Nelson, John (2012) *Novel optical resonators*. PhD thesis.

Novel Optical Resonators

John Nelson
BSc(HONS)

Submitted in fulfilment of the requirements for the Degree of PhD

School of Physics & Astronomy
College of Physical Sciences
University of Glasgow

June 13, 2011

Abstract

Optical resonators – two mirrors facing each other, separated by a distance – are a very well studied technology. However, even such a well understood technology can sometimes present surprises. The first part of this thesis investigates the surprising properties of some canonical optical resonators. The basic properties of resonators are introduced. The imaging properties of stable and unstable resonators are examined.

The second part of this thesis examines the potential use of grating-coupled cavities in gravitational wave detectors and describes an experiment carried out on a 10 m prototype of such a cavity.

Contents

List of Figures

Acknowledgements

I'd like to thank my supervisors, Johannes Courtial and Ken Strain, for all their help. I'd also like to thank my colleagues Bryan Barr, Sabina Huttner, Matt Edgar and Mike Plissi.

CHAPTER 1

Introduction

1.1 Resonators and interferometers - a very brief introduction

Near the end of the 19th century, Alfred Perot and Charles Fabry, both from the University of Marseilles, published a paper [1] in which they described a new device: an interferometer, where the interference was between waves successively reflected from two parallel, silvered glass plates (see figure 1.1). This device became known as a Fabry-Perot interferometer, and is an example of an optical resonator.

A few years earlier, Albert Michelson had invented the interferometer which

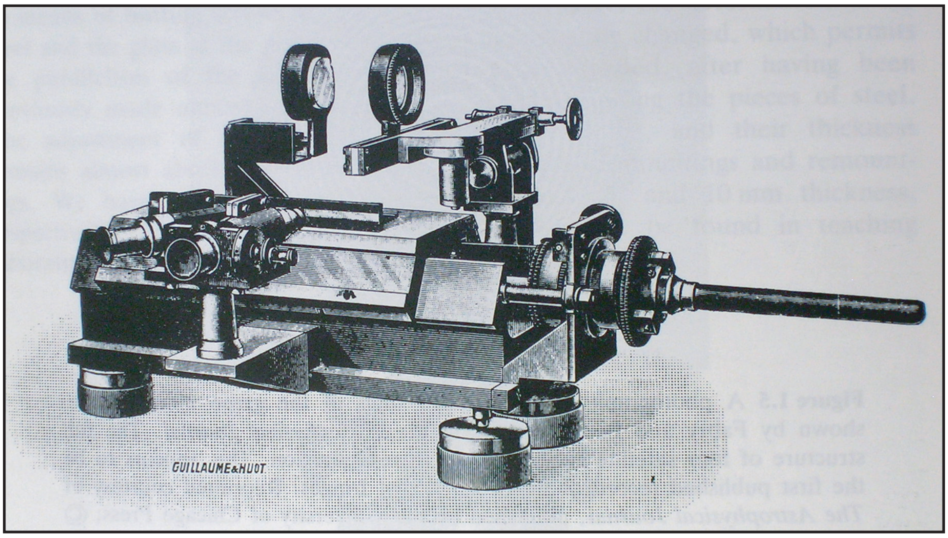

Figure 1.1: An interferometer of the same design as in [1], and built by M Jobin for Fabry and Perot in 1901.

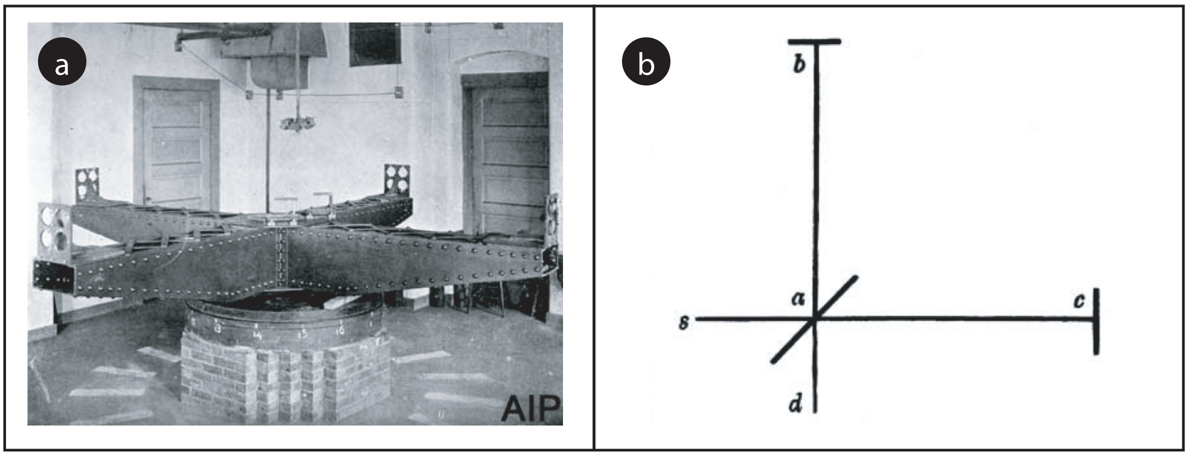

Figure 1.2: A photo (a) of a Michelson interferometer – in this case the Morley-Miller apparatus which repeated the original measurements between 1902 and 1906 (courtesy of AIP). A diagram (b) of a Michelson interferometer, taken from Michelson's 1887 paper [2].

bears his name (see figure 1.2). In this interferometer, coherent light from the source, s, falls onto a beamsplitter, a. Each of the beams is then reflected back to the beamsplitter (from the mirrors b and c), and the resulting superposition of these beams can be observed at d.

In order to improve their sensitivity, modern interferometric gravitational-wave detectors often combine the two devices. Instead of the basic Michelson interferometer layout shown in figure 1.2, the mirrors *b* and *c* are replaced with resonant Fabry-Perot interferometers [3].

This thesis details the research I have undertaken into novel optical resonators, and also into the novel properties of well known resonators.

1.2 Structure of this thesis

This thesis divides neatly into two main strands: simulations of the fundamental properties of novel and not-so-novel optical resonators; and experimental work towards the realisation of a diffractively-coupled suspended optical resonator.

Chapter 2 introduces the concept of an optical resonator and describes their division into two main classes – stable and unstable.

The concept of *ray-tracing* is introduced in chapter 3, and is used to illustrate the imaging properties of optical resonators (and also, briefly, optical billiards). Its usefulness in simulating novel optical systems is also explored.

Chapter 4 describes resonators from a new perspective – that of the Fox-Li method for determining the eigenmode of an arbitrary resonator. It then describes the application of these techniques to canonical optical resonators, and to a novel type of grossly spherically aberrated resonator which may be useful in high power applications. It concludes with a description of the surprising fractal properties of some unstable canonical resonators.

Chapter 5 gives a very brief introduction to some of the concepts of General Relativity and describes the solutions which lead to the proposed existence of gravitational waves. It then describes the astrophysical sources and anticipated strengths of gravitational waves. Finally it discusses the technological challenges of detecting these waves.

Chapter 6 introduces some of the theory necessary to describe diffraction gratings. It takes a new perspective on conventional optical resonators and then develops the theory of grating coupled resonators. It also discusses the challenges of controlling suspended optical cavities.

Chapter 7 describes experimental work performed using a 10 m prototype grating-coupled cavity, where – for purposes of vibrational isolation – all of the optical elements of the cavity are suspended on multi-stage pendulum systems.

Chapter 8 concludes with a discussion of the implications of results and suggestions for further investigation.

CHAPTER 2

Transverse Eigenmodes of Optical Resonators

2.1 Introduction

In this chapter we will give a brief introduction to the theory of optical resonators, with a more in-depth examination of a particular subset of general optical resonators, namely the stable two-mirror Gaussian optical resonators. Then we will explore methods for calculating the eigenmodes of optical resonators, concentrating on the Fox-Li method [4]. We have used this method to find the eigenmodes of a variety of novel optical resonators; results will be presented in the following chapter. We begin, however, by considering two

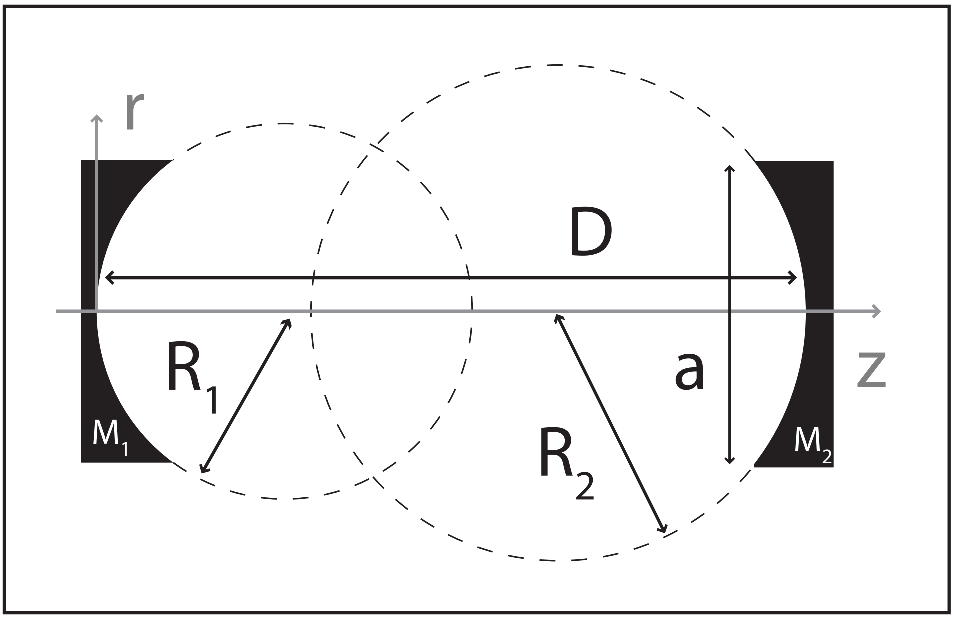

Figure 2.1: A two mirror optical resonator known, in this form, as a canonical optical resonator. Two circular mirrors M_1 and M_2 (with radii of curvature R_1 and R_2, respectively) are separated by a distance D. The diameter of the mirrors (or, equivalently, the aperture onto the mirrors) is a.

very basic questions: what do we mean by an “optical resonator”? And what do we mean by its “eigenmode”?

Optical resonators and their eigenmodes

The phrase “optical resonator” usually conjures up an image somewhat like that displayed in figure 2.1: two spherical mirrors, aligned on the same axis, separated by a distance, D. This family of resonators (the so called *canonical optical resonators* [5]) has achieved this pre-eminent recognition due to the fact that these are the resonators which can most easily be analysed purely

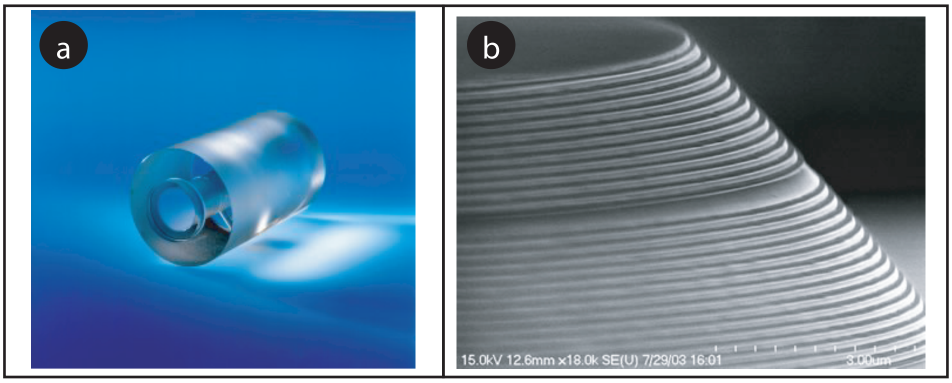

Figure 2.2: Two examples of optical resonators. In (a) we see a Fabry-Perot etalon [www.npl.co.uk/]. In (b) a Vertical-Cavity Surface-Emitting Laser [www.asu.edu/].

analytically, either by the application of the optics Gaussian beams [6] or by their imaging properties [7].

The canonical optical resonators represents one subset of an endless variety of possible resonators [8, 9, 10]. Indeed, any system where an electromagnetic field undergoes a "round-trip" periodic influence can be considered an optical resonator. That is, the light undergoes some periodic focussing. We can go further and state that whenever light undergoes a periodic *influence* — even when that influence is only an aperture — we have an optical resonator [6]. Figure 2.2 shows some optical resonator systems.

An interesting property of this type of system is the existence of transverse *eigenmodes*. These are fields which, when propagated around a resonator for one round trip, retain the same shape. More precisely we can say that although the overall intensity may — and in all real situations will — decrease, the relative intensities of the field will stay the same. This is also true of the relative phase, although the entire field may undergo some phase shift.

It is when we attempt to predict *ab initio* what form these eigenvalues might take that problems arise. While it is simple to write down the equations describing an optical resonator, an analytical solution to these equations is possible in only a few cases — for instance with the Gaussian eigenmodes of cavities with spherical mirrors. Because of this, numerical methods using computers have become very important in discovering the form of these eigenmodes. The Fox-Li method (see chapter 4) is one of the most important and widely used of these numerical procedures.

2.2 Optical Resonators

In this section we introduce some of the basics properties of optical resonators. We begin by examining in detail the properties of Gaussian resonators – in particular, the criteria for a Gaussian resonator to be stable. Although the majority of this thesis does not concern itself with Gaussian resonators, the ease with which these resonators can be analyzed will give us an intuitive insight into — and a reference for — the properties of more exotic resonators.

Gaussian Resonators

We begin by considering canonical optical resonators. This type of optical resonator consists of two curved (or flat) mirrors facing each other, separated by a distance. Resonators of this kind have many analogies to periodic lens waveguides, and we will exploit these similarities in our initial, ray optical description of optical resonators.

Consider a typical optical resonator consisting of two mirrors M_1 and M_2, separated by a distance D. Mirror M_1 has a radius of curvature R_1 and a focal length $f_1 = \frac{R_1}{2}$, and similarly for mirror M_2. We could represent one trip of a light ray around this resonator (for instance from M_1 to M_2 and then back to M_1) as the passage of a light ray through two lenses of a periodic lens waveguide, where the lenses have focal lengths $f_1 = R_1/2$ and $f_2 = R_2/2$ and are separated by the same distance (D) as the mirror of the optical resonator (see figure 2.3). Finally, before we begin to analyze the system it will be useful to subsume the curvature of each mirror (or the focal length) and their separation into a single, dimensionless parameter. Conventionally called *g-parameters*, these are defined as

$$g_1 \equiv 1 - \frac{D}{R_1} = 1 - \frac{D}{2f_1}, \tag{2.1}$$

$$g_2 \equiv 1 - \frac{D}{R_2} = 1 - \frac{D}{2f_2}. \tag{2.2}$$

For much of the following discussion we will assume that the mirrors of the optical resonator are of unlimited extent. When we come to consider realis-

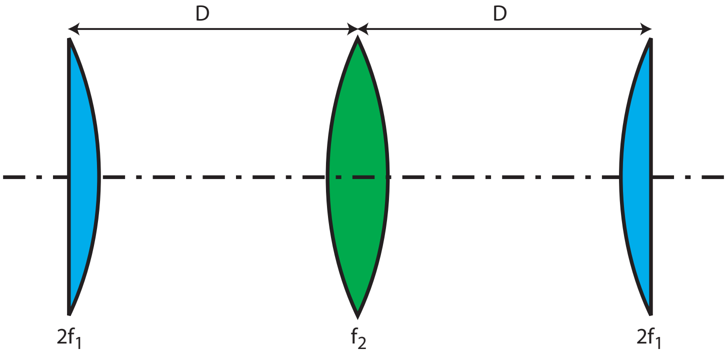

Figure 2.3: One section of a periodic lens waveguide equivalent for a two mirror optical resonator.

tic cavities where the mirrors are of finite size it is useful to define another dimensionless parameter, the Fresnel number, $\mathcal{N}$:

$$\mathcal{N} = \frac{a^2}{D\lambda}, \tag{2.3}$$

where $2a$ is the diameter of the mirror.

Stability Criteria

The first question we might ask is whether or not a particular optical resonator is stable: that is, after a large number of round trips will a light ray escape from the resonator. Representing this round trip as a periodic lens waveguide (figure 2.3) — starting midplane in the first lens — we consider the ray matrices of the total system:

$$\begin{bmatrix} r_1 \\ \theta_1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -\frac{1}{2f_1} & 1 \end{bmatrix} \begin{bmatrix} 1 & D \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -\frac{1}{f_2} & 1 \end{bmatrix} \begin{bmatrix} 1 & D \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -\frac{1}{2f_1} & 1 \end{bmatrix} \begin{bmatrix} r_0 \\ \theta_0 \end{bmatrix} \tag{2.4}$$

So the ray transfer matrix for one round trip is,

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} 1-\frac{D}{f_2} & d\left(2-\frac{D}{f_2}\right) \\ -\frac{1}{f_1}-\frac{1}{f_2}\left(1-\frac{D}{f_1}\right) & 1-\frac{2D}{f_1}-\frac{D}{f_2}+\frac{D^2}{f_1 f_2} \end{bmatrix}. \tag{2.5}$$

So, in general $r_{n+1} = Ar_n + B\theta_n$ and $\theta_{n+1} = Cr_n + D\theta n$. Therefore we arrive at,

$$\theta_n = \frac{1}{B}(r_{n+1} - Ar_n), \tag{2.6}$$

and

$$\theta_{n+1} = \frac{1}{B}(r_{n+2} - Ar_{n+1}), \tag{2.7}$$

which leads to,

$$Cr_n + D\theta_n = \frac{1}{B}(r_{n+2} - Ar_{n+1}). \tag{2.8}$$

If we substitute the θ_n terms we get,

$$r_{n+1} - (A + D)r_{n+1} + (AD - BC)r_s = 0. \tag{2.9}$$

It can be shown that $(AD - BC) = 1$. The simplest method being to note that the 4 original ray transfer matrices each had a determinant of 1, therefore the product matrix in equation 2.4 also has a determinant of 1. With the further substitution that

$$b = (A + D)/2 = \left(1 - \frac{d}{f_2} - \frac{d}{f_1} + \frac{d^2}{2f_1 f_2}\right) \tag{2.10}$$

we get,

$$r_{n+2} - 2br_{n+1} + r_n = 0, \tag{2.11}$$

which is equivalent to a second order differential equation. Thus, we are led to try a solution of the form $r_n = r_0 \exp{(inq)}$, so that

$$e^{2iq} - 2be^{iq} + 1 = 0, \tag{2.12}$$

and

$$e^{iq} = b \pm i\sqrt{1-b^2} = e^{\pm i\theta}, \tag{2.13}$$

where $\cos\theta = b$.

The general solution is then given by a linear superposition of $\exp(in\theta)$ and $\exp(-in\theta)$ terms. Equivalently,

$$r_n = r_{max} \sin{(n\theta + \alpha)}, \tag{2.14}$$

where $r_{max} = r_0/\sin\alpha$. If we then ask, what is the displacement of the initial ray after traversing n periods of the waveguide (n round-trips of the resonator)? After an arbitrary number of round trips the ray will remain inside the waveguide only if $|b| \leq 1$, i.e.,

$$-1 \leq \left(1 - \frac{d}{2f_1}\right)\left(1 - \frac{d}{2f_2}\right) \leq 1. \tag{2.15}$$

If we express this in terms of the previously introduced g-parameters, we have,

$$0 \leq g_1 g_2 \leq 1, \tag{2.16}$$

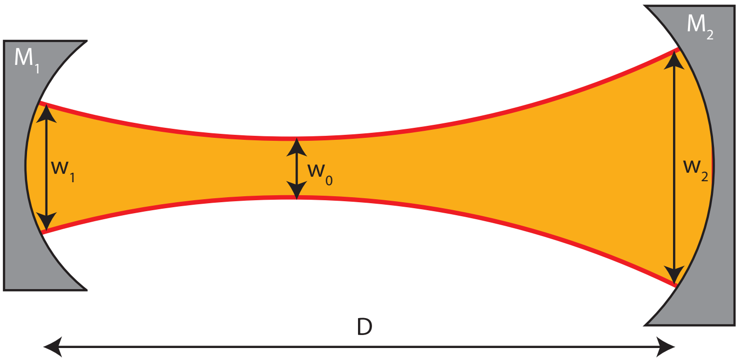

Figure 2.4: Mode properties of a stable optical resonator.

and it is under these conditions that a canonical optical resonator is stable.

Properties of Stable Optical Resonators

Restricting ourselves to the stable optical resonators ($0 \leq g_1 g_2 \leq 1$, see figure 2.5), we now ask what properties we might expect the lowest order modes of these systems to have. The lowest order modes of these systems are Gaussian beams, and the properties we are interested in are the spot sizes at the mirrors, w_1, w_2 and the waist size w_0 (figure 2.4).

In general, the easiest way to solve this problem is to start with the Gaussian beam in free space and then ask what mirrors match the curvature of the wavefronts at the required distances from the beam waist.

