Copenhagen University

Master Thesis

Cavity Enhanced Molecular Clocks and Their Applications

Christian Zhou Raahauge

Quantum Optics and Ultracold Atoms Group
Niels Bohr Institute, Copenhagen University
Denmark

Author: Christian Zhou Raahauge

Cavity Enhanced Molecular Clocks and Their Applications

Supervisor: Jan W. Thomsen

Submitted on the 6th of August 2018

Abstract

We use the Noise Immune Cavity Enhanced Optical Heterodyne Molecular Spectroscopy (NICE-OHMS) method to lock the frequency of a laser to the $P(16)v_1 + v_3$ transition in $13C_2H_2$ with a wavelength of 1542.4 nm. We manage to show an Allan deviation of the order of one kHz for averaging times of one second, but with a large drift for longer times.

Resume

Vi bruger Noise Immunce Cavity Enhanced Optical Heterodyne Molecular Spectroscopy (NICE-OHMS) metoden til at stabilisere frekvensen af en laser til $P(16)v_1+v_3$ transitionen i $13C_2H_2$. Vi opnår en Allan deviation i størrelsesordenen kHz for en midlingstid på 1 sekund, men med drift over længere tidsskalaer.

Acknowledgements

Endless thanks goes to my supervisor Jan Thomsen. Despite a busy schedule as the head of the Niels Bohr Institute he has always found time to help me when it was needed.

I would like to thank Martin Romme Henriksen for doing most of the work on with this project, Asbjørn Arvad Jørgensen for proofreading my thesis and for fixing our problem with thermal drift, and Stefan Alaric Schäffer for invaluable help in matters both theoretical and practical, and for letting me borrow all of his books.

I would like to thank the other members of my laboratory group, and I would like to thank the *JURA* coffee machine company for keeping me caffeinated.

Lastly I would like to thank my wife, Lunyi Raahauge, for her patience with my weird daily rhytm during the final days of this thesis.

Contents

Chapter 1

Introduction

Tic-Toc. A noise that has followed mankind since medieval times. Among all noises in the world, this one has a special property: It has the exact same frequency whereever it is encountered. It has a frequency of 1Hz, meaning there should be one second between the *tic* and the *toc*. How precisely a watch keeps time is determined by how exactly this frequency is one Hz.

An original mechanical clock would subdivide a day and a night into twenty four hours, each hour into sixty minutes and each minute into sixty seconds. Each second would be counted using some mechanical oscillator.

Rush forward to the modern time, and the mechanical oscillators, those swinging or hanging pendulums, get replaced by oscillating atoms. So precise has our ability become to measure those tiny oscillators, that the second now no longer is defined as 1/60 of 1/60 of 1/24 of a day, but instead as the time it takes the light stemming from a certain transition in a Caesium atom to vibrate 9,192,631,770 times.

And so also has the challenge for the fundamental clock maker change from finding precise pendulums and springs into finding ways of keeping a ray of light exactly on resonance with an atom or a molecule.

In this project, we will investigate making an infrared laser locked on resonance with a transition in Acetylene molecules at a wavelength of 1542.3837 nm, corresponding to a fequency of 194.4 THz.

This is not an attempt to put our precision into the range of the currently most precise atomic clocks, but instead to make a compact and simple frequency reference at an interesting wavelength.

Such a reference would have a number of possible applications. It could be used for high precision spectroscopy in the frequency range close to its own frequency, through the use of an optical comb. Since our clock can be made very compact, the volume of it is currently mainly determined by the electronics we use, the concept of it could be used for making frequency references for space travel. This thesis is structured in the following way:

First this introduction.

Chapter 2 is a short description of how the frequency of a laser can be locked to an error signal, followed by a description of how this error signal can be generated using heterodyne spectroscopy, and finishing of with an introduction to our method of Noise Immune Cavity Enhanced Optical Heterodyne Molecular Spectroscopy.

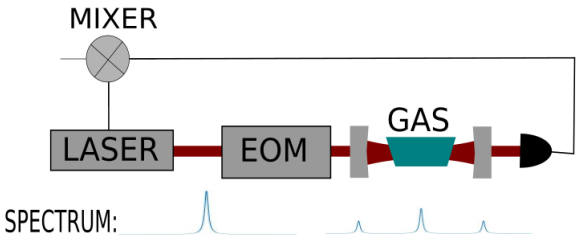

Figure 1.1: The concept we have used. Light from a laser with a narrow frequency is modulated in a way to produce three different frequencies. All three frequencies will produce a standing wave in an optical cavity, in which is placed a container with a specific gas. The output from the cavity is detected with a photo detector. The three different frequencies will make a beat note. A frequency mixer produces a DC electric signal which is proportional to this note. Interaction between the gas in the cavity and the different waves will change the amplitude of the beat note. By using the signal from the frequency mixer, the laser can be kept continuously on resonance with the molecules in the gas. And anything with a constant and precise frequency can in principle be used for making a clock.

Chapter 3 is a description of our actual setup, including both an experiment with a special designed cavity with an extra broad beam, and our current setup using simpler linear cavities with narrower beams.

Chapter 4 includes some of our results and discussions of those, with measurements of the linewidth of our transition and the slope of the generated error signals, along with a measurement of current precision.

Chapter 5 is a short conclusion, and an outlook to how the setup could be improved, and an estimation of possible future precision.

Chapter 2

Molecular Optical Frequency Standards

This chapter will deal with the concept of how spectroscopy can be used for making a frequency reference. It will start by telling how a laser can be tuned using an optical cavity. It will then deal with the concept of how any tuneable system can be stabilised if one can generate some sort of *error signal*. It will then quickly cover the concept of vibrational transitions in molecules, and spectroscopy of those, reaching the point of how an error signal can be generated with heterodyne spectroscopy, and introducing the concept of NICE-OHMS. At last there will be an explanation of the concept of *Allan deviation*, the standard way of comparing frequency references.

2.1 Optical Cavities

A central figure in our experiment is optical cavities, so we will start this thesis with a quick review of them.

An optical cavity is, to put a blunt point on it, two (or more) reflecting surfaces set up to face each other. (With a bit finer point on it, there do exist micro cavities for nanophotonics, which cannot really be said to have reflecting surfaces).

The magic of an optical cavity happens if at least one of the mirrors are made only partially reflective, and a beam of light is shone on it from the outside.

In that case a light ray will build up between the two mirrors, which can be much larger than the incident ray.

There are two principles that control how much of the incoming light can be caught between the mirrors. The first is whether the light can make an entire round trip between the mirrors in some whole number of wavelengths. The second is whether the incident light's shape matches the shape of a possible *modes* in the cavity.

An optical cavity can, based on its geometry, contain a number of modes. To contain a mode means that a beam of light with the geometry of that mode will be reflected around in the cavity, from mirror to mirror, without being reflected out(see figure 2.1).

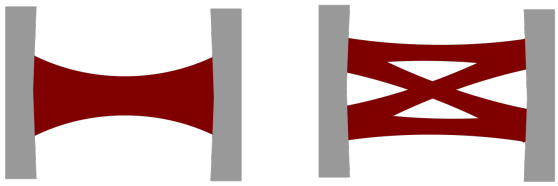

Figure 2.1: Two examples of optical cavities, with different possible light beams (modes) standing in them

When light matching the mode is steadily shone into a cavity, the field inside the cavity will grow to a level that depends on the detuning from the resonance frequency of the cavity of the light as a lorentzian function.

An optical cavity can be characterized by two numbers; its finesse and its Free Spectral Range, abbreviated FSR.

The Free Spectral Range, is, for a given mode, the difference in frequency between two possible standing waves. It can be calculated using the length of the cavity;

$$\text{FSR} = \frac{c}{2l}, \tag{2.1}$$

where l is the total length of the cavity, and c is the speed of light.

The finesse F of a cavity is fully given by how much light is lost pr round trip of a beam, either through imperfect reflection, misalignments of mirrors or other reasons:

$$F = \frac{\pi}{2\sin^{-1}(\frac{1-\sqrt{p}}{2\sqrt[4]{p}})}, \tag{2.2}$$

Where p is the percentage of the power in the beam which remains after having taken one round trip.

An optical cavity that loses 1% power pr round trip would have a finesse of around 600.

The finesse also gives the linewidth with which the optical cavity responds to different frequencies. The finesse times the FSR of the cavity gives the Full Width at Half Maximum (FWHM) of the lorentzian(see figure 2.2).

$$I_{cavity} \propto I_{in} \frac{1}{(\omega_{in} - \omega_{cavity})^2 + (\frac{1}{2}F \cdot FSR)^2} \tag{2.3}$$

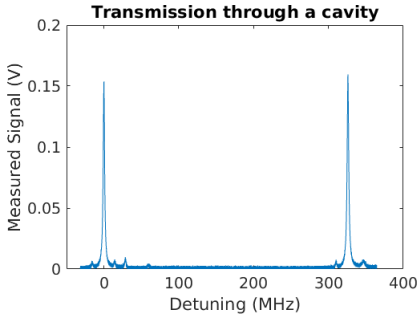

Figure 2.2: Example of a measurement taken by a photodetector on one side of a cavity with two semireflecting mirrors. Each mirror transmits 0.1% of the light shone on it. A laser is scanned over a broad range of frequencies. The measured signal are directly proportional to the intensity in the cavity. The two big peaks are the first transmission, and one more transmission after one entire FSR. The two small peaks on each of these are the sidebands. More about how they come about in the chapter on phase modulation. The smaller peaks in the middle are other possible modes, which have other frequencies.

2.2 Tuneable lasers

I assume the reader is familiar with the basics of lasers. A laser consists usually of a cavity with some gain medium. Alternatively it can be a fiber, which actually works like a cavity.

The exact length of the cavity determines the frequency of the laser, within a range determined by the medium. A lasers frequency can thusly be tuned by changing the length of the cavity. A good way of controlling tiny changes of length is by using a piezo-electric crystal.

A piezoelectric crystal is a crystal which can generate an electric potential when compressed. The word *piezin* means to squeeze in greek. Many disposable lighters use a small piezoelectric crystal to light up. Pressing on a button with ones thumb causes a hammer to hit the crystal and compress it, thusly creating a voltage high enough to let fly of a small spark.

Piezoelectric crystals are also reversible, in that an external voltage applied over them, causes them to change shape. This makes them useful for f.ex loudspeakers. We use piezo crystals different places in our setup in order to make very tiny and precise movements.

It is a piezoelectric crystal, expanding due to an electrical field, which allows us to control the cavity of the laser, and the cavity in which our spectroscopy will take place. (See figure 2.3)

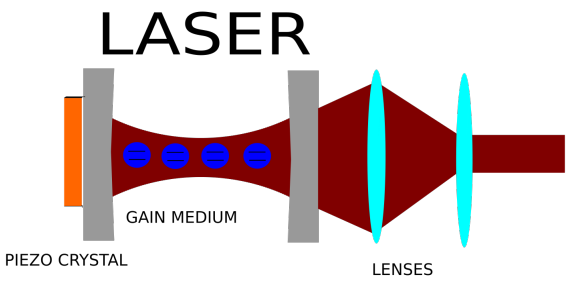

Figure 2.3: A simple schematic of a tuneable laser. The coherent Light is Amplified by the Stimulated Emission of the medium. The wavelenght of the light is chosen by the length of the optical cavity, which can be changed by a varying the elctric potential over a piezo crystal. The beam can be shaped to the desired mode with a pair of lenses.

2.3 Locking the Frequency

Think of how an old fashioned clock was stabilised to have exactly 24 hours of 60 minutes pr day.

Before atomic clocks, the main way of referencing the accuracy of clocks was the movement of the sun and the stars.

Before the pendulum clocks, in the medieval time, there existed clocks whose frequency was decided by the placement of some small weights on a lever. Imagine how a monk would tune such a clock using only the sun.

Our monk would first put a high pole somewhere. He knows that the shadow of the pole will point due north when the sun is in zenith at exactly noon. So the monk draws a long line on the ground pointing directly to the south.

Every day around noon, the monk goes out and looks at the shadow. At the moment when the shadow exactly covers the line on the ground, he looks up at the clock and notes the time down.

He could, if the difference from the measured time until twelve o'clock is

Figure 2.4: A medieval picture of an early clockmaker, Richard Wallingford, abbot of St Albans, pointing at the clock of the abbey.

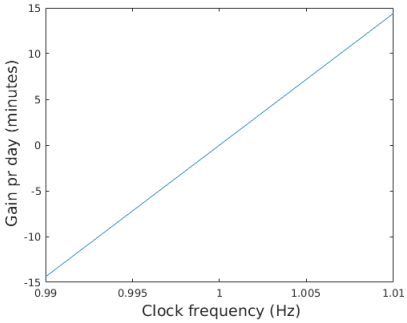

Figure 2.5: An old mechanical clock would be made to have a frequency of 1 Hz, one tic pr second. Here we can see the error signal obtained by comparing its noon with the sun.

not too big, just push the hands over to show twelve, and then let the clock continue running. That is how clocks were usually kept roughly in time. Our monk however wants not just to keep roughly the right time, but as far as possible let the clock have the right frequency by itself.

So, in addition to changing the hands, he notes down how many minutes the clock was away from noon. He then changes the position of the weights to change the frequency. If the clock was slow, he makes it run a bit faster. If the clock was too fast, he makes it run a bit slower.

Control theory is the theory of how the monk ought to change the position of the weights based on his measurements.

2.3.1 The error signal

The first concept in control theory is the error signal. For the case at hand, the error signal is the time measured at noon, minus twelve hours. When the time is less than twelve, the error signal is negative, and more than twelve the error signal is positive(see figure 2.5).

The essential thing about the error signal, is that it gives information of whether the clock is too slow or too fast. Any measurement, which is able to tell us this information, can be used as an error signal.

In the language of modern science, we would say that the monk is stabilising the clock at a frequency of 1 Hz using the error signal, which he obtains once pr day.

All the following thesis will be about obtaining an error signal for stabilising a laser at a frequency of 194.37 THz, and then locking the laser to this frequency with the error signal.

2.3.2 Locking the Frequency of a Laser

We can change the frequency of the laser by lowering or raising the voltage over a piezo crystal inside it. But in order to keep the laser on resonance we need a signal which gives us a negative number on one side of the setpoint and a positive number on the other side. For our purposes we will use an electronic

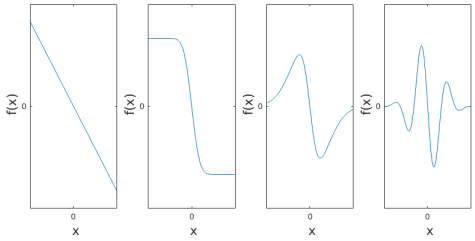

Figure 2.6: A number of possible error signals for locking the parameter x to the setpoint x=0. Notice that the two last will not push in the right direction if x starts too far away

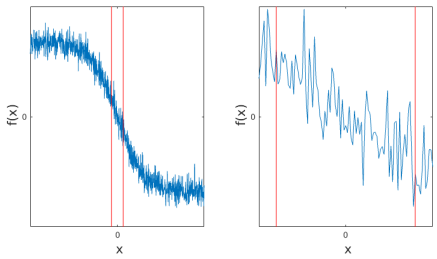

Figure 2.7: A hypothetical error signal with noise. The noise is changing randomly with time. All in all it is decreasing in the right way, but between the two red columns the noise masks out the slope. Therefore the parameter can only be locked somewhere in that area. A reater slope would give a narrower area.

signal which must be negative when the frequency of the laser is close to but lower than the frequency we want to hit, and positive if the frequency of the laser is close to but higher than the frequency we want. We call this signal the error signal.

Essentially any signal fulfilling this criteria could be used as an error signal.(see figure 2.6).

Some of these will however be better than others for locking to. In reality there will always be some sort of noise on the signals.

For precise control, one needs as steep a slope as possible around the center point, compared to the noise level.(see figure 2.7)

2.3.3 A Servo Loop

Assume that you have a laser which you want to lock to some angular frequency ω_0. Assume that you have some way to generate a useful errorsignal. We will send this error signal through an electronic circuit, which we call a servo circuit (see figure 2.8).

The frequency of the laser at time t is given as $\omega_1(t) = \omega_0 + f(t)$.

Here f(t), the difference between our desired frequency and our actual frequency, is a function caused by drifts from the environment that we have abso-

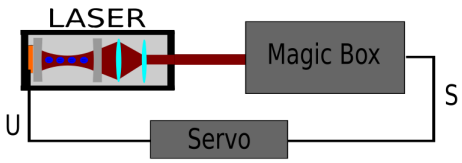

Figure 2.8: A feedback loop for locking the frequency of a laser. The laser is sent into a Magic Box, which produces an electronic error signal, S. This signal is sent through a circuit, which we will call the servo circuit, to produce a voltage U, which is used to tune the frequency of the laser.

lutely no control over.

Our magic box will generate an error signal for us, $s(f)$. For simplicity let us assume that s is linear for f not too big.

$$s = Cf(t). \tag{2.4}$$

Now we will fourier transform s, to obtain a function in frequency space, S(k). Here care must be taken. The frequencies we here are talking about, which I denote with k instead of ω, is the frequencies of our Fourier transform, and are not directly related to the actual frequencies f and ω_1 which they are describing (We are transforming a function describing frequency deviation as a function of time, into a a function describing frequency deviation as a function of frequencies).

$$S(k) = CF(k), \tag{2.5}$$

where F(k) is the Fourier transform of f(t),

$$F(k) = \int_{-\infty}^{\infty} f(t)e^{-ikt}dt. \tag{2.6}$$

Lets say that S is an electrical signal. We can then feed it through some circuit, which has a linear frequency dependent transfer function $g(k)$, so the output of the circuit, U will be:

$$U(k) = g(k)S(k) = g(k)CF(k). \tag{2.7}$$

This output signal can be added to the voltage over the piezo inside the laser, thus changing the frequency of the laser. Let us say that it does this with some proportionality constant D

Notice that the change from this input will become a part of F(k).

Let us define a new function, $F_0(k)$ which represents the frequency deviation if the circuit did not work back on the frequency of the laser.

We will see that:

$$F(k) = F_0(k) - DU(k), \tag{2.8}$$

where we have chosen a negative sign on D, since we are doing a negative feedback loop.

If we isolate $F_0(k)$ we will see that:

$$F_0(k) = F(k) + DU(k) = F(k) + Dg(k)CF(k) = F(k)(1 + Dg(k)C), \quad (2.9)$$

which can be reduced to:

$$F(K) = \frac{F_0(k)}{1 + Dg(k)C}. \quad (2.10)$$

Let us look at this equation and see what it tells us. The frequency deviation (in the Fourier transformed picture) for a laser with a negative feedback loop is smaller than the deviation without the loop by a factor $1 + Dg(k)C$, constants from the circuit and the setup.

If one wants to, this can be transformed back into the time domain.

2.3.4 PID Servo Circuits

A PID circuit is a special kind of servo circuit. It consists of different amplifiers, able to amplify a signal normally, integrate it, and differentiate it. All of these operations are linear operations.

PID is short for Proportional, Integrated, Differentiated.

A PID circuit can be characterised by three constants, determining how much each of these three components is amplified.

So, the output of a PID-circuit with a complex input $s = Ae^{ikt}$ would be:

$$u = c_P s + c_I \int s dt + c_D \frac{ds}{dt} = c_P s + c_I iks + \frac{c_D}{ik} s. \quad (2.11)$$

This means that the the transfer function $g(k)$ for such a circuit will be:

$$g(k) = c_P + ic_I k - i\frac{c_D}{k}. \quad (2.12)$$

Most of the servo loops used in this experiment is simple PID circuits, with $c_D = 0$.

2.4 Acetylene-13

Acetylene is a colourless gas made from two carbon and two hydrogen atoms. It is sometimes used for welding, because of its high combustion heat.

In the version of acetylene which we use, both carbon atoms are the Carbon-13 isotope, that is carbon with seven neutrons. Around 1% of carbon atoms found in nature are carbon-13. Acetylene-13 can be specially made for purposes such as ours.

The gas we used were kept at very low pressure (a few Pa) in some special made glass cells. A molecule, being made up of several atoms, have more degrees of freedom than a single atom. Most notably, it can vibrate. The bound between the atoms in that case works like springs.

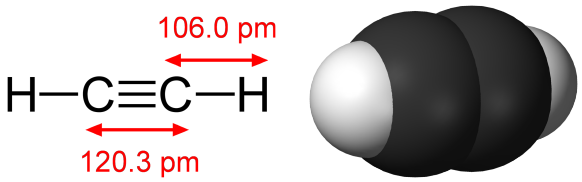

Figure 2.9: Acetylene. Two carbon atoms with a triple bond plus two hydrogen atoms

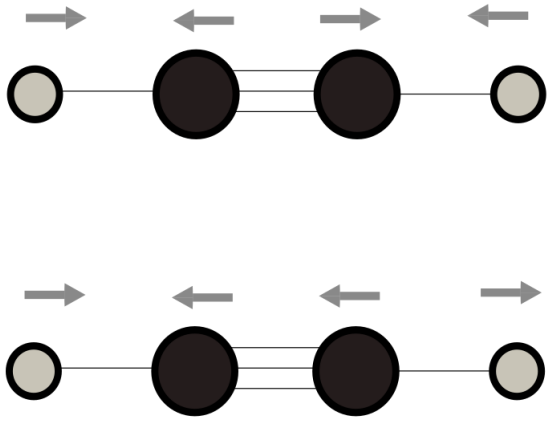

Figure 2.10: The v_1 and v_3 vibrational modes in Acetylene.

