# Laser Stabilisation using a Slow Light Cavity

**Teodor Strömberg**
**Supervisor: Lars Rippe**

Master's Thesis
Spring 2015


LUNDS UNIVERSITET

Department of Physics
Atomic Physics Division

## Populärvetenskaplig Sammanfattning

Det sägs ofta att laserljus är enfärgat, men detta är en sanning med modifikation då helt enfärgat ljus endast finns i tankeexperiment. Likväl kan lasrar komma väldigt nära till att producera enfärgat ljus och måttet på detta kallas linjebredd. Mer konkret är linjebredden ett mått på hur stora variationerna i antalet svängningar ljuset genomför per sekund är. I många av laserns användningsområden är det önskvärt att dessa variationer är så små som möjligt. Ett exempel på detta är den just nu pågående utvecklingen av atomur baserade på ljus.

De välkända och väldigt noggranna atomuren har hittills varit baserade på mikrovågor. I dessa ur jämförs en mikrovågs frekvens med en väldigt väldefinierad frekvens för en energiövergång hos en atom. Ju närmare mikrovågens frekvens ligger jämte atomens desto starkare svar kan utläsas. På detta vis kan man se till att svängningarna hos vågen är regelbundna även över lång tid. Svängningarna kan då användas som en referens för tid, i analogi med ett pendelur. Tanken bakom ljusbaserade atomur är ersätta mikrovågen med laserljus. Svängningarna i en ljusvåg är omkring 100 000 gånger snabbare, vilket möjliggör noggrannare tidmätningar. Detta förutsätter dock att svängningarna är tillräckligt regelbundna, och därmed återvänder vi till laserns linjebredd.

För att uppnå en snäv linjebredd måste oregelbundenheter i ljusets svängningar mätas och korrigeras. Denna process kallas stabilisering. Tillvägagångssättet för detta är vanligtvis att ljuset skickas in i en optisk kavitet, vilken består av två speglar positionerade så att ljuset studsar fram och tillbaka många gånger. En sådan uppställning tillåter endast ljuset passera om det har en viss frekvens, kallad resonansfrekvensen, och avvikelser från denna frekvens kan således mätas. Svårigheten med detta ligger i att om speglarna flyttar sig lite, t.ex. genom värmevibrationer, flyttar sig även kavitetens resonansfrekvens. Kaviteten är bara en bra referens för mätningen om den själv är stabil. För att nå den prestanda som krävs för att utnyttja all potential i t.ex. optiska atomur måste speglarna röra sig med mindre än radien på en proton.

Det framstår nog som självklart att sådan stabilitet är svår att uppnå. I detta examensarbete föreslår vi en metod för att minska känsligheten mot små förändringar i kaviteten. Detta kan åstadkommas genom att sakta ner ljuset. Då ljuset färdas långsammare kommer speglarna agera som om de vore längre ifrån varandra, men samtidigt kommer små förändringar i speglarnas position förbli opåverkade. Följden av detta blir att även om speglarna rör sig lika mycket, så rör de sig mindre i förhållande till avståndet dem sinsemellan. Genom att sakta ner ljuset över 10 000 gånger kan känsligheten mot speglarnas rörelser minskas lika mycket.

För att sakta ner ljusets används en kristall, till vilken en låg koncentration av jordartsmetalljoner har tillsats. Den i vanliga fall genomskinliga kristallen kommer då att absorbera ljus inom ett visst frekvensintervall. Med andra ord släpps somliga färger inte igenom. Sedan skärs ett så kallat spektralt hål i kristallen med hjälp av en laser. Det är inget fysiskt hål, utan ett hål i det frekvensintervall som kristallen absorberar. Det vill säga, inom detta intervall där ljuset inte släpps igenom finns ett mindre intervall där ljuset kan passera. Ljus med just den frekvens som krävs för att passera genom det här intervallet kommer att saktas ned avsevärt.

**Abstract**

In this thesis the application of slow light effects to the area of laser stabilisation is investigated. The frequency stability of lasers locked to external cavities is limited by the stability of the modes in the reference cavity. By using dispersive slow light effects the effective length of a cavity can be increased by several orders of magnitude, without a proportional increase in the sensitivity to fluctuations in cavity length, thereby yielding a net increase in stability. The effects are achieved through spectral hole burning in cryogenically cooled rare-earth ion doped crystals. Two sets of experiments are carried out, the first focusing on hole burning and the second focusing laser locking. Linewidths and frequency drift of the cavity spectrum are studied. Feedback control using an error signal generated from the slow light cavity is demonstrated, but successfully locking the laser is still yet to be achieved.

## Acknowledgements

I would like to thank, first of all, my supervisor Lars Rippe for always being more than willing to answer any and all questions, for being so enthusiastic and being the best supervisor I could have hoped to have. Secondly, I would like to thank head of the Quantum Information group, Stefan Kröll, not only for finding me this project on short notice but also for his great lecturing in the atomic physics courses that drew me to do my thesis on the subject. Additionally I would like to thank all the other members of the QI group, Andreas, Jenny, Adam, Aswath and Qian, as well as all my fellow thesis students in the group, for always being willing to help me in the lab, answer any questions I might have had and for making my time at the division a memorable and enjoyable one. Finally I would like to thank my sister, for her help and encouragement.

# Contents

# Abbreviations

**AOM** - Acousto Optic Modulator

**AWG** - Arbitrary Waveform Generator

**EOM** - Electro Optic Modulator

**$Eu^{3+}$** - Europium three times ionised

**FFT** - Fast Fourier Transform

**FSR** - Free Spectral Range

**FWHM** - Full Width Half Max

**HeNe** - Helium-Neon

**PDH** - Pound-Drever-Hall

**$Pr^{3+}$** - Praseodymium three times ionised

**PI** - Proportional-Integral

**RF** - Radio Frequency

**TEM** - Transverse Electric Mode

**VCO** - Voltage Controlled Oscillator

**$Y_2SiO_5$** - Yttrium Orthosilicate

# 1 Introduction

Lasers have played an important role in physics, technology and medicine ever since their invention in 1960, with uses ranging from surgery to high precision spectroscopy. In many applications the precision of the laser, the linewidth, is of critical importance. One example of this is the somewhat recent development of optical clocks.

For half a century atomic clocks have set the standard for precision timekeeping and have allowed technologies such as GPS to be realised. Over the years the precision of atomic clocks based on the hyperfine levels in Caesium that form the definition of the SI second has improved by several orders of magnitude, from a fractional uncertainty of $10^{-11}$ in the 1960's to $10^{-16}$ today [15]. The transition used in Cs clocks lies in the microwave regime, at around 9.19 GHz. By basing a clock around a transition in the optical range, with a frequency some five orders of magnitude higher, the potential exists for even more stable clocks [16]. The development of optical frequency combs, for which Theodor W. Hänsch and John L. Hall were awarded the Nobel prize in physics in 2005, has made it possible to measure the frequency of light with the precision required to realise optical clocks and in recent years optical clocks that surpass the performance Cs clocks have been constructed [20]. These clocks however, are limited in part by the stability of the lasers used to probe the atomic transitions in question. For this reason laser stabilisation will play an important role in the continued development of optical clocks.

High stability lasers are important in many other applications as well, including gravitational wave detection [14] and quantum information research [7]. Today, state of the art lasers are capable of linewidths below 50 mHz [17], a remarkable feat considering that the frequency of visible light is in the hundreds of THz. This is achieved through temperature controlled high finesse ultra-low expansion cavities. While this approach has yielded very good results, it is nevertheless a good idea to investigate other techniques that could be used independently of or in conjunction with existing ones. This thesis concerns one such possible technique.

## 1.1 Aim of the Thesis

The goal of this master's thesis project was to investigate how to lock a laser to a resonance peak in a slow light cavity, as well as to plan and perform experiments in which the locking of a laser using a $Pr^{3+}$:$Y_2SiO_5$ crystal was to be attempted. Areas of investigation include many specifics of the locking scheme, the limits on what resonance peak linewidths can be achieved, what an appropriate intensity of the light used to couple to the cavity is and the optimisation of hole burning pulses on a coated crystal.

## 1.2 Outline

This thesis consists of seven sections. Section 2 begins with a brief introduction to the optics concepts that form the basis for the experiments conducted, which

is followed by a discussion of the principles behind laser stabilisation. Section 3 discusses spectral hole burning, how this can be exploited to achieve a slow light effect and the benefits this could offer to laser stability. In section 4 the first set of cryogenic experiments, focusing on spectral hole burning in the crystal, are described and accompanying results are presented. Section 5 goes over the components used in assembling the locking system as well as a small test of the EOM. The final set of experiments, in which the laser was attempted to be locked to the slow light cavity, are presented in section 6. Finally, section 7 contains a brief summary of the results as well as a discussion on what could be done in the future.

# 2 Laser Locking

Before discussing the principles behind laser stabilisation I will first provide the reader with a brief background on optical resonators and lasers.

## 2.1 Resonator Optics

An optical resonator, or optical cavity, consists in its simplest form of two mirrors arranged in such a way that light bouncing back and forth between them will be trapped. Consider the thought experiment of two fully reflecting parallel plane mirrors in vacuum, as pictured in Fig. 1. A monochromatic beam of light reflecting back and forth between the two mirrors accumulates after each round trip a phase shift of:

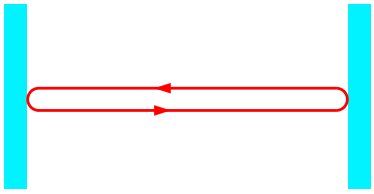

Figure 1: A plane mirror optical resonator. If the mirrors are properly aligned the beam will retrace its path each trip, the separation is only for illustrative purposes.

$$\varphi = \frac{4\pi d}{\lambda} \tag{1}$$

where $d$ is the length of the cavity and $\lambda$ the wavelength of the light. There is an additional $2\pi$ phase shift from the two reflections off the mirrors, but this can be neglected. If the wavelength satisfies the condition:

$$\lambda = \frac{2d}{m}, \qquad m = 1, 2, ... \tag{2}$$

we can see that the phase shift of the beam will be $2\pi m$; the light interferes constructively and forms a standing wave inside the cavity. A beam that matches this criterion is called a longitudinal mode of the cavity. Since a monochromatic beam must be infinitely long, the beam will at each point be an infinite sum of infinitesimal fields, all with equal magnitude. If their phases do not add constructively according to Eq. 2 the beam will vanish in every point [1]. In other words, only light that is resonant can exist inside the cavity. The spacing between two neighbouring modes supported by the cavity is constant in frequency. It is referred to as the free spectral range, and using that the frequency of light is $\nu = \frac{c}{\lambda}$ it can be written as:

$$\nu_{fsr} = \frac{c}{2d} \tag{3}$$

where $c$ is the speed of light in vacuum. The spectrum of a cavity, showing the allowed frequencies, is depicted in Fig. 2a.

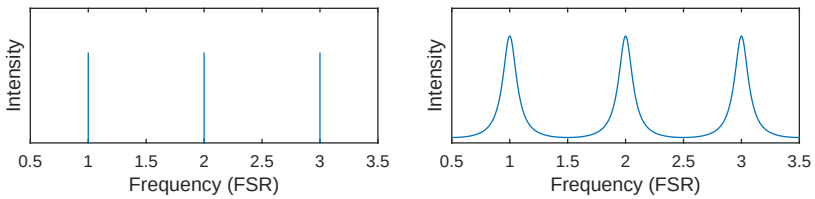


(a) Spectrum of a lossless cavity.

(b) Spectrum of a lossy cavity.

Figure 2: The imagined spectrum of a lossless cavity in (a) only has discreet frequencies, but in the presence of losses the spectrum broadens (b) and becomes continuous.

### 2.1.1 Lossy cavities

Since perfectly reflective mirrors do not exist, we will instead consider mirrors that have some degree of loss (through transmission, absorption or, more likely, both). This changes the situation. A beam retracing its own path inside a cavity will lose power each time it reflects off a mirror. When summing up the beams to get the total field we therefore only need to consider a finite subset of the reflected beams. As long as the phase shift during each round trip is small enough that these beams still interfere constructively they can form a standing wave inside the cavity. The result is that the spectrum in Fig. 2a broadens and instead looks like what is depicted in Fig. 2b. The width, or linewidth, $\delta\nu$ of a peak in a spectrum such as Fig. 2b is measured at half of the max height, and is referred to as the Full Width Half Maximum or FWHM for short. An important quantity for lossy resonators is the finesse [1]:

$$\mathcal{F} = \frac{\pi\sqrt{|r|}}{1-|r|} \tag{4}$$

Here $r$ is the the fractional amplitude loss for each round trip (the fractional intensity loss is $r^2$). For cavities with small losses the finesse is related to the linewidth through the free specral range:

$$\mathcal{F} = \frac{\nu_{fsr}}{\delta\nu} \tag{5}$$

Low losses (hence high finesse) and a long cavity leads to narrow peaks with a wide spacing.

### 2.1.2 Optical cavities as filters

A cavity's property of only permitting standing waves at very specific frequencies can be used to create a frequency filter. Consider a beam hitting the front cavity mirror from the outside, as in Fig. 3. Most of the light will be reflected back, but a small fraction of it will leak into the cavity. This light is trapped and will continue to build up until the leakage into the cavity (from the incoming

beam) and the leakage of the light inside reach an equilibrium. Suppose that the mirrors are dielectric mirrors with a thickness of $\lambda/4$ where $\lambda$ is a wavelength resonant with the cavity. The promptly reflected beam experiences a $\pi$ phase shift as it reflects off of a denser medium. The light that leaks into the cavity and then back out again, however, will receive a $\pi/2$ phase shift each time it passes through the front mirror, in addition to the $\pi$ phase shift from the reflection off of the rear mirror. If the light is resonant with the cavity, the phase accumulated on each round trip through the cavity is a multiple of $2\pi$. The result is that the promptly reflected light and the light leaking out are out of phase and cancel each other; only the leakage through the rear mirror remains. This means that in steady state all the incoming light is transmitted through the cavity. When the light is off resonance however, the light leaking into the cavity never builds up as it interferes destructively with itself, consequently most of the light is reflected back.

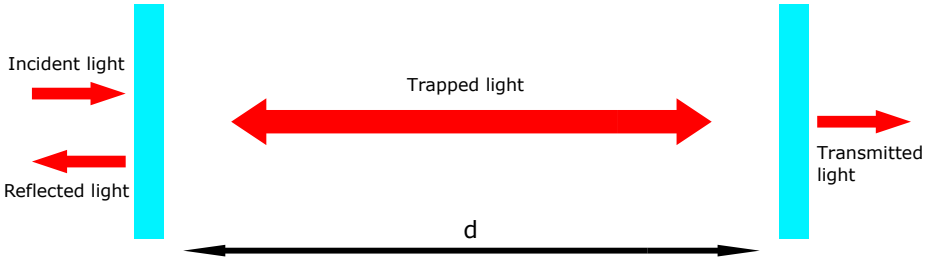


Figure 3: The reflected light is the sum of the initial reflection of the incident light and the leakage of the light from inside the cavity.

To express this mathematically we need to sum up the contributions from all the E-fields from all the reflected beams. The first transmitted beam accumulates, in accordance with Eq. 1, a phase shift of $2\pi d/\lambda$ during its trip through the cavity. The subsequent beams will have an additional $4\pi d/\lambda$ phase shift for each round trip:

$$\begin{aligned} E_t &= TE_0e^{i2\pi d/\lambda} + RTE_0e^{6\pi id/\lambda} + R^2TE_0e^{10\pi id/\lambda} + ... = \\ &= TE_0e^{i2\pi d/\lambda}\sum_{n=0}^{\infty}(Re^{i4\pi d/\lambda})^n = \frac{TE_0e^{i2\pi d/\lambda}}{1 - Re^{i4\pi d/\lambda}} \end{aligned} \tag{6}$$

where $E_t$ is the total E-field of the transmitted light, $E_0$ is the E-field of the incoming light and $R$ and $T$ are the mirror reflectance and transmittance respectively. Doing the same calculation for the reflected light gives:

$$E_r = \frac{\sqrt{R}E_0(e^{i4\pi d/\lambda} - 1)}{1 - Re^{i4\pi d/\lambda}} \tag{7}$$

### 2.1.3 Coupling and Mode Matching

To couple light into a cavity it is not enough to match the frequency of the light with the cavity's resonance frequency; the spatial properties of the light

also need to be considered. Different modes in a cavity can have different spatial distributions, these are called transverse modes (see Fig. 4). The coupling efficiency can be represented by an overlap integral over the incoming field and the various modes supported by the cavity [19]. This can be thought of as a projection of the beam onto the eigenmodes of the cavity. Unless the incoming light closely matches one of the cavity modes spatially, it will be smeared out over several of them [4]. A consequence of matching the transverse mode of a cavity is that the curvature of the wavefront must also match the curvature of the cavity mirrors at the point where the two meet. This is illustrated in Fig. 5.

