with detection pulse lengths of 1.5 ms each. There was only a weaker dependence on the repump pulse length (0.3 ms), background pulse length (4.0 ms) and free-fall time

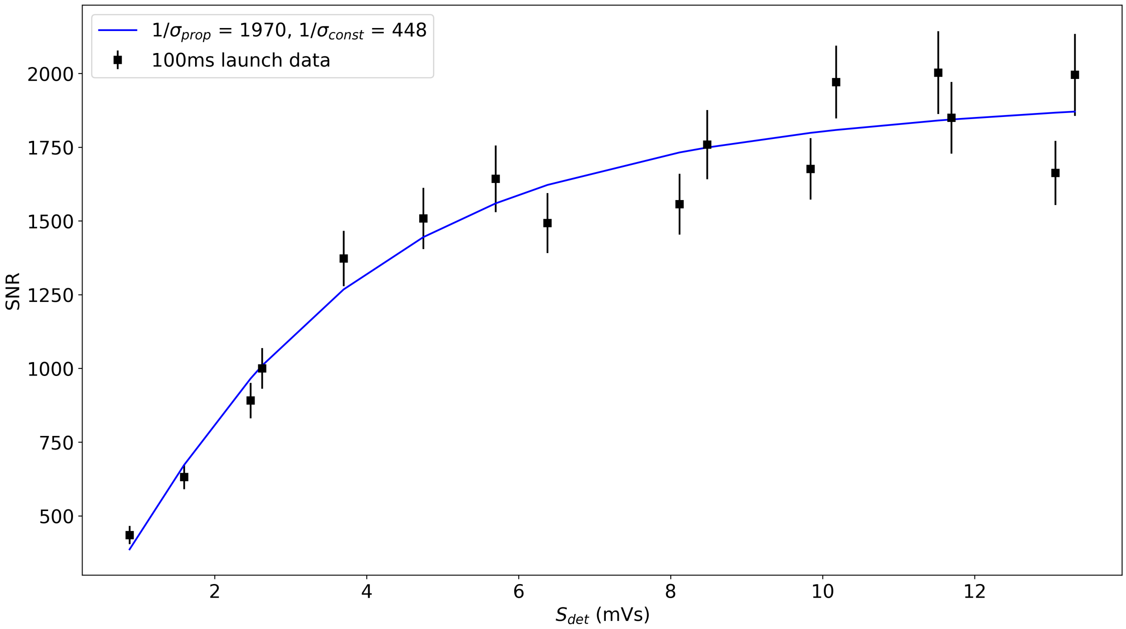

Figure 2.7: SNR vs $S_{\text{det}}$ for detection method tests in the demonstrator system after 100 ms launches at optimised detection pulse frequency and intensity.

before the background pulse (22 ms). The atom number was varied by adjusting the MOT loading time from 0.02 s to 0.7 s, and the resulting SNR vs returning signal $S_{\text{det}}$ is shown in figure 2.7.

While the large SNR $\approx$ 1,700 for $S_{\text{det}} \approx 10$ mVs (loading time of 0.5 s) is encouraging, at this relatively short launch time of 100 ms, almost all of the atoms are returning to be detected. A mini-fountain that maintains the 30 cm free-flight above the cavity may launch for a total of 600 ms or so, after which time the cloud would have expanded to a much larger size. Additionally, a state-selection method that only allows the $m_F = 0$ atoms to leave the MOT region means that the actual $S_{\text{det}}$ would only be 1/9th of the $S_{\text{det}}$ for the same loading time, which on figure 2.7 would correspond to an SNR below 500, even at this short launch time.

### 2.1.3 Further Tests with NPL-CsF3

To test the compact detection method in a more realistic setting, colder atoms are needed to see a returning signal after a longer launch. This can be achieved by using NPL-CsF3, a full-sized primary frequency standard. The design and operation of NPL-CsF3 is very similar to that of NPL-CsF2, which has been described extensively

[76, 84, 68].

In order to implement the compact detection method with the MOT beams, the fountain's original repump beam, which is overlaid with one of the cooling beams, needs to be blocked, and replaced by another repump laser which passes through an AOM, such that it can be quickly turned on and off, as the method requires. This laser was an 852 nm DFB laser, and it was locked to the $F = 3 \rightarrow F' = 3$ & 4 crossover resonance located 100.5 MHz below the repump transition, as the AOM can only fully intensity-modulate the diffracted beams. Similarly, a 2" diameter lens was placed as close as the vacuum chamber allowed from the centre of the MOT, such that fluorescence from 1% of the total solid angle was focussed onto a photodiode.

The down-going, vertical MOT beam in CsF3 passes through the Ramsey cavity, which has 11 mm diameter openings in the endcaps for the atoms and beams, so the size of the trapping region in the MOT is reduced. The loading rate is proportional to $r^4$, thus a reduced number of atoms is loaded for the same loading time. To compensate for this, the "normal" MOT loading time was increased to 1.2 s, to match the same fluorescence signal during loading (for the same collection solid-angle) as in the demonstrator system after 0.5 s of loading.

The temperature of the launched atom cloud in this system was measured by the standard time-of-flight method: a horizontal beam of vertical extent much smaller than the returning cloud results in a Gaussian voltage vs time plot where the width is proportional to the temperature (as in figure 1.12), which was found to be $T = 1.4 \pm 0.3\,\mu$K, which is what one would typically expect of a fountain clock.

With these colder atoms, a launch time of 600 ms, (height above the MOT of 44 cm) was achieved, and the detection pulse sequence was applied to the returning cloud. The optimal intensities were found to be 4.6 mW per beam, with detunings of $\delta = -1.8\,\Gamma$. The optimal detection pulse time was found to be 0.4 ms, much less than for the shorter launches in the demonstrator system. This is counter-intuitive, because one would assume that a larger area under the voltage vs time curve would result in a better SNR. There are two potential explanations for this. The first is because the off-resonant pumping is far from saturation, any intensity instability has a disproportionate effect on the amount of fluorescence collected later in the detection pulse, such that the later sections add disproportionately more noise than signal. The second is because the cloud is falling through the beams faster ($2.9\,\mathrm{ms}^{-1}$) than in the demonstrator system ($0.49\,\mathrm{ms}^{-1}$), there is less time in which the central, highest-density portion of the returning atom cloud is overlapped with the central, highest-intensity portion of the MOT beams, such that extending the pulse time collects more light that similarly adds disproportionately more noise to the integrated fluorescence than it does signal.

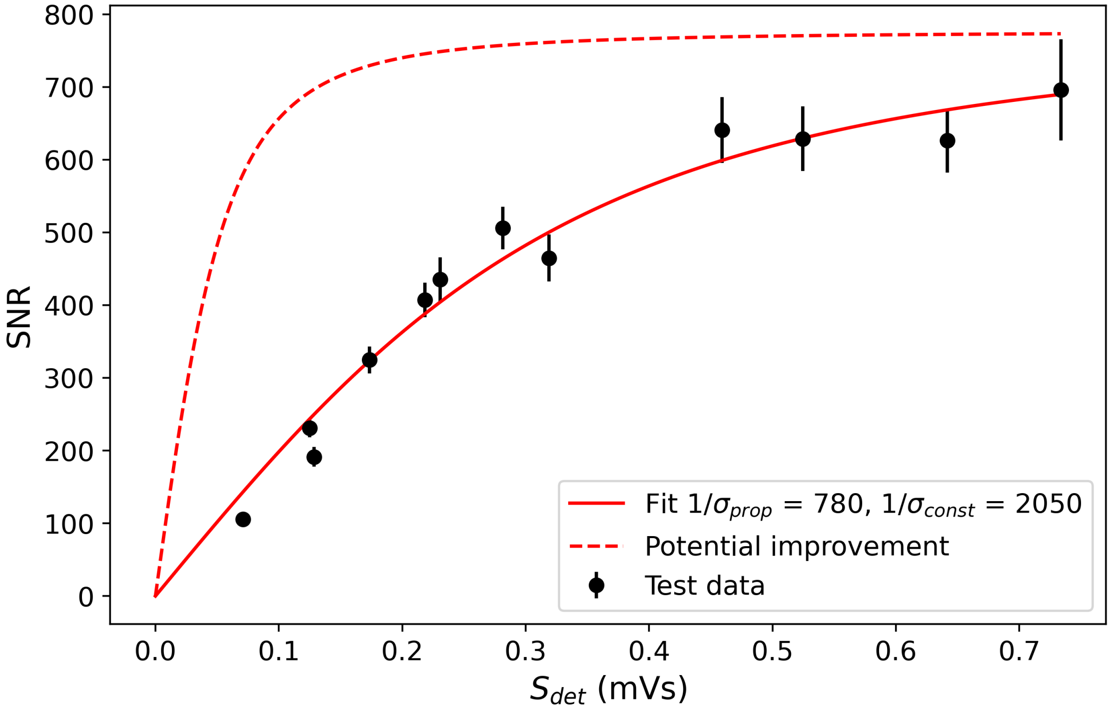

Figure 2.8: SNR vs $S_{\text{det}}$ by varying the MOT loading time between 0.1 s and 2.0 s.

There is not a simple, linear relationship between loading time and $S_{\text{det}}$ because the loaded atom number is only linear for loading times much smaller than the loss time constant, $\tau_L$, from the trap. However, loading for longer does generally result in more atoms.

The result, shown in figure 2.8, is that for higher $S_{\text{det}} > 0.5\,\text{mVs}$, the SNR approaches an upper limit of around 700, with the calculated fit of equation 2.7 resulting in $1/\sigma_{\text{prop}} = 780$. At lower $S_{\text{det}} < 0.2\,\text{mVs}$, the SNR increases linearly, as expected when the noise is dominated by technical sources of a constant size.

The 1.2 s loading time corresponded to $S_{\text{det}} \approx 0.5\,\text{mVs}$, such that, after state selection, the expected returning signal would only be 0.05 mVs, which has an SNR $\approx$ 100. This means that, while a suitable SNR has been demonstrated to be theoretically possible using just the MOT beams, some more design work is needed to realise this in practice.

## 2.2 Chamber Design

For the MOT chamber to perform trapping and detection, it is imperative that the MOT beam diameter of 24 mm be preserved, while maximising the fluorescence collection solid angle. The MOT chamber also serves as the central part to which all other parts of the fountain are attached, so the design choices made here will have an effect on many other parts of the system.

Both the Demonstrator system and CsF3 use commercially-available (SLS Optics) 32 mm diameter anti-reflection coated N-BK7 viewports through which the MOT beams pass to enter the vacuum chamber. To seal a glass viewport to the MOT chamber requires an indium wire seal, which is soft enough to not require an amount of pressure to seal that will crack the glass. A 1 mm diameter indium wire typically expands in width to around 3 mm after sealing, such that at least 3 mm of flat metal should be allowed beneath the edges of the viewports (as illustrated in figure 2.9). This means that, a 32 mm diameter viewport allows up to 26 mm of clear aperture. Thus, the same 32 mm viewports are suitable for the mini-fountain.

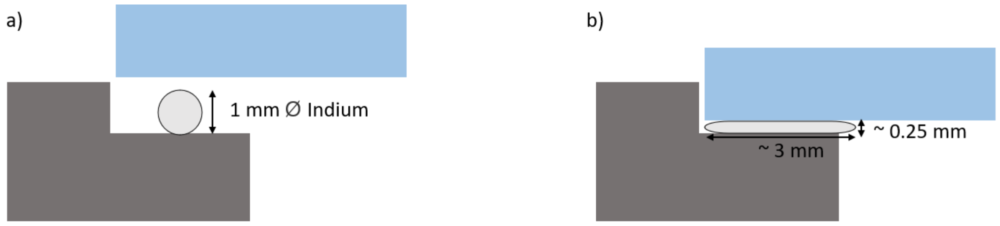

Figure 2.9: a) 1 mm indium wire beneath the viewport before sealing. b) Wire is compressed to approximately 3 mm width and 0.25 mm thickness.

The fluorescence-collection solid angle in both the Demonstrator system and CsF3 was approximately 1% of the total $4\pi$ steradians. Figure 2.8 shows that, after state selection reduces the returning atom number, it is expected that the SNR will be limited by constant technical noise, which can be improved by increasing the collected fluorescence. The geometry of this situation is shown in figure 2.10. The solid angle $\Omega$ subtended by a lens of radius $r$ placed $d$ from the centre of the MOT is calculated by finding the area of the sphere enclosing the lens that is blocked by the lens. For $d \gg r$, this area is approximately the area of the lens itself, such that $\Omega = \pi r^2/d^2$. However, for $r \sim d$, this is not a good approximation and slightly underestimates the total solid angle. In general, $\Omega = A/R^2 = \iint \sin\theta d\theta d\phi$, such

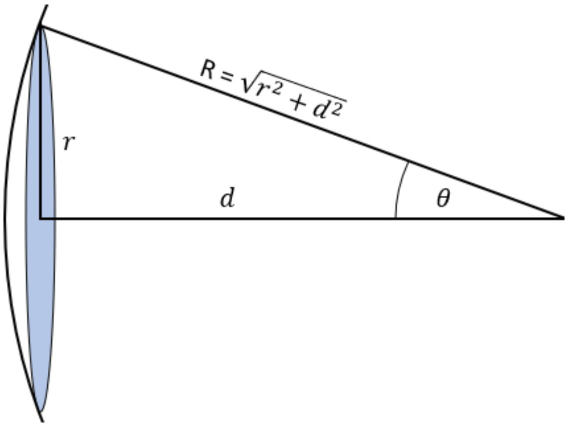

Figure 2.10: A lens of radius $r$ placed a distance $d$ from a point-source of illumination subtends an angle $\theta = \tan^{-1}(r/d)$.

that the area of the sphere of radius $R = \sqrt{r^2 + d^2}$ subtended by our lens is

$$A = \int_0^{2\pi} d\phi \int_0^{\tan^{-1}(r/d)} R^2 \sin\theta d\theta, \tag{2.8}$$

which can be evaluated as

$$A = 2\pi R^2 \left(1 - \frac{d}{\sqrt{r^2 + d^2}}\right). \tag{2.9}$$

This means that the fraction of collected fluorescence is

$$\frac{\Omega}{4\pi} = \frac{1}{2}\left(1 - \frac{d}{\sqrt{r^2 + d^2}}\right). \tag{2.10}$$

It is clear that minimising $d$ and maximising $r$, subject to the constraints imposed by the viewports and other aspects of the chamber, is the way to maximise the collected fluorescence, and thus obtain a better SNR during detection.

To collimate the light scattered from atoms in the centre of the chamber, one wants a focal length $f = d$. Another lens can then focus this light onto a photodiode. As is the case in most caesium fountains, the detection solid-angle can be doubled by placing another lens of the same focal length on the opposite side of the vacuum chamber, with a mirror to retro-reflect the scattered light back through the centre of the chamber and towards the detector, as illustrated in figure 2.11. Thus, the MOT (& detection) chamber requires a total of 10 ports: 6 for MOT beams, 2 for

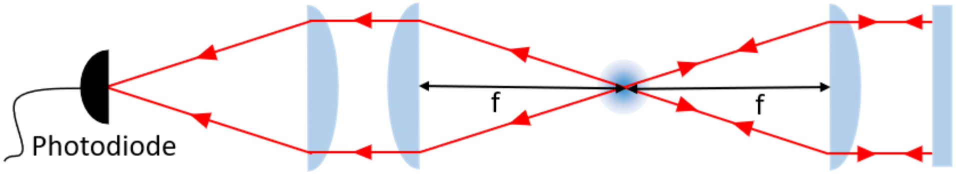

Figure 2.11: Fluorescence collection by two lenses of focal length $f$ placed $f$ from the centre of the chamber.

detection lenses, 1 for atom passage upwards to the microwave cavity and flight tube, and another for the introduction of alkali vapour into the chamber.

There are three potential geometries for atomic fountain clock cooling chambers. The (1,0,0) geometry has 2 pairs of horizontal beams and 1 pair of vertical beams that propagate through the entire interaction region, like CsF2 and CsF3. The (1,1,1) geometry has 3 pairs of beams inclined at 54° to the vertical, like the demonstrator. The other is the (0,1,1) geometry, which has 1 pair of horizontal beams and 2 pairs of beams inclined at 45° to the vertical. Both the (1,1,1) and (0,1,1) geometry have the advantage of not requiring a beam to propagate through the interrogation region, which will be useful in section 2.5. From experience with the demonstrator system, the (1,1,1) geometry does not lend itself well to accommodating both a vertical passage to the cavity and the proper MOT coil geometry (coil separation equal to the coil radius). For this reason, the (0,1,1) geometry is chosen for the mini-fountain's MOT/detection chamber. An advantage of this geometry is that the smaller angle to the vertical increases the effective height of the MOT beam region, which allows more time to perform state selection on the initially-launched cloud. A CAD model of the final design of the MOT/detection chamber is shown in figure 2.12. To the best of the author's knowledge, this geometry has not actually been implemented in a fountain clock before.

The lens used is 30 mm in diameter and has a back-focal length of 34 mm, which is close to the separation between the centre of the chamber and the lens's flat surface. The limiting aperture has $r = 12\,\text{mm}$, and is $d = 33.5\,\text{mm}$ from the chamber centre, such that, by equation 2.10, each lens collects 2.93% of the fluorescence, so the total collection efficiency of the system is 5.86%, nearly a six-fold increase from the demonstrator and CsF3 detection tests. While this is still much lower than the 30% achieved in NPL-CsF2's dedicated detection chamber [76], it is sufficient to greatly increase the SNR at small $S_{\text{det}}$. One thus expects an approximate SNR due

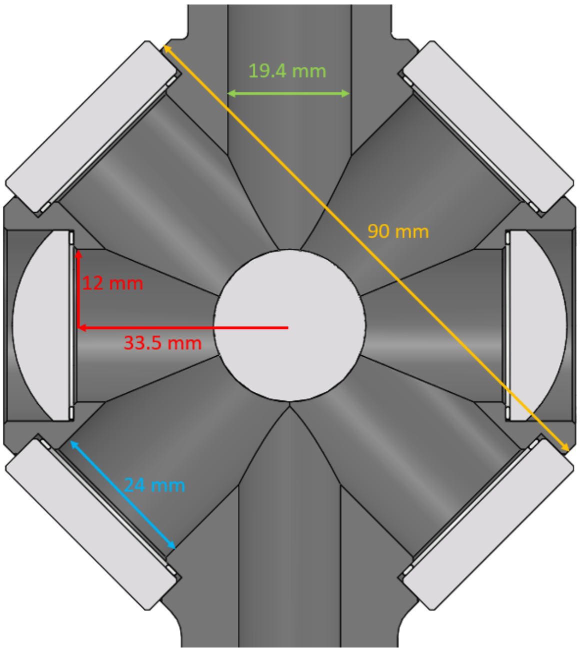

Figure 2.12: CAD drawing of the Mini-Fountain's MOT/detection chamber, with glass viewports and detection lenses indium-pressed onto the chamber.

to detection of 600 (as per the dashed line in figure 2.8). This is sufficient to expect the short-term stability to be limited by local oscillator phase noise.

### 2.2.1 Rubidium to Improve Clock State Population

One of the reasons for a low SNR in section 2.1.3 is that the upper hyperfine level of Cs has total angular momentum $F = 4$, such that there are nine $m_F$ states among which the atomic population is approximately evenly split. This factor of 9 degradation in the returning signal reduced the SNR results in figure 2.8 from 700 to 100. A way to improve this would be to use a different, but similar, atom with a smaller total angular momentum in the upper hyperfine level, such as $^{87}$Rb, which has $F = 2$ and thus only 5 $m_F$ states. An energy level diagram for $^{87}$Rb is shown in figure 2.13.

The radiative properties of rubidium are very similar to those of caesium. As can be seen from table 2.1, both of the laser-cooling transition wavelengths are in the near-IR, both have a natural linewidth of a few megahertz, both have saturation intensities of a couple of mW/cm$^2$, such that the Doppler and recoil limits are quite similar.

The slightly lower ground-state hyperfine transition frequency of 6.834 GHz does reduce the short-term stability by a factor of 1.34 compared to Cs, but this is outweighed by the factor of 1.8 increase due to the reduced number of $m_F$ states.

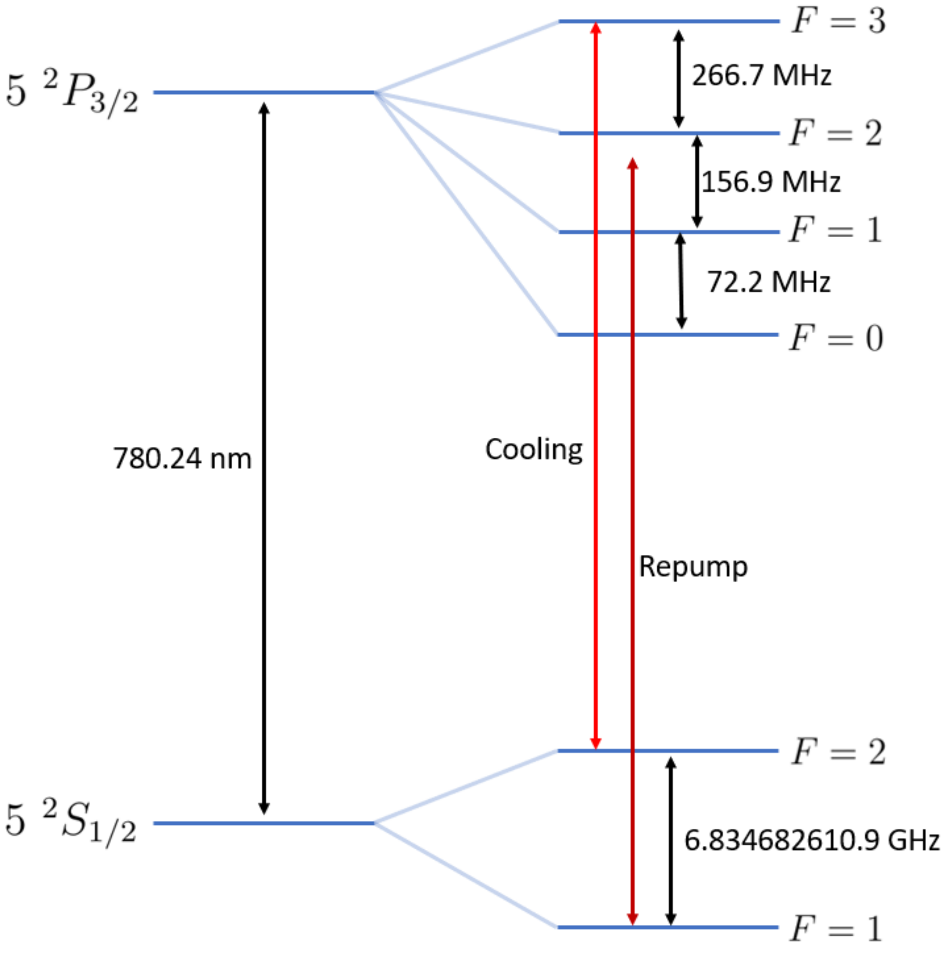

Figure 2.13: Energy level diagram for $^{87}$Rb [179], showing the fine-structure splitting (left) and hyperfine splitting (right).

Perhaps critically, because the collisional shift in $^{87}$Rb is smaller [152] and less dependent on the exact collision energy [154] than in Cs, a rubidium fountain may not need to spend a large fraction of its time operating in a low-density mode, such that, in practice, the SNR achieved for accurate clock operation can avoid being degraded, in the case of a local-oscillator-noise limited clock, by a factor of $\sqrt{10} = 3.16$, as per equation 1.83.

While these considerations potentially make rubidium the appropriate choice for a mini-fountain with a compact detection method, there is debate as to whether Cs

| **Parameter** | | $^{133}$**Cs** | $^{87}$**Rb** | **Unit** |
|---|---|---|---|---|
| $D_2$ wavelength | $\lambda$ | 852 | 780 | nm |
| Scattering rate | $\Gamma$ | 5.22 | 6.07 | MHz |
| Isotropic saturation | $I_{\mathrm{sat}}^{\mathrm{iso}}$ | 2.71 | 3.51 | mW/cm$^2$ |
| Cycling saturation | $I_{\mathrm{sat}}^{\mathrm{cyc}}$ | 1.10 | 1.67 | mW/cm$^2$ |
| Doppler limit | $T_D$ | 125 | 146 | $\mu$K |
| Recoil limit | $T_r$ | 198 | 362 | nK |
| Hyperfine splitting | $\nu_0$ | 9 192 631 770 | 6 834 682 610 .90429(9) | Hz |
| Excited state splitting | $\delta$ | 251.0 | 266.7 | MHz |
| Ground-state g-factors | $\lvert g_F \rvert$ | 1/4 | 1/2 | |

Table 2.1: Radiative properties of the Cs $D_2$ line [180] compared to $^{87}$Rb [179]. The hyperfine splitting of Cs is, by definition, exact, whereas the hyperfine splitting of Rb has been measured with a relative uncertainty of $1.3 \times 10^{-14}$ [181].

or Rb makes a better species overall; the second was defined by Cs before the advent of laser cooling, so the collisional shift considerations only arose after Cs was well-established as the primary standard. Similarly, the full-sized fountains' separated detection zones allow for a high signal-to-noise ratio to be extracted from Cs even without optical pumping to increase the $m_F = 0$ population.

### 2.2.2 State Selection to Maintain Fringe Contrast

State selection in an atomic fountain clock is typically implemented by a microwave cavity similar to the Ramsey cavity. This is located just below the Ramsey cavity, such that once atoms in the upper clock state have been transferred to the lower clock state, a laser pulse resonant with the cooling transition (usually from the down-going MOT beam in the (1,0,0) geometry) removes the remaining upper-level atoms with $m_F \neq 0$ from the cloud by radiation pressure. Minimising the time between state selection and interrogation reduces the opportunity for microwave leakage or Majorana transitions to affect the state of the atoms before interrogation.

State selection is necessary to remove the partition noise associated with inconsistent population of each $m_F$ state from shot-to-shot, and to maintain the full fringe contrast. In the mini-fountain, the aim is to keep the Ramsey cavity as close to the MOT chamber as possible, thus performing state selection in the MOT chamber is desirable. State selection requires a detuning between the clock transition and the $m_F \neq 0$ transitions, which is easily achieved with a non-zero background field $B_z(z)$.

For small magnetic fields, the linear Zeeman splitting detunes transitions by

$$\Delta\nu_{m_F m'_F} = \frac{g_F \mu_B}{h}(m'_F + m_F)B_z, \tag{2.11}$$

from the clock transition. It was found that, without deliberately adding any additional current-carrying coils, there existed a background magnetic field in the MOT chamber of CsF3 of about 1.7 $\mu$T, as shown by the transition probability vs frequency shown in figure 2.14, seen by applying a 1.0 ms microwave pulse with a microwave horn after the sub-Doppler cooling ramps. This field has similar vertical and horizontal components, as the $\pi$ and $\sigma$ transitions from vertically-oriented microwave radiation are excited with similar strengths.

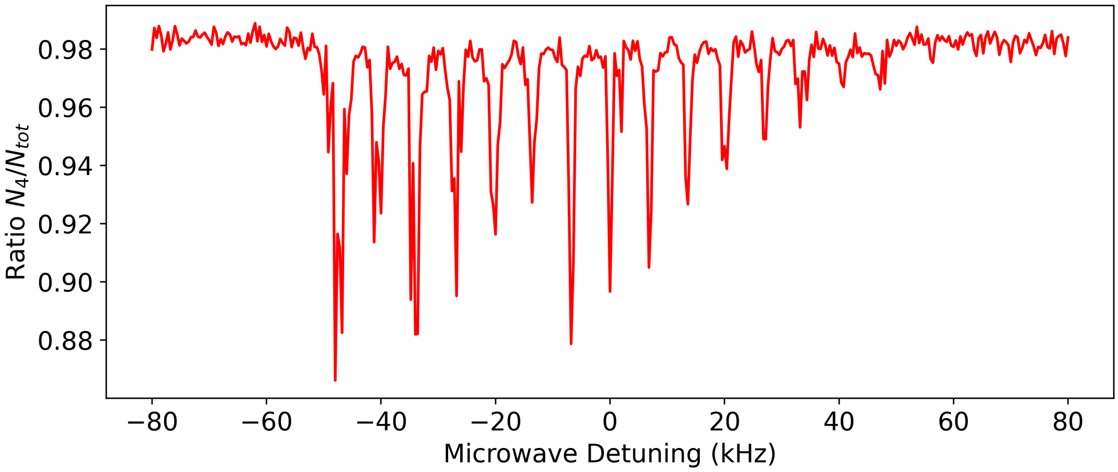

Figure 2.14: Frequency scan in the MOT chamber of CsF3 using a microwave horn. The small noise in the data away from resonance implies a high signal-to-noise ratio there, however the larger noise on the peaks, particularly on the furthest-detuned transitions, implies non-negligible microwave noise.

Clearly resolved in figure 2.14 are the 7 $\pi$-transitions and 8 $\sigma$-transitions allowed between the $|F, m_F\rangle$ states in Cs, albeit not with equal probability; the asymmetry is a result of the cooling process in a MOT utilising optical pumping between different $m_F$ states. The linewidth of the clock transition is Fourier-limited to 1 kHz by the 1 ms pule time, and the transitions with larger $m_F$ are broader due to increased sensitivity to field inhomogeneity. Thus, a 1.0 ms $\pi$-pulse resonant with $|4, 0\rangle \rightarrow |3, 0\rangle$ populates the lower clock state.

The next step is to use the MOT beams to push the remaining $|4, m_F \neq 0\rangle$

atoms out of the cloud. This requires that the laser frequency be brought from 71 MHz below resonance (which is where it finishes after a large detuning ramp for sub-Doppler cooling) to resonance within the time of the microwave pulse. The total time available for the microwave pulse and pusher pulse is limited by the intended angle of the mini-fountain's diagonal MOT beams and the launch velocity (no such restriction exists for a vertical beam). For a 600 ms flight time, the mini-fountain will launch atoms at 2.9 $\mathrm{ms}^{-1}$. For 24 mm diameter MOT beams at an angle of 45°, with the cloud starting out in the centre of the trap region, the maximum vertical space above the cloud that is still in the MOT beams is 16.9 mm, which the cloud will traverse in 5.8 ms. This suggests that a 1 ms microwave pulse will allow plenty of time for the subsequent pusher pulse, so long as the laser stays locked.

Laser unlocking is potentially an issue with both the demonstrator system and CsF3 because the frequency ramps for sub-Doppler cooling are implemented by adjusting the frequency to an AOM before the reference cell, such that the laser frequency itself is changing, and the shifted light is kept on resonance. Both master lasers are ECDLs, so the frequency is adjusted by moving a diffraction grating on a piezoelectric translation stage. It was found that, by allowing 0.8 ms in the dark before and after the microwave pulse, the ECDL in CsF3 consistently stayed locked. Allowing no extra time consistently caused the laser to unlock. It was found that a 1.0 ms resonant pulse at maximum intensity from the down-going MOT beam was sufficient to remove seemingly all of the $|4, m_F \neq 0\rangle$ atoms from the cloud. Thus, the entire state-selection process takes a total of 3.6 ms, well below the 5.8 ms limit. This process is illustrated in figure 2.15.

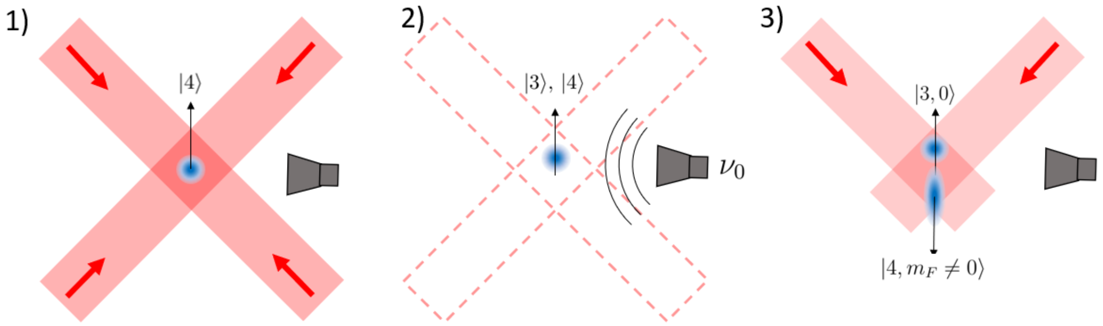

Figure 2.15: Mini-fountain state-selection process. 1) Atoms are launched and sub-Doppler cooled. 2) Beams are turned off and a microwave pulse populates the lower clock state. 3) Down-going beams remove $m_F \neq 0$ atoms from the cloud by radiation pressure.

## 2.3 Ramsey Cavity

The Ramsey cavity is perhaps the most important part of the physics package, so it requires careful design consideration. Its size is mostly set by the dimensions required to achieve the correct resonant frequency, so the focus isn't so much on miniaturisation as it is on ensuring reliable operation in situations with sub-optimal temperature control without requiring bulky apparatus to set the cavity temperature or balance the phase and amplitude of each feed.

### 2.3.1 TE$_{011}$ Cavity Geometry

Since the early days of atomic fountain clocks, a cylindrical cavity using the TE$_{011}$ mode has been employed for Ramsey interrogation. This mode is shaped according to

$$H_\rho = H_0 \frac{\beta}{k_c} J_0'(k_c \rho) \sin(\beta z), \tag{2.12}$$

$$H_\phi = 0, \tag{2.13}$$

$$H_z = H_0 J_0(k_c \rho) \cos(\beta z), \tag{2.14}$$

for a cavity length $d$ and radius $a$, where $\beta = \pi/d$, $k_c = p_{01}'/a$, $J_0$ is the Bessel function of the first kind, and $J_0'$ is its derivative . This mode has the advantage of maximising the uniformity of the longitudinal component and minimising the radial field component seen by the atoms. The resonant frequency can be calculated according to [121]

$$f_{011} = \frac{c}{2\pi n}\sqrt{\left(\frac{p_{01}'}{a}\right)^2 + \left(\frac{\pi}{d}\right)^2}, \tag{2.15}$$

where $n$ is the refractive index of the medium within the cavity. It is clear that, for a particular $f_{011}$, the necessary dimensions are not unique; a short, wide cavity can result in the same resonant frequency as a slim, long cavity. Neither extreme is suitable for a mini-fountain, so some intermediate case with $d \sim a$ will be appropriate. The aspect ratio also determines the unloaded quality-factor, $Q_0$ of the cavity, which relates the energy stored to the rate of energy dissipation; a cavity with a quality-factor of $Q$ can oscillate $Q$ times before the energy stored has dissipated by a factor of $e$. For a TE$_{011}$ mode, the unloaded $Q$ is [121]

$$Q_0 = \frac{(ka)^3 \eta a d}{4p'^2_{01} R_s} \frac{1}{\frac{ad}{2} + \left(\frac{\beta a^2}{p'_{01}}\right)^2}. \tag{2.16}$$

where $R_s = \sqrt{\omega \mu_0 / 2\sigma}$ is the frequency-dependent surface resistivity and $\eta = 377\,\Omega$ is the impedance of free-space. Since copper has $\sigma = 5.81 \times 10^7\,\mathrm{Sm}^{-1}$, then at 6.834 GHz, $R_s = 0.0215\,\Omega$. The unloaded Q-factor and cavity height are plotted vs cavity radius in figure 2.16.

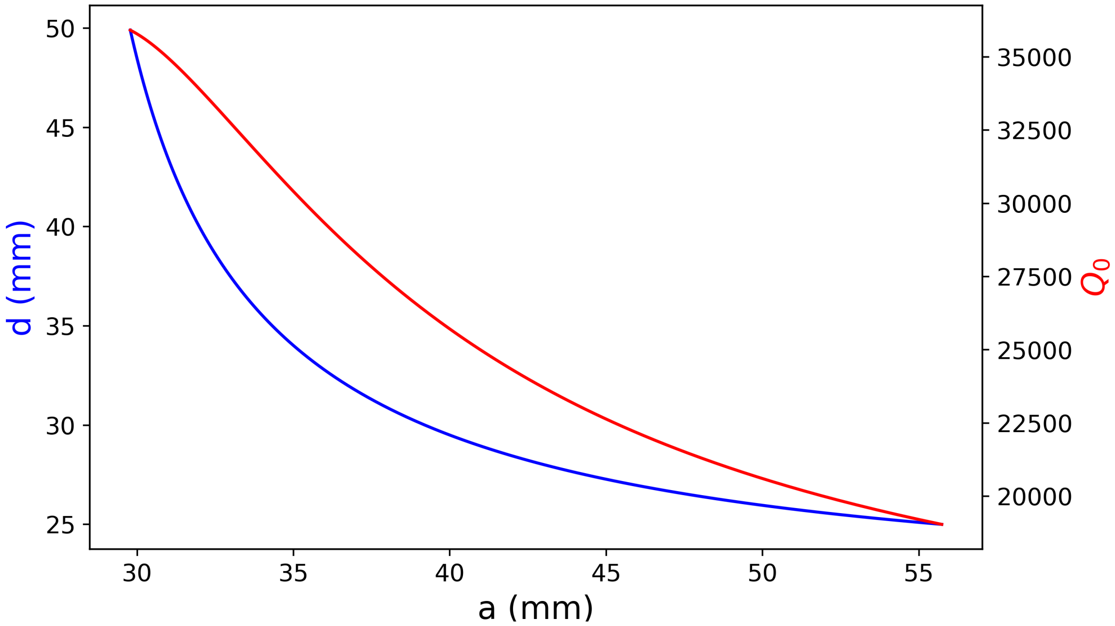

Figure 2.16: Cavity height $d$ and unloaded Q-factor $Q_0$ vs cavity radius $a$, for a fixed resonant frequency of 6.834 GHz, corresponding to the hyperfine splitting in $^{87}$Rb.

There is no specific optimum aspect ratio: while it is thought that a wider cavity reduces the phase gradients across the atom-passage holes, this also means it is shorter, thus the edge-effects associated with the atom-passage holes play a larger role in the phase structure. A range of aspect ratios, $d/a$, between 0.73 [124] and 1.18 [158] have been employed by different fountains. One may be tempted to assume that higher $Q_0$ is also better for reducing DCP, but the *loaded* Q-factor, which is determined by the coupling of the cavity to the feeds, is the relevant parameter. Additionally, a higher Q requires tighter temperature control to keep the cavity resonant with the atomic transition.

### 2.3.2 Feeding Waveguide

Another key aspect of the cavity is the way it is fed; placing probes or loops directly into the cavity perturbs the phase of the microwave field, so typically an external waveguide is fed and the mode of this waveguide is coupled into the Ramsey cavity via a small aperture.

Most fountains do this with two (or even four) opposing apertures covered by rectangular $TE_{101}$ cavities that are each fed by a straight antenna. This requires independent phase and amplitude balance of each feed to minimise any travelling wave component in the cavity. There is another approach developed for PTB's CSF1 and CSF2 [158]; where two opposing feeds are fed by the same bent waveguide operating on the $TE_{106}$ mode, such that the standing wave inside the waveguide ensures that each feed is fed with the same phase and amplitude. The downside is that it becomes harder to evaluate the DCP effect, as the feeds can't be independently disabled. These two set-ups are shown in figure 2.17.

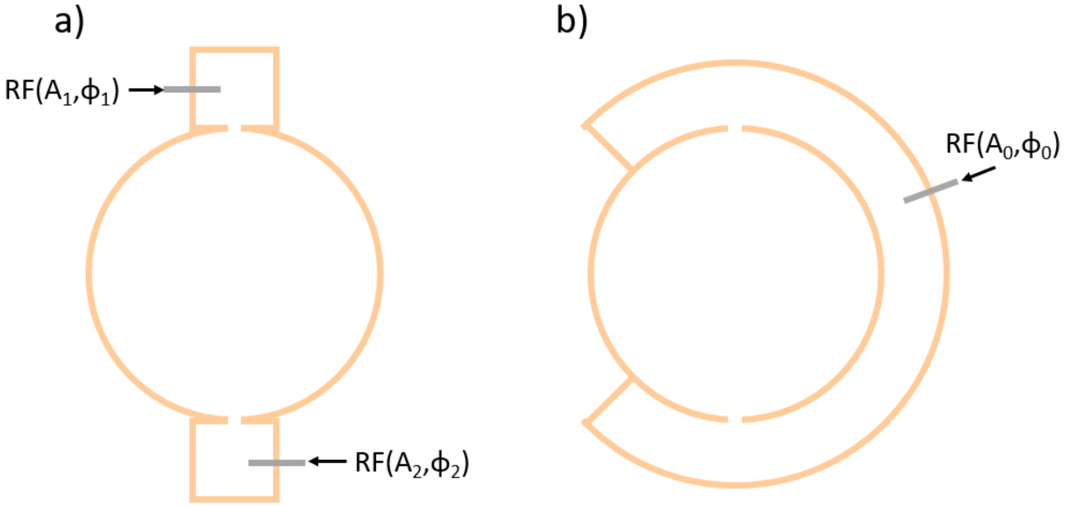

Figure 2.17: a) Two independent waveguides fed by antennas. b) Two coupled apertures from the same waveguide with one antenna feed.

Rectangular waveguide modes can be denoted as $TE_{qrs}$ modes. The magnetic field components of the rectangular $TE_{10s}$ modes for a cavity of dimensions $X$, $Y$, and $Z$ in the $x$, $y$, and $z$ directions, respectively, can be expressed as

$$H_x = \frac{-iE_0}{Z_{\mathrm{TE}}} \sin\left(\frac{\pi x}{X}\right) \cos\left(\frac{s\pi z}{Z}\right), \tag{2.17}$$

$$H_y = 0, \tag{2.18}$$

$$H_z = \frac{i\pi E_0}{k_{10s}\eta X} \cos\left(\frac{\pi x}{X}\right) \sin\left(\frac{s\pi z}{Z}\right), \tag{2.19}$$

where $k_{10s} = \sqrt{(\pi/X)^2 + (s\pi/Z)^2}$ and $Z_{\mathrm{TE}}$ is the waveguide impedance [121]. The resonant frequency of this mode is

$$f_{10s} = \frac{c}{2\pi n}\sqrt{\left(\frac{\pi}{X}\right)^2 + \left(\frac{s\pi}{Z}\right)^2}. \tag{2.20}$$

In both cases, the wavelength of the resonant mode inside the waveguides is

$$\lambda_g = \frac{\lambda_0}{\sqrt{1-(\lambda_0/2X)^2}}, \tag{2.21}$$

however, some modifications arise for a bent waveguide [182]. For the geometry of figure 2.18, with a radius of curvature $R$ and width $Y$, the wavelength becomes

$$\lambda_g' = \lambda_g\left(1 - \frac{1}{12}\left(\frac{Y}{R}\right)^2\left(-\frac{1}{2} + \frac{1}{5}\left(\frac{2\pi Y}{\lambda_g}\right)^2 + \dots\right)\right), \tag{2.22}$$

which, for the example of PTB's fountain cavity, where $R = 24.2\,\mathrm{mm}$ and $Y = 20\,\mathrm{mm}$, corresponds to an increase in wavelength of 0.1%.

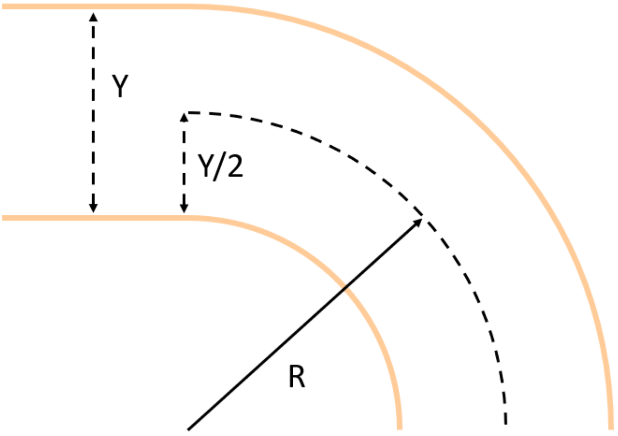

Figure 2.18: Waveguide bent in the E-plane, with central radius $R$ and width $Y$.

In either case, the feeding waveguides are coupled to the cavity by apertures. A small circular aperture of radius $r \ll \lambda_g$ in a wall parallel to the magnetic field of the feed waveguide can be treated as a dipole [183] with magnetic and electric polarisability

$$\alpha_m = \frac{4}{3} r^3, \tag{2.23}$$

$$\alpha_e = \frac{2}{3} r^3, \tag{2.24}$$

such that the waveguide magnetic field $\mathbf{H}_{\mathrm{wg}}$ and electric field $\mathbf{E}_{\mathrm{wg}}$ produce dipole strengths

$$\mathbf{P}_m = \mp \alpha_m (\mathbf{H}_{\mathrm{wg}} + \mathbf{H}_{\mathrm{wg,r}} - \mathbf{H}_{\mathrm{cav}} - \mathbf{H}_{\mathrm{cav,r}}), \tag{2.25}$$

$$\mathbf{P}_e = \mp \epsilon_0 \alpha_e (\mathbf{E}_{\mathrm{wg}} + \mathbf{E}_{\mathrm{wg,r}} - \mathbf{E}_{\mathrm{cav}} - \mathbf{E}_{\mathrm{cav,r}}), \tag{2.26}$$

where the subscript r denotes the reaction fields radiated by the dipole back into the waveguide and cavity, which are necessary to satisfy conservation of energy [184], and the cavity fields $\mathbf{H}_{\mathrm{cav}}$ and $\mathbf{E}_{\mathrm{cav}}$ are derived from the magnetic and electric current densities,

$$\mathbf{M} = i\omega\mu_0 \mathbf{P}_m, \tag{2.27}$$

$$\mathbf{J} = i\omega \mathbf{P}_e, \tag{2.28}$$

such that

$$\mathbf{H}_{\mathrm{cav}} = \sum_n a_n \mathbf{H}_n, \tag{2.29}$$

$$\mathbf{E}_{\mathrm{cav}} = \sum_n a_n \mathbf{E}_n, \tag{2.30}$$

where each mode $\mathbf{H}_n$, $\mathbf{E}_n$ is excited with an amplitude

$$a_n = \iint_S \mathbf{H}_n \cdot \mathbf{M} \, dS, \tag{2.31}$$

where the double-integral is over the surface of the aperture [184]. This situation is illustrated diagrammatically in figure 2.19.

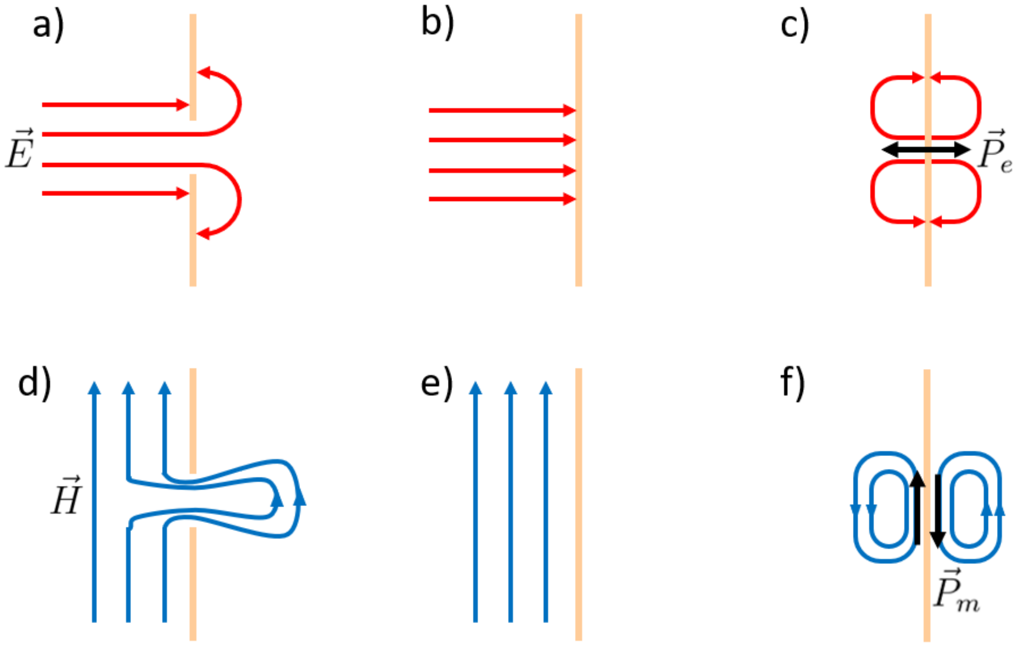

Figure 2.19: a) Electric field lines near an aperture in a conducting wall can be thought of as a combination of (b) the field without the aperture and (c) the field due to an electric dipole moment $\mathbf{P}_e$ at the aperture. d) Magnetic field near an aperture is a combination of (e) the magnetic field without the aperture and (f) the field due to two oppositely-directed magnetic dipole moments $\mathbf{P}_m$.

Because the direction of the magnetic field line radiated through the aperture is aligned with the incident field, the appropriate waveguide geometry can efficiently excite the cylindrical $\mathrm{TE}_{011}$ mode, and largely avoid the excitation of other unwanted modes. However, choke grooves in the end caps of the cavity are employed to detune the otherwise-degenerate $\mathrm{TM}_{111}$ mode.

The coupling between the feeding waveguide and the Ramsey cavity also determines the loaded $Q$ of the Ramsey cavity, which affects the cavity pulling and the distributed cavity phase shifts. As noted by figure 1.20, a lower cavity $Q$ results in a flatter cavity pulling shift across a wider range of detuning (and hence temperature). However, a lower $Q$ can also worsen the DCP shift. While section 1.2.3 treated the DCP shift in a generalised way, some specific results are notable.

For separate waveguides, the feed amplitudes have to be independently balanced. If the total losses from the cavity to the waveguides are represented by $Q_{\mathrm{ex}}$, then the loss to each feed is represented by $Q_F = 2(1 \pm \epsilon/2)Q_{\mathrm{ex}}$, where $\epsilon$ is the imbalance.

The imaginary part of the perturbing secondary field is then [128]

$$\mathrm{Im}(h_{z,m}) \propto \frac{\epsilon}{2}\left(\frac{Q_0}{Q} - 1\right). \tag{2.32}$$

However, for a single waveguide coupling both feeds, the power in through each aperture scales in the same way as the losses from each aperture, such that

$$\mathrm{Im}(h_{z,m}) \propto \frac{\epsilon}{2}, \tag{2.33}$$

which is independent of $Q$. From figure 2.16, the unloaded $Q_0$ of a 6.834 GHz $\mathrm{TE}_{011}$ cavity for some reasonable aspect ratio may exceed 30,000. One can also calculate that keeping the peak cavity pulling shift below around $10^{-16}$ requires a loaded $Q \lesssim 5,000$, such that $(Q_0/Q - 1) > 1$, and thus one derives a benefit to using a single waveguide to feed both apertures because $\mathrm{Im}(h_{z,m})$ is not amplified by $(Q_0/Q - 1)$.

An additional benefit to the coupled-feeds approach is that the phase imbalance between the feeds, $\delta\Phi_{\mathrm{feeds}}$, is kept small due to the standing wave formed in the waveguide. This is useful because, if the cavity is detuned from resonance by $\delta\omega_c$, then an additional DCP shift occurs according to [128]

$$\mathrm{Im}(h_{z,m}) \propto 2\frac{\delta\omega_c}{\omega_0} Q_0 \tan\left(\frac{\delta\Phi_{\mathrm{feeds}}}{2}\right), \tag{2.34}$$

such that suppressing $\delta\Phi_{\mathrm{feeds}}$ removes this contribution to the DCP.

### 2.3.3 Vacuum-Sealing to Fountain

In order to keep the physics package compact, the cavity forms part of the vacuum vessel itself, which requires that the bottom surface of the cavity is vacuum sealed to the MOT & detection chamber, and the upper surface sealed to the flight tube. Since copper is too soft for a standard ConFlat flange, indium wire seals are used instead. Only the Ramsey cavity is under vacuum, the feeding ring waveguide can be left in air (although still requires sealing against microwave leakage).

Most fountains use Ramsey cavities made from oxygen-free high-conductivity (OFHC) copper. There are some examples of cavities being made from titanium [185] or 3D printed and plated with gold or copper to increase the conductivity

[186] (and thus Q-factor). There is also an example of a Ramsey cavity made of 6061 aluminium [187], with approximately half the conductivity of OFHC copper, which is still high enough to keep the DCP shift manageable.

