

MASTER

Active stabilization of a high finesse cavity for the QND experiment

Domen, K.F.E.M.

Award date:
2002

[Link to publication](#)

Disclaimer

This document contains a student thesis (bachelor's or master's), as authored by a student at Eindhoven University of Technology. Student theses are made available in the TU/e repository upon obtaining the required degree. The grade received is not published on the document as presented in the repository. The required complexity or quality of research of student theses may vary by program, and the required minimum study period may vary in duration.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

**Active stabilization of
a high finesse cavity
for the QND experiment**

K.F.E.M. Dömen
Graduation report AQT/B 02-03
January 2002

Supervision: Prof.dr. K.A.H. v. Leeuwen
Prof.dr. H.C.W. Beijerinck

Summary

One of the current projects in the Atomic Physics and Quantum Electronics group at Eindhoven University of Technology concerns a direct demonstration of the existence of the quantization of the electromagnetic field through an atomic diffraction experiment. In the proposed experiment the diffraction of a beam of helium atoms interacting with the electromagnetic field in a high-finesse optical cavity will be studied.

For this experiment a highly sophisticated setup has been built to produce an extremely precise beam of helium atoms. The setup has been constructed, tested and used for preliminary experiments by two (former) graduate students. However, after the departure of the last graduate student, the machine has been idle for more than half a year. This report discusses the efforts to 'resurrect' the machine together with a second undergraduate student. Furthermore changes in the design, that proved necessary to improve reliability and operational ease, are described, as well as tests of new elements for the setup, needed for the real experiments.

In the first chapter a review of the beam machine and changes in the mechanical design are presented. The remaining part of this report deals with the design and improvement of the stabilization configurations for the laser and the high finesse cavity. The stabilization method used is based on the Pound-Drever-Hall scheme. Some aspects of feedback control theory and its consequences for our setup are discussed, as well as the theory of Pound-Drever-Hall.

As the high finesse cavity constitutes another fundamental part of the experiment, its optical and mechanical properties and resonance conditions are investigated. Some general aspects of cavity theory are presented. An expression for the internal, reflected, and transmitted light intensity and the relative phase of these beams is deduced using a purely classical model. The effect of polarization was not included in this model. However, in practice the cavity displays birefringence. The effect of this birefringence on the polarization of the cavity mode, is discussed. Experimental results on the cavity birefringence are presented.

Finally the route towards the stabilization of the high finesse cavity is explored. Exploiting the Pound-Drever-Hall scheme, the distance between the two mirrors of the cavity should be kept stable to within 1 pm. The stringent demands on the stabilization setup are discussed. The necessary electronics have been built and tested. In Chapter 7 the test results are presented and further improvements are proposed.

Contents

Summary	i
1 Introduction	1
2 QND setup and improvements	3
2.1 The QND setup	3
2.2 Guide line	6
2.3 MOC revisited	7
3 Feedback control	11
3.1 Elementary Feedback Control	11
3.2 PID control	13
3.3 Pound-Drever-Hall: a quantitative model	16
4 Laser stabilization	21
4.1 Stabilization using saturated absorption	21
4.2 Pound Drever Hall in saturated absorption	22
4.3 Stabilization setup	23
5 Cavity theory	25
5.1 Cavity transfer function	25
5.2 Finesse	28
5.3 Phase response	30
5.4 Cavity modes and Gaussian beams	31
5.5 Mode matching	32
5.6 Polarization	32
6 Cavity diagnostics	35
6.1 Cavity design	35
6.2 Diagnostics	35
6.3 Cavity test facility	36
6.4 Beam alignment	38
6.5 Cavity stabilization	41
6.6 Electrical and mechanical behavior	42

7 Cavity stabilization 45

7.1 Stability demands 45

7.2 Stable supply 46

7.3 Total scheme for cavity stabilization 48

7.4 Pound-Drever-Hall with a high finesse cavity 50

7.5 PI controllers 51

8 Conclusions 55

A Technology Assessment 57

Chapter 1

Introduction

Quantum measurements, even in the form of textbook schemes applied to very simple systems, are generally not simple. For example, position measurements on single particles cannot be performed in the form of Quantum Non Demolition (QND) measurements, i.e., these experiments cannot be performed without disturbing the measured quantity. In measuring the position, the particle's momentum is changed and hence the position at later times. Successive position measurements can therefore not lead to the same result. Alternatively, momentum measurements at the quantum limit may be QND. The basic and most interesting feature of a QND measurement is thus that it can be repeated; successive QND measurements on the same system will give identical results. Many QND measurements have been realized in quantum optics without achieving, however, sensitivity on the single photon scale and thus without addressing the basic axiom of field quantization.

In the AQT group at Eindhoven University of Technology, a setup is currently being developed that will allow such sensitive QND experiments through which we hope to provide a direct demonstration of the existence of the quantization of the electromagnetic field. In the proposed experiment the diffraction of a beam of atoms interacting with the electromagnetic field in a high-finesse optical cavity will be studied. The meter in this experiment is a single helium atom that crosses the light field in the cavity. After the interaction the atom exits the cavity, leaving the amplitude of the light field unchanged. The helium atom, however, carries away the information about the amplitude of the light field and hence the number of photons in the cavity, imprinted on its wave function. As the setup is designed such that the diffraction is observable even when the field in the cavity consists of only a single photon, the diffraction pattern will display a clear signature of quantization.

The interaction effects for a single atom and a single photon in the cavity will be extremely small. To be able to observe these effects with the required precision, the requirements for the helium beam and the cavity are accordingly stringent. In the beam setup four laser cooling sections are exploited to prepare the necessary ultra-bright, collimated and monochromatic beam of metastable helium atoms. Special stabilization techniques are necessary also to regulate the distance between the mirrors of the optical cavity to within the required 1 pm stability.

The development of the beam setup has already taken five years. Therefore, this precious QND setup is designed such that it can and will be used for a variety of other experiments studying atom-laser interactions. Koolen [2] already performed record-breaking Bragg scatter-

ing experiments at a standing light wave in 2000. Further discussion and numerical simulations of these and future Bragg scattering experiments can be found in Jansen [7]. In 2002, a start will be made with the construction of a very large area interferometer based on the coherent splitting and recombination of the atomic beam, as well as the above-mentioned diffraction-measurement based demonstration of field quantization, leading to the QND experiment after which our project group bears its name.

In this report a review of the operation of the helium beam setup is presented. The difficulties that we experienced using the machine and the improvements implemented in the design are discussed.

The stabilization of the lasers used for the experiments, as well as of the high-finesse cavity, are discussed in the rest of the report. Both stabilizations are based on the so-called Pound-Drever-Hall scheme, which is treated in some detail. The properties of the high-finesse cavity and its stringent stability demands are considered.

Finally, measurements of the electro-mechanical behavior of the cavity and of the polarization properties of the cavity are presented as well as the design of the sub-picometer stabilization system.

Chapter 2

QND setup and improvements

The present QND setup, a machine producing an extremely collimated and ‘monochromatic’ (single axial velocity) beam of metastable helium atoms, has been designed such that it is suitable for a plethora of different experiments investigating light-matter interactions. Yet, the initial motive to start the project was to provide a direct demonstration of the existence of the quantization of the electromagnetic field through an atomic diffraction experiment. In the proposed experiment the diffraction of a beam of atoms interacting with the electromagnetic field in a high-finesse optical cavity will be studied. The diffraction should be observable even when the field in the cavity consists of only one photon. In this case, the diffraction pattern displays a clear signature of quantization.

The choice for helium in our setup is mainly motivated by the fact that the interactions considered are rather small. In order to observe the diffraction effects on a 2D detector placed behind the interaction region, a light atom is preferred. For an atom such as helium relatively large deflection angles will be induced even by a single photon recoil. Using an atomic beam with a small axial velocity further increases the deflection angle.

Another necessary condition is that the initial angular spread of the atom beam is much smaller than the detection angle impact resulting from the absorption or emission of a single photon. Even for a light and slow atom, this implies that an extremely well-collimated beam is required. The application of laser cooling techniques is indispensable to obtain such a high quality slow atomic beam. Helium is both light and eminently suitable for laser cooling. It has therefore been selected for our QND setup.

2.1 The QND setup

At one end of the setup helium gas expands through a nozzle into the vacuum system. At the other end, nearly ten meters downstream from the source, we wish to have a high-quality helium beam that meets our stringent demands. To achieve the required beam collimation before the interaction region, the beam passes through two apertures of $60\text{ }\mu\text{m}$ and $25\text{ }\mu\text{m}$ in diameter respectively, positioned at 6 and 8 meter from the source. Between the source and the first aperture, the initially fast ($\approx 1300\text{ ms}^{-1}$) helium atoms are decelerated by a counterpropagating laser beam to a uniform and low velocity of 250 ms^{-1} . Running numerical calculations Knops [1] found that the resulting count rate in such a simple setup would however be unacceptably

low for the experiments. In order to increase the count rate extensive manipulation of the atomic beam has to be performed before the first aperture. This requires several laser cooling techniques.




Figure 2.1: Partial energy level scheme for helium

Figure 2.1 shows a partial energy level scheme for helium. Both the laser cooling sections and the experiment use the $2^3S_1 \rightarrow 2^3P_2$ transition at the infrared wavelength of $\lambda = 1083$ nm. It is a closed transition with a radiative lifetime of 98.8 ns for the upper state. The lower state 2^3S_1 has a lifetime of more than two hours and is therefore referred to as a metastable state. Absorption of a single photon at 1083 nm by a metastable helium atom He^* results in a velocity change $v_R = \hbar k / M = 0.092 \text{ ms}^{-1}$ (known as the recoil velocity), with k the wave vector of the light and M the mass of the helium atom.

In Figure 2.2 a schematic view of the current QND beam setup is presented. References [1,2] give a more detailed description, we confine ourselves to a short overview.

A gas inlet system is used to inject helium gas into the source. The gas expands supersonically through a boron-nitride nozzle with a 0.1 mm diameter, which is in contact with a liquid nitrogen reservoir. This cools the effective source temperature down to 200 K and reduces the average (longitudinal) velocity of the atoms down to $\langle v \rangle = 1300 \text{ ms}^{-1}$. Metastable helium atoms are formed in a DC gas discharge drawn from the nozzle to a skimmer that is located 10 mm downstream and has a 1 mm diameter aperture. Typical operating conditions of the source are 5-6 mA discharge current at a voltage of 760-800 V.

The metastable atoms passing through the skimmer still form a unacceptably divergent gas flow. To prevent the atoms from deviating too much from the central beam line after the first few centimeters, a first laser cooling stage is positioned directly after the skimmer. This so-called collimator consists of a two-dimensional standing light wave with slightly red-detuned light (2-3 MHz). In the collimator the atoms are redirected into a parallel beam, thus securing sufficient input flux for further laser cooling stages. The collimator is driven by a temperature and frequency stabilized Distributed Bragg Reflection (DBR) diode laser.

Next, a two meter long Zeeman slower reduces the broad thermal velocity distribution to a




Figure 2.2: Schematic overview of the QND beam setup

narrow velocity distribution of around 250 ms^{-1} . In the Zeeman slower the metastable helium atoms are decelerated by the radiation pressure of a counterpropagating laser beam. The changing Doppler shift upon deceleration is compensated for by the Zeeman shift in an inhomogeneous axial magnetic field. Unfortunately, during slowing the axial divergence and the diameter of the beam increase rapidly towards the end. To compensate for this effect a final two-stage beam compression section is added. The first stage, a magneto-optic lens (MOL), functions as a thin lens for the broad (approx. 25 mm diameter) atomic beam with a focal length of 60 cm. The lens consists of a two-dimensional standing light wave, as for the collimator, combined with a two-dimensional magnetic quadrupole field with the zero-field axis coinciding with the center line of the atomic beam. At this point we possess a slowed and narrow beam which however is still highly divergent. A magneto-optic compressor, or MOC, positioned in the MOL's focal point funnels the atoms into a narrow ($250 \mu\text{m}$ diameter) and collimated (angular divergence 10 mrad) beam. The compressor is essentially identical to the lens, however a stronger quadrupole field results in 'funnelled' atom trajectories [1,2]. This beam provides sufficient brightness (flux $5.0 \times 10^9 \text{ s}^{-1}$ through a 1 mm orifice) for the final mechanical selection by the pinholes.

In order to check the performance of each laser cooling section in the machine and to make sure that it operates at maximum performance, three diagnostic devices are present. First, the global performance of the source and collimator is measured with a wire scanner positioned two meter behind the collimator, half-way the Zeeman slower. The performance of the MOL is measured with a second wire scanner positioned 0.50 m behind the end of the MOL, i.e., near the entrance of the MOC. Both wire scanners consist of two crossed stainless steel wires of 0.1 mm diameter and are mounted on a computer-controlled translator. When a highly energetic metastable helium atom impinges on the surface of a steel wire, an Auger electron will be emitted from the metal. Measuring the current from the wire gives the number of impacting metastable atoms per second. The performance of the two stage beam compressor, finally, can be measured using a double knife-edge scanner that is situated 100 mm behind the end of MOC. The knife-edge scanners are made out of glass with a conducting coating. The current on the scanner's blade released by the impact of He^* atoms is a direct measure for the total flux. In practice, it is virtually always this knife-edge scanner that is used to determine the overall performance of the four laser cooling stages. The wire scanners are used when more detailed profiles of the atomic beam are needed.