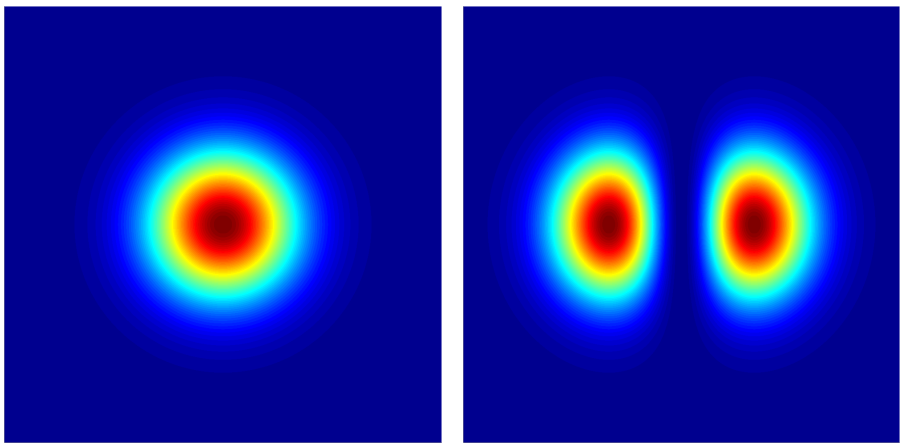

(a) A fundamental ($TEM_{00}$) Gaussian mode.

(b) A $TEM_{10}$ Hermite-Gaussian mode.

Figure 4: This figure shows the intensity distribution of two different cavity modes, perpendicular to the propagation direction of the beam.

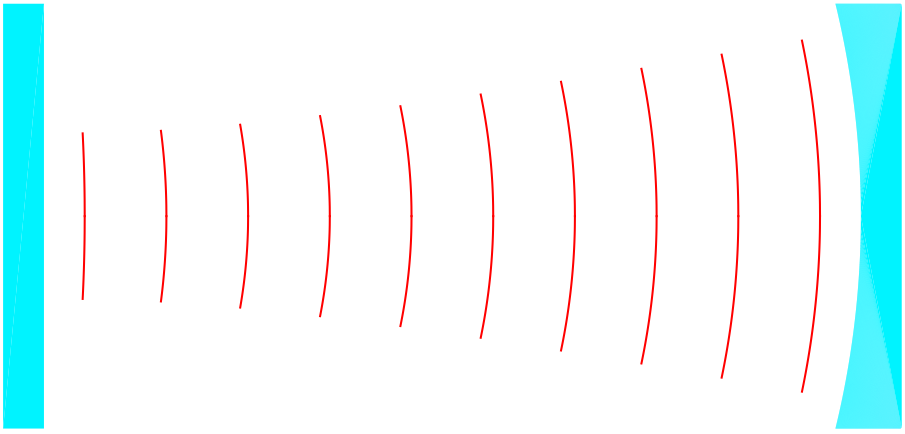

Figure 5: For resonant light inside a cavity the curvature of the wavefront matches the curvature of the mirrors at the point at which it is reflected.

## 2.2 Linewidth of a laser

In essence, a laser consists of a gain medium (that amplifies the light) placed inside an optical resonator. Hence the linewidth of laser light is related to the linewidth of the corresponding optical resonator. Depending on the type of laser in question the specifics of the cavity can vary quite a lot, but in general one cannot use very high finesse cavities for a laser since they only transmit a minimal amount of light. In addition to the limits imposed by the mirror reflectances, the output of a laser is also affected by various sources of noise, for example changes to the refractive index in parts of the cavity or mechanical vibrations in the setup. From Eq. 2 it follows that if the length of the cavity experiences a small change $\Delta d$ the wavelength and frequency change accordingly:

$$\frac{\Delta d}{d} = \frac{\Delta \lambda}{\lambda} = -\frac{\Delta \nu}{\nu} \tag{8}$$

This means that even small vibrations can significantly broaden the linewidth. As an example, consider a laser with a frequency of 600 nm which has a 20 cm long cavity. If it is subjected to mechanical vibrations with an amplitude of 1 µm the corresponding frequency fluctuations would be 5 GHz, whereas the linewidth of a laser cavity can easily be a few MHz, three orders of magnitude smaller.

## 2.3 Frequency Stabilisation

The idea behind active frequency stabilisation of a laser is to rapidly measure the frequency and then use mechanical or optical components to compensate for any fluctuations. For example, the cavity mirrors might be attached to a motor that can adjust their positions. The optical version of this, altering the optical path length instead of the physical length of the cavity, will be discussed in the next section. It turns out that the limiting factor in this kind of scheme is typically the accuracy with which the frequency error can be measured and that is also the aspect which this thesis concerns. With a good stabilisation setup it is possible to achieve laser linewidths much smaller than the linewidth of the laser cavity.

### 2.3.1 Electro-Optic Modulators

Certain materials exhibit a phenomena known as the Pockels Effect, which is a linear change in refractive index in response to an applied electric field. An Electro-Optic Modulator (EOM) is a component that uses this effect to allow for very rapid phase modulation of light [1]. EOMs have two important uses in laser stabilisation, the first being to rapidly correct the light inside the laser cavity and the second being the generation of sidebands, which will be discussed in Sec. 2.4.

Laser stabilisation setups often involve an EOM placed inside the laser cavity. By varying the voltage applied to the EOM the optical path length of the cavity changes. This can be used to counteract any changes in the physical

length of the resonator. It can, importantly, also be used to tune the frequency of the light inside the cavity. When light passes through the EOM it is delayed, since it passes through a medium with higher index of refraction. This time delay can equivalently be thought of as a phase shift. By applying a linearly increasing voltage to the EOM the refractive index, and in turn the phase shift, increases (approximately) linearly with time. A linearly increasing phase shift is equivalent to a change in frequency:

$$\sin\left(\frac{2\pi x}{\lambda} - \frac{x\varphi}{c}\right) = \sin\frac{2\pi x}{\lambda'} \tag{9}$$

Here $\varphi$ is a phase shift, that increases linearly with $x$. This is what happens in an EOM. It's worth looking at this frequency shift in a bit more detail. Suppose that the refractive index of the crystal in the EOM is $n_0$ and that the applied voltage makes the refractive index increases by $k$ s$^{-1}$:

$$n(t) = n_0 + kt \tag{10}$$

During a trip through the EOM, the average refractive index is:

$$\langle n\rangle = \frac{1}{2}(n(t_1) + n(t_0)) = \frac{1}{2}(2n_0 + kt_1 + kt_0) \tag{11}$$

Where $n_0$ is the refractive index in the crystal with no electric field applied, and $t_0$, $t_1$ are the initial and final times respectively. If the crystal inside the EOM has the length $L$ the time delay of a point on a wave traveling through the EOM with the voltage applied compared to one without is:

$$\tau = \frac{L\langle n\rangle}{c} - \frac{Ln_0}{c} = \frac{Lk}{2c}(t_1 + t_0) = \frac{Lk}{2c}(2t_0 + \Delta t) \tag{12}$$

where $\Delta t = t_1 - t_0$. The relative time delay between two points, after the EOM, is then (see Fig. 6):

$$\Delta\tau = \tau' - \tau = \frac{Lk}{2c}(2t_0' - 2t_0 + \Delta t' - \Delta t) \approx \frac{Lk}{2c}(2t_0' - 2t_0) = \frac{Lk}{c}\Delta T \tag{13}$$

Where $\Delta T = t_0' - t_0$ is the time delay between the points before the EOM and I used that $\Delta t' - \Delta t << 2t_0' - 2t_0$. This approximation says that the time it takes to propagate through the EOM ($\Delta t$) varies slowly with time. Since the maximum change in the refractive index induced by the Pockels Effect will generally be no more than on the order of $10^{-5}$ [1], and the fractional change in travel time will be similarly small, this is a good approximation unless the EOM is modulated exceptionally quickly (on the order of PHz).

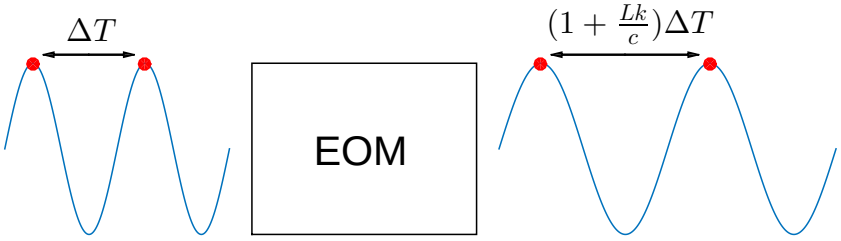


Figure 6: When the voltage applied to the EOM increases linearly, the added time delay between two different points on a wave after it has passed through the EOM ($\Delta\tau$) is proportional to the time delay between those points before the EOM ($\Delta T$). This linearly increasing time delay is equivalent to a (fixed) increase in wavelength.

If the spatial distance between the two points in Eq. 13 is taken to be one wavelength, that is $\Delta T = \lambda/c$, then the change in wavelength induced by the EOM can be written as:

$$\Delta\lambda = c\Delta\tau = \frac{Lk\lambda}{c} \tag{14}$$

If the total (optical) length of the cavity is $d$ then the total number of periods of the standing wave inside is $m = d/\lambda$. Hence the total change in length of the wave is:

$$m\Delta\lambda = m\frac{Lk\lambda}{c} = \frac{dLk}{c} \tag{15}$$

To see how much the length of the cavity changes during the same time period, we use that the time it takes for the entire wave to pass through the EOM is:

$$T = \frac{d}{c} \tag{16}$$

The total change in optical path length during this time is then:

$$LkT = \frac{dLk}{c} \tag{17}$$

Comparing Eq. 15 and Eq. 17 it is evident that the change in wavelength is matched by the change in optical path length of the cavity. If this weren't the case, the EOM would bring the light off resonance as it changed the frequency.

### 2.3.2 Acousto-Optic Modulators

Another way to shift the frequency of light is through an Acousto-Optic Modulator (AOM). Like the name suggests, an AOM works by scattering light off an acoustic wave inside a medium. The easiest way of describing this process is to consider it as a photon absorbing one or more phonons. The phonon energy is $E_m = hf_m$ where $f_m$ is the modulation frequency (the frequency of the acoustic wave) [6]. Energy conservation means that the frequency shift of the scattered light is $f_s = f + mf_m$ where $f$ is the frequency of the incoming light and $m$ the number of absorbed phonons. This process works for acoustic waves well inside the radio-frequency range (hundreds of megahertz) which allows for frequency shifts over a much broader range than an EOM would. However, because of momentum conservation the angle of the scattered light depends on the modulation frequency [1] and in some applications this can be a problem.

## 2.4 Pound-Drever-Hall Locking

By sending light into a high finesse cavity and measuring either the transmission or reflection it is possible to detect very small changes in frequency. However, since the both the transmission and reflection have even symmetry around the cavity's resonance frequency a slight increase in frequency is indistinguishable from a decrease in frequency. There are several ways around this problem, one of them being the so called Pound-Drever-Hall (PDH) method, first described by R.W.P. Drever and J.L. Hall in 1983 [18]. The technique produces a sign sensitive error signal by way of a phase measurement.

### 2.4.1 Sign Sensitive Error

Using Eq. 7, the ratio between the reflected and incoming E-field from a symmetric cavity (one whose mirror reflectances are equal) can be written:

$$\frac{E_r}{E_0} = \frac{re^{i\omega/\nu_{fsr}} - r}{1 - r^2 e^{i\omega/\nu_{fsr}}} = F(\omega) \tag{18}$$

where $r$ is the reflection coefficient of the mirrors, $E_r$ the amplitude of the reflected light, $\omega$ the angular frequency of the light and $\nu_{fsr}$ the free spectral range of the cavity. The complex phase of Eq. 18 is plotted in Fig. 7. What this figure shows is that the phase change in the reflected light is an odd function that jumps 180° around resonance. In a real cavity losses will make the phase shift continuous across resonance. The fact that it is an odd function means it can be used to determine which side of the resonance peak the incoming light is.

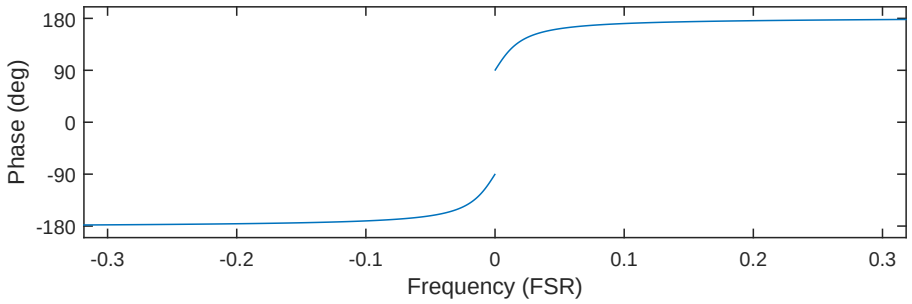


Figure 7: The phase shift upon reflection from outside the cavity as a function of frequency.

### 2.4.2 Sidebands

One would like to measure the phase in Fig. 7. Because photodetectors register light as photons, rather than being sensitive directly to the electric field, the absolute phase of light cannot be measured. Instead the phase of the reflected light must be measured by comparing it to the phase of another field. This is accomplished with the help of an EOM. By sending an AC voltage to the EOM the phase shift of the light passing through it will vary periodically:

$$E(t) = E_0 e^{it(w+\beta \sin \Omega)} = E_0 e^{iwt} e^{i\beta \sin \Omega t} \tag{19}$$

here $\Omega$ is the modulation frequency and $\beta$ is maximum phase modulation, called the modulation depth.

The second exponential factor in Eq. 19 can be written in a Fourier series expansion:

$$e^{i\beta \sin \Omega} = \sum_{k=-\infty}^{\infty} c_k e^{ik\Omega t} \tag{20}$$

The Fourier coefficients $c_k$ are:

$$c_k = \frac{\Omega}{2\pi} \int_{-\pi/\Omega}^{\pi/\Omega} e^{i\beta \sin \Omega t} e^{-ik\Omega t} dt = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{i\beta \sin \tau - k\tau} d\tau = J_k(\beta) = (-1)^k J_{-k}(\beta) \tag{21}$$

where $\tau = \Omega t$ and $J_k(\beta)$ is the k-th order Bessel function of the first kind. Inserting Eq. 20 and Eq. 21 in Eq. 19 yields:

$$E(t) = E_0 e^{iwt} \sum_{k=-\infty}^{\infty} J_k(\beta) e^{ik\Omega t} \tag{22}$$

In other words, applying a sinusoidal phase modulation with the EOM creates sidebands at frequencies $\omega \pm k\Omega$ with amplitudes $J_k(\beta)E_0$. When $k > \beta$ the

Bessel function is small, which means that most of the higher order sidebands can be neglected. In the PDH-method the modulation depth is chosen such that most of the power lies in the carrier (0-th order) and first order sidebands; the higher orders are neglected:

$$E(t) \approx E_0[J_0(\beta)e^{i\omega t} + J_1(\beta)e^{i(\omega+\Omega)t} - J_1(\beta)e^{i(\omega-\Omega)t}] \tag{23}$$

Using Eq. 18 and Eq. 23 the reflected E-field can be written as:

$$E_r = E_0[F(\omega)J_0(\beta)e^{i\omega t} + F(\omega+\Omega)J_1(\beta)e^{i(\omega+\Omega)t} - F(\omega-\Omega)J_1(\beta)e^{i(\omega-\Omega)t}] \tag{24}$$

which then gives the reflected power $P_r = |E_r|^2 = E_r E_r^*$ as:

$$\begin{aligned} P_r &= P_c|F(\omega)|^2 + P_s[|F(\omega+\Omega)|^2 + |F(\omega-\Omega)|^2] \\ &+2\sqrt{P_c P_s}\Big(Re\big(F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)\big)\cos\Omega t \\ &+Im\big(F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)\big)\sin\Omega t\Big) + P_{2\Omega} \end{aligned} \tag{25}$$

Here $P_c = J_0^2(\beta)P_0$ is the carrier power, $P_s = J_1^2(\beta)P_0$ is the sideband power, $P_0$ the total power and $P_{2\Omega}$ is a term containing the beat notes at $2\Omega$. The only relevant terms in Eq. 25 are the sine and cosine terms, as they contain information about the carrier signal's phase; the other terms will be filtered out.