NPL has lots of recent experience with machining and tuning OFHC copper cavities, so continuing with this material is an attractive option.

It is important to ensure the cavity dimensions are independent of the pressed-indium vacuum seals, because it is difficult to control exactly how far an indium wire is squashed. However, the design still needs to allow sufficient freedom to tune the frequencies of the cavity and waveguide.

One such design that achieves this is shown in figure 2.20. The Ramsey cavity walls and upper endcap are a single piece, to which the rest of the fountain is vacuum-sealed. The lower endcap is bolted into the bottom of the cavity body, and is the only part of the design to be completely in-vacuum. Fine tuning of the Ramsey cavity frequency is achieved by machining small amounts off of the lower endcap: machine down the upper-inner surface to lengthen the cavity and reduce the cavity frequency, and machine down the surface contacting the cavity body to shorten the cavity and increase the frequency.

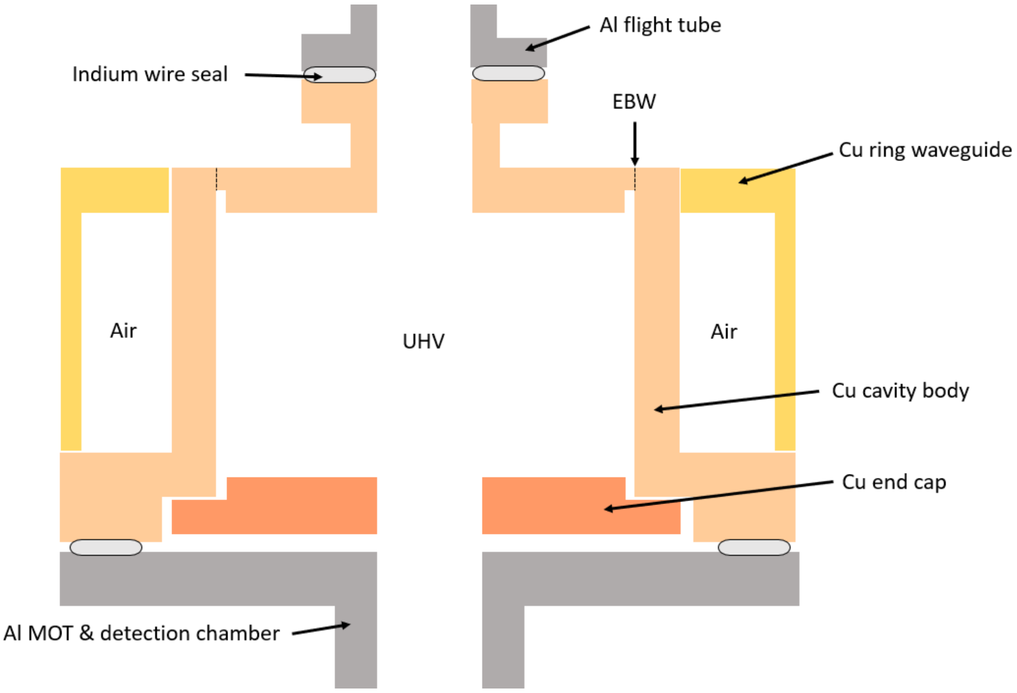

Figure 2.20: Cavity design comprised of three different copper parts, indium-sealed to the MOT & detection chamber below, and flight tube above.

One complication is that the upper choke grooves are too small to be machined into the upper endcap. The solution is to machine the body and upper endcap as two separate pieces, and then weld them together, leaving the small choke groove between them. Electron-beam welding (EBW) is chosen for this purpose, because the very localised heating minimises distortion to the rest of the cavity, and the large, controllable weld-depth ensures that no air pockets remain between the pieces after welding. An initial test piece was sent out to be welded, shown in figure 2.21, and the result is a clean weld penetrating all the way down the join. This was tested externally to be vacuum-tight, with a leak rate $< 10^{-10}\,\mathrm{mbar\,l\,s^{-1}}$.

Figure 2.21: Initial test of EBW on a similar geometry to the cavity body and endcap. The cut-out cross-section shows excellent weld penetration.

### 2.3.4 Vacuum-Sealing the Coupling Apertures

Since the Ramsey cavity is under vacuum and the ring waveguide is not, a dielectric viewport must be vacuum-sealed over the apertures. The challenge here is that the copper cavity has a thermal expansion coefficient $\alpha = 17 \times 10^{-6}\,\mathrm{K^{-1}}$, which is a few times higher than most types of ceramic/glass. This potentially causes issues during the high-temperature bake out required to achieve ultra-high vacuum inside the physics package: too much relative motion between the parts could cause a leak.

One way this can be done is by metallising one surface of a ceramic disk, such that it can then be soldered to another metal with an intermediate thermal expansion coefficient, to reduce stress on the parts during bake-out [185]. This was considered for the mini-fountain, however it was difficult to find companies who specialise in the metallisation of ceramics willing to accept such a small quantity of work.

Another approach may be to use a pressed-indium seal between the copper cavity and a glass viewport [188]. This requires no additional specialist processes, but would require modifications to the inner surface of the ring waveguide to accommodate screw threads for a clamp to press the indium wire with. This is not ideal, because alterations to the waveguide surface near the coupling aperture may affect the field being coupled into the cavity.

The optimal solution would neither require alterations to the waveguide, nor any specialised process that can't be performed locally. Optical glass windows are readily-available in appropriate sizes. Some relevant physical properties of various optical glasses are shown in table 2.2. One can see that the material with the closest thermal expansion coefficient to copper ($17 \times 10^{-6}\,\mathrm{K}^{-1}$) is $CaF_2$, and the materials with the lowest relative permittivity are UV fused silica and N-BK7.

| Material | $\alpha$ ($\times 10^{-6}$ K$^{-1}$) | $\epsilon_r$ |
|---|---|---|
| UV Fused Silica | 0.5 | 3.8 |
| Sapphire | 5.5 | 10.8 |
| N-BK7 | 0.7 | 3.8 |
| $CaF_2$ | 18.8 | 6.8 |

Table 2.2: Thermal expansion coefficient $\alpha$ and relative permittivity $\epsilon_r$ (sometimes called the dielectric constant) for various optical glass materials [189, 190].

The first sealing method investigated was indium solder: glass viewports with 1/2" diameter and 3 mm thickness were sealed to some copper flanges with a small aperture by placing a 0.5 mm indium wire between the copper and glass and using an electrical hotplate to heat the assembly until the indium melted. Some aluminium lens posts were placed on top of the glass during heating to ensure that the surface tension was overcome and the indium spread out between the glass and copper. Some example results of this are shown in figure 2.22. Interestingly, small bubbles appear in the indium, and only a small amount adheres to the surfaces, such that when a thicker 1 mm indium wire was used, as in figure 2.22d, the resulting seal is about the same, and the excess indium instead escapes out the side.

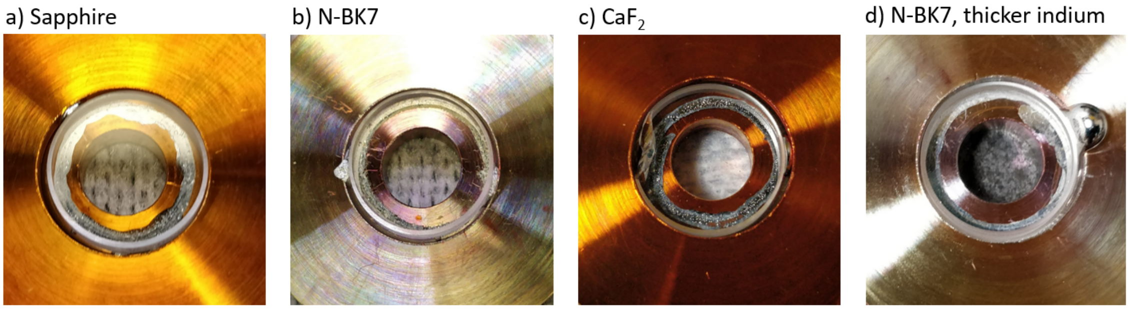

Figure 2.22: Viewports indium-soldered to copper flanges. The orange colour results from being heated to 180°C, and the gold-green from 230°C, as it was initially assumed that the weight of the glass alone would overcome the surface tension.

To test the vacuum-tightness of this sealing method, the copper flanges were sealed to a steel flange by a pressed-indium seal, which was attached to a turbo-pump (Agilent Technologies TPS-Compact) and residual gas analyser (Stanford Research Systems RGA-100) by a DN40 (CF1.33) knife-edge seal. This setup is shown in figure 2.23. The pressure inside the system is indicated by an ion gauge (Agilent Technologies FRG-700).

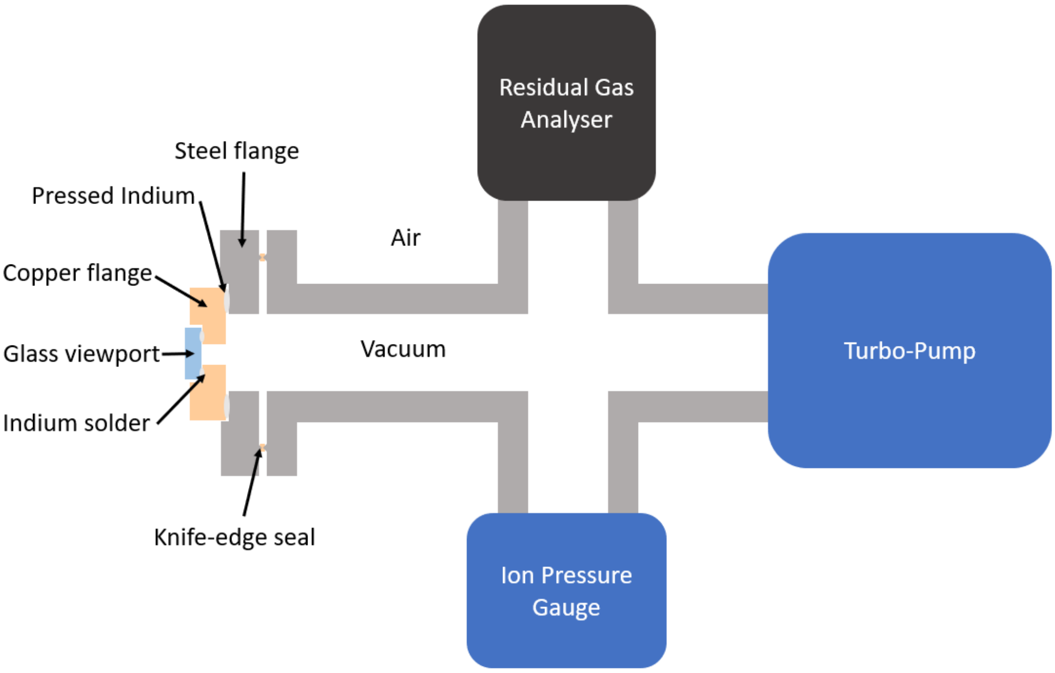

Figure 2.23: Set-up for testing the vacuum-tightness of glass-copper seals.

This sealing method was unsuccessful: the lowest pressure achieved was $1.7 \times 10^{-4}$ mbar, by both $CaF_2$ and sapphire, which is too high to even turn on the residual gas analyser to carry out a helium leak-check. Instead, a volatile solvent such as isopropyl alcohol was squirted onto the seal, and the presence of a leak was indicated by a subsequent rise in the displayed pressure.

The same method was attempted with another solder material called C-Solder, which has a much higher melting point, such that each copper flange was heated to 300°C. The results were similar to those of indium, with the best pressure being $1.8 \times 10^{-3}$ mbar. Solvent leak-checking confirmed that it was the glass-copper seals that were leaking.

Another method was to use a UHV compatible epoxy (Loctite Ablestik 2116, also known as KL-320K), which is marketed as a sealant and adhesive. There are examples of this being used as an adhesive inside a vacuum [191, 192], but no examples could be found of this epoxy being used to form a vacuum seal. None the less, a liberal amount of epoxy was applied to the copper flanges, and one of each type of glass viewport was attached. Each was allowed to cure at room temperature for 5 days under the weight of an aluminium lens post. This is shown in figure 2.24.

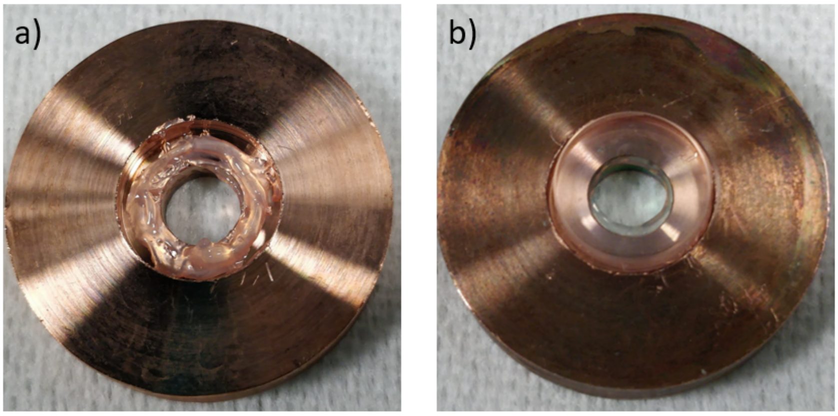

Figure 2.24: a) Epoxy applied to copper flange. b) N-BK7 glass viewport sealed to copper by epoxy.

When these flanges were installed, the results were much better: the pressures quickly dropped to below $10^{-5}$ mbar, and the RGA was activated. The $CaF_2$ flange got down to $2.5 \times 10^{-6}$ mbar, at which point a helium leak-check found that the epoxy seal was leaking. The rest continued to pump down as shown in figure 2.25, where pressures of around $10^{-8}$ mbar were reached after several days.

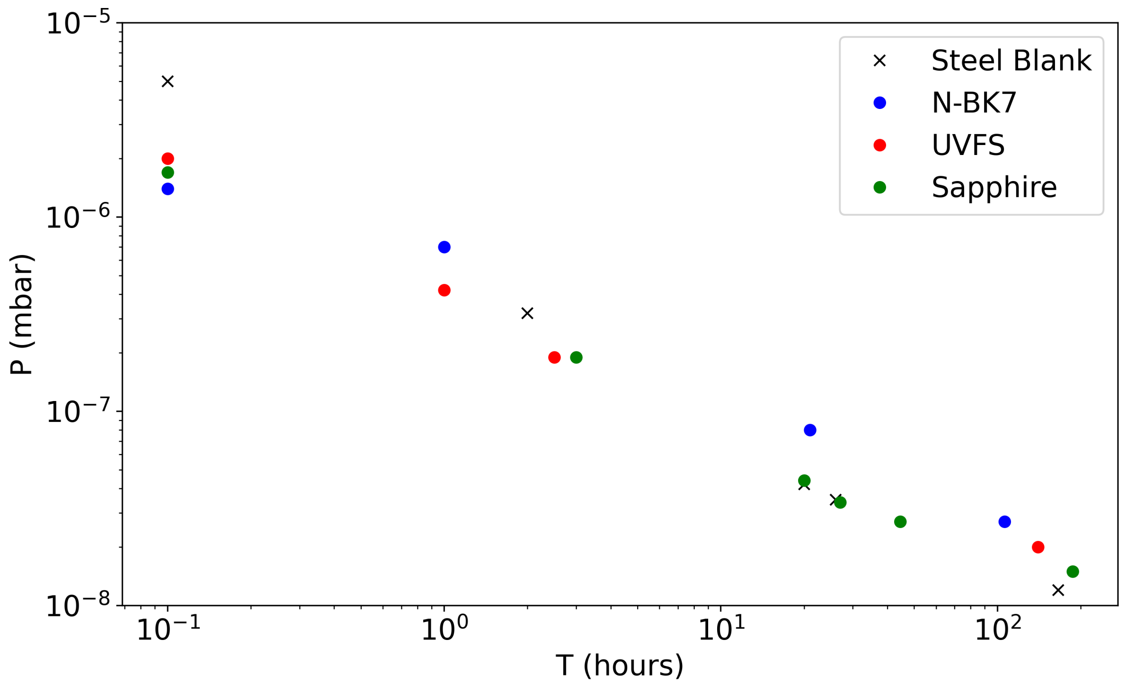

Figure 2.25: Pressure vs time for various glass viewports epoxied to copper flanges. The same data for a steel blank, in place of the glass-copper-steel flange, is included for comparison.

The next step is to check how the epoxy seals stand up to bake-out. The use of indium seals would restrict the mini-fountain bake-out to $\lesssim 140°\text{C}$, however the epoxy manufacturer specifies a 130°C maximum operating temperature, so this further reduces the upper limit on the bake-out temperature. This was tested by attaching a thermocouple to the copper surface, and wrapping some aluminium foil over the flange to increase the thermal homogeneity within, followed by an electrical heater tape around the flange, followed by another layer of aluminium foil. A layer of thermal insulation was then wrapped over the assembly, followed by a layer of reflective mylar foil. The heater tape was controlled by a variable transformer outputting between 0 and 240 V, set by a manual dial. The maximum temperature the heater tape is rated for is 200°C. However, because of the finite time taken for heat to travel from the core of the tape to the thermocouple, the temperature in the tape could be much higher than the temperature indicated by the thermocouple. To avoid problems, the temperature had to be increased gently from room temperature to around 100°C over a couple of hours.

The first viewport to be tested was sapphire. Initially, it was heated to 100°C and held there for around 30 minutes, after which it was allowed to cool. No leak was discerned with the RGA by helium leak-checking. The process was repeated for a maximum temperature of 125°C, after which a small leak was detected with the RGA. The process was then repeated again for a maximum temperature of 140°C, which exceeds the epoxy's maximum operating temperature. No increase in vacuum pressure was seen while holding at 140°C.

The same method was applied to the UV fused silica viewport and produced the same results. When attempting to test the N-BK7 flange, it was found that the indium seal between the copper and steel, which by this point had undergone several cycles of sealing and removal, was leaking. Thus, getting more information from this epoxy seal was not possible.

The final aspect to test was how to remove the viewports, in the event that something goes wrong in the sealing or bake-out stage and it needs to be re-done. It was found that the epoxy should be heated to above 145°C, and brute-force applied through the aperture to separate the glass from the copper. Even so, the copper surface may have some epoxy remaining on it, such that the cavity body would need to go back to the workshop to clean up the surface before a second attempt at epoxy sealing could be made. This is deemed an acceptable cost, as the cavity goes back and forth for alterations during the tuning process (section 3.1) anyway.

Thus, when the cavity is constructed in section 3.1, the coupling apertures are covered with sapphire viewports and sealed by epoxy, which introduces a maximum bake-out temperature of around 100°C. Sapphire was chosen over UV fused silica or N-BK7 because it has the highest thermal expansion coefficient, and is thus the least-dissimilar to copper. It is, however, the material with the highest relative permittivity, and thus the highest loss to microwaves. The effect of the viewport on the coupling between the ring and cavity will be investigated in section 3.1.

### 2.3.5 Final Cavity Design

The CAD drawings of the final cavity design are shown in figures 2.26 and 2.27. These also show the choke grooves in the end caps, of length $l_c = 5\,\mathrm{mm}$ and width $w_c = 0.5\,\mathrm{mm}$, as well as the waveguides-below-cut-off, of length $l_{\mathrm{wbc}} = 30\,\mathrm{mm}$ and radius $r_{\mathrm{wbc}} = 8\,\mathrm{mm}$, that attenuate microwave leakage from the endcaps.

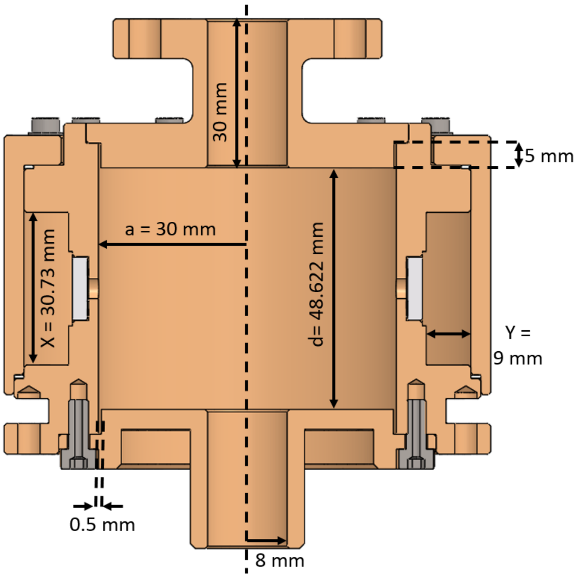

Figure 2.26: Side view of the Ramsey cavity.

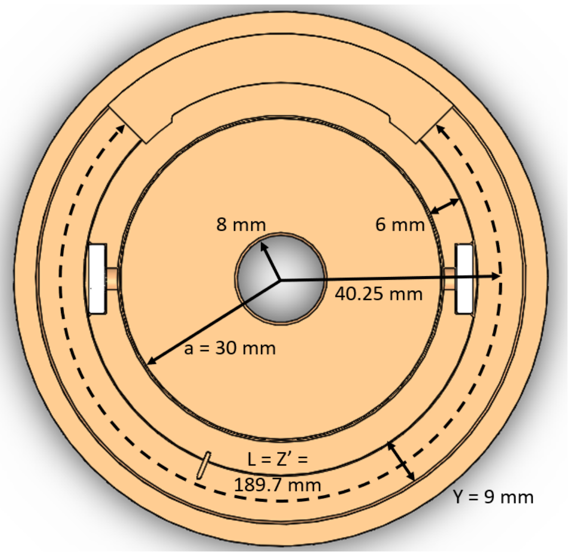

Figure 2.27: Top view of the Ramsey cavity.

The circular waveguide mode with the smallest cut-off wavelength is the TE$_{11}$ mode [121], which has $p'_{11} = 1.841$. The cut-off wavelength for a circular waveguide of radius $r_{\mathrm{wbc}}$ is $\lambda_c = 2\pi r_{\mathrm{wbc}}/p'_{11}$. The attenuation of the evanescent mode is [182]

$$\alpha_{\mathrm{wbc}} = -\frac{54.57}{\lambda_c}\sqrt{1-\left(\frac{\lambda_c}{\lambda}\right)^2}\ \mathrm{dB/m}, \tag{2.35}$$

such that a narrower waveguide produces greater attenuation, and thus less microwave leakage. However, the atom-passage holes must be wide enough to allow sufficient atoms to return after interrogation, so a compromise of $r_{\mathrm{wbc}} = 8\,\mathrm{mm}$ is reached, which has $\alpha_{\mathrm{wbc}} = -1,565\,\mathrm{dB/m}$. For a length of $l_c = 30\,\mathrm{mm}$, the attenuation is $-47\,\mathrm{dB}$, which, from experience with NPL CsF2 and CsF3, is thought to be sufficient to avoid microwave leakage during interrogation at the $10^{-16}$ level or better, as microwaves leaking from the cavity should be in-phase with the cavity field, and the atom's ballistic flight naturally approximates a leakage Rabi frequency vs time that is symmetric about $T/2$. The $\mathrm{TE}_{11}$ mode is only excited in the waveguides if the atom-passage holes are slightly off-centre [193], with the directly-excited $\mathrm{TE}_{01}$ mode having 2.5 times higher $\alpha_{\mathrm{wbc}}$. The flight tube will also be a waveguide-below-cut-off, to further attenuate microwave leakage.

The choke grooves act like a coaxial waveguide and modify the resonance condition for the $\mathrm{TM}_{111}$ cavity mode [193], and simulations with our cavity geometry suggest a downwards detuning of 370 MHz [194].

Another noteworthy feature of the cavity design is the method of tuning the ring-waveguide frequency: a 360° waveguide is machined into the cavity body, and a separate copper block is inserted to fill the required 90°. The frequency can then be decreased by machining down the end surfaces of this block, which is much easier and less risky then machining the entire cavity body. This assumes that the frequency is initially too high, which is made so by adding some additional leeway to the initial tuning block length. The sapphire windows and epoxy seal are dielectrics that will increase the cavity volume and thus lower the ring frequency, by about 6.8 MHz according to simulations [194].

The ring waveguide itself can be fed from a coaxial cable that terminates either as a straight probe if fed from the side, or as a loop if fed from above, either directly in the ring cavity or in the end of an E-plane T waveguide. Space constraints, specifically the fitting of the C-field coil around the cavity, require the ring cavity to be fed directly and from above, therefore a vertical loop is employed inside the ring cavity. The coupling can be increased (and thus broaden the ring waveguide resonance) by increasing the effective area of the loop. This is achieved by taking a straight probe and bending it into a recess in the cavity body, such that the complete loop is formed by the bent probe and the cavity body, which are in electrical contact.

## 2.4 Magnetic Shields

The largest systematic shift is the second-order Zeeman shift, so the background magnetic field (the "C-field") must be carefully considered.

In full-size fountain clocks, homogeneity on the order of 1% is achieved with a solenoid of length significantly longer than the interrogation region, so the edge-effects occur outside this critical zone. Additionally, three or more [69] layers of magnetic shielding typically separate the interrogation region from the outside, so any external magnetic fields do not disturb the homogeneity of the C-field. While this approach is very successful, it results in a very large physics package.

Initially, the plan was to use only two layers of shielding around the entire mini-fountain physics package; there would not be a shield specifically surrounding the interrogation region. This would have the advantage of lengthening the magnetic shields and C-field coil, such that edge-effects would occur away from the interrogation region. However, this would also require strict attention to the parts that make up the mini-fountain: anything even slightly magnetic could destroy the required field homogeneity. It was soon realised that this would entail large number of custom-made parts that swap out the commonly used A4 stainless steel for materials like aluminium or titanium, which manufacturers are not always willing to entertain in such small volumes.

Instead, the use of steel is minimised but not eliminated. A mu-metal shield is designed to encompass the interrogation region and C-field coil to attenuate the residual field produced by magnetic parts of the physics package. A second layer surrounds the entire physics package (except the ion pump) to attenuate external magnetic fields. However, transient magnetisation of the shields surrounding MOT coils has been known to produce an additional frequency shift, of the order of $10^{-15}$ [195], due to the finite time taken for the eddy currents in the shields to decay. Finite-element simulations using COMSOL are used to design the coils and shields.

### 2.4.1 C-Field Homogeneity

It has long been known that enclosing a solenoid inside a cylinder of high permeability can improve the homogeneity of the axial magnetic field inside the solenoid [196, 197, 198]. A solenoid of half-length $l$ and radius $a$, carrying a uniform current density equivalent to a current $I$ through $N$ turns and enclosed by a shield with a relative permeability $\mu_r \to \infty$ of half-length $L \geq l$ and $b \geq a$ results in an axial magnetic field

$$B_z(z) = \frac{\mu_0 N I}{L}\left(l + 2a\sum_{n=1}^{\infty}\left(K_1(nka) + \frac{K_0(nkb)I_1(nka)}{I_0(nkb)}\right)\cos(nkz)\sin(nkl)\right), \tag{2.36}$$

where $k = \pi/L$, $I_q$ and $K_q$ are the modified Bessel functions of the first and second kind, respectively. Some examples of $B_z(z)$ vs $z$ are shown in figure 2.28. One can see that, the smaller the gap between the end of the coil and the shield endcaps, the better the axial magnetic field homogeneity, to the extent that $l = L$ produces, in theory, a perfectly uniform field.

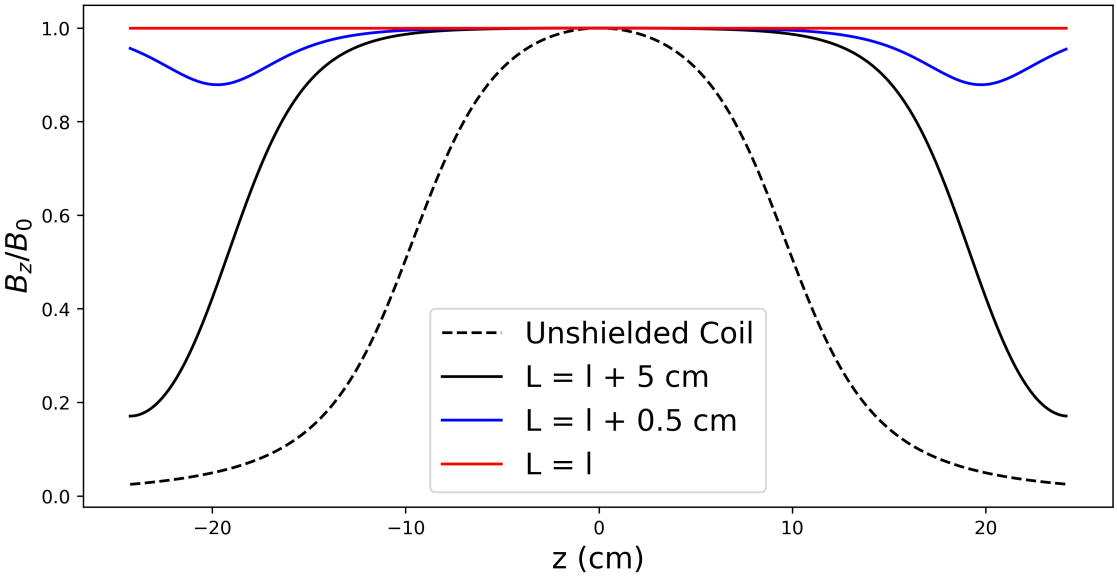

Figure 2.28: Normalised magnetic field strength vs $z$ for $a = 5.1\,\text{cm}$, $b = 5.5\,\text{cm}$, $l = 19.25\,\text{cm}$, and various $L$.

In practice, the C-field coil does not have a perfectly uniform current density. To best exploit the mirroring effect of the shield, the last of the $N$ discrete wire turns is placed a distance $\delta z = L/N$, half the usual coil separation, from the end cap [198]. To keep the weight of the coil down, each turn is separated by 5 mm. This requires 2.5 mm between the shield end caps and the final turns. The resulting axial magnetic field for a coil with $N = 77$ turns, carrying $I = 400\,\mu\text{A}$, for the same dimensions as in figure 2.28, is shown in figure 2.29.

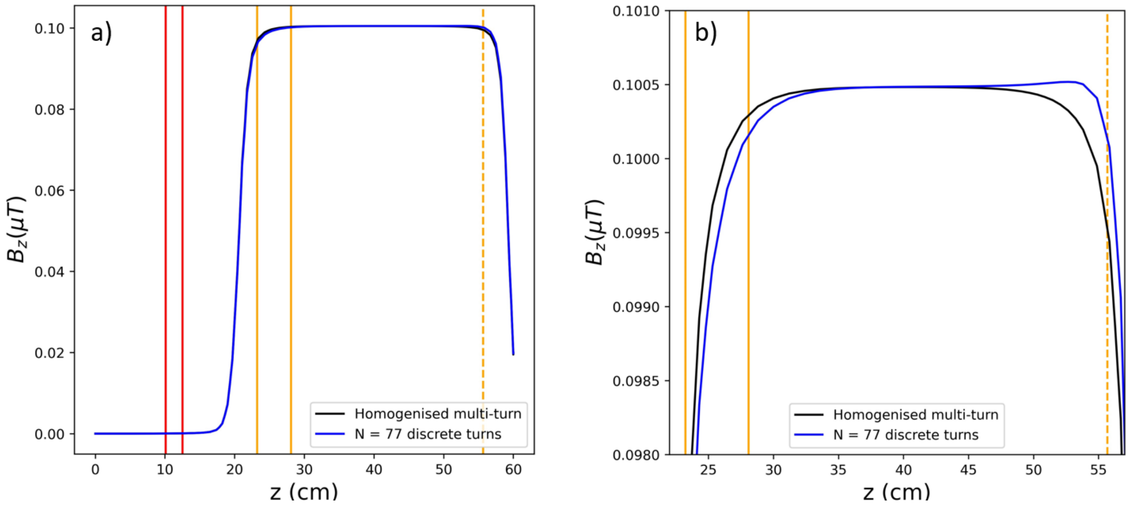

Figure 2.29: a) Simulated $B_z$ for the height of the mini-fountain. Red lines denote the MOT beam region, solid orange lines denote the Ramsey cavity, and the dashed line the apogee. b) Close-up of the interrogation region, with differences between the homogenised and discrete coils.

Inhomogeneities occur towards the ends of the interrogation region. In the discrete case, the average field strength in the Ramsey cavity is $\langle B_z \rangle_c = 99.02\,\mathrm{nT}$, and in the flight tube it is $\langle B_z \rangle_f = 100.46\,\mathrm{nT}$. By equation 1.48 with $|g_F| = 1/2$ for $^{87}$Rb, this would produce an offset between the centre of the Rabi envelope and the "central" Ramsey fringe of the first-order magnetically-sensitive transition of 2.5 Hz, which is far less than the width of the Rabi envelope, expected to be 50 Hz for a 20 ms cavity transit time. Thus, there should be no additional uncertainty associated with fringe misidentification in the evaluation of $\langle B_z \rangle_f$ for use in the second-order Zeeman shift calculation.

### 2.4.2 Shielding Factor

Because the C-field is expected to be very uniform, the majority of the magnetic field inhomogeneity may come from external fields, which are attenuated by mu-metal shielding. The transverse shielding factor of a single shield of radius $b$, relative permeability $\mu_r$, and thickness $t$ is [199]

$$S_t = 1 + \frac{\mu_r t}{2b}. \tag{2.37}$$

For a second layer of radius $b' > b$, the total shielding factor becomes

$$S_{t,2} = 1 + \frac{\mu_r t}{2b} + \frac{\mu_r t}{2b'} + \frac{\mu_r^2 t^2}{4bb'}\left(1 - \left(\frac{b}{b'}\right)^2\right), \tag{2.38}$$

which is dominated by the final term, and is independent of the lengths of the shields. In the axial direction, a single layer of half-length $L$ has a shielding factor

$$S_z = 1 + \frac{\mu_r t}{2b}\left(\frac{2K}{1 + a' + 0.85a'^3/3}\right), \tag{2.39}$$

where $a' = 2L/b$ and

$$K = 1.83\left(1 + \frac{1}{4a'^3}\right) - \frac{1}{2} + 2 \times 0.85\left(\ln\left(a' + \sqrt{1 + a'}\right) - 2\left(\sqrt{1 + \frac{1}{a'^2}} - \frac{1}{a'}\right)\right), \tag{2.40}$$

where the numerical pre-factors of 1.83 and 0.85 were determined experimentally [199]. For an additional layer with $L' > L$, the shielding factor becomes

$$S_{z,2} = 1 + S_z + S_z' + S_z S_z'\left(1 - \frac{L}{L'}\right), \tag{2.41}$$

which is also dominated by the final term. These shielding factors, for symmetrically-placed shields, are shown in figure 2.30.

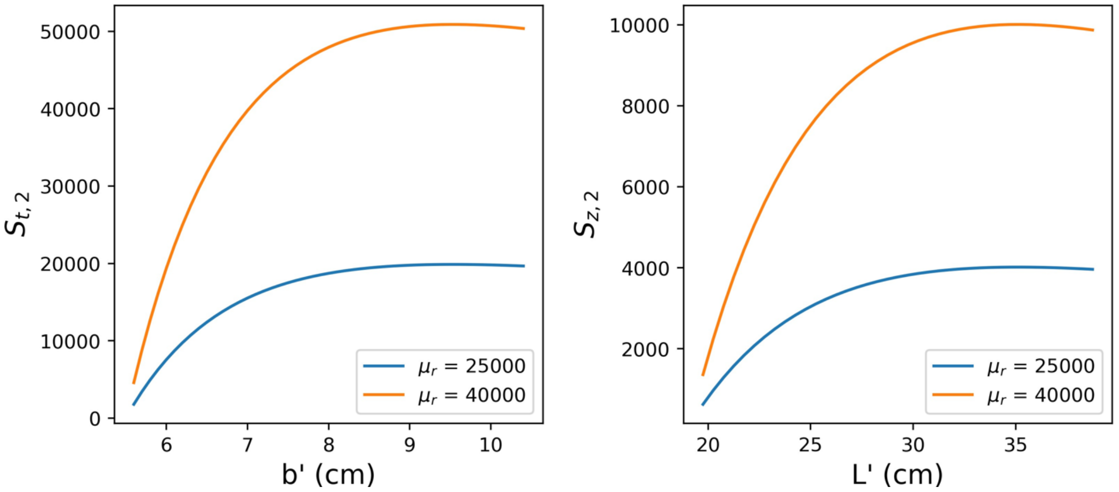

Figure 2.30: Transverse and axial shielding factors for $b = 5.5\,\mathrm{cm}$, $L = 19.25\,\mathrm{cm}$, and for the $S_{z,2}$, $b' = 10.25\,\mathrm{cm}$, and different $\mu_r$.

The shielding factor depends on the relative permeability of the mu-metal, which can be affected by many factors, such as field strength [200], annealing, and mechanical stress. Additionally, the mini-fountain shields aren't placed symmetrically; there is a larger gap between the inner and outer shields at the bottom, to accommodate the MOT chamber, than at the top. To model the effect of the asymmetric shield spacing, a magnetic field of 50 $\mu$T was applied through the $z$-direction in the COMSOL simulation. The resulting $S_{z,2}$ vs $z$ is shown in figure 2.31. The resulting standard deviation of the background magnetic field in the flight tube would be, for this worst-reasonable-case scenario, $\sigma_{B,f} = 5.8\,\mathrm{nT}$, which would result in a maximum fractional frequency uncertainty on the second-order Zeeman shift due to static field inhomogeneities of $\sigma_{Z2} = 3 \times 10^{-16}$, which is tolerable.

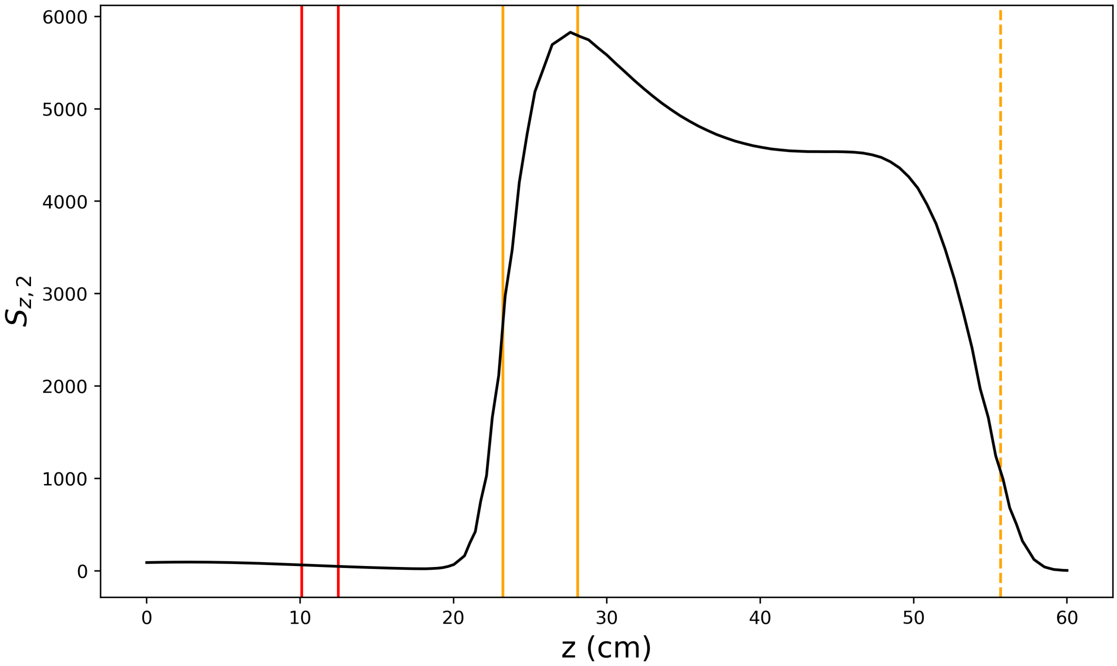

Figure 2.31: Axial shielding factor vs height for the finalised shield, and with $\mu_r = 25{,}000$. This includes upper and lower neck lengths of 1.5 cm.

### 2.4.3 MOT Background Field

In order to perform state selection, as discussed in section 2.2.2, a vertical magnetic field is required to lift the degeneracy between the clock transition and the nearby magnetically-sensitive transitions. For the tests in CsF3, this background field had a strength of 1.7 $\mu$T, which was large enough to clearly resolve each transition, while being small enough to seemingly not compromise the sub-Doppler cooling process. $^{87}$Rb has a $g_F$ twice as large as Cs, so a similar linear Zeeman frequency splitting could be achieved with a field of 0.85 $\mu$T. This will require additional current-carrying coils around the mini-fountain's MOT chamber, and it is important that these do no adversely affect the homogeneity of the magnetic field in the interrogation region. The simulated result of two coils placed near the MOT, each with a radius of 7 cm with 12 mA through 5 turns, in addition to the C-field coil inside the inner shield, is shown in figure 2.32.

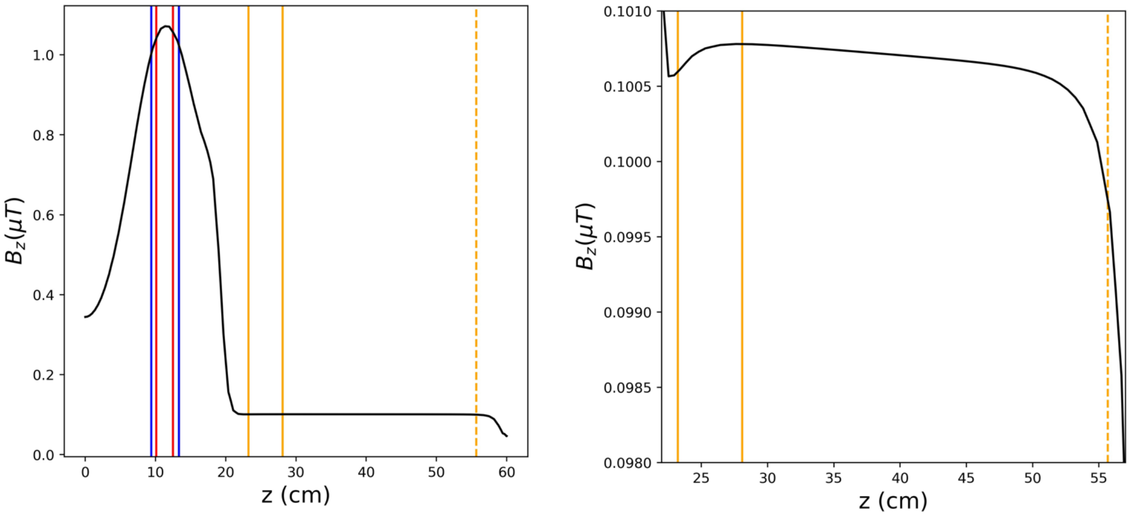

Figure 2.32: Simulated $B_z$ for the height of the mini-fountain. a) Blue lines denote the positions of the additional coils, such that the field in the centre of the MOT is 1.07 $\mu$T b) Close-up of the interrogation region, which is slightly affected by the additional coils.

While a close-up view of the interrogation region shows some small differences compared to figure 2.29, the result is that there is only 1 nT difference between $\langle B_z \rangle_c$ and $\langle B_z \rangle_f$, with the inhomogeneity of the field in the flight tube being $\sigma_{B,f} =$ 0.15 nT, which is much less than the expected inhomogeneity due to the finite shielding factor, and thus adds negligible additional frequency uncertainty.

The final shield geometry is shown in figure 2.33.

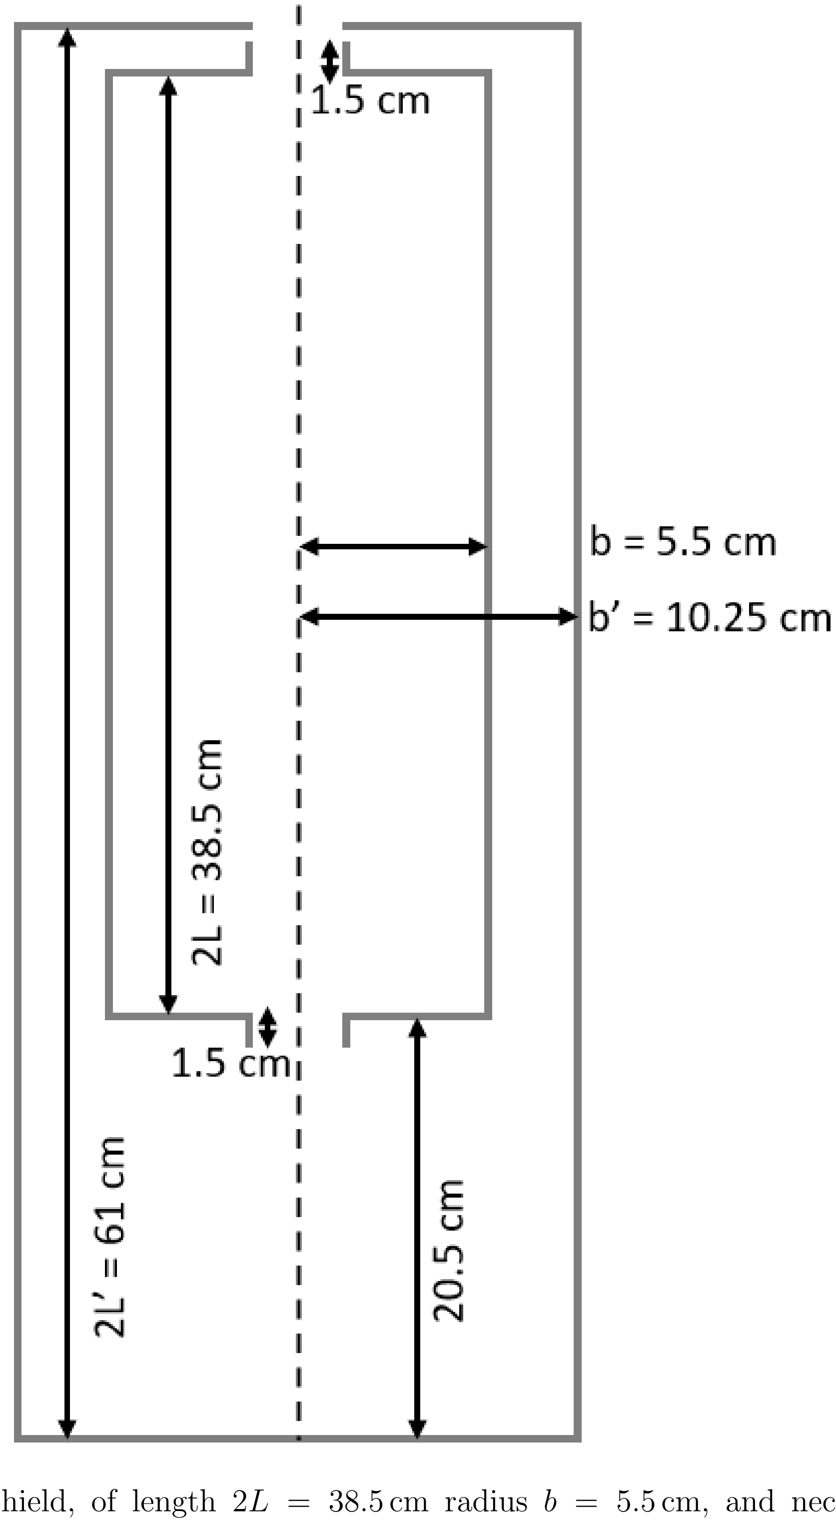

Figure 2.33: Inner shield, of length $2L = 38.5\,\mathrm{cm}$ radius $b = 5.5\,\mathrm{cm}$, and neck lengths of 1.5 cm, placed inside the outer shields, of length $2L' = 61\,\mathrm{cm}$ and radius $b' = 10.25\,\mathrm{cm}$, with a lower gap between them of 20.5 cm.

## 2.5 Optics

The aspect of fountains that requires the most maintenance is the optics. Beams are typically manipulated by free-space components that require meticulous alignment, and coupled into optical fibres, which requires precise positioning to achieve optimal efficiency. The mounts are subject to external vibrations and relaxation of the springs that hold the parts together, such that, over time, the alignment can change. Additionally, separate lasers are used for the cooling and repump transitions, which adds to the complexity of the optical system.

The mini-fountain makes greater use of fibre-coupled components, such that the amount of free-space alignment and fibre coupling is minimised. Additionally, the second repump laser can be eliminated by replacing it with an electro-optic modulator (EOM), which produces sidebands on the main beam. As well as reducing the complexity of the optics, this replaces a £20,000 component with a £5,000 component. Fibre-coupled AOMs replace the free-space double-pass AOMs traditionally used in the fountain optics.

### 2.5.1 Laser and Spectroscopy

The source of the laser light is a Vescent D2 780 nm distributed Bragg-reflector (DBR) laser. The power curve for our laser is shown in figure 2.34.

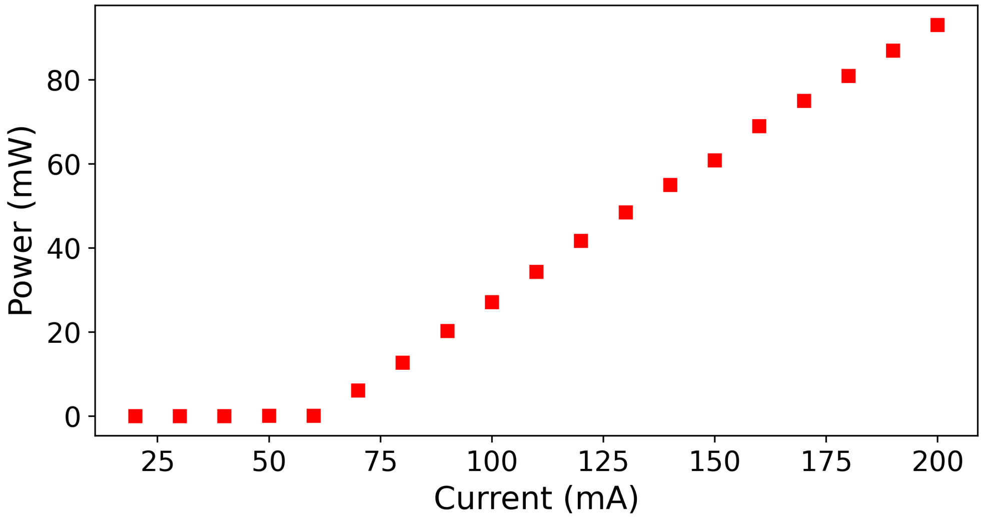

Figure 2.34: Output power vs input current for our DBR laser at 21.7°C. The threshold current is 60 mA, with a maximum power of 93.1 mW at 200 mA.

Most of the output power is coupled into a fibre EOM (50 mW incident, 25 mW coupled), and a small fraction is split off for saturated absorption spectroscopy. Similar to figure 2.3, a double-pass AOM down-shifts the laser light to be resonant with the $5\,^2S_{1/2}\,|F=2\rangle \rightarrow 5\,^2P_{3/2}\,|F'=3\rangle$ cooling transition, and is split into two probes and a pump. The laser current is directly modulated at 4 MHz to produce the error signal. This setup is shown in figure 2.35. The resulting spectrum and error signal, for 0.12 mW pump power and 0.012 mW probe powers are shown in figure 2.35.

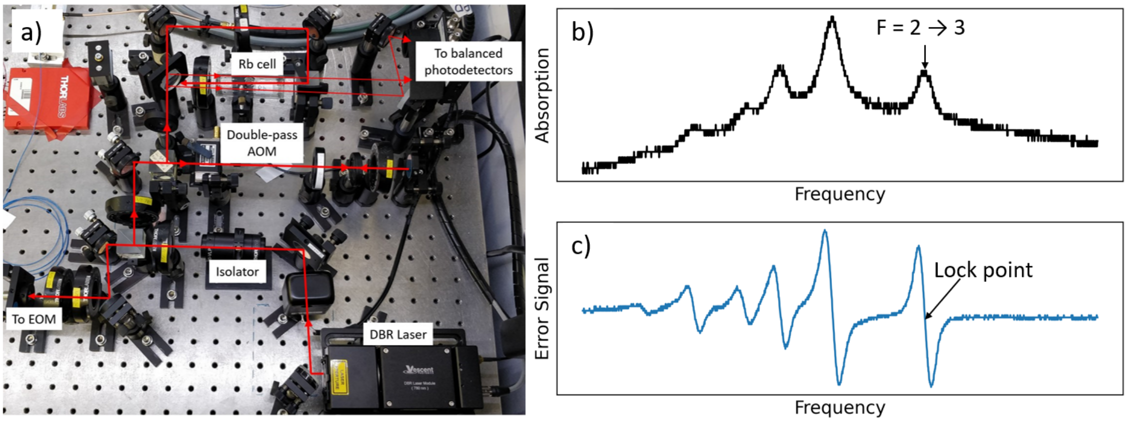

Figure 2.35: a) Photo of the laser system and saturated absorption spectroscopy setup for the mini-fountain optics, with the beam path traced in red. b) Absorption spectrum and c) error signal obtained with saturated-absorption spectroscopy on the $D_2$ manifold of $^{87}$Rb.