With any two mirrors, separated by a distance D, and with arbitrary radii of curvature R_1 and R_2 — as long as these radii comply with the requirement

$0 \leq g_1 g_2 \leq 1$ — we can reverse this solution process and discover what waist and spot sizes a particular resonator gives rise too.

Knowing the expression for the Rayleigh range [11], $z_R = \frac{\pi w_0^2}{\lambda}$, we can then use the gaussian beam formula,

$$R(z) = z + \frac{z_R^2}{z} \tag{2.17}$$

to find the wavefront curvature. Requiring this to match the mirror curvature leads us to the expressions:

$$z_1 = -D\frac{g_2(1-g_1)}{g_1+g_2-2g_1g_2}, \tag{2.18}$$

$$z_2 = D\frac{g_1(1-g_2)}{g_1+g_2-2g_1g_2}, \tag{2.19}$$

$$z_R = D^2\frac{g_1g_2(1-g_1g_2)}{(g_1+g_2-2g_1g_2)^2}. \tag{2.20}$$

It follows that the expressions for waist and spot sizes for any resonator are,

$$w_1^2 = \frac{D\lambda}{\pi}\sqrt{\frac{g_2}{g_1(1-g_1g_2)}}, \tag{2.21}$$

$$w_2^2 = \frac{D\lambda}{\pi}\sqrt{\frac{g_1}{g_1(1-g_1g_2)}}, \tag{2.22}$$

$$w_0^2 = \frac{D\lambda}{\pi}\sqrt{\frac{g_1g_2(1-g_1g_2)}{(g_1+g_2-2g_1g_2)^2}}. \tag{2.23}$$

We can considerably simplify these expressions if we restrict ourselves to a discussion of *symmetric* optical resonators, i.e. those where the focal lengths of both mirrors are equal, leading to $g_1 = g_2 = g$. Making this substitution, we find

$$w_0 = \sqrt{\frac{D\lambda}{2\pi}} \left(\frac{1+g}{1-g} \right)^{\frac{1}{4}}, \tag{2.24}$$

$$w_1 = w_2 = \sqrt{\frac{D\lambda}{\pi}} \left(\frac{1}{1-g^2} \right)^{\frac{1}{4}}. \tag{2.25}$$

These symmetric optical cavities lie on the line $g_1 = g_2$ on the stability diagram, between $g = -1$ (concentric resonator) and $g = 1$ (plane-plane resonator). It will be a useful reference point in later sections if we think about what happens to w_0, w_1 and w_2 as we travel along this line of symmetric resonators in the stable region (see figure 2.5).

One last point should be made. In this analysis we have assumed mirrors of infinite extent and thus that the edge of the mirror has no effect on the mode. In reality the transition between low and high loss modes is more gradual than the stability diagram might suggest. In fact the sharpness of the transition [12] increases with increasing Fresnel number, i.e. for $a^2 \gg D\lambda$ it approximates the case of mirrors of infinite extent.

These solutions to equation 2.17, form the lowest order set of *transverse eigenmodes* for the Gaussian resonators. We will discuss the transverse eigenmodes of optical resonators in greater depth in chapter 4.

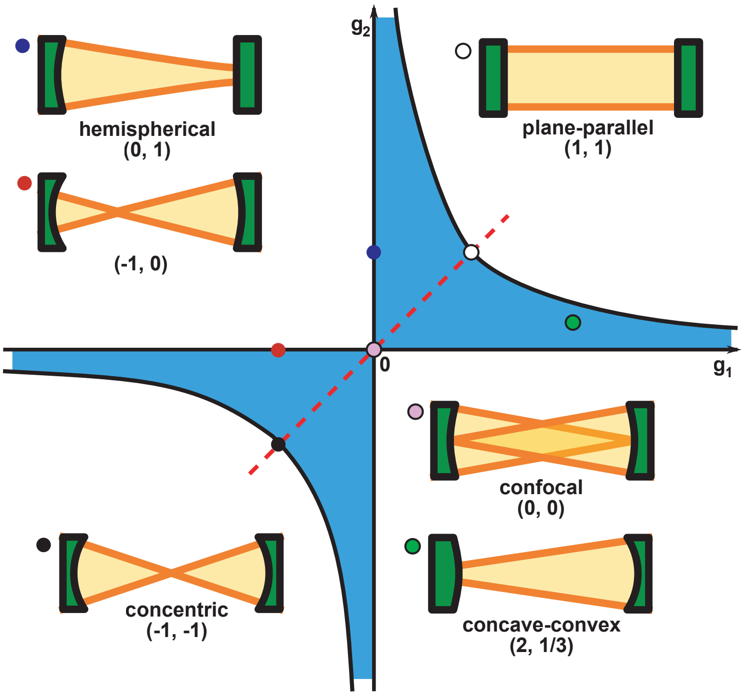

Figure 2.5: Stability criteria for canonical optical resonators. In this graph of g_1 vs. g_2, stable resonators are contained within the blue region. After an original by F. Dominec (Creative Commons license).

Non-Gaussian Resonators

Canonical optical resonators have transverse eigenmodes which are well represented by Hermite-Gaussian functions [13]. Can we then assume that we will be able to find functions to represent the eigenmodes of all resonators? The answer is no — some resonators are impossible to describe analytically. And how should we describe imperfect Gaussian resonators, with imperfectly spherical mirrors, and imperfect mirror alignment? We can approximate the solutions of certain resonators by finding a Gaussian resonator that resembles them. For instance, we can approach an analytical solution to the plane-plane cavity with finite mirrors by considering a symmetrical Gaussian resonator with mirror radius of curvature much greater than the mirror separation. But in order to solve more general resonators we must consider numerical methods using computers, and this is the subject of chapter 4.

CHAPTER 3

Photorealistic visualization of imaging in canonical optical resonators

This chapter looks again at optical resonators, but from the point of view of their imaging properties and is based upon work which was published as [14] and [15].

3.1 Introduction

Ray-tracing programs can create photo-realistic images of scenes described in the computer. One popular scene for such visualization comprises reflecting spheres; many examples can be found on the internet (e.g. [16]). Scientifically, such scenes contain many surprises; an arrangement of four mirrored spheres in a hexagonal close-packed structure, for example, exhibits chaotic scattering that can result in complex fractal images [17].

A canonical optical resonator (or cavity) [6] (see figure 3.1) consists of a pair of spherical mirrors facing each other. As spherical mirrors are simply segments of reflecting spheres, canonical optical resonators are closely related to the popular reflecting-spheres scene. To the best of our knowledge – and rather surprisingly – the inside of a resonator has never been visualized using ray tracing.

Here we use the freely available software POV-Ray [18] to investigate some of the ray-optical properties of resonators (see figure 3.1), specifically the different imaging characteristics of geometrically stable [7] and unstable [19] resonators, and the emergence of fractal eigenmodes in the latter [20]. Such investigations could form part of physics undergraduate computing labs with the beneficial side-effect of learning something about the imaging properties of resonators.

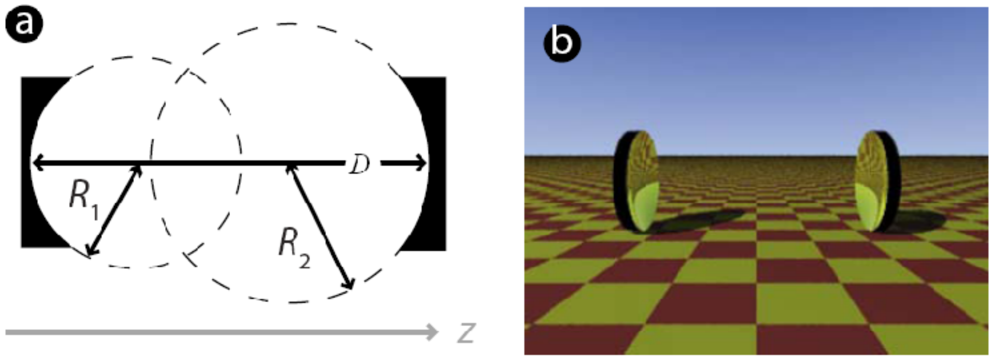

Figure 3.1: Example of a canonical optical resonator, in diagrammatic form (a) and rendered using the ray-tracing software POV-Ray [18] (b). Each mirror is a segment of a sphere; the respective radii of the spheres are R_1 and R_2, giving the mirrors focal lengths $f_1 = R_1/2$ and $f_2 = R_2/2$.

3.2 Peeking inside symmetric resonators

Resonators consist of segments of spheres of radii R_1 and R_2 (see figure 3.1). These radii of curvature can be positive or negative, respectively corresponding to concave / convex mirrors with the mirror surface being on the sphere's inside / outside. The two mirrors respectively have focal lengths $f_{1,2} = R_{1,2}/2$. For a mathematical analysis it is customary, even useful, to define the g parameters of a resonator as [6]

$$g_{1,2} = 1 - \frac{D}{R_{1,2}}, \tag{3.1}$$

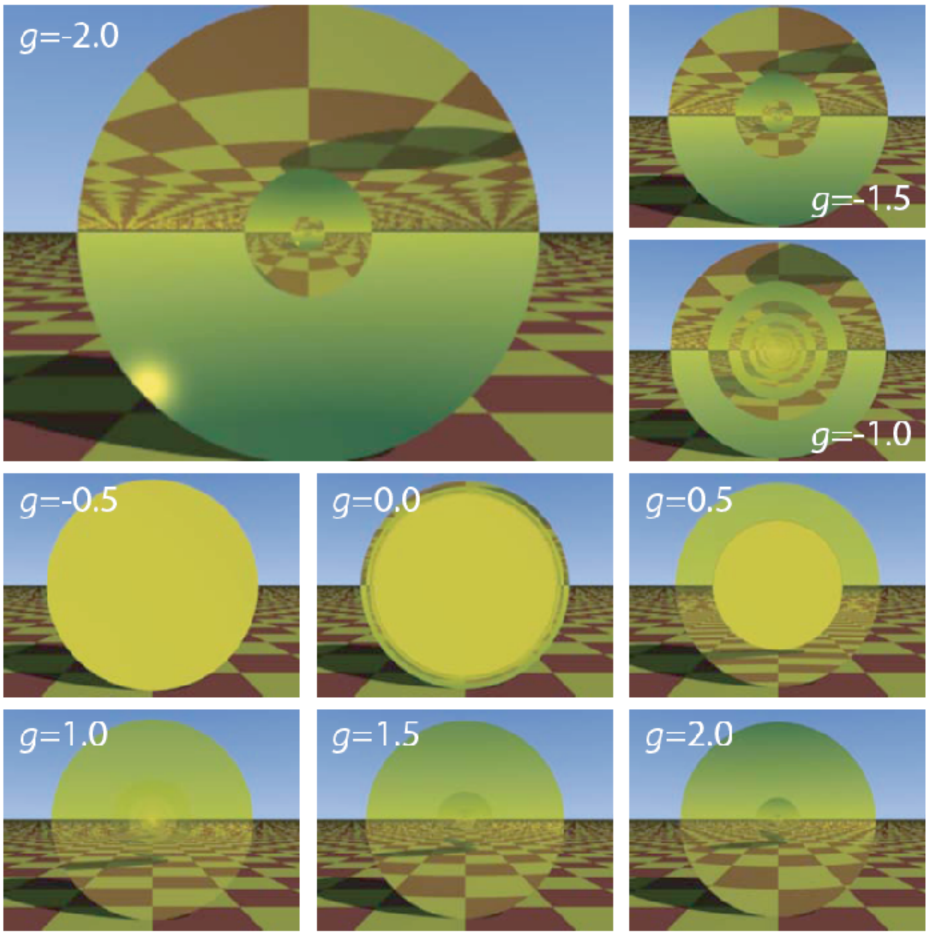

Figure 3.2: Views in symmetric resonators with g values between -2 and $+2$. This sequence of images was calculated for a fixed resonator length, $D = 5$, but different radii of curvature as calculated from equation (3.1). The mirror diameter is $d = 2$, as it is throughout the figures in this chapter. The camera was positioned on the resonator axis at $z = 1.8$. (As g is increased the mirror appears to get smaller. This is due to the edge of the mirror moving further away from the camera as its shape becomes more convex.)

where L is the separation between the mirror centres (see figure 3.1(a)). Throughout this and the following section, we restrict ourselves to symmetric resonators, that is resonators with $g_1 = g_2 = g$.

The top left frame of figure 3.2 shows the view from within a symmetric resonator with $g = -2$. It is essentially the simulation of a photograph taken with an invisible camera inside the resonator, pointing towards the first mirror. (Real cameras are, of course, not invisible, but in a real experiment a similar view could, in principle, be seen through the second mirror, provided it was partially transparent, in the limit of the transparency approaching zero, which would make the view very dim). In the centre, the first mirror can be seen, surrounded by the chequered plane it stands on and the sky. In the centre of this first mirror is a reflection of the second mirror, again surrounded by the chequered plane and the sky. The reflection is inverted, so the chequered plane appears to be above the second mirror and the sky. Each subsequent reflection inverts once more.

Subsequent frames of figure 3.2 show views in the same resonator, but with the radii of curvature changed such that g is increased. Between $g = -1$ and $g = +1$, a yellow spot can be seen in the centre. Throughout this paper we have chosen to give our mirrors a slightly golden surface (using POV-Ray's `pigment` and `finish` keywords), so that light coming from a specific direction that has experienced more mirror reflections has a more yellow tinge. The yellow spot corresponds to light rays that bounce endlessly between the

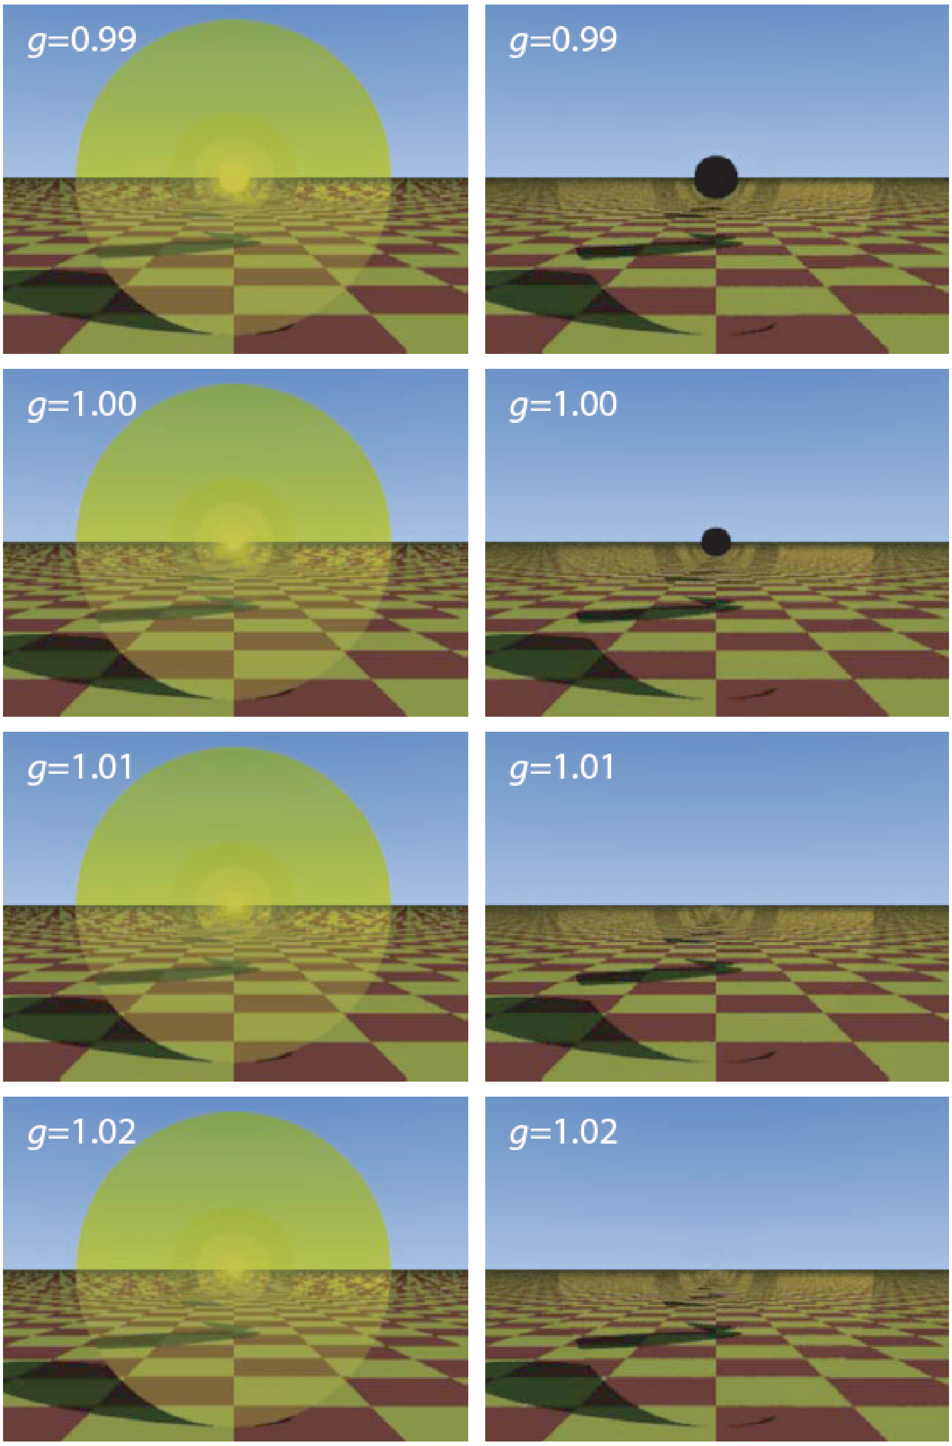

Figure 3.3: Details of the transition from instability to stability. In this sequence of images, g is increased from 0.99 to 1.02. Other parameters are the same as in Fig. 3.2. The left column shows the view inside a resonator with slightly golden mirrors, while the right column shows the view inside a resonator with perfectly reflecting, and therefore colourless, mirrors.

mirrors[1]. The existence of such “trapped” light rays is the definition of a resonator’s geometrical stability [22]. A mathematical analysis [23] reveals that resonators are geometrically stable provided that

$$0 \leq g_1 g_2 \leq 1. \tag{3.2}$$

In symmetric resonators this happens for $-1 \leq g \leq +1$, precisely the range of g parameters over which the yellow spot can be seen in our simulations. The emergence of the spot as g is increased from just below -1 to just above -1 is shown in figure 3.3. At $g = -1$, it begins to grow in size from zero.

Trapping of light rays does not happen in geometrically unstable resonators, that is resonators that do not satisfy the stability criterion, equation (3.2). In the family of symmetric resonators, this is the case for $|g| > 1$. In such resonators, every light ray will sooner or later miss a mirror and hit another object or the “sky sphere” [21] – as in some early models of the universe, POV-Ray represents the sky as the internal surface of a very large sphere. For our ray-tracing images this means that in the views calculated in unstable resonators in all directions we can see either the chequered floor or the sky.

In the following we use POV-ray to investigate a few striking imaging properties of geometrically stable and unstable resonators in more detail.

---

[1] To be more precise, the yellow spot corresponds to light rays that have bounced a finite, but large, number of times. So large, in fact, that the slight colour change upon each reflection (it gets slightly more golden) makes the light ray completely yellow. This uses POV-Ray’s Adaptive Depth Control and requires the POV-Ray parameter `max_trace_level`, which specifies the maximum number of intersections with surfaces after which the programme stops tracing a ray any further [21], to have values ≥ 50.

3.3 Imaging in geometrically stable resonators

Stable canonical resonators have very specific geometric imaging characteristics [7, 24]. We demonstrate a few of these characteristics here by visualizing the images of a semi-transparent object placed inside a resonator (see figure 3.4(a)).

One familiar example of resonator imaging is displayed by an optical resonator composed of two plane mirrors, shown in figure 3.4(b). This is the familiar“hall-of-mirrors” effect, which creates an infinite series of equidistant images of any object between the mirrors, here a semi-transparent letter “P”. Such a resonator is actually at the edge of stability ($g = 1$).

