Development of a Pulsed Optical Pumped Rubidium Clock.

by

Stuart KENNY

Doctor of Philosophy

University of York

Physics, Engineering and Technology

December 2023

Abstract

Keywords: Atomic clock, rubidium vapour cell, double resonance, Ramsey interferometry, waveguide, frequency stability, metrology.

Since the first ammonia clock in 1949, atomic clocks have been developed to provide improved frequency stability as well as miniature designs capable of being worn on the wrist. With a small number of notable exceptions, most research activity has been directed at improving clock accuracy without regard to power or size, or with very constrained power and volume goals. A design that offers improved frequency stability, within a moderate volume, and that can be used whilst moving, would prove a widely useful reference in laboratory, communication, and navigation applications.

To fill this gap, this thesis uses simplified apparatus to achieve the Pulsed Optical Pumping (POP) technique; proves the viability of waveguide as an alternative to a conventional cylindrical cavity; develops a higher-resolution Digital to Analog Convertor (DAC) than anything commercially available; and does so in a way that can fit within the volume of a shoebox. Limiting factors to the performance of a ^{85}Rb cell are calculated, and a central Ramsey fringe with 100 Hz Full Width at Half Maximum (FWHM) is achieved. A route to demonstrating long-term stability of 2×10^{-14} (τ=100sec) whilst maintaining -123 dBc/Hz (1Hz offset) close-to-carrier phase noise is provided. The clock's performance has not been measured as the laser and microwave locking loops are not fully functional.

List of contents

List of Tables

List of Figures

Acknowledgements

This thesis would not have been possible without the help and support of numerous people. I would like to thank my supervisors, Professor Jeremy Everard and Doctor Simon Bale, for their sustained generosity and guidance, Doctor John Dawson for his timely review and perspective, and Doctor Andrew Pratt for the use of his lab and Doppler-free Saturated Spectroscopy apparatus. Warm thanks are due to my sponsors, the Engineering and Physical Sciences Research Council (EPSRC) and Leonardo MW, and particularly to Professor Robert Lamb and Stephen Clark for their engagement, encouragement, and experience. Many thanks to Robin Collings and Sophie Hutchins for making the long trek to York, and to everyone who provided feedback on the high-resolution DAC.

I am hugely grateful to all the professional support staff in the School of Physics, Engineering and Technology that have made this thesis possible: Camilla Danese has expertly steered me through procedural hurdles and I have benefitted greatly from her unparalleled knowledge of the University. Mark Hough kindly allowed me to pick his brains, and has helped with the laser enclosure and numerous PCBs; Karen Smith fielded countless enquiries and helped with a great deal of laser cutting; David Hunter, Ian Helliwell, and Andy White were unfailingly generous when I've needed to borrow yet another piece of equipment. Jason Flatt and Brent Wilkinson machined two of the MW cavities; and Akiko Tabata provided warm and patient support with purchases and budget.
Doctor Iain Will and Doctor Kari Niemi provided guidance and constructive approach to laser safety.

I'm very grateful to Doctor Abigail MacQuarrie, to St.John Gilbert, and to Luke Dummott for a wonderful mix of motivation, understanding, and irreverence during tea-breaks and lunchtimes. And thanks too to my wife Helen and daughter Anna who have been endlessly kind, have helped me find time for this PhD, and without whose encouragement this thesis would not exist.

Author's declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as references.

1 Introduction

1.1 General introduction

This thesis documents the design and build of an ^{85}Rb clock using Pulsed Optical Pumping (POP) and Ramsey interferometry. It demonstrates that POP can be conducted with a travelling wave for Microwave (MW) interaction, and it is therefore possible to avoid the frequency pulling effect associated with a cavity. Traditional apparatus is further simplified by using a single reference cell for both laser and MW lock. The effects of varying the static magnetic field (C-field), which provides a quantized axis for the atomic transition, are observed, including operation in its absence. Lastly, an outline of how this design could to productised to fit into a volume no bigger than a shoebox.

1.2 Structure of this thesis

This document is intended to support a range of readers: from those who are not yet experts seeking context, and an accessible way to access the theory needed; to those intimately familiar with Rb or Cs clocks who wish to dive straight into implementation detail. The body of this document has been divided into chapters as follows:

1. Introduction. This chapter provides context for new readers, providing a brief historical timeline and outline of some applications, as well as detailing the motivation and starting point for this thesis.
2. Atomic Theory. How atomic clocks work, written for an audience not familiar with quantum theory.
3. Physics Package. Everything that has a physical effect which includes the microwave enclosure, field generation, heating, and the full optical path including the source of a laser beam and optical detection.
4. POP timing and the Control Platform. How the necessary control sequence and timing for Pulsed Optical Pumping (POP) is generated. It covers generation of the required MW signal, the STM32 Control Platform, and covers development of a high-resolution Digital to Analogue Converter (DAC).
5. Spectroscopy, Double Resonance and POP. How the Physics Package was confirmed working and power/signal/voltage levels were optimised, starting with a simple form of atomic excitation, then exploiting additional properties to improve the clock's accuracy, before achieving and optimising POP.
6. Control loops. Laser stability and the feedback loops used to lock the laser to the correct wavelength, and how the atomic clock is 'disciplined'.

7. Summary. An opportunity to expand on what did happen, what should have happened, and what could happen next.

Material of potential use to a limited readership, but that detracts from the logical flow of the thesis has been placed in Appendices.

A. A comprehensive list of the lab equipment used in this PhD research, including the major commercial products used in the construction of the clock.
B. The design and build of a 'traditional' cylindrical cavity that could support a standing MW field for ^{85}Rb and a theoretical analysis on how to measure the resonant bandwidth (B/W) of a one-port cavity.
C. The design, simulation and build of an amplified photodiode. This is not included in the main body of the thesis as it was superseded by a commercial part with superior specification.
D. Schematics for the Logical Interface (I/F) board that connects a Micro-control Unit (MCU) to the Physics Package.
E. Schematics for a high-resolution Digital to Analog Converter (DAC).

All code produced for this thesis has been checked into GitHub online software repositories, with links from the appropriate section.

1.3 What is an atomic clock?

An atomic clock is an accurate reference of time, based on the properties of individual atoms, rather than those of a material. Previous clocks were limited by factors such as the consistency of spring tension, the length of a pendulum, or how accurately quartz can be machined. In an atomic clock, by observing the behaviour of individual electrons, or the consequences of that behaviour, a measure of time can be based on an invariant property of atoms, largely independent of other physical factors.

1.4 Origins of atomic clocks

Contrary to Descartes' 1644 postulation that matter is infinitely divisible [1], Newton defended Atomism and wondered if there was "distinction between mathematical and physical divisibility" [2]. Perhaps starting with Thomas Young's double slit experiment c.1801, scientists observed behaviour that was inconsistent with Newtonian physics and that led to the realisation that matter, and energy, comes in small units. The field of Quantum mechanics emerged to understand the behaviour of individual atoms and their particles; the field is named from the Latin "quantum", or 'how much'.

As Lombardi notes [3], the concept of using atoms to keep time began with James Clerk Maxwell in the 1870s, with much formative work performed by Isidor Isaac Rabi and his colleagues in Columbia University in the 1930s and 1940s [4]. The first atomic clock was built in 1949 courtesy of Harold Lyons and the US National Bureau of Standards (NBS) [5], and was based on an ammonia absorption line at

23.87GHz, but never surpassed the stability of quartz oscillators. It was Louis Essen at NPL, building on the somewhat disrupted work at NBS, who produced the first Caesium Beam standard in 1955 [6]. This marked the first time that a lab-based clock could provide better accuracy than astronomical observation, and Essen's subsequent work with the US National Observatory (USNO) [7] eventually led to the second being defined as 9,192,631,770 transitional cycles of ^{133}Caesium [8]. Since then, there have been significant improvements to Caesium Beam Frequency Standards and as well as the successful demonstration of additional types of clock.

1.5 Active v.s. passive references

The Hydrogen maser (an acronym of Microwave Amplification by Stimulated Emission of Radiation) is the only active reference and works due to hydrogen atom emissions, inside a storage bulb, inside a very high-Q microwave cavity. As illustrated in US patent US5838206A [9], an electrical arc dissociates molecules of Hydrogen into atoms which are collimated and pass through a state selector, into a storage bulb. The natural 1.42 GHz resonance being detected is the v_{hfs} zero-field hyperfine frequency (strictly the v_{24} transition which is offset from v_{hfs} by the nearly-degenerate v_{12} and v_{23} Zeeman frequencies), which can be inductively coupled to a small loop inside the microwave cavity, amplified and divided to provide very low-noise outputs of a desired frequency [10].

All other references are passive and require atomic properties to be interrogated by an external source, generated by a Local Oscillator (LO). When the system is tuned correctly, a physical property may be detected as a maximum or minimum, dependent on the method used. In the case of a passive Hydrogen maser, the mechanism is the same as above but if the microwave cavity is fed from an external 1.420405752 GHz source [10], maximum output will be observed when the LO is aligned to the 'natural frequency' of the Hydrogen.

1.6 Microwave (MW) references

A microwave is defined as having a wavelength between 1mm and 1m, therefore a MW clock is one utilising transitions between 300MHz and 300GHz in frequency; in practice between 1.4GHz and 9.2GHz, as these are based on either Hydrogen, or alkali metals in Group 1 of the periodic table which have a single 'valence' electron in their outer shell. The Hydrogen maser, is one such device, others are based on the frequency of transitions between ground hyperfine states in the atoms of an alkali metal with ^{133}Cs, ^{85}Rb and ^{87}Rb being common choices [11, p. 17]. Even the highly reactive ^{23}Na [12] has been used, as all of the alkali metals have one free electron in their outer shell, though Francium is too unstable, radioactive and rare for it to have been considered as the basis of a clock to date.

1.6.1 Spectral lamps

Used in conjunction with a vapour cell and microwave cavity, spectral lamps are a very common form of Rubidium-based, Double Resonance (DR) reference due to a convenient coincidence involving the

two naturally occurring isotopes: one of the ^{87}Rb discharge lines corresponds to one of the transitions in ^{85}Rb. A combination of ^{87}Rb discharge lamp and ^{85}Rb isotopic filter can therefore be used to 'optically pump' [13], [14] a vapour cell containing ^{87}Rb. A MW field with frequency correctly aligned with the hyperfine transition will then produce a dip in the light output detected [15, pp. 40–41]. For a ^{87}Rb vapour cell, this field must correspond to the 6.834 GHz hyperfine transition frequency, resulting in approximately 0.5% less light passing through the vapour cell. The Stanford Research Systems PRS10 uses this method but avoids the need for a separate filter by carefully selecting buffer gases and including a mix of the ^{85}Rb isotope in both the lamp and vapour cell. [16, pp. 9–10].

1.6.2 Beam Tube

Caesium beam references are based on the apparatus shown in Figure 1, which is a form of the Stern-Gerlach experiment [17].

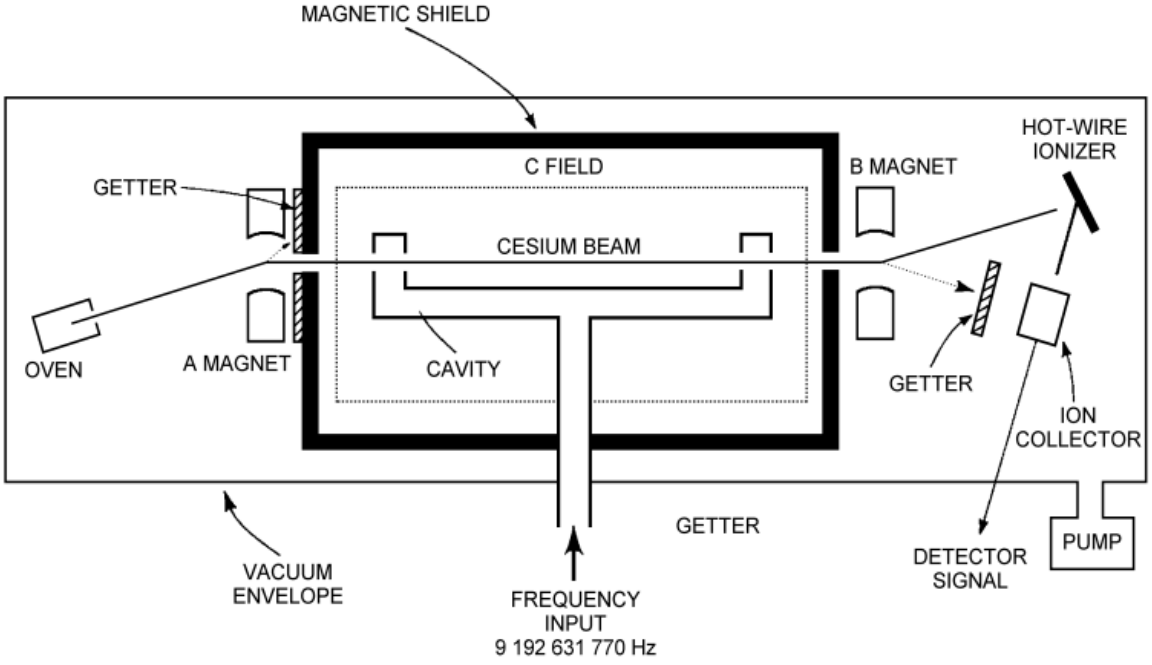

Figure 1. Diagram of a Caesium-beam frequency standard using magnetic state selection and detection [18, Fig. 1].

An oven heats Caesium to almost 100°C, giving off a vapour that is collimated and passed through a state selection gate: the first set of magnets seen in Figure 1. The inhomogeneous magnetic field allows only atoms in one specific ground level state to pass into the cavity; all other atoms are absorbed by the getter plate. Similarly at the exit of the beam, only atoms that are in a different ground state are directed to the detector: an ionizer in this case. Therefore, only atoms that have switched ground state due to Ramsey MW interaction will be detected [19].

1.6.3 Coherent Population Trapping (CPT)

Although initially called 'dark resonance' or 'non-absorption resonance', CPT has been a recognised technique since circa 1976 [20], relying on a laser (or lasers) as the source of excitation. In a material with two hyperfine ground levels (e.g. Rubidium), the laser will excite electrons in both levels simultaneously. When the exciting laser is tuned to one of the atomic D-lines [21], part of the light is absorbed, but if a microwave field is applied at exactly the hyperfine splitting frequency then some will no longer absorb incoming light [22, p. 259]. This is normally achieved by modulating the laser at half of the hyperfine splitting frequency so that two sidebands are separated by the splitting frequency e.g. [11, pp. 133–135]. The ready availability of inexpensive Vertical Cavity Surface Emitting Lasers (VCSEL) meant that it was possible to conduct CPT with very compact equipment and Goka notes considerable research from a number of groups since 2000 CE [23]. Burtichelov's ^{85}Rb CPT clock achieves 3.3% output contrast with a resonant Full Width Half Maximum (FWHM) of 800Hz [24].

1.6.4 Fountain clocks

Although the 'Zacharias fountain' was initially conceived in the 1950s, fountain clocks only became feasible in the late 1980s due to advances in laser cooling. As shown in Figure 2, they work by trapping a cloud of atoms between three orthogonal pairs of laser beams, which cool the atoms down to a few μK, before de-tuning the lasers so that atoms are released vertically upwards.

On the path upwards, in the dark, atoms pass into a magnetically shielded microwave cavity, through a state selection cavity and Ramsey interrogation cavity, before running out of kinetic energy and falling back down under gravity. Detection lasers are switched on to probe the atoms, utilising laser-induced fluorescence, on their descent from the microwave cavity, thereby determining the influence of the magnetic field within the microwave cavity. All of this happens in high vacuum which removes the influence of collisions e.g. with the buffer gases used in a vapour cell. Fountain clocks rely on Ramsey interaction to measure the absorption frequency between the ground and excited states and, as the atoms travel slowly up and down a long microwave cavity, take advantage of the long interrogation time to "sharpen" the resonance curve [25]. Guéna at al. have developed a Dual Alkali (Rb and Cs) fountain that projects atoms almost 1m upwards, providing circa 900ms of ballistic flight, and have demonstrated 0.82Hz linewidth using this method [26]. The narrow linewidth is achieved by using roughly 620ms of flight time for free precessing and is covered in Section 2.2 (see interaction time τ_m).

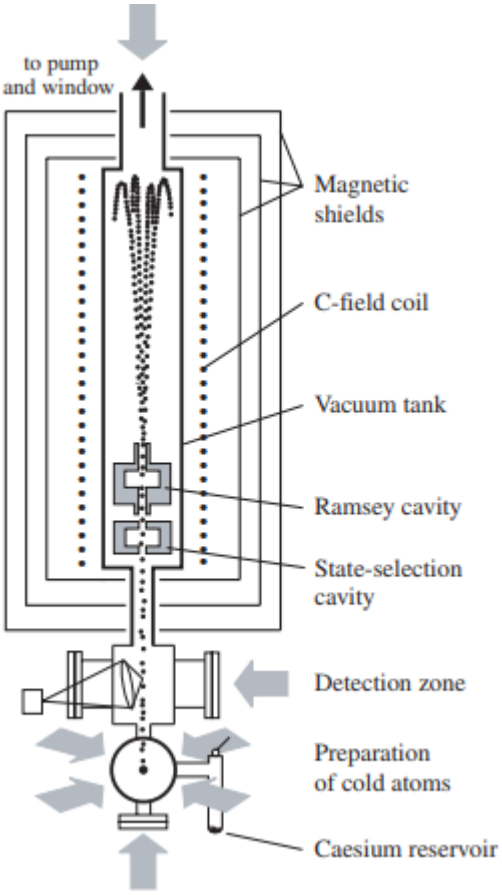

Figure 2: Simplified set-up of the atomic fountain clock [27, Fig. 3]

A traditional (or pulsed) design must trade a very lengthy cycle time for this narrow linewidth but a continuous variant of this is possible where the atoms have a have a parabolic trajectory. Wynands & Wyers note that this avoids the Dick effect [28] (where high frequency LO noise is aliased to a lower frequency through the process of sampling), but unfortunately loses the cancelling effect of phase shifts that would otherwise take place by atoms retracing their path through the cavity [27].

A further variation is to allow atoms to fall through a shielded magnetic cavity under gravity as per Ascarrunz et al. [29] and is sold by SpectraDynamics as a "cold Rubidium" reference [30]. As with all fountain clocks, trajectory must be vertical or the atoms may not pass through the detection zone. For the same reason, these devices are also very sensitive to acceleration and vibration – any change in physical position during the 78ms free-precessing time [31] can result in atoms not being detected.

1.6.5 Pulsed Optical Pumping (POP)

Pulsed Optical Pumping is the type most pertinent to this thesis: POP clocks use a three-phase Ramsey scheme to optically pump a vapour, then perform MW interrogation, then sample the population present in one of the ground states. As the three phases are separated in time, POP is largely immune to laser noise modulating the MW output. Apparatus is based on a vapour cell contained in a MW

enclosure with one or more lasers used for preparing and sampling the population. Long term stability is significantly improved over spectral lamps and CPT with Micalizio et al. reporting an Allan deviation of 1 $x10^{-14}$ ($\tau=10^5$s) [32]. Like a conventional fountain clock, operation is not continuous, and is based on periodic sampling.

1.7 Optical references

Optical clocks are based on electron transitions in the visible light spectrum i.e. in the range 430-770 THz. At the heart of each clock is an array of individual atoms, held in place by either an optical lattice or ion trap. Electrically neutral atoms are held in an optical lattice: an array created by the interference patterns resulting from counter-propagating laser beams. Charged atoms are held in an ion trap: a combination of magnetic and electric fields that can retain a single ion [33]. These supercooled atoms remain almost motionless in this structure and are therefore largely unaffected by Doppler broadening of absorption lines and collisions, which results in very narrow spectral lines. The LO in these clocks is a highly stable laser, typically an optical cavity-stabilised design, which uses a Fabry-Perot cavity made from Ultra Low Expansion (ULE) glass. The atoms held in the array are initially excited by one laser and then probed by a second, using fluorescence to determine which atoms remain in an active state, enabling the probing laser to be locked to the hyperfine transition. The final step, made possible by the advent of optical frequency combs, allows accurate division of this laser frequency so that it is practical to create a control circuit that can lock the laser to the hyperfine resonance. In 2000, a team based jointly at JILA and Bell Laboratories demonstrated one of the earliest successful femtosecond mode-locked lasers [34], and similar devices are now widely used as frequency combs.

Optical clock accuracy and stability now far exceed the best MW clocks, JILA's reported precision of $2.5x10^{-19}$ [35] mean that it is possible to detect a 2cm change in height due to relative dilation of time i.e. gravitational time dilation. In a summary of the field, Poli et al. note that working clocks have been demonstrated using a wide range of elements and ions: Al, Al^+, Ca, H, Hg^+, I, Sr, Sr^+, Yb, Yb^+ [36]; a list since widened to include In^+, Mg, and Th^+ [37]–[39]. Although this technology is expected to supersede the Hydrogen and Caesium references currently in use, these clocks do not yet reliably operate for extended periods of time [40]. Also, as noted by Gellesch et al. [41], [42], optical clocks are still physically large, typically greater than $1m^3$, with their own portable design a predicted 180l.

1.8 Allan Deviation and Noise

For some applications, a low-noise clock is of paramount importance. One example is that local oscillator noise present on the input of a superheterodyne receiver will also be present on the output, potentially causing an input to appear to spread into an adjacent communication channel. In a radar application, this effect degrades performance and appears as smearing of the input.

It is common practice to use a Phase Locked Loop (PLL) to multiply an external reference up to a required operating frequency with the output stability depending greatly on the how well the input phase adheres to a pure sine wave. Noise will manifest as a disturbance in phase and can cause corruption in high-speed digital communications where multiple devices each generate their own high-speed clock from a distributed reference signal, forming a plesiochronous system.
Noise can also be considered as jitter in the time domain. An Analogue to Digital Converter (ADC) sampling aperiodically will provide a poor representation of the measured signal.

In a 1966 paper, Allan states a need to understand "quality in an atomic frequency standard", differentiates between "flicker noise" and white noise, and proposes a method of determining the power spectral density [43]. By sampling an oscillator's output frequency at regular time intervals of τ, the average frequency y, Allen Variance (AVAR) $\sigma_y^2(\tau)$, and Allen Deviation (ADEV) σ_y can be calculated as [44]:

$$\bar{y} = \frac{1}{\tau}\sum y \tag{1}$$

$$\sigma_y^2(\tau) = \frac{1}{2} < [y(t+\tau) - y(t)]^2 > \tag{2}$$

$$\sigma_y(\tau) = \sqrt{\frac{1}{2} < [y(t+\tau) - y(t)]^2 >} \tag{3}$$

where <> denotes an average taken over infinite time.
Allan's method supports conversion of an N-sample variance measurement into its 2-sample equivalent, allowing direct comparison of measurements for any value of N. Figure 3 shows Weiss and Shenoi's use of ADEV measurements to distinguish between different sources of noise and illustrates how each is expected to contribute plotted against sample time τ [45].

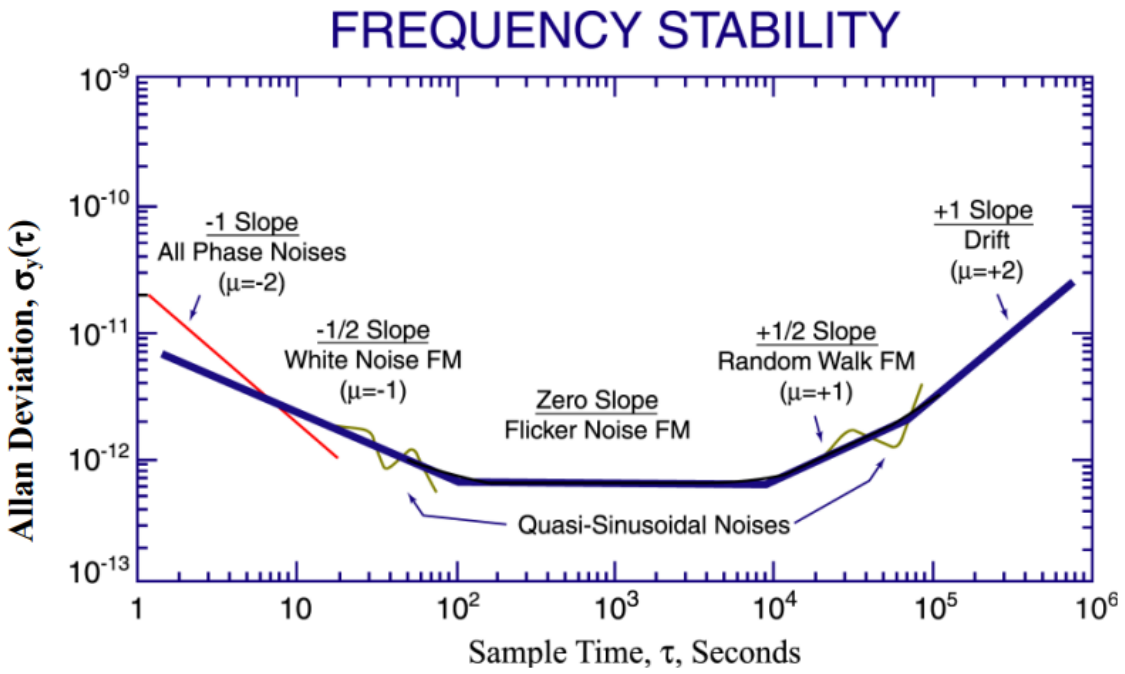

Figure 3. An example of frequency stability plotted as Allan Deviation against sample time [45].

1.9 Goals

This thesis has been jointly sponsored by the EPSRC and Leonardo MW Ltd. with the aim of producing the best low-noise clock possible inside the volume of half to one shoebox. Long-term stability should be comparable or better than current commercial alternatives, i.e. an Allan Deviation of $1x10^{-12}$ (τ=100sec), while also retaining high short-term accuracy. The reference should be portable, , maintain excellent accuracy over a period of several hours without the need for a GPS-derived time reference, and be able to tolerate vibration and acceleration that could be experienced on a moving vehicle or aircraft. The intent is to explore ideas, and prove the feasibility of emerging technology rather than productising and miniaturising it: specifically working at lower Technology Readiness Levels (TRLs) [46]. Lastly, there should be consideration for the engineering and design of a product that could be used in an extreme environment, either variation in temperature or vibration.

The need for portability, or specifically for accuracy whilst moving, rules out fountain clocks and the SpectraDynamics cRb-CLOCK which need perfect vertical alignment. As detailed in Section 1.6.4, , they are also ruled out as they can't operate when subject to acceleration or significant vibration. Optical clocks are currently ruled out due to volume requirements; high vacuum and cryogenic cooling requirements are also prohibitive. The use of a long Fabry Perot interferometer, or any linear device based on multiple MW wavelengths would fall foul of physical constraints. Lastly, any existing technology that would take significant development to miniaturise has been excluded.

1.10 The starting point.

Although the original intent was to improve on the performance of Burtichelov's CPT reference [11], it was determined that POP would offer significantly long-term better performance in a similar volume. The 10MHz OCXO developed by Everard at al. [11, Sec. 3.2], [47] was chosen as the Local Oscillator (LO) due to its excellent phase noise and short-term stability, as shown in Figure 4 and Figure 5. In these figures the shaded grey area shows the noise-floor of the test equipment (Symmetricom TSC 5120A) used for the measurement.

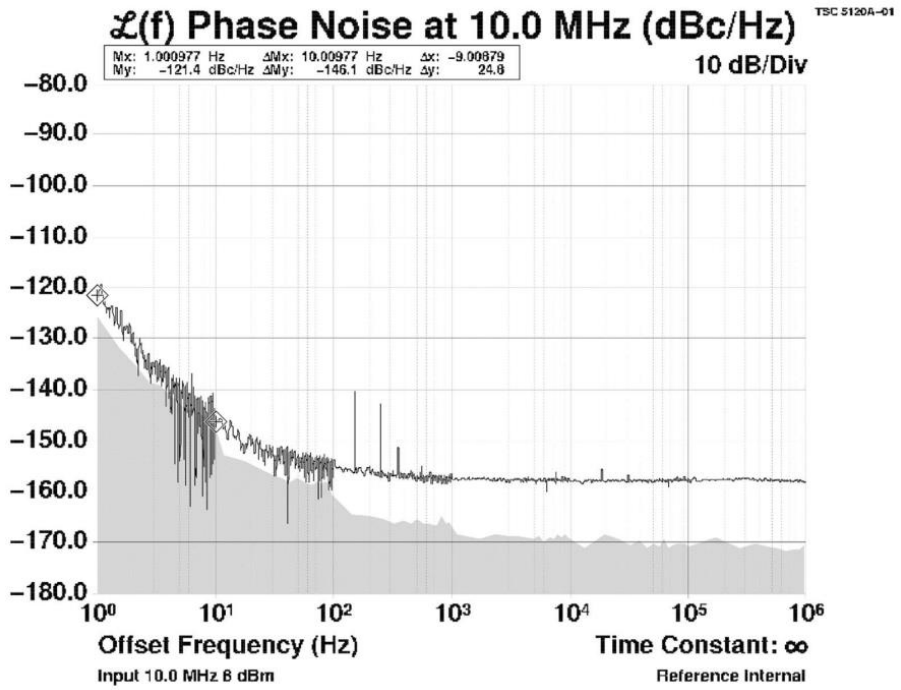

Figure 4. Burtichelov's phase noise measurement of the OCXO [11, Fig. 3.30].

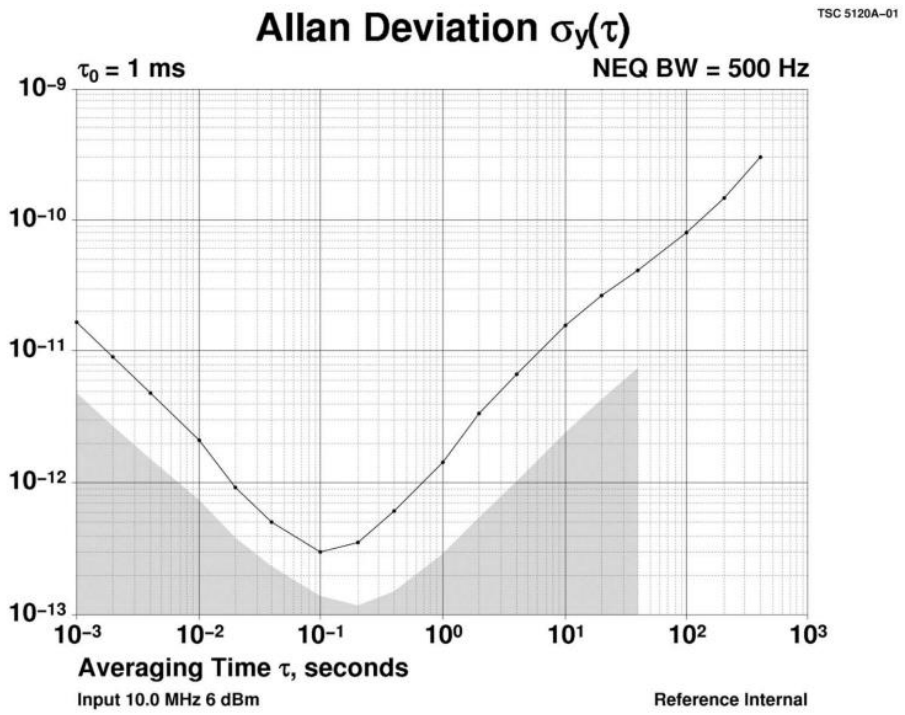

Figure 5. Burtichelov's Allan Deviation measurement of the OCXO [11, Fig. 3.31].

As Burtichelov's had already demonstrated CPT with his vapour cells, a spare ^{85}Rb cell, optimised around a turnover temperature of 50°C [11, Sec. 5.3.3] was used as the starting point of the POP Physics Package.

1.11 Performance goals

An updated comparison, modified from Burtichelov [11, p. 25], is presented below:

Table 1: Commercial clock specifications

		Safran LNRClok-1500	Stanford Research PRS10	Microchip SA.22c-LN	Microchip SA.35m	York Potential
Allan Deviation	τ=1s	2.0×10^{-12}	2.0×10^{-11}	1.4×10^{-11}	3.0×10^{-11}	2.0×10^{-13}
	τ=10s	2.0×10^{-12}	1.0×10^{-11}	8.0×10^{-12}	1.6×10^{-11}	5.0×10^{-14}
	τ=100s	1.0×10^{-12}	2.0×10^{-12}	2.5×10^{-12}	8.0×10^{-12}	2.0×10^{-14}
	τ=1000s					7.0×10^{-15}
	$\tau=10^4$s					6.0×10^{-15}
RF Output Phase Noise at 10MHz (dBc/Hz)	1Hz	-103		-100	-70	-123
	10Hz	-133		-130	-87	-149
	100Hz	-153	-130	-145	-114	-155
	1kHz	-158	-140	-150	-130	-160
	10kHz	-161		-155	-140	-160
	100kHz					-160
Volume (cc)		460	394	640	49.5	12000

The products chosen for comparison are restricted to those that could potentially meet the portability and size requirements and are therefore based on CPT or Rubidium vapour lamps. There are no

commercially available POP clocks at the time of writing. The short-term stability of a reference is measured as phase noise at specific offsets from the 10MHz output; any random fluctuations will produce jitter in the time domain which correspond to phase variations in the frequency domain. Long term stability is measured as an Allan deviation [48], [49], with observation time τ, and represents the frequency instability between two observations. The 'York potential' performance figures are based on the phase noise and short-term Allan Deviation of the group's existing Local Oscillator (LO) (Figure 4 and Figure 5) but with improved long-term ($\tau > 10s$) stability consistent with an atomic physics package limited by 1mK thermal stability.

2 Atomic Theory

2.1 An introduction to atomic theory for clocks

Anyone seeking to get to grips with atomic references for the first time would do well to read this document in conjunction with Kitching's 2018 overview of Chip-scale atomic devices [50] and Section 1 of Gharavipour's 2018 thesis [51].

Atomic physics is 'quantum' because each atom is composed of sub-atomic components and there are a small number of valid combinations for those components in the makeup of a given atom. For example, one atom of ^{85}Rb is composed of 48 neutrons, 37 protons and 37 electrons. As per the Bohr model, 36 of these electrons are tightly bound in 4 stable shells (containing 2, 8, 18 and 8 electrons) but the electron in the outer (5th) shell is only weakly bound and can exist in a number of states [52]. The atom's Hamiltonian $\hat{H}$, the total energy of the system, is the sum of the potential and kinetic energy of all the components and therefore the atom can have a number of different states, each with a well-defined energy level. By extension, an atom may move between these states if it gains or loses the correct amount of energy needed, either in the form of a photon received or emitted, or energy radiated.

2.1.1 Fine structure

Figure 6 shows the 'fine' and 'hyperfine structure' of 85Rb, the energy levels possible due to different states of its unbound (or 'valence') electron.

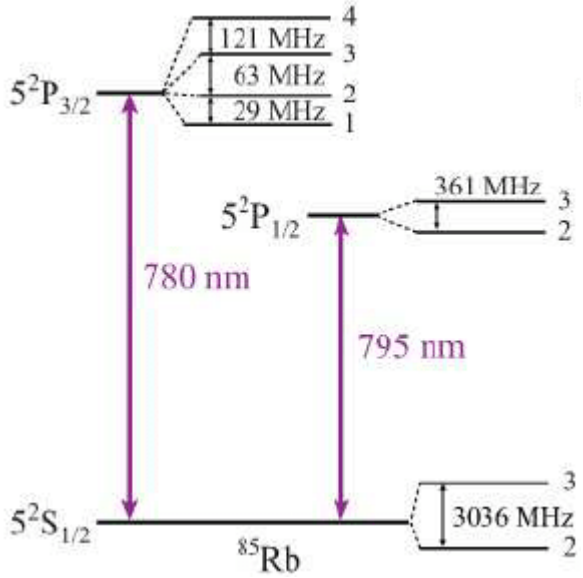

Figure 6. Fine and hyperfine structure for ^{85}Rb D1 (795 nm) and D2 (780 nm) transitions [53, Fig. 5].

The fine structure exists because the valence electron has orbital angular momentum **L**, as well as spin angular momentum **S**, and therefore the electron's total angular momentum **J** can have different

values. **S**, also called the quantum number, is ½ for all electrons. An atom's lowest energy states, its ground states, are represented by **L**=0 and the first higher-energy orbit possible is **L**=1.

As the electron can spin in either direction, **J** can be either ½ or 3/2, which results in two possible 'excited' energy states. These **L**=0 → **L**=1 transitions are called D lines as they were the 4th set of energy lines discovered during the advent of spectroscopy [54, p. 378]. The energy required for an electron to transition into one of these excited states is normally expressed in terms of the wavelength of the photon that can stimulate the transition.

$$E = hv = h\frac{c}{\lambda} \tag{4}$$

In equation (4) E is a photon's energy, h is Planck's constant, v is the frequency of radiation, c is the speed of light, and λ is the photon's wavelength. The D1 (lowest frequency) transition aligns with a wavelength of 794.979 nm in a vacuum, and D2 at 780.241 nm [55].

2.1.2 Hyperfine structure

In contrast with the fine structure, which is based on magnetic moments, the hyperfine structure is due to interactions between the valence electron and the nucleus. The total nuclear angular momentum **I** interacts with the valence electron's total angular momentum **J** to produce a total atomic angular momentum **F**.

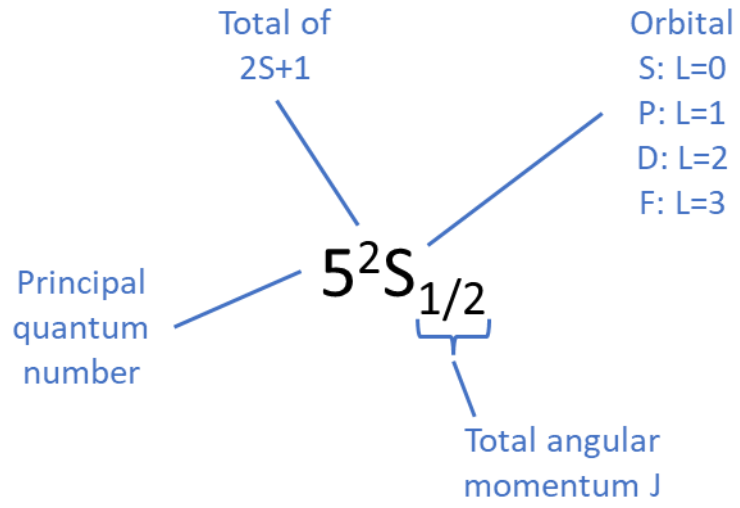

Figure 7. Components in an atomic term symbol.

Figure 7 shows how a 'term symbol' is used to label an atomic state based on the state of the valence electron(s). Using this terminology to identify hyperfine states - in ^{85}Rb's ground state ($5^2S_{1/2}$), **J**=1/2, **I**=5/2, and since F = J + I, **F** is either 2 or 3 [55]. For the excited state of the D2 line, $5^2P_{1/2}$, **F** can also be either 2 or 3. For the excited state of the D2 line, $5^2P_{3/2}$, **F** can be either 1, 2, 3 or 4. The two lowest

energy states, the ground states, are shown at the bottom of Figure 6 and are referred to as $5^2S_{1/2}$ F=2 and $5^2S_{1/2}$ F=3. A transition can be written e.g. $5^2S_{1/2} \rightarrow 5^2P_{3/2}$ which, in this example, represents the ^{85}Rb D2 transition.

2.1.3 Zeeman magnetic interaction

Each hyperfine energy state contains 2**F**+1 sublevels, which correspond to interactions between the valence electron's total magnetic moment, and that of the nucleus [15]. Kitching neatly illustrates these sublevels in Figure 8 for ^{87}Rb [50].

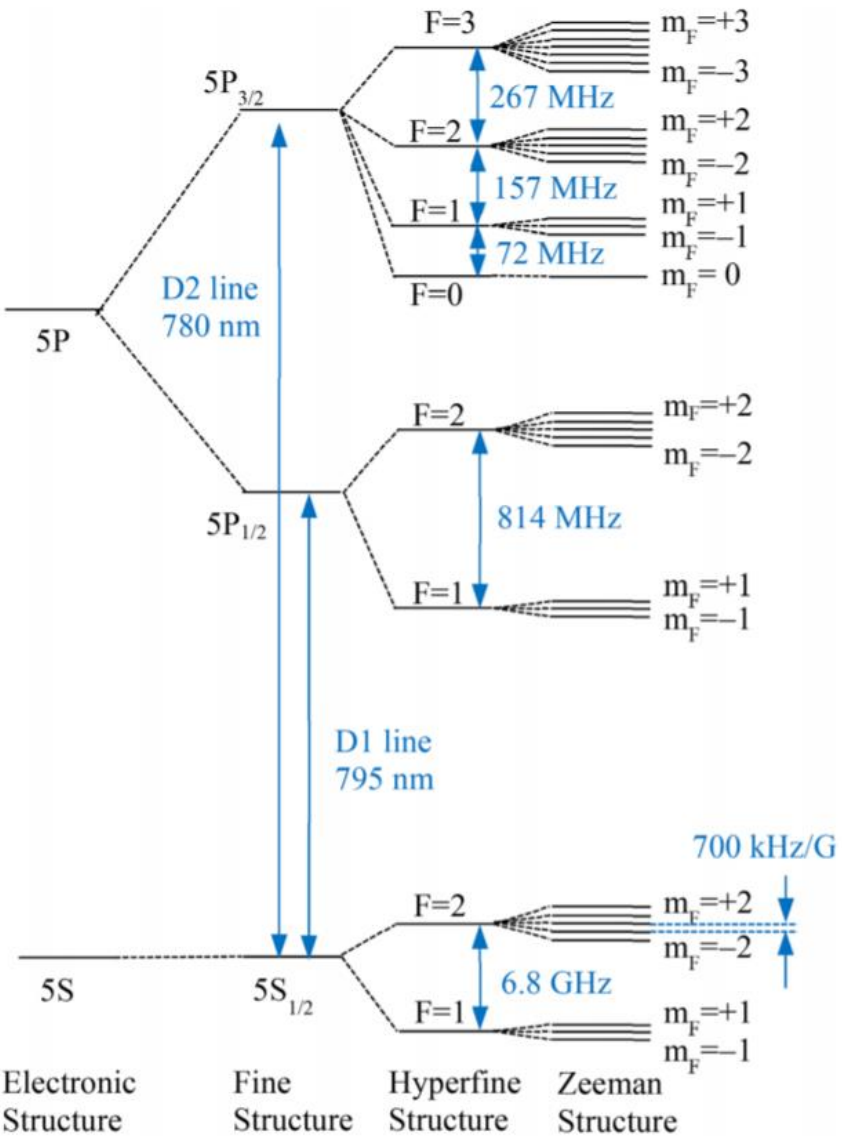

Figure 8: Energy spectrum of ^{87}Rb [50, Fig. 1]

In the absence of an external magnetic field these sub-levels are said to be 'degenerate' i.e. they share the same energy level and it would not be possible measure the energy to determine the quantum

state (value of m_F). However, in the presence of a static magnetic field each sub-level has a distinct energy level, sometimes referred to as being 'lifted out of degeneracy' or that 'degeneracy is broken'.

If the external field is weak (i.e. produces a much smaller frequency shift than the hyperfine structure) then these sub-levels are split linearly and called the Zeeman Effect [56]. Steck [55, Sec. 3.2.1] quantifies this change in quantum energy as:

$$\Delta E_{|F\ mF>} = \mu_B g_F m_F B_Z \qquad (5)$$

Where μ_B is Bohr's Magneton (the magnetic moment of an electron) [57], m_F is the Zeeman sub-level, B_Z is the magnetic field strength which is defined as aligned with the z-axis, and g_F is the hyperfine Landé g-factor which is calculated as:

$$g_F = g_J \frac{F(F+1) - I(I+1) + J(J+1)}{2F(F+1)} g_I \frac{F(F+1) - I(I+1) + J(J+1)}{2F(F+1)} \qquad (6)$$

Where g_I is the nuclear g-factor, and gJ is the fine structure Landé g-factor for a specific state, values for which can be found in [55, p. 18].

Figure 9 and Figure 10 show Zeeman frequency splitting as a function of magnetic field strength. Note that although these graphs extend beyond the Zeeman region (field strength <1000G) into the Paschen-Back region (field strength >1500G), atomic references require a very small C-field, typically 0.1G to 1G.

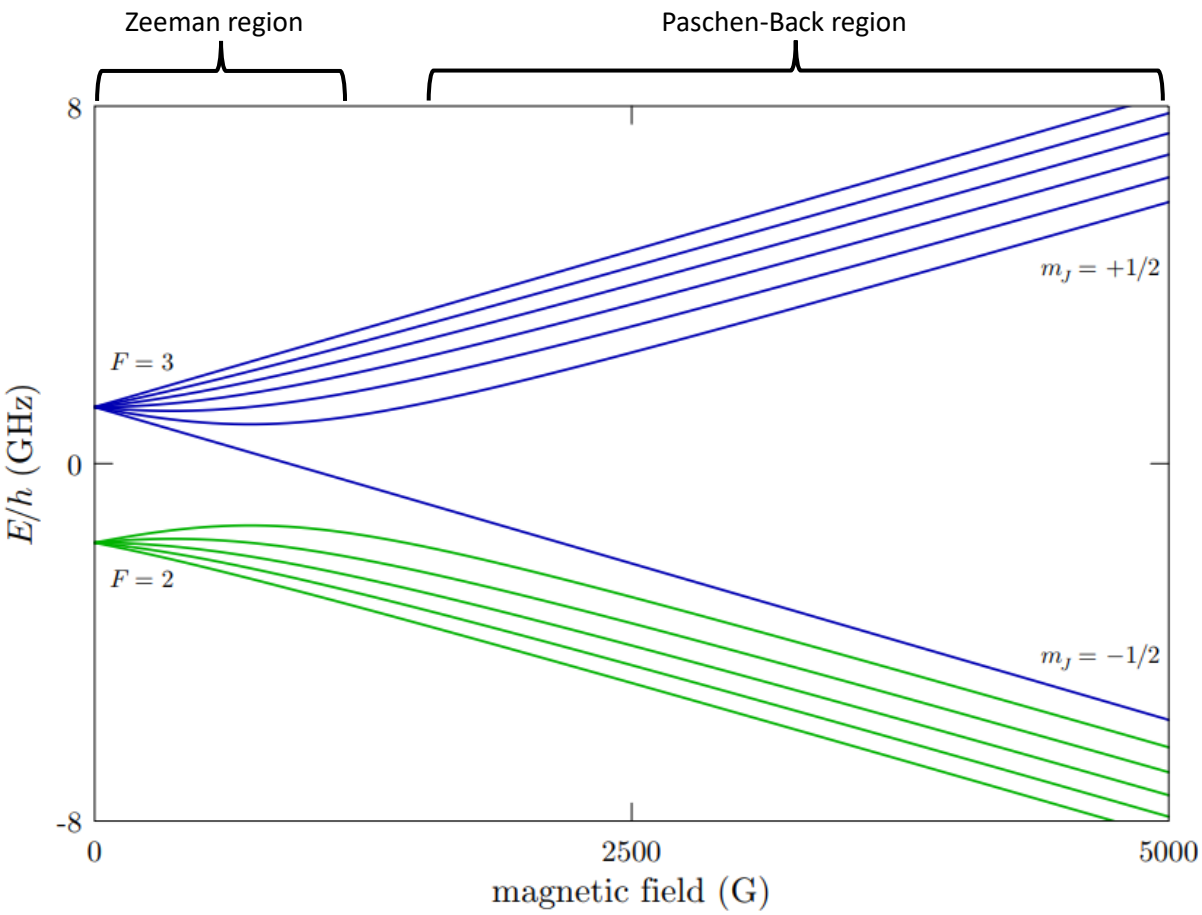

Figure 9: ^{85}Rubidium $5^2S_{1/2}$ (ground) level hyperfine structure in an external magnetic field [55, Fig. 4].

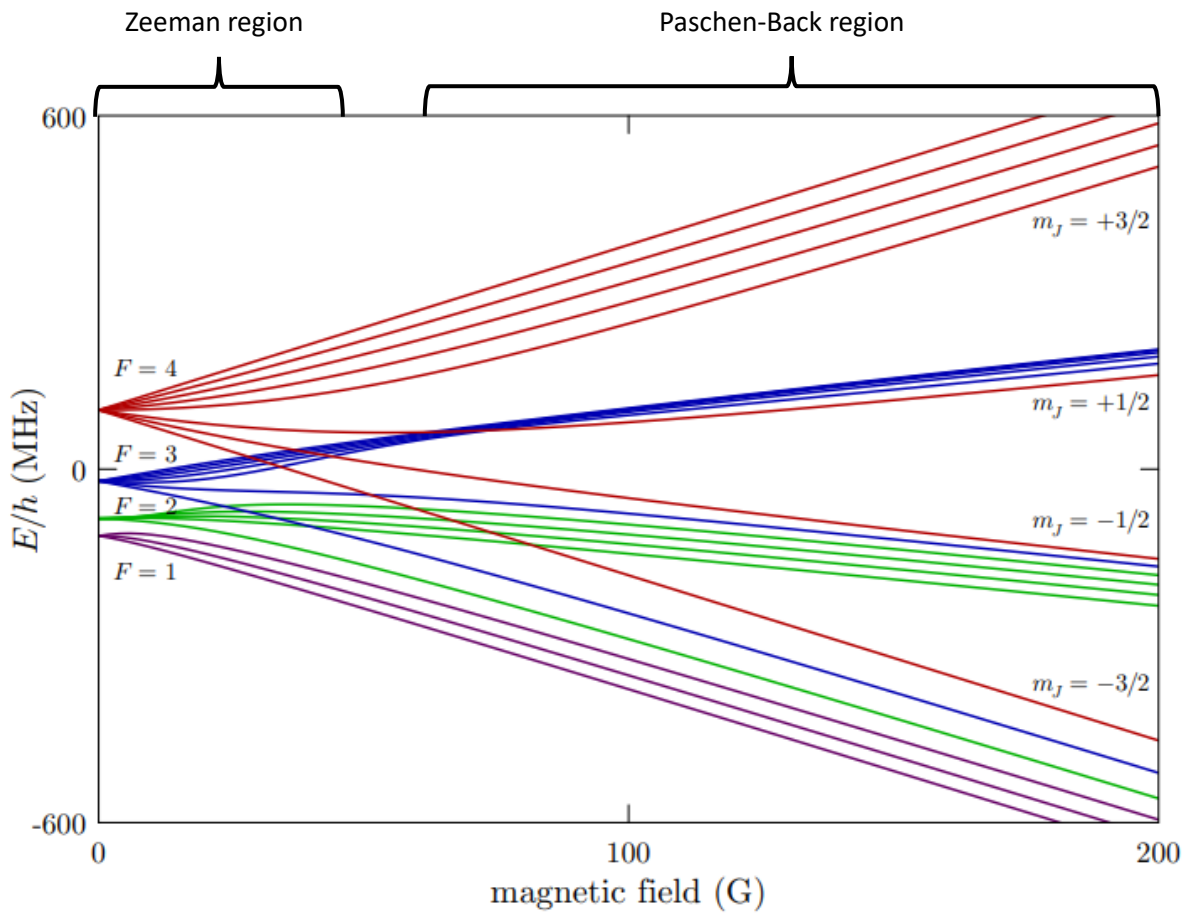

Figure 10. ^{85}Rubidium $5^2P_{3/2}$ D2-excited (i.e. 780nm) level hyperfine structure in an external magnetic field [24, Fig. 4].