### 2.4.3 The error signal

To generate the error signal, the output from the detector Eq. 25, along with the modulation voltage of the EOM, is sent to a mixer which gives the product of the two. The product of two sinusoidal signals contains one oscillation at the difference frequency and one at the sum of the constituents' frequencies:

$$\begin{aligned} 2\sin(\omega_1 t)\cos(\omega_2 t) &= \sin\big((\omega_1-\omega_2)t\big) - \sin\big((\omega_1+\omega_2)t\big) \\ 2\sin(\omega_1 t)\sin(\omega_2 t) &= \cos\big((\omega_1-\omega_2)t\big) - \cos\big((\omega_1+\omega_2)t\big) \end{aligned} \tag{26}$$

In this case $\omega_1 = \omega_2$, since the oscillation frequency $\Omega$ in Eq. 25 is the EOM modulation frequency. This means that the first term on the right hand sides in Eq. 26 is a constant term, which can be isolated using a low-pass filter. Furthermore the constant sine term vanishes since $\sin 0 = 0$, which means that only one of the two trigonometric terms in Eq. 25 contributes to the error signal. Either one can be selected by controlling the relative phase of the two signals sent to the mixer. With a clever choice of modulation frequency one can distribute most of the power in either the imaginary- or real-part terms in Eq. 25. For example, by choosing the modulation frequency much larger than the linewidth the sidebands will be fully reflected when the carrier is near resonance, $F(\omega \pm \Omega) \approx 1$ (see Fig. 8), which means that:

$$Re\big(F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)\big) \approx Re\big(F(\omega)) - F^*(\omega)\big) = 0 \qquad (27)$$

As a consequence the cos term in Eq. 25 vanishes, and the output from the low-pass filter after the mixer is:

$$\epsilon = LP[P_r \sin\Omega t] = 2\sqrt{P_c P_s}\, Im\big(F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)\big) \qquad (28)$$

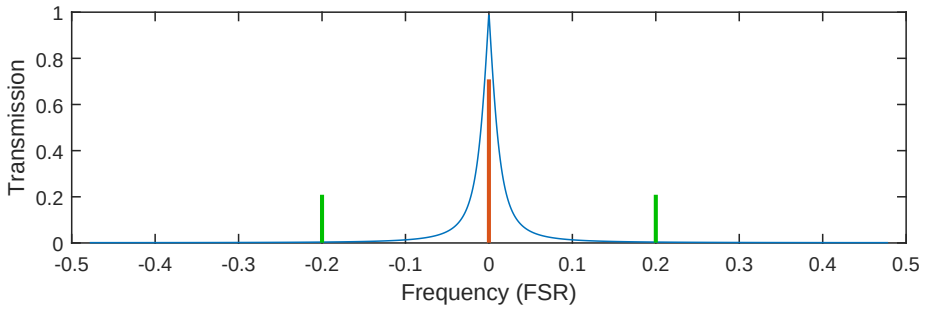


Figure 8: On resonance, the carrier (red) with frequency $\omega$ is transmitted while the sidebands (green) with frequencies $\omega \pm \Omega$ are fully reflected, if the modulation frequency is chosen higher than the linewidth of the cavity.

The shape of this error function depends on the modulation frequency as well as the specifics of the cavity, and can vary quite a lot. An example of the error signal is shown in Fig. 9. Crucially it changes sign depending on which side of the resonance the incoming light is on. It is desirable to maximise the slope of the error signal and aside from tuning the modulation frequency this can also be done by maximising the $\sqrt{P_c P_s}$ term. It is simple to find the maximum of this function numerically, and this results in a modulation depth of $\beta = 1.08$. The corresponding power distribution is that the carrier has 53% of the total power, each sideband has 21.5% and approximately 3% lies in the remaining higher order sidebands.

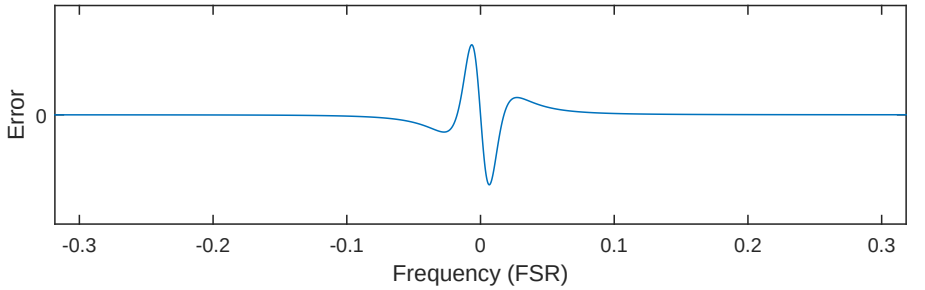


Figure 9: A PDH error signal for a cavity with a finesse of approximately 120 and using a modulation frequency a few percent of the FSR. Close to resonance, the sign of the error depends on whether the light is above or below resonance.

### 2.4.4 Shot noise

When attempting to measure light its quantised nature comes into play. The photons in a beam of light follow a Poisson distribution, which means that during any measurement of light there will some statistical variance in the number of photons detected. These variations in signal are referred to as shot noise. The Poisson distribution has the property that the standard deviation is the square root of the expected value ($\sigma = \sqrt{\mu}$). This means that the signal to noise ratio $\mu/\sigma$ grows as $\sqrt{\mu}$ and as a result shot noise is more important at low intensities.

In situations where shot noise is the dominant source of noise the frequency stability that can be achieved is limited by the shot noise, and this limit can be calculated; it depends on the averaging time, the slope of the error signal and the spectral density of the noise. The spectral density can be approximated as [5]:

$$S = 2\sqrt{\frac{hcP_s}{\lambda}} \tag{29}$$

The shot noise limited frequency stability is then given by:

$$\delta\nu_{sn} = \frac{S}{D\sqrt{\tau_a}} \tag{30}$$

where $D$ is the slope of the error signal in the linear regime near resonance and $\tau_a$ is the averaging time of the signal. This relationship is illustrated in Fig. 10.

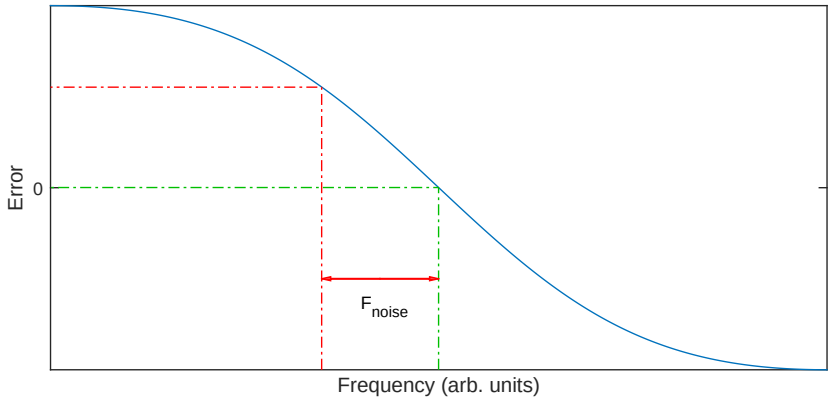


(a)

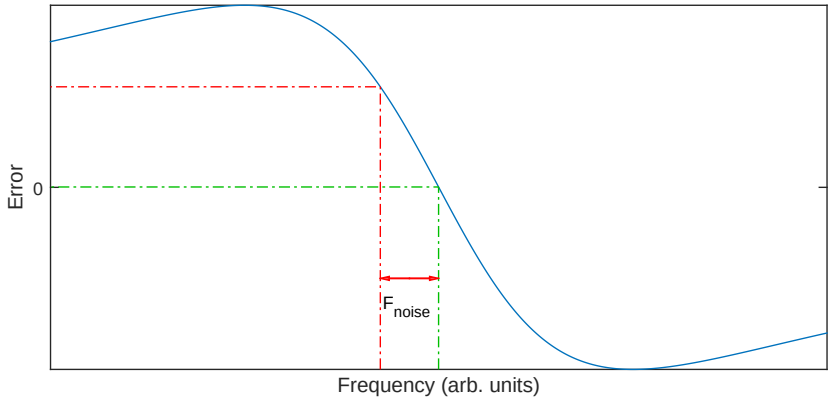


(b)

Figure 10: The blue curve is a PDH error signal. Measurement noise will make it impossible to get the control error to be zero. In the figure this is represented by the red curve. Shot noise of a given intensity corresponds to a certain frequency error (labelled $F_{noise}$). Increasing the slope of the error signal (b) will cause noise of a given intensity to correspond to a lower frequency error, since the red curve (the noise) intersects the blue curve (the error as a function of frequency) closer to the zero crossing (marked by the green curve). The intensity of the noise is the same in both figures.

# 3 Slow Light Cavities

In order to achieve narrow laser linewidths one needs a reference cavity which also has a narrow linewidth. Eq. 5 says that the linewidth is improved either by increasing the reflectivity of the cavity mirrors or making the cavity longer. However, making a long cavity that is also very well isolated from mechanical vibrations and thermal fluctuations is challenging and impractical. A potential way around this problem is using a slow light effect, which causes a short cavity to behave much longer than it is.

## 3.1 Inhomogeneous Broadening

When a crystal is doped with another element the dopant ions are randomly distributed throughout the crystal. The addition of a dopant slightly distorts the symmetry of the crystal which introduces small mechanical strains. These strains have the effect of shifting the energy levels of the ions in the crystal. Since their distribution (therefore strains) is random the corresponding energy shift is different for each ion. As illustrated in Fig. 11, the result is that the absorption spectrum of the ions is broadened. This type of broadening, where the effect on each individual ion is different, is called inhomegeneous broadening.

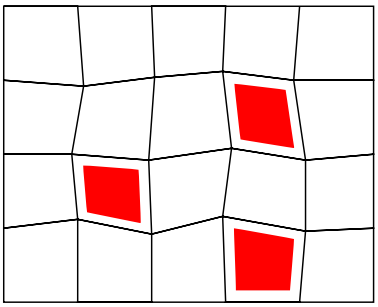

(a) Dopant ions (red) perturb the crystal and deform its structure.

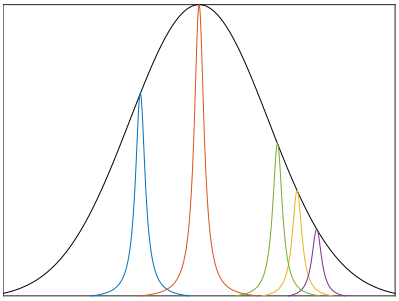

(b) The black envelope is the total absorption profile. The differently coloured curves represent the absorption of ensembles of ions with different energy shifts.

Figure 11: When dopant ions are added to a crystal they slightly deform its structure (a). The result is that the energy level of each ion is shifted (b) and the absorption profile of the entire crystal broadens.

## 3.2 Spectral Hole Burning

Consider the three level system shown in Fig. 12, where both atomic transitions are inhomegeneously broadened. By sending in a pulse of light some of the electrons will be excited, and after some time they will decay back into one of the two ground states with some probability that depends on the transitions. If an electron decays back into its original state it can be excited by another pulse, but if it decayed into the opposite state it won't interact with the light, as long as the bandwidth of the pulse is narrower than the frequency separation of the two ground states. By sending many such pulses all electrons can eventually be moved to atomic states in which they do not interact with the light. The result is called a spectral hole,

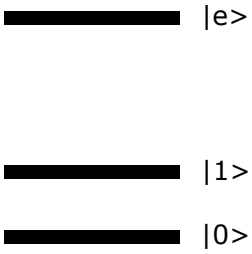


Figure 12: A three level atomic system, with a two lower states $|0\rangle$ and $|1\rangle$, as well as an excited state $|e\rangle$.

it is a part of the inhomegeneous absorption profile where no ions are absorbing. The process is illustrated in Fig. 13.

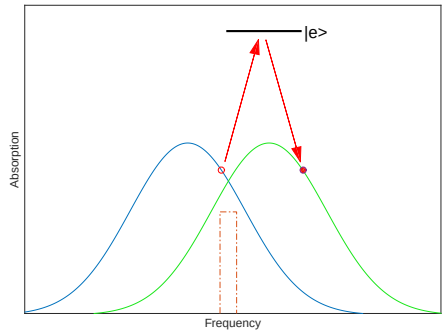


(a) The dashed red curve represents the region being excited by laser light. Electrons occupying states in this region will be excited to the upper state $|e\rangle$, from which they'll randomly decay back either into their original state, or a different state. If they decay into a different state they will remain there, as they are longer be excited by the laser.

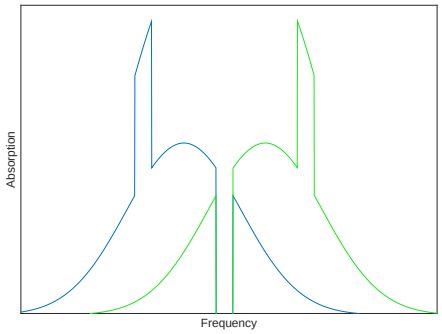


(b) When all the electrons in the region being excited by the laser have decayed into states other than their original one a spectral hole will have been formed. At the same time anti-holes, regions with increased absorption, are formed as well.

Figure 13: An illustration of spectral hole burning, using the same three level system as in Fig. 12. The blue curve shows the inhomogeneously broadened absorption of the $|1\rangle$ state and the green curve shows the same thing for the $|0\rangle$ state.

## 3.3 Dispersion in Specral Holes

Dispersion is a frequency dependence in the refractive index of light in a medium and it is strongly related to the absorption of said medium; a non-absorbing medium cannot be dispersive. If the absorption profile is known it is possible to calculate the corresponding dispersion profile. The index of refraction $n(\nu)$, absorption $\alpha(\nu)$ and the electric susceptibility of the material $\chi(\nu)$ are related via the following equation [1]:

$$n - i\frac{\alpha}{2k} = \sqrt{1 + \chi' + i\chi''} \tag{31}$$

where $\chi'$ and $\chi''$ are the real and imaginary parts of $\chi$ respectively and $k = \omega/c$. The susceptibility is an analytic function, which means that if either the real or imaginary part is known the other can be calculated from the known part through the Kramers-Kronig relations:

$$\begin{aligned} \chi' &= \frac{2}{\pi}\int_0^\infty \frac{s\chi''(s)}{s^2 - \nu^2} ds \\ \chi'' &= \frac{2}{\pi}\int_0^\infty \frac{\nu\chi'(s)}{\nu^2 - s^2} ds \end{aligned} \tag{32}$$

In a weakly absorbing medium the real and imaginary parts of Eq. 31 can be approximated:

$$\begin{aligned} n &\approx \sqrt{1+\chi'} \\ \alpha &\approx -\frac{k}{n}\chi'' \end{aligned} \tag{33}$$

If the medium is a doped crystal in which most of the absorption caused by the dopant atoms Eq. 33 becomes:

$$\begin{aligned} n_0 &\approx n_0 + \frac{\chi'_d}{2n_0} \\ \alpha &\approx -\frac{k}{n_0}\chi''_d \end{aligned} \tag{34}$$

where $n_0$ is the ambient refractive index of the crystal and $\chi_d$ is the susceptibility of the dopant. For a simple rectangular spectral hole the dispersion can be calculated analytically, an example of this is shown in Fig. 14. Outside the hole the refractive index is largely unchanged. Inside the hole however there is a steep gradient in the refractive index, which is linear over most of the hole. Performing the calculation for holes of different widths one finds that the magnitude of the change in the refractive index does not depend on the width, which means that dispersion increases with decreasing hole width.