Figure 3.4(c) shows the more complex imaging that can occur in stable resonators. Images of the letter “P” now appear in various sizes, orientations, and, as can be seen from figure 3.5, z positions. The positions appear complex, but are well-understood mathematically [7]. The z position of an image after n round trips, z_n, can be uniquely mapped onto a corresponding angle, Ω_n, with the property that each additional round trip advances the angle through a constant step angle δ, so [7]

$$\Omega_{n+1} = \Omega_n + \delta. \tag{3.3}$$

Two cases can then be distinguished: periodic and aperiodic imaging. *Periodic* imaging occurs if δ is a rational multiple of 2π, that is if it can be expressed

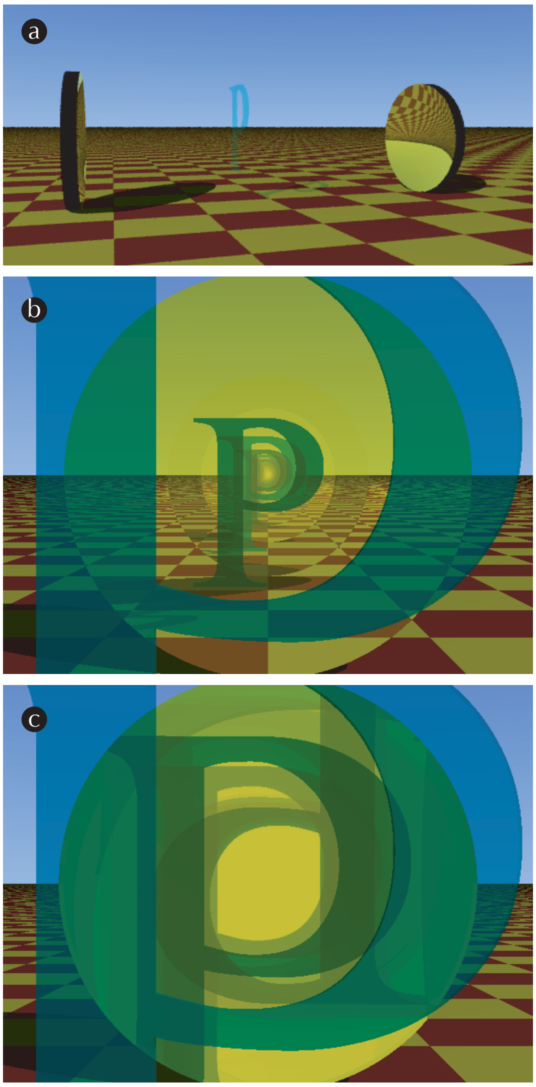

Figure 3.4: Imaging of an object placed inside geometrically stable resonators. Here the object is a semi-transparent letter “P” in the transverse plane in the centre of the resonator, shown in (a). A plane-plane resonator ($g = 1$, b) shows the familiar hall-of-mirrors effect, creating an infinite series of equidistant images. Other stable resonators, like the symmetric resonator with $g = 0.2$ shown in (c), create a more complex series of images. In both cases, the resonator length is $L = 4$, the semi-transparent letter “P” is positioned at $z = 1.5$, and the camera is positioned at $z = 0.1$.

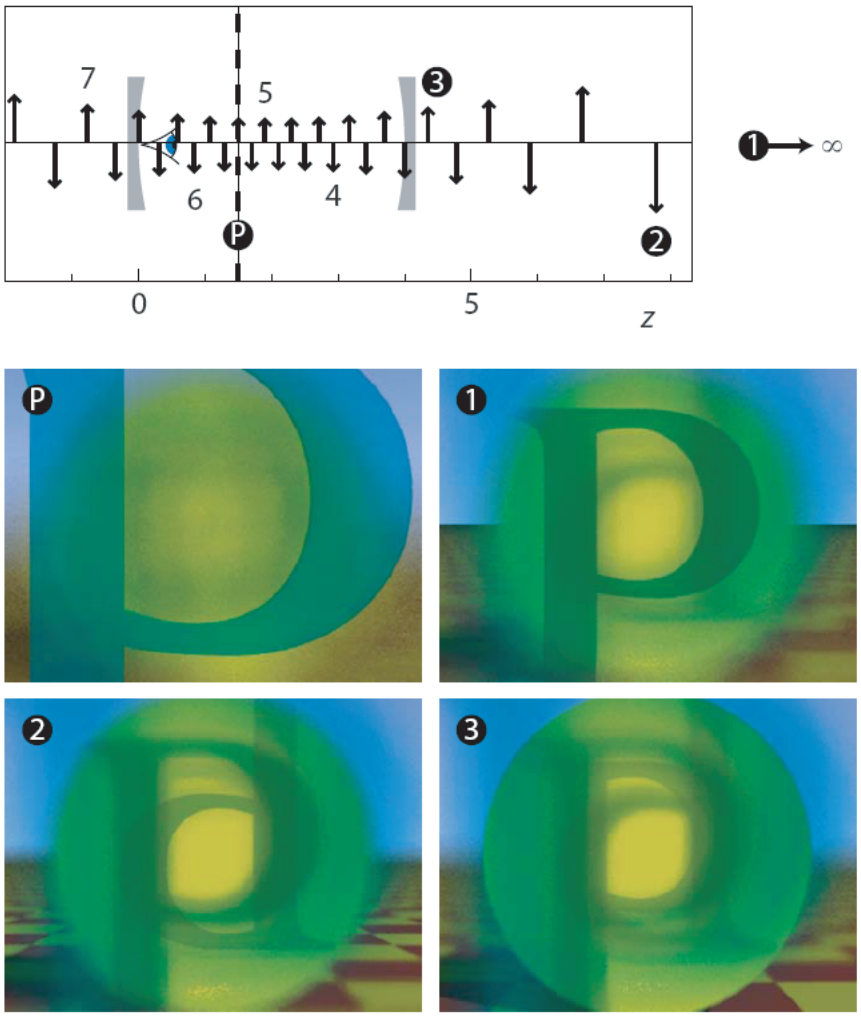

Figure 3.5: Location of the images of the letter “P” in Fig. 3.4(c). The diagram at the top indicates the z positions and relative magnification of the images due to the first 40 round trips in the resonator (some images are lying outside the shown z range and are not marked). Each image is marked by an arrow, whose horizontal position indicates the image’s z position and whose length indicates the image’s relative magnification. Numbers above or below the dot indicate the number of round trips after which the image is formed (always starting with light travelling to the right; image positions due to light travelling to the left are not shown, and no such images are visible anyway). The object is in the plane marked “P”. The frames below show the view from the camera position ($z = 0.1$, marked by an eye; the camera is looking towards the right mirror). Whereas in Fig. 3.4 the depth of field is infinite, it is reduced here so that individual planes come into focus. In POV-Ray, this is done by setting `aperture` to values above 0 (we use here 0.7). The different frames focus on the object plane ($z = 1.5$) and the planes that contain the first ($z = \infty$), second ($z = 7.7$), and third ($z = 4.3$) images. Note that the size at which the image is seen depends not only on its relative magnification, but also on its distance from the camera.

as

$$\delta = \frac{p}{q} 2\pi, \tag{3.4}$$

where p and q are integers. After q round trips, each plane is imaged back onto itself, via $(q-1)$ intermediate planes. The magnification is ±1, so any images of the same object that lie in the same plane are the same size (but they may be rotated through 180°, if the magnification is -1). *Aperiodic* imaging happens in all other cases, namely when δ is an irrational multiple of 2π. No plane is then ever imaged back onto itself; instead, it is imaged into *every other plane*. All intensity cross-sections through any eigenmode of such a resonator therefore have to be similar – in other words, the eigenmode is structurally stable.

Figures 3.4(c) and 3.5 are examples of aperiodic imaging in a geometrically stable resonator. They do not demonstrate all of the imaging properties described above, but restrict themselves to demonstrating different images in different positions. Directly visualizing some of the more intricate properties is considerably more challenging and computer-intensive, as large numbers of round trips through the resonator have to be visualized.

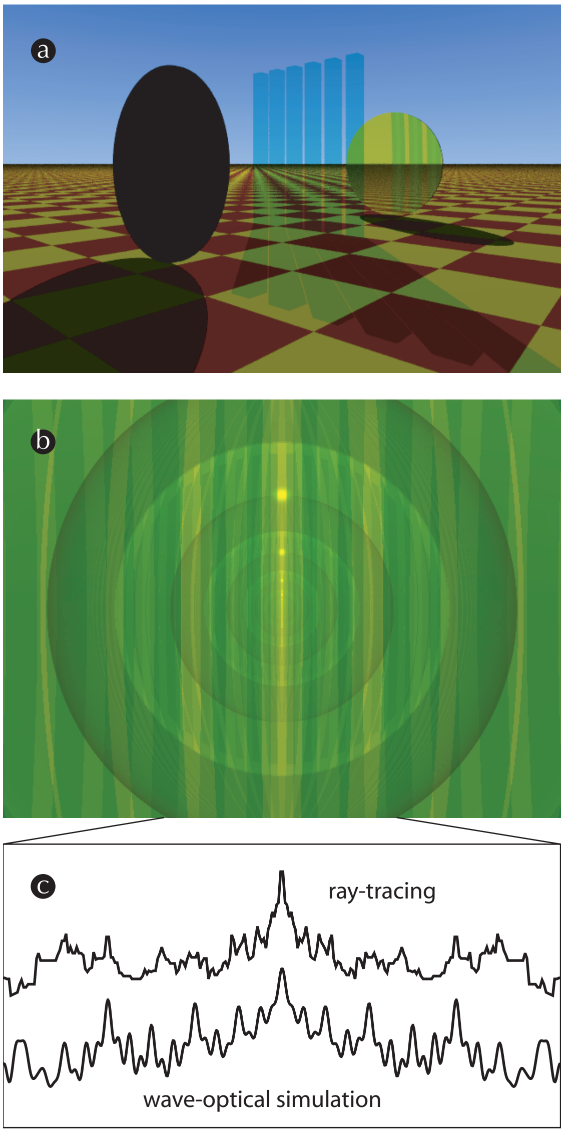

Figure 3.6: Creation of fractal eigenmodes inside a geometrically unstable resonator. (a) View of the confocal resonator with a semi-transparent grating in the confocal plane with magnification $M = -2$. The grating simulates the intensity pattern of the diffraction pattern from a slit aperture inside the resonator. For clarity, the number of grating lines is greatly decreased. (b) View from inside a confocal resonator similar to the one shown in (a), but with more (approximately 200) and thinner grating lines and – in the interest of symmetry - the light source moved to a position directly above the resonator axis and the floor removed. A brightness cross-section through the resulting pattern (c, top) shares the basic fractal character and many details with the intensity cross-section through the fractal eigenmode of a resonator with the same magnification and a slit aperture, calculated using a full wave-optical simulation (Fig. 7(b) in ref. [20]).

3.4 Imaging in geometrically unstable resonators

While geometrically stable resonators eventually image every plane into itself (or an arbitrarily close plane) with magnification $M = \pm 1$, geometrically unstable resonators possess special planes, called self-conjugate planes, that are imaged into themselves after each round trip, with magnification $|M| \neq 1$ [20]. In confocal resonators, i.e. those resonators where $D = f_1 + f_2$, there are two self-conjugate planes with round-trip magnifications $M = -f_1/f_2$ and $M = -f_2/f_1$, respectively [25]. (Confocal resonators with $f_1 = f_2$ have round-trip magnifications $M = -1$ and are on the edge of stability; all other confocal resonators are unstable.) We restrict ourselves here to confocal resonators with $f_1 > f_2$, and we study the self-conjugate plane with magnification $M = -f_1/f_2$. As this magnification is negative, the images are rotated through 180° with respect to the original; as $|M| > 1$, they are bigger than the original.

If a wave is circulating in an unstable resonator, the diffraction pattern due to any apertures in the resonator provides an intensity pattern that is magnified during subsequent round trips. At the same time – because the beam passes through the aperture again during each round trip – the aperture-diffraction pattern is re-created. This process leads to an eigenmode intensity cross-section in the transverse plane that is fractal [20, 25] (the eigenmode's intensity in the longitudinal direction can also be fractal – see chapter 4).

As ray-tracing does not include diffraction effects, this cannot be directly demonstrated here. However, the important functions of the diffraction pattern of the aperture (we use here a slit aperture) can be crudely counterfeited by placing a semi-transparent grating into the self-conjugate plane (see figure 3.6(a)). Like an aperture diffraction pattern, the grating provides a pattern that is magnified and re-created during subsequent round trips.

Figure 3.6(b) shows the view inside such a resonator. The pattern clearly has the expected fractal characteristics. In fact, a horizontal brightness cross-section through the centre (figure 3.6(c)) reveals close similarity with the results of a full wave-optical simulation of the intensity cross-section in the self-conjugate plane through the eigenmode of a confocal resonator with a slit aperture [20].

3.5 Other interesting cases: optical billiard and ray-rotation sheets

In this section, we describe briefly some other optical systems which we examined through ray-tracing. These included work on optical billiards with chaotic orbits and ray-optical metamaterial analogues; this later work led to the publication of [15], and was continued by other members of the optics group (see, for example, [26]).

Our work on optical billiards came from the idea of using ray-tracing to render a two-dimensional billiard as a three-dimensional mirror cabinet, with a finite height. A billiard is a dynamical system where a point particle is contained within a boundary (the boundary being piecewise smooth [27]); the particle moves at constant velocity except where it encounters the boundary, when it undergoes specular reflection. This makes them ideal systems to investigate using ray-tracing.

Billiards can be divided into two categories according to properties that are defined by the shape of the boundary. Those where the particle propagates along periodic or quasiperiodic orbits (such as square or equilateral triangle billiards) are called *regular* billiards. Those where the particle can follow either a periodic or a chaotic orbit (well known examples being the stadium or sinai billiard) are called *chaotic* billiards. We decided to investigate chaotic billiards, since stable resonators can, in some ways, be viewed as a kind of regular billiard, albeit only for particles on paths almost parallel to the optical axis.

Figure 3.7(a) shows the geometry of a chaotic billiard, known as the Bunimovitch stadium [29] after it's discoverer. It can be considered to be an unstable symmetrical strip-resonator (with $g = -3$) but with closed, reflective sidewalls. Figure 3.7(b) is from [28], created using the principle of a mirror cabinet map, outlined in that paper. One can think of it as being the virtual images of the cabinet walls as seen by an observer situated at the centre of the stadium (with the circles being virtual images of the observer – apart from the central one); the first 7 virtual walls are shown.

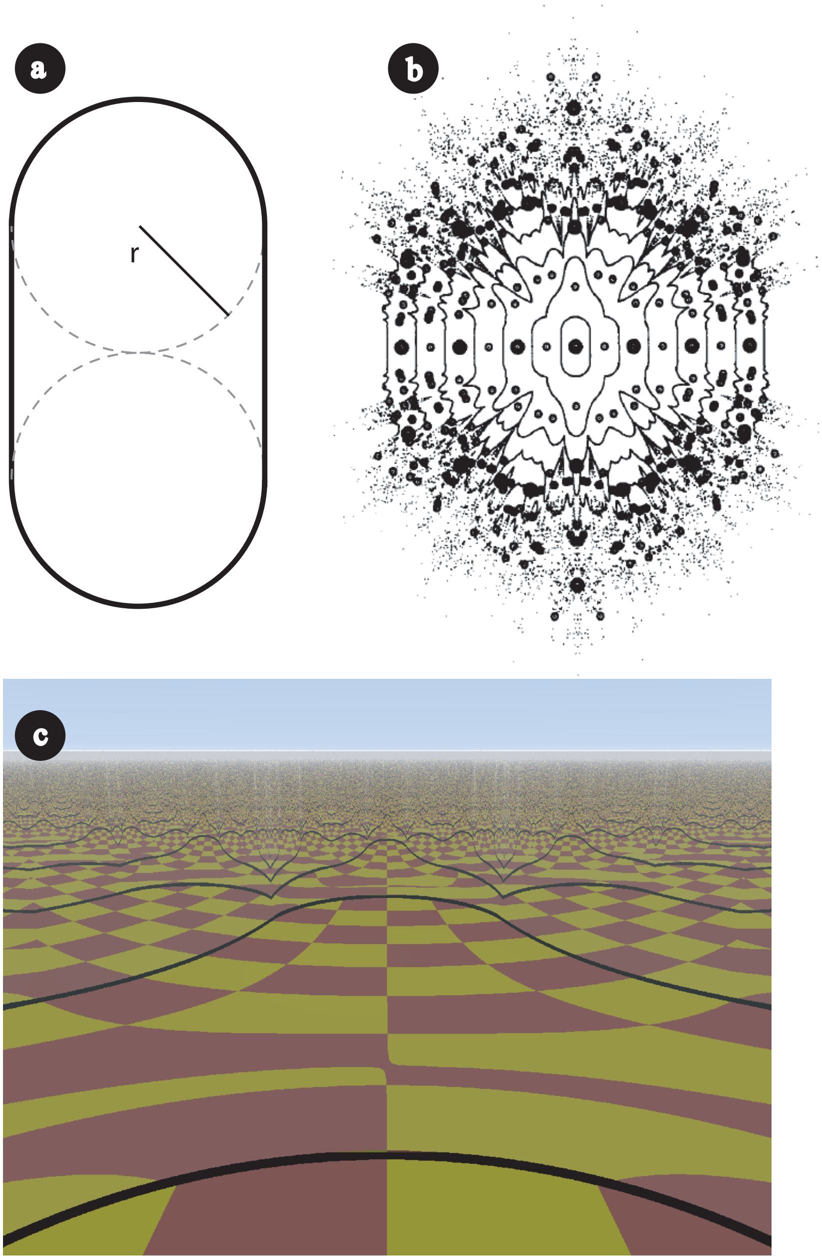

Figure 3.7: (a) shows the geometry of the Bunimovitch stadium optical billiard. (b) and (c) show to equivalent ways of looking at the virtual images of the walls of a mirror cabinet shaped like a Bunimovitch stadium. In both, the observer is at the centre of the stadium. (b) is a mirror cabinet map from [28] while (c) is a ray-traced rendering of the view inside a 3D mirror cabinet.

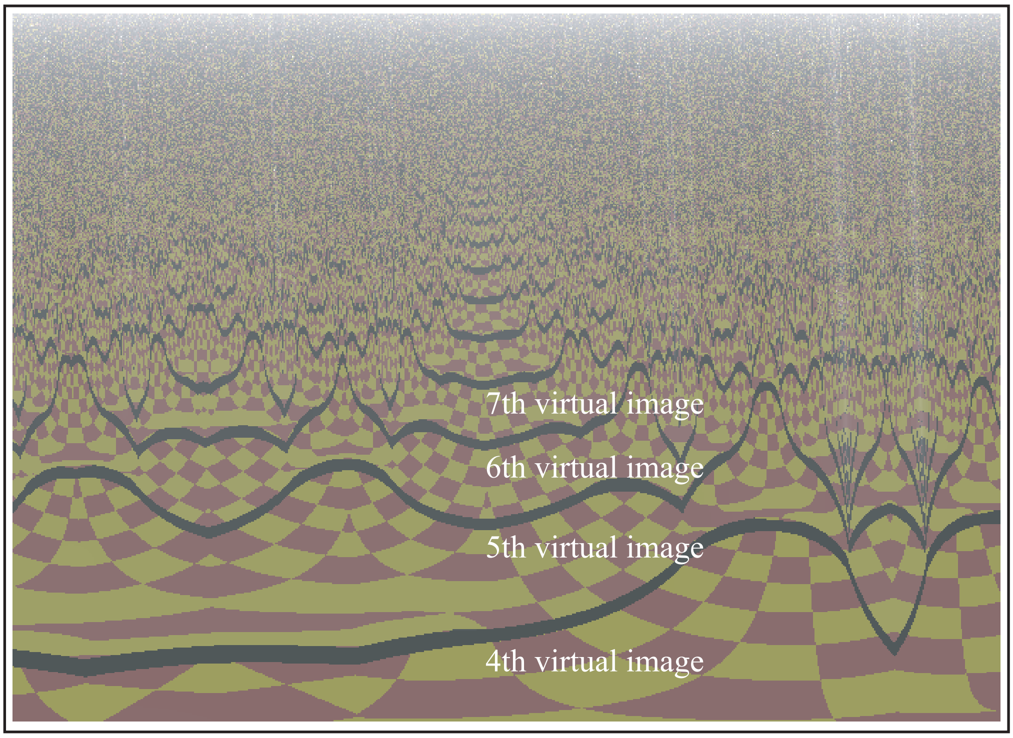

Figure 3.8: A detailed view of the inside of a Bunimovitch stadium optical billiard. The limit to the fractality of the boundary images is the resolution of the rendered image.

By creating a virtual scene of a closed strip-resonator (with a finite height), we can see exactly what the observer inside such a mirror cabinet would see, if only he were invisible. The virtual mirror cabinet was created with a narrow, black band at the bottom of the mirror. This makes it easier to see where each successive virtual image of the cabinet's walls appears (see figure 3.7(c)). What can also be seen (and seen more clearly in figure 3.8) is that each successive virtual image of the cabinet boundary becomes increasingly fractal. We can also see that the view in some directions is more fractal than in others. For instance, if the observer was oriented at 90° to that shown in figure, the image would be of the interior of a plane-plane resonator.

Our other use of ray-tracing visualization was to view scenes containing ray-optical metamaterial analogues. A sheet consisting of an array of small, aligned Dove prisms can locally (on the scale of the width of the prisms) invert one component of the ray direction. A sandwich of two such Dove-prism sheets that inverts both transverse components of the ray direction is a ray-optical approximation to the interface between two media with refractive indices $+n$ and $-n$.

Negative refraction has already been realized ray-optically in the form of lenslet arrays: pairs of lenslet arrays with a common focal plane bend light rays like the interface between optical materials with refractive indices $+n$ and $-n$. These have been realized in the form of standard and GRIN lenslet arrays , and their three-dimensional imaging properties, including pseudoscopic imaging, have been examined.