As the m_F=0 sub-level is the least affected by the C-field it is often chosen for both CPT and POP references [32], [51], [58], [59]. However, when interacting with light, some of the atomic population will migrate from this sub-level. If light is circularly polarised (as in CPT) then photons will transfer their angular momentum, driving the atomic population into the 'furthest' sub-level, with a positive or negative value of m dependent on the direction of polarisation. Linearly polarised light will remain in the sub-levels around m=0 as per Steck [55, p. 4.3.2] and Gao [60].

2.2 Lambda system

A simplified diagram of ^{85}Rb's energy states is provided in Figure 11. Either of these can be used as the basis for an atomic reference. In an unperturbed state, the atomic population will be evenly distributed between the two ground levels $5^2S_{1/2}$ F=2 and F=3. If atoms from one of these ground states are optically pumped, as per Kastler and Brossel [13, p. 73], they will transfer to the excited state. Taking the $5^2S_{1/2}$ F=2 → $5^2P_{3/2}$ transition as an example, the atom will remain in that state for some nanoseconds before decaying back to one of the ground states. There is an equal probability that the atom will decay into either ground state, so by continuing to pump atoms from the F=2 ground state, one is de-populating that state, and filling F=3.

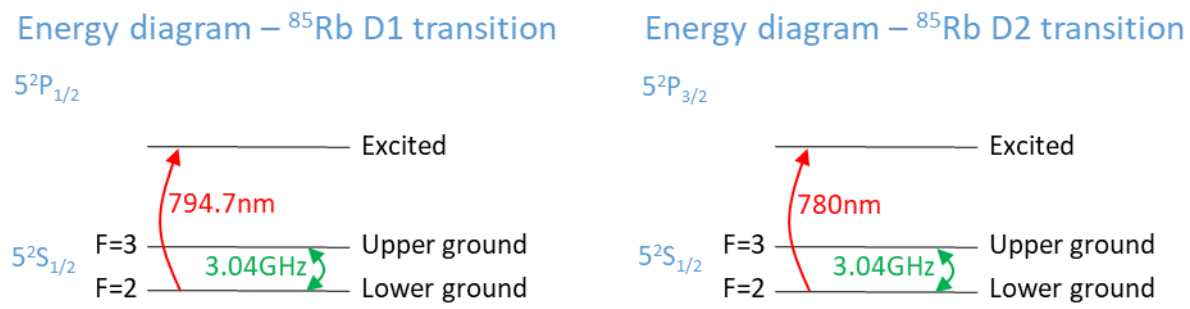

Figure 11. Energy diagram for two ^{85}Rb systems: one based on the D1 transition, one based on D2.

If an atom in either of the ground states is subject to a MW field oscillating at the hyperfine splitting frequency, the atom can transfer to the other ground state. In the example above, if the F=2 ground state is being depopulated by optical pumping, it can be re-populated by a MW field. This is the basis for Double Resonance (DR).

Rabi's molecular beam apparatus used magnetic selection to ensure that only one ground state was populated, then passed a constant beam of atoms through a rotating magnetic field, then through another magnetic selector that rejected all atoms that had not switched ground state [61]. The fixed path also fixed the interaction time, τ_m. Gharavipour [51] shows that:

> the probability of the atom swapping ground states is a (sinc (f))^2 function of MW frequency [46, eq. 1.13];
>
> the width of the central peak is inversely proportional to interaction time τ_m; [46, eq. 1.14]
>
> and that the amplitude of the central peak is also inversely proportional to interaction time τ_m [51, Sec. 1.3.2].

The interaction time needed to transition a population from one ground state to its complement has been retrospectively called a π pulse, after the complex representation of an atom's state on a Bloch sphere.

Ramsey's work takes this a step further and demonstrates that a π pulse can be replaced with two pulses, each half as long and separated by free-precessing time, and the width of the central peak is inversely proportional to the free-precessing (or Ramsey) time. Ramsey also demonstrates that it's possible to have the interrogation pulses separated in time rather than space [62]. To represent the shorter MW interrogation pulses Ramsey coins the expression $\pi/2$. Figure 12 shows an example of Ramsey fringes measured by Micalizio et al. [32].

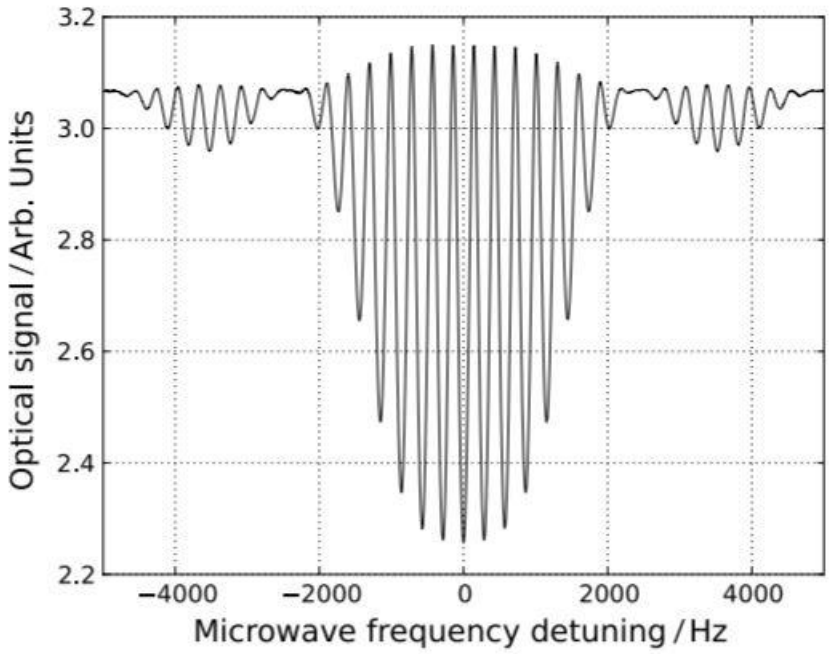

Figure 12: Ramsey fringes, taken from Micalizio et al. [32, Fig. 9].

Figure 12 has been captured by pumping an ^{87}Rb vapour cell with a 780nm laser, then exposing the cell to two microwave (6.834 GHz) interrogation pulses separated by Ramsey time T, followed by a low-intensity 'probe' laser beam to determine the atomic population remaining in the F=1 ground. The timing diagram is provided in Figure 13 and shows one complete atomic clock cycle.

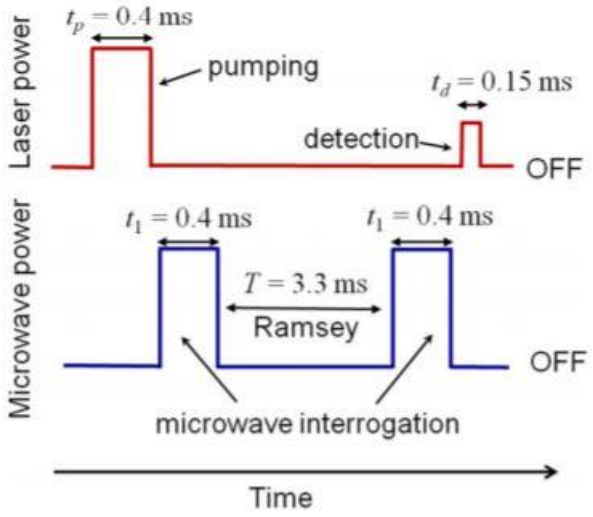

Figure 13: POP timing sequence, with optical detection [32, Fig. 7]

2.3 Hyperfine stability and the use of vapour cells

Although the hyperfine resonance provides an excellent frequency discriminating reference, the resonance is defined not only by the choice of element, but is weakly affected by temperature, pressure, magnetic fields, and light. CPT and DR take place in a sealed cell, typically glass, filled with a

small quantity of the required isotope. Although the first work on atomic spectroscopy used evacuated cells, i.e. filled in a vacuum so that no other elements were present, most references use a cell containing buffer gas that interacts only very weakly with the alkali atoms. Argon, Helium, Neon and Nitrogen are often used and cause Brownian, rather than ballistic motion, reducing collisions with the walls but suffering some frequency shift and linewidth broadening as described by Arditi and Carver [24]. Adjusting the pressure of the buffer gas allows some optimisation; greater pressure or a larger cell means the alkali atoms will collide more frequently with buffer gas atoms, lower pressure will mean more collisions with the cell walls. Using ^{85}Rb and a fixed temperature, Kitching illustrates how linewidth can be optimised through careful selection of cell size and buffer gas pressure [50, Fig. 2].

As described in the introduction, Burtichelov had already ordered a number of Rb cells, using Argon and Nitrogen as buffer gases, optimised around an inversion temperature of 50℃. No analysis or characterisation of their relaxation time was available prior to the writing of this thesis, but since this is crucial to clock performance it has been performed below.

Cell relaxation γ is defined as the rate at which atoms lose their polarisation, in units of s^{-1}. Rb atoms in a cell can lose polarisation in collisions with the buffer gas, with the cell walls, and through spin exchange with other Rb atoms. These are not a function of an external electromagnetic field and are therefore termed “intrinsic” relaxations [63].

As per Gharavipour [51, Sec. 1.6], total cell relaxation is:

$$\gamma_i = \gamma_{iCW} + \gamma_{iBG} + \gamma_{iSE} \tag{7}$$

where γ_{iCW} is the cell-wall relaxation rate, γ_{iBG} is buffer gas relaxation rate, γ_{iSE} is the spin exchange rate, and γ_1 and γ_2 are the relaxation rates for ground-state population and coherence respectively. As per Lindvall et al. [64], the overall coherence relaxation is calculated as:

$$\begin{aligned}\gamma_2 = &\left[\left(\frac{2.405}{r}\right)^2 + \left(\frac{\pi}{L}\right)^2\right] p_0 \frac{D_{N2} D_{Ar}}{D_{N2} p_{Ar} + D_{Ar} p_{N2}} \\ &+ \frac{N_0}{p_0}\left(\bar{v}_{Rb-N2}\sigma_2^{N2} p_{N2} + \bar{v}_{Rb-Ar}\sigma_2^{Ar} p_{Ar}\right) + \frac{6I+1}{8I+4} N_{Rb} \bar{v}_{Rb-Rb} \sigma_{se}\end{aligned} \tag{8}$$

The three terms in this equation represent cell wall collisions, pressure broadening, and spin exchange collisions, respectively, based on Vanier & Audoin [65].

2.3.1 Collisions with the cell wall

Collisions with an uncoated cell wall completely depolarise an atom and Vanier et al. confirm that population and coherence relaxation are equal [66]. As per (9), both γ_{1CW} and γ_{2CW} are therefore:

$$\gamma_{cw} = \left[\left(\frac{2.405}{r}\right)^2 + \left(\frac{\pi}{L}\right)^2\right] p_0 \frac{D_{N2} D_{Ar}}{D_{N2} p_{Ar} + D_{Ar} p_{N2}} \tag{9}$$

where Table 2 provides a list of the symbols and their numerical values

Table 2. Symbols used in (9) and their numerical values.

Symbol	Description	Value
r	Cell internal radius	0.0125 m
L	Cell internal length	0.025 m
p_0	Reference pressure	760 Torr
p_{N2}	N_2 partial pressure	7.3 Torr [58]
p_{Ar}	Argon partial pressure	12.7 Torr [58]
D_{N2}	Diffusion constant of ^{85}Rb in N2 at p_0 (60°C) Scaled to 50°C (as per [67])	1.59 x10^{-5} m^2/s [68] 1.52 x10^{-5} m^2/s
D_{Ar}	Diffusion constant of ^{85}Rb in Ar at p_0 (70°C) Scaled to 50°C (as per [67])	1.2 x 10^{-5} m^2/s [66] 1.1 x10^{-5} m^2/s

Cell wall relaxation is therefore γ_{cw} = 17.2 s^{-1}

2.3.2 Buffer gas relaxations

The use of a buffer gas greatly extends a cell's relaxation time by interrupting an atom's ballistic motion and increasing time before a collision with the cell walls. An atom can typically undergo hundreds of thousands of collisions with a buffer gas before loss of coherence, as there is little spin interaction between Rb and Ar/N2 [50]. The greatly reduced mean-free-path (less than 1µm at 20 Torr) also suppresses much of the Doppler broadening from Rb's thermal motion, in a process known as Dicke narrowing [69].

The RMS velocity v of an ideal gas is given by:

$$v = \sqrt{\frac{3kT}{m}} \tag{10}$$

where k is Boltzmann's constant, T is the absolute temperature and m is the mass of an atom or molecule. The average relative velocity of two gases can be calculated as the geometric mean of their individual velocities.

As per (9), buffer gas relaxation is calculated as:

$$\gamma_{1BG} = \frac{N_0}{p_0}\left(\bar{v}_{Rb-N2}\sigma_1^{N2}p_{N2} + \bar{v}_{Rb-Ar}\sigma_1^{Ar}p_{Ar}\right) \tag{11}$$

$$\gamma_{2BG} = \frac{N_0}{p_0}\left(\bar{v}_{Rb-N2}\sigma_2^{N2}p_{N2} + \bar{v}_{Rb-Ar}\sigma_2^{Ar}p_{Ar}\right) \tag{12}$$

where Table 3 provides a list of the symbols and their numerical values.

Table 3. Symbols used in (12) and their numerical values.

Symbol	Description	Value
k	Boltzmann constant	1.380649×10^{-23} J/K
p_0	Reference pressure	760 Torr
p_{N2}	N_2 partial pressure	7.3 Torr [58]
p_{Ar}	Argon partial pressure	12.7 Torr [58]
N_0	Loschmidt's constant	2.69×10^{25} m^{-3}
$\bar{v}_{Rb-N2}$	Average relative velocity between Rb and N2 molecules	619 m/s
$\bar{v}_{Rb-Ar}$	Average relative velocity between Rb and Ar atoms	545 m/s
$\bar{v}_{Rb-Rb}$	Average relative velocity between Rb and Ar atoms	436 m/s
σ_1^{N2}	Population cross section for ^{85}Rb-N2 collisions	8.3×10^{-27} m^2 [51]
σ_1^{Ar}	Population cross section for ^{85}Rb-Ar collisions	4.9×10^{-26} m^2 [66]
σ_2^{N2}	Decoherence cross section for ^{85}Rb-N2 collisions	7.43×10^{-26} m^2 [66]
σ_2^{Ar}	Decoherence cross section for ^{85}Rb-Ar collisions	37.1×10^{-27} m^2 [66]

Buffer gas relaxation rates are therefore γ_{1BG} = 13.3 s^{-1} and γ_{2BG} = 21.0 s^{-1}.

2.3.3 Spin exchange

Spin exchange happens when two Rb atoms with opposite spins collide and exchange their spin in doing so [51], [70]. These collisions can also result in spin destruction but as this effect has 1000x less impact (per the σ_{SE} and σ_{SD} values in Table 4) I am ignoring this process [71]. Spin exchange is calculated as [51]:

$$\gamma_{1SE} = N_{Rb}\bar{v}_{Rb-Rb}\sigma_{SE} \tag{13}$$

$$\gamma_{2SE} = \frac{6I+1}{8I+4}\gamma_{1SE} \tag{14}$$

Per Shang et al. [72], Rb number density is calculated assuming saturated vapour, and equals:

$$N_{Rb} = \frac{1}{T}10^{32.178-\frac{4040}{T}} \tag{15}$$

This figure is based on data from Alcock et el. [73] which does not distinguish between Rb isotopes; presumably the group characterised a naturally occurring mix of ^{85}Rb and ^{87}Rb. For a single isotope the calculated value of N may therefore be less accurate than the 1% experimental results suggest.

Table 4 provides a list of the symbols and their numerical values.

Table 4. Symbols used in (13), (14), and (15) and their numerical values.

Symbol	Description	Value
$\bar{v}_{Rb-Rb}$	Average relative velocity between Rb and Ar atoms	436 m/s
I	^{85}Rb nuclear spin	3/2
N_{Rb}	Rb number density (50°C)	1.47×10^{17} m^{-3}
σ_{SE}	Spin exchange cross section for ^{85}Rb	2.54×10^{-18} m^2 [66]
σ_{SD}	Spin destruction cross section for ^{85}Rb	1.6×10^{-21} m^2 [71]

Spin exchange relaxation rates are therefore γ_{1SE} = 162 s^{-1} and γ_{2SE} = 102 s^{-1}.

2.3.4 Relaxation totals

Cell relaxation is summarised in Table 5 with spin exchange being the dominant effect in loss of atomic spin.

Table 5. Intrinsic cell relaxation values.

	Ground state population γ_1(s^{-1})	**Coherence γ_2(s^{-1})**
Spin exchange	162	102
Buffer gas collisions	13.3	21.0
Cell wall collisions	17.2	17.2
TOTAL	193	140
Relaxation time (ms)	5.2	7.2

Due to the relatively low cell temperature, this is a favourable result even compared with two recent high-performance Ramsey-DR clocks [51], [74]. Keeping the temperature low also slows the rate at which Rb reacts with the cell walls, prolonging cell lifespan.

2.4 Source of frequency shifts

In his 2018 thesis, Gharavipour quantifies the sources of frequency shifts and their contribution to his clock based on the D2 transition of a ^{87}Rb Physics Package [51, Sec. 1.8]. Figure 14 lists the contribution of each effect.

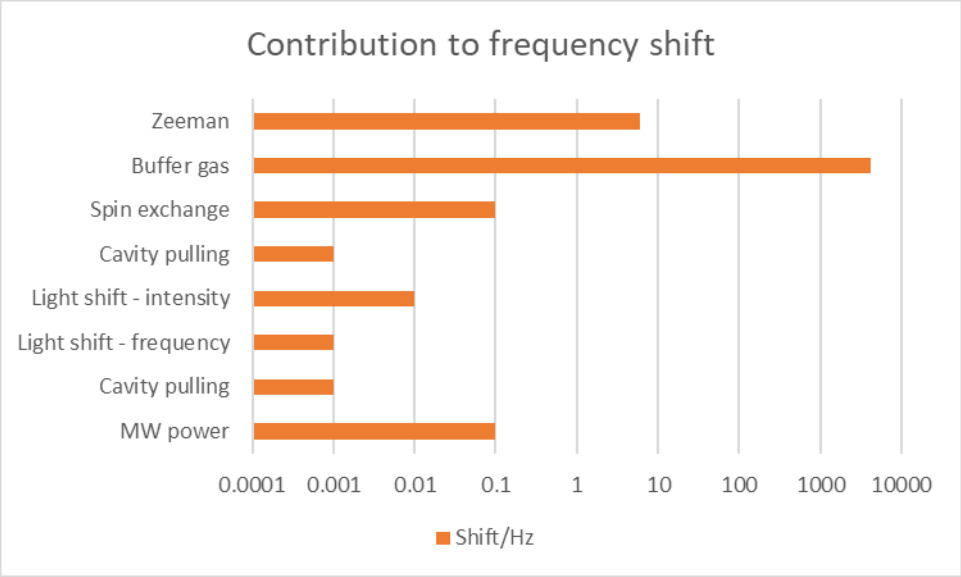

Figure 14. Summary of perturbing frequency shifts for Gharavipour's Ramsey-DR Rb clock.

Of these, the two greatest factors are the temperature dependent contribution from buffer gas pressure, and Zeeman splitting which is dependent on magnetic field strength.

2.4.1 Buffer gas shift

The use of a buffer gas is not entirely beneficial. Collisions with the buffer gas modify the interaction between the valence electron and nucleus. Van der Waals forces tend to pull the valence electron further from the nucleus, but collisions with a buffer gas force it closer, with a resulting change in the hyperfine coupling. Happer shows the change in contact interaction δA(r)I.S is inversely proportional to time between collisions, therefore proportional to gas pressure [14].

In Table 6, the pressure and temperature coefficients for N2 and Ar are known, and taken from Burtichelov [58, Sec. 5.3.3]. The buffer gases are present in a known ratio so ratio r = 1.74 and the 'mix' coefficients for the buffer gas are calculated as [51]:

$$x = \frac{x_{N2} + r x_{Ar}}{1 + r} \tag{16}$$

Where x is either the pressure coefficient β, linear temperature shift coefficient δ, or quadratic temperature shift coefficient γ. Results are presented in Table 6.

Table 6. Calculated pressure coefficients for ^{85}Rb cell at 50°C.

Buffer gas	**β (Hz Torr^{-1})**	**δ (Hz Torr^{-1} K $^{-1}$)**	**γ (HzTorr^{-1} K $^{-2}$)**
N2	-26.78 [58]	-0.1555 [58]	1.56x10^{-04} [58]
Ar	248.733 [58]	0.2399 [58]	5.77x10^{-04} [58]
Mix	73.8	-1.12x10^{-02}	3.10x10^{-04}

Given the operating conditions of P_0 = 20 Torr, and T = 323.15K, the absolute frequency shift due to buffer gas is therefore 2.05kHz, with a temperature sensitivity of 0.2Hz/K at the inversion temperature.

2.4.2 Second-order Zeeman shift

As per Section 2.1.3, the presence of a static external C-field both provides a quantisation axis for the atoms, and introduces Zeeman splitting. The Breit-Rabi formula can be used to calculate the frequency shift for the clock transition, and as m_F=0 there is no first-order Zeeman shift. The second-order shift can be calculated as [51, Sec. 1.7.3.1]:

$$\Delta\nu_{zeeman} = \frac{(g_J\mu_B)^2}{2hE_{HFS}}B_0^2 = A_0B_0^2 \qquad (17)$$

where Planck's constant h is 6.626 x10^{-34} J.s and the Bohr Magneton μ_B is 9.274 x10^{-24} J/T. For ^{85}Rb the Landé g-factor g_J ($5S_{1/2}$) is 2.0023 [55], and Hyperfine splitting energy E_{HFS} is 2.012 x10^{-24} J (calculated as h x 3.0357 GHz). A_0 is therefore 1290 Hz/G^2 which is in good agreement with the qualification of the ^{85}Rb quadratic Zeeman shift (5S1/2, F = 2, mF = 0 → 5S1/2, F = 3, mF = 0) by Li et al. [75] (1297 Hz/G^2) and Steck's figure of 1293 Hz/G^2 [55].

For a 100mG C-field this equates to a quadratic shift $\Delta\nu_{Zeeman}$ of 13Hz and a Magnetic coefficient $\delta_{B_0}^{\nu_{Zeeman}}$ of 519 Hz/G.

2.4.3 Cavity pulling

In an atomic reference where a cavity creates a standing wave for MW interrogation, it can suffer from two unhelpful effects. If Q is large then the time for the MW field to establish and decay may become significant, lengthening the effective π/2 pulse. This can be controlled by using a PIN diode to reduce the loaded Q.

If the resonant frequency of the cavity is not tuned with the hyperfine frequency then it will 'pull' the hyperfine frequency [76]. Since a waveguide dissipates all its energy each cycle, there no Q and no mechanism for a small change in length to cause a frequency shift $\Delta\nu_{CP}$, removing the need for varactor-based frequency tuning as used by Godone et al. [77]. Note that the cutoff frequency f_c and s-parameters will change if the waveguide is subject to thermal expansion therefore MW field strength inside the waveguide is sensitive to temperature.

2.5 Choice of clock transition

Many optical Rb clocks chose the 795nm D1 transition [51], [78]–[82] for reasons well explained by Stähler et al. [83], resulting in a reduction of optical contrast observed and greater resonant width. The D2 line has four excited states (F' = I ± 1/2, I ± 3/2) but only the (F' = I ± 1/2) states can relax into either of the ground states (as per selection rules of an electric dipole transition). The $5^2P_{3/2}$ F=4 state will always transition to $5^2S_{1/2}$ F=3, and $5^2P_{3/2}$ F=1 will always transition to $5^2S_{1/2}$ F=2. Figure 15 shows that when the D2 line is used to excite the ground states, only $5^2P_{3/2}$ F=2 and F=3 levels permit atoms to swap ground states in a single iteration of pumping and relaxation.

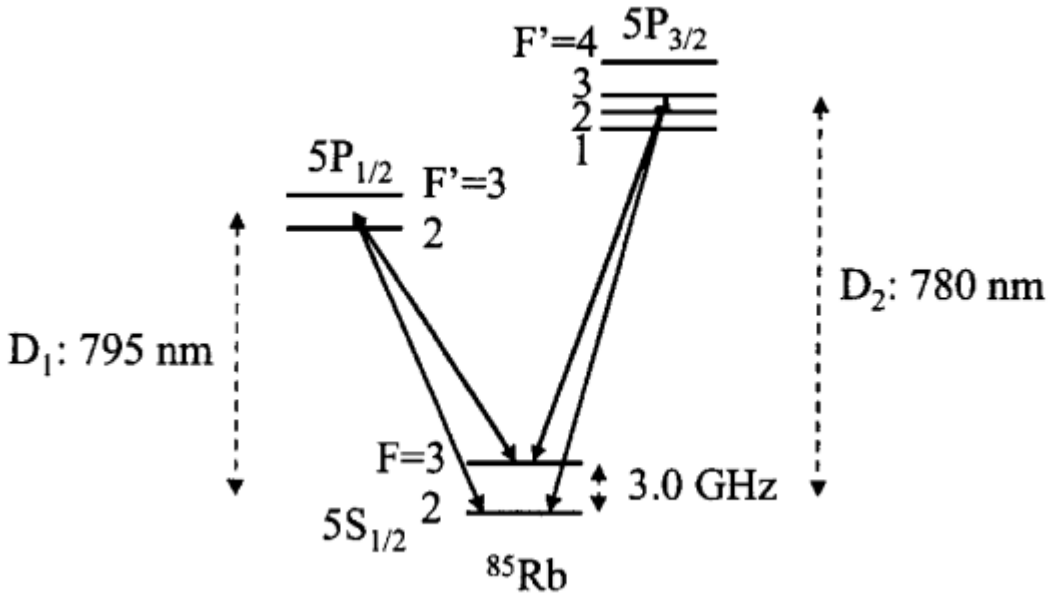

Figure 15. Energy diagram comparing D1 and D2 resonances for use in CPT [83, Fig. 1b]

For CPT this is a significant concern however POP does not use the Lambda system in the same manner. Although the presence of two extra excited states will reduce the pump beam's effectiveness at de-populating a ground level, this can be mitigated by a longer laser pulse. MW interrogation will be unaffected, and it may even prove a slight benefit during the probe phase, lengthening the time the ground state population can be sampled by a low-power laser pulse.

It's clearly entirely possible to build a reference using the D2 line - Lindvall et al. use this transition for a CPT reference [64] and Baryshev et al. use D2 for POP [84]. The decision to use a 780nm laser was further influenced by the availability of a collimator keyed to the 'pigtail' output from the laser diode (LD) covered in the next chapter.

3 Physics Package

3.1 A foreword about safety

A CAT 3B Near-infrared (NIR) laser is used to pump the Rb sample and since the output can cause retinal damage, a number of steps have been followed as per institutional requirements: all experimentation has taken place in a dedicated laboratory; the optical table used has been fitted with an inter-locked safety enclosure (as seen in Figure 16); the beam is normally enclosed in tubes ('entubed') to minimise the risk of light escaping; safety goggles are worn; training was undertaken; operational procedures written and followed; and periodic inspections ensure that equipment continues to meet the required standard.

Figure 16. Safety enclosure used to encapsulate the experiment, with laser driver and power supplies to left.

The enclosure has been partially disassembled in some of the photos, sometimes to aid significant realignment, sometimes to get a better photograph.

3.2 Overview

The Physics Package is normally defined as including anything that has a physical effect on the atomic sample. In this case that includes the Rb cell, the microwave enclosure, a pair of Helmholtz coils used to generate a C-field, heating, and the full optical path including a laser source and optical detection. Figure 17 shows the complete package with labelled components including the Laser Diode (LD), Acousto-optic Modulator (AOM), a Neutral Density Filter (NDF), a microwave generator (MW), and an Analogue to Digital Converter (ADC).

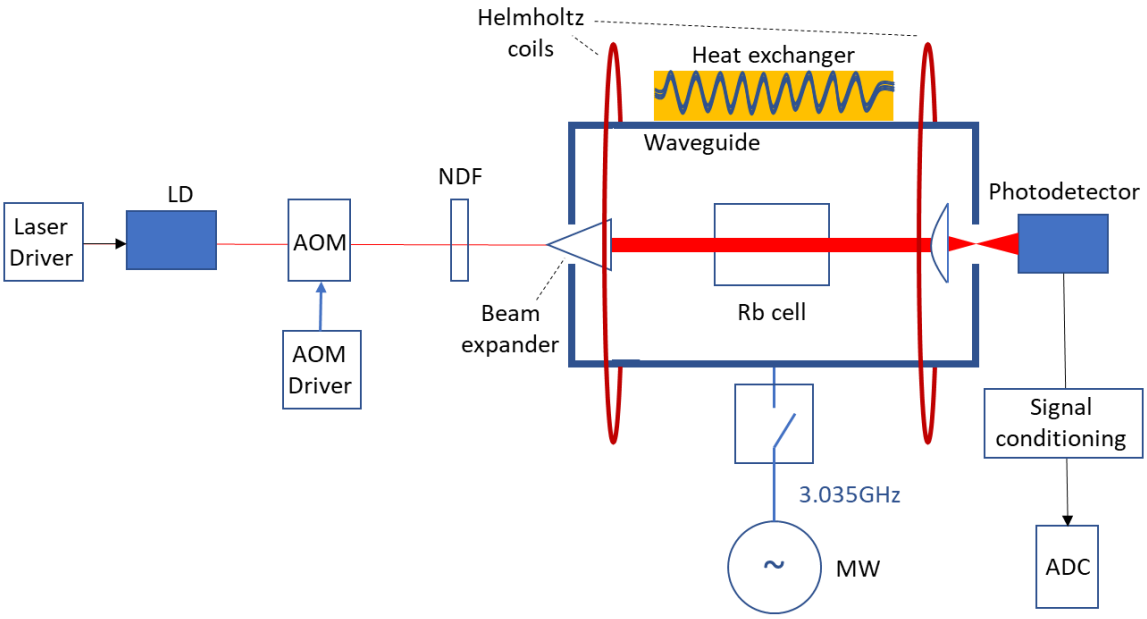

Figure 17. Diagram of the complete Physics Package

3.3 Rb Cell

As analysed in Section 2.3 and Section 2.4, one of Burtichelov's vapour cells is at the heart of the Physics Package and is pictured in Figure 18.

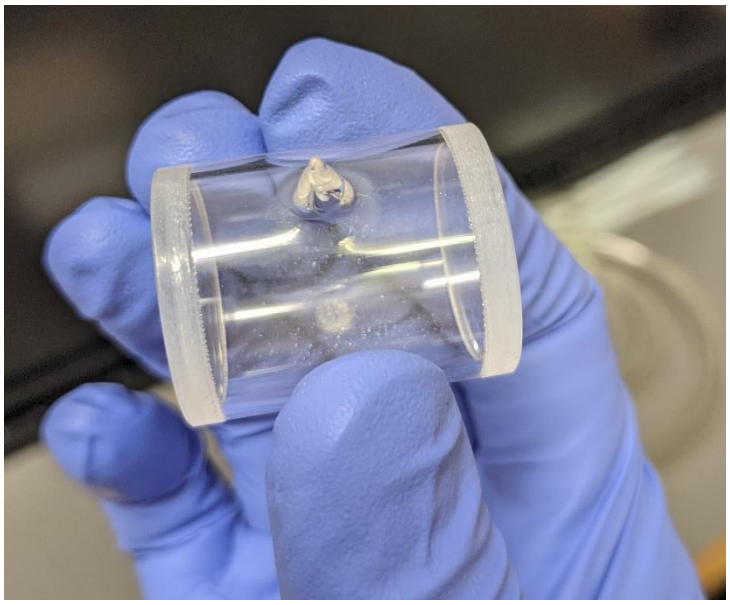

Figure 18. ^{85}Rb Pyrex-walled reference cell with Argon/Nitrogen buffer gas.

The end walls have been treated with an anti-reflective coating to aid NIR optical performance, but internal surfaces are uncoated.

3.4 Microwave (MW) enclosure

As a number of groups had successfully demonstrated POP using a MW cavity [77], [85], the same approach was initially employed, and a copper cavity was designed and constructed with dimensions appropriate for ^{85}Rb i.e. supporting a TE_{011} resonance at 3.035GHz. Full details are included in Appendix

B, including the field simulations and performance of the prototype. However, when faced with the choice between continuing to follow in others' steps by designing a TM_{111} mode choke for the cavity, and trying something completely different, further work on the cavity was shelved in order to make an original contribution to knowledge.

The concept of replacing the cavity with a waveguide proved very enticing, although it was not clear whether a travelling wave could be used for the π/2 pulses in place of the standing wave found in a cavity. As electromagnetic modes differ between rectangular section and a cylindrical body, the former can be used to suppress a number of unwanted modes. There looked to be potential for design simplification, as well as ease of construction, reduction of volume, and novelty. Although there an absence of any documented work conducting POP in a waveguide, Kasevich et al. used waveguide to enable MW interrogation in a fountain clock [86].

For a chosen frequency, and appropriately excited, the dominant mode in a rectangular waveguide is TE_{10}. By choosing dimensions appropriate for the MW frequency, TE_{10} can propagate but other modes, notably TE_{01} and TE_{11}, cannot. This can be used to expose a cell to an even field at 3.035 GHz. Pozar [87] derives the cut-off frequency for a particular mode as per the following equation (equation 3.84 in the source material)

$$f_{c_{mn}} = \frac{k_c}{2\pi\sqrt{\mu\varepsilon}} = \frac{1}{2\pi\sqrt{\mu\varepsilon}}\sqrt{\left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2} \qquad (18)$$

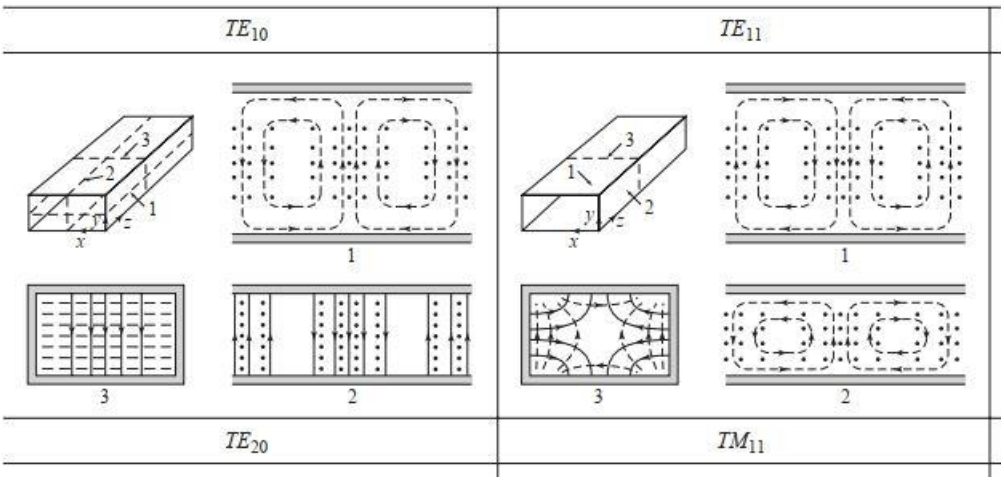

Figure 19. Field lines for TE_{01} and TE_{11} modes of a rectangular waveguide. Dashed lines are the B field. Solid lines represent the E field. [87, Fig. 3.14], [88, Sec. 8]

In his thesis, Burtichelov specified ^{85}Rb cells with an external length of 28.5mm diameter of 26mm [58, Sec. 5.3.3]. Since an ^{85}Rb cell was available to this specification it was used as a starting point for the design of a waveguide. Therefore, if a rectangular structure met the following constraints then it would serve the purpose:

i. Waveguide is sufficiently wide to hold the cell i.e. a > 0.0285m

ii. Waveguide is sufficiently tall to hold the cell i.e. b > 0.026m

iii. TE10 can propagate at 3.035 GHz i.e. cut-off frequency for TE_{10} < 3.035 GHz

iv. TE01 cannot propagate at 3.035 GHz i.e. cut-off frequency for TE_{01} > 3.035 GHz

v. TE11 cannot propagate at 3.035 GHz i.e. cut-off frequency for TE_{11} > 3.035 GHz

vi. Length of waveguide to be minimised

From Pozar's equation

$$f_{c_{01}} = \frac{1}{2\pi\sqrt{\mu\varepsilon}}\sqrt{\frac{\pi^2}{b^2}} \tag{19}$$

And

$$f_{c_{10}} = \frac{1}{2\pi\sqrt{\mu\varepsilon}}\sqrt{\frac{\pi^2}{a^2}} \tag{20}$$

So the minimum width (value of a) that will support TE_{10} at 3.035 GHz is

$$\frac{1}{2f_{c_{10}}\sqrt{\mu\varepsilon}} = 4.94cm \tag{21}$$

Ensuring that TE01 does not propagate at 3.035 GHz, (setting f_{c01} at 3.4 GHz to ensure adequate attenuation) means that a maximum height (value of b) is

$$\frac{1}{2f_{c_{01}}\sqrt{\mu\varepsilon}} = 4.41cm \tag{22}$$

Restating for TE_{11}

$$f_{c_{01}} = \frac{1}{2\pi\sqrt{\mu\varepsilon}}\sqrt{\frac{\pi^2}{a^2} + \frac{\pi^2}{b^2}} \tag{23}$$

Then expressing a in terms of b gives

$$a = \frac{b}{\pi}\sqrt{4\pi^2 f_{c_{11}}^2 \mu\varepsilon - \pi^2} \tag{24}$$

If f_{c11} is set to 3.4GHz, to ensure that TE_{11} is properly suppressed, and the width is greater than the height, then a is bounded by b < a < 22.67b.

To minimise the length of the waveguide, λ_g (the wavelength inside the waveguide) should be minimised by maximising propagation constant β, which means maximising waveguide width a. Restating the list of constraints with dimensions

i. Waveguide is sufficiently wide to hold the cell i.e. a > 0.0285m
ii. Waveguide is sufficiently tall to hold the cell i.e. b > 0.026m
iii. TE10 can propagate at 3.035 GHz i.e. a > 4.94 cm
iv. TE01 cannot propagate at 3.035 GHz i.e. b < 4.41 cm
v. TE11 cannot propagate at 3.035 GHz i.e. b < a < 22.67b
vi. Length of waveguide to be minimised i.e. maximise a

Selecting practical values of width and height, of 10 cm and 3 cm respectively, the calculated cut-off frequencies and wavelength are as follows:

$$f_{c_{10}} = 1.50\ GHz, f_{c_{01}} = 5.00\ GHz, f_{c_{11}} = 5.21\ GHz$$

Consulting Pozar [87] again for the following equations:

$$wavenumber\ k = \omega\sqrt{\mu\varepsilon} = 63.64 \quad (25)$$

$$cutoff\ wavenumber\ k_c = \sqrt{\left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2} = 31.42 \quad (26)$$

$$phase\ constant\ \beta = \sqrt{k^2 - k_c^2} = 55.35\ rad/m \quad (27)$$

$$wavelength\ \lambda_g = \frac{2\pi}{\beta} = 11.35\ cm \quad (28)$$

Assuming a full wavelength between the launcher and the cell to ensure an even field, and a 2-port design with an injector at each end, the waveguide would need to have an internal length of at least 31.4 cm.

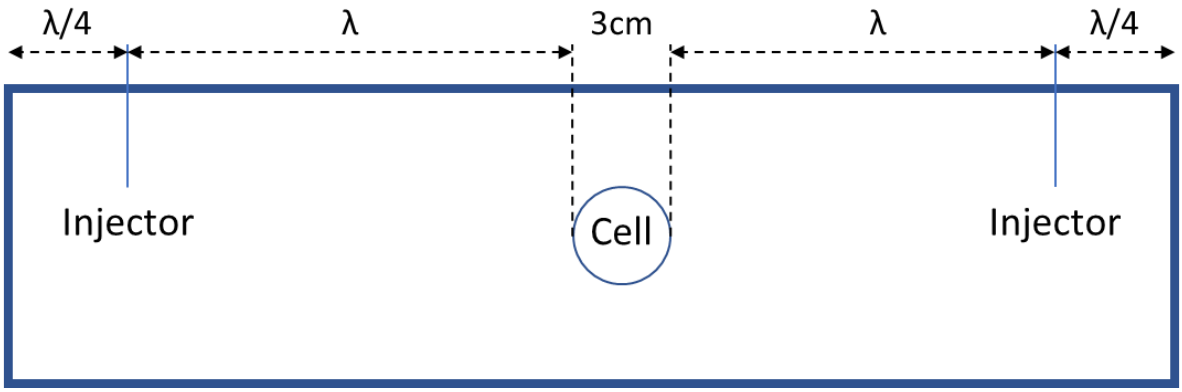

Figure 20. Diagram showing minimum length of the waveguide including cell and injector positions.

As these dimensions were not close to any waveguides or launchers readily available, it was necessary to construct the waveguide. Some aluminium box section was identified as being suitable for operation at 3.035 GHz. A source of 6082T6 aluminium section was found and it was determined that the manufacturer's 4" x 1 ¾" x 10 swg variant provided internal dimensions of 95.2 x 38.1 mm, meeting the 6 constraints identified earlier. The wall thickness of 3.2mm proved rigid and provided sufficient material to tap screws into. Table 7 shows that only TE_{10} can propagate, as TE_{01} and TE_{11} are both well below their cutoff frequency.

Table 7. Cutoff frequencies for 6082T6 box-section as a waveguide.

			width	height	cut-off				
Mode	**m**	**n**	**a**	**b**	**fc/GHz**	**k**	**kc**	**Beta**	**lambda.g**
TE10	1	0	0.0952	0.03805	1.574068	63.64333	32.99992	54.41947	0.115458
TE01	0	1	0.0952	0.03805	3.938273	63.64333	82.56485	#NUM!	#NUM!
TE11	1	1	0.0952	0.03805	4.241189	63.64333	88.91541	#NUM!	#NUM!

CST Studio Suite was used to simulate the performance of the design before commencing build. In the first instance, the waveguide, one injector and its associated end-wall were present. Port 1 was defined as being the coaxial tail attached to the injector, and port 2 was a waveguide port covering the opening at the right in Figure 21.

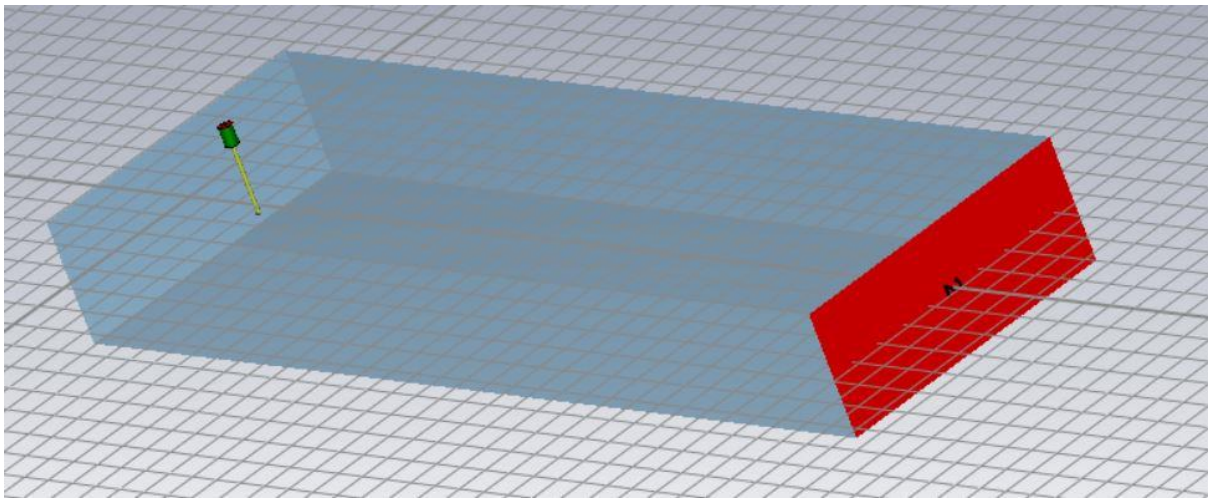

Figure 21. Model of the waveguide with one injector. Walls are transparent. Waveguide port 2 is visible.

Figure 22 shows the s-parameters returned from the time-domain solver. In a real measurement it would only be possible to measure reflection but s_{21} can be established in a simulator by considering the energy radiating into a TE_{10} modal termination.

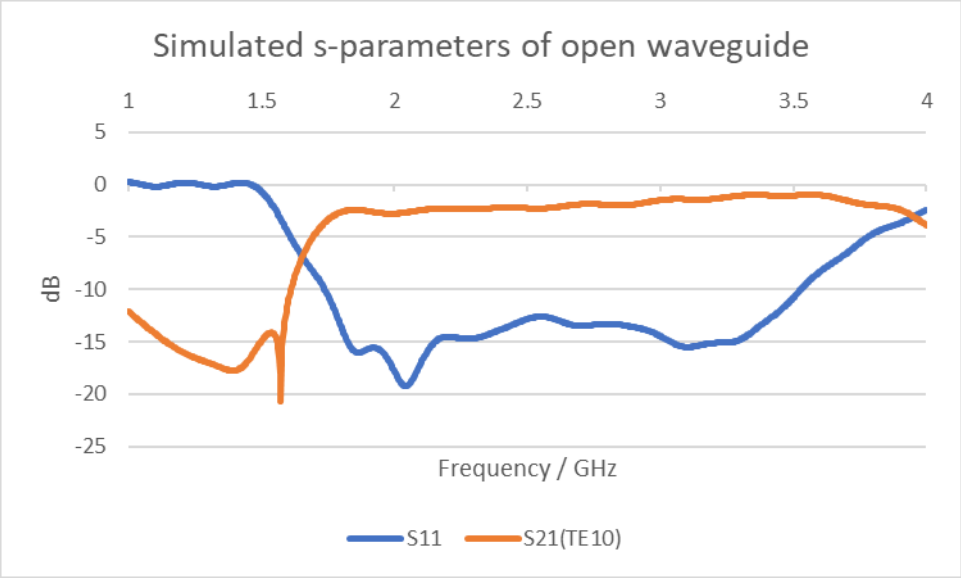

Figure 22. s_{11} and s_{21} coefficients of open waveguide.

Table 8 shows the dimensions used to build the parametric model, and also the injector dimensions following optimisation (maximising s_{21} and minimising s_{11} at 3.035 GHz). The antenna performs slightly better if its length, and the distance between the antenna and the endwall, is somewhat less than $\lambda_g/4$ (28.9 mm).

Table 8. Critical dimensions following optimisation of the injector.

Name	Expression	Value	Description
wg_l	= 326.69058411513	326.69058411513	internal length of waveguide
wg_w	= 95.2	95.2	internal width of waveguide
wg_h	= 38.1	38.1	internal height of waveguide
wg_wall	= 3	3	waveguide wall thickness
ant_l	= 19.720961987537	19.720961987537	length of antennae
ant_r	= 1.27/2	0.635	radius of antennae
coax_r	= 4.115/2	2.0575	radius of coaxial cable
ant_wall_l	= 23.443044891071	23.443044891071	distance from centre of antenna to endwall

Building on that model, a 2nd injector was added and the open end sealed, as per Figure 23.

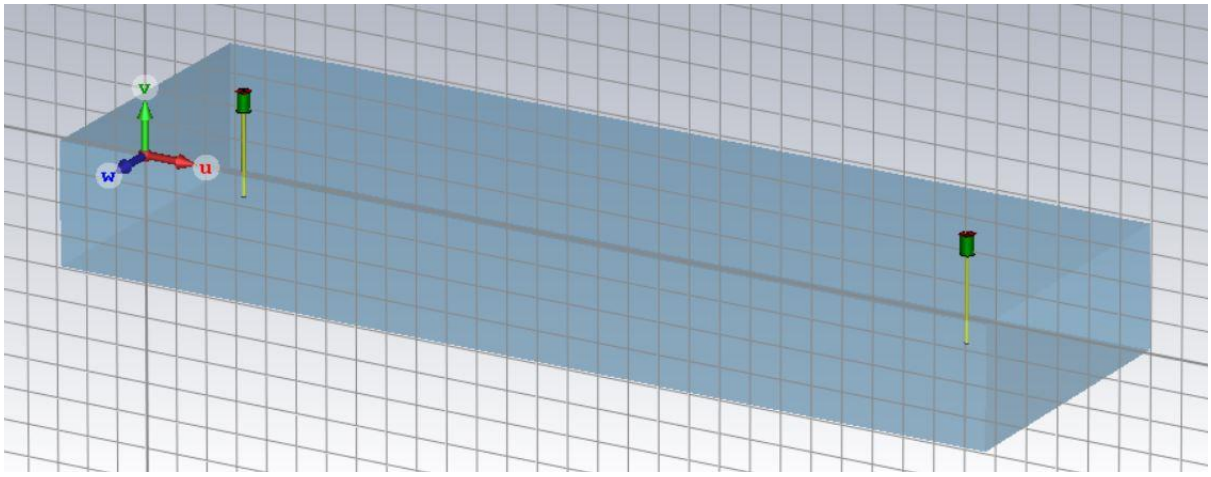

Figure 23. Model of the 2-port sealed waveguide. Injectors and their coaxial tails are shown but aluminium walls are shown transparent.