A molecule with two atoms, like for example oxygen, O_2, have only one mode of vibrating, both atoms vibrating back and forth, with the center of mass unaffected. This vibration works like a quantum mechanic oscillator. For small vibrations, it can usually be approximated as a harmonic oscillator.

Molecules with more atoms, such as acetylen, have more complex vibrational states. One can count up all the different ways it is possible for it to vibrate.

Our transition is based on exciting a combination of two vibrational states at the same time. They are

v_1; where the two carbon atoms move inwards and outwards like the case for a two atomic molecule while the hydrogen atoms follows their movements, and

v_3; where the two carbon atoms follow each other, and the two hydrogen atoms follow each other, but the carbon atoms vibrate against the hydrogens.(See figure 2.10)

In dipole interaction with light, the conservation of angular momentum means that the total angular momentum of the molecule, J can only change by -1,0 or 1.

$$\Delta j = -1, 0, 1. \tag{2.13}$$

The coupling of rotation and vibration gives rise to three branches. In one, an excitation of a vibrational mode is combined with a change in rotation of

-1, in one the rotation is not changed, and in one the rotation along our chosen axis is increased by +1. These are named the P, the Q and the R branch.

The convention is to name a transition where j changes from n to n-1 the $P(n)$ transition. Our specific transition is fully known as the $P(16)v_1 + v_3$ transition, meaning a photon excites one level of the $v_1 + v_3$ vibrational mode while decreasing J from 16 to 15.

The wavelength of this transition is 1542.3837 nm, equivalent to a frequency of 194.4 THz.

2.5 Spectroscopy

2.5.1 Basics of Spectroscopy

When a coherent light ray such as a laser is shone through a medium, it can change intensity and phase.

The loss of intensity due to the medium can be expressed by the mediums attenuation coefficent α which in general depends on the frequency of the light (it also depends on the intensity of the light, the temperature and the pressure of the medium and other things, but we will cover them later).

The intensity of light passing through a completely stationary medium with attenuation coefficent α will after a distance l be:

$$I(l) = I(0)e^{-\alpha l}. \tag{2.14}$$

If one is in the area of a resonance of the medium, with other resonances sufficiently distant not to play a role, α will have a Lorentzian dependence on the frequency of the light.

$$\alpha(\omega) = \frac{\alpha_0}{(\omega - \omega_0)^2 + \gamma^2}, \tag{2.15}$$

where ω_0 is the resonance frequency of the medium.

The breadth of this function is set by the parameter γ, which depends on the resonance, and the pressure of the medium. γ is called the Half Width at Half Maximum. 2 times γ is the Full Width at Half Maximum, abbreviated FWHM, and in this thesis denoted as Γ.

If the light is measured on the other side of the medium, the intensity will now depend on the frequency, and one can thereby determine the resonance frequency ω_0. (see figure 2.11)

2.5.2 Doppler broadened spectroscopy

Above, we assumed that the medium was completely stationary. But in for example a gas at room temperature, the molecules will be moving in all directions with high speed.

Any object that is moving will observe all waves around it to have a different frequency than if it had been stationary. This is known as the doppler effect.

A molecule moving with a speed v along the direction of a passing wave sees an effective frequency ω_{eff}, given by:

$$\omega_{eff} = \omega(1 + \frac{v}{c}), \tag{2.16}$$

where ω is the frequency a stationary molecule would have experienced, and c is the speed of the wave, in this case the speed of light.

So, for one single molecule with a definite speed component along the wave, the attenuation coefficent would then be:

$$\alpha(\omega, v) = \frac{\alpha_0}{(\omega(1 + \frac{v}{c}) - \omega_0)^2 + \gamma^2}. \tag{2.17}$$

For any single molecule in a gas cloud in thermal equilibrium, the chance of it having a velocity v along any axis will be given by a gaussian probability distribution,

$$p(v) = Ae^{\frac{-v^2}{2\sigma^2}}, \tag{2.18}$$

where σ depends on the temperature of the gas.

The total attenuation coefficent for the gas will then be:

$$\alpha(\omega) = \int p(v) \frac{\alpha_0}{(\omega(1 + \frac{v}{c} - \omega_0))^2 + \gamma^2} dv. \tag{2.19}$$

This integral over a Gaussian function times a Lorentzian function is called a Voigt function.

One can calculate that in the limit where $\sigma >> \gamma$, the Lorentzian function can be treated as a Dirac delta function, and the entire thing becomes a gaussian function of ω.

On the other hand, if $\gamma >> \sigma$ one can treat the Gaussian function as a Dirac delta function around $v = 0$, in which case the entire thing becomes a Lorentzian function again.

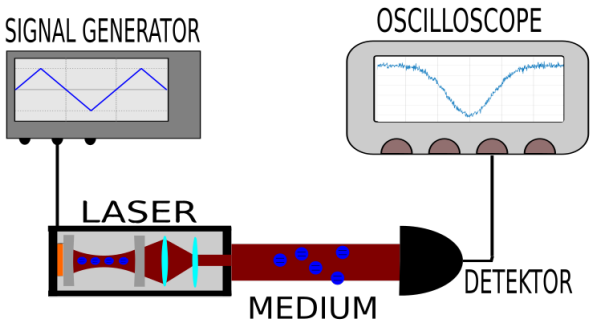

Figure 2.11: Absorbtion spectroscopy. The frequency of the laser is scanned up and down with a signal generator. The light is shone through an absorbing medium into a photodetector, and the output of the detector is plotted in an oscilloscope. We used this method for rough-tuning our lasers to the right frequency.

2.6 Heterodyne Spectroscopy

The oldest kind of spectroscopy is absorbtion spectroscopy, where light is absorbed when on resonance with some medium.

Another way of performing spectroscopy is detecting the change of phase in the signal due to the medium.

The simplest way to perform this spectroscopy is called homodyne spectroscopy. You split a beam of light, send one part of it through a medium, and then combine the two rays again on a photo diode. In that case the intensity measured by the diode will depend on the phase between the two rays, that is whether they are interferring constructively or not.

Our experiment instead works by putting three different frequencies in the same beam and looking at the beat between them. This is called heterodyne, because one uses different frequencies.

When light hits a detector, the detector reads out the intensity of the light averaged over some time. A normal detector cannot directly turn an electromagnetic wave oscillating at optical frequencies into a likewise oscillating electric current.

But, if it is hit by two lightrays with different frequencies, it might be able to detect the beat between the two frequencies.

This is the same phenomenon as when the human ear is able to detect the beat tone between two tones.

More formally we could say, if we have a field that is a superposition of two waves travelling in the same direction, at the point of the detector, we can write the total field as:

$$E(t) = E_1 \cos(\omega_1 t) + E_2 \cos(\omega_2 t) \tag{2.20}$$

$$B(t) = \frac{1}{c} E_1 \cos(\omega_1 t) + \frac{1}{c} E_2 \cos(\omega_2 t). \tag{2.21}$$

The poynting vector will become:

$$S(t) = \frac{1}{\mu_0 c}(E_1^2 \cos^2(\omega_1 t) + E_2^2 \cos^2(\omega_2 t) \tag{2.22}$$

$$+2E_1 E_2 \cos(\omega_1 t)\cos(\omega_2 t)), \tag{2.23}$$

which can be rewritten as

$$S(t) = \frac{1}{\mu_0 c}(E_1^2 \cos^2(\omega_1 t) + E_2^2 \cos^2(\omega_2 t) \tag{2.24}$$

$$+E_1 E_2 \cos((\omega_1 + \omega_2)t) \tag{2.25}$$

$$+E_1 E_2 \cos((\omega_1 - \omega_2)t)). \tag{2.26}$$

If ω_1 is close to ω_2 the last term will oscillate much slower than the three others. If the detector averages over all the fast terms, the measured intensity will be

$$I = \frac{1}{\mu_0 c}(\frac{1}{2}E_1^2 + \frac{1}{2}E_2^2 + E_1 E_2 \cos((\omega_1 - \omega_2)t)). \tag{2.27}$$

Heterodyne detection is a smart way of using this principle of beats.

2.6.1 Phase Modulation

An electro optical modulator, which I from now on will only talk of as an EOM, is a piece of optics designed to change the phase of light passing through it, based on an electric signal.

Let us look at a beam of light moving throught the air. At some point where it is passing it can be described as an electric field swinging in the direction of its polarisation with a size (here given as the complex amplitude plus it's complex conjugate)

$$E = \frac{E_0}{2}e^{i\omega t} + c.c. \tag{2.28}$$

where we have chosen our time starting point so there are no phase. But, if somehow the path of the light were to get longer, say by moving the light source a tiny bit away from our point, a phase would apper;

$$E = \frac{E_0}{2}e^{i\omega t + i\phi} + c.c. \tag{2.29}$$

So, slowly moving a laser forward and backwards by a few nano-meters (slowly, so you don't change the frequency of the light through the doppler effect), would be seen as a slow change in the phase of the light at some later point.

But there are other ways of changing the phase of light. One would be, instead of changing the distance the light travels, instead to change somehow the speed of the light at some point. Which can be accomplished with a crystal that changes its refractive index based on the Electric Field.

Imagine a crystal of length l, with a refractive index n. Then imagine making the refractive index larger by an amount Δn. In that case the extra time it takes a ray of ligth to pass through the crystal, Δt would be

$$\Delta t = \frac{l\Delta n}{c}, \tag{2.30}$$

where c is the speed of light in vacuum. This will lead to a phase change ϕ depending on the frequency of the light ω,

$$\phi = \omega \Delta t. \tag{2.31}$$

In general any sort of change can be made to the phase of a light ray using an EOM. However, a particularly funny thing happens when the EOM is affected by a harmonic oscillating electric field.

Let us say that this field introduces a change of refractive index:

$$\Delta n = b \cdot sin(\Omega t), \tag{2.32}$$

where b is a constant, which depends on the amplitude one modulates with, then

$$\phi = \omega b \cdot sin(\Omega t). \tag{2.33}$$

Let us define:

$$\omega b = \phi_0. \tag{2.34}$$

So at the point of interest, the electric field from the light ray will be

$$E = \frac{E_0}{2} e^{i\omega t + i\phi_0 sin(\Omega t)} + c.c. \tag{2.35}$$

If the amplitude of the electric field over the EOM crystal is sufficiently small, $\phi_0 << 1$. In that case we could taylor expand the equation for E in ϕ_0 up to the first order, and have a relatively good approximation:

$$E \simeq \frac{E_0}{2}(e^{i\omega t}(1 + i\phi_0 sin(\Omega t))) + c.c. \tag{2.36}$$

$$E \simeq \frac{E_0}{2}(e^{i\omega t}(1 + i\phi_0 \frac{e^{i\Omega t} - e^{-i\Omega t}}{2i})) + c.c. \tag{2.37}$$

$$E \simeq \frac{E_0}{2}(e^{i\omega t} + \frac{\phi_0}{2} e^{i(\omega+\Omega)t} - \frac{\phi_0}{2} e^{i(\omega-\Omega)t}) + c.c. \tag{2.38}$$

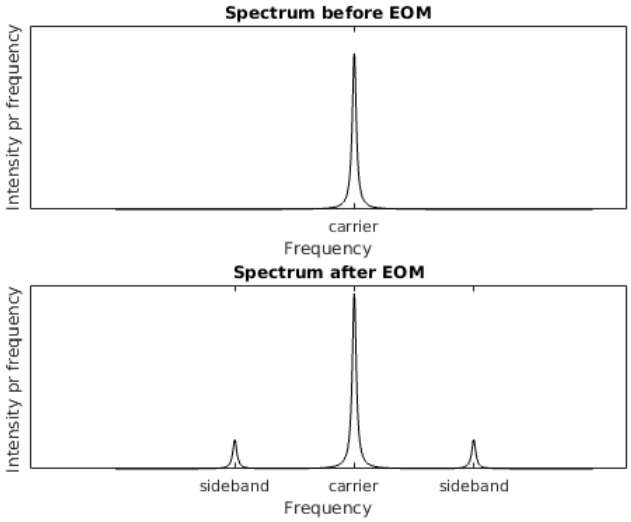

Figure 2.12: Spectrum of light with a sinusoidally oscillating phase

$$E \simeq E_0(cos(\omega t) + \frac{\phi_0}{2} cos((\omega + \Omega)t) - \frac{\phi_0}{2} cos((\omega - \Omega)t)) \tag{2.39}$$

So, what comes out is a superposition of three waves with frequencies ω, $\omega + \Omega$ and $\omega - \Omega$. We will (just for the sake of conventionality) call the wave with the original frequency the carrier, and the two other waves the sidebands.(see figure 2.12)

2.6.2 Heterodyne Spectroscopy with Sidebands

In our setup, we use a light signal with sidebands, (such as can be generated with an EOM). If the sidebands are generated with phase modulation they will both have a beat with the central carrier, but their beats will be exactly out of phase.

Going through it step by step:

The electric field at the photo detector is:

$$E(t) = E_0(\cos(\omega t) + \frac{\phi_0}{2} \cos((\omega - \Omega)t) - \frac{\phi_0}{2} \cos((\omega + \Omega)t)). \tag{2.40}$$

This can easily be rewritten to:

$$E(t) = E_0(\frac{1}{2} \cos(\omega t) + \frac{\phi_0}{2} \cos((\omega - \Omega)t)) + E_0(\frac{1}{2} \cos(\omega t) - \frac{\phi_0}{2} \cos((\omega + \Omega)t). \tag{2.41}$$

So two parts, each of which will produce a beat with exact opposite sign of the other.

This means that the beat signal between the lower sideband and the carrier will cancel out with the beat between the higher sideband and the carrier. There will however come beats between the two sidebands, with frequency 2Ω.

However if anything happens that causes a change of phase in one of the sidebands, a beat will appear with frequency Ω.

If the electric signal from the photo detector is demodulated using a mixer with another signal of frequency exactly Ω, any component with frequency Ω will give rise to a dc signal.

Then if one knows the phase change as function of frequency caused by the medium, the dc signal after the mixer can be computed as a function of frequency.

Now imagine our perfect triplet of carrier and two sidebands passing through some material. This material can have different refractive index n for each of the three frequencies. This will (in general) lead to three different phases for the three sidebands.

If the medium is l long, the change in phase due to it will be for the different components

$$\phi_1 = ln(\omega), \phi_2 = ln(\omega - \Omega), phi_3 = ln(\omega + \Omega). \tag{2.42}$$

Then the total electric field at the photo detector will be:

$$E(t) = E_0(cos(\omega t + \phi_1) + \frac{\phi_0}{2}\cos((\omega - \Omega)t + \phi_2) - \frac{\phi_0}{2}\cos((\omega + \Omega)t + \phi_3)) \tag{2.43}$$

The terms of the intensity from this field oscilatting with frequency Ω is

$$I_\Omega = \frac{E_0^2}{\mu_0 c}\frac{\phi_0}{2}(cos(\Omega t - \phi_1 + \phi_2) - cos(\Omega t + \phi_1 - \phi_3)). \tag{2.44}$$

Now, let us assume that the phase shift is cause by a resonance with a two level system in the medium somewhere close to the frequency of the carrier, and that we have placed the sidebands so far from the carrier in frequency that they are not affected by this resonance, and also that they are not affected in other ways by the medium. In that case,

$$\phi_2 = \phi_3 = 0, \tag{2.45}$$

which gives us a signal oscillating at Ω;

$$I_\Omega = \frac{E_0^2}{\mu_0 c}\frac{\phi_0}{2}(\cos(\Omega t - \phi_1) - \cos(\Omega t + \phi_1). \tag{2.46}$$

This can be rewritten using a trigonomic identity to

$$I_\Omega = \frac{E_0^2}{\mu_0 c}\frac{\phi_0}{2}\sin(\Omega t)\sin(\phi_1). \tag{2.47}$$

We could just proceed with this result. But, if we assume that the phase shift is small, $\phi_1 << pi$ then we can use the small angle approximation: $\sin(\phi_1) \approx \phi_1$ to write

$$I_\Omega = \frac{E_0^2}{\mu_0 c} \frac{\phi_0}{2} \sin(\Omega t)\phi_1. \tag{2.48}$$

If we demodulate the electric signal from the photodetector with frequency Ω we will get an electric signal S,

$$S = K \cdot \phi_1, \tag{2.49}$$

where K is a constant including all our electrical amplification and losses and also the factor from the phase which we demodulate with.

Let us just recap what all this means. If we send a laser beam through an EOM where it is phase modulated, and then through a medium which imparts a (small) phase ϕ_1 to the carrier, while not interfering with the sidebands, we can get an electric signal S which is directly proportional to ϕ_1 by demodulating the electric signal from the photodetector with the same frequency we used for the phase modulation.

2.6.3 Phase Shift from a Simple Two Level System

Now we must ask us self how does the medium change the phase of the carrier. Of course for a given arbitrarily complex medium, this change be arbitrarily complex.

But, in the case where phase shift is caused by the resonance between only two energy eigen states, we can actually describe this change pretty well.(see figure 2.13)

Lorentzian lines In the case where the absorbtion line would be Lorentzian (since the doppler broadening caused by the movement of molecules in the medium can be neglected) the phase shift is given as:

$$\phi_1 = -ClN \frac{\omega - \omega_0}{(\omega - \omega_0)^2 + \gamma^2}, \tag{2.50}$$

where C is a constant that depends on the medium(and possibly the intensity of the ligth), l is the length of the medium, N is the density, ω is the frequency of the carrier, ω_0 is the resonance frequency of the transition and γ is the HWHM linewidth of the transition including all Lorentzian broadening mechanisms (natural linewidth, pressure broadening, power broadening etc).

This would for example be the case for a very cold medium with high pressure broadening.

Gaussian lines In the case where the doppler broadening is dominating (In f.ex a low pressure gas at room temperature) the absorbtion line would be a gaussian function.

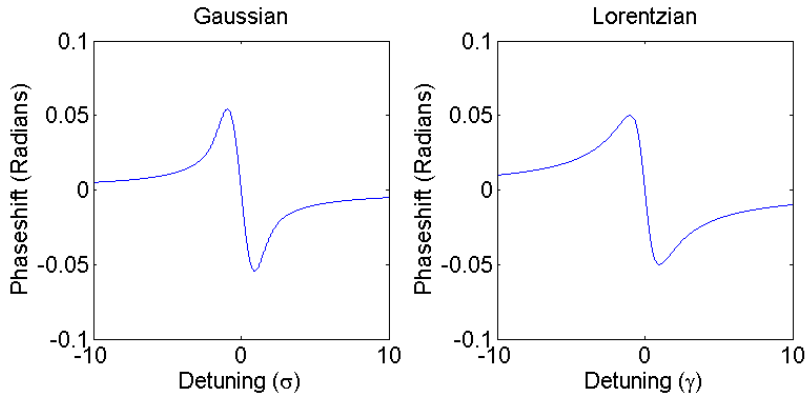

Figure 2.13: Examples of possible phaseshifts. On the left the shift from a gaussian line, and on the right the shift from a purely Lorentzian line.

Here the phase change can be given as:

$$\phi_1 = ClNe^{-\frac{(\omega-\omega_0)^2}{\sigma^2}} \int_0^{\omega-\omega_0} e^{\frac{s^2}{\sigma^2}} ds, \tag{2.51}$$

Where again C is a constant depending on the medium. σ is the Gaussian width of the transition, determined by the speed of the molecules.

Voigt lines In areas in between the mentionedd above, the absorbtion lineshape is a combination of a Lorentzian and a Gaussian shape, called a Voigt profile. In that case the phase shift cannot be written down in a closed form, and must be calculated using the complex errorfunction of a complex input.

2.6.4 The Heterodyne signal as an error signal

All of these three functions have one thing in common. $\phi_1 = 0$ for exact resonance ($\omega = \omega_0$), and they are odd symmetric functions, which means they have a slope around the resonance.

So, if we scan a laser across resonance with the medium, we will get an electric signal with the same characteristics.

If we recall the chapter on control theory, we will realize that this signal can be used as an error signal for locking the frequency of our laser to the resonance frequency of the medium.

And there we have our first possible way of generating an error signal.

2.7 Lineshape and Susceptibility

2.7.1 Susceptibility

Now would be a good time to define exactly what we mean by the susceptibility of some medium, χ.

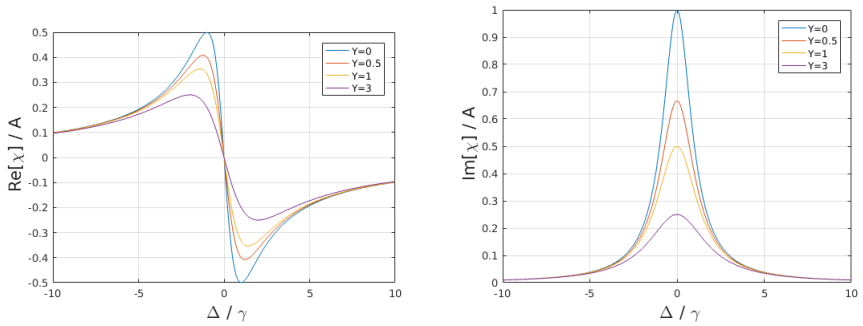

Figure 2.14: The real and imaginary parts of χ, plotted for different values of the parameter $Y = \frac{I}{I_s}$, given in values of $A = \frac{N|\mu|^2}{\epsilon_0 \hbar \gamma}$

If a medium is affected by some timedependent electric field $\vec{E}(t)$, and has a dipole polarisation $\vec{P}(t)$, then the susceptibility χ is a tensor fulfilling

$$\vec{P}(t) = \epsilon_0 \chi \vec{E}(t). \tag{2.52}$$

Where ϵ_0 is the vacuum permitivitty.