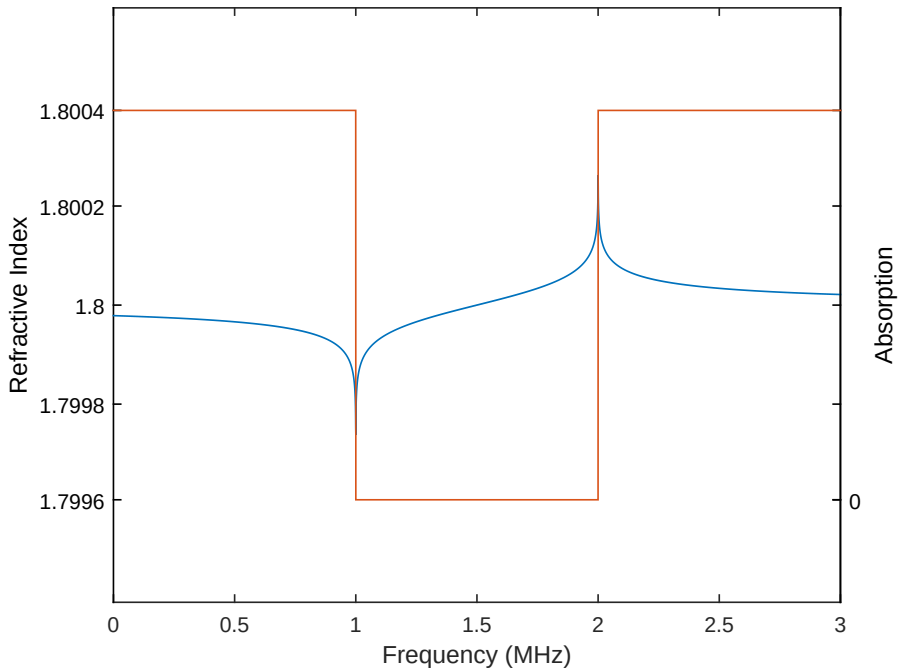


Figure 14: The refractive index (blue) as a function of frequency around a 2 MHz spectral hole. The red curve is absorption.

One way to conceptualise this change in refractive index is to consider how the light interacts with each atom. The E-field of the light in the crystal causes the electron cloud of the atoms it is interacting with to oscillate at the same frequency as that of the light. The oscillation of the electron cloud gives rise to a second electric field which interferes with the incoming light. The phase relationship of these two fields depends on the detuning, which is the frequency difference between the light and the resonance frequency of the electron transition. When the light is resonant with the transition the second E-field is 180° out of phase with the first; they interfere destructively and this corresponds to absorption. A negative detuning causes the emitted field to be ahead in phase, which in turn causes resulting field's phase to be advanced (see Fig. 15). In the case of a positive detuning the situation is reversed and the subsequent successive phase retardation of the light is equivalent to a lower phase velocity, hence higher refractive index. Mathematically this process is described by the Maxwell-Bloch equations [7].

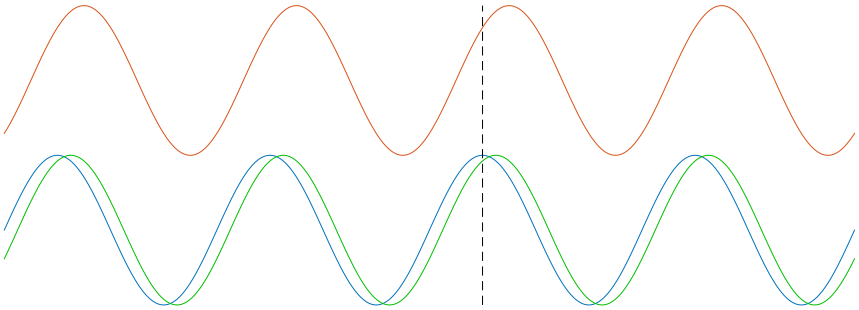

Figure 15: The green wave is the sum of the red and blue waves. Because the red wave is slightly ahead in phase of the blue wave, the sum of the waves is also ahead in phase of the blue wave.

## 3.4 Group Velocity and Slow Light

The propagation speed of the envelope for a pulse of light is called the group velocity and is denoted $v_g$. If all frequency components of the pulse propagate at the same velocity then the pulse will also propagate at this velocity. In a dispersive medium, where the frequency components propagate at different speeds, this is not the case and one can show that the group velocity depends on the slope of the refractive index (see Fig. 16) [5]:

$$v_g = \frac{c}{n - \lambda \frac{dn}{d\lambda}} = \frac{c}{n_g} \tag{35}$$

where

$$n_g = n - \lambda \frac{dn}{d\lambda} = n + \nu \frac{dn}{d\nu} \tag{36}$$

is called the group index, or sometimes group refractive index. A consequence of it being proportional to the slope of the refractive index is that even though the fractional change in refractive index inside a spectral hole might be very small the group velocity can still be strongly affected, as long as the hole is quite narrow so that the refractive index changes rapidly with respect to frequency.

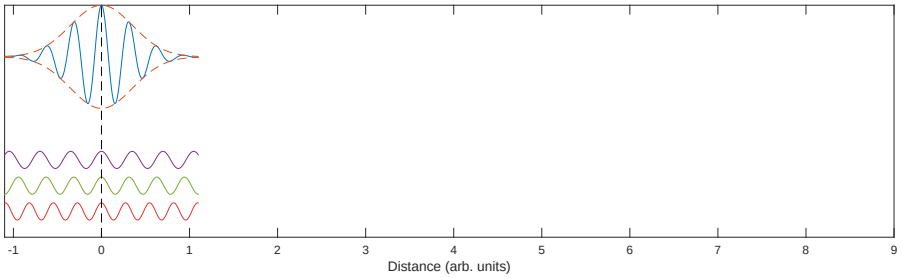


(a) A pulse (blue) at an initial time $t_0$. The dashed orange line is the pulse envelope.

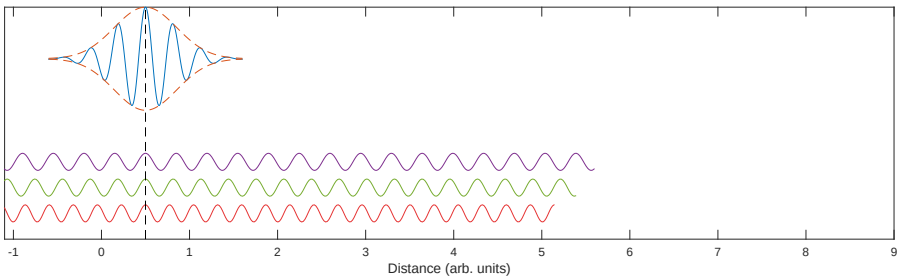


(b) The same pulse as in (a) at a later time $t_1$.

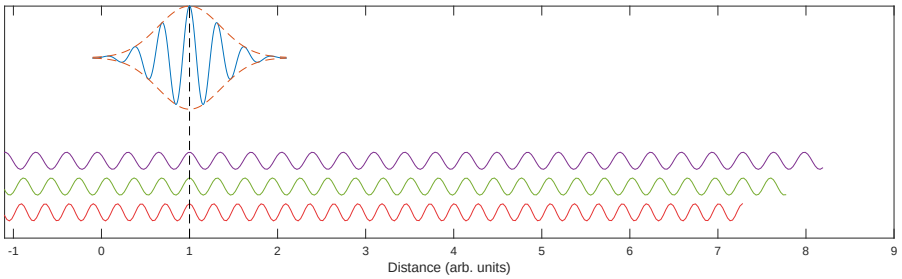


(c) The same pulse again at a third time $t_2$.

Figure 16: The figure shows a pulse and and three of its (possibly infinitely) many frequency components at three different times. Each frequency component propagates at a different speed. The result is that the envelope of the pulse (the points at which the frequency components interfere constructively) propagates slower than the individual frequency components.

## 3.5 Frequency Stability

Eq. 8 says that a change in cavity length gives a proportional change in the resonance frequencies of the cavity. This makes the benefits of using a slow light cavity unclear, as it would appear that the fractional change in cavity length would not be affected by the dispersive effects. A slightly more detailed analysis, however, reveals this not to be the case [23]. By taking the refractive index into account in Eq. 2 and Eq. 3 the resonance condition can be written:

$$\nu = m\nu_{fsr} = \frac{mc}{2Ln(\nu)} \tag{37}$$

where $L$ is the physical length of the cavity, $m$ is an integer representing the mode number and $n(\nu)$ is the refractive index as a function of frequency. Rearranging the terms and differentiating with respect to frequency gives:

$$\frac{dL}{d\nu} = -\frac{mc}{2\big(\nu n(\nu)\big)^2}\left[n(\nu) + \frac{dn(\nu)}{d\nu}\right] \tag{38}$$

Using Eq. 36 this can be written as:

$$\frac{dL}{d\nu} = -\frac{mcn_g}{2\big(\nu n(\nu)\big)^2} = -\frac{mcn_g}{2}\frac{1}{\nu n(\nu)}\frac{1}{\nu n(\nu)} \tag{39}$$

Identifying that $L = \frac{mc}{2\nu n(\nu)}$ simplifies Eq. 39 to:

$$\frac{dL}{d\nu} = -\frac{Ln_g}{\nu n(\nu)} \tag{40}$$

Rearranging the terms one final time yields:

$$\frac{d\nu}{\nu} = -\frac{dL}{L}\frac{n(\nu)}{n_g} \tag{41}$$

What Eq. 41 says is that the fractional change in resonance frequency of a cavity mode, induced by a change in cavity length, is dampened by the dispersion. When the group velocity is low, and the group index consequently high, the cavity will become less sensitive to fluctuations in length. To illustrate the reason for this result more clearly, consider the resonance condition Eq. 37, written in a different form:

$$m\frac{\lambda_0}{n(\lambda_0)} = m\lambda_c = 2L \tag{42}$$

where $\lambda_0$ is the wavelength in vacuum and $\lambda_c$ is the wavelength in the cavity. Suppose that the length of the cavity is perturbed slightly (see Fig. 17, then the resonance condition becomes:

$$m\frac{\lambda_0 + \Delta\lambda_0}{n(\lambda_0 + \Delta\lambda_0)} = m(\lambda_c + \Delta\lambda_c) = 2L + 2\Delta L \tag{43}$$

If the dispersion is strong, the refractive index $n$ changes rapidly with respect to $\lambda_0$. This means that a small change in vacuum wavelength gives a very large change in the wavelength inside the cavity. Consequently, a large change in the resonance condition can be corrected by a small change in vacuum wavelength. This is illustrated in Fig. 18.

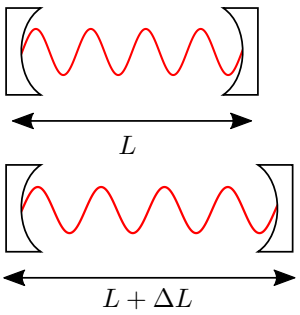


Figure 17: A small change in cavity length gives a corresponding change in resonance wavelengths.

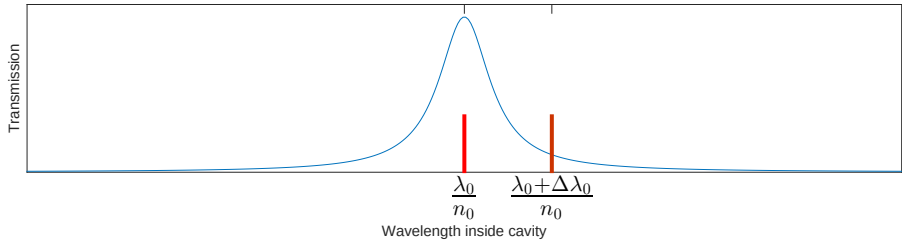


(a) The effect of a small change of vacuum wavelength in a non-dispersive cavity on the wavelength inside the cavity.

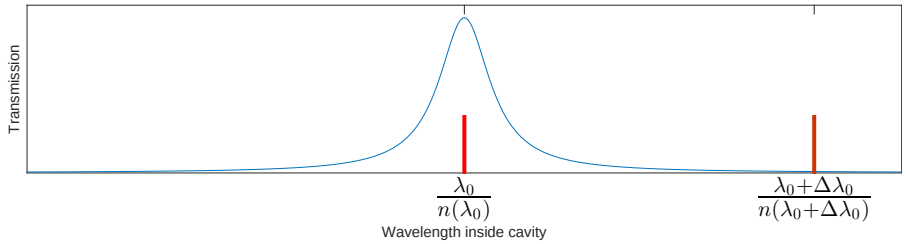


(b) The effect of a small change of vacuum wavelength in a dispersive cavity on the wavelength inside the cavity.

Figure 18: Figure (a) shows how small change in vacuum wavelength affects the resonance condition for a non-dispersive cavity (frequency independent refractive index). (b) shows the same for a dispersive cavity. The wavelength inside the cavity $\lambda_c$ is shown on the x-axis, and on the y-axis transmission peak of the cavity. Because the refractive index decreases rapidly with increasing wavelengths in the dispersive case (see Fig. 14) the wavelength in the cavity changes much faster than the vacuum wavelength. This has the effect of both decreasing the linewidth of the cavity, as it becomes more sensitive to changes in vacuum wavelength, and simultaneously reducing the sensitivity to changes in cavity length, since a big change in the resonance condition can be corrected by a small change in vacuum wavelength. In the figure the effect is greatly understated.

# 4 Measurements on $Pr^{3+}$:$Y_2SiO_5$ Crystal

As preparation for the laser locking experiments to be performed later, a set of hole burning experiments were carried out. The purpose of these experiments was to study how narrow the holes could be made, what resonance linewidths could be expected and what an appropriate readout intensity during the locking process would be, as well as to optimise the burning pulses and build some familiarity with the whole procedure. The experiments in this project were all conducted using a 6 mm long $Y_2SiO_5$ (Yttrium Orthosilicate) cylindrical crystal doped with Praseodymium at a relative concentration of 0.05 %. Except for a small area both ends of the crystal have a 95 % reflective dielectric mirror coating, letting the crystal act as a cavity.

## 4.1 Atomic Levels

The $Pr^{3+}$ ions can be modeled as a six level system. As shown in Fig. 19. there are three hyperfine levels in the ground state, and three hyperfine levels in the optically excited state. There are relaxation processes between these states, which can limit the lifetime of a spectral hole. In the presence of a weak magnetic field (a few mT) however, the time constants for these relaxation processes can be increased from around 100 s to several thousand seconds [12].

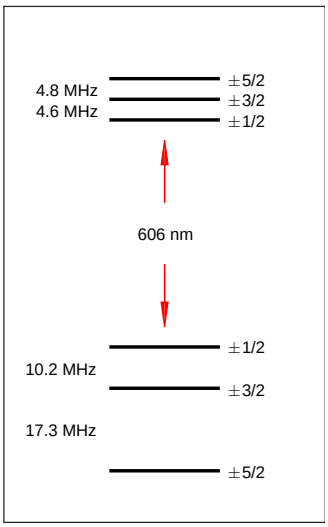


Figure 19: The hyperfine levels in the Praseodymium-doped crystal. The levels are labelled according to the quantum number $I_z$.

The energy levels in the system set a limit for how wide a spectral hole can be, which for $Pr^{3+}$ can be calculated to 18.1 MHz. To efficiently burn a spectral hole one needs to consider the transition strengths of the different transitions, the shape of the pulse etc. Efficient pulse sequences for burning an 18 MHz hole have already been developed by the Quantum Information group in Lund [13]. The pulses for burning narrower pits were based on the work of Philip Dalsbecker, another master's thesis student in the group [8]. These pulses were optimised using a three level system however. In both cases the pulse is a complex hyperbolic secant (sechyp) pulse, which is a hyperbolic secant in time and a hyperbolic tangent in frequency. Different ions with different resonance frequencies (from the inhomogeneous broadening) will have different responses to a pulse, and the sechyp pulse is good at suppressing this effect.

## 4.2 The Experimental Setup

An illustration of the experimental setup is shown in Fig. 20. A key component is the 606 nm tuneable dye laser (shown in Fig. 38), which was used to burn and

probe spectral holes in the crystal. This laser is stabilised to an external ultra-low expansion cavity, permitting laser linewidths down to 10 Hz. Laser pumping is done using a 6 W frequency doubled Nd:$YVO_4$ solid state laser. The output power from the dye laser was typically around 350 mW.

Before being sent to a second optical table using a single mode fibre the light passes through two AOMs. The first is one is a double pass AOM at 210 MHz and the second one is single pass at 360 MHz. Both of these are controlled by an arbitrary waveform generator that can be programmed using a computer in the lab. The AOMs provide the capability to sweep the laser frequency necessary for burning spectral holes.

The second optical table houses a cryostat, which was used to cool the crystal to temperatures as low as 2 K. This is necessary, not only for removing the thermal broadening that occurs at room temperature, but also for increasing the lifetime of the hyperfine states. For details on how the crystal was mounted in the cryostat, see Fig. 21. On this table the beam first passes through a zoom beam expander, which enables tuning of both the size and location of the beam waist. Since the crystal mirrors are flat the wavefronts of the beam also need to be flat when they hit the crystal if the light is to couple into the cavity effectively. For a Gaussian beam the wavefronts are plane in the beam waist. In order to center the beam waist over the crystal the spot size of the beam was measured an equal distance before and after the cryostat. Since the beam is symmetric around the waist it was tuned so that these measurements matched. The beam diameter (at $1/e^2$ fractional intensity) was 1.16 mm, measured 20 cm from the crystal, as close as the cryostat would permit. This corresponds to a 1.15 mm beam diameter at the beam waist.