When the s-parameters are measured the periodic frequency is akin to that of a tuned cavity, especially between 2 GHz and 4 GHz. By again optimising the launcher dimensions and cavity length in the simulator, it is possible to obtain significantly better s_{11} and s_{21} figures than in the open waveguide in Figure 22. The results in Figure 24 were obtained for an internal length of 360mm, with dimensions listed in Table 9. Transmission approaches 0 dB at 3.035GHz, and reflection drops below -30 dB, indicating that TE01 can efficiently transfer power through the waveguide.

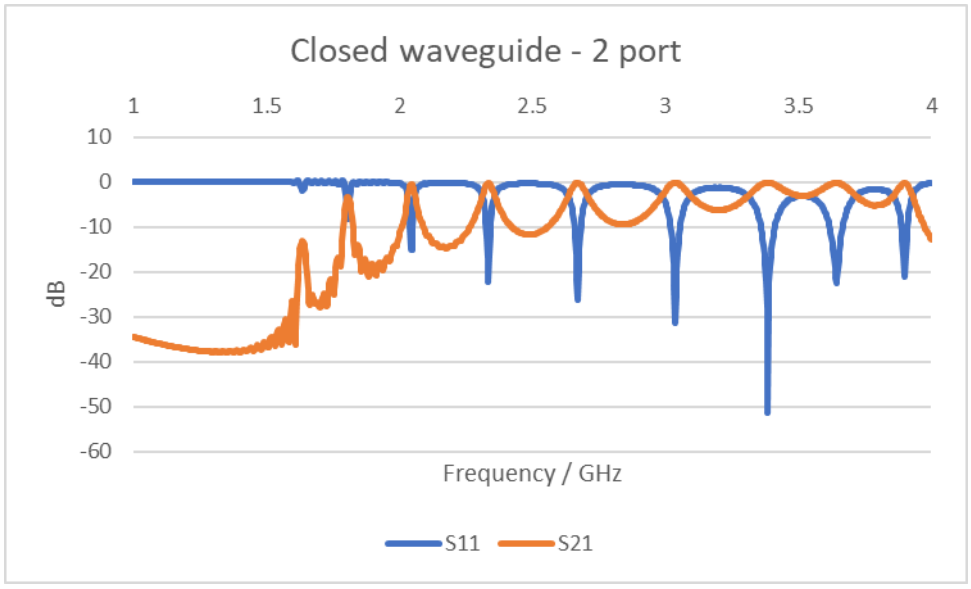

Figure 24. S-parameters for the closed waveguide.

Again, optimal performance is obtained with antenna length slightly shorter than $\lambda_g/4$, but this time the antennas are slightly further away from the end walls than $\lambda_g/4$. A cavity supports resonance when an integer number of wavelengths match the length of an out-and-back path. Therefore it is possible to lengthen the cavity by a multiple of $\lambda_g/2$ (57.8mm), and still see optimal performance at the 3.035 GHz operating point.

Table 9. Critical dimensions following optimisation of the cavity length and injectors.

Name	Expression	Value	Description
wg_l	= 360.02653082698	360.02653082698	internal length of waveguide
wg_w	= 95.2	95.2	internal width of waveguide
wg_h	= 38.1	38.1	internal height of waveguide
wg_wall	= 3	3	waveguide wall thickness
ant_l_tx	= 26	26	length of tx antenna
ant_r	= 0.635	0.635	radius of antennae
coax_r	= 2.0575	2.0575	radius of coaxial cable
ant_wall_l_tx	= 32.8	32.8	distance from centre of tx antenna to endwall
ant_l_rx	= 26.77	26.77	length of rx antenna
ant_wall_l_rx	= 32.8	32.8	distance from centre of rx antenna to endwall

The s_{11} and s_{21} responses for 6 resonant lengths are as per Figure 25 and Figure 26:

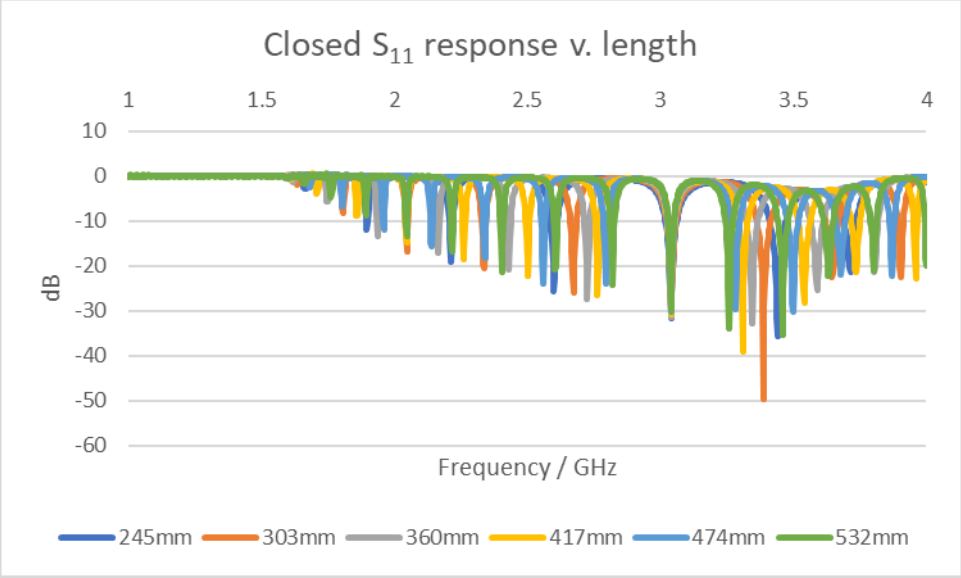

Figure 25. s_{11} responses overlaid for different lengths of closed waveguide.

Figure 26. s_{21} responses overlaid for different lengths of closed waveguide.

For all 6 lengths identified, performance remained optimal at the 3.035 GHz operating point. Shorter lengths were not explored as placing the cell close to the injectors would have introduced

near-field effects. 303mm was selected as shortest length that provided at least λ_g between the injector and cell, as pictured in Figure 20. Plotting the h-field inside the volume of the vapour cell (in the 303mm waveguide) shows there to be good uniformity of field as per Figure 27.

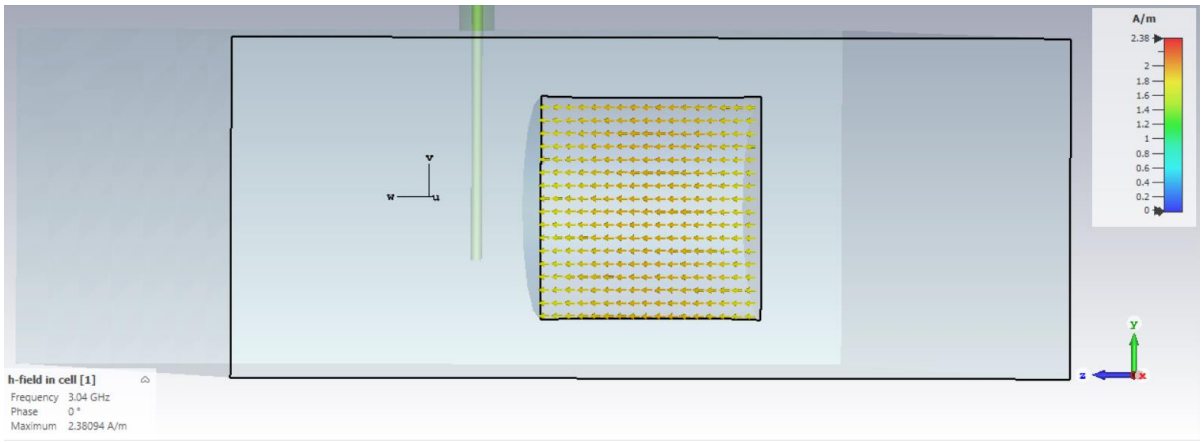

Figure 27. h-field inside the volume of the vapour cell, in the waveguide.

Looking at the s-parameters, as shown in Figure 28, reveals s_{11} is -30.7 dB at resonance, s_{21} is -0.11 dB, and the -3dB bandwidth is 110MHz.

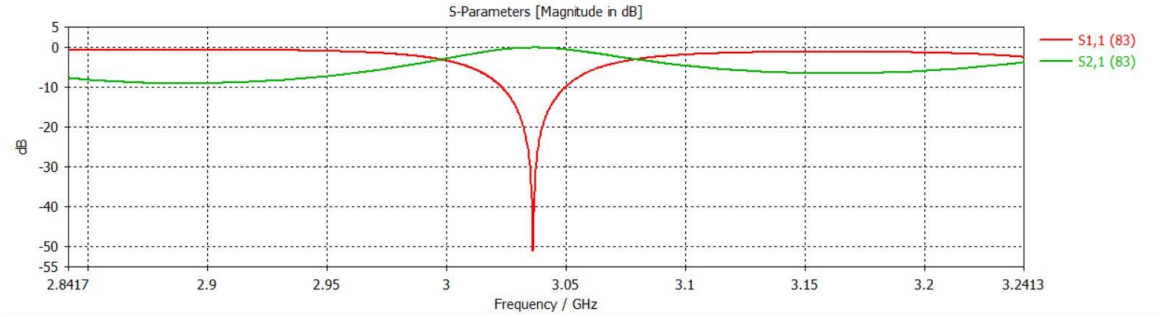

Figure 28. Waveguide s-parameters.

After promising results from a hand-built prototype, a second prototype was locally machined with internal length 329.8mm and neatly fitting end-caps as per Figure 29 and Figure 30.

Figure 29. Locally machined waveguide.

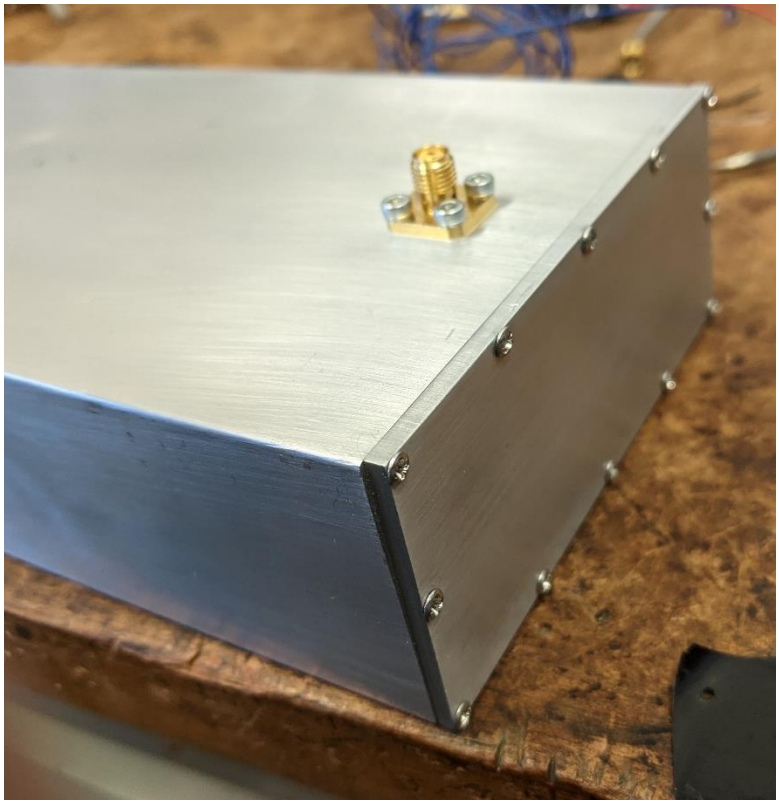

Figure 30. Close view of machined end-cap.

The performance was characterised on a VNA and plotted against simulated behaviour as per

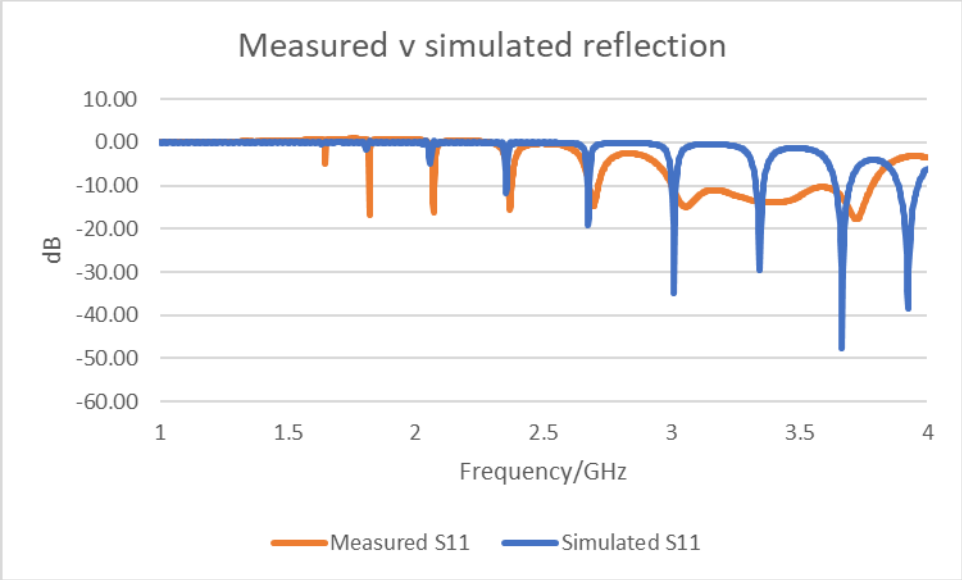

Figure 31 and Figure 32.

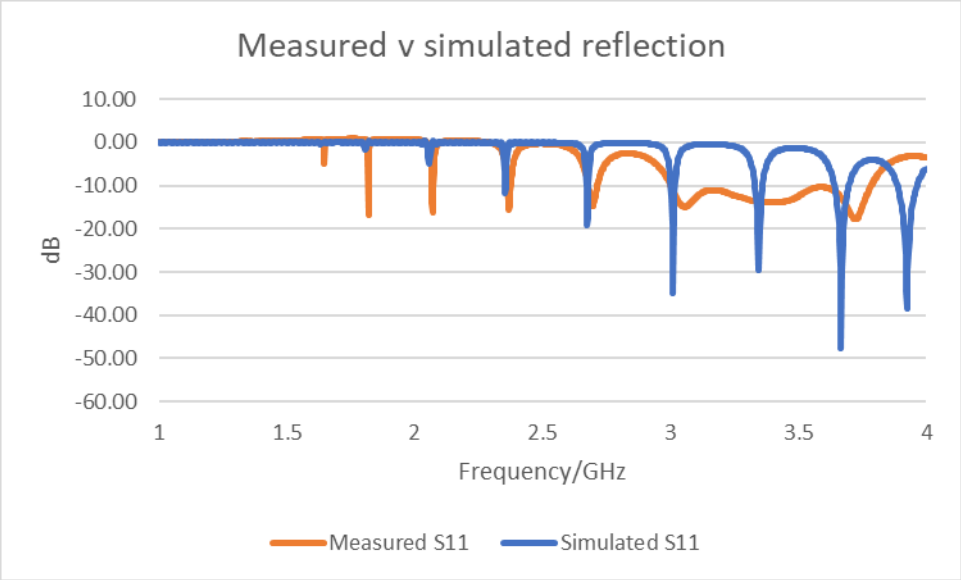

Figure 31. Reflection coefficient of machined waveguide prototype.

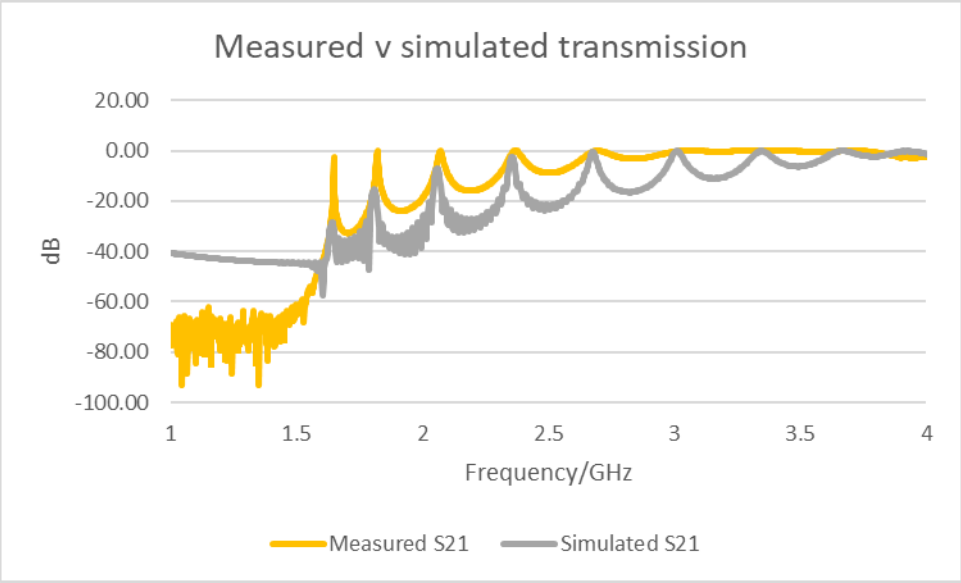

Figure 32. Transmission coefficient of machined waveguide prototype.

The measured reflection coefficient diverges significantly from simulation between 3 and 4 GHz. This is probably due to the unpolished internal surface, the irregular internal shape of a poor-quality extrusion, and exacerbated by the fit of end-caps, and hand-made nature of the antennae. Although the transmission coefficient is slightly lower than simulated at 3.035 GHz, it is also much less frequency dependent and provides a suitable environment for the Rb cell. This ^{85}Rb cell was then fitted inside as shown in Figure 33.

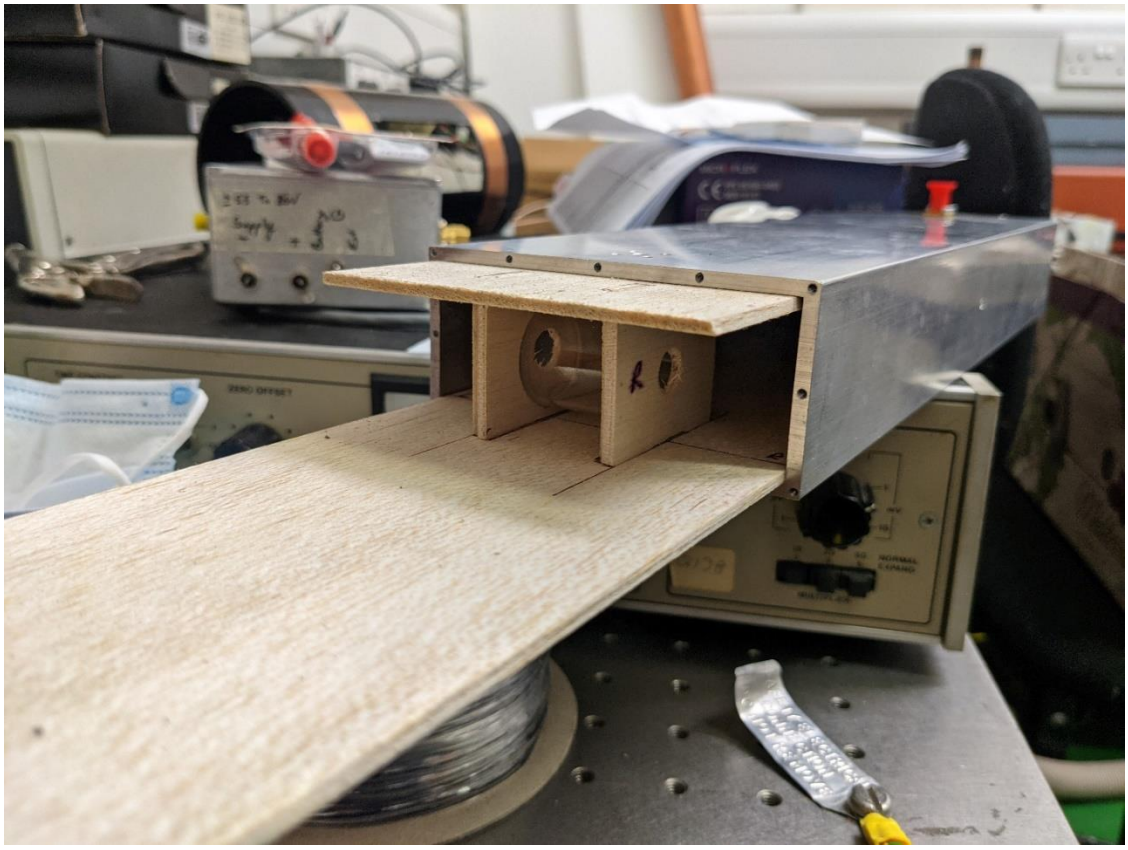

Figure 33. Cell mounted in balsa sled, partially inserted into waveguide.

This assembly was tested again on a VNA, in case the cell or support structure significantly changed the S-parameters. Fortunately, the S_{11} and S_{21} results in Figure 34 and Figure 35 are very similar at 3.035GHz. The reflection and transmission coefficients were measured as -13.52dB and -0.75dB respectively, which remain acceptable. In both figures 'Empty' represents the waveguide without contents, and 'Final' includes Rb cell and support structure.

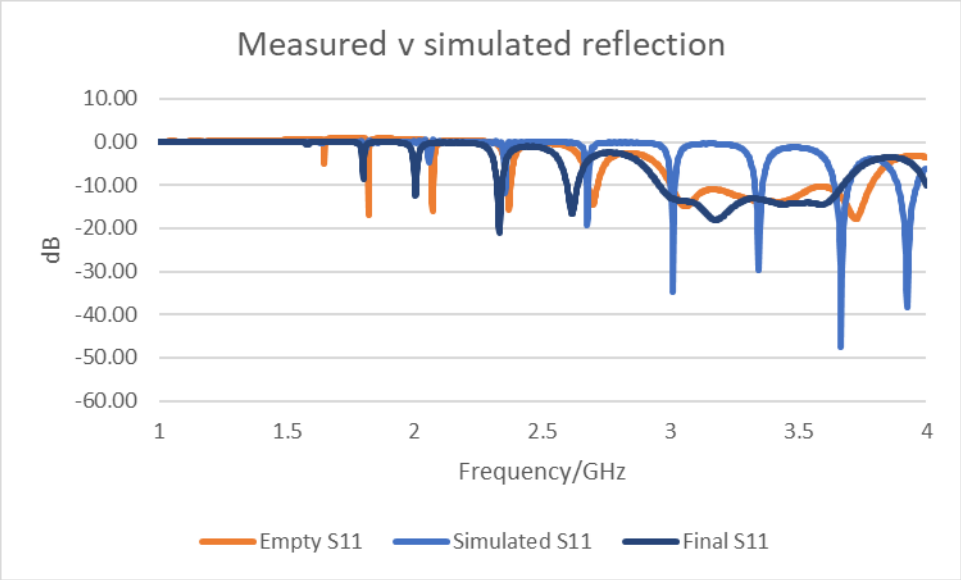

Figure 34. Comparison of simulated reflection v.s. measured waveguide.

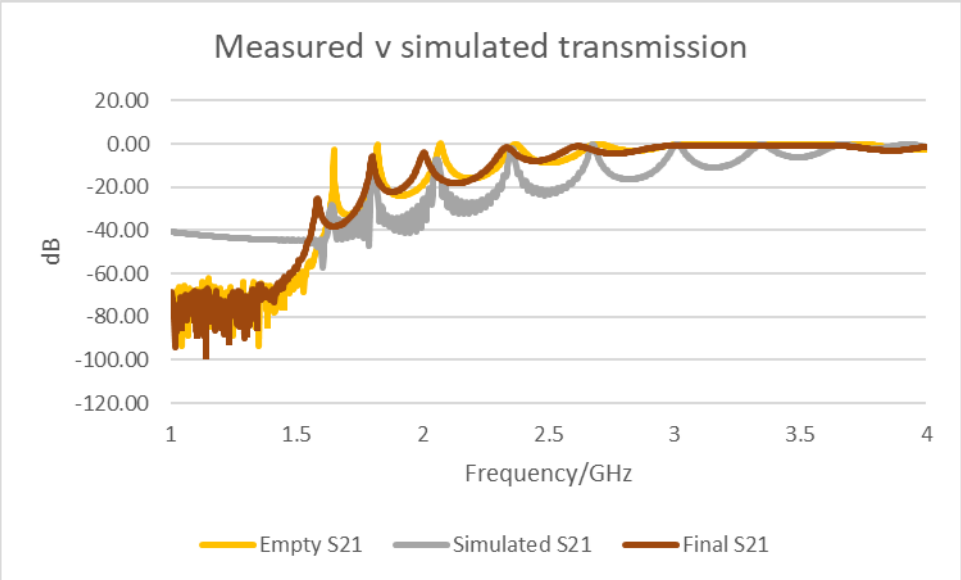

Figure 35. Comparison of simulated transmission v.s. measured waveguide.

3.5 Heating

Atomic spectroscopy can only work when atoms are available for interaction i.e. in a vapour, and so the cell needs to be significantly above ambient temperature. Burtichelov's cell was optimised for operation at 50°C [58, Sec. 5.3.3] and so a means of regulating the temperature of the physics package was required. Initial thoughts were that bi-filar windings wound around the outside of the waveguide could provide a good mechanism for doing this: the magnetic field produced by current passing in each winding is mostly cancelled by that from its neighbour, and heating is distributed across the length of the coil.

Initial experimentation with bi-filar windings was conducted using an oxide-coated Nickel-Chromium heating wire that provides a resistance of 1.73 Ω/m (Block RD100/0,6). A 10m coil was

wrapped around the waveguide and although it was able to provide almost 40W of heat with 1.5A flowing, the waveguide reached a steady state temperature of 39°C, significantly short of the requisite operating temperature. Although it would have been possible to increase the current and add extra coils, concerns around generating a magnetic field so close to the cell prompted a move to a water-heated solution. Four aluminium heat exchangers were clamped to the top and bottom of the waveguide and were fed in series from a Grant W6 1.5kW circulating water bath.

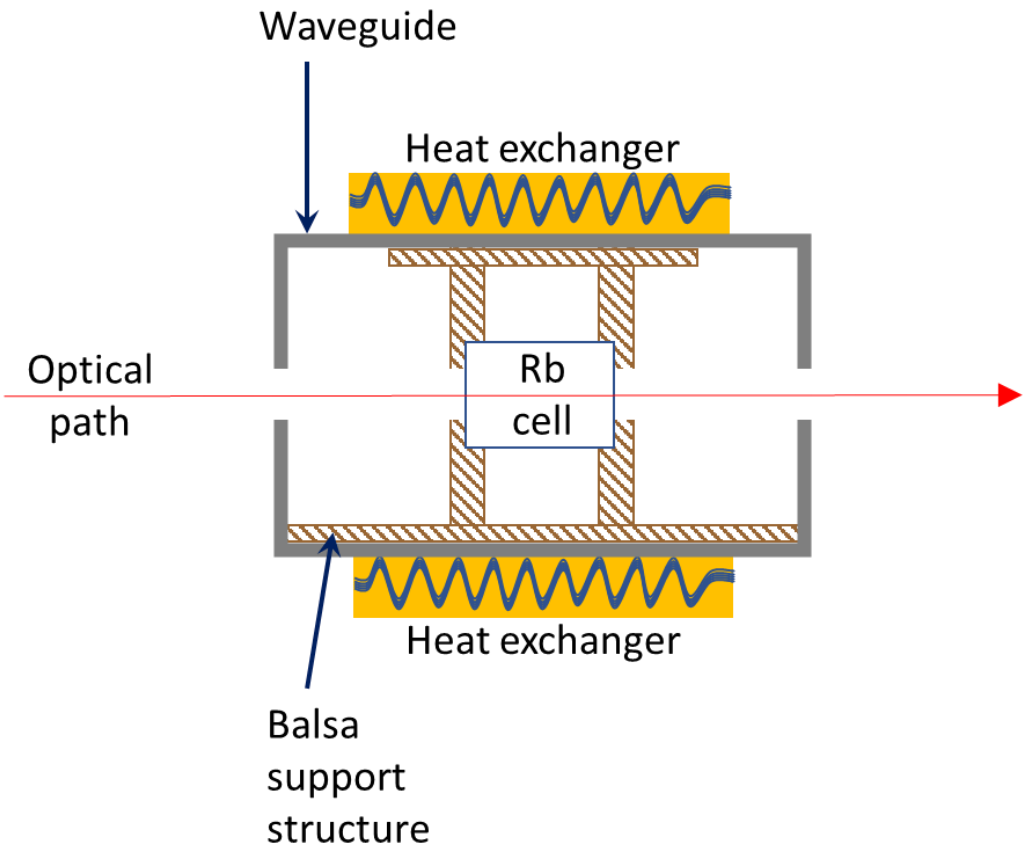

Figure 36. Diagram of the Rb Physics Package showing the position of the cell, heat exchangers and optical path.

PVC tubing was used to connect the components and a thermocouple was used to measure the temperature at the cell. Figure 37 maps the relationship between cell temperature and water temperature for an ambient lab temperature of 21.4°C.

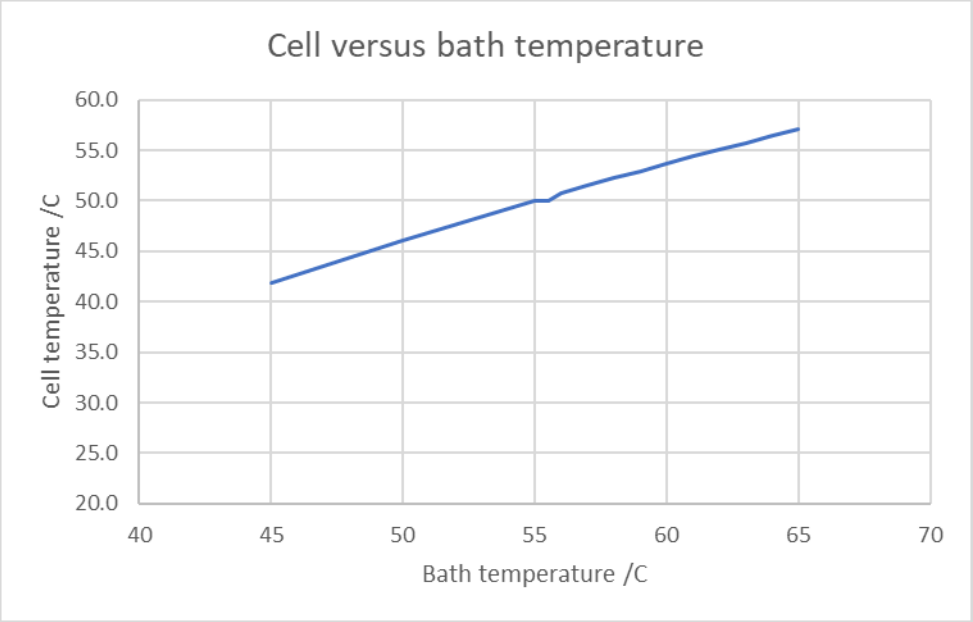

Figure 37. Measured steady-state cell temperature for varying bath temperature.

A 5.5°C temperature difference was measured close to the required cell temperature, then confirmed that steady state bath temperature of 55.5°C maintained the cell at 50.0°C. It was further noted that the PVC tubing becomes extremely soft if the water temperature is raised significantly over 60°C, greatly increasing the possibility of a leak. Aluminium's relatively high thermal conductivity of 237 W m^{-1} K^{-1}, and the four heat exchangers distributed around the waveguide, meant that there was no measurable thermal gradient across the surface of the waveguide, for measurements within 5 cm of the cell with a thermometer with 100mK resolution.

This solution regulates the temperature in the water tank, relying on the large thermal mass and steady state conditions in the lab (also within the safety enclosure) to ensure the cell reaches the required temperature. It would be possible to add a thermostat much closer to the physics package but appeared to work sufficiently well to proceed. As health precautions, a lid was fitted on the water bath, the water was periodically refreshed, and a biocide (Fernox F7) was used to guard against Legionella or any other organisms that might otherwise thrive in warm water.

3.6 C-field

One further addition was made to the physics package before installing it onto the optical bench. Helmholtz coils were installed as per Figure 38, where green PVC tubing connecting the 4 heat exchangers is visible. The Helmholtz coils are used to provide the necessary C-field for the atomic quantization axis and Zeeman splitting, and are aligned with the laser beam. As there is no magnetic screening, the C-field is imposed on top of any ambient magnetic field.

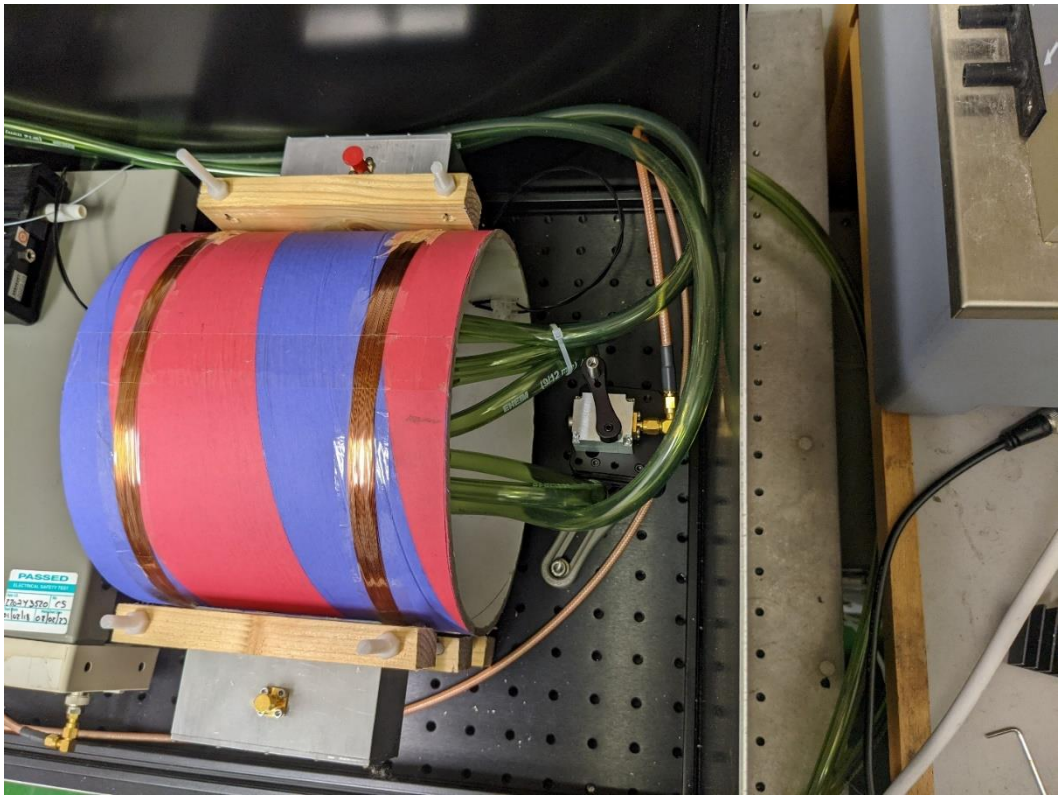

Figure 38. Conclusion of optical path; physics package and photodiode..

The field strength at the centre of the coils' axis is calculated as:

$$B = \frac{8\sqrt{5}}{25}\frac{\mu N I}{a} \tag{29}$$

Each coil is 70 turns of 0.25mm ⌀ enamelled copper wire of radius 0.107m, which can sustain a current of 500mA and therefore produce a peak field strength of 294µT (2.94G). The coils were initially driven from a lab power supply in current-limited mode and correct connectivity verified using a Gaussmeter with transverse probe. The power supply was then replaced with an Analogue DC1531A current source development board to better regulate the current.

The unusual construction seen in Figure 38 is partly due to thermal considerations: a wooden structure is used to rigidly support without conducting heat away from the waveguide. A very rigid clay-composite former has been used for the Helmholtz coils for the same reason. Interference with the magnetic field is the other key concern. As far as possible steel has been avoided near the physics package e.g. nylon bolts used to secure the WG supports.

3.7 MW switch

A Renesas 'IDT F2913' absorptive switch was selected as it offered sensible insertion loss and power-handling capabilities, and was available on an evaluation board. The 'absorptive' nature of the device refers to poles being terminated to ground when not connected, which helps provide > 60dB isolation

at 3GHz. The waveguide was connected to a MW source (initially an HP signal generator) via the absorptive switch, then to a spectrum analyser.

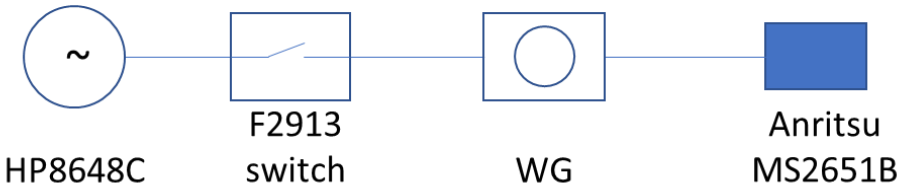

Figure 39. MW connectivity

The switch was characterised in-circuit to ensure that it would perform correctly. In Figure 39 the blue line is received power versus the input, showing a 6.7dB loss in the cables, switch and WG. The orange line records the consistent 0.9dB loss in the switch, accounting for the cable and WG loss. Although ideally conducted at the 3.035GHz operating frequency, these measurements were conducted at 2.9GHz due to a limitation with the spectrum analyser. It is also worth noting that the switch has a maximum CW power handling capability of 23 dBm but the characterisation was capped at 20dBm to avoid testing to destruction.

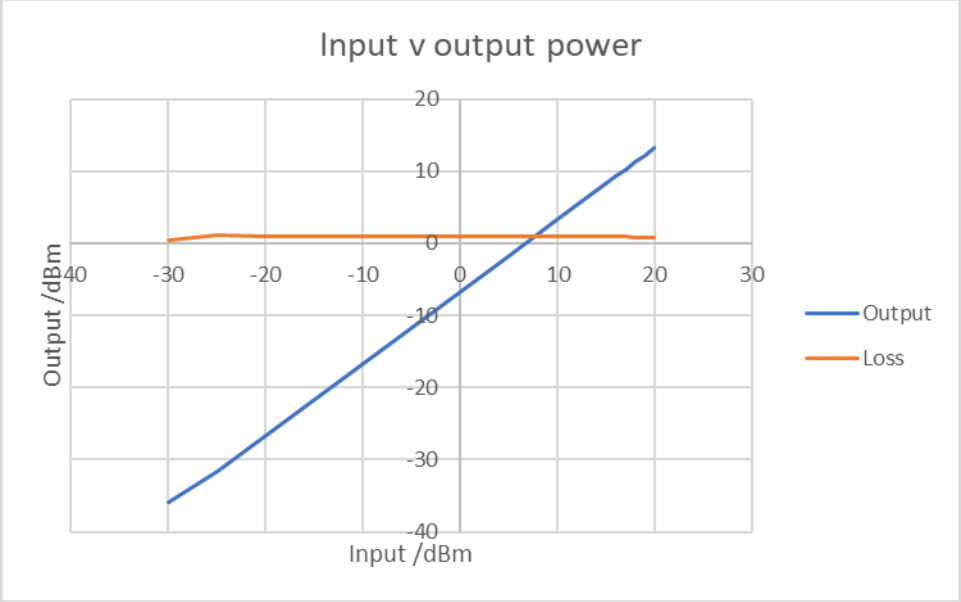

Figure 40. Input vs output power at 2.9GHz

The switch was measured as providing 71.7dB of isolation on a 20dBm input signal. Lastly, harmonics were measured using a 20dBm 1GHz input signal. The fundamental was measured at 16.95dBm, and 2nd harmonic at -15.14dBm i.e. 32dB down.

3.8 Laser

A stable, controllable source of light at the correct wavelength is needed to pump and probe the atomic population for POP. After comparison and analysis, a mixed-vendor solution was chosen as it provided 1MHz linewidth, adequate power for experimentation (without tipping into Class 4), built-in isolator,

a 10µA driver step size (translating to a 9.8MHz frequency step), and was within the available budget. The purchased system was composed of a Thorlabs DBR780PN Distributed Bragg Reflector (DBR) laser diode; Stanford Research Systems LDS-501 250 mA laser driver and O500C3 and O500C4 cables; a Newport LDM-4984 LD mount; and Thorlabs F240APC-780 fibre collimation lens. As this is not a turnkey solution it required construction, test and characterisation as covered Section 5.2.

As is typical of laser diodes, output wavelength can be tuned by varying either the temperature or drive current. The laser diode's sensitivity to current ($\Delta\lambda/\Delta I$) is specified as 0.002 nm/mA whereas sensitivity to temperature ($\Delta\lambda/\Delta T$) is 0.06 nm/°C. These can be restated as $\Delta f/\Delta I$ = 986 MHz/mA and $\Delta f/\Delta T$ = 29.6 GHz/°C respectively which makes it clear that laser frequency is vastly more sensitive to temperature than current, and since the laser controller current source has a bandwidth of ~1MHz (versus a temperature response of roughly 1.6 °C/s), it makes most sense to stabilise the temperature whilst adjusting current.

3.9 Photodetection

Once the optical path has passed through the cell it strikes a photodiode whose electrical output can be amplified and used to determine how much light has been absorbed. In the first instance an Osram BPX61 large-area photodiode was connected to a SR570 transimpedance amplifier, providing an optical response of ~1.06 V/µW at 780nm, which was adequate for early CW measurements including basic spectroscopy and double resonance.

This was then upgraded to a bespoke amplifier that incorporated the BPX61 PIN photodiode into a small package, providing a solution with only 17.0 µV noise into 50kHz B/W. The design, simulation and build of this device is contained in Appendix C, but that solution was prone to oscillation and subsequently replaced with a Thorlabs amplified photodiode. Figure 41 shows the final detector used: an unpackaged design whose purchase was driven by post-Covid supply issues.

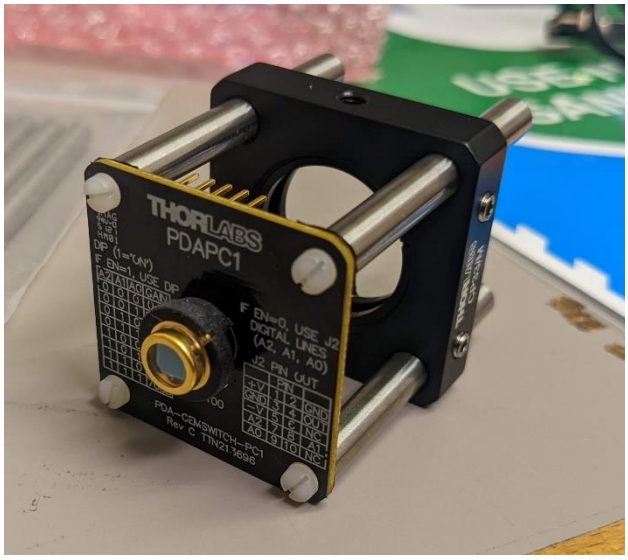

Figure 41, Thorlabs PDAPC1 module mounted on a 30mm cage.

This solution uses its own external, double-isolated power supply, eliminating a previously problematic ground loop, and providing 4.75 x10^5 V/A open-loop gain with 28kHz B/W and only 3.25 x 10^{-12} W/√Hz Noise Equivalent Power (NEP).

This output is then passed into a 'conditioning' block that optimises the output for sampling by an ADC. Although the photodiode output is only sampled during the probe phase, the detector is exposed to both the pump and probe beams, and the signal chain must be tolerant of the full output range. Sampling the output directly with an ADC would mean that less than 20% of the available dynamic range was useful for sampling the probe beam so as an alternative to mechanically blocking the pump beam, the output of the photodiode is passed through a block that electrically limits the response to a pump pulse. Figure 42 shows an OPA699-based voltage limiting circuit which was built on Analog's OPA-SO-1A evaluation platform. The circuit was designed to terminate the Thorlabs amplified photodiode in 70Ω, whilst providing 5x gain, and limiting the output voltage to 2.1 V (which was a limiting factor for one of the ADCs being considered). The output is then buffered through a unity-gain amplifier to drive a low impedance load. This block has independent, regulated power supplies and

has been constructed in a die-cast box for screening. It has also been electrically isolated in a plastic box to avoid creating multiple conductive paths to the ADC.

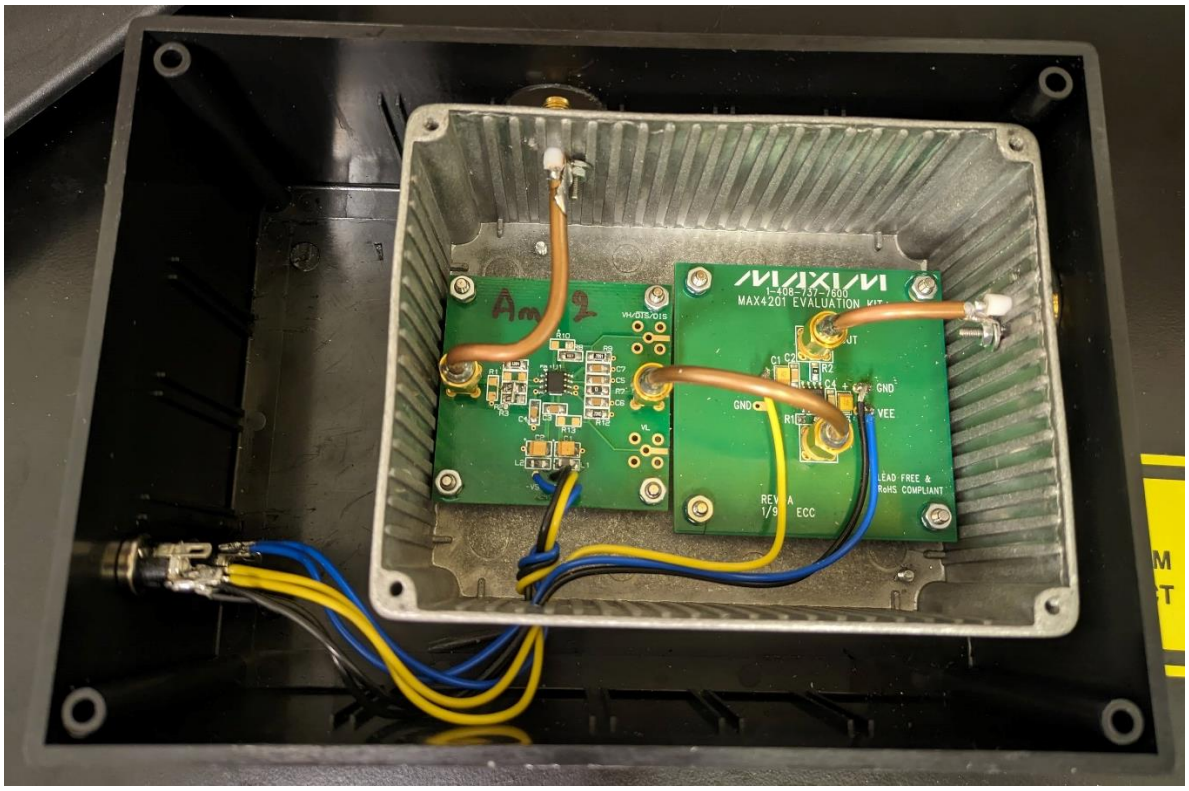

Figure 42. Signal conditioning block.

The circuit was confirmed to work correctly, and output noise was measured as 5.5mV pk-pk (in a 2MHz B/W) when measuring the photodiode dark state: significantly higher than the 236µV that the RMS sum of individual datasheet noise figures, implying that this result could be improved by combining all of the components onto a single PCB with careful layout, but also that alternative schemes could be a potential improvement.

Lastly, this signal is fed to an ADC for conversion to a digital value, as per Figure 43.

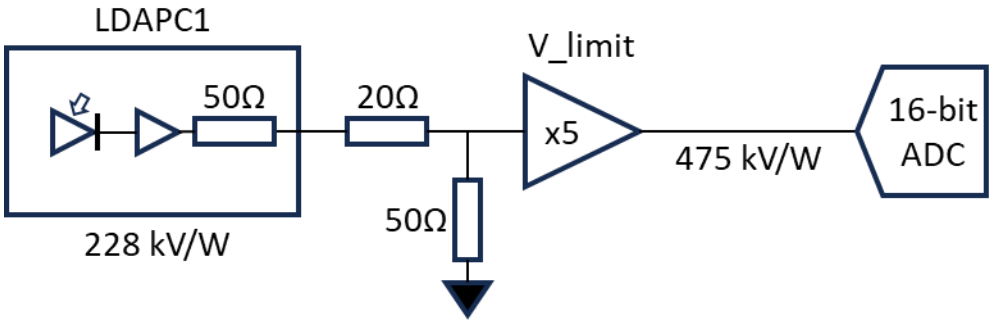

Figure 43. Full optical detection path: amplified photodiode, signal conditioning, and ADC.

It is crucial that the probe pulse is sampled at the correct time. The mechanism for this is covered in detail in the next chapter.

First, an analogue sampling mechanism was employed. A design was built around a Texas Instruments LF398M sample and hold amplifier, and schematics for this can be seen in Appendix D, as it was part of the Logical I/F board. This approach was shelved due to performance limitations and voltage spikes on the output each time the hold was released, but can be seen in Figure 44.

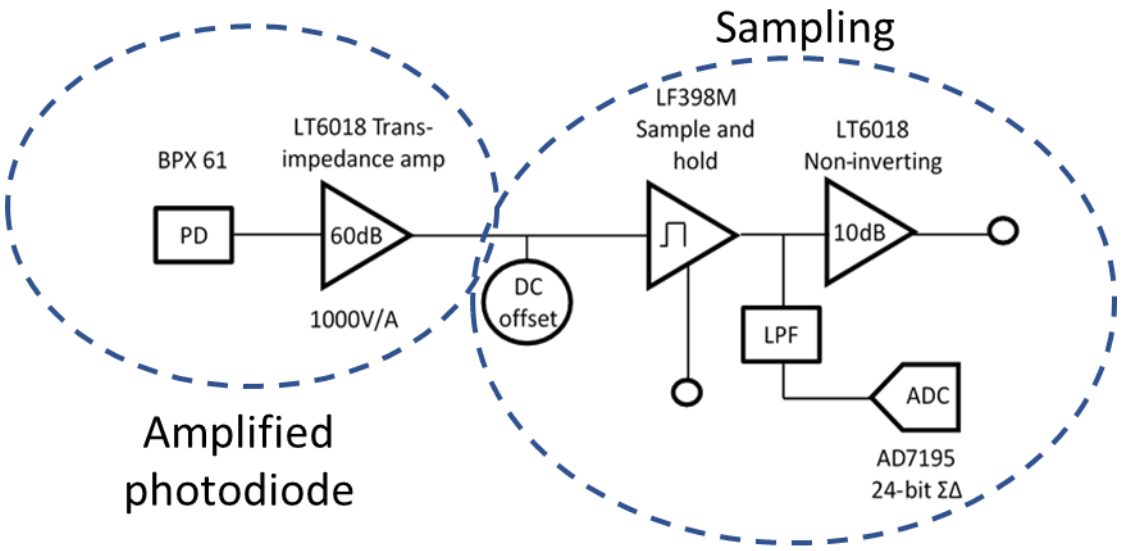

Figure 44. An early implementation of the optical detection path.