For a two level system, with the dipole approximation, affected by a linearly polarised monochromatic field the susceptibility reduced to only a complex scalar, which will be a function of the frequency of the light.

The susceptibility in that case gives us full knowleddge of both the absorbtion and refractive index of the medium.[12]

The attenuation coefficient is given as:

$$\alpha(\omega) = \frac{2\omega}{c} \mathrm{Im}[\sqrt{1 + \chi(\omega)}]. \tag{2.53}$$

and the refractive index is given as:

$$n(\omega) = \mathrm{Re}[\sqrt{1 + \chi(\omega)}] \tag{2.54}$$

A two level system dipole interacting with an electric field can be characterised by its linewidth γ, its resonance frequency ω_0, its saturation intensity I_s and $|\mu|$, the absolute value of the off diagonal elements of its dipole operator.

Its susceptibility is then given as

$$\chi(\omega) = \frac{N|\mu|^2}{\epsilon_0 \hbar \gamma} \left(\frac{-\frac{\Delta}{\gamma}}{1 + \frac{\Delta^2}{\gamma^2} + \frac{I}{I_s}} + i \frac{1}{1 + \frac{\Delta^2}{\gamma^2} + \frac{I}{I_s}} \right) \tag{2.55}$$

where $\Delta = \omega - \omega_0$ is the detuning, and I is the intensity.

You can see the real and imaginary part of this function in figure 2.14

For small values of the susceptibility ($\chi \ll 1$), the attenuation and refractive index will depend approximately linearly with the susceptibility.

$$\alpha(\omega) \approx \frac{\omega}{c}\mathrm{Im}[\chi(\omega)]. \tag{2.56}$$

$$n(\omega) \approx 1 + \mathrm{Re}[\frac{\chi}{2}]. \tag{2.57}$$

This approximation is valid when the medium for example is a very diluted gas, such as in our experiment.

2.8 Saturated spectroscopy

We have already mentioned how the random movement of warm molecules (or temperate molecules. Really anything over -200 degrees celsius is pretty warm by these standards) leads to a gaussian broadening of the absorbtion and dispersion lineshapes.

There is however smart ways of removing this broadening, which doesn't require cooling the molecules down.

To understand the concept, we must start at the concept of optical saturation.

Actually the attenuation constant of a medium does not only depend on the frequency of light passing through it, but also on the the intensity of the light.

If one held the frequency constant, and varied the intensity instead, one finds an effective attenuation given as:

$$\alpha(I) = \frac{\alpha_0}{1 + \frac{I}{I_s}}. \tag{2.58}$$

Where α_0 is the attenuation constant at low intensity, and I_s is a parameter that depends on the medium.

Now, imagine what would happen if a beam of light was sent through a medium onto a mirror and then back again, and into a photodetector.(see figure 2.15) Assume the light have completely well defined frequency, and that the loss of intensity trough the medium is small, i.e. $\frac{I_{in}-I_{out}}{I_s} << 1$

Then assume that the doppler broadening is much bigger than the natural linewidth. In that case, if the laser is detuned from the resonance frequency by an amount Δ,

$$\Delta = \omega - \omega_0, \tag{2.59}$$

it will on its first trip only be affected by molecules with a velocity along the beam

$$v = \frac{\Delta c}{\omega_0}. \tag{2.60}$$

While when it is passing back it will only be affected by molecules with a velocity component

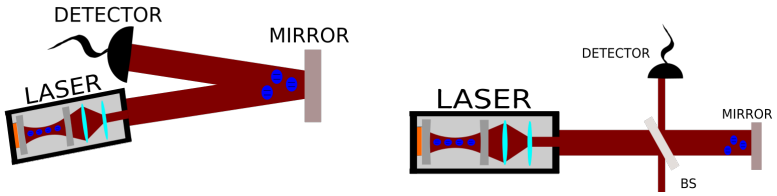

Figure 2.15: Saturated spectroscopy. Light needs to pass through the medium in both directions, and then onto a photodetector. This can be done by shining it at a slightly angled mirror (first picture). More commonly one uses a beamsplitter (BS), which reflects some of the returning beam onto the photodetector (second picture)

$$v = -\frac{\Delta c}{\omega_0}. \tag{2.61}$$

Since the velocity distribution is symmetric it will lose the same amount of intensity going both ways.

But, if $\Delta = 0$, it will in both ways be affected by the exact same molecules, that is the molecules that has v=0 along the beams direction.

This means again, that those molecules will be affected by the beam travelling both ways. And that means that they experience a total intensity of 2I instead of I, assuming that the molecule is placed at a an antinode in the standing wave resulting from two travelling waves passing each other(see figure 2.16).

(Note, in reality two opposite waves will create a standing wave, with different powers at different positions. On average however the intensity experienced by a molecule will be twice that of one of the two waves).

Mathematically speaking we can say that if the doppler broadenend profile has a gaussian width σ then

$$\alpha(I, \Delta) = \alpha_0 e^{-\frac{\Delta^2}{2\sigma^2}} \frac{1}{1 + \frac{I}{I_s}}, \tag{2.62}$$

if Δ isn't equal to zero, and

$$\alpha(I, \Delta = 0) = \alpha_0(\Delta = 0) \frac{1}{1 + 2\frac{I}{I_s}}, \tag{2.63}$$

when Δ is equal to zero. Here $\alpha_0(\Delta = 0)$ in this case means the attenuation constant for exact resonance and low intensity.

This means that the attenuation is smaller for exact resonance than for a small detuning away from the resonance.

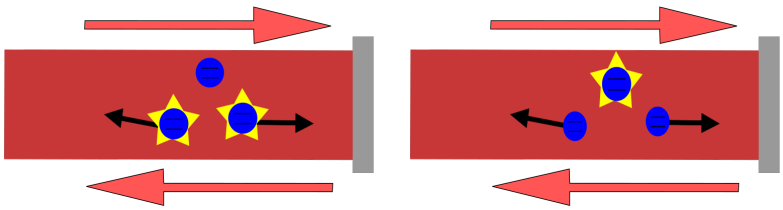

Figure 2.16: A semiclassical wiev of saturated spectroscopy in a cloud of three molecules. In the picture on the left the light is off resonance with the molecules in their restframe. Travelling to the left it excites the molecule moving towards the right. The light is reflected from the mirror, and travelling to the right it excites the molecule moving towards the left. In the picture on the right, the light is on exact resonance with the molecules. First it excites the molecule that is standing still, and then coming back it meets no new molecules it can excite.

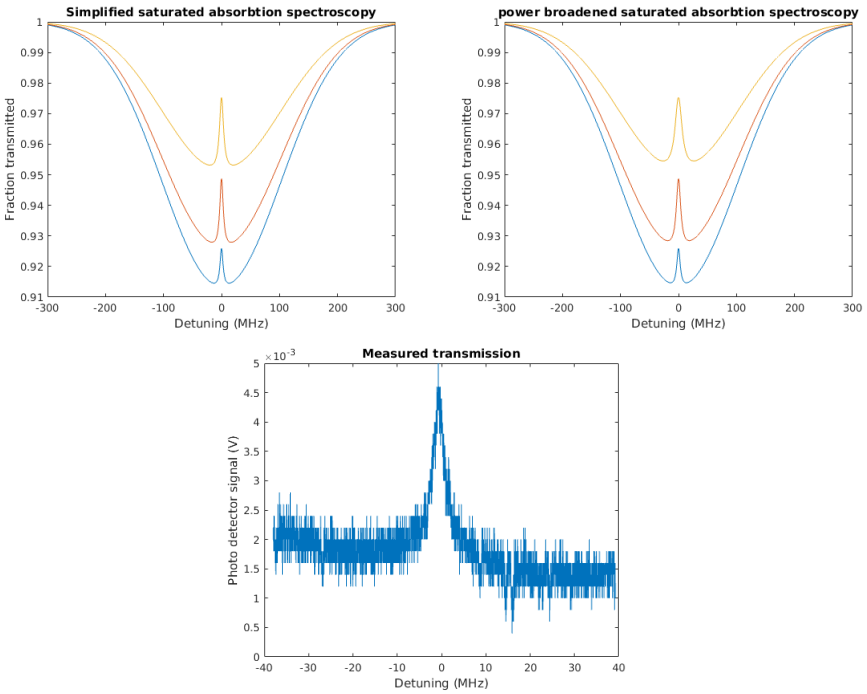

Figure 2.17: Simulation of the fraction of power passing back and forth through some medium which has a doppler broadend gaussian width of $\sigma = 100$ MHz, and a Lorentzian width $\gamma = 3$ MHz, for three different intensities: $I = 0.1 \cdot I_s$ (blue), $I = 0.3 \cdot I_s$ (red) and $I = I_s$ (yellow). In the picture on the left, we just that the width of the saturation peak is a Lorentzian with width γ, and no other effects, as described here. In the picture on the right we have included the power broadening of γ, which will be described later. Below is a picture of an actual lamb dip we have measured. Note that here the scan is too small to show the very broad Gaussian profile

In the real world, there will also still be some lorentzian linewidth, so what one will actually see would be a lorentzian hole in the function of the attenuation constant around resonance. (see figure 2.17)

This hole will mean *higher* transmission through the medium, so the signal measured at the photo detector will be higher at exact resonance than slightly detuned.

2.8.1 The Lamb dip

The name for the peak in intensity right when the light is on resonance with the medium is called the Lamb dip. The reason for this misnomer is that it was first experimentally encountered in lasers, were the saturation meant a loss of gain power, and hence a dip in the power output when the laser was on resonance with its gain medium.

Since however, that the Lamb dip is not broadened by the movement of the molecules, it is a much more precise way to measure the exact resonance frequency than non-saturated spectroscopy (at least for warm gasses).

2.8.2 Heterodyne Saturated Spectroscopy

Now we have seen how passing light twice through a medium can give a smaller absorbtion right on resonance. For our purpose, it would however be more interesting to see what saturation means for the phase change induced by the medium.

This could be treated in an analogous way. When the medium is saturated, it causes a smaller phase shift. This means that when a doppler broadened signal is obtained, there will be a small signal around the resonance with opposite sign. This signal could then be used as the error signal of our locking mechanism, since changing the sign of the signal is only a simple operation of changing the phase of the demodulation.See figure 2.18)

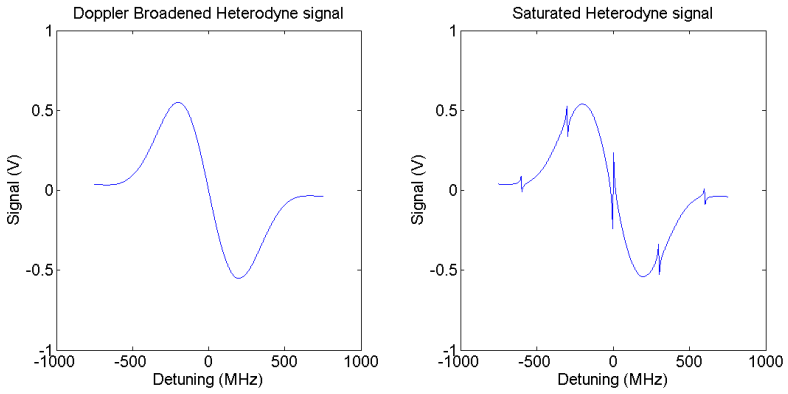

Figure 2.18: A Saturated Heterodyne Signal. On the left is a standard heterodyne signal, which would be recorded for a gaussian line with width σ =250 MHz. On the right is a signal which would be obtained if the laser were to saturate the medium. Here we have assumed that the medium has a Lorentzian HWHM of γ =2.5 MHz, and the laser is modulated with 600 MHz. The signal in the middle is from the carrier interacting with the stationary molecules. The signal at 300 MHz detuning comes from the sidebands interacting with the same, non stationary subset of the molecules. The furthest signals at 600 MHz detuning comes from the sidebands interacting with the stationary molecules.

2.9 Broadening

2.10 Line broadening for saturated spectroscopy

Even with saturated spectroscopy, there will still be effects that broadens the absorbtion line. These are summarised here.

2.10.1 The natural linewidth

From [7], we have the dipole moment related to the $P(12)v_1+v_3$ line in $^{13}C_2H_2$, of $\mu = 1.01$ Cm (Coulomb times meter).

This should give a natural linewidth of:

$$\gamma_0 = \frac{2}{3}\frac{\omega_0^3}{h\epsilon_0 c^3}\mu^2, \tag{2.64}$$

where ω_0 is its angular resonance frequency. This gives an expected natural linewidth of 0.78 Hz. We highly expect the natural linewidth of the $P(16)$ transition to be close to this value. In that case, it will be so much smaller than other broadening effects that we do not expect to be able to measure it.

2.10.2 Pressure broadening

The molecules in the gas will interact with each other, which causes a broadening of the line. The higher the pressure, the more interaction and hence the broader the line.

The change from the pressure will be given as

$$\gamma_p = C_p p \tag{2.65}$$

Where p is the pressure and C_p is a constant, which depends on the gas. The constant for our gas, $13C_2H_2$ has been measured to be roughly [**?**]:

$$C = 113\text{kHz/Pa} \pm 5\text{kHz/Pa} \tag{2.66}$$

for pressures in the range we are using, also in saturated spectroscopy. (In the cited article, the FWHM pressure broadening (so twice the HWHM) is given as $0.96 \cdot 234$ kHz/Pa, which I, to stay in my own notation, has recalculated directly.)

In addition to the pressure from acetylene, there might be other gasses present, each of which would add their own contribution, with their own constant.

2.10.3 Power broadening

This broadening mechanism was discussed in the chapter on theory of a two level system. It arises due to the saturation of the medium at high intensities.

If a line has a natural linewidth γ, and a saturation intensity at zero detuning I_s, and is affected by a monochromatic electric field with an intensity I, then the powerbroadened linewidth will be

$$\gamma_{power} = \gamma\sqrt{1 + \frac{I}{I_s}} \tag{2.67}$$

2.10.4 Transit time broadening

When we talked about saturated spectroscopy, we assumed that the molecules would continually be in an oscillating electric field with a constant amplitude. In reality the laserbeam has a finite width, and even the molecules with zero velocity along the line of the beam is still moving orthogonally to it. Hence they will drift in an out of the beam.

For the point of view of a single molecule, it would seem that the amplitude of the field is changing in time. As such, it must be expressed as an integral over different frequencies.

Take a molecule which is travelling with a speed component v orthogonal to our laser beam, which has a gaussian width w. It will then take it a time $\tau = \frac{v}{w}$ to move a distance w.

If we set our start time, so the molecule is in the middle of the beam at $t = 0$, the amplitude experienced by the molecule would be

$$E = E_0 e^{\frac{-r^2}{w^2}} = E_0 e^{\frac{-v^2 t^2}{w^2}} = E_0 e^{\frac{-t^2}{\tau^2}}, \tag{2.68}$$

where E_0 is the amplitude at the center of the beam.

This translates into a total time dependent electric field:

$$E(t) = E_0 e^{\frac{-t^2}{\tau^2}} cos(\omega t). \tag{2.69}$$

We can make a fourier transform of this to see what amplitude at a given frequency ω_i this would translate into:

$$F(\omega_i) = \frac{1}{2}\int_{\text{inf}}^{\text{inf}} E(t) e^{-i\omega_i t} dt. \tag{2.70}$$

If we evaluate this integral we would find

$$F(\omega_i) = \frac{E_0\sqrt{\pi}\tau}{2}\left(e^{-\frac{(\omega-\omega_i)^2\tau^2}{4}} + e^{-\frac{(\omega+\omega_i)^2\tau^2}{4}}\right). \tag{2.71}$$

This is two gaussian functions, with peaks at $\omega_i = \pm\omega$, both of which will affect the molecule in the same way. They have a gaussian width of $\frac{\sqrt{2}}{\tau}$.

When the molecule sees our monochromatic light as a gaussian package of frequencies, we will see the reaction of the light as a gaussian distribution with the same width.

This broadening is only for a molecule with a given speed v moving straight through the center of the beam. In reality, we would have to work with a spread of orthogonal velocities and different paths through the beam, after which we can calculate the mean broadening; [9]

$$\Delta\gamma_T = 2\sqrt{2ln(2)}\sqrt{\frac{\pi k_b T}{2m}}\frac{1}{w}. \tag{2.72}$$

Where w is the width of the beam, m is the mass of the molecules, T is the temperature and k_b is Boltzmanns constant.

2.10.5 Total Linewidth

These are the main broadening mechanisms in our experiment. Neglecting the unknown but small natural linewidth, the total linewidth should be

$$\gamma_{total} = (\gamma_p + \gamma_{tt})\sqrt{1 + \frac{I}{I_s}} \tag{2.73}$$

2.11 Laser Stabilised with Hetrodyne Saturated Spectroscopy

With just the concept of saturated heterodyne spectroscopy, one can actually make a well stabilised laser to our chosen resonance. This has been done, as is described in [6], using methods somewhat different from ours.

Here a setup is used in which a clever positioning of mirrors means the beam passes through the molecules a total of four times.

This laser has an uncertainty of frequency of less than 5kHz.

2.12 Cavity Enhanced spectroscopy

When doing spectroscopy, the experimenter usually wants as big and as narrow a signal as possible. The size of the signal scales with the number of moleculse or atoms encountered. But the width of the line is increased with greater pressure (more about pressure broadening in the section on broadening mechanisms).

If our medium then is a gas of molecules, the pressure will be directly proportional to the density of molecules, as will the number of molecules encountered. So, by just increasing the density, the size of a signal will increase, but at the same time it will broaden.

In order to let the light pass through more molecules without increasing the pressure, one could have a gas cell that was longer. But this strategy will not work if one wants a compact frequency reference.

But there is another way of increasing the interaction time between the light and the molecules; placing the medium inside an optical cavity.

For the purpose of simple absorbtion spectroscopy, this would mean that the light, travelling back and forth in the cavity, would seem to interact with many more molecules.

When light in a cavity with a finesse F interacts with a medium, both the attenuation and the phase shift will be increased by a factor $\frac{2F}{\pi}$. [4]

Note that in order to use this approximation, the absorbtion from the medium must be so small that the change in the finesse from it can be neglected, i.e. the loss from the medium must be much smaller than the loss from other components in the cavity.

The problem is however, if one scans a laser over the resonance of the medium, one migth scan it out of resonance with the cavity. In order to perform real cavity enhanced spectroscopy, it is necessary for the cavity to change length in order to always remain on resonance with the laser.

2.13 PDH Locking of the Cavity

There is one great problem with making spectroscopy in a high finesse cavity: In spectroscopy one wants to be able to scan the laser to the right wavelength, but a high finesse cavity will only allow a narrow band of wavelengths to stand in it.

For this use, where we already know precisely the wavelength we want to hit, one could of course try to build a cavity to exactly match the desired wavelength. But a physical cavity will always be expanding or contracting a bit because of thermal fluctuations.

For a ten centimeters long cavity, with a finesse of a hundred with light in it with a wavelength around one micrometer, it would only require a change in length of 10 nanometers to bring the light off resonance with the cavity.

There does exist very stable cavities. But in fact, if one had a cavity which always could stay exactly and sharply on resonance with ones system, you might as well throw away the molecules and instead use the cavity as a reference (and some people do use super stable cavities as frequency references).

But for our purpose we will instead use a not-quite-so-stable cavity, and put a piezo electric crystal behind one of the mirrors, just like the cavity inside the laser.

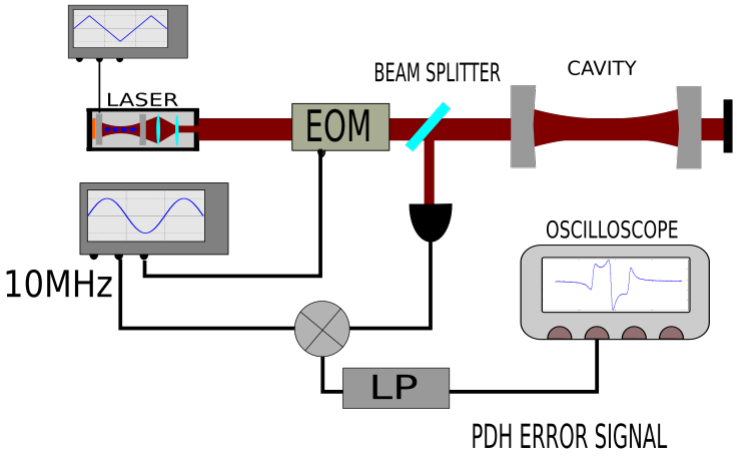

Figure 2.19: Generation of a PDH error signal. A tuneable laser is slowly scanned, and the light is phase modulated with some frequency, for example 10MHz. Some of the light reflected from a cavity is directed into a photodetector using a beam splitter. The signal from the photodetecot is demodulated with a mixer.

