NPL Report CLM 8

September 1999

Laser Linewidth at the sub-Hertz Level

M. Roberts

P. Taylor

P. Gill

Centre for Length Metrology
National Physical Laboratory
Teddington, Middlesex, TW11 0LW

Abstract

A number of fields require state of the art low noise laser systems. The report has been written to aid the development of a laser source with a sub-hertz linewidth. The report brings together the most recent work in the field of laser stabilisation, along with a description of the relevant basic principles. There are explanations of the parameters used to describe laser noise and stability (power spectral density, Allan variance, linewidth, modulation index). The effects of perturbation on a laser (plus reference cavity) system are discussed, including vibration, acoustic noise, thermal effects and opto-electrical considerations. The report concludes with a recommended laser system, drawing on the facts learnt during the process of writing this report.

ISSN 1369-6777

Authorised for publication on behalf of the Managing Director of NPL by
Dr David Robinson

Contents

Of course, the real test of the achieved performance can only be accomplished with a second, independent detector system. Disappointment is the experimenters' first reward for this measurement.

J. L. Hall
1986

Chapter 1

Introduction

The purpose of this document is to provide information to aid the development of a sub-hertz linewidth laser. It contains a brief review of the current state of the art, a tutorial exposition of laser noise theory and a discussion of various reference cavity considerations necessary to obtain a narrow linewidth laser. These considerations include environmental, optical and electronic issues, vibration isolation, beam pointing and power stability. In addition, a reasonably comprehensive list of references has been collected.

1.1 Assumptions

Some knowledge of the field of stabilising lasers to reference cavities is assumed, due to the necessarily short nature of this report. A good starting point for those unfamiliar with the field is the review article by Hamilton [1].

Improvement of the passive stability of a laser, though useful, can only reduce the laser linewidth so much. In order to progress further active stabilisation is required. A prerequisite for active stabilisation is a frequency discriminator. Either a molecular absorption or a reference cavity can be used. A reference cavity has two advantages, firstly a comb of resonances allowing access to anywhere in the optical spectrum. Additionally the signal to noise for the control signal can be increased, almost without limit, without power broadening the resonance.

Before this type of laser stabilisation can take place, the laser source must be running with a single spatial and temporal mode. It is also assumed that there are actuators of sufficient bandwidth to encompass the intrinsic noise of the laser. These actuators can either act on the laser cavity itself (piezo mounted mirror, intra-cavity Brewster plate....) or on the light outside the cavity (Acousto-optic modulator – AOM, Electro-optic modulator – EOM).

During the 1980's there were a number of technical developments which made the construction of a one hertz laser a possibility. One of the main problems with the use of a reference cavity is thermal length change. This was considerably reduced by the use of special spacer materials of low thermal expansion, such as ULE or Zerodur. Next the development of very high finesse mirrors ($\sim 10^5 - 10^6$), led to extremely narrow reference cavity linewidths. This in turn provides a very steep discriminator (volts/hertz) and makes tight servo control possible. Finally a detection scheme capable of exploiting the narrow cavity linewidth was developed. The Pound-Drever-Hall detection scheme [2], has clearly been shown as the most effective method for stabilising to a high finesse reference cavity.

Around this time it was also fully appreciated that achieving a tight servo lock to a cavity was only a necessary and not sufficient condition to achieving a narrow linewidth laser. Perturbations to the reference cavity itself, such as vibration, simply do not show up in the error signal of a tightly locked servo system. In order to assess the 'real' linewidth, a second, independent reference cavity is needed. The reduction of perturbations to the cavity has been the focus of much of the work over the past 15 years.

1.2 Omissions

A number of topics are not discussed in this report.

An area where there is little discussion, due to lack of time rather than importance, is the Pound-Drever-Hall technique [2] and the construction of appropriate servo electronics. The best explanation of how and why the Pound-Drever-Hall works is to be found in the original reference. Some general information about servo systems is contained in section 1.4. There are also two rather different stabilisation techniques which are not mentioned: the PTB self oscillating EOM [113] and optical feedback stabilisation of laser diodes (see for example [24]). The EOM scheme provides an alternative to electronic servos. The optical feedback scheme is only applicable to a particular laser source, and cannot on its own narrow the linewidth to an arbitrarily low level.

Another area not discussed, though certainly useful should you have a one hertz laser, is the topic of "optical time transfer". That is transferring narrow linewidth light from one place to another either in free space [37] or by optical fibre[14].

Lastly another area which is not discussed in any detail is cryogenic optical resonators [42, 41, 40, 65]. Here a laser is locked to a Fabry-Perot cavity made from a single crystal of sapphire. At cryogenic temperatures (2–3 K) sapphire exhibits a thermal expansion coefficient 1000 times smaller than room temperature spacers like ULE or Zerodur. This provides an optical cavity with excellent long-term stability. "Other sources of cavity length instability, such as mechanical deformations induced by seismic noise, support vibrations and tilt...will not be reduced. In fact, some of these are likely to be larger in initial experiments, due to mechanical noise generated within the cryostat"[40]. This report is primarily interested in solving problems like seismic noise, vibrations and tilt, which currently limit laser linewidth. Designing for low vibration and excellent beam pointing, within a cryostat introduces a number of specific problems, and as such is considered beyond the scope of this report.

1.3 Laser Noise Requirements

There are three main measures of the noise of a laser: Linewidth, Spectral Density and Allan variance. All three parameters are related to each other mathematically (see chapter 2). However these measurements are really used to show different aspects of the laser noise. Having a good Allan variance is *not* the same as having a narrow linewidth. For example in Schiller's work [42], he achieves an excellent long(ish) term stability of 2×10^{-15} @ 20 s, but a relatively poor linewidth of 420 Hz. Similarly a good spectral density is *not* the same as having a narrow linewidth. This report focuses on achieving a narrow linewidth.

Two major areas of research which require state of the art laser stabilisation are optical frequency standards and interferometric gravity wave detectors. Though these two fields have much in common, there is a subtle difference in the required laser noise. A laser in a gravity wave detector requires low noise spectral density across the intended detector frequency range (say 100–1000 Hz). An optical frequency standard requires a narrow linewidth (say 1 Hz), for which low frequency components of the noise spectral density are far more important that high frequency components. This has profound consequences when designing a stabilisation scheme for the laser. The development of an optical frequency standard is the ultimate motivation of this work.

As discussed in this report (see section 2.3), the linewidth of a laser source depends on the amplitude of the frequency noise and its Fourier frequency. It is the ratio of these two quantities, the modulation index, that determines the power in the laser spectrum away from the carrier frequency, and hence the linewidth. The consequence of this is that relatively large noises at high frequencies contribute less to linewidth than low amplitude noise at frequencies close to DC. It is because of this fact that ground vibration is considered as one of the greatest concerns for making a very narrow linewidth laser. To reduce this effect isolation platforms with very low resonant frequencies are required, such as a long pendulum.

This is to be contrasted with the requirements of gravitational wave detection. In this work ground vibration is also a critical issue, though here Fourier frequencies in the 100–1000 Hz range matter. At this frequency it is possible to get excellent vibration isolation by cascading multiple pendulums (see section 4.2.2). Multiple pendulums are not useful for narrow linewidth however, since they do not reduce the resonant frequency of the isolation system. Another even more dramatic difference is that some gravity wave workers [29] are using a reference cavity without a spacer. The spacerless cavity can have excellent high frequency performance by individually isolating each mirror, though its low frequency performance is rather poor making it unsuitable for attaining narrow linewidth.

These examples illustrate the need to be aware of the differing requirements of other experiments, and may mitigate against the wholesale transfer of ideas and technologies from one field to another.

1.4 Choice of Laser/ Servo Electronics

> The first question is whether, for high-stability experiments, one must use a laser with special properties, for example a massive invar construction. The answer both from theory and experiment is no, absolutely no. The only consequence of the choice of a fast evolving [phase error] source is that the control loop design and actuator must be of significantly greater bandwidth and lower time delay.

·J.L. Hall [13]

The picture painted by Hall, though true, perhaps understates the difficulty of making a high bandwidth servo. There are three main classes of laser in which we are interested: the monolithic Nd:YAG laser (linewidth 5 kHz), the Ti:Sapphire laser (linewidth 100 kHz) and the extended cavity diode laser (linewidth 300 kHz). The linewidths stated here are the typical free running rms linewidths observed on a few millisecond timescale. The range of Fourier frequencies over which these fluctuations occur is also important to consider. The Nd:YAG noise is less than 1 radian (see section 2.3) for Fourier frequencies above ~10 kHz, the Ti:sapphire from above ~50 kHz whilst significant noise at around ~1 MHz occurs for an extended cavity laser. Given a choice it is *much* easier to design a servo loop for the intrinsically quiet Nd:YAG laser, than for the rather noisy extended cavity diode laser.

> In the theory of servo control systems it is easy to show that the errors within the servo loop can be driven below any small value one wishes, merely by increasing the closed loop gain. Of course more gain can be usefully employed, considering loop stability issues, only if we can also increase the bandwidth and reduce the time delay if it approaches more than 1/6 of the inverse bandwidth†.
>
> —J.L. Hall [11]

A very clear exposition of control theory applied to a laser system is given in Day *et.al.* [63], and Day's thesis [75]. This analysis clearly shows that with sufficient gain both laser frequency noise and servo noise can be made negligible within the control bandwidth. Another example of control theory applied to a laser is given in [110]. Some useful detail of actual servo loops are given in the following papers: Dye laser [2, 4], YAG laser [63, 75] and diode laser [64]. A good text for detailed information on control theory is given in [112].

1.5 The State of the Art

Direct comparison of results in the literature is often not straight forward, as multiple types of noise measurements are used. The choice of noise measurement in a paper often reflects the goal of the work. Table 1.1 shows the best results achieved to date by various groups in a number of categories. In terms of laser linewidth the current state of the art is a 0.84 Hz beat from Jim Bergquist's group at NIST [119]. For more details of this work see section 3.5.

1.6 Report summary

The remaining chapters of this report cover several topics: laser noise and noise theory (chapter 2), studies of various groups laser systems (chapter 3) and reference

†If the time delay was equal to 1/4 of the inverse bandwidth an additional phase lag of 90 degrees [6] is added to the 90 degree lag of the integrating servo loop, and the servo oscillates. For a reasonable servo lock, the phase shift must be a reasonable margin less than 180 degrees, hence this factor of 1/6

cavity issues viz. attaining a narrow linewidth (chapter 4). The report's conclusion is in terms of a recommended system, based on the information presented elsewhere. Each chapter is self contained, and so the report can be read in any order, provided that the terminology of chapter 2 is understood. An attempt is made throughout to use consistent symbols. These are listed in a glossary on page 61.

Group (laser)	Spec. density $\sqrt{S_\nu(f)}$ [Hz/$\sqrt{\text{Hz}}$]		Beat Linewidth $\Delta\nu_l$ [Hz]		Allan variance $\sigma_y(\tau)$	
	(1L-1C)	(1L-2C)	(2L-1C)	(2L-2C)	(2L-1C)	(2L-2C)
Ohtsu (Nd:YAG)	$\sim$ **10^{-4}** [26], (1994)	>1 ($f<1$ kHz) **7×10^{-3}** ($f>1$ kHz), [20], (1995)	—	30 [36, 20] (1995)	—	2×10^{-14} @10ms [20],(1995)
Ueda (Nd:YAG)	$\sim$ **3×10^{-4}** [30],(1997)	—	—	16 [31],(1993)	—	6.7×10^{-14} @3ms [31], (1993)
Day/Byer (Nd:YAG)	—	—	0.33	~ 10 (*est*) [63], (1993)	—	1×10^{-14} (0.4-2s) [74],(1993)
NIST (Dye)	—	2.2 $f<15$ Hz [37], (1991)	<1 [37],(1991)	20 [37],(1991) 8 [38],(1998) **0.8** [in prep], (1999)	—	—
Hall (He-Ne)	—	—	**0.05** [5], (1988)	—	**8×10^{-17}** @200 s [13],(1989)	—
Schiller (Nd:YAG)	—	—	—	420 [42],(1997)	4×10^{-16} @ 50 s [43],(1997)	**2.3×10^{-15}** @20 s [42],(1997)
Walter (Dye)	$<10^{-2}$ $f<10$ kHz [81],(1992)	~ 5 [81],(1992)	—	10* [81],(1992)	$\sim 2\times10^{-17}$* @0.1s [81],(1992)	$\sim 10^{-14}$* (0-1s) [81],(1992)

Table 1.1: State of the Art in Stable Lasers

Numbers in **bold** indicate the state of the art for that measurement. L = laser, C = cavity, *e.g.* 2L-2C indicates that two lasers are locked to two cavities. * The results of Walter in the last three columns were actually observed with only one laser in each case. The Allan variance of 1L-2C was used to calculate the linewidth.

1.7 Reference List

The following list collects together relevant papers on constructing a narrow linewidth laser. It is organised both by author and subject.

- Introduction to stabilised lasers [1]
- Hall's papers [2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19]
- Ohtsu's papers [20, 21, 22, 23, 24, 25, 26, 27, 28]
- Ueda's papers [29, 30, 31, 32, 33, 34, 35]
- Bergquist [36, 37, 38]
- Schiller [39, 40, 41, 42, 43, 44]
- Gravity wave [45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60, 61]

Key Papers

- Hall [2, 5, 3]
- Bergquist [37]
- Kimble [62]
- Day/Byer [63]
- Clairon [64]
- Blair [65]

Diode stabilisation

- Clairon [64, 66, 67]
- Helmcke/Riehle [68, 69]
- Other groups [70, 71, 72, 69, 73]
- Ohtsu [21, 22, 23]

YAGS

- Ohtsu [20, 26, 27, 28]
- Ueda [32, 35, 31, 30, 34, 33]
- Byer [74, 63, 75]
- Others [76, 77, 65, 78]
- Schiller [39, 40, 41, 42, 43, 44]

Other Lasers

- Ti:Sapphire [62, 18]
- Dye laser [79, 80, 81]
- Argon [82]

Vibration issues

- Main text book for pendulum, vibration *etc.* [83]
- Other text books for pendulum, vibration *etc.* [84, 85]
- Simple engineering of structures: stress, strain, moment of area *etc.* [86]

- Seismic background [46, 49, 48, 61, 47]
- Active vibration isolation [48, 61, 87, 88, 89, 7]
- Negative stiffness [90, 55, 56, 59, 60]
- Multiple pendulums and Magnetic damping [51, 20]
- Coulomb damping [91, 83]
- Active damping [57]
- Internal resonances [85, 52]
- Cable and point suspension – effective mass [52]

Intensity fluctuations/Beam Pointing

- Suppressing beam jitter [53]
- Measurements of beam jitter [53, 54, 92]
- Problems with EOM's [58, 5]
- Intensity sevro [93]
- Angular vibration of cavity [35]
- Tilt [18]

Diagnostics Theory

- Window functions in FFT [94]
- NIST Tech Note [95]
- Getting from Power Spectral Density to Lineshape [96]
- Allan variance [97, 98, 99, 100, 25]
- Power spectral density [101, 25]
- Fundamental limit to length measurement [102]
- Generally useful Byer[63, 67]
- Modulation index [103]

Materials

- ULE [104]
- Zerodur [105]
- General materials [106, 107]
- Materials in gravity wave detectors [49]
- Rubber [108]

Random Topics

- Sound insulation [10, 109]
- Pound-Drever Hall Stabilisation [2, 110, 111, 4, 9, 12]
- Servo system theory [63, 75, 110, 112]
- Fibres [14, 17]
- Doppler cancelation [14, 17, 37]
- Telle's EOM locking [113]
- Mirror birefringence [114, 115]

Chapter 2

Noise

This chapter discusses various aspects of noise measurement, with regard to laser stabilisation. Its primary aim is to gain some physical insight into the connection between laser linewidth and the power spectral density - the main noise measurement technique. The modulation index approach provides the greatest physical insight into this relationship.

The chapter begins with a discussion of the various configurations that can be used to measure laser noise and attempts to make clear the differences between, and the limitations of, each configuration. The following sections cover the quantities that are used to characterise laser noise: linewidth, power spectral density, root power spectral density, modulation index and Allan variance. Some effort has been made to draw out the connections between all of these measurements. The shot-noise limit to laser linewidth is also presented.

2.1 Measurement techniques

There are two basic laser diagnostic techniques which are illustrated in figures 2.1 & 2.2: error signal analysis and beat note analysis. In error signal analysis there is one laser beam, in beat note analysis there are two laser beams. As well as the two diagnostic techniques, there are also two measurement types: either using one or two reference cavities. Though it is common to apply error signal analysis to one cavity and beat note analysis to two cavities, these are not the only measurement configurations.

To illustrate this, examine figure 2.1. For error signal analysis there are two configurations. Either the laser is locked to the reference cavity shown (*i.e.* using one cavity), or the laser is stabilized to a cavity (not shown) and the reference cavity in the figure is used as an optical spectrum analyser (*i.e.* using two cavities). In both configurations, one can either look directly at the error signal, or spectrum analyse it.

Now consider figure 2.2. Here two lasers beams are used to perform a beat note analysis†. The laser beams can be locked to two independent cavities, or

†In the discussion of beat note analysis we have described two laser sources. In practice these do not have to be derived from two separate lasers locked to two cavities. Instead a single laser output can be split into two components which are independently locked to the cavities using external

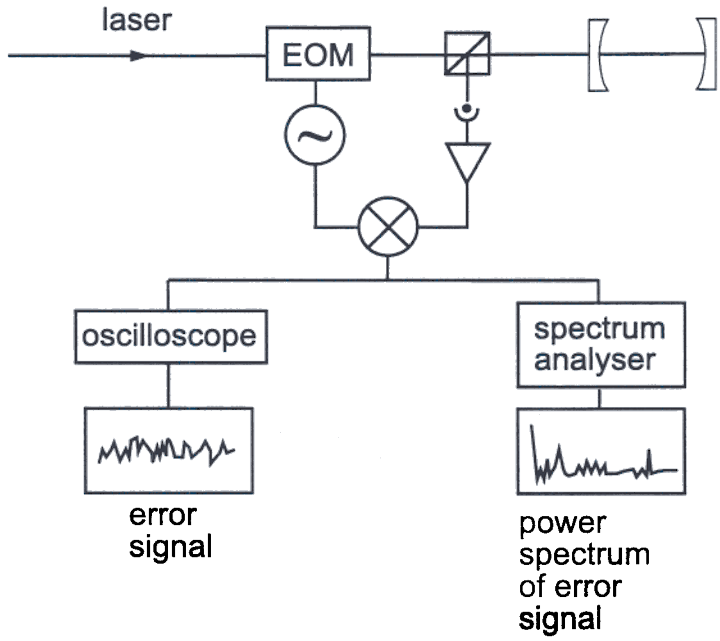

Figure 2.1: Error Signal Analysis

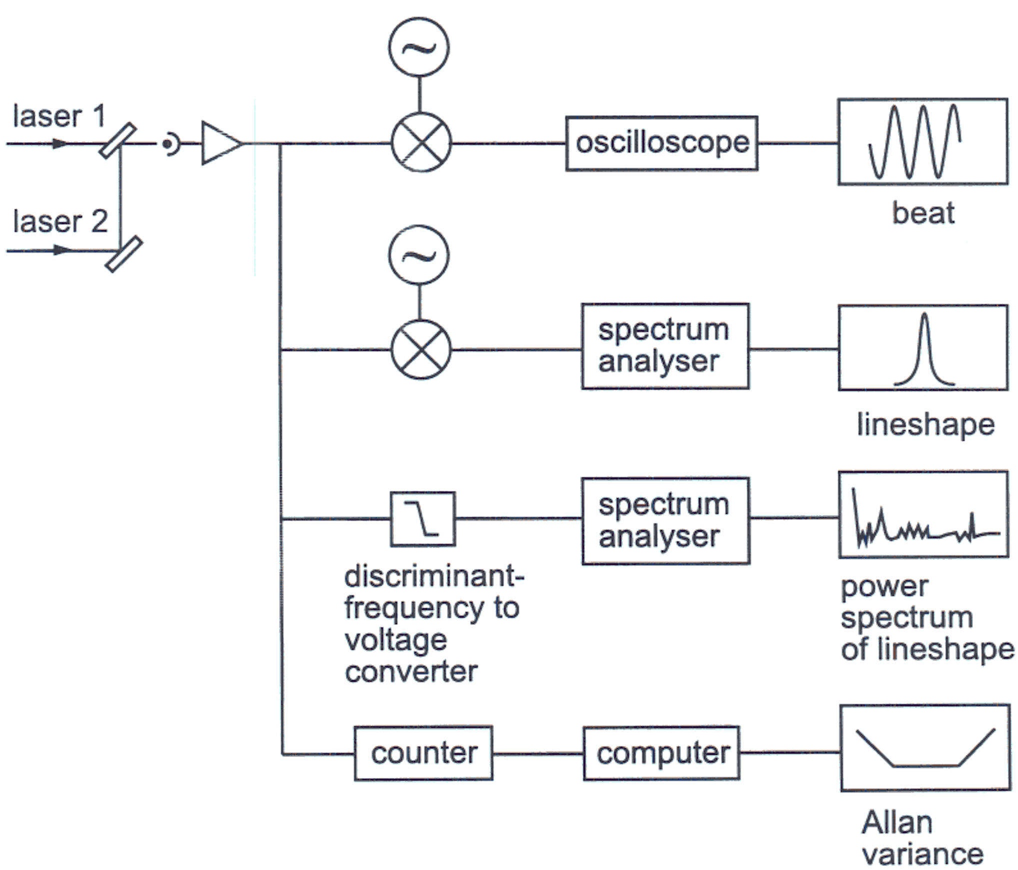

Figure 2.2: Beat Note Analysis

locked to a single cavity. In the latter case the lasers are usually locked to adjacent longitudinal modes so that the beat frequency is away from DC. The beat signal can be analysed in a number of ways: it can be downshifted to near DC and monitored on a storage oscilloscope; the beat lineshape and width can be measured by looking at it on a spectrum analyser; the Fourier frequencies of the noise making up the beat linewidth can be analysed using a frequency-voltage convertor (*e.g.* an rf spectrum analyser with the sweep turned off) and a spectrum analyser; finally the Allan variance between the two laser sources can be measured by downshifting the beat to a countable frequency and analysing it.