3.10 Optical modulation

As POP benefits from different power levels for 'pumping' and 'probing' the Rb population, and requires complete darkness for free precessing, a high bandwidth device capable of modulating optical power is ideal. The piezo-electric effect of an Acousto-Optical Modulator (AOM) allows a VHF electrical signal to create an acoustic wave across a birefringent crystal, altering the index of refraction, which can then deflect a laser beam by a fixed angle [89]. Figure 45 shows the principle of operation; though the incident and diffraction angles have been exaggerated and are typically around 0.5°. As these cells can produce Bragg diffraction, they are sometimes referred to as Bragg cells.

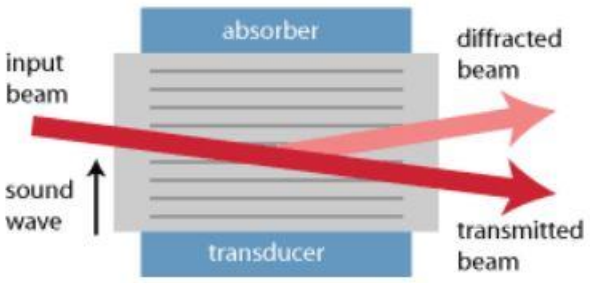

Figure 45. Diagram depicting light diffraction in an Acousto-Optic Modulator (AOM) [90, Fig. 1]

Baryshev et al. have achieved encouraging results using an Isomet 1205-2-804 device and have reported an extinction ratio of 49dB when switching 200mW power in single pass configuration [84]. Light passing through an AOM is modulated by the VHF driver signal and will therefore be ~ 85MHz lower/greater in frequency unless a double-pass arrangement is employed, such as that demonstrated by Schwenger and Higbie [91]. In practice, an 85MHz change in frequency alters the wavelength by 0.00002% and is well within a laser diode's electrical tuning range. The quoted 18.3mrad separation angle means that at least 11cm separation is needed between the AOM and a 2mm aperture to completely discriminate between the incident and deflected beam.

Both laser output power and frequency are functions of current, so both will be affected if using current to tune one parameter. Micalizio et al. [32] found that laser powers of 15mW and 1mW were optimum for the preparation and detection phases respectively. Distributed Feedback (DFB) and DBR laser diodes can readily operate over these wide conditions, but linewidth and SideMode Suppression Ratio (SMSR) are typically optimal only when approaching their full rated power. AOMs can be used as variable attenuators by modulating the VHF power, which will change the percentage of light deflected. The laser frequency can therefore be fixed by the diode current, with optical attenuation used to achieve the desired power level without needing physically switched paths.

Figure 46 shows the optical path used for initial experiments with the AOM, including provision for a Quarter-Wave Plate (QWP), not fitted, but available for experimentation. The laser diode is mated to a polarisation-maintaining fibre represented by the wiggly blue line between DFB and Collimator.

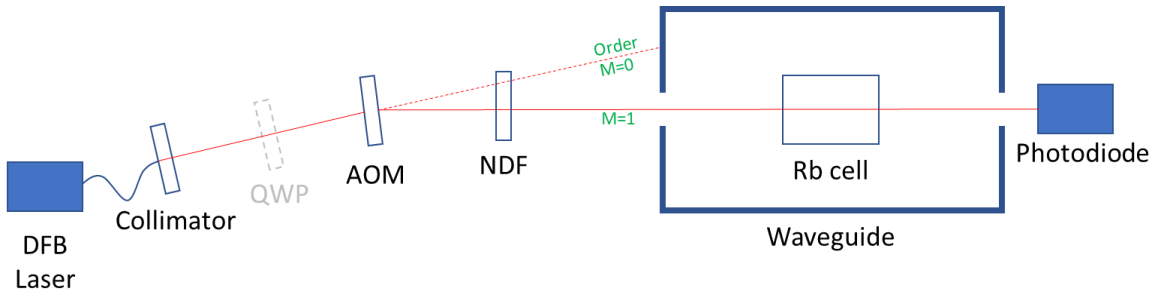

Figure 46. Optical path for POP experiment

The aperture of the waveguide is aligned with the M=1 order beam which means no (or very minimal) light incident on the cell when the AOM is not being driven by a VHF source. Increasing the VHF power deflects increasing levels of light through the cell, up to a maximum of around 50% of the incident beam. Some power will always be present in the M=0 order beam, and some will be deflected into the higher order beams e.g. M=2 and M=3.

The AOM was paired with an Isomet 630C-80 driver which comprises a self-contained oscillator, and provides external inputs for VHF frequency tuning and modulation depth. It was then characterised by measuring the percentage of optical power that could be deflected along the M=1

order beam. A 35mW incident beam was passed through the AOM, and the output measured at a point downstream of the M=1 order beam (and only that beam). Figure 47 shows some non-linear behaviour at very low VHF power levels, requiring extra care during setup of the probe beam.

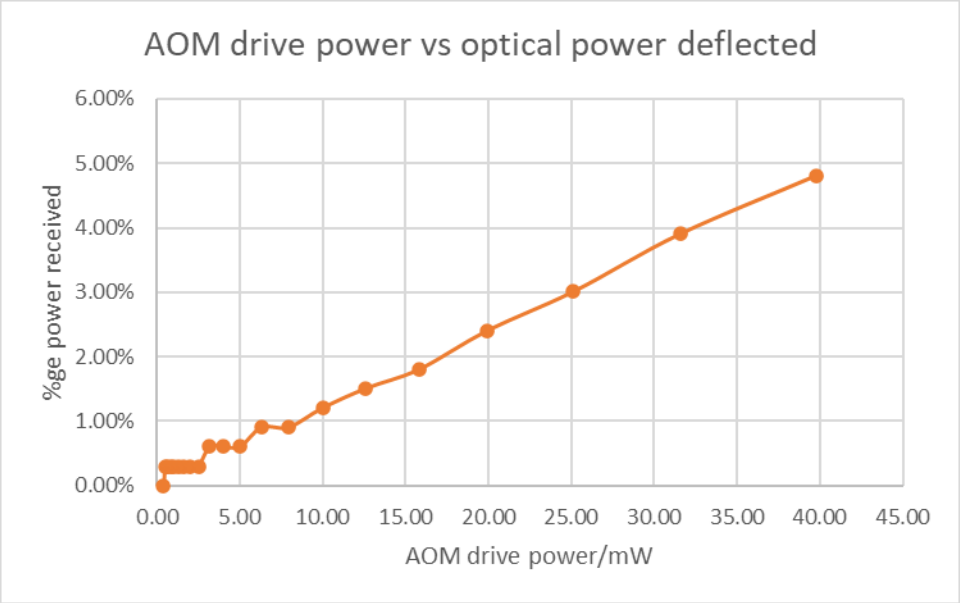

Figure 47. AOM drive power vs optical power deflected.

3.11 Beam expander

The addition of a beam expander completes the diagram shown in Figure 17. By increasing the diameter of the beam that passes through the cell, photons can interact with a greater number of Rb atoms thereby improving the effective signal to noise ratio. A number of solutions were considered to minimise the impact on the uniform field inside the waveguide, as shown in Figure 48. From the top:

- The simplest solution: beam expansion outside the waveguide, but which would require enlarged apertures in the waveguide with associated MW leakage and non-uniformity of field.
- An improved solution: beam expansion outside the waveguide with elongated ports. Simulation of these elongations reduced the leakage by ~23dB per length/diameter of the port and would have been a viable solution for a 10mm aperture (since the aperture is much smaller than λ_g = 11.35cm).
- A Galilean-configuration expander inside the waveguide. This would allow the existing 3mm apertures to be unchanged but require 3 lenses inside the waveguide. Beam magnification of x3 would be possible.

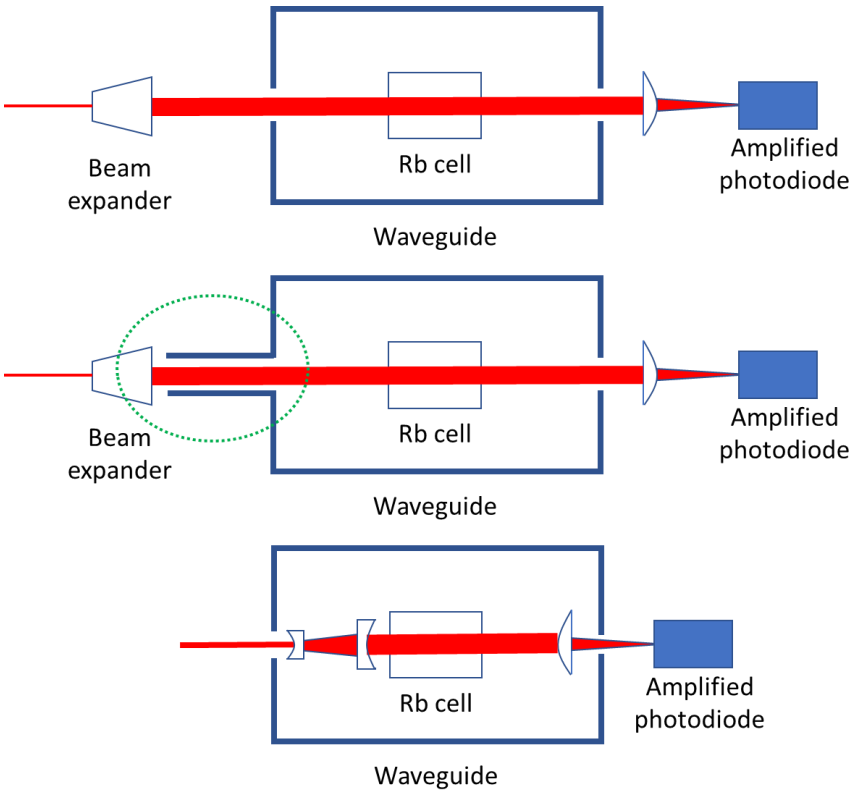

Figure 48. Possible beam expansion solutions. From top: Beam expander outside waveguide; beam expander with port elongation; Galilean-configuration expander inside the waveguide.

The Galilean configuration could be made more compact by incorporating the lenses into the cell structure and embedding the concave lens into the waveguide, however, the chosen solution was to place a beam expander partially inside the waveguide, as shown in Figure 49. As with the Galilean expander, the existing 3mm apertures were retained whilst increasing beam diameter through the cell.

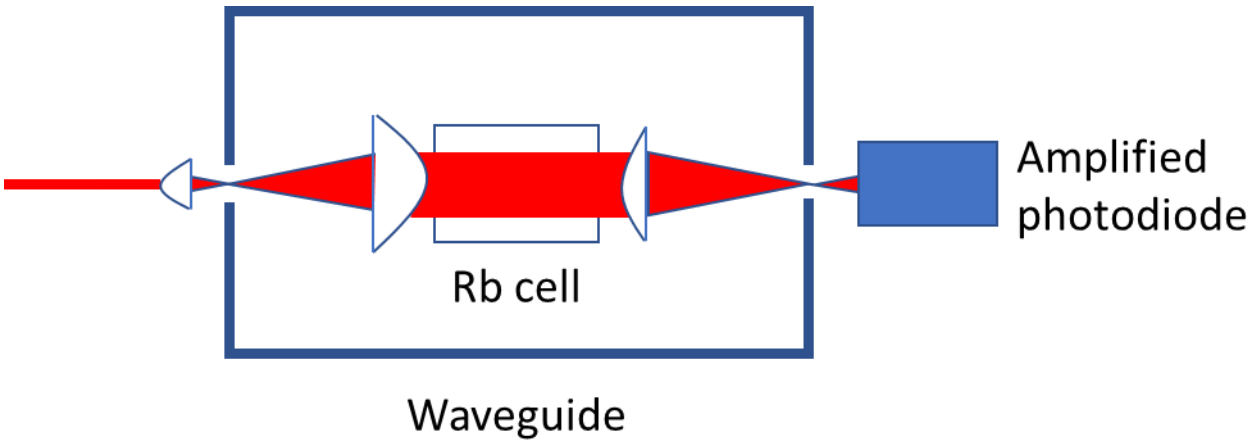

Figure 49. A Keplerian-configuration beam expander partially inside the waveguide.

This required two lenses inside the waveguide, whilst the objective lens sits outside the structure. Beam expansion is limited by the length available in the waveguide but x4.17 was achieved by using an objective lens with focal length of -6mm, and a plano-convex image lens with 25mm focal length.

This means that the beam diameter is expanded to 8.3mm, and that beam area expands from 3.1mm^2 to 54.5mm^2. Optical power density in the cell is therefore 17.4 times lower than the 2mm ⌀ beam.

Figure 50. ^{85}Rb cell and 2x Thorlabs LA1270-B Plano-Convex lenses mounted in a balsa support structure.

Figure 50 shows the cell mounted in a new balsa structure. The balsa has been laser cut to achieve accurate lens placement, and each piece is a light interference fit with the next, and fits snugly inside the waveguide. The waveguide was reassembled and measuring the S-parameters on a VNA confirmed that the lenses (and extra balsa support) appear to make little difference – at the 3.035 GHz operating point there is 0.8dB more reflection but 0.1dB greater transmission as depicted in Figure 51 and Figure 52.

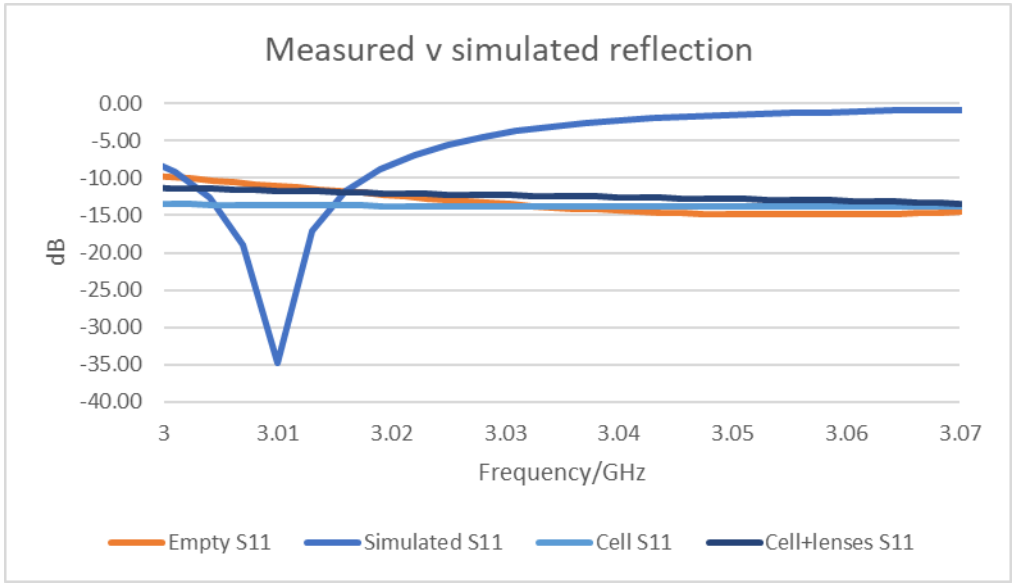

Figure 51. Comparison of waveguide reflection: simulation v.s. measured (empty, cell, cell & beam expander)

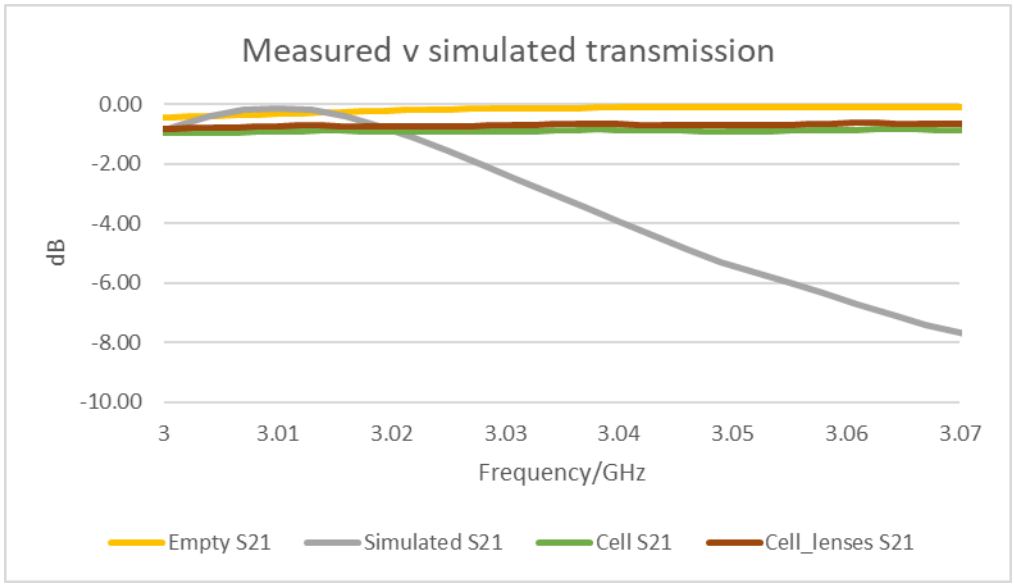

Figure 52. Comparison of waveguide transmission: simulation v.s. measured (empty, cell, cell & beam expander)

4 POP Cycle Timing and the Control Platform

4.1 Overview

This chapter covers the components surrounding the Physics Package, and generation of the Pulsed Optical Pumping (POP) timing signals, as seen in the diagram of the complete clock in Figure 53.

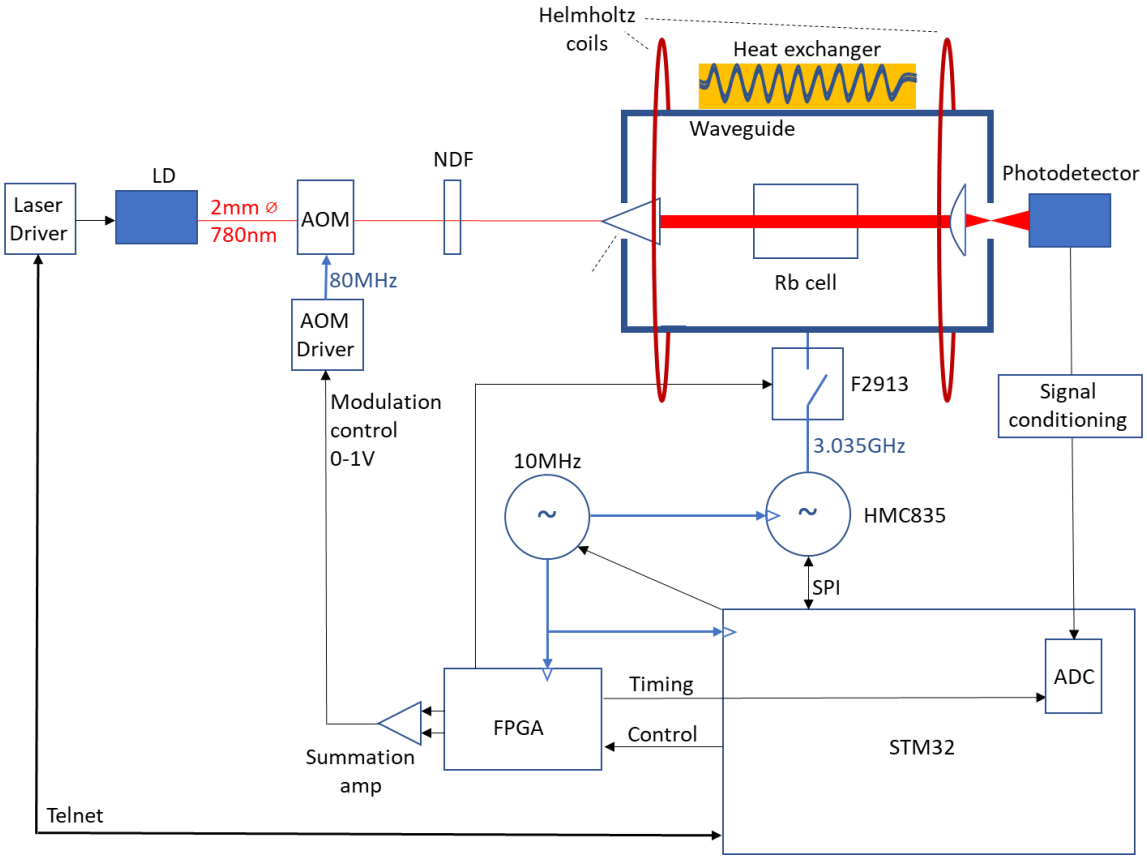

Figure 53. Diagram of the complete clock

This includes an FPGA used to generate the timing-critical POP control signals, a 10MHz OCXO, an MCU development board (labelled STM32), an HMC835 Phase Locked Loop (PLL), and interface circuitry including the summation amplifier in the diagram above.

4.2 POP timing control

An FPGA is responsible for generating all the digital timing-critical control signals needed for POP as seen in Figure 54.

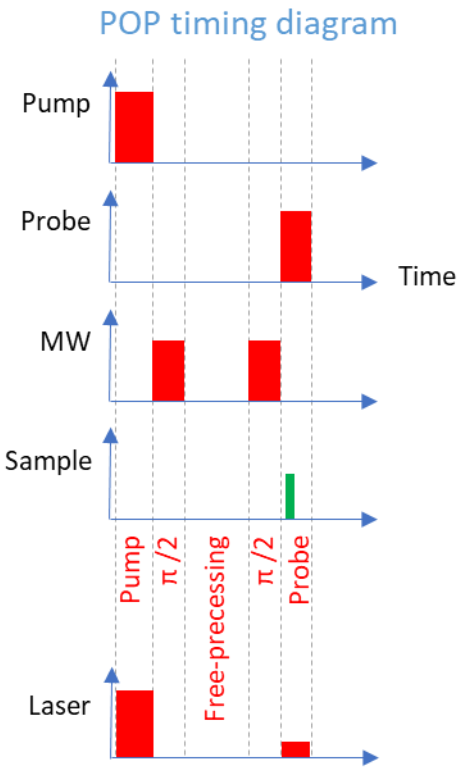

Figure 54. POP timing diagram

This functionality was originally intended to be provided by a SAML22 microcontroller but was unfortunately beyond the capabilities of the onboard Timer/Counter (TC) blocks. A first, and successful, alternative was to use a bank of Arbitrary Waveform Generators (AWG) as seen in Figure 55. The Pump and MW control signals can be seen on the oscilloscope in the foreground.

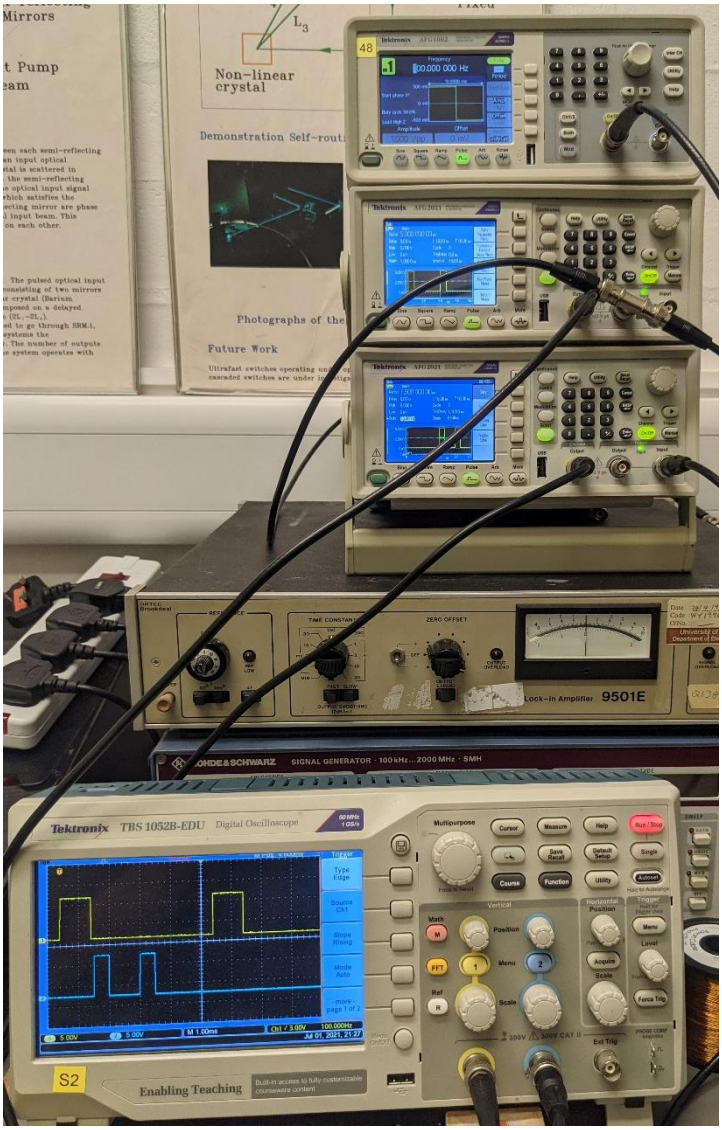

Figure 55. Bank of 2 single channel, plus one dual-channel Arbitrary Waveform Generator (AWG), being used to generate POP timing..

Although correct functionality was achieved (replicating the timing values used by Micalizio and team [74, Fig. 4]), it required synchronising on initial power-up each day, was fiddly to adjust the individual pulse timing settings, and didn't provide a mechanism for switching between laser frequency tuning and POP. The latter was needed at least hourly as the laser frequency drifted.

Given the flexibility and control offered by an FPGA, a small Lattice Semiconductor Mach-X02 development board was purchased, allowing development off-site when lab access was restricted during the Covid-19 pandemic. Figure 56 shows the mechanism used to generate the necessary control signals. The heart of the design is a counter that increments periodically. A bank of comparators each evaluates the counter value with a preset figure, which is used to ensure that each output will turn on and off at the correct time in the cycle. The pump, probe, and sample signals each use two comparators, whilst the MW signal uses 4 as there are two π/2 pulses needed per cycle. One final comparator is used to reset the counter and restart the cycle.

An external 10MHz source is divided (by the internal PLL) to produce a 2.5MHz clock for a 16-bit counter; providing a 400ns tick, for a cycle time of up to 26.2ms.

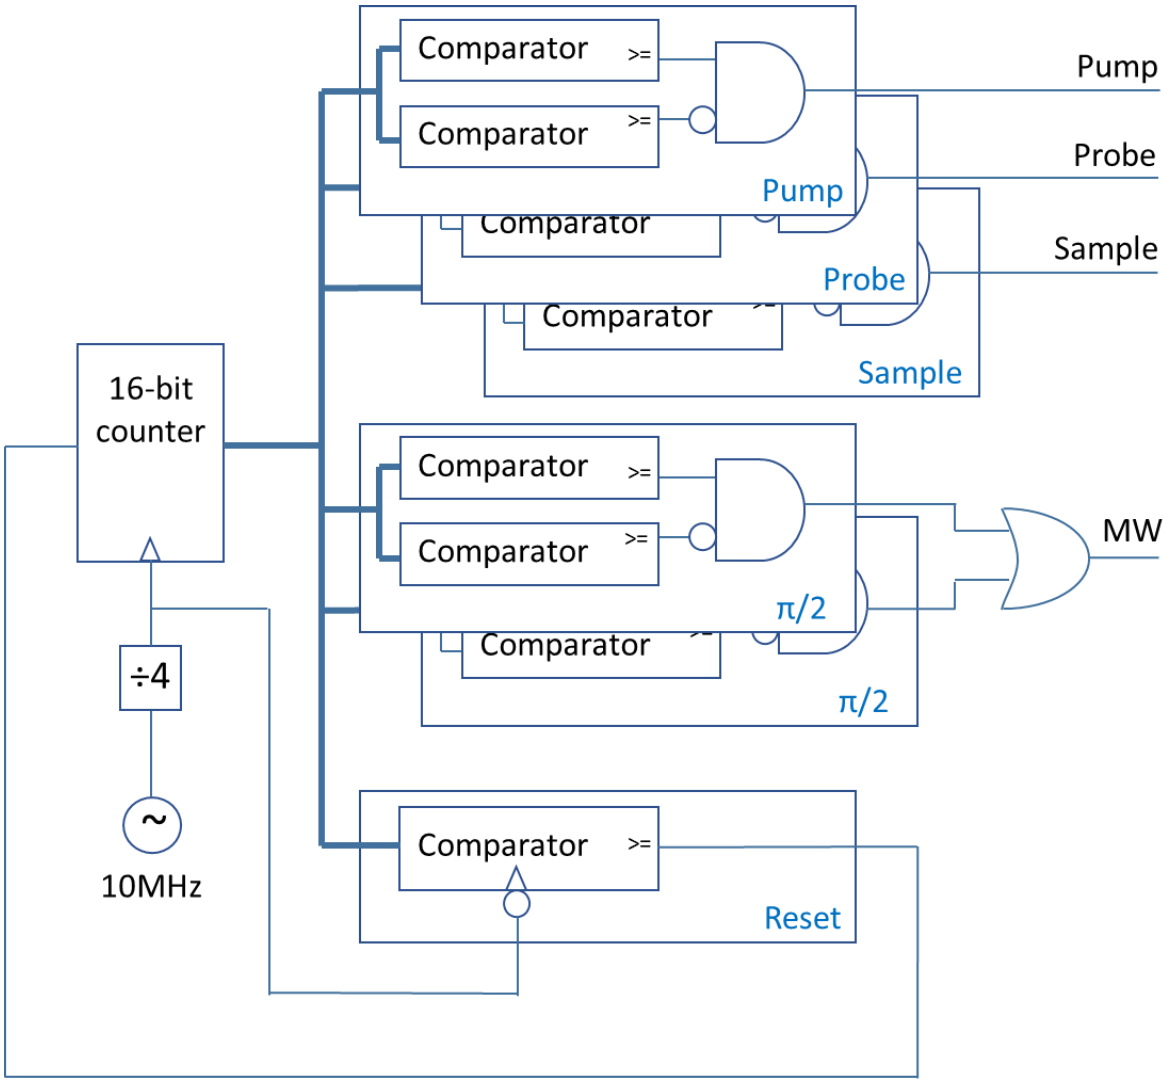

Figure 56. FPGA-based POP cycle timing generation

This timing mechanism is then fed into a state machine with outputs as specified in Table 10. This permits fast changes between laser alignment and POP, and provides other modes useful for experimentation and calibration of the clock.

Table 10. POP timing Field Programmable Gate Array (FPGA) state machine

		Outputs				
Mode	**Description**	**Pump**	**Probe**	**MW**	**Sample**	**LED**
0	MCU control	POP cycle timing / laser frequency setup				Control status
1	Laser frequency setup	0	1	0	1	Slow flash
2	POP cycle	POP cycle timing				On
3	Dark	0	0	0	1	Fast flash
4	Pump beam calibration	1	0	0	0	Off
5	Double resonance	0	1	1	1	4 quick flashes
6	Probe with pulsing MW	0	1	Pulsed	1	3 dots & 1 dash
7	POP with continuous sample	POP cycle timing			1	2 quick flashes

Mode 0 is the power-on default and allows the MCU to select between laser calibration and POP. It additionally allows the MCU to suppress the sample output, e.g. when the MW frequency is changing, and the output of the Physics Package is not valid. 6 pushbuttons were added as per Figure 57, allowing fast adjustment of functionality and timing.

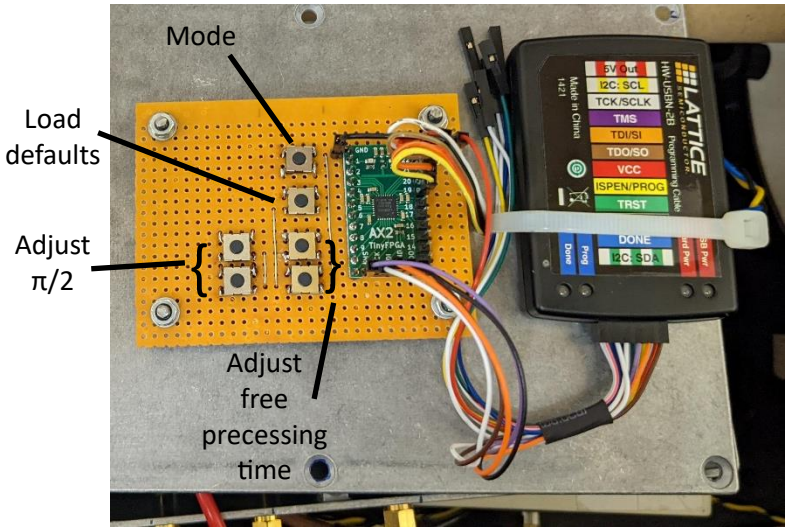

Figure 57. Mach X02 1200HC FPGA and JTAG programmer/debugger.

After debugging, the FPGA's behaviour was verified and the control sequence necessary for one POP cycle is shown in the outputs seen in Figure 58.

Figure 58. FPGA outputs as used to control a single POP cycle.

All of the RTL (Verilog) and testbenches for the FPGA were written from scratch, and are available on the Github online software repository - POP_timing_FPGA.

4.3 Logical Interface (I/F) circuit

The FPGA outputs are all digital so it is necessary to have some additional circuitry to generate the analogue (0-1V) output needed for laser modulation control. It is also desirable to be able to clean up the digital outputs; as the FPGA is not especially low-noise. A current summation amplifier is used to generate the AOM modulation output voltage according to Table 11.

Table 11. Summation amplifier Karnaugh map

Pump	Probe	State	Output
0	0	Dark	-10mV
0	1	Probe	100mV
1	0	Pump	800mV
1	1	Invalid	N/A

The MCU logic outputs are first buffered by TSX0101 level translation devices to act as a logic-dependent voltage sources, providing a clean 5V or 0V output. These sources then feed a 3-stage current summation amplifier as per Figure 59, which allows independent tuning of the 'dark', 'pump', and 'probe' output voltage by adjusting RV1, RV2, and RV3 respectively. Please note that the voltage figures in Table 11 are test values: these are not the final values.

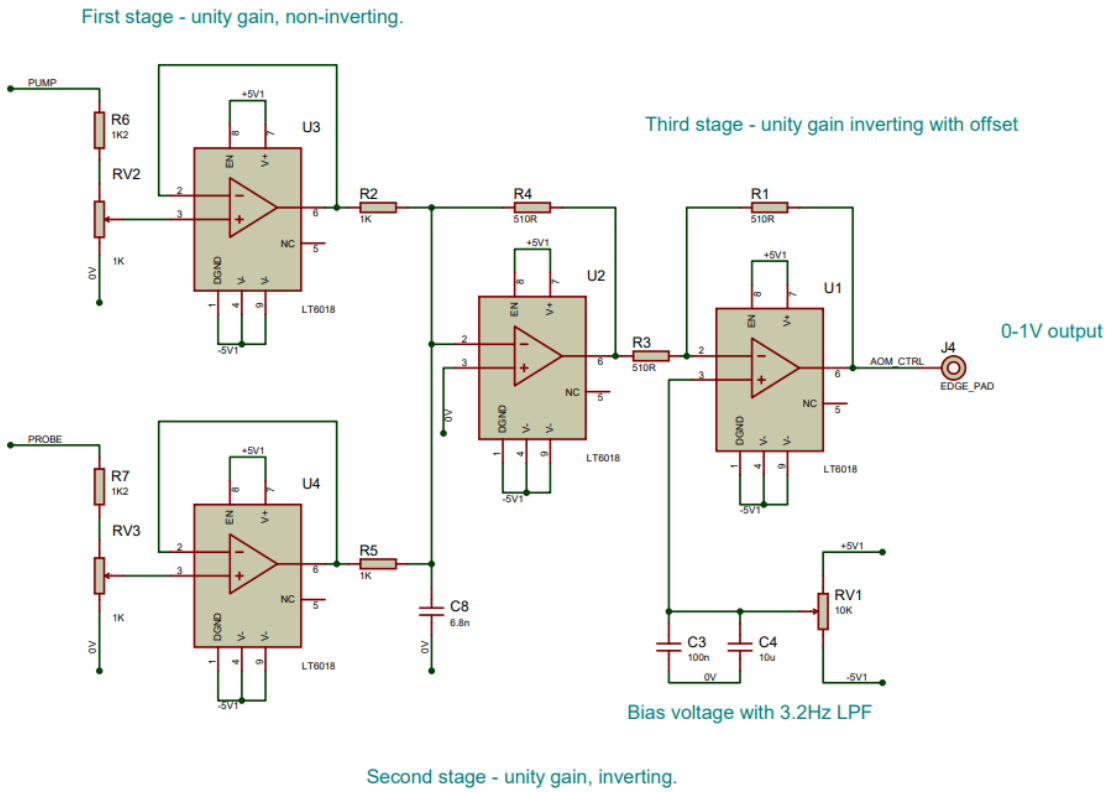

Figure 59. AOM modulation control circuit.

This functionality was combined onto a single PCB along with three well-regulated power supplies; a dual-rail (+/- 5.1V) supply and +3.3V for the FPGA. This interface board was then connected to the FPGA as shown in Figure 60.

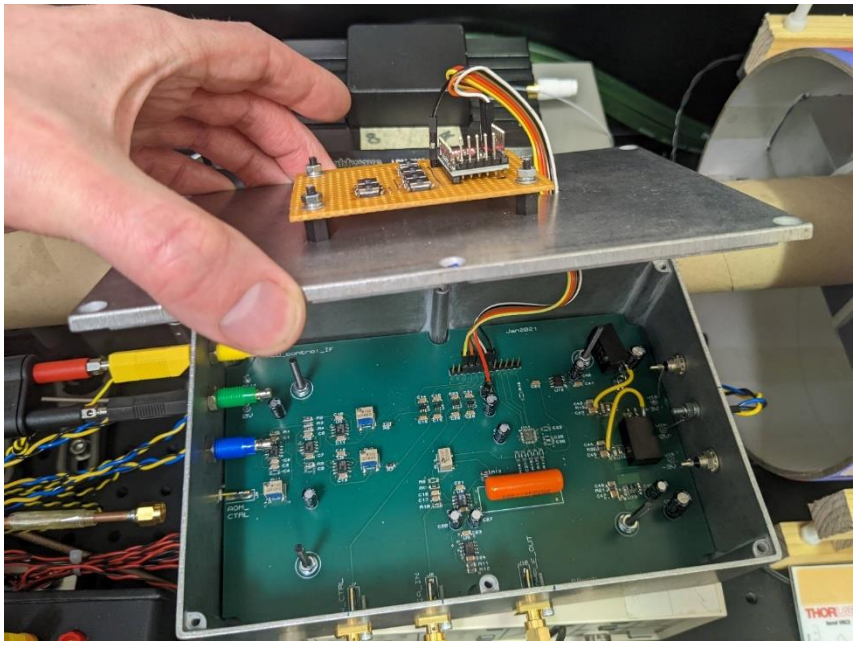

Figure 60. FPGA and interface board in situ.

Figure 61 shows the outputs measured from this assembly during a POP cycle. Cycle time remains 4.5ms but it now has an 8ms pause before repeating.

Figure 61. Captured traces from the outputs of the interface board.

This was reasonable assurance that the circuit was a sensible implementation of the Karnaugh map in Table 11, isolating the output from MCU noise, and with the initial voltage spikes being largely due to the scope leads.

4.4 Control platform (STM32)

Central to the control of the entire clock is a 32-bit MCU development board. It is responsible for maintaining laser lock, performing MW sweep, and disciplining the main 10 MHz oscillator. It therefore needs extensive connectivity as per Figure 53. A Nucleo STM32H743ZI2 platform was selected for this as it has sufficient external interfaces (including Ethernet) and generous onboard flash memory, and because there was experience of this platform within the department. The STM32Cube Integrated Development Environment (IDE), availability of example code, and online support forums were intended to ease the learning curve and shorten time to functional code. Unfortunately, the IDE does not always generate functionally correct code for a requested H/W configuration. Debug was painful and very time consuming.

Not all code had to be produced from scratch. Much of the initialisation code for the platform is generated by the IDE and is provided by ST on an AS-IS basis (is not licensed). The Lightweight TCP/IP stack (LwIP) was written by the Swedish Institute of Computer Science and is available under a BSD-style license. Some of the code needed for MW generation was already available within the department and Dr. Simon Bale has been acknowledged as the author in the source file. This consists of an SPI driver, initialisation function, and low-level functional calls for the HMC835 PLL.

In the previous chapter, the conclusion of the optical path was a 16-bit ADC: a component which currently resides inside the STM32 System on Chip (SoC). ADC3 is a successive approximation device that has been configured for single channel operation with 16-bit resolution, in single-ended mode, and begins capture and conversion on the rising edge of an external line (fed from the FPGA). The performance of this onboard ADC is compromised by the large number of video-frequency peripherals and buses on-chip, and has an Effective Number of Bits (ENOB) of 13.3.

Two onboard 12-bit Digital to Analog Converters (DAC) are employed in measurements in the next chapter. The chip-level outputs are extremely noisy, with significant 10MHz interference and further odd harmonics requiring the use of a Low Pass Filter (LPF). A 1st order device, built with 1kΩ and 470nF components, provided a 340Hz cut-off and 92dB of attenuation to the 10MHz breakthrough. The output was very stable and worked well for low-B/W analogue monitoring.

For ease of development, one of the DACs continuously mirrored the ADC sampled value. This is purely for initial proof of concept and not the intended mode of operation once functional.

Code has been developed in C, and as with the FPGA, has been checked into GitHub: SKs_Rb_clock

4.5 10MHz Oven Controlled Crystal Oscillator (OCXO)

Two different clock sources have been used during this development. As per the introduction, the reference for the clock is a 10MHz OCXO built by Burtichelov as part of his thesis as pictured in Figure 62.

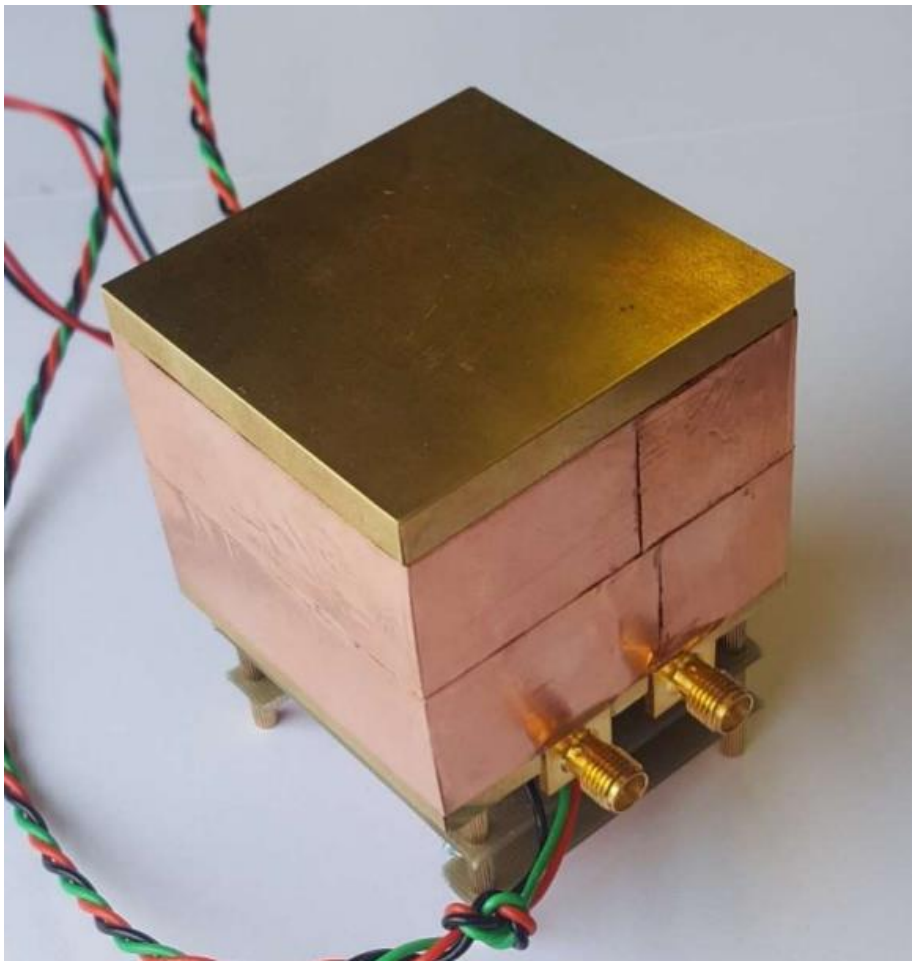

Figure 62. Burtichelov's double-oven 10MHz OCXO including temperature controllers [11, Fig. 3.29].

However, as only one OCXO was fully built and characterised, an alternative source was used on a day-to-day basis, as pictured in Figure 63 and Figure 64 below. This is a simple lab-built power supply feeding a Stanford Research PRS10: a lamp-based atomic clock, but with much higher phase noise close to carrier. This allowed development to proceed without undue risk.

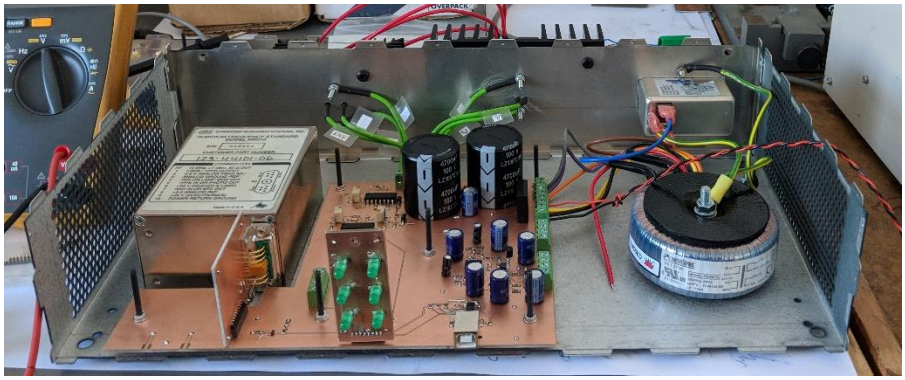

Figure 63: Partially-built PRS10-based reference clock.

Figure 64: Finished view of the PRS10-based reference clock

4.6 MW generation

For initial experiments a Hewlett Packard HP8648C 3.2GHz signal generator was used to supply the 3.035GHz signal needed to excite the hyperfine response. The output frequency was swept by feeding the signal generator's modulation input with a sawtooth waveform generated by an AWG. The first encouraging signs of Ramsey fringes were observed using this MW source, even though each π/2 pulse would be at slightly different frequencies within a single POP cycle.

This was replaced with a more controlled solution which used an Analog Devices HMC835 fractional-N PLL to multiply up the main 10MHz OCXO. The input clock was first multiplied by 5 to create a 50MHz board-level reference frequency. An on-chip, MW frequency VCO then had its output divided and compared with this 50MHz reference. The chip's internal 24-bit resolution allows it to support a 2.98Hz step in MW frequency. As expected with a fractional-N PLL, there are unwanted spectral contributions close to the output frequency as shown in Figure 65. Performance looks acceptable with the desired MW output more than 50 dB above the next closest peak. Control is through an SPI interface connected to the STM32 platform.

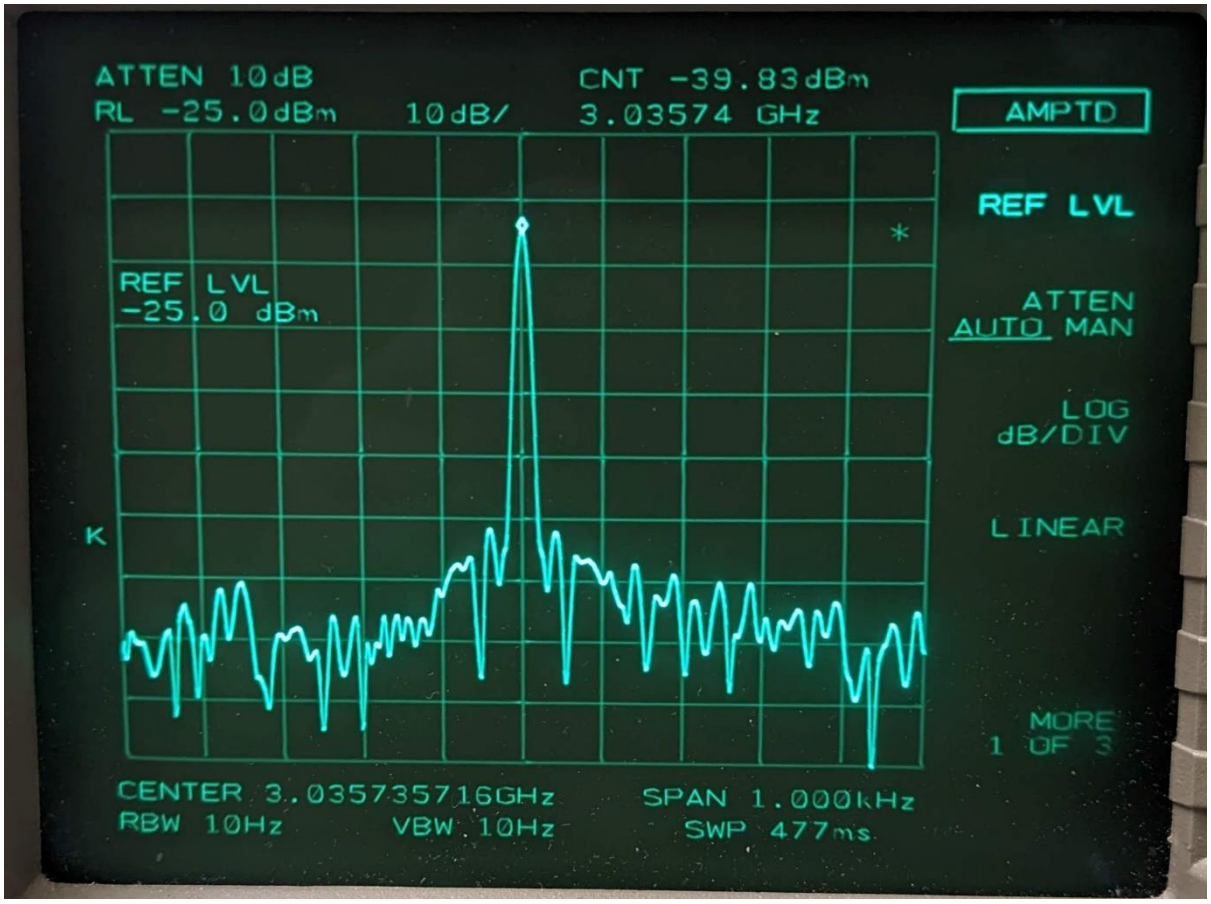

Figure 65. MW generator output, measured on a 26GHz spectrum analyser..