The experiments performed thus far with the QND machine took place in the interaction chamber just behind the $25 \mu\text{m}$ pinhole, eight meter downstream from the source. As the experi-

ments aim to measure the transverse velocity changes of the atoms, we need a two-dimensional position sensitive detector. This detector consists in essence of a large array (25 mm diameter) of microscopic electron-multipliers (15 μm diameter) forming a microchannel plate (MCP) over which typically a voltage of 800 volts is applied. Upon impact of a metastable helium atom a surface electron will be released and accelerated in the electric field, releasing more electrons on each subsequent collision with the channel walls. Three Z-stacked layers of these MCP's increase the signal gain. At the back side of the microchannel plates roughly 10^7 electrons are released for each initial electron. This electron cloud then hits a square surface resistor with charge sensitive preamplifiers attached to each of the four corners. The amount of charge that flows to each corner contains information on the position of the electron cloud and thus on the position where the metastable atom impacted on the detector. The charge signals are processed electronically and stored for later retrieval together with timing information with respect to a trigger pulse. This timing information can be used to perform time-of-flight measurements as an extra diagnostic tool.

2.2 Guide line

The machine described above has been constructed and tested by the former graduate students Knops [1] and Koolen [2]. They realized the envisioned highly collimated beam and were able to determine its main characteristics with several diagnostic tools. Koolen then was the first to perform atomic Bragg scattering experiments with the beam. It was concluded that the beam seemed to meet all of our requirements. The main properties of the beam are summarized in table (2.1).

beam diameter	25 μm .
velocity	250 ms^{-1}
flux	250 s^{-1}
collimation	$\sigma v_{\perp} = 9 \times 10^{-3} \text{ ms}^{-1} < 0.1 v_R$
monochromaticity	$\sigma v_{\parallel} / v_{\parallel} = 1.5 \times 10^{-2}$

Table 2.1: Properties of the helium beam as determined by Koolen [2]

The QND setup for preparing a cold beam has thus been proven effective and ready to use. However, operating this highly complex and sensitive setup at its maximum capabilities requires not only thorough knowledge and skills but also the right *Fingerspitzengefühl*, qualities that were for this machine in the end mainly concentrated in one person. Considering the long term continuity of the project the need rose to at least preserve the present knowledge in our group. A 'reference manual' has been written by Koolen, offering a step-by-step guidance through the startup procedures that need to be followed in order to produce the desired beam. It also contains setup details and target figures for diagnostic comparison. Unfortunately it is not possible to record information down to the smallest details on mirror positions and magnet configurations nor to pass over the practical experience, let alone to put the *Fingerspitzengefühl* adequately on paper.

January 2001, after a six months' period of hibernation due to the lack of a graduate student in charge of the setup, the machine was breathed new life into by two undergraduate students

(myself and Maarten Jansen) in their graduation project. Thanks to the well-considered design and good maintenance most of the essential elements in the setup appeared to be still functioning. After a period in which we had to acquire the right touch to manipulate every separate optical device and had to replace some parts that were broken in the meantime, the setup was ready for fine tuning. Using the available methods for diagnostics and following the target figures summarized in the reference manual we tuned the different laser cooling sections. In this way we gradually realized values for the flux of cold atoms on the knife-edge detector that closely approached the target figures. In other words, at the last diagnostics before the 2D detector we were capable of generating an atom beam with the desired qualities. However, bridging the last obstacle towards the detection of a highly collimated cold beam on the 2D detector was hindered for two reasons: lack of reproducibility and difficulty of adjustment. Both problems were traced to the two-stage beam compressor.

2.3 MOC revisited

After a necessary overhaul of our beam machine, during which a high voltage source, a vacuum pump and some electronics were replaced, it turned out to be very difficult to re-attain the earlier measured wire-scan profiles and atom flux values on the knife-edge scanner downstream of the MOC. The setup's adjustment clearly suffered from irreproducibility. The question rose which part of the setup wasn't working properly. Since the two-stage beam compressor is far more difficult to align than the collimator and Zeeman slower, the first was suspected to cause most of the problems. From a series of time-of-flight measurements currently done, it seems indeed that both the collimator and Zeeman section are functioning properly, although perhaps not yet optimally. The MOL and MOC systems thus require more attention.

The operation of the MOC (and the MOL) is based on a force that depends on both the transverse position and the transverse velocity. The position dependency is a result of the magnetic quadrupole field which is zero on the beam axis and increases linearly with the distance from the axis. Unless an atom moves perfectly along the axis $B = 0$ with zero radial velocity $v_r = 0$, it experiences a force towards the center line. The MOL now acts as a thin lens, focussing the beam in the capture region of the MOC. In the MOC the atoms are funnelled towards the center line. The transverse motion is similar to that of an overdamped harmonic oscillator.

In practice, however, the MOC may *not* be perfectly aligned, as was assumed in the aforementioned description. One of the possible errors to occur is a small error in the alignment of the angle of incidence of the incoming light field. However, the light field of the MOC extends over a length of more than 100 mm, so that the retroflected beam will still coincide with the incoming beam. The resultant light field therefore is not influenced much, except that we have now introduced a very small component of the force on the atom along the direction of the center line. This will not noticeably influence the atom's transverse motion.

Secondly, the positioning of the quadrupole field can be slightly off-center or tilted over a small angle. This has important consequences for the atom's trajectory. As we saw, the atoms are guided towards the centerline, which is defined by the $B = 0$ axis of the quadrupole field. In other words, by adjusting the magnets of the MOC we should aim the atom beam through the pinholes.

In our current setup the four quadrupole magnets are firmly attached to a ring-shaped plexiglass mount which in turn is suspended to a ring-shaped metal mount. Both plexiglass and metal mounts have the same outer diameter of 40 cm and an inner diameter that is about 20 mm wider than the vacuum tube around which they are attached. Four screws, inwardly pierced through the metal ring, clamp the mount to a vacuum flange. When tightened, the whole construction thus rests only at the four screw tips. By turning these screws the whole suspension and in consequence the axis of the quadrupole field can be shifted in a plane perpendicular to the centerline. Symmetrically around the rings and along the vacuum tube's axis four more screws are pierced through both mounts, such that with extra bolts the distance between the two rings can be fixed. In other words, when adjusting the screw at each 'corner' the plexiglass disc can be tilted over a small angle relative to the metal mount. On paper, we thus have at our disposal a construction that allows us to tune the required four degrees of freedom. In practice, however, this is not sufficient to be able to align the helium beam.

The collimated and slowed atoms that leave the MOC still have another four meters to go before they can be detected on the 2D detector. In this part the 60 μm and 25 μm apertures pick out those atoms that fulfill the extreme restrictions on the allowed transverse momentum spread. Apart from measuring the total flux we lack possibilities to monitor the beam's position after the last laser cooling stage. In other words, we don't possess any information on the beam's position and angle with respect to neither the vacuum pipeline's center-line nor to the positioning of the moveable pinholes: we are literally groping around in the dark. Therefore we are looking at the feasibility of adding another diagnostic tool just in front of the second pinhole in the form of a second position sensitive detector. This detector will either have to have a small center hole or have to be moveable, in order to allow the beam to reach the second pinhole. However, it needn't be extremely sensitive nor have extreme resolution (and be accordingly expensive) compared to the 2D detector. At this point in the machine the helium flux will be relatively large and we are not interested in detecting one photon recoil effects. At present a final solution has not yet been chosen.

In the absence of this extra diagnostic tool, one could of course hope to find the optimal configuration by systematically adjusting the light field and the magnetic field in both MOL and MOC and monitoring the signal from the knife-edge scanner. Unfortunately the design of the suspension of the four permanent magnets that produce the quadrupole field as described above, doesn't offer flexible and reproducible adjustments. The main reason is that none of the components of the MOL and the MOC have scaled adjusting knobs that could provide any coordinates to hold on to. Furthermore the fragile suspension of the heavy permanent magnets at four screw tips doesn't allow displacements in a controlled manner. The solution to this problem was to design a more stabile and controllable way of suspending the magnets. The diameter of the beam exiting the MOC is estimated to be 200 - 280 μm . The divergence is on the order of 10 mrad rms, which translates into a velocity spread of 2.5 ms^{-1} rms. These measurements serve to estimate the required precision and range of the new suspension design. Logically the displacement should be adjustable to within a fraction of the beam's diameter, say 50 μm over a range of at least as large as the adjustment range of the position of the first pinhole (5 mm). For the tilting angle we set the tolerance to 1 mrad. These demands are not extremely hard to meet. Further restrictions on the construction were imposed by limitations in available space on the optical table. The design is shown in Fig. 2.3 and has been recently implemented.




Figure 2.3: Front and rear view of the new quadrupole magnets suspension design. A square aluminum block (grey) is attached to the vacuum flange (not shown) and forms the rigid mount for the MOC suspension. The four quadrupole magnets (light grey) are rigidly attached to a second aluminum plate (black). Through a suspension system of stiff springs this second plate is firmly but not immovably attached to the suspension mount. Micrometer screws allow the controlled transverse displacement of the four magnets with respect to the beam line, with micrometer precision over a 10 mm range. The three micrometer screws in triangle configuration in the front, allow for a stable adjustment of the magnets' tilting angle with 0.01 mrad precision.

Chapter 3

Feedback control

In the framework of the QND project several experiments are proposed that study the interaction between excited atoms and resonant light fields. This interaction is studied by looking at the diffraction of a narrow, collimated beam of helium atoms by a transverse field. The major part of the ten meter long QND setup is occupied by a series of laser cooling stages that are used to create the necessary ultra-high precision metastable helium beam. The laser power needed for the laser cooling stages and the experiments is provided by four diode lasers. For the experiments it is crucial that these lasers can be accurately locked to the helium $2^3S_1 \rightarrow 2^3P_2$ transition. The cavity that is used in the QND experiment will in turn be resonantly locked to the laser frequency. Since the stability demands for all four lasers and especially for the cavity are extremely stringent, as we will see later, we can not passively rely on their inherent behavior. An active stabilization scheme needs to be designed to control the output of the lasers and the cavity in order to reduce their drift and fluctuations. A feedback controlled system offers the solution. In this chapter we will discuss some basic properties of feedback systems and the scheme used to obtain the error signal in our implementation.

3.1 Elementary Feedback Control

Suppose there is a process (or system) H we wish to control. Call the input to the process u and the output from the process y . In many cases there is one particular input, \bar{u} , that would result in the desired output value known as the ‘setpoint’, designated \bar{y} . Should we now have exact knowledge of the behavior of H , i.e., we know the output y for any given input value, we could simply supply the system with input \bar{u} . This way of controlling the process is often called ‘open-loop control’. However, in reality we never possess a complete description of the process, especially when there are external disturbances acting on the system. In such cases differences between the actual and desired output are likely to occur. This difference, $\bar{y} - y$, is often referred to as the system error ϵ . The performance of the process would improve if we could control it based on observations of this system error. In feedback (or closed-loop) controlled systems this error, suitably amplified, is then used as their input to the process. If the system is functioning well, the system error should be forced close to zero by the feedback.

The block diagram in Figure 3.1 shows a basic feedback control system, with blocks representing the dynamic elements of the system and arrows representing the flow of information,




Figure 3.1: Feedback control systems determine their output by observing the error ϵ between the setpoint and the actual process output

generally in the form of an electrical signal.

Limiting ourselves to linear control theory, we suppose that the operation of the process can be described with the simple algebraic relation $y = H u$ and that for the amplifier a similar representation $u = G \epsilon$ holds, where G is the open loop gain of the amplifier. In general the feedback network then subtracts a fraction of the output signal from the input. For simplicity we choose to set this fraction equal to one. It then follows that

$$y = H G \epsilon = H G (\bar{y} - y) \quad (3.1)$$

or, solved for the output y ,

$$y = \frac{H G}{1 + H G} \bar{y}. \quad (3.2)$$

The closed-loop gain, A , is just

$$A = \frac{y}{\bar{y}} = \frac{H G}{1 + H G}. \quad (3.3)$$

From this result it is clear that the output y can never be equal to the desired output \bar{y} , but that it can approach the desired value as closely as we wish by increasing the amplifier's (or total closed-loop) gain.

Such feedback systems have several important effects such as changed input and output impedance. To us however, with the eye on process stabilization, the most significant are predictability of gain and reduction of distortion. It is easy to show, by taking the partial derivative of A with respect to H (i.e., $\partial A / \partial H$), that relative variations in the system's open-loop gain $H G$ are reduced by a factor $1 + H G$ which is sometimes called the return difference or desensitivity:

$$\frac{\partial A}{A} = \frac{1}{1 + H G} \frac{\partial H}{H} \quad (3.4)$$

Thus, for good performance it is essential to have the loop gain $G H$ larger than 1. In other words: the open-loop gain should be much larger than the closed-loop gain. An important consequence of this is that nonlinearities in the process H , which can be described as gain variations that depend on the signal level, are reduced in exactly the same way.

3.2 PID control

In the system discussed so far the amplifier G was nothing more than a simple proportional controller (or P-controller) in a linear system. The proportional behavior means that the controller's output is based on the current value of the error. As its gain is increased the system responds faster to changes in the setpoint or fluctuations in the output signal. Moreover, according to Equation 3.2, we should generate the largest possible gain for the best result. However, the same formula also implies that although the actual output y can approach the desired setpoint \bar{y} as closely as we wish, it can never really become exactly equal to this value. A small offset cannot be avoided.

In general however, the black box that we called an amplifier with gain G , is usually more than just an amplifier. It also changes the shape of the signal that passes through it. Inserting an integrating circuit, we create a proportional + integral or PI controller. The effect of the integral term is to adjust the output signal until the time-averaged value of the error is zero. The integral part thus eliminates the offset of the P-controller as can be seen as follows.