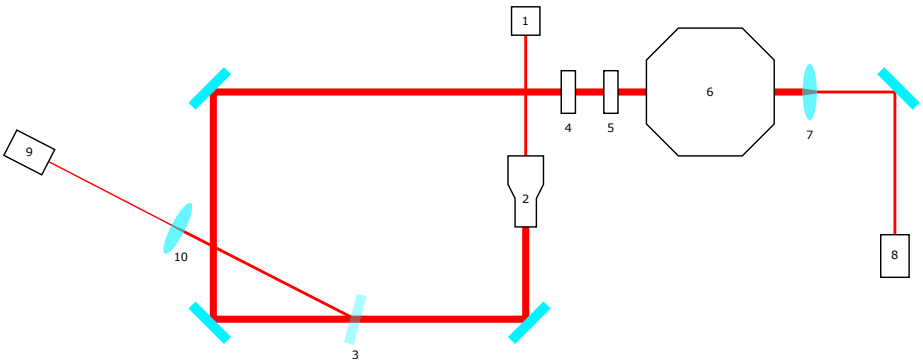


Figure 20: An illustration of the experimental setup. The beam enters the table through the optical fibre collimator (1), then passes through a zoom-beam expander (2) and a beamsplitter (3), from which the first reflection is disregarded. It is then sent through a half-wave plate (4) and a polariser (5) before entering the cryostat (6). The transmitted light is focused by a lens (7) and sent to a detector (8), while the reflected light is picked off by the beamsplitter (3), focused by a lens (10) and sent to a second detector (9).

After being focused, the beam passes through a low effectiveness beam splitter. The exact effectiveness of this beam splitter depends on the angle of incidence

as well as the polarisation and was not measured, but it was assumed to be in the range of $5 - 10\,\%$. The beam then goes through a quarter waveplate that rotates the polarisation and then finally through a polariser before entering the cryostat. The transmission through the crystal is focused and send to a detector after the cryostat, and the reflection is picked off by the beamsplitter and sent to a second detector. The purpose of the waveplate is to rotate the polarisation of the light with respect to the crystal so that it can be aligned along the desired axis. The reason for this is that the absorption is different for the different crystal axes [8]. To maximise the absorption (and therefore dispersion) the light should be linearly polarised and aligned along the axis of maximum absorption [21]. Although the light is already linearly polarised when it comes out of the single mode fibre, the purity of the purity of the polarisation might not be sufficient, hence the need for a polariser.

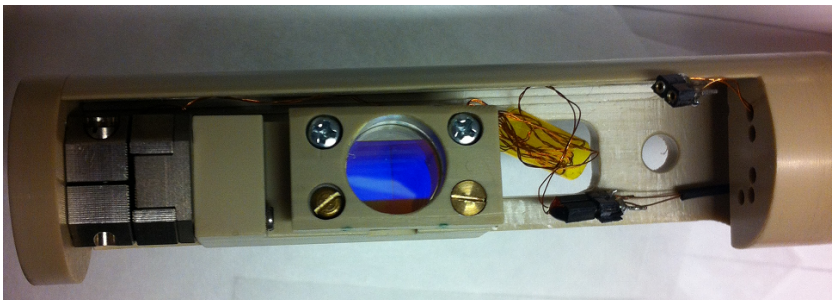

(a)

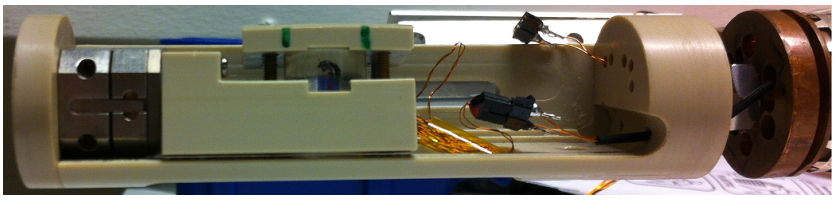

(b)

Figure 21: Top view (a) and side view (b) of the crystal holder. In (a) the blue tinted part of the crystal is the area which has a reflective coating and the smaller area at the top is the uncoated part. The small crystal holder is attached to an Attocube (the metallic object on the left), which allows for translational movement of the crystal along the horizontal axis in the picture. The entire holder is attached to a rod which is lowered into the cryostat.

## 4.3 18 MHz spectral hole

One complication with using a crystal that acts as a cavity is that the transmission through the cavity doesn't correspond directly to the absorption in the crystal. This makes it hard to see the structure of for example a spectral hole. It is for this reason that part of the crystal does not have a reflective coating. To make sure the hole burning was working our experiments began on this part of the crystal. Fig. 22 shows a readout of an 18 MHz spectral hole burned in the

crystal. The oscillation inside the low absorption region is the resonator modes resulting from the inherent 8% reflectivity on the crystal surface combined with the slow light effect which compresses the spectrum.

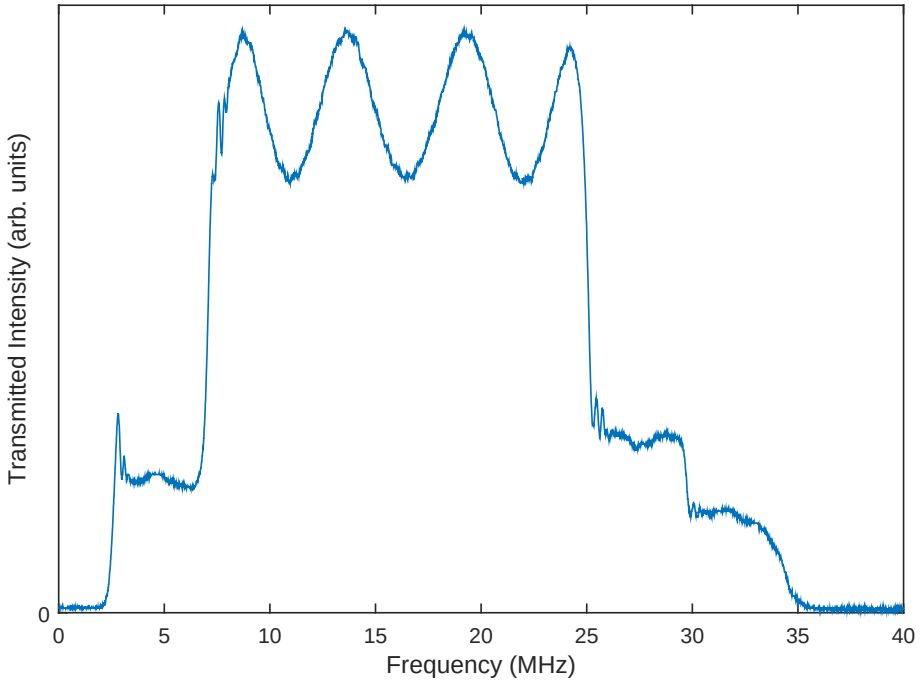


Figure 22: Transmission through the uncoated part of the crystal. The center area with the highest transmission is the spectral hole.

The next step was to perform hole burning on the coated area of the crystal. Before the spectral hole has been burned light entering the crystal will be absorbed before it reaches the rear mirror. This means that during the hole burning process the crystal does not act as a cavity and 95% of the intensity in the hole burning pulses will be reflected. As a result the number of hole burning pulses as well as their intensity needed to be tuned. This was largely a trial and error process. An example of a cavity spectrum captured using an 18 MHz spectral hole is shown in Fig. 23.

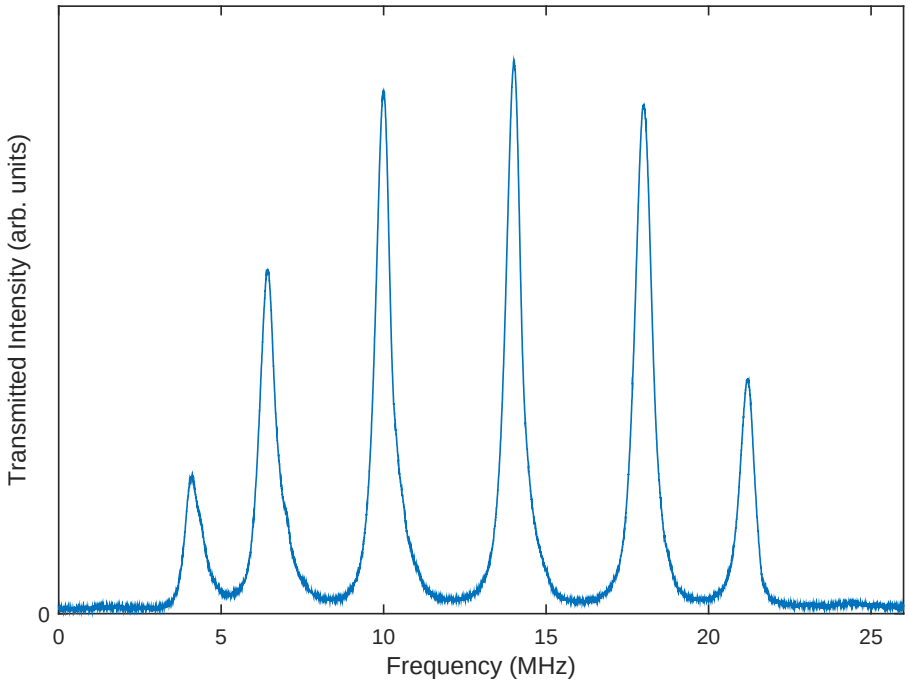


Figure 23: Transmission through the coated part of the crystal, showing the spectrum of the cavity for the 18 MHz hole.

The free spectral range of the spectrum in Fig. 23 is 4 MHz and the FWHM linewidth of the narrowest peak is 0.53 MHz. The corresponding values for this crystal in the non-dispersive regime have previously been measured to 13 GHz for the FSR and 1 GHz for the linewidth [3], which means the dispersive effects have reduced both by over three orders of magnitude.

The reflected signal is of course also of interest, since that is the one used for the control system. The corresponding reflection for the spectrum in Fig. 23 is shown in Fig. 24. Due to the low effectiveness of the beamsplitter used the reflected signal is more noisy than the transmission. It is worth observing that the signal does not fully vanish on resonance. One possible reason for this is that the cavity mode is not matched well enough. Another contributing factor could be the losses due to residual absorption inside the crystal. Also seen in Fig. 24 is a slight asymmetry in the peaks. This is most likely from higher order modes in the cavity [3].

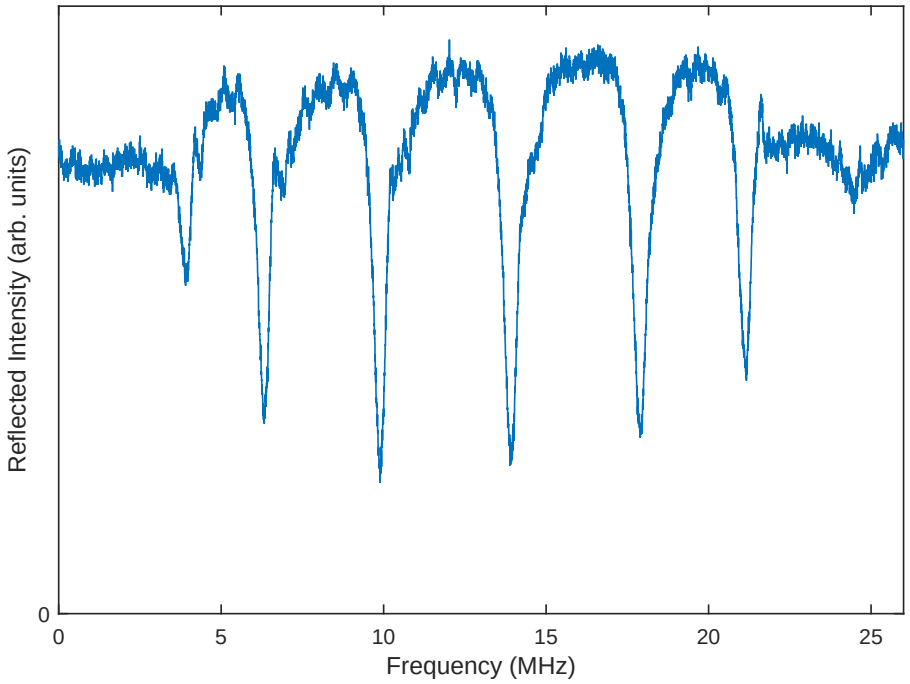


Figure 24: Signal showing the light reflected off of the coated part of the crystal after an 18 MHz pit had been burned.

## 4.4 Narrower Spectral Holes

Since the dispersion grows stronger the narrower the hole is, it is natural to attempt to create smaller structures. One fundamental restriction on the minimum width of a spectral hole comes from the linewidth of the atomic transition. The absorption profile of a single ion has a Lorentzian shape, which means that it is non-negligible even quite far away from resonance. This is illustrated in Fig. 25 and Fig. 26.

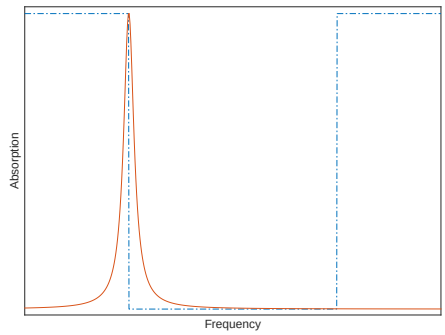


(a) Absorption profile of a crystal (blue) and a single ion (red).

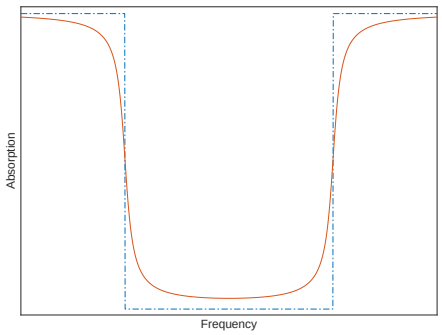


(b) Idealised spectral hole (blue) and linewidth limited hole (red).

Figure 25: (a) shows how the Lorentzian absorption profile of a single ion can extend into the spectral hole when it is sufficiently narrow. (b) shows the shape of a narrow (relative to the ion's linewidth) spectral hole when the linewidths of the individual ions are taken into account.

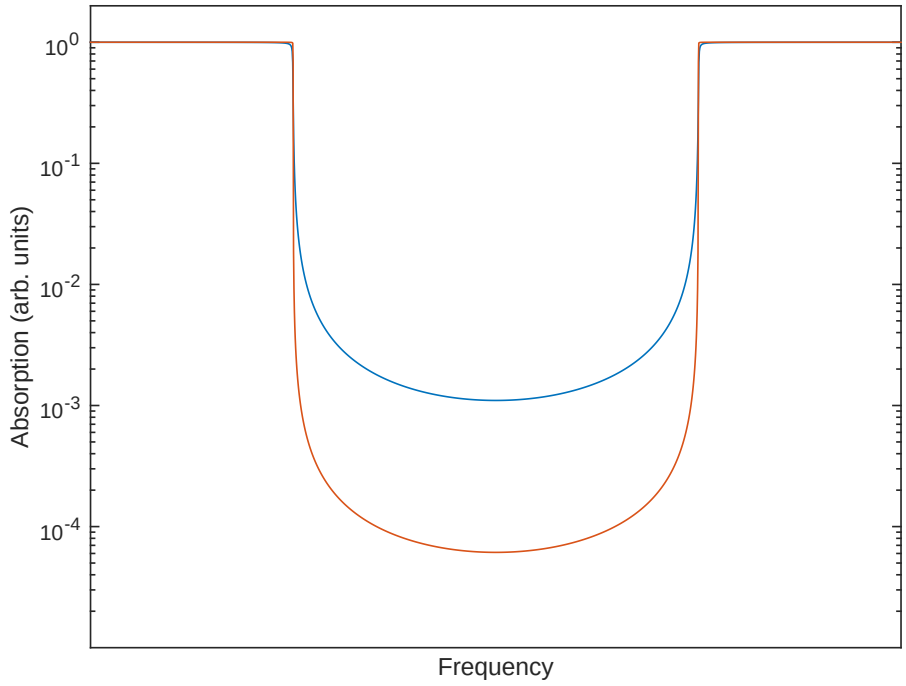


Figure 26: Comparison between the linewidth limited absorption in a 18 MHz hole (red) and a 1 MHz hole (blue). The latter has been stretched in frequency by a factor of 18 when plotting to make the comparison easier. The difference in absorption is also factor 18. This relationship can be found analytically as well.