Another way of achieving ray-optical negative refraction uses combinations of miniaturized Dove prisms. Our combinations of Dove prisms consist of two periodic Dove-prism arrays we call Dove-prism sheets, where one sheet is rotated with respect to the other by 90°. These Dove-prism-sheet sandwiches work differently from the lenslet arrays described above: the lenslet arrays work by forming an intermediate image; the Dove-prism-sheet sandwiches work by successively inverting the ray vector's x and y components.

The basic building block of a Dove-prism sheet is a Dove prism. With the coordinate system chosen as in figure 3.9, a Dove prism inverts the y direction of any transmitted light ray. It also offsets the rays, whereby the offset is on the

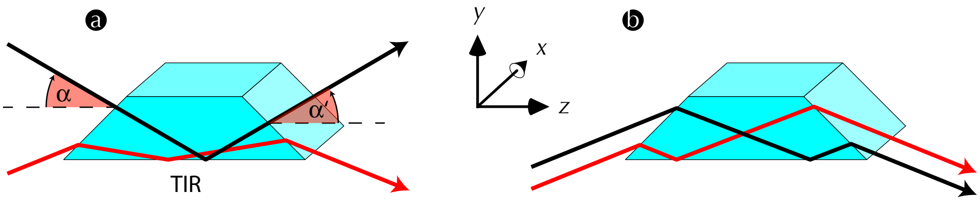

Figure 3.9: Ray optics of Dove prisms. A Dove prism orientated as shown in (a) inverts the y component of the direction of individual light rays in a light beam passing through it. This happens during total internal reflection (TIR) from the prism's bottom surface. This corresponds to an inversion of the angle the light ray has with respect to the prism's bottom surface: $\alpha = -\alpha'$, where α and α' respectively are the angles before and after passage through the prism, as shown in (a). The prism also offsets the y position of each ray; in (b), the prism swaps the positions of the red and black rays.

scale of the prism diameter. We are considering here the limit of small Dove prisms, so small in fact that we can ignore this offset. Clearly, wave-optically this limit breaks down as the prism diameter reaches the wavelength of the light. Acceptable compromises for visual purposes could be prism diameters of between 10μm and 1mm.

Dove prisms that are stretched in the x direction (again with the choice of coordinate system shown in figure 3.9) and stacked on top of each other form a Dove-prism sheet (figure 3.10). Note that the prisms need to be separated by a few wavelengths to ensure that total internal reflection at the long side (see figure 3.9(a)) is not frustrated.

The ray optics of such a sheet are simple: in the limit of small Dove prisms the sheet flips the y direction of individual light rays in a beam passing through it. This implies that for light rays incident in a plane parallel to the (y, z)

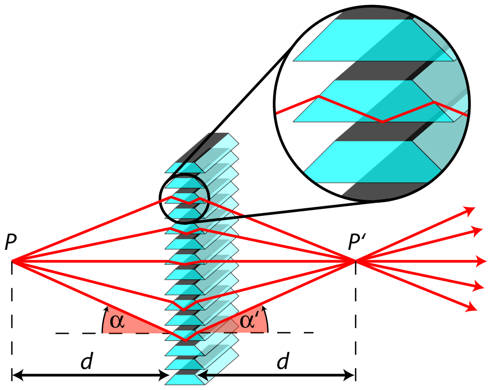

Figure 3.10: Structure and ray-optics of a Dove-prism sheet. The inversion of the angle with respect to the prism's bottom surface (Fig. 3.9) now becomes an inversion of the angle with respect to the Dove-prism sheet normal. In the limit of small prism apertures, the sheet does not significantly offset light rays (red). A point light source, P, placed a distance d in front of the sheet, creates light rays that intersect again in a point P' the same distance d behind the sheet.

plane, the angles of incidence, α_1, and refraction, α_2, are related through the equation

$$\alpha_1 = -\alpha_2. \tag{3.5}$$

It is particularly interesting to combine a Dove-prism sheet with another, parallel, Dove-prism sheet that is rotated around the z direction through 90°, and which therefore flips the x direction of light rays passing through it. Such Dove-prism-sheet sandwiches then flip both transverse ray directions (x and y), and invert the angle of incidence for any plane of incidence. When the two crossed Dove-prism sheets are close together, they lead to no additional ray offset. They therefore act like the interface between two optical media with equal and opposite refractive indices, $+n$ and $-n$: Snell's law, written for this situation, states that

$$n \sin(\alpha_1) = -n \sin(\alpha_2), \tag{3.6}$$

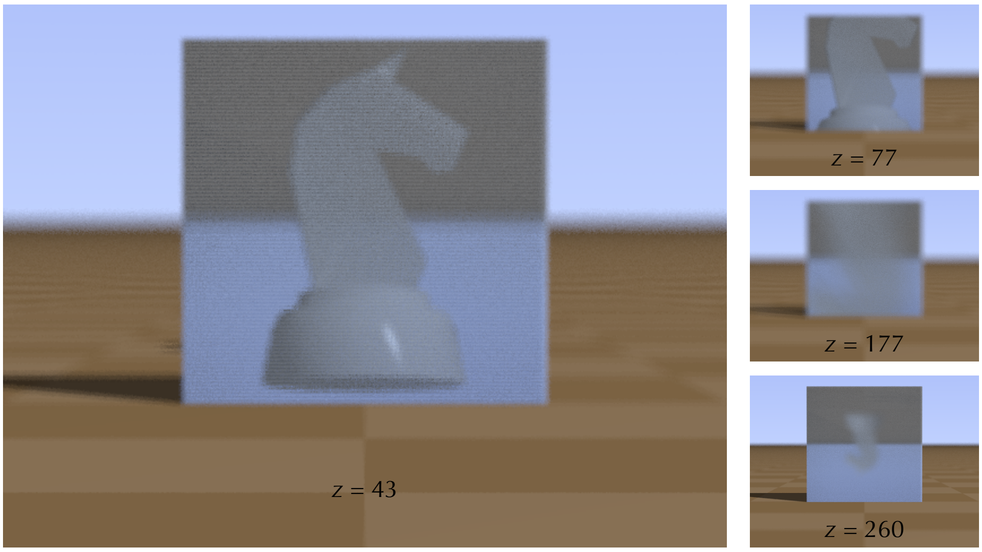

Figure 3.11: Relationship between object and image distance for crossed Dove-prism sheets. A chess piece – the object – is positioned a distance z behind the sheets; the crossed Dove-prism sheets image it to a position a distance z *in front of* the sheets. The different frames show the image of the chess piece for various object distances; the sheets and the camera are stationary. In the first ($z = 43$) and second ($z = 77$) frames the image becomes larger and larger as it moves towards the camera, positioned a distance 120 units in front of the sheets. The image then moves through the camera plane and behind it, where it re-appears upside-down and getting smaller. In the first two frames, $z = 43$ and $z = 77$, the camera is focussed onto the image of the chess piece; its image can be gleaned by inspection of the position of the focus on the chequered floor, which has a square length of 20 units. In the second two frames simple focussing is not possible as the chess piece is behind the camera, which is roughly focussed onto the sheets. The frames are from a movie (MPEG-4, 256 KB, available in the supporting online material) calculated by performing ray tracing through the detailed prism-sheet structure, using the freely-available software POV-Ray [18].

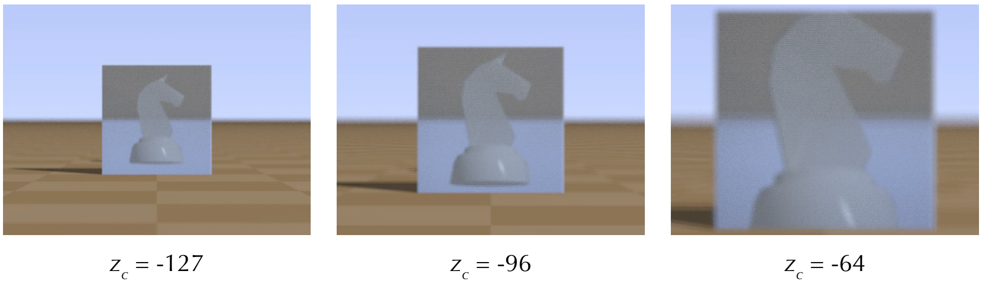

Figure 3.12: Image of a chess piece created by crossed Dove-prism sheets as seen from different distances. The sheets are in the $z = 0$ plane; the image is at $z = -40$ (that is 40 units in front of the sheet); z_c is the position of the camera. From left to right, the frames show the simulated view as seen with a camera moving closer to the Dove-prism sheets; both the sheets and the chess piece are stationary. Because the distance between camera and image is less than that between camera and sheets, a decrease in both distances by the same absolute amount, that is moving the camera in the direction of image and sheets, decreases the distance to the image by a larger factor than that to the sheets. This means that the angle under which the image of the chess piece is seen grows more than the angle under which the sheets are seen. The frames from a POV-Ray [18] movie (MPEG-4, 204 KB, available in the supporting online material).

which (provided that $-90° \leq \alpha_1, \alpha_2 \leq +90°$) is equivalent to equation (3.5).

Pseudoscopic imaging

Images produced by single lenses are orthoscopic: if two objects at longitudinal positions z_1 and z_2 are imaged into positions z_1' and z_2', and if the first object is in front of the second, i.e. if $z_1 < z_2$, then the image of the first object will be in front of the image of the second, so $z_1' < z_2'$. The opposite is true in pseudoscopic imaging, where the image of the second object is in front of that of the first, so $z_1' > z_2'$.

The effect of the inversion of the angle of incidence by crossed Dove-prism sheets is to image any object a distance d behind the sheets to the same distance in front of the sheets (figure 3.10). In other words, if the longitudinal coordinate z is chosen such that the sheets are at $z = 0$, then an object distance z corresponds to an image distance $-z$. For the two longitudinal object positions with $z_1 < z_2$ discussed above this results in image positions $z'_{1,2} = -z_{1,2}$, and therefore the inverted relationship between the longitudinal image positions $z'_1 > z'_2$. Crossed Dove-prism sheets therefore produce pseudoscopic images.

Figures 3.11 and 3.12 demonstrate this pseudoscopic imaging with ray-tracing simulations performed using the software POV-Ray [18]. Both figures visualize imaging of a chess piece through crossed Dove-prism sheets, each comprising 200 Dove prisms. In figure 3.11 the distance of the chess piece behind this Dove-prism-sheet sandwich is varied; in figure 3.12 the distance of the (simulated) camera from the sheet sandwich is varied.

The inversion of the z coordinate during imaging implies that crossed Dove-prism sheets produce pseudoscopic images. Figure 3.13 demonstrates various properties of these pseudoscopic images. Specifically, it shows that pseudoscopic images appear to be “inside out”; the pseudoscopic image of a convex chess piece, for example, is concave. When looking at this image from different directions, the image appears to have rotated, just like the hollow face mask in the famous hollow-face (or “Bust of the Tyrant”) illusion [30]. In the case of the chess piece shown in figure 3.13, looking at the pseudoscopic image of one of the chess pieces placed behind the Dove-prism sheets from the left lets

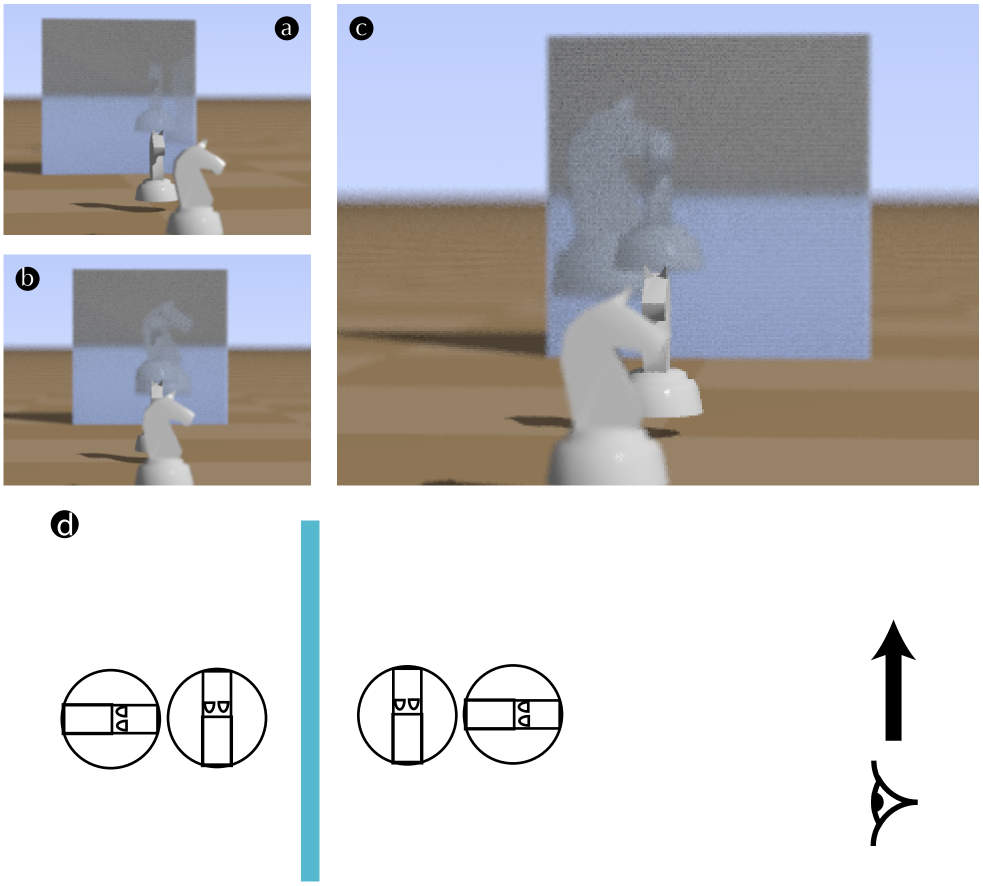

Figure 3.13: Orthoscopic objects and pseudoscopic images. The frames show two chess pieces in front of crossed Dove-prism sheets, and the images of two chess pieces behind the sheets, as seen from three different viewing positions. The pieces are arranged such that one image is at the same distance as one of the chess pieces in front of the sheet, the other image is at the same distance as the other piece in front of the sheet. This can be seen by one chess piece always being below one image, independent of viewing angle, which means they are always undergoing the same parallax, which in turn implies that they are the same distance from the camera. However, while the left side of the front piece is visible from the left-most viewing point (a) and the right side from the right-most viewing point (c), the opposite is true for the pseudoscopic images. Also, while the piece in front (which, of course, appears bigger) obscures the piece behind it, the image in front (again the bigger image) *is obscured* by the image behind it. (d) shows a schematic overhead view of the scene. Frames (a)-(c) are from a movie (MPEG-4, 848 KB, available in the supporting online material of reference [15]) calculated by performing ray tracing through the detailed prism-sheet structure, using the freely-available software POV-Ray [18].

us see the right side of the chess piece, not the left side, as is the case with the chess piece placed in the same longitudinal position for comparison. Figure 4 in [30] shows the same effects in the hollow-face illusion.

Figure 3.13 also demonstrates another striking property of pseudoscopic images. If two objects are placed behind one another, the object in front obscures the object behind. In the pseudoscopic images of two objects placed behind one another, the image *behind* obscures the image in front.

3.6 Conclusions

Ray-tracing can directly visualize important properties of optical resonators and other interesting ray optical systems. Running such simulations requires only a computer with modest power and ray-tracing software, which can be obtained for free.

We have demonstrated here only very few imaging properties of optical resonators. We hope that many more, like for example the intricate imaging properties in stable resonators mentioned above, will be directly visualized. Hopefully that this will help students to understand optical resonators better, and perhaps even researchers to gain completely new understanding.

We have also investigated the effect of miniaturizing and repeating an optical component with interesting ray-optical properties, the Dove prism. Using ray-tracing simulations, we have demonstrated that the resulting Dove-prism sheets can ray-optically act like the interface between optical media with refractive indices of the same magnitude but opposite sign. We have also demonstrated some of the unusual properties of their pseudoscopic imaging.

CHAPTER 4

Resonator eigenmodes: the Fox-Li method

In this chapter we deal with the problem of finding the eigenmode of an arbitrary resonator using computer simulations. Fortunately there is a well-known and widely used method for solving this problem: the method developed by Gardner Fox and Tingye Li and first reported in their paper of 1961[31].

Using this method (often referred to as the Fox-Li method), we examine the mode properties of spherically aberrated cavities, with a view to their use as an alternative to Mexican Hat mirrors [32] within interferometric gravitational wave detectors. We also use the Fox-Li method to investigate the fractal geometry of the eigenmodes of unstable confocal resonators.

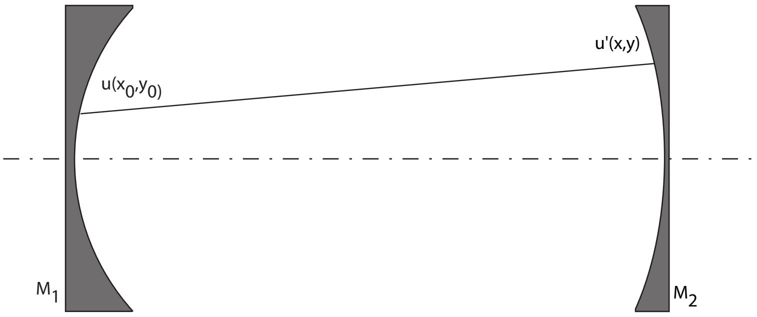

Figure 4.1: Half of a round trip in an optical resonator.

4.1 The Fox-Li method

In the early 1960s A. G. Fox and Tingye Li developed a numerical method for finding the lowest loss eigenmodes of optical resonators. Their initial paper [31] concerned itself with plane-plane Fabry-Perot cavities as this was a special case of the canonical resonators, the solution of which could only be approximated by Gaussian analysis [23].

Figure 4.1 shows half a round-trip of an optical resonator. With an initial wave amplitude of $u(x_0, y_0)$ at the first mirror, M_1, we have

$$u''(x,y) = \iint \widetilde{L}(x, y : x_0, y_0) u(x_0, y_0) dx_0 dy_0 \tag{4.1}$$

where $\widetilde{L}(x, y : x_0, y_0)$ is the round-trip propagator and $u''(x,y)$ is the wave amplitude after one round trip (and $u'(x,y)$ would be the wave amplitude after half of a round-trip, i.e. at the other mirror, M_2).

As we know, an eigenmode of the resonator should have the property that this wave amplitude should not change shape, although the absolute amplitude can – and, in an empty resonator, will – decrease.

These integral equations are, save in a few special cases, not amenable to analytical solution. However, beginning in 1961, A. G. Fox and T. H. Lee [31] published an influential series of papers in which they discussed an iterative computer solution to equations of this nature. The procedure is simple: an initial wave amplitude $u^0(x, y)$ is chosen. The wave amplitude is sent for one round trip around the simulated resonator. This procedure is repeated until eventually we have a wave amplitude where $u^{(n+1)}(x, y) = \gamma u^n(x, y)$, i.e. the system has reached an eigenmode. The success of the process does not depend on what form this initial amplitude takes, with some caveats: for instance, using a uniform plane wave as the initial amplitude can never give rise to an eigenmode with an odd symmetry [33]. Our standard initial amplitude was a uniform intensity with a pseudorandom phase.

The Fox-Li method filters out higher loss modes and converges to the lowest loss mode. The rate of the convergence depends on the resonator's loss per round trip and on its *mode discrimination*. If one or more modes have loss almost as low as the lowest loss mode it can take many thousands of iterations for them to die away.

At its most basic level, the Fox-Li method [31, 4] can be understood as a method for finding the eigenmode of an optical resonator by simulating the effects of sending some arbitrary initial beam through a finite number of periods of an infinite periodic lens waveguide (see section 2.2) which is equivalent to the resonator. After a certain number of periods (or the same number of round-trips of the equivalent resonator) the shape of the beam (in amplitude) ceases to change: this is the fundamental mode of the resonator.

How does this differ from what happens in a real resonator? The most obvious difference is the insensitivity of the Fox-Li method to length tuning. For an optical resonator to be on-resonance we have the requirement that $D = m\lambda/2$; looking at the output transverse mode of the resonator while the mirror separation is slowly changed by $\pm\lambda/2$ reveals this. However a Fox-Li simulation of a resonator where $D = D_0 = m\lambda/2$ will converge to (almost) the same mode as a simulation where $D = D_0 \pm x$ (and $x \in (0, \lambda/2)$). The sneaky ‘almost’ in the previous sentence covers the fact that the tiny length change x will have an effect on the g-parameters of the resonator. The reason for this is obvious, when we consider what the basic Fox-Li method actually simulates: it really *is* a simulation of a lens waveguide, not an optical resonator (though by summing the field at the mirrors over successive round trips the output field of a resonator from a given input field can be obtained). D enters the simulation only as the propagation distance between the lens l_1 and l_2. The phase difference between the propagating wavefront at, say, l_1 during the n^{th} round trip and the $n + 1^{th}$ round trip has no cumulative effect upon the mode – indeed it is completely irrelevant. In an optical resonator on-resonance, the phase difference would be 0, causing the output mode to build up.