Still assuming a linear system, the output y of a PI controller and its input ϵ are related by the equation:

$$y = G_p \epsilon + G_I \int_0^t \epsilon d\tau \quad (3.5)$$

where G_p is the proportional tuning constant and G_I the integral gain parameter also known as the controller 'reset level'. Knowing that the error signal ϵ is equal to the difference $\bar{y} - y$, we find:

$$y = G_p (\bar{y} - y) + G_I \int_0^t (\bar{y} - y) d\tau \quad (3.6)$$

Often the setpoint \bar{y} has a constant value. Differentiating y with respect to time t we find the time dependence of the output:

$$\frac{dy}{dt} = G_p \left(-\frac{dy}{dt} \right) + G_I (\bar{y} - y) \quad (3.7)$$

Our interest goes to the stationary solution of this equation for which $dy/dt = 0$. That leads to:

$$0 = G_I (\bar{y} - y) \quad \Rightarrow \quad y = \bar{y} \quad (3.8)$$

Under the condition that we have a constant setpoint \bar{y} an integral controller is clearly capable of eliminating the offset that occurs when a P-controller is used. Yet, the P-controller remains a very useful controller in those situations where it is necessary to regularly adjust or modulate the setpoint value. In such cases the behavior of the integral controller may simply be too slow for the process. The faster linear gain of a proportional controller than reduces the setpoint's dynamic deviations. Combining the advantages of both the proportional and integral controller into a PI controller provides us with a powerful tool that is indeed extremely common in feedback systems.

Unfortunately feedback is not always as simple as the above analysis suggests. The problem is that because of the process dynamics, the operation of the process under control can usually not simply be represented by a linear relation. A notorious problem is the system that exhibits a time lag. Such a system will, with increasing feedback gain, gradually become undamped and eventually unstable. The stability and overshoot problems that arise when the P-controller is used at a high gain can be warded off by adding a damping term that is proportional to the derivative of the error signal. This is what we call a differential or D-controller. The combination of a proportional + integral + differential (PID) controller is probably the most commonly used standard form of dynamic compensation in practice. By properly adjusting the damping parameters of the system a critically damped response to changes can be achieved. Sometimes this can however be extremely difficult: too little damping clearly doesn't settle the overshoot, but too much causes an unnecessarily slow response. Particularly in cases when the process is susceptible to noise in the form of very sudden jumps or spikes, derivative action can cause the system to fluctuate wildly. Therefore, in those systems where accurately stabilized output signals are required, as is ours, it is often best to use a P-plus-I controller. PD control then, finds application in robot control where dynamic behavior and rapid responses have highest priority.

In the description sofar we have only discussed feedback control in the time domain. The frequency dependence of the gain and phase shifts that may be imposed on the signal by the controller haven't been taken into account. This feature however is of vital importance. For example, if high-frequency noise is superimposed on the input, how can the system be designed to respond well to the input but still filter out the noise? Phase shift is important when it comes to stability considerations. These problems are tackled much easier when described in the frequency domain.

The amplitude of the error signal ϵ will in general not have a constant value but may fluctuate in time: $\epsilon(t)$. Feedback control is then described more conveniently by Fourier transforming our signals to the frequency domain. Our fluctuating error signal will now consist of a infinite spectrum of frequencies with matching frequency dependent amplitudes:

$$\tilde{\epsilon}(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \epsilon(t) e^{-i\omega t} dt \quad (3.9)$$

For sake of simplicity for the moment only one component of specific frequency ω of our signal is considered. The error signal can then be represented as (omitting the tilde):

$$\epsilon(\omega) = \epsilon_{\omega} e^{i\omega t} \quad (3.10)$$

where ϵ_{ω} is the Fourier amplitude at this frequency component.

If the controller G in Figure 3.1 is a proportional plus integral controller, the relation between its output u and input ϵ now become in the frequency domain (analogous to Equation (3.6):

$$\begin{aligned} y(\omega) &= G_p \epsilon(\omega) + G_I \int_0^t \epsilon(\omega) d\tau \\ &= G_p \epsilon_{\omega} e^{i\omega t} + \frac{G_I \epsilon_{\omega}}{i\omega} e^{i\omega t} \end{aligned} \quad (3.11)$$

where G_p still is the proportional gain parameter and G_I the integral gain parameter. From this formula the so-called transfer function for the PI controller, that expresses the relationship between the system output and input in the frequency domain is found:

$$G_{PI}(\omega) = G_p + \frac{G_I}{i\omega} \quad (3.12)$$

The controller's main characteristics, its open-loop gain and phase response, are best depicted in a Bode plot (a log-log plot of gain and phase versus frequency) as in Figure 3.2. For high frequencies the gain approaches asymptotically to the value G_p whereas the gain increases proportional to $1/\omega$ for frequencies lower than G_I/G_p . The gain actually reaches infinitely high values for zero frequency. This is perfectly in agreement with the earlier found property of an integrator that it eliminates any offset in the output signal (see equation (3.8)).




Figure 3.2: Bode plots for a PI controller. Left the gain versus the frequency is sketched, right the phase versus the frequency

Having derived the transfer function G_{PI} for the PI controller, we will next look at the performance of a closed loop system. Let's therefore assume that the process under control, H , is an amplifier. The gain of this amplifier is for this example fixed to unity and is constant over a wide range of frequencies. The transfer function G_{ampl} for such amplifier is then simply equal to unity. Inserting G_{PI} and G_{ampl} in Equation 3.3, the transfer function for the system as a whole (unity amplifier plus PI controlled feedback) then becomes:

$$G_{system} = \frac{G_{PI}}{1 + G_{PI}} \quad (3.13)$$

For low frequencies we now see that the closed-loop gain, or the ratio of output signal to input setpoint, is equal to one (Figure 3.3). This simply means that for these frequencies the output exactly follows the input. For higher frequencies this behavior clearly breaks down as the gain of the controller drops.

The importance of this result can perhaps be understood even more clearly from Figure 3.3 where the desensitivity of the feedback loop versus frequency is depicted. The desensitivity is a measure for the fraction of the fluctuations in the amplifier's (open loop) gain that remain after the feedback loop has been closed. Should we thus wish to eliminate the noise of our

amplifier, we have to make the desensitivity as low as possible. Unfortunately, Figure 3.3 at the same time indicates that we can not completely suppress fluctuations at any finite frequency. Always some fraction survives. This fraction increases rapidly with increasing frequency, until it becomes nearly one (depending on the value for G_p) which means that the noise propagates freely through the circuit. In other words, the *bandwidth* of feedback is limited. For every system one thus has to assess over what range and by what factor noise should be suppressed.




Figure 3.3: Absolute value of the closed-loop gain (left) and the desensitivity (right) versus frequency

As has become clear from the above discussion, controllers will play an essential role in the realization of high quality stabilization circuits for the laser and the cavity. In the next chapters we will discuss the criteria for our feedback systems in more detail and assign real values to the controller's parameters G_I and G_p .

3.3 Pound-Drever-Hall: a quantitative model

Feedback control thus (at least on paper) promises to stabilize our lasers and cavity to the right frequency by nullifying the system error signal. But before we can start to think about the exact design of our controllers, another part of the feedback loop demands our attention: the error signal ϵ itself. Although quickly defined as the difference between the desired output \bar{y} and the realized output y , it is not always that trivial to measure $\bar{y} - y$ and thus to provide a decent error signal. Often, the 'error detection' scheme measures a symmetric function of $\bar{y} - y$. In this case, the controller has no way of deciding whether to adjust the output one or the other in order to reduce the error. In our case, the output variable of the system is the frequency of the laser or the resonance frequency of a cavity mode (adjusted through the distance between the cavity mirrors). The desired error signal in both cases is a frequency deviation that, however, cannot be measured directly. Absorption of the light by the atoms in a gas cell or transmission through the cavity are easily measured and depend strongly on the frequency deviations, but are symmetric functions of ϵ . The scheme of Pound-Drever-Hall provides us with a beautiful scheme that breaks this symmetry. It provides an antisymmetric function of $\bar{y} - y$ that, for a certain range around zero, is indeed proportional to ϵ . The key to Pound-Drever-Hall (in this

discussion abbreviated to PDH) is that it looks at the signal's phase instead of its amplitude. In order to do so, we must modulate the system's signal, which in our case is a laser beam.

Generally experimentalists modulate the phase of laser light by using, e.g., an electro-optic modulator (EOM). However, we will modulate the instantaneous frequency of the light directly in the source by adding a AC signal to the laser driver's DC current. The results remain essentially the same for both approaches as is easily seen.




Figure 3.4: Schematic representation of the signals in the Pound-Drever-Hall scheme: (a) shows the sidebands in the frequency spectrum, (b) shows the shape of the error signal

Consider the field of a simple laser beam of carrier frequency ω_c that can be represented by $E(t) = 1/2\tilde{E}(t) + c.c.$ with

$$\tilde{E}(t) = E_0 e^{i\phi_c}, \quad (3.14)$$

where E_0 is the electric field amplitude and ϕ_c the phase.

When modulating the laser's phase with a sinusoidal signal the total phase of the laser output is given by:

$$\phi(t) = \omega_c t + \beta \sin \omega_m t, \quad (3.15)$$

where β is the modulation depth or peak phase deviation in radians, and ω_m the modulation frequency. Then the instantaneous frequency of this signal can be represented, using the usual definition, by:

$$\omega(t) \equiv \frac{d\phi(t)}{dt} = \omega_c + \beta \omega_m \cos \omega_m t. \quad (3.16)$$

This is equivalent to saying that the frequency of the laser output signal oscillates around ω_c at the modulation frequency with an amplitude $\beta \omega_m$.

We now assume that we are phase modulating the laser source with a frequency ω_m . The electric field $E(t)$ of the laser beam can then be described by $E(t) = 1/2\tilde{E}(t) + c.c.$ with

$$\tilde{E}(t) = E_0 e^{i\phi(t)} = E_0 e^{i(\omega_c t + \beta \sin \omega_m t)}. \quad (3.17)$$

Using one of the many standard Bessel function identities:

$$e^{i\beta \sin \omega_m t} = \sum_{n=-\infty}^{\infty} J_n(\beta) e^{in\omega_m t} \quad (3.18)$$

we can approximate the electric field by:

$$\begin{aligned} \tilde{E}(t) &= E_0 e^{i(\omega_c t + \beta \sin(\omega_m t))} \\ &\approx E_0 [J_0(\beta) + 2i J_1(\beta) \sin(\omega_m t)] e^{i\omega_c t} \\ &= E_0 [-J_1(\beta) e^{i(\omega_c - \omega_m)t} + J_0(\beta) e^{i\omega_c t} + J_1(\beta) e^{i(\omega_c + \omega_m)t}] \end{aligned} \quad (3.19)$$

where E_0 is the electric field amplitude of the original beam and $J_n(\beta)$ the n^{th} order Bessel function. The modulation depth or amplitude β is assumed to be small ($\beta < 1$) so that the approximation is justified. The obtained field is a FM optical spectrum with sidebands in the frequency domain at $\omega_c \pm \omega_m$.

Let P_0 now be the total power of the beam. The power in the carrier itself, P_c , and in each side band, P_s , are then given by:

$$P_c = J_0^2(\beta) P_0 \quad (3.20)$$

$$P_s = J_1^2(\beta) P_0 \quad (3.21)$$

Since we assumed that the modulation amplitude β is small, and that higher frequency sidebands can be neglected, we can also conclude that the total power is confined in the carrier and the first sidebands:

$$P_c + 2P_s \approx P_0. \quad (3.22)$$

The modulated beam is sent into the optical active element, which can, e.g., be a cavity or a gas cell as we will later see. With a photodetector the outgoing signal is detected.

We now make a plausible assumption for the properties of the inserted element: it can alter the amplitude of the electric field, introduce a certain phase shift, but does not influence the wavelength. Both phase shift and amplitude change may be functions of the optical frequency. This means that we can find a relation between incoming and outgoing signal of the form $E_{out} = F E_{in}$, or in other words: we represent the action of the optical element in the frequency domain by a complex transfer function F that is frequency dependent.

To calculate the output beam when there is not only the carrier beam but also several sidebands impinging on the optical active element, we have to multiply each of the incident beams by the transfer function $F(\omega)$ at the appropriate frequency. In the approximation of only two sidebands this becomes:

$$\begin{aligned} E_{out} &= E_0 [-F(\omega_c - \omega_m) J_1(\beta) e^{i(\omega_c - \omega_m)t} + \\ &\quad + F(\omega_c) J_0(\beta) e^{i\omega_c t} + F(\omega_c + \omega_m) J_1(\beta) e^{i(\omega_c + \omega_m)t}] \end{aligned} \quad (3.23)$$

With our photodetector we now measure the power of this outgoing beam, which is proportional to the squared amplitude $P_{out} = E_{out} E_{out}^*$:

$$\begin{aligned}
P_{out} = & P_c |F(\omega_c)|^2 + P_s [|F(\omega_c + \omega_m)|^2 + |F(\omega_c - \omega_m)|^2] \\
& + 2\sqrt{P_c P_s} [Re\{F(\omega_c)F^*(\omega_c + \omega_m) - F^*(\omega_c)F(\omega_c - \omega_m)\} \cos \omega_m t \\
& + Im\{F(\omega_c)F^*(\omega_c + \omega_m) - F^*(\omega_c)F(\omega_c - \omega_m)\} \sin \omega_m t] \\
& + (2\omega_m \text{ terms}).
\end{aligned} \tag{3.24}$$

This signal is subsequently mixed with the original modulation frequency ω_m . By properly adjusting the demodulation, we can specifically extract the term proportional to either $\cos \omega_m$ or $\sin \omega_m$ which then serves as our operational error signal ϵ' :

$$\epsilon' = 2\sqrt{P_c P_s} \frac{Re}{Im} [F(\omega_c)F^*(\omega_c + \omega_m) - F^*(\omega_c)F(\omega_c - \omega_m)] \tag{3.25}$$

Both a dispersive phase behavior of $F(\omega)$ around the setpoint-frequency and an asymmetry in sideband absorption when ω_c varies across this frequency, can lead to the desired antisymmetric character of ϵ' . Then, ϵ' can be used effectively as error signal for the feedback circuit.