### 2.5.2 EOM for Repump Generation

An EOM is a device which modulates the phase or amplitude of an optical wave, such that the resulting spectrum looks like a carrier with sidebands at the modulation frequency. Lithium niobate is a commonly used material for near-IR EOMs, and it exploits the variability of the material's refractive index with applied electric field to modulate the phase of the beam.

To produce a sideband resonant with the $5\,^2S_{1/2}\,|F=2\rangle \rightarrow 5\,^2P_{3/2}\,|F'=2\rangle$ repump transition from cooling light requires a modulation frequency of the ground-state splitting minus the excited-state splitting, such that the input frequency to our EOM (phase modulator, iXblue NIR MPX800-LN-05) must be 6.8347 GHz − 0.2667 GHz = 6.568 GHz. This is illustrated in figure 2.13.

In the demonstrator system, approximately 2 mW of repump light was used alongside 6×10 mW cooling beams, so about 3% of the light is repump light. A Fabry-Perot interferometer was used to see the frequency-composition of the EOM output, and the sideband heights vs RF input power are shown in figure 2.36. The $m = +1$ sideband exceeded 3% of the carrier when the RF input (after cable losses of $-4$ dB) was $+2$ dBm, without any noticeable excitation of the $m = \pm 2$ sidebands.

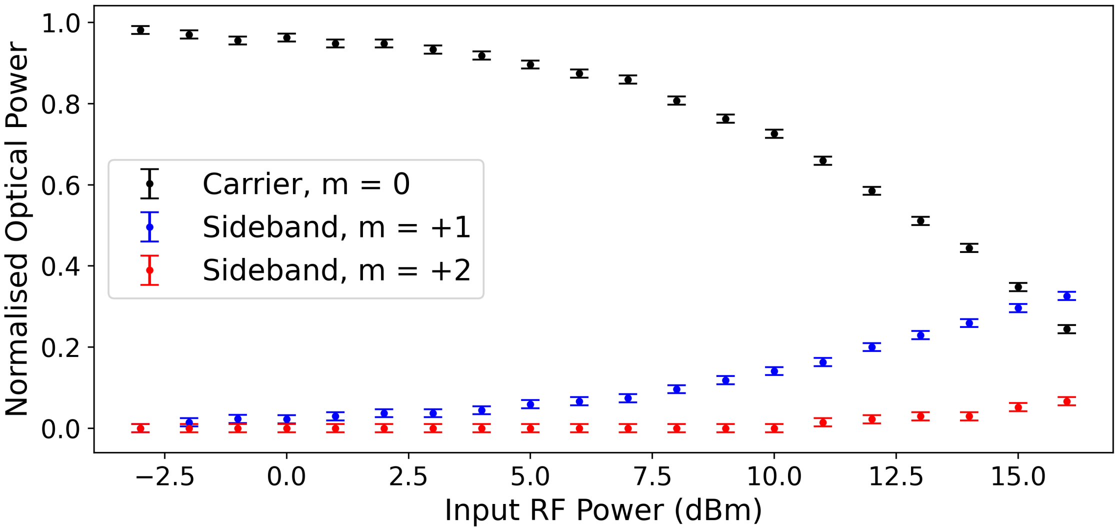

Figure 2.36: Optical output power vs RF input power for the EOM.

To extinguish the repump sideband, the microwave signal to the EOM can be switched off by an RF switch (RF Lambda SPSTA0612G), which claims over 80 dB of isolation between the input and output at 6.8 GHz. Furthermore, the RF synthesiser can be detuned when the repump is not in use.

### 2.5.3 Control Electronics

The (0,1,1) MOT beam geometry requires three separate MOT beam frequencies: horizontal beams that are not detuned for launch, up-going beams that are blue-detuned for launch, and down-going beams that are red-detuned for launch. Furthermore, the fibre-coupled AOMs that will be described on the next section are single-pass, rather than double-pass as in the demonstrator system and CsF3, so the nominal input frequencies are twice as large.

In principle, the required RF generation system is the same as that of NPL's full-sized Cs fountain clocks, which also use three separate frequencies for the (1,0,0) geometry. Unfortunately, the direct digital synthesis (DDS) boards used in that design do not exist in a model that will output the required 200 MHz for the mini-fountain's fibre AOMs, so a different DDS is used in a new, but similar, design.

One suitable DDS board is the Analogue Devices AD9910, which accepts a clock input of up to 1 GHz and has a maximum output frequency of 400 MHz. Each board comes with a 25 MHz quartz oscillator and a phase-locked loop for frequency multiplication. The quartz oscillator from the board that seeds the horizontal ("H") beams is connected to the clock input of the up-going ("U") and down-going ("D") boards, such that all boards are using a common reference. Each board's PLL is then populated to optimise for 40× multiplication, such that the maximum clock frequency is used. The programming of the DDS is done manually using PC software via a USB microcontroller. However, the fast, synchronised frequency switching required to launch the cold-atom cloud will be difficult to achieve via USB. Instead, control over which frequency is active is done by a signal from the PC that bypasses the microcontroller and goes straight to the "profile" pins on the DDS chips, that selects which of the two pre-programmed frequencies to output.

The spectroscopy AOM that determines the laser frequency is still a free-space double-pass AOM, with a peak efficiency of 70 % and a full-width-half-maximum frequency bandwidth of 50 MHz in the cat's-eye configuration [72]. Agile and continuous frequency control is required to set the detuning for the MOT and molasses phases and to ramp up the detuning during the sub-Doppler cooling phase. The same VCO as used in the demonstrator is thus appropriate here, and its output frequency vs control voltage is show in figure 2.38. Since the maximum total detuning occurs at 72.4 MHz away from resonance, which is larger than the AOM bandwidth, then the transmission to the saturated absorption spectroscopy is reduced by a factor of around 10 at this large detuning. This reduced power is seen to be sufficient to keep the laser locked.

Control of the amplitude of the RF signals is achieved with four separate voltage-controlled attenuators (VCAs) (Mini Circuits ZX73-2500-S+) after the output of the DDS boards and VCO. The control electronics box is shown in figure 2.37.

The maximum output power from each of the three DDS boards is 3.3 dBm, and the subsequent maximum powers out of the VCAs were found to be −2.7 dBm for H, −3.1 dBm for U, and −3.3 dBm for D. No notable additional attenuation

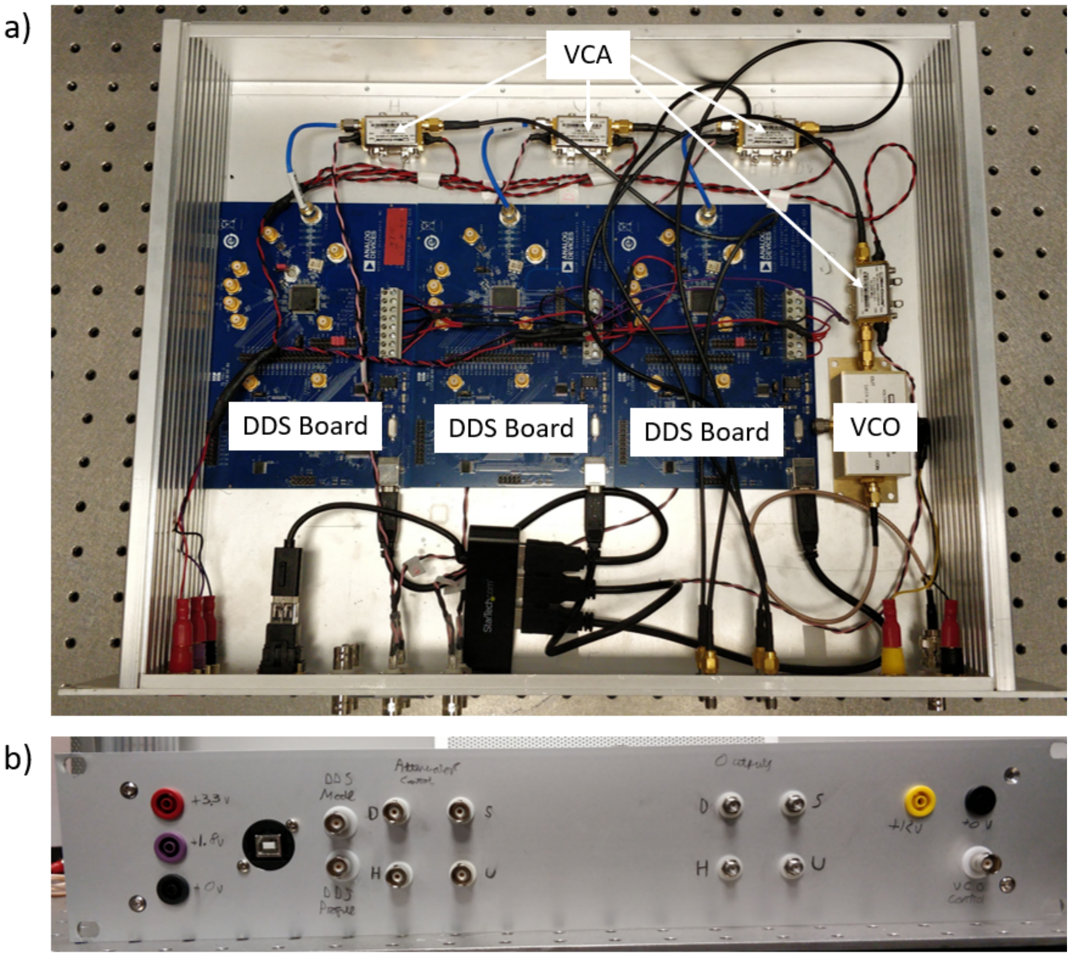

Figure 2.37: a) Control electronics box for all-in-fibre control optics, showing the DDS boards as the source of the fibre AOM RF signal, and VCO as the source of the spectroscopy AOM RF signal. b) Front panel.

occurs until the control voltage is below 5 V, and at 0 V (maximum attenuation) the output signal is entirely swamped by the noise in the spectrum analyser. The VCA datasheet claims −60 dB of attenuation for 200 MHz. Each of the H, U, and D RF signals are then amplified by +34 dB by broadband RF amplifiers, such that the resulting power levels of ≈ +31 dBm are sufficient to efficiently drive the fibre-coupled AOMs. The free-space double-pass AOM is amplified by +28 dBm, and optimum diffraction efficiency occurred at +29.6 dBm.

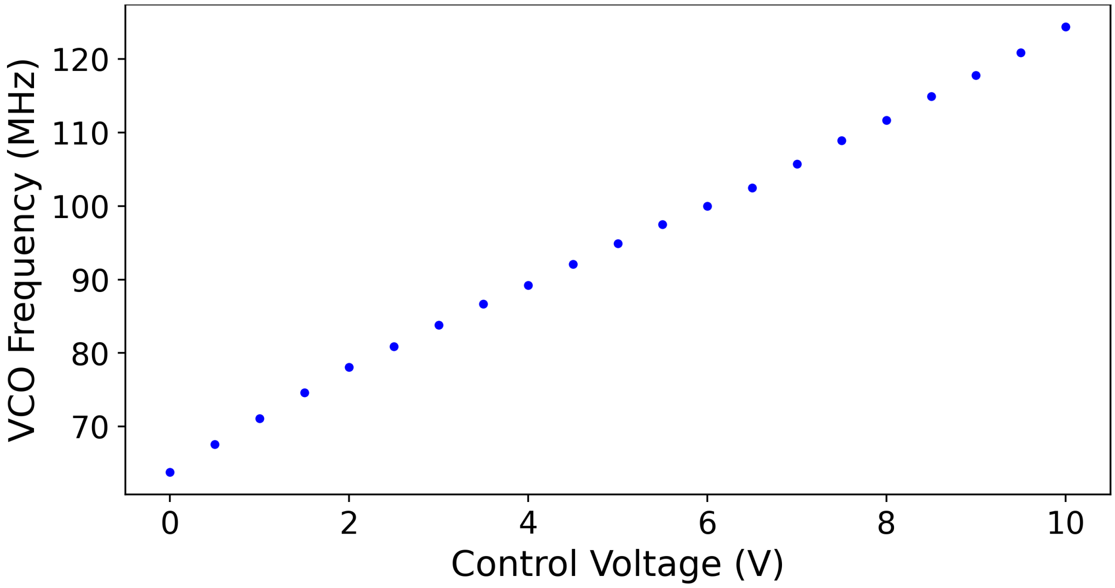

Figure 2.38: VCO output frequency vs control voltage. Because of the double-pass configuration, resonance occurs at 6.0 V, and the maximum negative detuning occurs at 0 V which is 72.4 MHz ($-11.9\,\Gamma$). Error bars of less than 0.1 MHz have been omitted for clarity.

### 2.5.4 All-In-Fibre Splitter Optics

The EOM has a maximum input power of 25 mW, and an insertion loss of $-3.5$ dB, such that the maximum output power is 11 mW. This is sufficient (but close to the minimum required) to seed the tapered amplifier (Toptica BoosTA). The optimised TA output power vs input current is shown in figure 2.39.

A 1x3 33:33:33 fibre-coupled non-polarising, polarisation-maintaining beamsplitter (Oz Optics FOBS 13P-111) then splits the TA output power into three separate fibres, with splitting ratios in each of the three outputs measured to be 33.2%, 34.3%, and 32.5%, respectively. Next, three single-pass fibre-coupled AOMs (AA Opto-electronic MT200-NIR10-Fio-PM0,5-J1-A-(s-)VSF) downshift the light by 200 MHz, such that the actual laser detuning from the cooling transition is

$$\delta = 2\nu_{\text{VCO}} - 200 \text{ MHz}, \tag{2.42}$$

where $\nu_{\text{VCO}}$ is the frequency of the spectroscopy VCO. The output from each AOM is then split by a 1x2 50:50 splitter (ThorLabs) into the six MOT beams. This set-up is shown in figure 2.40.

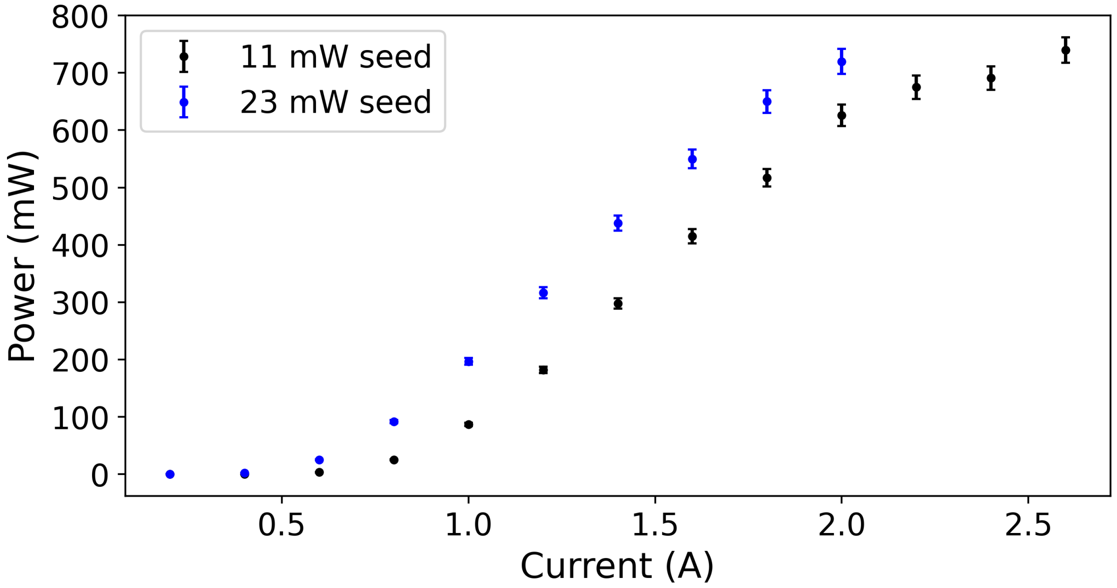

Figure 2.39: Fibre-coupled output power vs input current for the TA, with the EOM seed of 11 mW and a higher power seed for comparison. The slopes are the same in each case, but the higher seed power has a lower current threshold.

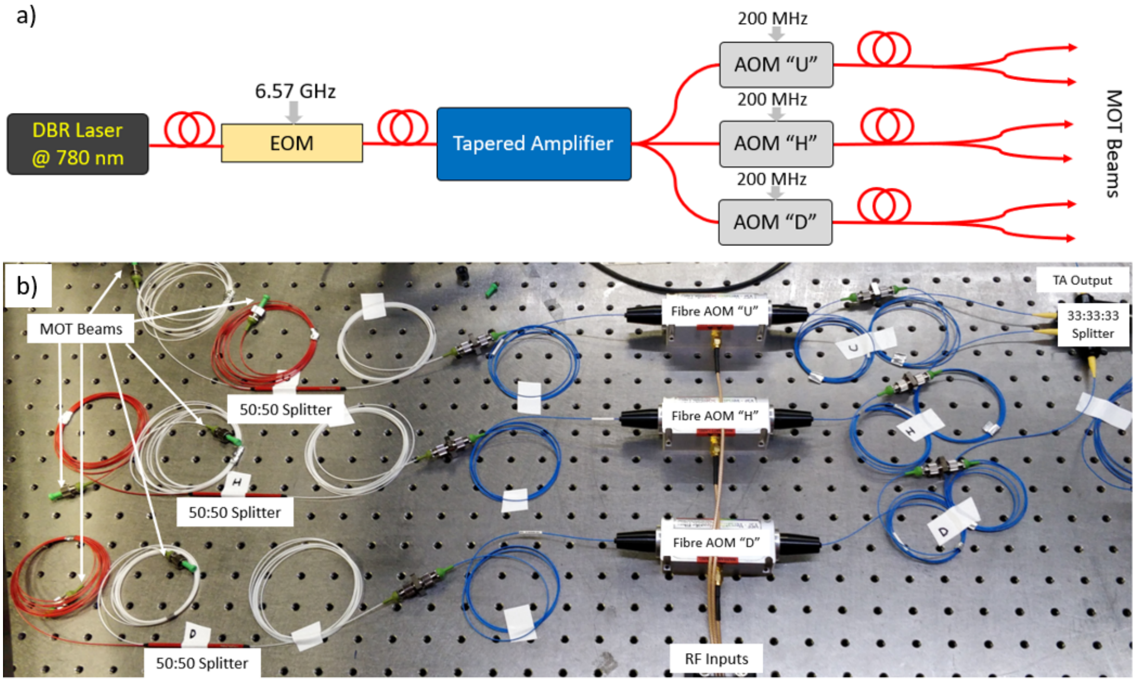

Figure 2.40: a) diagram of, and b) photo of, all-in-fibre AOMs and splitter optics. All fibre components and splitters are connected by mating sleeves.

The intensity of each MOT beam is controlled by controlling the input RF power to the fibre AOMs. The power in each fibre vs VCA control voltage is shown in figure 2.41, for a TA current of 1.12 A. The maximum efficiency of each AOM is around 50%, with some variation.

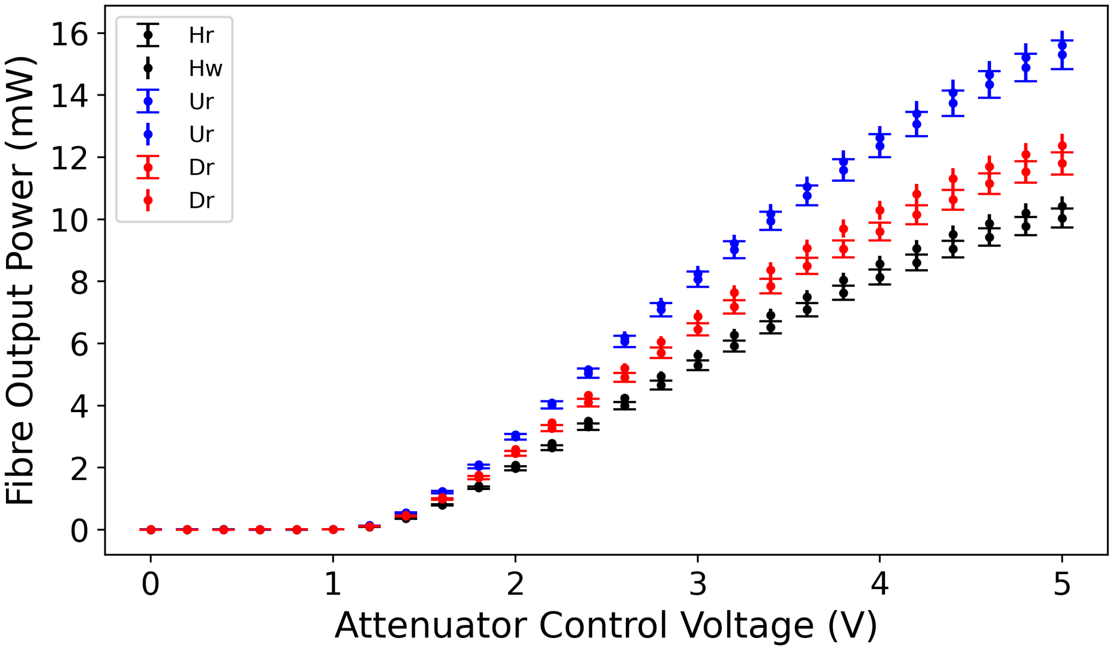

Figure 2.41: Output power of each of the six MOT beams after the 50:50 splitter vs VCA control voltage, which controls the RF power to the fibre AOMs.

It was found that the AOM diffraction efficiency continued to increase, albeit at a decreasing rate, even at the maximum available RF input power. This suggests that the optimal input power is slightly above 31 dBm, which is consistent with the manufacturer's datasheet, which states a maximum input power of 31.7 dBm (1.5 W). The TA current was deliberately set such that the weakest pair of MOT beams, the horizontal beams, were just above 10 mW, being 10.04 mW in the red splitter output and 10.42 mW in the white splitter output. The down-going beams reached maximum powers of 11.79 mW and 12.40 mW, and the up-going beams reached 15.29 mW and 15.60 mW in the red and white fibres, respectively. The discrepancies are probably due to inconsistent coupling efficiencies in the fibre mating sleeves. This issue could be removed by splicing the fibres together, which would result in decreased losses across the whole system. However this would also make reconfiguring the system much more difficult. In order to match the powers in each beam, the D and H AOMs should have their RF input attenuated. It was found that all six MOT beams were at approximately 10 mW when the VCA voltages to the H, U, and D AOMs were 5.0 V, 3.6 V, and 4.2 V, respectively.

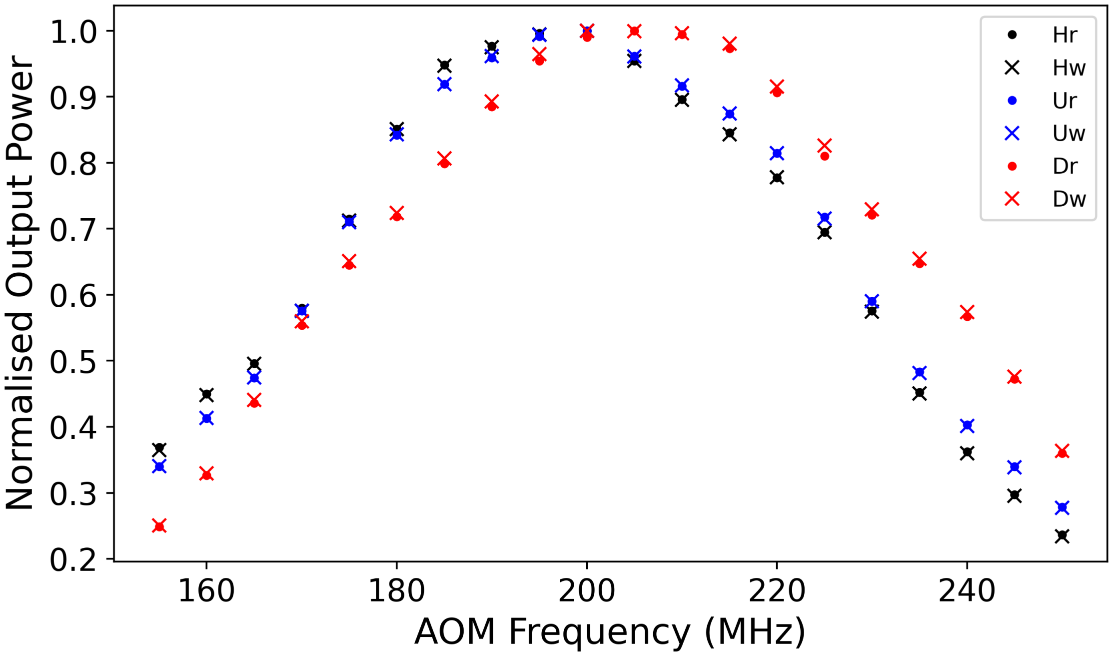

Figure 2.42: Normalised output power of each of the six MOT beams, having been approximately balanced via the attenuator voltage, vs the AOM modulation frequency. Error bars of $\lesssim 3\,\%$ have been omitted for clarity.

The next aspect of the fibre optic control system to check is the frequency modulation bandwidth of the fibre AOMs. In the (0,1,1) geometry, the launch velocity is determined by the beam detuning $\delta\nu$ according to

$$v = \pm\sqrt{2}\lambda\delta\nu, \tag{2.43}$$

where the positive detuning applies to the up-going beam and the negative to the down-going beam. The extra factor of $\sqrt{2}$ relative to equation 1.42 comes from the fact that the detuned beams are at 45° to the vertical. For example, to launch to a height of 44 cm above the MOT requires an initial velocity $v = 2.94\,\mathrm{ms}^{-1}$, and thus a detuning of $\delta\nu = \pm 2.66\,\mathrm{MHz}$. Thus, it is required that the AOMs have sufficient bandwidth to not appreciably imbalance the intensities during the launch. The output power vs frequency for each fibre AOM is shown in figure 2.42. Each one seems to have a FWHM bandwidth of approximately 70 MHz, much larger than the launch detuning, thus very little imbalance is imparted during launch. This large bandwidth also opens up the possibility of replacing the free-space AOM in the spectroscopy section with a fibre-coupled AOM, as the change in VCO frequency between normal MOT operation and the end of the sub-Doppler cooling ramp is only 36.2 MHz, which results in a 72.4 MHz shift in the double-pass configuration.

The final aspect to check is the extinction ratio. In section 1.2.2, the light shift due to near-resonant optical light was discussed:

$$\frac{\Delta\nu_{\rm ls}}{\nu_0} = \frac{\Gamma^2}{8\nu_0\delta}\frac{I}{I_{\rm sat}}, \tag{2.44}$$

such that the light shift is minimised by minimising the intensity, and maximising the detuning, of the laser light. In NPL's primary frequency standards, mechanical shutters in front of the fibres ensure total attenuation of the MOT beams. This is not practical in the mini-fountain, as the aim is to minimise the number of free-space fibre couplings to improve the overall robustness of the system.

To test the acceptable residual MOT beam intensity, NPL-CsF3 was run with the horizontal MOT beams left partially on, between 0 mW and 5 mW, during the interrogation time. The same can not be done with the vertical beams because these pass directly along the atoms' trajectory. The resulting frequency shift to the clock transition is shown in figure 2.43.

In the mini fountain, all six MOT beams are "partially on" during interrogation (so additional factor of 1.5 in intensity), the limiting aperture of the cavity will be approximately 4× closer to the centre of the MOT (additional factor of 16 in solid-angle), and it will be 16 mm in diameter, as opposed to 11 mm in CsF3 (additional factor of $\sim$2 in solid-angle). The different $\Gamma$ and $I_{\rm sat}$ from Cs to Rb approximately cancel each other out. The gradient of light shift vs beam power in the mini-fountain may thus be approximately 50 times larger than in CsF3, which would be $2.0\times10^{-12}$ mW$^{-1}$. To keep the shift below $10^{-16}$, each MOT beam's residual intensity must be $\leq 50$ $\mu$W per beam, which from 10 mW beams requires $-23$ dB of attenuation. The extinction ratios of each fibre AOM when turned fully-off were found to be $-49$, $-46$, and $-49$ dB for the H, U, and D AOMs, respectively. This should be sufficient to eliminate the light shift.

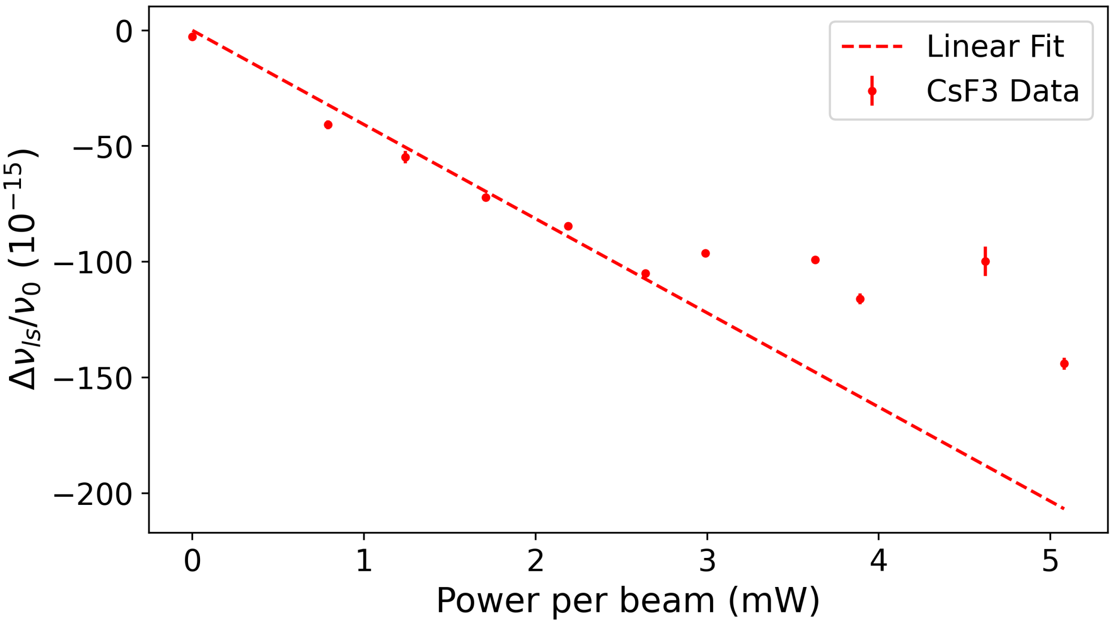

Figure 2.43: Light shift vs MOT beam power in CsF3, for $\delta = -71\,\text{MHz}$. The line of best fit (to the first six linear points) passes through the origin and has a gradient of $-4.07 \times 10^{-14}\,\text{mW}^{-1}$. At the higher-power data points, it might be the case that more than just the light shift is changing if the atom number and fringe contrast reduce due to excessive scattering.

### 2.5.5 MOT Beam Tubes

In many Primary Frequency Standards the cooling beams are delivered by optical fibres, and allowed to freely diverge until they are collimated by a single converging lens. This means that the length of the beam collimator is set by the focal length of the lens, which itself is determined by the divergence angle out of the fibre and the desired collimated beam diameter. In the demonstrator system and in NPL's Cs fountains, this results in $L_0 = 10\,\text{cm}$ long collimators, which result in Gaussian beam profiles with a $1/e^2$ radius of around 12 mm. The total useful beam radius is $r = 12\,\text{mm}$, limited by the retaining ring that secures the lens, such that the divergence half-angle of the useful light is $\alpha \approx 12\,\text{mm}/10\,\text{cm} = 6.8°$.

To reduce the size of the physics package, it is useful to shorten the beam collimators. This is achieved by introducing a strong diverging lens at the fibre output, such that the beam expands to the desired size in less space. This requires that the focal points of both the diverging and collimating lenses are coincident.

### Geometric Ray Optics

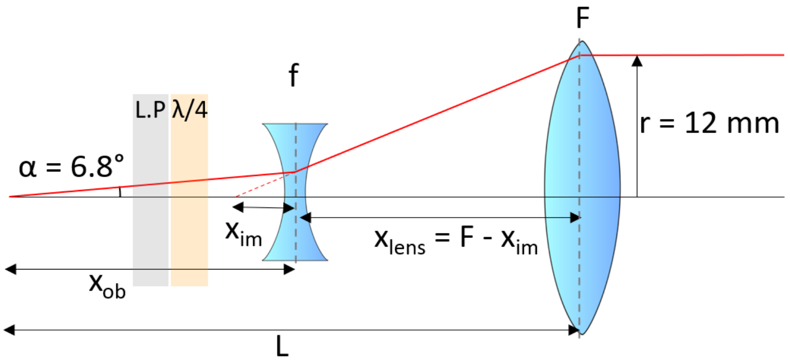

Figure 2.44: Diverging lens of negative focal length $f$ and converging lens of positive focal length $F$ with coincident focal points take light diverging from the fibre and collimate it to radius $r$. Circular polariser consists of a linear polariser and a $\lambda/4$.

The free parameters one wishes to find are the diverging focal length $f$, converging focal length $F$, fibre-lens separation $x_{\mathrm{ob}}$, and lens-lens separation $x_{\mathrm{lens}}$. The constraints are $r$, which is set by the clear aperture of the lens tubes, and the divergence angle out of the fibre $\alpha$. Another constraint is seeking that the total length $L = x_{\mathrm{ob}} + x_{\mathrm{lens}}$ be minimised. For a diverging lens, the lens equation can be written as

$$\frac{1}{x_{\mathrm{ob}}} - \frac{1}{x_{\mathrm{im}}} = -\frac{1}{f}, \tag{2.45}$$

where all lengths are positive. By parametrising the object distance as $x_{\mathrm{ob}} = af$, the lens equation results in an image distance

$$x_{\mathrm{im}} = \frac{af}{a+1}. \tag{2.46}$$

By defining the angle at which the edge of the beam leaves the diverging lens as $\theta$ and the distance from the axis at which it does so as $r' = af \tan\alpha$, we have

$$\tan\theta = \frac{r'}{x_{\mathrm{im}}} = \frac{r}{F}, \tag{2.47}$$

such that the focal length $F$ is given in terms of the parameter $a$ by

$$F = \frac{L_0}{a+1}. \tag{2.48}$$

The total collimator length is thus

$$L = x_{\mathrm{ob}} + x_{\mathrm{lens}} \tag{2.49}$$

$$= af + \frac{L_0}{a+1} - \frac{af}{a+1}, \tag{2.50}$$

which can be differentiated with respect to $a$ to find an optimum:

$$\frac{\partial L}{\partial a} = f - \frac{L_0 + f}{(a+1)^2}, \tag{2.51}$$

such that setting $\partial L/\partial a = 0$ results in

$$a = \sqrt{1 + \frac{L_0}{f}} - 1. \tag{2.52}$$

The system is solved by selecting a value of $f$. There is no optimum value: a shorter $f$ always results in a shorter $L$, as shown in figure 2.45.

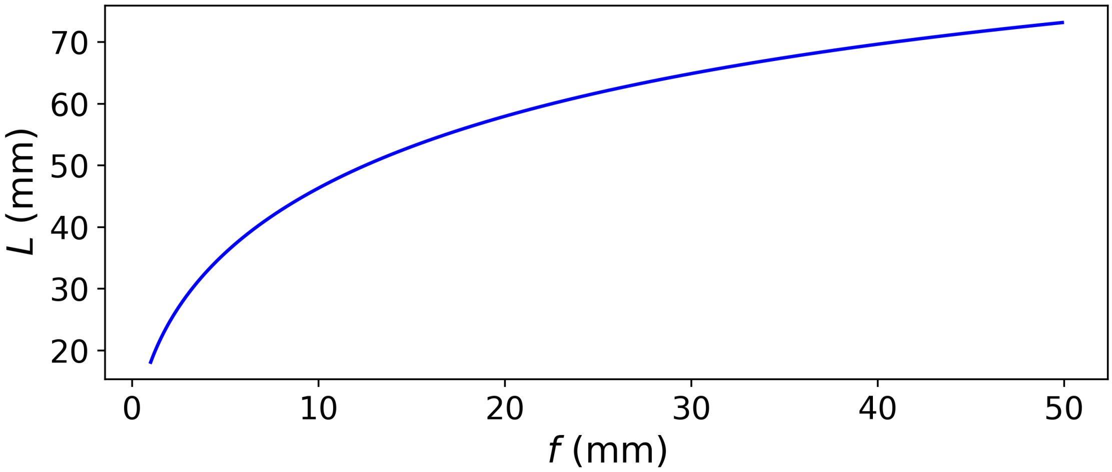

Figure 2.45: Optimum collimator length $L$ vs diverging lens focal length $f$.

Thus, the final constraint is to use the shortest focal length that is commercially available. There are $f = -6$ mm lenses available, which results in $a = 3.2$ such that $F = 23.8$ mm, and $x_{\mathrm{ob}} = x_{\mathrm{lens}} = 19.2$ mm. This means that the total collimator length becomes $L = 38.4$ mm, a significant reduction on the original 100 mm.

In practice, a lens with a back-focal length of $F = 24$ mm will have to do. The collimator length is 1.51": a 1" lens tube, a 1/2" lens tube, and a kinematic mount to hold the fibre tip will be sufficient. The resulting beams are shown in figure 2.46 at a distance of 5 cm from the converging lens, which is the profile of the beam atoms in the centre of the chamber will see. The diameter of the beams, from hard edge-to-edge, is 18 mm, slightly less than the intended 24 mm.

The overall beam profile appears Gaussian, but with a few types of aberration. Striations appear across the diameter of the beams, probably as a result of the resins used to bind the various layers of the circular polarisers together. The intensity also seems to increase at the very edge of the beam, possibly as a result of reflections from the side of the strong diverging lens.

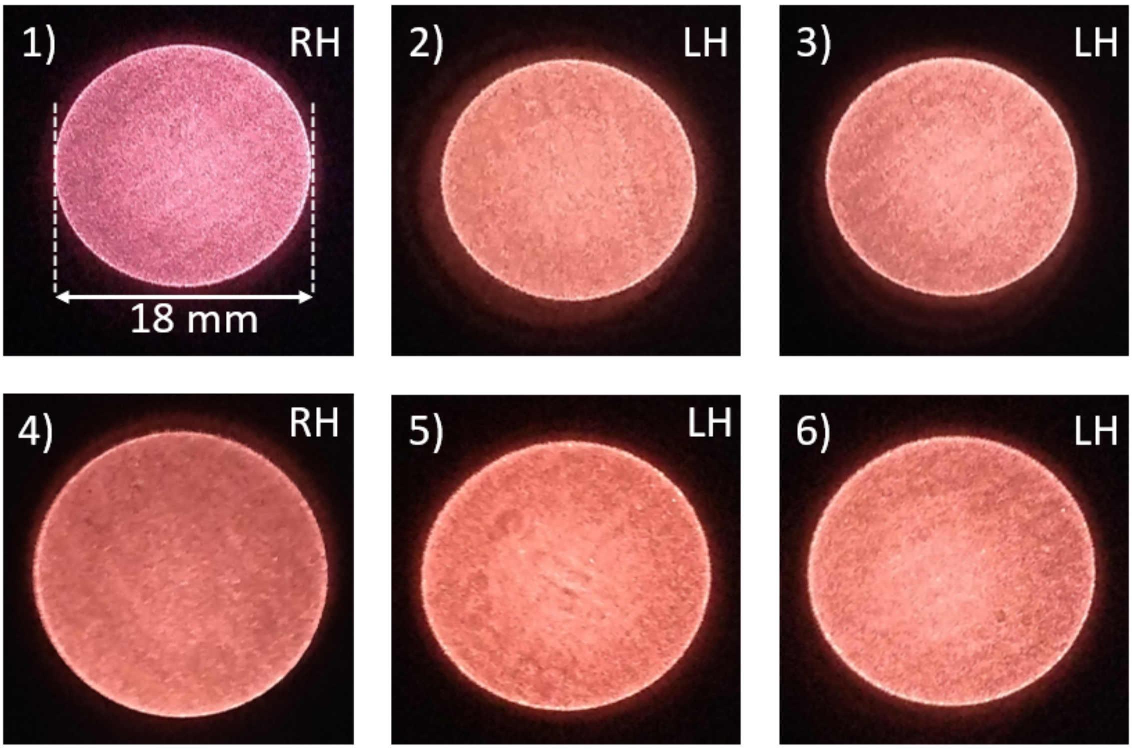

Figure 2.46: Profiles of all six MOT beams, labelled 1-6 and with the handedness of the circular polariser. Taken with a near-IR sensitive camera. Transverse alignment was achieved by maximising the power detected through a 10 mm diameter aperture.

However, when the MOT beams are viewed further away, it is clear that they are not particularly well-collimated, being focused to a slightly smaller diameter about 50 cm away, after which the beam expands again and diffraction takes over, with a single bright spot in the centre and rings making up the rest of the profile. This is shown in figure 2.47.

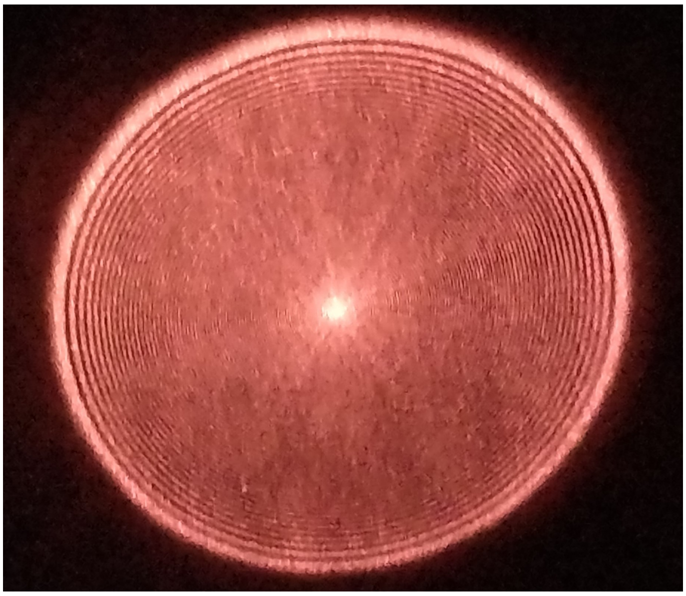

Figure 2.47: Image of the far field of one of the MOT beams, where the diffractive behaviour is evident in the ring structure. Beam diameter is approximately 10 cm.

## Gaussian Optics

To approximate the output of the fibre as a cone with a single divergence angle is overly-simplistic, and instead the details of Gaussian beam propagation must be taken into account: the lowest-order electromagnetic mode propagating in free-space is the $\mathrm{TEM}_{00}$ mode, which has a Gaussian intensity profile and a radius $r(z)$ that changes with propagation:

$$r(z) = w_0\sqrt{1+\left(\frac{z}{z_R}\right)^2}, \tag{2.53}$$

where $w_0$ is the beam waist (the size of the beam at a focus) and $z_R$ is the Rayleigh range (the distance over which the beam area doubles from the waist). This propagation profile is shown in figure 2.48. The Rayleigh range and the waist size are related for this mode by

$$z_R = \frac{\pi}{\lambda}w_0^2. \tag{2.54}$$

In this regime, a beam waist can be treated like an object in traditional ray optics, but with a modified lens equation that takes into account the propagation

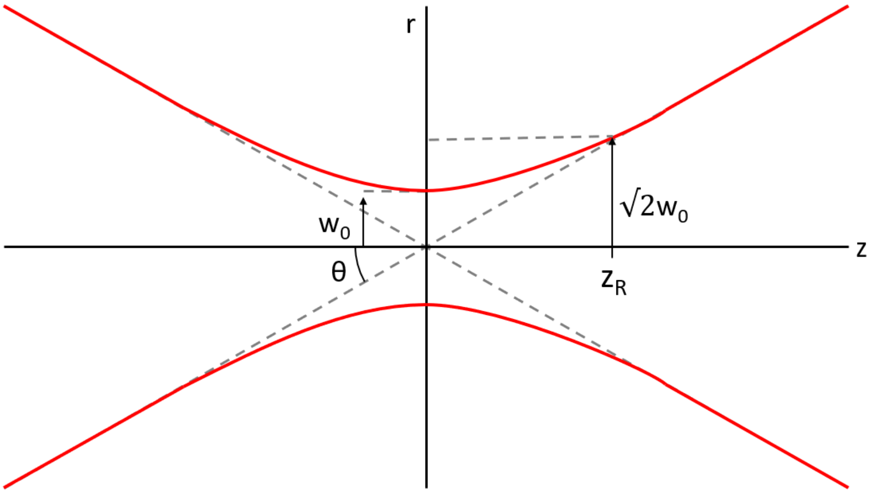

Figure 2.48: A Gaussian beam focused to a waist $w_0$ with a Rayleigh range $z_R$.

of the TEM$_{00}$ (Gaussian) mode [201]:

$$\frac{1}{s'} = \frac{1}{s + \frac{z_R^2}{s+f}} + \frac{1}{f}. \tag{2.55}$$

Perfect collimation is, of course, impossible. Instead, the closest one can get is to minimise the far-field divergence angle, $\theta = \lambda/\pi w_0$ of the outgoing beam, which occurs at a maximum of outgoing beam waist (or Rayleigh range). The reciprocal of equation 2.55 can be used to find the image waist size [201]:

$$\frac{1}{s} = \frac{1}{s' + \frac{z_R'^2}{s'-f}} - \frac{1}{f}. \tag{2.56}$$

The magnification can thus be derived as

$$M = \frac{w_0'}{w_0} = \frac{1}{\sqrt{\left(1 - \frac{|s|}{f}\right)^2 + \left(\frac{z_R}{f}\right)^2}}, \tag{2.57}$$

such that $w_0'$ is maximised when $1 - |s|/f = 0$, such that $|s| = f$; the incoming beam waist should be placed at the focus of the converging lens, as is the case in figure 2.44. For a positive $f$, this means that $s = -f$.

The incoming beam waist to the converging lens is the outgoing (virtual) beam waist from the diverging lens. Equations 2.55 and 2.56 can be used to find the image distance $s'$ and resulting outgoing Rayleigh range $z'_R$ from the diverging lens, with the results shown in figures 2.49 and 2.50 for an initial beam waist out of the fibre $w_0 = 2.65\,\mu$m, which corresponds to the radius of the fibre core.

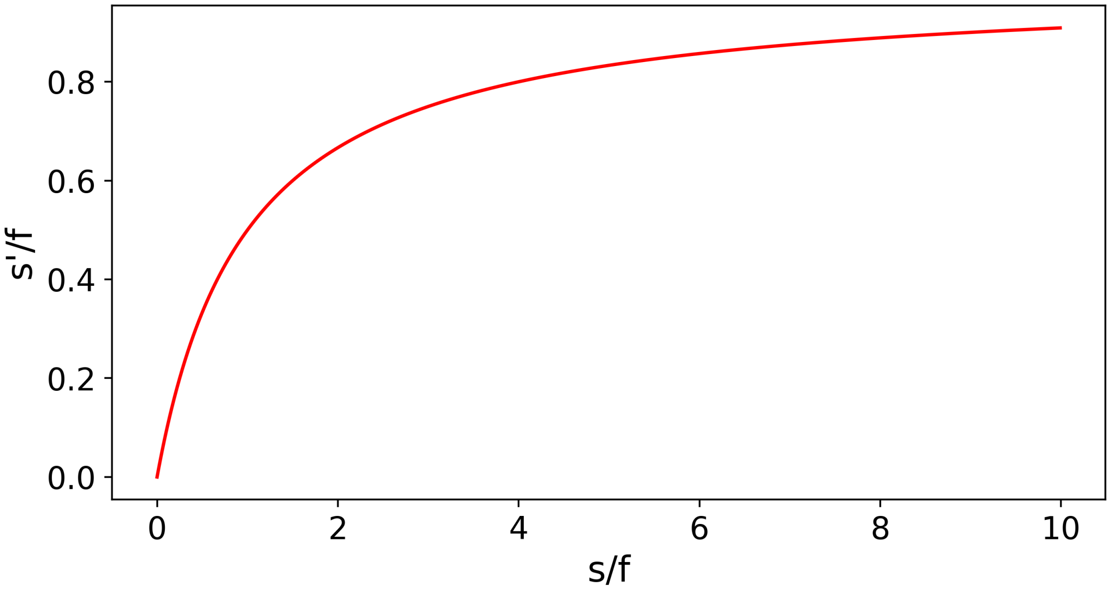

Figure 2.49: Normalised image distance $s'/f$ vs normalised object distance $s/f$.

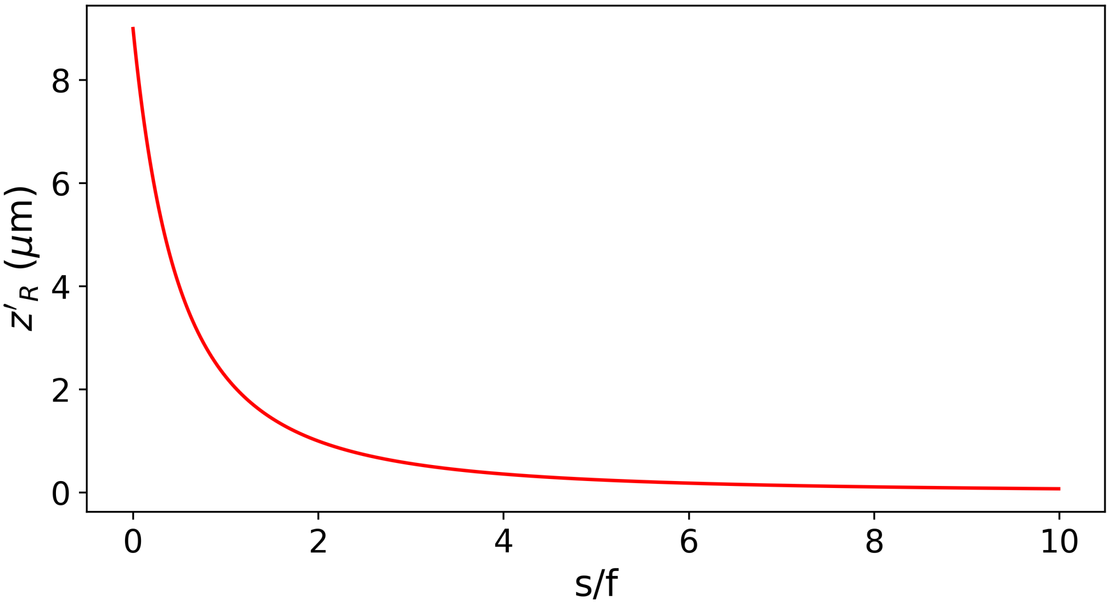

Figure 2.50: Output Rayleigh range from the diverging lens vs normalised object distance $s/f$, for an initial Rayleigh range of $9\,\mu$m.

By placing the diverging lens closer to the fibre output, the image distance decreases rapidly, especially below $s = f$. However, this also has the effect of increasing the output Rayleigh range, which is counterproductive to expanding the beam. The resulting collimated beam size as a function of $s$ is shown in figure 2.51, for a $-6$ mm diverging lens and $+30$ mm converging lens. To achieve $r > 12$ mm thus requires $s > 7$ mm. This is a slightly shorter lens-fibre separation than was derived using geometric ray optics, but results in a similar lens-lens separation.

Since "collimation" here simply means the outgoing beam has a minimised divergence angle, the resulting beam profile vs distance between the ray-optics and Gaussian-optics designs are very similar; it's impossible to remove all divergence in each case, and as the beam has to pass through small apertures diffraction is also inevitable. Thus, the diagonal MOT beams will stay with the original 19.2 mm fibre-lens and lens-lens separation, whereas the horizontal beams, which need to be short enough to ensure sufficient clearance between the optical fibres and outer shield, will be shortened by shortening the fibre-lens separation, while keeping the lens-lens separation unchanged.