4.7 Clock distribution

Except for the 80 MHz oscillator driving the AOM, all the timing signals used in the clock are derived from the main 10MHz OCXO. This is necessary to ensure that POP timing does not drift, causing a change in the output of the physics package which would be mistaken for a change in MW frequency.

Figure 66 shows how the clock is routed, including a clock buffer used to convert a 50Ω sine-wave source to the 3.3v square wave needed for the CMOS-based FPGA and STM32 board.

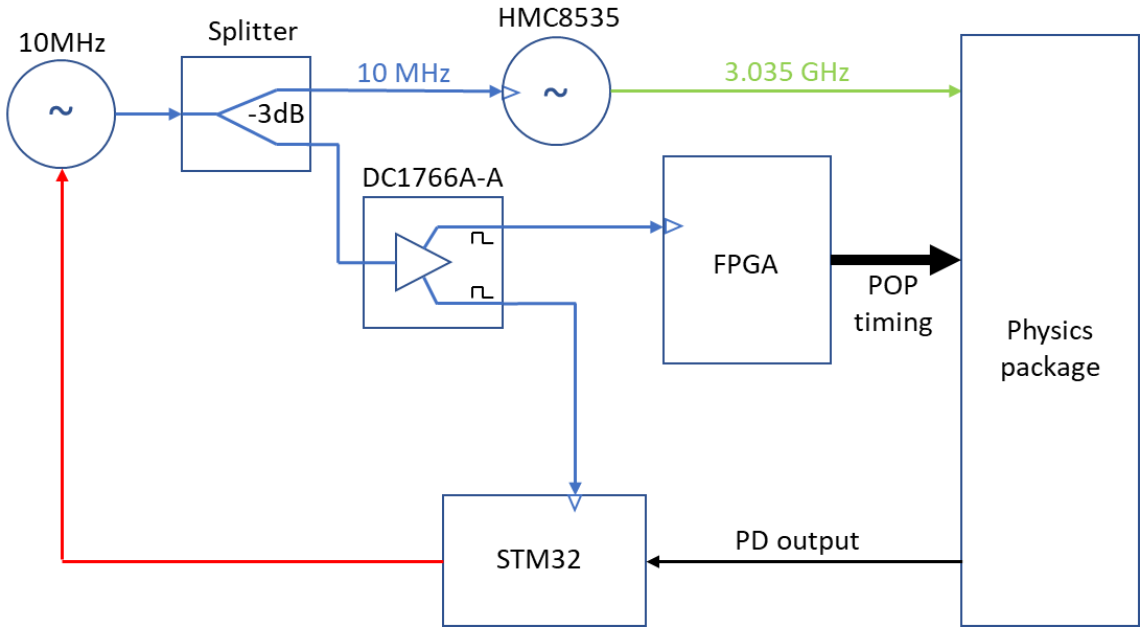

Figure 66. Clock distribution

By default, the STM32 Nucleo board uses an onboard 8 or 8.33 MHz clock but was configured to use 10 MHz from an external oscillator. This allows the onboard timers to be synchronised to the main output.

4.8 24-bit Digital to Analogue Converter (DAC)

In Burtichelov's CPT clock, the LO is disciplined by a small DC voltage which feeds a varactor-based phase shifter in the OCXO. He generates this with a 20-bit Digital-to-Analogue Converter (DAC), Analogue Devices' AD5791, which is specified as having 1 part per million (ppm) stability. However, calculated noise levels of 12.5µV are compromising performance and the current DAC resolution is limiting accuracy to 10^{-13}. By increasing this to at least 23 bits, and improving the noise figure, it should be possible to discipline the OCXO to an accuracy of 10^{-14}. This is significantly beyond the capability of commercially available DACs but, in an application note for the AD5791 device, Egan refers to the architecture shown in Figure 67 below.

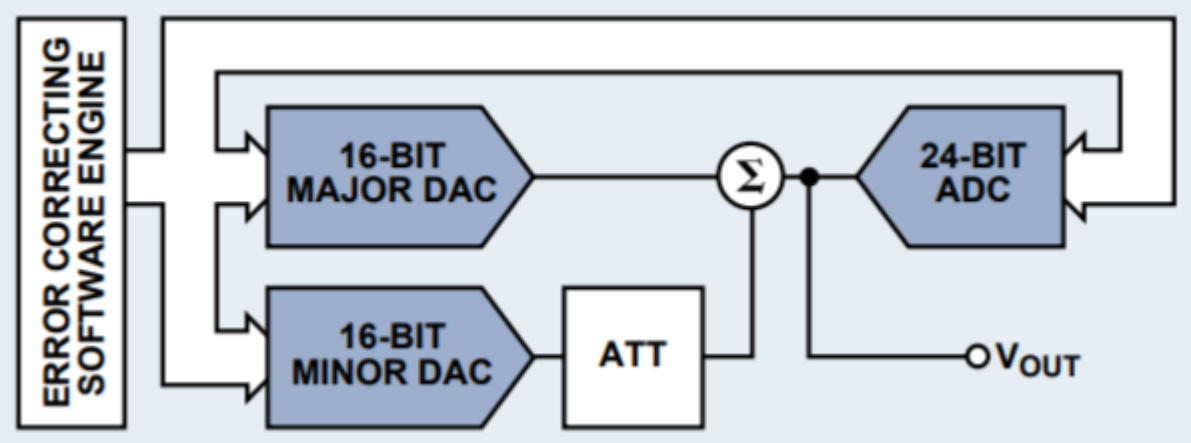

Figure 67: A scheme for increasing DAC resolution [92, Fig. 8]

By utilising a high-resolution Analogue-to-Digital Convertor (ADC), e.g. the 24-bit AD7195, it is possible to control the combined output of the DACs as accurately as it can be measured. Measurement accuracy can be further increased by oversampling: any non-systematic white noise can be greatly reduced by averaging 4^n samples, effectively improving the resolution by n-bits as described in [93], [94].

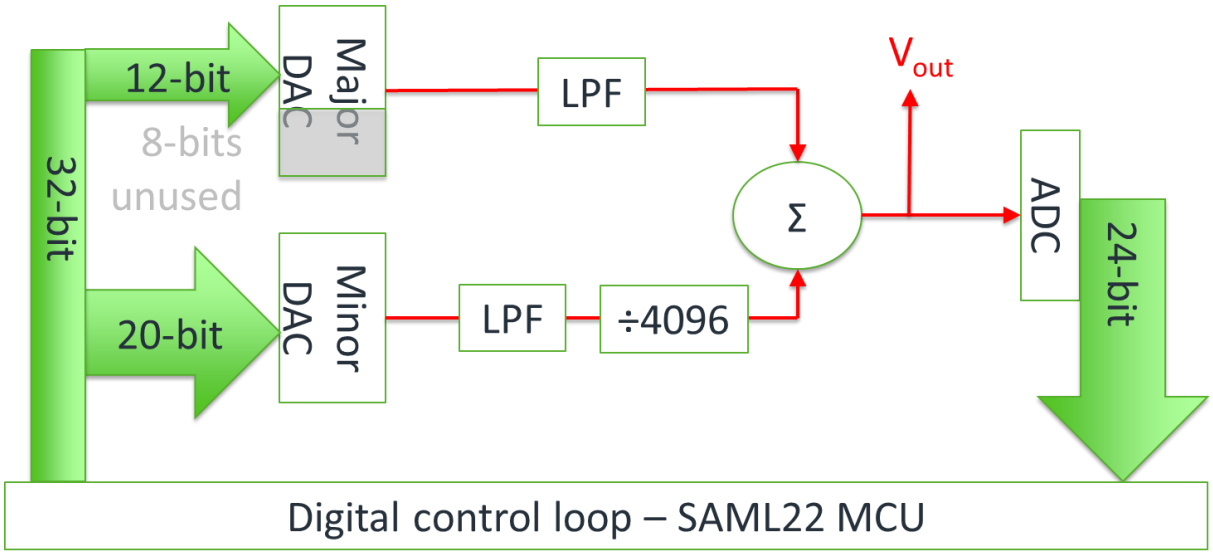

Figure 68: High resolution DAC block diagram

Schematics were drawn up, and are available in Appendix E, and three prototype boards constructed, one of which can be seen in Figure 69 and Figure 70.

Figure 69: High resolution DAC - first prototype PCBA.

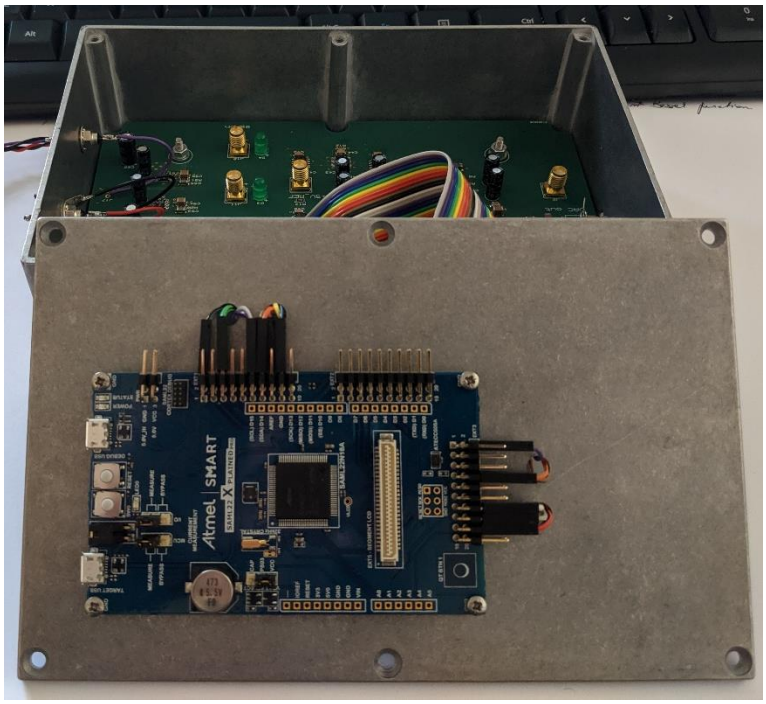

Figure 70: High resolution DAC with MCU connected.

As this building block was developed early in proceedings, the MCU seen in Figure 70 is an Atmel SAML22 device, and not the final choice of STM32. Connectivity between the MCU and the two DACs, and the ADC, is via SPI. Code has been written in C, but using the AtmelStart platform, and their ASF4 software framework. It is available from the following repository: [DAC_proto_ucode.]

5 Spectroscopy, Double Resonance, and POP

5.1 Overview

This chapter covers a number of disparate steps needed to show that the Physics Package and Control Platform (detailed in Chapters 3 and 4) are capable of supporting POP. It describes the process of laser bring-up, initial spectroscopy, how physical parameters were optimised, and how additional complexity was added to demonstrate POP.

5.2 Laser setup

As mentioned earlier, the Laser Diode (LD) did not come as a part of a pre-packaged product, so it was necessary to establish a sensible 'working envelope'. Thorlabs supplied the LD with an individual characterisation sheet suggesting a diode operating temperature of 15C and current of 240mA would provide an output suitable for 85Rb (λ=780.03nm in air), but this would mean operating the device at 96% of its full (45mW) power rating. Instead, 33mW output power was selected as the desired operating point as this would prolong diode life, whilst still providing near-optimal laser linewidth and sideband suppression. Next, a suitable LD operating temperature had to be determined that could allow (a current-based) frequency adjustment without encountering a mode-hop. A Thorlabs CCS100 350 – 700 nm CCD spectrometer was used to analyse the output while slowly sweeping the diode current by 50mA. Figure 71 shows the typical output expected with one peak of minimal width.

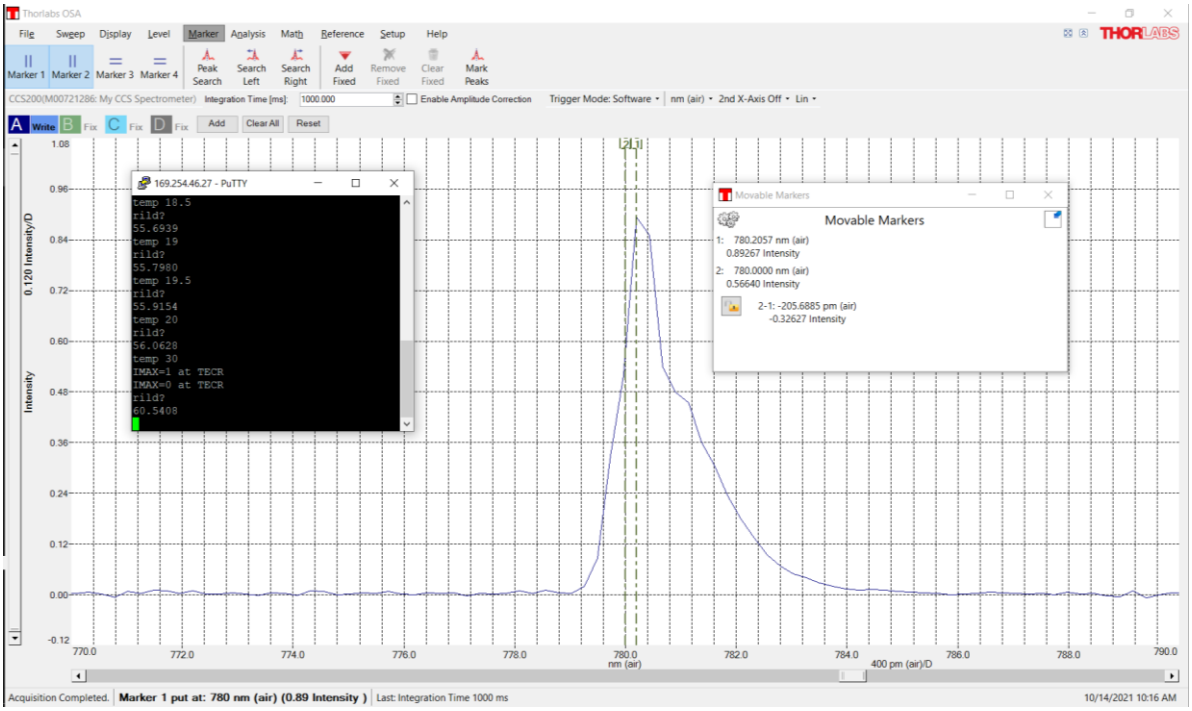

Figure 71. Beam spectrum as measured with a Thorlabs CCS100 – intensity plotted against wavelength/nm.

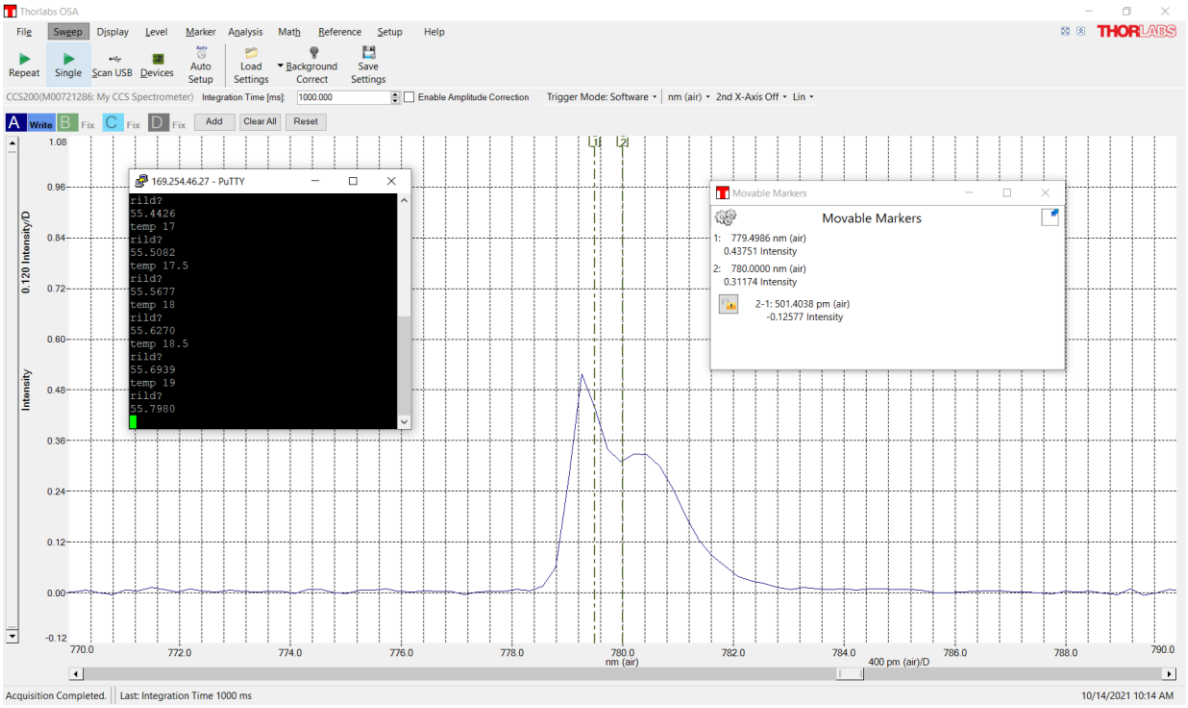

Figure 72. Beam spectrum during a mode-hop, with characteristic drop in power and double peak – intensity plotted against wavelength/nm.

Figure 72 is an example of the double peak seen when a mode-hop is in progress i.e. when the standing wave developed in the laser cavity undergoes a longitudinal mode shift. Although it is not possible to calculate the expected size of a hop from information provided in the datasheet, a 0.5nm hop is not an unexpected result for this type of diode. Each mode oscillates at its own frequency and if operating conditions are just on the upper limit of one mode, and just on the lower limit of another, noise will increase, and the output power will fluctuate whilst being spread across the two frequencies. A suitable operating range was found to be 160mA to 180mA at 21.15°C. No mode hops were detected in this region and a current of 173mA produced 33.4mW of optical power centred on a wavelength of 780.0nm. It was noted that the current required to maintain this wavelength increased by up to 4mA if the diode had been powered up for 60 minutes.

5.3 Initial spectroscopy

To control the number of new elements being introduced at one time, a mixed-isotope (72% ^{85}Rb and 28% ^{87}Rb) vapour cell was taken from a known-working experiment along with its water heating apparatus. It is shown in Figure 73 at the centre of two (unused) sets of perpendicular Helmholtz coils, sitting on an optical bench, with cardboard tubing to act as shielding from ambient light.

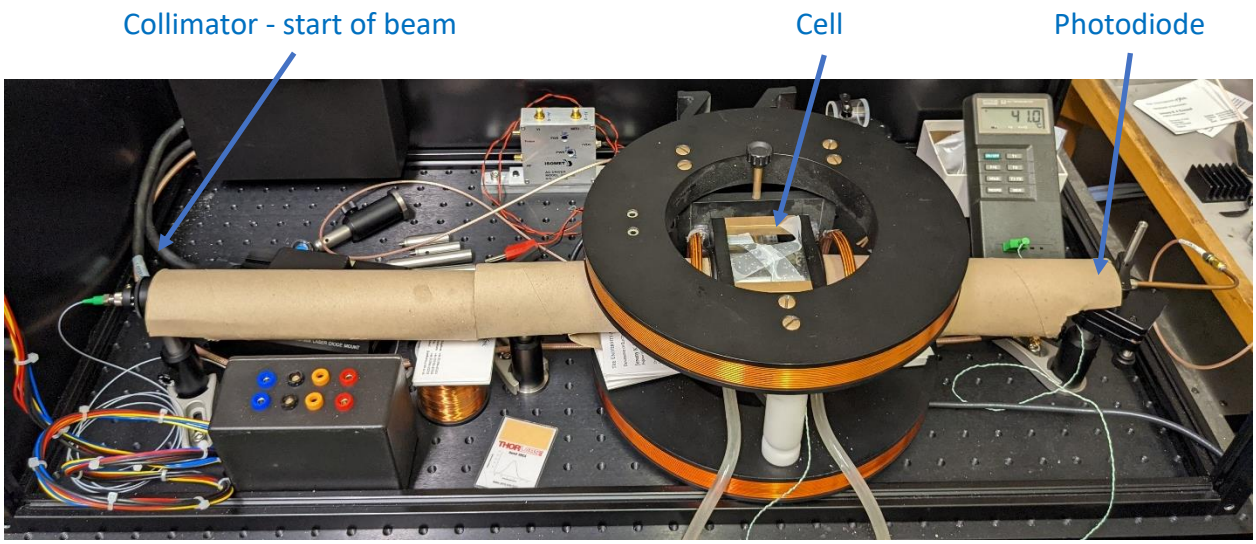

Figure 73. Basic absorption setup on optical bench

The optical path was terminated with a large-area photodiode connected to a Stamford Research SR570 transimpedance amplifier, with gain set to 500 nA/V, providing an optical response of ~1.06 V/µW at 780 nm. Laser output power was set to 1 mW (58 mA drive current) for alignment: this is equivalent to a CAT 1 laser which will not cause retinal damage. Alignment was established from the laser collimator to the photodiode, first using a fluorescing photocard to determine the position of the IR beam, then maximising the electrical output once roughly lined up with the photodiode.

Figure 73 also includes 1x NE13A-B and 1x NE30A-B absorptive neutral density filters placed in series between the laser and cell. Laser current was increased to 170mA and the two filters in series provided an optical density (OD) of ~2.94 at 780nm i.e. transmission of 0.1% of the incident beam. The beam was enclosed in cardboard tubes, using the photodiode output to check that alignment was still good.

5.3.1 Temperature sweep

By setting the laser diode current at 170mA, and sweeping the temperature it is possible to perform a broad frequency sweep, with moderately constant power output (assuming one avoids a mode-hop).

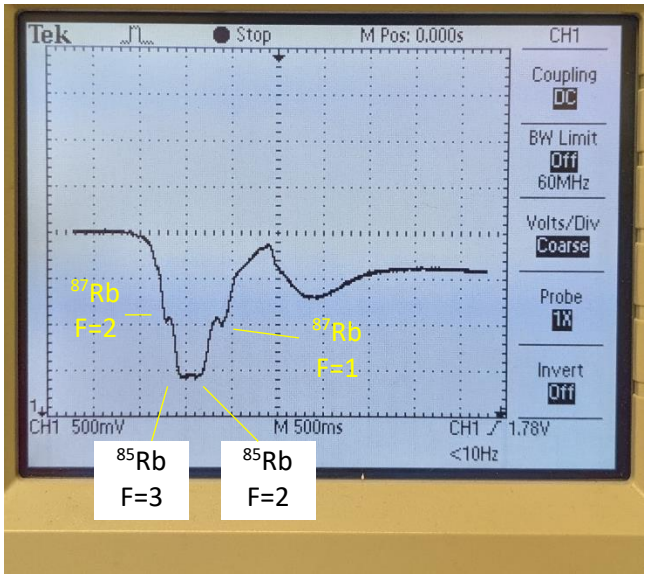

Figure 74. Mixed isotope Rb absorption dips: an increase in laser frequency has been stimulated by an increase in temperature.

Figure 74 shows the photodiode response as the temperature rises from 20.5°C to 21°C. The dip contains the 4 peaks associated with the D1 and D2 absorption lines associated with each naturally occurring isotope of Rb i.e. the cell contains both 85Rb and 87Rb. The x-axis is time, and is not a linear scale of temperature; the rightmost peak and dip are caused by a temperature overshoot, and undershoot, before the diode temperature stabilises. The vertical scale represents optical output in arbitrary units.

5.3.2 Current sweep

The mixed-isotope cell was then replaced with the waveguide and ^{85}Rb cell, as per Figure 75.

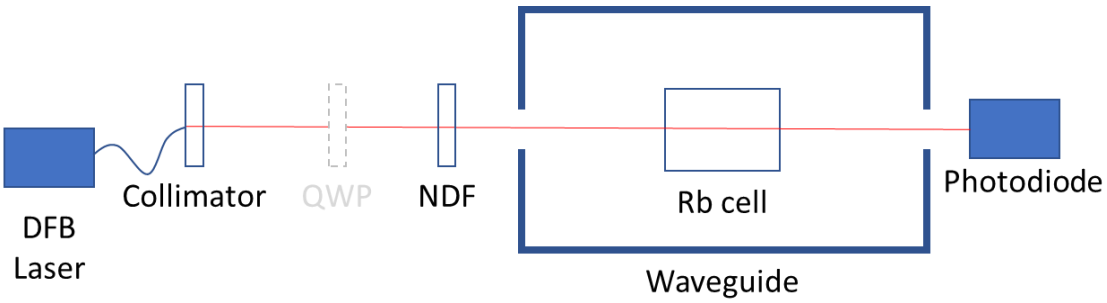

Figure 75. A simple optical configuration supporting basic absorption and double resonance.

As per section 3.8, the laser diode temperature was fixed at 21.15°C and, in contrast to Section 5.3.1, frequency is tuned by adjusting current. Using the Laser Driver's built-in SCAN function, it is possible to ramp current from 150mA to 200mA as shown in

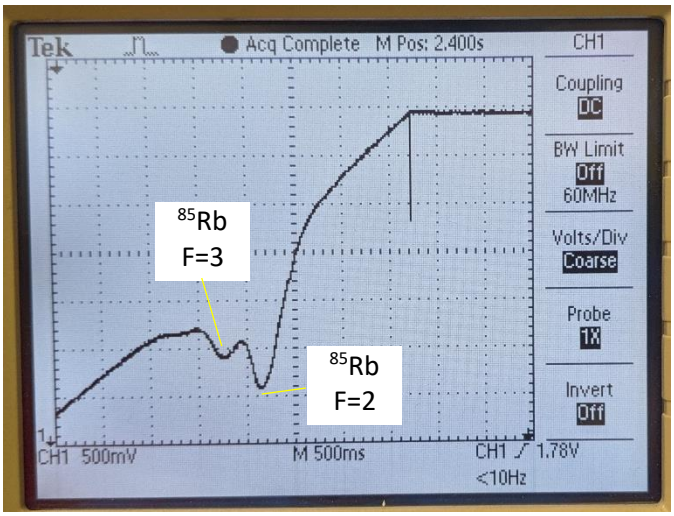

Figure 76.

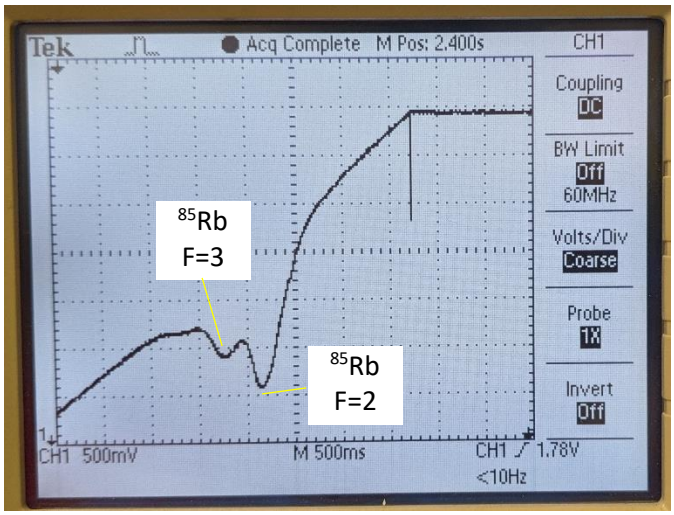

Figure 76. Photodiode output with increasing laser current.Note that power increases with current but that the ^{85}Rb lines can still be clearly distinguished. The x-axis represents current increasing at a rate of 5mA/s, and therefore corresponds to an increase of 4.9GHz for each 500ms division. This was not a precision measurement but was reassuringly consistent with the expected 3GHz ^{85}Rb hyperfine spacing.

5.3.3 Optimising laser power

By placing Neutral Density Filters between the laser and cell, it was possible to reduce the optical power received without affecting frequency. The swept current measurement was repeated with a combination of different filters, and with optical power levels recorded. In Figure 77 to Figure

81, the laser temperature remains fixed at 21.15°C and laser current is ramped from 160 to 185 mA over 5s, therefore the horizontal scale is 2.47GHz per 500ms division. It should be noted that no C-field, nor MW power is present yet, and beam is linearly polarised (the QWP was present but with the fast axis aligned with that of the emerging beam). As laser power varies with current, the optical power figures below have been stated when the frequency is aligned with the F=3 dip.

After optimizing laser power for best optical contrast, the trace seen in Figure 77 was observed, corresponding to 15µW incident on the cell.

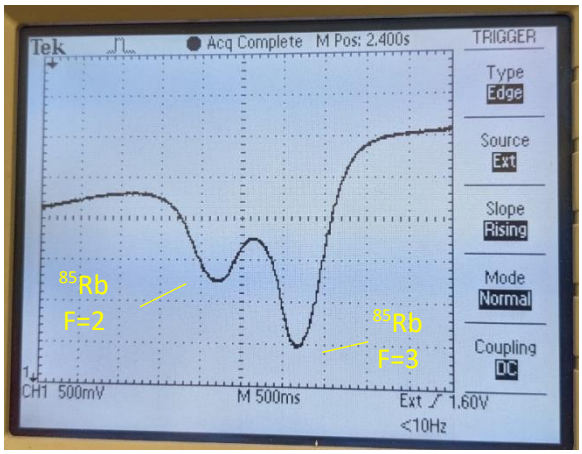

Figure 77. Optical output plotted against increasing laser frequency. P_{cell} = 15µW

Figure 78 to Figure 81 show examples of varying laser power and the effect on optical contrast.

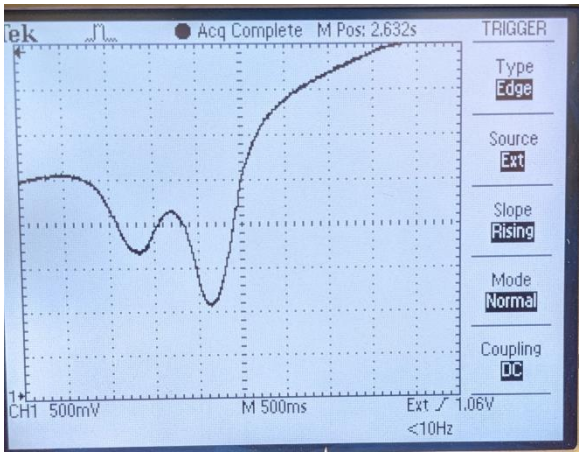

Figure 78. Optical output plotted against increasing laser frequency. P_{cell} = 38µW

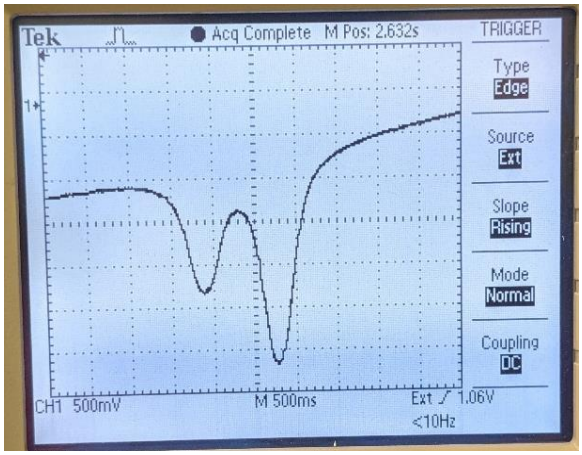

Figure 79. Optical output plotted against increasing laser frequency. P_{cell} = 19μW

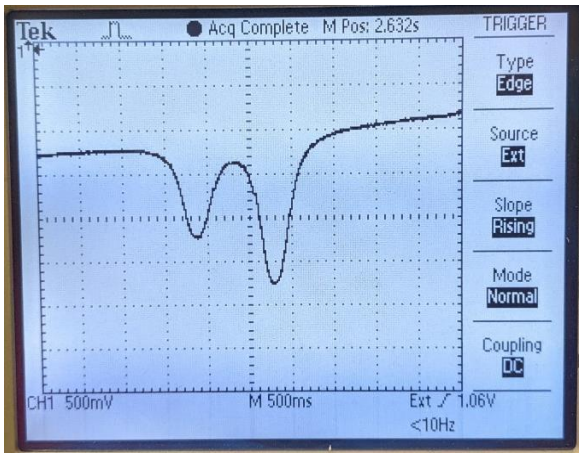

Figure 80. Optical output plotted against increasing laser frequency. P_{cell} = 9.6μW

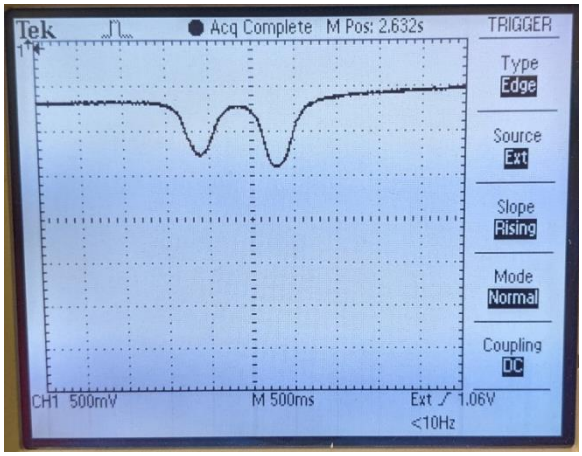

Figure 81. Optical output plotted against increasing laser frequency. P_{cell} = 3.6μW

5.4 Double resonance

Having established basic spectroscopy the next logical step was to show that the waveguide arrangement, shown in Figure 33 and Figure 36, would support double resonance. Figure 82 shows the energy states associated with the D2 transition when the laser is tuned to the F=2 dip. As the laser

de-populates atoms in the F=2 state, a 3.035GHz MW field allows atoms in the F=3 ground state to transition to F=2, and leaves them available for interaction with the laser beam again.

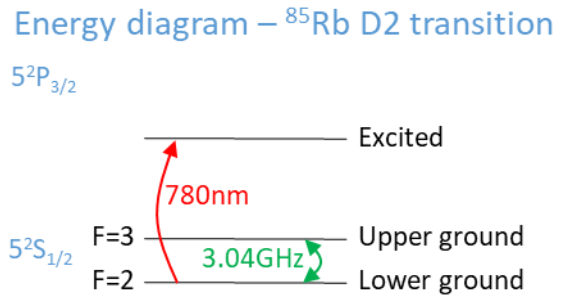

Figure 82. Energy diagram of the ^{85}Rb D2 transition

Still using linear polarisation, and with no C-field, the F2913 switch was used to pulse a 10dBm microwave source. After losses in the switch and cables, this translated to +3.3dBm 'on' power, and -68dBm 'off' power as measured passing through the waveguide. The laser was tuned to the ^{85}Rb F=2 absorption dip. Figure 83 shows the Physics Package's optical response to this pulsing MW field with peaks correspond to -68dBm MW power, and dips to +3.3dBm received power.

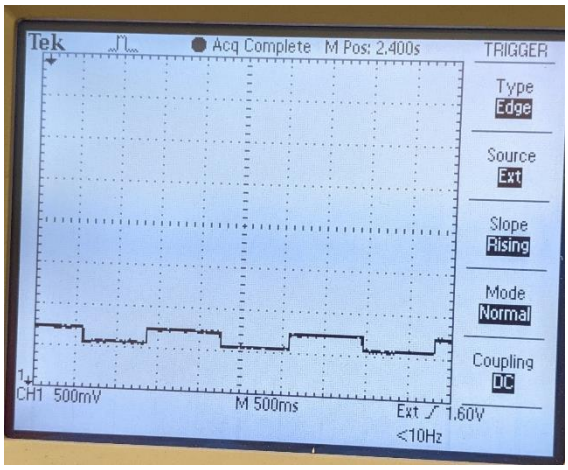

Figure 83. Optical output whilst a pulsing MW field is applied.

5.4.1 Measuring the Double Resonance (DR) width with varying MW power

The laser was aligned to the ^{85}Rb F=3 dip and the Helmholtz coils were fed with a 165mA current to generate a calculated 97µT C-field. The MW frequency was swept across range 3.035732439 GHz +/- 25kHz with fixed power levels ranging from +7.8dBm to -22.7dBm (power levels are again specified at the output of the waveguide). Figure 84 to Figure 86 show a selection, with higher microwave power offering increased optical contrast at the expense of broadening the FWHM. The horizontal scale is 5kHz per division.

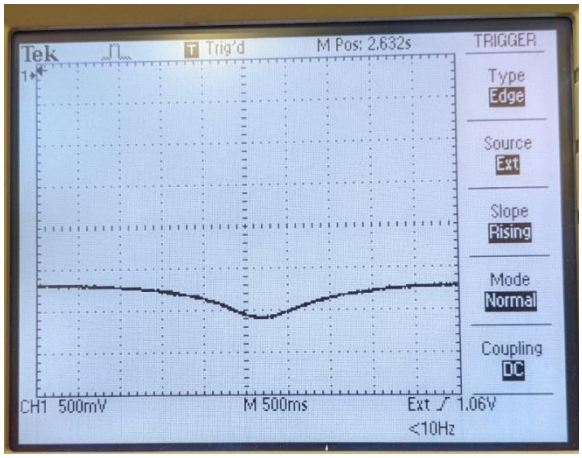

Figure 84. Optical response plotted against increasing MW frequency. +3.3dBm MW power

Figure 85. Optical response plotted against increasing MW frequency. -12.7dBm MW power.

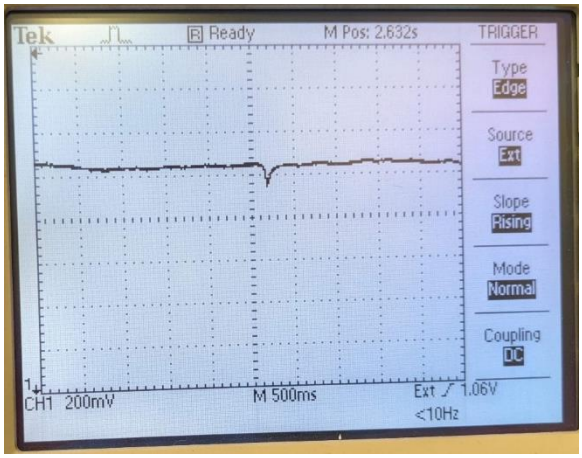

Figure 86. Optical response plotted against increasing MW frequency. -20.7dBm MW power.

A generator power of -10dBm, corresponding to a received power of -16.7dBm, was selected as being a suitable compromise with FHMW of <2kHz.

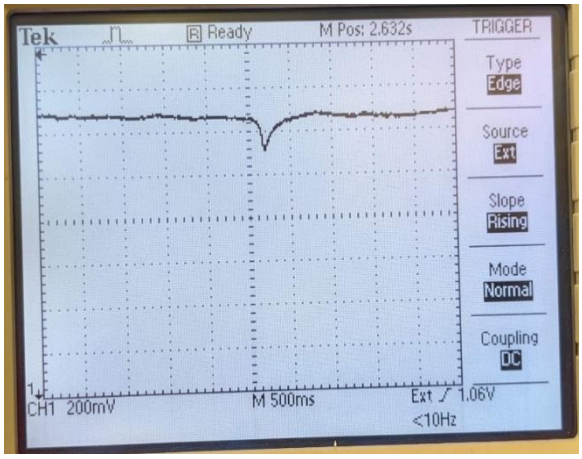

Figure 87. Optical response plotted against increasing MW frequency. -16.7dBm MW power.

5.4.2 Measuring the DR width with varying C-field

This time the laser was aligned to F=3 dip and MW power was fixed (at -10dBm generated, -16.7dBm received). The MW frequency was again swept across range 3.035732439 GHz +/- 25kHz whilst the C-field was varied. Optical output was then plotted against swept MW frequency: examples of which are provided in Figure 88 and Figure 89. The horizontal resolution remains 5kHz per division.

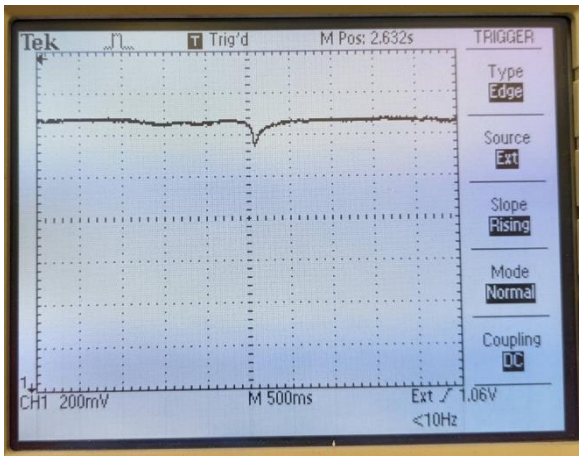

Figure 88. Optical response plotted against increasing MW frequency. 47μT C-field.

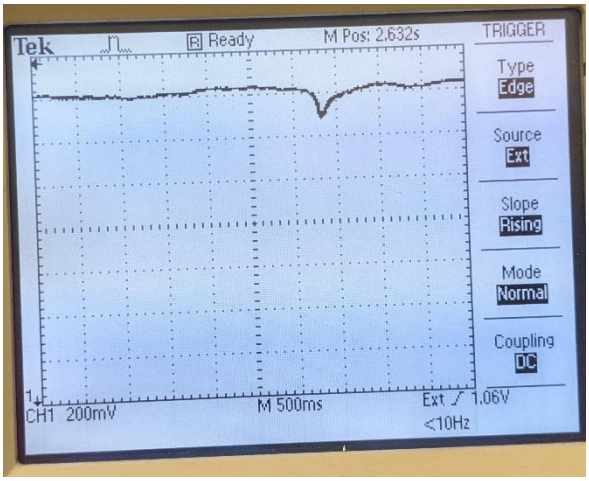

Figure 89. Optical response plotted against increasing MW frequency. 265µT C-field.

5.5 Choice of F=2 versus F=3 dips absorption lines

Figure 90 and Figure 91 offer a comparison of the FWHM when the laser is tuned to the ^{85}Rb F=2 dip and when it is tuned to F=3. The figures below, the MW power is fixed at -6dBm, with 59µT C-field, and 5kHz increase in MW frequency per horizontal division

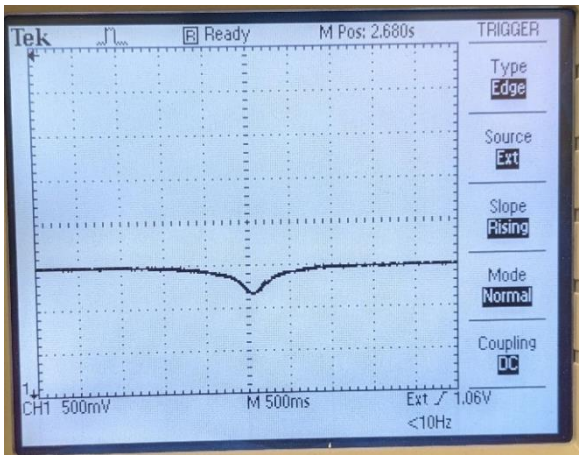

Figure 90. Optical response plotted against increasing MW frequency. Laser tuned to F=2 dip.

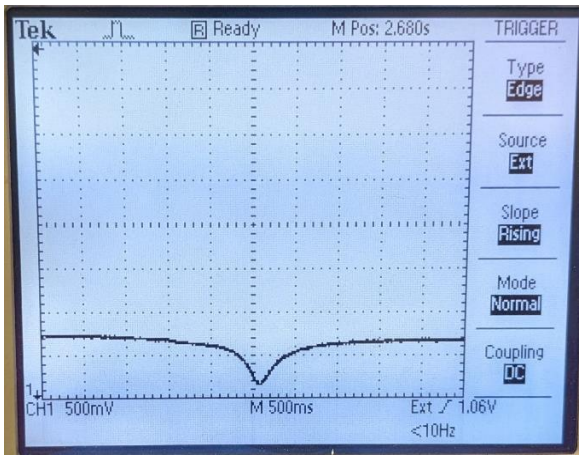

Figure 91. Optical response plotted against increasing MW frequency. Laser tuned to F=3 dip.

The F=2 line offers slightly narrower FWHM, though the optical contrast of the F=3 dip is better. Either dip is entirely suitable for continuing work until more accurate measurements can be conducted with automated laser frequency stabilisation. Figure 92 shows a control experiment that has been performed with the laser tuned to the peak between F=2 and F=3 dips and is therefore not aligned with either of the D-lines. It confirms that there is no hyperfine response, and no response to TE01 field, unless the laser is tuned to one of the dips.

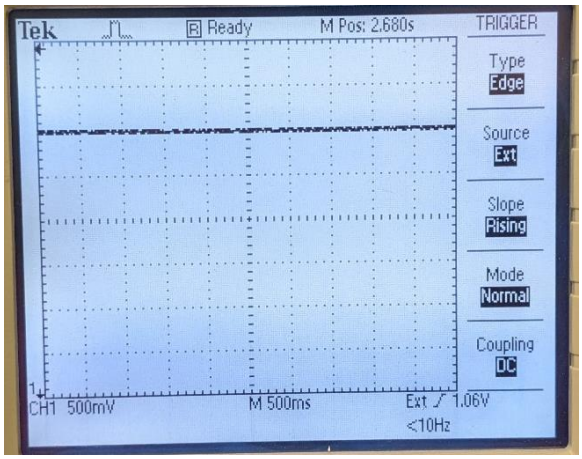

Figure 92. Optical response plotted against increasing MW frequency. Laser is not aligned with a D-line.

5.6 Further preparation for POP

At this point the simple optical path in Figure 75 was replaced with the full Physics Package as per Figure 17, and previously optimised values were used as starting points when looking for POP. It was assumed that probe beam power was likely to be similar to the optimal value for absorption spectroscopy, and that the MW power that provided a good balance of DR contrast and width was likely to also work for POP.

5.6.1 Ground state depopulation

When spectroscopy was performed with a pulsing probe beam it was noted that the ground state could be seen de-populating in real time. Using a continuous beam, the laser was first tuned to the F=3 dip. Figure 93 shows the photodiode output as the probe beam is pulsed on and off. The photodetection bandwidth is sufficiently high that there is a crisp response to the laser being turned off. The delay visible on the rising edge is due to a physical effect, i.e. depopulation of the F=3 ground state. Although not shown here, the identical effect was observed for F2 dip.

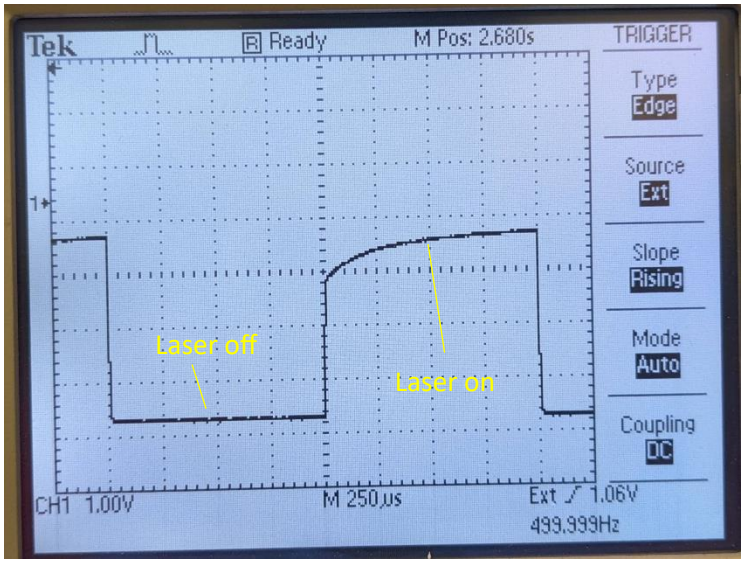

Figure 93. Optical response to the pulsing probe beam.

5.6.2 Optical response to pump and probe beams

As a means of checking the optical response during a POP cycle, the cell was exposed to pump and probe beams and the output recorded. There was no MW power present, and in the first instance, the laser frequency was tuned significantly below either D-line. Figure 94 shows the output clipping when the pump beam is active, but that the probe beam is well within the available range of measurement. This clipping is due to the output of the PDAPC1 transimpedance amplifier saturating when the pump beam is incident on the photodiode. The rising and falling edges are sharp and consistent with the calculated 28kHz B/W.

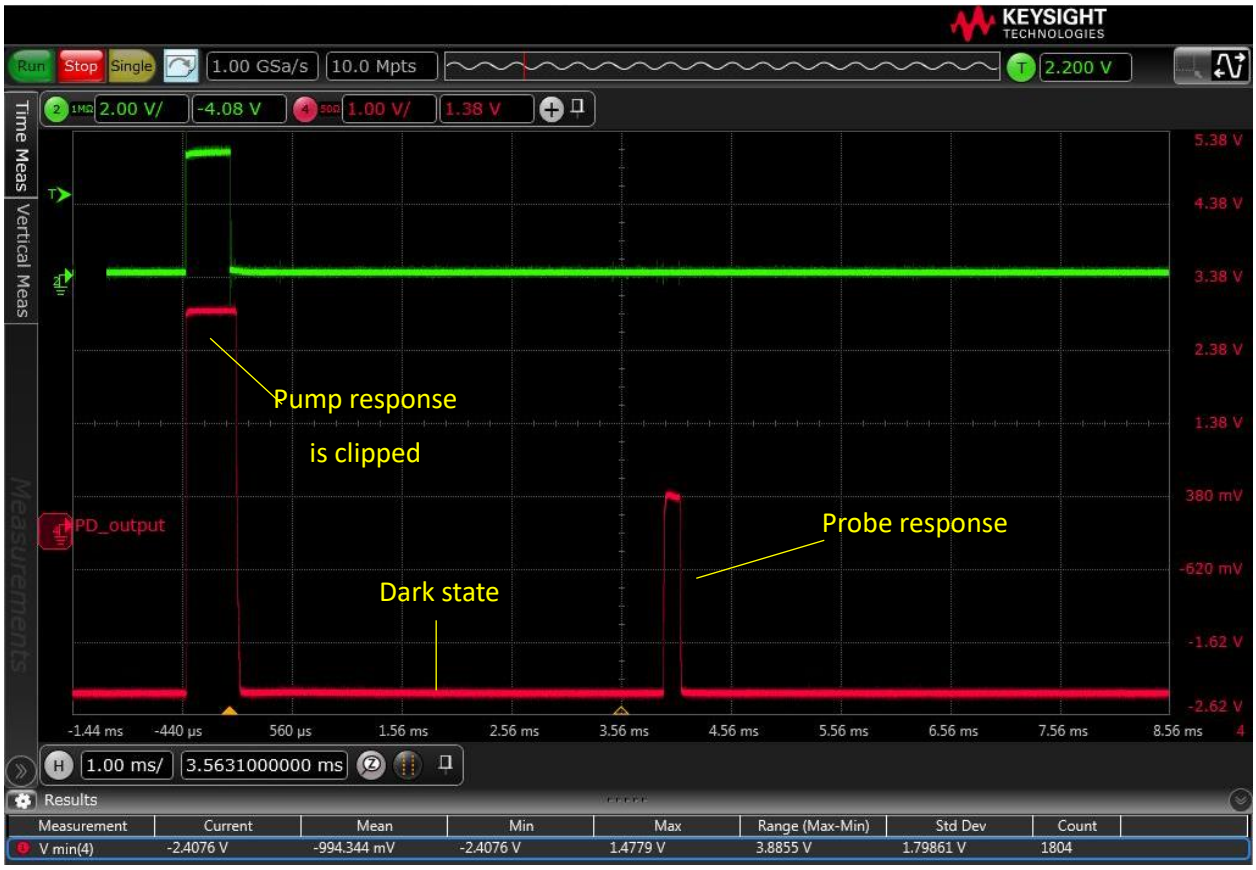

Figure 94. Optical output of physics package with laser frequency significantly below D-lines.