Chapter 4

Laser stabilization

The laser power needed for the laser cooling stages is provided by three 1083 nm Distributed Bragg Reflection (DBR) diode lasers that are locked to the $2^3S_1 \rightarrow 2^3P_2$ transition of helium. The output power is typically tens of milliwatts at working conditions. A fourth laser is needed to supply the laser power for the actual experiments, such as the Bragg scattering experiments, and later to feed light into the optical cavity.

The drivers that are used for all four lasers are homemade and stabilize the temperature of the lasers within 1 mK. The free running linewidth of these DBR lasers is typically 3 MHz, which is larger than the natural linewidth $\Gamma/(2\pi) = 1.6$ MHz of the transition used for laser cooling. This will reduce the efficiency of the laser cooling section, especially in the compressor where low intensities are needed. For the QND experiments a laser linewidth broader than the cavity linewidth is also undesirable. The photon number spread due to fluctuations in the laser frequency should be kept smaller than the spread of the coherent state. Therefore the linewidth of the lasers will have to be reduced. Linewidth reduction is achieved by optical feedback from an external mirror. This is a well-known technique [8] and will not be discussed here. However, this optical feedback does not eliminate short-term frequency fluctuations or long-term drift due to inevitable ambient thermal fluctuations. These fluctuations must be compensated for by an external electronic feedback loop. This has been realized, based on saturated absorption spectroscopy in an RF helium discharge.

4.1 Stabilization using saturated absorption

The lasers for the collimator and the two-stage compression section (MOL/MOC combination) require only a small detuning from resonance (-1 to -10 MHz), while the laser used for the actual experiments needs no detuning. These three lasers can be locked using saturated absorption spectroscopy. This technique uses two laser beams, a pump and a probe beam, that propagate through a helium RF discharge along a common axis but in opposite direction (see Figure 4.1). The laser power is simply derived by splitting off a small fraction of the main laser beam with a beam splitter. Only the intensity of the transmitted probe beam is measured.

The technique is based on the property that both counterpropagating beams saturate atoms with opposite velocity components for a given laser frequency provided the laser frequency is within the Doppler width of the spectral profile. However, around resonance those atoms that




Figure 4.1: Setup for saturated absorption spectroscopy. The intensity of the probe beam is measured with a photodetector.

have zero velocity, are addressed to twice and thus saturated more. This results in a lower total absorption from the probe beam and a peak is detected when a transmitted power profile is made of the probe beam.

For the first two (laser cooling) lasers the saturated absorption peak is transformed into a dispersive signal by applying a small oscillating magnetic field over the helium gas cell. The resulting modulation of the transition frequency causes a modulation of the detector signal. Using a lock-in amplifier this signal can be transformed such that the amplifier's output is proportional to the derivative of the absorption profile. Feeding back this error signal, we can lock the laser to its zero crossing with a frequency stability of around 500 kHz. For detuning the laser from resonance, a DC magnetic field is added over the gas cell, hereby shifting the zero crossing in frequency.

The scheme that will be used to actively stabilize the interaction laser (and later also the cavity) is based on the Pound-Drever-Hall scheme [9]. The regular Pound-Drever-Hall scheme locks the frequency of a laser to the resonance frequency of a stable high-finesse cavity with very high precision. Instead of the stable cavity we once more apply the saturated absorption technique. The arrangement of our setup is thus mainly identical to the one used for stabilizing the other two lasers. It consists of a combination of a pump and a probe beam, and a helium gas discharge. The main difference now is that not the transition frequency of the atoms but the frequency of the laser itself is modulated.

4.2 Pound Drever Hall in saturated absorption

As we saw, the key to PDH is to modulate the light instead of the magnetic field around the gas cells. We now have to find the transfer function $F(\omega)$ for a modulated beam that passes through the helium gas discharge cell.

The modulated beam passes in the gas cell through an optical active sample that is assumed to be of length L and to have intensity absorption coefficient α and refractive index n . These properties are all functions of the optical frequency. In order to write the field in terms of

attenuation and phase shift, we now define the optical transmission T in terms of attenuation and phase shift for each spectral component:

$$T_j = e^{(-\delta_j - i\phi_j)}, \quad (4.1)$$

with

$$\delta_j = \alpha_j / 2L, \quad (4.2)$$

$$\phi_j = \frac{n_j L}{c}(\omega_c + j\omega_m), \quad j = 0, \pm 1 \quad (4.3)$$

where $j=0, \pm 1$ indicates the carrier signal and the side bands respectively. We have used δ_j to describe the amplitude attenuation and ϕ_j to describe the optical phase shift. The resulting field $E_T(t) = \frac{1}{2}\tilde{E}_T(t) + c.c.$ emerging after the interaction with the gas sample now becomes:

$$\tilde{E}_T(t) = E_0[-T_{-1}\frac{\beta}{2}e^{i(\omega_c - \omega_m)t} + T_0e^{i\omega_c t} + T_1\frac{\beta}{2}e^{i(\omega_c + \omega_m)t}] \quad (4.4)$$

The photo-detector measures the impinging electrical field resulting in an output voltage which is proportional to the intensity. The light intensity is related to the electric field by $I_T(t) = \epsilon_0 c |\tilde{E}_T(t)|^2$. Dropping terms of order β^2 and assuming that the differences $|\delta_0 - \delta_{\pm 1}| \ll 1$ and $|\phi_0 - \phi_{\pm 1}| \ll 1$, we can approximate this intensity by:

$$I_t(t) = \epsilon_0 c E_0^2 e^{-2\delta} [1 + (\delta_{-1} - \delta_1)\beta \cos \omega_m t + (\phi_1 + \phi_{-1} - 2\phi_0)\beta \sin \omega_m t]. \quad (4.5)$$

The cosine term is proportional to the difference in attenuation for both side bands while the sine term contains the information about the phase shifts for both modulation and carrier signal. Mixing the photodetector's signal with the original modulation frequency and adjusting the phase of this reference signal, we can separate both effects. Since we are not interested in displaying the amplitude of each single effect but in maximizing the error signal's amplitude, we choose a combination of both effects by experimentally optimizing the phase.

4.3 Stabilization setup

The stabilization setup for our laser is schematically shown in Figure 4.2. A function generator produces a 10 MHz sine wave, with an amplitude of roughly 5 Volt. This signal is split in two: one branch passes through an electronic circuit with which the phase of the signal can be adjusted, the other is directly added to the laser current input in order to frequency modulate the laser light.

The laser output beam passes through a system of two lenses, creating a small and parallel laser beam. An optical isolator, positioned directly after the first lens, prevents back reflections back into the diode laser that could cause unwanted feedback. An adjustable half wave plate then ensures that the light is vertically polarized in order to avoid phase shifts upon reflection from the mirrors in the setup.




Figure 4.2: Setup for the laser stabilization based on Pound-Drever-Hall

A first beam splitter then branches part of the modulated beam off towards the saturated absorption setup. With a second beam splitter this branch is split further: the transmitted signal will serve as the pump beam, the reflected beam as the probe beam. The vertically polarized pump beam fully reflects on a polarizing beam splitter cube (PBS) and passes through a $\lambda/4$ wave plate. The now circularly polarized beam is then sent through the helium gas discharge. The partially transmitted signal is not used any further.

Via two gold coated mirrors the second beam approaches the gas cell from the opposite direction. Before passing through the gas cell, this probe beam is also circularly polarized through a $\lambda/4$ wave plate. The transmitted part of the probe beam then passes once more through a $\lambda/4$ wave plate such that it is now horizontally polarized and thus able to pass the PBS. A final lens focusses the transmitted probe beam on a photodetector.

The detector's output is amplified, filtered and then sent to a mixer. With the phase shifted signal from the function generator now serving as the reference signal, the signal that was obtained from the saturated absorption section is demodulated. The signal that emerges from the mixer's output is the discussed dispersive error signal. To reduce the high frequency noise on this signal, it is led through a low pass filter. The filtered error signal is now ready to be used in the active stabilization of the laser frequency through proportional and integral control.

The error signal is split in two. The first part is integrated and sent to a high voltage piezo amplifier that adjusts the laser's feedback mirror. The integration constant is on the order of milliseconds because of the limited bandwidth of the PZT. Fluctuations in the frequency faster than a few hundred Hertz can consequently not be suppressed in this way. Therefore a second feedback loop is formed, based on proportional control. The second split off part of the error signal is amplified and then directly added to the laser's fast current input, which can adjust the laser frequency with a much higher bandwidth for fast error correction. The described system has been constructed and tested. The system is at present operational, but the performance in terms of frequency stability has not yet been evaluated.

Chapter 5

Cavity theory

As the high finesse cavity constitutes a fundamental part of the QND experiments, we are interested in its optical properties and resonance conditions. To obtain an expression for the internal, reflected and transmitted light and the relative phase of these beams, it suffices to examine a simple purely classical model. That is, we assume plane waves, a steady state in the cavity and no quantization of the light field.

5.1 Cavity transfer function

We thus consider a simple cavity consisting of an input mirror M_1 and output mirror M_2 , positioned with an intermediate distance d , as schematically depicted in Figure 5.1. In this model the mirrors are assumed infinitely thin, and their optical properties are described with amplitude reflection coefficients r_1 and r_2 , and transmission coefficients t_1 and t_2 respectively.




Figure 5.1: In general a cavity consists of two highly reflective mirrors M_1 and M_2 , positioned at an intermediate distance d that may be filled with an absorbing medium

The coefficients r and t may be complex numbers to include a π phase shift between the transmitted and reflected beam. This phase shift is introduced when going from a low-index to a high-index material in a dielectric interface. For lossless and reciprocal two-port networks, as are the mirrors in our model, the following conditions apply to t and r :

$$\begin{aligned} |r|^2 + |t|^2 &= 1 \\ r t^* + t r^* &= 0, \end{aligned} \quad (5.1)$$

where the star denotes the complex conjugate of the coefficient. If we were to follow most textbooks, we could fulfill these conditions and allow for the required phase shift by stating that the reflection coefficients have opposite signs depending on the direction from which the light wave approaches the interface. The coefficients r and t can then be conveniently chosen purely real. This is, however, only one of the possible ways in which the two conditions of (5.1) can be satisfied. As we now wish to assign the reflection coefficients r to the (infinitely thin) mirror as a whole, symmetric coefficients are preferred. The price for doing so is that, when calculating with the transmission coefficients, we now have to take into account an additional complex factor of i , or a phase shift of $\pi/2$, in the transmission coefficient. In the following discussion we will always write the additional factor i explicitly so that the coefficients r and t themselves are again purely real but now also symmetric for each mirror.

Since we often consider the effects of losses and phase shifts after one round-trip, we also define the so-called perimeter or round-trip path length p which is in our simple model twice the cavity length $p = 2d$.

In a general approach the cavity may be filled with a material that causes a signal, that is going around inside, to be (slightly) attenuated. The reduction of the signal amplitude per round-trip is in general given by a factor $e^{-\alpha p}$ with α the absorption coefficient of the attenuating material.

Circulating inside a finite size cavity the light signal will also pick up a phase shift. This propagation factor equals e^{ikp} or $e^{i\frac{\omega}{c}p}$ per round-trip path p , where $k = \omega/c$ is the wave vector, ω the frequency of the signal and c the velocity of light.

After covering one perimeter, a circulating signal will return to its reference point with a net round-trip transmission factor, or net complex round-trip gain, which for a passive lossy cavity is given by:

$$\tilde{g}_{RT}(\omega) = r_1 r_2 e^{-\alpha p - i\frac{\omega}{c}p} \quad (5.2)$$

Ignoring initial transient effects, we will suppose that the cavity is in a steady-state, i.e., all considered beams are in equilibrium. The complex amplitudes of the optical signals that are incident on and reflected from the first mirror M_1 , as measured just outside this mirror, are designated \tilde{E}_{inc} and \tilde{E}_{refl} respectively. We will use \tilde{E}_{circ} to denote the circulating signal amplitude inside the cavity as measured just after mirror M_1 .

In our steady state approach the field amplitude inside the cavity consists of two contributions. The first part is formed by the portion of the incoming beam that is transmitted through the input mirror M_1 and has a value $i t_1 \tilde{E}_{inc}$. The second part originates from the already existing cavity amplitude. The contribution represents the circulating signal that left the same position one round-trip time earlier, and which has thus completed one full round-trip path length p . The effect of reflection on both mirrors and losses in the cavity are accounted for by the transmission factor \tilde{g}_{RT} . The cavity amplitude can now be written as an recursive expression:

$$\tilde{E}_{circ} = i t_1 \tilde{E}_{inc} + \tilde{g}_{RT}(\omega) \tilde{E}_{circ}. \quad (5.3)$$




Figure 5.2: Elementary model for the incident, reflected, transmitted, and circulating waves in a resonant optical cavity

Rewriting this equation, we can relate the circulating signal inside the cavity to the incident signal outside the cavity:

$$\frac{\tilde{E}_{circ}}{\tilde{E}_{inc}} = \frac{it_1}{1 - \tilde{g}_{RT}(\omega)} \quad (5.4)$$

From this formula it is clear that cavities with low internal and reflective losses (i.e., α very small and $r_{1,2}$ close to 1) display a strong resonance behavior each time the round-trip phase shift $\omega p/c$ equals an integer multiple of 2π . In our case the frequency of the laser light is fixed, so we should vary the cavity length such that $d \equiv p/2 = n \times \pi/k = n \times \lambda/2$, with n being an integer. For these conditions the circulating intensity inside the resonator becomes many times larger than the intensity of the beam incident on the cavity. For a symmetric cavity with small internal losses the peak value of the internal field on resonance is approximated by:

$$\left. \frac{\tilde{E}_{circ}}{\tilde{E}_{inc}} \right|_{p=p_{res}} = \frac{it}{1 - r_1 r_2 e^{-\alpha p}} \approx \frac{it}{1 - r^2} = \frac{i}{t} \quad (5.5)$$

where we have used $t_1 = t_2 = \sqrt{1 - r^2}$. In other words, the ratio of the intensity inside the cavity to the incident intensity scales with:

$$\left. \frac{I_{circ}}{I_{inc}} \right|_{p=p_{res}} \approx \left| \frac{1}{t} \right|^2 = \frac{1}{T} \quad (5.6)$$

with $T \equiv t^2$ the power transmission of the mirror which, in our cavity, amounts to a mere few part per million!