Error signal versus beat note analysis

So what is the difference in using error signal or beat note analysis? In beat note analysis it is easy to obtain all of the measures of laser noise (the purity of the beat sine wave, the laser linewidth and lineshape, spectral density and Allan variance). In error signal analysis only the spectral density is readily available, though the rest of the information can be obtained after mathematical manipulation of the data (see later sections). Beat note analysis has traditionally been used as the 'Rolls-Royce' diagnostic technique. The need for a second actuator and servo electronics does however carry an economic cost. If confidence in the mathematical extraction of the lineshape from error signal data were established, this would provide a significant time and effort saving.

One versus two reference cavities

Measurements using only *one* reference cavity (whether using error signal or beat note analysis) gives information about the quality of the electronic servo loop and other locking accuracy issues. It does not provide any information about 'real' laser linewidth since this is not only determined by the electronics, but also by what the reference cavity is doing. In order to make measurements of the real laser linewidth, a second independent reference cavity is necessary.

Infact to completely avoid common mode rejection it may be necessary go even further and mount the cavities on separate vibration isolation platforms, and use cavities of differing physical dimensions. This is to fully exclude the possibility of seismic common mode rejection.

> It is not sufficient for evaluating the frequency noise to use the same type of cavities. When a reference cavity and a monitoring cavity are of the same type and mounted on the same bench, the observed frequency noise is smaller than the absolute frequency noise because the noise from the outside gives almost the same influences on these two cavities and those kinds of noise are canceled and not observed. On the other hand, when the monitoring cavity is not identical to the reference cavity and

---

actuators (e.g. AOMs). Provided the servo loop has sufficient gain, it will make the laser light accurately track the reference cavity discriminant, and so effectively the reference cavity becomes the laser source. The laser itself is merely used to 'top up' the cavity with light.

mounted on a different bench, the observed frequency noise is close to the absolute frequency noise.

Musha *et.al.* –[29].

It is not completely clear how necessary such precautions are. For example would two cavities of different dimensions, but mounted on the same table, provided sufficiently independent measurements?

2.1.1 The four measurement configurations

To clarify these ideas, let's summarise the combinations of the techniques of error signal analysis and beat note analysis, along with the choice of one or two cavities.

Error signal analysis using one cavity

The error signal gives information about the quality of the servo loop (loop gain, loop bandwidth and loop stability). It is also useful to compare the free running laser noise, the locked laser noise and the detector noise. The detector noise is measured by taking the laser off resonance and recording the spectrum. It is important to check that the detector noise is less than the expected shot noise limit. This ensures that there is not excess electronic noise in the servo loop. Discriminator noise does not show up in this measurement. As pointed out in the introduction, it is possible to drive the error signal to as small a value as one wishes (even below the shot noise limit), given enough gain and servo bandwidth. Hence this diagnostic **cannot** tell you anything about the real laser linewidth.

EXAMPLES: Ohtsu-[20, 36, 26, 28]

Beat note analysis of one cavity

In this system two laser beams are independently locked to adjacent modes of a single cavity, using duplicated actuators and servo electronics (or using two completely separate laser systems). This measurement contains slightly more information than the previous case. Discriminator noise is now measured, and the DC locking accuracy can now be tested by locking the lasers to adjacent modes and then swopping them over (see [5]). This can reveal problems with modulation accuracy and optical feedback. This measurement **cannot** be used to give real laser linewidths. This is because the laser systems track the vibration of the cavity together and hence the cavity noise imposed on the laser is common to both sources. The observed noise will be suppressed by a common mode rejection factor [35, 63], which is equal to the longitudinal mode order of the cavity $N = \nu_0/\nu_{FSR} \sim 200\ 000$.

Error signal analysis of two cavities

In this system the laser is locked to one cavity and the second cavity is used as an optical spectrum analyser. This measurement configuration gives the complete information about the laser, *i.e.* a true power spectral density is measured. From this it is possible to calculate lineshape, linewidth and Allan variance, though this information can only be obtained after mathematical manipulation of the data.

EXAMPLES: Ohtsu-[20, 36]

Beat note analysis of two cavities

This set-up gives the complete information about the laser, in a readily accessible form - lineshape, linewidth, power spectral density and Allan variance simply read off the screen of spectrum analysers etc. It is also the most complex system, requiring two cavities, two actuators and two sets of servo electronics. To ensure there is no occurrence of common-mode rejection of the noise due to the vibrations of the two cavities, it may prove necessary to mount the two cavities on separate vibration isolation platforms.

EXAMPLES: Ohtsu -[20] NIST - [37].

2.2 Characterisation of Frequency Stability

Oscillators are characterised in the frequency domain by power spectral density and in the time domain by the Allan variance. The time and frequency characterisation of oscillators is an extensive subject, and only a brief overview is given here. NIST technical note 1337 [95] contains a collection of key papers in this field. The information in the following is based extensively on references [98, 100].

It is common to use dimensionless units when expressing stability. The dimensionless frequency fluctuation $y(t)$ is defined as:

$$y(t) \quad \frac{\nu(t) - \nu_0}{\nu_0}$$

where $\nu(t)$ is the instantaneous frequency and ν_0 is the nominal (average) frequency.

2.2.1 Power Spectral Density

Spectral density measures the spread of signal energy in the Fourier frequency spectrum. In the notation used here a quantity x, say, has a spectral density $S_x(f)$, which has dimensions of (dimension of $x)^2$/Hz. The per Hertz bit means the scale has been normalised to unit measurement bandwidth. $S_x(f)$ is implicitly an rms quantity.

From the spectral density it is possible to calculate the rms variation of the quantity using:

$$x_{rms}^2 \quad \int_0^\infty S_x(f)df$$

For some types of noise (i.e.white noise, $1/f$ noise) this integral diverges. In these situations it is often the case that there is an implicit bandwidth to the problem, which puts an upper and/or lower limit on the integral. For example, for white noise there is often an upper frequency limit or bandwidth, B, in which case $x_{rms}^2 = S_x B$. As another example, in modulation index analysis (see section 2.3), you require the amount of noise in a band around a frequency, f, say. To estimate this you could do the integral $x_{rms}^2 = \int_{0.5f}^{1.5f} S_x(f)df$.

To convert between the spectral density S_x, and the spectral density of the time derivative $S_{\dot{x}}$, use $S_x = (2\pi f)^2 S_{\dot{x}}$. This type of relation is used to get between phase and frequency, or acceleration and velocity.

In frequency standards the spectral density is generally used to characterise the short term behaviour. It is usual to work with the spectral density of frequency fluctuations $S_\nu(f)$ [$\mathrm{Hz}^2/\mathrm{Hz}$]. To add confusion, publications often quote the linear or root spectral density $\sqrt{S_\nu(f)}$, which is measured in $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$. When interpreting data in papers it is important to be aware of which version of spectral density the authors are using. In addition, the normalised spectral density $S_y(f) = S_\nu(f)/\nu_0^2$ and phase spectral density $S_\phi(f) = S_\nu(f)/f^2$ are occasionally used.

Spectral densities are measured with a spectrum analyser. It is necessary to convert the vertical scale, measured in dBm (*i.e.* power), into a frequency scale (Hz^2) via a calibration of the frequency discriminant. It is also necessary to divide the measured spectrum by the resolution bandwidth of the spectrum analyser, to then convert the scale into $\mathrm{Hz}^2/\mathrm{Hz}$. The power spectrum always shows a spike at DC, this is an artifact of the spectrum analyser.

To convert from the spectrum analyser output P [dBm] to noise spectral density $S_\nu(f)[\mathrm{Hz}^2/\mathrm{Hz}]$ frequency units we need to know the discriminator slope $D[\mathrm{Hz}/\mathrm{V}]$, the conversion from dBm to Volts and the spectrum analyser bandwidth $B[\mathrm{Hz}]$. First:

$$P[\mathrm{dBm}] = 10 \log \frac{V^2}{50\Omega \times 1\mathrm{mW}} \tag{2.3}$$

if the analyser is used with a 50 Ω input And

$$S_\nu(f) \quad \frac{\mathrm{noise}[\mathrm{Hz}^2]}{B} = \frac{V^2 \times D^2}{} \tag{2.4}$$

In the case of error signal analysis, the discriminant is simply the slope of the reference cavity. In this case the easiest way to calculate the discriminator slope D [Hz/V], is by measuring this slope. That is measure the peak to peak height of the Pound-Drever signal (in Volts) and divide this by the cavity linewidth. However, unless done carefully, this tends to be inaccurate due to the difficulty in making good finesse measurements.

A more accurate calibration of the spectrum analyser signal can be deduced by applying a calibrated modulation to the laser light, and observing the effect of this modulation on the error signal[75, 64]. The modulation frequency must either be outside the unity gain frequency of the servo loop or the servo must be operating with a very low gain lock, so that the servo does not reduce the size of the modulation spike. If the modulation has a frequency f_0, and a modulation index β (look at the transmitted light from the cavity) then the amplitude of the signal will be $f_0\beta$. Therefore the height of the spike on the spectrum analyser will be $(f_0\beta)^2$ [Hz^2]. (Again you must then divide by the resolution bandwidth of the analyser, to get the scale in $\mathrm{Hz}^2/\mathrm{Hz}$.) Finally when connecting a spectrum analyser to servo electronics "it is important to use a directional coupler or other appropriate isolator so that ac coupling at the input of the spectrum analyser will not introduce measurement error"[8].

Spectral densities of high stability frequency standards can often be approximated by a power law:

$$S_y(f) = \sum_{\alpha=-2}^{+2} h_\alpha f^\alpha$$

for $0 \leq f \leq f_h$, where f_h is an upper cutoff frequency. The names given to the different types of noise are illustrated in figure 2.3.

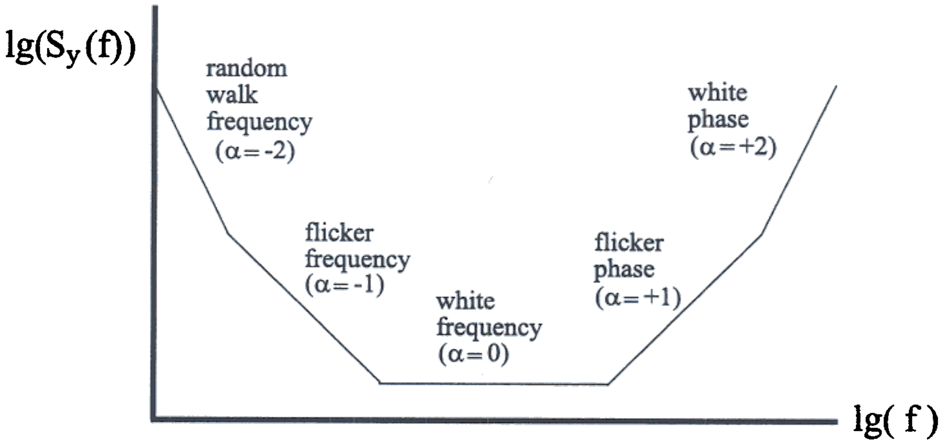

Figure 2.3: Spectral density of frequency fluctuations

2.2.2 Allan variance

The Allan variance is used to assess the stability of a frequency standard over a time interval τ. The Allan variance was introduced because it was found that the classical variance diverged for some types of noise commonly encountered in frequency standards.

The basic tool for measurement of the Allan variance is a frequency counter. A series of m measurements of the (dimensionless) frequency $\bar{y}_i$ are made, each of which is averaged over a time τ. The Allan variance is estimated [100] using the formula:

$$\sigma_y^2(\tau) \quad \frac{1}{2(m-1)} \sum_{i=1}^{m-1} (\bar{y}_{i+1} \quad \bar{y}_i)^2$$

The Allan variance can be modeled by a power law, in a similar way to the spectral density, *i.e.*

$$\sigma_y^2(\tau) \sim \tau^\mu$$

It is only for the case of white frequency noise ($\alpha = 0$) that the Allan variance is equal to the classical variance [98]. Changes in the average frequency over long times do not bias the characterisation of short term frequency stability [100]. The Allan variance can cope with linear drift (no tractable model exists for the spectral density). If the subject drifts as $y(t) = dt$, then the root Allan variance is $\sigma_y(\tau) = d\tau/\sqrt{2}$ [100]. The Allan variance is not useful for distinguishing between white phase noise ($\alpha = +2$) and flicker phase noise ($\alpha = +1$), as the slope is almost identical for both cases. This limitation can be overcome by the use of the modified Allan variance [98, 100].

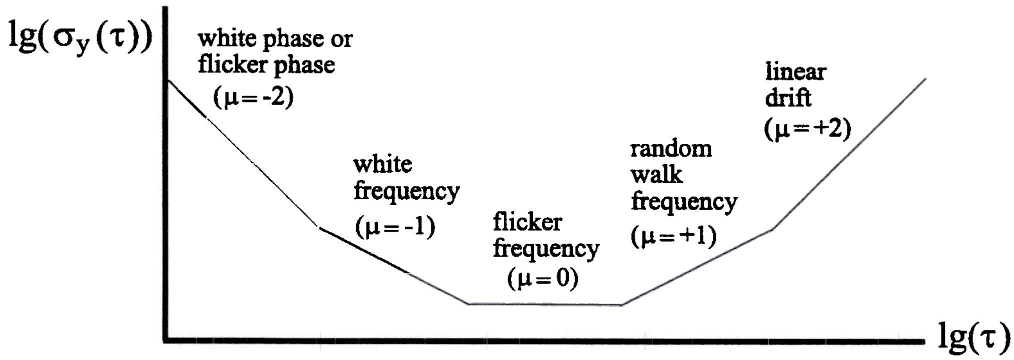

Figure 2.4: The Allan Variance

Allan 'variance' pictures are infact almost always plots of the Allan deviation $\sigma_y(\tau)$, as is the case here. In this picture though the exponent μ shown is the exponent for the Allan variance, equation 2.7.

2.2.3 Relating Spectral Density to Allan variance

The spectral density contains the maximum information about the process. The Allan variance is related [100] to spectral density via:

$$\sigma_y^2(\tau) \quad 2\int_0^\infty S_y(f)\frac{\sin^4(\pi f\tau)}{(\pi f\tau)^2}df \tag{2.8}$$

Some information is lost in the conversion process[100], and so it is not in general possible to convert back from Allan variance to spectral density.

However, when the noise present can be approximated by a power law, a two way relationship does exists between spectral density and Allan variance. The spectral density exponent α and Allan variance exponent μ are related by [98]:

$$\begin{array}{llll} \mu & +2 & & \alpha \quad \text{undefined} \\ \mu & -\alpha & 1 & -3<\alpha\leq 1 \\ \mu & -2 & & \alpha\geq 1 \end{array} \tag{2.9}$$

The relationship between the coefficients of the power laws is shown in table 2.1. Taking the case for white frequency noise, the Allan variance $\sigma_y^2(\tau)$ has a slope of -1 and a value $a_0 = h_0/2$ at $\tau = 1$ second. Comparing this to the equation for spectral density (equation 2.5) shows that the normalised spectral density $S_y = 2a_0$ in this case. Using the relationship to convert into ordinary spectral density gives $S_\nu = 2\nu_0^2 a_0$.

2.2.4 Relating Spectral Density to Linewidth

Elliot *et al* [96] derive a key relationship between spectral density and laser linewidth. They consider a laser with a white noise spectral density S_ν up to a frequency or bandwidth B, and no noise above this frequency. If the rms frequency fluctuations

Name of noise	α	$S_y(f)$	$\sigma_y^2(\tau)$	a_α
White phase	2	$h_2 f^2$	$a_2\tau^{-2}$	don't know
Flicker phase	1	$h_1 f^1$	$a_1\tau^{-2}$	$(1.038 + 3\ln(2\pi f_h\tau)h_1)/(2\pi)^2$
White frequency	0	$h_0 f^0$	$a_0\tau^{-1}$	$h_0/2$
Flicker frequency	-1	$h_{-1}f^{-1}$	$a_{-1}\tau^0$	$2\ln(2)h_{-1}$
Random walk frequency	-2	$h_{-2}f^{-2}$	$a_{-2}\tau$	$(2\pi)^2 h_{-2}/6$

Table 2.1: Relationship between Allan variance and spectral density
Assuming the power law approximation. f_h - measurement system bandwidth. a_α in the last column shows the coefficient for the Allan variance in column 4. Adapted from reference [98]

$\delta\nu_{rms} \ll B^\dagger$, then laser lineshape is Lorentzian with linewidth:

$$\Delta\nu_l = \frac{\pi(\delta\nu_{rms})^2}{B} \quad \pi S_\nu \tag{2.10}$$

For a white noise spectral density, but with $\delta\nu_{rms} \gg B$ (ie $\beta \gg 1$), a Gaussian lineshape is predicted, with linewidth

$$\Delta\nu_l = 2.35\,\delta\nu_{rms} \tag{2.11}$$

These formula are used to get an idea of the laser linewidth from the spectral density. If the spectral density has been derived from two independent systems (for example a laser locked to one cavity, and spectrum analysed by a second cavity) and the noise is white, then the formula will give an accurate value for the linewidth.

In practice at low frequencies the spectral density tends to show flicker frequency noise, invalidating the formula. A crude attempt to allow for a non-white noise is to just integrate the spectral density to calculate the rms frequency fluctuation ($(\delta\nu_{rms})^2 = \int_0^\infty S_\nu(f)df$), and then just use the formula. However the proper method of dealing with this problem is to do a full numerical integration of the rather complicated formula in Elliots paper [96].

The big difficulty with these formula is that they apply only to white noise. White noise is often predominates above a certain frequency, but as already mentioned, at low frequency $1/f$ noise dominates. This problem is compounded since, as will be explained in the next section, low frequency components of noise have much more influence on linewidth than high frequency components. A theory for specrtal density and lineshape/linewidth when $1/f$ noise is present would be extremely useful.

2.3 Frequency Modulation Analysis

The frequency modulation (FM) analysis of a physical situation is a very powerful tool, as it quantifies the effect of noise processes at different Fourier frequencies.

---

†This condition is equivalent to a modulation index $\beta \ll 1$ (see section 2.3). To show this first find $(\delta\nu_{rms})^2 = \int_0^B S_\nu df = S_\nu B$. Next substituting this relation into the inequality gives $\sqrt{S_\nu/B} \ll 1$. Finally compare to equation 2.20 with $f = B$ gives $\beta \ll 1$.

This section develops FM analysis in a general way which can be applied to many problems. The treatment here follows reference [103], p579, though similar treatments can be found in many other texts. The application of this technique in laser stabilisation, involves finding a relationship between the modulation index and the spectral density; the subject of the following section. To show the versatility of the technique a selection of other examples are also given.

Consider a general oscillation:

$$E(t) \quad E_0 \exp i[\phi(t)] \tag{2.12}$$

where the phase $\phi(t) = \int \omega \, dt$. If the oscillation is modulated at a Fourier frequency f [Hz] with an amplitude $2\pi\delta\nu$ [radHertz]†, then the instantaneous angular frequency ω_i is:

$$\omega_i = \omega_0 + (2\pi\delta\nu)\cos(2\pi f t) \tag{2.13}$$

Hence:

$$E(t) \quad E_0 \exp i[\omega_0 t + \frac{\delta\nu}{f}\sin(2\pi f t)] \tag{2.14}$$

The ratio $\delta\nu/f$ is called the modulation index β, and has units of radians. We will see later that this is a most useful relation, and is worth writing out again:

$$\beta = \frac{\delta\nu}{f} = \frac{\text{Modulation Amplitude}}{\text{Fourier frequency}} \tag{2.15}$$

From equation 2.14, use a Bessel function identity, take the Fourier transform and then square to get then intensity of the wave as a function of angular frequency:

$$I(\omega) \propto \sum_{n=-\infty}^{\infty} |J_n(\beta)|^2 \, \delta(\omega \quad [\omega_0 \quad 2\pi n f]) \tag{2.16}$$

Thus the effect of the modulation is to impose n sidebands on the carrier to form a comb of peaks at frequencies $\omega_0 \pm 2\pi n f$. The intensity of these sidebands is given by $|J_n(\beta)|^2$.