Please ignore the green trace: the FPGA's 'pump' output has been used for synchronisation purposes.

The laser was tuned to the ^{85}Rb F=2 dip, and the response to the same stimulus is shown in Figure 95, with the seemingly slow rise time due to depletion of atoms in the F=2 state.

Figure 95. Optical output of physics package with laser tuned to F=2 absorptive dip.

5.6.3 Phenomena around switching MW on and off.

The system response to switching the MW on and off was investigated. It was measured as a secondary effect by tuning the laser to the F=2 absorptive dip, pulsing the MW switch, and observing the optical output. Figure 96 and Figure 97 show the optical output as the MW power is turned on and off respectively. Each π/2 pulse is likely to be around 400µs long (based on [95]) so very brief effects were

unlikely to matter significantly. However, at elevated MW power, significant transition effects could be observed, lasting as long as 1ms.

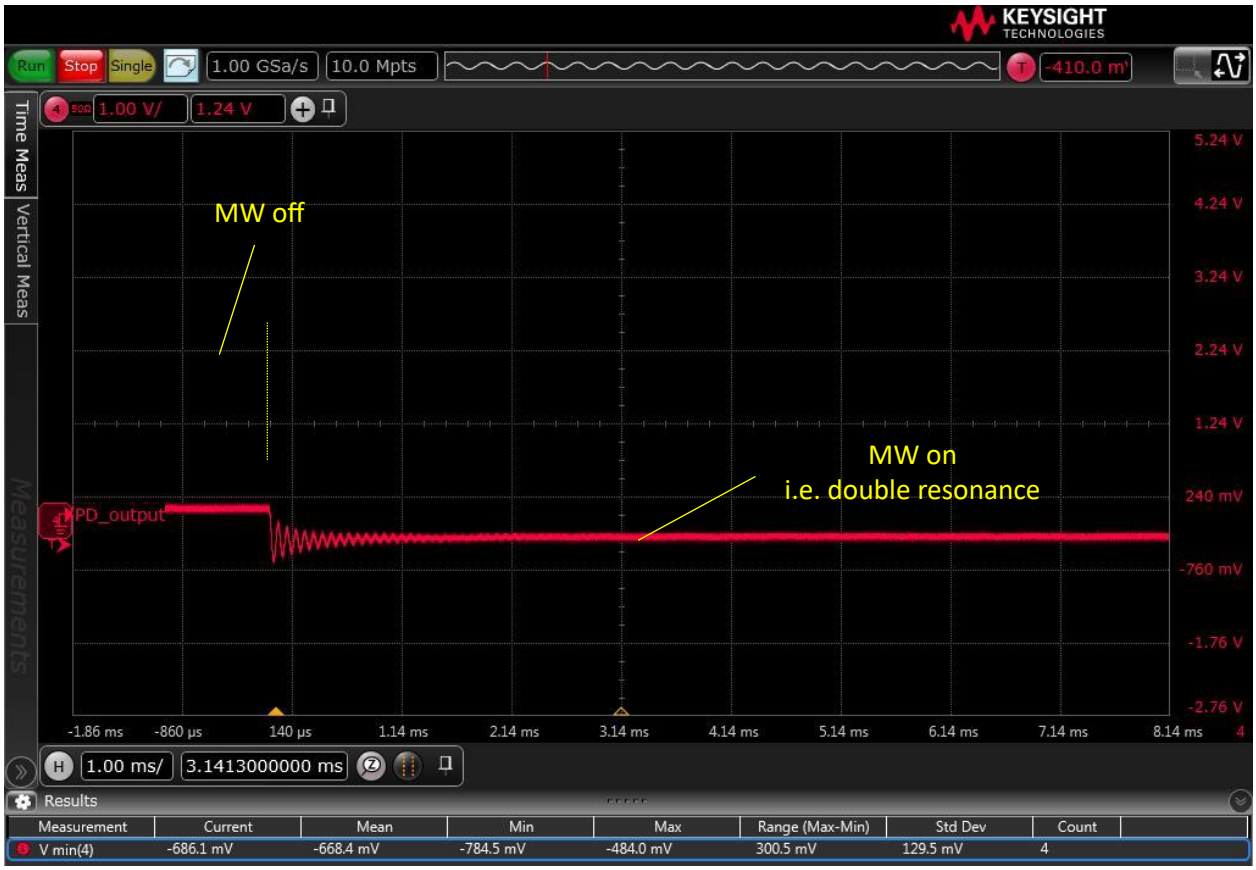

Figure 96. Optical output as MW power is turned on (21dBm input power to switch)

Figure 97. Optical output as MW power is turned off.

The ringing seen in Figure 96 is thought to be partly due to the length of the waveguide, but probably exacerbated by the length of cable between the absorptive MW switch and the waveguide. Locating the two of these as close as possible should alleviate matters somewhat. Fortunately, this effect is only prominent at power levels ~30dB higher than the expected operating point.

5.6.4 MW settling time

One further concern was that the MW generator output should take 30 µs to stabilise after a frequency change is requested, however, this is highly dependent on the loop filter design and initial observations of the development platform appeared to be taking much longer than specified. This was quantified in much the same manner as before, i.e. measured as a secondary effect by tuning the laser to the F=2 absorptive dip, tuning the MW generator to the hyperfine frequency, requesting a small drop in frequency, and observing the optical output.

The experiment showed that a requested rise in MW frequency showed a heavily damped response that stabilised after 3ms. A drop in frequency caused lightly damped oscillations that required longer to settle. In Figure 98 and Figure 99, the Y-axis is the optical output in V, the X-axis is time in ms, the blue trace is a single sample, and the black trace is the average of the last 16 samples (offset slightly for visibility).

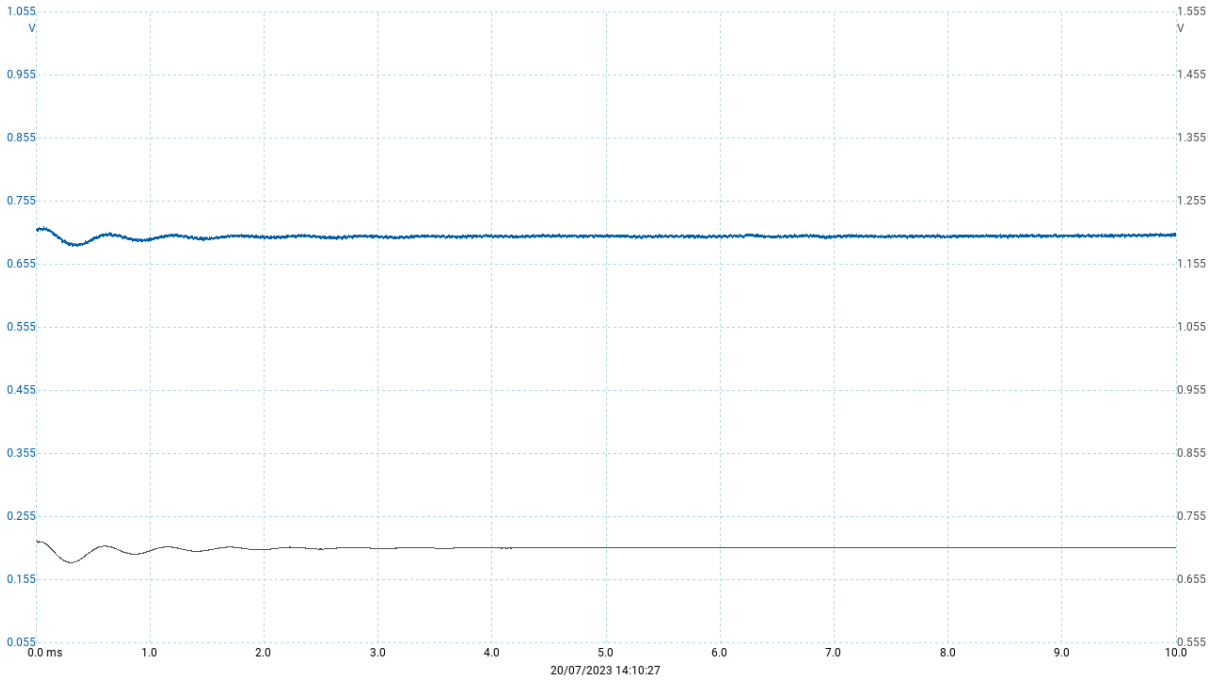

Figure 98. Optical response to a 9Hz drop in MW frequency.

The same measurement, zoomed in.

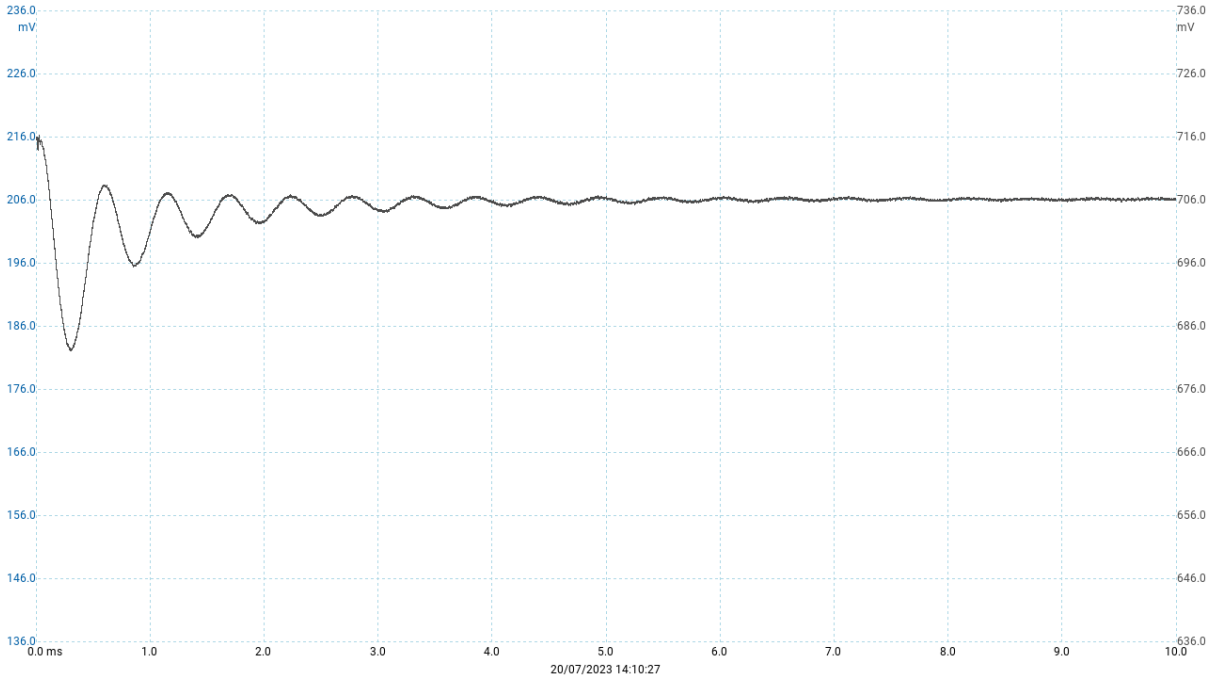

Figure 99. Optical response to a 9Hz drop in MW frequency (zoomed in).

As a simple workaround, STM32 code was amended to allow 6ms between requesting a MW frequency change and starting a POP cycle. Optimising the PLL's loop filter would be a more elegant solution but requires further work.

5.7 First POP results

Initial timing values were taken from Micalizio et al. [95, Fig. 3] as a sensible starting point. These values were 400µs pump pulse, 400µs π/2 pulse, 7.5ms free precessing time, and 150µs probe pulse. The Probe beam and MW power were set to the optimal values established earlier: 15µW incident on the cell and -16.7dBm received power respectively. After some careful exploration of the state space, the first promising results were observed as shown in Figure 100. This measurement was taken with the amplified photodetector with its output fed through the signal conditioning block, then through the sample and hold block. A Picoscope 4262 provided 16-bits resolution.

The free precessing time has been reduced to only 1ms to avoid any possibility of exceeding the relaxation time but a central fringe can be seen, as can the voltage spikes when the 'hold' is released each time.

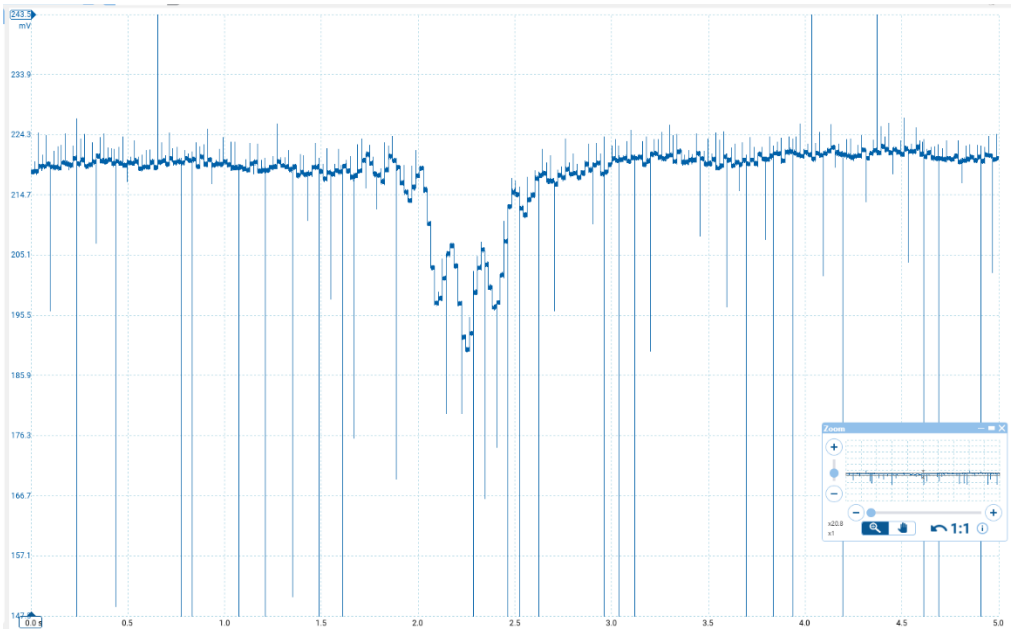

Figure 100. Optical output of POP cycle as MW frequency is swept over 10kHz. -28.2dBm MW power, 58.8µT C-field.

Table 12 shows the values used to achieve this.

Table 12. POP timing values

	Number of 400ns periods	Length/s
Pump	2000	800µ
π/2	795	318µ
Free precessing time	2500	1m
Probe	2500	1m
Probe to sample delay	2000	800µ
Sample time	50	20µ
Laser – MW gap	10	4µ
Post cycle delay	40000	16m

5.8 Optimising POP

The complete Physics Package and control loop were connected, and basic spectroscopy repeated. Figure 101 shows the optical response on improved measurement equipment (a 12-bit, 4 channel oscilloscope).

Figure 101. Photodiode response to increasing laser diode current showing both $5^2S_{1/2}$ -> $5^2P_{3/2}$ transitions. 14µW optical power on cell.

As per Section 5.4, double resonance was measured again to ensure that laser and MW power were appropriate, with results plotted in Figure 102. Although the optical contrast was greatest for optical powers below 10 µW, it offered poor SNR. Values of 14.2µW probe beam power, and -17dBm MW power (received at the waveguide) were chosen as providing a good compromise between contrast and SNR.

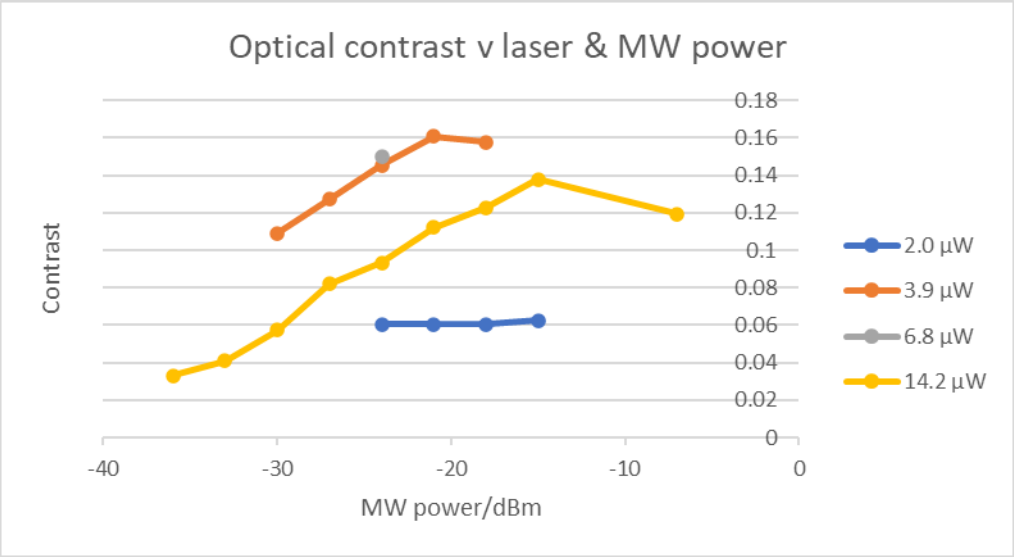

Figure 102. Optical contrast of double resonance plotted against MW power. The legend shows different levels of optical power.

The use of the C-field was explored again. Optical power was fixed at 14.2µW, MW power was -7dBm measured at the waveguide, and the MW frequency was swept by 10kHz. Figure 103 and Figure 104 show the double resonant response without, and with a C-field respectively. Feeding the output to a spectrum analyser shows large spikes at 50Hz and 150Hz: double resonance is clearly very susceptible to mains-borne interference in the lab. The C-field generated by the Helmholtz coils reduces the effect of interference, but to the detriment of a broadened resonance.

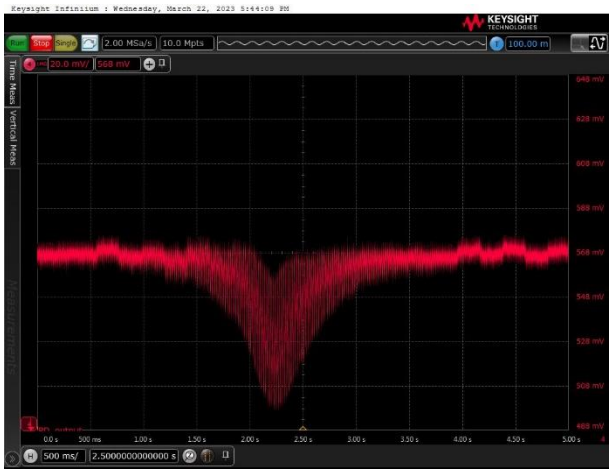

Figure 103. Optical output plotted against increasing MW frequency. No C-field.

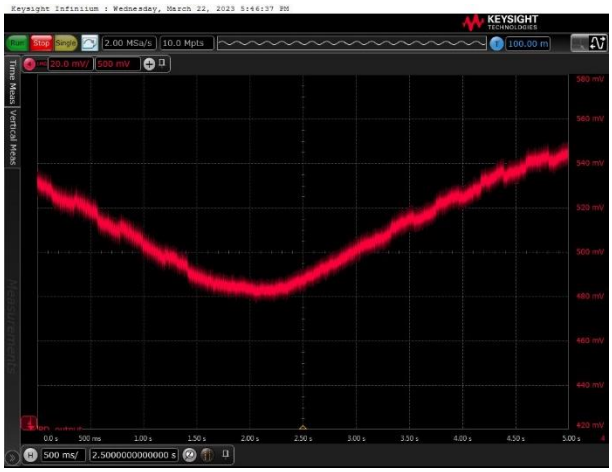

Figure 104. Optical output plotted against increasing MW frequency. 58.8µT C-field.

For all remaining POP measurements in this thesis, the Helmholtz coils have been fed with 5.9 mA, producing a 10µT / 100mG C-field, This is measurably smaller than the ~20µT ambient field, and highlights the need for magnetic screening.

5.9 Improved POP results.

Figure 105 to Figure 107 are plots of optical output plotted against a 10kHz sweep in MW frequency. The noise is largely due to the use of an ADC on-board the STM32: initial tests have shown that this is Gaussian in nature and can be cleaned up significantly by multi-point averaging. A 4-point average has been used as this improves the SNR by a factor of 2, though each measurement therefore takes 4 times the POP cycle time.

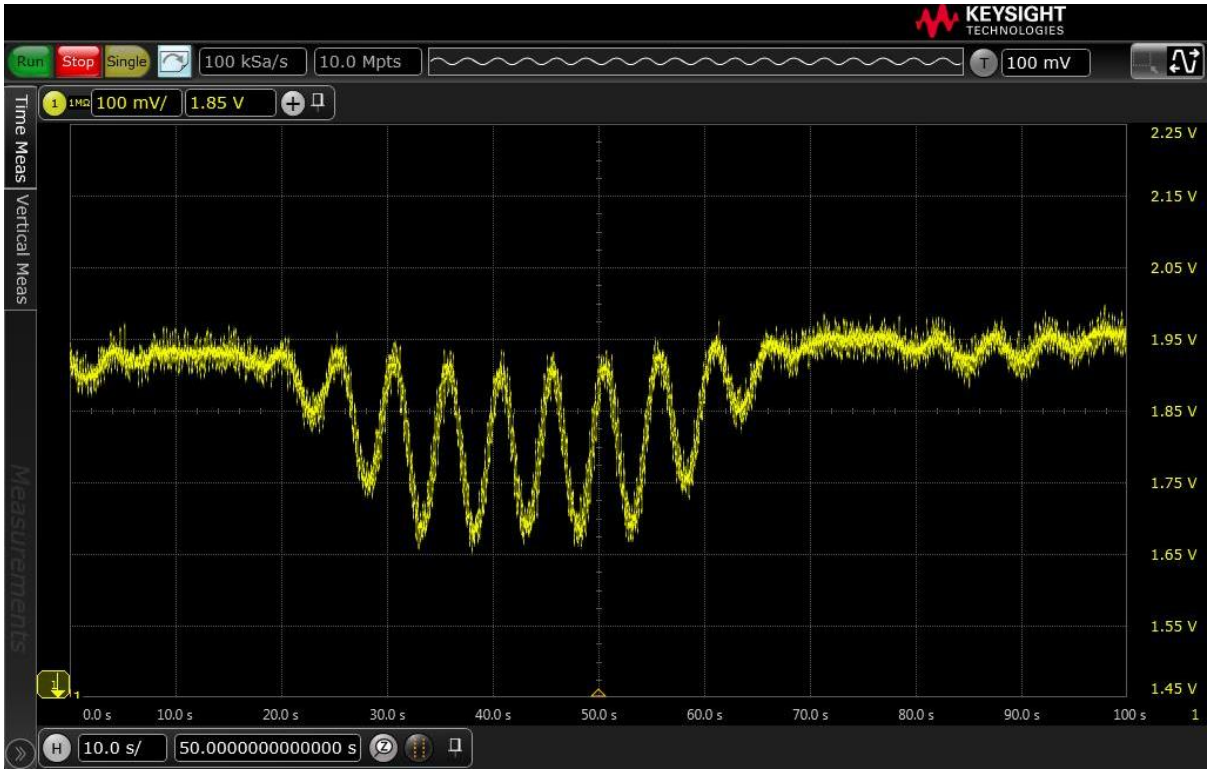

Figure 105. Optical output plotted against increasing MW frequency: 2ms free-precessing time. Horizontal scale is 1kHz / division.

As per Section 2.2, both the width and the amplitude of the central fringe are reduced as the free-precessing time is increased.

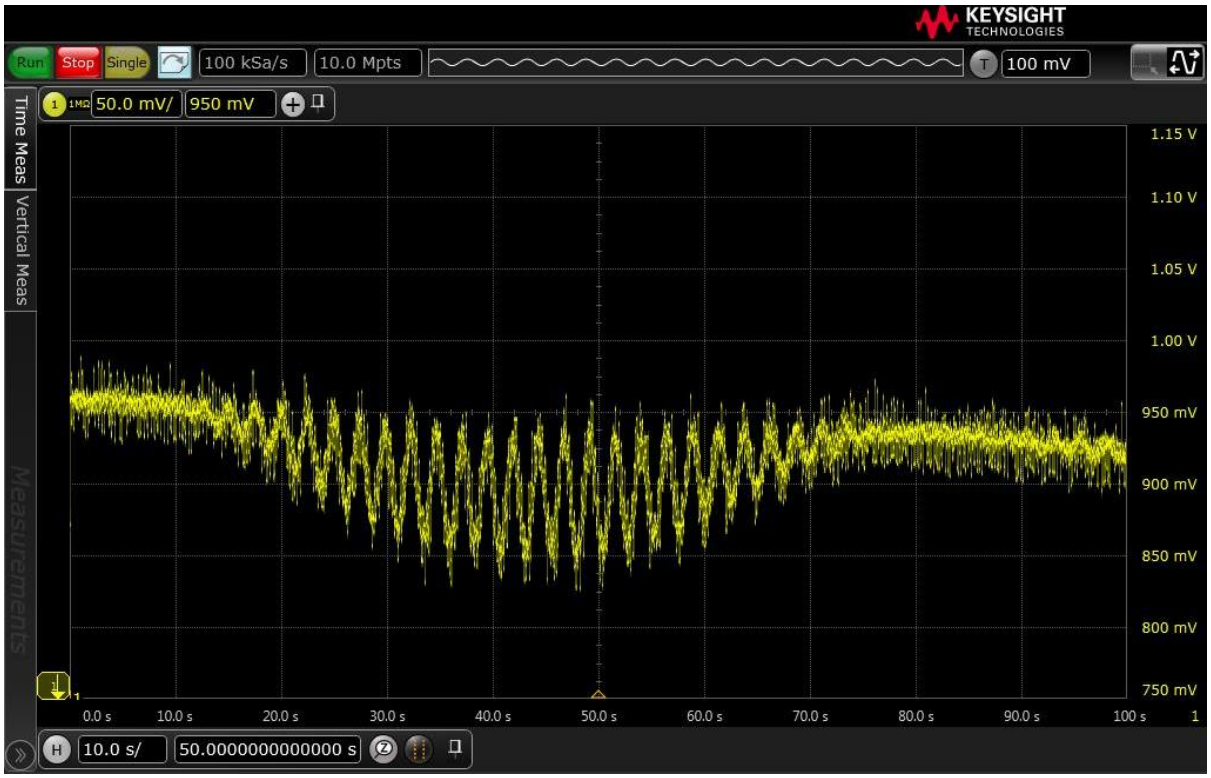

Figure 106. Optical output plotted against increasing MW frequency: 4.8ms free-precessing time. Horizontal scale is 1kHz / division.

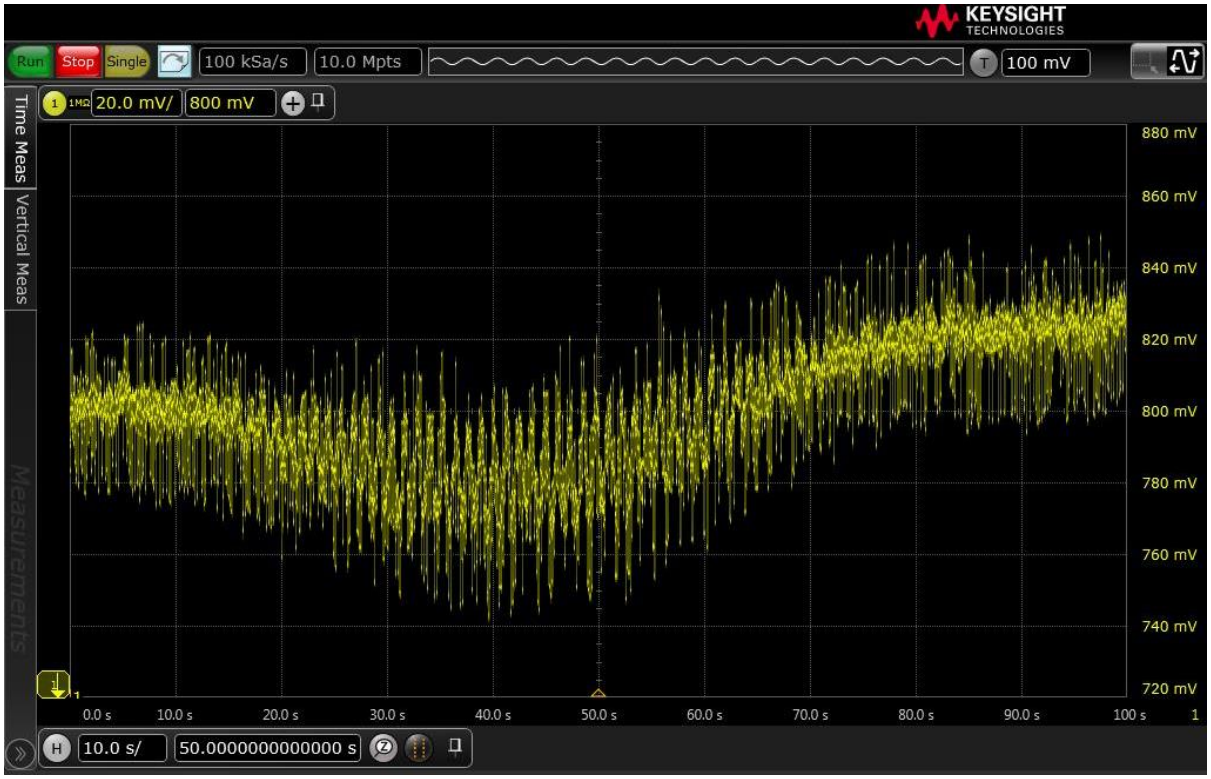

Figure 107. Optical output plotted against increasing MW frequency: 10ms free-precessing time. Horizontal scale is 1kHz / division.

Figure 107 is included to show the system response when pushed beyond the 7.2ms coherence relaxation time calculated in Section 2.3.4. As the amplitude of the response diminishes, the SNR has degraded to the point where the central fringe cannot be reliably distinguished using the 'onboard' ADC. A Ramsey time of 5.9ms results in a central fringe with 100Hz FWHM. Until a lower-noise ADC can be used, this is the limit of the prototype.

5.10 Summary

This work shows that the Physics Package is capable of supporting POP but that reduced sampling noise, and magnetic screening are required. The laser and MW generator were not subject to closed-loop control, which is the next step in Chapter 6.

6 Control Loops

6.1 Overview

This chapter covers the feedback loops used to stabilise the laser wavelength and MW frequency. The feedback paths are shown in blue in Figure 108, superimposed onto a diagram of the whole clock.

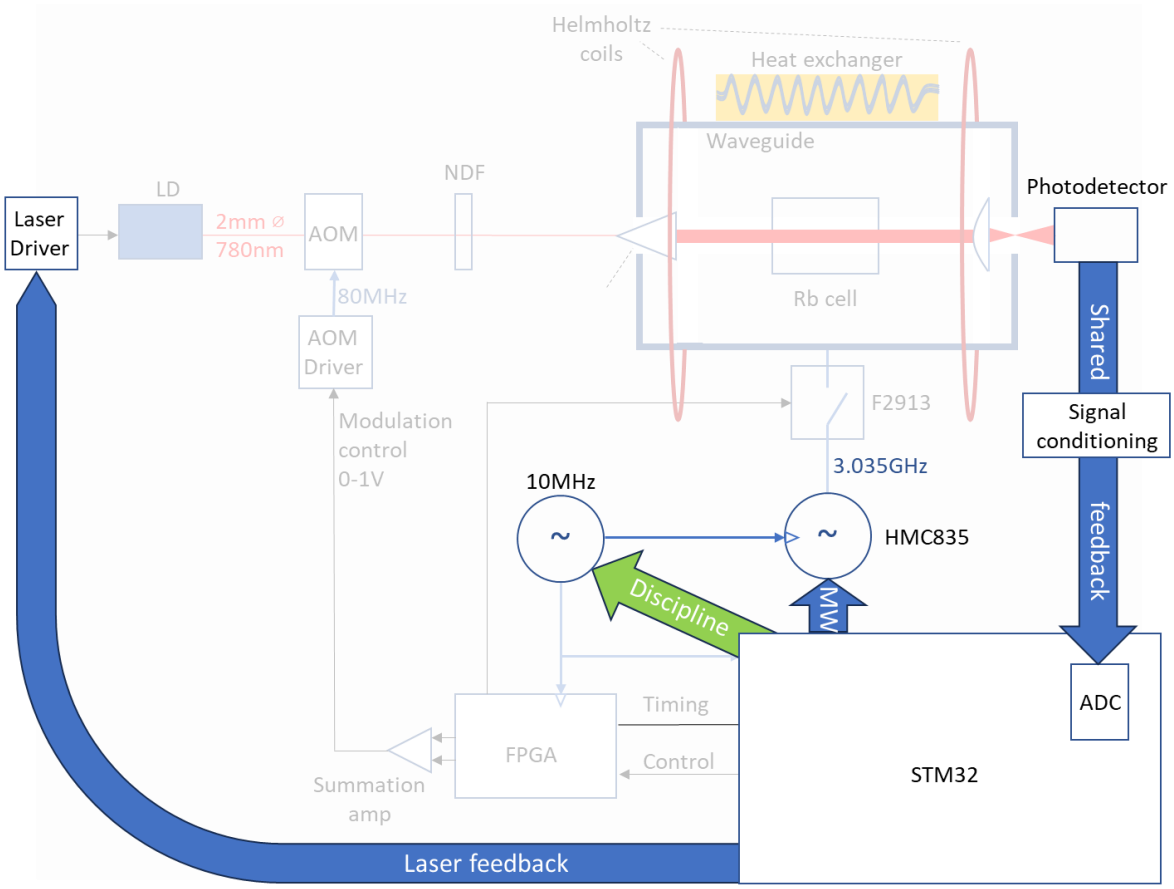

Figure 108. Laser and MW feedback loops superimposed onto diagram of the whole clock.

6.2 Laser stability

All of the experiments described so far have been conducted with the laser 'open-loop', i.e. with the driver stabilising the diode temperature and current, but without any other means of compensating for a change in the laser's output frequency. Schmidt-Eberle [96] provides a number of common contributors to changing output, along with the timescales on which each factor might affect the laser. I have compiled these in Table 13 .

Table 13. Sources of laser noise and drift.

Source	**Timescale /s**
Laser diode current noise	1µ - 50m
Acoustics & vibrations	10µ - 100m
Temperature control	100m – 50
Air pressure and humidity	20 – 200k
Environmental temperature	20 – 200k
Piezo drifts	50 – 200k

Layered vibration countermeasures have been employed: the Newport laser mount used was fitted with four anti-vibration feet; the laser safety enclosure is built on a rigid aluminium base of mass 14.5kg and was also fitted with anti-vibration feet; and this was mounted on a very substantial optical bench.

Diode temperature was maintained to an accuracy of ±0.001K, representing a frequency change of 59.2MHz. However, it was noted that the greatest laser frequency drift is during the first 30 minutes of operation: in this time the frequency drops by roughly 4GHz after initial power-on. The bandwidth needed to stabilise laser drift is therefore well below 1Hz.

6.3 Laser lock

A number of comparable clocks have taken the approach of using a 2nd Rb cell to stabilise the laser frequency [51], [80], [82]: Gozzelino et al. have recently published an example [97, Fig. 2], whose experimental setup I have reproduced as Figure 109. The block marked as SAS is a Saturated Absorption Spectroscopy module and contains its own Rubidium cell and photodiode, the feedback from which is used to stabilise the laser. This has the advantage of allowing continuous tuning of the laser frequency, but at the cost of an additional cell and photodiode. It does not compensate for the video-frequency shift through the AOM.

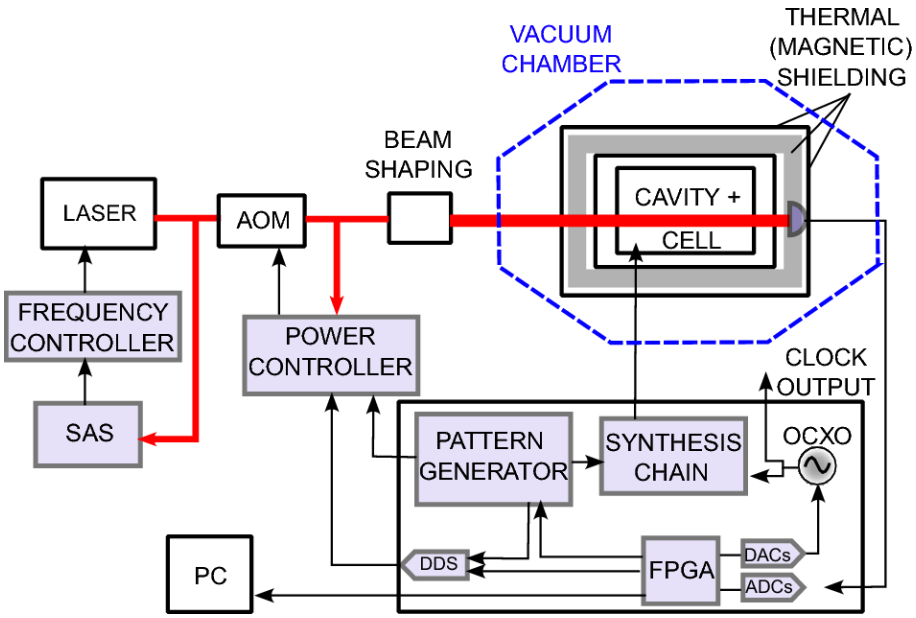

Figure 109. Experimental setup from Gozzelino et al. [97, Fig. 2].

As a hardware simplification, a single cell is used to maintain both laser and MW frequency locks. As each POP cycle is a discrete event, it is possible to alternate between cycles of POP and performing absorption spectroscopy. One full cycle is shown in Figure 110.

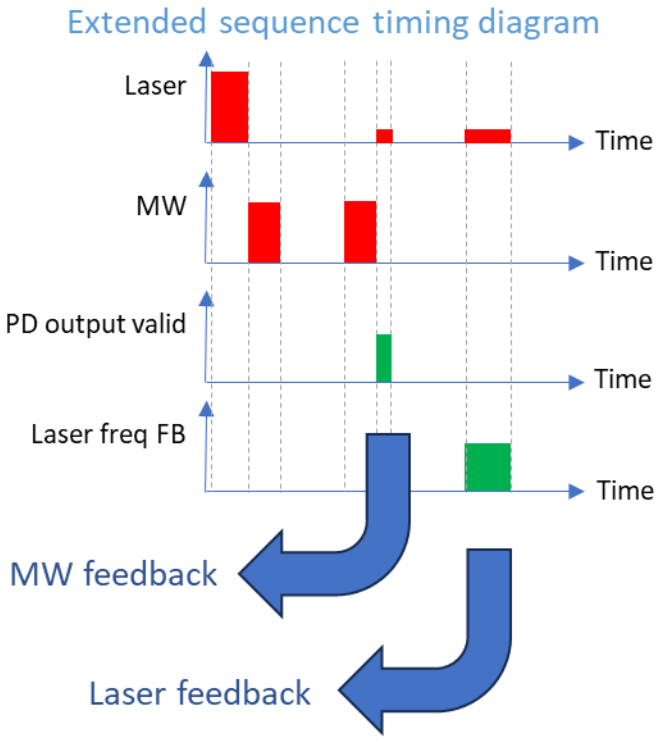

Figure 110. Extended timing sequence

For this to work effectively, there must not be a coupling mechanism between the MW and laser frequency feedback loops, as this would undermine a key benefit of POP. To avoid spectroscopy being influenced by the last π/2 pulse it is either necessary to wait for longer than the cell's relaxation time, or to continuously sample until spectroscopy reaches a steady-state value. Micalazio et al. note a

'memory effect' between POP cycles that can be particularly pronounced for low laser powers, or when the pump pulse is short [98]. By ensuring that the POP cycles are well separated, or that the (power x time) 'area' of the pump pulse is high, it is possible to greatly minimise the 'flattening' of the central Ramsey fringe, and broadened resonance associated with atomic preparation from the last POP cycle. This effect will be further mitigated if a POP cycle commences after CW spectroscopy: e.g. if the laser is tuned to the F=3 ground state, this state will already be partly de-populated, and the F=1, F=2 and F=4 ground states will have had significant time for relaxation since the last POP cycle.

The chosen mechanism for communicating with the laser is via the LDC501 laser driver's built in telnet server. This interface provides access to all functions available in the driver, including checking the enclosure interlock status, starting the diode's Thermo Electric Cooler (TEC), setting the diode current (with 10µA resolution), and querying the status of anything the driver controls. This is accessible via a 100Mb Ethernet link which has been configured with a static IP and connected directly to the STM32 onboard Ethernet. At the time of writing, a telnet client and LDC501 driver has been implemented in C and, although functional, is not robust, generates LwIP (an open-source TCP/IP stack) assertion failures and requires debugging.

The laser driver also provides an analogue input for modulating the laser current, which could be used for laser lock, but does not provide access to any other functionality.

Sampling one point on each side of the centre frequency has been chosen as the most efficient mechanism based on the relatively symmetrical nature of each F=2/3 dip as seen in the left hand side of Figure 111.

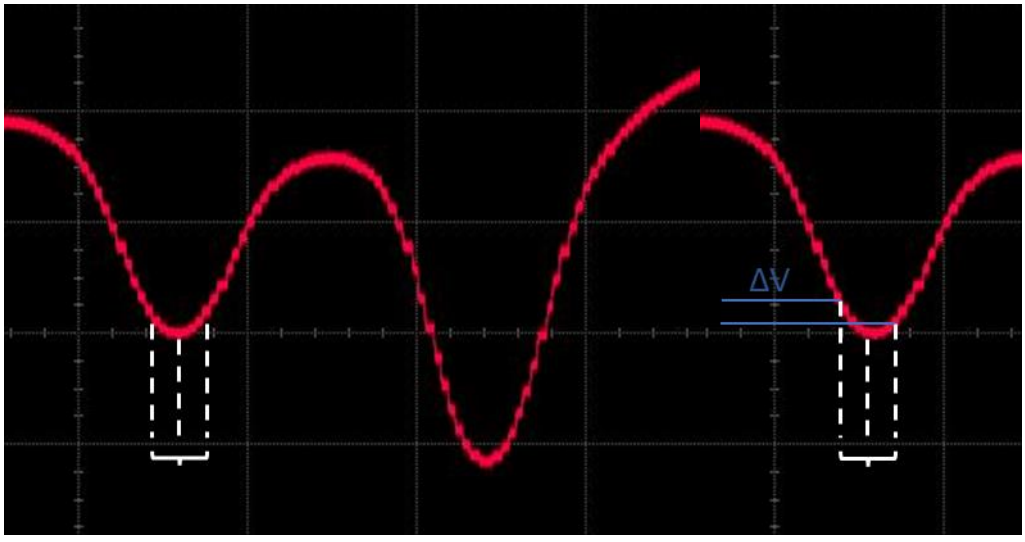

Figure 111. Rb absorption dips showing symmetrical polling around F=2 absorption dip; on frequency (left) and after drift (right).

After initial tuning to the optical minimum associated with the F=2 or F=3 dip, laser current is incremented by 50µA, the optical output sampled, then current reduced by 100µA to measure to optical output on the other side of the dip. If, as per the left hand side of Figure 111, the two readings

are comparable then the current is returned to its original value for the start of the next POP cycle. The right-hand side of Figure 111 shows a case where the laser has drifted and there is a measurable voltage difference. The original current is decremented by 10µA (the smallest current change supported by the laser driver) to compensate for the drift and the next POP cycle commences.

The 100µA current change for laser tuning constitutes a 0.06% change in laser diode power and does not appear to upset the thermal stability of the diode and TEC. Lastly, due to the relatively high noise figure associated with the STM32 onboard ADC, the output voltage is sampled 4 times at each operating point, greatly reducing Gaussian noise and effectively providing one extra bit of ADC resolution. Swapping this onboard DAC for a much lower noise solution, e.g. an external 16-bit device, could provide lower noise measurements, increasing the SNR, and shortening sample time by removing the need for oversampling.

6.4 MW feedback loop

As per Section 2.2, the central Ramsey fringe can be expressed as a $(\text{sinc}(x))^2$ function, and is therefore symmetrical as shown in Figure 112, taken from Gozzelino et al. [97].

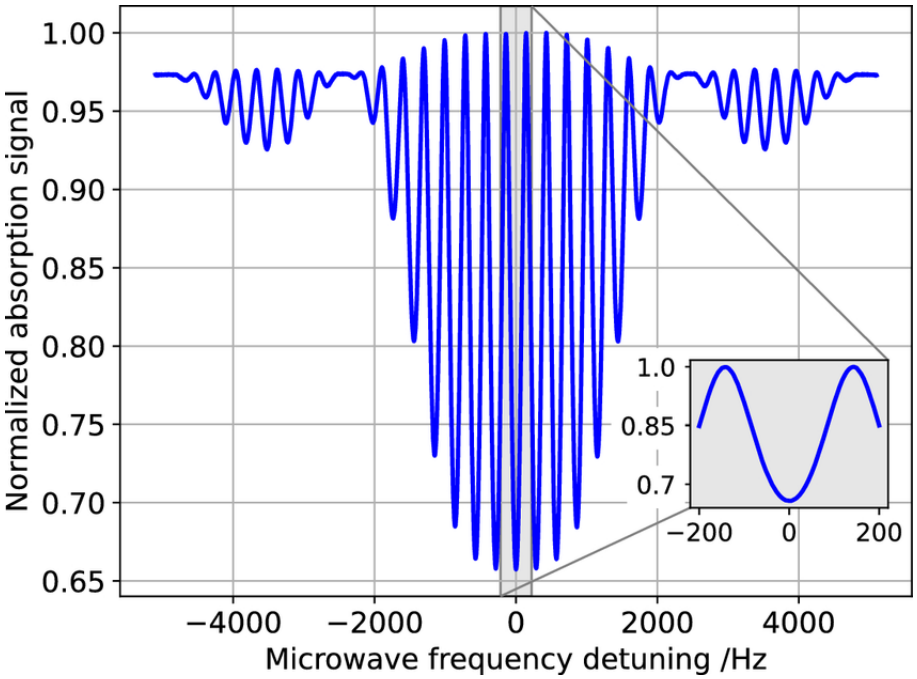

Figure 112. Typical experimental Ramsey fringes measured with the engineered physics package [97, Fig. 3].

This allows the MW response to be treated as per the absorption dips shown in Figure 111, but with MW frequency as the x-axis instead of laser current/frequency. The feedback loop uses the same 2-point polling mechanism (as in Section 6.3) which steps the MW frequency in multiples of 2.98Hz (set by the HMC835 PLL's VCO control resolution). By setting the frequency 23.8 Hz (i.e. 8 steps) below the central frequency, then measuring again 23.8 Hz above, the comparison can be performed in a region with steep gradient, maximising sensitivity to changing frequency (calculated as 1.00 mV/Hz). The

onboard DAC resolution of 0.3mV should therefore be able to maintain a frequency lock with an accuracy of 0.3Hz but in practice this will be greater due to high levels of noise. A low-noise, 16-bit, external DAC should be able to lock to ±46 mHz.

7 Summary

7.1 Successful outcomes

This thesis has shown it possible to perform POP using a travelling wave for MW interrogation. Using a waveguide as the MW enclosure effectively removes any influence due to cavity pulling. The use of a single cell to perform both POP and stabilise laser lock is both novel and a hardware simplification that should aid construction of a product based on this design. A central Ramsey fringe with FWHM of 100Hz has been established, which should allow for a frequency lock accurate to 0.3Hz with the current apparatus, and with a route forward to ±46 mHz accuracy.

7.2 Achievable product size

As per the initial goals, time was not spent optimising the size of the prototype, however the potential volume of a product has been calculated from the sum of the parts, as shown in Figure 113.