The length of this cavity can then be tuned just like the length of the laser cavity.

Then we need to find a way to always keep the cavity tuned to the laser, while the laser tries to tune to the molecules.

For this we will use the Pound-Drever-Hall locking scheme, or PDH scheme for short.

2.13.1 Generating the PDH Error Signal

When the laser is off resonance with the cavity, all ligth directed against the first mirror will be reflected back. When it is on resonance, a part of it will be transmitted through the cavity, and so not be reflected back. Some of the reflected ligth can be sent to a photo detecor using a beam splitter.(see figure 2.19)

Similarities between cavities and molecules Actually, the light reflected off from an optical cavity seems in a lot of ways to behave like light reflected through some absorbing medium. If one scans the frequency of the light, there will be a lorentzian dip around the resonance frequency of the cavity.

Now one might wonder if this similarity also hold for the phase of the reflected light? The case is, it does. The phase of the reflected light is shifted when the light is close to resonance with the cavity, in almost exactly the same way as when it is close to resonance with some molecules.

Light reflected from a simple two mirror cavity, where both mirrors have the same coefficent of reflection r is given as[10]

$$\tilde{E}_{ref} = F\tilde{E}_{in}, \tag{2.74}$$

where $\tilde{E}$ is the complex field, so the real field is given as $E = \tilde{E} + \tilde{E}^*$.
F is a complex number given by

$$F(\omega) = \frac{r(e^{i\phi} - 1)}{1 - r^2 e^{i\phi}}, \tag{2.75}$$

where ϕ is the total phase picked up by the light on its way through the cavity and r is the reflection coefficent of the mirrors.

ϕ can be calculated from the length of the cavity L and the angular frequency ω of the light as

$$\phi = \frac{\omega 2L}{c}, \tag{2.76}$$

where c is the speed of light.

Now imagine that we send in light with two sidebands, generated with phase modulation, so the incoming electric field is:

$$\tilde{E}_{in} = \tilde{E}_0 e^{-i\omega t} + \tilde{E}_1 e^{-i(\omega+\Omega)t} - \tilde{E}_1 e^{-i(\omega-\Omega)t}. \tag{2.77}$$

The reflected electric field is just the reflection of each component;

$$\tilde{E}_{ref} = F(\omega)\tilde{E}_0 e^{-i\omega t} + F(\omega+\Omega)\tilde{E}_1 e^{-i(\omega+\Omega)t} - F(\omega-\Omega)\tilde{E}_1 e^{-i(\omega-\Omega)t}. \tag{2.78}$$

The intensity a photodetector would measure is then

$$I_{ref} = |\tilde{E}_{ref}|^2 = I_0|F(\omega)|^2 + I_1(|F(\omega+\Omega)|^2 + |F(\omega+\Omega)|^2) \tag{2.79}$$

$$+2\sqrt{I_0 I_1}(\mathrm{Re}(F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega))\cos(\Omega t) \tag{2.80}$$

$$+2\sqrt{I_0 I_1}(\mathrm{Im}(F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega))\sin(\Omega t) \tag{2.81}$$

$$+\text{Terms} \propto \sin(2\Omega), \tag{2.82}$$

where I_0 and I_1 is the intensities of the carrier and sidebands respectively.

By demodulating with the right phase one can select out only the part that oscillates as $\sin(\Omega t)$. This then gives an electric signal as

$$S(\omega) = S_0\sqrt{I_0 I_1}(\mathrm{Im}(F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)) \tag{2.83}$$

Where S_0 is a constant depending on the electric equipment.

This signal can be readily computed, and yields a nice, useful error signal.(see figure 2.20)

The same calcultion could be made also with mirrors with different transmission coefficents, and more than two mirrors, yeilding essentially the same results.

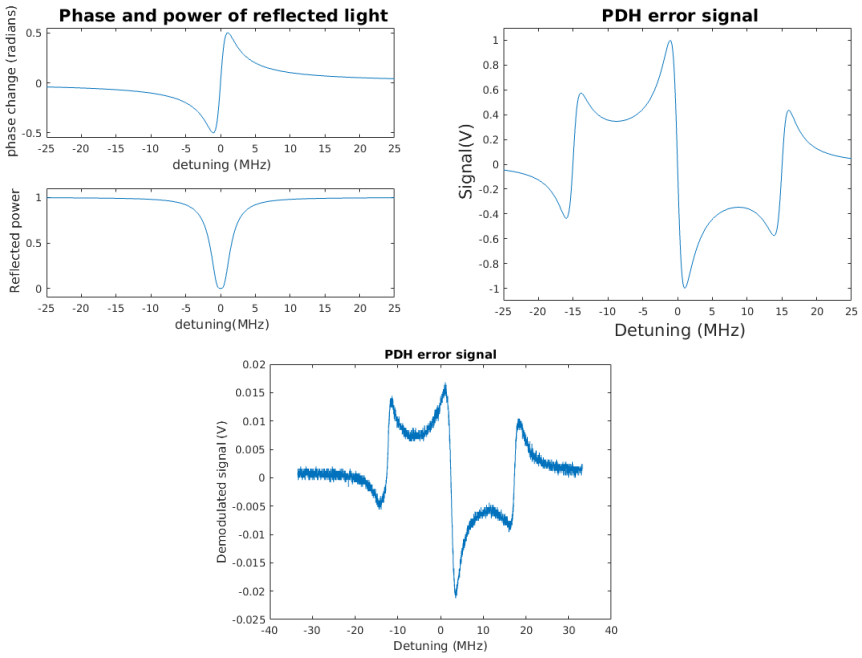

Figure 2.20: On the left: Phase change and reflected power(given as fraction of incoming power) from a two mirror cavity, where both mirrors have reflection coefficent r=0.995 and the Free Spectral Range of the cavity is 650MHz. On the right: The obtained Pound Drever Hall error signal using sidebands 15MHz from the carrier. Note that the scale of the y-axis is arbitrary, and will depend on how much one amplifies or filters the signal. Below: Actually measured Pound Drever Hall error signal from our cavity.

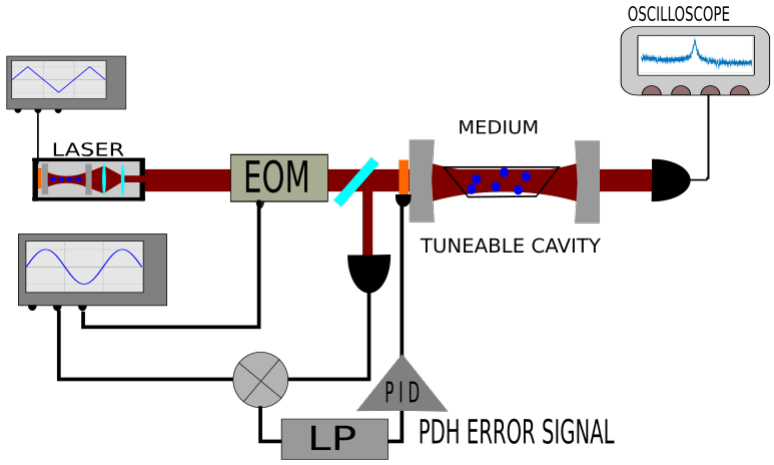

Figure 2.21: Cavity Enhanced saturated absorbtion spectroscopy. The laser is slowly scanned over a range of frequencies. The cavity's length can be changed by adding a voltage over a piezo crystal. This voltage is generated by the PDH error signal sent through a PID circuit. The cavity then follows the laser during the scan. When the laser is on resonance with the medium, the Lamb dip appears.

2.13.2 Locking the Cavity With the PDH Signal

Once the error signal can be generated, one can use it for locking the cavity to the laser's frequency.

In order to do this, the error signal is sent through a PID circuit.

A more thorough description of the PID circuits we used for our setup will be given in the relevant section.

For now we can suffice to say that we were able to lock our cavity with only a simple PI-circuit, that is

$$S_{out} = k_p(S_{in} + k_o) + k_I \int (S_{in} + k_o)dt. \tag{2.84}$$

Using this setup we could perform saturated spectroscopy on our Acetylene molecules, and measure the width of the Lamb dip.(See figure 2.21)

2.13.3 A Gas Cell for Cavity Enhanced Spectroscopy

When performing normal spectroscopy, the method used for keeping ones medium in place is not very essential. If a gas is kept in for example a glass cell, where a small amount of the light is lost in the ends, either through absorbtion or reflection, it will only mean a slight decrease in the final signal.

For Cavity Enhanced spectroscopy the finesse of the cavity depends directly on all the losses inside the cavity, and the signal will be directly proportional to the finesse.

For example: A loss of only 1% when passing through the side of a glass cell will mean a total loss 4% for a round trip of the light. If a cavity normally had a round trip loss of 2%, it would have a finesse of roughly 300. With a 6% round trip loss, the finesse would only be 100. So the signal would decrease to one third.

In order to minimise the losses in our cavity, a special cell was made, which has brewster-angled windows in both sides, for a given polarisation.

2.14 NICE-OHMS

We now have all the ingredients for explaining NICE-OHMS: Noise Immune Cavity Enhanced Optical Heterodyne Spectroscopy.

This is heterodyne spectroscopy with cavity enhancement. In order to have all sidebands transmitted through the cavity at the same time, the modulation frequency must be chosen to be equal to th FSR of the cavity. In that case, each sideband can make a standing wave in the cavity.(See figure 2.22)

The name Noise Immune derives from the fact that when all sidebands are standing in the caveity at the same time, the phase noise imparted by the cavity will affect all three equally much, and so will cancel in the final heterodyne signal

NICE-OHMS can both be used for saturated and unsaturated spectroscopy. In our case we use saturated spectroscopy.

The demodulated signal can be used as an error signal for locking the laser to the transition frequency of the medium. In that case, the slope-to-noise level of the signal, and the drift of the signal will decide how accurately it can be locked.

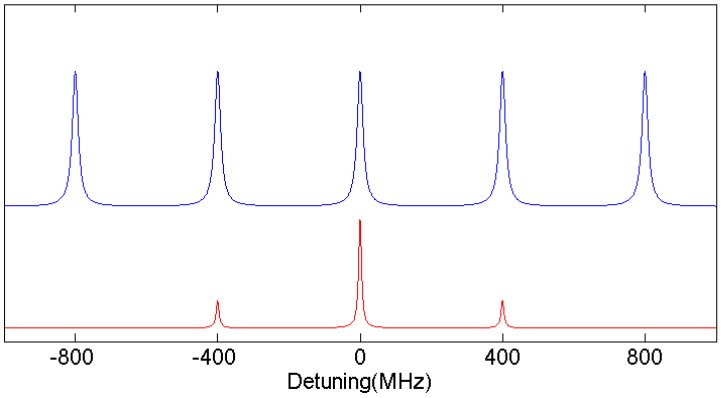

Figure 2.22: The blue line represents the transmission of a cavity with FSR=400 MHz. The red line shows the needed spectrum of the laser to perform NICE-OHMS in that cavity, that is a modulation with Ω=FSR

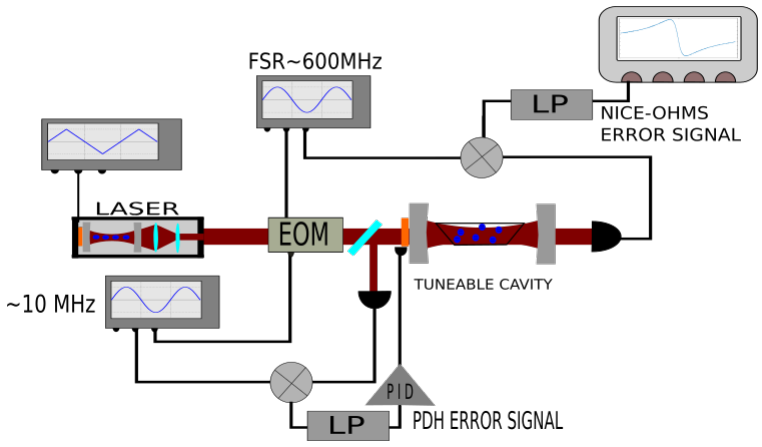

Figure 2.23: Generation of the nice-ohms signal. It is essentially the same as normal heterodyne spectroscopy, with the cavity adding extra interaction between the medium and the light.

2.15 RAM

Residual Amplitude Modulation

In the chapter on Heterodyne spectroscopy, we made heavily use of the fact that phase modulation of a laser beam leads to a perfectly balanced triplet, i.e. two sidebands whose beats with the central carrier exactly cancels.

This assumption is however not always good.

A multitude of effects can lead to either a change in the phase of one the sidebands compared to the other, or to a change in their respective amplitudes.

Or, the very phase modulation might also slightly modulate the amplitude. If the passage of light through the EOM depends on the refractive index of the crystal, the amplitude will change when an electric field is put over the crystal.

Alternatively, if the change in refractive index is different for different orientations of polarization, the EOM will work as a waveplate, which modulates the polarisation in some way, i.e. it is different for the ordinary and extraordinary axis of the crystal. This will lead to amplitude modulation if the light later encounters polarization filters of some sort.(Such as, for example, our brewster angled gas cell). Thorough calculation of this effect can be found in [11]

Amplitude modulation leads to two new sidebands, which always are in phase with each other, and so gives a beat signal out at the modulation frequency.

In that case, there will be a signal with the modulation frequency present, even when the laser is not on resonance with the medium.

This problem is known as Residual Amplitude Modulation, or RAM for short.

RAM will cause our error signal to have a drifting offset.

One can tackle RAM in two ways. One is to try to minimise the different causes of it. Work with EOM's with crystals cut to exact brewster angles with the incoming beam, temperature stabilise the EOM's, temperature stabilise everything else that could have a possible effect, etc.

Alternatively, one can try to cancel the RAM out with other effects.

It is an experimental fact that one can change the size of the RAM by putting a constant electric field over the EOM. We measure this effect by putting a variable offset voltage on the EOM, while measuring the beat between the sidebands.(See figure 2.24).

The obtained data suggests that there always exist some DC-offset value that will drive the beat to zero.(See figure 2.25) Measurements on different days gave different values for this, as internal parameters in the EOM or the laser might have drifted. We remain agnostic about the precise physical effects that gives us our data. (If it was only the effect of polarization change given in [11] we should have observed a signal varying sinusoidally with the offset).

But, if we add a control mechanism, which takes the amplitude of the beat before the molecules, and use this as an error signal. The error signal can be fed to a PID circuit, which puts an offset on the EOM, driving the error signal to zero.

So, add one more control mechanism into the machine.

To generate an error function for this, a beam splitter placed before the cavity sends a small percentage of the beam into a detector. The signal from the detector is demodulated, and this error signal is fed to a PID circuit, which output is used as the DC offset on the EOM.

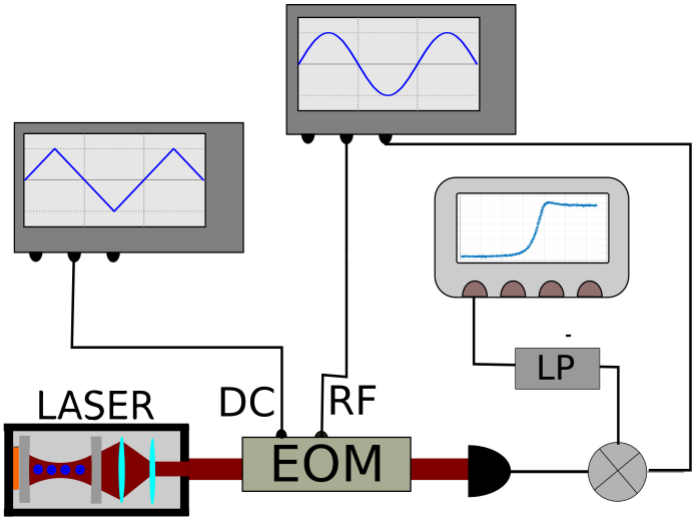

Figure 2.24: Measurement of RAM. The beat between the sidebands are measured without any medium to cause a phase change, while a DC-offset is scanned over the EOM. The RF input on the EOM is the same we use in the actual spectroscopy, around 650MHz.

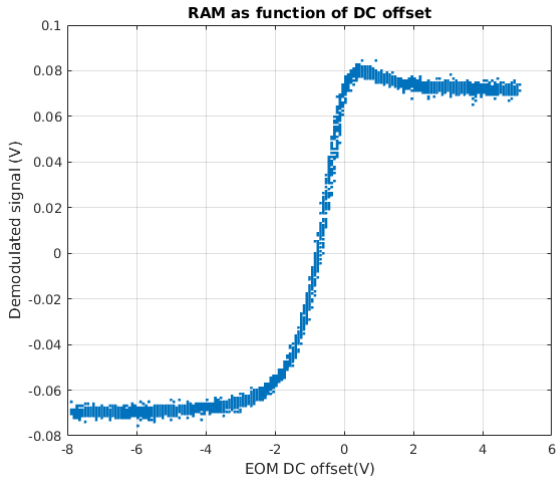

Figure 2.25: The measured beat of the triplet before passing through our molecules, when the EOM is driven with our usual modulation frequency and a variable offset. As can be seen there exists an offset, which is different from zero, that makes the beat between the sidebands and the carrier zero.

Notice that with this control mechanism it is not one particular offset we are interested in finding, but just the actual driving of the error signal to zero.

2.16 Large waist Cavities

The width of the NICE-OHMS error signal is broadened by transit time broadening. Since the molecules are moving orthogonally to the beam, they will have some finite time they actually interact with it.

This broadening effect can be reduced by having a broader beam in the cavity. But here another important consideration comes into play, which we skipped over in the first part about cavities; the stability of the cavity.

Cavity Stability Two completely flat mirrors will not make a stable cavity, since even the tiniest change in their alignment will make the light wander out of the cavity. Therefore one usually uses concave mirrors. Then in order to not change the beamshape when it is reflected, the light beam in the cavity needs to be expanding when it meets the mirror.

Now, a gaussian beam of light is defined fully by its direction, the position of its waist and the size of that waist compared to the wavelength of the light. So, in order for the beam to have the correct shape when it meets the mirror it needs to have the correct placement of its waist. If both mirrors are concave and otherwise identical, the beam waist needs to be between them, and the size the waist must have is given by the curvature of the mirrors.

Formally we can say, a gaussian light beam travelling along the z-axis with a waist a $z = 0$ of size w_0 will have a curvature at the point z given by

$$R(z) = z\left(1 + \left(\frac{\pi w_0^2}{\lambda z}\right)^2\right) \tag{2.85}$$

So here comes the problem: If one just sends in a very broad gaussian beam, the waist will be very broad, which means that the curvature of the beam will be very large. So, a normal two mirror cavity will be very sensitive to small isalignments if the beam it contains is too broad.

Large waist cavity Another approach could be to put a telescope inside the cavity. Then the beam can be made larger while it is interacting with molecules, but still have a sufficient curvature at the first mirror to give some stability(see figure 2.26).

In that case one gets both a smaller transit time broadening, and interaction with more molecules.

This is paid for by a smaller intensity pr molecule, so less saturation, and by more losses in the extra optical components inside the cavity, leading to a smaller finesse, and so even less light in the cavity.

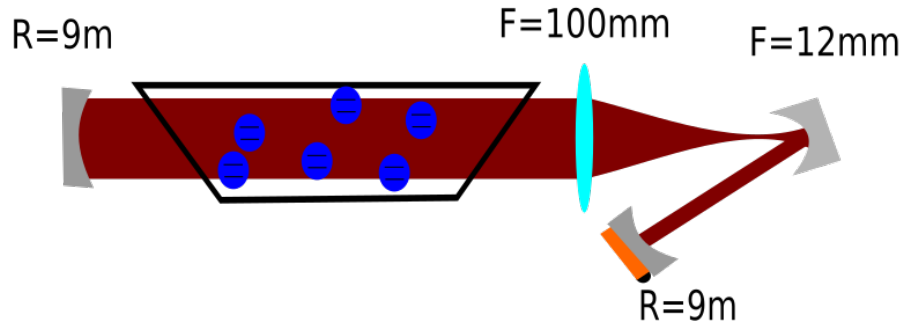

Figure 2.26: A design for a cavity with an internal telescope, made with a concave mirror and a lens to achieve a wider beam in the area where the light interacts with the medium.

2.17 Allan Deviation

At some point we will want to see if our frequency reference actually is precise, accurate and stable. Or rather, we want to say *how* precise, accurate and stable it is.

In order to give a quantitive measure of how good our frequency reference is, we need to define a new concept, the Allan deviation.

Recall that for any set of numerical data, one can find a variance, and a standard deviation, which is the square root of the variance. You might also be familiar with formulas that tells you that the expected deviation on a value which is calculated as the mean of N data points should be equal to the deviation of the N data points divided by the square root of N. Which essentially means, that to get an arbitrarily precise estimation of something all one needs to do is make arbitrarily many measurements and then take the mean.

This last formula however are only really true for data that stems from a source with a totally constant mean value with perfectly normal distributed noise added on. In reality it only too often breaks down.