The reason for this becomes apparent when we consider that, at no point is the electromagnetic field inside the resonator as a whole simulated. Instead we simulate what Siegman calls ‘a recirculating pancake of light’ [34] – that is, we consider the electromagnetic field only in a single plane and then propagate it by means of plane wave decomposition and FFT propagation (see Appendix A).

As we know, an eigenmode of the resonator should have the property that this wave amplitude should not change shape (though the absolute amplitude can — and, in the absence of a gain medium, will — decrease).

One difference between our simulations and the original simulations of Fox and Li is that ours are 3D, where theirs were 2D; that is, in their simulation the electric field was represented as a 1D array, while in ours it is represented as a 2D array.

Some practical considerations

The best way to describe the Fox-Li method is to relate it to the reality of the transient response of an optical resonator. If we consider a wave pulse at one end of the resonator (we need not worry about the problem of injection). The form of the electromagnetic field in this pulse is unimportant. As this field is (numerically) propagated around the resonator, or through the equivalent lens waveguide, some of the light will be guided by the periodic influence — be it mirror or lens — to remain close to the axis, and some will fall outside the aperture and be diffractively lost. Under this influence the beam will naturally shape itself to have the lowest diffractive loss possible. As we have seen, these transverse mode shapes in canonical resonators tend to be Gaussian.

One way of understanding what is happening as the Fox-Li method is carried out is to look at the loss per round trip. When an eigenmode of a resonator is propagated around that cavity the loss should be the same on each round trip and equal to γ as calculated above. This, in fact, gives us a way of deciding when our Fox-Li method has reached an eigenmode: if the loss per round trip becomes constant, the system is in an eigenmode. Figure 4.2 shows the behaviour of loss per round trip for two resonators. The simple behaviour shown in figure 4.2 (upper graph), which comes from a stable symmetric Gaussian resonator, can be explained as follows: what we are seeing is the effect of the slow dying away of the next-lowest-loss mode. As the lowest-loss mode and the next-lowest loss mode are propagated around the resonator they acquire a phase shift relative to each other — this is the origin of the beating behaviour. As parts of the field which are not in either of these modes are lost the overall loss per round trip decreases leaving only the beating between these two modes which eventually dies away as the next-lowest loss mode eventually fades. Figure 4.3 shows the time-evolution of the intensity at one of the mirror surfaces typical of such a graph. The much more complex trace shown in the lower graph is the result of a non-canonical resonator with very poor mode discrimination. Many modes with almost equal loss are being propagated around the resonator.

The gaussian beam analysis of the previous section allowed us to gain an insight into the properties of optical resonators with spherical mirrors, as long

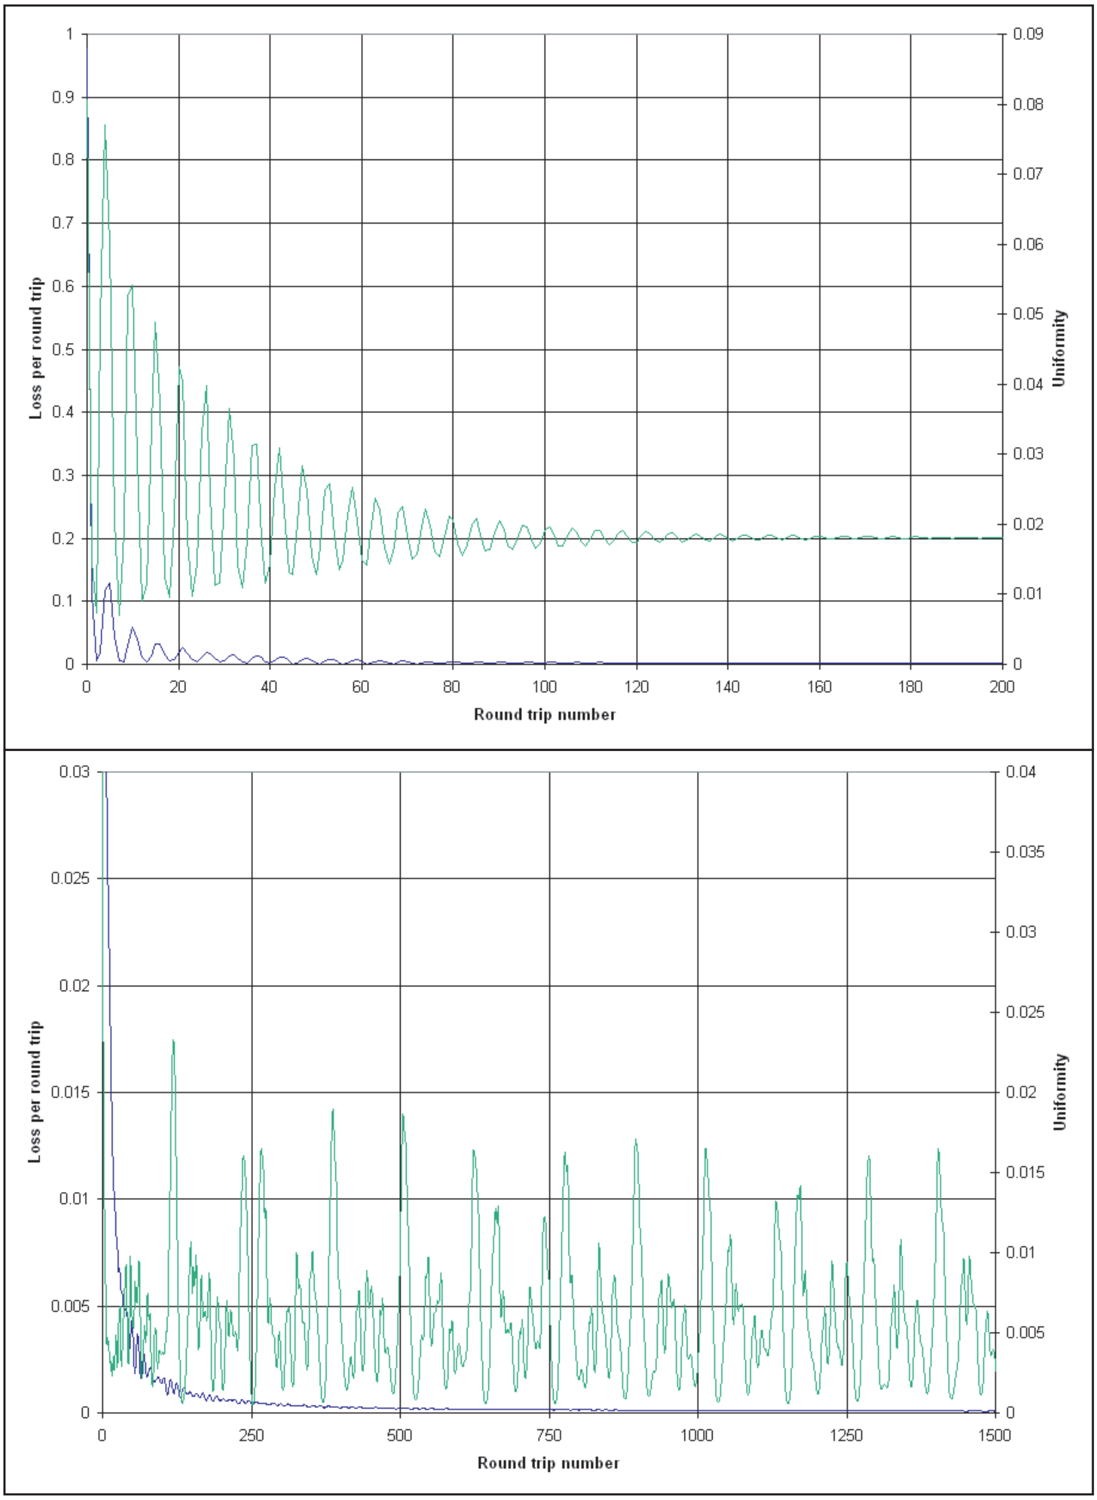

Figure 4.2: The upper graph shows the typical behaviour of a Fox Li simulation of a spherically aberrated cavity. Several modes exist in the cavity (shown by the oscillation of the uniformity (green trace) and loss (blue trace). The higher loss modes decay quickly, leaving only the fundamental mode of the cavity. The lower graph shows a similar simulation, but of a cavity with very poor mode discrimination. Several modes with very similar losses exist in the cavity and the higher order modes can take many thousands of iterations to decay.

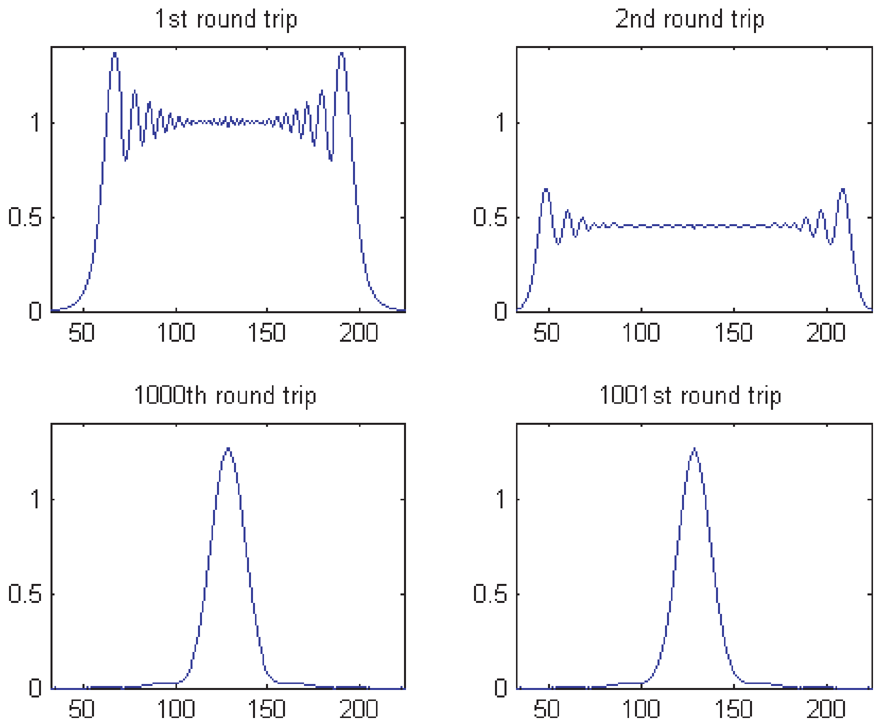

Figure 4.3: Intensity cross-sections after 1,2,1000 and 1001 round trips. Note the clear diffraction ripples in the mode shape of the 1^{st} round trip and that by the 1000^{th} round trip the resonator has settled into it's eigenmode, shown by the mode shape of the next round trip being identical.

as $0 \leq g_1 g_2 \leq 1$. Now we ask how we can find the eigenmode of an optical resonator with mirrors of any shape, and of finite extent[35].

Indeed the Fox-Li method has found uses far beyond that originally envisaged for it. Most use of the Fox-Li method has been concerned with the final solution at which the simulation arrives, but some researchers have also been interested in the transient response of the system. An instance of this is given in reference [36] where the Fox-Li method is used to examine the mode evolution in a spherical microdroplet of liquid when illuminated by laser light, and to discover how long it took to reach a steady state.

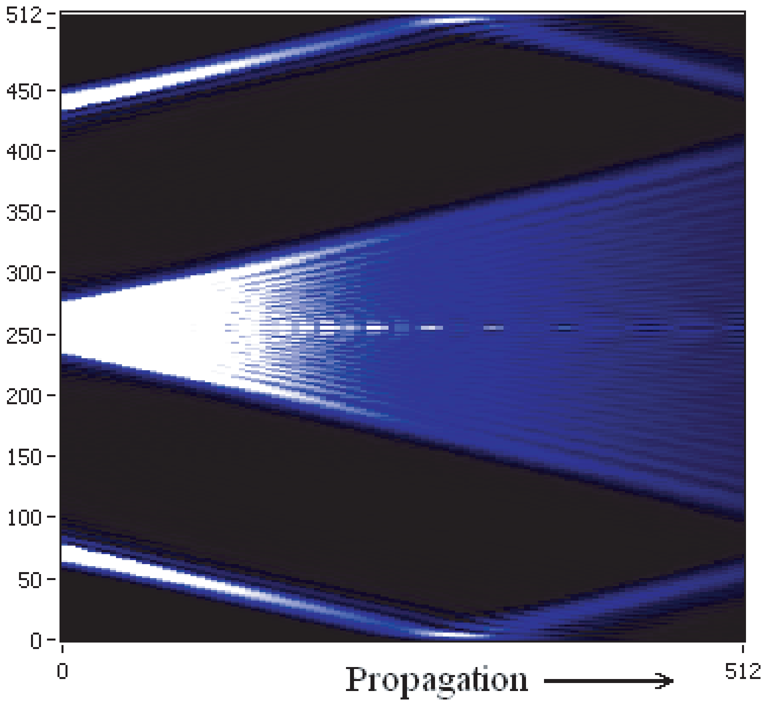

Figure 4.4: Cross section through a beam showing apparent reflection from the edge of the simulation area as the field is propagated in the $+z$ direction in 512 discrete FFT steps, which is an artefact caused by the FFT propagation method. Simulation is on a 512×512 grid.

Our Fox-Li simulations were carried out using *WaveTrace*, a software suite developed by the Optics Group. For a more in depth description of this software refer to appendix A.

The simulation takes place on a discrete grid of points representing some 2D height and width. We can propagate (non-paraxially) some initial field, $u^0(x, y)$, a distance, d, using fast Fourier transforms (FFTs), and find the resulting field $u(x, y)$. We work with the resonator unfolded into a (potentially infinite) periodic lens waveguide, where the lenses are represented as phase holograms [37].

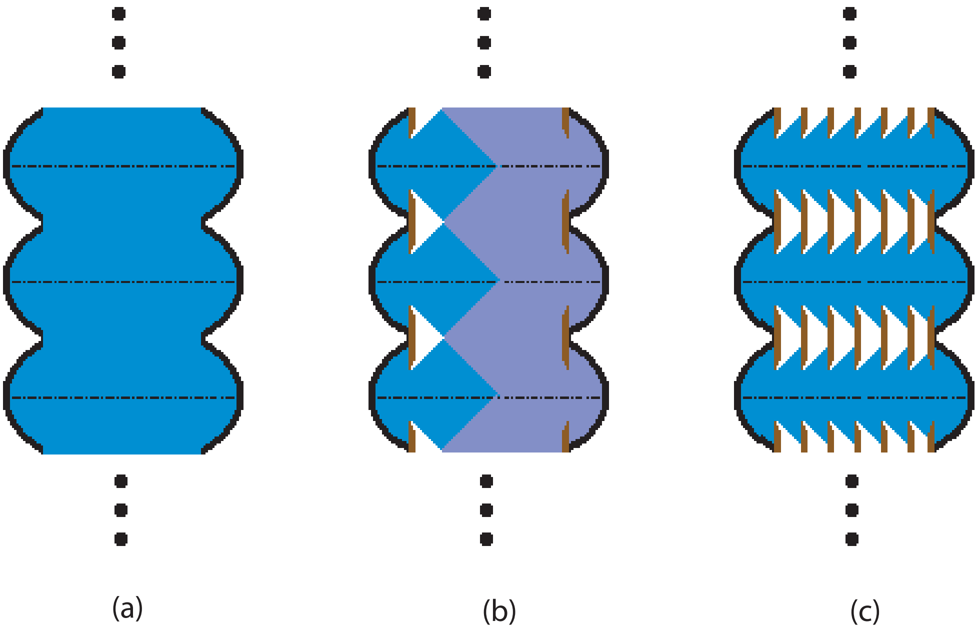

Figure 4.5: Example of FFT propagation for $\frac{1}{2}$ of one round trip. Propagation is to the right. (a) shows the effect of having neither apertures at the mirror, nor along the length of the cavity. (b) shows how diffraction from apertures at the mirrors alone can lead to spuriously low loss. (c) shows a simulated cavity that will give physically correct results at the expense of increased computing time; the $\frac{1}{2}$ round-trip is broken up into 6 FFT steps.

An input wave is created at one lens and then sent through the lens waveguide using a FFT propagation method. Due to the nature of the FFT propagation, our simulation area has periodic boundary conditions, i.e. any light that we might expect to lose, since it has gone outside the simulation area, will reappear from the opposite edge of the area (see figure 4.4). If this effect is not corrected for by including absorbing boundaries (see figure 4.5), the results of the simulation will be spurious (figure 4.6).

Now we must deal with some of the weaknesses of the Fox-Li method. The

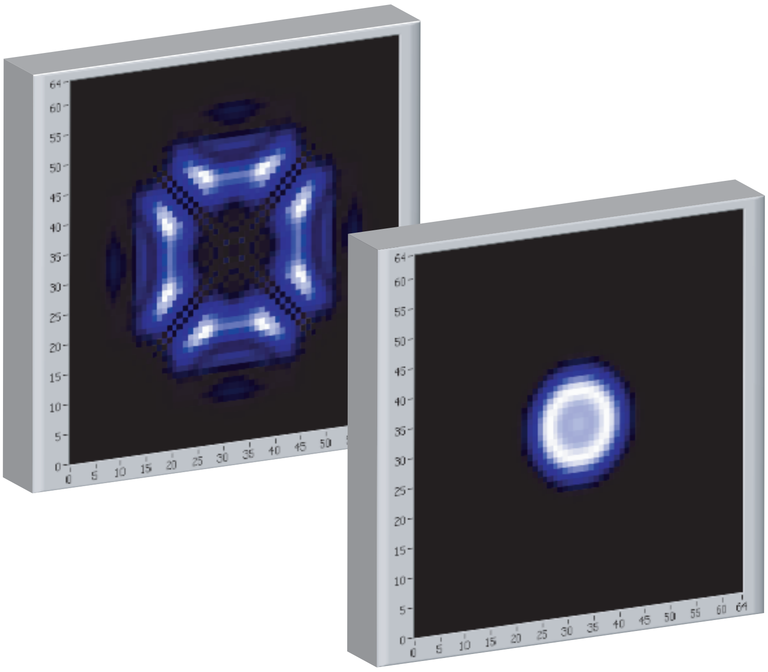

Figure 4.6: The effect of the periodic boundary conditions on beam intensity, if not corrected for by the inclusion of baffles. The front plane is propagated into the rear plane by the *FFT* propagation method with unphysical results.

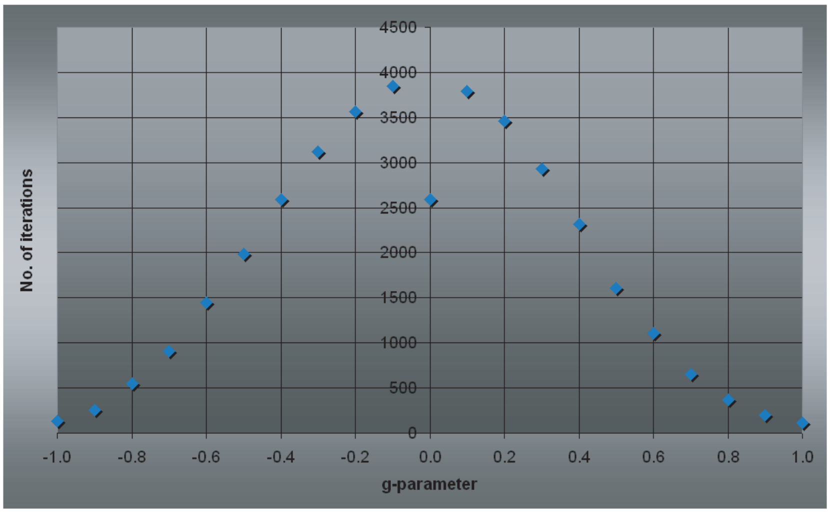

Figure 4.7: This plot shows the number of iterations taken to reach a solution for symmetric stable Gaussian resonators.

first and most obvious of these is an unavoidable consequence of the nature of the Fox-Li method: its dependence upon loss. Figure 4.7 shows the number of round trips taken to converge to a solution for a stable symmetric Gaussian resonator. The loss in question is not usually that of the lowest loss mode, but that of the next lowest loss mode or lowest loss modes. It is the time taken for these to die away that is the chiefly behind the long convergence times of certain modes. A further problem is that this poor *mode discrimination* happens mainly in resonators that are already low loss.

However, looking at the confocal case in the same figure illuminates another problem with the method — and a general weakness of all iterative methods — which can be stated as follows: how do you know when you've reached the solution? Our simulation method used a measure of certain qualities of the beam at the mirror (diffractive loss, uniformity) and when these quantities ceased to vary by a required amount, the simulation had reached a solution. The problem with the confocal case — and as we discovered, certain other cases (see next chapter) — is that the final solution is approached asymptotically. Long before the simulation has reached a solution, the qualities of the beam that we measure have ceased to vary appreciably, leading the simulation to halt far too quickly.

Another problem we must address, and with which we are particularly concerned, is the competition between the robustness of our method and the amount of computation time taken to find an eigenmode. This manifests itself in two ways:

1. The rigour with which we set the termination conditions.
2. How many baffles, and therefore how many FFT steps, will our resonator contain.