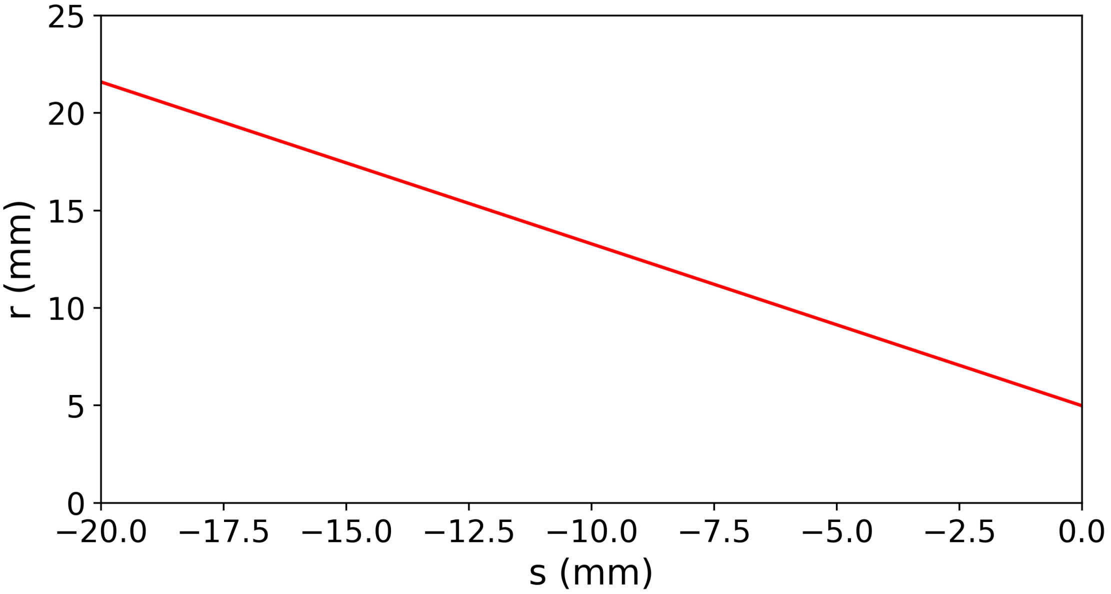

Figure 2.51: Collimated beam size vs diverging lens-fibre separation.

### 2.5.6 Repump Laser

While it has been demonstrated that the EOM can provide the necessary repump sideband to the main cooling light, it would, by necessity, be operating at close the the maximum allowed input power before optical damage starts occurring to the modulator. To hedge against potential problems that may arise due to this, a distributed-feedback (DFB) laser was also acquired for the purpose of generating the repump light. The fibre-coupled optical power vs input current for the DFB is shown in figure 2.52.

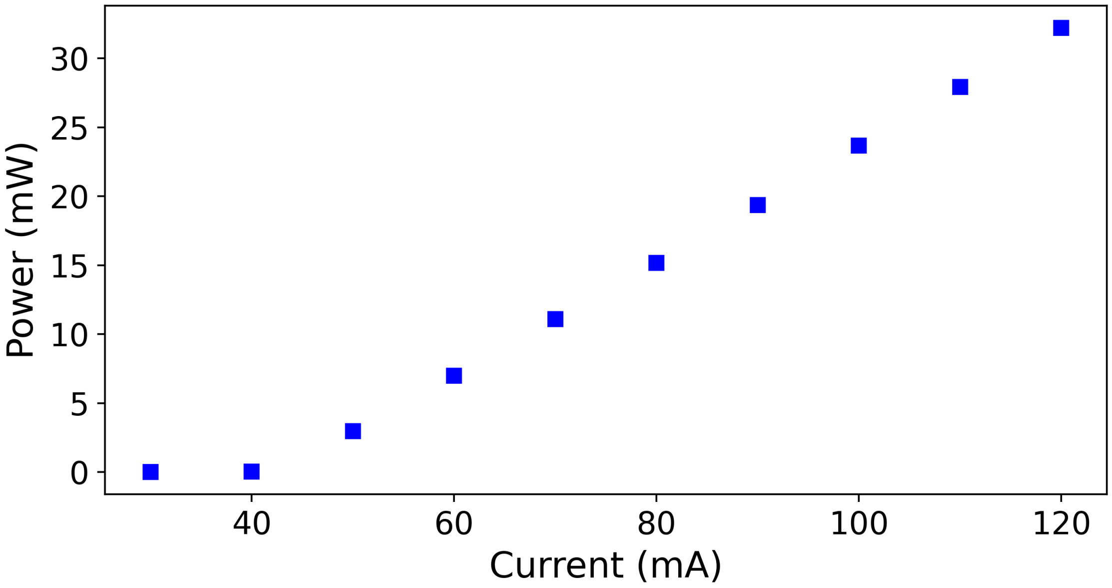

Figure 2.52: Fibre-coupled output power vs input current for the DFB repump laser. Error bars of $\lesssim 2\%$ have been omitted for clarity.

The laser output is split by a 90:10 fibre splitter, with 10% going directly to saturated absorption spectroscopy: the laser frequency is locked to the $F = 1 \rightarrow F' = 1$ & 2 crossover peak, which occurs 78.5 MHz below the $F = 1 \rightarrow F' = 2$ repump transition. As with the demonstrator system, the repump will need to be switched on and off quickly to facilitate detection with the MOT beams. The 90% splitter output is collimated into free-space and passes through an AOM, which brings the diffracted light into resonance with the repump transition. The optimised AOM diffraction, with RF power from a 9.1 dBm VCO amplified by +28 dBm, is 50% efficient and the subsequent fibre coupling into a 2x2 50:50 splitter, which is shared by the horizontal MOT beams, is 63% efficient, such that a total 4.77 mW is delivered to the MOT chamber when the laser current is 85 mA. The repump setup and $F = 1$ transition manifolds from saturated absorption spectroscopy are shown in figure 2.53.

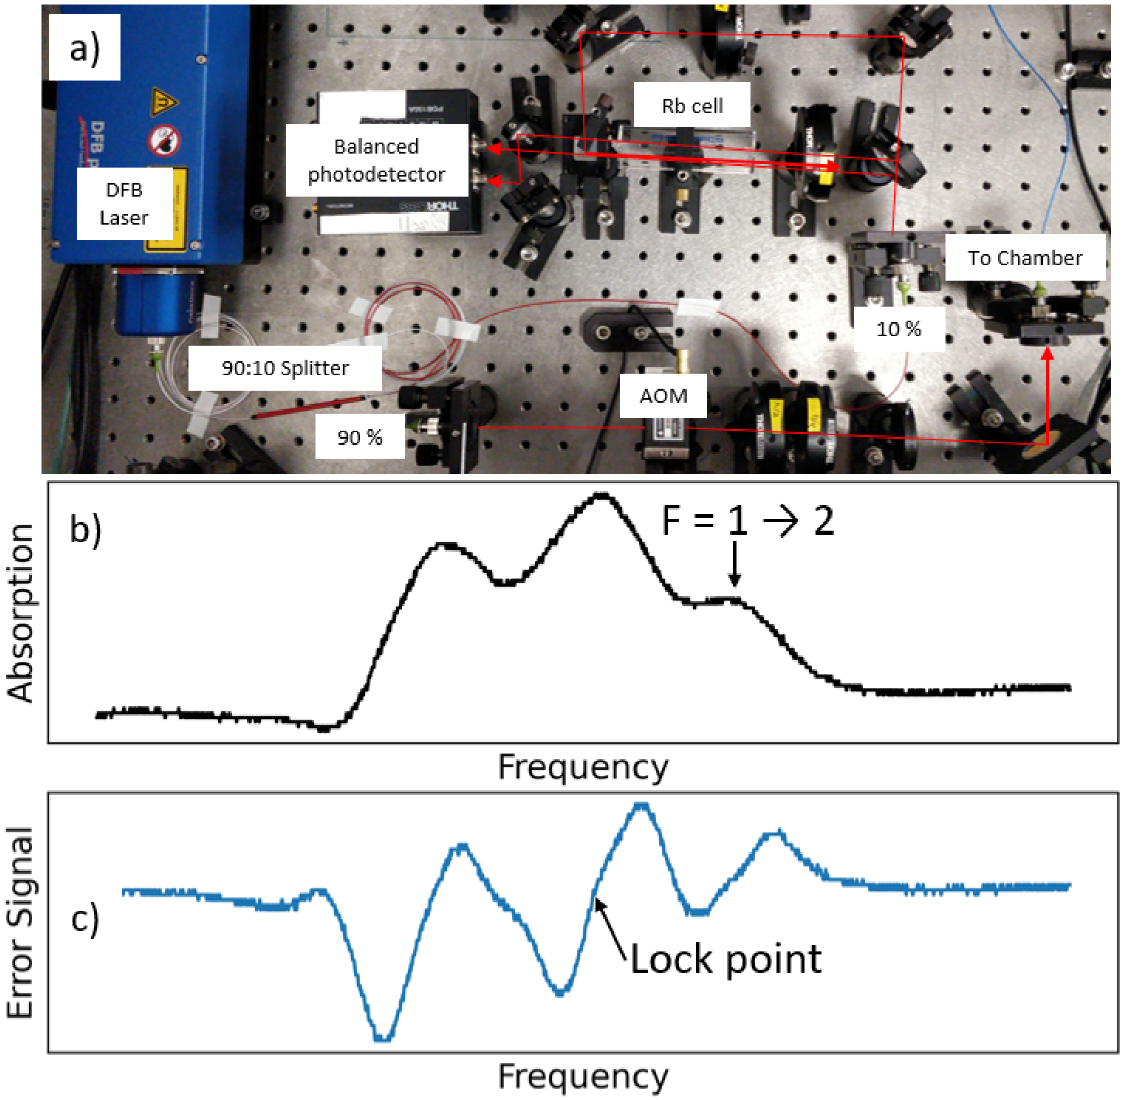

Figure 2.53: a) Spectroscopy setup for locking the DFB laser, and the AOM for fast extinction of the resonant light. b) Spectrum for the $F = 1$ manifold of $^{87}$Rb. c) Error signal for locking to the $F = 1 \rightarrow F' = 1$ & 2 crossover transition, which occurs at a laser current of 85 mA and a diode temperature of 30.4°C.

# Chapter 3

# Fountain Construction

## 3.1 Cavity Tuning

With the Ramsey cavity manufactured, its frequency needs to be measured and brought it into resonance with the clock transition. The coupling between the ring waveguide and the Ramsey cavity can also be investigated. These measurements are made with a vector network analyser (VNA) that measures the reflection coefficient, $S_{11}$, of the cavity-waveguide assembly. An additional loop antenna can be inserted through the atom-passage holes to make a measurement of the (perturbed) transmission coefficient, $S_{21}$. and the reflection coefficient from the additional feed, $S_{22}$.

The Ramsey cavity has radius $a = 30.0\,\mathrm{mm}$ and height $d = 48.622\,\mathrm{mm}$. Its frequency is adjusted by changing $d$. From equation 2.15, the tuning sensitivity is

$$\frac{\partial f_c}{\partial d} = -\frac{c^2}{4d^3 f_c} = -28.6\ \mathrm{MHz\,mm^{-1}}. \tag{3.1}$$

The waveguide frequency is tuned by lengthening its length $Z$. From equation 2.20 for a $\mathrm{TE}_{106}$ mode, the tuning sensitivity is

$$\frac{\partial f_{\mathrm{wg}}}{\partial Z} = -\frac{9c^2}{Z^3 f_{\mathrm{wg}}} = -17.0\ \mathrm{MHz\,mm^{-1}}. \tag{3.2}$$

Later in the tuning process the coupling apertures will be covered with viewports, so the Q-factor must be tuned before then. The theoretical unloaded Q of our cavity is $Q_0 = 36,100$. The aim is to make the coupling sufficiently strong to reduce the loaded Q to around 5,000, to reduce the size of the cavity pulling shift, which has a steep gradient near resonance. To determine which aperture size to use, a finite-element model was used to simulate the electromagnetic fields in the cavity and waveguide [194]. Additionally, the Q-factor was calculated analytically from a parallel, equivalent-circuit model of our cavity and the feeding waveguide [121, 182, 184]: the external Q-factor $Q_{\mathrm{ex}}$ due to the loading can be written, for the case of $f_c = f_{\mathrm{wg}}$, as

$$Q_{\mathrm{ex}} = \left( \frac{\pi f_c}{4c} \frac{XY}{\alpha_m k_{10} \beta} \right)^2 , \tag{3.3}$$

where $XY$ and $k_{10} = \sqrt{k^2 - (p'_{01}/X)^2}$ are the cross-sectional area and transverse wavevector for the ring waveguide, $\beta = \pi/d$ is the longitudinal wavevector for the Ramsey cavity, and $\alpha_m = 4/3\, r^3$ is the magnetic polarisability of the coupling aperture of radius $r$. The loaded cavity Q-factor $Q_c$ is then given by

$$\frac{1}{Q_c} = \frac{1}{Q_0} + \frac{1}{Q_{\mathrm{ex}}}. \tag{3.4}$$

Both models are shown in figure 3.1. The remainder of this section will discuss the experimental data points: some agree well with the simulations, and others agree better with the analytical model.

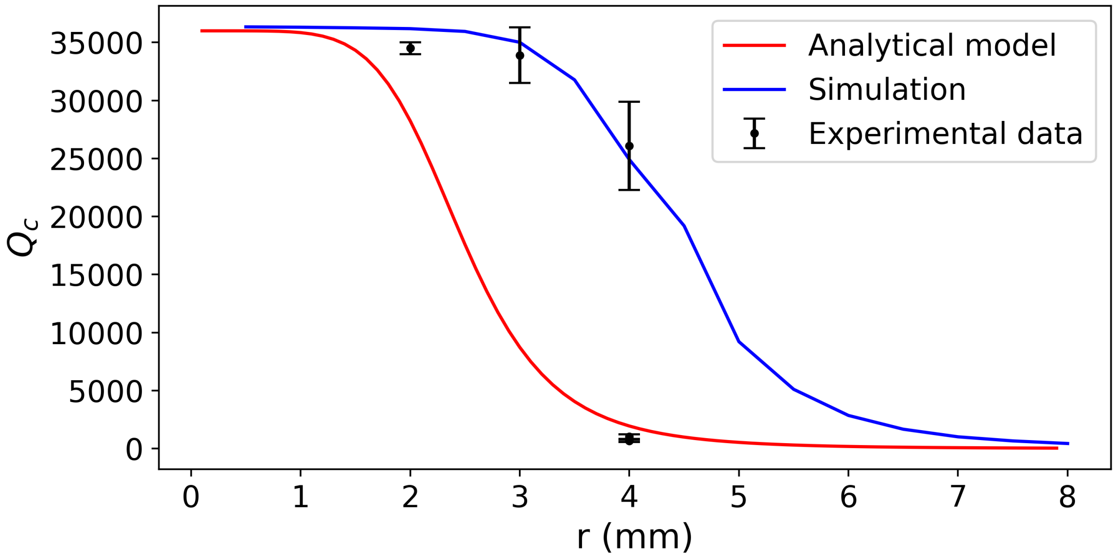

Figure 3.1: Analytical and finite-element simulations for $Q_c$ vs $r$, and experimental data. Multiple points at $r = 4\,\mathrm{mm}$ correspond to different waveguide detunings.

Initially, the coupling holes were 4 mm in diameter, and the $S_{11}$ vs frequency is shown in figure 3.2.

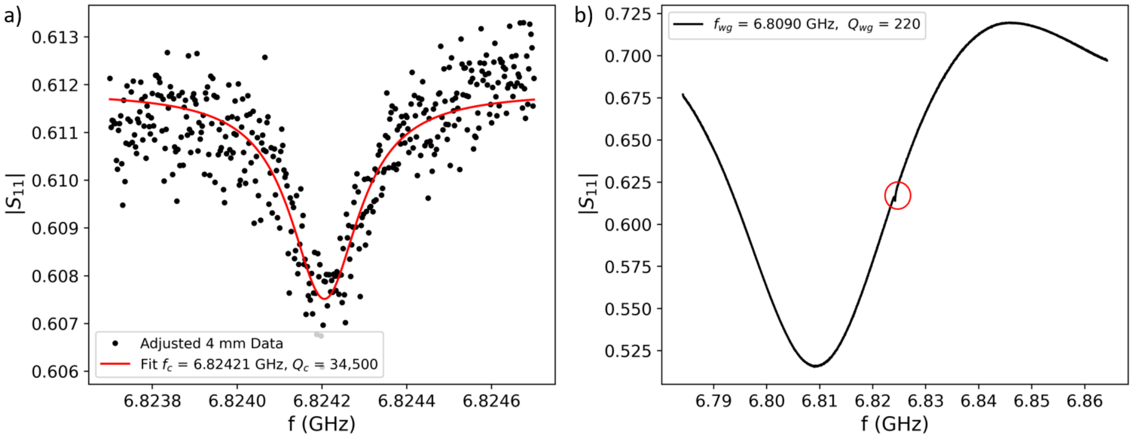

Figure 3.2: a) $|S_{11}|$ vs $f$ for the Ramsey cavity with 4 mm diameter coupling apertures, after compensation for the slope of the waveguide background. b) $|S_{11}|$ vs $f$ for the feeding waveguide, with the cavity feature highlighted.

Initially, the cavity was 10.47 MHz too low, and the waveguide was 25.7 MHz too low. The $Q_{\rm wg}$ is much lower than the unloaded Q of a $\mathrm{TE}_{106}$ waveguide [121], which with our dimensions is $Q_{\rm wg0} = 8,700$. This means that the feed loop is strongly-coupling the feedline to the waveguide. The cavity feature is very weak, $|\Delta S_{11}| = 0.004$, owing to the small coupling between the cavity and waveguide. Thus, the aperture diameters were widened to 6 mm, and the new $|S_{11}|$ is shown in figure 3.3.

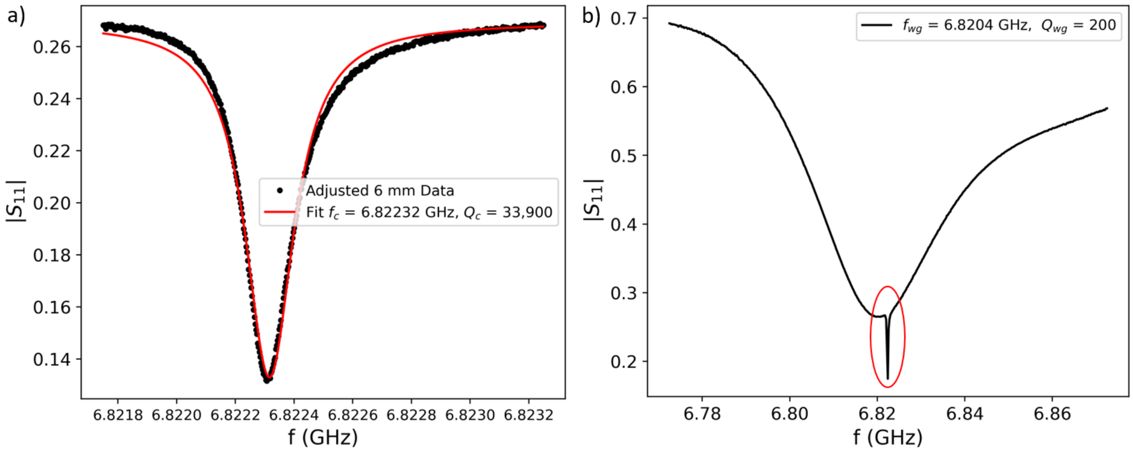

Figure 3.3: a) Ramsey cavity $|S_{11}|$ for 6 mm diameter coupling apertures. b) $|S_{11}|$ for the feeding waveguide.

The change in aperture size affected the resonant frequencies of the cavity and ring. The cavity frequency decreased by 1.89 MHz, as the larger coupling holes increase the effective volume of the cavity. The waveguide frequency increased by 11.6 MHz, possibly due to a slightly different size of the feed-loop, which must be fully removed to when making adjustments to the cavity apertures. While the size of the cavity feature had increased to $|\Delta S_{11}| = 0.136$, and the detuning between the cavity and waveguide decreased to 1.9 MHz, the cavity resonance was still very narrow, with $Q_c = 33{,}900 \pm 2{,}400$. Thus, even wider holes were deemed necessary.

The maximum allowable diameter of the coupling apertures is 8 mm. This is because at least 2 mm must be allowed around the edge for the epoxy to seal the sapphire viewports to the cavity. Thus, the apertures were widened, for the final time, to 8 mm. Additionally, since the waveguide frequency was still 14.3 MHz too low, a new tuning block that was total 1.2 mm longer, which may increase the waveguide frequency by 20.4 MHz, with the expectation that small amounts could be taken off later to reduce the waveguide frequency towards resonance. The results are shown in figure 3.4.

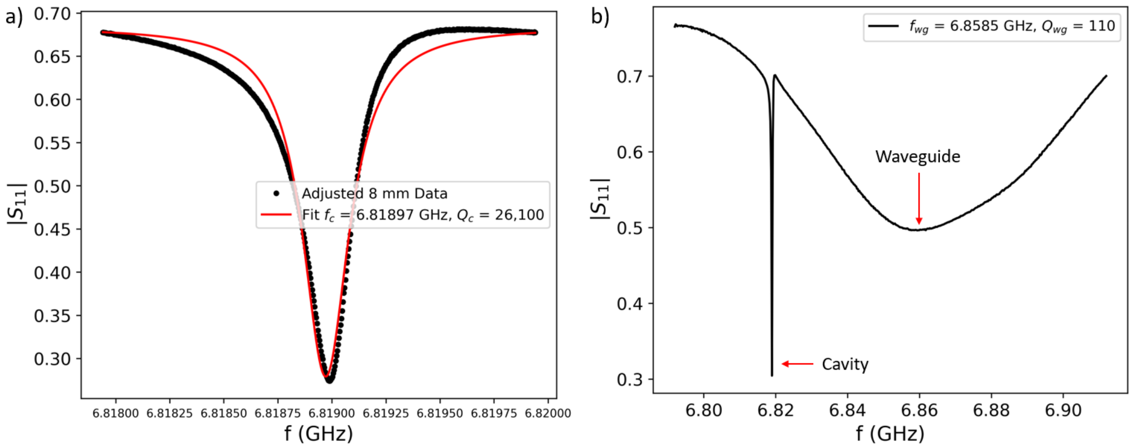

Figure 3.4: a) Close-up of the Ramsey cavity $|S_{11}|$ for 8 mm diameter coupling apertures. b) $|S_{11}|$ for a broader frequency range that covers the cavity and waveguide.

The cavity had $Q_c = 26{,}100 \pm 3{,}800$, which is quite high and in-keeping with the simulated $Q_c$ vs $r$ on figure 3.1. Its frequency decreased by 3.35 MHz, again because of the increased effective cavity volume. The waveguide did not behave as expected: the frequency increased by 38.1 MHz, almost twice as much as anticipated. Upon closer inspection, it seems there was excess length built into the new block, resulting in a shorter waveguide than expected. The introduction of the sapphire viewports and epoxy decreases the waveguide frequency, as the dielectric material increases the effective size of the cavity. As such, the excessively-high waveguide frequency was less pressing than the slightly-low cavity frequency. It was also expected, based on the model of the dielectric constant of air [202] and experience with NPL's other Cs fountain clock cavities, that the frequency will be at least 2 MHz higher when the system is under vacuum, such that the desired increase in cavity frequency was approximately 13 MHz, which corresponds to 0.45 mm machined off of the mating surface of the lower endcap to the cavity body. The result of this adjustment is shown in figure 3.5.

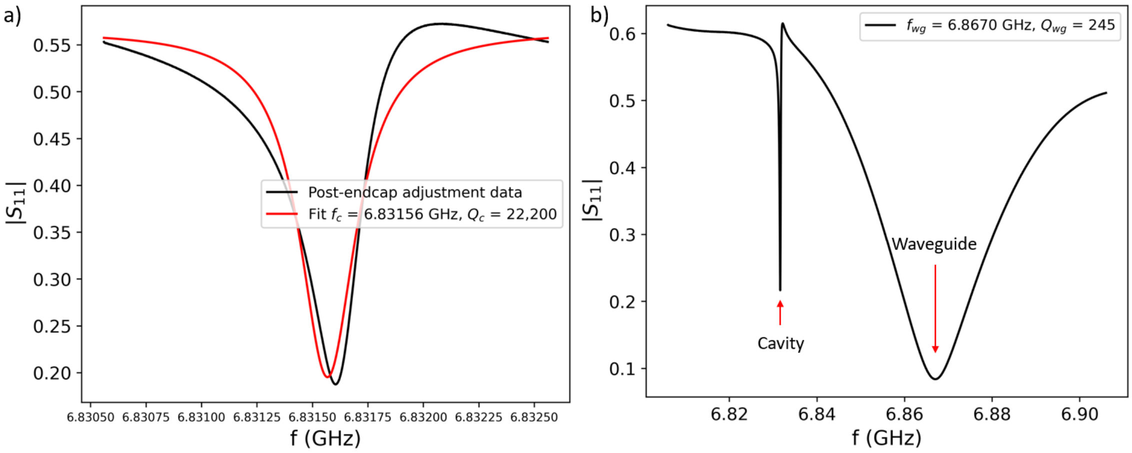

Figure 3.5: a) Close-up of the Ramsey cavity $|S_{11}|$ for 8 mm diameter coupling apertures after a 0.45 mm decrease in the cavity length $d$. b) $|S_{11}|$ for a broader frequency range that covers the cavity and waveguide.

The cavity frequency increased to 3.12 MHz below resonance. The waveguide shape became less dispersive, probably due to another slightly different loop feed position with the measured $Q_{\mathrm{wg}}$ more than doubling and the $f_{\mathrm{wg}}$ increasing by 8.5 MHz.

Since $f_c$, when accounting for the decrease due to the refractive index of air, is quite close to resonance, it is clear that only small further adjustments to the frequency, via the lower endcap, are necessary. Thus, the next step was to electron-beam weld (EBW) the upper endcap to the cavity body. The resulting weld is shown in figure 3.6, and the subsequent frequency measurements are shown in figure 3.7.

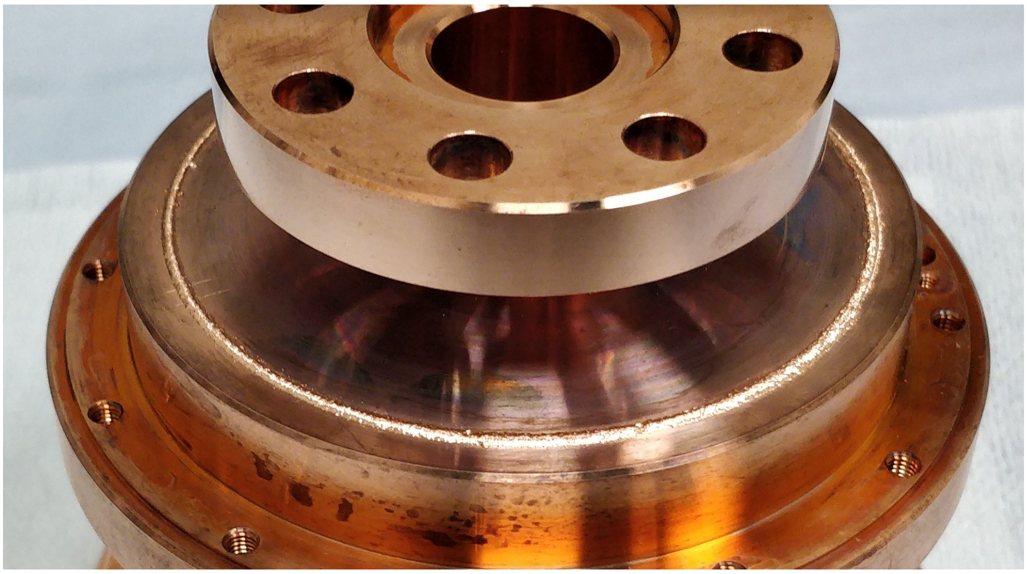

Figure 3.6: Result of electron-beam welding the upper endcap to the cavity body.

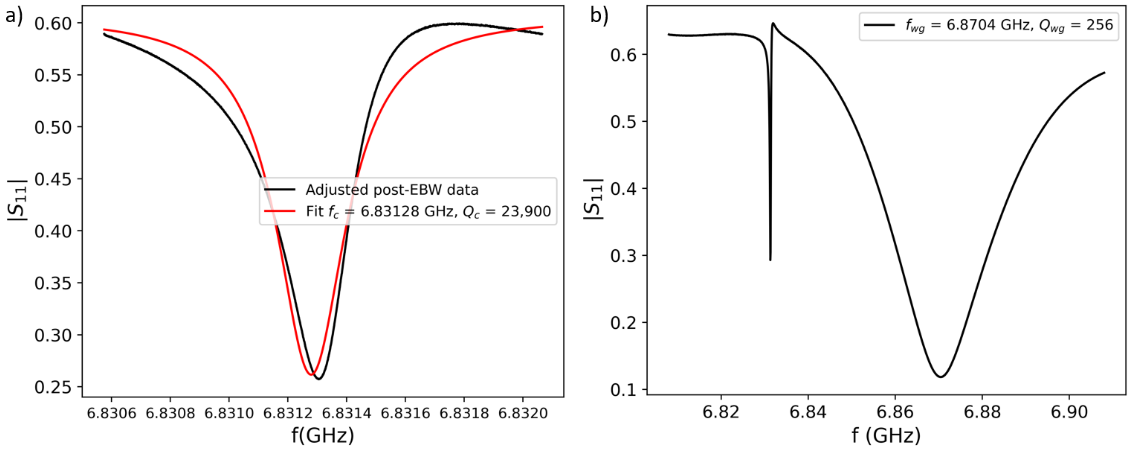

Figure 3.7: Post-electron-beam welding results of $|S_{11}|$ vs $f$ for a) Ramsey cavity and b) Ramsey cavity and waveguide.

Figure 3.5 and figure 3.7 look very similar, as expected. The cavity frequency appeared to shift downwards by only 280 kHz, with $Q_c$ being similar to before. The next step is the addition of the epoxy-sealed sapphire windows, which were expected, according to simulations, to reduce the frequency of the waveguide by around 7 MHz [194]. They were done one at a time, as several days were allowed for the epoxy to cure at room temperature under the weight of an aluminium lens post. The cured viewports are shown in figure 3.8, and the frequency measurements in figure 3.9.

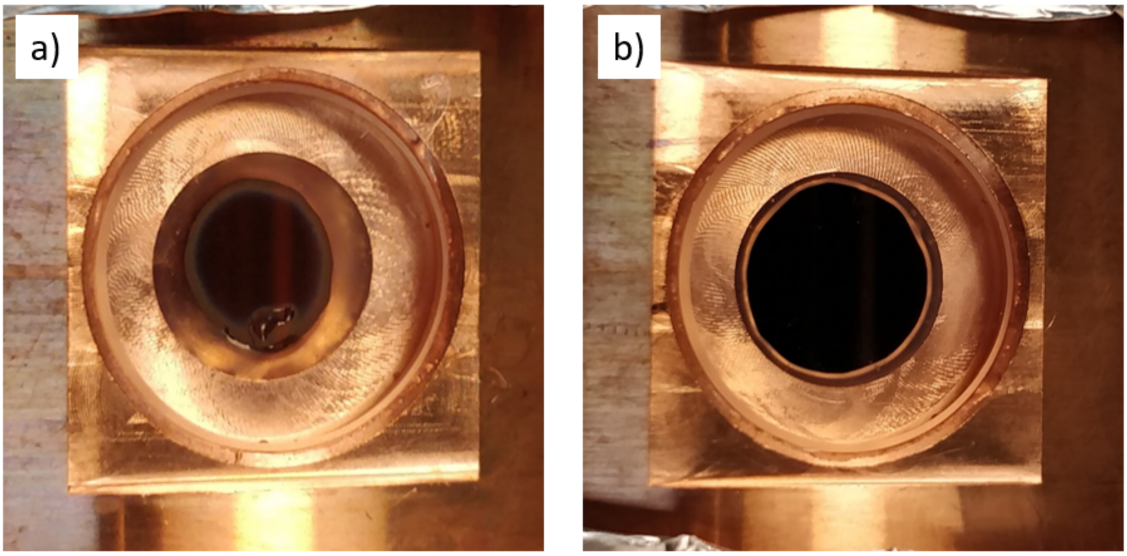

Figure 3.8: Epoxy sealing of 1/2" sapphire viewports to the cavity body. a) First attempt, where clearly too much epoxy was used, resulting in a large amount of excess obstructing the edge of the aperture. b) second attempt, where less epoxy was used to achieve a much cleaner result.

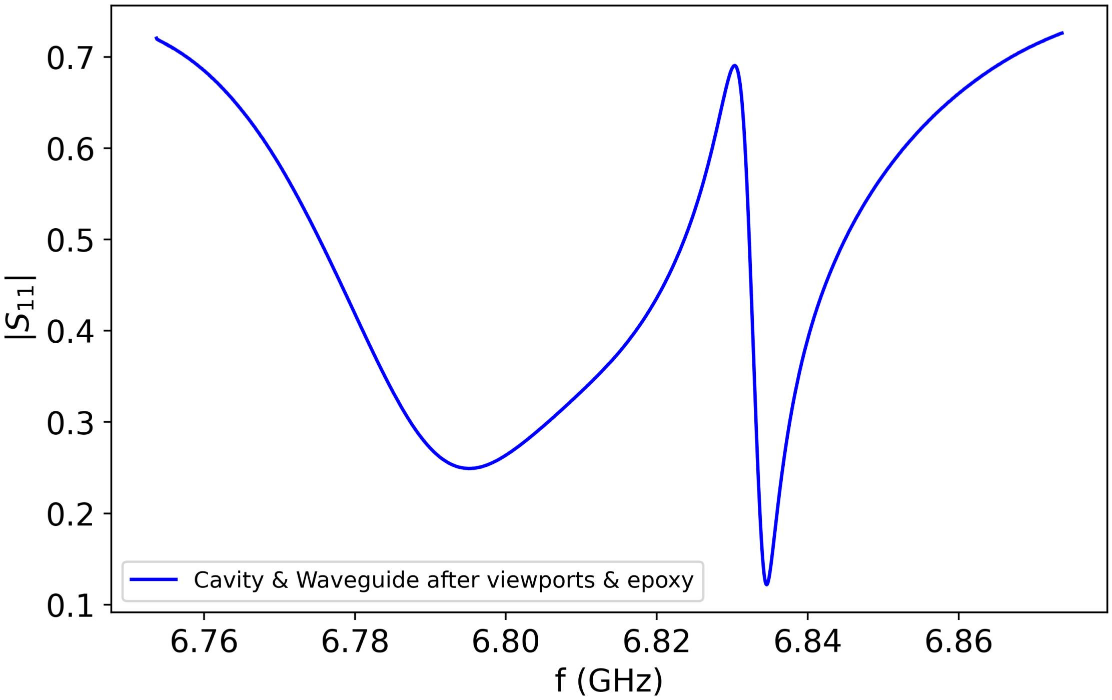

Figure 3.9: $|S_{11}|$ vs $f$ for the cavity and waveguide after sealing sapphire viewports with epoxy.

Determining exactly the width of the cavity and waveguide features is more difficult from figure 3.9, because the features are overlapping each other in a non-trivial way, due to the phase changes that occur in driven harmonic oscillators across resonance. The cavity and waveguide were much closer together in frequency in figure 3.3, and yet the cavity feature appeared as a clear superposition on top of the waveguide feature. This, along with the width of the feature near 6.83 GHz in figure 3.9, implies that the cavity linewidth is much wider, and thus $Q_c$ much lower.

The waveguide did decrease in frequency due to the addition of the viewports and epoxy, but by 65 MHz, almost 10 times as much as expected from simulations [194]! One potential explanation might be that the simulations accounted only for the relative permittivity of the sapphire, which is approximately $\epsilon_r = 10.8$, and not the epoxy itself, which the manufacturer states has $\epsilon_r = 4.5$, albeit at 1 kHz. Perhaps the dielectric constant of the epoxy is much higher at microwave frequencies.

To get better information about the cavity resonance, an additional feed loop is placed in the atom-passage holes to make $S_{21}$ measurements. The loop is inserted enough to see a clear transmission feature, and then gradually retracted until no feature is discernable. This is done with the new and initial waveguide tuning blocks, to see how the cavity behaves at different waveguide detunings. The resulting $f_c$ and $Q_c$ measurements are shown in figure 3.10. A loop depth of $d_L = 0\,\mathrm{mm}$ corresponds to the lower edge of the loop coincident with the upper surface of the cavity, $d_L > 0$ means the loop is partial or fully inside the cavity, and $d_L < 0$ means the loop is not protruding beyond the upper cut-off waveguide.

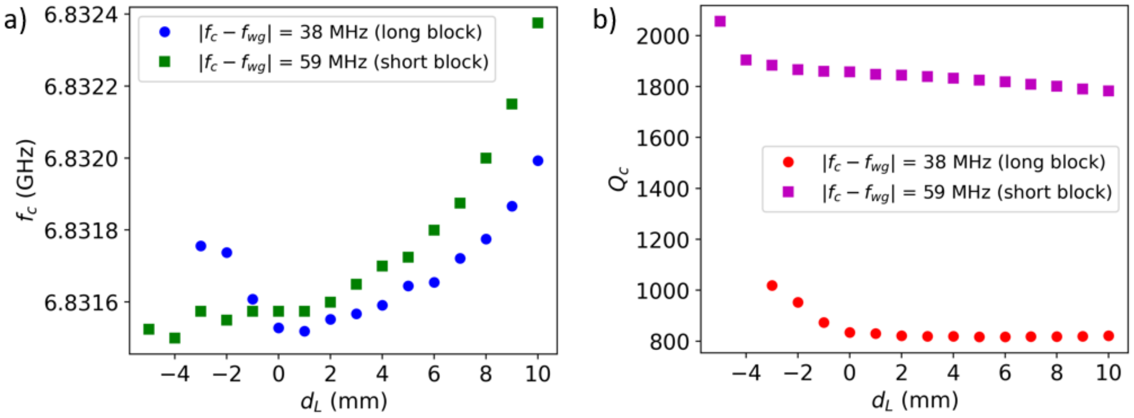

Figure 3.10: a) $f_c$ vs loop depth from $|S_{21}|$ measurements, for each waveguide detuning. b) $Q_c$ vs loop depth for the same $|S_{21}|$ datasets.

From figure 3.10a, one can see that in both cases $f_c$ initially decreased as the loop was removed, and were approximately coincident at $d_L = 0\,\mathrm{mm}$, where $f_c = 6.83152\,\mathrm{GHz}$. For $d_L < 0$, the small waveguide detuning began to show an increase in $f_c$, while the larger waveguide detuning continued to decrease. One would expect the lowest $d_L$ to result in the smallest perturbation due to the extra loop, and thus the most accurate measurement, which was $f_c = 6.83157\,\mathrm{GHz}$ at $d_L = -5\,\mathrm{mm}$ with the further-detuned waveguide, which is close to the $f_c$ value obtained before the viewports were epoxied on.

From figure 3.10b, one can see that a different waveguide detuning results in a different $Q_c$. This makes sense, because at $Q_c < Q_0$, the losses should be dominated by the coupling aperture to the waveguide, and thus a smaller detuning between the cavity and waveguide should couple them more strongly, with a correspondingly lower $Q_c$. However, both waveguide detunings resulted in much lower $Q_c$ values than for similar detunings before the epoxied viewports. For example in figure 3.7, $f_c - f_{\mathrm{wg}} = -39\,\mathrm{MHz}$ is large, so $Q_c = 23{,}900$ is also large. This suggests some broadening occurred specifically due to the insertion of the epoxied viewports. At the smaller detuning of 38 MHz, the cavity Q-factor stayed steady at around $Q_c = 820$ while $d_L > 0$ and rose once $d_L < 0$ until a maximum of $Q_c = 1,020$ was visible at $d_L = -3\,\mathrm{mm}$. This implies a FWHM of $\Gamma_c = 6.7\,\mathrm{MHz}$, which is larger than the detuning of the cavity from the clock transition.

The next aspect to address is the large waveguide detuning. While an even longer tuning block could be made to decrease the waveguide length and increase the frequency closer to resonance, an easier way may be to adjust the shape of the loop antenna feeding the waveguide. The feed loop consists of a stripped semi-rigid co-axial cable where the silver-plated copper core is bent into an “L” shape such that the end sits inside a small indent in, and in electrical contact with, the waveguide wall of the cavity body, as shown in figure 3.11a. Because $Q_{\mathrm{wg}} \ll Q_{\mathrm{wg0}}$, it is clear that the waveguide losses are dominated by losses induced in the feed loop. Much as the cavity mode can radiate back out into the waveguide through the coupling apertures, energy in the waveguide mode is lost when it induces current in the feed loop. The resonant frequency of a damped, driven oscillator is always less than the natural frequency:

$$f_{\mathrm{wg}} = \sqrt{f_{\mathrm{wg0}}^2 - \frac{8}{Q_{\mathrm{wg}}^2}}, \tag{3.5}$$

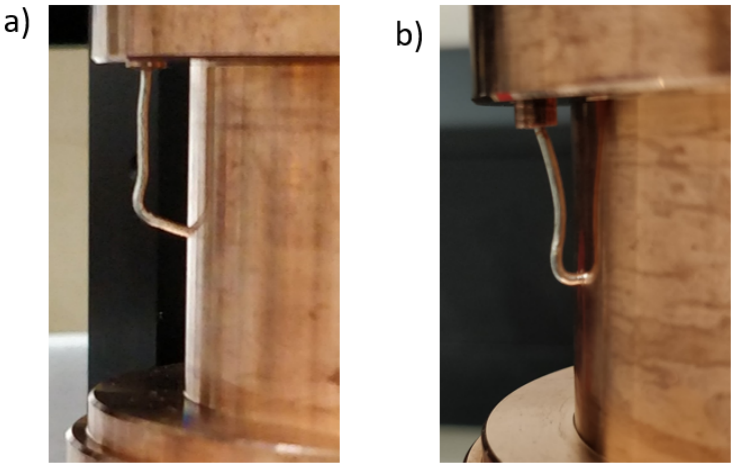

Figure 3.11: a) Initial and b) adjusted waveguide feed loops.

such that a higher $Q_{\mathrm{wg}}$ corresponds to a higher $f_{\mathrm{wg}}$. To try and increase $Q_{\mathrm{wg}}$, the size of the feed loop was reduced by bending the long straight section closer towards the wall, as shown in figure 3.11b. The resulting $|S_{11}|$ measurements are shown in figure 3.12.

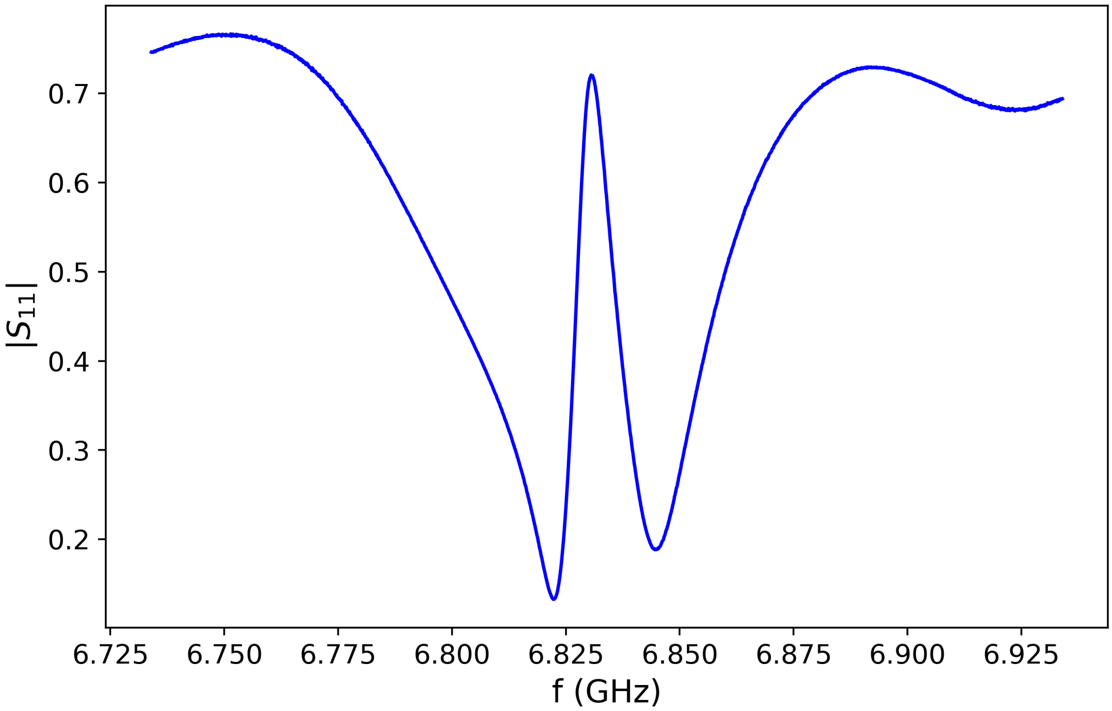

Figure 3.12: $|S_{11}|$ vs $f$ after decreasing the area of the waveguide feed loop.

As expected, the waveguide frequency has increased such that it is now very close to coincident with that of the cavity. It is difficult to tell exactly what $f_{\text{wg}}$ is from figure 3.12 because of the relative phase changes that take place across the cavity and waveguide resonances: the cavity feature is now in the opposite direction to the waveguide feature. This is also the case on NPL's Cs fountain clock cavities. By removing the data that correspond to the cavity, a Lorentzian can be fit to the remaining data that has $f_{\text{wg}} = 6.8311\,\text{GHz}$ and $Q_{\text{wg}} = 155$. The waveguide will remain accessible even once the fountain is under vacuum, so finer adjustments can be made after bake-out if necessary. However, the waveguide is only 0.1 linewidths away from the atomic resonance, which is good enough.

The cavity peak occurs at $f_c = 6.83055\,\text{GHz}$. To get more information about the new cavity resonance, the additional feed loop was inserted through the atom-passage holes and $|S_{21}|$ measurements were made at various $d_L$, and the results are shown in figure 3.13. The frequency results are similar to those already obtained by the $S_{11}$ data. The $Q_c$ results are quite different to the previous measurements: the waveguide is very close to resonance with the cavity, so stronger coupling is expected. At $d_L = -1\,\text{mm}$, $Q_c = 670$, which corresponds to a cavity linewidth of 10.2 MHz. Thus, the cavity being 4.13 MHz below resonance with the clock transition is not a concern, particularly because the cavity resonant frequency should further increase by 2 MHz when under vacuum. A summary of all the cavity tuning steps described in this section is shown in table 3.1.

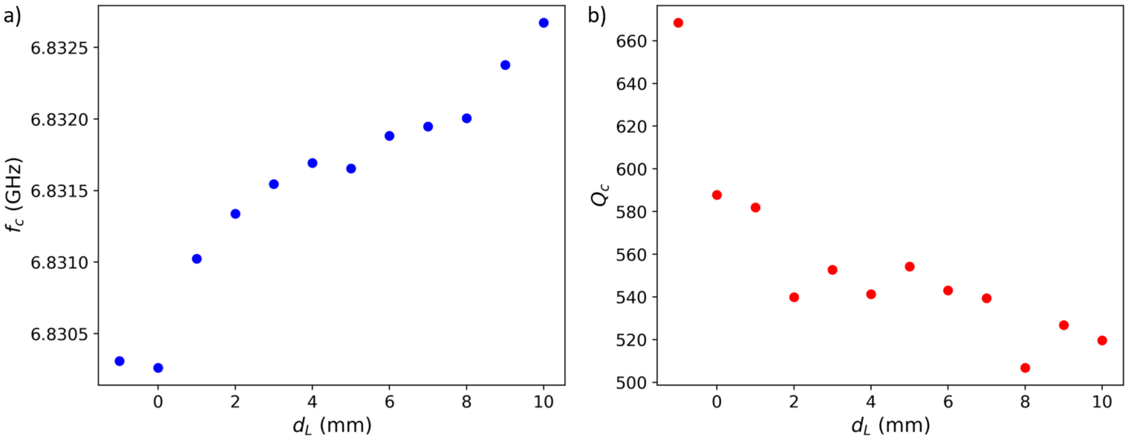

Figure 3.13: a) Cavity frequency vs loop depth for $|S_{21}|$ data. b) Cavity quality-factor vs loop depth for $|S_{21}|$ data.

| Stage | $f_c$ (MHz) | $Q_c$ | $f_{\mathrm{wg}}$ (MHz) | $Q_{\mathrm{wg}}$ |
|---|---|---|---|---|
| $r = 4$ mm | 6824.22 | 34,500 | 6809.0 | 220 |
| $r = 6$ mm | 6822.32 | 33,900 | 6820.4 | 200 |
| $r = 8$ mm, longer waveguide block | 6818.97 | 26,100 | 6858.5 | 110 |
| Tune lower endcap | 6831.56 | 22,200 | 6867.0 | 245 |
| After EBW | 6831.28 | 23,900 | 6870.4 | 256 |
| Epoxied viewports | 6831.57 | $> 1020$ | 6795.0 | 145 |
| Adjust feed loop shape | 6830.55 | $> 670$ | 6831.1 | 155 |

Table 3.1: Summary of the steps taken to tune the cavity and waveguide.

The $Q_c$ values obtained from figures 3.10 and 3.13 imply that, in the regime of strong cavity-waveguide coupling, a smaller cavity-waveguide detuning results in a stronger coupling, and thus a lower $Q_c$. The discrepancy between the analytical model and simulations of figure 3.1 are likely due to the fact that some finite detuning was inherent to the simulated geometry (up to about 6 MHz), whereas the analytic calculations assumed from the beginning that $f_c = f_{\mathrm{wg}}$. Thus, neither approach was "wrong" as such: the simulations were valid when $|f_c - f_{\mathrm{wg}}| > f_c/Q_c$, as in the initial points for small coupling holes, and the analytical approach was broadly correct when $|f_c - f_{\mathrm{wg}}| < f_c/Q_c$, as for the final case.

However, $Q_c$ was only really increased by the addition of the epoxy and viewports, which implies that there is some energy loss from the cavity mode to the dielectric. This may become clearer in the future when the DCP shift is analysed; if the losses are uneven, perhaps because of the uneven amount of epoxy on each viewport, then a large imbalance in the power in (and out) through each viewport will amplify the DCP effect.

## 3.2 Vacuum Preparation & Sealing

Ultra-high vacuum (UHV, $< 10^{-9}\,\mathrm{mbar}$) is necessary to allow the launched atom cloud to return and be detected, as collisions with background thermal gases can degrade the number of cold atoms present at every stage of fountain operation. At these pressures, the gas composition inside the system bears little resemblance to air. Instead, the dominant species are hydrogen gas, which resides in the bulk of the metal and diffuses outwards into the vacuum, and water vapour, which is sorbed onto the surface of metals and desorbs under low pressure.

### Cleaning

Achieving UHV requires stringent processes to be followed: inclusion of a material with a high vapour pressure, either as part of the vacuum system or as an unwanted contaminant, is unacceptable. Potential contaminants include lubricating grease present on metal parts after machining, and organic grease that can be accidentally transferred during handling and construction. Similarly, any faulty seal that permits the passage of air will result in a higher equilibrium pressure.

The first step to attaining UHV is a thorough clean of all parts of the vacuum system. There are many ways in which this can be done, which depend on the required vacuum level and the desired function of the vacuum vessel. Many cleaning techniques have been developed specifically with particle accelerators in mind [203, 204], as the emission of high energy photons (Bremsstrahlung radiation) during operation can induce outgassing by the vacuum vessel, and increase the rate of collisions between the particles of the beam and the background gas, which is undesirable.

For atomic fountain clocks at NPL, in general a four-step process is used. Firstly, the parts are rinsed in acetone to remove grease and hydrocarbons. Next, the parts are immersed in a 30°C solution of neutral, organic detergent and subjected to ultrasonic agitation at 37 kHz for 30 minutes. Afterwards, the remaining detergent is rinsed off with deionised water. The parts are left to dry overnight in a laminar flow cabinet, which avoids significant re-contamination with dust.

For the full-sized primary frequency standards, this procedure only applied to parts small enough to fit in the ultrasonic bath, of length around 40 cm and width and depth around 20 cm. The larger parts, such as the aluminium flight tube and detection chamber, are usually cleaned by a particularly zealous rinse in acetone only, which has been shown to produce vacuums of perfectly acceptable quality, with pressures in the low $10^{-11}$ mbar range achieved in the latest generation of NPL's Cs fountain clocks. All parts of the mini-fountain vacuum vessel fit inside the ultrasonic bath, thus the full process is applied to all parts.