For small modulation index, β, only the first sideband is important and its intensity I_1 relative to the carrier I_0 is approximately:

$$\frac{I_1}{I_0} = \left(\frac{\beta}{2}\right)^2 \qquad \text{for} \qquad \beta \lesssim 1 \tag{2.17}$$

For large modulation index many sidebands occur, spaced symmetrically around the carrier. The number of sidebands that contribute to the sum is roughly:

$$\beta \lesssim n \lesssim \beta \qquad \text{for} \qquad \beta \gg 1 \tag{2.18}$$

These results can be easily transferred to a number of problems. All that is required, is a calculation of the modulation index for the particular problem. The spectrum

†note $\delta\nu = \sqrt{2}\delta\nu_{rms}$

is then given exactly by equation 2.16. In practice either of the approximations, eqn. 2.17 or 2.18, can by used to get a good idea of what's going on.

The analysis presented here is exact for a sinusoidal modulation. In many practical problems though the modulation may be made up of a broadband noise represented by a spectral density. In this case to get $\delta\nu$ for equation 2.15, integration of the spectral density over a small band, around the frequency of interest f, gives $(\delta\nu)^2 = 2(\delta\nu_{rms})^2 = 2\int_{0.5f}^{1.5f} S_\nu(f)df$. If this procedure is followed, the modulation analysis is now approximate and can only be used to give order of magnitude estimates.

2.3.1 Modulation index, Power Spectral Density and Linewidth

It is possible to get an idea of the laser linewidth from FM analysis. The spectral density measures the amount of noise away from the carrier frequency. At a particular frequency, f say, this spectral density will cause a modulation index β. The laser linewidth is going to be roughly the frequency where this modulation index is big enough to put significant power into the first sideband of the carrier, *i.e.* when $\beta \approx 1$. In terms of equations, combine eqn 2.17 and eqn. 2.15, to give:

$$\text{sideband power} \approx \frac{(\delta\nu)^2}{f^2} = \frac{(\text{FM amplitude})^2}{(\text{Fourier frequency})^2} \tag{2.19}$$

This result has a profound result on designing a narrow linewidth laser. It is now clear that low frequency perturbations contribute far more significantly to laser linewidth, because of the $1/f^2$ factor. It thus follows that seismic noise on the reference cavity at 0–10 Hz is going to contribute much more to the laser linewidth than, say, acoustic noise at 1 kHz.

J.L. Hall is the biggest exponent of FM theory applied to laser linewidth, and perhaps explains the ideas best of all:

> If the FM has a small amplitude and occurs at a high rate.....the resulting small phase modulation index $\beta \ll 1$ puts nearly no power into the sidebands. The laser linewidth will be approximately equal to the Fourier frequency at which unity phase-modulation first occurs.
>
> —Hall [5]

Essentially the modulation index is a measure of phase excursion from the ideal sinusoid. As mentioned earlier modulation index has units of radians. Hence Hall's frequent statements about phase excursions of less than one radian. For example:

> ...we can see that small frequency excursions are not too serious if they occur at a high frequency because there will not be enough time to integrate a radian phase error before the fluctuation has reversed sign and begins to integrate the phase error back to zero. When we reach ~1 radian phase error the phase modulation sidebands do become prominent (19%), so it is reasonable to take $\beta = 1$ as the division point between phase coherence and incoherence.
>
> —Hall [13]

It is useful to make the connection between modulation index and power spectral density. From equation 2.15, $\beta = \sqrt{2}\delta\nu_{rms}/f$. Now the rms frequency fluctuations due to white noise in a small band around a frequency, f say, is: $(\delta\nu_{rms})^2 = \int_{0.5f}^{1.5f} S_\nu df = S_\nu f$. Therefore the modulation index in terms of spectral density is

$$\beta \quad \sqrt{\frac{2S_\nu}{f}} \tag{2.20}$$

Using equation 2.20 you can draw the line $\beta = 1$ on a plot of spectral density. Noise features below this line will contribute little to the resulting laser linewidth (ie the sideband power $\ll 1$). Noise features at or above this line will have an important effect on the laser linewidth. To get an estimate of laser linewidth it is reasonable to take the frequency when the noise curve crosses the $\beta = 1$ line.

2.3.2 Modulation Index - other examples

The Lamb-Dicke regime

An ion held in an rf (Paul) trap experiences a harmonic potential well with a natural frequency $2\pi f$ and performs simple harmonic motion of amplitude x_a within the well. This motion modulates the absorption or emission spectrum of the ion such that:

$$E(t) = E_0 \exp i[\omega_0 t + kx_a \sin(2\pi ft)] \tag{2.21}$$

where $k = 2\pi/\lambda$ and λ is the wavelength of the absorbed or emitted light. Comparing to equation 2.14 we see that the modulation index of the sidebands is $\beta = kx_a$ and hence the carrier is essentially unmodulated for $\beta < 1$ or $x_a < \lambda/2\pi$. If this condition is fulfilled then the first order Doppler effect has been eliminated. The requirement, that the amplitude of the ion motion is less than the wavelength of the incident light, is known as the Lamb-Dicke regime.

Zeeman effect on a single ion spectrum

Consider an atom subject to an AC magnetic field $B_{mag}(f)$. What quantitative effect does this field have on the atom absorption and emission spectrum?

The sensitivity of a transition to magnetic field is given by $\delta\nu = g_J\mu_B B_{mag} m_J/h$. A magnetic field variation at frequency f therefore gives a modulation index of:

$$\beta - \frac{g_J\mu_B B_{mag} m_J}{hf} \tag{2.22}$$

As a numerical example consider AC magnetic field variations at the 'mains' frequency of $f = 50$ Hz. A typical value measured in our lab gives $B_{mag} = 0.2\ \mu$T. Taking $g_J = 2$ and $m_J = 1/2$ gives $\beta = 56$! Using equation 2.18 this means that the transition is now made up of over 100 sidebands with an end-to-end spread of $2\beta f = 5.6$ kHz. Magnetic shielding would have to reduce the field to 4 nT to get $\beta < 1$ at 50 Hz.

Contrast that with a transition free from the linear Zeeman effect such as the 411 nm transition in $^{171}Yb^+$. Here $\delta\nu = 0.3\ \mathrm{Hz}/(\mu T)^2$ and a 50 Hz field of 0.2 μT at a DC field of around 2 μT gives $\beta = 1/500$, *i.e.* completely negligible.

Inertial frames and the Doppler effect

Consider two vibration isolation platforms designed to isolate a structure from all ground vibration, for example two tables suspended from a ceiling using weak springs. If the isolation is perfect, these two surfaces become inertial frames and a laser beam passed from one to the other experiences no Doppler effect as the relative motion between the tables is zero. However in practice the isolation platform will have a resonant frequency below which no isolation occurs. Oscillation of the platform at the resonant frequency will modulate the laser light via the Doppler effect. How large a problem is this?

The Doppler shift to the laser frequency is given by:

$$\delta\nu = \frac{\nu_0 v}{c} \tag{2.23}$$

where v [m/s] is the relative velocity of the two tables and c is the speed of light. For the light not to be greatly modulated, we require:

$$\beta = \frac{\nu_0 v}{cf} < 1 \tag{2.24}$$

Hence we need to estimate a likely value for v. The spectral density for floor vibration velocity in our labs is given by equation 4.11 on page 37. To make an estimate, approximate $(v_{rms})^2 = \int_{0.5f}^{1.5f} S_v(f)df$. For a typical oscillation frequency of $f = 1$ Hz this gives $v_{rms} = 7 \times 10^{-6}$ ms^{-1}. Hence the modulation index is $\beta = 7$. This means that Doppler shifts due to the swinging motion of the two vibration isolation platforms relative to each is a problem. To counter this effect and avoid unnecessary linewidth broadening of the laser when transferring it from table to table, some kind of Doppler cancellation technique is necessary [14, 17, 37].

2.4 Shot Noise Limit

It is relatively straight forward to calculate the expected root power spectral density, $\sqrt{S_\nu(f)}$ for the shot noise limit of a side of fringe lock (see for example [5]). The shot noise photo-current on the detector is $\delta I_{rms} = \sqrt{2eIB}$. This is converted into frequency noise by the reference cavity discriminant $\delta\nu_{rms} = (\partial f/\partial I) \times \delta I = \Delta\nu_c \delta I/I$, where $\Delta\nu_c$ is the reference cavity linewidth. This relationship can be written in terms of optical power, P, instead of photocurrent using $I/e = P\eta_d/h\nu$, where η_d is the quantum efficiency of the detector. Now for white noise, the spectral density is related to the frequency noise by $\sqrt{S_\nu(f)} = \delta\nu_{rms}/\sqrt{B}$, therefore:

$$\sqrt{S_\nu(f)} = \Delta\nu_c \sqrt{\frac{2h\nu}{P\eta_d}} \tag{2.25}$$

The root power spectral density for the Pound-Drever-Hall locking scheme, derived in reference [63], is:

$$\sqrt{S_\nu(f)} = \Delta\nu_c \sqrt{\frac{h\nu}{8J_0^2(M)P\eta_d}} \sqrt{1 + \left(\frac{2f}{\Delta\nu_c}\right)^2} \tag{2.26}$$

The equation has the same form as the side of fringe lock. The second term now contains a factor $J_0(M)$ to do with the rf modulation scheme. A value of $M = 1.08$ for the rf modulation index optimises the discriminator slope. For this value the Pound-Drever-Hall scheme leads to an improvement of $\sqrt{8.5}$ in spectral density over the side of fringe lock. The third term represents the roll off of the cavity response at high frequencies (this effect is ignored in the analysis of the side of fringe lock).

A more complex formula for the Pound-Drever-Hall scheme, which allows for the effect of absorption in the cavity mirrors, is quoted in reference [26]. This formula allows for the fact that in real life the fringe contrast is not 100%. It is probably worth using when you really want to know exactly what the shot noise limit is. It is however rather complicated, and so is not repeated here.

Shot noise limited linewidth

Consider a realistic example, using equation 2.26, and the cavity parameters of section 4.1. Assuming $\eta_d = 0.9$, $M = 1 \Rightarrow J_0^2(M) = (0.765)^2$ even a small power of only 10 μW gives excellent spectral density.

$$P \quad 10\ \mu\mathrm{W} \implies \sqrt{S_\nu} = 1 \times 10^{-3}\ \mathrm{Hz}/\sqrt{\mathrm{Hz}} \tag{2.27}$$

Using equation 2.10, this gives a tiny linewidth of only:

$$\Delta\nu_l \quad \pi S_\nu = 3\ \mu\mathrm{Hz} \tag{2.28}$$

This illustrates that even relatively low incident powers, on a high finesse cavity, should give extremely narrow linewidths. There is no need for milliwatts of incident power, which is just as well as high intracavity power can cause large degradations in laser linewidth (see section 4.3.1). It is also possible to calculate the expected Allan variance. Remembering that shot-noise is white, this will give an Allan deviation with a slope of $1/\sqrt{\tau}$. The coefficient, worked out at the end of section 2.2.3, is $a_0 = S_\nu/2\nu_0^2$, giving:

$$\sigma_y(\tau) = \sqrt{\frac{S_\nu}{2\nu_0^2\tau}} = \frac{2 \times 10^{-18}}{\sqrt{\tau}} \tag{2.29}$$

Again this is unlikely to form any real limit to the performance of the system.

What these example really show, is that shot noise should never be a limiting problem in achieving a narrow linewidth. The real difficulty is controlling the perturbations to the reference cavity.

Schawlow-Townes limit

The Schawlow-Townes limit to laser linewidth is not relevant to the linewidth obtained from a laser locked to a cavity.

The Schawlow-Townes limit refers to the minimum linewidth of a laser oscillator. It is (see for example [63])

$$\sqrt{S_\nu(f)} = \Delta\nu_{lc}\sqrt{\frac{2h\nu}{P}} \tag{2.30}$$

where $\Delta\nu_{lc}$ is the laser cavity linewidth, and P the laser output power.

For a laser under tight servo control to a high finesse cavity, it is more proper to think of the cavity as the laser source. The formula above has the same form as the shot noise limit of previous sections. However as already seen, when locking to a cavity, a number of factors enter the equation to do with the particular locking scheme. Several publications (for example [33]) demonstrate noise below the Schawlow-Townes limit.

2.5 Conclusion

A lot of theory has been presented in the chapter, and this is an attempt to reiterate the most important results.

- Spectral Density section 2.2.1. This is the key noise measurement for making a narrow linewidth laser.
- Modulation index, Spectral Density and Linewidth section 2.3.1. This information of this section allows one to interpret the noise on spectral density data, and assess which noise components are contributing to the linewidth
- Relating Spectral Density to Linewidth, section 2.2.4. This section discusses the key mathematical results of Elliot *et.al.* [96], which provides a firm mathematical relationship between spectral density and linewidth.

The example at the end of the shot noise section 2.4 illustrates how to calculate linewidths and Allan variances, when the spectral density has white frequency noise. Unfortunately in practice stabilised lasers always seem to have $1/f$ noise, which is not possible to handle by these simple analytic formulas. In this case either the approximate modulation index method can be use to get a rough idea of what's going on, or a full numerical integration of the complicated formula in Elliots paper [96] could be performed.

Chapter 3

Case Studies

The aim of this chapter is to analyse the experimental outputs of a few of the leading groups in laser stabilisation. It is interesting to compare and contrast the various approaches used. These are dictated by the different research goals, and hence act as an illustation of the care needed when interpreting another group's results for inclusion in one's own work. For the sake of clarity, the use of a single researcher's name is used to indicate the work of the whole research group.

3.1 Hall

Hall's interest in laser stabilisation has concentrated on the attainment of narrow linewidth, good locking accuracy and high stability. This work is therefore very relevant to the achieving a laser linewidth at the hertz level. Hall does not seem to have any ulterior motive for his work – he is motivated primarily by the stabilization of the laser itself. Not only was he a co-inventor of the so-called Pound-Drever-Hall laser stabilisation technique [2], but his work on accurate servo control has yet to be bettered [5].

The development of the Pound-Drever-Hall stabilisation technique [2] in 1983 ushered in the modern era of laser stabilisation. In this first paper, a dye laser and He-Ne laser were stabilised to a single reference cavity and showed a sub-100 Hz beat linewidth. Already in this work many of the important issues necessary for attaining a narrow linewidth were identified (vibration, optical feedback, etalon effects, beam pointing and amplitude noise on the laser light.)

Following this [3], two completely *independent* cavity stabilised dye lasers were compared and showed a linewidth of 750 Hz. This linewidth was limited not by the very tight lock to the reference cavity, but by the reference cavity itself. This was one of the first demonstrations that making a laser with a linewidth of less than 1 Hz was not solely an electronic servo problem.

Between 1984 and 1988 Hall produced a number of stabilisation related papers [15, 6, 11] which either discussed general stabilisation methods or the latest results but with few details. This work was brought together in 1988 in the paper "Laser stabilization at the millihertz level" [5]. The paper summarises the relevance of the choice of laser source and the importance of the Fourier frequency of the noise on laser linewidth. Also included is a discussion of the effects of discriminator noise and the importance of feedback isolation and modulation accuracy. To put the ideas into practice, two He-Ne lasers were locked to adjacent longitudinal modes of a single cavity. A locking accuracy of $< 2.5 \times 10^{-5}$ linewidths was achieved, an amazing feat that has not yet been matched. A corresponding linewidth of 50 mHz for each source was measured, which has not been beaten either. The relevance of locking accuracy can be seen from the results of locking two lasers to one cavity. The record Allan variance that Hall achieved (8×10^{-17}) was only possible once all (slow) lock-point shifting effects had been removed. Hence these effects are important for any system needing stability in the 1–10 000 second region. This *is* the case for a narrow linewidth laser, as measurement times approaching a minute are needed if beat notes are to be recorded with high resolution.

Hall's subsequent papers to 1994 reiterate the same lessons [13, 7, 8, 19, 18]. An exception to this is the work Hall did on active vibration control. An brief analysis of the expected frequency shift for a cavity in a gravitational field of $1g$, shows that 'micro-g' control of the cavity's acceleration is needed [19]. He attempted this by actively controlling the air to a commercial air suspension table (see section 4.2.2). Although he achieved impressive results, the fact that not all six-degrees of freedom were controlled meant that vibrational noise still crept into the system. He measured a sensitivity to tilt of 10 Hz/μrad, which ultimately prevented him (possibly along with other effects) from achieving a sub-hertz linewidth laser.

The most recent papers we have from Hall in this field discusses transferring the narrow linewidth of a laser from one place to another [14, 17]. In this work the optical phase noise that occurs due to fibre delivery is discussed along with a technique to remove it at the millihertz level. This important paper is beyond the scope of this report. The technique is analogous, if not identical, to the Doppler cancellation technique used (though not discussed) by NIST [37].

3.2 Ohtsu

Ohtsu's interest in high quality laser stabilisation, and in particular that of Nd:YAG lasers, lies in the field of gravitational wave detection. Ohtsu has concentrated on producing a low power master laser with a noise floor close to the design specification of 10^{-6} Hz/$\sqrt{\text{Hz}}$ @ 1 kHz. He has, however, also made measurements of absolute linewidths that are of more relevance to our goal.

Work on stabilised Nd:YAG lasers started in 1992 [28]. Here a single laser was stabilised to one reference cavity. A 30 kHz feedback loop to the piezo was implemented with a 3 stage lag-lead servo. The cavity linewidth was measured using the decay of the cavity beat method [28]. An error signal analysis of this one cavity system showed a noise floor of $\sim1.5\times10^{-3}$ Hz/$\sqrt{\text{Hz}}$ for Fourier frequencies in the 1 Hz to 1 kHz range.

By 1994 [26] he had reduced this noise level on the error signal to $<1\times10^{-4}$ Hz/$\sqrt{\text{Hz}}$ for 10 Hz to 10 kHz which was the shot noise limit. The reduction in noise was obtained by using a high bandwidth servo actuator (an external electro-optic modulator) giving a combined unity gain servo frequency of 600 kHz. This gave over 100 dB of gain at 1 kHz.

Once a good level of laser stabilisation to the cavity had been achieved, he went on to look at the absolute stability of the laser, by analysing the light with a second cavity [27]. Using this method he showed that although below 10 kHz his error signal noise was around 1×10^{-4} Hz/$\sqrt{\text{Hz}}$ the noise as observed by the second cavity was well over 1 Hz/$\sqrt{\text{Hz}}$ for f <10 Hz . To get a good measure of a beat linewidth he then used two lasers locked to the separate cavities. A beat of under 30 Hz was observed with a 50 Hz/second drift. He ascribes this drift to a lack of temperature control or evacuation of the vacuum chamber. However, high light levels were used in the cavity and could have caused some of the drift. The reference cavities were suspended in the vacuum chambers by 5 wires. There is no mention of any additional vibration isolation. Measurements of the power spectrum of the relative intensity noise showed that high modulation frequencies (>20 MHz) are needed to achieve the shot-noise-limited sensitivity he requires. This is because laser intensity noise extends out to these high frequencies, and because his shot noise requirements are so stringent.

The final paper on this work is from 1995 [20]. The main change from previous work is the use of a double pendulum cavity suspension inside the vacuum chamber. Eddy-current (velocity) damping was achieved with 1 Tesla Nd-Be-Fe permanent magnets. The advantage (and limitations) of a double pendulum suspension are explained in section 4.2.2. It is interesting to note that no vertical isolation was used except for the optical table on air-legs. Figure 3 in this paper has a rather beautiful comparison of the spectral density for various levels of isolation. With the two cavities not suspended, floating the air table reduces the noise by about 10 dB but only in the 10–200 Hz frequency range. (Which is consistent with our accelerometer measurements, figure 4.1, page 38). Hanging the cavities via a single suspension stage reduces the spectral density by a further 10 dB right from 1 Hz to 1 kHz. Adding a double pendulum gives around half an order of magnitude further improvement over this same range. The final result with the double pendulum is a spectral density of 7×10^{-3} Hz/$\sqrt{\text{Hz}}$ @ 1 kHz. Measurements of a beat linewidth with 2 lasers locked to the 2 separate cavities showed that the linewidth had not been reduced significantly (still 30-50 Hz). This was thought to be due to residual drift of the cavities because of heating the cavity mirrors, but again the vacuum chamber was still not temperature controlled in these measurements.

3.3 Ueda

Ken-ichi Ueda of the Institute for Laser Science (ILS) has studied the stabilisation of Nd:YAG lasers to high finesse reference cavities with a view to building gravitational wave detector. Not only has he worked on building a quiet master oscillator, but also on high power slave lasers. The latter will be required to reach the desired shot-noise limited sensitivity needed for detection of gravitational waves.

The first paper we have from this group is in 1992 with the stabilisation of two lasers to a single reference cavity [32]. A beat linewidth of 193 mHz was achieved. From the error signal of each laser a lock noise of $\sim 1\times10^{-2}$ Hz/$\sqrt{\text{Hz}}$ was measured. His other paper of 1992 reported the same results along with design considerations for high power Nd:YAGs [33].