The trouble with point (1) can be seen particularly well in confocal resonators. Instead of the normal beating behaviour of the loss per round trip, instead the loss asymptotically approaches the final value. Long before the correct value has been reached the loss is changing so slowly that the simulation terminates. We could make more demanding requirements of the loss, but this would adversely affect the number of iterations required to converge to a solution in other cases.

In relation to point (2): as we have already seen (figure 4.6) failure to include baffles can result in unphysical results, but the inclusion of a single baffle doubles the number of calculations required for a round trip. Since our aim is to map the eigenmodes of a family of resonators (see next chapter), including the minimum number of baffles which still gave the correct result was crucial. We found we could reduce the number of baffles even further if we made all apertures in the resonator, including the baffles, “soft-edged”. That is, having a perfectly absorbing boundary, and a perfectly transmitting centre to the aperture; but instead of a discontinuity between them having a smooth graduation from one to the other. This is achieved by using a function of the form,

$$a(x) = \begin{cases} 1 & : \quad x < r_i \\ \cos^2 \sqrt{\frac{\pi}{2}\frac{x-r_i}{r_o-r_i}} & : \quad r_i \leq x \leq r_o \\ 0 & : \quad x > r_o \end{cases} \tag{4.2}$$

where r_o is the radius of the perfectly absorbing aperture and r_i is the radius of the perfectly transmitting aperture.

Despite this, choosing the correct number of baffles remains more art than science. As can be seen in figure 4.8, it is necessary to have two aperture profiles: one for the mirrors and one for the propagation baffles. The r_o of the mirror aperture must be smaller than the r_i of the baffle aperture otherwise resonators with a g-parameter close to -1 would show a spuriously higher loss than their counterparts at $+1$.

Symmetric canonical resonators

A nice way to check the validity of any simulation is to compare it with analytic solutions, where such solutions exist. By simulating symmetric canonical resonators we can compare our simulation results with the predicted gaussian mode properties outlined in Chapter 2. Specifically, knowing (from equation 2.25) the expected size of the gaussian beam on the mirrors of the resonator; and knowing (from our simulation parameters) the size of the mirror; we can compare the loss per round trip of the evolved Fox-Li eigenmode with the loss per round trip of the predicted gaussian mode.

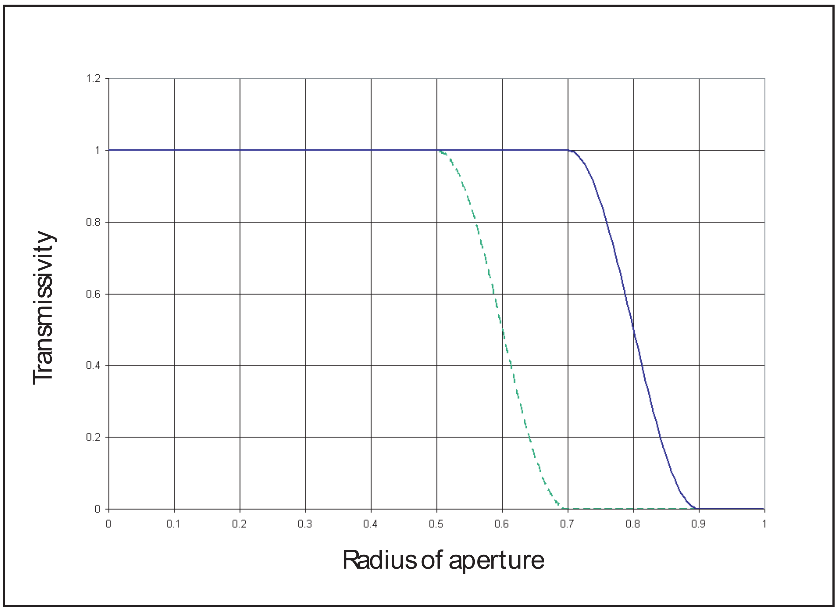

Figure 4.8: Cross-section of the soft-edged aperture that was found to be useful in reducing the required number of baffles. The green dashed line is the profile of the mirror apertures while the blue line is the profile of the propagation baffle.

Section 4.2 shows that our simulation result agreed very well with the predictions in most cases, and discusses some situations where our simulations fail to give the correct results.

Fox-Li method: when do you stop?

So far we have not discussed a fundamental issue of Fox-Li simulations. When do we stop? The case of resonator outputs (discussed above) is simple: the number of lens-waveguide periods to propagate through is defined by the resonator's finesse. But when looking for the fundamental mode, how far down the (in effect) infinite lens-waveguide do we propagate before the 'pancake' is ready?

One of the defining characteristics of a mode of a stable resonator is its structural stability [7], i.e. if we examine the transverse mode-shape at any plane within the cavity and compare the intensity profile with the intensity profile of any other plane in the resonator, the two should (aside from scaling) be the same. Thus some figure of merit based upon the transverse mode-shape, or upon the variation of the transverse mode-shape from one round trip to the next, is ideal. Our original simulations used the diffractive loss per round trip as a figure of merit, but as is shown in the next section the intensity uniformity (defined in equation 4.3) proved to be a more sensitive figure of merit. When the figure of merit remains the same (within preset limits) from one round trip to the next, the simulation has reached a solution, which should be the eigenmode of the equivalent resonator.

4.2 Effects of spherically aberrated mirrors on cavity mode shapes

We decided to explore the effect on mode shape of modifying the mirror profiles with primary and secondary spherical abberation terms. The motivation for this was to find resonator eigenmodes with the same power at the mirror and the same diffractive loss as a reference gaussian eigenmode, but with a lower peak intensity (similar to mesa beam resonators [38]). This would decrease the thermal loading of the mirror in such high-power applications as 2^{nd} generation gravitational wave detectors (outlined in greater detail in chapter 5). The intensity distribution of the eigennmode would have greater uniformity, where uniformity is defined as,

$$U = \frac{1}{N^2} \frac{(\sum i)^2}{\sum(i^2)}, \tag{4.3}$$

summed over the pixels of the simulation, where N^2 is the total number of pixels, and the individual pixel intensities are i. This definition of uniformity is a modified participation function, where having a flatter intensity profile and a more spread-out beam spot gives a greater uniformity. Obviously, the presence of an aperture at the point where we measure the intensity (i.e. the size of the mirror) means that some pixels will always be dark – and more subtly, due to the cosine edge, some will be artificially darker than they should be – but for the purposes of our simulation, simply maximizing the uniformity while maintaining the same diffractive loss is sufficient. It should be noted that the uniformity function can be combined with the mirror aperture function so that a normalized, uniform intensity across the mirror (modulated by the cosine edge) produces $U = 1$. This gives the absolute value of U some physical meaning, but is computationally more costly (important, since U is calculated once per resonator round-trip) and was unnecessary for our purposes.

The Zernike polynomials form a complete set on the unit circle and are often used in the analysis and correction of optical abberations [39]. Our code represents the mirror surface as a superposition the first four cylindrically symmetric Zernike polynomials.

The first of these is conventionally known as *piston*,

$$z_0^0 = 1. \tag{4.4}$$

This simply moves the surface forward or backward without changing the shape. In our simulation, it could be thought of a length tuning of the resonator by $\pm\pi$.

Defocus,

$$z_0^2 = 2x^2 + 2y^2 - 1, \tag{4.5}$$

is the equivalent of focal length. In essence, it defines the spherical mirror which the next two terms deform.

Primary spherical aberration,

$$z_0^4 = 6x^4 + 6y^4 + 12x^2y^2 - 6x^2 - 6y^2 + 1, \tag{4.6}$$

and secondary spherical aberration,

$$z_0^6 = 20x^6+20y^6+60x^4y^2+60x^2y^4-30x^4-30y^4-60x^2y^2+12x^2+12y^2-1, \tag{4.7}$$

are then used to reshape the spherical mirror defined by equation 4.6. The mirror shape is defined by the superposition of these four terms, i.e.

$$S = \alpha_0 z_0^0 + \alpha_1 z_0^2 + \alpha_2 z_0^4 + \alpha_3 z_0^6 \tag{4.8}$$

and the resulting shape is used as the basis for a phase hologram for a mirror. For simplicity, the cavity was set to consist of two identical aberrated mirrors a distance L apart, or a single aberrated mirror a distance $L/2$ from a plane mirror. A parameter space of $21 \times 21 \times 21$ resonators was mapped out from $g = -0.8$ to $g = 0.8$ corresponding to the z-axis, primary spherical aberration (applied both positively and negatively) to the x-axis, and secondary spherical aberration (again applied both positively and negatively) corresponding to the y-axis. This procedure took of the order of 4 weeks of computing time. A representative plane of the resulting data is shown in figure 4.9. This plane might be called the '$g = 0.3$ plane' although the g-parameter has a meaning only for the central (unaberrated) resonator. Indeed there are resonators in this plane (top left and bottom right) which are effectively unstable. The features of each plane are similar and consist of a low loss 'valley' running from bottom left to top right, caused be primary and secondary spherical abberation effectively correcting for each other.

One would intuitively expect that smoothly varying some resonator parameter would result in a smooth variation in the mode properties of the resonator – unless, of course, some fundamental limit were crossed, such as the transition

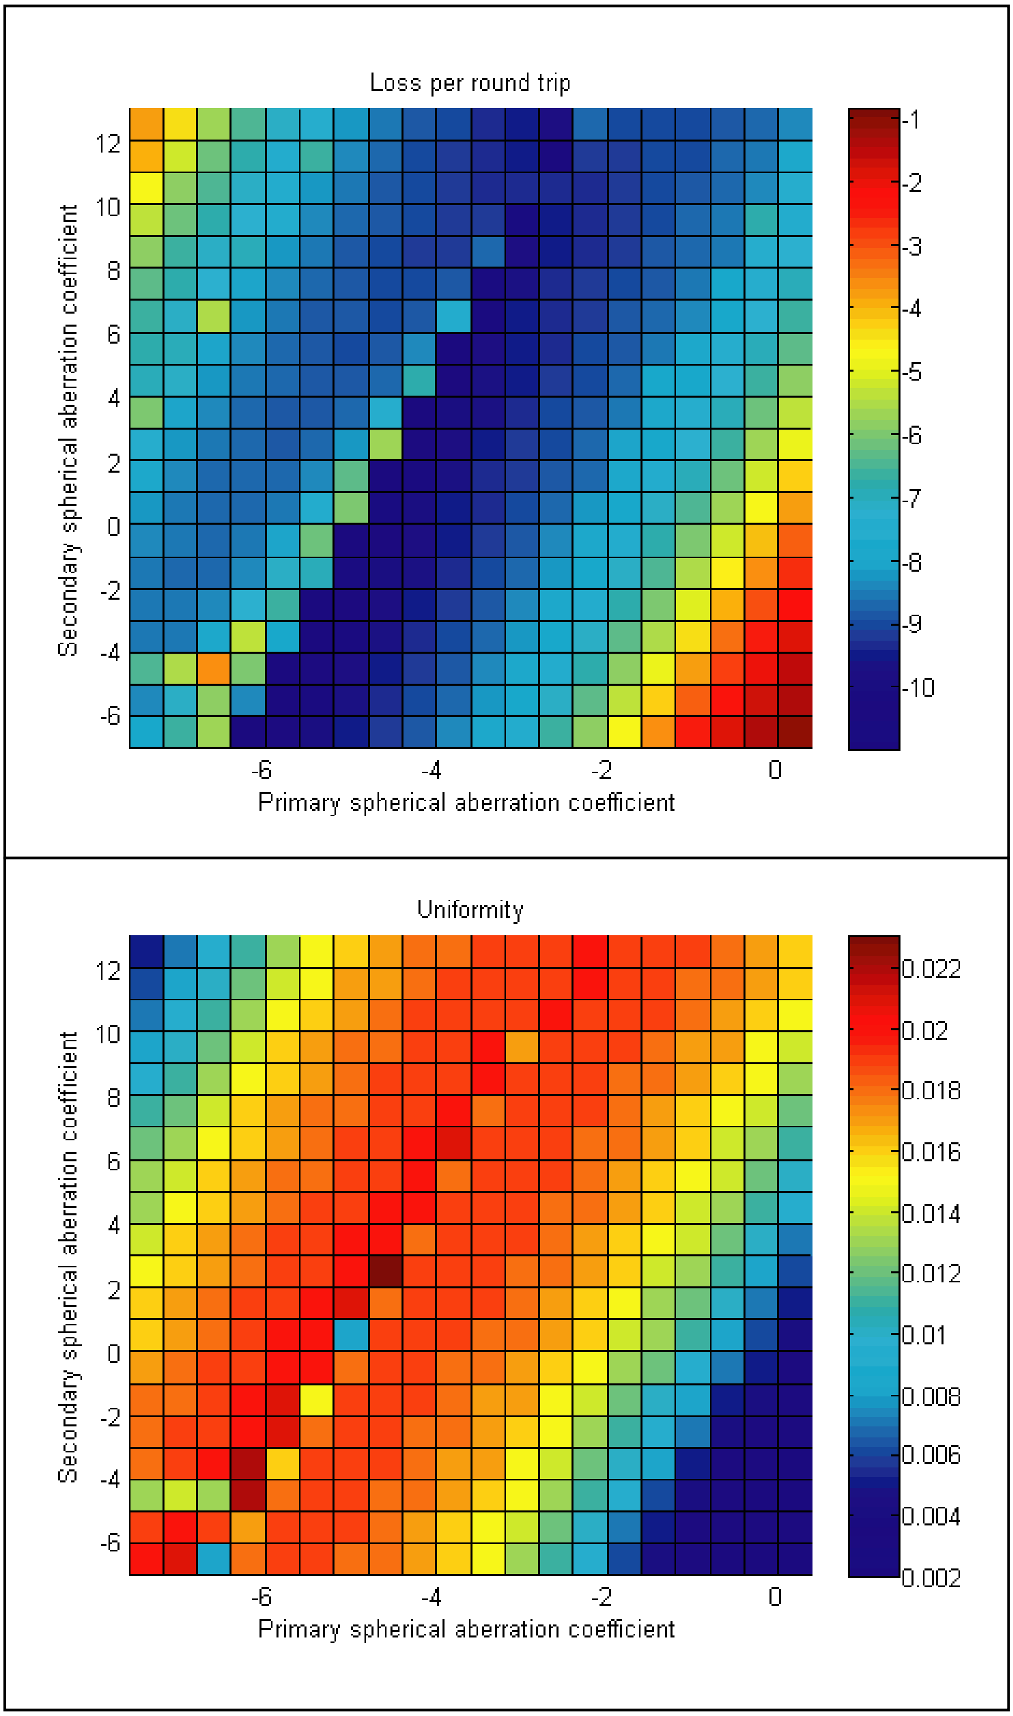

Figure 4.9: A contour map of loss and uniformity variations (with log(loss) displayed to highlight the variations) as primary and spherical aberration are applied to a $g = 0.3$ cavity. The z-axis (not shown here) varies the initial (unaberrated) g-parameter of the cavity. This is one plane of 21×21 values from a series of 21 planes, which can be imagined as stacked above and below this one.

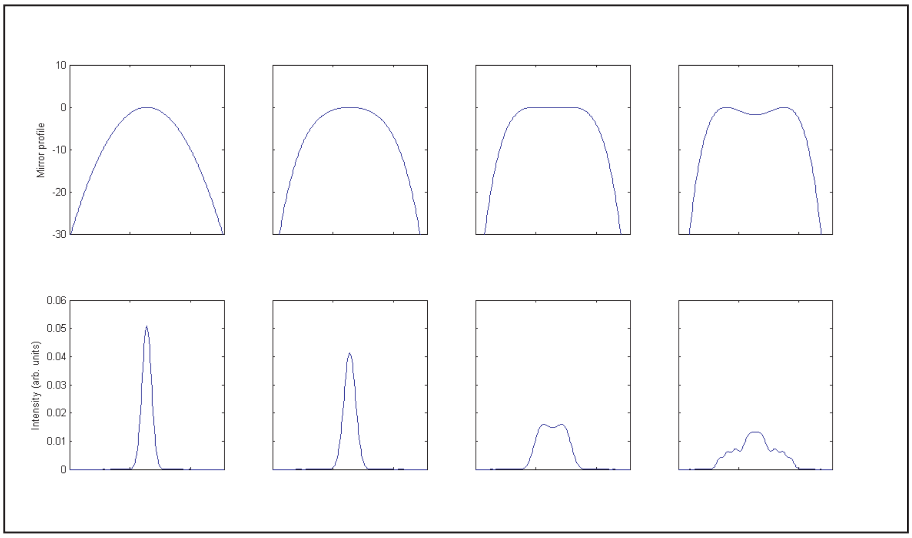

Figure 4.10: The top row shows the effect of increasing primary spherical aberration upon the shape of a $g = 0.6$ cavity mirror. The bottom row shows the associated intensity pattern of the fundamental mode of the cavity.

from stability to instability. Although figure 4.9 shows hints of this, we also see certain resonators behaving quite differently from their neighbours. We believe these to be due to problems with the Fox-Li simulation itself, rather than properties of the resonators (see section 4.2). Disregarding these isolated points, we disappointingly see a monotonic relationship: increasing the uniformity of an eigenmode by aberrating one or both cavity mirrors results in an increase in the diffractive loss per round trip.

Many of what might be termed *re-entrant* mirrors (i.e. mirror with a central bulge of opposite curvature to the main body of the mirror, as with the right hand-most mirror profile of figure 4.10) had, perhaps unsurprisingly Laguerre-Guassian fundamental modes. These mirrors essentially have a g-parameter that is greater than unity close to the optical axis, decreasing to less than unity as we approach the mirror's edge - they form a resonator that is unstable on the optical axis, but stable at the edges. This is perhaps equivalent to occluding the central portion of one or both mirrors of a stable resonator with some non-reflective material; or to an axicon resonator [40, 41].

One question might be asked: what are the benefits of performing these simulations with $2D$ mirrors, with the greatly increased time and computational costs? Why not examine the behaviour of strip resonators? The answer is that, if a promising mode was found, its sensitivity to tilt and misalignment of the cavity could be investigated. Since none were found, and there remained unanswered questions as to the reliability of the simulations, this was never systematically undertaken.

Pathological cases

During these many cavity simulations, we observed two particular cases of our Fox-Li simulations producing unexpected, and perhaps incorrect, results.

The first of these cases occurs when simulating stable canonical resonators. Examining figure 4.11 we see that, for resonators with certain values of g, the simulated loss does not match the theoretically predicted value. That is, for $g = 0$ and $g = \pm 0.5$ (and, although not shown $g = \pm\sqrt{2}/2, \cdots$), the simulated loss is lower than predicted. Various attempts were made to remedy this behaviour: for instance, allowing these particular resonators an order of magnitude more round trips to settle into a mode, changing the Fresnel number of the resonator, etc. All were unsuccessful.

One possible explanation for the disagreement, however, can be outlined by considering the imaging properties of stable canonical resonators. In general, every transverse plane within a stable resonator is imaged into every other plane within the resonator – this is the key to the ray-optical explanation of stable resonator eigenmodes [7]. But there exist a subset of stable canonical resonators in which every plane is imaged into only a finite number of planes before being imaged back into itself; and it is exactly these resonators which display the lower than predicted loss. In the perfectly aligned resonator, represented in the simulation, it takes a longer time, by several orders of magnitude, for the field within the cavity to reach an eigenmode.

The second and perhaps more worrying case occurred during the simulation of cavities with spherically aberrated mirrors: an example is shown in figure 4.12, and similar results were observed when simulating many other similar resonators. All share a typical initial evolution, of appearing to converge stably to an eignenmode. If one examines only the loss per round trip one could be forgiven for thinking that this was the case. The modes uniformity, however, tends to show an oscillatory behaviour which increases until the mode starts to 'flip' into another, lower loss, mode. On occasions this mode, too, will appear stable but then flip to a lower loss mode. The final mode appears to be physically reasonable (for instance, the mode never flips into a higher loss state), but this behaviour raises the question of the reliability of our Fox-Li simulations. As was mentioned above (section 4.1), it is usual to stop the simulation when the loss per round trip (or equivalently the normalization factor) is the same

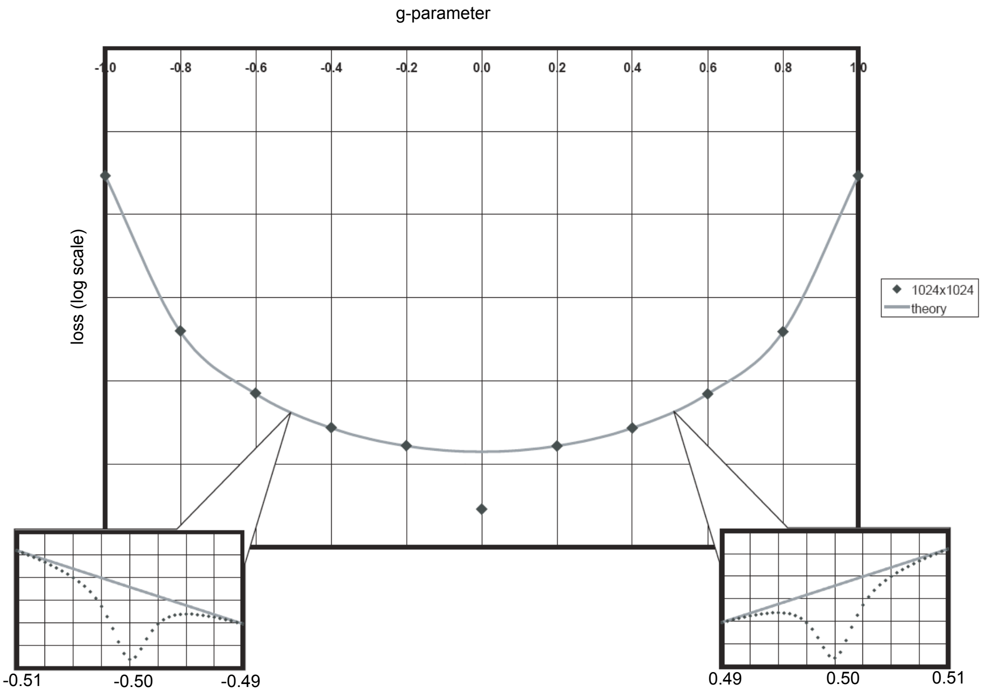

Figure 4.11: Loss per round trip vs. g-parameter for stable symmetrical canonical resonators. The theoretical prediction is for the assumed gaussian mode of such a resonator while the simulated results are from Fox-Li simulations, carried out on a 1024 × 1024 grid. There is disagreement at $g = \pm 0.5$ (shown in greater detail below the main graph) and at $g = 0$. One possible explanation is that these are 'magic number' resonators.