One issue that was identified was that the mini-fountain parts are made predominantly of aluminium and copper, which are much softer metals than the stainless steel used throughout much of the full-sized fountains. It was discovered that 30 minutes of ultrasonic agitation is not appropriate for these parts. This is because the bubbles that form as part of the ultrasonic cleaning process [205] can also cause erosion on the metal surface [206]. On aluminium, this results in dull, white patches on the surface, which are not acceptable on the indium sealing surfaces. To avoid this, the ultrasonic cleaning time for all parts was reduced to 5 minutes, and the parts were re-oriented in the bath every 60 seconds, to avoid any particular part of the surface being exposed to antinodes of the standing wave pattern for too long.

The only parts which are not cleaned before assembly are the graphite cylinders, which are placed just above and below the Ramsey cavity's cut-off waveguides to keep the partial pressure of thermal rubidium down in the interrogation region (graphite can absorb large quantities of alkali metals [207, 208]), the ion-NEG combination pump, and the rubidium dispensers. The graphite is instead "cleaned" by thoroughly purging a vacuum furnace with inert gas to reduce the oxygen content in the graphite, before undergoing ex-situ baking at 450°C under vacuum for 4 hours. The SAES ion-NEG pump is assumed to be UHV-clean from the manufacturer, and it is shipped under vacuum. The rubidium dispensers are also assumed clean, and are shipped sealed in an inert atmosphere.

## Assembly

Once all parts are clean and dry, assembly can begin. This primarily involves pressed-indium seals, as the usual knife-edge and copper gasket method employed in many steel vacuum systems is not appropriate here, due to the softness of the aluminium and copper parts, and the fragility of glass viewports.

To seal two pieces together with an indium wire seal, both surfaces and the wire must be clean and dust free. The first part of the system to seal is the MOT chamber, to six glass viewports and two glass lenses. Firstly, the cleanliness of the aluminium surface is ensured by blowing dust off with compressed air, and soaking some non-

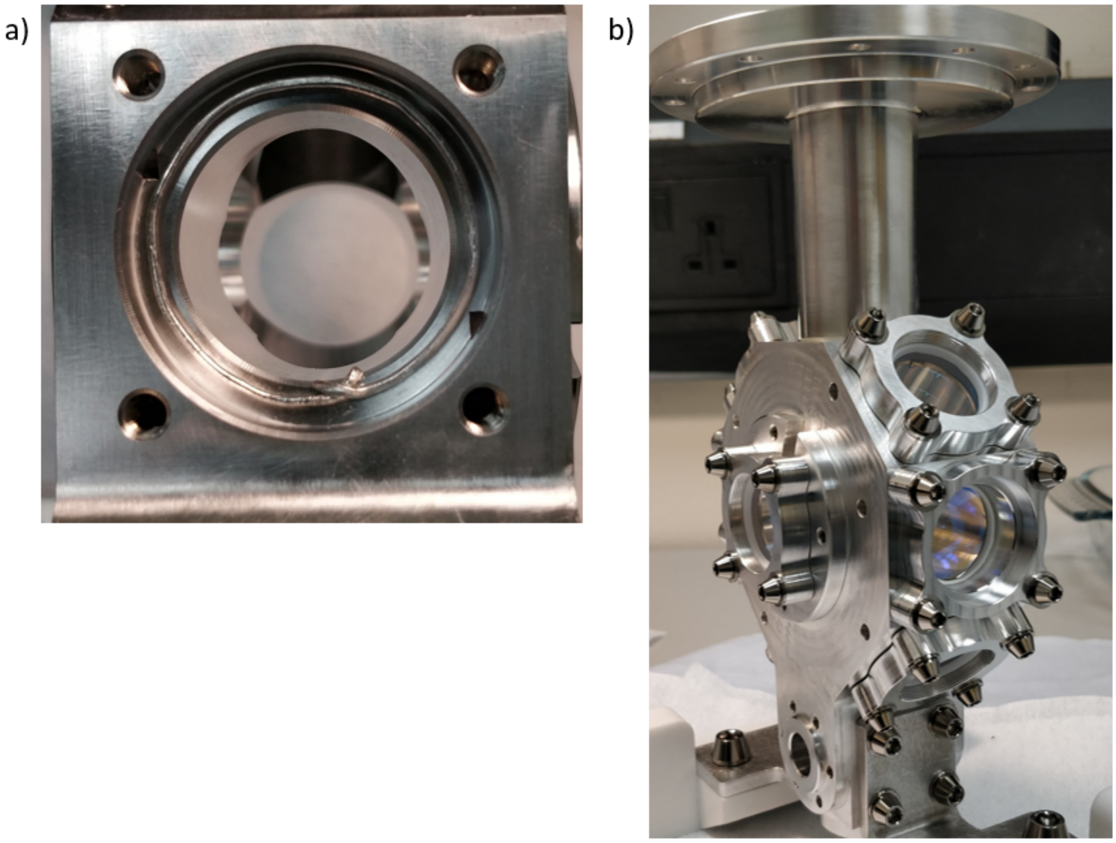

Figure 3.14: a) Close-up of indium wire used to seal glass viewport to the MOT chamber. b) Resulting MOT chamber with all viewports and lenses sealed

lint-producing cloth in acetone and gently wiping the surface down, followed by a second compressed air blow. Next, the 1 mm diameter indium wire is cleaned by rinsing it in acetone, taking care to avoid any impacts or stresses that may cause undesired deformation to the shape of the wire. Finally, the glass surfaces, which are assumed to be optically clean from the manufacturer, are removed from their protective packaging and simply given an air blow. Next, the indium wire is placed on the flat, smooth sealing surface of the chamber. An example is shown in figure 3.14a, with the ends crossing over and overlapping by several millimetres. Next, the glass viewport is lowered onto the wire, taking care not to move the wire in the process. Once in position, the glass can be given a firm manual press to slightly squash the indium and set its position. Next, plastic spacers and aluminium clamps are positioned over the glass, such that four screws are used to squash the indium wire. The resulting MOT chamber is shown in figure 3.14b.

The next pieces to connect to the chamber are the SAES alkali metal dispensers, which each contain 6.4 mg of rubidium, and the copper pinch tube, which will be used to connect the mini-fountain to the turbopump and seal the vacuum shut after bake-out. These are commercial-off-the-shelf parts with stainless steel DN16 flanges, which have been modified to remove the knife edge to allow pressed-indium sealing to the aluminium MOT chamber.

A DC current of the order of a few Amps will be passed through the dispensers, to heat them to several hundred degrees and cause the emission of rubidium vapour. Each dispenser assembly requires an input and return path for the heating current. Such a current may produce a magnetic field in the MOT region of a few microtesla, which is not expected to compromise cooling. The magnetic shields surrounding the MOT should ensure that this current also does not compromise clock operation. Initially, it was thought that the dispensers could be connected to the power-feedthroughs by beryllium-copper barrel connectors, so long as the magnetic steel screws are swapped for titanium screws of the same size. Unfortunately, the combination of barrel connector and screw head did not allow to fit two dispensers in the lower part of the chamber. Instead, two dispensers were spot-welded together, and the resulting two free ends were spot-welded to the copper conductors of the power feedthroughs. Two such arrangements were made, such that a total of 4 rubidium dispensers are present in the system, for a total rubidium capacity of 25.6 mg. The same indium sealing method as for the glass viewports was used to seal the stainless steel power feedthroughs and modified end of the copper pinch-tube to the aluminium chamber, and afterwards a multimeter was used to verify that the dispenser assemblies were still connected, and not short-circuiting to each other or the chamber walls.

Next, the lower graphite gettering sleeve is inserted into the upper neck of the MOT chamber, and the (tuned) Ramsey cavity can be sealed to the upper flange of the chamber. The pressing is done by 12 titanium screws that pass through the copper cavity body and into two threaded semi-circular rings, also made of titanium. This is because OFHC copper is even softer than aluminium, and the risk of stripping threads in the Ramsey cavity is best avoided wherever possible. This is shown in figure 3.15.

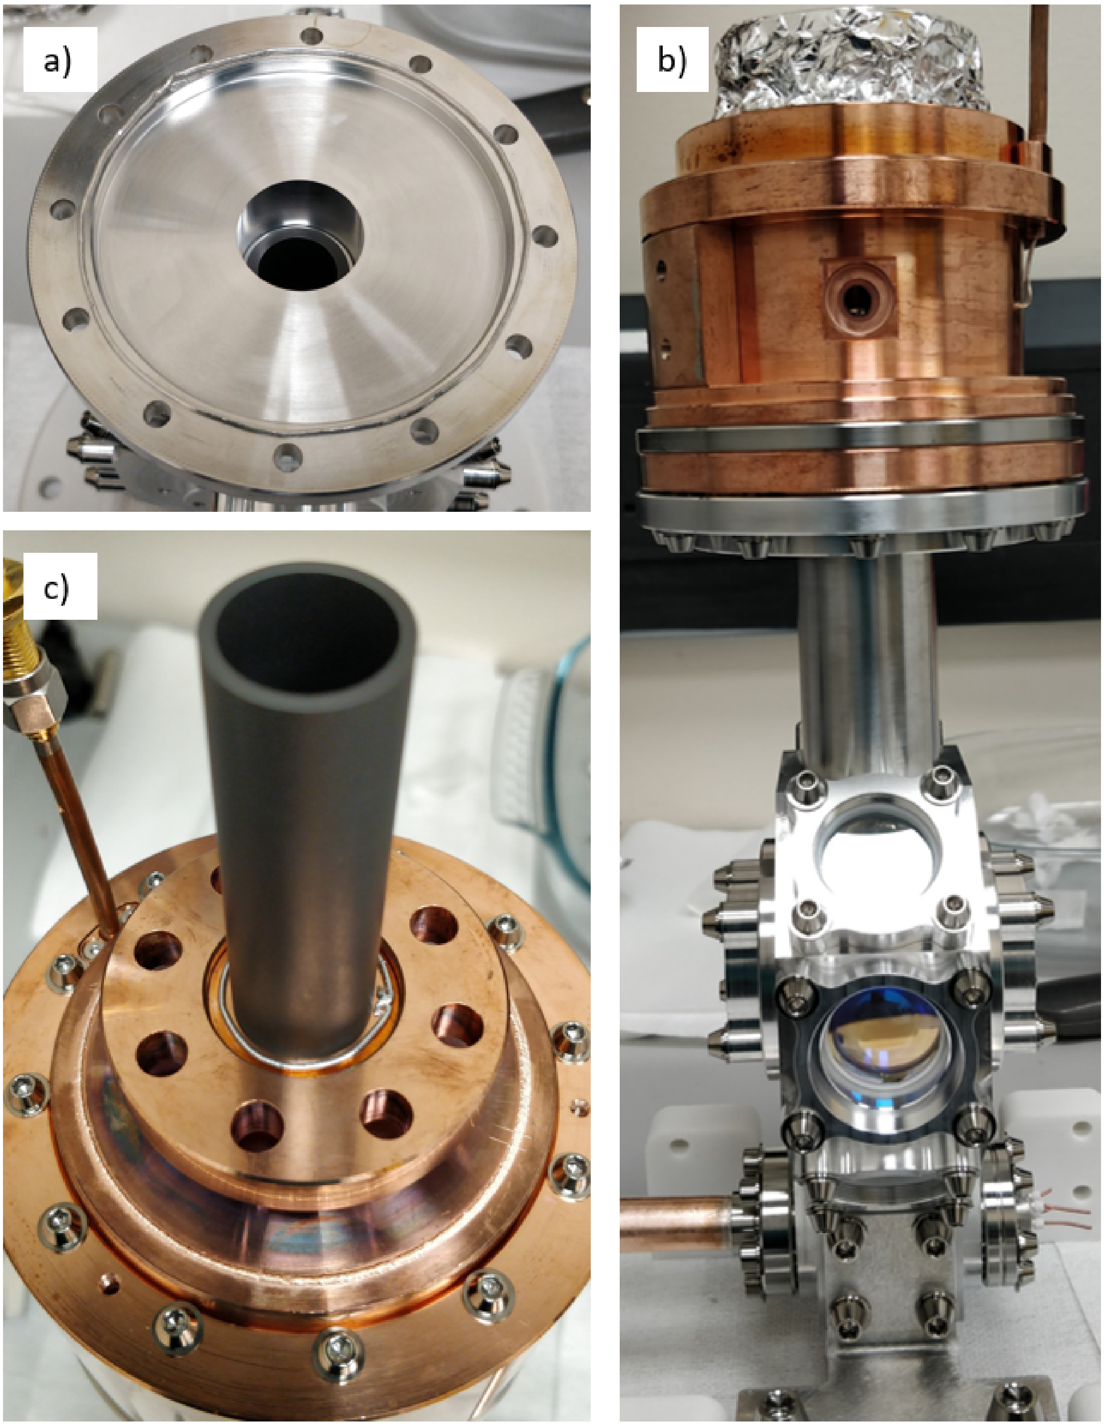

Figure 3.15: a) Lower graphite getter inserted into the neck of the MOT chamber to form the lower waveguide-below-cut-off, and indium wire placed around the edge. b) Ramsey cavity sealed to the MOT chamber. c) Upper graphite getter on top of the cavity, along with the indium wire for sealing to the flight tube.

The lower graphite getter is held in position once the cavity is sealed, however the upper graphite getter is not; it is simply placed on top of the cavity and the aluminium flight tube goes around it. The inside of the 33 cm long flight tube was manufactured by wire-EDM (electrical discharge machining). This is shown in figure 3.15c.

As with the chamber-cavity seal, the screws sealing the flight tube to the cavity by a pressed-indium seal screw into titanium nuts, not the copper flange of the cavity. Once the flight tube is sealed to the top of the cavity, the last pieces to connect are the ceramic break, where one end is a flat steel flange and the other is a standard DN40 knife-edge, and the ion-NEG combination pump. The lower, flat surface of the ceramic break is pressed-indium sealed to the top of the flight tube, and the upper surface is sealed by a DN40 knife edge seal and a copper gasket to the pump. Because the copper gasket requires more force to squash than the indium wire, the pump was attached to the ceramic break first, and then the ceramic break/pump assembly was then indium-pressed onto the top of the flight tube. This is shown in figure 3.16.

With that, the fountain's vacuum system assembly was complete. The full mini-fountain vacuum system is shown in figure 3.17.

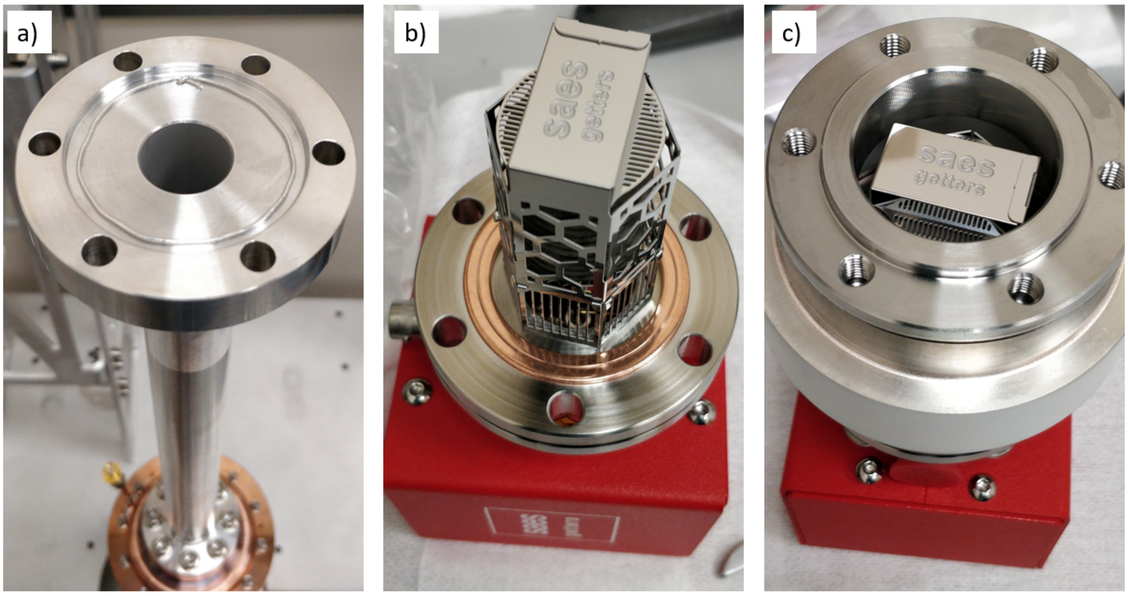

Figure 3.16: a) Indium wire for sealing the flight tube to the ceramic break. b) Ion-NEG combination pump. c) Ceramic break sealed around NEG cartridge.

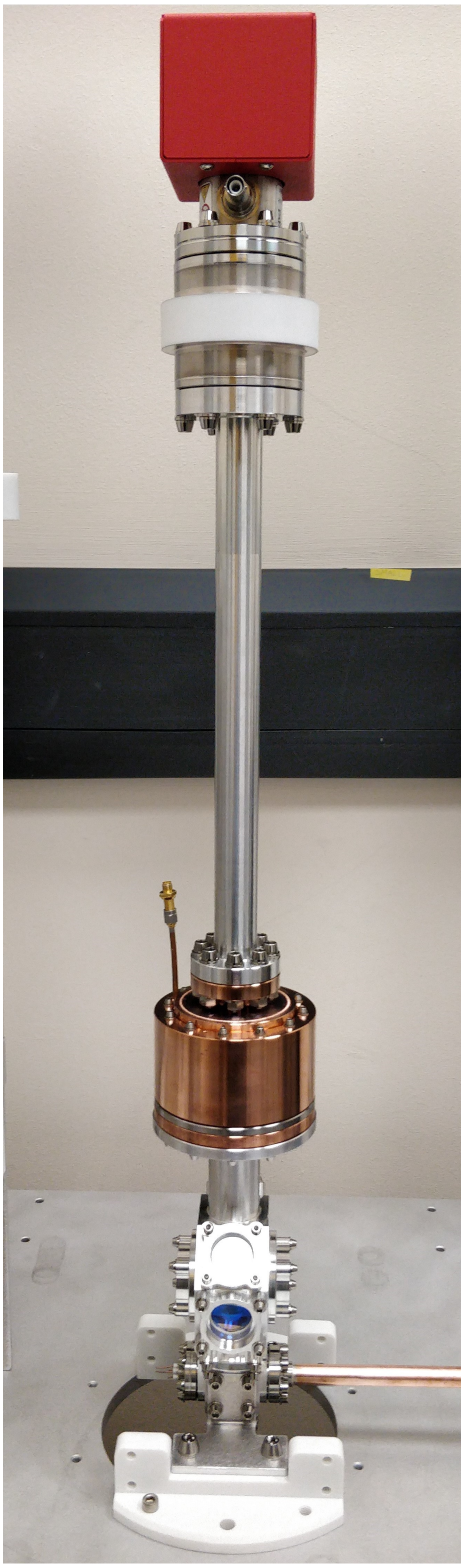

Figure 3.17: The full mini-fountain vacuum system.

## Leak Checking

The next step is to leak-check the vacuum system. This is done by connecting the other (DN16 flange) end of the copper pinch tube to steel bellows, which connects to a turbo-molecular pump, a residual gas analyser (RGA) and an ion gauge. This set-up is shown in figure 3.18. Once the turbopump has reduced the pressure to below $10^{-5}$ mbar or so, as read by the gauge, then the RGA can be used to gain information about the gas inside the system.

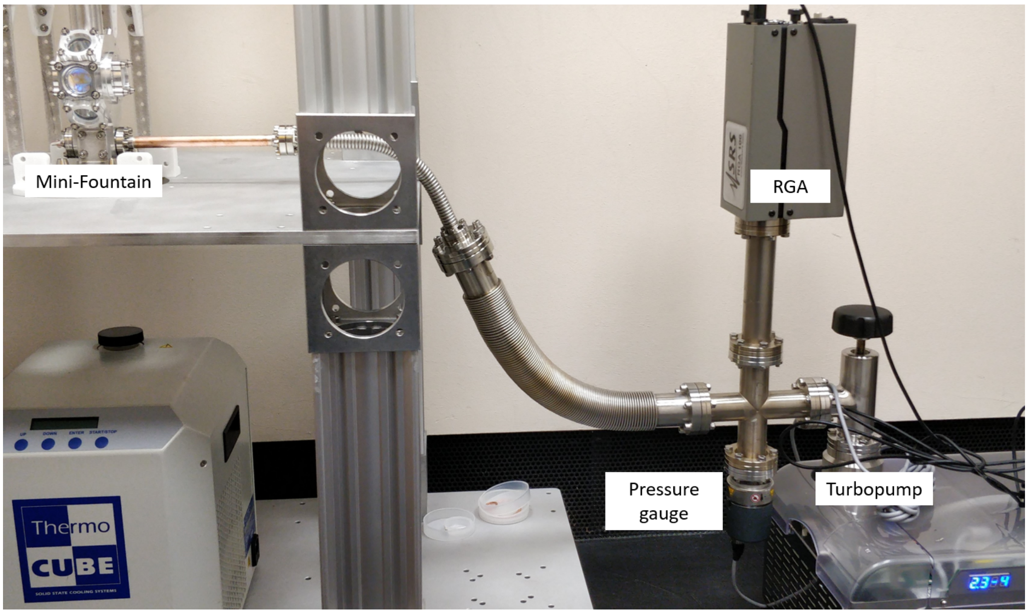

Figure 3.18: Equipment used to pump the mini-fountain, measure the pressure, and check for leaks. The right-hand side equipment is the same as in figure 2.23.

The best way to check for leaks is to set the RGA to look at gas species of mass $4\,m_p$, as only helium has this mass, and helium is not naturally present in the atmosphere. Thus, when helium is deliberately released near a leaky seal, a large spike in partial pressure is detected by the RGA.

After a couple of hours of pumping, the pressure was $3 \times 10^{-6}$ mbar, which was sufficiently low to turn on the RGA and perform a helium leak-check. The initial results were surprising; all of the aluminium-indium-glass seals (six MOT beams and two lenses), as well as both dispensers and the chamber-side of the copper pinch tube, were leaking. Further attempts at tightening the screws were not effective at reducing the leaks. The rest of the metal-indium-metal seals, and also the sapphire-epoxy-cavity seals, were seen to be leak-tight.

There were several possible reasons why the chamber-indium-glass seals were unsuccessful. The most obvious was surface finish: because of the complex shape of the MOT chamber, it was not possible to form every sealing surface by annular machining on a lathe, such that there were groves with radial components. These have been thought to enable leaks [209]. However, the copper to which the sapphire viewports are epoxied was also machined in this way, and these didn't leak. Another reason may have been insufficient radial clearance between the chamber and the diameters of the glass parts. Gradually squashing an indium wire by alternately tightening several screws does inherently require some additional clearance, as the glass must "see-saw" its way down, and thus its non-zero thickness presents a slight increase in its radial size when it is not flat. The viewports and lenses were removed from the chamber, and the diameters of the pockets that house the viewports were increased by 1 mm, and the diameters of the lenses were decreased by 0.6 mm. Then, the aluminium sealing surfaces were polished to a mirror finish, which has been shown to be an acceptable alternative to spiral-groove machining [209].

Upon closer inspection of the power feedthroughs and copper pinch tube, it was noticed that a weld seam of around 2 mm width and 1 mm height protruded from the steel surface. This seam is a result of the processes used to form the components in the first place, but they were not present on the CAD models provided by the manufacturers, and thus not considered in the design of the lower chamber ports. Measurements with a micrometer revealed that there was approximately 0.3 mm of overlap between these weld seams and the inside of the chamber sealing surfaces, which may explain why these seals were all leaking. Thus, the diameters of the through-holes on the lower chamber ports were increased by 1 mm.

Additional plastic spacers were used in the re-construction of the chamber, such that, at the appropriate compression, there was roughly a 1 mm gap between the chamber and the aluminium clamps, which was verified by the insertion of 1 mm thickness spacers between the chamber and clamps, to ensure consistency.

Once the fountain was re-assembled, pumping was restarted. This time no leaks were detected, even on the most sensitive RGA setting. The fact that the epoxy seals were leak tight in the presence of radial grooves, but the more tried-and-tested method of indium sealing can not tolerate such grooves, was very interesting, and suggests that perhaps epoxy sealing should be a more widespread vacuum technique than it currently is.

## 3.3 Bake-Out & NEG Activation

The next step in achieving UHV is bake-out: because the pressure is limited by outgassing from the walls, it is preferable to reduce this outgassing rate as far as possible. Heating the system increases the rate of outgassing, and this gas can be pumped away. Afterwards, there will be less gas sorbed inside the system and thus less outgassing.

Initially, the outgassing of a clean metal vacuum system is mostly water vapour from the metal surface. The outgassing rate is inversely proportional to the pumping time [210]. The proportionality coefficient depends on the temperature but does not vary much amongst UHV compatible metals [211]. Once the water has been pumped away, the dominant gas species is hydrogen. Atomic hydrogen gets trapped in liquid metal during its formation [210], particularly in steel, and slowly migrates until it reaches the surface. The outgassing rate doesn't vary with pumping time, but can be increased by increasing the temperature, such that afterwards, at a lower temperature, the hydrogen outgassing rate is reduced.

Hydrogen is not recharged in the bulk metal when exposed to normal atmospheric conditions. For this reason, it is common to perform an ex-situ bake of the individual parts of a vacuum system, typically at 350°C for a few hours, to permanently deplete the hydrogen content in the part. Unfortunately, NPL does not currently have a functioning vacuum furnace, so an ex-situ bake of every individual part was not possible. However, most of the mini-fountain is made of aluminium and copper, with only small amounts of steel. These metals, either due to the passive oxide layers that form on the surface, or because of different metallurgical processes, tend to have an order of magnitude lower hydrogen outgassing rates than a steel part of the same treatment [212]. Thus, the in-situ bake-out temperature of 100°C is expected to be sufficient to achieve UHV. The mini-fountain is baked for 4 weeks, as is the case with the full-sized Cs fountains at NPL.

To bake the mini-fountain, the support structure of a future full-sized fountain was assembled, as there exists a custom-made insulation tent for this. Two heaters are attached on one side of the structure and each consist of a pulsed heating element, fan, and internal thermocouple. They can be connected to another thermocouple, via their electrical control units, that is attached to the fountain, such that a feedback loop is formed for autonomous operation around a pre-programmed temperature.

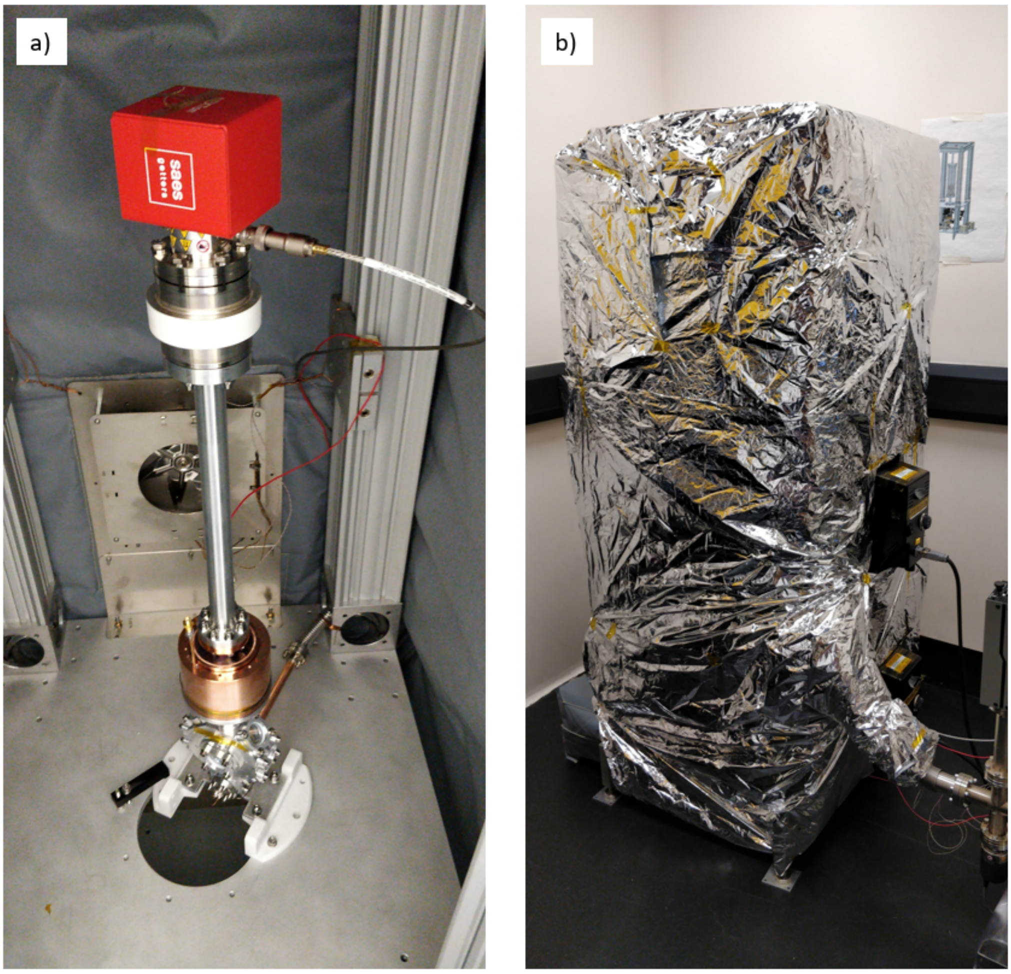

Figure 3.19: a) Inside of the bake-out tent, showing the fountain and upper heater. b) Outside of the tent, with additional mylar foil insulation.

Some additional thermocouples, without feedback to the heaters, are attached to the MOT chamber, cavity, and ion pump. This setup is shown in figure 3.19. The resulting pressure and temperature vs time for the bake-out is shown in figure 3.20.

Before heating, the pressure had been reduced to $7.7 \times 10^{-8}$ mbar. The heat was turned on gradually, to avoid accidentally overshooting the intended temperature. After around 48 hours, the physics package had reached the target temperature of 100°C, and the pressure continued to decrease. The fountain temperature oscillated between 102°C and 108°C for another 200 hours, until the heater settings were adjusted to reduce the oscillation amplitude. At 350 hours, the pressure had reduced

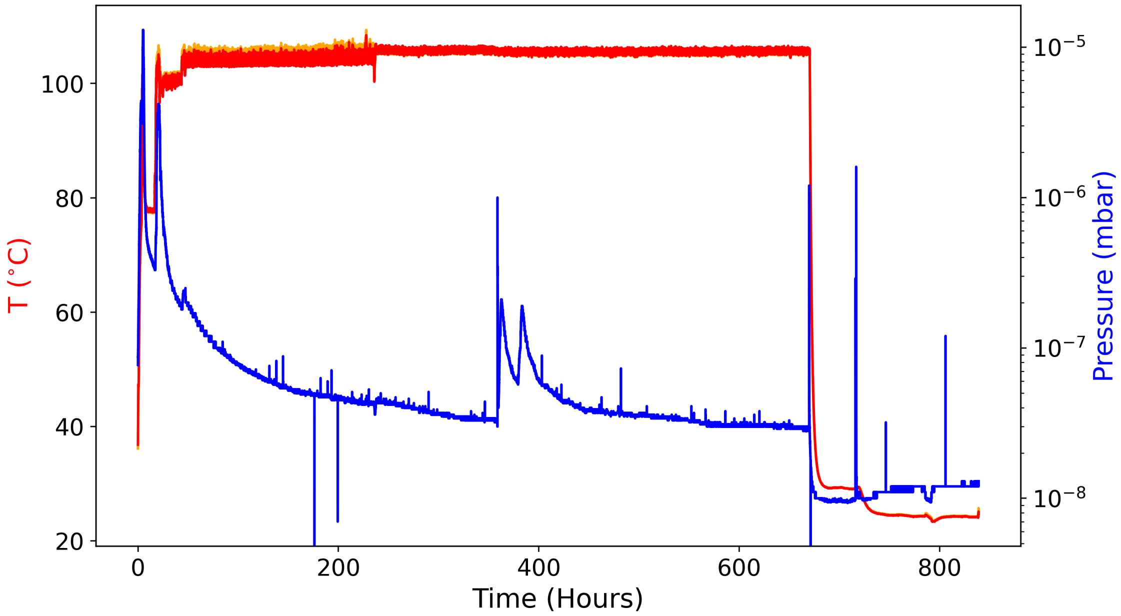

Figure 3.20: Temperature of the MOT chamber (red) and cavity (orange), and pressure reading of the ion gauge (blue) vs time for a 4-week bake-out.

to $3.3 \times 10^{-8}$ mbar. At this point, the RGA was activated to check the composition of the vacuum, which corresponds to the large, narrow spike in pressure at 350 hours. It was decided that the hydrogen outgassing of the unbaked steel part was dominating the pressure reading, so additional heaters and insulation were added to the steel bellows, and the two decaying peaks between 350 and 400 hours are a result of the temperature of the steel bellows increasing gradually to 105°C. This state persisted until 670 hours, when all heaters were turned off. The pressure dropped during cooling from $2.9 \times 10^{-8}$ mbar to $9.0 \times 10^{-9}$ mbar, where it continued to hover while further helium leak-checks were conducted.

The ion pump, as well as providing pumping to improve the vacuum level, acts as a pressure gauge by measuring the ionisation current. Unlike the NEG pump, the ion pump requires no particular process to activate, the high-voltage power supply simply needs to be connected. After about 15 minutes of reading 0 nA ($P_{\text{pump}} < 2.1 \times 10^{-11}$ mbar), the plasma discharge was ignited and the pressure reading initially exceeded $10^{-5}$ mbar, before slowly coming down. The ion pump was turned off and on a couple of times, which is recommended by the manufacturer before starting the NEG activation process, and the third time it was left on to pump down the pressure. The resulting pressure vs time is shown in figure 3.21.

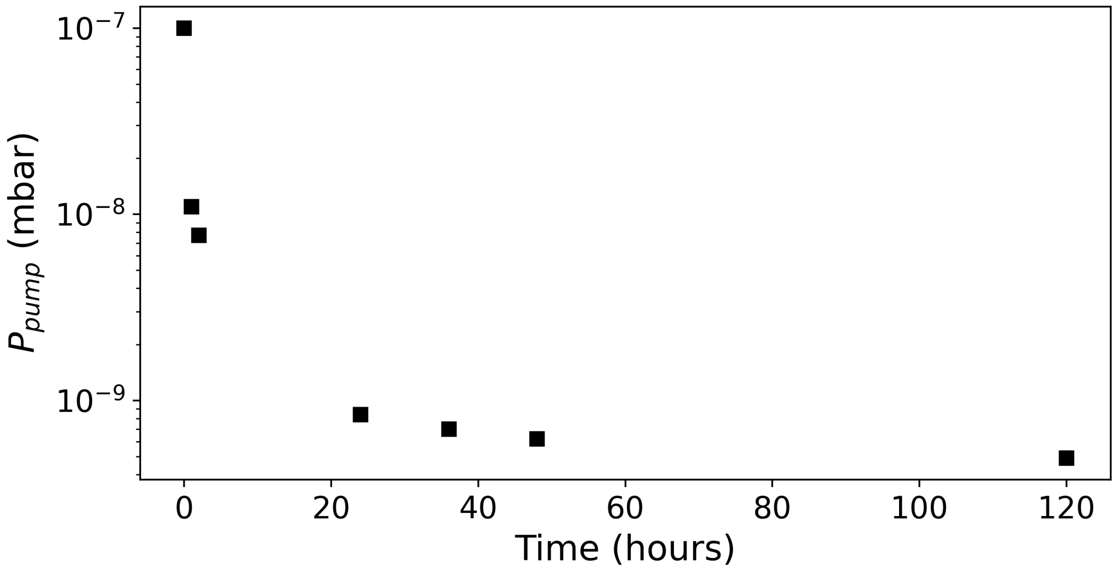

Figure 3.21: Ion pump pressure reading, $P_{\text{pump}}$, vs time, after bake-out for 4 weeks at 100°C. After 5 days, the pressure at the top of the flight tube was $4.9 \times 10^{-10}$ mbar. The pressure reading on the ion gauge remained steady at around $1 \times 10^{-8}$ mbar.

Once it was clear that the ion pump alone was not going to reduce the pressure much further, it was time to activate the non-evaporable getter (NEG), which provides the majority of the pumping speed of the combination pump. The manufacturer recommends delivering 45 W of electrical heating power to the NEG for 1 hour to activate. To do this, the ion pump must be switched off, so the only pressure reading from the system comes from the ion gauge located near the turbopump. This is because the NEG gets very hot ($\sim$ 550°C) and will produce a lot of outgassing, so it is best to pump the resulting gas out of the system via the turbopump than have it embedded in the ion pump filaments.

A thermocouple was attached to the top of the flight tube, near the indium seal, to check that the temperature did not approach the melting point of the indium seal. Initially, the gauge pressure was $1.3 \times 10^{-8}$ mbar, and increased as the heating power was increased, until a maximum pressure of $3.2 \times 10^{-5}$ mbar was reached when the heating power was 19.7 W, less than half the maximum. As the power was increased further, the pressure decreased. Once the maximum power had been delivered for 1 hour, the power was shut off, and an hour later the gauge pressure reading was close to where it was before the process started. The ion pump was then restarted, and after about 18 hours the ion pump was reading a pressure of

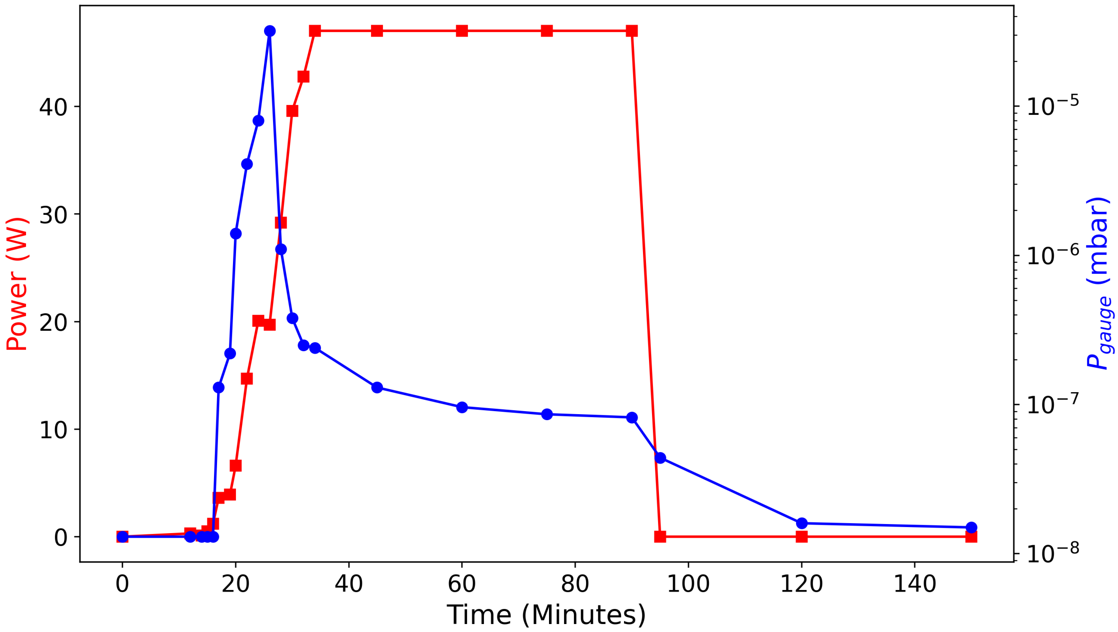

Figure 3.22: Heating power and ion gauge pressure vs time for the NEG activation process.

$P_{\text{pump}} < 2.1 \times 10^{-11}$ mbar (0 nA of current). This makes sense, as the majority of the pumping speed comes from the NEG, especially for hydrogen gas, so after activating the NEG the pumping speed would be much higher than for the ion pump alone, such that a pressure drop of a factor of 20 or so is to be expected.

The next step is to switch on the rubidium dispensers. This also involves heating them to a high temperature, and since the dispensers have some NEG material incorporated into them to reduce the outgassing of unwanted reaction products, a similar strategy to the NEG activation was used, this time with the ion pump running. The dispenser current and ion pump pressure vs time for the activation of the lower set of dispensers is shown in figure 3.23, with very similar results to the NEG activation.

The last step is to separate the mini-fountain from the turbopump assembly by closing the copper pinch tube. This is done by a hydraulic vice that squashes the tube flat against itself, such that a new seal is created. Before the pinch off, the ion pump, which was reading $< 2.1 \times 10^{-11}$ mbar, was turned off and the turbopump valve was closed. This is to protect the ion pump and turbopump from a sudden influx of air, should the pinch process go wrong. The resulting pinch tube end is

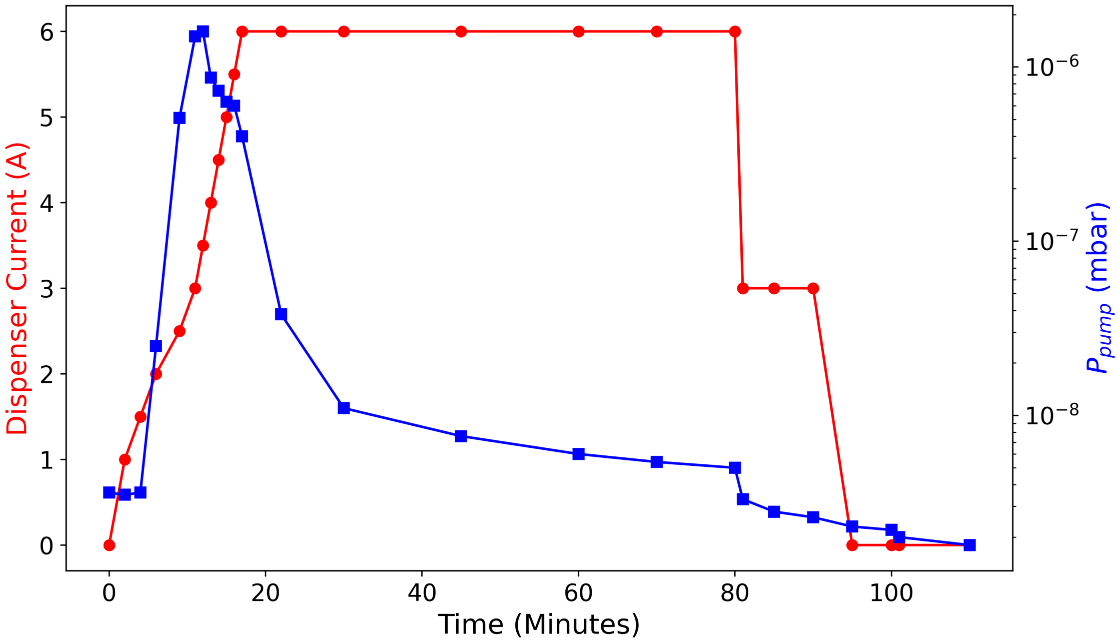

Figure 3.23: Dispenser current and $P_{\mathrm{pump}}$ vs time for dispenser activation.

shown in figure 3.24, and when the ion pump was restarted the pressure reading was still $< 2.1 \times 10^{-11}$ mbar, indicating a successful seal.

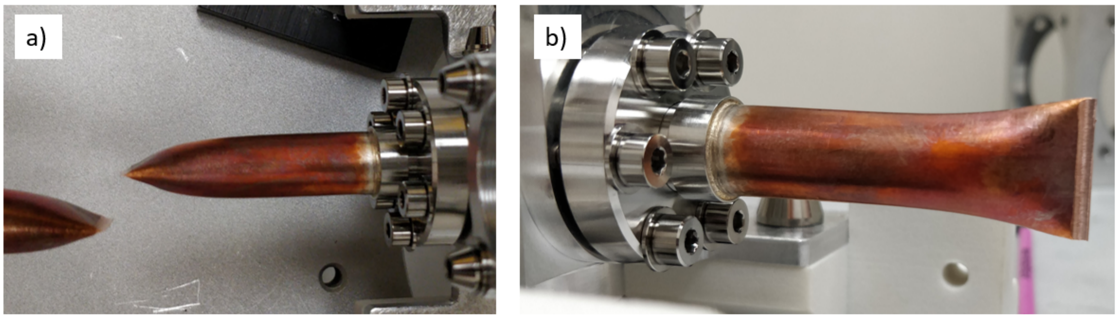

Figure 3.24: a) Top-down view and b) side-on view of the resulting pinch tube after closing by hydraulic vice.

# Chapter 4

# Operation & Evaluation

With the mini-fountain physics package fully assembled, the quest for clock operation can begin. There are many pre-requisite steps to take before its frequency stability can be measured. The MOT beams and coils must be set up, and the correct laser lock settings must be found to achieve cold atoms. The atom temperature then needs to be minimised by adjusting the sub-Doppler cooling stage of the sequence. Then, the atoms can be launched to the appropriate height, and the detection signal-to-noise ratio can be maximised by adjusting the detection parameters (pulse timings, intensity, detuning . . . ), as well finding the correct microwave power to the Ramsey cavity to implement a nominal $\pi/2$-pulse per cavity transit. Only then can the clock's frequency stability be measured by comparing it to a more stable clock, in this case NPL's hydrogen maser clock responsible for the national timescale, which is in turn disciplined by the primary frequency standard NPL-CsF2.

## 4.1 Initial Cold Atoms

To verify the presence of the rubidium background thermal vapour, a small fibre collimator was installed to produce a small diameter beam, which could be viewed perpendicularly through another port with an IR viewer. When the dispenser current was set to 4.0 A, and the cooling and repump lasers locked, a diagonal column of fluorescence was clearly visible, confirming the presence of the thermal vapour. This is shown in figure 4.1a.

When all six MOT collimators were installed and connected to the fibre optic control system, the MOT coils, with inner radius 3.08 cm, radial thickness 1.25 cm, coil-to-coil separation 4.8 cm, axial thickness 1.4 cm, and 120 turns of 1 mm diameter enamelled-copper wire, were connected in the orientation shown in figure 4.1b, which was verified with a magnetometer. The MOT coil current was set to pulse periodically, and bright spot was observed in the centre of the MOT chamber that coincided with the MOT current pulse, which verifies the presence of cold atoms. This is shown in figure 4.1c.

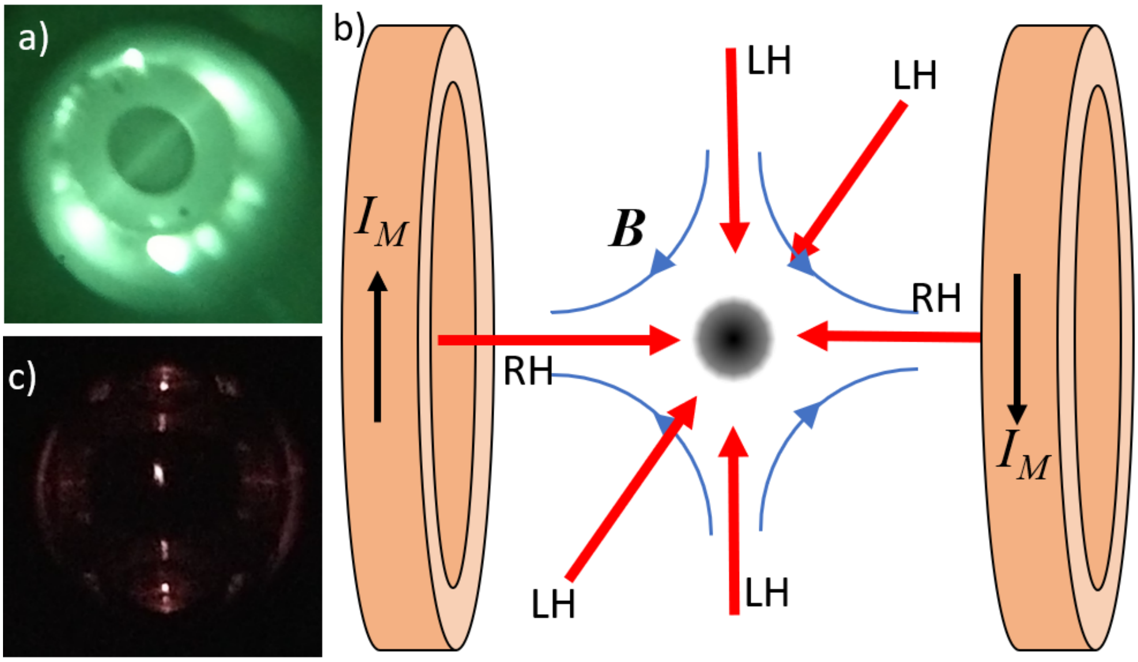

Figure 4.1: a) Diagonal beam of fluorescence from thermal background vapour. b) Orientation of the beams relative to the coils, with the right-handed beams along the axis of the coils. c) Initial MOT fluorescence.

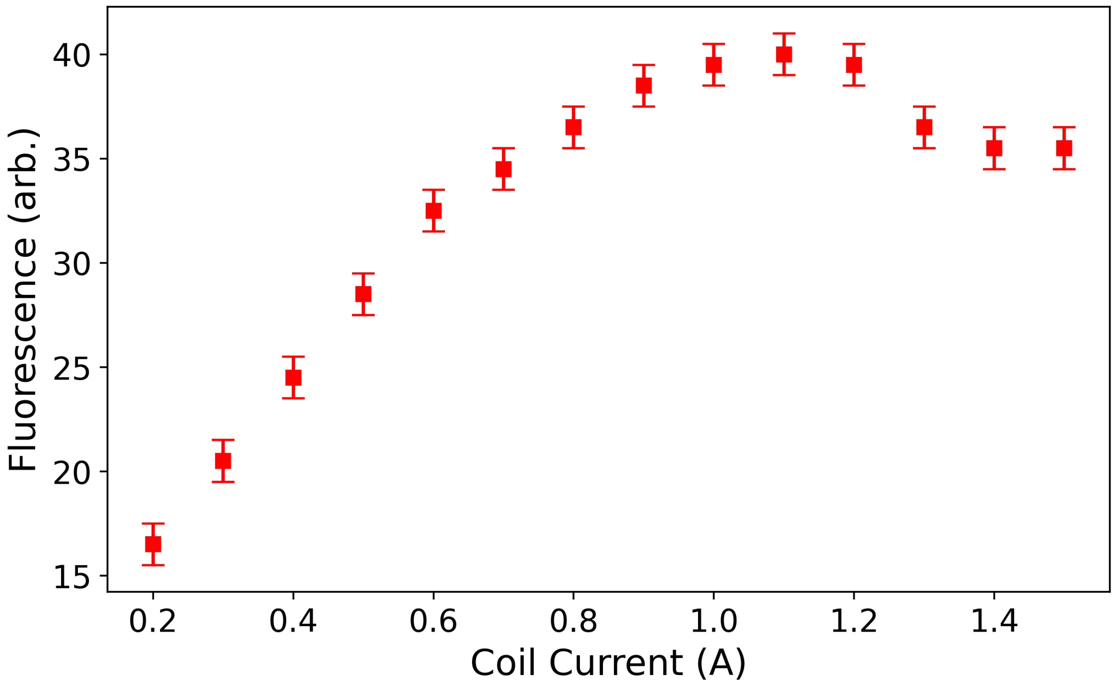

Figure 4.2: Fluorescence after 0.45 s loading time vs MOT coil current.

To find good MOT parameters, a CCD camera was installed to observe the cold-atom cloud, and a photo was taken after 0.45 s of loading, which was used to measure the fluorescence. The first parameter to optimise was the MOT coil current, by finding the current at which the fluorescence is maximised. The results are shown in figure 4.2, and the optimal MOT coil current was found to be $I_M = 1.1\,\mathrm{A}$.

Next, the laser intensity, via the TA current, was optimised. This time, the quantity to maximise is the fluorescence divided by the laser power, because, unless all atoms in the cloud are saturated, more power will of course result in more scattered light, but not necessarily more cold atoms. The result is shown in figure 4.3, and the result is that the fluorescence/power is maximised when the TA current is 1.1 A, which corresponds to an intensity of $4.8\,\mathrm{mW/cm^2}$ over the central 1 cm diameter of each collimated beam, which is a little more than the isotropic saturation intensity for the cooling transition of $3.51\,\mathrm{mW/cm^2}$.