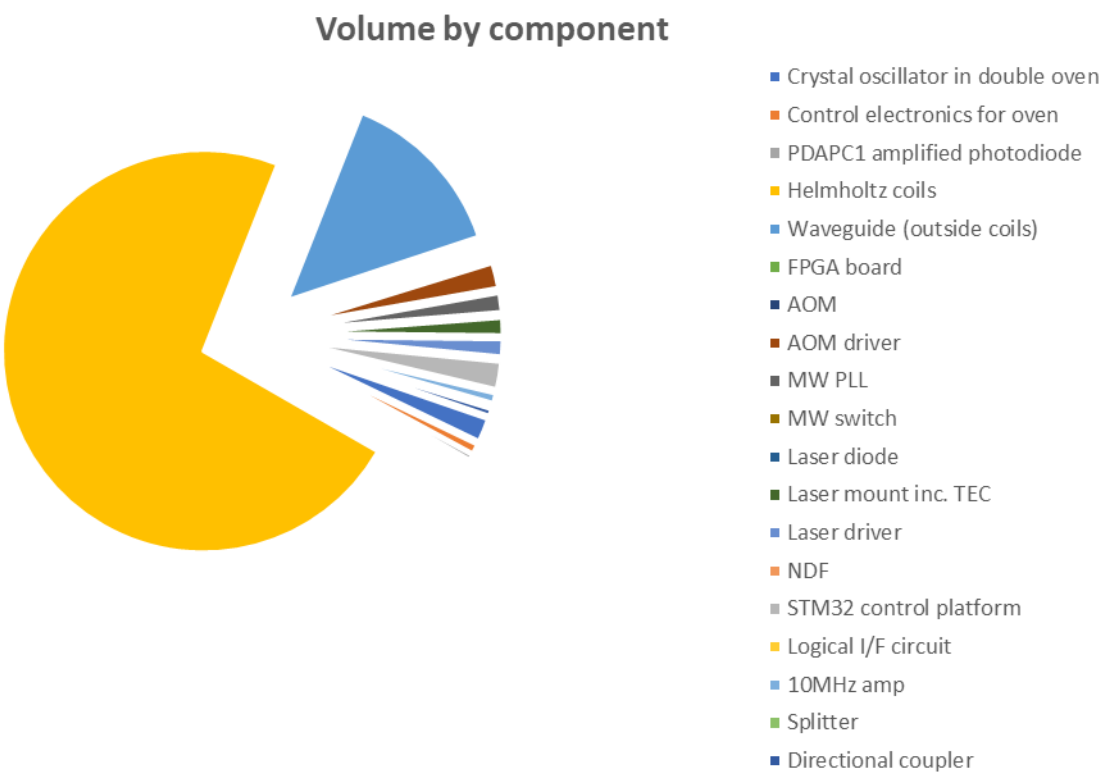

Figure 113. Potential volume by component.

Using a nominal figure of 0.09 l to represent production laser driver electronics, this adds up to a total volume of 6.4 l, of which 73% is the volume enclosed by the Helmholtz coils. Overall length is constrained by the optical path, chiefly the need for at least 110mm separation between the AOM and the waveguide aperture, to discriminate between M=0 and M=1 order beams. If this proves

problematic, then a mirror can be used to fold the path on itself to allow for more convenient packaging.

The height is constrained by the 214mm diameter of the Helmholtz coils, which is determined by the width of the waveguide. A reduction of 35% would be possible without significant change to the TE01 mode, though this would either require a shorter cell, or prohibit the inclusion of the beam expander inside the waveguide which would reduce the Signal-to-noise ratio (SNR). The 333mm length of the current waveguide could be reduced by at least one wavelength (λ_g = 113.5mm) without significant risk of exposing the cell to near-field effects close to the injectors.

7.2.1 Possible volume with ^{87}Rb

As ^{87}Rb has a much higher hyperfine frequency, a physics package constructed around this isotope could be much more compact. The waveguide could be scaled to have internal dimensions of 151.1mm x 48.3mm x 22.9mm and use a slightly smaller cell than at present. The Helmholtz coils would have a radius of 51mm, enclosing a corresponding volume of 1 l. The entire clock could then have a volume of less than 2l.

7.3 Next steps

This thesis has been able to show POP, but not demonstrate a performant clock in the time available. To bring this work to a rewarding conclusion would require:

1. A functional laser lock: either by debugging the Telnet connectivity between the STM32 and laser driver, or using the laser driver's analogue modulation input.
2. Migration of the high-resolution DAC drivers to the STM32 platform, so that it can be used to discipline the OCXO.
3. Measurement of the clock's performance to allow the OCXO 'disciplining' bandwidth to be optimised.

Numerous improvements are possible, and in roughly priority order:

1. Install magnetic screening around the Physics Package. At its greatest, the Earth's magnetic field varies between 0.25 and 0.65 Gauss which will pull the Hyperfine frequency by between 1.3kHz and 3.4kHz as per the 2nd order Zeeman shift calculations.
2. Swap onboard ADC for 16-bit external ADC with SPI interface. This must support external triggering and would ideally provide a line out to signal when conversion was complete. Doing this would provide lower noise measurements, increasing the SNR, and shortening sample time by removing the need for oversampling.

3. Place a temperature sensor on the waveguide and use that to regulate the water bath temperature.
4. Use a 795nm laser to compare optical contrast with the D2 transition. This will require either a 795nm collimating lens, or a small correcting lens in series with the current collimator.
5. Better electrical screening, shorter connecting leads between equipment, and consolidation of functionality onto the fewest PCBs possible.

7.4 Reflections

Development of this prototype has required numerous different skills as well as the theoretical background for atomic clocks. Refreshing and learning these skills has been fascinating, though time consuming, and development of this prototype could readily be done efficiently by a small team. The PhD has been a rich learning experience developing my skills in:

- 3D CAD (Autodesk Inventor)
- Circuit simulation (PathWave Design ADS and SIMetrix Elements)
- Schematic capture and layout (Proteus Design Suite)
- RTL, verification, and synthesis for FPGA (Verilog, Lattice Diamond)
- Programming and debug on an embedded platform (C, AtmelStart, Atmel Studio, STM32CubeIDE, ASF4 Software Framework, MicroChip Advanced Software Framework, STM32 Standard Peripheral Libraries, LwIP)
- Wave theory and MW Simulation (Simulia CST Studio Suite)
- Mathematical simulation (Wolfram Mathematica & MathWorks MATLAB)
- Use of lab equipment, each with its own I/F and software (Signal Generators, Vector Network Analysers, Spectrum Analysers [MW and optical], Phase Noise and Allan Deviation Testers, Power Meters [MW and optical], Oscilloscopes, Power Supplies, Laser Drivers, Lock-in Amplifiers)
- High (video and MW) frequency electronics
- Lasers and optical benches
- Quantum mechanics
- Atomic clocks

Bibliography

[1] R. Descartes, *Renati Descartes Principia philosophiae*. apud [et per] Ludovicum Elzevirium, 1644.

[2] A. Janiak, 'Space, atoms and mathematical divisibility in Newton', *Stud. Hist. Philos. Sci.*, vol. 31, no. 2, pp. 203–230, Jun. 2000.

[3] M. A. Lombardi, 'The evolution of time measurement, part 3: atomic clocks [Recalibration]', *IEEE Instrum. Meas. Mag.*, vol. 14, no. 6, pp. 46–49, Dec. 2011.

[4] I. I. Rabi, J. R. Zacharias, S. Millman, and P. Kusch, 'A New Method of Measuring Nuclear Magnetic Moment', *Phys. Rev.*, vol. 53, no. 4, pp. 318–318, Feb. 1938.

[5] D. B. Sullivan, 'Time and frequency measurement at NIST: the first 100 years', in *Proceedings of the 2001 IEEE International Frequncy Control Symposium and PDA Exhibition (Cat. No.01CH37218)*, 2001, pp. 4–17.

[6] L. Essen and J. V. L. Parry, 'An Atomic Standard of Frequency and Time Interval: A Cæsium Resonator', *Nature*, vol. 176, p. 280, Aug. 1955.

[7] W. Markowitz, R. G. Hall, L. Essen, and J. V. Parry, 'Frequency of Cesium in Terms of Ephemeris Time', *Phys. Rev. Lett.*, vol. 1, pp. 105–107, Aug. 1958.

[8] 'Definition of the second', *BIPM*. [Online]. Available: https://www.bipm.org/en/publications/si-brochure/second.html. [Accessed: 23-Jul-2018].

[9] G. Busca, L.-G. Bernier, and P. Rochat, 'Active hydrogen maser atomic frequency standard', 5838206A, 17-Nov-1998.

[10] M. A. Humphrey, 'Precision measurements with atomic hydrogen masers', PhD, Harvard University, 2003.

[11] T. Burtichelov, 'Developing an Ultra-Low Phase Noise Atomic Frequency Reference using Coherent Population Trapping (CPT) in Rubidium', PhD, University of York, 2017.

[12] M. Arditi, 'Gas Cell "Atomic Clocks" Using Buffer Gases and Optical Orientation', in *12th Annual Symposium on Frequency Control*, 1958, pp. 606–622.

[13] C. Cohen-Tannoudji and A. Kastler, 'Optical Pumping', in *Progress in Optics*, vol. 5, E. Wolf, Ed. Elsevier, 1966, pp. 1–81.

[14] W. Happer, 'Optical Pumping', *Reviews of Modern Physics*, vol. 44, no. April 1972, Apr. 1972.

[15] E. Arimondo, M. Inguscio, and P. Violino, 'Experimental determinations of the hyperfine structure in the alkali atoms', *Rev. Mod. Phys.*, vol. 49, no. 1, pp. 31–75, Jan. 1977.

[16] 'PRS10 Rubidium Frequency Standard Operation and Service Manual', Stanford Research Systems.

[17] D. Castelvecchi, '100 Years Ago, a Quantum Experiment Explained Why We Don't Fall through Our Chairs', *SA Space & Physics*, vol. 5, Apr-2022.

[18] D. B. Sullivan *et al.*, 'Primary Atomic Frequency Standards at NIST', *J. Res. Natl. Inst. Stand. Technol.*, vol. 106, no. 1, pp. 47–63, Feb. 2001.

[19] R. E. Beehler, R. C. Mockler, and J. M. Richardson, 'Cesium Beam Atomic Time and Frequency Standards', *Metrologia*, vol. 1, no. 3, p. 114, Jul. 1965.

[20] B. J. Dalton and P. L. Knight, 'The effects of laser field fluctuations on coherent population trapping', *J. Phys. B At. Mol. Opt. Phys.*, vol. 15, no. 21, p. 3997, 1982.

[21] The Editors of Encyclopædia Britannica, 'D-lines', *Encyclopædia Britannica*. Encyclopædia Britannica, inc., 20-Jul-1998.

[22] E. Arimondo, 'Coherent Population Trapping in Laser Spectroscopy', in *Progress in Optics*, vol. 35, E. Wolf, Ed. Elsevier, 1996, pp. 257–354.

[23] S. Goka, '^{85}Rb D1-line CPT atomic clock with low power consumption', in *2009 IEEE International Frequency Control Symposium Joint with the 22nd European Frequency and Time forum*, 2009, pp. 578–581.

[24] M. Arditi and T. R. Carver, 'Pressure, Light, and Temperature Shifts in Optical Detection of 0-0 Hyperfine Resonance of Alkali Metals', *Phys. Rev.*, vol. 124, no. 3, pp. 800–809, Nov. 1961.

[25] N. F. Ramsey, 'A New Molecular Beam Resonance Method', *Phys. Rev.*, vol. 76, no. 7, pp. 996–996, Oct. 1949.

[26] J. Guéna *et al.*, 'Demonstration of a dual alkali Rb/Cs fountain clock', *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 57, no. 3, pp. 647–653, Mar. 2010.

[27] R. Wynands and S. Weyers, 'Atomic fountain clocks', *Metrologia*, vol. 42, no. 3, p. S64, Jun. 2005.

[28] G. J. Dick, J. D. Prestage, C. A. Greenhall, and L. Maleki, 'Local oscillator introduced degradation of medium-term stability in passive atomic frequency standards', in *Proceedings of the 22th Annual Precise Time and Time Interval Systems and Applications Meeting December 4 - 6, 1990 Sheraton Premiere Hotel Vienna, Virginia*, Vienna, Virginia, 1990, pp. 487–508.

[29] F. G. Ascarrunz *et al.*, 'A Portable Cold 87Rb Atomic Clock with Frequency Instability at One Day in the 10−15Range', in *2018 IEEE International Frequency Control Symposium (IFCS)*, 2018, pp. 1–3.

[30] 'SpectraDynamics cRb- Clock - a portable cold Rubidium microwave atomic clock', *SDI*. [Online]. Available: https://spectradynamics.com/products/crb-clock/. [Accessed: 13-Aug-2020].

[31] 'Portable Microwave Cold Atomic Clock - Dimensions'. [Online]. Available: https://app.dimensions.ai/details/grant/grant.4382146. [Accessed: 21-May-2024].

[32] S. Micalizio, C. E. Calosso, A. Godone, and F. Levi, 'Metrological characterization of the pulsed Rb clock with optical detection', *Metrologia*, vol. 49, no. 4, p. 425, May 2012.

[33] T. Rosenband *et al.*, 'Frequency ratio of Al+ and Hg+ single-ion optical clocks; metrology at the 17th decimal place', *Science*, vol. 319, no. 5871, pp. 1808–1812, Mar. 2008.

[34] D. J. Jones *et al.*, 'Carrier-envelope phase control of femtosecond mode-locked lasers and direct optical frequency synthesis', *Science*, vol. 288, no. 5466, pp. 635–640, Apr. 2000.

[35] M. Vengalattore, 'Viewpoint: A Boost in Precision for Optical Atomic Clocks', *Physics*, vol. 11, no. 22, 05-Mar-2018.

[36] N. Poli, C. W. Oates, P. Gill, and G. M. Tino, 'Optical atomic clocks', *arXiv [physics.atom-ph]*, 10-Jan-2014.

[37] N. Ohtsubo, Y. Li, K. Matsubara, T. Ido, and K. Hayasaka, 'Frequency measurement of the clock transition of an indium ion sympathetically-cooled in a linear trap', *Opt. Express*, vol. 25, no. 10, pp. 11725–11735, May 2017.

[38] S. M. Brewer *et al.*, 'Measurements ofAl+27andMg+25magnetic constants for improved ion-clock accuracy', *Phys. Rev. A*, vol. 100, no. 1, Jul. 2019.

[39] M. Piotrowski, J. Scarabel, M. Lobino, E. Streed, and S. Gensemer, 'Studies of thorium and ytterbium ion trap loading from laser ablation for gravity monitoring with nuclear clocks', *OSA Contin.*, vol. 3, no. 8, p. 2210, Aug. 2020.

[40] J. Yao *et al.*, 'Optical-Clock-Based Time Scale', *Phys Rev Appl*, vol. 12, no. 4, 2019.

[41] M. Gellesch *et al.*, 'Transportable optical atomic clocks for use in out-of-the-lab environments', *Advanced Optical Technologies*, vol. 9, no. 5, pp. 313–325, Oct. 2020.

[42] M. Aldous *et al.*, 'Route to a Portable Optical Clock', in *2018 IEEE International Frequency Control Symposium (IFCS)*, 2018, pp. 1–3.

[43] D. W. Allan, 'Statistics of atomic frequency standards', *Proc. IEEE*, vol. 54, no. 2, pp. 221–230, Feb. 1966.

[44] '1139-1999 - IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology-Random Instabilities', *IEEE Std 1139-1999*. IEEE, pp. 1–40, 21-Jul-1999.

[45] M. Weiss and K. Shenoi, 'TDEV Then and Now', in *ISTF 2015*, Edinburgh, 2015.

[46] T. Mai, 'Technology readiness level', *NASA*, 2015. [Online]. Available: http://www.nasa.gov/directorates/heo/scan/engineering/technology/txt_accordion1.html. [Accessed: 11-Sep-2018].

[47] J. Everard, T. Burtichelov, and K. Ng, 'Ultralow Phase Noise 10-MHz Crystal Oscillators', *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 66, no. 1, pp. 181–191, Jan. 2019.

[48] D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls, 'Characterisation of Clocks and Oscillators, NIST technical note 1337', NIST, Mar. 1990.

[49] 'Allan Variance', *Iowa State University*. [Online]. Available: http://home.engineering.iastate.edu/~shermanp/AERE432/lectures/Rate%20Gyros/Allan%20variance.pdf. [Accessed: 11-Sep-2018].

[50] J. Kitching, 'Chip-scale atomic devices', *Applied Physics Reviews*, vol. 5, no. 3, p. 031302, Sep. 2018.

[51] M. Gharavipour, 'Ramsey Spectroscopy in a Rubidium Vapor Cell and Realization of an Ultra-Stable Atomic Clock', PhD, UNIVERSITÉ DE NEUCHÂTEL, 2018.

[52] R. M. Eisberg and R. Resnick, *Quantum Physics of Atoms, Molecules, Solids, Nuclei, and Particles*. Wiley, 1985.

[53] R. C. Pooser, A. M. Marino, V. Boyer, K. M. Jones, and P. D. Lett, 'Quantum correlated light beams from non-degenerate four-wave mixing in an atomic vapor: the D1 and D2 lines of 85Rb and 87Rb', *Opt. Express*, vol. 17, no. 19, pp. 16722–16730, Sep. 2009.

[54] W. H. Wollaston, 'XII. A method of examining refractive and dispersive powers, by prismatic reflection', *Philosophical Transactions of the Royal Society of London*, vol. 92, pp. 365–380, Jan. 1997.

[55] D. A. Steck, 'Rubidium 85 D Line Data', no. 2.14. 23-Dec-2010.

[56] P. Zeeman, 'XXXIII. Doublets and triplets in the spectrum produced by external magnetic forces.—(II.)', *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 44, no. 268, pp. 255–259, Sep. 1897.

[57] R. Shankar, *Principles of Quantum Mechanics*. Springer US, 1995.

[58] T. Burtichelov, 'Design and Construction of a Rubidium Atomic Clock', University of York, Jun. 2012.

[59] T. Zanon, S. Guerandel, E. de Clercq, D. Holleville, N. Dimarcq, and A. Clairon, 'High contrast Ramsey fringes with coherent-population-trapping pulses in a double lambda atomic system', *Phys. Rev. Lett.*, vol. 94, no. 19, p. 193002, May 2005.

[60] B. Gao, 'Effects of Zeeman degeneracy on the steady-state properties of an atom interacting with a near-resonant laser field: Analytic results', *Phys. Rev. A*, vol. 48, no. 3, pp. 2443–2448, Sep. 1993.

[61] N. F. Ramsey, 'Experiments with separated oscillatory fields and hydrogen masers', *Rev. Mod. Phys.*, vol. 63, no. 3, pp. 541–552, Jul. 1990.

[62] N. F. Ramsey, 'A molecular beam resonance method with separated oscillating fields', *Physical Review*, vol. 78, no. 6, Jun. 1950.

[63] T. Scholtes, S. Woetzel, R. IJsselsteijn, V. Schultze, and H.-G. Meyer, 'Intrinsic relaxation rates of polarized Cs vapor in miniaturized cells', *Appl. Phys. B*, vol. 117, no. 1, pp. 211–218, Oct. 2014.

[64] T. Lindvall, M. Merimaa, I. Tittonen, and E. Ikonen, 'ALL-OPTICAL ATOMIC CLOCK BASED ON DARK STATES OF 85Rb', in *Frequency Standards and Metrology*, WORLD SCIENTIFIC, 2002, pp. 183–190.

[65] J. Vanier and C. Audoin, *The Quantum Physics of Atomic Frequency Standards*. A. Hilger, 1989.

[66] J. Vanier, J.-F. Simard, and J.-S. Boulanger, 'Relaxation and frequency shifts in the ground state of Rb85', *Phys. Rev. A Gen. Phys.*, vol. 9, no. 3, pp. 1031–1040, Mar. 1974.

[67] Y. Ito, D. Sato, K. Kamada, and T. Kobayashi, 'Optimal densities of alkali metal atoms in an optically pumped K-Rb hybrid atomic magnetometer considering the spatial distribution of spin polarization', *Opt. Express*, vol. 24, no. 14, pp. 15391–15402, Jul. 2016.

[68] K. Ishikawa and T. Yabuzaki, 'Diffusion coefficient and sublevel coherence of Rb atoms in N2 buffer gas', *Phys. Rev. A*, vol. 62, no. 6, p. 065401, Nov. 2000.

[69] O. Firstenberg, M. Shuker, A. Ben-Kish, D. R. Fredkin, N. Davidson, and A. Ron, 'Theory of Dicke narrowing in coherent population trapping', *arXiv [quant-ph]*, 02-Jan-2007.

[70] S. Micalizio, A. Godone, F. Levi, and J. Vanier, 'Spin-exchange frequency shift in alkali-metal-vapor cell frequency standards', *Phys. Rev. A*, vol. 73, no. 3, p. 033414, Mar. 2006.

[71] J.-H. Liu *et al.*, 'The polarization and the fundamental sensitivity of 39K (133Cs)-85Rb-4He hybrid optical pumping spin exchange relaxation free atomic magnetometers', *Sci. Rep.*, vol. 7, no. 1, p. 6776, Jul. 2017.

[72] H. Shang *et al.*, 'Design of a Measuring Device and Experimental Study into the Relationship between Temperature and the Density of Alkali Metal-Vapor', *Photonics*, vol. 10, no. 2, p. 112, Jan. 2023.

[73] C. B. Alcock, V. P. Itkin, and M. K. Horrigan, 'Vapour Pressure Equations for the Metallic Elements: 298–2500K', *Can. Metall. Q.*, vol. 23, no. 3, pp. 309–313, Jul. 1984.

[74] F. Levi, C. E. Calosso, A. Godone, and S. Micalizio, 'Pulsed optically pumped Rb clock: A high stability vapor cell frequency standard', in *2013 Joint European Frequency and Time Forum International Frequency Control Symposium (EFTF/IFC)*, 2013, pp. 599–605.

[75] R. Li, L. Zhou, J. Wang, and M. Zhan, 'Measurement of the quadratic Zeeman shift of 85Rb hyperfine sublevels using stimulated Raman transitions', *arXiv [physics.atom-ph]*, 01-Dec-2008.

[76] A. Godone, S. Micalizio, C. E. Calosso, and F. Levi, 'The pulsed rubidium clock', *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 53, no. 3, pp. 525–529, Mar. 2006.

[77] A. Godone, S. Micalizio, F. Levi, and C. Calosso, 'Microwave cavities for vapor cell frequency standards', *Rev. Sci. Instrum.*, vol. 82, no. 7, p. 074703, Jul. 2011.

[78] T. Burtichelov and J. Everard, 'Latest results in the development of an ultra-low phase noise Rb CPT vapour cell atomic clock', in *2017 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFC)*, 2017, pp. 837–841.

[79] S. Micalizio, A. Godone, C. Calosso, F. Levi, C. Affolderbach, and F. Gruet, 'Pulsed optically pumped rubidium clock with high frequency-stability performance', *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 59, no. 3, pp. 457–462, Mar. 2012.

[80] Q. Shen, H. Lin, J. Deng, and Y. Wang, 'Pulsed optically pumped atomic clock with a medium- to long-term frequency stability of 10−15', *Rev. Sci. Instrum.*, vol. 91, no. 4, p. 045114, Apr. 2020.

[81] Q. Su-Ping, Z. Yi, and G. Si-Hong, 'Dependence of the 85Rb coherent population trapping resonance characteristic on the pressure of N2 buffer gas', *Chin. Physics B*, vol. 22, no. 9, p. 099501, Sep. 2013.

[82] P. Yun *et al.*, 'High-Performance Coherent Population Trapping Clock with Polarization Modulation', *Phys. Rev. Applied*, vol. 7, no. 1, p. 014018, Jan. 2017.

[83] M. Stähler *et al.*, 'Coherent population trapping resonances in thermal (85)Rb vapor: D(1) versus D(2) line excitation', *Opt. Lett.*, vol. 27, no. 16, pp. 1472–1474, Aug. 2002.

[84] V. N. Baryshev, M. S. Aleynikov, G. V. Osipenko, and I. Yu. Blinov, 'Technique of pulsed optical pumping and pulsed excitation of microwave resonances using the Ramsey scheme in a 87Rb cell with a buffer gas', *Quantum Electron.*, vol. 48, no. 5, p. 443, May 2018.

[85] V. Baryshev, M. Aleynikov, G. Osipenko, V. Epikhin, I. Blinov, and S. Donchenko, 'Pulsed Optically Pumped Rb Clock: Laboratory Prototype at VNIIFTRI', in *Proc. of 2016 IEEE Int. Frequency Control Symposium*, Roosevelt Hotel, New Orleans, Louisiana, USA, pp. 205–208.

[86] M. A. Kasevich, E. Riis, S. Chu, and R. G. DeVoe, 'RF spectroscopy in an atomic fountain', *Phys. Rev. Lett.*, vol. 63, no. 6, pp. 612–615, Aug. 1989.

[87] D. M. Pozar, *Microwave Engineering*. John Wiley & Sons, 2011.

[88] S. Ramo, J. R. Whinnery, and T. Van Duzer, *Fields and waves in communication electronics*. John Wiley & Sons, 1965.

[89] H. Eklund, A. Roos, and S. T. Eng, 'Rotation of laser beam polarization in acousto-optic devices', *Opt. Quantum Electron.*, vol. 7, no. 2, pp. 73–79, Mar. 1975.

[90] R. Paschotta, 'Acousto-optic Modulators', *RP Photonics Consulting GmbH*, 05-Jul-2019. [Online]. Available: https://www.rp-photonics.com/acousto_optic_modulators.html. [Accessed: 19-Oct-2019].

[91] W. J. Schwenger and J. M. Higbie, 'High-speed acousto-optic shutter with no optical frequency shift', *Rev. Sci. Instrum.*, vol. 83, no. 8, p. 083110, Aug. 2012.

[92] M. Egan, 'The 20-Bit DAC Is the Easiest Part of a 1-ppm-Accurate Precision Voltage Source | Analog Devices', *Analog Dialogue*, vol. 44, no. 4, pp. 1–4, Apr-2010.

[93] R. Y. Richard Lyons, 'Reducing ADC Quantization Noise', *Microwaves & Radio Frequency*, 17-Jun-2005. [Online]. Available: http://www.mwrf.com/components/reducing-adc-quantization-noise. [Accessed: 16-Apr-2018].

[94] SiliconLabs, 'AN118 - Improving ADC resolution by oversampling and averaging'. Jul-2013.

[95] S. Micalizio, F. Levi, C. E. Calosso, M. Gozzelino, and A. Godone, 'A pulsed-Laser Rb atomic frequency standard for GNSS applications', *GPS Solutions*, vol. 25, no. 3, p. 94, Apr. 2021.

[96] S. Schmidt-Eberle, ‘Phase and frequency locking of diode lasers’, TOPTICA Photonics AG, Apr. 2021.

[97] M. Gozzelino *et al.*, ‘Realization of a pulsed optically pumped Rb clock with a frequency stability below [Formula: see text]’, *Sci. Rep.*, vol. 13, no. 1, p. 12974, Aug. 2023.

[98] S. Micalizio, C. E. Calosso, F. Levi, and A. Godone, ‘Ramsey-fringe shape in an alkali-metal vapor cell with buffer gas’, *Phys. Rev. A*, vol. 88, no. 3, p. 033401, Sep. 2013.

Glossary

Analogue to Digital Converter (ADC)

Allan Deviation (ADEV)

Acousto-Optic Modulator (AOM)

Anti-reflective (AR)

Allan Variance (AVAR)

Arbitrary Waveform Generator (AWG)

Bandwidth (B/W)

Coherent Population Trapping (CPT)

Continuous Wave (CW)

Digital to Analog Converter (DAC)

Dielectric Resonator Oscillator (DRO)

Direct Digital Synthesizer (DDS)

Distributed Bragg Reflector (DBR)

Double Resonance (DR)

Effective Number of Bits (ENOB)

Field Programmable Gate Array (FPGA)

Firmware (F/W)

Full Width at Half Maximum (FWHM)

Hardware (H/W)

Integrated Development Environment (IDE)

Interface (I/F)

Laser Diode (LD)

Local Oscillator (LO)

Low-pass Filter (LPF)

Lightweight TCP/IP stack (LwIP)

Microwave amplification by stimulated emission of radiation (Maser)

Micro-control Unit (MCU)

Microwave (MW)

Neutral Density Filter (NDF)

Near-infrared (NIR)

Negative Temperature Coefficient (NTC)

Noise Equivalent Power (NEP)

Oven Controlled Crystal Oscillator (OCXO)

Phase Locked Loop (PLL)

Pulsed Optical Pumping (POP)

Quarter Wave Plate (QWP)

Root Mean Squared (RMS)

Saturated Absorption Spectroscopy (SAS)

Side-Mode Suppression Ratio (SMSR)

Signal-to-Noise Ratio (SNR)

System on Chip (SoC)

Software (S/W)

Thermo Electric Cooler (TEC)

Technology Readiness Level (TRL)

Ultra Low Expansion (ULE)

Voltage Controlled Oscillator (VCO)

Vertical Cavity Surface Emitting Laser (VCSEL)

Vector Network Analyser (VNA)

Waveguide (WG)

Appendix A. Equipment

A list of the major components used to construct, and test, the atomic clock.

Test equipment

Anritsu MS2561B 3GHz spectrum analyser

Anritsu 37377C vector network analyser

Fluke 52 thermometer & NTC thermistor

GM08 Gaussmeter with transverse probe

HP 3561A signal analyser

HP 8563E spectrum analyser

HP 8648C 3.2GHz signal generator

HP 8720B network analyser

Keysight DSO-S 054A oscilloscope

Keysight E8257D 20GHz signal generator

PicoScope 4262 16-bit USB oscilloscope

Rhode & Swartz SMB100A signal generator

Symmetricom (Microsemi) TSC5120A-01

Tektronix AFG2021 & AFG1062 arbitrary waveform generators

Tektronix TDS1002 oscilloscope

Thorlabs CCS100 CCD spectrometer

Thorlabs PM100D optical power meter and S120C sensor

Clock and physics package

Analog Devices DC1531A current source

Analog Devices DC1766A-A clock buffer and CMOS converter

Analog (Hittite) HMC835LP6GE PLL & VCO development platform

Grant W6 1.5kW circulating water bath

Isomet 630C-80 AOM Driver

Isomet D804-P80L-2 Acousto-Optic Deflector (formerly 1205C-2-804B Modulator)

Minicircuits ZRL-3500+ MW amplifier

Newport LDM-4984 LD mount

Osram BPX61 photodiode

Photonics Technologies SC-RB85-(25x25-Q)-AR-AR(12.7T)-N2(7.3T) 85Rb cell

Renesas F2913EVBI absorptive microwave switch

Stanford Research Systems LDC-501 laser driver with O500C3 and O500C4 cables

Stanford Research Systems SR570 transimpedance amplifier

ST Microelectronics Nucleo STM32H743ZI2 MCU development board

Thorlabs DBR780PN laser diode (SN 46115)

Thorlabs F240APC-780 1.5mm dia FC/APC collimator

Thorlabs MB4590/M optical bench and 1” mounts

Thorlabs NExxA-B 25mm absorptive neutral density filters with AR coatings

Thorlabs PDAPC1 amplified photodiode

Thorlabs WPQSM05-780 zero-order quarter-wave plate

TI OPA699 voltage limiting amplifier on OPA-SO-1A evaluation platform

10mm internal diameter PVC hose

40x80x12.5mm aluminium heat exchangers

Waveguide based on 4”x1.75” aluminium extrusion, constructed by Lowton Engineering

Appendix B. Cylindrical cavities

Anyone seeking to construct a microwave cavity should read Godone et al. 2011 [77]. It's an excellent reference, and details everything involved in the design of a 6.83 GHz cavity for ^{87}Rb. This appendix covers the process of designing a cavity that supports TE_{011} at 3.035 GHz for ^{85}Rb, followed by a short section on how to characterise a single-port cavity that I've not seen elsewhere. All of this has been relegated to an appendix due to the decision to build the Physics Package around waveguide.

A TE_{01} cavity

The starting point for this PhD was a CPT clock, so no microwave enclosure existed and building a 3.035 GHz equivalent to Micalizio's very successful cavity was a logical first step. Following their example, dimensions were calculated for the smallest possible cavity that would support TE_{011} at 3.035 GHz. For ease of prototyping, the choice of diameter was restricted to European Standard EN 1057 pipe sizes and meant that 133mm Outer Diameter (OD) pipe would be suitable. A specified internal diameter of 130mm, and length of 131mm mean that the dominant modes as listed in Table 14 can be present in the cavity.

Table 14. Dominant modes supported in a '133mm' cavity

Mode	Frequency f/Hz	Multiple of f	k/ 1/m	Q
TM010	1,778,935,014	0.59	37.28	206,275
TE111	1,779,091,822	0.59	37.29	146,033
TM011	2,115,346,535	0.70	44.33	134,369
TE210	2,259,324,401	0.74	47.35	64,816
TM012	2,899,126,506	0.96	60.76	61,412
TM110	2,834,474,922	0.93	59.41	59,983
TE011	3,056,849,565	1.01	64.07	53,404
TE112	2,663,721,315	0.88	55.83	53,422
TM111	3,056,849,565	1.01	64.07	49,759

Figure 114 shows the same information graphically, highlighting both that TE_{011} is not the highest-Q mode possible, and that the cavity is capable of supporting other modes quite close in frequency, as well as the degenerate TM_{111}.

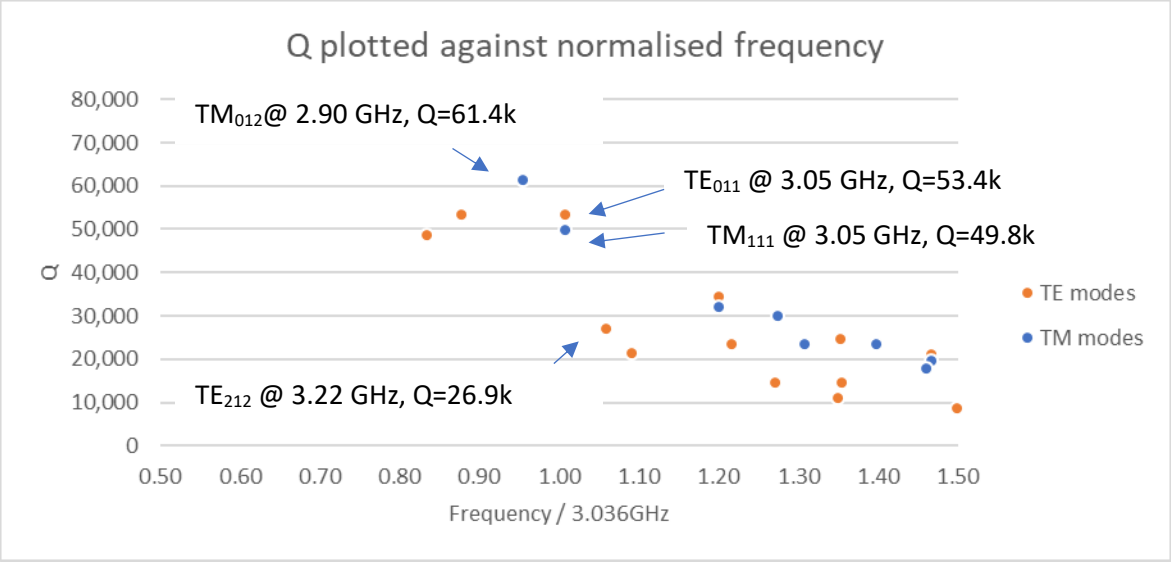

Figure 114. Resonant modes present in '133mm' cavity

It is possible to suppress TM_{111} by fitting a mode-choke – non-conducting slots in the end-walls which present the necessary electrical short circuit for TE_{011}, but at right-angles, and therefore an open circuit for TM_{111}, preventing that mode from propagating. However, it is also possible to use the position and directional qualities of a loop antenna to couple energy into modes whose h-field aligns with that of the antenna, at the site of the antenna. As per Figure 115, by inserting the probe mid-way up the side-wall, and intruding slightly into the cavity it is possible to couple efficiently to TM_{011} whilst discriminating against TM_{111}. Note that any misalignment of the loop antenna can couple energy weakly into other modes, and that any irregularity in the cavity (e.g. a vapour cell) can couple energy into other modes.

Figure 115. A view looking down into the '133mm' cavity with top removed. The loop probe is visible half-way up the sidewall.

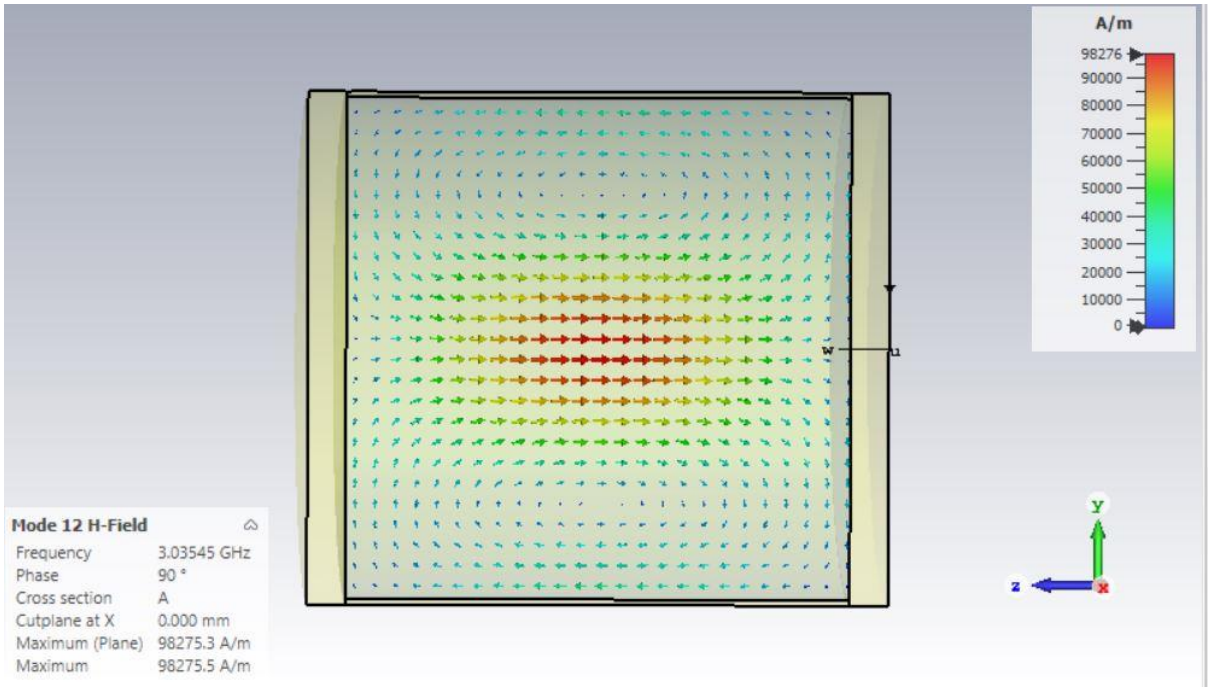

Figure 116. Cross-section of the 133mm cavity showing the simulated TE011 h-field

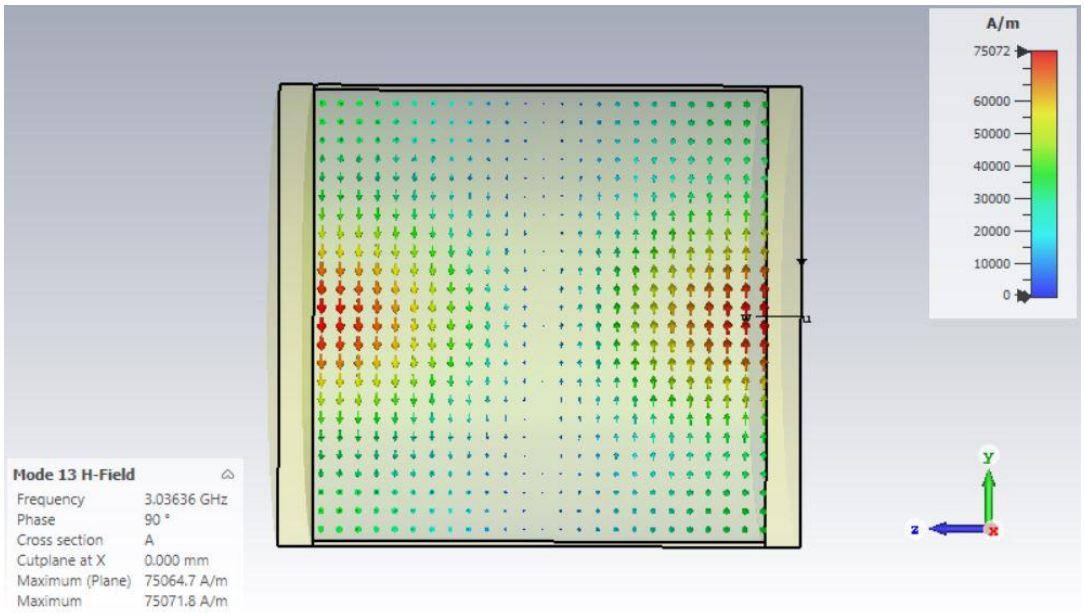

Figure 117. Cross-section of the 133mm cavity showing the simulated TM111 h-field

Figure 116 and Figure 117 show the h-field for TE_{011} and TM_{111} modes respectively. The latter illustrates that the loop is positioned in an area where the TM_{111} h-field strength is zero.

Using an Anritsu 37377C Vector Network Analyser, the cavity was characterised with the probe in three different positions: inserted into the lid with the loop aligned with the radius of the cavity, inserted into the lid with the loop at right angles to the previous position, and inserted into the side with the loop parallel to the bottom-plate – with Figure 118 and Figure 119 showing results from the latter position. TE_{011} is centred on 3.03GHz, with 650kHz BW, as measured at the -7db point. This is a measured Q of 4700 which is significantly down on the calculated 53k mainly due to very loose contact

between the walls and lid, but more than adequate for POP, even when cell loading is taken into account.

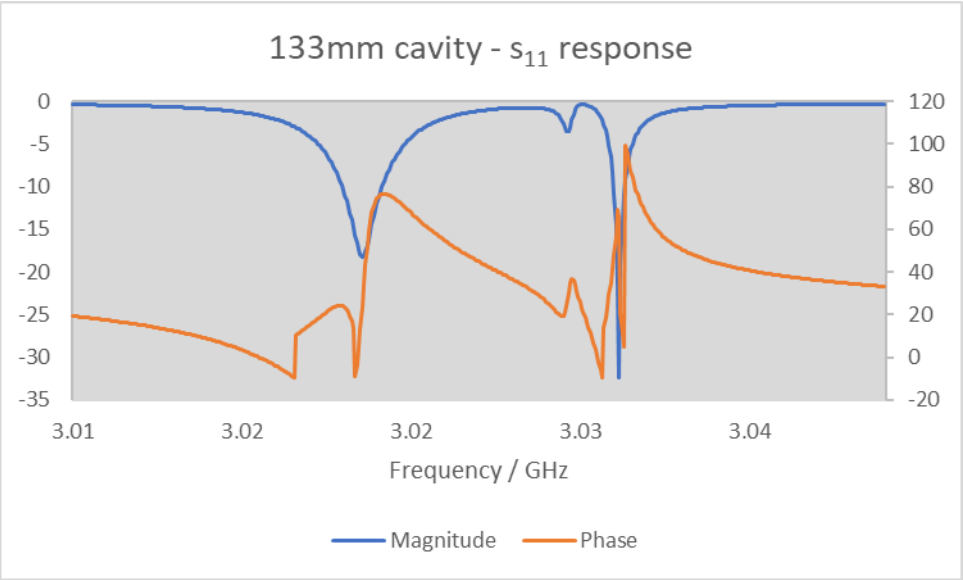

Figure 118. Measured reflection coefficient of the 133mm cavity.

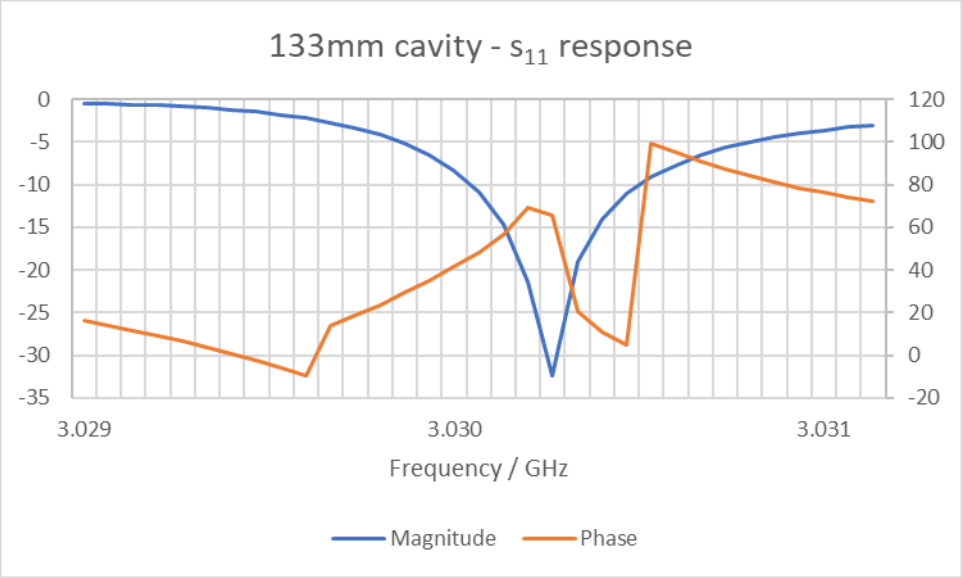

Figure 119. Reflection coefficient of the 133mm cavity close to resonance.

By way of comparison, If the cell is removed from the Turin cavity [32], the calculated TE_{011} resonance would provide a Q_0 of 14k at 7.67 GHz – the presence of the cell dominates the loaded Q (~1000) and lowers f_0 by 12%.

Characterising a single-port cavity

As the cavity only has a single port, characterising it requires the use of a VNA to measure the S_{11} reflection coefficient.

$$\Gamma = \frac{Z_L - Z_0}{Z_L + Z_0} \tag{30}$$

Where, Γ is the measured reflection coefficient (S_{11}), Z_0 is the characteristic impedance of the source, and Z_L is the impedance of the MW resonator cavity. If the cavity is modelled as a series RLC resonator then Z_L can be expressed as:

$$Z_L = R_L + j\omega L + \frac{1}{j\omega C} \tag{31}$$

As Z_0 is purely resistive then Γ can be expressed as:

$$\Gamma = \frac{1 - \omega^2 LC + j\omega C(R_L - R_0)}{1 - \omega^2 LC + j\omega C(R_L + R_0)} \tag{32}$$

At resonance $Q = \frac{\omega L}{R_L} = \frac{1}{CR_L\omega}$ so it's possible to eliminate L and C and express the reflection coefficient as follows:

Equation 33. Reflection coefficient expressed as a function of ω, Q, R_L and R_0

$$\Gamma = \frac{1 + \frac{j\omega}{QR_L\omega_0}(R_L - R_0) - \frac{\omega^2}{\omega_0^2}}{1 + \frac{j\omega}{QR_L\omega_0}(R_L + R_0) - \frac{\omega^2}{\omega_0^2}} \tag{34}$$

A single-port copper microwave cavity has been designed to be resonant at 3.035 GHz, and constructed with probe impedance ~50Ω (set by the physical dimensions of RG402 semi-rigid cable). R_s is 50Ω, set by the Anritsu 37377C VNA used to measure the cavity, and the resonance can be identified by noting a minimum in Γ. R_L is assumed to be 50Ω and it can be seen that the measured s_{11} tallies quite well with the modelled behaviour.

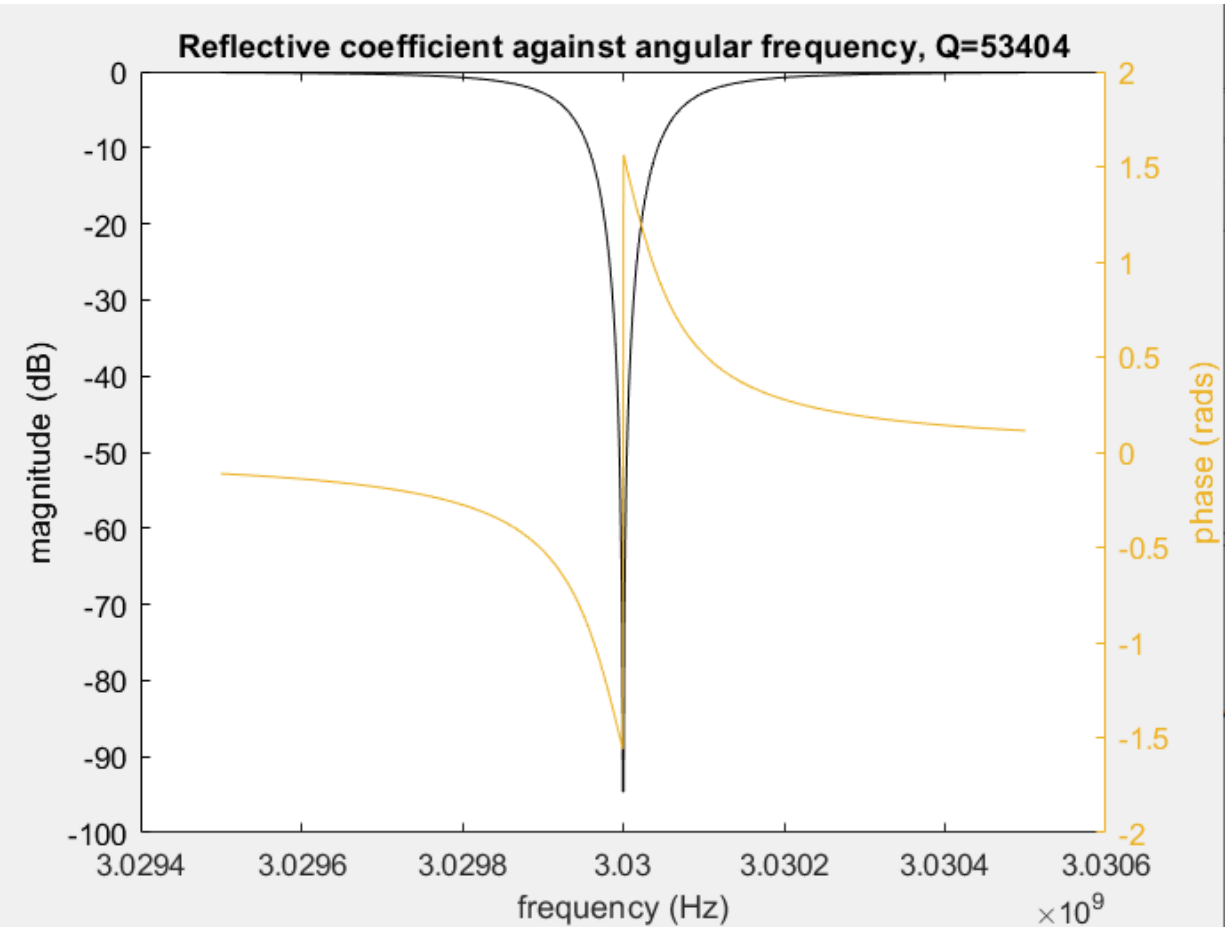

Figure 120: Reflection coefficient close to resonance

A secondary definition of Q is that $Q = \frac{f_0}{resonant\ BW}$, making it possible to calculate Q by measuring the resonant bandwidth, if one knows the reflection coefficient at the resonant bandwidth.