For the analysis and stabilization of the cavity we however can't use the internal field. Therefore we need to examine the reflected and transmitted field. Using the same method as before, we find that the transmitted signal amplitude leaking out through mirror M_2 will be given by:

$$\tilde{E}_{trans} = i t_2 \tilde{g}_{RT} \tilde{E}_{circ}. \quad (5.7)$$

from which we simply derive a similar equation for the net transmission through the cavity, from input to output:

$$\frac{\tilde{E}_{trans}}{\tilde{E}_{inc}} = -\frac{t_1 t_2}{\sqrt{r_1 r_2}} \frac{\sqrt{\tilde{g}_{RT}(\omega)}}{1 - \tilde{g}_{RT}(\omega)} \quad (5.8)$$

The minus sign expresses a phase shift of π that arises from our definition of r and t . As a consequence of our definition of the reference planes for the transmitted and incident beam, Equation 5.8 results in a phase shift $-e^{-i\frac{\omega}{c}d}$ between the two beams even in the absence of a cavity (i.e., inserting $t_1 = t_2 = 1$, $r_1 = r_2 = 0$, and $\alpha = 0$). This phase is clearly the propagation factor associated with the distance d between the reference points.

The calculation of the field amplitude for the reflected wave, finally, is also rather straightforward now when we realize that it again has two contributions: one component is due to the reflection of the incoming beam on the first mirror surface, the other reflects the leakage of the internal signal \tilde{E}_{circ} through this mirror into the same direction. Note however that the latter component emerges from the circulating signal that left its starting position one round-trip time earlier: in this time it has travelled around but instead of bouncing off mirror M_1 it is now transmitted through it. The total reflected wave thus consists of

$$\tilde{E}_{refl} = r_1 \tilde{E}_{inc} + i t_1 \frac{\tilde{g}_{RT}(\omega)}{r_1} \tilde{E}_{circ} \quad (5.9)$$

Substituting the earlier expression for \tilde{E}_{circ} , the total reflection coefficient for the input mirror is:

$$\frac{\tilde{E}_{refl}}{\tilde{E}_{inc}} = r_1 - \frac{t_1^2}{r_1} \frac{\tilde{g}_{RT}(\omega)}{1 - \tilde{g}_{RT}(\omega)} = \frac{1}{r_1} \frac{r_1^2 - \tilde{g}_{RT}(\omega)}{1 - \tilde{g}_{RT}(\omega)}. \quad (5.10)$$

In the last step we used the expression $r_1^2 + t_1^2 = 1$ for a lossless mirror. If we now furthermore assume no internal losses inside the cavity and identical properties for both mirrors, we recover the formula that is found following from the standard textbook approach of adding an infinite number of partial waves:

$$\frac{\tilde{E}_{refl}}{\tilde{E}_{inc}} = \frac{r(1 - e^{-ikp})}{1 - r^2 e^{-ikp}} \quad (5.11)$$

5.2 Finesse

Let's look again at the transmitted signal for resonant optical cavities, formula (5.8). For the calculations in this section it is more convenient to rewrite the complex round trip gain in the form $\tilde{g}_{RT}(\omega) = g_{RT} e^{-i\phi_{RT}}$, where $g_{RT} = |\tilde{g}_{RT}(\omega)|$ and $\phi_{RT} = \frac{\omega}{c}p$ the round trip phase.

The ratio of the transmitted intensity to the incident intensity for a symmetric cavity is thus given by:

$$\frac{I_{trans}}{I_{inc}} = \left| \frac{\tilde{E}_{trans}}{\tilde{E}_{inc}} \right|^2 = \frac{t^4}{r^2} \left| \frac{\sqrt{g_{RT}} e^{-i\phi_{RT}/2}}{1 - g_{RT} e^{-i\phi_{RT}}} \right|^2$$

$$= \frac{t^4}{r^2} \frac{g_{RT}}{1 + g_{RT}^2 - 2g_{RT} \cos(\phi_{RT})} \quad (5.12)$$

With a little algebra the expression then becomes:

$$\frac{I_{trans}}{I_{inc}} = \frac{t^4}{r^2} \frac{\frac{g_{RT}}{(1-g_{RT})^2}}{1 + \frac{4g_{RT}}{(1-g_{RT})^2} \sin^2\left(\frac{\phi_{RT}}{2}\right)} \quad (5.13)$$

The transmitted signal displays the strong resonance behavior each time the round-trip phase shift $\omega p/c$ equals an integer multiple of 2π , similar to the behavior we already saw for the cavity field. The amplitude of the transmitted intensity at resonance, or maximum transmission T_{max} , is:

$$T_{max} = \frac{t^4}{r^2} \frac{g_{RT}}{(1 - g_{RT})^2} = \frac{t^4 e^{-\alpha p}}{(1 - r^2 e^{-\alpha p})^2} \quad (5.14)$$

Using the relation of Eq. 5.13 we can determine the full width at half maximum (FWHM) linewidth $\Delta\omega_{FWHM}$ of the resonance peaks:

$$\begin{aligned} \Delta\omega_{FWHM} &= \frac{4c}{p} \sin^{-1} \left(\frac{1 - g_{RT}}{2\sqrt{g_{RT}}} \right) \\ &\approx \frac{2\pi c}{p} \frac{(1 - g_{RT})}{\pi\sqrt{g_{RT}}} = \frac{\Delta\omega_{FSR}}{\mathcal{F}} \end{aligned} \quad (5.15)$$

where $\Delta\omega_{FSR} \equiv 2\pi c/p$ is the mode spacing or free spectral range between two resonances. The approximation in the second line assumes that the gain amplitude g_{RT} is close to unity. The resonance bandwidth is only a fraction of the free spectral range and becomes narrower as the round-trip gain approaches closer to unity. The finesse \mathcal{F} is now defined as the ratio of the free spectral range to the FWHM linewidth:

$$\mathcal{F} \equiv \frac{\Delta\omega_{FSR}}{\Delta\omega_{FWHM}} = \frac{\pi\sqrt{g_{RT}}}{1 - g_{RT}} \quad (5.16)$$

Inserting these parameters in Eq. 5.13, we find the usual Airy function for the power transmission through a symmetric etalon:

$$\left| \frac{\tilde{E}_{trans}}{\tilde{E}_{inc}} \right|^2 = \frac{T_{max}}{1 + (2\mathcal{F}/\pi)^2 \sin^2(\pi\omega/\Delta\omega_{FSR})} \quad (5.17)$$

In standard literature, losses in the cavity are generally accounted for by assuming absorption loss A in the mirrors. Instead of the relation $R+T=1$, we then assume that $T+R+A=1$. In this model, the maximum transmission is written as

$$T_{max} = \frac{T^2}{(1-R)^2} = \left(1 - \frac{A}{1-R} \right)^2 \quad (5.18)$$

It can easily be shown that this expression is equivalent to Eq. 5.14 for a high-finesse cavity.

5.3 Phase response

The cavity introduces a certain phase shift between the incoming, the circulating, the reflected, and the transmitted fields. Using the result of Equation 5.10 we can calculate the phase shift of the reflected field, which will be used for the stabilization scheme, with respect to the phase of the incoming beam:

$$\phi_{refl} = \arctan \left(\frac{r_2 e^{-\alpha p} (1 - r_1^2) \sin \phi_{RT}}{r_1 (1 + r_2^2 e^{-2\alpha p}) - r_2 e^{-\alpha p} (1 + r_1^2) \cos \phi_{RT}} \right) \quad (5.19)$$

For a symmetric cavity this expression simplifies to:

$$\phi_{refl} = \arctan \left(\frac{e^{-\alpha p} (1 - r^2) \sin \phi_{RT}}{(1 + r^2 e^{-2\alpha p}) - e^{-\alpha p} (1 + r^2) \cos \phi_{RT}} \right) \quad (5.20)$$

where $\phi_{RT} = \frac{\omega}{c} p$ is the round-trip phase delay. On resonance the phase shift term for the reflected beam is equal to zero. Around resonance the phase diagram shows a dispersive character: there is a π phase change right across the cavity resonance. It is exactly this feature that is exploited in the Pound-Drever-Hall scheme.

The incoming light is modulated and the reflected light is detected. The two sidebands experience almost no phase shift. Interference with the reflected light at the fundamental frequency results in signals at the modulation frequency which are the sum of the two beat signals between the sidebands and the central laser frequency. On resonance these two beat signals are out of phase and cancel each other. For a slightly detuned cavity however, the difference in phase shifts will prevent this cancelation and a beat signal appears at the modulation frequency. After demodulation with the original modulation signal, this beat signal provides a dispersive error signal around the cavity resonance, providing the basis for the Pound-Drever-Hall scheme.




Figure 5.3: Cavity response. Phase and absolute value of the reflected beam plotted versus cavity length detuning. The reflectivity is $|r|^2=0.99$, the absorption $\alpha=0$ (solid) and $\alpha=0.05$ (dashed)

5.4 Cavity modes and Gaussian beams

In a previous section we already discussed some features of a cavity, and deduced the transmission and reflection coefficients for a cavity consisting of two plane highly reflective mirrors. We found that resonances occur when the cavity separation d is equal to an integral multiple n of half wavelengths:

$$d = n \left(\frac{\lambda}{2} \right) \quad (5.21)$$

At each resonance the amplitude of the electromagnetic wave is enhanced due to the fact that under this condition a standing wave establishes such that it constructively interferes with itself after each reflection. Each of these resonant waves is referred to as a longitudinal mode because they are associated with the longitudinal direction of the cavity. In principle an infinite number of resonant frequencies ν_n and thus different modes would fit within a cavity:

$$\nu_n = n \frac{c}{2d}. \quad (5.22)$$

However, the strongly wavelength dependent reflectivity of the cavity mirrors limits this number.

The actual cavity is constructed with spherical mirrors. In a stable cavity configuration this leads to a transverse mode structure when solving the wave equation for the electric field in the cavity. If we restrict ourselves to small angles around the beam axis, we can use the paraxial approximation to the wave equation:

$$\nabla^2 E + ik \frac{\partial E}{\partial z} = 0 \quad (5.23)$$

The simplest free-space solution of Eq. 5.23 is the Gaussian beam, which has a circular symmetry around the optical axis and is in general a good representation for a laser beam. There are also many other spatial solutions with different intensity profiles. A particular set of solutions is formed by the Hermite-Gaussian beams. The mode function (or wave amplitude) U for the transverse electric field then becomes in two dimensions [3]:

$$U_{lm}(x, y) = H_l \left(\frac{\sqrt{2}x}{w(z)} \right) H_m \left(\frac{\sqrt{2}y}{w(z)} \right) e^{-(x^2+y^2)/w(z)^2} \quad (5.24)$$

where the waist $w(z) = w_0 / \sqrt{1 + (z/L)^2}$, $L = \pi w_0^2 / \lambda$ and w_0 equals to the minimum of $w(z)$. In this solution H_l and H_m are the Hermite polynomials of order l and m , respectively. Every set (l, m) represents a specific wave profile at one of the cavity mirrors. In other words, the shape of the transverse field profile of the modes is determined by a Gaussian function multiplied with the Hermite polynomials. Each of the Hermite-Gaussian solutions U_{lm} leads to a specific stable transverse mode of the cavity, each with its own series of longitudinal modes, spaced by $c/2d$ in resonance frequency, with d the on-axis distance between the curved mirrors. The full mode is thus specified by the mode numbers (l, m, n) . For a cavity with two identically curved mirrors with radius of curvature R_c , the resonance frequencies ν_{lmn} of these modes are given by:

$$v_{lmn} = \delta_l + \delta_m + n \frac{c}{2d}, \quad (5.25)$$

with the transverse mode off-sets δ_l and δ_m can be derived to be given by [4]:

$$\delta_l + \delta_m = (1 + l + m) \frac{c/2d}{\pi} \left[\arccos\left(1 - \frac{d}{R_c}\right) \right] \quad (5.26)$$

We designate these transverse mode distributions (Eq. 5.24) as transverse electromagnetic or TEM modes. The lowest-order mode is the TEM₀₀ mode which has a Gaussian distribution of the form $E_0 e^{-(x^2+y^2)/w^2}$ for the electric field with E_0 the electric field amplitude at the center of the beam.

5.5 Mode matching

The cavity has its own geometrical mode that is determined by the curvature of the two mirrors. The incoming light wave can only fully interact with the cavity if the spatial properties of this incoming wave match with the internal mode. The wave front curvature should therefore exactly coincide with the curvature of the mirrors which in turn is the curvature of the internal mode. In such case we would have a mode matched system. In practice, the wave front of the laser light incident on the cavity will not be identical to that of the internal mode. The cavity does then not only have a resonance on one fundamental mode, but also on a whole series of higher order spatial modes. These modes will generally not be resonant simultaneously: a small shift in the mirror position is required for the resonance condition of higher order modes.

On resonance the cavity will select the part of the external field which corresponds to the resonant internal mode, an eigenmode of the cavity. Only this light will enter the cavity and contribute to the resonant power build-up. The remaining non-matching part of the beam will be reflected. Thus the equations for the transmitted and reflected power should actually be modified by including a factor that accounts for the fraction of the total power that corresponds to the overlap with higher order modes.