Ueda went on to look at the system of two lasers and two independent reference cavities [31]. A beat linewidth of 16 Hz was reported between the two systems which were not suspended. The only vibration isolation that appears to be present is the optical table which is not described. At this stage the error signal at the cavity was reduced to the shot noise of 3.7×10^{-4} Hz/$\sqrt{\text{Hz}}$ using 4 mW of light incident on the cavity.

Ueda realised that the 16 Hz beat was likely to be due to the cavity, and investigates this in more detail in his 1994 paper [35]. In this paper, short term drift of a beat between two systems was ascribed to heating of the mirror coatings of the cavity because the drift rate decreased the longer the laser had been locked to the cavity. Long term drifts were seen due to the cavity not being temperature controlled. An attempt to quantify the effect of vibrations to the cavity was made by placing it on a platform that could be vibrated with piezo actuators. A linear sensitivity of 3.4 (mHz/$\sqrt{\text{Hz}}$)/nrad (for a modulation frequency of 5 Hz) was observed although there is no link with this number and the observed linewidth of 16 Hz. (cf Hall measured a sensitivity of 10 Hz/μrad for his whole system [18]).

The goal of Udea's research is to build a gravity wave antenna with a sensitivity of around 10^{-7} Hz/$\sqrt{\text{Hz}}$ at Fourier frequencies of 100-1000 Hz. To achieve this shot noise requires relatively high incident power on the cavity. His paper in 1994 discussed the accurate measurement of absorption in mirrors [34] and on his visit to NPL he discussed trying to manufacture mirrors in Japan with world beating low loss.

Our most recent paper from Ueda's group is from 1997 [30]. Here steps are taken to actually produce a high power (2 W) and low noise source. To reduce the cavity noise at around 1 kHz a double-suspension system was used. In order to reduce the error signal noise to these very low levels, an extra-cavity EOM was needed to give the necessary bandwidth of 1 MHz. In his visit to NPL in 1998, Ueda indicated that he would no longer be working in this field, although obviously the gravitational wave antenna work is still on going. The final results of this work are very comparable to that of Ohtsu. From the acknowledgements in the latest Ueda's papers it appears that a collaboration may have taken place, where Ohtsu's lasers were used in Ueda's work.

3.4 Byer

The contents of Day's thesis and the paper he co-authored with Byer describe the stabilisation of an Nd:YAG laser to a single high finesse (order 20 000) cavity. Initial work analysed the closed loop error signal to assess the system performance which was limited by detection noise problems to around 2 Hz. This was verified by using two lasers locked to one cavity and monitoring the beat. This linewidth was reduced to 330 mHz by increasing the laser power into the cavity to 40 mW from 2 mW, thus increasing the signal-to-noise of the discriminator. The observed beat linewidth was limited by variations in the cavity used. This was proved by reducing the discriminator linewidth from 200 kHz to 50 kHz (by increasing the cavity length) with no observed improvement in system performance [75, 63]. The very high input power levels to the cavity were offset by the fact the cavity had a relatively low finesse. The fact that the Allan variance of this work is flat for time scales between

0.4 and 2 seconds [74] indicates that the mirror heating effects were not as severe as for Ohtsu [20] and Ueda [31]. Additional improvement came also from including an AOM in addition to an optical isolator for feedback isolation.

A two stage integration servo was used, giving 12 dB/octave at low frequencies reducing to 6 dB/octave at unity gain ($\sim$ 60-100 kHz) by a lead stage (high pass filter). A second lead stage was incorporated when using the low linewidth cavity to correct for the roll off of response due to the cavity decay time. Gains of 120 dB at low frequencies were obtained.

His work also includes a careful assessment of the free running performance of a Nd:YAG laser at various Fourier frequencies, showing it to be Shawlow-Townes limited above 100 kHz.

The performance of the system described seemed to be due to using a poor reference cavity, giving reduced stability performance even for two lasers locked to the same cavity. This work does not address absolute stability issues, such as vibration or heating of the cavity mirrors.

3.5 NIST

The work from NIST on narrow linewidth lasers has led the field on and off for several years. However a relatively small number of papers have been published [76, 37, 38]. The goal of the NIST group, like ourselves, is the development of a laser with sub-hertz absolute linewidth. Their aim is the construction of an optical frequency standard based on a single mercury ion [37]. The original published work [37] will be discussed first, followed by more recent results presented at 1998's CPEM conference [38] and in a private communication [119].

Set-up circa 1994

The following is based solely on reference [37]. The system is based on a frequency doubled dye laser. The system is composed of one laser and two cavities with the laser output locked to each of them. The dye laser is pre-stabilised to below 1 kHz using a finesse 800 scanning cavity and a 2 MHz PZT+EOM servo bandwidth. The output is then split into two parts and sent via fibre optics to the two cavities. The light is post-stabilised at the cavities via acousto-optics, so fibre induced noise is not a concern. In early work both cavities were made of Zerodur and are around 300 mm long giving a free-spectral range of 500 MHz - which is useful as any necessary frequency offsets can be achieved via AOMs. Each cavity was suspended by short wires (natural frequency 1.4 Hz) inside the vacuum chamber. The chamber itself was sat on a Vee-block via Viton strips. A doubled passed AOM is used with a 100 kHz servo bandwidth to get a tight lock, the slow drift of the error signal going to the laser cavity. The light inside the cavity is controlled by monitoring the input power and feeding back to the AOM. Acoustic enclosures surround the vacuum chambers to reduce acoustic perturbations.

At the time of publication one cavity was fixed rigidly to a table sat in a sandbox and the other suspended from the ceiling via 3 m long latex tubes. The suspended table gives good vertical isolation compared to the rigidly mounted table. The height

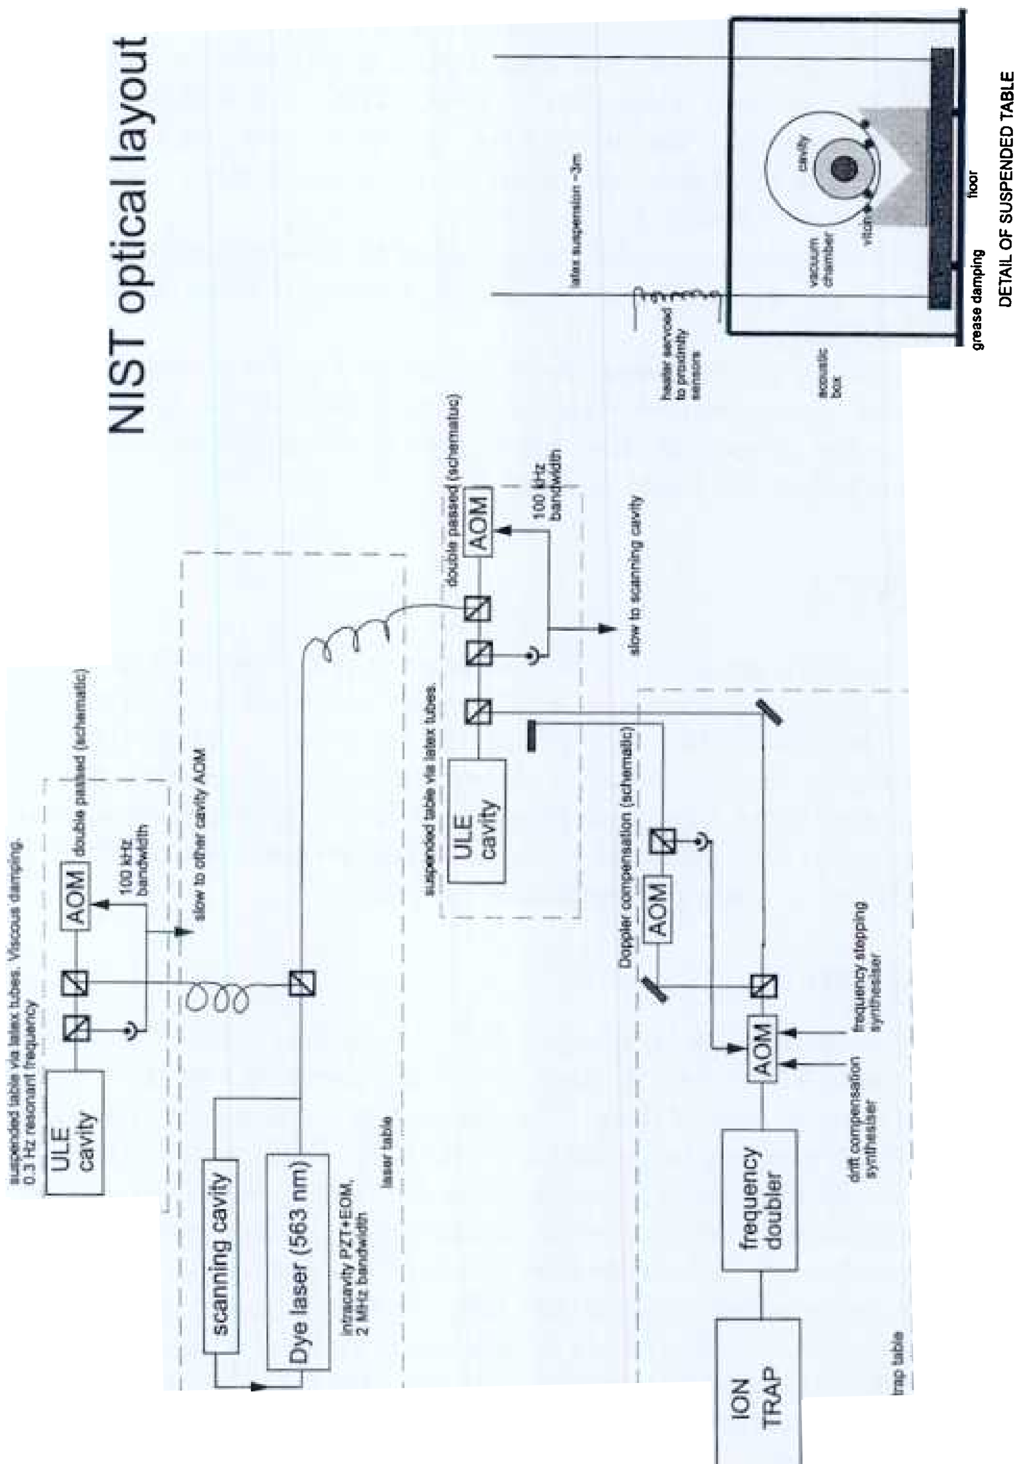

Figure 3.1: NIST Optical Layout

of the table is slowly servoed via a heater wrapped around the latex, monitored via proximity sensors. The problems of inadequate active control observed by Hall [19, 18] are not seen here, since the height and tilt are servoed only slowly on the ~100 second timescale. The 0.3 Hz resonance of the pendulum is damped by 2 dabs of grease.

The light is sent to the ion trap from the sandbox table in free space. Another AOM is used to step the laser frequency via a synthesiser and remove cavity drift via a swept synthesiser. The light is then frequency doubled and sent to the ion trap. Doppler shifts between the trap and cavity tables are removed. A portion of the beam is picked off near the ion trap, offset by an AOM and sent over to the cavity and back again. A homodyne beat is detected and half the frequency sent to the drift compensation AOM (see also the work of Hall [17, 14]).

A 30 Hz beat was achieved which was thought to be limited by the sandbox table mounted cavity. The paper suggests that vibration is the dominant effect on the cavity, floor resonances being seen in the beat. The effects of vertical vibration was shown by driving a noise via a loudspeaker on the table. Horizontal motion was small due to the wire suspension, but vertical motion was large due to the lack of a spring. This lead them to mount one cavity on the latex platform.

Latest results

A number of modifications to the system have been made, giving a sub-hertz linewidth beat [38, 119]. The most major of these is to suspend the sand-box mounted table on latex supports in an similar way to their other cavity. This reduced the beat to 8 Hz [38]. The next major improvement was to replace the Zerodur cavities with ULE. This reduced the beat to 0.84 Hz, measured with a 40 second integration time. This change also reduced the drift rate of the beat to 0.1 Hz/sec from 1.2 Hz/sec. The reason for the change to ULE was that it had been observed with the Zerodur cavity that the beat was narrow, but jumped around on the second timescale.

A power dependent frequencies shift of 1 Hz/μW has been measured and a low(ish) incident power of 100 μW is used on the cavity. The *output* power from the cavity is controlled to 0.1%, which suggests laser power noise contributes about 0.1 Hz to the observed beat linewidth. Inside the vacuum chamber the cavities now sit on 4 Viton pads. This is to remove any beam pointing issues which were problematic when the cavities were suspended by thin wires in the vacuum chamber. The Viton pads did not reduce the thermal isolation of the cavity from its environment. Both cavities are of an American football shape, the idea being to increase the resonant frequency of bending modes of the bar, although it is not obvious whether this played a major role in achieving the reduced linewidth.

It is not known what limits the laser linewidth at this (low) level. An understanding may be forthcomming when this latest work is published.

Chapter 4

Reference Cavity Issues

This chapter discusses the various perturbing effects to a reference cavity. An attempt is made to quantify these effects, where possible, to allow the design of a laser system capable of exhibiting a linewidth in the hertz region.

The chapter begins with effects of temperature and pressure changes to the cavity resonance frequency, including a discussion of the choice of cavity materials.

The next section discusses the subject of vibration. For some time the importance of vibration on the reference cavity has been appreciated as an effect limiting laser linewidth, at about the 100–1000 Hz level. This section includes: characterising the magnitude of the vibration problem and a discussion of vibration reduction techniques.

The third section introduces the less well known limitations to laser linewidth. These are a collection of related effects to do with intra-cavity power stability and beam pointing. The changes of light power within the cavity can effect its length. These changes can come from either amplitude noise on the laser, or beam pointing issues. These issues have an important bearing on exactly how the vibration isolation system is to work. In trying to achieve a narrow linewidth many groups have attempted to isolate only the laser cavity from vibration (for example by hanging it from wires within the vacuum chamber). This unfortunately introduces rather poor beam pointing stability, and consequent intra-cavity power variability due to slight variations in mode-matching efficiency. It is believed that these effects limit the linewidth at the 10–100 Hz level if they are not dealt with. To achieve stable intra-cavity power and good beam pointing stability to the cavity, it is much better to rigidly fix the laser and cavity relative to each other. This is perhaps only the first step, and it is probably also necessary to actively stabilise the intra-cavity power. It would also be wise to use as little laser power as possible to stabilise the laser (bearing in mind the shot noise limit). Fixing the laser and cavity relative to each other does however have ramifications to the vibration isolation system. To achieve vibration isolation for the cavity, it is now necessary to isolate the entire system. †

†Isolation of the entire system is perhaps an overstatement, as it should be possible to use optical fibres to connect the laser to the isolation platform. However, optical fibres introduce frequency noise at the kilohertz level [14]. This problem with optical fibres should not however limit the system, as the fibres would only be used to transfer the 'unstabilised' light around. There are two possibilities. If the frequency actuators (say AOM's) are on the isolation system, the noise introduced by the fibre is just servoed away with the rest of the intrinsic laser noise. If the

The final section of this chapter discusses other issues which effect the overall laser system. These are: frequency modulation considerations, optical feedback and the birefrigence of the cavity mirrors.

Environmental considerations

cavity

The work contained in this report concerns the locking a laser to a Fabry-Perot etalon. The etalon has length L, free-spectral range ν_{FSR}, having equal lossless mirror reflectivities R, linewidth $\Delta\nu_c$ and finesse $\mathcal{F}$. These parameters are related by the following relationships:

$$\nu_{FSR} = \frac{c}{2L} \tag{4.1}$$

$$\mathcal{F} = \frac{\pi\sqrt{R}}{1-R} \approx \frac{\pi}{1-R} \tag{4.2}$$

$$\Delta\nu_c = \frac{\nu_{FSR}}{\mathcal{F}} \tag{4.3}$$

The choice of material used to make the spacer of the cavity is critical. It is essential that the material has a low coefficient of thermal expansion, so that temperature fluctuations do not effect the laser frequency. For a room temperature system there are two possible materials, ULE or Zerodur. ULE is the preferred choice (see the NIST case study). The cavity is contained within a vacuum chamber to eliminate pressure effects, and equally importantly to help with temperature control. The outside of the chamber is temperature controlled to 1–10 mK.

Another material with very low thermal expansion is sapphire when used at 2–3 K. This type of cryogenic system is a good method to get excellent medium to long term stability. However as mentioned in the introduction it is likely that the difficulties introduced by the cryogenic environment will make it very difficult to achieve a narrow linewidth from such a system.

Length changes in the cavity cause frequency changes of the laser locked to the cavity according to the formula:

$$\Delta\nu = \frac{2\nu_{FSR}\Delta L}{\lambda} = \frac{\Delta L \nu_0}{L} \tag{4.4}$$

In the examples given in this chapter a cylindrical ULE cavity of length $L = 0.1$ m and diameter 50 mm is assumed. The mirrors are assumed to be 5 mm thick and 10 mm radius, with a finesse of $\mathcal{F}$=100 000. The wavelength is $\lambda = 1\ \mu$m, the free spectral range is $\nu_{FSR} = 1500$ MHz.

---

frequency actuator (say piezo mirror) is on the laser, the fibre still poses no problem, as the noise source is within the servo loop and is again controlled. In both cases however, it is important to realise that the narrow linewidth light exists only after both the fibre and the actuator (ie at the cavity); any pick off for an experiment must be made here.

ULE vs Zerodur

Zerodur is a mixed phase glass-ceramic made by Schott[105]. ULE is a true glass formed of SiO_2 and TiO_2, and is made by Corning[104]. Both materials have very low thermal expansion coefficients of about $\alpha \sim 10^{-8}\mathrm{K}^{-1}$.

Since both materials are glassy rather than crystalline solids, they both exhibit the phenomenon of "creep" – a slow long term change in dimensions. Zerodur has a typical creep rate [13] of $\Delta L/L = 2 - 4 \times 10^{-15}$ s^{-1}. ULE has a smaller creep rate [119, 13] of $\Delta L/L = 0.2 - 0.5 \times 10^{-15}$ s^{-1}. An important issue here is the predictability of the drift rate, as this puts an upper limit on the locking cycle time in a frequency standard. Hall [13] estimates a frequency predictability of roughly 100 better than the total drift in a particular time period.

Once machined, the spacer is acid etched. This (unpublished) process, performed by REO, allegedly relieves micro-stresses in the material's surface.

The real advantage of ULE lies in the nature of the creep. ULE seems to creep in a smooth fashion, whereas Zerodur makes small random jumps in length every few seconds [119]. This random jumping of the Zerodur limits the achievable linewidth at about 10 Hz. This discovery by the NIST group has allowed them to achieve a real laser linewidth of less than 0.8 Hz. This is the narrowest laser linewidth to date. It is for this reason that we are using ULE in our work.

ULE glass properties [104]

- Youngs Modulus $E = 67$ GPa
- Poisson Ratio = 0.17
- Density = 2200 kg m^{-3}
- Coefficient of thermal expansion (specification) $\alpha = 0.00 \pm 0.03 \times 10^{-6}$ K^{-1} (5° to 35°C)
- Coefficient of thermal expansion (for calculations) $\alpha = 1 \times 10^{-8}$ K^{-1}

Temperature control

Temperature control of the cavity is an important issue, and considerable care needs to be taken to achieve a good result.

Temperature affects the length of the cavity according to the formula

$$\Delta L \quad \alpha L \Delta T$$

Therefore for the cavity mentioned above, a temperature variation of $\Delta T = 1$ mK gives a frequency shift of $\Delta\nu = 3$ kHz. This is a very serious shift if one wants to obtain a laser linewidth of ~ 1 Hz.

Fortunately the effects of temperature changes to the cavity are moderated slightly. The cavity is contained in a vacuum chamber, and so has excellent insulation from the outside world. This means that the time constant for a temperature change can be very long, say 24 hours [37], which provides a natural smoothing of temperature fluctuations. The suppression of the amplitude of the temperature fluctuations is roughly:

$$\frac{\Delta T_{in}}{\Delta T_{out}} \approx \frac{1}{f\,\tau}$$

where T_{in} and T_{out} are the inside and outside temperatures, τ is the time constant and f is the frequency of the fluctuations. Thus for a time constant of $\tau = 24$ hours, a temperature fluctuation on the 10 second timescale is suppressed by a factor of 10 000. What is still very important is the long term temperature stability, since for a slow temperature variations, say between day and night, there is hardly any time constant suppression. In practice one does see a component that follows the temperature change more quickly [19], which may be due to radiative rather than conductive heating.

Say a temperature variation of 1 mK occurs over a 12 hour period, and ignoring any time constant suppression, a frequency drift of 0.07 Hz s^{-1} occurs. This size of (random) drift should not be too difficult to cope with. The ideal aim here is to reduce this random temperature drift, to a small fraction of the cavity creep rate. As already mentioned this creep rate is something like 0.15 Hz s^{-1}, though it is predictable to something like 0.0015 Hz s^{-1}.