While attempting to burn narrower holes we found that going below 1 MHz gave

no appreciable reduction in linewdith, although the free spectral range did keep decreasing. An explanation for this is that the hole is still well defined enough to have a strong dispersion gradient, however the absorption inside the hole broadens the resonance peaks, counteracting the effects of the dispersion. Fig. 27 shows a comparison between spectra captured in 1 MHz and 0.5 MHz holes. Finally, shown in Fig. 28 is the reflected signal from a 1 MHz hole. Although the laser is coupling very well into the cavity, as seen by the reflection almost vanishing on resonance, we found that in general it was hard to couple efficiently into the narrow holes.

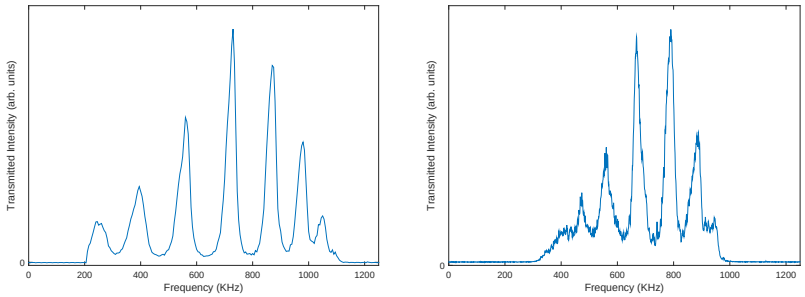


(a) Transmission spectrum of a 1 MHz hole.

(b) Transmission spectrum of a 0.5 MHz hole.

Figure 27: A comparison between a 1 MHz and a 0.5 MHz hole. The FSR in the wider hole is 170 kHz measured for the center peaks, and 120 kHz in the narrower hole. However, the linewidth in both holes is approximately 35 kHz.

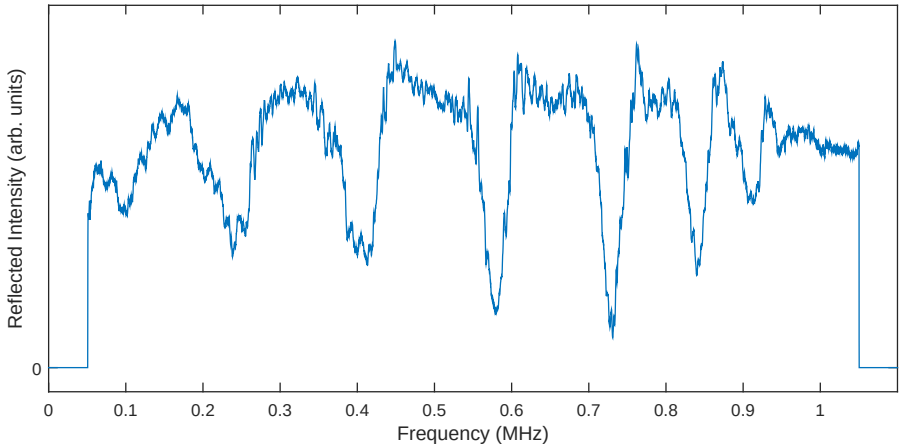


Figure 28: Reflected signal from a 1 MHz spectral hole. The maximum coupling efficiency is almost 90 %.

## 4.5 Time Evolution of the Spectrum

Due to relaxation processes between the hyperfine levels in the ground states the hole will not be truly persistent, but will instead decay over time. Additionally, some of the light used to probe the hole will be absorbed and this may change the burned structure and in turn the resonance peaks. It is therefore important to ensure that the intensity of the light is low enough that it does not affect the hole much on a timescale similar to the lifetime of the hole.

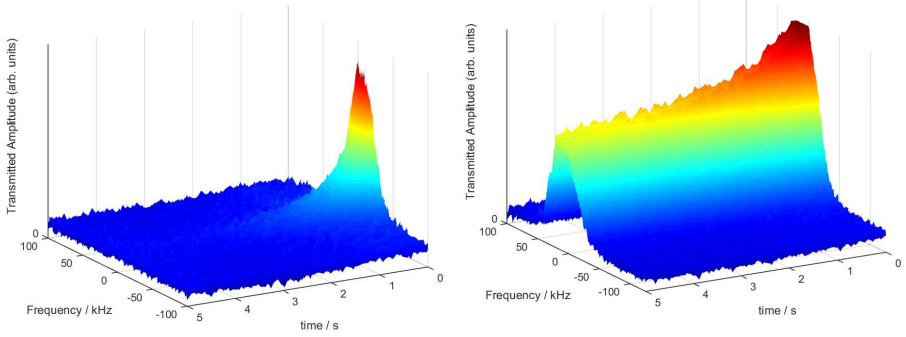


(a) Amplitude of a transmission peak over time with no external magnetic field.

(b) Amplitude of a transmission peak over time with a 10 mT external magnetic field applied.

Figure 29: Comparison between how quickly the peak transmission through a 1 MHz spectral hole decays over time with (blue) and without (red) a 10 mT external magnetic field applied. The figures show that the lifetime of the peaks is increased significantly by the presence of a magnetic field.

Fig. 29 shows a comparison between how the resonance peaks decay with and without a small external magnetic field applied. As expected, the hole decays much slower in the presence of the magnetic field. The lifetime of the peak appears to have increased by factor of around 50. This can be compared with the expected increase in the hyperfine level lifetime, which is approximately a factor of 60. Some readouts showed larger increases in lifetime, which most likely comes from the effect of the probing light on the hole. To measure the effect of the readout pulses on the hole we began scanning a subsection of the hole repeatedly over a shorter period of time. Fig. 30 shows an example of such a measurement, where the intensity of the probing light is too high.

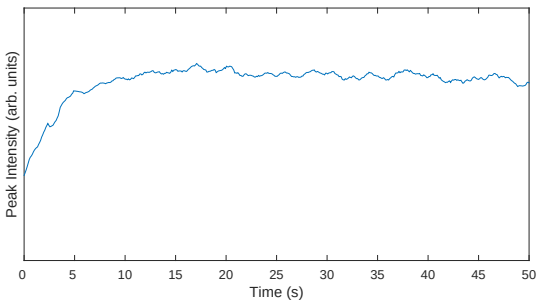


Figure 30: Peak transmission through the hole when the probing intensity is too high. Each probe pulse re-burns the hole. In this measurement a readout was done every 100 ms and the magnetic field was 10 mT.

The result is that the hole is continuously re-burned and doesn't decay. The problem with this is that if one were to attempt to actually lock a laser to a resonance peak using this intensity the hole would not be re-burned, but most likely deformed in some way (and thereby shifting the resonance peaks), since that light would not be scanned over the hole but rather centered on a single frequency. To determine an appropriate readout intensity we gradually lowered it until we were able to scan the hole continuously (starting a new readout as the previous one ends) and still observe that the peaks decayed when the magnetic field was not applied. An example of a readout that satisfies this is shown in Fig. 31.

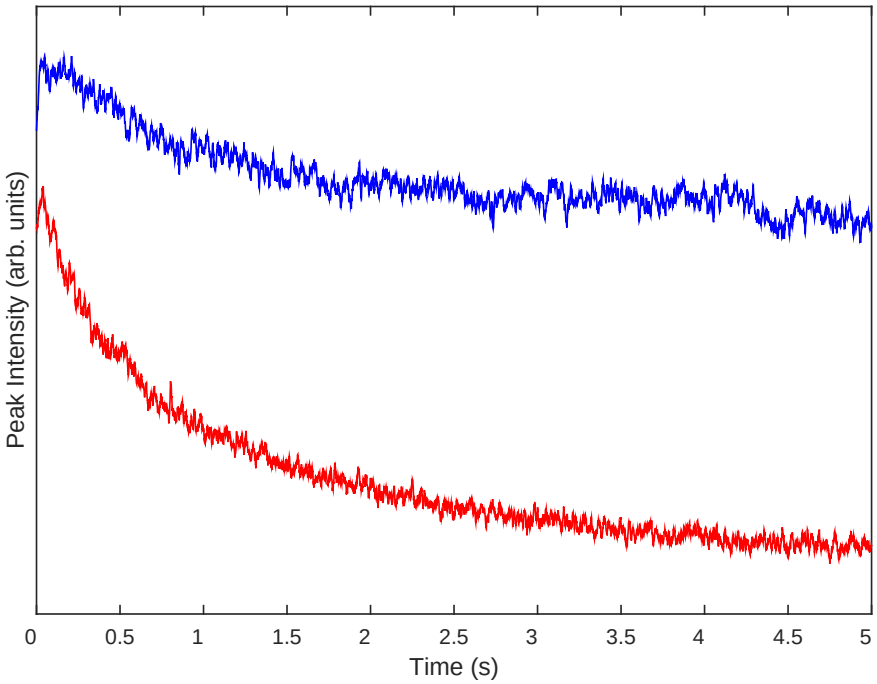


Figure 31: Comparison between how quickly the peak transmission through a 1 MHz spectral hole decays over time with (blue) and without (red) a 10 mT external magnetic field applied, during a continuous readout. When the magnetic field is switched off the transmission rapidly decays. This shows that the probing light does not continuously re-burn the hole. The optical power was 290 nW.

### 4.5.1 Linewidth and Drift

As the spectral hole changes through relaxation or due to the interaction with the probing light, the shape and position of the resonance peaks may also change. In our measurements where the hole was scanned continuously we consistently observed a linear drift in the position of the peaks. An example of this is shown in Fig. 32. One would expect the reflection and the transmission to exhibit the same drift. A possible explanation for the slight discrepancy could be the

relative time delay between the transmitted and reflected light due to the slow light effect. During a readout the laser frequency is increased linearly and this allows us to correlate detection time to frequency. However, if the group velocity, and hence the time delay, in the hole changes over time this would be indistinguishable from a change in frequency.

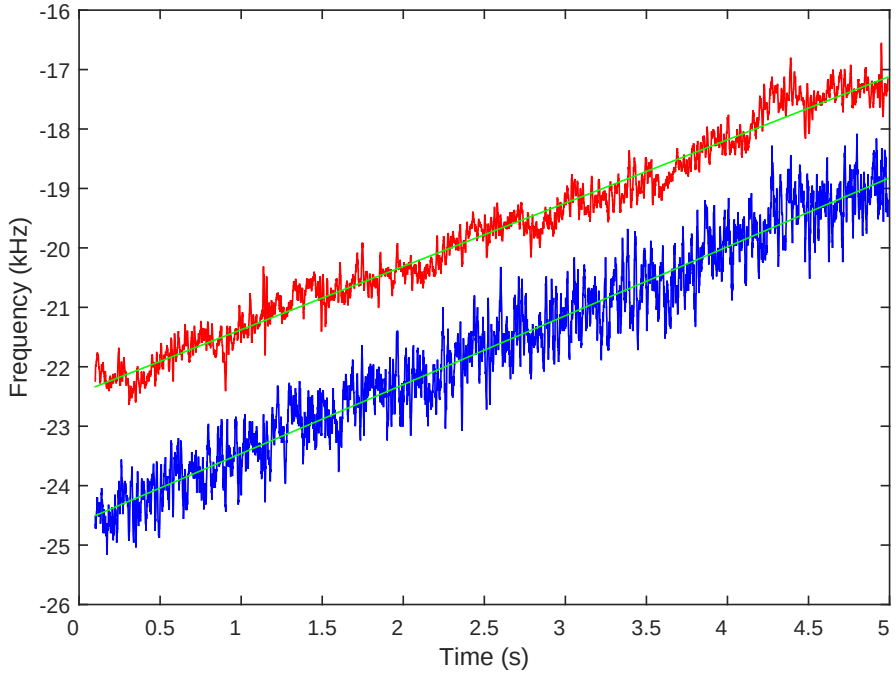


Figure 32: The resonance frequency of a single peak in a 1 MHz hole, as observed in the transmission (blue) and the reflection (red) with a 10 mT magnetic field applied. The green lines represent linear fits to the respective plots. The drift in the transmission is 1.06 Hz/ms and 1.16 Hz/ms in the transmission.

Since the linewidth of the resonance peaks in large part govern the stability that can be achieved it is relevant to see how this property changes over time. If the absorption in the hole increases over time this could lead to a broadening of the peaks. In our measurements, however, we did not observe any noticeable broadening. Fig. 33 shows the transmission around a resonance peak at different times during a continuous readout. Although the intensity of the transmitted light decreases, the shape of the peak remains largely unaffected. A plot of the FWHM linewidth over time is shown in Fig. 34. Any broadening is below the noise threshold of the signal.

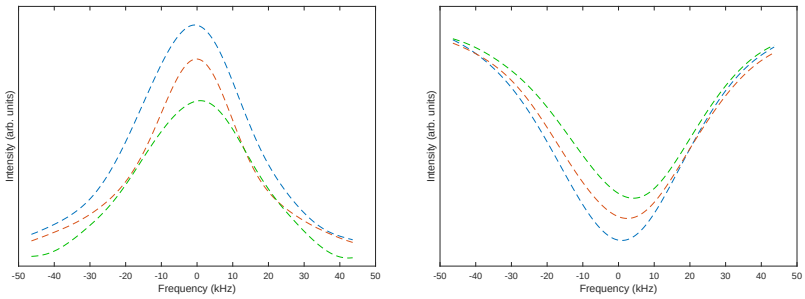


(a) Snapshots of the transmission. (b) Snapshots of the reflection

Figure 33: Snapshots of the transmission (a) and reflection (b) over a single peak during a continuous readout. These snapshots are taken at 160 ms (blue), 1400 ms (red) and 3000 ms (green). The amplitude of the peaks decreases over time but their shape is preserved.

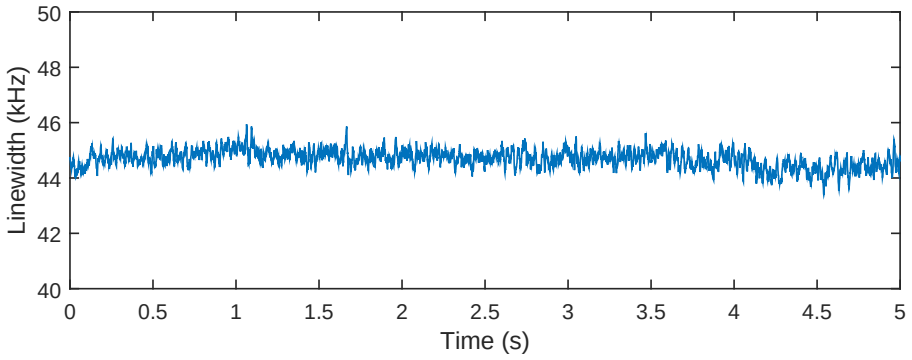


Figure 34: The time averaged FWHM linewidth of a peak in a 1 MHz pit as a function of time, measured in the reflection. The linewidth does not appear to increase over time.

## 4.6 Shot Noise Limit

Using the captured spectra and knowing roughly the intensity needed to not perturb the spectral holes too much, it is possible to calculate what the shot noise limited frequency would be at a given optical power and averaging time. This is shown in Fig. 35. The slope of the error signal was taken from an error signal simulated using Eq. 18 and Eq. 28. The cavity parameters were the same FSR and linewidth as the 1 MHz hole in Fig. 27a. The total optical power used in the calculation was 290 nW, the modulation frequency was 0.5 FSR and the modulation depth was 1.08. The slope of the error signal is:

$$\frac{d\epsilon}{d\nu} = 2\sqrt{P_c P_s}\frac{d}{d\nu} Im\big(F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)\big) = 2.7\,\mathrm{nW\,Hz^{-1}} \quad (44)$$

At 606 nm and $P_s = J_1^2(1.08) \cdot 290\,\text{nW}$ the spectral density of the shot noise is:

$$S = 2\sqrt{\frac{hcP_s}{\lambda}} = 2.9\,\text{pW}/\sqrt{\text{Hz}} \tag{45}$$

The shot noise limit is then found using Eq. 30; this is what is shown in Fig. 35.