The proper alignment of a mode matched cavity is an extensive job requiring patience and good quality components. In general there are six degrees of freedom that have to be optimized. Two degrees of freedom for the tilt of the input beam in respect to the cavity, two degrees for the position of the center line of the beam and two degrees for the wave front curvature. In our setup, where the light source and the cavity have more or less fixed positions due to the lack of space on the optical tables, we tackle the alignment problem with an optical system consisting of two lenses and two mirrors.

5.6 Polarization

The effect of polarization was not included in the above discussion as for a perfectly isotropic cavity the polarization of the light does not matter. However, in practice the cavity mirrors will, to some extent, be anisotropic. The anisotropy may, e.g., introduce birefringence. This makes life much more complex.

In case of birefringence the losses and the effective distance between the mirrors can differ for different states of polarization. In the description of the cavity response we now have to search for the polarization eigenstates of the cavity. Light in these states will return with the same polarization after one round-trip in the cavity. It should always be possible to find two such orthogonal eigenstates.

The description of the cavity now requires two sets of equations, one for each polarization state. As a result, each polarization eigenstate has its own resonance condition. The resonant modes of the cavity are split in two, with different resonance frequencies. If this splitting is larger than the width of the resonance, it is not possible any more to produce light with an arbitrary polarization in the cavity. As the QND experiment is to be performed with circularly polarized light, this constitutes an unwanted effect. Measurements of the polarization properties of the cavity are presented in the next chapter.

Chapter 6

Cavity diagnostics

In the proposed QND experiments the photon statistics of the optical field in a high-finesse cavity is to be studied by measuring the diffraction of a traversing beam of metastable helium atoms. In the interesting configuration where the cavity field consists of a very low number of photons, the diffraction angles will be very small. In order to be able to measure these effects a slow and highly collimated atomic beam is a logical necessity. Just as crucial for the QND experiments is the high-finesse cavity. During the time span that it takes an atom to pass the cavity the interaction strength needs to be approximately constant. For our cavity this is equivalent to saying that the decay time of the cavity at resonance should be longer than $0.4 \mu\text{s}$. The realization of such a high quality cavity poses a challenging technical problem.

6.1 Cavity design

Before we can perform any experiment with our cavity, its precise properties and behavior need to be known. Our high-finesse cavity consists of two highly reflective mirrors, supplied by Research Electro-Optics, Boulder. The mirror diameter is 7.75 mm, and the radius of curvature is 50 mm. The mirrors are glued to cylindrical holders, made from glass. The distance between the mirrors was determined by Knops [1] to be 0.93 mm. For diagnostic experiments however we wish to be able to scan this distance. The glass holders are therefore clamped into a tubular piezo electric transducer (PZT) that allows us to adjust the mirror separation in a controlled manner. Later the same PZT will be deployed to stabilize the position of the mirrors such that cavity remains resonant with the laser light during the experiment. Holes in the mirror holder and the PZT allow the passage of the atom beam and the transverse light field. In Fig. 6.1 a cross-section of the cavity construction is shown.

6.2 Diagnostics

In 1998 Knops [1] already performed some diagnostic measurements on the cavity. The cavity has been kept well-protected from the environment ever since in an airtight container with optical access through vacuum windows. However, we cannot exclude deterioration over the last three years.




Figure 6.1: Schematic view of the cavity

New measurements are therefore necessary to verify the cavity's optical quality. Furthermore, extra information on the polarization properties and on the mechanical and electrical response of the cavity are needed.

Knops determined the finesse of the cavity by measuring the ring-down times for various modes. In his setup two methods were alternatively used to rapidly shut down the laser beam feeding the cavity (a fast acousto-optic modulator switching the laser beam, and a fast FET switching the laser current respectively). In both methods the cavity was scanned by the PZT and the laser beam switched off as soon as a transmission peak was detected. From the decay of the transmitted light through the cavity the decay time of the light in the cavity and hence the finesse can be determined. This resulted in a finesse of $\mathcal{F} = 6.5 \times 10^5$.

This method doesn't provide any information of the polarization properties of the cavity. The design of the suspension of the cavity mirrors is such to avoid any stress which could lead to stress-induced birefringence in the mirrors. However, as in a high-finesse cavity even the smallest mirror birefringence can lead to a polarization-dependent mode splitting, polarization effects can be expected anyway.

Another important criterion for the design of the cavity control circuit is the maximum bandwidth available for the active stabilization. The bandwidth is strongly constrained by the cavity's electric and mechanical behavior at different frequencies. In practice this means that we have to look for the first resonances - either electronic or mechanical - that appear in the construction.

6.3 Cavity test facility

The laser used is a 1083 nm Distributed Bragg Reflection (DBR) diodelaser (type SDL-6702-H1). A small part of the output power is used to stabilize the laser at the right frequency using saturated absorption spectroscopy based on the Pound-Drever-Hall scheme, as discussed in a previous chapter. The main part of the laser output is coupled into a single mode polarization preserving optical fiber. This fiber transports the beam from the optical table, where the laser is stabilized, to our test facility. This transportation of the light through an optical fibre has two major advantages. First, the stabilization setup for the laser now stands completely separated

from the optical table on which the cavity analysis facility is placed. Small movements in the latter setup that are inevitable during the build-up or vibrations produced during the experiments cannot disturb the alignment of the laser stabilization setup at the other table. Secondly, the optical fibre not only preserves the polarization of the laser light, but also functions as an effective spatial filter filter, drastically improving the spatial structure of the beam.




Figure 6.2: Experimental setup used for mode matching the beam to the cavity and to perform polarization experiments. L: lens, M: mirror, BS: beam splitter, BW: Brewster window, PD: photodetector, CCD: CCD camera

Because of the fibre's narrow diameter the light at the output is highly divergent. A 16x microscopic objective (L1) placed a few millimeters behind the fibre's output recollimates the beam into a parallel beam.

In general, as discussed before, an infinite number of electromagnetic modes can exist in a cavity. For the experiments with the cavity however, we are primarily interested in the fundamental Gaussian TEM_{00} mode. This mode has to be selectively excited by matching the wave front of the incident beam to the curvature of the cavity mirrors and to the shape of the desired mode.

This process is referred to as mode-matching of the input beam to the cavity mode. For our Gaussian beam this implies that the waist position, the waist size, and the axis of the input beam must coincide exactly with those of the fundamental cavity mode. Via an optical system containing two extra lenses (L2 and L3) and two gold coated mirrors (M1 and M2) the light beam is led to the cavity. The two lenses (focus lengths 76.2 mm) create the required waist size of $40\ \mu\text{m}$ at the center of the cavity, while the mirrors provide means of adjusting the angle and position of the axis of incidence. The approximate positioning of the lenses was determined using the results of a paraxial optics calculation of the transport of the Gaussian beam through the system, implemented in Maple. An overview of the setup is shown in Figure 6.2. Depending on the measurements carried out and thus the polarization needed, additional optical elements as $\lambda/4$ and $\lambda/2$ wave plates and polarizing beam splitter cubes (PBSC) can be inserted.

After proper alignment the cavity will transmit part of the incoming light beam if the laser frequency is resonant with a TEM_{00} mode of the cavity. This output signal is then split in two by a 50/50 beam splitter. One beam passes a lens and is projected onto a CCD camera for the

direct monitoring of the spatial distribution of the transmitted light. This makes it possible to verify that the cavity is resonant with the fundamental mode during the measurements. The camera image also provides information on the intensity of the beam, which proves to be an indispensable tool in the alignment of the cavity. The other beam, finally, is focussed by a lens on a fast GaAs photodetector with integrated amplifier. This signal can be displayed on an oscilloscope.

6.4 Beam alignment

When the mirror distance is scanned while the laser is not perfectly mode-matched to the fundamental mode, many higher-order modes will become visible in succession. However, for the experiment, we are interested only in the fundamental mode. Thus we have to suppress the higher order modes by careful alignment. For this purpose the total power in the higher order modes can be used as a (rough) measure for the quality of the alignment: by decreasing the overlap of the incident light wave with the higher-order modes, more power is available for the TEM_{00} mode. A perfect alignment is in practice unobtainable, but also not strictly necessary for our diagnostic measurements. Typically we find profiles as depicted in Fig. 6.3 which are satisfactory for our purpose.

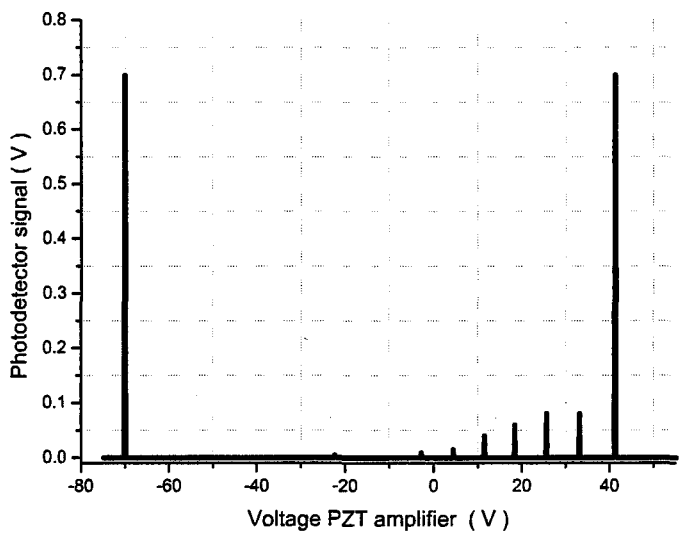


Figure 6.3: When the cavity mirror distance is scanned, many higher order modes are found. The two large peaks indicate the fundamental TEM_{00} mode

Although the intensity of the transmitted signal itself does not reveal any knowledge about the cavity mode, we can verify the mode by looking at the beam profile as registered with the CCD camera.

Let's now have a better look at the TEM_{00} mode. Unfortunately, we cannot have the cavity permanently on resonance as the linewidth is a factor of 100 smaller than the fluctuations in the mode resonance frequency introduced by the voltage fluctuations of our supply. We therefore

apply a DC voltage to the piezo transducer such that the cavity is as close to the resonance as possible. If we now add a small AC signal, the cavity will periodically sweep through resonance. Connecting the photodetector to an oscilloscope that is triggered on the same sweep signal, the intensity profile of the cavity sweeping through resonance will be revealed. Figure 6.4 shows the TEM_{00} resonance peak. As the measurement was performed with an unstabilized diode laser, the peaks are broadened much wider than the cavity linewidth.




Figure 6.4: Sweeping the cavity through resonance, two transmission peaks are visible: our cavity is birefringent.

It is immediately obvious that there is not one transmission peak but two peaks. Knowing the amplitude and frequency of the sweeping signal, the relative frequency separation for both resonances is calculated to be about 5 MHz. This separation is significantly smaller than for consecutive higher-order modes. Both peaks thus clearly represent the transmission signal of the TEM_{00} cavity mode, as was also easily verified on our CCD camera. This would indicate that our cavity is slightly birefringent.

To investigate the hypothesis of birefringence the transmission intensity is measured for different angles of polarization of the incoming light beam. To rotate the angle of polarization we used two methods. First we inserted a $\lambda/4$ wave plate (QWP) and a polarizer beam splitter cube (PBSC) in the setup just before the cavity. Later, these two components were replaced with a $\lambda/2$ wave plate.

A QWP is a polarization-selective phase retarder. Such a retarder introduces a phase shift $\Delta\phi$ between the two linear polarization components of the incident light oriented along and perpendicular to the optical axis of the retarder. If the net phase difference $\Delta\phi = 90^\circ$, the retardation plate is called a quarter wave plate (QWP), if the difference is 180° it is called a half wave plate (HWP). When linearly polarized light is incident on a QWP at an angle of 45° with respect to the plate's optical axis, the light is decomposed into two orthogonal beams of equal intensity. On emerging the net phase difference of 90° results in circularly polarized light.

For linearly polarized light the HWP has the effect of a rotator: it rotates the direction of

polarization by some particular angle. This angle is simply twice the angle between the HWP's optical axis and the direction of polarization of the incident light. By rotating the polarization axis of the wave plate the polarization angle of the light that passes through is varied. For each angle we then measure the intensity of the transmitted signal at both maxima of Fig. 6.4 separately. The results are shown in Fig. 6.5.




Figure 6.5: The amplitude of the transmission as a function of the angle of polarization of the incident beam. The two graphs indicate the two transmission peaks of Figure 6.4.

The apparent difference in maximum intensity for the two transmission peaks is due to the fact that the light has to pass two Brewster windows. At each window surface about 15% of the horizontally polarized component of the light beam is reflected. After two Brewster windows (four surfaces) we thus in total lose 50% of the horizontal polarization whilst the vertical component is fully transmitted. Furthermore the amplifier of the photodetector that is used to measure the intensity of the transmitted light saturates quickly for higher intensities. The detector's output signal is then no longer linear to the incident light intensity.

Taking these effects into account, we are inclined to conclude from this profile that the two peaks in the transmission spectrum indeed correspond to two (approximately orthogonal) linear polarized cavity modes. The cavity thus displays birefringence.

At first sight, the polarization eigenmodes do not seem to be exactly orthogonal. This is rather counter-intuitive. However, the apparent non-orthogonality is largely caused by the Brewster windows as well. The optical axis of the Brewster windows (the polarization direction which is not attenuated) does not coincide with one of the polarization eigenmodes. The attenuation of the polarization component perpendicular to the Brewster windows' optical axis squeezes the plot shown in Fig. 6.5 in the horizontal direction, causing the apparent non-orthogonality of the rotated modes.

The birefringence of our cavity is inconvenient, since we strongly prefer a circularly polarized cavity field for the QND experiments. With the lifting of the degeneracy of the two linearly polarized modes, it is not possible to resonantly produce circularly polarized light in the cavity. The interaction strength for the atom-cavity system depends on the Clebsch-Gordan coefficients of the atomic transition. As this coefficient is smaller for linearly polarized light

than for circularly polarized light, the interaction strength will be accordingly lower.