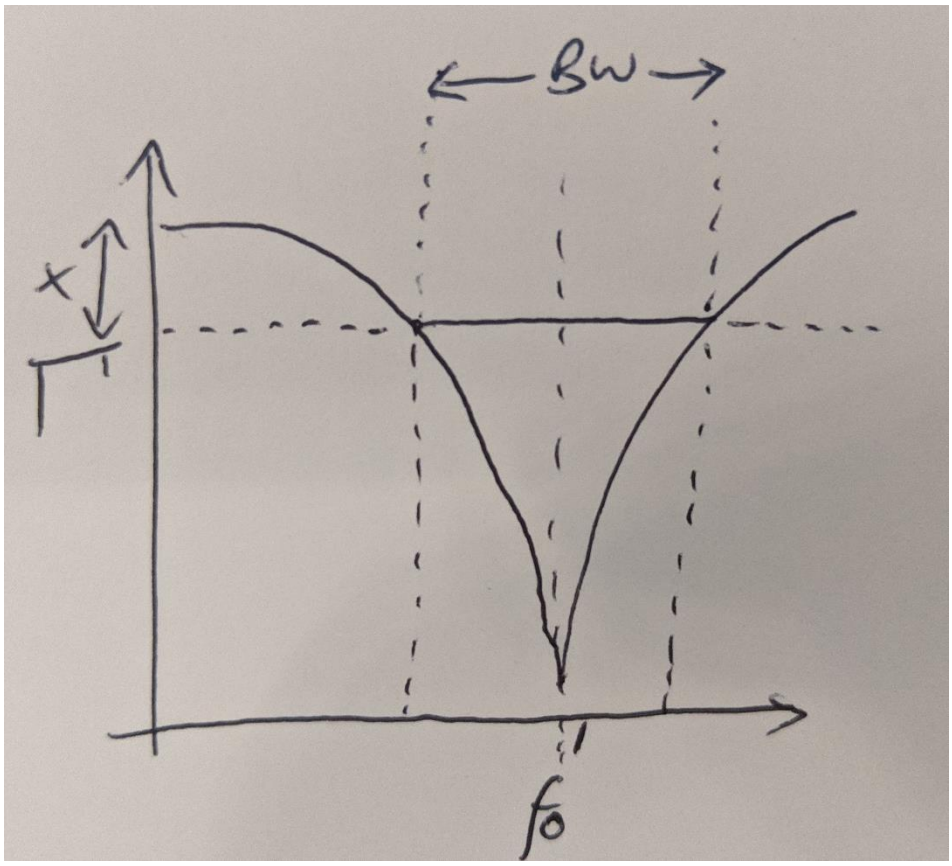

Figure 121. Reflection coefficient close to resonant frequency

Figure 121 shows a symmetrical response as excitation frequency is swept across f_0. By letting ω equal $\omega_0(1 + \frac{1}{2Q})$, all ω terms drop out of (34), resulting in (35).

$$\Gamma = \frac{R_L(1 - \frac{1}{4Q}) + j(1 - \frac{1}{2Q})(R_L - R_0)}{R_L(1 - \frac{1}{4Q}) + j(1 - \frac{1}{2Q})(R_L + R_0)} \tag{35}$$

As all instances of Q are as part of denominators, it can be seen that Reflection Coefficient Γ is not dependent on Q for Q > 10.

$$\Gamma \rightarrow \frac{R_L + j(R_L - R_0)}{R_L + j(R_L + R_0)} \tag{36}$$

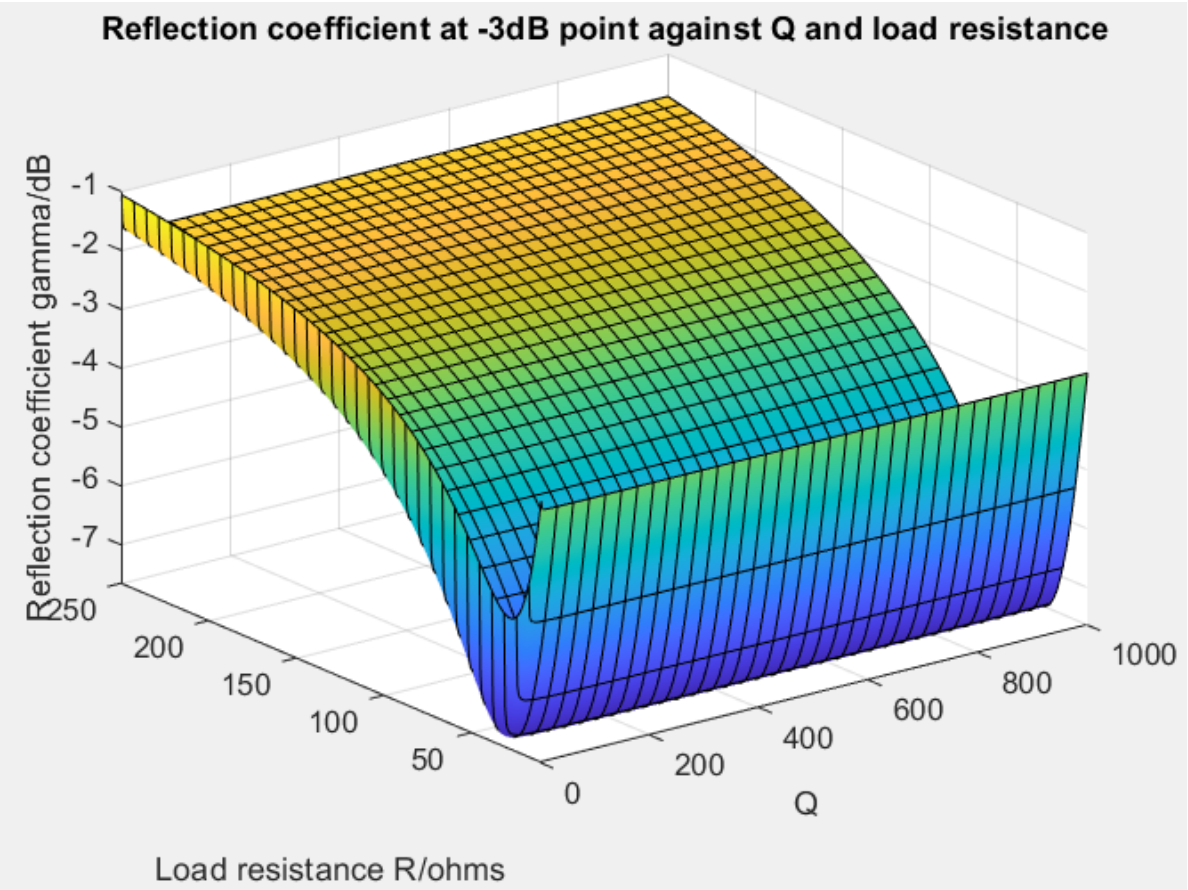

Figure 122. Reflection coefficient plotted against load resistance and Q

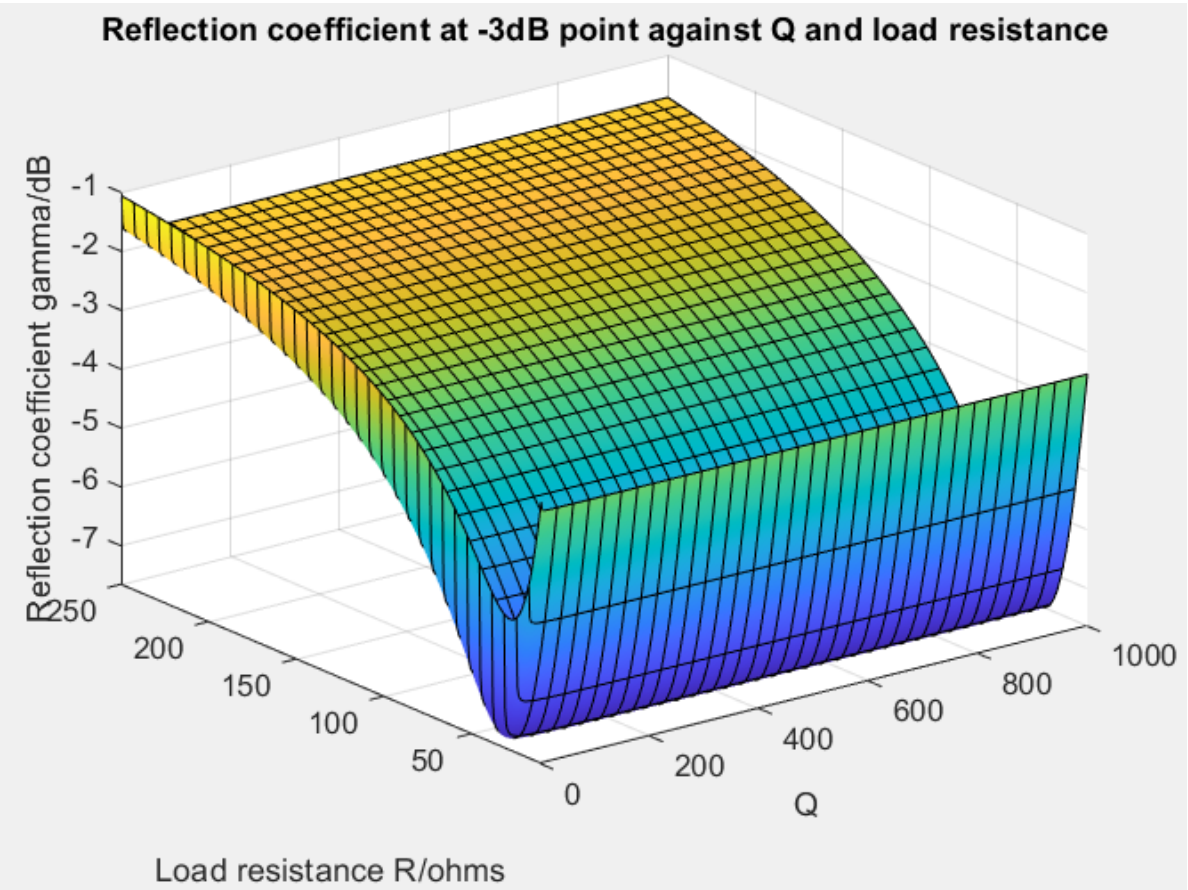

Figure 122 plots the magnitude of the reflection coefficient against load resistance and Q, shows relative independence between |Γ| and any practical value of Q. Figure 123 is calculated with matching 50Ω source and load resistance and more clearly shows Γ converge with increasing Q.

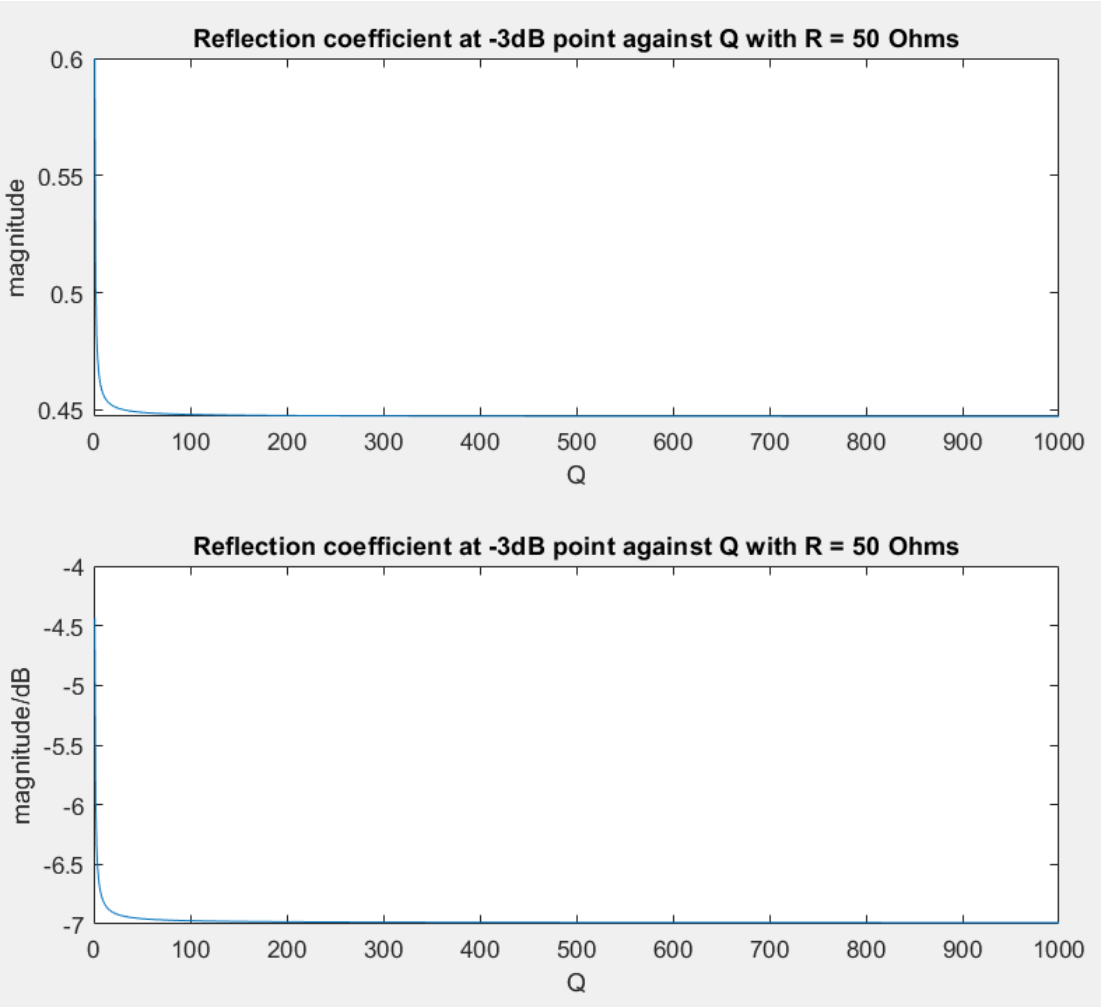

Figure 123. Reflection coefficient against Q (50Ω source and load impedance)

Perhaps less intuitively, Figure 124 shows that the -3db point occurs when source and load resistance are in the ratio of $1:\sqrt{\frac{1}{2}}$. When the source and load resistance are matched (50Ω), the resonant bandwidth can be measured by noting the frequency range where |Γ| dips below -7dB.

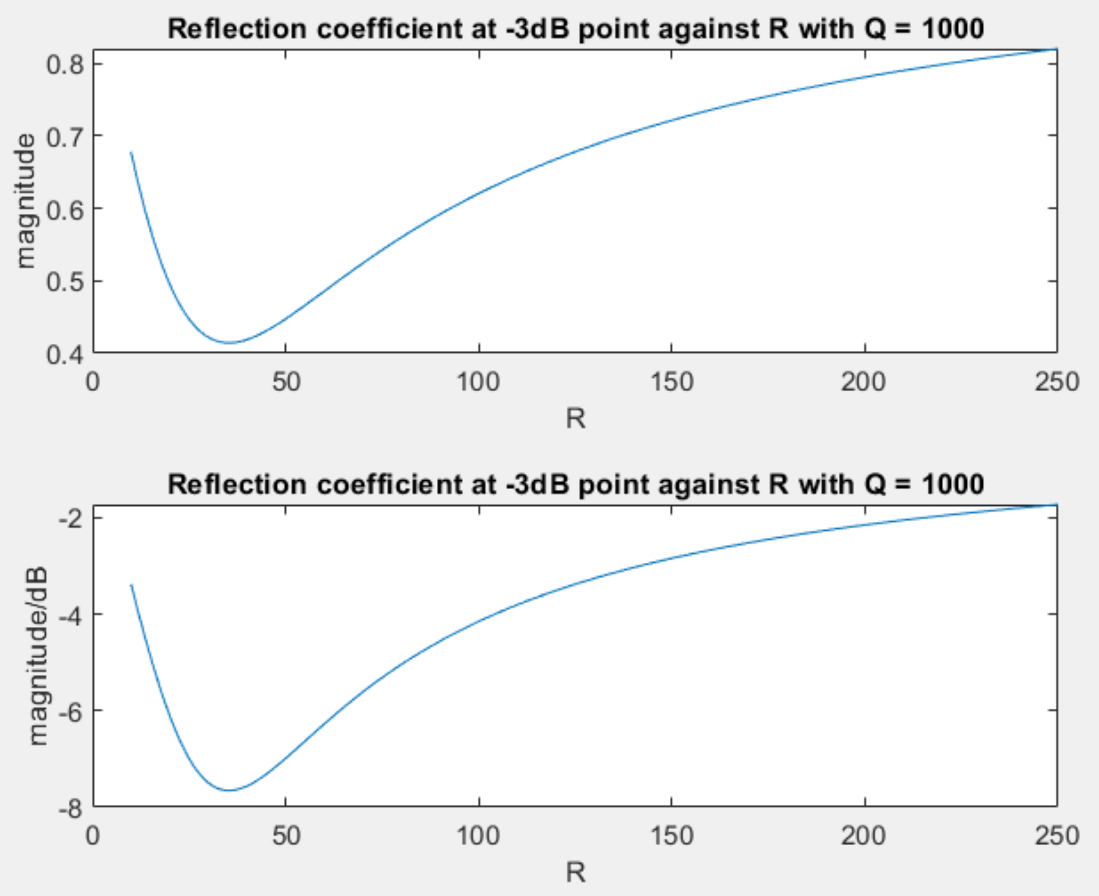

Figure 124. Reflection coefficient against load resistance (50Ω source impedance and Q fixed at 1000)

Appendix C. Design, simulation and build of an amplified photodiode.

An improvement was made to optical detection at this point, and instead of connecting the photodiode, via a length of cable, to a generic transimpedance amplifier, a bespoke amplifier was designed that would incorporate the Osram BPX61 PIN photodiode into a small package, providing a much lower noise solution. It was desirable that the photodiode be able to cope with a 4mW optical beam, which, at 795nm equates to 1.6×10^{16} photons/s. The photodiode's peak spectral sensitivity is 850nm, the current output is approximately 95% of that specified at 795nm – this equates to a quantum yield η of approx. 0.86 electrons per photon or 2.57mA. The device is extremely fast and exhibits a 20ns rise and fall time when operated with a reverse bias of 5V. This was modelled as a 2.5mA current source, 20pF capacitor and 250MΩ shunt resistance (to represent the specified 20nA dark current at 50°C operating conditions, with a 5V reverse bias).

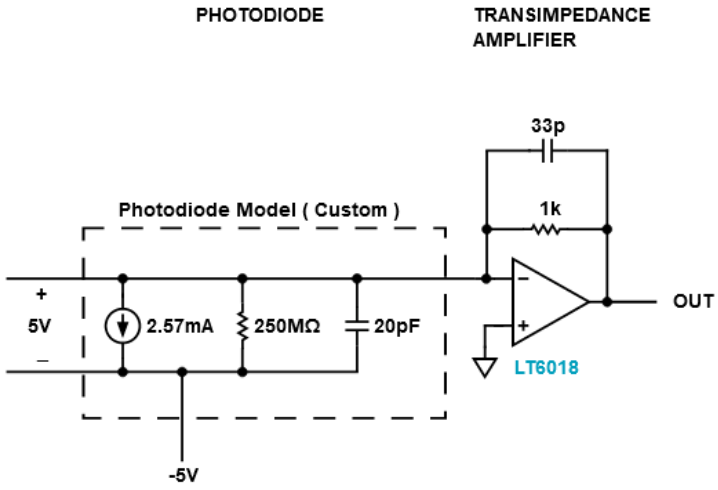

Figure 125. Photodiode and transimpedance amplifier model

Figure 125 shows the model, with 5V reverse bias, connected to an LT6018 op-amp, modelled in Analog Devices online simulator (https://tools.analog.com/en/photodiode/). Although the manufacturer does not recommend using an LT6018 in this application, due to a relatively high output offset (there's potentially a bias current of 150nA), this op-amp is extremely low noise and has already proven successful as a pre-amplifier in the York lab. The inverting amplifier configuration and 1kΩ feedback

resistor set a gain of -1000 V/A and a 33pF feedback capacitor improves the amplifier response time at the expense of the little overshoot, as visible in Figure 126.

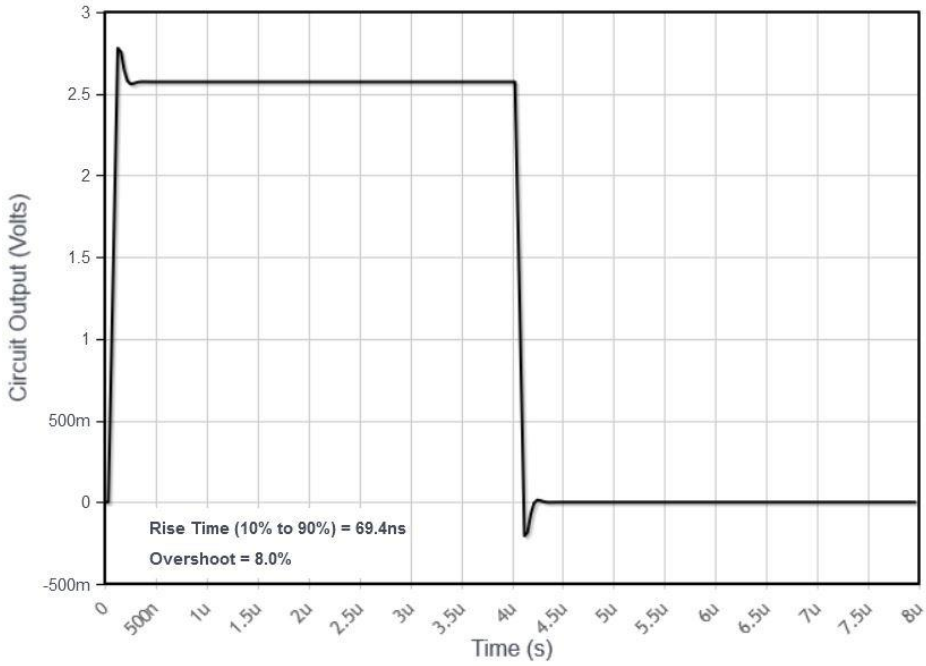

Figure 126. Photodiode and TIA pulse response.

Figure 126 and Figure 127 show that the roll-off frequency is vastly higher than the 50kHz requirement and the first stage could be used in much higher frequency applications if required. Unfortunately, the online simulator cuts off the low-end frequency response (and doesn't model the 1/f noise) but the 60dB transimpedance gain does extend to DC.

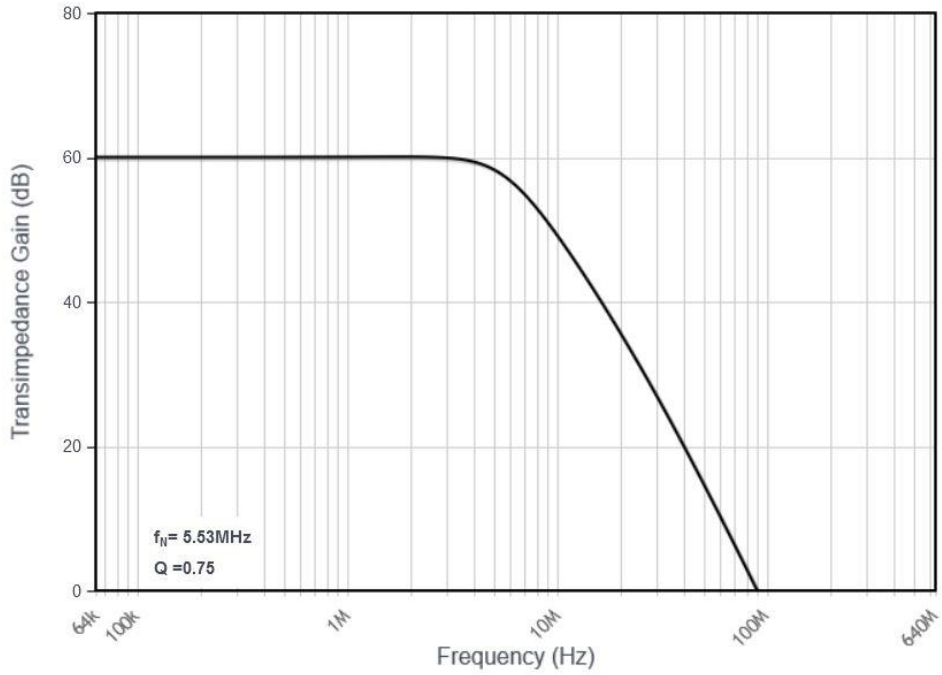

Figure 127. Combined photodiode and transimpedance amplifier frequency response.

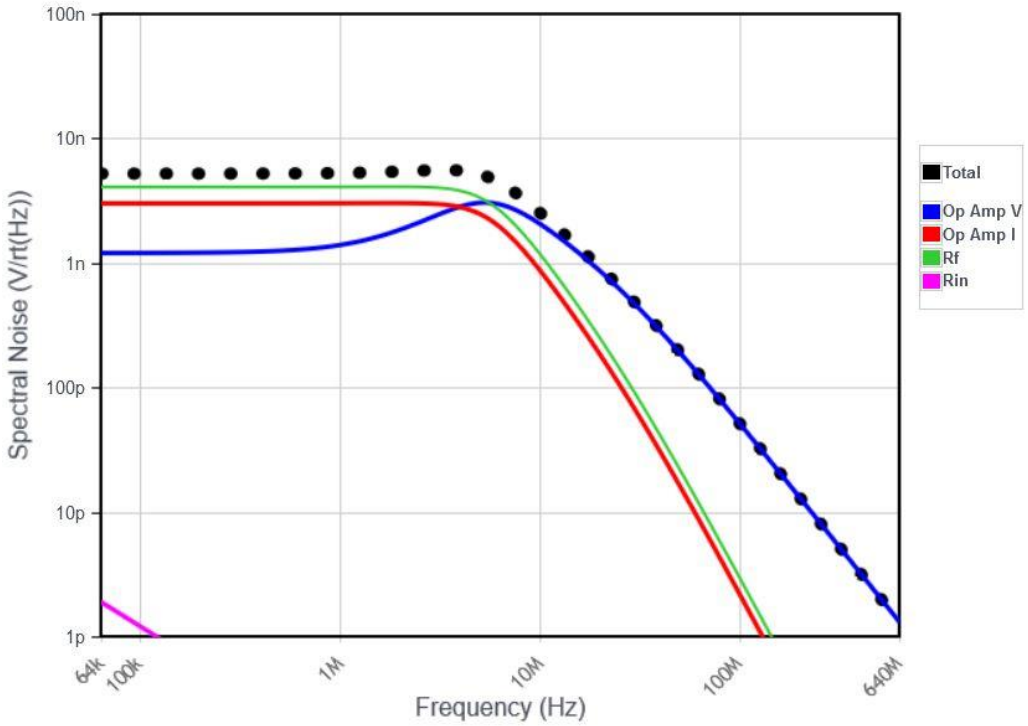

Figure 128. Photodiode and transimpedance amplifier spectral noise.

Figure 128 shows that the overall noise is dominated by feedback resistor thermal noise, followed by op-amp current noise, producing a figure of ~4nV/√Hz below 1.5MHz. The 1kΩ feedback resistor provides a convenient 2.55V output with 15.8µV noise floor and signal to noise ratio of 95.1dB. This simulated figure appears not to include shot noise, which would be an additional 6.3µV noise contributor (into a system with 50kHz bandwidth), bringing the noise-floor up to 17.0µV.

Appendix D. Logical Interface schematics

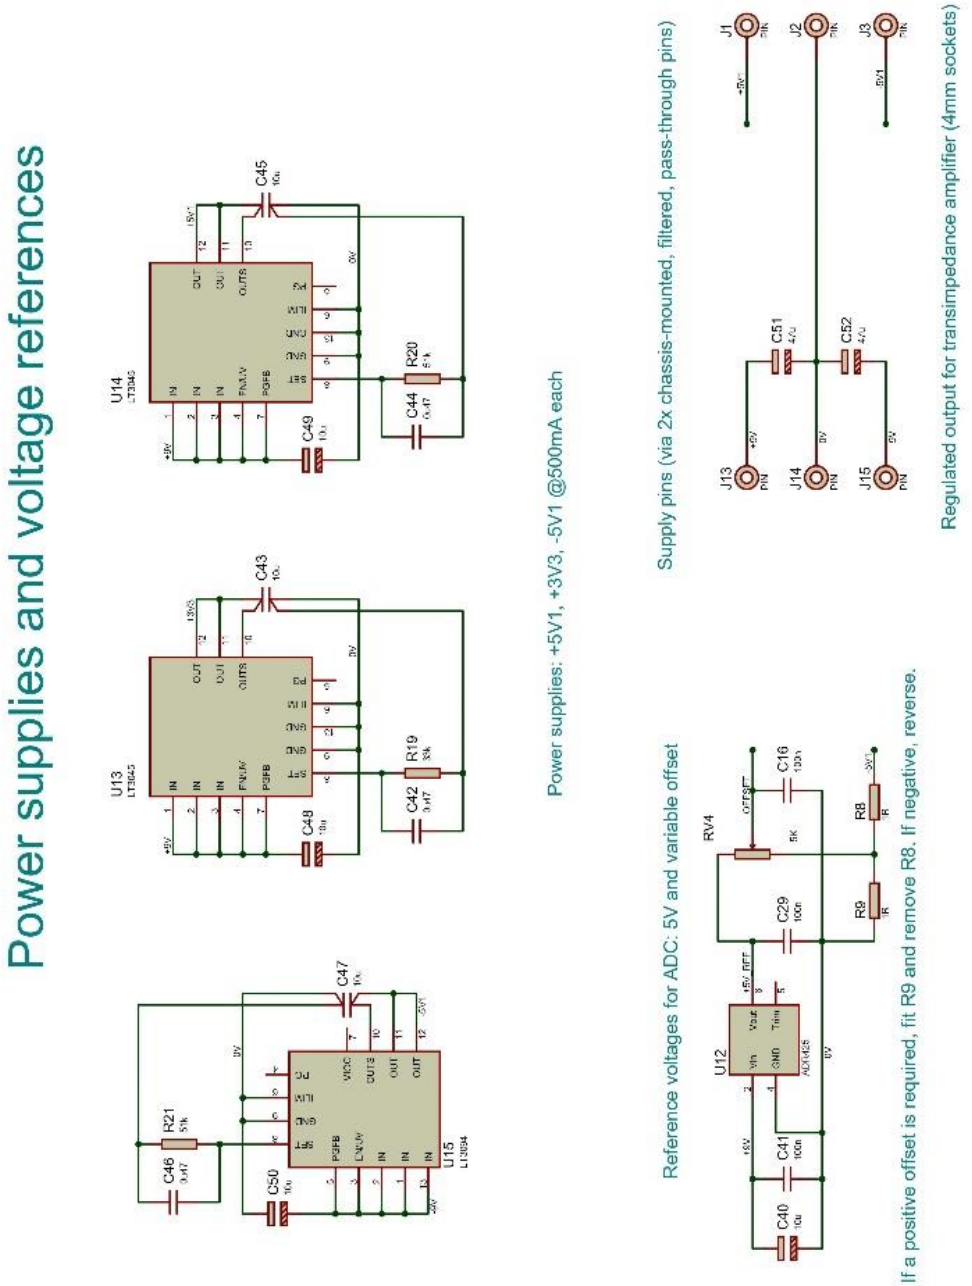

Sample & hold, ADC

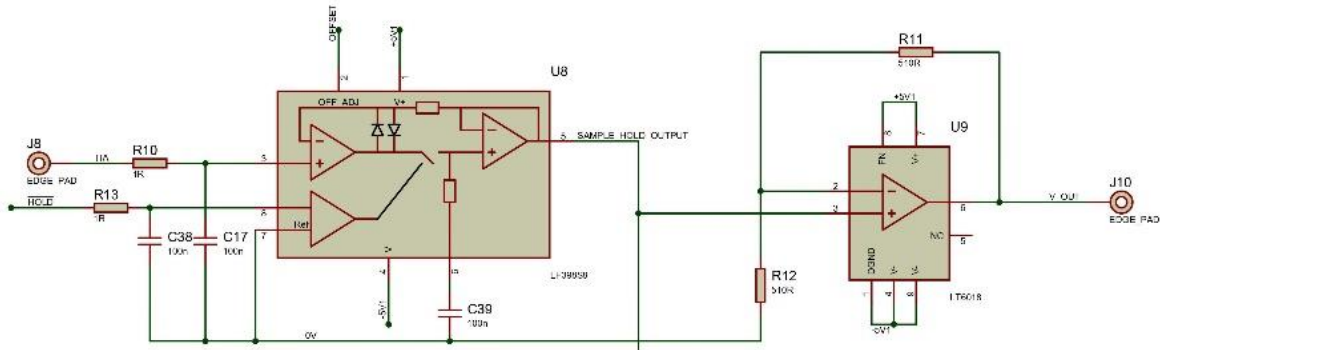

U9 prevents loading of the sample and hold circuit and provides a x2 voltage gain

Hold capacitor C39 to be low leakage, low dielectric absorption
i.e. mica, polystyrene or polypropelene
Reduce board leakage by using the output signal to create a guard ring

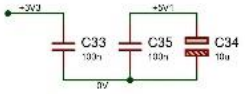

Decoupling

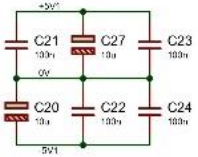

Low Pass RC filters
R14, C37 - 150kHz
Provision only
R10, C17 - do not fit C17
R13, C38 - do not fit C38

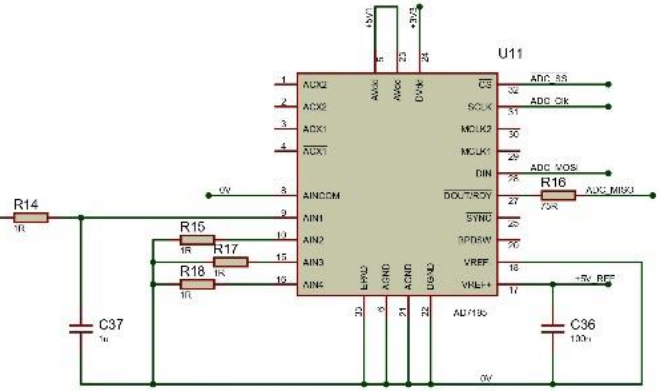

The reference voltages must not be greater than AVdd

AOM modulation control

First stage - unity gain, non-inverting.

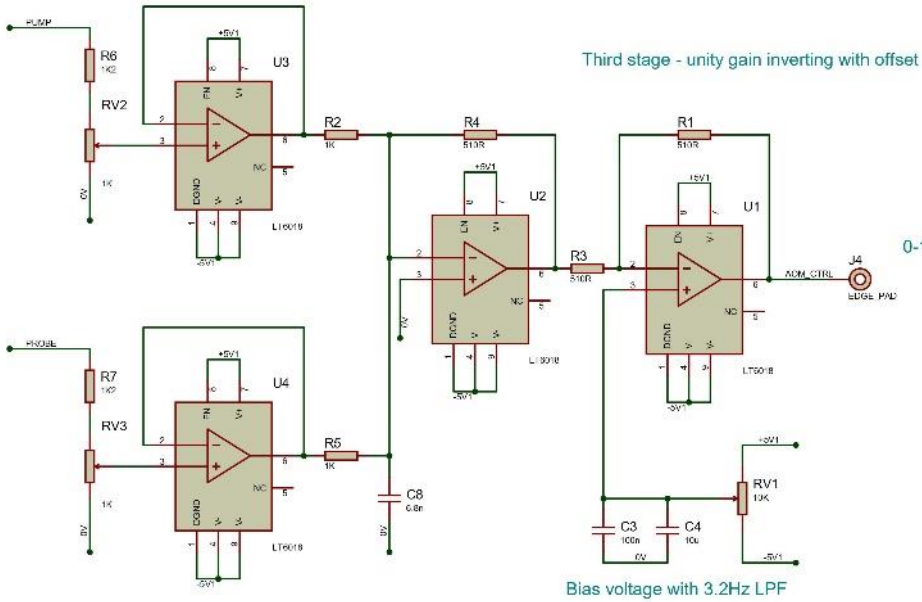

Decoupling

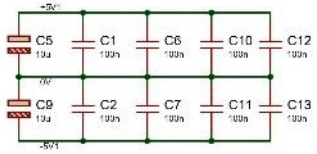

MCU connectivity

To Atmel SAML22 dev board

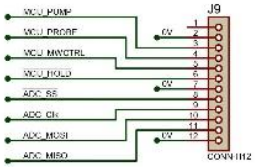

Decoupling

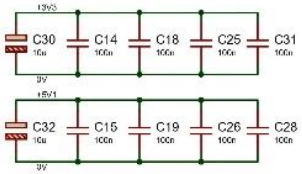

Mounting posts and ground stitching

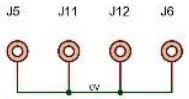

MCU to TTL voltage translation

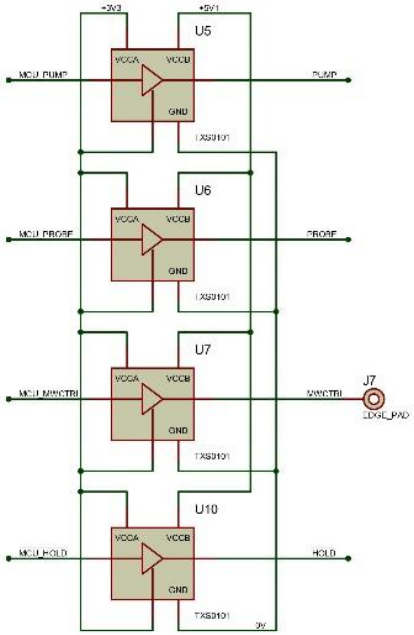

Appendix E. High-resolution DAC schematics

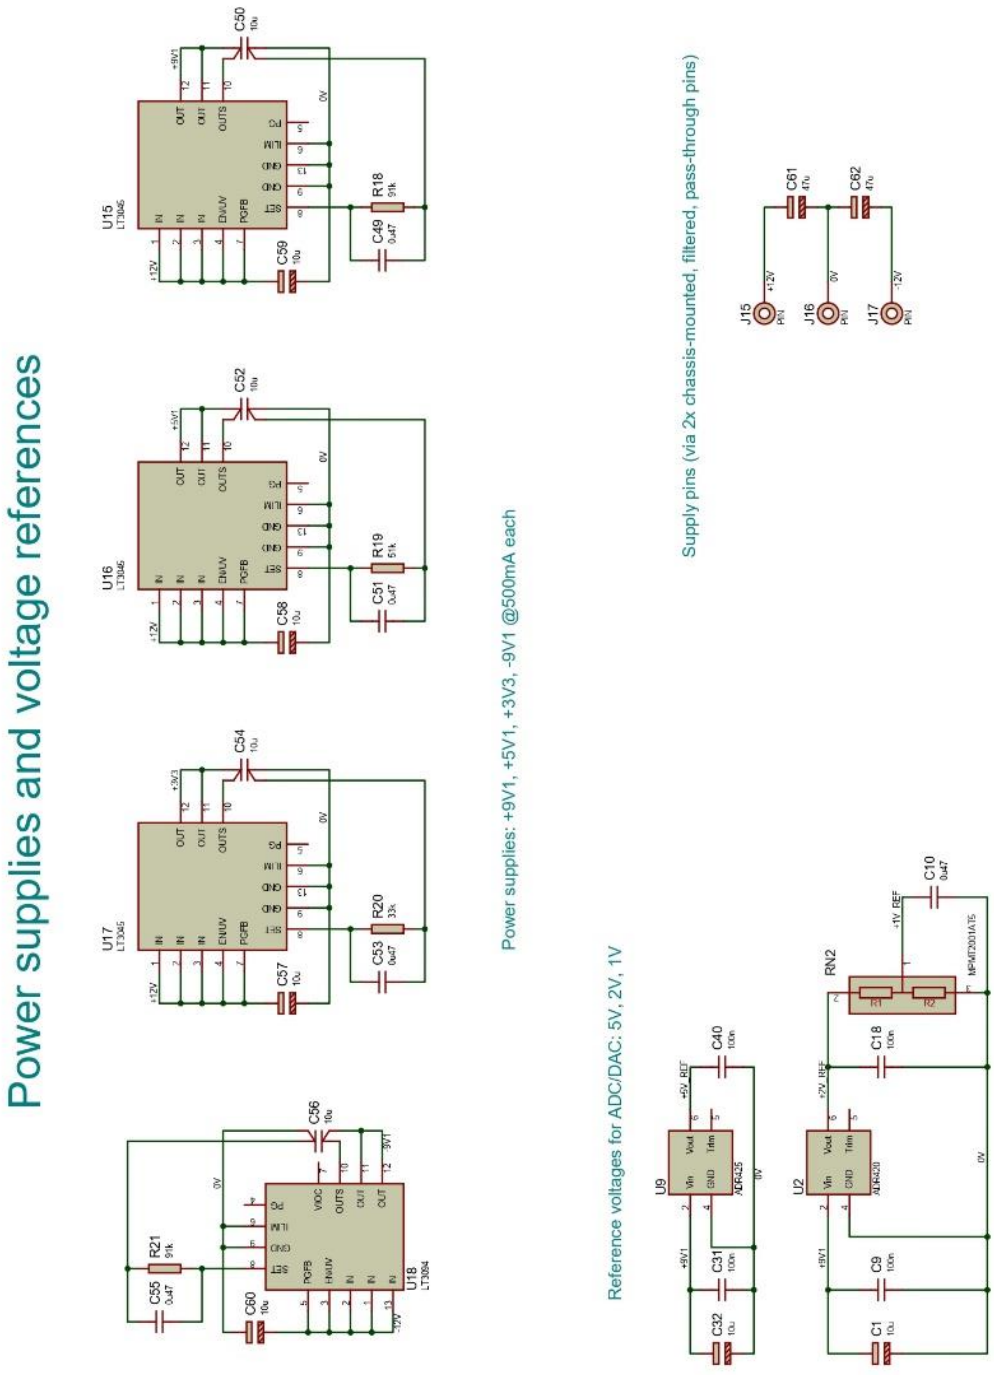

Major-minor DAC

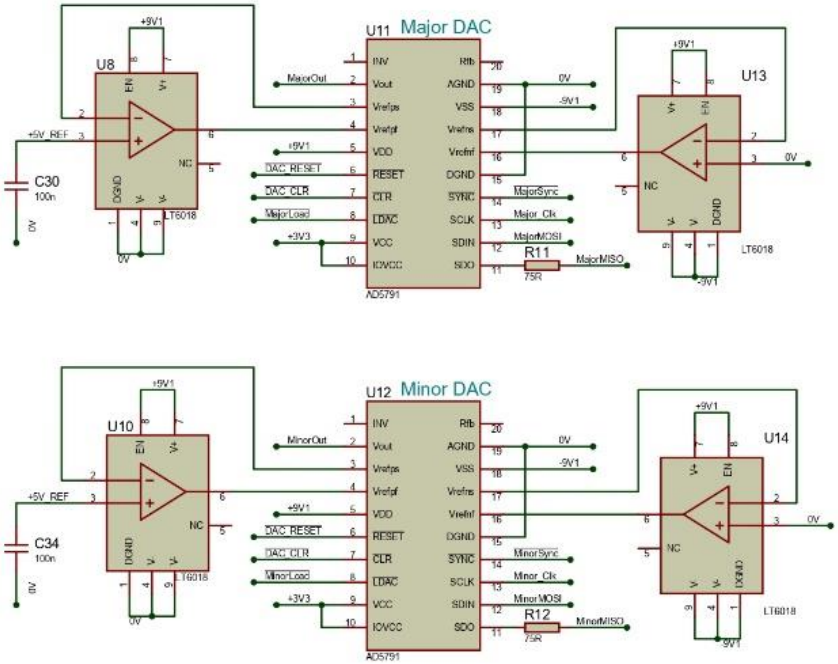

Decoupling

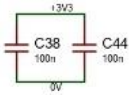

100nF close to each DAC and amp (both +ve and -ve supplies)

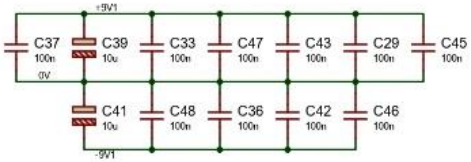

Summation amplifier

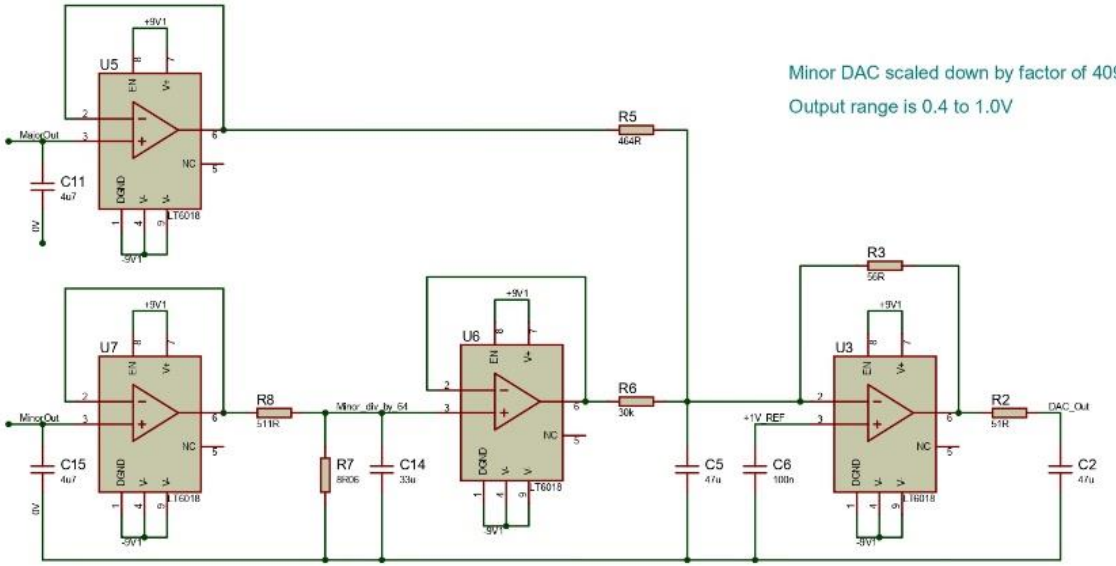

100nF close to each amp (both +ve and -ve supplies)

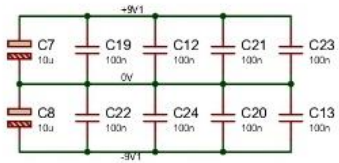

ADC

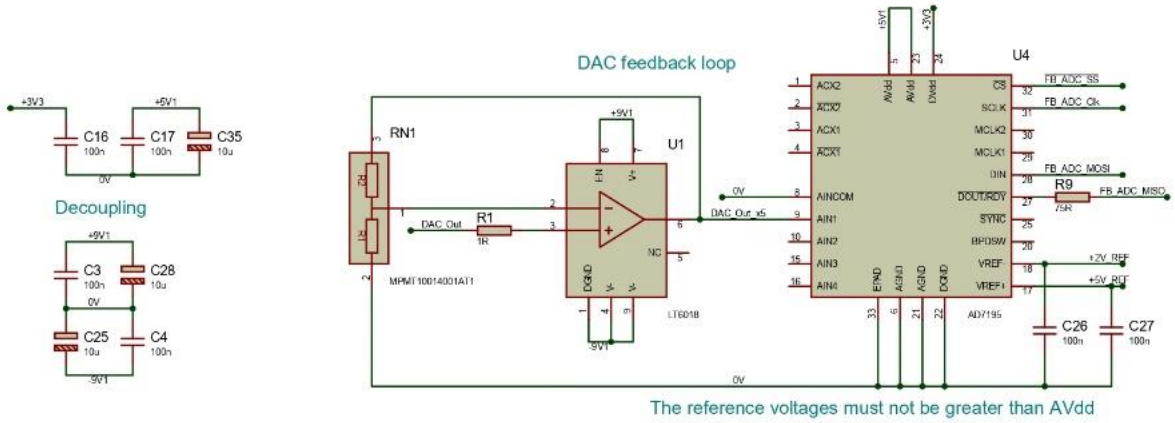
Decoupling
+3V3
+5V1
C16
100n
C17
100n
C35
10u
0V
+9V1
C3
100n
C28
10u
0V
C25
10u
C4
100n
-9V1
DAC feedback loop
RN1
MPMT10014001AT1
DAC_Out
R1
1R
+9V1
U1
EN
V+
DGND
V-
NC
LT6018
-9V1
0V
DAC_Out_x5
U4
+5V1
+3V3
AVdd
DVdd
ACX2
ACX1
ACX0
AINCOM
AIN1
AIN2
AIN3
AIN4
CS
SCLK
MCLK2
MCLK1
DIN
DOUT/RDY
SYNC
BPDSW
VREF
VREF-
EGND
AGND
DGND
AD7195
FB_ADC_SS
FB_ADC_Clk
FB_ADC_MOSI
FB_ADC_MISO
R9
75R
+2V_REF
+5V_REF
C26
100n
C27
100n
0V
The reference voltages must not be greater than AVdd

Debugging

Voltage monitoring points

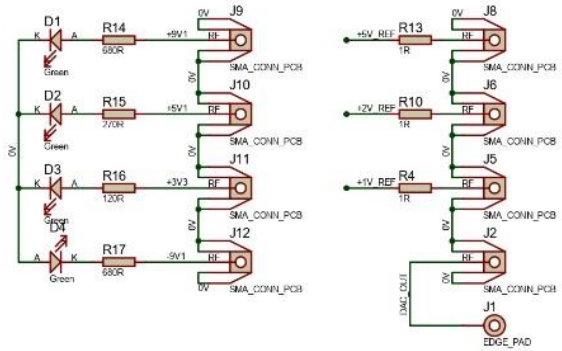

MCU connectivity

To Atmel SAML22 dev board

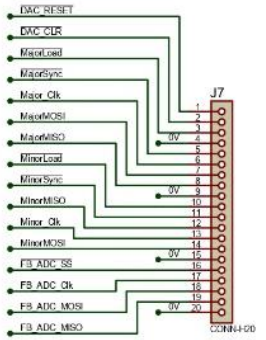

Mounting posts and ground stitching