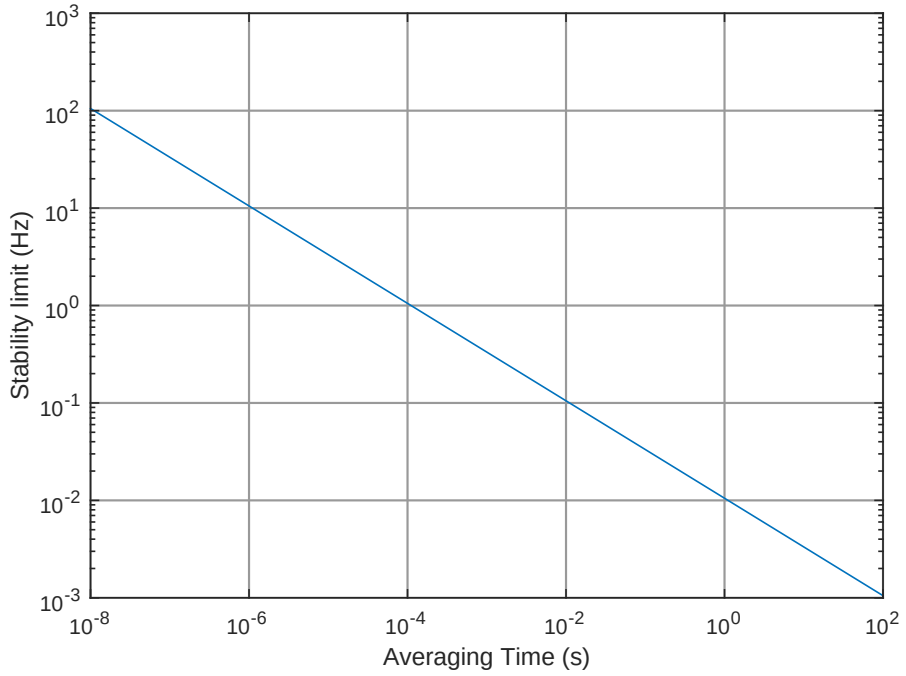


Figure 35: The theoretical shot noise limited frequency stability using different averaging times of the measured signal. As the averaging time increases the effect of shot noise becomes less significant and the stability limit is lowered.

# 5 The Locking System

This section details the electronic components necessary to put together the locking system, and how they were tested.

## 5.1 The components

### 5.1.1 PI-controller

The controller used for our locking setup was an LB1005 PI-controller from New Focus. It can be operated in three modes: 1) with the integrator disabled, giving proportional feedback 2) limited gain mode, where the max gain of low frequency signals is clamped 3) full PI-mode. A knob on the front panel allows for a DC offset to be added to the error signal and two potentiometers on the

back panel give control over the limits of the output voltage. Additionally the integrator can be paused using an external port, it can unfortunately not be reset that way however.

### 5.1.2 Arbitrary Waveform Generator

To generate the sidebands for the EOM a waveform generator was needed. For this we used a TGA1244 arbitrary waveform generator from TTi. A necessary feature of this AWG was the ability to output two phased locked signals with an arbitrary phase offset. This ability is needed to ensure that the two signals fed into the mixer have the correct phase difference, as discussed in Sec. 2.4.3.

### 5.1.3 Voltage Controlled Oscillator

This component outputs an RF signal that can be modulated using an external port. This is the signal that controls the laser frequency. The original plan was to use a VCO constructed by a previous master student in the group, Anders Rönnholm. A limitation of this VCO however was that its output frequency could only be tuned over 250 kHz [9]. This is sufficient to lock a laser to a peak in the narrow 1 MHz hole, but not in the wider hole since linewidth of the corresponding peaks is more than twice the tuning range of the laser.

Due to this limitation we eventually switched to a different component, namely an SML01 signal generator from Rhode & Schwarz, which at 260 MHz had a tuning range of 1 MHz.

### 5.1.4 Low-Pass Filter

For the low-pass filter a Stanford Research Systems SR560 was used. It contains a first order filter with a tunable frequency cut off.

### 5.1.5 High Voltage Amplifier

The high voltage amplifier used to boost the modulation signal sent to the EOM was a New Focus 3211. One limitation of this amplifier was that its output range was limited to $\pm 200$ V.

### 5.1.6 Mixer

The mixer used was a ZAD-3 from Mini-Circuits. It has a 25 kHz to 200 MHz operating range, which covers the modulation frequencies we used.

### 5.1.7 EOM

The EOM used was a PM-25 from Linos. At 633 nm it has a $\lambda/2$ voltage of 200±20 V which means that the phase shift of a beam passing through the EOM will at this voltage be $2\pi/10$.

## 5.2 Testing the Locking System

Since time in the lab is limited, it was necessary to test that these components could be used to assemble a working feedback control system. Initially the plan was to construct an optical cavity and lock a HeNe laser to it. After about two weeks' work however it became apparent that the modulation on the HeNe laser was not working correctly, which meant that it could not be stabilised externally. Instead, the VCO was locked to the AWG. This was a much simpler task, but was nonetheless helpful in learning how to configure all the components.

## 5.3 Testing the EOM

In 2.4.3 the ideal modulation depth of the sidebands was calculated to 1.08 rad. At 606 nm and 200 V the PM-25 EOM produces a phase shift of $\frac{2\pi}{10} \cdot \frac{633\,\text{nm}}{606\,\text{nm}} = 0.656\,\text{rad}$ [2]. The EOM is capable of withstanding higher voltages, but unfortunately the high voltage amplifier could not output more than 200 V. At this modulation depth:

$$\frac{P_c}{P_s} = 8.3 \tag{46}$$

In decibel this is:

$$10 \cdot \log_{10} 8.3 = 9.2\,\text{dB} \tag{47}$$

This means that most of the power is in the carrier. To make sure that the EOM was working as specified we decided to test it. A simple setup, shown in Fig. 36, was constructed. The beam from a HeNe laser is split using an AOM operating at 80 MHz. The AOM is adjusted so that most of the optical power is concentrated in two beams, one being the unshifted transmitted beam and the other being the first order frequency shifted beam. The first order beam is passed through the EOM which generates several sidebands (in this experiment at 80 kHz). The unshifted beam bypasses the EOM and is then recombined with the first order beam using a beamsplitter. The beating between the sidebands and the carrier is then analysed using the FFT function of an oscilloscope. The results of this is shown in Fig. 37. The carrier has around $8.6 - 9\,\text{dB}$ more power than the sidebands, which was slightly lower than expected and taking the average gives a corresponding modulation depth of 0.683 rad. This is within the specified 10% (voltage) range specified by the EOM manufacturer.

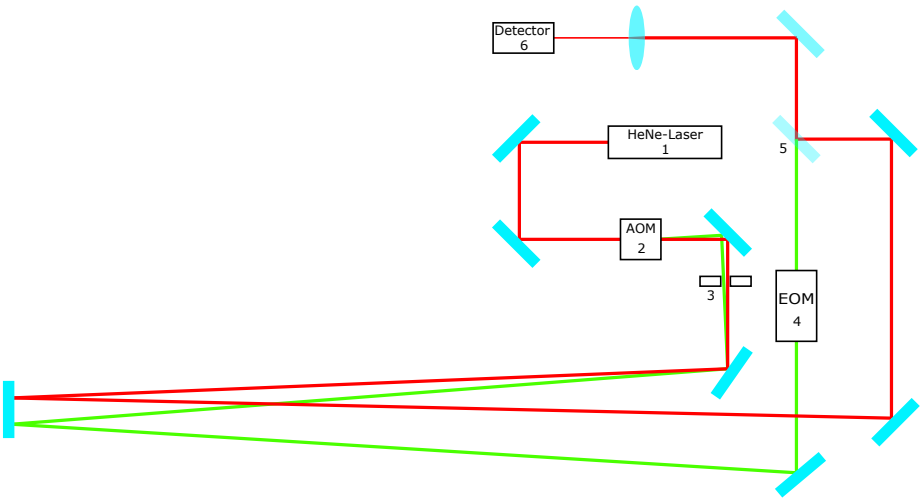


Figure 36: The setup used to test the EOM. A beam from a HeNe laser (1) is passed through an AOM (2). An aperture (3) blocks all but two of the resulting beams, the unshifted (red) and one of the first order beams (green). These are allowed to propagate some distance so that they can be separated spatially. The first order beam is then passed through the EOM (4) while the unshifted beam bypasses it. They are then recombined using a beamsplitter (5) and focused on the detector (6).

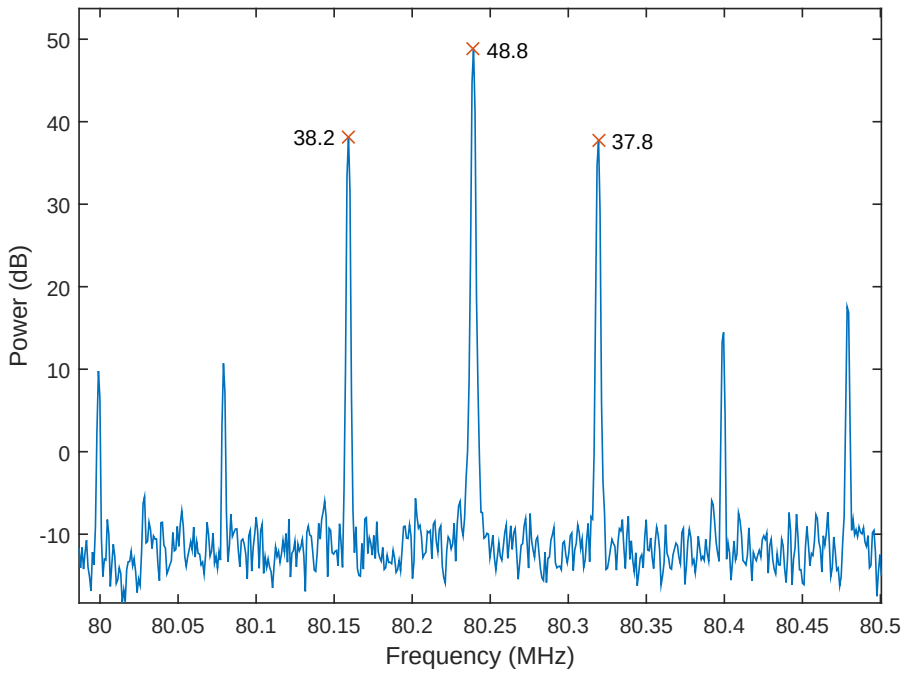


Figure 37: FFT spectrum of the beating between the beam that passed through the EOM and the unmodulated beam. The center peak corresponds to the carrier, and the two adjacent peaks are the first order sidebands. As expected the separation between each peak is exactly 80 kHz.

# 6 Laser Locking Experiments

This second run of experiments was carried out over a week and a half with the intended goal of eventually locking the laser to a resonance peak in the crystal. The first few days were spent arranging and calibrating the optical setup and several days were spent putting together the electronics as well as trying to get a good signal from the crystal. Unfortunately, only the last day was spent trying to lock the laser. The crystal was the same as the one used in the previous experiments.

## 6.1 Experimental Setup

For the second set of cryogenic experiments the setup was slightly more complicated; it is illustrated in Fig. 39 and Fig. 40. The beam followed more or less the same path as in the first set of experiments, the main difference being that it was passed through the EOM (3) before going into the cryostat (7). Additionally, the beam diameter used was smaller, about 0.96 mm, and instead of using a low effectiveness beam splitter, a 50/50 beamsplitter (4) was used to split the beam (the unwanted reflection was blocked (10)). This was done in order to increase the power of the reflected light hitting the detector (11). Unwanted reflections from the polariser (6) and wave plate (5) were blocked by an aperture (31). A strong focusing lens (29) was also added for the second detector (11). The output signal from this detector was fed into into a mixer (14), where it was mixed with a signal from the second AWG (12), phased locked to the EOM driving signal that was passed to the high voltage amplifier (13). The output from the mixer was then passed through a low-pass filter (15), before being sent to the PI-controller (16).

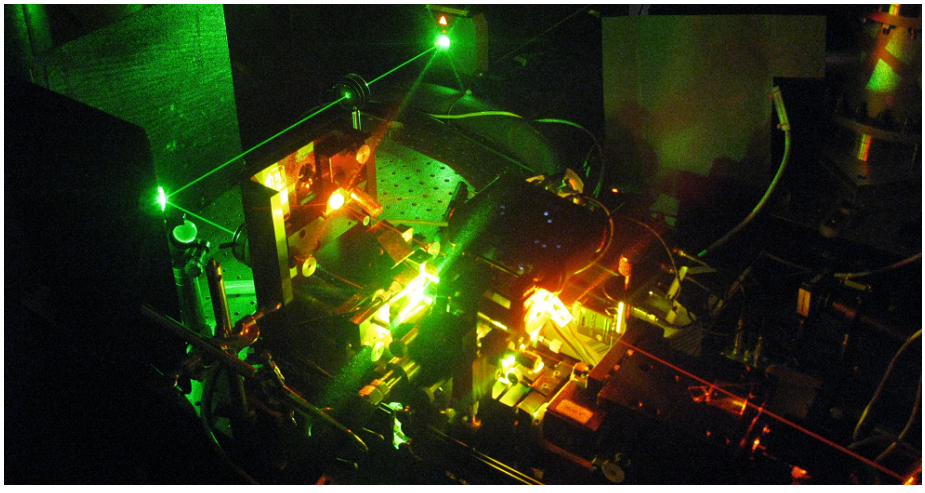

Figure 38: The two lasers used in the setup. The green beam comes from the Nd:$YVO_4$ pump laser, and the more faint red beam is that of the dye laser. Because the setup is contained in a clean room the visible scattering from the beams is not off dust, but is actually Rayleigh scattering.

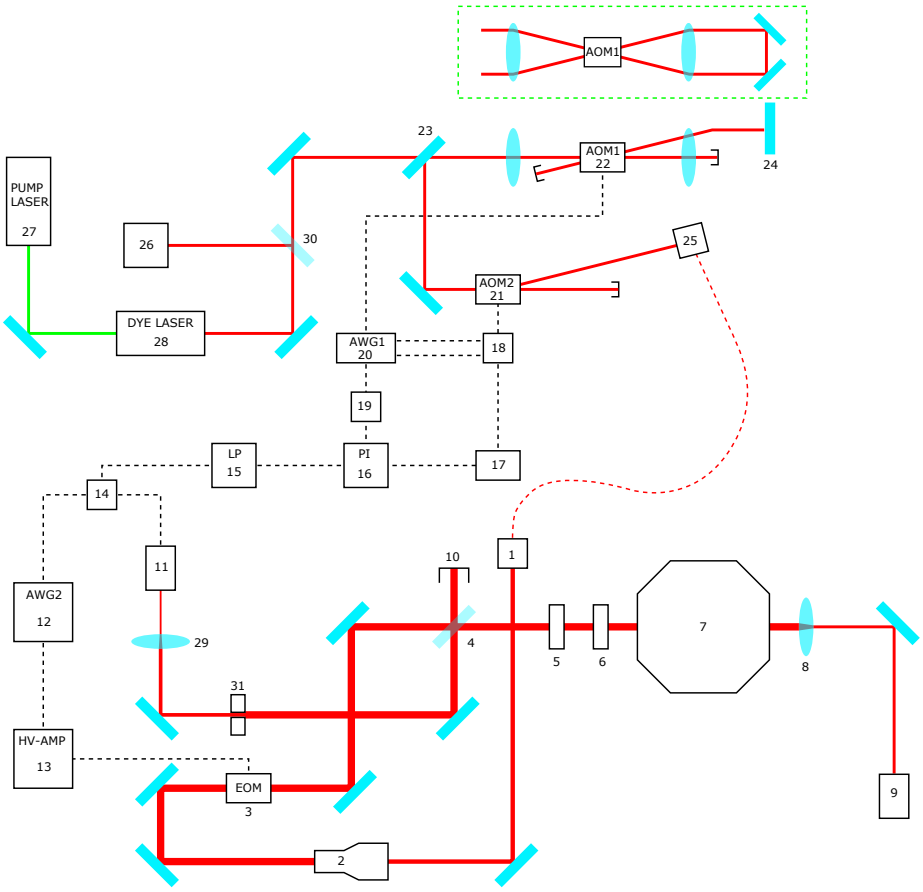


Figure 39: A sketch of the experimental setup used in the second set of experiments. The Nd:$YVO_4$ laser (27) pumps the dye laser (28) which is used for the experiments (see Fig. 38). A beam from this dye laser is split off (30) and sent to a locking cavity (26) that is used in the stabilisation system for the laser. Another beam, the one used for the experiments, passes under a mirror (23) and then makes the first of two passes through the 210 MHz AOM (22). The green cut out shows the beam path from the side as opposed from the top through this part of the setup. The first order beam from the AOM is reflected off two vertically aligned mirrors (24), before going through the AOM a second time, in the opposite direction. It is then picked off by the mirror it passed under the first time (23), passed through the second AOM (21) and is finally sent into a single mode optical fibre (25) which connects to the second optical table and exits through (1). See the text for the rest of the details.