One cautionary example, that is often told immediatly after presenting the one-over-squareroot-N rule, is a way of measuring either the emperor of China, or the current chairman of the chinese communist party. Let's say, instead of finding the chairman and measuring him with a measuring tape, we call every single inhabitant of China and aske them to give a guess of the height. Let's say you can guess the height of a person you only have seen on the television with an uncertainty of 10 cm. The population of China currently is 1403.5 million people, so the expected deviation you would get would be less than 3 micro meters.

To understand what exactly Allan deviation is, let us start with an example. Assume that we are measuring a constant electric signal over time, but on top of the signal there is a lot of white noise. Let us say that we take one measurement every second for ten thousand seconds. We get ten thousand datapoints. Now assume that we decide to take the data and put it in groups of ten, and take the mean of each group. So the first ten seconds are used for one data point,

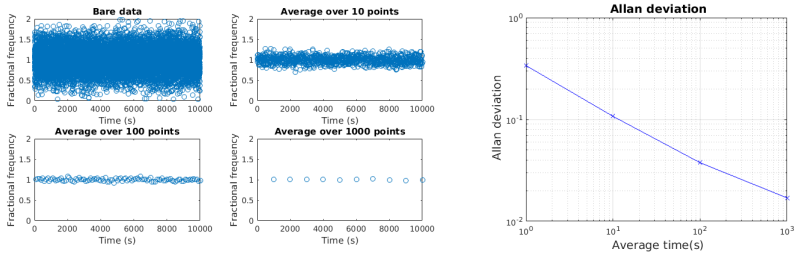

Figure 2.27: Allan deviation of simulated data with only normally distributed noise. A simulation of ten thousand data points with mean one and deviation 0.3, where we imagine we are making a measurement of a frequency, divided by some standard frequency. On the left you can see the raw simulated data, and the averages over 10, 100 and 1000 points. On the right you can see the Allan deviation for each of these averaging times, plotted in a double logarithmic system. Notice that as the averaging cancels out the noise, the Allan deviation steadily decreases.

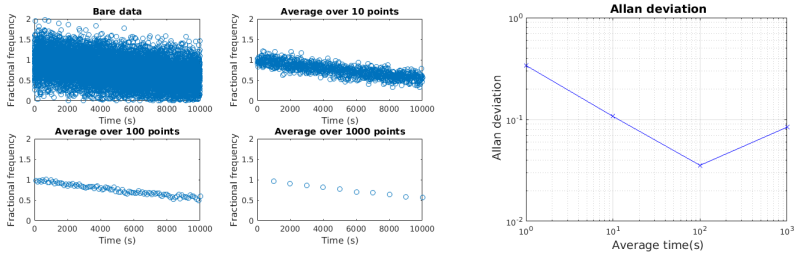

Figure 2.28: Simulated measurements of an frequency signal, where the mean of the signal is slowly decreasing with time, in addition to the noise. Again one measurement is taken pr second. On the left you can see the allan deviation of each set. Here the drift means that the difference between two consecutive points is larger for a large average, hence the Allan deviation is higher at the last point.

the next ten seconds are used to make one data point and so on. This second data set has only one thousand data points. We can do exactly the same to this dataset and get a set with only one hundred data points. And we could of course continue and get a set with only ten points in.

Since the entire standard deviation in this example is caused by normal distributed noise, the deviation of the second dataset will most likely be smaller than the first, and so on.

Now, instead imagine that there is some sort of drift causing a slow change in the mean of the signal. In that case there will remain a standard deviation, even when one makes a dataset that is made by averaging over a lot of original datapoints.

The Allan deviation is a special way of handling these situations for the particular case of a frequency measurement.

Allan deviation for two frequency references

Here we will describe exactly what is meant by the Allan deviation between two frequency references.

Let us say that we have two frequency references. In general we are not sure which is the most precise. By beating the output of them together on a photodetector we can get the difference between their frequencies.

We start out by assuming one of the references to have some angulat frequency ω_0. We feed the beat signal from the photodetector into a frequncy counter, which takes the signal over some time interval τ_1 and estimates the angular frequency of the beat $\Delta\omega$.

Let us say that we leave the experiment running for a some time T, during which we take N frequency measurements, i.e.

$$T = N\tau_1 \tag{2.86}$$

We will then have a data set of N estimated frequencies, which we can call $\Delta\omega_i$, for i between 1 and N.

We then define a new data set derived from this, which we call y_{i,τ_1} , where

$$y_{i,\tau_1} = \frac{\Delta\omega_i}{\omega_0} \tag{2.87}$$

So, y_{i,τ_1} is the dimensionless fractional difference in frequency.

Using this, we can define a third data set σ^2_{i,τ_0} as

$$\sigma^2_{i,\tau_1} = \frac{1}{2}(y_{i,tau_0} - y_{i+1,tau_0})^2 \tag{2.88}$$

This set contains N-1 points. We can then finally define one number as the mean of this last set

$$\sigma^2_y(\tau_1) = \frac{1}{N-1}\sum_{i=1}^{N-1}\sigma^2_{i,\tau_1} \tag{2.89}$$

This number is known as the Allan variance, for measuring times τ_1. From this we can again define Allan deviation as

$$\sigma_y(\tau_1) = \sqrt{\sigma^2_y(\tau_1)} \tag{2.90}$$

using the same data we can define the allan deviation for any other time τ_n where $\tau_n = n\tau_1$ and $n < N$.

We do this by making a new dataset, y_{i,τ_n}, where

$$y_{i,\tau_n} = \frac{1}{n}\sum_{j=(i-1)n+1}^{in} y_{j,\tau_1} \tag{2.91}$$

For example

$$y_{1,\tau_4} = \frac{1}{4}(y_{1,\tau_1} + y_{2,\tau_1} + y_{3,\tau_1} + y_{4,\tau_1}) \tag{2.92}$$

Notice that the dataset of y_{i,τ_2} will have roughly $\frac{N}{2}$ datapoints, depending on whether N is odd or even, and so forth. We will just call the length of the data set for time τ_n for N_n.

For each of these new datasets we can quickly define the Allan variance and Allan deviation as for τ_1

$$\sigma^2_{i,tau_n} = \frac{1}{2}(y_{i+1,\tau_n} - yi, \tau_n)^2 \tag{2.93}$$

$$\sigma^2_y(\tau_n) = \frac{1}{N_n - 1} \sum_{i=1}^{N_n - 1} \sigma^2_{i,\tau_n} \tag{2.94}$$

$$\sigma_y(\tau_n) = \sqrt{\sigma^2_y(\tau_n)} \tag{2.95}$$

Now, notice that the very concept of Allan deviation is quite distinct from the notion of a standard deviation. Standard deviation deals with the spread of values around a mean, where as Allan variance tells you the mean of the squared difference between two consecutive datapoints.

One important difference to note is that the ordering of datapoints makes no difference for standard variance and deviation, whereas it is essential for the Allan deviation.

Allan deviation can be calculated for a range of τ's and compared. Usually one will find that it falls when τ gets bigger. This happens since if the inital measurement time τ_1 is short, there will be a considerable uncertainty in the frequency measurement. But this uncertainty will average to zero for larger τ_n's.

But, usually one will see that the Allan deviation will start increasing again for big enough averaging times.

This happens because of slow drift in frequency. For the first measurement, if the drift is much slower than the timescale of τ_1, the difference between two measurements due to the drift will be negligible.

However, once one starts averaging over long stretches of time, the drift will cause bigger jumps from datapoint to datapoint. So, in general, the Allan deviation will not only stop decreasing, but might even increase dramatically for long averaging times.(see figure 2.29)

Of course, the concept of Allan deviation can be used for any sort of data, not just the fractional frequency difference between two oscillators.

2.17.1 Reading the Allan deviation

Now we have shown how the number 'Allan deviation' actually is a function of a dataset, the time interval the set was taken with, and the number of points averaged over. However, when scientists usually talk about the 'Allan deviation'

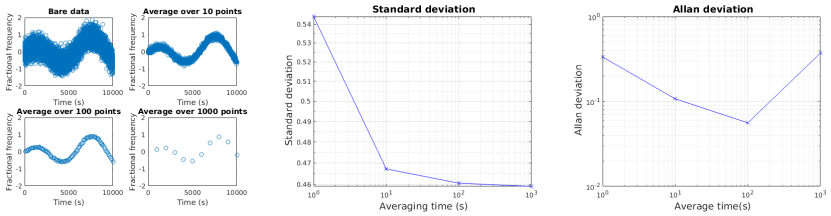

Figure 2.29: Here both the standard deviation and Allan deviation is found for a simulated frequency signal that varies slowly with time. To the left is the signal, averaged over different times. It can be seen how the standard deviation just stops decreasing as the drift becomes the main contributor to the deviation. The Allan deviation instead decreases first, as the noise is filtered out, and then increases again to around the original level.

of some thing, they actually mean the set of a number of Allan deviations for the same data, with enough different averaging times to show the general behaviour.

Let us return to our two frequency references, and their frequncy-difference Allan deviation. This difference can be attributed to three different factors: Frequency instability in the first reference, frequency instability in the second reference and measurement errors in the frequency counter.

Look at the frequency counter first. For short measurement times, there will always be an uncertainty in the measurement of frequencies. If the frequency counter is good and non biased, this error will average to zero for larger averages. If the counter is not perfect, it might be the cause for some of the measured Allan deviation, even for larger averages.

Then the two references. If we assume that the frequency counter is accurate and precise, and one of the references is much better than the other, then the entire Allan deviation can be ascribed to the least precise reference. The Allan deviation in that case can be used to show the accuracy and precision of this one. A low Allan deviation for short averages means a reference with high precision. A sharp rise in Allan deviation for higher averaging times will mean that the reference drifts on such a timescale.

If we instead assume that both references are equally accurate, they could for example be two identically built machines, then one might assume that they add the same amount to the deviation. In that case, the Allan deviation of one of them can be estimated to be the measured Allan deviation divided by $\sqrt{2}$.

For fundamental frequency references (such as the one we are building), one often beats two in principle identical references against each other to determine their Allan deviation. This is a necessity, since there aren't any better references at the same frequency, and it can be very complicated to measure a beat signal with a reference of a much different frequency.

Chapter 3

Experimental Setup

This chapter gives the details of two different setups we used during this project. First will be a description of two linear cavities, and then a short description of an earlier cavity with a broader waist. At last there is a description of our problems with drift in the length of our cavities caused by changes in temperature.

3.1 Linear Cavities

The final setups used to test the Allan Deviation was called Acetylene 2 and 3. They had all free space optical components placed in a 50cm x 35cm x 15cm large airtigth metal box.

The lasers used in the experiments were Koheras Basik fiber coupled lasers from NKT photonics. They were temperature stabilised, had a maximal output of 40 mW.

The laser was connected to an optical fiber with a 99:1 beamsplitter, which split some of the light out to be matched with the other setup. The main beam continued through a fiber coupled EOM, and through a fiber into the box containing the optics.

The fiber between the EOM and the optics was a polarisation maintaining fiber from Thorlabs. The fiber between the laser and the EOM was not polarisation maintaining.

In the box, some of the light was diverted to the RAM detector, with a beam splitter we experimentally measured to be ca 96:4 in intensity.

Then an optical isolator was placed to prevent reflected light from being coupled back into the fiber.

Two lenses was used for shaping the beam to match the cavity, and two adjustable mirrors were used to match the position and direction of the cavity.

A $\lambda/2$ plate was placed before the cavity to align the polarisation of the light.

A beamsplitter that also was measured to be around 96:4 diverted some of the reflected light into a detector for generating the PDH signal.

The detector for the transmission after the cavity and the PDH signal were DET08CLM InGaAs Biased Detectors from Thorlabs. We used different detectors for the RAM signal in the two setups.

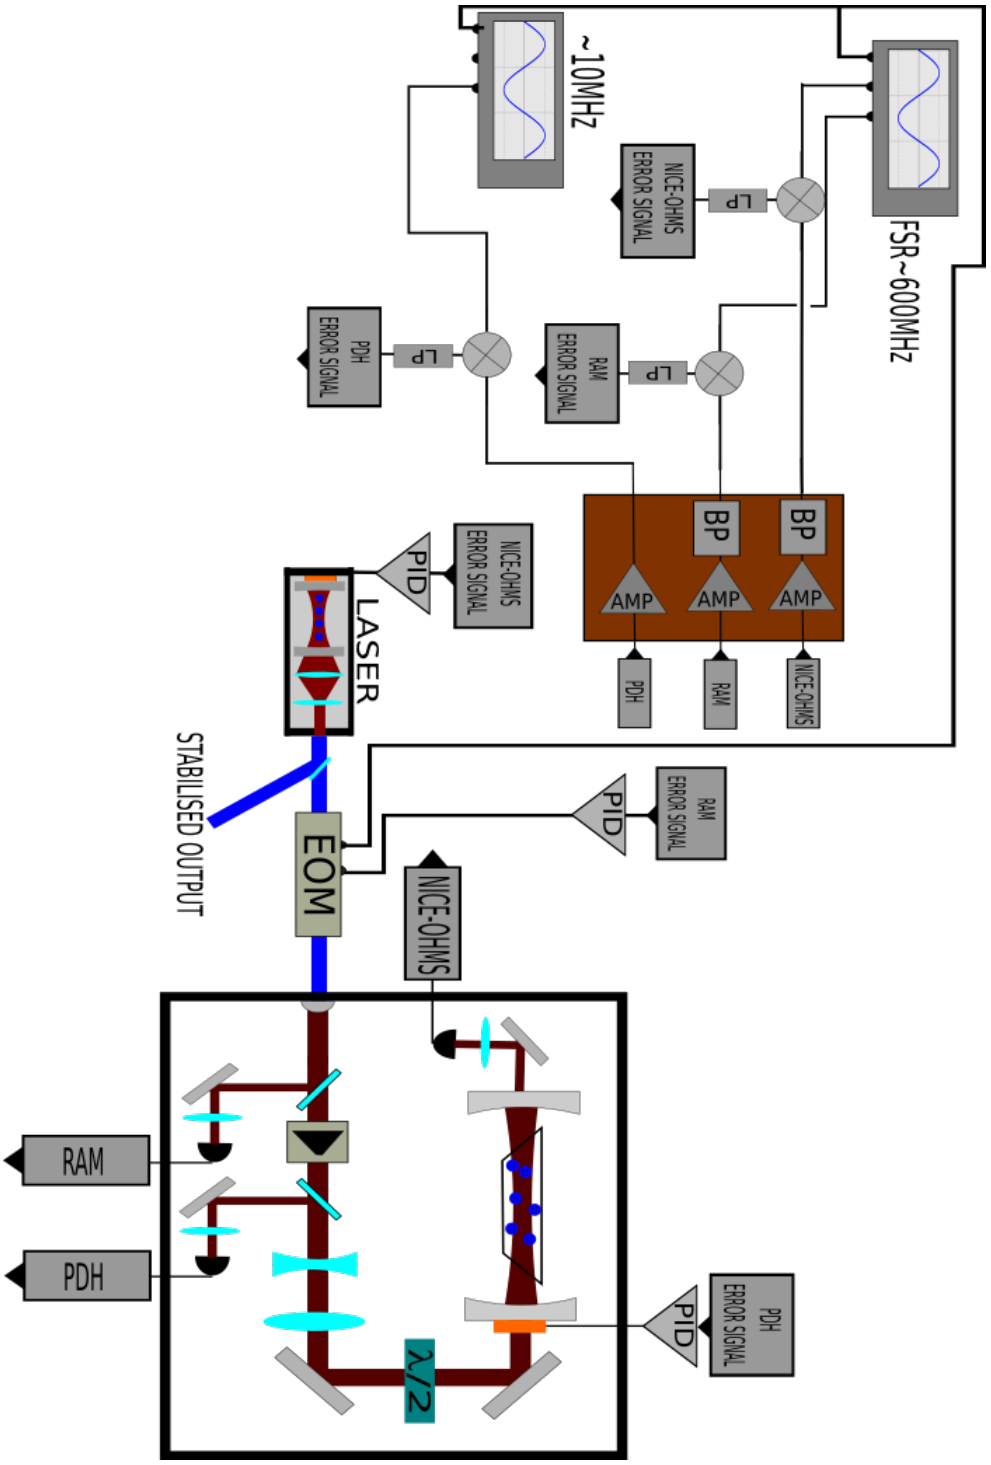

Figure 3.1: The final design used in the experiments called Acetylene 2 and 3.

The modulation frequency for the NICE-OHMS signal was generated using an AD9959 Evaluation board. It was filtered through a CBP-670F 670MHz BandPass filter from MiniCircuits, and amplified using a ZBL-700+ 12V amplifier, also from MiniCircuits.

The modulation frequency for the PDH signal was generated using a Rigol DG1022 signal generator.

The laser were controlled with an Vescent Photonics D2-125 Laser Servo, which also had an inbuilt PID circuit for locking the laser frequency.

All signals from the detectors were amplified using ZX60-P103LN+ 5V amplifiers from minicircuits. The NICE-OHMS and RAM signal were also filtered through a 670 MHz BandPass filter.

All these 5V amplifiers wer placed in one closed metal box, with a solid copper bottom, which was used as a common ground.

The different signals were passed using thick co-axial SMA cables.

The cavity was locked using an IPRE High Speed Low Delay $(P-I)^2D$ Loopfilter. The signal from this was passed throug a x5 amplifier, were we also could add an offset to find the resonance with the cavity.

3.2 Bent Cavity

An earlier setup was made with a telescope in the cavity, to broaden out the laser beam.

This setup was only used to measure Lamb dips. It never had the desired stability to lock a laser to it.

Our cavity was made with two end mirrors with a radius R=9m, and a concave mirror with a focal length of 12mm along with a lens with a focal length of 100mm. The mirrors was set up so the incoming beam would hit the concave mirror with an incident angle of 8 degrees (see figure 3.2).

The lengths involved was: Between the first mirror and the concave mirror: 65mm. Between the concave mirror and the length: 110mm. Between the lens and the last mirror: 240mm.

It had a measured finesse of around 125, but this number could drift with up to 10% during a day, due to thermal expansion or contraction in the plate it was mounted on.

This cavity was not mounted on zerodur, but on a plate of aluminium.

We could achieve a total incoupling into the cavity of around 50-60% of the incoming beam.

The PID for locking the cavity was a simple PI circuit, we had built our selves using operational amplifiers. The PDH signal had to have an offset added to it in order to lock to the centre of the resonance of the cavity. This phenomenon was not explored further.

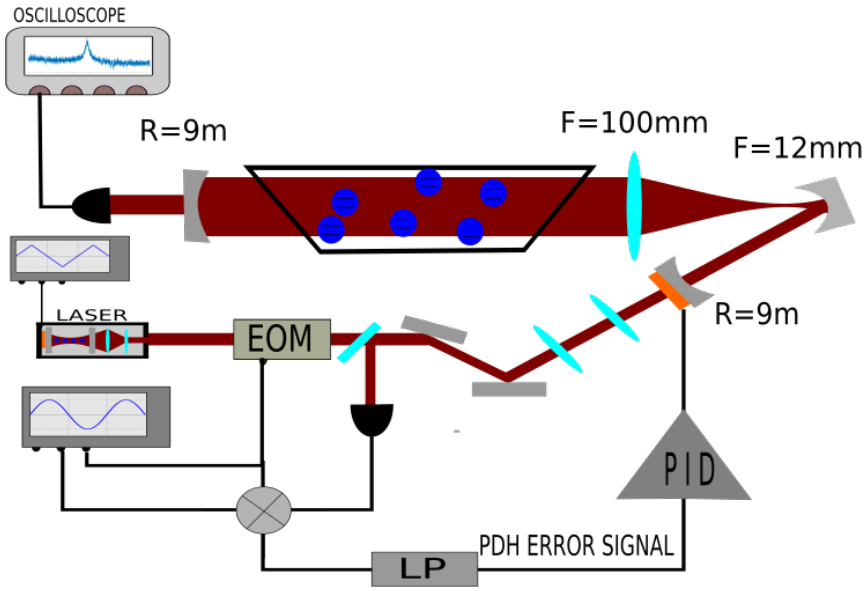

Figure 3.2: Simple schematic of the Acetylene-1 setup. Two mirrors and two lenses were used for shaping the beam to the desired mode. The beam was expanded with a F=12 mm concave mirror and a F=100 mm lens. The incident angle with the concave mirror was 8 degrees. The beam had a gaussian width of 3 mm during interaction with our molecules.

3.3 Our Cells

The acetylene-13 samples used in this experiment were kept in three special made cells. The cells are cylindrical with the two ends angled in exactly the brewster angle for the light we use(see figure 3.3).

I label these cells cell 1, cell 2 and cell 3, according to which of our experimental setups they were used in.

We could not be sure of the exact pressure in the cells, but we initially believed that it was close to 1Pa, perhaps somewhat more.

A fourth cell seemed to be contaminated with more gasses than acetylene-13, since it seemed to show more resonances than the others, and had a much broader linewidth. In addition it seemed to contain much more acetylene than the others, since it needed a much higher intensity before saturation could be seen. The cell was therefore not used in the main experiment.

In addition to the three brewsterangled cells, we had a cylindrical cell with normal sides, containing 1 kPa of acetylene-13. This meant that by just shining a laser directly through it, the absorbtion would be great enough to show us if the laser was close to the resonance or not.

Since our laser had two exits, one containing the main beam, and another with only a few mW of intensity, we could simoultaneously measure our signal from the cavity and from the 1 kPa cell.