Pressure Fluctuations

Provided that the cavity is under some sort of acceptable vacuum, pressure fluctuations are not nearly as serious as temperature fluctuations. Pressure effects the frequency of the laser locked to the cavity as follows: A change in the optical length of the cavity nL, occurs because of changes in the refractive index, n, of the gas inside the cavity (ie $d(nL)/dp = L\ dn/dp$). The change of refractive index with pressure is [37] roughly

$$\frac{dn}{dp} = 3 \times 10^{-9}\ \mathrm{Pa}^{-1} \tag{4.7}$$

Therefore, combining this stuff with equation 4.4 gives

$$\Delta\nu = (3 \times 10^{-9}\ \mathrm{Pa}^{-1})\nu_0 \Delta p \tag{4.8}$$

Changes in pressure can occur because of a change in temperature (assume it is an ideal gas). For a pressure of $p = 10^{-5}$ mbar $= 10^{-3}$ Pa, at room temperature $T = 293$ K, a temperature fluctuation of ΔT=1 mK implies a pressure fluctuation of $\Delta p = 3 \times 10^{-9}$ Pa. The frequency shift due to this pressure fluctuation is only $\Delta\nu$=3 mHz. (cf the 3 kHz shift caused by the same temperature change directly effecting the length).

The second (fundamental) way in which change in pressure occurs is just a statistical number fluctuation. So if in a particular volume there are N gas atoms, the rms fluctuation in the number is $\sqrt{N}$ which will give a fluctuation in the pressure. Assuming that the laser beam is 2 mm diameter, then the volume of gas is approximately 3×10^{-7} m^3. Assuming room temperature, and a pressure of $p = 10^{-3}$ Pa, this gives a pressure fluctuation of $\Delta p = 3 \times 10^{-9}$ Pa, ie again $\Delta\nu$ =3 mHz.

4.2 Vibration

Low frequency components of the spectral density have a bigger influence on laser linewidth than high frequency components. It is this fact that leads to concern about the effect of seismic noise on the reference cavity, and consequent attempts at vibration isolation. It is reasonable to expect that ground vibration in the range say 0.1 Hz to 100 Hz are most important in achieving a narrow linewidth laser.

4.2.1 Seismic noise

> The spectrum of seismic noise is very dependent on site location and time. Below 50 mHz, the main source of noise is atmospheric pressure fluctuations caused by wind turbulence and infrasonic waves. From 50 mHz to 500 mHz, the main source of noise is ocean waves generating surface and body waves in the earth's crust at coastal areas, with most of the energy coming from coastlines in the vicinity of large storms. In this frequency range, there are two dominant features, one at roughly 60-90 mHz, corresponding to the fundamental ocean wave frequency. Another is at twice the fundamental (120-160 mHz), coming from the interaction of the reflected and incoming ocean waves, known as the microseismic peak. Above 500 mHz, along with oceanic and coastal body waves, the main sources of noise are local, such as wind blown vegetation and human activity.
>
> —D.B. Newall *et.al.* [61]

Above 0.5 Hz the spectral density for acceleration can be approximated by whit(ish) noise. Typical measured values in a "quite place" [49, 46, 48, 61] are in the range $\sqrt{S_a} \simeq 10^{-7} \rightarrow 5 \times 10^{-6}$ ms$^{-2}/\sqrt{\text{Hz}}$. For the microseismic peak "typical amplitudes are in the range 1 μm–10 μm with a spectral width of 0.1 Hz FWHM."[49]. "The actual frequency [of microtremor] varies from place to place and is not isotropic. It is relatively broadband and attenuation lengths are several hundred kilometers"[47]. Further the numbers quoted for both terms "are representative values only, and vary considerably with weather conditions, e.g. wind velocity and swell height respectively."[49]

Measurements of the vibration levels in room 39, building 3 at the NPL have been made. The results of these measurements are shown in figure 4.1. The measured acceleration rms root spectral density $\sqrt{S_a(f)}$, of vibration on the lab floor is over an order of magnitude bigger than the largest number quoted in the literature. This is not that surprising since their measurements are at a "typical remote site" rather than in a busy city.

Though the measured acceleration spectrum is not completely 'white', for the purposes of calculation it is reasonable to approximate it by:

$$\sqrt{S_a(f)} \simeq 4 \times 10^{-5}\ \text{ms}^{-2}/\sqrt{\text{Hz}} \tag{4.9}$$

$$\sqrt{S_v(f)} \simeq 6 \times 10^{-6}\ \frac{\text{Hz}}{f}\ \text{ms}^{-1}/\sqrt{\text{Hz}} \tag{4.10}$$

$$\Longrightarrow \quad \sqrt{S_d(f)} \sim 1 \times 10^{-6}\ \frac{\text{Hz}^2}{f^2}\ \text{m}/\sqrt{\text{Hz}} \tag{4.11}$$

where $\sqrt{S_v(f)}$ and $\sqrt{S_d(f)}$ are the root spectral density for velocity and displacement respectivly.

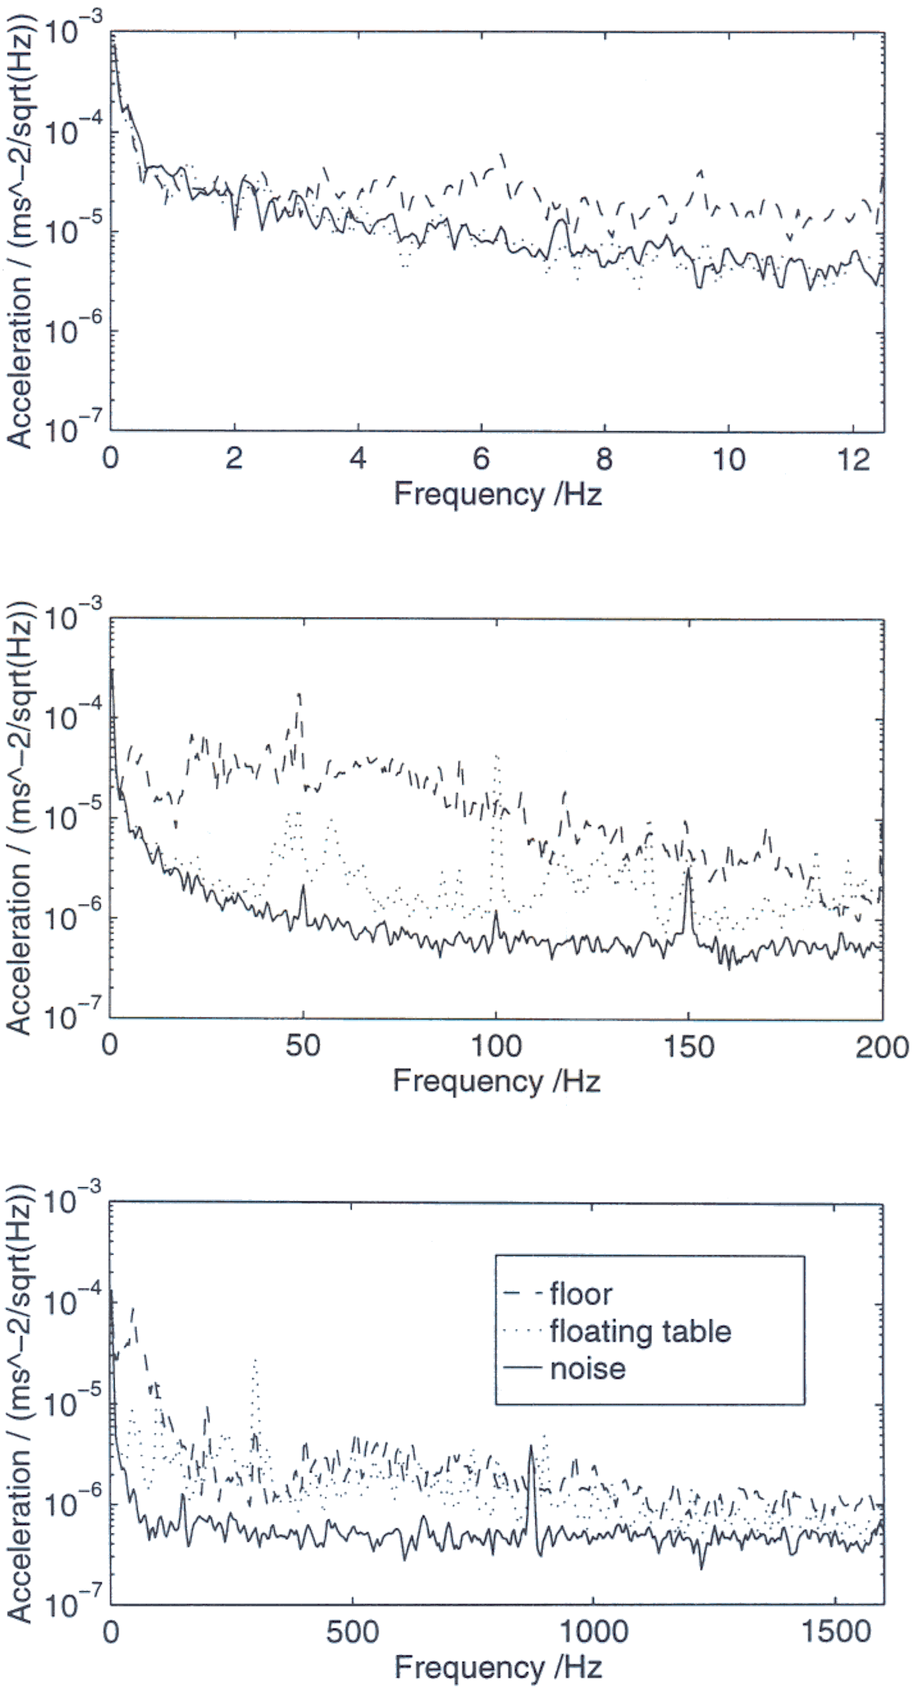

Figure 4.1: Vibration levels in room 39, building 3 at the NPL

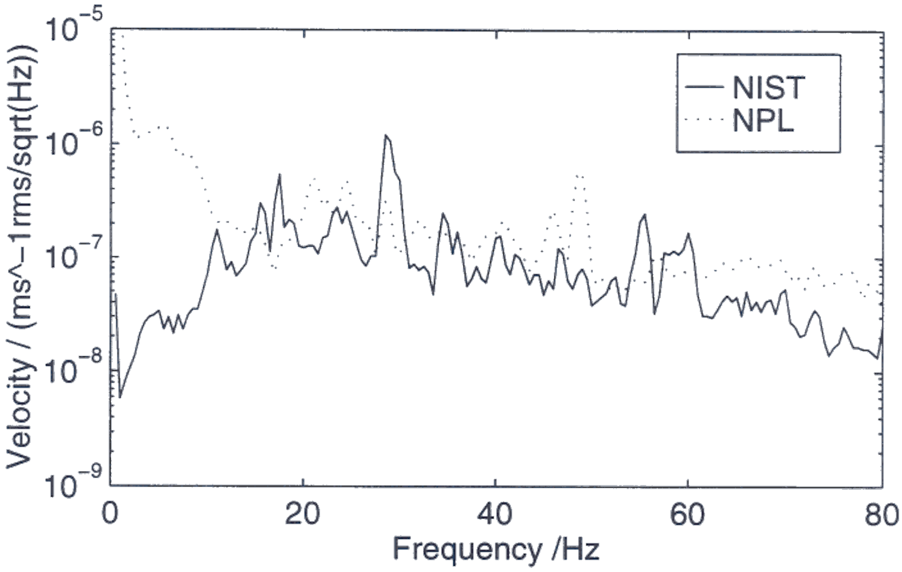

Figure 4.2: Comparison with NIST data - vertical velocity

The measurements also show the vibration on the surface of an air suspended optical table. From this data it can be seen that the table begins to provide measurable isolation above about 3 Hz. The vibration level is in fact at the noise floor of the detector up to about 40 Hz, so it is impossible to say *how much* isolation is being provided up to this frequency. In theory a table like this should provide seismic isolation $\propto 1/f^2$ above its resonant frequency (which preliminary measurements indicate is about 2.5 Hz). The higher frequency data ($f \geq 50$ Hz) does not show this kind of behavior, and is in fact far worse than this simple theory predicts. It is probable that most of the reduced performance is due to internal resonances of the table (see section 4.2.2), and possible that acoustic noise causes some of the degradation.

Vibration levels of the labs at NIST have kindly been supplied by Jim Bergquist [120]. His data is of velocity root spectral density, which suggests it was taken with a seismometer. This data is shown in figure 4.2, plotted with data taken from our lab for comparison. It shows in the 15–80 Hz region, the vibration levels at NIST are similar to those measured here. However at low frequencies the NIST data shows a rapid fall off in amplitude. It is suspected that this is not a real effect, but is rather roll-off in the detector response.

4.2.2 Vibration isolation

Simple theory of passive isolation

Vibration isolation in vertical and horizontal directions can be provided by springs and pendulums respectively. For a single stage of vibration isolation the transmissibility T is (see for example [84, 83]):

$$T = \frac{\omega_0^2}{\sqrt{(\omega_0^2 - \omega^2)^2 + 4\gamma^2\omega_0^2\omega^2}} \tag{4.12}$$

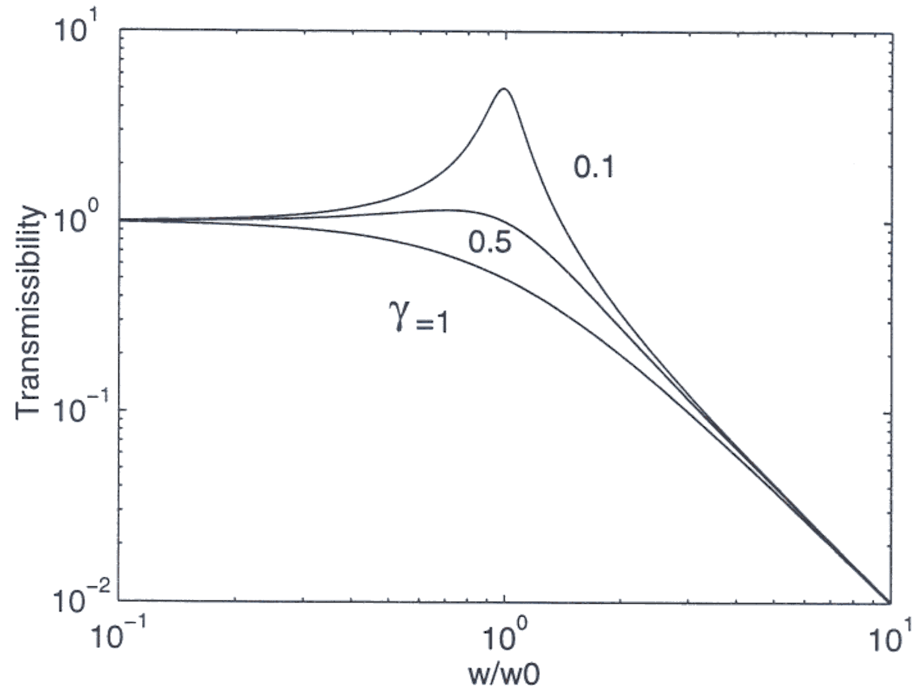

Figure 4.3: Transmissibility of a single degree of freedom velocity damped system

where ω_0 is the resonant frequency and γ is the damping factor. ($\gamma > 1$ overdamped, $\gamma = 1$ critical damping, $\gamma < 1$ underdamped). A plot of the transmissibility for various damping factors is shown in figure 4.3. In practice the damping factors tend to be quite small.

Vibration isolation is provided only above the resonant frequency of the system:

$$T \simeq \left(\frac{f_0}{f}\right)^2 \qquad \text{for} \qquad f \gg f_0 \tag{4.13}$$

therefore to provide vibration isolation at low frequencies it is necessary to make the resonant frequency as low as possible. The resonant frequency is given by:

$$f_0 = \frac{1}{2\pi}\sqrt{\frac{g}{l}} \qquad \text{for a pendulum} \tag{4.14}$$

$$f_0 = \frac{1}{2\pi}\sqrt{\frac{k}{m}} \qquad \text{for a spring} \tag{4.15}$$

For a 3 m long pendulum (like the NIST system [37]), the resonant frequency is 0.3 Hz. The reason for this choice of length is probably that this is the height of the ceiling. It is not practical to increase the length of the pendulum much beyond 3 m, and the improvement only goes like $l^{-1/2}$, therefore 0.3 Hz represents the practical limit to a simple pendulum resonant frequency.

The equation for the resonant frequency of the spring can be written in a much more revealing form by relating it to the static extension of the spring δl. Since the static restoring forces is $k\delta l = mg$, the resonant frequency is

$$f_0 = \frac{1}{2\pi}\sqrt{\frac{g}{\delta l}} \qquad \text{for a spring} \tag{4.16}$$

Thus to get a low resonant frequency in a spring system it is necessary to have a large static extension. For example, a spring with a 1 m static extension has a resonant frequency of 0.5 Hz. This extension could be achived by stretching out a 2 m rubber band, to give a 3 m total length. This resonant frequency is pretty much the lower limit for a simple spring.

Another way to look at the equation for the spring system is to write it in terms of material properties. For a stretched wire type of spring $k = AE/l$, where A is the cross sectional area and E is Youngs modulus. The necessary cross sectional area of the spring is determined by the strength of the material and the weight of the load, ie $\sigma = mg/A$, where σ is the strength of the material. Putting the equations together gives

$$f_0 = \frac{1}{2\pi}\sqrt{\frac{gE}{l\sigma}} \qquad \text{for a spring} \qquad (4.17)$$

To make a spring with low resonant frequency it is necessary to find materials with small (E/σ), *i.e.* very stretchy for a given strength. The properties of materials are catalogued in this way in reference [85], which show that the elastomers have the best performance, with soft butyl rubber perhaps being the best of these.

Damping

In vibration isolation systems it is important to have damping, to avoid significant amplification of the motion at the resonant frequency.

In equation 4.12 it is tacitly assumed that velocity dependent (viscous) damping is used. This is attractive since velocity damping does not degrade the high frequency performance of the system (and also because the problem simple to solve). Eddy current damping provides an excellent approximation to velocity damping and has been used in several systems [20, 51]. It does however require rather large magnets (~1 T) to make it effective, which may not be desirable in a frequency standard application. Jim Berquist at NIST suggests that a small dab of grease, can provide very effective damping without affecting the high frequency performance [37].

It is also possible, though rather difficult, to analyse a friction (Coulomb) damped system [91, 83]. The problem with this type of damping is for low amplitude excitation, a 'slip-stick' motion occurs which limits the performance of the system.

Internal resonances

A limit to high frequency vibration isolation is caused by internal resonances of the isolation structure or the object itself [85, 52, 45]. At low frequencies the transmissibility is given correctly by the simple theory, equation 4.12, but once the first internal resonance is reached, the isolation does not improve any more. Typically the first internal resonance occurs somewhere in the acoustic frequency range, a few hundred hertz say†.

†A 1200 × 2400 × 305 mm Newport table has a first resonant frequency of 160 Hz [121]

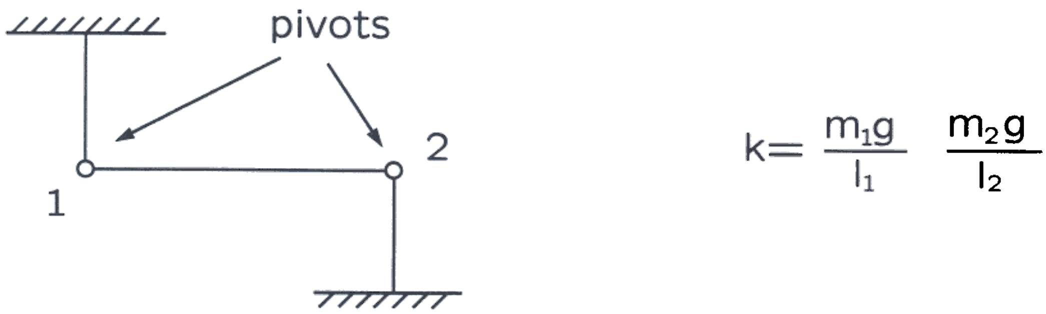

Figure 4.4: Illustration of negative k concept

The 'ordinary' pendulum 1, has a spring constant $k_1 = m_1g/l_1$. It is connected by a rigid rod to the upside down pendulum 2, which has a spring constant $k_2 = -m_2g/l_2$. Together the total system has a spring constant $k = m_1g/l_1 - m_2g/l_2$. With appropriate choices of masses and pendulums k can be very small, so giving a system with very low resonant frequency.

Multi-stage isolation

A multiple stage isolation system can be used to improve performance at high frequency. For an n-stage system the transmissibility is (see for example [45, 46]):

$$T \quad \left(\frac{f_0}{f}\right)^{2n} \qquad \text{for} \qquad f \gg f_0 \tag{4.18}$$

In such a system the resonant frequency, f_0, is roughly equal to the resonant frequency of an individual component, and so multistage isolation does not improve the low frequency performance. Multiple stage isolation systems are very common in gravity wave detector work, where large isolation at high(ish) frequencies, say $\sim$ 1kHz, is required. Their relevance for narrow linewidth lasers is far less obvious however, since it is low frequency vibrations which are most troublesome in this application. A double pendulum design has been used in some of Ohtsu's work [20].

Negative k systems

The idea in a "negative k" system is to balance the action of two springs or two pendulums, so as to produce a net system which effectively has a very small spring constant. This idea is illustrated in figure 4.4.