The control signal from the PI-controller (16) was then sent into the modulation port of a signal generator (17), which was driving the single-pass AOM (21). Due to initially choosing a different component for modulating this AOM it had been calibrated to operate at 260 MHz. In addition to being slightly less efficient at this frequency, the modulation range of the signal generator was also halved to 1 MHz. During the hole burning pulses both AOMs (21,22) were controlled by the primary AWG (20), which after completing its sequence sent a trigger

signal to a switch (18) that changed the input of the single-pass AOM to be the signal generator, as well as a signal to unpause the integrator on the PI-controller. This last signal was passed through an inverter (19) before reaching the PI-controller. The setup also included a lens (8) and detector (9) to study the transmission, just as in the previous set of experiments.

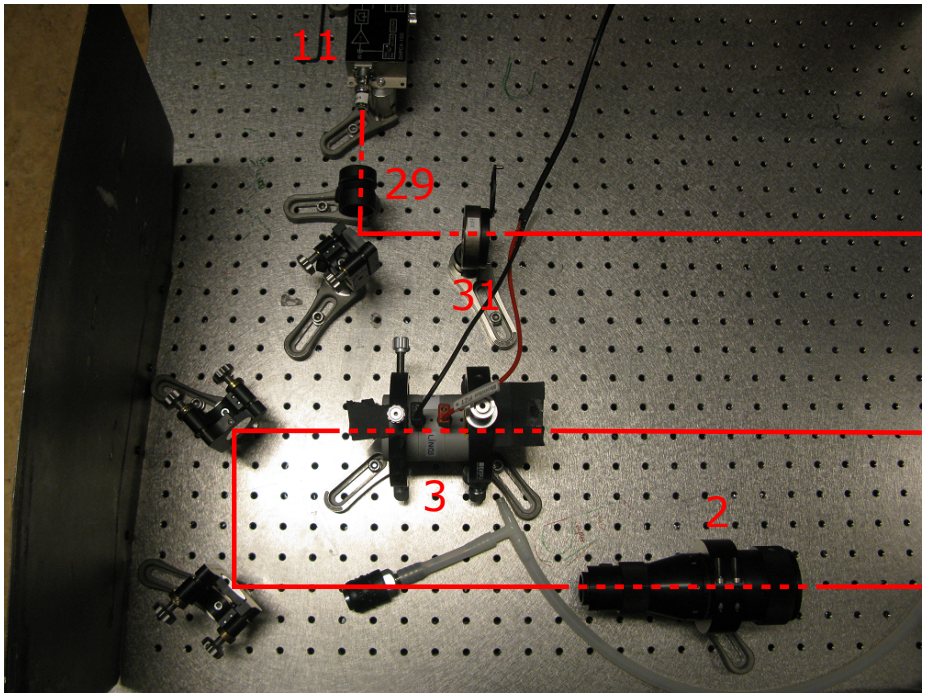


Figure 40: A photo showing part of the setup as well the beam path through it. (2) is the zoom-beam expander, (3) is the EOM, (31) is the aperture, (29) is a focusing lens and (11) is the detector used to pick up the reflection. The small pieces of black paper on either side of the EOM were put there to block reflections from the Brewster windows. For more details see the text and Fig. 39.

## 6.2 PDH Error Signal

Before attempting to lock the laser we wanted to study the error signal. The first step in doing this was to obtain good reflection spectra, similar to Fig. 24. This proved to be more difficult than anticipated. We struggled to get a usable reflection from the 1 MHz pit and this was the reason the VCO had to be swapped out, as discussed in 5.1.3.

To get the error signal, the phase offset between two signals sent out by the secondary AWG was tuned manually until a good signal was achieved. When working with the 18 MHz pit we found that it was difficult to get a good error signal with the sidebands closer than 400 kHz. This was problematic for two reasons; the first was that the cable from the RF amplifier to the EOM was

not designed specifically to accommodate such high frequencies, and as a consequence the modulation signal is attenuated. The second reason is that, as discussed in 4.5, one would like to use as low an optical power as possible when locking the laser. This necessitates increasing the transimpedance on the detector amplifier (to amplify the signal more), but this simultaneously decreases the bandwidth of the detector. On the detector we used to measure the reflected signal the bandwidth at its highest transimpedance setting was 220 kHz. That means that a beating at 400 kHz will be quite strongly attenuated.

In Fig. 41 two examples of recorded PDH error signals are shown. It is worth noting that the signal in Fig. 41b is not centered around zero as one would expect. This was not a problem specific to this sideband frequency however. We noted previously that the shape of the resonance peaks was slightly asymmetric. This, coupled with the fact that the reflected signal does not vanish on resonance, will mean that the error signal does not vanish on resonance either. However, this is not enough to explain why the entire signal is offset.

The control system will of course lock to the zero crossing of the error signal, it is therefore necessary to add a DC offset to cancel the intrinsic offset on the error signal. The PI-controller we used had this option, but it is still far less than ideal as the offset of the error signal was not entirely consistent.

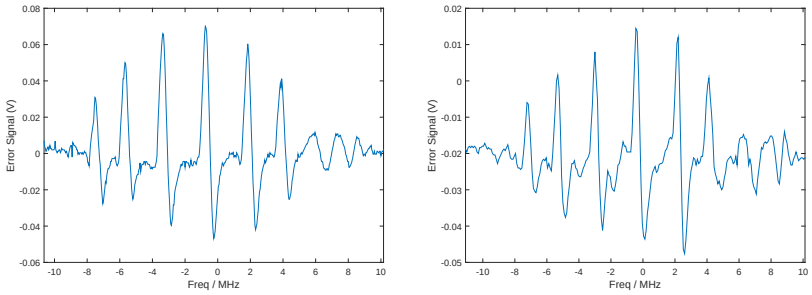


(a) (b)

Figure 41: PDH error signals taken with sidebands at 600 kHz (a) and 400 kHz (b) in an 18 MHz wide spectral hole. The higher frequency sideband signal shows a nicer behaviour between the resonance peaks. For frequencies lower than 400 kHz the signal started to deteriorate.

Although the 1 MHz hole proved to be very uncooperative we still captured an error signal, shown together with the reflection in Fig. 42. The sidebands were placed at 60 kHz for this measurement. The error signal in Fig. 42b has the wrong sign, but this can be corrected simply by changing the phase offset of the signals sent to the mixer by 180°.

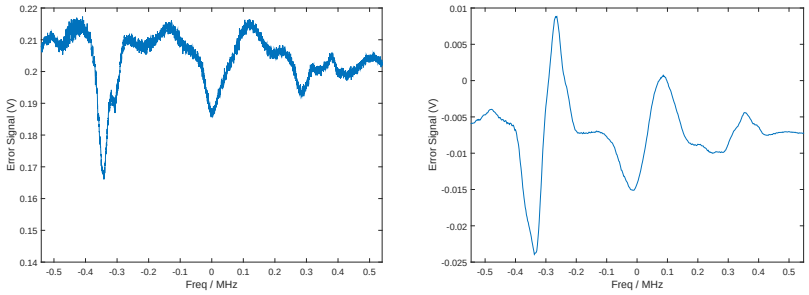


(a) (b)

Figure 42: Direct reflection (a) and PDH error signal (b) in a 1 MHz wide spectral hole.

## 6.3 Error Signal with Control Applied

The final part of the experiments was to attempt to lock the laser to one of the resonance peaks. One of the limitations of the setup used was the ±500 kHz modulation range on the signal generator controlling the 260 MHz AOM. Since the FSR in the 18 MHz pit was around $3 - 4$ MHz the signal generator could not be reliably used to find a resonance peak, and if one were found the control action would potentially be very limited. After several unsuccessful attempts at locking the laser using this approach, we decided to sweep the laser using the primary AWG. The advantage of this was that the laser could be swept over a large enough range to guarantee that a peak would be found, but a significant downside was that the sweep pulses had to be pre-programmed. In practice this meant that the sweep would continue past a resonance peak and the locking system, limited to its ±500 kHz modulation range, would not be able to keep up.

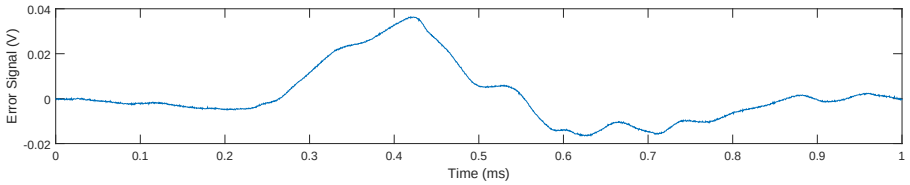


(a)

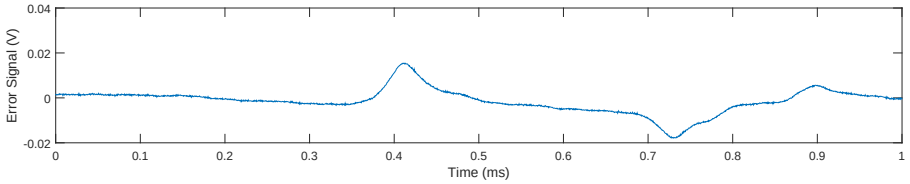


(b)

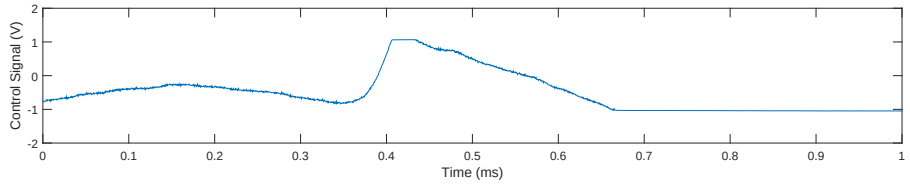


(c)

Figure 43: A comparison between how the error signal looks like during a scan, with (b) and without (a) control applied, as well as the control signal (c) during the former scan.

An example of such a scan with control applied is shown in Fig. 43. Although the control system is not able to keep the error at zero throughout the scan, it is clear that it is still able to drastically reduce the magnitude of the error. At 4 ms the PI-controller reaches its saturation value, corresponding to a 500 kHz and shift can't compensate any further. The AWG keeps scanning the laser, and as it finds a resonance peak the controller begins applying control action in the opposite direction; the linearly decreasing control signal from $4.5 - 6.5$ ms counteracts the linear scan by the AWG. Finally, the laser is scanned too far away from resonance and the control signal is saturated once again.

# 7 Conclusions and Outlook

A locking system was constructed and generation of PDH error signals from modes in a slow light cavity as well as working control action from said signals was demonstrated. The laser was never fully locked to one of the transmission peaks, but we were not given reason to believe this could not be done given more time and perhaps a slightly more suitable choice of electronics. This includes a signal generator with a broader modulation range, a controller that can be reset remotely, as well as a high voltage amplifier sufficiently powerful to achieve the desired modulation depth with the EOM.

Resonance peak linewidths were decreased by a factor 30 000 compared to the non-dispersive case and were improved by burning narrower spectral holes, but we found that once going below 1 MHz the absorption from the ions outside the hole deteriorated the spectrum. A way to go below this limit would be to use a crystal doped with $Eu^{3+}$ instead of $Pr^{3+}$. $Eu^{3+}$ has dramatically (almost two orders of magnitude [11],[10]) longer excited state lifetimes, something which allows for narrower transition linewidths. $Eu^{3+}$ also offers the benefit of significantly longer hole lifetime, up to several days [22].

Narrower holes would, however, reduce the bandwidth of the control system.The reason for this is that the sidebands need to be inside the spectral hole, hence a narrower hole means the modulation frequency on the EOM needs to be lowered. This in turn means that the low-pass filter must be set to a lower cut off frequency to filter out the $2\Omega$ term in the output from the mixer. One possible way around this would be two additional holes, at intervals that match the modulation frequency.

Another way of further lowering resonance peak linewidths would be using a crystal with much higher reflectivity mirrors. Even with 95 % reflectivity mirrors, however, the amount of burn pulses (therefore time) required started to become an issue, due to most of the power in the burn pulses never entering the cavity. When locking a laser to a crystal with a longer hole lifetime, such as $Eu^{3+} : Y_2SiO_5$, this would not be a problem, as the hole would only need to be re-burned infrequently. However, for rapid iteration during experiments it could be solved by using a second beam, incident from the side where the crystal does not have a reflective coating, to burn the spectral holes.

Ways to couple into the cavity more efficiently and reliably need to be investigated. This could include studying the mode structure in the cavity, to allow for a more thoughtful optimisation of the beam parameters, or manufacturing a crystal with curved rather than plane mirrors so that its modes would more closely resemble Gaussian modes. Although much work remains to be done both in assessing and developing this technique for the application to ultrastable lasers, these early results indicate that the potential exists.

# References


[1] Saleh, B.E.A., Teich, M.C. 2009. *Fundamentals of Photonics.* 2nd ed. Wiley.

[2] Zhao, X. 2013. *Diode laser frequency stabilization onto an optical cavity.* Master's thesis. Division of Atomic Physics, LTH.

[3] Sabooni, M., Li, Q., Rippe, L., Mohan., R.K., Kröll, S. 2013. *Spectral Engineering of Slow Light, Cavity Line Narrowing, and Pulse Compression.* Phys. Rev. Lett. 111, 183602

[4] Siegman, A.E. 1986. *Lasers.* Oxford University Press.

[5] Black, E.D. 2001. *An introduction to Pound-Drever-Hall laser frequency stabilization.* Am. J. Phys. 69, 79

[6] Kittle, C. 2005. *Introduction to Solid State Physics.* 8th ed. Wiley.

[7] Rippe, L. 2006. *Quantum computing with naturally trapped sub-nanometre-spaced ions.* PhD thesis. Division of Atomic Physics, LTH.

[8] Dalsbecker, P. 2015. *Development of narrow-bandwidth filters for the suppression of scattered light for optical and ultrasound analysis of tissue.* Master's thesis. Division of Atomic Physics, LTH.

[9] Rönnholm, A. 2012. *Fiber noize [sic] cancellation.* Master's Thesis. Division of Atomic Physics, LTH.

[10] Equall, R.W., Cone, R.L. 1995. *Homogeneous broadening and hyperfine structure of optical transitions in* $Pr^{3+} : Y_2SiO_5$. Phys. Rev. B. 52, 3963

[11] Böttger, T., Thiel, C.W., Sun, Y., Cone., R.L. 2006. *Optical Decoherence and Spectral Diffusion at* 1.5 µm *in* $Er^{3+} : Y_2SiO_5$ *versus Magnetic Field,Temperature, and* $Er^{3+}$ *Concentration.* Phys. Rev. B 73, 075101

[12] Ohlsson, N., Nilsson, M., Kröll, S. 2003. *Experimental investigation of delayed self-interference for single photons.* Phys. Rev. A 68, 063812

[13] Amari, A., Walther, A., Sabooni, M., Huang, M., Kröll, S., Afzelius, M., Usmani, I., Lauritzen, B., Sangouard, N., de Riedmatten, H., Gisin, N. 2010. *Towards an efficient atomic frequency comb quantum memory* J. Lumin. 130, 1579

[14] Abott, B.P. et al. 2006. *LIGO: the Laser Interferometer Gravitational-Wave Observatory.* Rep. Prog. Phys. 72, 076901

[15] Poli, N., Oates, C.W., Gill, P., Tino, G.M. 2013. *Optical Atomic Clocks.* Riv. del Nouvo Cim. 36, 555

[16] Diddams, S.A., Bergquist, J.C., Jefferts, S.R., Oates, C.W. 2004. *Standards of Time and Frequency at the Outset of the 12st Century.* Science. 306, 1318

[17] Kessler, T. et al. 2012. *A sub-40-mHz-linewidth laser based on a silicon single-crystal optical cavity.* Nature Photon. 6, 687

[18] Drever, R.W.P. et al. 1983. *Laser Phase and Frequency Stabilization Using an Optical Resonator.* Appl. Phys. B 31, 97

[19] Uehara, N., Gustafson, E.K., Fejer, M.M., Bye, R.L. 1997. *Modeling of efficient mode matching and thermal-lensing effecton a laser-beam coupling into a mode-cleaner cavity.* Proc. SPIE. 2989, 57

[20] Hinkley, N. et al. 2013. *An Atomic Clock with $10^{-18}$ Instability.* Science. 341, 1215

[21] Sabooni, M., Nilsson, A.N., Kristensson, G., Rippe, L. 2015. *Wave propagation in birefringent materials with off axis absorption or gain.* arXiv:1507.01076

[22] Könz, F. et al. 2003. *Temperature and concentration dependence of optical dephasing, spectral-hole lifetime, and anisotropic absorption in $Eu^{3+}$ : $Y_2SiO_5$* Phys. Rev. B 68, 085109

[23] Rippe, L. 2015. Personal communication.