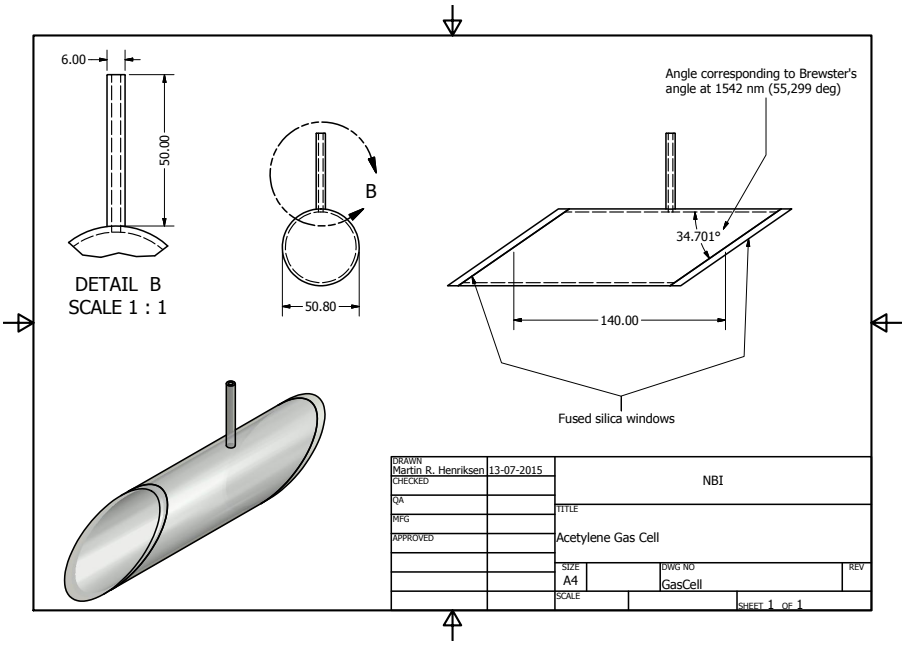

Figure 3.3: Schematics for the cell we used to contain our Acetylene molecules. Norice the precisely angled windows.

3.4 Thermal Drift

In an earlier version of the setups Acetylene 2 and 3 we had not placed the cavity on a piece of zerodur, and we kept the amplifiers for the PDH and NICE-OHMS signal inside the box with optics, using the aluminium board as a common ground.

With this setup, we had a lot of drift of the cavity length of the cavity. We could measure the drift in the cavity by measuring the output from the PDH locking PID, which kept it on resonance with the laser. In addition, the finesse of the cavity could change with the drift, causing a drift in the intensity transmitted through it.

In order to measure if this was caused by a thermal expansion of the aluminium bread board, we made a setup where we would keep the cavity locked to the laser, while at the same time measuring the temperature of the bread board in the middle of the cavity.

The temperature was measured using a thermistor placed in a lump of thermal gel on the bread board, held in place with a piece of tape.

At the same time we measured the output of the PDH PID.

Both temperature and output were taken with data measured once pr second.

A number of measurement series were made with the amplifiers in place, and normal air in the box, for five minutes at a time. The amplifiers were turned on shortly before the measurements began.

The amplifiers were then taken out of the box and the air pumped down to ca 0.1 Atmosphere. Another series of measurements over five minutes were taken.

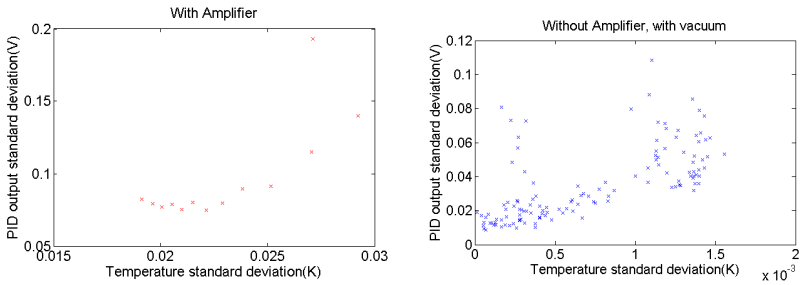

Figure 3.4: The standard deviation of the temperature and the PID-output. All points are standard deviations of measurements taken over five minutes. The graph on the right is taken with amplifiers placed on the breadboard. The graph on the left is with the amplifiers removed, and vacuum inside the box. Notice the difference on the x-axis

For every measurement, the standard deviation of the temperature, and the output of the PID was calculated.(See figure 3.4)

It can be seen that the average deviation of the temperature fell with a factor hundred after removing the amplifiers.

It can also be seen that the change in the PID output correlates with the temperature change, from which we could infer that the cavity was indeed affected by thermal expansion and contraction.

To give an image of roughly how much this change was, we can mention that a change of around 10 Volts from the PID circuit would give a change of one FSR of the cavity, and hence would equal a change in length of 1542nm

After this, we decided to remove all amplifiers from the optics and place it in a separate box. We also decided to mount the cavity on a piece of zerodur.

The length of the bent cavity in the setup Acetylene 1 was even stronger than the one in the linear cavities. The finesse of this cavity would also drift a lot. This could probably also be explained by thermal expansion. If we were to proceed with a bent cavity, we would also have to mount that on zerodur, or some other invariant material.

Chapter 4

Results and Discussion

This chapter contains a quick calculation of expected linewidths. After that is an estimation of our linewidth from lamb dips measured with a bent cavity, after which there is an estimation of linewidths from our NICE-OHMS signal measured using two linear cavities. Then we present a measurement of our current Allan deviation, which appears to be affected by frequency drifts on time scales of seconds.

The reason for our drift is investigated. It seems to be correlated with drift in intensity, but the cause remains unclear.

At last there is a discussion of why we chose to use a small waisted cavity, followed by a discussion of the possible gains from changing to a cavity with a broader beam.

4.1 Expected Linewidths

Using saturated spectroscopy, we are not affected by doppler broadening. The remaining broadening mechanisms will be dealt with here

4.1.1 Natural Linewidth

We estimate the natural linewidth of the $P(12)v_1 + v_3$ transition to be slightly less than one Hz, using data from [7]. Assuming that our $P(16)v_1 + v_3$ line has a comparable natural width, we can neglect it here, since other broadening mechanisms are of an order 100 000 times larger than it.

4.1.2 Pressure Broadening

The pressure broadening coefficent of Acetylene-13 for the P(16)v_1+v_3 transition has been measured to be [6] 225 ± 10 kHz/Pa for the FWHM.

We are not sure of the exact pressure in our cells, but think it is close to 1 Pa.

4.1.3 Transit Time Broadening

We have used two different cavities at different times in the experiment, one with a waist of 0.7mm and one with a waist of 3.0mm.

The molecular mass of acetylene-13 is 28 atomic units, and the temperature of our laboratory was around 21 degrees celsius, equivalent to around 294K.

Since the Transit time broadening for the FWHM is [9]

$$\Gamma_{tt} = \frac{\sqrt{2\ln(2)}}{2\pi}\sqrt{\frac{\pi k_b T}{2m}}\frac{1}{w} \tag{4.1}$$

we get for the narrow waist

$$\Gamma_{tt} = 99.3\text{kHz} \tag{4.2}$$

and for the large waist

$$\Gamma_{tt} = 23.2\text{kHz} \tag{4.3}$$

4.1.4 Wavefront Broadening

We used a cell that was roughly 16cm long to contain our molecules. Assuming that the beamwaist is located in the middle of the cell, the furthest distance from waist to a molecule will be 8cm.

If we use the narrow waist, this will give a maximal radius of beam curvature any molecule can meet of $R(8\text{cm}) = 12.5\text{m}$.

The total of transit time and wavefront broadening is

$$\Gamma_{tt+wf} = \Gamma_{tt}\sqrt{1 + \left(\frac{\pi w^2}{R\lambda}\right)^2} \tag{4.4}$$

Which at maximum curvature gives

$$\Gamma_{tt+wf} = \Gamma_{tt} \cdot 1.0032 \tag{4.5}$$

So a change of less than one percent. Since this effect is so small (much smaller than the uncertainty on the estimation of our waist size) we will neglect it completely.

4.1.5 Table of Expected Linewidths

We expect the pressure in our cells to be 1Pa, perhaps somewhat more.

According to [6] the total FWHM of our transition at these pressures and for the intensity $I \ll I_s$ is given as

$$\Gamma = C_p p + \Gamma_{tt} \tag{4.6}$$

where C_p is the pressure coefficent and p is the pressure

The expected FWHM of our transition are then

Waist=0.7 mm	324 kHz
Waist=3.0 mm	248 kHz

4.2 Pressure Shift

In addition to broadening the linewidth, pressure can shift the resonance frequency of a molecule.

The pressure shift coefficent for the $P(16)v_1 + v_3$ line in Acetylene-13 has been found to be 0.3 ± 0.1kHz/Pa [13].

So a change of one Pa in pressure will only give a fractional change in frequency of $1.5 * 10^{-12}$.

4.3 Linewidth Estimated from Lamb Dips

To measure the linewidth, and possibly the saturation intensity, we measured the lamb dips in a bent cavity (the Acetylene 1 setup) for different intensities. In all in this experiment we made 120 measurements over two days.

The finesse of the cavity with the cell was around 120, and the incoupling was measured to be around 70%. The finesse was however not totally constant for different voltages over the piezo crystal, different lengths of the cavity had slightly different finesses.

Through the experiment the length of the cavity drifted quite strongly. It drifted out of the range of the PID over roughly one hour, so every hour the lock had to be turned off, and the length of the cavity turned down one FSR where after it was locked again to the laser.

The waist of the beam inside the part of the cavity containing the cell was around 3mm. It was impossible to measure this exactly. It was 3.1mm before the first incoupling mirror was installed, but after this some minor adjustments had to be made to couple into the cavity, and it was impossible to measure the waist inside the finished cavity.

The lamb dips were fitted with a lorentzian profile in order to estimate their FWHM. This was however made harder by them often being located on a non-linear background, due to the coupling between the length of the cavity and the intensity in the cavity. Vibrations and drift in the cavity during the measurements might also have cause more non constant background.(See figure 4.1).

In these figures the intensity experienced is defined as

$$I = \frac{P}{\pi w^2} \tag{4.7}$$

Where P is the total power in the cavity, and w is the waist of the beam.

The width of the dips did not seem to be in any way correlated with the intensity in the cavity. (See figure 4.2) It was however, due to the low finesse not possible to put more than ca 450mW power inside the cavity, even with the laser at full output, which coupled with the large waist made for a very low intensity compared with the linear, small-waisted cavity. So the intensity might be much too low compared with the saturation intensity, I_s, to show any power broadening.

The mean of the different FWHM was found to be 1.74 MHz. If the cell only contained Acetylene, and subtracting the transittime broadening, this would translate into a pressure of 7.6 Pa.

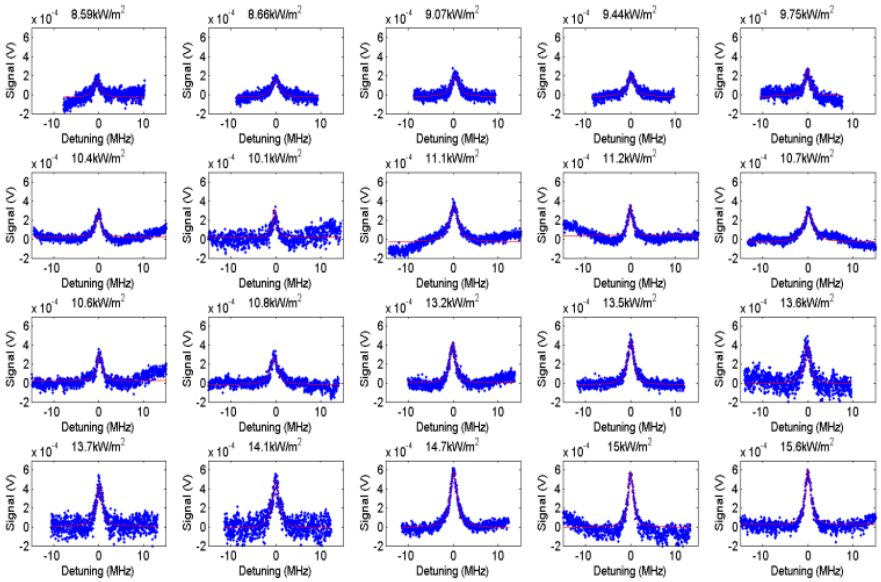

Figure 4.1: A selection of 20 out of the 120 measured dips. The background has been subtracted from the signal, to make it easier to fit a Lorentzian lineshape to the lamb dip.

However, this does not mean that the cell necessarily contains 7.6 Pa of Acetylene. There might be other gasses inside the cell, who can have other pressure broadening coefficents. There is simply no way to know how much of the broadening is caused by such unknown gasses.

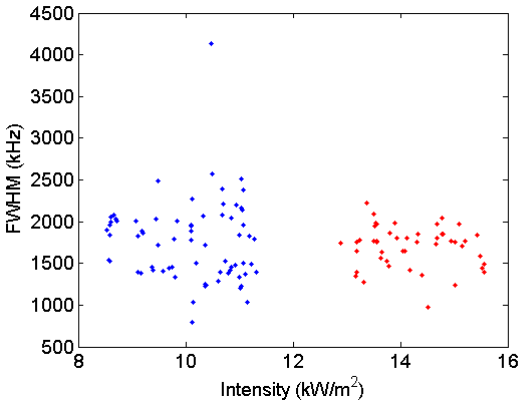

Figure 4.2: The fitted FWHM of the measurements. The blue and red dots are from two series with one night between. The finesse or incoupling was slightly higher during the second series, which is the reason for the sudden jump in intensities. As can be seen, the line does not broaden out at the used intensities

4.4 Linewidth Estimated from NICE-OHMS

A number of measurements of the nice OHMS signal was taken for both of the linear cavity setups, named Acetylene 2 and Acetylene 3. After locking the cavity and finding the resonance, a total of nine measurements were made with one second between for each of fourteen different powers.

The data first had their linear background subtracted, and then were fitted to a standard lorentzian dispersive function

$$S(\Delta) = A\frac{\Delta}{\Delta^2 + \gamma^2} + C \tag{4.8}$$

Where A is a fitting constant, and γ would equal the HWHM of a standard absorbtion lorentzian.(See figure 4.3 and 4.4).The value 2γ is referred to here as the FWHM, despite not directly being the width at half maximum of a dispersive function.

The intensity used here is again defined as the total power divided by πw^2. We were able to reach total cavities powers above 2 W with these cavities.

The mean was taken of the FWHM of the 9 measurements with the same intensity, as an estimate of the true FWHM. The standard deviation of the nine values were used as an estimate on the uncertainty.

In addition, the slope at zero detuning was calculated from the fitted function, including the contribution for a linear background.

This was compared with the RMS noise, which was measured after the measurements, for the laser detuned far from the saturated NICE-OHMS signal. The noise at maximum power and the same electronic setup was found to be 1.2mV for Acetylene 2, and 4.2mV for Acetylene 3.

We fitted a standard power broadening profile to the measurements of the

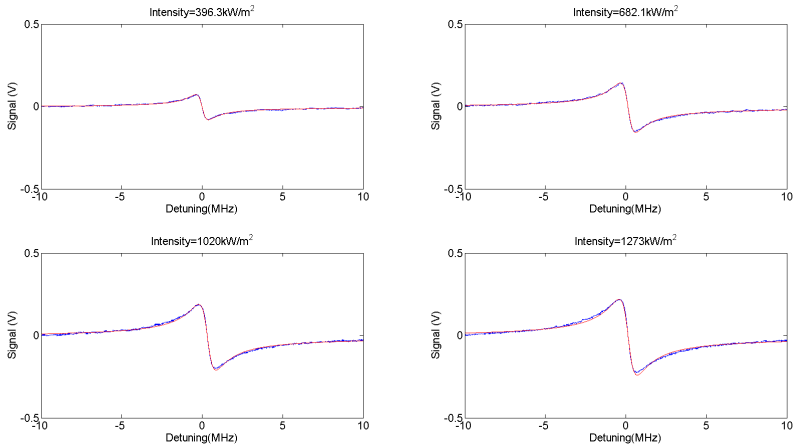

Figure 4.3: Four of the measurements from the Acetylene 2. The red line is the fit of a lorentzian dispersive function.

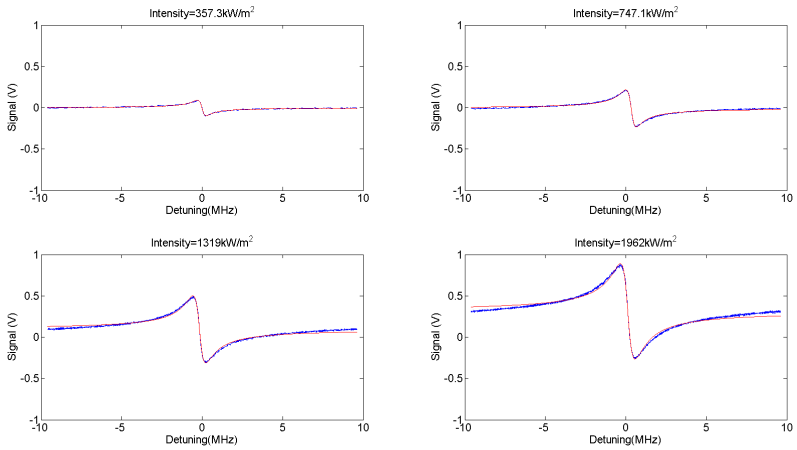

Figure 4.4: Fits of four of the measurements from the Acetylene 3 setup.

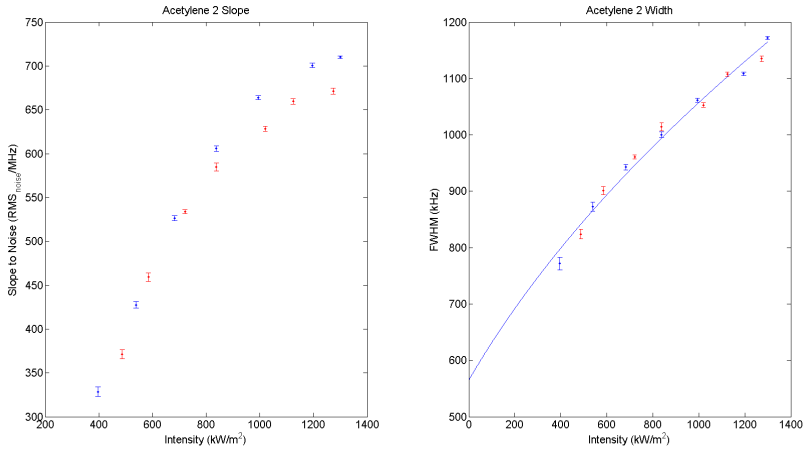

Figure 4.5: Measurement on setup Acetylene 2. The slope divided by the noise level, and the FWHM width of the dispersion signal for different intensities. The blue and the red dots are for two different sets of measurements taken one hour after each other.

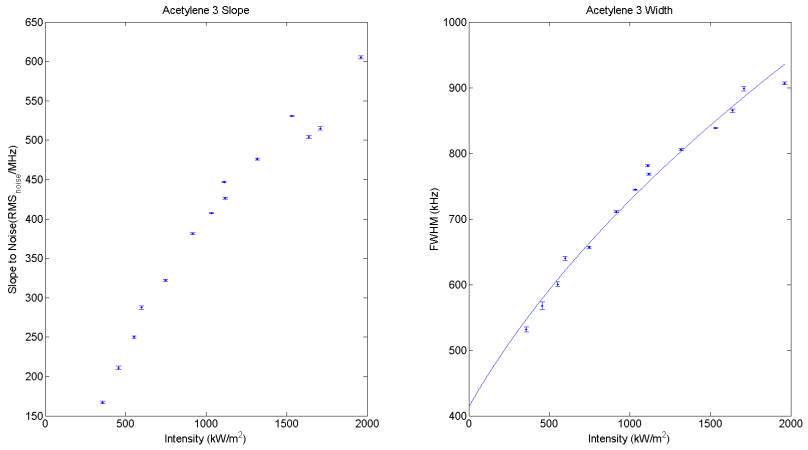

Figure 4.6: The same data for Acetylene 3

FWHM

$$\Gamma = \Gamma_0 \sqrt{1 + \frac{I}{I_s}} \tag{4.9}$$

Where Γ is the FWHM, I is the intensity and I_s is the saturation intensity.(See figure 4.5 and 4.6). Γ_0 will then be our estimate of the FWHM for zero power.

The results of the fit are here summarized

Acetylene 2	Acetylene 3
$\Gamma_0 = 565 \pm 24$kHz	$\Gamma_0 = 414 \pm 17$kHz
$I_s = 400 \pm 49$kW$/m^2$	$I_s = 477 \pm 56$kW$/m^2$

Here the uncertainties are our 68% confidence intervals. We expect the value of I_s to be the same, depending only on the molecules themselves, which it is within uncertainties.

If we assume the cells only contain Acetylene, we will estimate, using the FWHM at zero power, the pressure in the cells to be:

Acetylene 2: 2.07 Pa

Acetylene 3: 1.40 Pa

In reality, the difference from our expected 1 Pa might be because of pressure from other gasses, with unknown pressure broadening coefficents.