A number of devices have been developed: 1-D pendulums [55, 56, 60], 2-D pendulums [59] and spring system [35]. The 2-D pendulum [59] has achieved impressive results with resonant frequencies of less than 50 mHz for both x and y motions. Though as the authors point out this is still a prototype device, and is not currently suitable for suspending further structures.

A commercial version of a negative k system has been developed [122], though its performance is no better than the long spring/pendulum system described earlier.

Active vibration isolation

Active vibration isolation is to many people the ultimate method for extending the low frequency isolation capabilities of a system. It is however formidably difficult.

Single degree of freedom systems were demonstrated many years ago [87, 88, 89]. However such systems are of little use in a practical device because "a non-isolated degree of freedom reintroduces the seismic noise even if the other degrees of freedom are isolated"[46]. This fact is graphically illustrated by a servo system developed by J. L. Hall [7, 19], who did a vertical servo on a 4×8 foot optical table by controlling the feed/exhaust of the airlegs. Unfortunately the table did not have a very effective tilt servo and so began rocking at about 0.16 Hz, the known centre of the microseismic band. This caused an excursion of about 10 Hz in the beat between two independent laser systems.

A six degree of freedom system has only recently been demonstrated (1997) [61]. In this system isolation begins at 0.1 Hz and reaches a factor of 100 at 1 Hz. The servo system, and many of the problems associated with active vibration isolation are discussed in detail in reference [48].

Commercial systems do exist [123], though their performance is inferior to the research systems mentioned above. In fact they are slightly worse than the long spring/pendulum mentioned earlier.

Conclusion

In terms of good performance at a reasonable cost, hanging a table off the ceiling on rubber bands gives the best solution. A system like this, with a 3 m long pendulum mentioned in the examples, will have resonant frequencies below 0.5 Hz. This is at least as good as if not better than commercial negative k and active systems, and it is much cheaper. It is far superior to an optical table on air legs.

In terms of ultimate performance a 'home made' negative k or active system would be the best. It would however take a long time to develop, and would be extremely difficult and costly.

4.2.3 How vibration couples to cavity

Hall [7, 19] has attempted to calculate the effect of vibration on the cavity. His idea goes like this: consider the cavity stood on its end and shrinking under its on weight. The length change due to this effect is

$$\frac{\Delta L}{L} = \frac{\rho L}{2E} g \tag{4.19}$$

which gives a shift of about 20 MHz/g. Therefore to get a linewidth of 2 Hz say you need to reduce the ground acceleration noise to about 0.1 μg. From equation 4.10 you can calculate the rms acceleration in the 0–10 Hz band to be about 100 μg, so according to Hall a factor of at least 1000 in vibration isolation is needed to get to 2 Hz.

Blair [65] also considers a similar situation. The cavity is held vertically, but suspended from its centre. He suggests that this will provided some rejection of changes in length caused by vertical acceleration.

Neither Hall or Blair's theories account properly for the spectrum of acceleration noise that is present. This next section makes some attempt to address this issue.

Longitudinal motion

Length changes of the reference cavity result in frequency changes of the laser locked to it. Thus longitudinal vibration of the bar is extremely undesirable, since it leads to linewidth broadening of the laser.

Now consider the (slightly unusual) case of a cavity stood vertically on the floor. The vibration of the floor $d_0 \sin(\omega t)$ will induce a movement $D \sin(\omega t + \phi)$ on the top of the cavity. The amplitude of the differential motion of the ends of the cavity is $\Delta L = D - d$ (where have assumed that $\phi \approx 0$, which is true for $f \ll$ resonant frequency of cavity). Now the two amplitudes of the motion are related to each other by the transmissibility (which is the same equation as the transmissibility of a pendulum):

$$T = \frac{D}{d} = \frac{\omega_1^2}{\sqrt{(\omega_1^2 - \omega^2)^2 + 4\gamma^2\omega_1^2\omega^2}} \tag{4.20}$$

where γ is the damping factor, and $f_1 = \omega_1/2\pi$ is the first resonant frequency for longitudinal motion of the cavity. After making the approximation that $f \ll f_1$ and doing some algebra find that

$$\Delta L = d(1 - 2\gamma^2)\frac{f^2}{f_1^2} \tag{4.21}$$

It is probably reasonable to assume that the damping is quite small, i.e. $\gamma \ll 1$, so can simplify this equation a bit more.

$$\Delta L = d\frac{f^2}{f_1^2} \tag{4.22}$$

At first sight it appears that this equation 'blows up' at high frequencies, which would lead to an infinite linewidth. This is ignoring the fact that the amplitude of the floor displacement, d, is a function of frequency. Measurements of the floor displacement are in terms of spectral density (equation 4.11), so at this point it is appropriate to recast equation 4.22 in terms of spectral density:

$$\sqrt{S_{\Delta L}} = \frac{f^2}{f_1^2}\sqrt{S_d} \tag{4.23}$$

In fact you can recast equation 4.4, in terms of spectral density as well, to give

$$\sqrt{S_\nu} = \frac{\nu_0}{L}\sqrt{S_{\Delta L}} = \frac{\nu_0 f^2}{L f_1^2}\sqrt{S_d} \tag{4.24}$$

Now $\sqrt{S_d} \propto 1/f^2$ (see equation 4.11), so equation 4.24 does not diverge and a finite linewidth is predicted.

The final piece of information needed is the resonant frequency for longitudinal motion of the bar, f_1. For an elastic bar with constant area cross section [83] the resonant frequencies are:

$$f_n = \frac{n}{2L}\sqrt{\frac{E}{\rho}} \qquad \text{where} \qquad n = 1, 2, 3, \tag{4.25}$$

where L is the length of the bar, E is Youngs modulus and ρ is the density of the material. For the standard ULE spacer of section 4.1, a bar of length 0.1 m has fundamental longitudinal resonant frequency of $f_1 = 27.5$ kHz.

Putting this information together predicts a spectral density of $\sqrt{S_\nu} = 6.3\ \mathrm{Hz}/\sqrt{\mathrm{Hz}}$. Using equation 2.10, this gives a linewidth of $\Delta\nu = 124$ Hz. This number seems quite reasonable for an estimate of the linewidth due to vibration (*i.e.* is is similar to other reported linewidths).

There are however some problems. First, the cavity is stood on its end, rather that laying on its side as in the real case. Hall [7] guesses at a factor of 100 suppression in spectral density for this effect, which would give a linewidth of $\Delta\nu = 0.01$ Hz. Even if the suppression factor is less that this, say only a factor of 10, the linewidth it predicts is still uncomfortably small.

It is reasonable to expect that horizontal motion couples into the cavity, as predicted (more or less) by the theory. We have made no measurements of horizontal motion at this stage, but data from NIST suggests that the horizontal velocity is a factor of 10 less than the vertical motion. This would give a linewidth of $\Delta\nu = 1.2$ Hz. This number is a bit small.

Veitch [52] treats a similar problem to this, to do with a resonant bar gravitational antenna. This paper is very hard to understand, but he seems to be talking about how vertical vibration couples into horizontal motion of a bar. It is not completely clear to me how to apply this work to our situation.

Another question area in the theory presented here relates to equation 4.20. This is the transmissibility of an oscillator (in this case the spacer material) with viscous damping. It is not known whether viscous or structural damping is the appropriate model materials like ULE. (See reference in Blairs paper [65]).

Bending motion

A detailed theory of how bending motion contributes to length changes of the cavity does not exist. It is reasonable to expect that an effect exists at some levels though. What can be said is that if bending does contribute to frequency shifts of the cavity, then making the resonant frequency of the first bending mode higher will reduce the effect.

The frequencies of the bending modes of the motion [83] are given by:

$$f_n \quad \frac{n^2\pi}{2L^2}\sqrt{\frac{EI}{\rho A}} \tag{4.26}$$

where ρ, is the density of the material and I is the second moment of area. From [86] the second moment of area I, for a cylinder is:

$$I = \frac{\pi r^4}{4} \tag{4.27}$$

where r is the radius of the bar.

For the standard ULE cavity, this gives a first bending mode of frequency of 10 kHz. Note the $1/L^2$ dependence though – a 0.3 m cavity (like the NIST cavities) has a first resonant frequency of only 1.2 kHz.

Because of the lack of a theory about how bending effects the frequency, it is difficult to assess its importance. However, NIST (based on some research at the VIRGO gravity wave detector [no reference unfortunately]) have recently invested in an American football shaped cavity. The idea of this is to increase the frequency of the first bending mode. Brent Young [119] at NIST, did not seem to think this had contributed significantly to the recent success of this work though.

Intra-Cavity Power/Beam Pointing

Intra-cavity power

A typical cavity used in a stabilised laser application may have a finesse of $\mathcal{F} =$ 100 000. If say 1 mW of light power is incident on the cavity, then the intra cavity power is 100 W ! The radiation pressure of this light can have an effect on the locked laser frequency, by bending the cavity mirrors and stretching the spacer material. Bergquist [37] treats these effects in detail. Assuming the cavity of section 4.1, a 1 μW change in incident power causes a frequency shift of 10 mHz due to mirror bending and 1.5 mHz due to cavity stretching. These shifts are quite small.

Far more significant is the effect of power fluctuations on the dielectric mirror coatings. "Most of the power is absorbed in the first few layers of the dielectric stack where the intensity is highest. When the light amplitude fluctuates, there is a transient response followed by a relaxation to a steady state condition. For a radiation mode size of 200 μm, we have measured a 2 Hz/μW shift of the cavity resonance...the magnitude has been measured to be as much as 20 Hz per μW." [37]. Schiller [41] has also measured a power dependent shift in frequency of 6 Hz/μW. In light of this fact both of these groups are rather sensibly using relative low incident powers on the cavities of 100 μW and 30 μW respectively. It would probably be a good idea to use even less power than this, given that the shot noise limited linewidth could still be extremely small.

Further the NIST group have implemented intracavity power stabilisation, which is effective to the 0.1% level [119]. This is almost certainly a very good idea for achieving a 1 Hz linewidth. Only low(ish) frequency intensity fluctuations matter, since the high finesse of the cavity provides averaging at the higher frequencies. This cut-off frequency f_{cut} is infact equal to half the cavity linewidth. For the standard cavity considered here $f_{cut} = 7.5$ kHz.

$$f_{cut} = \frac{\Delta\nu_c}{2} = \frac{c}{4L\mathcal{F}} \tag{4.28}$$

Another related type of effect has been noted by both Ueda [35] and Ohtsu [20]. Both these groups see instabilities which decrease with increased locking time. For example 1 hour after switch on Ueda sees frequency fluctuations with standard deviation of 42 kHz over merely 10 seconds. After 18 hours of operation this has reduced to about 1.3 kHz over 10 seconds. Ueda suggests "this effect is caused by the relaxation of a heating effect of the dielectric coatings of the FP cavities by resonant laser powers" [35]. A significant difference between this work and the quest for a narrow linewidth does exist though. In their work rather larger powers (10–100 mW) are incident on the cavity to achieve the very low shot noise requirements, consequently the circulating powers are rather bigger.

There are a number of causes of intra cavity intensity fluctuations. Firstly intrinsic intensity fluctuations of the laser. Secondly if an optical fibre is used, then polarisation fluctuations caused by the fibre lead to intensity fluctuations through subsequent polarising elements in the beam path. Finally beam pointing instability causing a variation in the modematching to the cavity (see section 4.3.3).

4.3.2 Beam Jitter

The effect of lateral and angular beam jitter on Fabry-Perot cavity stabilization has been investigated theoretically by Calloni et al [111].

They derive the expected spectral density which consists of two terms. This formula is only partially repeated here, because the second term is rather complicated and in practice is negligible:

$$\begin{aligned} \sqrt{S_\nu(f)} &= 2(\epsilon_s\tilde{\beta} + \beta_s\tilde{\epsilon})f \\ &+ \frac{c}{\mathcal{F}^2}\left\{\frac{factor\ same\ size}{as\ first\ term}\right\}\left\{\frac{factor\ to\ do\ with}{off\ axis\ modes}\right\} \end{aligned} \tag{4.29}$$

where $\epsilon = \delta x/w_0$ and $\beta = \delta\theta\pi w_0/\lambda$ are lateral and angular beam displacements terms. β_s and ϵ_s are static (slowly varying) terms, and $\tilde{\beta}$ and $\tilde{\epsilon}$ are root spectral densities (quickly varying). δx is the lateral beam displacement, $\delta\theta$ the angular beam displacement and w_0 is the beam waist.

The second term of their equation represents the effect of 'linepulling'. To our knowledge it is the only proper treatment of this often mentioned effect. Now for a high finesse cavity, provided the cavity is not close to confocal, this second term is negligible. This is because of the $c/\mathcal{F}^2$ term. Even though the speed of light c is big, the high finesse of a modern cavity $\mathcal{F} = 100\ 000$ suppresses this term by about a factor of about 100 relative to the first term. This statement must be qualified though. The full formula contains a factor to do with off axis modes. For a confocal cavity this factor 'blows up', i.e. tends to infinity. This condition should therefore be avoided. For a 'typical' cavity this factor is about 1. For a carefully chosen cavity this off axis mode factor can infact be made very small. The point is that this very careful choice of cavity length and mirror curvature is not really necessary (provided to confocal condition is avoided).

The first term is not that easy to interpret. It may be effectively 'seeing a different length of cavity' because of misalignment. Even though this is a second order effect it dominates for the reasons just described.

It is possible to estimate the size of the effect. In the paper [111] they suggest "static misalignment of $\epsilon_s \approx \beta_s \approx 0.16$, which corresponds to 95% of light coupled into the fundamental mode". Though it is almost certainly possible to initially align the cavity much better than this, this does not seem unreasonable for the alignment after the system has 'relaxed' for a few days. The linear jitter has been measured by Hough [54] for a YAG laser, who found $\delta x(100\ \text{Hz})/w_0 \approx 1\times10^{-6}/\sqrt{\text{Hz}}$, $\delta x(10\ \text{Hz})/w_0 \approx 3\times10^{-5}/\sqrt{\text{Hz}}$, $\delta x(1\ \text{Hz})/w_0 \approx 3\times10^{-3}/\sqrt{\text{Hz}}$. This corresponds to a root spectral density of $\sqrt{S_\nu(100\text{Hz})} = 3\times10^{-5}\text{Hz}/\sqrt{\text{Hz}}$, $\sqrt{S_\nu(10\text{Hz})} = 1\times10^{-4}\text{Hz}/\sqrt{\text{Hz}}$, $\sqrt{S_\nu(1\text{Hz})} = 1\times10^{-3}\text{Hz}/\sqrt{\text{Hz}}$. No measurements are available for the angular deviation in this frequency range. In terms of making a 1 Hz linewidth, these sizes of spectral density are not that important. (They would be a worry for a laser for a gravitational wave detector).

4.3.3 Beampointing and power fluctuations

Before the effect of beam pointing is completely dismissed, there is another way in which it can effect linewidth. Variation of beam pointing will lead to a variation in the modematching to the cavity, and a consequent power fluctuation of the circulating power. As already mentioned, power dependent frequency shifts have been noted by several groups of workers.

It should be possible to completely calculate the effect on modematching efficiency of an angular or lateral displacement. We however do not know how to do this. Instead to get an idea of what's going on you can extrapolate from some data in Calloni*et.al.*[111]. They suggest that for $\epsilon_s \approx \beta_s \approx 0.16$ you get 95% modematching efficiency. Assuming a linear relationship between linear displacement and modematched power gives $P = 1 - 0.31\epsilon$. From Houghs data [54] you can calculate $\epsilon_{rms} = 2\times10^{-3}$. This value would give a 0.1% variation in modematched power. This fluctuation though small is not a completely negligible consideration for stabilisation at the 1 Hz level.

This type of effect has been recognised by the NIST group: "it may be better to stabilize the intensity of the light circulating in the resonator by using the transmitted light from the cavity. This should give a first-order insensitivity of the frequency of the laser to power fluctuations caused by relative motion between the cavity and the injected light beam." [37]

Compliance

In catalogues for optical tables it is usual to see graphs for the compliance. The compliance, C, is equal to the displacement per unit force (i.e. apply a force to a thing and look how much it moves). For an ideal rigid body $C = 1/m\omega^2$, where m is the mass of the object and ω is the angular frequency of the applied force. Real bodies follow the ideal law at low frequencies, but at higher frequencies show resonances, corresponding to bending modes of the object. For a larger optical table the first resonance may be at 150 Hz say.

It is important to realise that compliance is *not* telling you anything about the vibration isolation properties of the system. Compliance information is primarily about the stiffness of the table, and also about the internal damping of the table. Compliance tell you about the *relative* motion on the table, not the overall vibration level on the table.

The Newport website [121] contains information about how to get beam pointing stability from compliance data and a knowledge of the floor acceleration. Unfortunately this information is not referenced, and time has proved insufficient to verify it.

Tilt

Tilting the reference cavity, even by a very small angle, can have a significant effect on the frequency of laser locked to it.

Consider a cavity tipped up by an angle θ from horizontal. The force acting along the cavity due to its own weight is $mg\sin\theta$, where m is the mass of the cavity. The mass can be rewritten in terms of density, to give a force $(AL\rho)g\sin\theta$. Now Youngs modulus is defined as E =(stress)/(strain)=$(F/A)/(\Delta L/L)$. Thus the frequency shift of the cavity due to this tilt is

$$\Delta\nu = \frac{\Delta L\nu_0}{L} = \frac{\nu_0 F}{AE} = \frac{\nu_0 L\rho g\sin\theta}{E} \tag{4.30}$$

Inserting the numbers for a standard cavity (and using $\sin\theta \approx \theta$ for small angles) gives a frequency shift of

$$\Delta\nu = 10\ \text{Hz}/\mu\text{rad} \tag{4.31}$$

This size shift is very similar to that measured by J.L.Hall on his servoed table [7]. Infact the inadequate tilt servo of this active vibration isolation, was one of the limiting factors in this work.

The only other explicit measurement of tilt sensitivity is by Ueda [35]. The results here are slightly more difficult to interpret. He drives one end of the cavity up and down with a piezo actuator, at a frequency of 5 Hz. The measured sensitivity is $\sqrt{S_\nu} = 3.38$ mHz/$\sqrt{\text{Hz}}$/nrad. This must be the noise measured on the spectrum analyser at 5 Hz (not white noise). To get an idea of the linewidth, use the modulation index $\beta = \sqrt{2S_\nu/f} = 2.5$ rad/μrad @ 5 Hz. So a μrad of tilt is giving more than a 5 Hz linewidth. This suggests the result are consistent.

4.4 Other issues

4.4.1 Modulation issues

The following considerations were discussed in depth in the paper by Hall [5]. These issues can seriously effect the observed level of noise, and also the accuracy of which the discriminator zero-crossing is followed. The latter can lead to frequency excursions on timescales of above ~ 0.1 seconds. This can lead to broadening of a beat note, or poor performance in an Allan variance.

AM at the detection frequency

Any noise sources occurring on the laser light at the modulation frequency will be interpreted as FM noise by the servo [75, 5]. Two candidates for this are electrical pick-up at the modulation frequency by the detector, and residual AM produced by the modulation technique. The first of these can be observed on a spectrum analyser with no light on the detector. The second can be seen by observing AM level on a photodiode. To avoid AM at the modulation frequency, it is desirable to decouple the modulation mechanism from the laser itself. One of the most effective ways of doing this is to use an extra-cavity EOM to provide the modulation sidebands. However, unless care is taken, an EOM can also give AM at the modulation frequency. If the EOM is poorly aligned - the polarisation of the light is rotated at the drive frequency. Polarising elements in the system then convert this rotating into amplitude modulation. Looking at the photodiode output on a spectrum analyser allows one to reduce the AM by orders of magnitude by careful alignment of the EOM. A residual AM of 50 dBc was observed using a Leysop modulator. The required level of AM appears to be set by the accuracy of the lock requirement and linewidth of the discriminator. Hence to achieve a 10 mHz lock with a 10 kHz linewidth requires AM lower than -60 dBc [5] (note Hall achieves -100 dBc). The effect of misalignment of the EOM is modulation of the orthogonal polarisation in the modulator. The same effect can occur due to other orthogonally polarised back reflections, which may explain why Hall has an isolator in front of the EOM [5].

Modulation accuracy

The modulation accuracy and stability is also brought into question by Hall - "..even order sidebands clearly will lead to a systematic frequency shift of the lock point." [5]. These can be thermally varying and affected by light scattering. This goes to explain the care with which Hall modulated his EOM (-80 dBc waveform distortion achieved by using a crystal locked sine-wave source), and demodulated the signal (notch and synchronous filters to give 50 dB isolation of 2nd and 3rd harmonics and a 7-pole Chebyschev filter to remove higher harmonics).

This care of modulation is suggested by other groups by using function generators to drive the EOM, and often notch filters are used to remove the dangerous 2nd harmonic.

4.4.2 Optical Isolation

Optical isolation is universally regarded as a problem in stabilising a laser system. The optical feedback can upset the laser oscillation - but more generally feedback between components can cause a shift in the lock point. This second effect occurs because the small counter-propagating beam adds coherently with the main beam causing a phase shift.