Next, the laser detuning, via the VCO control voltage, was optimised, by maximising the MOT fluorescence. The results are shown in figure 4.4, where fluorescence increased until $\delta/\Gamma \approx -2.05$ (4.8 V on VCO), after which continued decreases in detuning rapidly dropped the MOT fluorescence.

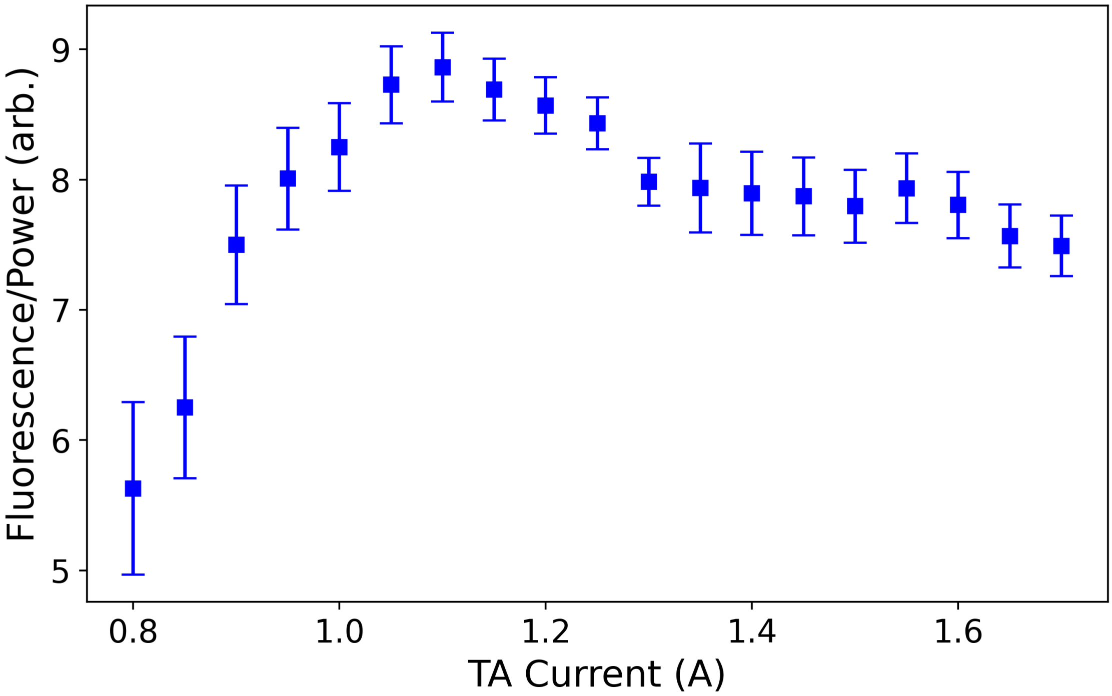

Figure 4.3: MOT fluorescence normalised to laser power vs TA current.

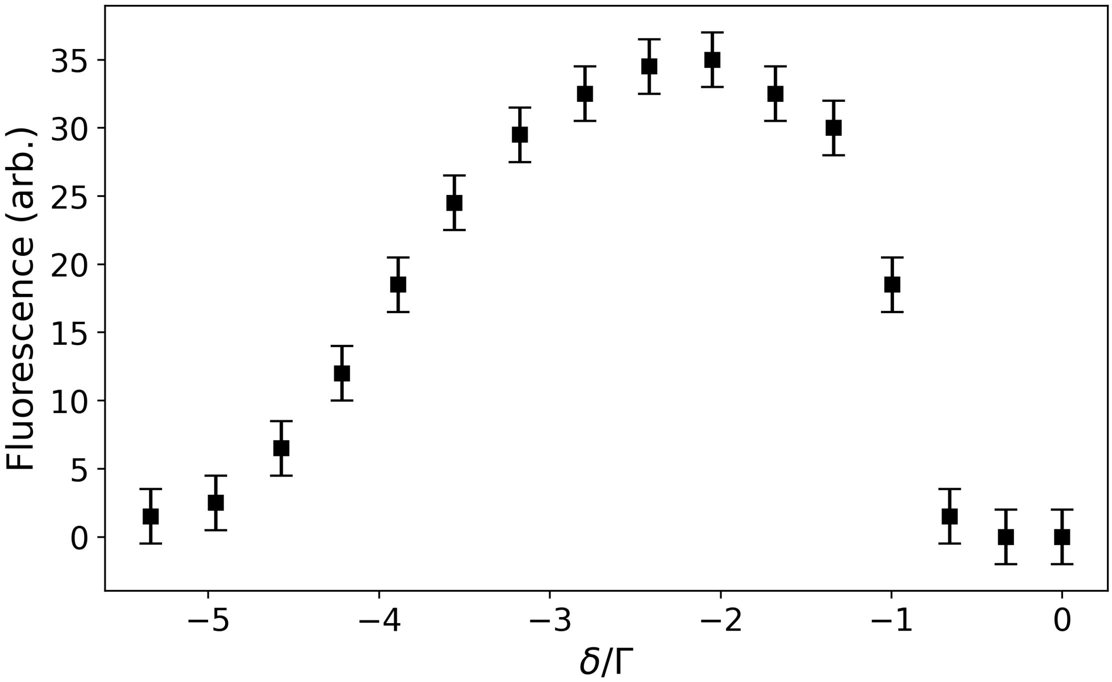

Figure 4.4: MOT fluorescence vs beam detuning.

## 4.2 Atom Number & Background Pressure

The rubidium pressure can be estimated from loading rate $R_L$ and loss time constant $\tau_L$. To determine these, the fluorescence was focussed onto a photodiode with a sensitivity $\eta = 0.57\,\mathrm{A/W}$, and the resulting photocurrent was amplified by a transimpedence amplifier with a gain of $g = 10^5\,\mathrm{V/A}$. Since the fraction of the total solid angle subtended by a single detection lens is 0.029, the atom number relates to the detector voltage $v$ by

$$N_{\mathrm{at}} = \frac{v}{g\eta h\nu R_s(\delta, I) \times 0.029}, \tag{4.1}$$

where $R_s(\delta, I)$ is the scattering rate from equation 1.24. In a six-beam MOT, each beam has a different scattering rate which depends on the atomic velocity and the $m_F$ state. A reasonable approximation is to assume the isotropic saturation intensity of $I_{\mathrm{sat}}^{\mathrm{iso}} = 3.51\,\mathrm{mW/cm^2}$, and to use the total intensity of all six MOT beams, which at $4.8\,\mathrm{mW/cm^2}$ per beam (in the central 1 cm diameter of the trap) is much greater than the saturation intensity. With a detuning of $\delta = -2.05\,\Gamma$, the scattering rate (per atom) in the MOT is $R_s = 9.6 \times 10^5\,\mathrm{s^{-1}}$, such that the atom number in the MOT is $N_{\mathrm{at}} \approx 2.5 \times 10^9\,\mathrm{V^{-1}}$.

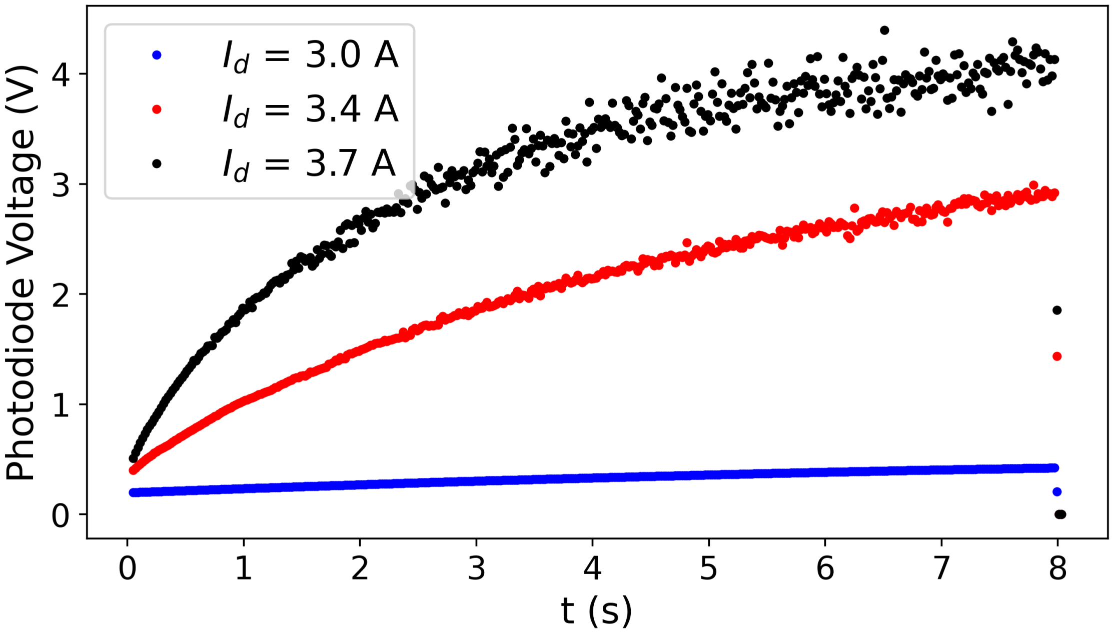

Figure 4.5: Voltage vs time for loading the MOT for different dispenser currents.

Some example loading curves are shown in figure 4.5. These exponential curves can be fit to

$$N(t) = R_L \tau_L (1 - e^{-t/\tau_L}), \tag{4.2}$$

to obtain $R_L$ and $\tau_L$. Plots of $R_L$ and $\tau_L$ vs dispenser current are shown in figure 4.6. Between each dispenser current value, the MOT was allowed to continue pulsing until the fluorescence value had stabilised, usually after only a few minutes.

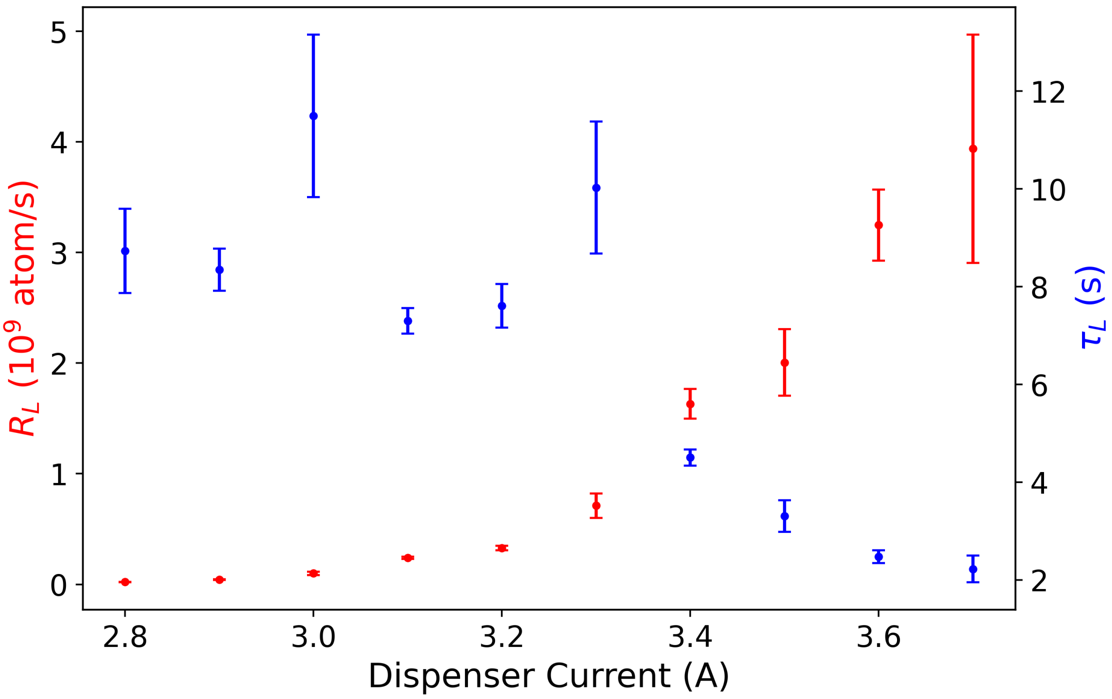

Figure 4.6: Loading rate and loss time constants vs dispenser current. As expected, higher current increases $R_L$, via a larger density of atoms available to trap, and decreases $\tau_L$, via a larger density of atoms with which cold atoms can collide.

The best way to determine the background pressure is via the loss time constant $\tau_L$ via equation 1.34, as, unlike $R_L$, it is independent of many of the parameters of equation 4.1, and instead comes from a simple fit to a function of time. There may be multiple background gas species, each with a different thermally-averaged collisional-loss cross-section. However, the ion pump at the top of the fountain is reading the lowest possible value ($2.1 \times 10^{-11}$ mbar), which suggests that residual gases in the vacuum are negligible. Thus, only the loss due to collisions with other rubidium atoms is considered here.

The collisional cross-section depends on the $C_6$ coefficient of the interatomic Van der Waals potential for ground-state atoms, which for rubidium-rubidium collisions is $C_6 = 9.5 \times 10^{-76}\,\mathrm{Jm}^6$ [213], and on the $C_3$ coefficient for an excited-state cold atom [42], which for rubidium is $C_3 = 5.4 \times 10^{-48}\,\mathrm{Jm}^3$. Thus, the effective cross-section is typically somewhere between the smaller Van der Waals value and the larger dipole value [49]. Using the formula [42]

$$\sigma_n = \left( \frac{\pi f_n C_n}{m v_c \bar{v}} \right)^{2/n}, \tag{4.3}$$

where $f_3 = 20.2$ and $f_6 = 8.08$, which accounts for glancing collisions in the trap of escape velocity $v_c$ (equivalent to equation 1.36 for the capture velocity) at the mean thermal velocity $\bar{v} = \sqrt{8k_BT/\pi m}$ ($270\,\mathrm{ms}^{-1}$ at $300\,\mathrm{K}$). Typically the excited-state population in a MOT of moderate optical intensity is around 0.3 [42], thus the appropriate cross-section is

$$\sigma_c \approx 0.3\,\sigma_3 + 0.7\,\sigma_6 \approx 1.6 \times 10^{-17}\ \mathrm{m}^2. \tag{4.4}$$

Thus, the rubidium vapour pressure in the MOT chamber is approximately

$$P_{Rb} \approx \frac{9.6 \times 10^{-9}}{\tau_L[\mathrm{s}]}\ \mathrm{mbar}. \tag{4.5}$$

The inferred rubidium pressure vs dispenser current is shown in figure 4.7. It is important to note that $P_{Rb}$, since it comes from the loss rate, is for the total partial-pressures of both $^{87}$Rb and $^{85}$Rb, however only the $^{87}$Rb pressure contributes to the loading rate. Even for the highest dispenser current used so far, the background $^{87}$Rb pressure is only $1.1 \times 10^{-9}\,\mathrm{mbar}$. The finally-used dispenser current will be a trade-off between the loading rate and the collisional loss during the time-of-flight through the interrogation and detection region.

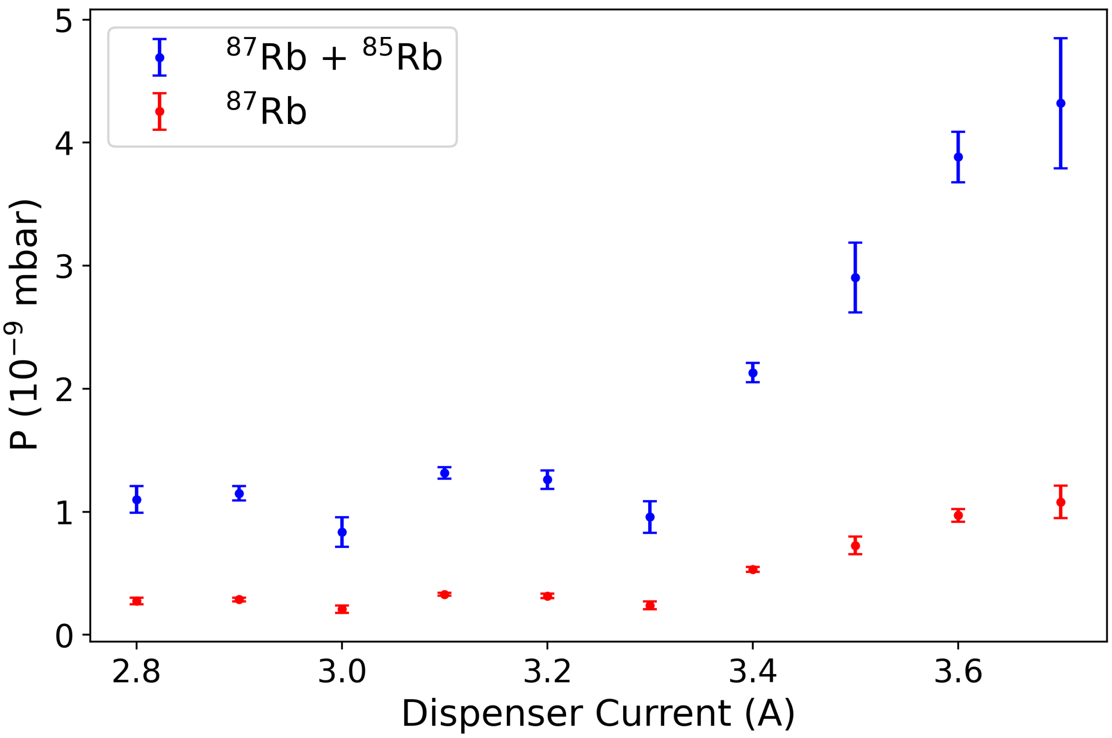

Figure 4.7: Pressure vs dispenser current in the MOT chamber. The values for $^{87}$Rb are the total pressure values, multiplied by the relative abundance of 0.25.

## 4.3 Sub-Doppler Cooling and Launching

Part of what is required to achieve the best cloud temperatures is to ensure zero residual magnetic field at the centre of the trap. Full mini-fountain operation will do this by shielding out the background field, but at that point the fountain physics package becomes inaccessible. Until then, an interim method is to compensate the background magnetic field, by wrapping three orthogonal pairs of square Helmholtz coils, of 5 turns each, around the MOT chamber. The MOT fluorescence was unaffected by the compensation current, because the MOT field gradients of $1000\,\mu\text{T/cm}$ along the coil axis and $500\,\mu\text{T/cm}$ in the plane of the coils are large compared to the background fields of $\lesssim 50\,\mu\text{T}$.

After the MOT phase, which lasts 0.5 s, the laser light is turned off for 2.0 ms, such that the MOT coils are turned off while the atoms are in the dark. This is to avoid erratic atom motion in the eddy currents caused by the sudden, large change in MOT coil current. Then, the laser light is restored, at a slightly lower intensity

($3.13\,\mathrm{mW/cm^2}$) and slightly increased detuning ($-2.4\,\Gamma$), where the cloud is further cooled without the influence of the large magnetic field gradients. This phase lasts for 30.0 ms. The lengths of both the eddy-current decay time and the molasses time were not seen to affect the final atom temperature.

To view the background magnetic field and set the correct compensation coil currents, Zeeman spectroscopy is carried out using the different $m_F$ components of the hyperfine transition. For this, a microwave waveguide is directed into the chamber. After the molasses phase, the laser light is then turned off for 5 ms, during which time a pulse of microwave power is applied. The resulting transition probability is then inferred by the normalised detection method described in section 2.1, albeit the pulse timings, intensity, and detuning need not be carefully optimised yet; most of the cold atoms are detected after this short release time such that the signals are large. The frequency of the microwave pulse can be scanned to obtain a plot of detected ratio vs frequency, where dips in the ratio correspond to a specific $|2, m_F\rangle \rightarrow |1, m_{F'}\rangle$ transition. The Zeeman shift of a state $|F, m_F\rangle$ due to a small magnetic field with magnitude $B$ is

$$\Delta E_{|F,m_F\rangle} = g_F \mu_B m_F B, \tag{4.6}$$

where, in the $F = 2$ and $F = 1$ ground-state hyperfine levels of $^{87}$Rb, $g_F = \pm 1/2$, respectively. The resulting frequency shift from the clock transition is then

$$\delta\nu_{m_F,m_{F'}} = \frac{1}{2}\frac{\mu_B}{h}(m_{F'} + m_F)B \; = 6.99\,\mathrm{kHz}/\mu\mathrm{T}. \tag{4.7}$$

The frequency scan with the compensation currents that achieved the minimum fringe spacing (to the nearest 0.01 A) is shown in figure 4.8, where the $|2, 0\rangle \rightarrow |1, \pm 1\rangle$ transitions occur at $\delta\nu = \pm 2.0\,\mathrm{kHz}$, implying a residual field of only $B = 0.3\,\mu\mathrm{T}$. As noted in section 2.2.2, a small but non-zero $B_z$ component is required to lift the degeneracy between the clock transition and the other $m_F$, $m_{F'} \neq 0$ transitions during state selection. This can be achieved by increasing the vertical compensation coil current, $I_z$, such that there are several kilohertz between the $|2, 0\rangle \rightarrow |1, 0\rangle$ and $|2, 0\rangle \rightarrow |1, \pm 1\rangle$ transitions.

After the molasses phase, the profile registers of the DDS boards that generate the RF for the up and down-going beams are shifted, such that the up-going beams become blue-detuned and the down-going beams become red-detuned by a

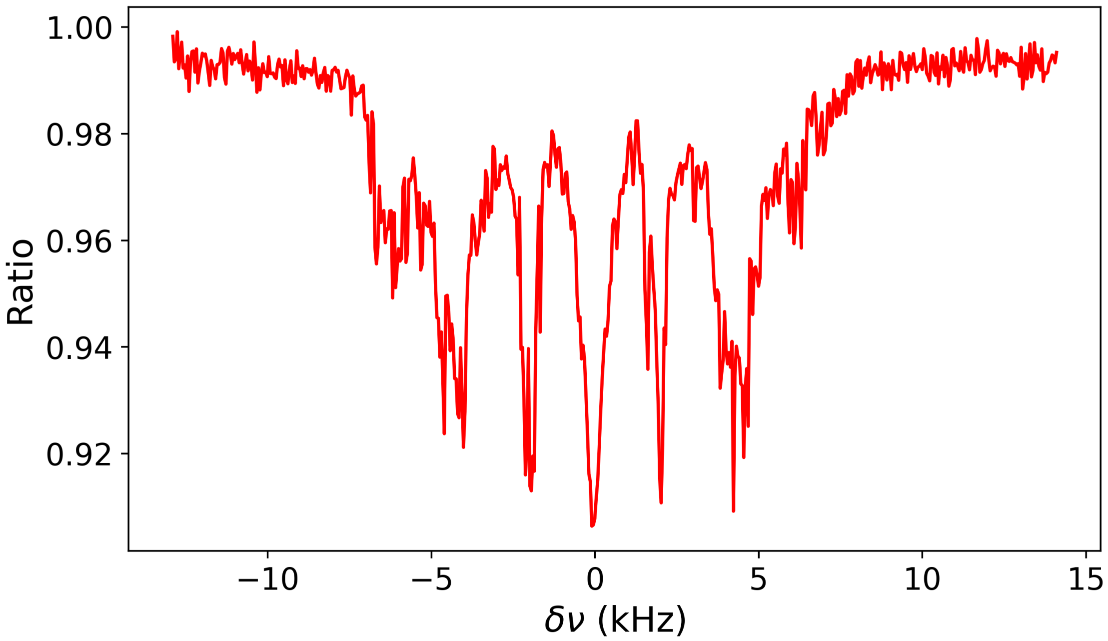

Figure 4.8: Ratio vs frequency detuning from the clock transition for close-to-optimal background field compensation: $I_x = 0.69\,\mathrm{A}$, $I_y = 0.05\,\mathrm{A}$, and $I_z = 0.96\,\mathrm{A}$, results in a residual field of $0.3\,\mu\mathrm{T}$.

few megahertz, depending on the intended launch velocity $v$, height $h$, and time $t$:

$$\delta\nu = \frac{v}{\sqrt{2}\lambda} = \frac{gt}{2\sqrt{2}\lambda} = \frac{\sqrt{gh}}{\lambda}. \tag{4.8}$$

All six beams are left on during this phase. After 0.5 ms in the moving frame, the detuning is increased and the intensity decreased in three stages; the first two lasting 0.5 ms each, and the final stage lasting 1.0 ms. However, the resulting cloud temperature seemed insensitive to the exact details of the ramp, so a exponential approximation where the intensity halves, and the detuning doubles, every 0.5 ms, is used. By then end of the ramps, the intensity in each cooling beam is only a few $\mu\mathrm{W/cm^2}$, and the detuning is $-11.9\,\Gamma$.

One can attempt to measure the cloud temperature by using the CCD camera and the release-recapture method described in section 2.1.2. However, if the approximation that the cloud's atom number density is Gaussian is not a good one, for example, if the cloud initially has a multi-lobe structure, then the results can be misleading, and the cloud's temperature is significantly over-estimated.

To accurately measure the final cloud temperature, a time-of-flight method similar to the detection method in Primary Frequency Standards is used, where the cloud width is measured as a function of launch time. The result is a series of Gaussian voltage vs time graphs, where the widths are a convolution of the returning cloud's radius, and the horizontal beam's Gaussian ($1\,\sigma$) radius. Some examples are shown in figure 4.9, where the detector gain has been increased to $g = 10^6\,\mathrm{V/A}$, and the horizontal beams are half of the MOT intensity, such that off-resonant pumping to the $F = 1$ states is minimal and the actual cloud profile is visible. To turn the time-width of the cloud into a spatial-width requires dividing by the launch velocity (which is also the velocity of the returning cloud) for each launch time. The resulting plot of detected radius vs launch time is shown in figure 4.10, and the integrated returning cold-atom photodetector signal vs time is shown in figure 4.11.

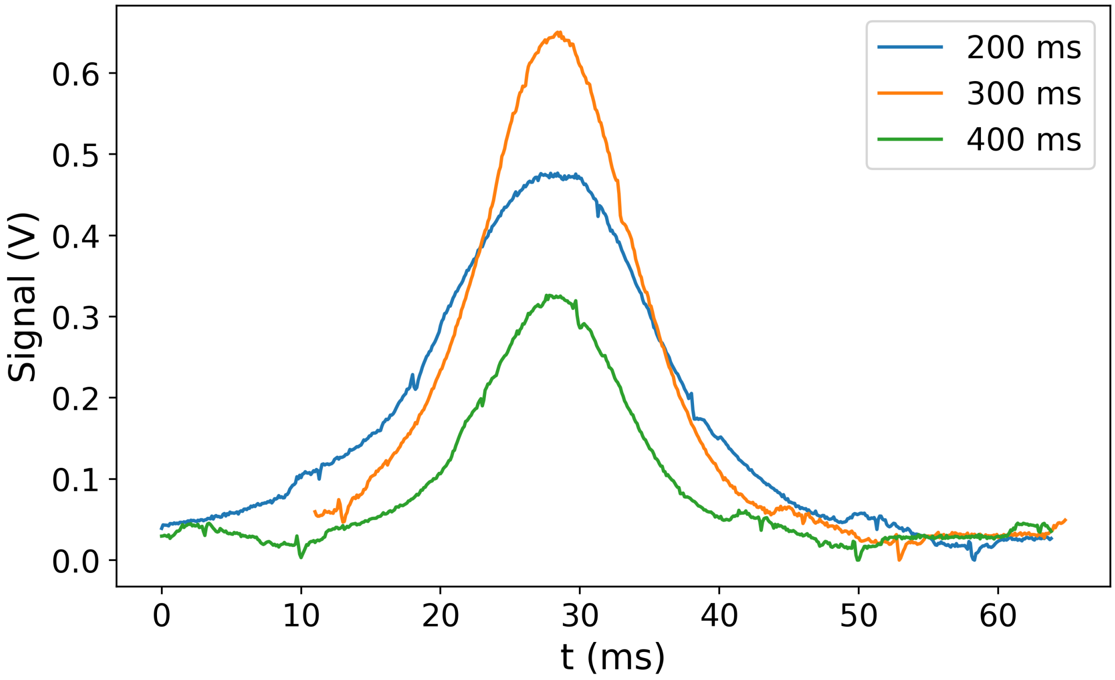

Figure 4.9: Gaussian time-of-flight signals for several launch times.

The cloud's temperature is found to be $T = 3.6 \pm 0.4\,\mu\mathrm{K}$, which is within the expected range from good polarisation gradient cooling. The cloud's temperature is robust to small variations (a few microtesla) in the background magnetic field, which bodes well for state selection.

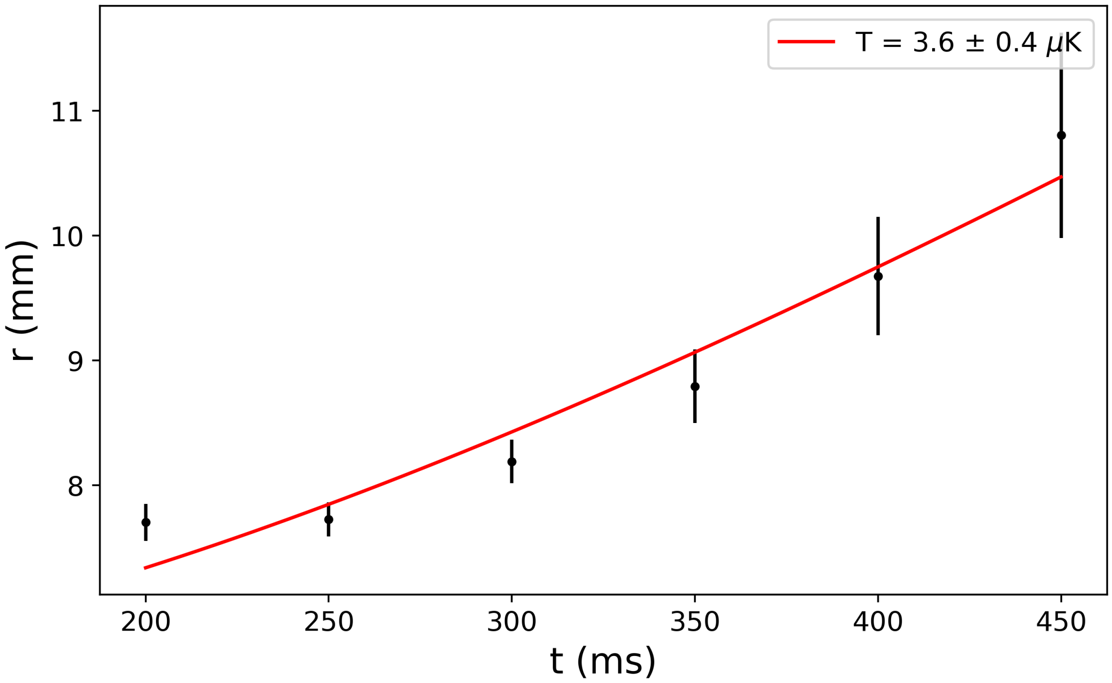

Figure 4.10: Detected cloud radius vs launch time, with a fit to $T = 3.6 \pm 0.4\,\mu$K.

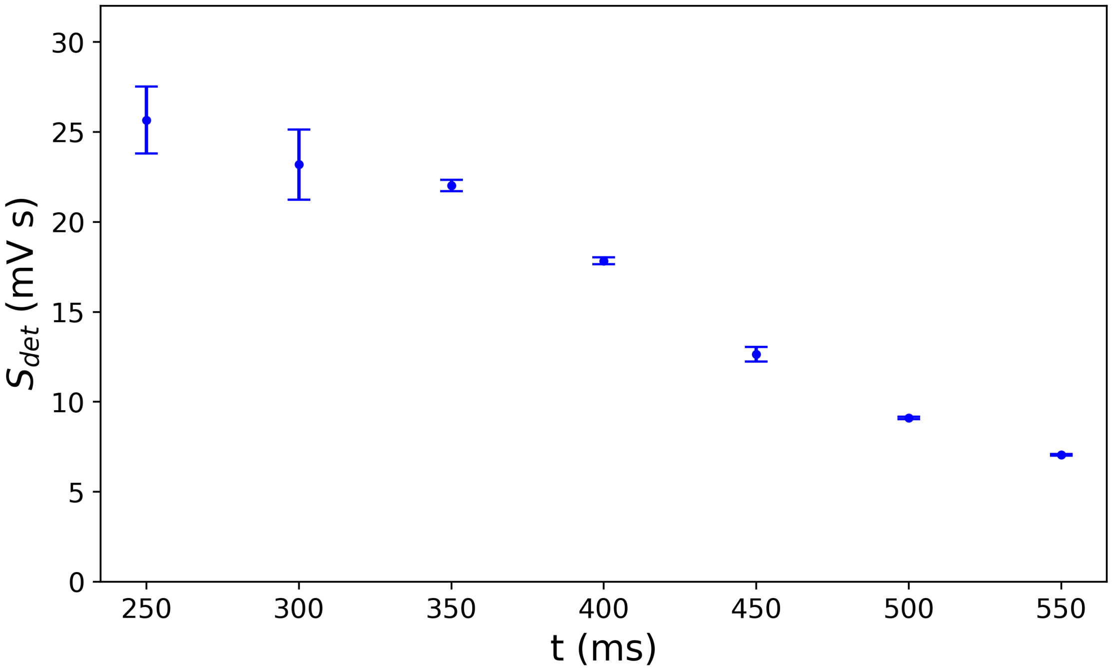

Figure 4.11: Returning total atom signal (measured in millivolt-seconds) vs launch time for Gaussian time-of-flight detection.

## 4.4 State Selection

Removing atoms not in the clock state ($m_F \neq 0$) from the launched atom cloud is important for two reasons: to maximise the Ramsey fringe contrast; and to remove the potential for additional detection noise due to inconsistent partitioning between clock/non-clock state populations from shot-to-shot.

In the mini-fountain, this isn't done with a dedicated cavity. Instead, microwaves are directed into the MOT chamber immediately after the sub-Doppler cooling. The state-selection microwaves are generated by a versatile synthesiser (Rhode & Schwartz SMB-100A), with up to +30 dBm output power. Amplitude modulation (to keep the microwaves "off" before and after state selection) is done by a broadband microwave switch which claims over 80 dB of isolation at 6.8 GHz. The microwaves are directed into the chamber by a coax-to-waveguide adapter. In the demonstrator system, and for the tests done in CsF3, this was followed by a microwave horn, which improves the impedance matching between the waveguide and free space. However, this would greatly increase the necessary diameter of the outer-shield, and hence the entire physics package. Thus, the waveguide is attached directly to the chamber, behind the mirror that retro-reflects the detection fluorescence, as in figure 2.11. This is shown in figure 4.12. The H-plane of the waveguide is aligned vertically, to optimally drive the clock $\pi$ transition in a vertical background magnetic field.

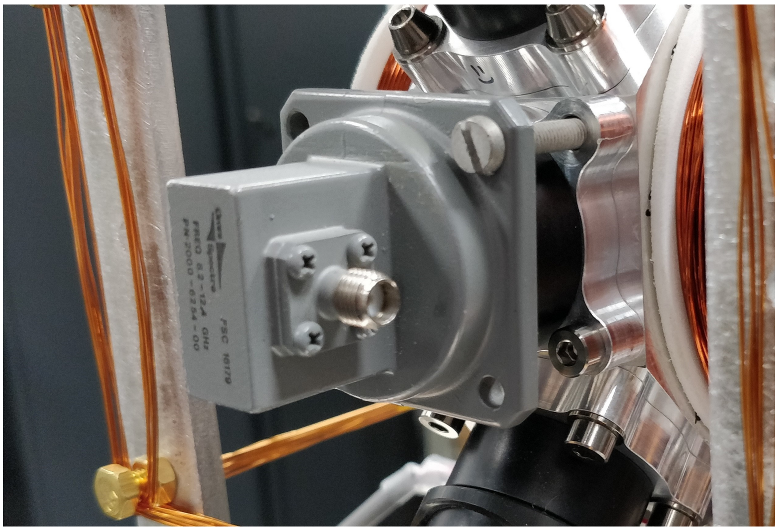

Figure 4.12: Coax-to-waveguide adapter attached behind the lens tube that houses the (dielectric) mirror for retro-reflecting detection fluorescence.

To ensure that the field in the chamber is vertical, the background field is first compensated, such that increasing the current through the vertical coils by 0.02 A results in a positive, vertical background magnetic field of 2.0 $\mu$T, which means a spacing of 14 kHz between the clock and $|1,0\rangle \rightarrow |2,\pm1\rangle$ transitions.

The state-selection pulse must be brief because of the need to clear the non-clock state atoms before the launched cloud leaves the beam region. Thus, 1.0 ms microwave pulses are used. The result is that the transferred atom fraction is highest at +24 dBm of synthesiser power. This is shown in figure 4.13.

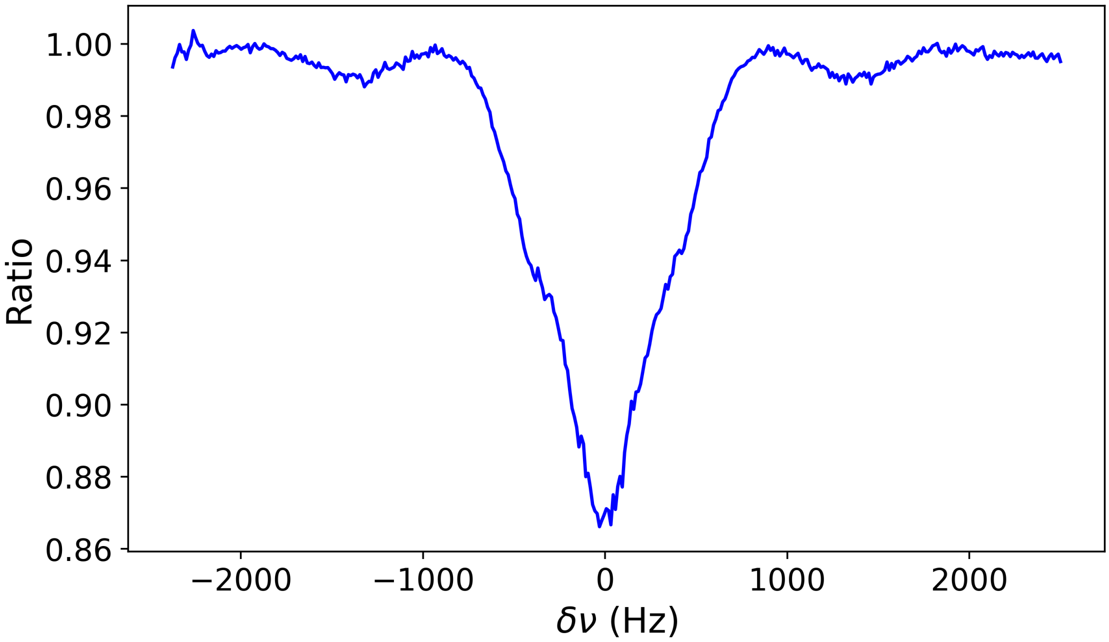

Figure 4.13: Detected ratio vs microwave frequency for 1.0 ms pulses of +24 dBm microwaves in the MOT chamber, aligned with the atom's vertical quantisation axis.

The detected ratio dipped from around 1.0 away from resonance down to 0.87 on resonance, implying that 65% of the clock-state atoms were transferred to the lower level. This implies that perhaps more microwave power is necessary, however increasing the power further, even by several decibels, had a negligible effect on the depth. The full-width-half-maximum is close to the expected 1.0 kHz.

Before and after the state-selection microwave pulse, it is necessary to detune the microwave frequency to avoid accidentally exciting Ramsey transitions with the residual microwave field in the chamber; the 80 dB of attenuation from the switch was seen to not be sufficient by itself. The exact detuning is not critical, so 5 kHz was chosen as it puts the residual radiation partway between the clock transition and the $|1, 0\rangle \rightarrow |2, +1\rangle$ transition, yet well-outside of the expected Rabi envelope of the Ramsey pattern. This modulation is done with the synthesiser's built-in frequency modulation function.

After the microwave pulse, the next step in the state-selection method is to use the down-going cooling beams to apply a short (0.5 ms), resonant pulse of maximum intensity to push the atoms remaining in the upper level out of the cloud. The state-selected atoms in the lower clock state do not see this light, so continue upwards at their launch velocity. Thus, only the clock state atoms are detected, allowing (theoretically) for a Ramsey fringe contrast of close to 100%. It was found that the laser lock system had no problem in going from far-detuned (at the end of sub-Doppler cooling) to on-resonance (during the pushing-pulse) in just over 1 ms.

## 4.5 Unshielded Ramsey Operation

Before the mini-fountain can operate as a clock, the Ramsey cavity's final microwave feed must be installed, which has a much longer length of rigid copper-core co-axial cable than the short feed used in section 3.1 for tests (to ensure that microwave leakage from the connector occurs away from the interrogation region). The result is that, when the new feed is appropriately shaped (as in figure 3.11b), the $|S_{11}|$ coefficient vs frequency still looks like figure 3.12, such that the oxidation of the copper waveguide that occurred during the bake-out has not appreciably changed the frequency of the waveguide. With the final feed installed, indium wire seals were placed around the top and bottom surfaces of the ring waveguide, to further reduce the cavity's microwave leakage.

To generate the Ramsey interrogation field, a frequency synthesiser (Spectra Dynamics RB-1) is installed, and can be controlled remotely by the same LabView software as the rest of the fountain sequence. This can be phase-locked to NPL's local hydrogen maser via a 5 MHz reference input for maximal frequency accuracy.

For a 550 ms launch (22.7 cm above the centre of the cavity), the atoms spend 23 ms inside the Ramsey cavity. When the synthesiser output power was set to +2.6 dBm, −60 dB of fixed attenuators are used, and the power is delivered to the cavity by a cable with −4 dB of loss, the contrast of the resulting Ramsey fringes was roughly maximised. These Ramsey fringes are shown in figure 4.14.

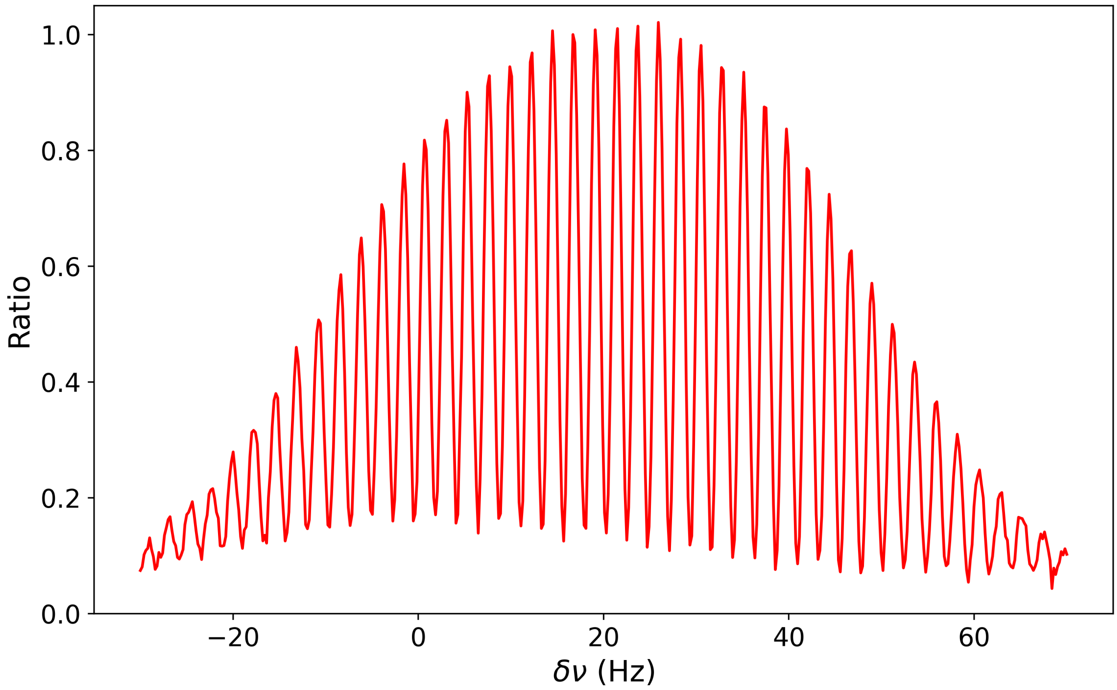

Figure 4.14: Ramsey fringes obtained for a 550 ms launch. The FWHM linewidth of each fringe is 1.2 Hz.

The centre of the Rabi envelope is shifted to approximately 20 Hz due to the second-order Zeeman shift, which implies a background field in the cavity, as per equation 1.47 with $g_J \approx 2$ for $S$ states, of 19 $\mu$T.

The signal-to-noise ratio does not seem to depend critically on the exact detection parameters; here the SNR on the peak of the "central" Ramsey fringe was around 100. This was done with just the horizontal beams operating at 3.6 mW/cm$^2$, a detuning of $\delta = -1.2\,\Gamma$, and detection pulse times of 1.0 ms. Adding in the diagonal beams actually degraded the SNR considerably.

Preliminary clock operation can now begin. In order to determine the difference between the local oscillator frequency and the atomic transition, the transition must be interrogated sequentially on each side of the Ramsey fringe, rather than on the peak, because the gradient of the ratio vs frequency, $dR/d\nu$, is zero on the peak but maximised on the slope. Thus, the frequency error requires two shots to determine:

$$\delta\nu = \left(\frac{dR}{d\nu}\right)^{-1} \frac{(R_{-} - R_{+})}{2}. \tag{4.9}$$

Without the C-field coil and magnetic shields to provide stability to the second-order Zeeman shift, the clock's instability is dominated by drift for any averaging time. This is shown for Ramsey operation in figure 4.15.

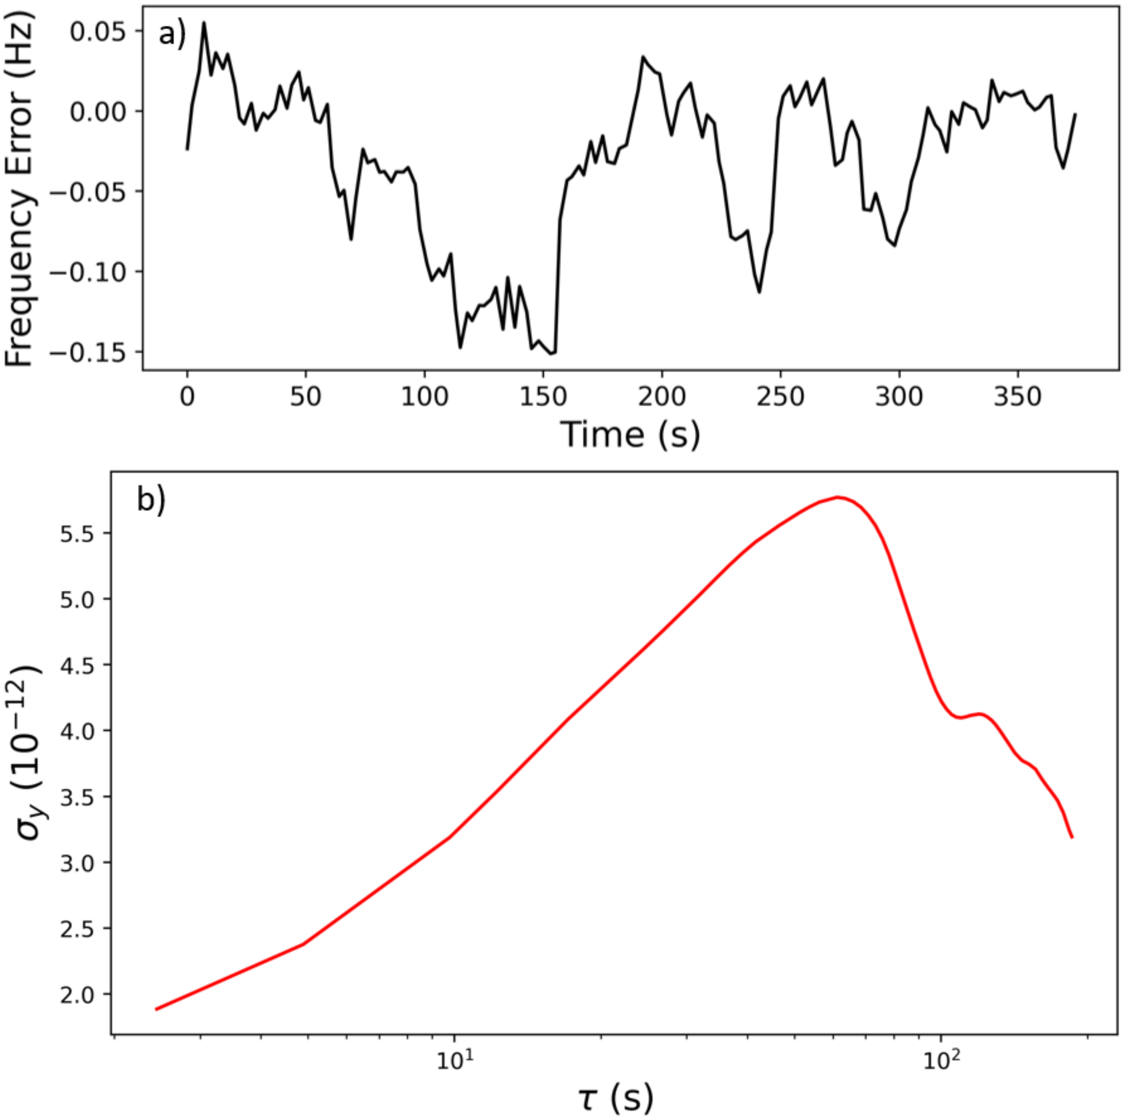

Figure 4.15: a) Frequency error vs time. b) Allan deviation vs averaging time for unshielded Ramsey operation. While $\sigma_y(1\,\mathrm{s}) \lesssim 2.0 \times 10^{-12}$, the clock's frequency stability is clearly dominated by drift due to instabilities in the second-order Zeeman shift.

## 4.6 Shield Integration

Clearly, no meaningful SNR optimisation or frequency stability measurements can occur until the magnetic field is well-controlled throughout the entire height of the interrogation region. To achieve this, a set of magnetic shields has been fabricated to the specifications discussed in section 2.4: an inner cylindrical mu-metal shield to surround the cavity and flight tube of length 38.5 cm and diameter 11 cm, and an outer shield of length 61 cm and diameter 20.5 cm. A C-field coil has been wound using a plastic tube cut down to fit tightly inside the inner shield, of length 38.5 cm and diameter 10 cm. The coil was wound with a total of 77 turns, spaced 5 mm apart by winding with 10 mm spacings on the way up and down, such that the wires for the input and output of the coil can be twisted together to reduce stray fields. The stages of installing the C-field coil and shields are shown in figure 4.16. At each stage, the background magnetic field, and thus the necessary compensation coil currents, varied massively, particularly in the vertical direction where up to 1.2 A ($\approx 120\,\mu$T) was necessary to observe good MOT fluorescence with just the inner shields present.

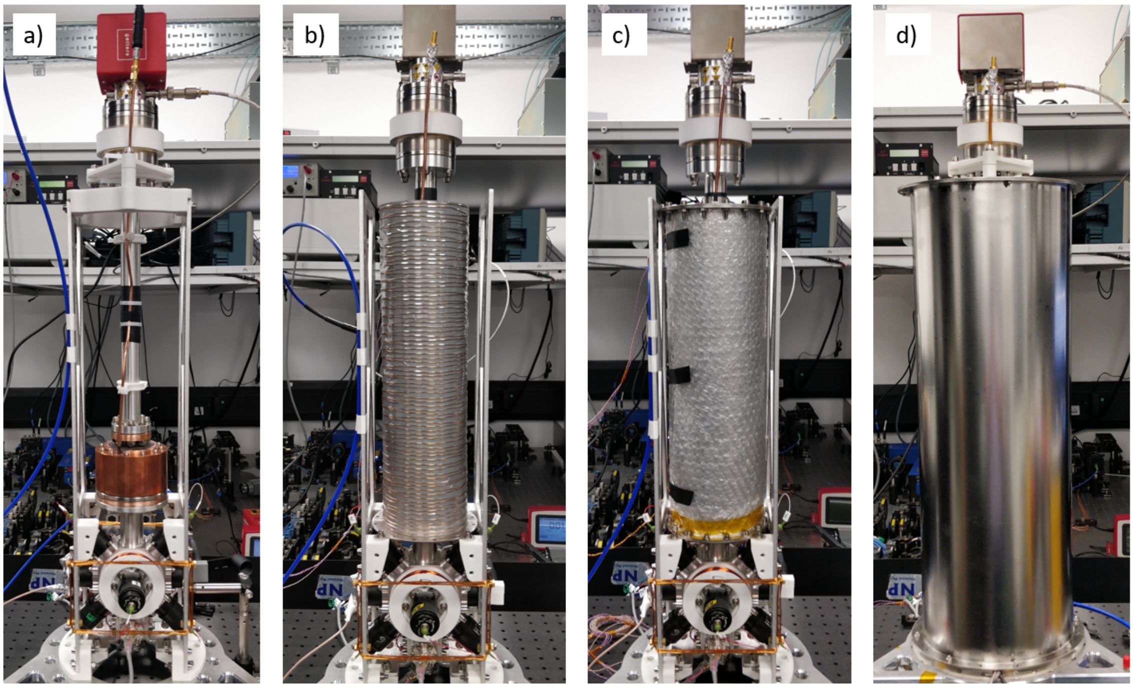

Figure 4.16: Mini-fountain physics package with a) no shields installed, b) lower-inner endcap to support the C-field coil, c) full inner-shield assembly and some additional thermal insulation, and d) full outer shield assembly installed.