(to within some tolerance) from one iteration to the the next. If this procedure were followed for the example in figure 4.12, the simulation would terminate somewhere around the 200^{th} iteration. We have seen (from figure 4.2) that the uniformity can be a better indicator of convergence to an eigenmode than loss per round trip – particularly in resonators with poor mode discrimination.

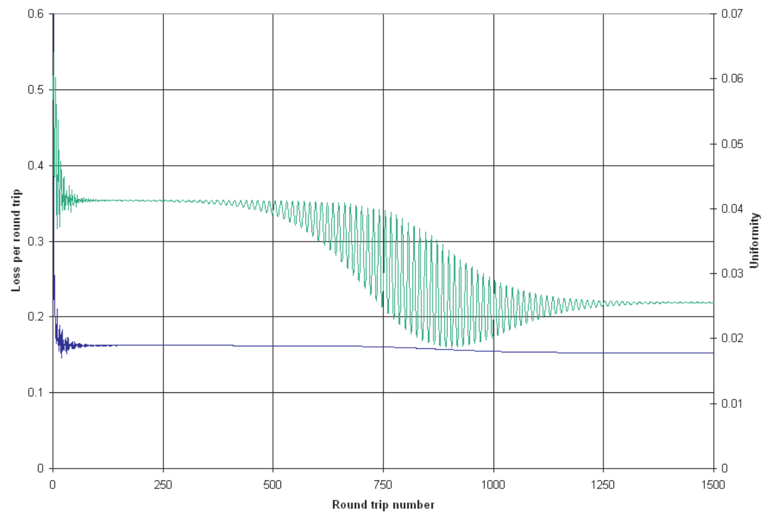

Figure 4.12: Untypical behaviour during a Fox Li simulation of a spherically aberrated cavity. The simulation appears to converge to a solution after only one hundred iterations. The uniformity then begins to oscillate and the simulation then converges on a lower loss mode than the first.

4.3 Fractal Eigenmodes of Unstable Resonators

That the transverse eigenmodes of unstable canonical resonators could possess fractal structure was an intriguing discovery [42, 43, 44]. Further simulation showed that manipulation of the parameters of unstable resonators (mirror curvature, resonator length, aperture and transverse offset of the aperture) could give rise to classic fractal patterns in the resonator's eigenmodes [25]. One question that remained to be answered was whether careful choice of parameters might reveal a resonator with a 3D classic fractal pattern as it's eigenmode. As a first step towards answering this, we examined the on-axis intensity of Fox-Li simulated unstable resonators and compared them with the results of a simple algorithm based on the Monitor-Outside-a-Monitor effect [45], described in the next section.

Mathematical fractals are structures which display self-similarity at all scales [46]; fractals in the real world always fail to be self-similar at some scale level. Our fractals provide an elegant demonstration of this since, in the reverse of the usual method, they are built up from successively magnified portions of the original pattern. Thus, the minimum scale is already set at the beginning of the fractals construction.

During each round trip through a canonical optical resonator, light is imaged geometrically by the resonator's two curved mirrors. Geometrical imaging of planes in canonical optical resonators is important both in unstable and in stable resonators: in unstable resonators it can explain the formation of fractal intensity structure in transverse cross-sections through the eigenmodes [20]; in stable resonators it explains the eigenmodes' structural stability [7].

The Monitor-Outside-a-Monitor (MOM) effect

Directing a video camera at the monitor which displays a live picture of what the camera sees results in video feedback [47]. The *monitor-inside-a-monitor* effect is the better known example of this; it occurs when the magnitude of the overall magnification of the camera and monitor, M, is less than 1 (see the

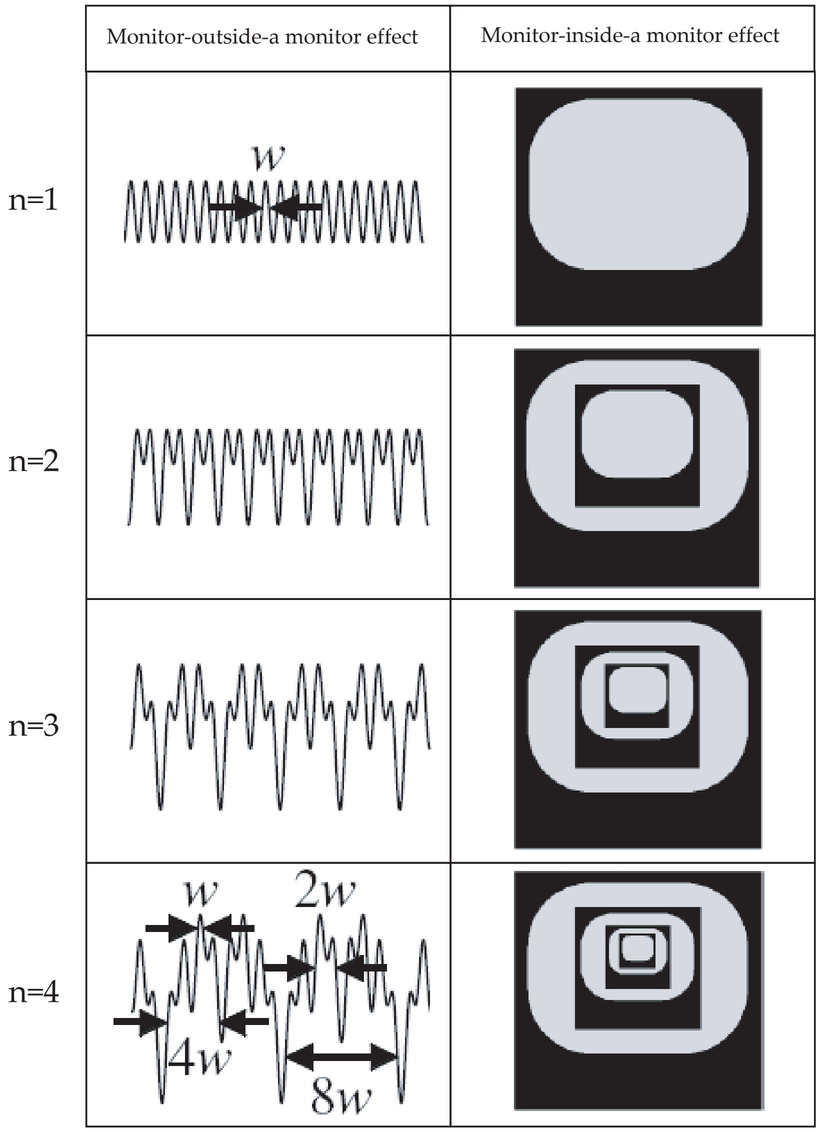

Figure 4.13: The first 4 iterations of the monitor-inside-a-monitor effect are shown in column (a), and the first 4 iterations of the monitor-outside-a-monitor effect are shown in column (b).

right hand column of figure 4.13. Less well known is the *monitor-outside-a-monitor* effect. This occurs when $|M| > 1$ and with the further proviso that the camera must be capable of resolving the individual pixels of the monitor. Then the pixel pattern of the monitor is repeatedly imprinted on the initial pattern at greater magnification, and from this smallest scale the system forms structures M, M^2, M^3, etc. times larger [48]. The left hand column of figure 4.13 shows that this quickly results in a fractal pattern, albeit one which possesses a smallest-possible scale, w. In 1D we can say that the pattern after n iterations arising from an initial periodic function, $p(x)$, is given by,

$$f_n(x) = \sum_{i=0}^{n-1} p(\frac{x}{M^i}). \tag{4.9}$$

Unstable canonical resonators can also exhibit an analogue of the monitor-outside-a-monitor effect. The resonator itself acts as the magnifier, and the periodic function, $p(x)$, is provided by the diffraction pattern from one of mirror apertures. We can ensure that the pattern comes from only one of the mirror apertures by making one of the mirrors much larger than the other. The role of the 'monitor' is taken by the *self-conjugate plane* – this is the plane within the resonator that is imaged back onto itself after every round-trip. The fractal nature of the transverse mode shape in the self-conjugate plane has already been investigated [25]. We suspected that there might also exist resonators which possess a fractal on-axis intensity. We also suspected that it might be possible to engineer an unstable resonator which had, as it's eigenmode, an intensity pattern related to one the classic 3D fractals, such as a Sierpinski gasket [46]. As a first step towards this, and in order the

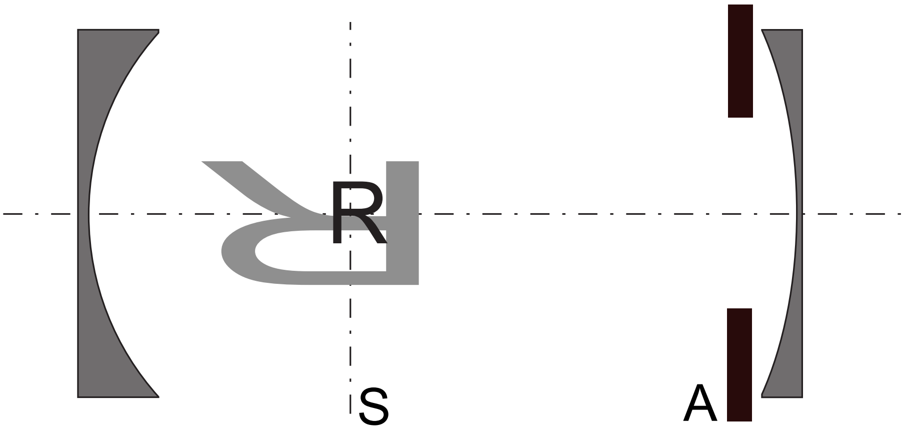

Figure 4.14: In an unstable confocal resonator, the volume around the self-conjugate plane, S, is imaged back onto itself with a transverse magnification, M_T, and a longitudinal magnification, $M_L = M_T^2$. In this illustration $M_T = -2$ and the upright "R" is imaged back onto itself inverted and stretched. The aperture, A, of one mirror is much smaller than the other mirror; as a result, only the diffraction pattern from this mirror contributes to the fractal pattern.

gain an understanding of the fractalicity displayed by some unstable resonator eigenmodes, we decided to investigate unstable resonators which have fractal patterns in their eigenmode's intensity, both in the self-conjugate plane and along the optical axis of the resonator.

We restricted ourselves to unstable confocal resonators, i.e. those resonators where $L = F + f$ and $F > f$. Our reasons where that the transverse and longitudinal magnifications of the volume around the self-conjugate plane are well understood in these resonators, being given by,

$$m_T = -\frac{F}{f}, \tag{4.10}$$

and

$$m_L = m_T^2 = \frac{F^2}{f^2}, \tag{4.11}$$

respectively (see figure 4.14). Thus, for example, to achieve a longitudinal magnification of 2 requires that $m_T = \sqrt{2}$, leading to resonators of the form,

$$L = F + \frac{F}{\sqrt{2}}. \tag{4.12}$$

Our other requirement for the resonator is that there should be a sufficiently regular intensity pattern on-axis in the region of the self conjugate plane (and also a sufficiently regular transverse intensity patten in the self-conjugate plane, but this is a much more easily fulfilled requirement). Ideally we would like a perfectly regular variation in intensity (for example, an intensity that varied like a sine wave). Some (numerical) experimentation revealed the patterns shown in figure 4.15: (a) shows a cross-sectional view of the first half-round-trip of an unstable confocal resonator, with the light propagating from left to right, and from 0 to L. The incoming plane wave encounters a circular aperture at 0. In (b), the intensity on the optical axis of the resonator is shown to the same scale; the region around the self conjugate plane (for a confocal resonator with $m_L = \sqrt{2}$) is outlined in red, and shown with greater magnification in (c). By appropriate choice of aperture radius, a, and resonator length D, we can make the on-axis diffraction pattern arbitrarily regular. We are, in effect, looking at geometrically the same resonator with different equivalent Fresnel numbers [34],

$$N_{eq} = \frac{a^2}{\lambda D}\sqrt{\frac{g_1(g_1 g_2 - 1)}{g_2}}, \tag{4.13}$$

and, indeed, the same effect could be achieved by keeping a and D constant while varying λ.

In order to gain this information about the on axis intensity (and also to make sure the problems outline in figure 4.6 do not occur) it is necessary to propagate the light field around the resonator in many steps. For comparison, a stable symmetrical resonator might be simulated with 4 to 16 propagation steps per half-round-trip. The information in figure 4.15 comes from propagation with 2048 steps per half round trip. The resolution of the simulation grid in the transverse plane is also large (typically 1024×1024) in order to see the details of the fractal intensity pattern in the self conjugate plane. As might be imagined, simulating such a resonator in this way is computationally very expensive – fortunately, due to the high loss nature of unstable resonators, the simulation typically converges to an eigenmode within a few tens of round trips.

Figure 4.16 shows the on-axis intensity pattern in an unstable confocal resonator where $m_L = 2$ as the simulation proceeds. The 0th round trip is simply the initial aperture diffraction pattern, which will act as the periodic function which will be magnified and superposed on subsequent round trips.

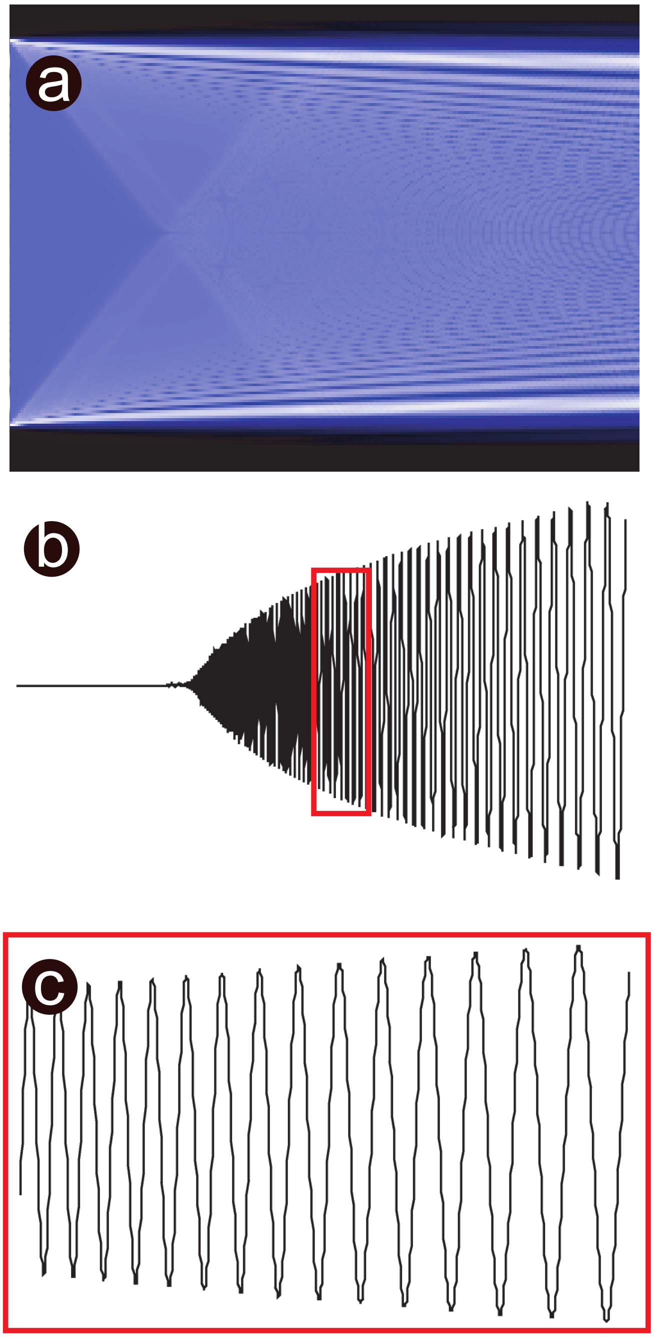

Figure 4.15: A plane wave, propagating from left to right, encounters a circular aperture resulting in the diffraction pattern shown in (a). The intensity in the optical axis is shown in (b), and a close up of the diffraction pattern in the region of the self-conjugate plane is shown in (c).

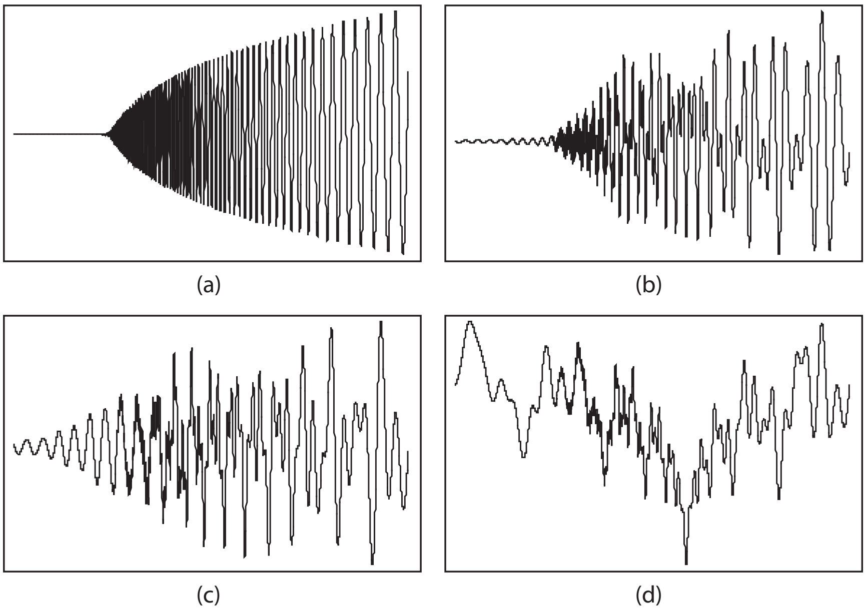

Figure 4.16: The on axis intensity after (a) 0, (b), 1, (c) 2, and (d) 10 round trips of the resonator. The 0^{th} round trip is essentially the aperture diffraction pattern.

Figure 4.17 shows the comparison of a monitor-outside-a-monitor effect performed upon an initial sine wave, with a Fox-Li simulation of the on-axis intensity of an unstable resonator in the region of the self-conjugate plane. The agreement of such a simple model with a complex wave-optical simulation of a complex resonator leads us to believe that a clever choice of parameters might yield an unstable resonator with a classic 3D fractal eigenmode.