If we compare our results to the results obtained with the older cell in the large waist cavity, we can see why it is not surprising we could not measure the power broadening. The saturation intensity apparently is somewhere above 400kW/m^2 and in that experiment we only reached a maximum of 15kW/m^2.

A naive interpretation of the slope-to-noise level, which seems highest for high intensities reaching more than 500 noise-levels pr MHz would be that we could, using a perfect locking mechanism, stabilise the frequency of the laser within $\frac{1}{500}$MHz = 2kHz of the desired frequency. Since our transition is pretty close to 200THz, this would mean a fractional stability of around 10^{-11}.

Despite the naivity of this calculation, we are to see that it comes pretty close to our short term stability.

4.5 Measured Allan Deviation

We measure the Allan deviation of our stabilised lasers by taking, using a fiber beam splitter, a small amount of the laser beam from both lasers and beating these two on a photodetector.

Since the two setups are identical, we can assume they have the same stability. In that case the deviation attributable to one of them will equal $1/\sqrt{2}$ times the total deviation.

The experiment was made by first locking both cavities to the lasers, then finding the resonances of the molecules. Data was taken of the NICE-OHMS signal and the noise level before the lasers was locked to the signal.

The beam was split up from the main beam with a 99:1 fiber beam splitter, and coupled in 1 50:50 fiber beam coupler.

For this experiment both of the lasers was at their maximal power output of 40mW. Because the metal boxes around the cavity was closed during the

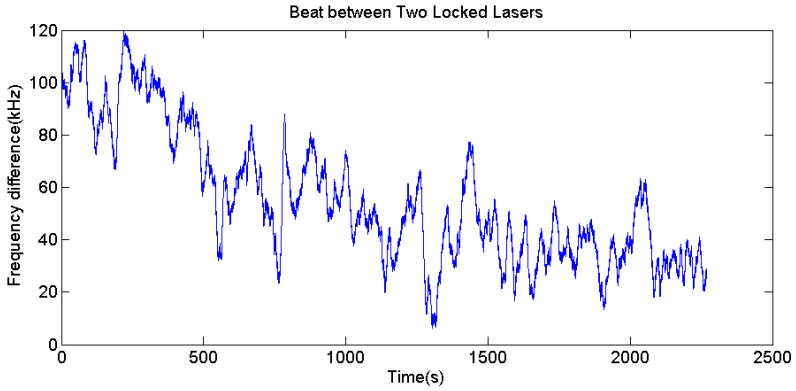

Figure 4.7: The absolute frequency difference between our two stabilised lasers, while both were locked. The measurement was made by slightly detuning both lasers, so there initially was a 100kHz difference, and then beating the output of both lasers on one photodetector.

experiment, it was not possible to measure exactly what the intensity in the cavities were, but our earlier experiments suggests that it would be somewhere close to 2W.

The boxes were not emptied of air during the experiment, there were normal atmospheric pressure in them.

The Lasers themselves were both locked using a D2-125 Laser Servo from Vescent photonics. RAM was supressed using two simple PI-circuits we had built ourselves. During the experiment, the output of the different PID circuits was monitored using a National Instruments evaluation board.

The beat note was registered with an Agilent Multimeter, with a shutter interval of 0.01 seconds.(See figure 4.7)

The Allan Deviation lies around $5 \cdot 10^{-12}$ for an averaging time of one second, and then drift upwards towards 10^{-10}.(See figure 4.8) It might well become higher for higher averaging times.

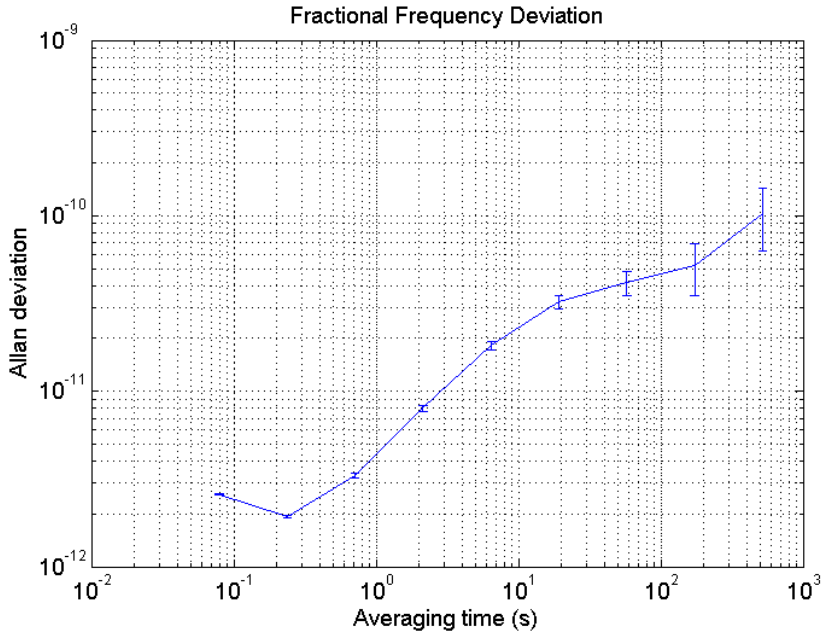

Figure 4.8: The Allan Deviation of the fractional frequency difference.

4.6 Intensity and Frequency Drift

We see that our Allan deviation is limited by drift on a time scale of one second or more.

In the current setup we experience some intensity drift. The reason is mainly our use of not-polarization-maintaining fibers. Before our cavity there is an optical isolator, which has a polarization filter. If the polarization of our laser changes in the fiber, it will lead to a change of intensity after the isolator.

We seem to experience a change in offset of our error signal, when the intensity of the laser changes.

Here we will investigate these effects.

4.6.1 Offset from the Cavity

In NICE-OHMS, if both carrier and sidebands are transmitted through the cavity with exactly the same detuning from the cavity, there should be no measureable signal, unless they interact with the molecules in the cavity.

This is since any change in phase or amplitude stemming from the cavity will affect all three components in exactly the same way, and so should not cause any beat between them.

This can only happen when the modulation frequency is exactly equal to the FSR of the cavity.

We generate our modulation frequency with an Analog Devices 9959 evaluation board, clocked with a signal of 500MHz, which in principle should give us a the ability to control the modulation frequency with a precision of

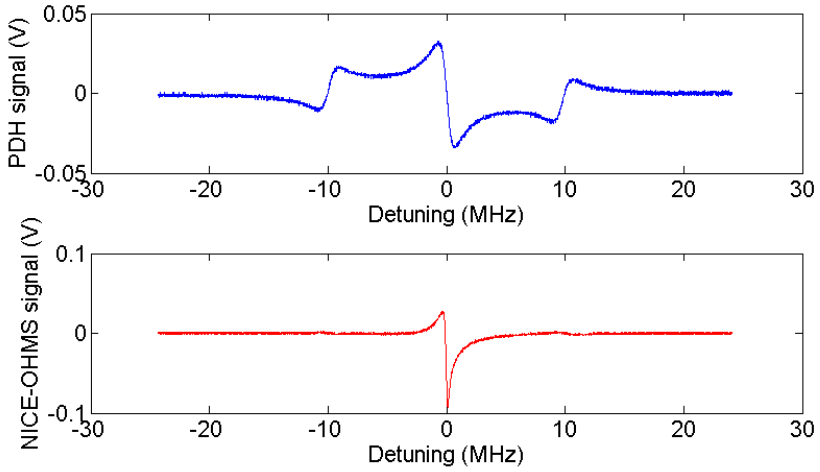

Figure 4.9: A scan over resonance of the cavity in Acetylene-2 where the resonance has been detuned far away from resonance with our molecules. The lower signal is slightly asymmetric since the modulation phase has been optimised to the signal from the molecules, and so is not quite matching the signal from the cavity.

$$\Delta\Omega = 500\text{MHz}\frac{1}{2^{31}} = 0.2\text{Hz} \tag{4.10}$$

This precision is quite high enough to be able to align the modulation frequency with the FSR

When scanning the frequency of the laser over resonance with the cavity, without locking the cavity, we can pick up a signal on the NICE-OHMS diode.(see figure 4.9)

This signal can be made larger by changing the modulation frequency away from one FSR. We have however not been able to make it completely diasppear, not matter how precisely we try to tune the modulation frequency. The illustration shows the smallest signal we were able to get.

With these conditions, it is possible to change the offset of the NICE-OHMS signal to the molecules by changing the setpoint at which the cavity locks to the PDH signal.

4.6.2 Cavity Offset versus Intensity

If the cavity is locked to the laser in a way that gives an offset to our NICE-OHMS signal, we would expect that offset to be proportional to the intensity of the light.

In order to test this hypothesis, we installed an extra photo detector in the setup. This photodetector detects the light that is redirected from the beam when it first passes the beamsplitter for making the PDH signal. (The light

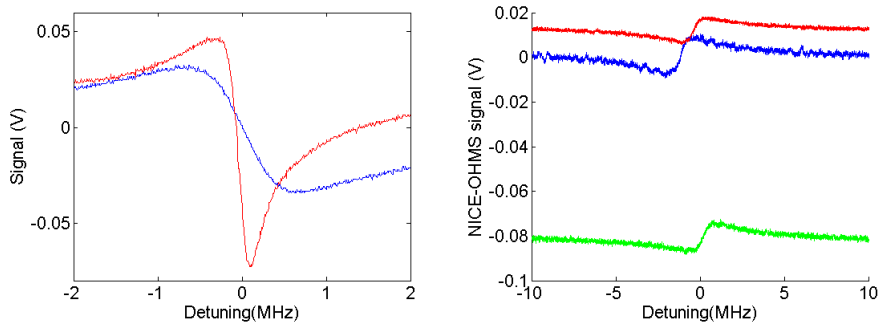

Figure 4.10: On the left, a closeup from figure 4.9, where both signals have been plotted on the same axis. The NICE-OHMS signal from the cavity (the red graph) have had 20mV added to it to make the connection between the signals clearer. The cavity can in principle be locked to any point on which the PDH signal has a negative slope. On the right, the NICE-OHMS signal to the molecules, with the setpoint of the PDH lock in three different poistions. This is taken with the cavity locked to the laser. As can be seen, the cavity can create an offset that is much larger than the signal itself. From the Acetylene-2 setup.

for the PDH signal is taken when the beam returns after being reflected from the cavity). The intensity of this light should be directly proportional to the intensity of the light reaching the cavity.

We then made an experiment, where the laser were detuned far from the molecules, and the cavity was locked to the laser. The setpoint of the cavity lock was adjusted to make for as large a (negative) offset as possible.

A number of measurements were then made, where the intensity of the laser was gradually lowered from 40mW to 12mW. At every intensity, the voltage of the diode was noted down, and the intensity of the NICE-OHMs offset was measured, using an oscilloscope over 0.01 second(see figure 4.11).

The data seems to confirm that the offset depends on the intensity, and so any intensity drift would be converted into offset drift, and hence into frequency drift if the laser was locked to the NICE-OHMS signal.

To make sure that this is not actually an artefact of RAM, the RAM was locked during the experiment, and the RAM error signal was measured. The RAM error signal, and the PDH error signal can be seen in figure 4.12.

4.6.3 Initial Offset of Zero

If the offset on our NICE-OHMS signal was caused entirely by the cavity, and the drift in this was then caused by drift of intensity, we should then in principle be able to negate any drift by locking the caity in a way which sets the offset to zero.

This we tried on Acetylene-2. First the laser was set to maximal output, the cavity was locked, and the setpoint on the PDH signal was chosen so the offset was ca zero.

Then we gradually lowered the intensity like before. The result was however still a drift in intensity, only somewhat smaller than in the earlier measurement

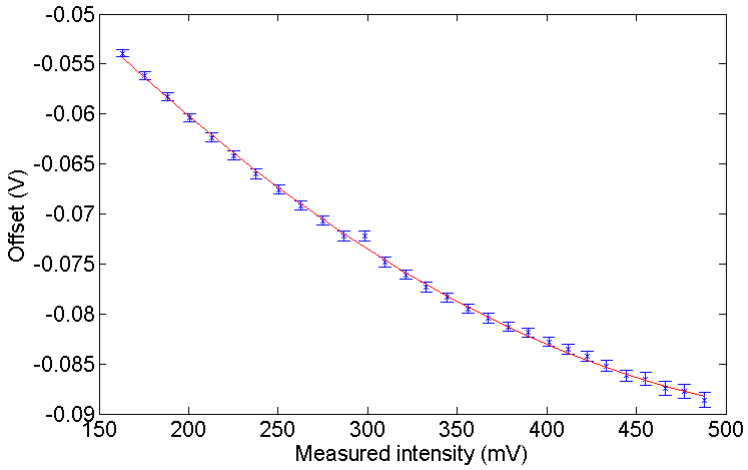

Figure 4.11: Measurement of the DC-Offset of the NICE-OHMS signal, far detuned from resonance with the molecules as a function of the intensity. Made in the Acetylene-2 setup. The output of the laser was changed from a max of 40mW to a minimum of 12mW. The intensity here given is the measurement of a photodetector placed before the cavity, and so includes any possible intensity drift. The offset changes with almost 40mV. In comparison, the peak to peak value of the NICE-OHMS signal to the molecules is ca 20mV. The red line is a second order polynomial fitted to the data. The errorbars is the standard deviation of the offset measured over 0.01 second.

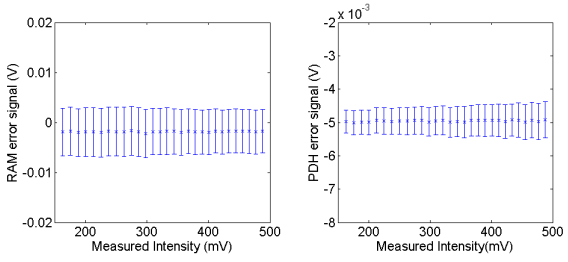

Figure 4.12: RAM and PDH error signals measured while both were locked while scanning the intensity, at the same time as the data in figure 4.11 was taken. As can be seen, the RAM does not change with intensity, and so can not be the cause of the intensity dependent offset

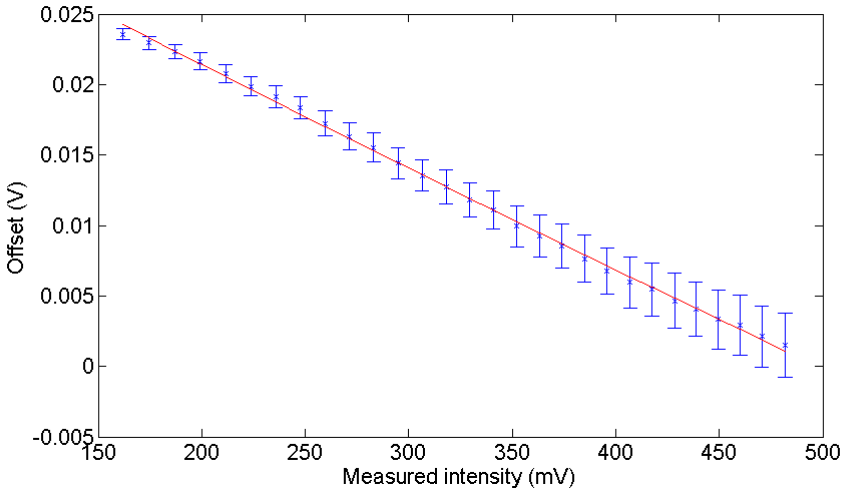

Figure 4.13: A measure of NICE-OHMS offset far from resonance as a function of intensity. Here the offset was initially set to roughly zero (the measurement with the greatest intensity). The offset however still shows a clear correlation with the intensity. Measurement done on the Acetylene-2 setup.

(with a total change of around 25mV instead of 35mV)(see figure 4.13).

Again during this experiment, the RAM was locked, and a measurement of the RAM error signal shows that this did not change with intensity.

Apparently, locking the cavity in a way where it does not cause an offset is not possible.

One possible explanation for this could be that the connection between intensity and offset is not soleley because of the offset from the cavity.

Another explanation could be that the change in intensity also changes the amplitude of the PDH signal. That means that a setpoint, which has been chosen by adding some offset to the PDH signal before feeding it into a PID, will no longer lock the cavity at the same exact detuning with the light if the intensity changes.

If the second explanation holds, the change in offset might actually be because the PDH signal changes, and so the setpoint of the cavity changes.

4.6.4 Noise Dependence on Setpoint

Another thing which can be inferred from figure 4.11 and 4.13 is that the noise on the signal is larger when the cavity is locked in a place where the NICE-OHMS offset is close to zero.

It should here immediatly be mentioned that the noise measured here cannot be directly compared to the noise in the measurements of the slope-to-noise value, since in those measurements additional filters, inbuilt in the servo for locking the laser was included.

The change in noise based on the PDH setpoint can be directly seen when

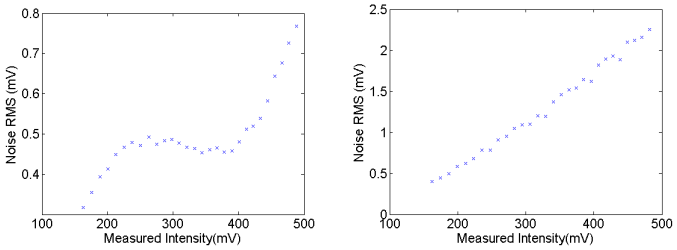

Figure 4.14: The standard deviation of the dc-offset, measured over 0.01 second with a 1.9MHz filter. The one on the left is for the measurement with a starting negative offset, the one on the right is for the measurement with an offset close to zero for maximal power. Notice that there is different y-axes. The noise with PDH setpoint positioned so the NICE-OHMS offset is zero has roughly three times as much noise as the other for each intensity.

looking at the signal on an oscilloscope.

We have plotted the standard deviation of the offset signal as a function of intensity for both the measurement in Acetylene-2 where the offset was negative and the measurement where it was close to zero for maximal power(see figure 4.14). The measurements were taken with a 1.9MHz lowpass filter, to exclude high frequency noise.

4.6.5 Intensity to Offset correlation in Acetylene 2

Using the data seen in figure 4.11 and 4.13, we can estimate the possible drift of the frequency due to drift in intensity.

To get an idea of typical drift, I let the laser stand on maximal output for 75 minutes, while measuring the voltage over our intensity detctor every five minutes. The data had a standard deviation of 0.96mV. (This should be compared to a measured voltage of roughly 482mV for maximal power. So during this measurement the drift was of a size 0.2% of the total power).

By making a linear regression for both of the measurement of offset vs intensity the total offset drift is found to be 0.99mV using data from the first experiment, and 0.67mV for the second experiment.

Using the same parameters, the slope of the NICE-OHMS signal was measured to be 23.4mV/MHz

Using this, the standard deviation on the drift over five quarters of an hour we would expect would be:

Experiment 1	Experiment 2
426Hz	297Hz

Of course we can never be sure if this intensity drift was typical for the drift one would typically encounter.

Still the possible drift in intensity is much too small to account for the effects we see when our laser is locked to the resonance. Here we see drift of the order 50kHz. Based on our current data, this would require an intensity drift of around 25% of our maximal power.

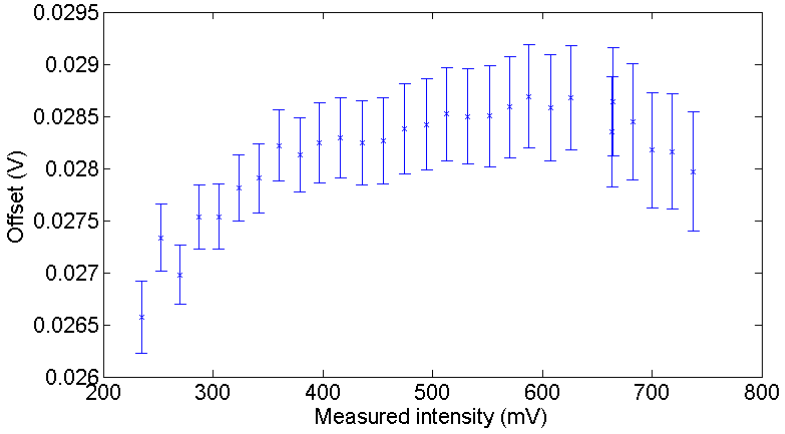

Figure 4.15: Measurement of the dc-offset on the NICE-OHMS signal as a function of intensity in Acetylene-3. The laser was changed from maximal to minimal output. At the last measurements, the error signal from RAM became too large for our PID to fully supperess it. The correlation between offset and intensity was much smaller for Acetylene-3 than for Acetylene-2.

4.6.6 Intensity to Offset Correlation in Acetylene 3

We also tried to establish the correlation between intensity and offset in the setup Acetylene 3. There, we also had an offset from a mismatch between modulation frequency and FSR, on the same order of size as the peak to peak value of the NICE-OHMS signal to the molecules.

Using the same method as for Acetylene-2 we locked the cavity in a way that gave an offset, and measured this offset while changing the output from our laser.

In this measurement we found no clear correlation between intensity and offset (see figure 4.15).

During this measurement, the RAM error signal became too big for our PID circuit to drive it to zero. The possibility of this happening is a problem that must be dealt with in the future.