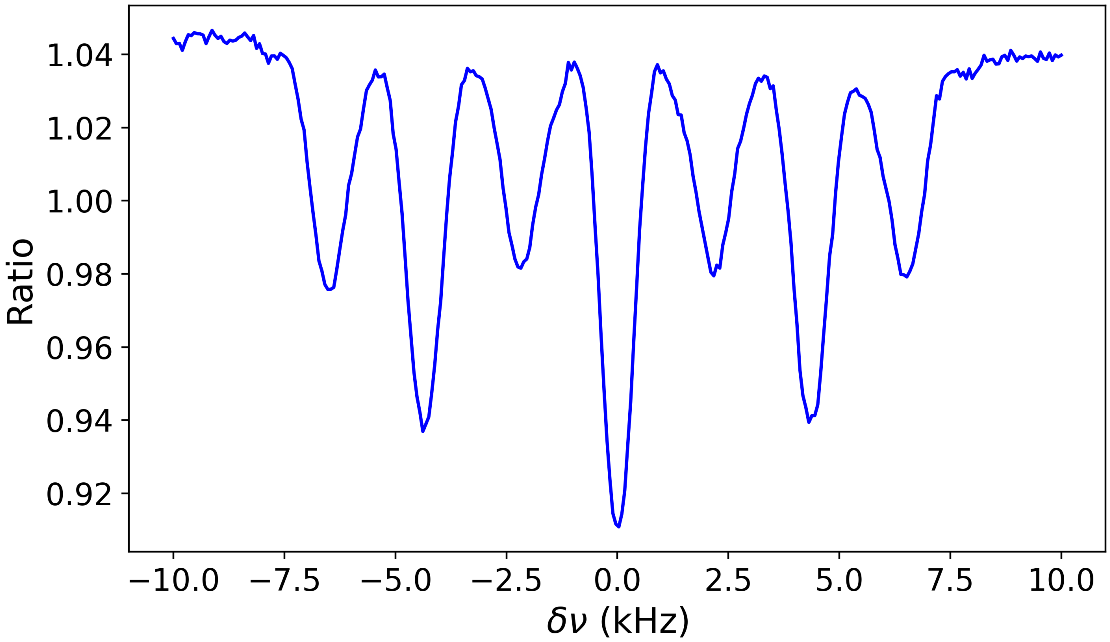

Figure 4.17: Frequency scan in the MOT with all magnetic shields installed and no compensation current.

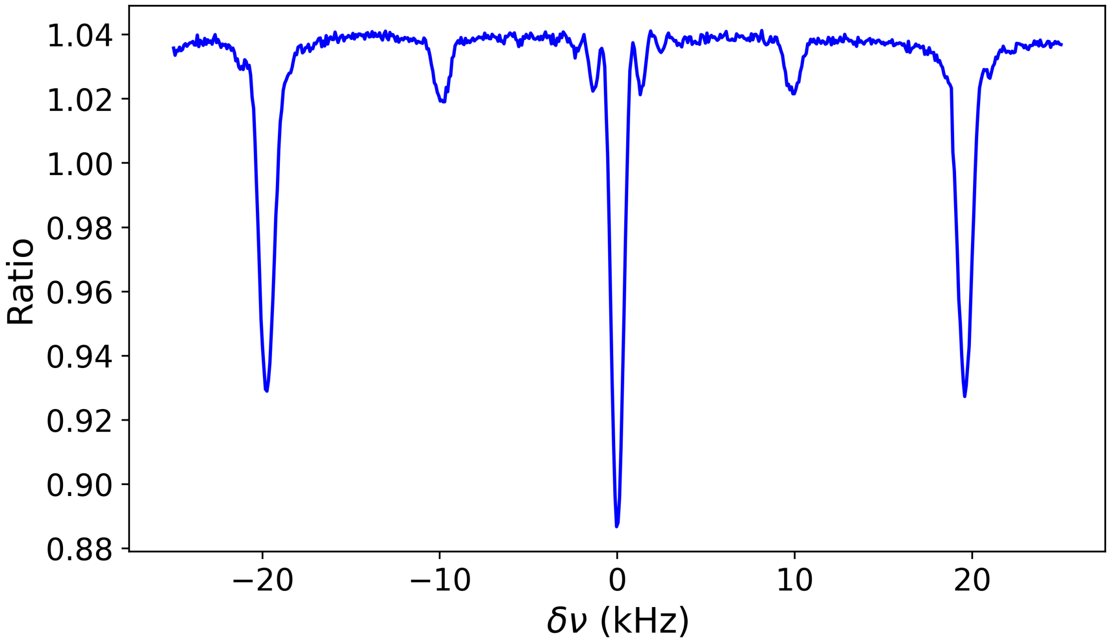

Figure 4.18: Frequency scan in the MOT with all magnetic shields installed and 0.2 A of vertical "compensation" current to increase the separation between adjacent transitions to 10 kHz.

With the full shield assembly installed, the microwave antenna was used to determine the residual magnetic field in the MOT chamber, with no compensation current. The resulting frequency scan is shown in figure 4.17, where the $|2, \pm 1\rangle \rightarrow |1, 0\rangle$ transitions occurs at $\pm 2.2\,\text{kHz}$, which corresponds to a residual magnetic field magnitude of $0.3\,\mu\text{T}$. Additionally, the width of each peak appears very similar, which implies that the magnetic field gradients in the MOT chamber are negligible.

With 0.2 A of vertical compensation current, the resulting magnetic field is dominated by a positive $B_z$ component. the resulting split between the clock transition and the $|2, \pm 1\rangle \rightarrow |1, \pm 1\rangle$ transitions is $\pm 20\,\text{kHz}$, which is sufficient separation to perform state selection with a 1 ms microwave pulse. This is shown in figure 4.18.

However, to ensure that atoms in the clock state stay in the clock state throughout their flight, and do not get re-distributed across all the other $m_F$ states by Majorana transitions, requires that the background magnetic field maintain its vertical orientation. Inside the inner shields this is accomplished by the C-field coil: according to the simulations of section 2.4, $400\,\mu\text{A}$ through our C-field coil should produce a uniform 100 nT background field. In the space between the MOT and the cavity, this is done by the vertical compensation coils. However, the inner shields have a neck of length 1.5 cm to aid with the attenuation of magnetic fields between the MOT region and the interrogation region. Thus, the field produced by the vertical compensation coils must be strong enough to keep the magnetic field vertical throughout the entire length of the shield neck. If it is too weak, then a change in sign of the $B_z$ component will cause Majorana transitions (as described in section 1.2.6) to occur on the atoms' way up, which is a far larger problem for the mini-fountain than on the way down: the $m_F$ states of the atoms will be re-distributed after state selection but before Ramsey interrogation, which will ruin the Ramsey fringe contrast. To demonstrate this effect, the atoms were launched to the centre of the cavity and underwent a $\pi$-pulse, and the resulting contrast of the Rabi oscillation, which now occurs at $\Delta\nu = 0$, is plotted vs coil current in figure 4.19.

With and without the C-field current, a steep change in contrast occurs at 300 mA of vertical current. Thus, avoiding Majorana transitions on the way up is achieved by setting $I_z \geq 300\,\text{mA}$. The fact that the fringe contrast is also good even without the C-field current means that there is a vertical background field present in the interrogation region that is also aligned with the compensation current and C-field. This background field was found to be mostly independent of $I_z$: with no C-field

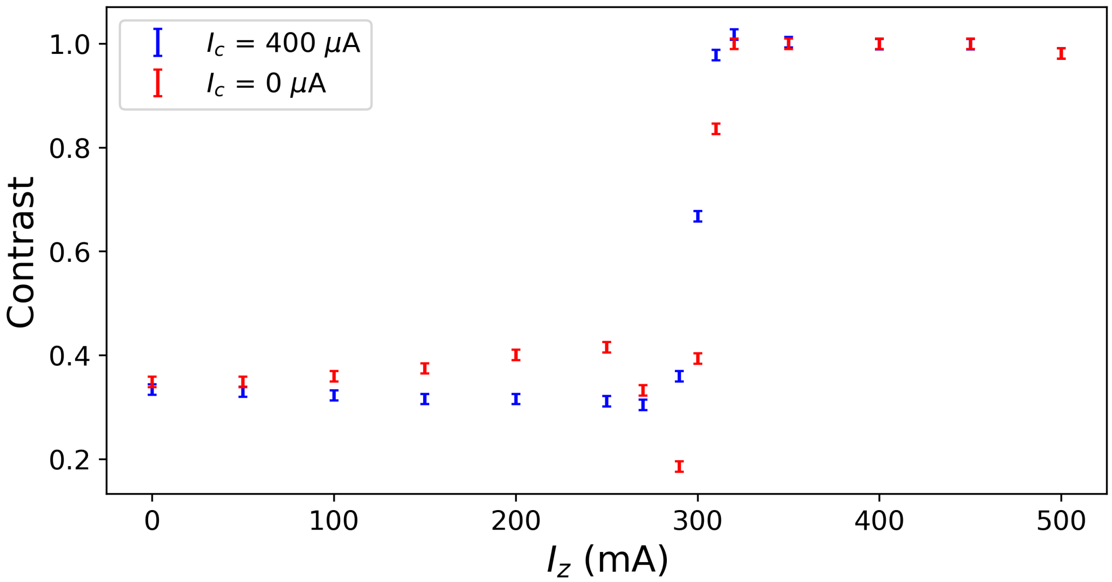

Figure 4.19: Vertical compensation current $I_z$ vs Rabi fringe contrast, for C-Field current $I_c = 400\,\mu$A and $0\,\mu$A. A step up to contrast $\approx 1.0$ can be seen when $I_z$ is large enough to create a vertical magnetic field throughout the entire shield neck.

current and $I_z = 350\,$mA, the 1-1 transition occurs around 560 Hz, which corresponds to a background field of 40 nT. With $I_z = 0$, this is shifted only by $-10\,$Hz to 550 Hz, thus the portion of the background field inside the cavity that is due to the compensation coils is minimal. The background field could be due to not having yet demagnetised the magnetic shields, or could simply be the residual Earth-field leaking in (40 nT would imply a longitudinal shielding factor of around 1000, which is somewhat consistent with the section 2.4, although on the low end; perhaps the real shielding factor is limited by the presence of holes in the endcaps).

Alignment between $I_z$ and $I_c$ was verified by flipping the direction of $I_c$: with the opposite polarity, no reasonable amount of $I_z$ can avoid a zero-crossing and thus the fringe contrast remains poor.

With the C-field at $I_c = 400\,\mu$A, the 1-1 transition occurs, in the cavity, at 1980 Hz, which corresponds to a total field of 140 nT; this means that the C-field coil is producing the expected 100 nT of vertical magnetic field, and is aligned with the residual 40 nT field. To verify the magnetic field throughout the atom's entire ballistic flight will require careful tracking of the 1-1 transition's central Ramsey fringe with incrementally increasing launch height.

## 4.7 Full Ramsey Operation & SNR Optimisation

With the magnetic shields and appropriate compensation and C-field coil currents in place, Ramsey operation around the clock frequency can now be achieved. The Ramsey fringes from a 550 ms launch time are shown in figure 4.20.

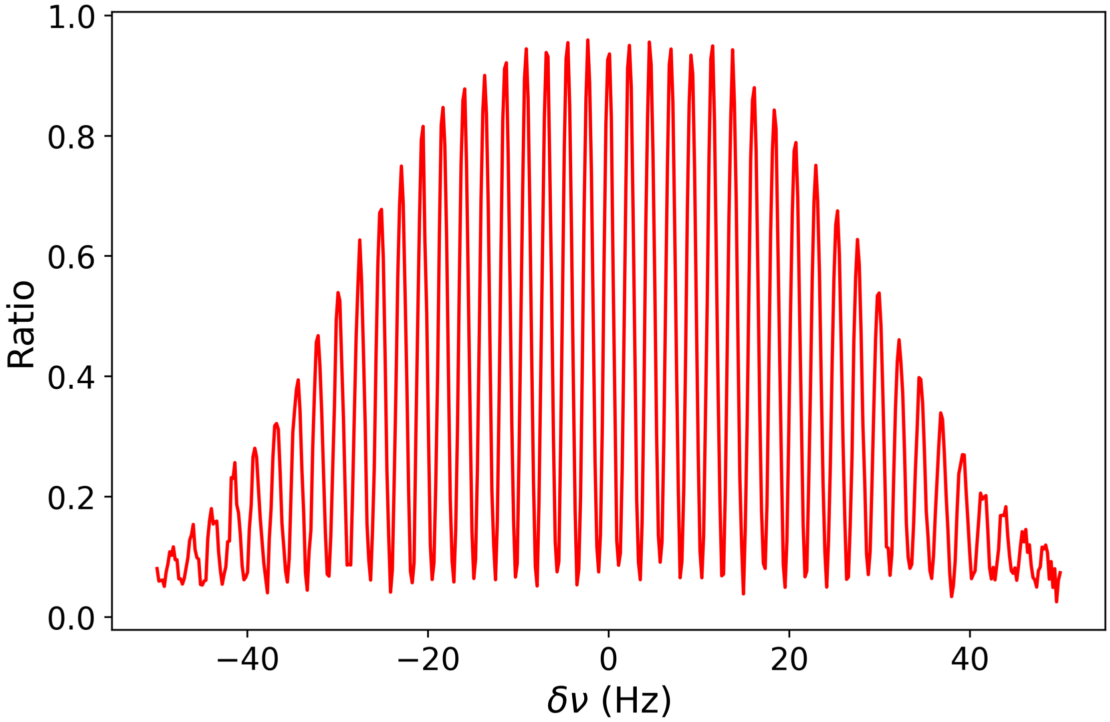

Figure 4.20: Detected transition probability vs frequency for a 550 ms launch, resulting in 1.2 Hz (FWHM) fringes centred on the clock frequency. The signal-to-noise ratio on the peak of the central fringe is 100 and the fringe contrast is 91%.

Now that there is consistently expected to be a fringe at $\delta\nu = 0$, this opens up the possibility of performing a more accurate scan of peak transition probability vs microwave amplitude in the Ramsey cavity. This is shown in figure 4.21. This clearly shows that the fringe contrast is optimised, and thus approximates a $\pi/2$-pulse per cavity transit, at $770 \pm 5$ mV (peak-to-peak out of the synthesiser), and that a $\pi$-pulse, which corresponds to the first minimum, occurs at twice this. Thus, 770 mV is the appropriate microwave amplitude for 550 ms launches.

Now, finally, the mini-fountain can begin actual clock operation, and the preliminary Allan deviation is shown in figure 4.22 for measuring vs NPL's hydrogen maser.

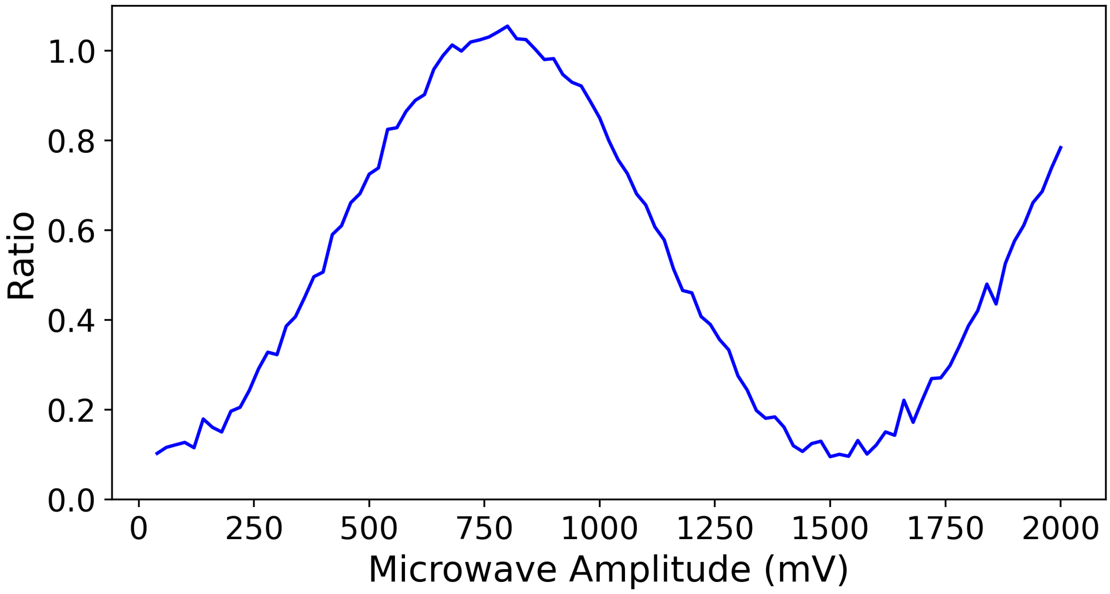

Figure 4.21: Transition probability vs peak-to-peak microwave amplitude out of the RB-1 synthesiser, which is attenuated by $-60\,\mathrm{dB}$ and coupled into the cavity.

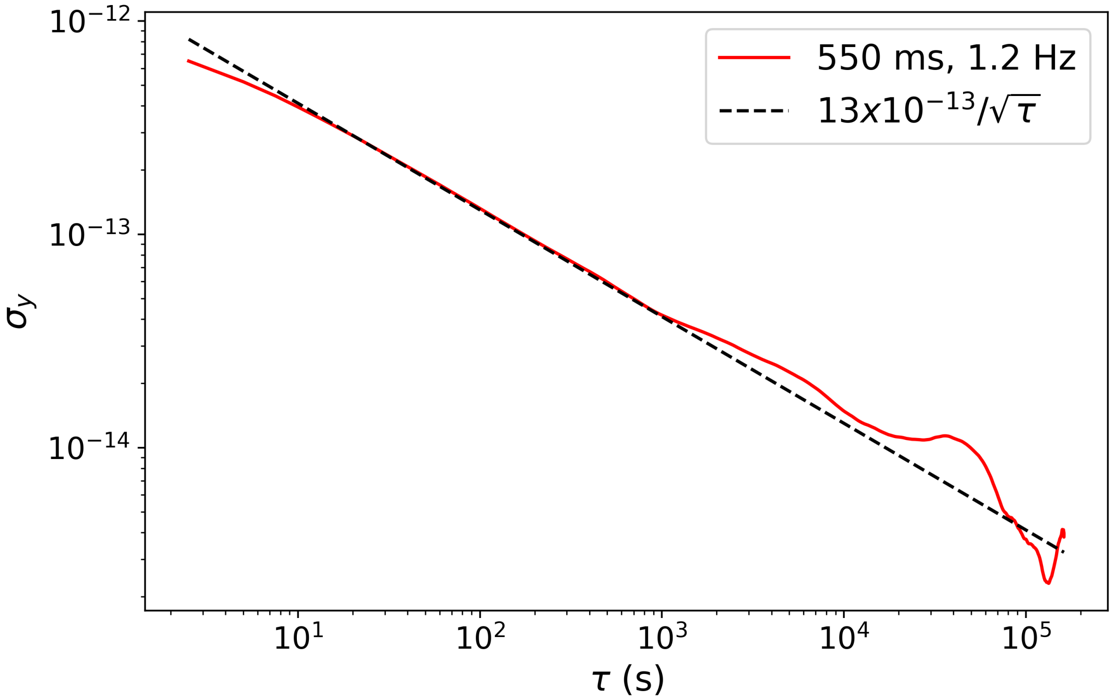

Figure 4.22: Allan deviation, $\sigma_y$, vs averaging time, $\tau$, for the frequency difference between the mini-fountain and the local hydrogen maser, for a launch time of $550\,\mathrm{ms}$ and a Ramsey fringe width of $1.2\,\mathrm{Hz}$. Error bars of size $\lesssim 10^{-15}$ are omitted for clarity.

The Allan deviation averages down with $\tau^{-1/2}$, as expected for a drift-free frequency standard. The short-term stability is improved by the presence of the shields to $13 \times 10^{-13}/\sqrt{\tau/\mathrm{s}}$, which can be further improved by analysing the signal-to-noise ratio (SNR).

The SNR due specifically to detection noise is optimised by varying the detection parameters; beam intensity, beam detuning, pulse time, repump time, background gap, and background time. From the previous experience with the demonstrator system and CsF3, the exact gap between the end of the detection pulse and the beginning of the background pulse is not critical, so long as the returning cold atoms have indeed left the beam region. The detection SNR is measured by taking a run of ~300 shots to calculate the shot-to-shot Allan deviation of the resulting transition probability, and dividing the Ramsey fringe contrast by this noise.

The beam intensity is determined by the RF power to the fibre AOMs; more RF power diffracts more optical power into the output. This RF power is controlled by a voltage controlled attenuator (VCA). The SNR vs VCA voltage is shown in figure 4.23. Here, the maximum control voltage of 5 V corresponds to an optical intensity of $3.8\,\mathrm{mW/cm^2}$ in each of the two horizontal beams only, as the preliminary results suggested that recruiting the diagonal beams to assist with "trapping" the returning atoms actually degrades the SNR. Perhaps this is because the cooling transition is already close to saturation with the horizontal beams alone, whereas the background fluorescence, and off-resonant pumping, are not. In the demonstrator system, there was no option but to use all 6 MOT beams. In CsF3, there were 4 horizontal beams, such that the additional 2 vertical beams made little difference.

The result is that there is no single "best" intensity to use: the SNR becomes flat after around 3.6 V. The peak at 4.8 V is dubious in its repeatability, but a perfectly fine choice none the less.

The beam frequency is determined by the frequency of the double-pass AOM in the spectroscopy section of the optics; the laser is locked above resonance by twice this AOM frequency. This frequency is controlled by adjusting the control voltage to the VCO, such that SNR vs beam detuning is shown in figure 4.24.

Here, the peak is much clearer: this somewhat resembles figure 4.4, but with the peak slightly closer to resonance. Thus, a VCO control voltage of 5.5 V seems the best choice, where the SNR is around 100.

In the previous tests in CsF3, it was found, perhaps counter-intuitively, that

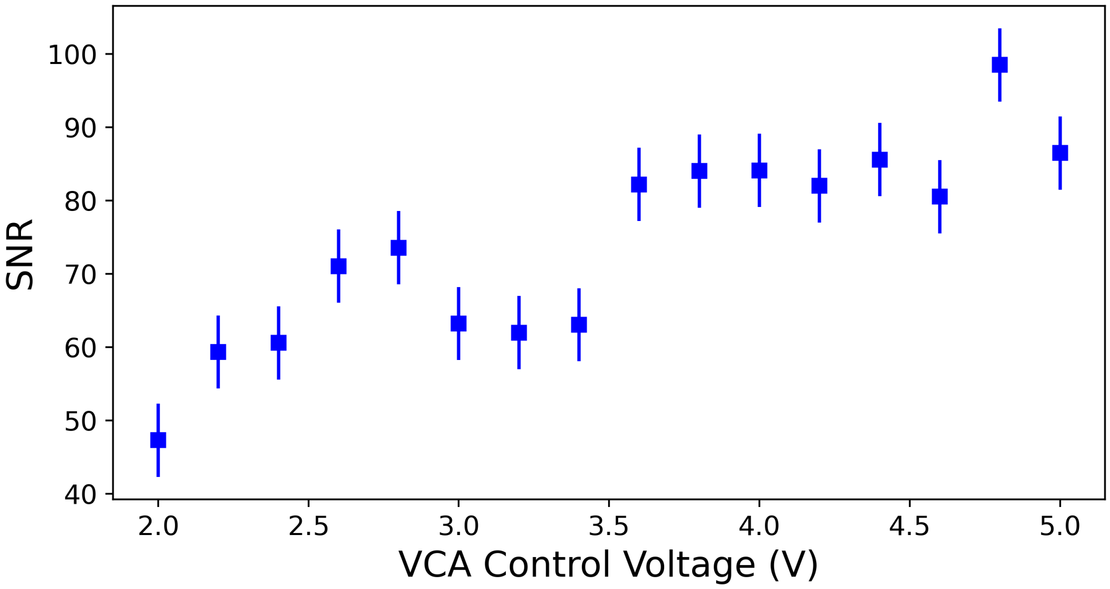

Figure 4.23: SNR vs attenuator control voltage, for a VCO voltage of 5.2 V and pulse lengths of 1.0 ms. The maximum beam intensity corresponds to 5.0 V.

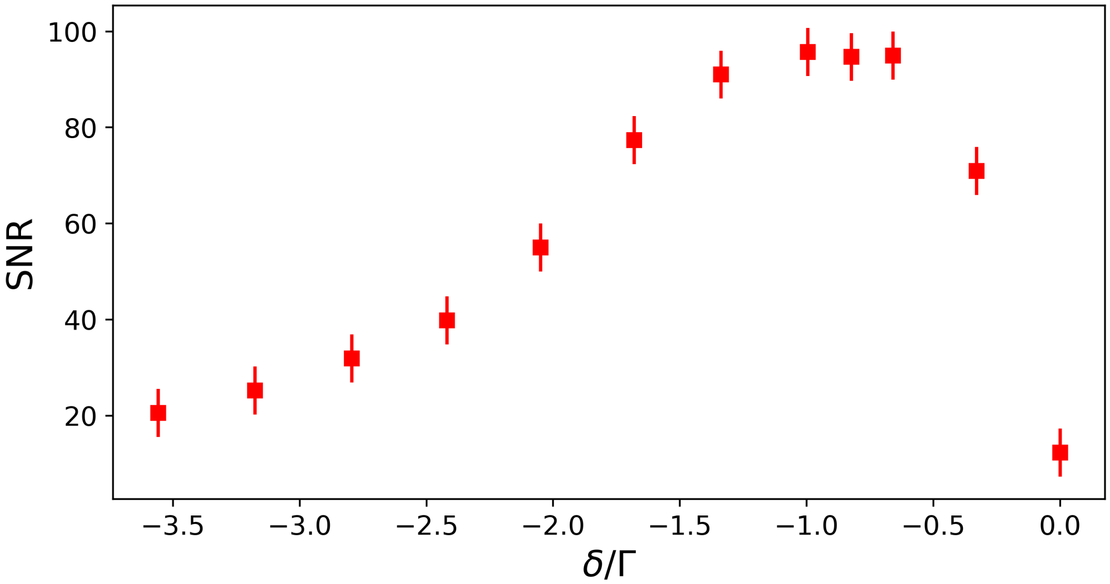

Figure 4.24: SNR vs detection beam detuning. The VCA voltage was 4.8 V and the pulse lengths 1.0 ms.

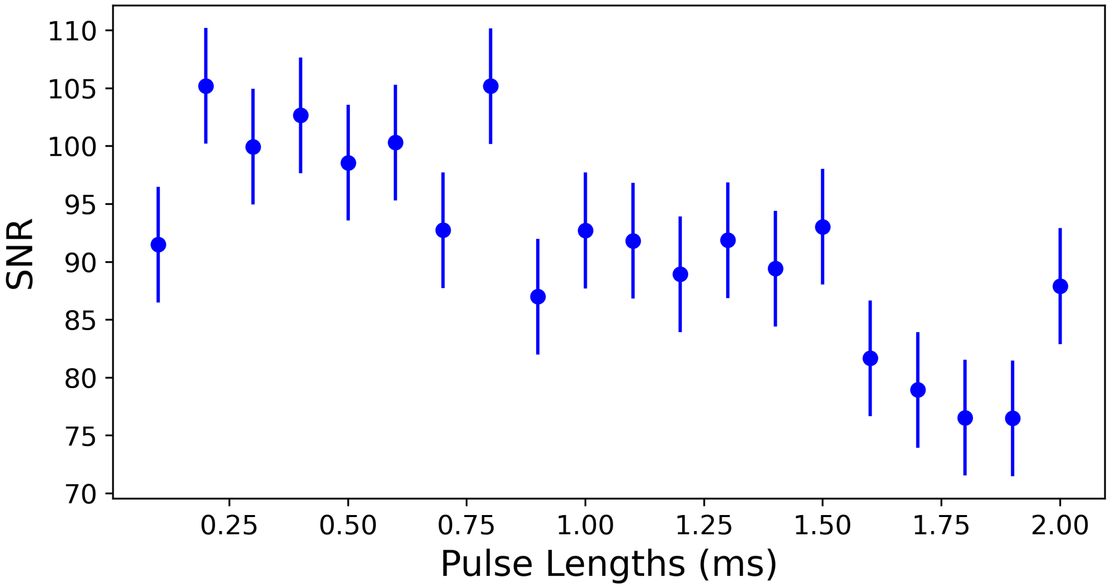

Figure 4.25: SNR vs detection (and background) pulse lengths, for a VCA voltage of 4.8 V and a VCO voltage of 5.5 V.

shorter detection pulses resulted in better SNR. So far, 1.0 ms has been used for each detection pulse and the background pulse (it is important to match the length of the background pulse to the detection pulses because of the large ringing in the photodiode signal). Using the current optimum values for intensity and detuning, the SNR vs pulse length is shown in figure 4.25.

Again, there is no clear best detection pulse length, only that shorter is perhaps slightly better. Even the shortest possible pulse, 0.1 ms, results in a reasonable SNR. This implies that the slight overshoot in the photodetector signal may not be as consistent as hoped, so probing it for longer gives more chance for the ringing to affect the total area under the detection curve Some example voltage vs time graphs are shown in figure 4.30. Thus, it is (somewhat arbitrarily) decided that 0.6 ms should be used going forwards, as it is a slight improvement over 1.0 ms while still maintaining most of the returning signal size.

Another avenue of investigation into the SNR is to vary the MOT loading time, to see how it behaves with varying numbers of atoms. As with sections 2.1.2 and 2.1.3, it is expected that SNR will initially rise linearly with increasing atom number, before reaching an upper limit. To vary the atom number, two strategies were employed: to vary the MOT loading time, and to forego state selection, pushing, and Ramsey interrogation. The former, in normal Ramsey operation, is sensitive to microwave phase and frequency noise, which is assumed negligible compared to the current levels of detection noise. The latter is not; without state-selection, pushing, or Ramsey interrogation, the SNR is due entirely to fluorescence detection. This also allows to more-easily access higher atom numbers. The signal-to-noise ratio vs returning fluorescence signal, $S_{\text{det}}$, is shown in figure 4.26. The result is that even with all the atoms returning (and no microwave noise) the SNR is limited to around 300, which is a lot less than was the case with CsF3 or the demonstrator system. This could be because of a different type of laser being used for the detection light (780 nm DBR vs 852 nm ECDL). In both CsF3 and the mini-fountain, the repump laser was a DFB laser. That the blue and red points in figure 4.26 appear to lie on the same curve implies that local oscillator noise is indeed negligible at this level.

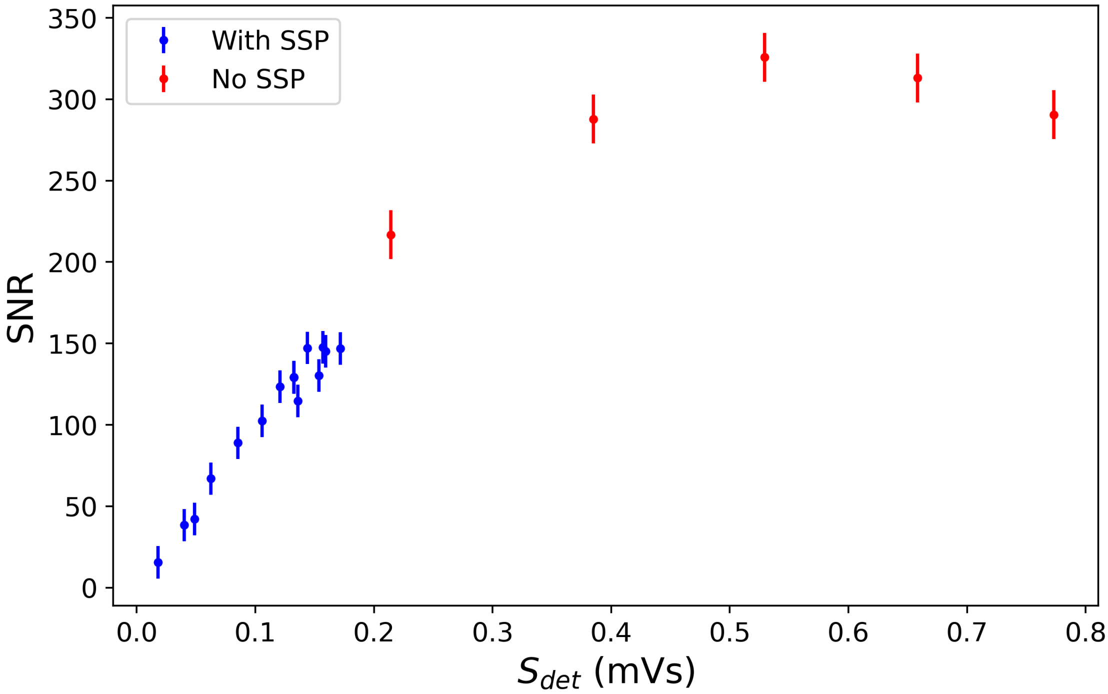

Figure 4.26: SNR vs $S_{\text{det}}$, for 550 ms launches and close-to-optimal detection parameters. Blue squares denote regular clock operation at the top of the central 1.2 Hz Ramsey fringe, with MOT loading times between 20 ms and 1.0 s. Red dots denote points where state selection, pushing, and Ramsey interrogation were absent, to reach higher returning atom numbers, while assuming the same (hypothetical) fringe contrast as for the normal data.

Normal operation with 0.5 s of MOT loading with a dispenser current of 4.4 A corresponds to around 0.13 mVs of signal, and an SNR of just over 100. Since this is firmly in the initial linear section of figure 4.26, more returning atoms would clearly be at least slightly beneficial. This could be achieved with a higher dispenser current. On the one hand, a higher background density of $^{87}$Rb atoms should result in a higher MOT loading rate $R_L$. and thus more loaded atoms (so long as the loss time constant $\tau_L \gtrsim 0.5\,\mathrm{s}$). On the other hand, an increased background density of both Rb isotopes will increase the collisional loss during the atom's free-flight, such that one may expect to reach a dispenser current where the additional loss outdoes the increased loading rate, and the returning signal reaches a maximum. Plots of returning signal and SNR vs dispenser current are shown in figures 4.27 and 4.28.

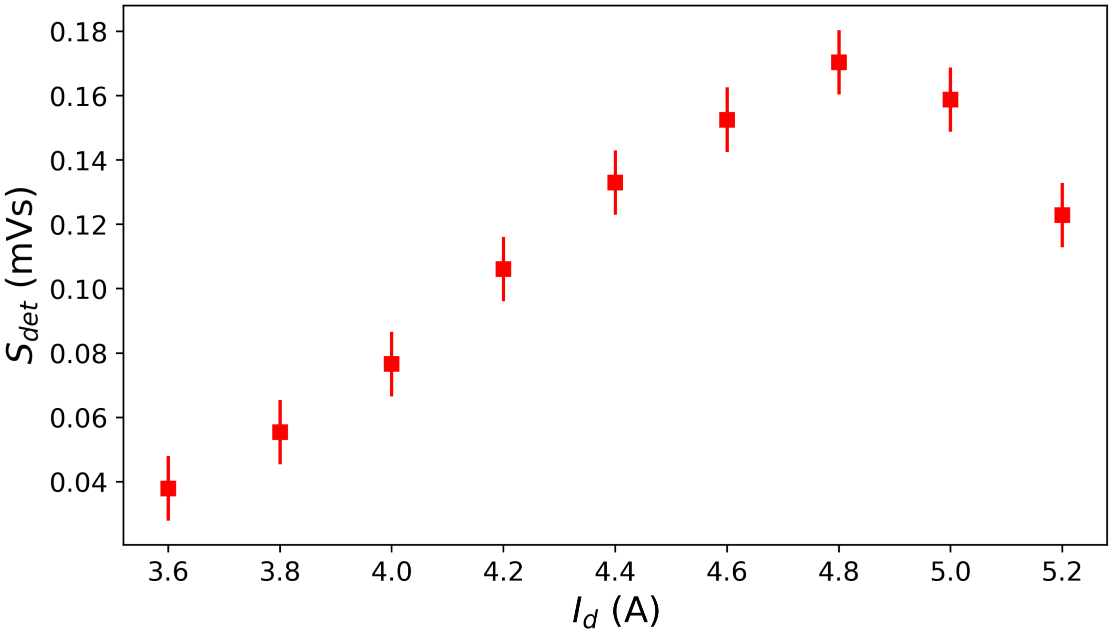

Figure 4.27: Returning signal $S_{\mathrm{det}}$ vs dispenser current $I_d$, for 0.5 s of MOT loading.

Figure 4.27 shows that increasing the dispenser current can only improve the returning signal to a limited extent; a peak is visible at 4.8 A, after which more background gas contributes more to collisional loss than it does to MOT loading. However, figure 4.28 shows that there is also a dispenser current that corresponds to a peak SNR; at 4.4 A the SNR was 126. This is because the background $^{87}$Rb also fluoresces during detection, so even if more background gas doesn't result in more cold-atom signal, it can still result in more noise. Thus, the optimum dispenser current was already being used to obtain the frequency stability data of figure 4.22

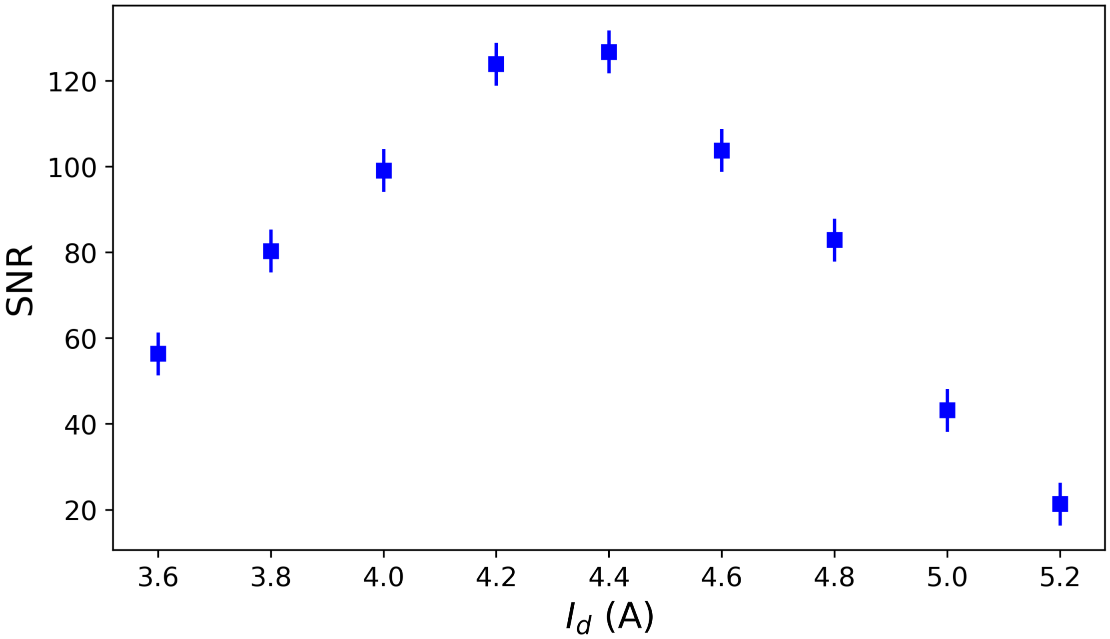

Figure 4.28: SNR vs dispenser current $I_d$, for 0.5 s of MOT loading.

and SNR data of figure 4.26.

Another tool for improving the clock's short-term stability is to vary the launch height, as a shorter launch should result in more returning signal and thus more SNR, even if this may be at the expense of a broader linewidth, $\Delta\nu_{\text{Ramsey}}$. The clock's short-term stability can be calculated from

$$\sigma_y(\tau) = \frac{2}{\pi}\frac{\Delta\nu_{\text{Ramsey}}}{\nu_0}\frac{1}{\text{SNR}}\sqrt{\frac{T_c}{\tau}}, \tag{4.10}$$

which is a just different form of equation 1.46. The relevant parameters obtained for a few different launch times are shown in table 4.1.

| $t_L$(ms) | $\Delta\nu_{\text{Ramsey}}$ (Hz) | $S_{\text{det}}$ (mVs) | SNR | Calculated $\sigma_y(\tau = 1\,\text{s})$ |
|---|---|---|---|---|
| 400 | 2.3 | 0.177 | 180 | $1.3 \times 10^{-12}$ |
| 450 | 1.7 | 0.171 | 160 | $1.1 \times 10^{-12}$ |
| 500 | 1.3 | 0.135 | 180 | $7 \times 10^{-13}$ |
| 550 | 1.2 | 0.131 | 120 | $1.0 \times 10^{-12}$ |
| 600 | 1.1 | 0.064 | 100 | $1.1 \times 10^{-12}$ |

Table 4.1: Ramsey fringe width and SNR (measured on the side-of-fringe) obtained for a few different launch times.

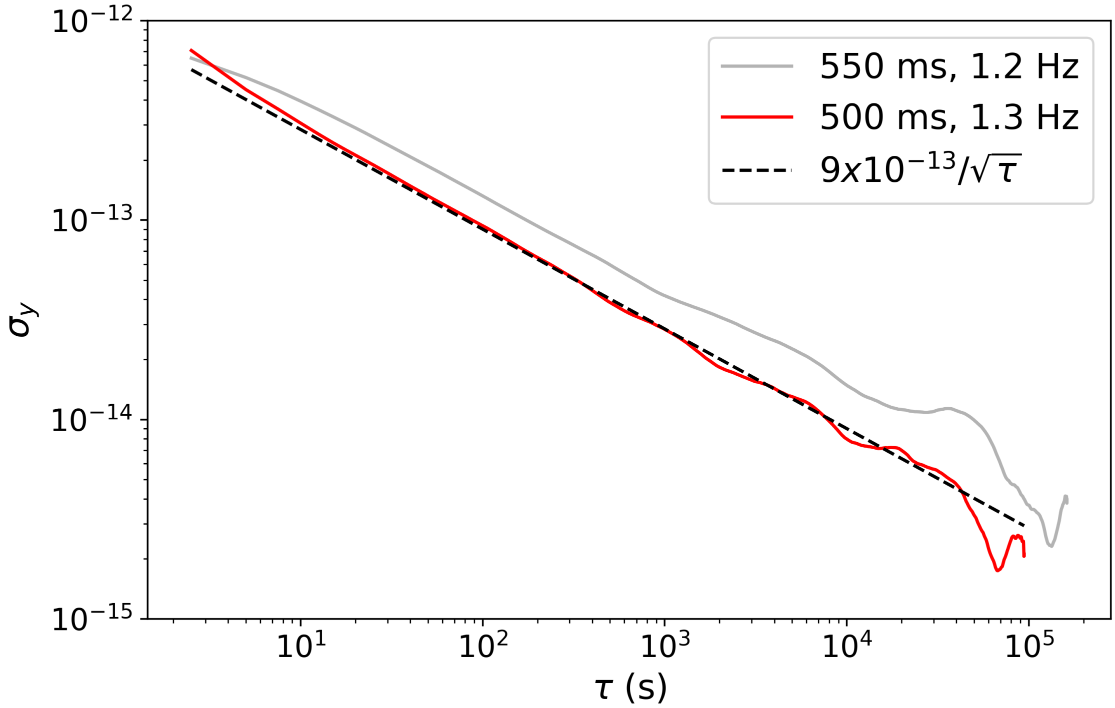

Figure 4.29: $\sigma_y$ vs $\tau$ when measuring the local oscillator frequency when referenced to the maser, for both 500 ms and 550 ms launch times. Error bars of size $\lesssim 10^{-15}$ are omitted for clarity.

From table 4.1, it is clear that using a 500 ms launch should result in an improved frequency stability, as, despite resulting in a similar $S_{\text{det}}$ to 550 ms, the SNR is increased by 50% while only degrading the linewidth by 8%. That the linearity between SNR and $S_{\text{det}}$ is not maintained at different launch heights is curious: perhaps a different optimum combination of detection intensity, detuning, and pulse time exists for each launch height.

The clock was left to operate at this shorter launch time for several days, and the resulting Allan deviation is shown in figure 4.29. This shows a clear, if slight, improvement in the clock's short-term stability, to $\sigma_y = 9 \times 10^{-13}/\sqrt{\tau/\text{s}}$. The inconsistency between the calculated $\sigma_y$ and the measured $\sigma_y$ could be due to the returning atom number varying slowly, over the course of minutes and hours, particularly during working hours, resulting in a time-varying SNR.

## 4.8 SNR Analysis

So far, adjusting the clock's operating parameters has only made incremental improvements in the frequency stability. Measurements of signal-to-noise ratio vs $S_{\text{det}}$ (figure 4.26) show that a higher number of returning atoms (e.g. by a colder atom temperature or longer MOT loading time) could make a slight improvement, however the best one could do with the current system seems to be SNR $\approx 300$, likely not approaching the limit due to local oscillator noise.

The inverse of the SNR is the quadrature sum of multiple components: electrical noise in the detection system $\sigma_{\text{elec}}$, due to instability in the photodiode bias voltage and amplification; noise from background light scattered from the walls of the chamber, $\sigma_{\text{bg,c}}$, which will depend on the cooling laser's amplitude stability; noise from background light scattered from the thermal $^{87}$Rb vapour, $\sigma_{\text{bg,t}}$, which will depend on the cooling laser's amplitude stability and the background gas density; amplitude and frequency noise from the repump laser, $\sigma_{\text{rp}}$, which will effect the efficiency of the repump pulse the occurs between the two detection pulses; and frequency noise from the cooling laser $\sigma_{c,\nu}$, which will affect the fluorescence from the returning cold atoms, but not from the background.

Quantum projection noise has been neglected. Typically, the detected voltage due to the returning cold atoms is 0.08 V, such that, from equation 4.1, an estimate of the returning atom number is $N_{at} \approx 3 \times 10^6$, such that the quantum limit for our SNR should be $\sqrt{N_{\text{at}}} \approx 1,700$: much higher than the limit due to other noise sources.

To measure $\sigma_{\text{elec}}$, $\sigma_{\text{bg,c}}$, $\sigma_{\text{bg,t}}$, $\sigma_{\text{rp}}$, and $\sigma_{c,\nu}$, fluorescence data can be analysed in several regimes: $\sigma_{\text{elec}}$ can be isolated by shutting off all laser light, and manually adding in an artificial cold-atom signal of appropriate size to the resulting series of voltage vs time plots, such that the "transition probability" is perturbed only by electrical noise inherent to our measurement device. The same can be done, but with the cooling laser on-but-detuned, such that the metal walls of the chamber, but not the background $^{87}$Rb vapour, are scattering. The resulting SNR here is $1/\sqrt{\sigma_{\text{elec}}^2 + \sigma_{\text{bg,c}}^2}$. The same can be done with the cooling laser locked to resonance (but no repump light), to get SNR $= 1/\sqrt{\sigma_{\text{elec}}^2 + \sigma_{\text{bg,c}}^2 + \sigma_{\text{bg,t}}^2}$. And finally, analysis of the real cold-atom signal results in the SNR due to the quadrature sum of all terms. This is shown in figure 4.30

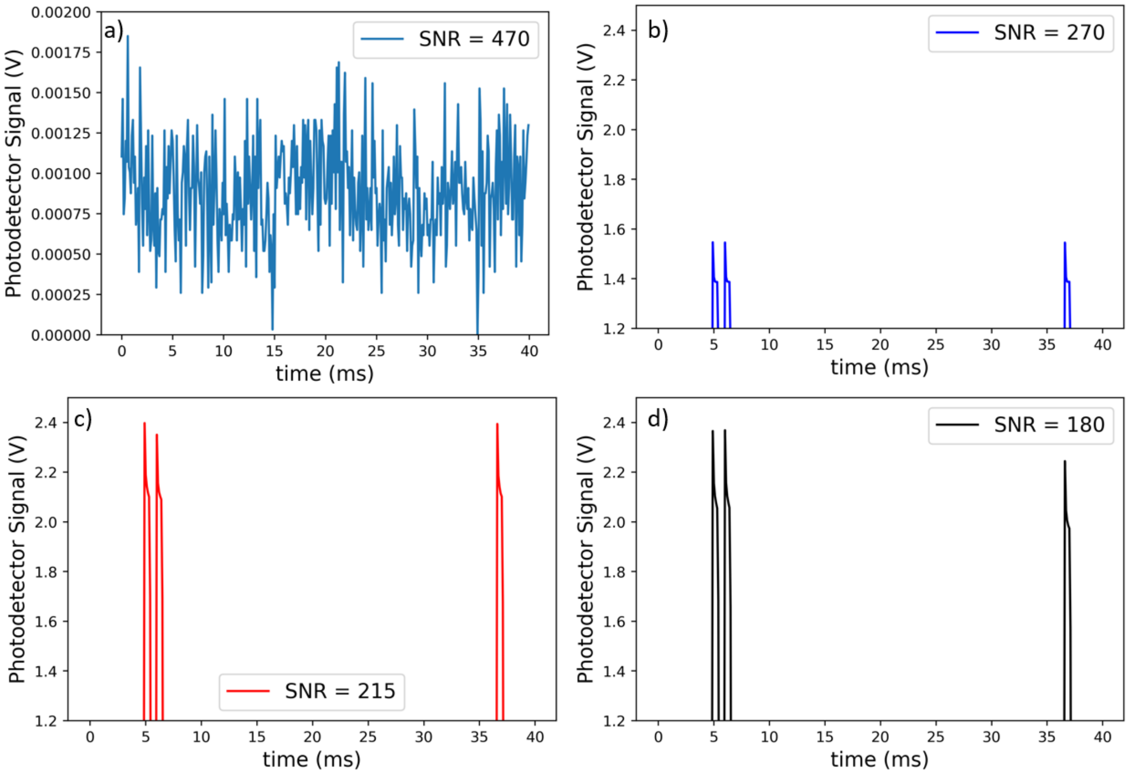

Figure 4.30: Examples of voltage vs time plots for a) no laser light, b) detuned cooling laser light, c) resonant cooling laser light, with the resulting SNRs from adding a 0.1 mVs artificial signal and d) real cold-atom signal with the actually-measured SNR.

From the SNR results of figure 4.30, we have

$$1/\sigma_{\text{elec}} = 470, \tag{4.11}$$

$$\sqrt{\frac{1}{\sigma^2_{\text{elec}}} + \frac{1}{\sigma^2_{\text{bc,g}}}} = 270, \tag{4.12}$$

$$\sqrt{\frac{1}{\sigma^2_{\text{elec}}} + \frac{1}{\sigma^2_{\text{bc,g}}} + \frac{1}{\sigma^2_{\text{bc,t}}}} = 215, \tag{4.13}$$

$$\sqrt{\frac{1}{\sigma^2_{\text{elec}}} + \frac{1}{\sigma^2_{\text{bc,g}}} + \frac{1}{\sigma^2_{\text{bc,t}}} + \frac{1}{\sigma^2_{\text{rp}}} + \frac{1}{\sigma^2_{c,\nu}}} = 180, \tag{4.14}$$

$$\tag{4.15}$$

such that

$$1/\sigma_{\mathrm{elec}} = 470, \tag{4.16}$$

$$1/\sigma_{\mathrm{bg,c}} = 330, \tag{4.17}$$

$$1/\sigma_{\mathrm{bg,t}} = 355, \tag{4.18}$$

$$\sqrt{\frac{1}{\sigma_{\mathrm{rp}}^2} + \frac{1}{\sigma_{c,\nu}^2}} = 330. \tag{4.19}$$

The quadrature sum of the cold-atom specific terms, $\sigma_{\mathrm{rp}}$ and $\sigma_{c,\nu}$, is around 300, thus is consistent with the high-signal limit of figure 4.26. Cooling laser intensity noise is presumed to be negligible here because the total intensity of the beams is $I \gtrsim I_{\mathrm{sat}}$, such that small intensity fluctuations should not result in similar-size fluctuation in the light scattered by the cold atoms. Separating out the contributions due to the repump and cooling lasers in the real cold-atom noise may require some more thought, although it is likely that $\sigma_{\mathrm{rp}}$ is dominating over $\sigma_{c,\nu}$, because of sub-optimal repump-laser lock settings. The background fluorescence terms are each very similar, as they each have their origins predominantly in laser intensity noise. The chamber walls result in 1.4 V of detected background level, and the thermal atoms an additional 0.7 V. As expected, electrical-noise-only results in the highest SNR, although the detector ringing, which is a large feature of all the other plots, is absent from our simulated “perfect” cold-atom signal.