The sensitivity of the laser to feedback depends on the type of laser. It is generally accepted that 2 mirror cavity lasers (He-Ne or diode) are much more susceptible than ring cavities (Ti:Sapphire, Nd:YAG or Dye). This is because in a ring cavity the counter-propagating beam is not oscillating (due usually to a Faraday rotator and Brewster surfaces in the cavity). As a practical comparison, Ohtsu uses a single AOM for isolating an Nd:YAG and achieves a real linewidth of 30 Hz [20, 26], whereas Hall uses 2 AOMs and two 33dB Faraday isolators for a HeNe system and reports that shifts due to optical feedback are well below 1 Hz[5]. So the difference in the level of isolation is clear, although it may be questioned whether Ohtsu does not have a slight feedback problem.

The effect of feedback between optical components is often referred to as an etalon effect. The weak reflection interferes with the main beam to form an interference pattern background to the main signal. As this etalon is air spaced and unstable, the background (and therefore the lock point) is shifted. It is easy to quantify this. If the width of the cavity resonance is 10 kHz and one requires a locking accuracy of 10 mHz, all feedback beams must be reduced by 10^6 in amplitude or 120 dB intensity. Slight misalignment of optical components may give most of this factor, but as an example Hall uses an isolating AOM *after* the cavity in front of the detector [5, 3]as it was found that light reflected from the detector increased the background noise by 20-30 dB. The problems of etalon effects are especially important if it is not possible to misalign the components. An example of this is the use of a fibre optic cable. For a plane-plane cleaved fibre, the ends must necessarily be parallel to the laser and reference cavity. It is therefore important to make sure any fibre optic has very low back reflections by making it angle cleaved at both ends.

4.4.3 Birefringence

> We observed a strong dependence of the resonance frequency on the polarisation of the incoupled laser light.... The frequency difference between orthogonal polarisations is 1.2 kHz. It is not due to power modulation or electronic offsets but to Birefringence in the mirrors.
>
> —Schiller [41]

The resonance frequency for a Fabry-Perot cavity has a slight dependence on the polarisation of the light. Both Schiller [41] and Hall [5] put this effect at the 1 kHz level for 180° rotation of the input light polarisation. This effect is ascribed to the Birefringence in the mirrors. The are two ways that this effect can be reduced. First one could use linearly polarised light set with a polarisation such that " ... the cavity frequency is an extremum" [41] (*i.e.* the frequency shift/degree slope is at a minimum). Else circularly polarised light can be used [5] reducing the effect by a factor of 40 or so. More information on the effects of Birefringence can be found in the reference [114, 115].

4.4.4 Conclusion

A number of issues have been addressed in this chapter:

- Materials choice – ULE.
- Thermal and pressure effects – thermal problems difficult but soluble, pressure easy.
- Vibration – follow the NIST scheme which is certainly good enough to get a 1 hertz linewidth. Unclear if will get you much better than this
- Intra cavity power stability – use low input power, and stabilise the output power from the cavity
- Beam pointing – mount cavity and laser (or optical fibre output) on same isolation stage
- AM at drive frequency,modulation accuracy, optical isolation and Birefringence – take care over all of these effects

The other point to make is that some effort has gone into trying to calculate the size of the effects limiting linewidth. In general it seems to be the case that in this subject calculations always seem to underestimate the size of the problem. It may well be that this is just because the are lots of small, and often related, things going on which need to be calculated. Ultimately performance of a real system is all that matters. If at every stage of development the diagnostics are capable of telling you WHAT the problem is, then achieving a 1 Hz linewidth will be easy.

Chapter 5

Conclusion

The conclusion is in the form of a recommended system and is shown in figure 5.1

5.1 The project goal

Optical frequency standards require a laser source of ultra-high spectral purity. Linewidths of under 1 Hz ($< 10^{-15}$ of the laser frequency) are needed. Our goal is to initially demonstrate a laser source with a linewidth of a few hertz. Achieving and demonstrating this requires two systems that can be compared. Ideally the two systems would be completely independent (*e.g.* separated by many miles), failing this, obvious sources of common mode rejection between the systems should be removed. To this end, it is planned to compare two cavity stabilised light sources on two independent vibration isolation platforms.

The philosophy of the solution is to decouple the two major technical problems: servo electronics and reference cavity issues. The idea is to use a laser source (Nd:YAG laser) which has low intrinsic noise which will make the servo problem easy to solve. This will allow us to investigate and solve all of the reference cavity problems like vibration, intracavity power etc. and demonstrate a narrow linewidth in this model system. The reference cavity problems are common to any type of system, so subsequent work on more difficult (but more useful) laser sources would be a matter of only solving the servo problem (which is unique for each type of laser source anyway).

The laser source

The laser source chosen to perform these experiments is a non-planar ring oscillator Nd:YAG laser (from Lightwave electronics). This laser has the lowest free running noise of any source, with a 50 ms linewidth of a few kilohertz at Fourier frequencies extending to only a few tens of kilohertz. This means that the servo loop required to reduce the laser noise to the shot noise of the system should be easier to design than for, say, an extended cavity diode laser system.

Although it is conceptually easier to work with two separate laser sources, this is not necessary (and in our case not financially possible). If the laser output is split into two portions and each is closely servoed to two independent reference

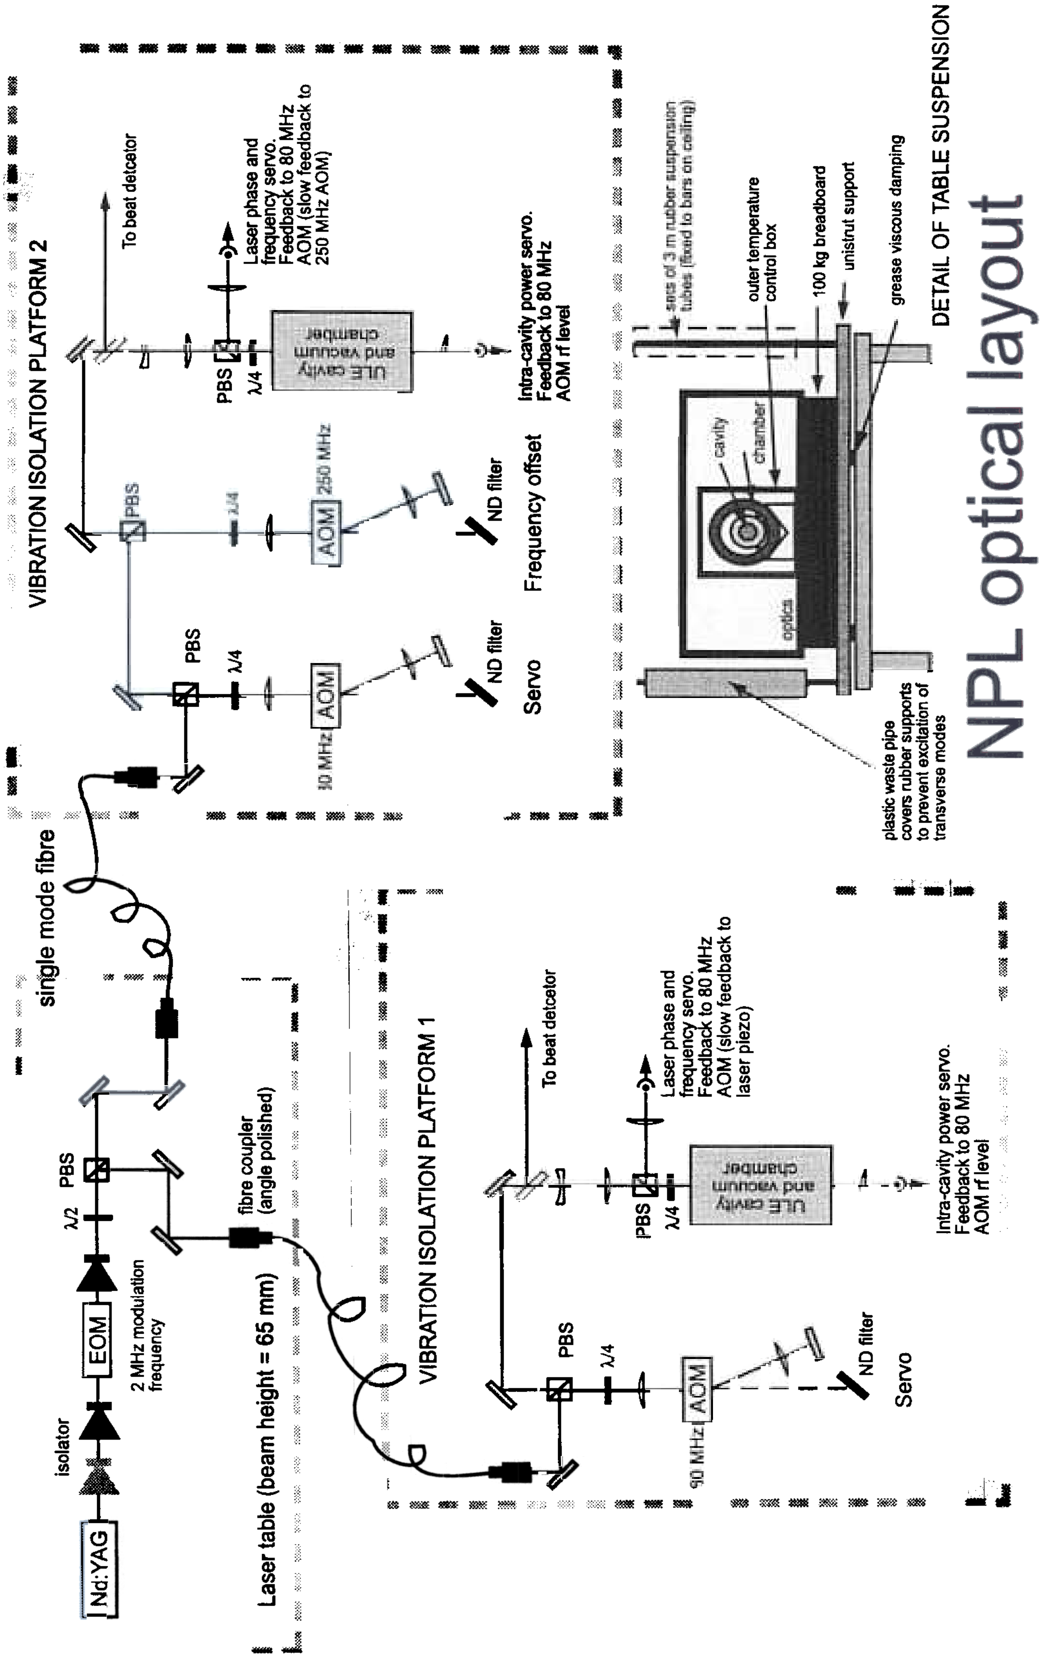

Figure 5.1: Proposed NPL system

cavities, the noise in the two laser beams becomes that of the cavity (cavity noise + discriminator noise). In our case we will split the output from the Nd:YAG laser and externally stabilise the two portions of light with AOMs (see setup in figure 5.1). The AOM was chosen, as it should be possible to achieve a servo bandwidth of 200 kHz, given enough gain at low frequencies to suppress the laser noise to less than 1 radian. The AOM is used in the double passed configuration to remove any beam pointing issues due to the frequency dependent Bragg angle of the AOM. Low frequency (80±5 MHz) yet high efficiency (85%) AOMs are used for frequency locking, whilst (a number) of high frequency (250±50 MHz), low efficiency (40%) AOMs are used to provide the frequency shift between the cavities.

The other use for an AOM is as an intensity stabilisation unit. By varying the rf power to the modulator, the Bragg efficiency is changed. This effect will be used to stabilise the *output* power of the cavity to around the 0.1% level. This is necessary as any intensity noise on the laser, or slight beam pointing fluctuation will lead to the intra-cavity power fluctuating in the reference cavity (see section 4.3.3). This induces frequency changes at the 1–20 Hz/μW level [37, 41]. The effect is reduced by using a low input power (say 10 μW or less), but intensity stabilisation will most likely prove to be necessary.

The reference cavities

The cavities will be made of Corning 7971 ULE material. Only Zerodur has a comparable room temperature coefficient of expansion ($\sim 10^{-8}$–10^{-9}). ULE is chosen in preference to Zerodur because of the experiences of the NIST group (see section 4.1 and section 3.5). The cavity's length will be chosen such that the free spectral range is low enough so that the frequency difference between two cavities can be bridged by broadband (250±50 MHz) AOMs. It is not clear whether the best route is to have two cavities of the same, or different lengths. There is a suggestion that for two cavities of the same length, some common-mode rejection of noise may occur (section 2.1). Almost certainly one cavity will be length 100 mm (FSR = 1500 MHz) and the other either 200 mm (FSR = 750 MHz) or 100 mm.

The cross section to length ratio will be as high as possible so that the first bending mode frequency is as high as practical (typically 60 mm circular diameter for a 100 mm length cavity). This is important to minimise the effects of vibrationally and thermally induced length changes. The super polished mirrors will have low loss to give high transmission (at least 50%) and low heating due to absorption. The finesse is high enough ($\sim$ 200 000) to give a narrow linewidth (5 kHz) but low enough so that the unstabilised laser noise is not significantly larger than the linewidth.

The cavities will be housed in vacuum chambers evacuated to around 10^{-5} mbar so that pressure variations are insignificant. The temperature of the chamber will be controlled to better than 10 mK, and it will be desirable to achieve a long time constant for thermal transfer to the cavity, in order that short term temperature fluctuations are damped. Temperature control will most likely be achieved with a two layer control system. Ideally daily temperature fluctuations will be reduced to the point where the frequency excursions they induce are small compared to the monotonic creep of the ULE material. This is so that the drift is as low as possible, but also as predictable as possible. Predictable drift can be taken out by an opposing sweep applied to the AOM from a synthesiser.

Vibration Isolation

The importance of vibration isolation in achieving a narrow linewidth laser has been recognised for a long time. There are two basic approaches: isolate the cavity or isolate the entire laser system. If vibration were the only important issue, then isolating the cavity should be sufficient to achieve good results. However as has been pointed out there are other issues which are important/critical: intra cavity power fluctuations and beam pointing (section 4.3.3). Obviously if the cavity is well isolated against vibration but the input beam is not, then the relative motion between the two will lead to both beam pointing problems and power fluctuations (via variation in mode matching). Hence we will fix the laser beam relative to the cavity by fibre coupling the light onto the suspended table holding the cavity. The cavity will be sat rigidly in the vacuum chamber (as opposed to being suspended in it) for the same beam pointing reasons. The cavity will sit inside the (aluminum) vacuum chamber on viton (or similar) pads to decrease the thermal coupling to the chamber.

The vibration isolation platforms for the two systems will follow the approach of the NIST group. That is a 100 kg breadboard suspended from the ceiling by a number of 3 m long rubber tubes. This gives a pendulum frequency of less than 0.5 Hz both vertically and horizontally. Unlike NIST, we do not plan, at least initially, any active position control on the tables. This is necessary only if there is a critical beam pointing issue for getting the light off the table and onto another one. In our case there will be a requirement to mix the outputs from both systems on a photodiode to observe the beat between the two systems. It is hoped that since the systems are close together and the beat measurements do not last very long beam pointing effects will not be an issue.

Crude height control will be achieved (as with NIST) by adding or removing weights from the table to keep it a few millimetres above the table frame. The pendulum frequency of the system will be velocity damped with blobs of grease between the table and the frame it sits on when not suspended. In order to lift 100 kg of the breadboard plus all the optics, it will be necessary to use a number of rubber tubes. These will be grouped in the four corners of the table. As for NIST, the 3 m long rubber 'springs' will be surrounded by a plastic tube to attempt to stop the excitation of transverse modes of the rubber by air currents.

The breadboard will be covered with a wooden acoustic isolation box. Varnished plywood is chosen for its stiffness to weight ratio [109] to reflect as much acoustic noise from the box as possible. Any transmitted acoustic energy will be further damped by a layer of lead lined foam. The lead sheets act as a non-impedance matched acoustic load. The inner foam layer also acts to damp any low frequency acoustic resonances occurring within the box. The vacuum will provide a large amount of acoustic isolation, but acoustically induced vibration of the whole table will not be reduced (apart from the fact that the table is so massive).

Modulation accuracy and optical feedback

Studying the results of Hall [5], shows that care will be needed when setting up the optics for the system.

An electro-optic phase modulator is chosen to produce the sidebands on the laser light for the Pound-Drever-Hall stabilisation technique. A modulation frequency of ~2 MHz is chosen being sufficiently higher than the required servo bandwidth. At present the EOM is driven by a resonant circuit and a power transistor oscillator. The spectral purity of the system may need to be improved by driving the EOM with an amplified synthesiser. This is to ensure higher order harmonic content on the light is minimised, as it will shift the lock point of the discriminator. At the rf beat detector, the signal will need to be notch filtered to remove 2nd harmonic. Further filters to remove 3rd and higher harmonics may be necessary. All rf pick up from the EOM drive will have to be removed from the detection electronics to remove base-line shifts to the discriminator signal. All mains-locked 50 Hz noise, and harmonics, will also have to be reduced to below the shot noise, as the low Fourier frequency will cause large problems. The orientation of the phase modulator will have to be carefully set to minimise AM at the drive frequency due to the misaligned EOM causing the polarisation fluctuations. An EOM is chosen over driving the laser piezo at ~ 700 kHz [77] or using an AOM because of these AM problems. It is not clear whether reducing the AM to -50 dBc will be enough to remove this effect or not (see Hall [5]).

We wish to reduce all lock point shifting problems to the 0.1 Hz level or less. From a 10 kHz linewidth, this means optical feedback (etalon effects) between components will have to be reduced to the 10^{-5} level in amplitude, or 10^{-10} level in intensity [5]. This is particularly important for elements which can not be slightly misaligned (the reference cavity, fibre optic cable). The fibre optic connection will be chosen to have very low back reflection (by using angle polished fibre at both ends). The EOM will have an isolator in front of it to protect it from any feedback (inducing polarisation rotation at the drive frequency). The laser will have at least 1 isolator between it and the EOM giving a total of at least 60 dB isolation (two isolators of 30 dB isolation each). In addition, the light reflected from the reference cavity will be picked off by a polarising beam splitter giving an addition ~20 dB of isolation to the laser. All transmitting optics will have antireflection coatings. Reflections from the photodiodes will have to be either isolated or misaligned. The zeroth order from the AOMs will be absorbed in a neutral density filter set at Brewster's angle.

5.2 The plan of action

Although the end goal is two independent systems, it is wise to build the experiment up in steps, making sure each stage works before carrying on to the next more difficult experiment. The following discusses the likely order of research, including what should be learnt from each optical set-up.

The first step is to lock the Nd:YAG laser to one cavity and improve the servo until the noise is suppressed to the shot noise over the servo bandwidth. The laser and cavity should be on the same table to avoid any beam pointing issues. The laser power incident on the cavity will be set such that the shot noise will not significantly limit the performance of the system, yet low enough that thermal loading of the cavity mirrors is minimised. From these considerations a laser power of ($\sim$ 10 μW) incident on the cavity seems reasonable. Work on the servo will be achieved by looking at the noise at the lock point of the cavity whilst improving the servo loop filter, so that the relevant noise is suppressed. The electronic detection noise will also have to be reduced so that it is significantly less than the shot noise. This is the only real drawback of using a low optical power.

Once the servo system has been shown to work, it will be necessary to observe the real performance of the system. This can be done initially by placing another cavity on the same isolation table and analysing the stabilised light with the second cavity acting as an optical spectrum analyser. This will give a fair indication of any vibration/acoustic disturbances without introducing Doppler shifts between the two cavities. The power spectral density observed can be used to estimate the linewidth of the locked laser. The effect of the aforementioned problems (feedback to the laser, intensity noise etc.) should show up in this measurement as noise in the power spectral density. Effects that shift the lock point (AM at the drive frequency, etalon effects, thermal drift, laser power induced thermal drift etc.) will show up on the DC coupled error signal, although they will be indistinguishable from cavity noise.

In order to untangle these optical effects and cavity effects, it would be desirable to lock the two portions of laser light to adjacent modes of a single cavity and perform an Allan variance (this will require duplicate servo electronics). This will detect any slow shifting of the beat frequency due to variable offsets, whilst suppressing cavity noise by a factor of $\sim$ 20 000. This test can look for problems like: effect of slight misalignment of EOM, effect of polarisation rotation of light into cavity (birefringence test) and effect of additional optical isolation at various points in the system.

Once the system appears to be working as well as possible, it is time to obtain some definitive data. The second beam from the laser will need to be stabilised to the second cavity using a similar servo to the first system. Again both systems should be on one table to avoid Doppler effects. The beat between the two stabilised outputs will show the true laser linewidth. Thermal and laser induced drifts will show up and need to be minimised. The effect of stabilising the power leaving the two reference cavities can be investigated. Allan variance data can be taken of the beat. This stage is the acid test for the system. Effects that can be investigated: effect of DC light intensity change into cavity, effect of cavity output power stabilisation, effect of extra vibration on laboratory floor and effect of additional acoustic noise.

If after this work it is clear the laser is working as well as possible it is time to remove the laser source from the table. This is because in most systems the laser cannot reside on the table as it is a vibration source. Even in the case of an Nd:YAG laser, the effect of stiff cables may be to compromise (slightly) the vibration isolation. The light will need to be fibre optic launched onto the vibration isolation table to ensure no beam pointing or Doppler shift issues arise. Any amplitude noise caused by polarisation noise in the fibre will have to be dealt with by the intensity servos.