Although it seems that only linearly polarized light can be excited in the cavity resonantly, this does not mean that there cannot be any circularly polarized light present inside the cavity. For example, by feeding the right elliptically polarized light, it should always be possible to excite both linear modes equally strong and with a 90° phase difference. By setting the distance of the cavity mirrors such that the cavity finds itself exactly in between the two resonances, circularly polarized light can be produced in the cavity, albeit with reduced field strength for a fixed input beam intensity.

However, it is not exactly clear how the off-resonant circularly polarized light in the cavity behaves quantum-mechanically, as the usual quantization procedure is based on resonant modes only. We are presently looking into the problem.

If this approach turns out not to be feasible, we will either have to perform the experiment with reduced interaction strength or remedy the birefringence of the cavity. The birefringence is most likely stress-induced, i.e., strain in the mirror suspension causes an anisotropy in the mirror surfaces. By applying extra forces on the mirrors in a controlled way, it should in principle be possible to compensate for the original strain in the holder. One could thus think of a suspension design where additional PZT elements, placed around the mirror holders, serve to minimize birefringence.

6.5 Cavity stabilization

The laser that will be used for the experiments is frequency modulated at 10 MHz by direct current modulation, as was necessary for the stabilization on the helium discharge cell. It is thus not surprising that we will also try to exploit the same modulation to stabilize the cavity on the right frequency. The setup for locking and stabilizing our high-quality cavity to the laser light will thus be based on the Pound-Drever-Hall scheme, analogous to the stabilization of the laser to helium transition, as described in Chapter 4.

As a consequence of the high finesse of our cavity, its bandwidth is proportionally small: approximately 250 kHz. This means that when our cavity is on resonance with the modulated laser light (or just slightly detuned from resonance), only the central carrier frequency can pass through. The sidebands at 10 MHz are so far from the cavity resonance (40 bandwidths) that they effectively experience total reflection. Yet these sidebands are essential for a stabilization circuit that is based on the Pound-Drever-Hall scheme. We can therefore not use the cavity's transmission signal, but have to look at the light that reflects from it. In the situation that the cavity is far off-resonance, reflection of the incoming beam is virtually 100%. The amplitude of the reflected signal sharply drops as soon as the cavity approaches resonance, as is clear from the cavity theory discussed earlier. Basically, on resonance, the field leaking out of the cavity into the direction of the incoming beam is in anti-phase with the directly reflected beam. For a lossless cavity the reflected and the leakage field will have the same amplitude, thus exactly nullifying each other through destructive interference. For any realistic lossy cavity, the depth of this drop is less trivial. Looking at Eq. 5.18 we see that the maximum transmittance is equal to the incoming intensity except for a term that is directly related to the absorption in the mirrors. As we assume a symmetric cavity, the same field will leak back through the first mirror. The expected minimum of the intensity reflection R_{min} is equal to:

$$R_{min} = \left(\frac{A}{1 - R} \right)^2 \quad (6.1)$$




Figure 6.6: Transmitted and reflected signal at resonance

6.6 Electrical and mechanical behavior

For the stabilization circuit it is also important to know how the cavity behaves electronically and mechanically. Mechanical and electrical resonances in the cavity construction may reduce the bandwidth available for feedback control.

To investigate the cavity's electronic behavior, the dynamical impedance for frequencies from 100 Hz to 200 kHz was measured. We applied a sinusoidal signal of about 1 Volt amplitude, provided by a function generator, over the cavity PZT. A 1 kΩ resistance was added in series with the cavity over which we measured the voltage in order to determine the current through the PZT. Comparing the amplitude and phase of this output signal to the initial drive voltage over the cavity, the complex impedance of the cavity can be determined.

Electronically speaking, the cavity consists of a loaded piezo electromechanical transducer. As PZTs are made of non-conducting ceramics, they can in a first approximation be considered as pure capacitors. Figure 6.7 shows the results of the measurements. For frequencies up to 10 kHz the behavior of the electronic circuit is indeed predominantly that of a capacitor. For this frequency range the capacitance value is about 7 nF.

Small perturbations in the complex impedance already appear at frequencies of about 4 kHz and more serious resonances are seen for frequencies of 50 kHz. The small disturbances are mainly due to mechanical resonances in the cavity construction as Jansen [7] also showed in other measurements. This is supported by the fact that, in the measurement as described, the electrical perturbations at 4-5 kHz were accompanied by clearly audible resonances. It is unclear whether these resonances are specific to the current construction and may disappear when the cavity is installed in the final interaction chamber of the QND beam machine. Nevertheless,




Figure 6.7: Electrical response of the high finesse cavity versus frequency. For frequencies lower than 5 kHz the behavior is predominantly that of a capacitor. At higher frequencies strong resonances are present.

we will decide to limit the bandwidth of the feedback control to about 4 kHz. If absolutely necessary, the PI controllers can later always be upgraded to higher bandwidths.

Chapter 7

Cavity stabilization

In the proposed experiment the diffraction of a beam of atoms interacting with the electromagnetic field in a high-finesse optical cavity will be investigated. It is therefore necessary to have a light source that is resonant with the proper helium transition. In Chapter 4 we extensively discussed the stabilization scheme that is used to keep our diode laser at the right frequency. We wish to investigate the properties of an electromagnetic field consisting of only one or a few photons on average. However, to be able to detect the diffraction effects a strong interaction is needed. A high-finesse cavity is then indispensable for the storage and power build-up of this field. In the previous chapter we already had a closer look at the design of this cavity and performed some diagnostic tests that displayed the main optical and electronic behavior of the cavity. In order to reach maximum interaction strength of a helium atom with the electromagnetic field inside the cavity, the cavity should also be resonant with the atom's transition. Having a locked laser, this is equivalent to saying that the cavity should be resonant with the laser light.

7.1 Stability demands

As we saw, our high-finesse cavity is only resonant with the transition frequency of helium atoms when its mirror separation is an integer multiple of $\lambda/2$, with λ the radiation wavelength, which is on the order of half a micron. It is however extremely difficult to build a cavity that has its mirror separation to this precision in absolute distance and even more, that holds that distance constant for the duration of the measurements.

We can easily make an estimate of the stability demands. In the previous chapter we found that the finesse \mathcal{F} of a cavity is defined as the ratio of the free spectral range (FRS) to the linewidth of the cavity resonance. As one free spectral range of our cavity corresponds to a shift of $\lambda/2$, the cavity linewidth clearly corresponds to a motion of $\lambda/2\mathcal{F}$. Setting the finesse value $\mathcal{F} = 6.5 \times 10^5$ as determined in ring-down experiments [1], we find that we have to stabilize the mirror spacing to an accuracy on the order of 1 pm.

Piezo transducers are particularly useful when displacements of these magnitudes are required. Therefore the mirrors of our high finesse cavity are suspended in a tubular piezo-electric transducer (PZT): by applying a voltage to the PZT we can accurately adjust the distance between the mirrors and consequently adjust the cavity's resonance frequency. If we gradually

vary the applied voltage over the PZT, it will always possible to find the cavity resonance with the atom's transition. The maximum range over which we may have to change the separation distance is now conveniently less than $\lambda/2$.

On resonance however, the demand for a 1 pm stable mirror spacing still stands. In order to minimize sensitivity to vibrations and thermal fluctuations the design of the cavity is kept compact and relatively simple. But even when these precautions are respected, thermal expansion of the construction can cause the cavity to drift from resonance. This drift must be compensated for, using an additional feedback loop, exploiting once more the suspension design. Once resonance is found, it allows us to use the PZT to actively stabilize the cavity, securing its resonance at the right frequency during the experiments.




Figure 7.1: Basic control loop for the stabilization of a cavity. PD: Photodetector, PI: controller, PZT: Piezo transducer

7.2 Stable supply

In the previous paragraph the stability demands for the cavity mirrors were deduced. Since the distance between these mirrors should be directly controlled with the PZT, the stability demands have also direct implications for the piezo driver. We will now look at the stability requirements for the piezo voltage supply.

Ab initio we don't have any knowledge about the relative positioning of the cavity's mirrors. We therefore need to be able to adjust one of the mirrors over at least one free spectral range guaranteeing to find a resonance when scanning the cavity. In Figure 6.3 we found that one FSR corresponds to approximately 110 Volts. Besides this we have to take into account the drift of the cavity which can amount to a few volts per hour. In total this sets the demands for the high voltage source to an available output range of minimally 120 volts, a requirement which itself is easily fulfilled and leaves open a whole range of possible high voltage drivers.

However, if we wish to stabilize the cavity to within a fraction of its linewidth, the output noise of the voltage source should be extremely low. This means that, if we set the desirable stability to one tenth of a the cavity linewidth, the maximum allowed amplitude of the output noise should not exceed $0.1 \times V_{FSR} / \mathcal{F} = 15 \mu\text{V}$. This corresponds to a signal-to-noise ratio of 10^7 which is a lot harder to accomplish. The low noise level itself is feasible in precision electronics, but the combination with the required output range makes it more difficult. We

found no commercially attractive solution that met our stability demands. Therefore a home-made design was proposed. This is not an easy task, yet it is technically possible to obtain such levels of stability ratios in electronic circuits using a PI controlled feedback system. Luckily we now don't have to worry about the generation of an error signal for the stabilization of the piezo driver: it is basically the difference between the desired output (setpoint) and the realized output signal, possibly contaminated with noise.




Figure 7.2: Feedback system for the stabilization and noise suppression of the high voltage piezo driver

As a basis we purchased a high voltage source (Thorlabs, MDT694) which is sold as a precision, low noise (1.5 mV rms) controller, especially suitable for piezo actuators. Its output voltage can be manually controlled with a potentiometer over a range of 150 Volt, in compliance with our required range. In the final feedback setup we wish to control the drive voltage of the high voltage supply externally which can be done by applying an analog voltage (from 0 to 10V) to its front panel BNC input. This external control voltage will be part of our feedback loop. The driver's input signal will then be the sum of the setpoint value supplied by a DAC output plus the value of the error signals. This total input voltage is then multiplied by a gain of 15 and summed with a possible manual control voltage. The driver's bandwidth is specified to 40 kHz.

According to the specifications the noise should be lower than 1.5 mV rms, which is two orders of magnitude larger than allowed in our setup. Before the piezo driver can be installed, the noise level has to be reduced with a factor of 100 using feedback control. For the design of the control system it is essential to know the nature of the noise. Therefore the piezo source's actual DC behavior is tested with an oscilloscope as load (capacitance of several pF) at a voltage of 75 volts. A typical oscilloscope image is shown in Figure 7.3.

The indicated rms noise level on the oscilloscope did indeed not exceed 2 mV and appeared independent of applied voltage. But it is clear that high amplitude noise is present in the form of sharp spikes. The most striking spikes have a frequency of 50 kHz and an amplitude of 10 to 15 mV due to the switching design of the piezo driver. With an additional low pass filter it is possible to suppress these frequencies to within our specifications. A passive second order filter with a cut-off frequency of a few hundred Hertz should be sufficient for this purpose. The proposed PI controller should take care of the lower frequency noise components, especially 50 Hz and its higher harmonics.

In order to examine the driver's gain as a function of frequency we applied a small test




Figure 7.3: Output signal of the high voltage source, containing 50 kHz spikes

signal to the external input. The test signal was generated using a function generator and thus consisted of a single frequency sinusoidal signal of known amplitude (several mV). Measuring the amplitude of the driver's output and phase shift relative to the input test signal, we could determine the driver's response, as shown in Figure 7.4. The profile shows the typical behavior for a second order system with its resonance frequency around 100 kHz.

7.3 Total scheme for cavity stabilization

As we have seen, we need at least two (nested) feedback systems, one for the cavity stabilization loop and one for the piezo driver's noise reduction loop.

Nesting two feedback systems is in generally a tricky thing to do and can be especially disadvantageous when the properties of the controllers diverge widely. For instance, in our case the bandwidth of the PZT amplifier loop is cut down to roughly 100 Hz by the additional second order filter which is significantly less than the bandwidth of several kHz that is available for the cavity stabilization. The piezo driver loop, now being the slowest link in the ring, will determine the bandwidth of the whole setup which is clearly not wanted.

The easiest way out of this problem is to add a third loop that is independent of the existing system and that takes over the fast feedback in the cavity stabilization. Combining and restructuring all three control loops necessary for the cavity stabilization, we obtain the total scheme as depicted in Figure 7.5.

Loop 1 and 2 essentially serve to stabilize the cavity at the resonance frequency. The first loop is a slow integrator-controlled loop that sets the piezo driver voltage such that the cavity is sufficiently close to resonance that the Pound-Drever-Hall error signal has the right sign.




Figure 7.4: Gain (solid line) and phase (dotted line) characteristics of the piezo high voltage supply

Therefore the voltage source needs to be adjustable with a precision of at least half the laser modulation which accounts to a 3 mV resolution over a 150 Volt range. Planning to set the loop's setpoint digitally with the aid of a low noise Digital-to-Analog Converter (DAC), this resolution implies that we need a 16 bit DAC. The DAC actually used is constructed by summing two 12 bit DAC's with a weighing factor of 1/32, creating an effective 17 bit DAC.