4.6.7 Intensity to Frequency Drift

To test how much of our frequency drift, when the laser is locked to the resonance, can be explained by the drift in intensity, we tried locking both of the lasers at the same time, and measured the beat frequency between them, while measuring the intensity after the optical isolator, with the forementioned detector.

The drift in intensity was comparable to the one discussed earlier, in the section about our Allan deviation.

These data shows that the beat between the two lasers correlate only with the drift in intensity in Acetylene-2 (see figure 4.16).

The correlation however seems to be around one hundred times greater than would be expected from our measurements of intensity-to-offset correlation. Ac-

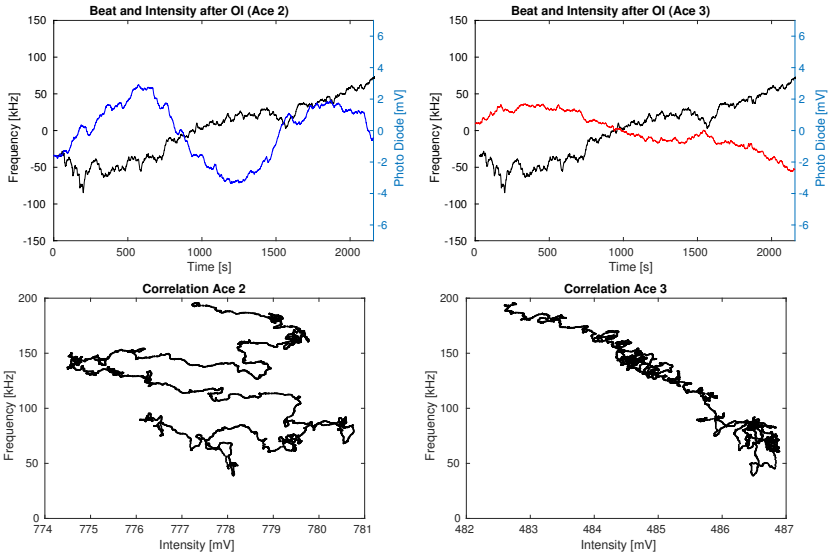

Figure 4.16: Measurements of the beat between the two setups and the intensity after the Optical Isolator, taken by Martin Romme Henriksen. HERE THE NAMES ACETYLENE-2 AND ACETYLENE-3 HAS BEEN SWITCHED because of crossed lines during data aquisition. We see that the beat frequency is correlated only with the drift in frequency in acetylene-2.

cording to that a drift in measured intensity of 4mV as we see, should only have given a change in frequency of around 1.5KHz, whereas we see a drift of 150kHz.

4.6.8 Ways of Avoiding Cavity Effects

All in all we have seen that if the modulation frequency is even slightly detuned from the FSR of the cavity, the effect will be an offset greater than our normal NICE-OHMS signal.

The best way to avoid this would be to always keep the modulation frequency exactly aligned to the FSR.

So far we have been adjusting the modulation frequency by hand, using a computer program which controls our function generator, before locking the laser to the signal.

Generally, the fine tuning has been done by scanning the laser with the cavity unlocked, and trying to get the resulting demodulated signal from the NICE-OHMS detector to be as small as possible.

There problems with this method is: Firstly, we have not been able to, using this method, to completely eliminate the signal from the cavity, we do not know why.

Secondly, even if one could tune the modulation frequency exactly to the FSR, if this was not done while the laser was on the exact right frequency, the cavity would change FSR while the laser was moved to the correct frequency.

To see this effect we must explain that usually the cavity is first locked to the laser, and the laser is then tuned to find the exact resonance. While the

lock is turned off, the laser can drift away from resonance. Usually it can drift over night something on the order of one FSR, or 650MHz.

If the cavity is locked to this frequency, and then tuned along with the laser afterwards, it will change length by one wavelength, or 1.542μm. Since the cavity is roughly 23cm long, the change in FSR due to such a change in length will be:

$$\Delta\text{FSR} = \frac{c}{2 \cdot 23\text{cm}} - \frac{c}{2 \cdot (23\text{cm} + 1.542\mu\text{m})} = 4.37\text{kHz} \tag{4.11}$$

In comparison, the FWHM of one resonance of the cavity is around 2MHz.

The transmitted light through a cavity has, just like the reflected light, a phase change which around resonance is directly proportional to the detuning. So, a detuning of the sidebands from resonance with only a few kHz can lead to an appreciable phase change.

In order to optimally tune the modulation frequency to the FSR I would propose the following procedure:

First lock the cavity to the laser.

Second, scan the laser slowly, and tune it to resonance with the molecules.

Third, put a slowly varying offset on the PDH signal, so small that the cavity can remain locked to the laser. The offset of the NICE-OHMS signal should now be seen to oscillate with the offset of the PDH signal. (The size and the noise will normally also oscillate with the setpoint of the cavity).

Fourth, fine tune the modulation frequency until the offset of the NICE-OHMS signal goes to constant zero.

Fifth, turn off the oscillation of the PDH offset, and set the offset to some chosen point.

Once all this has been done, the triplet should be exactly aligned with the cavity, except for the effect caused by the molecules. The laser can then be locked to the NICE-OHMS signal.

When the modulation frequency has been exactly matched to the FSR while the laser is on resonance with the molecules, it should not be necessary to readjust it again. However, if something causes the cavity to drift more than one FSR in length, it will be necessary either to find a way to tune the cavity back to the former length, or readjust the modulation frequency once more.

4.7 Comparison of Bent and Linear Cavities

Despite the possibility of a narrower line with a bent cavity, we chose to use linear cavitites for our final stability measurements.

The reason is that it is not only the width of the NICE-OHMS signal that is relevant, but the slope of the signal. If one were to cut both the width and the peak-to-peak height of a singal in half, the total slope would remain unchanged.

We had consistenly worse finesse in our bent cavities, and we could only couple a smaller amount of our intensity into them. This lead to a smaller overall signal, as the saturation, I/I_s, was much smaller. All in all we could not getter a better slope on the NICE-OHMS signal with them(see figure 4.17).

This might however be caused mainly by our bant cavities having consistently worse Finesse than our straight cavities.

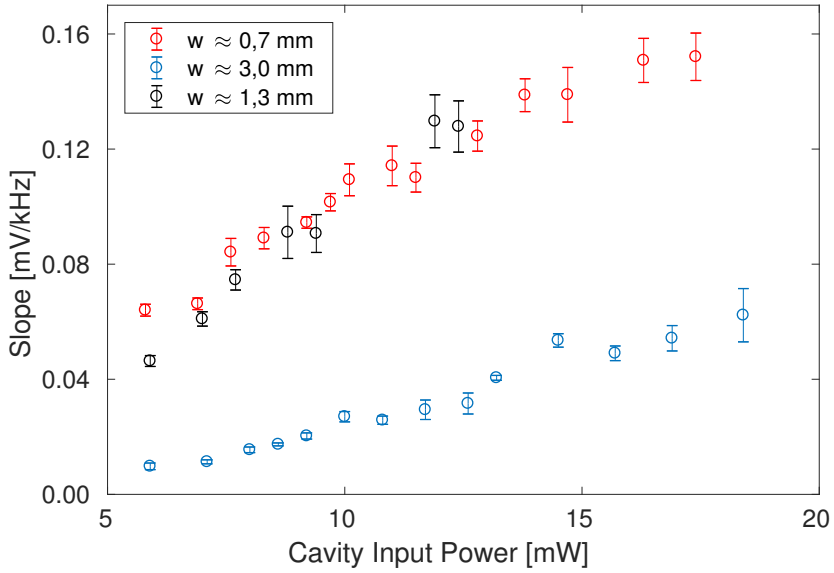

Figure 4.17: Measurements of the slope og the NICE-OHMS signal for three different cavity designs, made by Martin Romme Henriksen. The cavities are a linear with waist ca 0.7mm, identical to Acetylene-2 and 3 except for the cavity not being mounted on zerodur, a bent cavity with a waist of 3mm similar to Acetylene 1, and another design with a waist of 1.3mm, not further described in this thesis. The power on the x-axis is the power directly incident on the cavity.

The size of the NICE-OHMS signal is proportional to $2F/\pi$, where F is the Finesse. So any gains from a wider waist is zeroed out by the effects of a smaller finesse and a smaller total cavity power.

In that case there might still be something to be gained from a large waist cavity, if we could find a way to construct it with the same finesse as our straight cavities (which should be theoretically, albeit perhaps not practically, possible).

The next section will cover my calculations on possible gains from such a cavity.

4.8 Theoretical Gains from a Wider Beam

This section is quite speculative. I have used my best understanding of the problem at hand, which might be wrong. I therefore try to make sure to write down all assumptions I have used.

In order to get a proper understanding of the possible gains to be gotten from a larger waist cavity we must be able to calculate the expected signal size.

The size of the NICE-OHMS signal will be proportional to the phase shift induced by the medium.

According to [4], the Peak-to-Peak value of the phase shift from passing once through a medium of length L and absorbtion constant at low intensities α_0 is:

$$\Phi = 0.45\frac{\alpha_0 L}{2}\frac{8}{w^2}\int_0^\infty \frac{I/I_s e^{-4(r/w)^2}}{1+2I/I_s e^{-2(r/w)^2}} r dr. \tag{4.12}$$

Here w is the radius of the beam, assuming a gaussian beam. The overall result does however not depend on the beam size. This integral can be solved numerically, and the solution is plotted in figure 4.18. This has been experimentally verified.

For a fixed pressure and linewidth, the saturation intensity, I_s can be measured experimentally.

According to [7] and [8] (same authors of different articlse), the saturation intensity for experiments as ours can be calculated as:

$$I_s = \rho(\Gamma_{tt} + \Gamma_p)^2, \tag{4.13}$$

where ρ is a parameter to be measured experimentally, and Γ_{tt}, Γ_p is the broadening from a transit time and pressure, so $\Gamma_{tt} + \Gamma_p$ is the total FWHM of the line at low intensity.

We measured the width of the NICE-OHMS signal for different intensities in our cells. The data is repeated here, along with the value of ρ calculated from them

Acetylene 2	Acetylene 3
$\Gamma_0 = 565 \pm 24$kHz	$\Gamma_0 = 414 \pm 17$kHz
$I_s = 400 \pm 49\mathrm{kW}/m^2$	$I_s = 477 \pm 56\mathrm{kW}/m^2$
$\rho = 1.25 \pm 0.2\frac{\mathrm{W}}{\mathrm{m^2 kHz^2}}$	$\rho = 2.78 \pm 0.4\frac{\mathrm{W}}{\mathrm{m^2 kHz^2}}$

The great disrepancy between these two values might suggest an error in one or both of the experiments.

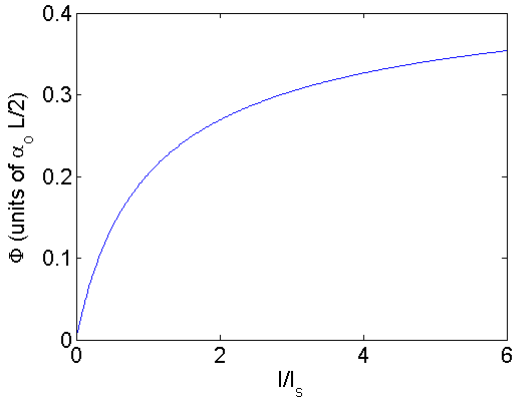

Figure 4.18: The peak-to-peak value of the dispersive function describing the phase-shift as a function of detuning, from passing through a medium, as a function of I/I_s. The Peak-to-Peak value of our NICE-OHMS signal is directly proportional to this. The value goes asymptotically towards a value of $0.45\alpha_0 L/2$ for I/I_s going to infinity.

4.8.1 Comparison of our Saturation Intensities with Others'

In [6] the saturation of our transition is measured. In the article, a beam with a waist of 0.58mm and a total power of 30mW is passed through a sample of Acetylene-13 a total of 4 times. This gives a mean power of 120 mW. If we use the definition we have used so far of calculating the intensity of the beam as:

$$I = \frac{P}{\pi w^2}, \tag{4.14}$$

(any other definition will only vary from this by some constant), the intensity of the beam would be $I = 113\text{mW/mm}^2 = 113\text{kW/m}^2$.

At very low pressure, the FWHM extrapolated to low intensity is reported to be $\Gamma_0 = 167\text{kHz} \pm 22\text{kHz}$, and the saturation is reported as $I/I_s = 0.75$.

This means a saturation intensity of $I_s = 151\text{mW/mm}^2$, corresponding to a total power of 4x40mW, as is stated in the article.

Calculating ρ from this, we find

$$\rho = \frac{I_s}{\Gamma_0^2} = 5.42\frac{\text{W}}{\text{m}^2\text{kHz}^2} \tag{4.15}$$

This is larger than our measured values. Our values however have a very large uncertainty, and my understanding of the exact setup in [6] might be flawed.

4.8.2 Slope vs Beam Waist

We use the following assumptions;

$$\Gamma = (\Gamma_{tt} + \Gamma_p)\sqrt{1 + I/I_s}, \tag{4.16}$$

$$I_s = \rho(\Gamma_{tt} + \Gamma_p)^2, \tag{4.17}$$

where we use the mean of our two values of ρ, to give us an estimate:

$$\rho = 2.01 \frac{\mathrm{W}}{\mathrm{m^2 kHz^2}}. \tag{4.18}$$

and the Peak-to-Peak value of our NICE-OHMS signal, Q being proportional to:

$$Q \propto p\Phi(I/I_s), \tag{4.19}$$

where p is the pressure and Φ can be seen in figure 4.18,
and at last that the total Slope, S is proportional to:

$$S \propto Q/\Gamma \tag{4.20}$$

The signal peak-to-peak value divided by its width.

In that case we can say that we hold all parameters execpt for waist (w) and pressure (p) constant, and calculate the slope. (So we assume we can have the same Finesse no matter the waist and the pressure).

In our cavity we could typically reach a total power of 2 W.

And investigation of the slope for this can be seen in figure 4.19.

Here it would seem that any gains from having a larger waist and lower pressure could equally be gained by having a waist slightly smaller than 1mm, and a slightly higher pressure.

Also, if we compare the area we are in now (a waist of 0.7mm and a pressure somewhere between 1 and 2 Pa), it does not seem that we could gain much more than 50% in total slope.

As noted, the precision with which we can claim to know ρ is very low. However, a change in the parameter ρ is essentially equivalent to a change in the total power. I therefore have made the same simulation for a number of different powers (see figure 4.20).

For all these different powers it still seems that the optimal slope can be found for a smaller waist at a suitable pressure.

Judging by these simulations, there does not seem to be any reason to try and make a cavity with a larger waist, unless by doing so one could get a higher finesse.

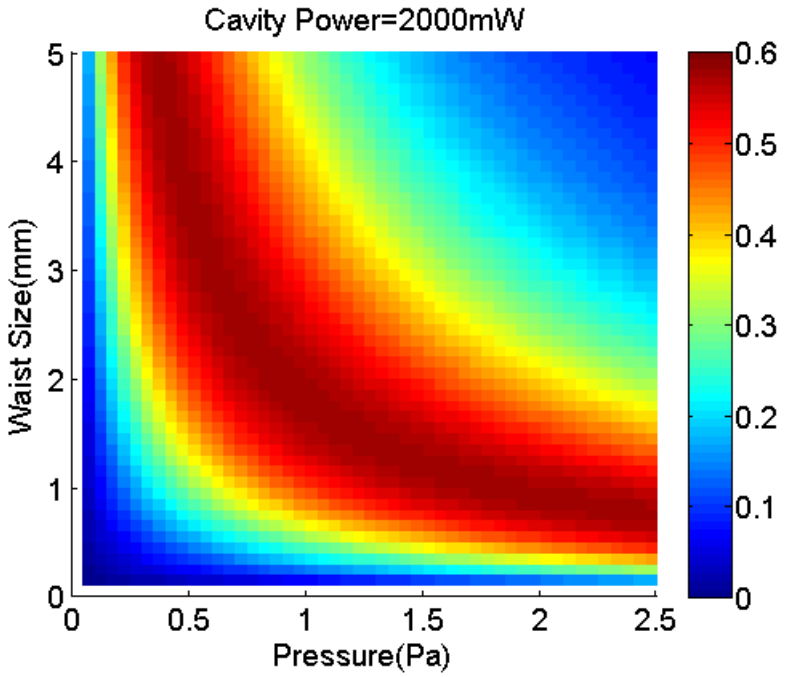

Figure 4.19: The slope, in arbitrary units, as a function of waist and pressure. Our current experiment has a waist of 0.7mm and a pressure around 1.5Pa.

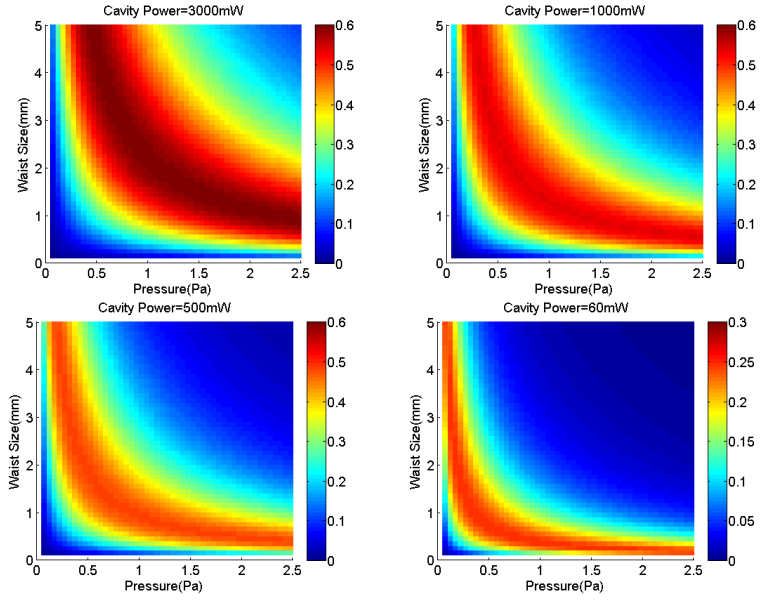

Figure 4.20: The slope for a total power of 3W, 1W, 500mW and 60mW. Notice that the color scale on the last has been changed to show the details along the others.

Chapter 5

Conclusion and Outlook

5.1 Conclusion

We have been able to frequency lock two lasers to the same transition in Acetylene-13, but the Allan Deviation obtained was only around $5 \cdot 10^{-}12$ for average times of one second, and radically worse for longer averaging times. Current references at this level can reach a ten times better deviation at one second, and down below 10^{-13} for longer averaging times.

5.2 Outlook

Our data seems to suggest that the drift in frequency we observe might be caused by intensity drift in one of our two setups.

If this is the case, then a stabilisation of the intensity might lead to a better long term stability. This is the first method we will try.

A better way of tuning our modulation frequency to the FSR of our cavity might also lead to a smaller drift.

If however, the drift we observe is mainly caused by one of our setups, then the other setup might have a much better Allan deviation than the one reported here. This hypothesis could be tested if we were to beat it with a third frequency reference, with known deviation.

We might possibly be able to get a better slope on our error signal, and so a better Allan deviation, if we could make a cavity with an optimised beam waist and gas pressure.

Furthermore, we might be able to get a system with less noise by using vacuum, albeit there still are some technical difficulties with this.

Chapter 6

Bibliography

Bibliography

[1] Fritz Riehle
Frequency Standards, Basics and Applications
WILEY-VCH verlag 2004

[2] Peter W. Milonni and Joseph H. Eberly
Laser Physics
Wiley 2010

[3] Robert W. Boyd
Nonlinear Optics (third edition)
Academic Press 2007

[4] Axner, Ehlers, Foltynowicz, Silander and Wang
NICE-OHMS Frequency Modulation Cavity-Enhanced Spectroscopy. Principles and Performance

[5] RP-photonics encyclopedia of Laser Physics
https://www.rp-photonics.com/encyclopedia.html

[6] Jan Hald, Lars Nielsen, Jan C. Petersen, Poul Varming and Jens E. Pedersen
Fiber laser optical frequency standard at 1.54 μm
Optical Society of America, 2011

[7] Long-Sheng Ma, Jun Ye, Pierre Dubé, John L. Hall
Ultrasensitive frequency modulation spectroscopy enhanced by a high finesse optical cavity: Theory and application to overtone transitions of C_2H_2 and C_2HD.
Optical Society of America, 1999

[8] Jun Ye, Long-Sheng Ma and John L. Hall
Ultrastable Optical Frequency Reference at 1.064 μm Using a C_2HD Molecular Overtone Transition
EEE TRANSACTIONS ON INSTRUMENTATION AND MEASUREMENT, VOL. 46, NO. 2, APRIL 1997

[9] Wolfgang Demtröder
Laser Spectroscopy, Basic Principles
Springer 2008

[10] Eric Black
Notes on the Pound-Drever-Hall technique
A working note from LIGO, 1998

[11] N.C. Wong and J.L. Hall
Servo control of amplitude modulation in frequency-modulation spectroscopy: Demonstration of shot-noise-limited detection
Optical Society of America 1985

[12] Daniel Adam Steck
Atom and Quantum Optics
available online at http://steck.us/teaching (revision 0.12.2, 11 April 2018).

[13] P. Balling and P. Krěn
Development of Wavelength Standard at 1542nm: Acetylene Stabilized DFB Laser
Czech Metrology Institute 2005