Finally, to completely rule out any possible common mode rejection, one of the systems will be placed on a second isolation table. In this case the two beams of light from the laser are fibred onto the two independent isolation platforms. Unfortunately using two isolation platforms will introduce Doppler shifts between the systems. Ideally these Doppler shifts would be corrected for. However solution to such an effect may well lie outside the scope of this project, but could form in interesting follow on project.

5.3 Future research

Although it is difficult to attempt to guess what will be in the future (as the research discussed in this report has barely begun), it might be worth pointing to a couple of ideas.

Once a stable Nd:YAG laser has been built, it will be desirable to try to stabilise noisier systems, such as extended cavity diodes or Ti:Sapphire lasers. These lasers can then be used optical frequency standard experiments.

The transfer of a stable frequency from one point to another needs to be investigated. This delivery of narrow linewidth light is certainly an important topic. This could be either in free space or via fibre optic cable.

Experiments relying on good Allan variance data between two systems (e.g. a test of special relativity) could be performed (*e.g.* see Hall [16]).

5.4 Linewidth: How low can you go?

This report has identified the perturbation limits to the production of a sub-hertz laser system. These include: seismic, acoustic and thermal perturbations to the cavity, as well as laser beam pointing and power fluctuations. It also identifies accuracy of lock issues such as optical feedback, modulation accuracy and mirror birefringence.

It is interesting to consider what will be the ultimate limiting factor to achieving a narrow linewidth laser, and how low this linewidth could be. At present NIST have a world record laser linewidth of less than 0.6 Hz. It is not known what limits the laser linewidth at this low level. An understanding may be forthcoming when this latest work is published.

To get a flavour of how much further there is to go it is amusing to consider the quantum limit. The quantum limit for this case was rigorously derived by Caves [102] and is also well explained by Bergquist*et.al.* [37]. "...the measurement of the length of the cavity to which the laser is locked brings about the inevitable competition between measurement precision and the perturbation of the measurement system. The measurement precision can be improved by a factor of $1/\sqrt{N}$ by increasing the number of N (signal) photons in the measurement interval, whereas the shot noise of radiation pressure on the mirrors increases as $\sqrt{N}$. The optimum flux of signal photons....is given when both effects are in equal magnitude"[37]. The result of this analysis is an optimum power for the measurement (which depends on the mirror finesse and other factors) and an uncertainty in the measurement of the length of the cavity (for optimum power) of

$$\Delta L = \sqrt{\frac{4\hbar B}{(2\pi f_1)^2 m}} \tag{5.1}$$

where B is the measurement bandwidth, f_1 the resonant frequency for longitudinal vibrations of the cavity and m the mass of the cavity. Now the frequency fluctuation corresponding to this length uncertainty is $\delta\nu = \nu_0 \Delta L/L$. Therefore applying equation 2.10, gives a linewidth of

$$\Delta\nu_l = \frac{\nu_0^2 \hbar}{\pi L^2 f_1^2 m} \tag{5.2}$$

For the standard cavity parameters this gives a linewidth of $\Delta\nu_l = 1 \times 10^{-12}$ Hz!!!!

Fortunately we have already discovered an atomic reference of suitably narrow linewidth to go with this quantum limited laser; the electric-octupole transition in Yb^+ has a linewidth of only 0.5×10^{-9} Hz [124].

This is all rather fanciful though, since there is an interrogation time problem. Heisenberg's uncertainty principle gives a time, τ, to realise a linewidth of $\tau > 1/2\pi\Delta\nu_l$. So to realise the quantum limited laser linewidth would require an interaction time of 5000 years. Even to reach the rather less ambitious electric octupole linewidth would require a 10 year interaction time.

Returning to more realistic limits, though still speculating, it seems possible that a materials problem may become the practical limit. The realisable linewidth from a ULE cavity may be limited by creep. Given that these cavities have creep rates in the 0.5 Hz/s range, and that maybe this creep is predicable to 100th of this, then a best linewidth of 5 mHz might be possible. (This is already a factor of 10 better than the best servo lock though). This also assumes that the creep happens in a smooth way; the case of Zerodur has already shown that this is not necessarily the case.

If the ULE material really were the limiting factor, then cavities with better dimensional and thermal stabilities must be found. One such possibility is cryogenically cooled sapphire, mentioned in the introduction. However the difficulties of the cryogenic environment (noisy helium and nitrogen fills, positional instability of the cavity ....) means that this technology will have advance considerably before the absolute linewidths of these systems match even current ULE systems.

In conclusion, throughout this report there has been numerous reference to very small quantities. It is easy to get complacent about what a part in 10^{15} means, or what environmental control one needs to achieve a sub-hertz linewidth.

> ...if the spacer for an optical cavity were the Earth, a human hair added to the diameter would cause a frequency shift of about 300 Hz!
>
> —Berquist *et al* [37]

Symbols

This glossary lists the symbols *commonly* used within the text. Other symbols not listed are defined locally within the text. Δ and δ are used infront of a symbol to represent a small change in that quantities value.

y	Dimensionless frequency
ν, ν_0	Laser frequency [Hz]
ϕ	Phase [rad]
f	Fourier frequency [Hz]
$S_x(f)$	Spectral density of x, [(dimensions of x)^2/Hz]
$S_\nu(f)$	Power Spectral Density of frequency fluctuations [Hz^2/Hz]
$S_\phi(f)$	Power Spectral Density of phase fluctuations [/Hz]
$S_y(f)$	Power Spectral Density of dimensionless frequency fluctuations [/Hz]
P	Power [W] (usually laser power incident on cavity)
V	Voltage [V]
B	Bandwidth [Hz]
β	Modulation index [rad]
$\sigma_y(\tau)$	root Allan variance of dimensionless frequency fluctuations
τ	averaging time [s]
$\Delta\nu_l$	laser linewidth (FWHM) [Hz]
$\delta\nu$	frequency fluctuation [Hz]
λ	wavelength [m]
$\Delta\nu_c$	reference cavity linewidth [Hz]
$h = 2\pi\hbar$	Planck constant [Js]
c	speed of light [ms^{-1}]
g	acceleration due to gravity [ms^{-2}]
e $\lvert e\rvert$	electron charge [C]
I	Detector photocurrent [A]
$J_0(M)$	Zeroth order Bessel function, with (rf) modulation index M
η_d	detector quantum efficiency

ν_{FSR}	Free spectral range [Hz]
$\mathcal{F}$	Finesse
L	cavity length [m]
ΔL	change in cavity length [m]
$\sqrt{S_a(f)}$	Root Spectral Density of acceleration fluctuations [$\mathrm{ms^{-2}}/\sqrt{\mathrm{Hz}}$]
$\sqrt{S_v(f)}$	Root Spectral Density of velocity fluctuations [$\mathrm{ms^{-1}}/\sqrt{\mathrm{Hz}}$]
$\sqrt{S_d(f)}$	Root Spectral Density of displacement fluctuations [$\mathrm{m}/\sqrt{\mathrm{Hz}}$]
T	Temperature [K] *or* Transmissibility
γ	Damping factor
f_0	Resonant frequency of pendulum or spring [Hz]
f_1	First longitudinal or bending frequency of a bar [Hz]
E	Youngs modulus [$\mathrm{Nm^{-2}}$]
A	Area [$\mathrm{m^2}$]
m	Mass [kg]
ρ	Density [$\mathrm{kgm^{-3}}$]

References

[1] M. W. Hamilton, Contemporary Phys. **30**, 21–33 (1989).

[2] R. W. P. Drever *et al.*, Appl. Phys. **B31**, 97–105 (1983).

[3] J. Hough *et al.*, Appl. Phys. **B33**, 179–185 (1984).

[4] J. Helmcke, S. A. Lee, and J. L. Hall, Appl. Opt. **21**, 1686–1694 (1982).

[5] C. Salomon, D. Hills, and J. L. Hall, J. Opt. Soc. Am. **B5**, 1576–1587 (1988).

[6] J. L. Hall, in *Quantum Optics IV*, edited by J. D. Harvey and D. F. Walls (Springer-Verlag, 1986), pp. 273–284.

[7] J. L. Hall, SPIE **1837**, 2–15 (1992).

[8] M. Zhu and J. L. Hall, J. Opt. Soc. Am. **B10**, 802–816 (1993).

[9] D. Hills and J. L. Hall, Rev. Sci. Instrum. **58**, 1406–1412 (1987).

[10] D. Hils, J. E. Faller, and J. L. Hall, Rev. Sci. Instrum. **57**, 2532–2534 (1986)

[11] J. L. Hall *et al.*, in *Laser Spectroscopy VIII*, edited by W. Persson and S. Svanberg (Springer-Verlag, 1987), pp. 376–380.

[12] M. D. Rayman, C. G. Aminoff, and J. L. Hall, J. Opt. Soc. Am. **B6**, 539–549 (1989).

[13] D. Hils and J. L. Hall, pp. 162–173, in [117].

[14] L.-S. Ma *et al.*, SPIE **2378**, 165–175 (1995)

[15] J. L. Hall and T. W. Hänsch, Opt. Lett. **9**, 502–504 (1984)

[16] D. Hils and J. L. Hall, Phys. Rev. Lett. **64**, 1697–1700 (1990).

[17] L.-S. Ma *et al.*, Opt. Lett. **19**, 1777-1779 (1994)

[18] J. L. Hall, Laser Physics **4**, 306–318 (1994).

[19] M. Zhu and J. L. Hall, in *Proc. 1992 IEEE frequency control symposium* (PUBLISHER, 1992), pp. 44–55.

[20] K. Nakagawa, T. Katsuda, and M. Ohtsu, Appl. Phys. **B60**, 489–494 (1995).

[21] C.-H. Shin and M. Ohtsu, Opt. Lett. **15**, 1455–1457 (1990).

[22] K. Nakagawa, M. Kourogi, and M. Ohtsu, Opt. Lett. **17**, 934–936 (1992)

[23] M. Kourogi, C.-H. Shin, and M. Ohtsu, IEEE Photonics Technol. Lett. **3**, 496–498 (1991).

[24] M. Ohtsu *et al.*, J. Appl. Phys. **73**, R1–R17 (1993).

[25] M. Ohtsu, *Coherent quantum optics and technology* (KTK Scientific Publishers, 1992).

[26] K. Nakagawa *et al.*, Opt. Commun. **109**, 446–450 (1994)

[27] K. Nakagawa *et al.*, Appl. Opt. **33**, 6383–6386 (1994)

[28] K. Nakagawa *et al.*, SPIE **1837**, 40–48 (1992)

[29] M. Musha *et al.*, Opt. Commun. **140**, 323–330 (1997)

[30] M. Musha, K. Nakagawa, and K. Ueda, Opt. Lett. **22**, 1177–1179 (1997)

[31] N. Uehara and K. Ueda, Laser Phys. **4**, 396–400 (1994)

[32] N. Uehara and K. Ueda, Opt. Lett. **18**, 505–507 (1993).

[33] K. Ueda and N. Uehara, SPIE **1837**, 336–345 (1992).

[34] N. Uehara and K. Ueda, Appl. Phys. **B61**, 9–15 (1995).

[35] N. Uehara and K. Ueda, Jpn. J. Appl. Phys. **33**, 1628–1633 (1994)

[36] J. C. Bergquist *et al.*, in *Light Induced Kenetic Effects on Atoms, Ions and Molecules*, edited by L. Moi *et al.* (ETS Editrice, 1991), pp. 291–299.

[37] J. C. Bergquist, W. M. Itano, and D. J. Wineland, in *Proceedings of the International School of Physics ≪Enrico Fermi≫, Course CXX*, edited by T. W. Hänsch and M. Inguscio (North Holland, 1994), pp. 359–376, also in [116].

[38] B. C. Young *et al.*, in *1998 Confreence on Precision Electromagnetic Measurement Digest*, edited by T. L. Nelson (IEEE, 1998), pp. 153–154.

[39] S. Seel *et al.*, in *Proceedings of the 5th Symposium on Frequency Standards and Metrology*, edited by J. C. Bergquist (World Scientific, 1996), pp. 347–352.

[40] R. Storz *et al.*, in *Proc. Workshop on Frequency Standards based on Laser Manipulated Atoms and Ions*, edited by J. Helmcke and S. Penselin (PTB, 1996), Vol. PTB-Opt 51, pp. 97–104.

[41] R. Storz *et al.*, Opt. Lett. **23**, 1031–1033 (1998).

[42] S. Seel *et al.*, Phys. Rev. Lett. **78**, 4741–4744 (1997).

[43] G. Ruoso *et al.*, Opt. Commun. **133**, 259–262 (1997).

[44] S. Schiller *et al.*, SPIE **2378**, 138–146 (1995).

[45] P. F. Micelson *et al.*, pp. 15–56, in [118]

[46] C. Bradaschia *et al.*, pp. 110–135, in [118]

[47] D. E. McClelland, pp. 207–210, in [118]

[48] P. R. Saulson, pp. 248–270, in [118].

[49] J. Ferreirinho and D. G. Blair, pp. 333–367, in [118].

[50] D. G. Blair *et al.*, pp. 415–426, in [118]

[51] K. Tsubono *et al.*, Rev. Sci. Instrum. **64**, 2237–2240 (1993).

[52] P. J. Veitch, Rev. Sci. Instrum. **62**, 140–155 (1991).

[53] A. Rüdiger *et al.*, Optica Acta **28**, 641-658 (1981).

[54] K. Skeldon and J. Hough, Rev. Sci. Instrum. **66**, 2760–2762 (1995).

[55] N. Kanda, M. A. Barton, and K. Kuroda, Rev. Sci. Instrum. **65**, 3780–3783 (1994).

[56] M. A. Barton and K. Kuroda, Rev. Sci. Instrum. **65**, 3775–3779 (1994)

[57] P. J. Veitch *et al.*, Rev. Sci. Instrum. **67**, 633–640 (1996).

[58] S. Kawamura, A. Abramovici, and M. E. Zucker, Rev. Sci. Instrum. **68**, 223–229 (1997).

[59] J. Winterflood and D. G. Blair, Phys. Lett. **A222**, 141–147 (1996)

[60] D. G. Blair *et al.*, Phys. Lett. **A193**, 223–226 (1994).

[61] D. B. Newell *et al.*, Rev. Sci. Instrum. **68**, 3211–3219 (1997)

[62] T. L. Boyd and H. J. Kimble, Opt. Lett. **16**, 808–810 (1991)

[63] T. Day, E. K. Gustafson, and R. L. Byer, IEEE J. Quantum Electron. **28**, 1106–1117 (1992).

[64] L. Hilico *et al.*, Rev. Sci. Instrum. **65**, 3628–3633 (1994)

[65] C. T. Taylor, M. Notcutt, and D. G. Blair, Rev. Sci. Instrum. **66**, 955–960 (1995).

[66] K. Szymaniec *et al.*, Opt. Commun. **144**, 50–54 (1997)

[67] D. Touahri, Ph.D. thesis, Universite de Paris, 1996.

[68] A. Celikov *et al.*, Opt. Commun. **107**, 54–60 (1994).

[69] V. Vassiliev *et al.*, Opt. Lett. **23**, 1229–1231 (1998).

[70] T. Kurson, J. Ishikawa, and N. Ito, Appl. Phys. **B63**, 265–275 (1996).

[71] Y. Yoshikuni, in *Coherence, Amplification and Quantum effects in Semicondutor Lasers*, edited by Y. Yamamoto (John Wiley, 1991), pp. 97–136.

[72] C. Bianchini *et al.*, Appl. Phys. **B66**, 407–410 (1998).

[73] A. S. Zibrov *et al.*, Appl. Phys. **B59**, 327-331 (1994)

[74] N. M. Sampas, E. K. Gustafson, and R. L. Byer, Opt. Lett. **18**, 947–949 (1993).

[75] T. Day, Ph.D. thesis, Stanford University, 1990.

[76] D. Shoemaker *et al.*, Opt. Lett. **14**, 609–611 (1989).

[77] G. Cantatore *et al.*, Rev. Sci. Instrum. **66**, 2785–2787 (1995)

[78] A. M. DeRiva *et al.*, Rev. Sci. Instrum. **67**, 2680–2684 (1996)

[79] R. Kallenbach *et al.*, Opt. Commun. **70**, 56–60 (1989)

[80] J. Helmcke *et al.*, Appl. Phys. **B43**, 85–91 (1987)

[81] J. Dirscherl *et al.*, Opt. Commun. **91**, 131–139 (1992)

[82] G. A. Kerr *et al.*, Appl. Phys. **B37**, 11–16 (1985).

[83] S. Timoshenko, D. H. Young, and W. Weaver, *Vibration Problems in Engineering*, (Wiley, 1974, 4th ed.).

[84] M. L. Boas, *Mathematical Methods in the Physical Sciences*, (Wiley, 1983, 2nd ed.)

[85] L. L. Beranek, *Noise and Vibration Control*, (McGraw-Hill, 1988, revised ed.

[86] J. E. Gordon, *Structures or why things don't fall down* (Pengiun, 1986).

[87] A. Giazotto, D. Passuello, and A. Stefanini, Rev. Sci. Instrum. **57**, 1145–1151 (1986).

[88] P. R. Saulson, Rev. Sci. Instrum. **55**, 1315–1320 (1982).

[89] N. A. Robertson *et al.*, J. Phys. **E15**, 1101–1105 (1982).

[90] D. L. Platus, SPIE **1619**, 44–54 (1991).

[91] J. P. D. Hartog, Trans. ASME **53**, 107–115 (1931)

[92] F. Cleva and J. P. Coulon, Acousto-optic effects in air, virgo report NTS95-011

[93] H. P. Layer, Appl. Opt. **18**, 2947–2949 (1979)

[115] D. Jacob *et al.*, Opt. Lett. **20**, 671–673 (1995)

[116] *NIST Technical Note* (US Govt Printing Office, 1996), Vol. 1380.

[117] in *Proceedings of the 4th Symposium on Frequency Standards and Metrology*, edited by A. De Marchi (Springer-Verlag, 1989).

[118] in *Gravitational Astronomy: Instrument design and astrophysical prospects*, edited by D. E. McClelland and H.-A. Bachor (World Scientific, 1991).

[119] B. C. Young – private communication

[120] J. C. Bergquist – private communication.

[121] Newport website at: http://www.newport.com.

[122] Minus k technology, 420 S. Hindry Ave.,Unit B, Inglewood, California 90301, USA.

[123] Integrated Dynamic Engineering Inc., 84 Cummings Park, Eoburn, MA 01801, USA.

[124] M. Roberts *et al.*, Phys. Rev. Lett. **78**, 1876–1879 (1997)

[94] F. J. Harris, Proc. IEEE **66**, 51–83 (1978).

[95] *NIST Technical Note "Characterisation of Clocks and Oscillators"* (US Govt Printing Office, 1990), Vol. 1337.

[96] D. S. Elliot, R. Roy, and S. J. Smith, Phys. Rev. **A26**, 12–18 (1982).

[97] D. W. Allan, Proc. IEEE **54**, 221–230 (1966).

[98] D. W. Allan, IEEE Trans. Ultrason. Ferroelec. Freq. Contr. **34**, 647–654 (1987).

[99] J. A. Barnes *et al.*, IEEE Trans. Instrum. Meas. **20**, 105–120 (1971).

[100] J. Rutman and D. F. Walls, Proc. IEEE **79**, 952–959 (1991)

[101] D. B. Percival, Proc. IEEE **79**, 961–972 (1991).

[102] C. M. Caves, Phys. Rev. Lett. **45**, 75–79 (1980).

[103] B. I. Bleaney and B. Bleaney, *Electricity and Magnetism (volume 2)*, (Oxford University Press, 1989, 3rd ed ed.).

[104] *ULE Titanium Silicate*, Advanced Produacts Department, Corning Glass Works, Corning, New York 14831, 1985.

[105] *Zerodur*, Schott Glaswerk, Geschaftsbereich Optik, Verkauf Optisches Glas, Hattenbergstrasse10, W6500 Mainz.

[106] G. W. C. Kaye and T. H. Laby, *Tables of Physical and Chemical Constants*, (Longman, 1986, 15th ed.).

[107] N. A. Waterman and M. F. Ashby, *Elsevier Materials Selector Vol I* (Elsevier Science Publishers, 1991).

[108] L. R. G. Treloar, *The Physics of Rubber Elasticity*, (Clarendon Press, 1975, 3rd ed.)

[109] R. Taylor, *Noise*, (Penguin, 1979, 3rd ed.)

[110] O. Mor and A. Arie, IEEE J. Quantum Electron. **33**, 532–540 (1997).

[111] E. Calloni *et al.*, Opt. Commun. **142**, 50–54 (1997).

[112] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback control of dynamic systems*, (Addison-Wesley, 1994, 3rd ed.).

[113] H. R. Telle, in *Proc. 11th European Frequency and Time Forum*, edited by FRSM (Swiss Foundation for Research in Microtechnology, 1998), pp. 109–111.

[114] F. Brandi *et al.*, Appl. Phys. **B65**, 351–355 (1997)