The clock’s performance is currently limited by noise specific to the detection method. To achieve further improvements in the clock’s detection SNR, and thus frequency stability, will require further investigation into the detection electronics, and the light scattered by the thermal background vapour.

# Conclusion

The aim of the work reported in this thesis has been to design, build, and test a prototype mini-fountain cold-atom clock physics package. This has been successful: a plan was devised to reduce the size, complexity and cost of the physics package; a new method of performing normalised detection with the MOT beams was implemented; and a novel all-in-fibre optical control system improved the robustness of the system. Furthermore, the Ramsey cavity was successfully integrated into the vacuum structure of the clock using a novel design and a novel epoxy sealing method. The clock was constructed and baked-out to achieve a residual vacuum pressure so low that the ion pump struggles to read a non-zero pressure. Atoms are cooled in our chamber to only a few microkelvin above absolute zero; Ramsey fringes with a linewidth similar to those of a full-sized Primary Frequency Standard have been achieved; and the clock's frequency stability has been measured for several days of continuous operation, limited only by the need to interrupt the clock operation to perform more tests.

The initial tests of the new detection method were reported in [214], and subsequent advances in the project are being reported in [215], with a view to publish an accuracy evaluation in the near future.

The prototype mini-fountain physics package is shown in figure 4.31, next to the current-generation of caesium fountain clocks, for scale.

The short-term stability of the clock still has scope for improvement: tests in chapter 2 suggested that the compact detection method with the MOT beams could achieve a signal-to-noise ratio of around 600, whereas in practice the best result so far is just 180. Potentially this situation could be improved by adding some form of intensity-stabilisation to our all-in-fibre optical control system, to reduce the noise associated with background fluorescence, and by revisiting the MOT beam collimators to improve the trap diameter and atom temperature, and thus improving

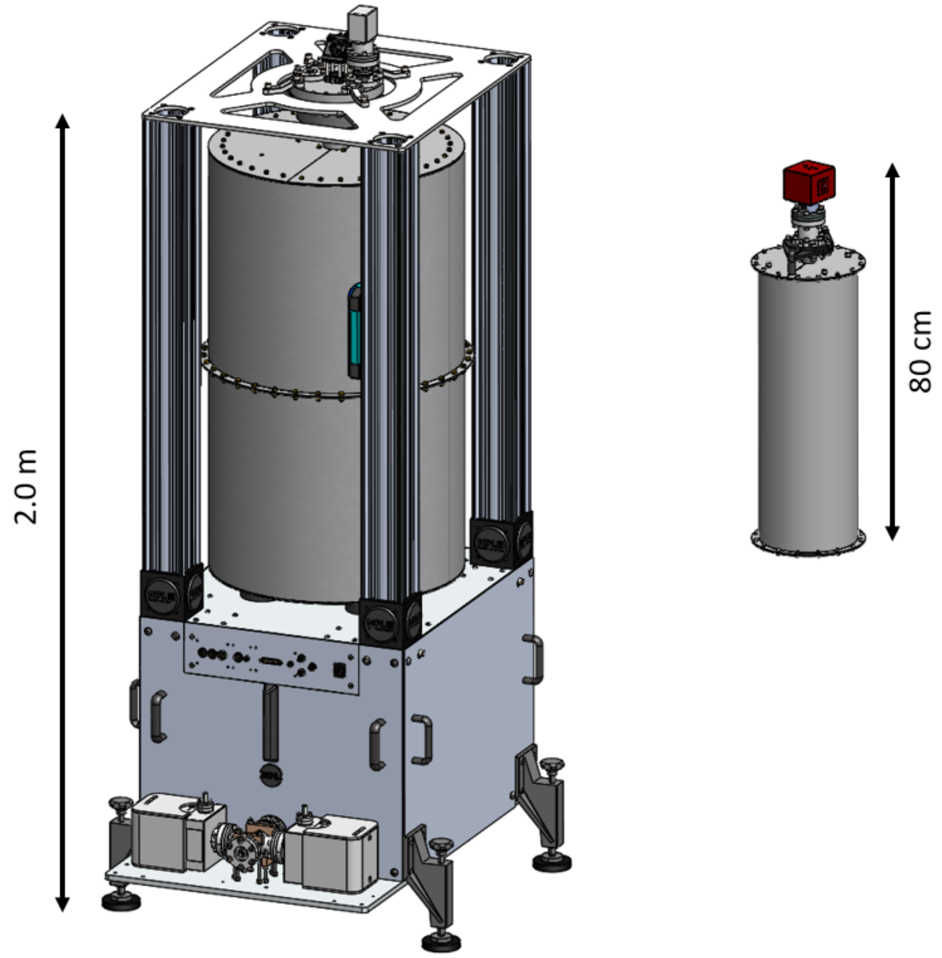

Figure 4.31: Size comparison between the current generation of caesium fountain clocks (left) and the prototype mini-fountain (right). CAD designs courtesy of Peter Lovelock, NPL.

the returning atom number, or even investigating the use of high-power frequency-doubled telecoms lasers, to exploit all of the technical development that has gone into 1550 nm lasers. Different photodetector/amplifier combinations may also allow for a reduction in ringing and electrical noise.

With an improved short-term frequency stability, systematic shift measurements below the $10^{-15}$ level become feasible. Part of the attractiveness of the fountain design is the accuracy with which the systematic shifts can be measured and accounted for. Thus, the major systematic effects, such as the second-order Zeeman shift, and blackbody shift, should be measured. The mini-fountain is expected to be capable of achieving sufficient long-term stability such that the more subtle shifts, such as

DCP, cavity pulling, and collisional shift, can be evaluated.

While the fountain physics package itself has been greatly reduced in size, currently some of the optics and most of the electronics take up similar sized tables and racks to the full-sized Cs fountains. Reducing the size of these axillary systems is mostly and engineering problem, rather than an atomic physics problem, hence why the focus has been on the physics package. Part of the plan for the optics was to remove the additional repump laser and replace it with a sideband on the main laser, generated with an electro-optic modulator. In section 2.5.4, it was reported that sufficient power was available out of the EOM to seed the tapered amplifier, yet at the beginning of chapter 4.1, it is noted that the DFB laser is indeed being used for the repump: about one year had passed between the tests of the EOM and the completion of the physics package, and the EOM behaved perfectly for about 99% of that time, before succumbing at the last moment before victory to damage due to the high optical throughput power. It may still be possible to replace the DFB laser with a sideband, perhaps with a different arrangement the doesn't see the EOM used as the TA seed, which opens up more possible combinations of optical throughput power and microwave modulation power.

If the stability, accuracy, and compactness of the mini-fountain can be addressed, it would be competitive in price and performance with existing commercial cold-atom clocks, and at a time where the trade-offs used to miniaturise the system become less and less critical. Caesium fountain primary frequency standards are becoming commercial products in their own right, particularly at NPL: any potential customers who demand $10^{-16}$ level accuracy, and are willing to pay for it, can have it. Similarly, the redefinition of the SI second in terms of an optical transition will mean that, one day in the not too distant future, microwave clocks will no longer be primary frequency standards, and the task of steering timescales will be carried out by optical clock physicists. This leaves more scope for the efforts of microwave clock physicists to be directed towards engineering efficiency and robustness into their systems.

# Bibliography

[1] Conference Générale des Poids et Measures and Comité International des Poids et Measures. Definitions of basic SI units (translated from the french). *Metrologia*, 4(3):147–147, 1968.

[2] M. A. Lombardi. An evaluation of dependencies of critical infrastructure timing systems on the global positioning system (GPS). 2021.

[3] U. Kumar. Synchronisation methods for the converging telecom and data centre infrastructures. In *2023 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS)*, pages 1–6, 2023.

[4] J. Levine. Timing in telecommunications networks. *Metrologia*, 48(4):S203, 2011.

[5] T. Saidkhodjaev, J. Voas, R. Kuhn, J. DeFranco, and P. Laplante. Aggregating atomic clocks for time-stamping. In *2020 IEEE International Conference on Service Oriented Systems Engineering (SOSE)*, pages 1–6, 2020.

[6] S. Reynaud, C. Salomon, and P. Wolf. Testing general relativity with atomic clocks. *Space Science Reviews*, 148(1-4):233–247, 2009.

[7] J. D. Prestage, R. L. Tjoelker, and L. Maleki. Atomic clocks and variations of the fine structure constant. *Phys. Rev. Lett.*, 74:3511–3514, 1995.

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

[9] J. P. Uzan. Varying constants, gravitation and cosmology. *Living Reviews in Relativity*, 14(1):2, 2011.

[10] Microchip Technology Inc. MHM-2020 active hydrogen maser. 2024.

[11] D. W. Allan, A. Lepek, L. Cutler, R. Giffard, and J. Kusters. The impact of the HP 5071A on international atomic time. In *Proc. 27th Annual PTTI*, pages pp. 235–256, 1995.

[12] Muquans. MuClock - a high-performance frequency standard based on cold atoms. 2019.

[13] Spectra Dynamics Inc. cRb-clock datasheet. 2023.

[14] A. Tofful, C. F. A. Baynham, E. A. Curtis, A. O. Parsons, B. I. Robertson, M. Schioppo, J. Tunesi, H. S. Margolis, R. J. Hendricks, J. Whale, R. C. Thompson, and R. M. Godun. $^{171}\mathrm{Yb}^+$ optical clock with $2.2\times10^{-18}$ systematic uncertainty and absolute frequency measurements. *Metrologia*, 61(4):045001, 2024.

[15] D. Rooney and J. Nye. Greenwich observatory time for the public benefit: Standard time and victorian networks of regulation. *The British Journal for the History of Science*, 42(1):5–30, 2009.

[16] H. M. Smith. International time and frequency coordination. *Proceedings of the IEEE*, 60(5):479–487, 1972.

[17] I. I. Rabi. On the process of space quantization. *Phys. Rev.*, 49:324–328, 1936.

[18] I. I. Rabi. Space quantization in a gyrating magnetic field. *Phys. Rev.*, 51:652–654, 1937.

[19] J. M. B. Kellogg, I. I. Rabi, and J. R. Zacharias. The gyromagnetic properties of the hydrogens. *Phys. Rev.*, 50:472–481, 1936.

[20] I. I. Rabi, S. Millman, P. Kusch, and J. R. Zacharias. The molecular beam resonance method for measuring nuclear magnetic moments. the magnetic moments of ${}_3\mathrm{Li}^6$, ${}_3\mathrm{Li}^7$ and ${}_9\mathrm{F}^{19}$. *Phys. Rev.*, 55:526–535, 1939.

[21] E. Majorana. Atomi orientati in campo magnetico variabile. *Il Nuovo Cimento (1924-1942)*, 9(2):43–50, 1932.

[22] P. Kusch, S. Millman, and I. I. Rabi. The radiofrequency spectra of atoms hyperfine structure and Zeeman effect in the ground state of $Li^6$, $Li^7$, $K^{39}$ and $K^{41}$. *Phys. Rev.*, 57:765–780, 1940.

[23] N. F. Ramsey. A molecular beam resonance method with separated oscillating fields. *Phys. Rev.*, 78:695–699, 1950.

[24] C. J. Foot. *Atomic physics.* Oxford master series in atomic, optical, and laser physics. Oxford University Press, Oxford, 2007.

[25] L. Essen and J. V. L. Parry. An atomic standard of frequency and time interval: A cæsium resonator. *Nature*, 176(4476):280–282, 1955.

[26] G. Breit and I. I. Rabi. Measurement of nuclear spin. *Phys. Rev.*, 38:2082–2083, 1931.

[27] S. Goudsmit and R. F. Bacher. The Paschen-Back effect of hyperfine structure. *Phys. Rev.*, 34:1499–1500, 1929.

[28] L. Essen and J. V. L. Parry. The caesium resonator as a standard of frequency and time. *Phil. Trans. R. Soc. A*, 250(973):45–69, 1957.

[29] F. G. Major. *The Quantum Beat: The Physical Principles of Atomic Clocks.* Springer, 1998.

[30] A. Bauch. The PTB primary clocks CS1 and CS2. *Metrologia*, 42(3):S43–S54, 2005.

[31] A. Kastler. Quelques suggestions concernant la production optique et la détection optique d'une inégalité de population des niveaux de quantifigation spatiale des atomes. Application à l'expérience de Stern et Gerlach et à la résonance magnétique. *J. Phys. Radium*, 11(6):255–265, 1950.

[32] G. Avila, V. Giordano, V. Candelier, E. de Clercq, G. Theobald, and P. Cerez. State selection in a cesium beam by laser-diode optical pumping. *Phys. Rev. A*, 36:3719–3728, 1987.

[33] J. Vanier and C. Audoin. The classical caesium beam frequency standard: fifty years later. *Metrologia*, 42(3):S31–S42, 2005.

[34] L. L. Lewis, F. L. Walls, and D. J. Glaze. Design considerations and performance of NBS-6, the NBS primary frequency standard. *J. Phys. Colloques*, 42:C8–241–C8–246, 1981.

[35] K. Hagimoto, S. Ohshima, Y. Nakadan, and Y. Koga. Accuracy evaluation of the optically pumped Cs frequency standard at NRLM. *IEEE Trans. Instrum. Meas.*, 48(2):496–499, 1999.

[36] A. Makdissi and E. De Clercq. Evaluation of the accuracy of the optically pumped caesium beam primary frequency standard of BNM-LPTF. *Metrologia*, 38(5):409–425, 2001.

[37] J. H. Shirley, W. D. Lee, and R. E. Drullinger. Accuracy evaluation of the primary frequency standard NIST-7. *Metrologia*, 38(5):427–458, 2001.

[38] T. W. Hänsch and A. L. Schawlow. Cooling of gases by laser radiation. *Optics Communications*, 13(1):68–69, 1975.

[39] P. J. Ungar, D. S. Weiss, E. Riis, and S. Chu. Optical molasses and multilevel atoms: theory. *J. Opt. Soc. Am. B*, 6(11):2058–2071, 1989.

[40] D. S. Weiss, E. Riis, Y. Shevy, P. J. Ungar, and S. Chu. Optical molasses and multilevel atoms: experiment. *J. Opt. Soc. Am. B*, 6(11):2072–2083, 1989.

[41] S. Stenholm. The semiclassical theory of laser cooling. *Rev. Mod. Phys.*, 58:699–739, 1986.

[42] A. M. Steane, M. Chowdhury, and C. J. Foot. Radiation force in the magneto-optical trap. *J. Opt. Soc. Am. B*, 9(12):2142–2158, 1992.

[43] D. M. Brink and G. R. Satchler. *Angular Momentum, by D.M. Brink and G.R. Satchler.* Oxford library of the physical sciences. Clarendon Press, 1962.

[44] P. D. Lett, W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook. Optical molasses. *J. Opt. Soc. Am. B*, 6(11):2084–2107, 1989.

[45] P. D. Lett, R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould, and H. J. Metcalf. Observation of atoms laser cooled below the Doppler limit. *Phys. Rev. Lett.*, 61:169–172, 1988.

[46] J. Dalibard and C. Cohen-Tannoudji. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *J. Opt. Soc. Am. B*, 6(11):2023–2045, 1989.

[47] J. Dalibard, S. Reynaud, and C. Cohen-Tannoudji. Potentialities of a new $\sigma^+\sigma^-$ laser configuration for radiative cooling and trapping. *Journal of Physics B: Atomic and Molecular Physics*, 17(22):4577–4594, 1984.

[48] T. Weber, J. Herbig, M. Mark, H. C. Nägerl, and R. Grimm. Bose-Einstein condensation of cesium. *Science*, 299(5604):232–235, 2003.

[49] C. Monroe, W. Swann, H. Robinson, and C. Wieman. Very cold trapped atoms in a vapor cell. *Phys. Rev. Lett.*, 65:1571–1574, 1990.

[50] A. Cable, M. Prentiss, and N. P. Bigelow. Observations of sodium atoms in a magnetic molasses trap loaded by a continuous uncooled source. *Opt. Lett.*, 15(9):507–509, 1990.

[51] S. L. Kristensen, E. Bohr, J. Robinson-Tait, T. Zelevinsky, J. W. Thomsen, and J. H. Müller. Subnatural linewidth superradiant lasing with cold $^{88}$Sr atoms. *Phys. Rev. Lett.*, 130:223402, 2023.

[52] G. M. Tino, M. Barsanti, M. de Angelis, L. Gianfrani, and M. Inguscio. Spectroscopy of the 689 nm intercombination line of strontium using an extended-cavity InGaP/InGaAIP diode laser. *Applied Physics B*, 55(4):397–400, 1992.

[53] V. V. Vassiliev, S. A. Zibrov, and V. L. Velichansky. Compact extended-cavity diode laser for atomic spectroscopy and metrology. *Review of Scientific Instruments*, 77(1):013102, 2006.

[54] S. Wang. Principles of distributed feedback and distributed Bragg-reflector lasers. *IEEE Journal of Quantum Electronics*, 10(4):413–427, 1974.

[55] O. Ishida, H. Toba, and Y. Tohmori. Pure frequency modulation of a multielectrode distributed-Bragg-reflector (DBR) laser diode. *IEEE Photonics Technology Letters*, 1(7):156–158, 1989.

[56] P. H. Lee and M. L. Skolnick. Saturated neon absorption inside a 6238-$\AA$ laser. *Applied Physics Letters*, 10(11):303–305, 2004.

[57] R. L. Barger and J. L. Hall. Pressure shift and broadening of methane line at 3.39 $\mu$ studied by laser-saturated molecular absorption. *Phys. Rev. Lett.*, 22:4–8, 1969.

[58] T. W. Hänsch, M. D. Levenson, and A. L. Schawlow. Complete hyperfine structure of a molecular iodine line. *Phys. Rev. Lett.*, 26:946–949, 1971.

[59] S. Haroche and F. Hartmann. Theory of saturated-absorption line shapes. *Phys. Rev. A*, 6:1280–1300, 1972.

[60] J. L. Hall, M. Zhu, and P. Buch. Prospects for using laser-prepared atomic fountains for optical frequency standards applications. *J. Opt. Soc. Am. B*, 6(11):2194–2205, 1989.

[61] M. A. Kasevich, E. Riis, S. Chu, and R. G. DeVoe. RF spectroscopy in an atomic fountain. *Phys. Rev. Lett.*, 63:612–615, 1989.

[62] C. Salomon, J. Dalibard, W. D. Phillips, A. Clairon, and S. Guellati. Laser cooling of cesium atoms below 3 $\mu$K. *Europhysics Letters (EPL)*, 12(8):683–688, 1990.

[63] A. Clairon, C. Salomon, S. Guellati, and W. D. Phillips. Ramsey resonance in a Zacharias fountain. *Europhysics Letters (EPL)*, 16(2):165–170, 1991.

[64] A. Clairon, P. Laurent, G. Santarelli, S. Ghezali, S. N. Lea, and M. Bahoura. A cesium fountain frequency standard: preliminary results. *IEEE Trans. Instrum. Meas.*, 44(2):128–131, 1995.

[65] A. N. Luiten, A. G. Mann, M. E. Costa, and D. G. Blair. Power stabilized cryogenic sapphire oscillator. *IEEE Trans. Instrum. Meas.*, 44(2):132–135, 1995.

[66] G. Santarelli, Ph. Laurent, P. Lemonde, A. Clairon, A. G. Mann, S. Chang, A. N. Luiten, and C. Salomon. Quantum projection noise in an atomic fountain: A high stability cesium frequency standard. *Phys. Rev. Lett.*, 82:4619–4622, 1999.

[67] S. Weyers, V. Gerginov, M. Kazda, J. Rahm, B. Lipphardt, G. Dobrev, and K. Gibble. Advances in the accuracy, stability, and reliability of the PTB primary fountain clocks. *Metrologia*, 55(6):789–805, 2018.

[68] K. Szymaniec, S. N. Lea, K. Gibble, S. E. Park, K. Liu, and P. Głowacki. NPL Cs fountain frequency standards and the quest for the ultimate accuracy. *Journal of Physics: Conference Series*, 723:012003, 2016.

[69] J. Guena, M. Abgrall, D. Rovera, P. Laurent, B. Chupin, M. Lours, G. Santarelli, P. Rosenbusch, M. E. Tobar, R. Li, K. Gibble, A. Clairon, and S. Bize. Progress in atomic fountains at LNE-SYRTE. *IEEE Trans. UFFC*, 59(3):391–409, 2012.

[70] L. Devenoges, A. Stefanov, A. Joyet, P. Thomann, and G. D. Domenico. Improvement of the frequency stability below the Dick limit with a continuous atomic fountain clock. *IEEE Trans. UFFC*, 59(2):211–216, 2012.

[71] V. Gerginov, N. Nemitz, S. Weyers, R. Schröder, D. Griebsch, and R. Wynands. Uncertainty evaluation of the caesium fountain clock PTB-CSF2. *Metrologia*, 47(1):65–79, 2009.

[72] E. A. Donley, T. P. Heavner, J. W. O'Brien, S. R. Jefferts, and F. Levi. Laser cooling and launching performance in a (1,1,1)-geometry atomic fountain. In *Proc. 2005 IEEE IFCS and Exposition.*, page 5 pp, 2005.

[73] A. Clairon, S. Ghezali, G. Santarelli, P. Laurent, E. Simon, S. Lea, M. Bahoura, S. Weyers, and K. Szymaniec. The LPTF preliminary accuracy evaluation of cesium fountain frequency standard. In *Tenth EFTF 96 (IEE Conf. Publ. 418)*, pages 218–223, 1996.

[74] S. Weyers, A. Bauch, U. Hubner, R. Schroder, and C. Tamm. First performance results of PTB's atomic caesium fountain and a study of contributions to its frequency instability. *IEEE Trans. UFFC*, 47(2):432–437, 2000.

[75] F. Levi, L. Lorini, D. Calonico, and A. Godone. Systematic shift uncertainty evaluation of IEN CSF1 primary frequency standard. *IEEE Trans. Instrum. Meas.*, 52(2):267–271, 2003.

[76] K. Szymaniec, S. E. Park, G. Marra, and W. Chałupczak. First accuracy evaluation of the NPL-CsF2 primary frequency standard. *Metrologia*, 47(4):363–376, 2010.

[77] Z. T. Lu, K. L. Corwin, M. J. Renn, M. H. Anderson, E. A. Cornell, and C. E. Wieman. Low-velocity intense source of atoms from a magneto-optical trap. *Phys. Rev. Lett.*, 77:3331–3334, 1996.

[78] E. A. Donley, T. P. Heavner, and S. R. Jefferts. Optical molasses loaded from a low-velocity intense source of atoms: an atom source for improved atomic fountains. *IEEE Trans. Instrum. Meas.*, 54(5):1905–1910, 2005.

[79] J. Guena, P. Rosenbusch, P. Laurent, M. Abgrall, D. Rovera, G. Santarelli, M.E. Tobar, S. Bize, and A. Clairon. Demonstration of a dual alkali rb/cs fountain clock. *IEEE Trans. UFFC*, 57(3):647–653, 2010.

[80] K. Dieckmann, R.J.C. Spreeuw, M. Weidemüller, and J. T.M. Walraven. Two-dimensional magneto-optical trap as a source of slow atoms. *Phys. Rev. A*, 58:3891–3895, 1998.

[81] G. Dobrev, V. Gerginov, and S. Weyers. Loading a fountain clock with an enhanced low-velocity intense source of atoms. *Phys. Rev. A*, 93:043423, 2016.

[82] V. Gerginov, N. Nemitz, D. Griebsch, M. Kazda, R. Li, K. Gibble, R. Wynands, and S. Weyers. Recent improvements and current uncertainty budget of PTB fountain clock CSF2. In *EFTF-2010 24th European Frequency and Time Forum*, pages 1–7, 2010.

[83] S. E. Park and K. Szymaniec. Accumulation of a clock state population by optical pumping in a caesium fountain. In *CPEM 2010*, pages 470–471, 2010.

[84] K. Szymaniec, H. R. Noh, S. E. Park, and A. Takamizawa. Spin polarization in a freely evolving sample of cold atoms. *Applied Physics B*, 111(3):527–535, 2013.

[85] J. E. Gray and D. W. Allan. A method for estimating the frequency stability of an individual oscillator. In *28th Annual Symposium on Frequency Control*, pages 243–246, 1974.

[86] D. W. Allan and J. A. Barnes. A modified "Allan variance" with increased oscillator characterization ability. In *35th Annual Symposium on Frequency Control*, 1981.

[87] W. M. Itano, J. C. Bergquist, J. J. Bollinger, J. M. Gilligan, D. J. Heinzen, F. L. Moore, M. G. Raizen, and D. J. Wineland. Quantum projection noise: Population fluctuations in two-level systems. *Phys. Rev. A*, 47:3554–3570, 1993.

[88] H. T. Mitchell. Ageing of quartz crystal vibrators. *Nature*, 174(4418):41–42, 1954.

[89] E. Seydel. The different ageing behaviour of fundamental and overtone modes. *IET Conference Proceedings*, pages 78–80(2), 2004.

[90] G. J. Dick. Local oscillator induced instabilities in trapped ion frequency standards. In *In Proc. of Precise Time and Time Interval*, pages 133–147, 1987.

[91] C. A. Greenhall. A derivation of the long-term degradation of a pulsed atomic frequency standard from a control-loop model. *IEEE Trans. UFFC*, 45(4):895–898, 1998.

[92] G. Santarelli, C. Audoin, A. Makdissi, P. Laurent, G. J. Dick, and A. Clairon. Frequency stability degradation of an oscillator slaved to a periodically interrogated atomic resonator. *IEEE Trans. UFFC*, 45(4):887–894, 1998.

[93] C. Audoin, G. Santarelli, A. Makdissi, and A. Clairon. Properties of an oscillator slaved to a periodically interrogated atomic resonator. *IEEE Trans. UFFC*, 45(4):877–886, 1998.

[94] W. Zhao, W. Qian, D. Lu, and R. Wei. Improvement of average magnetic field measurement based on magnetic-field-sensitive Ramsey fringes. *Opt. Lett.*, 47(8):2073–2076, 2022.

[95] C. Shi, R. Wei, Z. Zhou, D. Lu, T. Li, and Y. Wang. Magnetic field measurement on $^{87}$Rb atomic fountain clock. *Chin. Opt. Lett.*, 8(6):549–552, 2010.

[96] S. R. Jefferts, J. Shirley, T. E. Parker, T. P. Heavner, D. M. Meekhof, C. Nelson, F. Levi, G. Costanzo, A. De Marchi, R. Drullinger, L. Hollberg, W. D. Lee, and F. L. Walls. Accuracy evaluation of NIST-F1. *Metrologia*, 39(4):321–336, 2002.

[97] Z. Zhou, R. Wei, C. Shi, T. Li, and Y. Wang. Magnetic field measurement based on a stimulated two-photon Raman transition. *Chinese Physics B*, 20(3):034206, 2011.

[98] K. Moler, D. S. Weiss, M. Kasevich, and S. Chu. Theoretical analysis of velocity-selective Raman transitions. *Phys. Rev. A*, 45:342–348, 1992.

[99] P. G. H. Sandars. Differential polarizability in the ground state of the hydrogen atom. *Proceedings of the Physical Society*, 92(4):857, 1967.

[100] J. R. P. Angel, P. G. H. Sandars, and B. Bleaney. The hyperfine structure Stark effect I. theory. *Proc. R. Soc A*, 305(1480):125–138, 1968.

[101] R. D. Haun and J. R. Zacharias. Stark effect on cesium-133 hyperfine structure. *Phys. Rev.*, 107:107–109, 1957.

[102] T. F. Gallagher and W. E. Cooke. Interactions of blackbody radiation with atoms. *Phys. Rev. Lett.*, 42:835–839, 1979.

[103] M. Plank. Über das gesetz der energieverteilung im normalspektrum. *Annalen der Physik*, 309(4):553–563, 1901.

[104] W. M. Itano, L. L. Lewis, and D. J. Wineland. Shift of $^2S_{\frac{1}{2}}$ hyperfine splittings due to blackbody radiation. *Phys. Rev. A*, 25:1233–1235, 1982.

[105] M. Mizushima. Theory of resonance frequency shift due to radiation field. *Phys. Rev.*, 133:A414–A418, 1964.

[106] A. Bauch and R. Schröder. Experimental verification of the shift of the cesium hyperfine transition frequency due to blackbody radiation. *Phys. Rev. Lett.*, 78:622–625, 1997.

[107] E. Simon, P. Laurent, and A. Clairon. Measurement of the Stark shift of the Cs hyperfine splitting in an atomic fountain. *Phys. Rev. A*, 57:436–439, 1998.

[108] S. Zhang. *Frequency shift due to blackbody radiation in a cesium atomic fountain and improvement of the clock performances.* Theses, Université Pierre et Marie Curie - Paris VI, 2004.

[109] P. Rosenbusch, S. Zhang, and A. Clairon. Blackbody radiation shift in primary frequency standards. In *Proceedings of the Joint IEEE FCS-EFTF 2007*, pages 1060–1063, 2007.

[110] E. Simon, G. Santarelli, P. Laurent, S. Ghezali, P. Lemonde, and A. Clairon. Measurement of the frequency shift induced by thermal radiation in a Cs atomic fountain frequency standard. In *1998 Conference on Precision Electromagnetic Measurements Digest*, pages 126–127, 1998.

[111] F. Levi, D. Calonico, L. Lorini, S. Micalizio, and A. Godone. Measurement of the blackbody radiation shift of the $^{133}$Cs hyperfine transition in an atomic fountain. *Phys. Rev. A*, 70:033412, 2004.

[112] S. Micalizio, A. Godone, D. Calonico, F. Levi, and L. Lorini. Blackbody radiation shift of the $^{133}$Cs hyperfine transition frequency. *Phys. Rev. A*, 69:053401, 2004.

[113] K. Beloy, U. I. Safronova, and A. Derevianko. High-accuracy calculation of the blackbody radiation shift in the $^{133}$Cs primary frequency standard. *Phys. Rev. Lett.*, 97:040801, 2006.

[114] E. J. Angstmann, V. A. Dzuba, and V. V. Flambaum. Frequency shift of the cesium clock transition due to blackbody radiation. *Phys. Rev. Lett.*, 97:040802, 2006.

[115] S. R. Jefferts, T. P. Heavner, T. E. Parker, J. H. Shirley, E. A. Donley, N. Ashby, F. Levi, D. Calonico, and G. A. Costanzo. High-accuracy measurement of the blackbody radiation frequency shift of the ground-state hyperfine transition in $^{133}$Cs. *Phys. Rev. Lett.*, 112:050801, 2014.

[116] N. F. Ramsey and H. B. Silsbee. Phase shifts in the molecular beam method of separated oscillating fields. *Phys. Rev.*, 84:506–507, 1951.

[117] A. Bauch, T. Heindorff, and R. Schroeder. Measurement of the frequency-shift due to distributed cavity phase difference in an atomic clock. *IEEE Trans. Instrum. Meas.*, 34(2):136–138, 1985.

[118] R. F. Lacey. Phase shift in microwave Ramsey structures. In *22nd Annual Symposium on Frequency Control*, pages 545–558, 1968.

[119] A. De Marchi, J. Shirley, D. J. Glaze, and R. Drullinger. A new cavity configuration for cesium beam primary frequency standards. *IEEE Trans. Instrum. Meas.*, 37(2):185–190, 1988.

[120] G. Vecchi and A. De Marchi. Spatial phase variations in a $TE_{011}$ microwave cavity for use in a cesium fountain primary frequency standard. *IEEE Trans. Instrum. Meas.*, 42(2):434–438, 1993.

[121] D. M. Pozar. *Microwave engineering.* Fourth edition. Hoboken, NJ : Wiley, 2012.

[122] A. Khursheed and A. De Marchi. Spatial phase distributions within cavities for cesium fountain standards. In *Proc. 8th EFTF*, pages 65–70, 1994.

[123] A. Khursheed, G. Vecchi, and A. De Marchi. Spatial variations of field polarization and phase in microwave cavities: application to the cesium fountain cavity. *IEEE Trans. UFFC*, 43(2):201–210, 1996.

[124] S. R. Jefferts, R. E. Drullinger, and A. DeMarchi. Nist cesium fountain microwave cavities. In *Proc. 1998 IEEE IFCS*, pages 6–8, 1998.

[125] R. Li and K. Gibble. Phase variations in microwave cavities for atomic clocks. *Metrologia*, 41(6):376, 2004.

[126] P. Lemonde, G. Santarelli, P. Laurent, F.P. Dos Santos, A. Clairon, and C. Salomon. The sensitivity function: a new tool for the evaluation of frequency shifts in atomic spectroscopy. In *Proc. 1998 IEEE IFCS*, pages 110–115, 1998.

[127] R. Li and K. Gibble. Distributed cavity phase and the associated power dependence. In *Proc. 2005 IEEE IFCS and Exposition*, page 6, 2005.

[128] R. Li and K. Gibble. Evaluating and minimizing distributed cavity phase errors in atomic clocks. *Metrologia*, 47(5):534, 2010.

[129] J. Guéna, R. Li, K. Gibble, S. Bize, and A. Clairon. Evaluation of Doppler shifts to improve the accuracy of primary atomic fountain clocks. *Phys. Rev. Lett.*, 106:130801, 2011.

[130] S. R. Jefferts, J. H. Shirley, N. Ashby, E. A. Burt, and G. J. Dick. Power dependence of distributed cavity phase-induced frequency biases in atomic fountain frequency standards. *IEEE Trans. UFFC*, 52(12):2314–2321, 2005.

[131] F. Chapelet, S. Bize, P. Wolf, G. Santarelli, P. Rosenbusch, M. E. Tobar, P. Laurent, C. Salomon, and A. Clairon. Investigation of the distributed cavity phase shift in an atomic fountain. In *Proc. 20th EFTF 2006*, volume 1, pages 160–165, 2006.

[132] T. P. Heavner, E. A. Donley, F. Levi, G. Costanzo, T. E. Parker, J. H. Shirley, N. Ashby, S. Barlow, and S. R. Jefferts. First accuracy evaluation of NIST-F2. *Metrologia*, 51(3):174–182, 2014.

[133] S. Weyers, V. Gerginov, N. Nemitz, R. Li, and K. Gibble. Distributed cavity phase frequency shifts of the caesium fountain PTB-CSF2. *Metrologia*, 49(1):82, 2011.

[134] N. Nemitz, V. Gerginov, R. Wynands, and S. Weyers. Atomic trajectory characterization in a fountain clock based on the spectrum of a hyperfine transition. *Metrologia*, 49(4):468, may 2012.

[135] K. Burrows, R. J. Hendricks, K. Szymaniec, K. Gibble, S. Beattie, and B. Jian. Measuring atom positions in a microwave cavity to evaluate distributed cavity phase shifts. *Metrologia*, 57(6):065003, 2020.

[136] P. J. Leo, C. J. Williams, and P. S. Julienne. Collision properties of ultracold $^{133}$Cs atoms. *Phys. Rev. Lett.*, 85:2721–2724, 2000.

[137] E. Tiesinga, B. J. Verhaar, H. T. C. Stoof, and D. van Bragt. Spin-exchange frequency shift in a cesium atomic fountain. *Phys. Rev. A*, 45:R2671–R2673, 1992.

[138] B. J. Verhaar, J. M. V. A. Koelman, H. T. C. Stoof, O. J. Luiten, and S. B. Crampton. Hyperfine contribution to spin-exchange frequency shifts in the hydrogen maser. *Phys. Rev. A*, 35:3825–3831, 1987.

[139] W. Kolos and L. Wolniewicz. Potential-energy curves for the $X^1\Sigma_g^+$, $b^3\Sigma_u^+$, and $C^1\Pi_u$ states of the hydrogen molecule. *The Journal of Chemical Physics*, 43(7):2429–2441, 1965.

[140] J. Vanier and C. Audoin. *The quantum physics of atomic frequency standards.* Adam Hilger, Bristol, 1989.

[141] P. J. Leo, P. S. Julienne, F. H. Mies, and C. J. Williams. Collisional frequency shifts in $^{133}$Cs fountain clocks. *Phys. Rev. Lett.*, 86:3743–3746, 2001.

[142] J. M. V. A. Koelman, S. B. Crampton, H. T. C. Stoof, O. J. Luiten, and B. J. Verhaar. Spin-exchange frequency shifts in cryogenic and room-temperature hydrogen masers. *Phys. Rev. A*, 38:3535–3547, 1988.

[143] K. Gibble and S. Chu. Laser-cooled Cs frequency standard and a measurement of the frequency shift due to ultracold collisions. *Phys. Rev. Lett.*, 70:1771–1774, 1993.

[144] K. Szymaniec, W. Chałupczak, E. Tiesinga, C. J. Williams, S. Weyers, and R. Wynands. Cancellation of the collisional frequency shift in caesium fountain clocks. *Phys. Rev. Lett.*, 98:153002, 2007.

[145] K. Szymaniec, W. Chałupczak, S. Weyers, and R. Wynands. Prospects of operating a caesium fountain clock at zero collisional frequency shift. *App. Phys. B*, 89(2):187–193, 2007.

[146] F. Levi, D. Calonico, L. Lorini, and A. Godone. IEN-CsF1 primary frequency standard at INRIM: accuracy evaluation and TAI calibrations. *Metrologia*, 43(6):545, 2006.

[147] C. Vian, P. Rosenbusch, H. Marion, S. Bize, L. Cacciapuoti, S. Zhang, M. Abgrall, D. Chambon, I. Maksimovic, P. Laurent, G. Santarelli, A. Clairon, A. Luiten, M. Tobar, and C. Salomon. BNM-SYRTE fountains: recent results. *IEEE Trans. Instrum. Meas.*, 54(2):833–836, 2005.

[148] K. Szymaniec, W. Chałupczak, P. B. Whibberley, S. N. Lea, and D. Henderson. Evaluation of the primary frequency standard NPL-CsF1. *Metrologia*, 42(1):49, 2005.

[149] W. Chałupczak and K. Szymaniec. Collisions in a ballistically expanding cloud of cold atoms in an atomic fountain. *J. Phys. B*, 40(2):343, 2007.

[150] S. J. J. M. F. Kokkelmans, B. J. Verhaar, K. Gibble, and D. J. Heinzen. Predictions for laser-cooled Rb clocks. *Phys. Rev. A*, 56:R4389–R4392, 1997.

[151] C. Fertig and K. Gibble. Measurement and cancellation of the cold collision frequency shift in an $^{87}$Rb fountain clock. *Phys. Rev. Lett.*, 85:1622–1625, 2000.

[152] Y. Sortais, S. Bize, C. Nicolas, A. Clairon, C. Salomon, and C. Williams. Cold collision frequency shifts in a $^{87}$Rb atomic fountain. *Phys. Rev. Lett.*, 85:3117–3120, 2000.

[153] J. P. Burke, J. L. Bohn, B. D. Esry, and C. H. Greene. Impact of the $^{87}$Rb singlet scattering length on suppressing inelastic collisions. *Phys. Rev. A*, 55:R2511–R2514, 1997.

[154] Y. Sortais. Construction d'une fontaine double à atomes froids de $^{87}$Rb et $^{133}$Cs ; Etude des effets dépendant du nombre d'atomes dans une fontaine, 2001.

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

[156] K. Szymaniec, S. N. Lea, and K. Liu. An evaluation of the frequency shift caused by collisions with background gas in the primary frequency standard NPL-CsF2. *IEEE Trans. UFFC*, 61(1):203–206, 2014.

[157] S. Bize, Y. Sortais, C. Mandache, A. Clairon, and C. Salomon. Cavity frequency pulling in cold atom fountains. *IEEE Trans. Instrum. Meas.*, 50(2):503–506, 2001.

[158] R. Schroder, U. Hubner, and D. Griebsch. Design and realization of the microwave cavity in the PTB caesium atomic fountain clock CSF1. *IEEE Trans. UFFC*, 49(3):383–392, 2002.

[159] M. Kumagai, H. Ito, M. Kajita, and M. Hosokawa. Evaluation of caesium atomic fountain NICT-CsF1. *Metrologia*, 45(2):139–148, 2008.

[160] A. de Marchi, G. D. Rovera, and A. Premoli. Pulling by neighbouring transitions and its effects on the performance of caesium-beam frequency standards. *Metrologia*, 20(2):37, 1984.

[161] L. S. Cutler, C. A. Flory, R. P. Giffard, and A. De Marchi. Frequency pulling by hyperfine $\sigma$ transitions in cesium beam atomic frequency standards. *Journal of Applied Physics*, 69(5):2780–2792, 03 1991.

[162] H. S. Lee, T. Y. Kwon, H. S. Kang, Y. H. Park, C. H. Oh, S. E. Park, H. Cho, and V. G. Minogin. Comparison of the Rabi and Ramsey pulling in an optically pumped caesium-beam standard. *Metrologia*, 40(5):224, 2003.

[163] V. Gerginov, N. Nemitz, and S. Weyers. Initial atomic coherences and Ramsey frequency pulling in fountain clocks. *Phys. Rev. A*, 90:033829, 2014.

[164] F. Bloch and I. I. Rabi. Atoms in variable magnetic fields. *Rev. Mod. Phys.*, 17:237–244, Apr 1945.

[165] R. Wynands, R. Schroder, and S. Weyers. Majorana transitions in an atomic fountain clock. *IEEE Trans. Instrum. Meas.*, 56(2):660–663, 2007.

[166] K. Dorenwendt and A. Bauch. Spurious microwave fields in caesium atomic beam frequency standards: symmetry considerations and model calculations. In *Proc. 1999 EFTF and IEEE IFCS*, volume 1, pages 57–61 vol.1, 1999.

[167] B. Boussert, G. Theobald, P. Cerez, and E. de Clercq. Frequency shifts in cesium beam clocks induced by microwave leakages. *IEEE Trans. UFFC*, 45(3):728–738, 1998.

[168] S. Weyers, R. Schröder, and R. Wynands. Effects of microwave leakage in caesium clocks: Theoretical and experimental results. In *Proceedings of the 20th European Frequency and Time Forum*, pages 173–180, 2006.

[169] J. H. Shirley, F. Levi, T. P. Heavner, D. Calonico, D. H. Yu, and S. R. Jefferts. Microwave leakage-induced frequency shifts in the primary frequency standards NIST-F1 and IEN-CSF1. *IEEE Trans. UFFC*, 53(12):2376–2385, 2006.

[170] N. F. Ramsey. Resonance transitions induced by perturbations at two or more different frequencies. *Phys. Rev.*, 100:1191–1194, Nov 1955.

[171] C. Audoin, M. Jardino, L. S. Cutler, and R. F. Lacey. Frequency offset due to spectral impurities in cesium-beam frequency standards. *IEEE Trans. Instrum. Meas.*, 27(4):325–329, 1978.

[172] F. Levi, J. H. Shirley, T. P. Heavner, D. Yu, and S. R. Jefferts. Power dependence of the frequency bias caused by spurious components in the microwave spectrum in atomic fountains. *IEEE Trans. UFFC*, 53(9):1584–1589, 2006.

[173] T. P. Heavner, J. H. Shirley, F. Levi, D. Yu, and S. R. Jefferts. Frequency biases in pulsed atomic fountain frequency standards due to spurious components in the microwave spectrum. In *2006 IEEE IFCS*, pages 273–276, 2006.

[174] P. Wolf, S. Bize, A. Clairon, P. Laurent, P. Lemonde, and C. J. Bordé. Recoil effects in microwave atomic frequency standards: preliminary results. In *2001 IEEE IFCS*, page 37, 2001.

[175] P. Wolf and C. J. Bordé. Recoil effects in microwave Ramsey spectroscopy, 2004.

[176] K. Gibble. Difference between a photon's momentum and an atom's recoil. *Phys. Rev. Lett.*, 97:073002, 2006.

[177] N. K. Pavlis and M. A. Weiss. The relativistic redshift with $3 \times 10^{-17}$ uncertainty at NIST, Boulder, Colorado, USA. *Metrologia*, 40(2):66, apr 2003.

[178] J. Vanier and C. Audoin. *The quantum physics of atomic frequency standards.* Adam Hilger, Bristol, 1989.

[179] D. Steck. Rubidium 87 D line data. *Los Alamos National Laboratory*, 2003.

[180] D. Steck. Cesium D line data. *Los Alamos National Laboratory*, 2010.

[181] S. Bize, Y. Sortais, M. S. Santos, C. Mandache, A. Clairon, and C. Salomon. High-accuracy measurement of the $^{87}$Rb ground-state hyperfine splitting in an atomic fountain. *Europhysics Letters*, 45(5):558, 1999.

[182] N. Marcuvitz. *Waveguide Handbook.* Institution of Electrical Engineers, 1986.

[183] H. A. Bethe. Theory of diffraction by small holes. *Phys. Rev.*, 66:163–182, 1944.

[184] R. E. Collin. *Field Theory of Guided Waves.* IEEE-Press, second edition, 1991.

[185] F. Zheng, F. Fang, W. Chen, K. Liu, S. Dai, S. Cao, and T. Li. Microwave-vacuum integrated cavity with a low temperature sensitivity for Cs fountain clocks. In *2021 Joint EFTF and IEEE IFCS*, pages 1–4, 2021.

[186] G. P. Le Sage and R. Stringfellow. 3D printed, self-temperature compensated microwave cavity notch filters. *IEEE Access*, 9:32969–32978, 2021.

[187] G. W. Hoth, B. Patla, N. Ashby, and V. Gerginov. Initial study of the distributed cavity phase shift for the new microwave cavities of Cs fountains at NIST. *Proc. IEEE IFCS and EFTF Conf.*, 2022.

[188] S. Lee, G. W. Choi, H. G. Hong, T. Y. Kwon, S. B. Lee, M. S. Heo, and S. E. Park. A compact cold-atom clock based on a loop-gap cavity. *Applied Physics Letters*, 119(6):064002, 2021.

[189] T. Konaka, M. Sato, H. Asano, and S. Kubo. Relative permittivity and dielectric loss tangent of substrate materials for high-$T_c$ superconducting film. *Journal of Superconductivity*, 4(4):283–288, 1991.

[190] R.G. Geyer, J. Baker-Jarvis, and J. Krupka. Dielectric characterization of single-crystal LiF, $CaF_2$, $MgF_2$, $BaF_2$, and $SrF_2$ at microwave frequencies. In *17th IEEE LEOS, 2004.*, pages 493–497, 2004.

[191] L. Alaribe, C. Disch, A. Fauler, R. Engels, E. Keller, A. Cecilia, T. dos Santos Rolo, E. Hamann, and M. Fiederle. $SrI_2:Eu^{2+}$-scintillators for spectroscopy and X-ray imaging applications. *IEEE Trans. Nuc. Sci.*, 59(5):2193–2195, 2012.

[192] L. Pitre, F. Sparasci, D. Truong, A. Guillou, L. Risegari, and M. E. Himbert. Measurement of the Boltzmann constant $k_B$ using a quasi-spherical acoustic resonator. *International Journal of Thermophysics*, 32(9):1825, 2011.

[193] N. Ashby, S. Romisch, and S. R. Jefferts. Endcaps for $\mathrm{TE}_{01}$ cavities in fountain frequency standards. In *Proc. IEEE IFCS and 17th EFTF*, pages 1076–1083, 2003.

[194] M. Knapp. *Progress towards a compact cold-atom microwave clock.* PhD thesis, University of Oxford, 2024.

[195] S. Peil, S. Crane, T. B. Swanson, and C. R. Ekstrom. Rubidium-fountain characterization using the USNO clock ensemble. In *2007 IEEE International Frequency Control Symposium Joint with the 21st European Frequency and Time Forum*, pages 473–476, 2007.

[196] R. J. Hanson and F. M. Pipkin. Magnetically Shielded Solenoid with Field of High Homogeneity. *Review of Scientific Instruments*, 36(2):179–188, 2004.

[197] R. H. Lambert and C. Uphoff. Magnetically shielded solenoid with field of high homogeneity. *Review of Scientific Instruments*, 46(3):337–337, 1975.

[198] C. Y. Liu, T. Andalib, D. C. M. Ostapchuk, and C. P. Bidinosti. Analytic models of magnetically enclosed spherical and solenoidal coils. *Nuclear Instruments and Methods in Physics Research Section A*, 949:162837, 2020.

[199] T. J. Sumner, J. M. Pendlebury, and K. F. Smith. Convectional magnetic shielding. *Journal of Physics D: Applied Physics*, 20(9):1095, 1987.

[200] P. Arpaia, P. N. Burrows, M. Buzio, C. Gohil, M. Pentella, and D. Schulte. Magnetic characterization of mumetal for passive shielding of stray fields down to the nano-Tesla level. *Nuclear Instruments and Methods in Physics Research Section A*, 988:164904, 2021.

[201] S. A. Self. Focusing of spherical Gaussian beams. *Applied optics*, 22 5:658–61, 1983.

[202] J. M. Rüeger. Refractive index formulae for radio waves. *UNSW Sydney NSW*, 2002.

[203] C. Benvenuti, G. Canil, P. Chiggiato, P. Collin, R. Cosso, J. Guerin, S. Ilie, D. Latorre, and K. S. Neil. Surface cleaning efficiency measurements for UHV applications. *Vacuum*, 53:317–320, 1999.

[204] R. J. Reid. Cleaning for vacuum service. *CERN Accelerator School: Vacuum Technology*, 1999.

[205] D. G. Shchukin, E. Skorb, V. Belova, and H. Möhwald. Ultrasonic cavitation at solid surfaces. *Advanced Materials*, 23(17):1922–1934, 2011.

[206] J. R. Laguna-Camacho, R. Lewis, M. Vite-Torres, and J. V. Méndez-Méndez. A study of cavitation erosion on engineering materials. *Wear*, 301(1):467–476, 2013.

[207] M. Caragiu and S. Finberg. Alkali metal adsorption on graphite: a review. *Journal of Physics: Condensed Matter*, 17(35):R995, 2005.

[208] N. D. Bhaskar, C. M. Kahla, and L. R. Martin. Absorption of cesium by polycrystalline graphite-sticking coefficient studies. *Carbon*, 28(1):71–78, 1990.

[209] R. R. Turkington and R. F. Harris-Lowe. Note on the design of simple indium O-ring seals. *Review of Scientific Instruments*, 55(5):803–805, 1984.

[210] P. Chiggiato. Outgassing properties of vacuum materials for particle accelerators, 2020.

[211] H. F. Dylla, D. M. Manos, and P. H. LaMarche. Correlation of outgassing of stainless steel and aluminum with various surface treatments. *Journal of Vacuum Science and Technology A*, 11(5):2623–2636, 1993.

[212] C. Benvenuti, J. P. Bojon, P. Chiggiato, and G. Losch. Ultimate pressures of the large electron positron collider (LEP) vacuum system. *Vacuum*, 44(5):507–509, 1993.

[213] A. Derevianko, S. G. Porsev, and J. F. Babb. Electric dipole polarizabilities at imaginary frequencies for hydrogen, the alkali–metal, alkaline–earth, and noble gas atoms. *Atomic Data and Nuclear Data Tables*, 96(3):323–331, 2010.

[214] S. Walby, M. Knapp, J. Whale, A. Wilson, R. Hendricks, C. J. Foot, and K. Szymaniec. Normalised detection of clock states by cold atom recapture method. In *2022 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS)*, pages 1–3, 2022.

[215] K. Szymaniec, R. J. Hendricks, J. Whale, A. Wilson, S. Walby, and C. J. Foot. Robust design and performance of NPL Cs fountain clocks. In *The 9th Symposium on Frequency Standards and Metrology*, 2023.