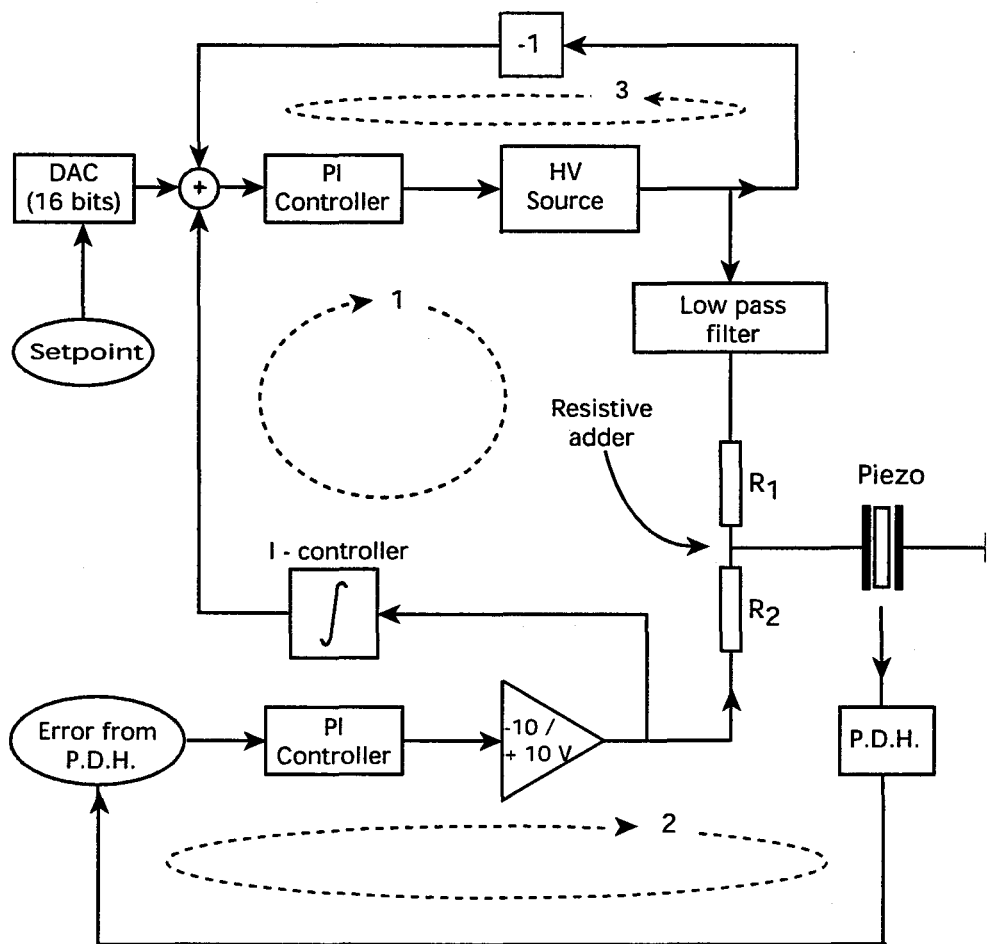


Figure 7.5: Cavity stabilization configuration including three feedback control loops. Loop 1 compensates for the slow cavity drift, loop 2 takes care of the fast cavity fluctuations, loop 3 reduces the noise of the high voltage piezo driver.

The second control loop takes care of the cavity's rapid fluctuations around resonance. To this end the error signal is processed and amplified in a PI controller. As this loop does not need to take care of the cavity's drift, its output range can be kept small (-10 to $+10$ V). Via a passive resistive adder this output is then added to the voltage of the first (slow) loop.

The third loop, finally, locks the piezo driver to the setpoint and guarantees the reduction of the driver's noise spectrum to the required microvolt levels as explained in the previous paragraph.

7.4 Pound-Drever-Hall with a high finesse cavity

As for the laser stabilization based on saturated absorption, we have to find an expression for the transfer function $F(\omega)$ for the cavity in order to understand the operation of the Pound-Drever-Hall stabilization scheme.

When a low-intensity monochromatic laser beam is sent into a (high-finesse) cavity, most of the incoming light will be reflected. The properties of the reflected beam can readily be inferred:

its frequency ω_{ref} is identical to that of the incident beam, and its amplitude E_{ref} is assumed proportional to the incident beam's amplitude E_{inc} . Upon reflection some phase shift may be introduced. Thus we can write: $E_{refl} = F(\omega)E_{inc}$ where the function $F(\omega)$ is the reflection coefficient which depends on the properties of both the beam and the cavity. This coefficient is a complex function, hereby taking into account the phase shift between the two beams. For a symmetric cavity (i.e., the reflection coefficients of the mirrors are equal) and no losses we found:

$$F(\omega) = \frac{r(e^{i\phi} - 1)}{1 - r^2 e^{i\phi}} \quad (7.1)$$

where r is the amplitude reflection coefficient of both mirrors, and ϕ is the phase the light has picked up after one round trip inside the cavity. This phase equals $\phi = 2\omega nd/c$ where d is the cavity length, $\omega/2\pi$ is the laser frequency and n is the refraction index which for a cavity placed in vacuum equals 1. It is convenient to rewrite this phase in terms of the free spectral range (FSR) and the wavelength:

$$\phi = \frac{\omega}{c/2nd} = \frac{\omega}{FSR} = 2\pi \frac{2nd}{\lambda} \quad (7.2)$$

Thus, the reflection coefficient is clearly periodic in both laser frequency and cavity length d . Since we are about to lock the cavity to the laser, we choose the latter view. We see that the coefficient F approaches zero when the optical path nd is an integer number of half wavelengths.




Figure 7.6: Amplitude of the error signal versus frequency for a symmetrical cavity.

For a cavity with negligible losses and identical mirrors, the error signal derived in Equation 3.25 can now be calculated. The result is shown in Figure 7.6, using realistic values for our system. Note that the error signal crosses zero when any of the beams (i.e., the carrier beam or either of the sidebands) resonates in the cavity.

7.5 PI controllers

A PI controller nearly always involves operational amplifiers that themselves are involved in a feedback loop. The open loop gain of an operational amplifier (opamp) is in general so high

that it can be considered perfect (i.e., infinitely high) for our closed-loop designs. A schematic view of the design of a simple opamp based PI controller is depicted in Figure 7.7.




Figure 7.7: Basic PI controller design.

A simple analysis of the design tells us that our PI controller can be considered as an inverting amplifier with voltage gain equal to $-Z_1/Z_2$ where Z_2 is the impedance of resistor R_2 and Z_1 the impedance of resistor R_1 and capacitor C_1 in series. The impedance of the components is frequency dependent. We write the gain of the controller in the frequency domain as the transfer function G_{PI} :

$$G_{PI} = -\frac{Z_1}{Z_2} = -\frac{R_1 + \frac{1}{j\omega C_1}}{R_2} \tag{7.3}$$

By choosing the resistances R_1 and R_2 and capacitance C_1 we control two important parameters. First, we need to set the proportional amplification coefficient G_p that determines the high frequency gain. Secondly the integrator’s time constant τ scales inversely with the bandwidth of the controller. The parameters are determined as follows:

$$\tau = R_1 C_1 \qquad G_p = \frac{R_1}{R_2} \tag{7.4}$$

The criterion for stability against oscillation for a feedback amplifier is that its open-loop phase shift should be kept less than 180° at the frequency at which the loop gain is unity (negative feedback becomes positive feedback at that frequency). We thus need to set the high frequency gain G_p to unity (or less) by choosing R_1 equal to (or smaller than) R_2 . Secondly, the bandwidth of both PI controllers is limited by the properties of the high voltage piezo driver and the cavity. Resonances in the cavity restrict the bandwidth to maximally 5 kHz. We choose values for the circuit components as indicated in following table.

Component	PI controller 1	PI controller 2
Capacitance C1	1 nF	10 nF
Resistor R1	10 kΩ	10 kΩ
Resistor R2	9.9 kΩ	9.4 kΩ

Both PI controllers have been built and were tested in a simple feedback setup, shown in Figure 7.8, involving three components. The output from the PI controller is subtracted from

a test signal by a differential amplifier. The result is used as input signal for the controller. The sinusoidal test signal is produced by a function generator and has known frequency and amplitude.




Figure 7.8: Feedback setup used to test the PI controllers. The controller's output is fed directly back to its input port.

The output of the controller should follow the test signal as it attempts to nullify the input signal. The closed loop gain for this system now becomes:

$$G_{closed} = \frac{G_{PI}}{1 + G_{PI}} = \frac{1 + j\omega R_1 C_1}{1 + j\omega C_1(R_1 + R_2)} \quad (7.5)$$

Measuring the ratio of output and input signal the frequency dependency of the controllers gain is determined as shown in Figure 7.9.

At low frequencies the closed-loop gain for both PI controllers is expected to be constant at a value of 1, which is exactly what we see. For the first PI controller the gain then drops for frequencies higher than 3 kHz, indicating a first time constant τ_1 at this frequency. PI controller number 2 clearly has a higher time constant τ_2 at a frequency of about 400 Hz.

For higher frequencies then we expect to find a frequency independent gain $G = 0.5$. While the second controller shows this behavior, even though on a limited range, it is hardly noticeable for the first controller. Clearly the system has a second time constant, appearing at a frequency of 70 kHz, that cuts off the gain. Replacing resistors and capacitors for different values didn't influence the frequency at which the second time constant appears. A closer look at the electronic circuit learns that the effect may be due to the presence of a light induced transistor in the amplification circuit. This transistor, with an estimated capacitance of $C \approx 1nF$, may short-circuit the amplification loop for high frequencies, resulting in zero amplification. Time constants in the operational amplifier itself might also attribute to the effect.

These measurements indicate that the built PI controllers display the expected behavior although some fine tuning is necessary to optimize the gain and bandwidth. However, the value for the open-loop gain cannot be inferred from these results. In order to be able to suppress the apparent noise levels at low frequencies in the high voltage piezo driver, this open-loop gain should be sufficiently large.




Figure 7.9: Gain of the PI controllers versus frequency. The dotted line shows the closed loop gain calculated from Eq.7.5

Chapter 8

Conclusions

The QND experiment concerns the direct demonstration of the existence of the quantization of the electromagnetic field. In the experiment the diffraction of a He^* beam interacting with electromagnetic field in a high-finesse cavity will be studied. Both the machine to produce the high-quality atom beam as well as the high-finesse cavity had already been constructed, but idle for a long period of time.

In this report we discussed the ‘resurrection’ of the beam and evaluated the problems that were encountered during the start. The beam machine appeared still functioning and, using Koolen’s reference manual, we were able to obtain his target figures at the last diagnostic tool before the 2D detector. Lack of reproducibility and difficulty of adjustment hindered further progress. The first problem appeared due to the quality of the suspension for the MOC quadrupole magnets. Therefore, a new, stable suspension has been designed. This design has already been implemented but not yet tested.

To solve the second difficulty, it will be necessary to implement extra diagnostics in the form of a second position sensitive detector just before the second pinhole.

As the QND experiment imposes stringent demands on the electromagnetic field, we looked at the stabilization of the diode laser and the high-finesse cavity based on the Pound-Drever-Hall scheme.

The laser system is at present operational. Through PI controlled feedback fluctuations in the laser are suppressed for frequencies to a least several hundred Hertz. The performance in terms of frequency stability has not yet been evaluated.

In order to be able to properly design a stabilization scheme for the cavity, we built a test facility to explore its electrical and mechanical behavior. Measurements using polarized light showed that our cavity is somewhat birefringent, meaning that only linearly polarized light can be excited in the cavity resonantly. As the QND experiment actually required circularly polarized light, we are currently investigating whether we can still produce circularly polarized light with reduced field strength, or that we have to remedy the birefringence.

Measuring the dynamical impedance of the of the cavity showed that resonances appear at frequencies of about 4 kHz, mainly due to mechanical resonances in the cavity construction.

Finally, we designed the stabilization scheme that should stabilize the cavity mirror spacing to an accuracy on the order of 1 pm. The scheme consists of three separate feedback loops. The necessary PI controllers were built and tested, and showed good response. In the total stabi-

lization setup, however, still too much noise was picked up by the feedback loops themselves to reach the required microvolt noise levels. The electronics are therefore currently redesigned more rigorously.

Appendix A

Technology Assessment

Quantum optics is the field of research that studies quantum mechanical properties of light fields. The key ingredient of quantum optics is given by the quantum of electromagnetic field energy, also known as the photon.

Although in modern (quantum) physics the concept of quantization enjoys a unassailable status, proving the existence of photons is far from obvious. The QND project now concerns the direct demonstration of the existence of the quantization of the electromagnetic field. Our approach is completely unique as we will attempt to directly study the momentum transfer from the light to the atoms through an atomic diffraction experiment. To observe the effects with the required precision, the setup requires state of the art technology from numerous disciplines. The project thus appeals to the deep fundamentals of both applied and theoretical physics.

The same sophisticated QND setup can and will be used for a variety of other high-precision experiments. We plan to make a Very Large Area Atom Interferometer using large-angle Bragg beam splitters [2], [7]. With little modification this Large Interferometer will be a practical tool for high-precision rotation sensing (Sagnac effect) or ultra-sensitive spectroscopy allowing the detection of very small shifts in atomic energy levels.

Bibliography

- [1] Knops, R., Quantum Optics with a Cold Helium Beam, Thesis TU/e (1998)
- [2] Koolen, A., Dissipative Atom Optics with Cold Metastable Helium Atoms, Thesis TU/e (2000)
- [3] Silfvast, W.T., Laser fundamentals, Cambridge University Press (1996)
- [4] Siegman, A.E., Lasers, Oxford University Press (1986)
- [5] Van de Vegte, J., Feedback control systems, Prentice-Hall (1994)
- [6] Friedland, B., Control system design : an introduction to state-space methods, McGraw-Hill (1986)
- [7] Jansen, G.T., Bragg diffraction of He^* by a standing light wave, Graduation report TU/e (2001)
- [8] Prevedelli, M., e.a., Opt. Commun. **125**, 231 (1996)
- [9] Drever, R.W.P., Hall, J.L., and Kowalski, F.V., Appl. Phys. B. **31**, 97 (1983).
- [10] Bachor, H.A., A guide to experiments in quantum optics, Wiley-VCH (1998)