4.4 Conclusion

In this chapter we have shown that our version of the Fox-Li method can recreate the expected transverse eigenmode of canonical optical resonators,

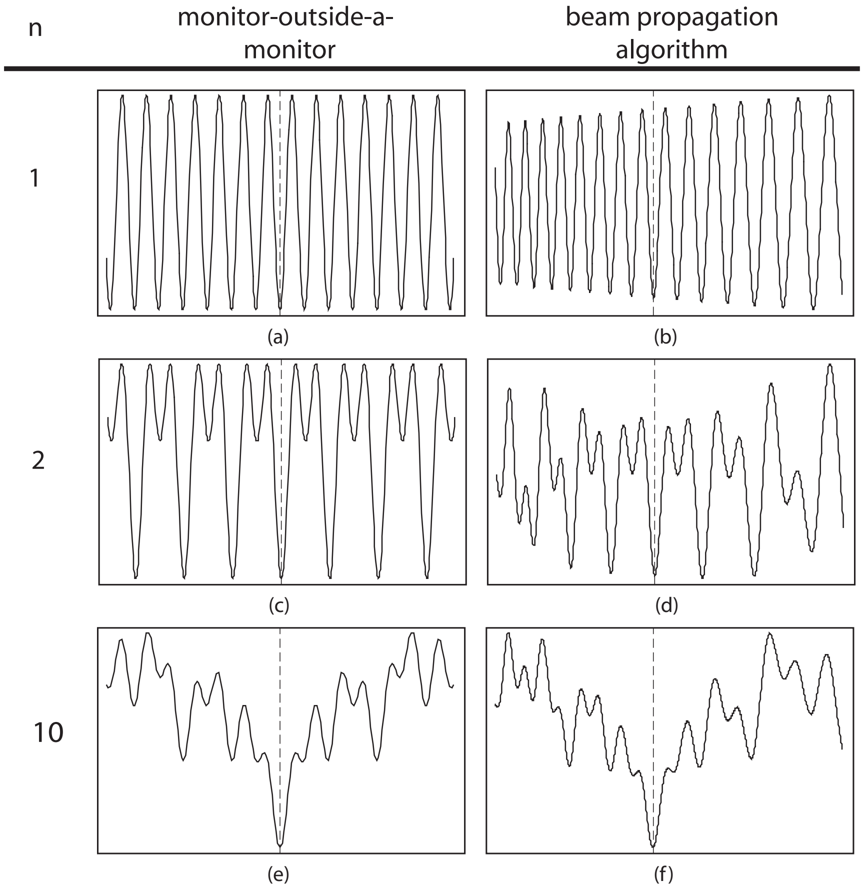

Figure 4.17: Comparison of MOM effect calculations and Fox-Li simulations.

with the exception of 'magic number' resonators. This *caveat* has not been reported before and therefore may be due to our beam-propagation method. We attempted to discover whether spherically aberrated resonators would support an eigenmode of greater uniformity (i.e. a flatter intensity profile) and similar diffractive loss, but due to difficulties with the simulation we were unable to answer this question.

CHAPTER 5

Gravitational Waves: Sources and Detection

In this chapter we briefly discuss the underlying physics of gravitational wave generation and propagation. We also catalogue some of the astrophysical objects which are likely to generate detectable (by 2^{nd} generation detectors – see section 5.3) gravitational radiation. Finally, we look at some of the proposed methods for detecting gravitational waves, and discuss the sensitivity limits of interferometric gravitational-wave detectors.

All equations in this chapter assume the Einstein summation convention.

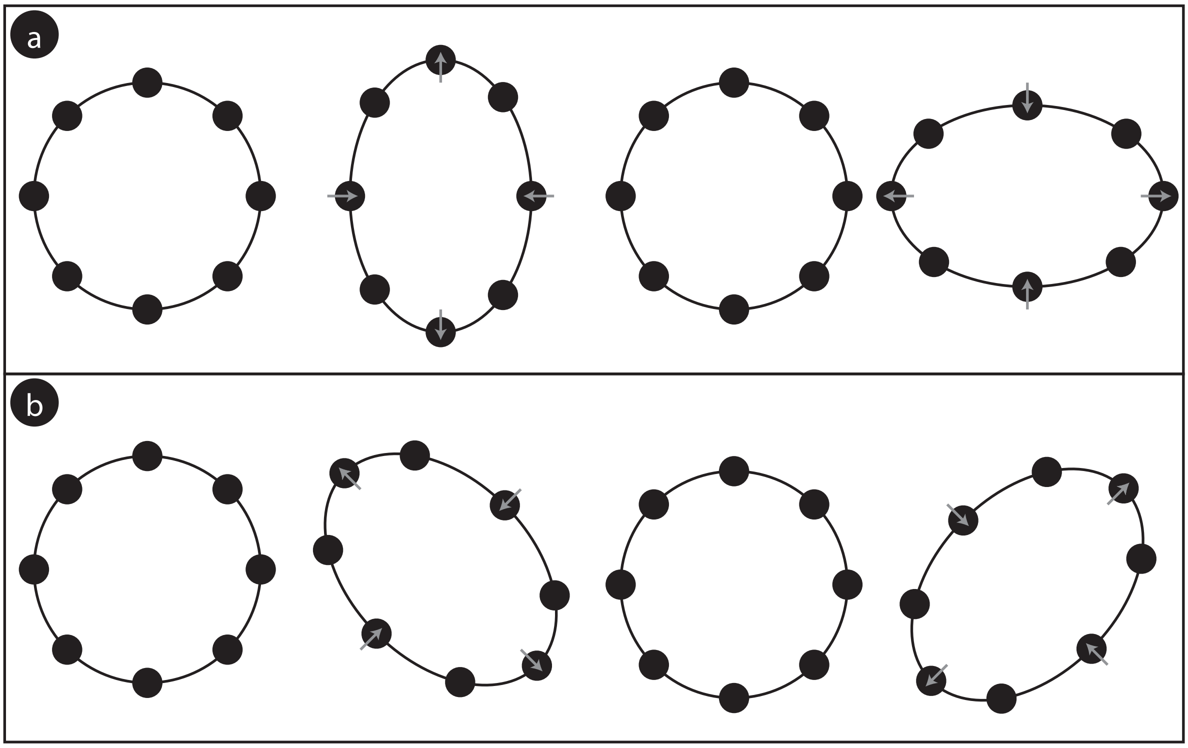

Figure 5.1: Time evolution of a number of proof masses, under the influence of a gravitational wave propagating into the page with (a) h_{22} (or +) polarization, and (b) h_{23} (or $\times$) polarization.

5.1 Gravitational Waves

Gravitational waves can be thought of as ripples in the curvature of space-time. If we imagine two spatial dimensions of a locally flat spacetime (see figure 5.1) then a suitably polarized gravitational wave has the effect of stretching one axis while contracting the other. The magnitude of this stretching effect is governed by the amplitude of the gravitational wave, h. This is physically observable by measuring the change in separation in two freely moving proof-masses, i.e.

$$h = \frac{2\Delta L}{L}, \tag{5.1}$$

where L is the separation of the proof-masses and ΔL is the change in that separation.

Detecting this length change arising from the gravitational-wave induced strain in space has proven to be a somewhat tricky proposition: the predicted amplitudes at the Earth's surface from even the most violent astrophysical events are of the order of 10^{-21}. So if $L = 1\,\mathrm{km}$ we are attempting do detect a length change of $10^{-18}\,\mathrm{m}$ — less than the radius of an atomic nucleus.

The equations of gravitational waves are the equations of General Relativity: the Einstein Field Equations (EFE)[49]. This set of equations, which describe how objects curve spacetime and the resulting gravitational interaction, can be succinctly written in tensor form as,

$$R_{\mu\nu} - \frac{1}{2}\mathbf{R}g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}, \tag{5.2}$$

where $R_{\mu\nu}$ is the Ricci tensor, $\mathbf{R}$ is the scalar curvature, $g_{\mu\nu}$ is the metric tensor, G is the gravitational constant, c is the speed of light in a vacuum, and $T_{\mu\nu}$ is the stress energy tensor.

Represented in this form – as Thorne [50] notes – gravity is nonlinear and it is therefore not possible to precisely separate the contributions of gravitational waves to the curvature of spacetime from the background curvature caused by massive objects. In practice this is not a problem. An analogy can be instructive: imagine the surface of a stormy sea, with large waves – this represents the geometry of space-time. The waves and troughs represent the curvature created by local massive objects. Now imagine dropping a large stone into the sea: the outward propagating ripples are the gravitational waves, and are different from the larger waves only in scale. But this difference of scale in the various phenomena does give us a clue to how to attack the otherwise intractable mathematical problem – the shortwave approximation, which is closely related to semiclassical WKB approximation [51]. It works because in typical astrophysical situations the length scale of the gravitational waves is significantly smaller than the length scales over which the other curvatures vary. Thus we can separate the Ricci tensor into the background component (which is the average curvature over many wavelengths of the gravitational wave) and the gravitational wave component,

$$R_{\mu\nu} \equiv g_{\mu\nu} + h_{\mu\nu}. \tag{5.3}$$

A further simplification can be obtained by setting the background term to be the term for flat space. Analytic solutions to equation 5.2 remain difficult; but with the further assumption that the magnitude of $h_{\mu\nu}$ is small we can make a linear approximation to a solution of the form,

$$(\bigtriangledown^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2})h_{\mu\nu} = 0, \tag{5.4}$$

i.e. a wave equation, where the wave propagates with velocity c. In the transverse traceless gauge [52] a gravitational wave propagating in the $z-$direction can be represented as the superposition of two polarisations,

$$\hat{h}_+ = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \tag{5.5}$$

and

$$\hat{h}_\times = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{5.6}$$

The invariant spacetime interval between two points is given by

$$ds^2 = g_{\mu\nu}dx^\mu dx^\nu. \tag{5.7}$$

As the name suggests, this is an invariant quantity, but it's projection along certain axes can be changed.

The Earth is sufficiently far away from astrophysical gravitational-wave sources for us to consider the gravitational waves reaching us to be planar. Consider a Cartesian coordinate system located on the (locally flat) surface of the Earth with the z-axis normal to the surface (i.e. in the opposite direction to $\vec{g}$), with the (plane) gravitational wave propagating in the $-z$ direction. Here on Earth, the gravitational waves reaching us from distant astrophysical sources should (hopefully [53]) always be weak, and hence we can consider them as perturbations of an otherwise flat spacetime. If the the gravitational wave is of the h_+ polarisation, we will see the $x-$ and $y-$ axes change length differentially (see figure 5.1(a)). If the wave is of the $h_\times$ polarisation then, although the space around our axes will be distorted by their passage (see figure 5.1(b)), we will see no length change of our $x-$ and $y-$axes. The question remains as to how to measure these possible length changes. As we shall see, a long-arm-length Michelson interferometer provides an excellent method of measuring differential length changes.

Very few analytic solutions of the Einstein Field Equations for the generation and propagation of gravitational waves exist. The mathematical difficulties of the generation problem and the propagation problem are very different, and thus require different approximations of the underlying theory in order to deal with them. This is typically accomplished by dividing the space being modeled into three zones (see figure 5.2). (This can be thought of as similar to the problem of solving Maxwell's equations for the near field of some optical element and propagating an optical wave using, say, scalar diffraction theory, but without the comfort of a rigourous mathematical link between the two.) One set of mathematical tools is used in the wave-generation zone (where $r \leq r_I$), and the local wave zone (where $r_I \leq r \leq r_O$); another set of tools entirely is used in the distant wave zone ($r \geq r_O$), but with the requirement that both techniques give matching results in the local wave zone. In the absence of a rigourous link between the two formalisms, some confidence can be gained by applying this technique to problems for which exact, analytic solutions are known, e.g. [54].

5.2 Sources of Gravitational Waves

After this brief description of the generation and propagation of gravitational waves, let us consider the real astrophysical objects which are the source of the waves. It is customary to divide these sources into three classes, grouped by there behaviour over a certain time-scale (for instance, that of a gravitational wave detector's observing run). *Burst sources* last only for a few cycles. *Periodic sources* consist of superpositions of sinusoidal waves, whose frequencies are more or less constant during the course of an observing run. *Stochastic sources* give rise to waves which fluctuate stochastically and last for a long time compared with an observing run.

Type II supernovae occur when the core of a massive star collapses to become a neutron star, while type I supernovae are believed result from nuclear explo-

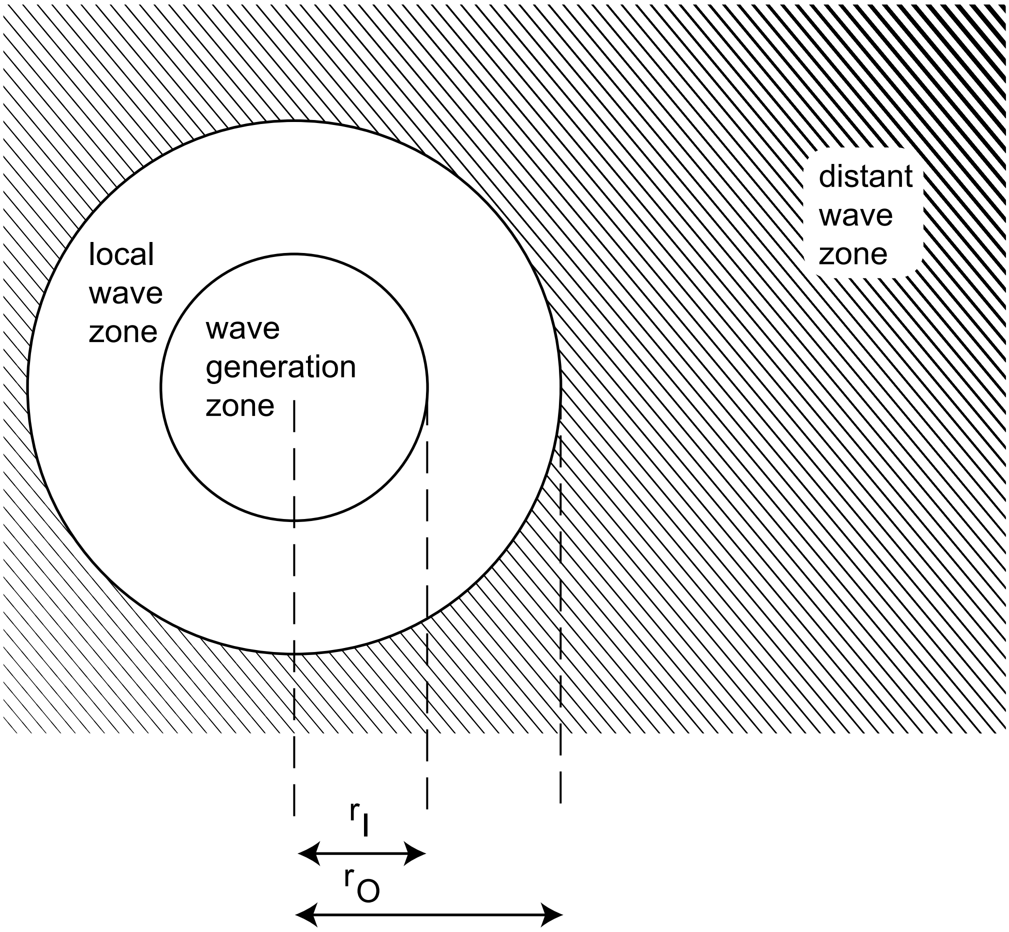

Figure 5.2: Different mathematical formalisms are used in describing the generation of gravitational waves (wave generation zone) and the propagation of the resulting waves (distant wave zone). The techniques are matched together in the local wave zone.

sions caused by the accretion of matter from a nearby large companion star by white dwarf stars. Observation leads us to expect around 2 type I supernovae and 2 type II supernovae per century, in the Milky Way (i.e. within a distance of $\sim 0.02\,\mathrm{Mpc}$). Extending the observational distance to $10\,\mathrm{Mpc}$ (including the centre of the Virgo Cluster) raises the event rate to a few per year for both type I and II – a rate that gives the putative gravitational wave observer something to work with. The asphericity of the event has a large effect on how strong the emitted gravitational waves will be – a completely spherically symmetric stellar collapse will produce no gravitational waves, while a highly aspherical collapse will produce strong waves. In the case of type II supernovae, the resulting neutron star may itself be highly asymmetric and may continue to emit gravitational waves until it reaches a ground state of perfect sphericity.

The strain observable on Earth from a this type of source, a distance $r\,$Mpc away, is estimated to be,

$$h \simeq 5 \times 10^{-22}\left(\frac{E}{10^{-3}M_{\odot}c^2}\right)\left(\frac{15Mpc}{r}\right)\left(\frac{1kHz}{f}\right)\left(\frac{1ms}{t}\right), \tag{5.8}$$

where E is the amount of gravitational energy emitted at the frequency f during an observation time t.

The coalescence of compact binary systems is another source of bursts of gravitational waves. These black hole or neutron star binaries will spend a long time spiralling towards each other with a decreasing orbital period before coalescing . An example of this type of system is the well-known PSR B $1913+16$ (see figure 5.3) which will coalesce in approximately $300,000,000$ years from now. Numerical relativity simulations are required to understand the final stages of binary inspirals, but it is clear that the emitted gravitational waves become stronger as the system comes closer to coalescence. Again, taking the system to be a distance $r\,$Mpc from Earth, the estimated strain is,

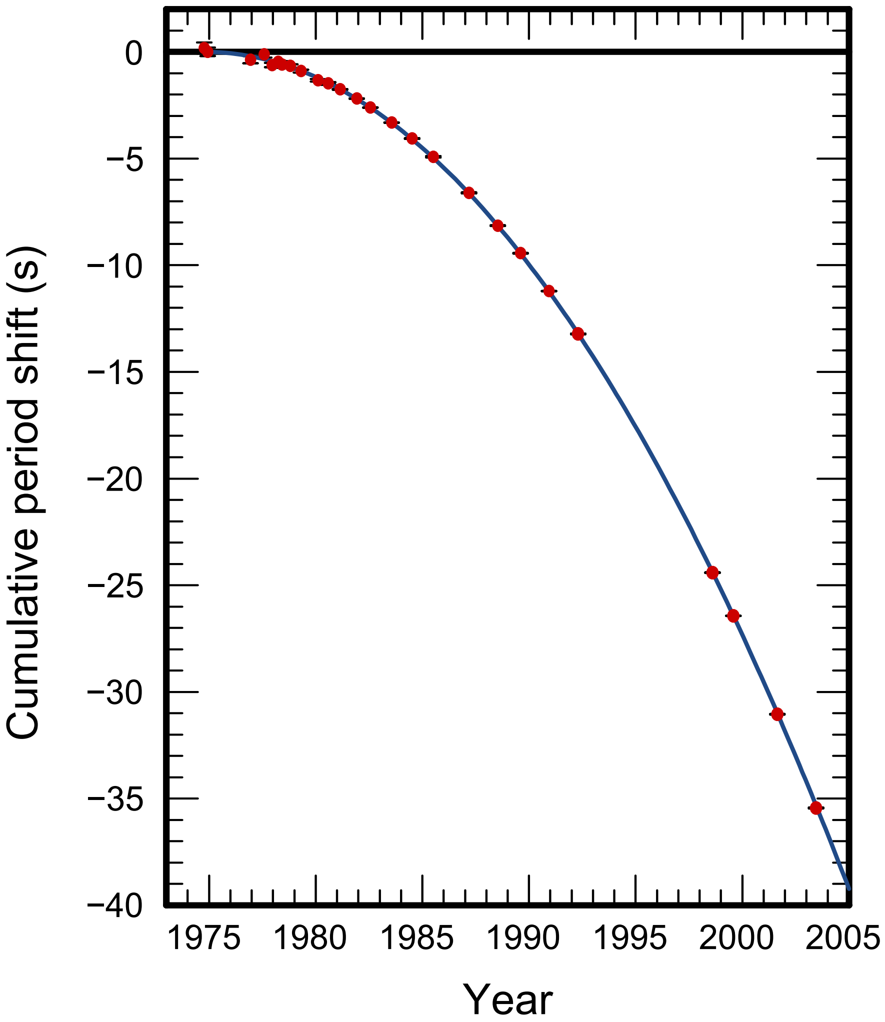

Figure 5.3: The orbital decay of the binary pulsar system $PSRB1913+16$ over a 30 year period. The solid line shows what would be expected if the system were losing energy by emitting gravitational waves, while the data points come from observations gathered at the Arecibo $305m$ antenna.

$$h \simeq 1 \times 10^{-23}\left(\frac{100Mpc}{r}\right)\left(\frac{M_b}{1.2M_\odot}\right)^{\frac{5}{3}}\left(\frac{f}{200Hz}\right)^{\frac{2}{3}} \tag{5.9}$$

where $M_b = (M_1 M_2)^{\frac{3}{5}}/(M_1+M_2)^{\frac{1}{5}}$ is a mass parameter describing the binary system.

Periodic sources emit gravitational radiation at a specific frequency (for a time length longer than an observing run). A binary system far from coalescence (for example the famous Hulse-Taylor binary [55], see figure 5.3) will radiate gravitational waves, in the plane of rotation, at twice the rotation frequency. Indeed, any rotating system that is not spherically symmetric, will radiate gravitational waves in this way.

The stochastic background of gravitational waves, which is expected to be isotropic, stationary and unpolarised, is to gravitational wave astronomy what the cosmic microwave background is to radio astronomy, i.e. of great cosmological significance [56] and hard to detect.

5.3 Gravitational Wave Detection

As has already been discussed in section 5.1, detection of gravitational waves is (or, rather, will be [57]) a matter of detecting a length change ΔL in a length L. Several methods have been proposed to accomplish this challenging task. Cryogenic resonant bar detectors [58] display excellent sensitivity to gravitational waves ($\sim 10^{-21} m/\sqrt{Hz}$[59]), but only at their particular resonant frequency. Precise Doppler tracking of distant spacecraft [60, 61] could provide a means of detecting low frequency gravitational waves, where the spacecraft is one test mass and the Earth (a distance $L \sim 1.5 \times 10^{12} m$ away) is the other. A third strand of research into gravitational wave detection – and in recent years, perhaps the main strand – is that of ground-based interferometric detectors. These provide a high sensitivity to a broad band of gravitational-wave frequencies [62].

Michelson interferometers [2, 63] provide a method of detecting a length change particularly suited to the detection of gravitational waves. The quadropole nature of a + polarised gravitational wave propagating in the $-z$ direction will cause a maximum change in the distance, L, between the origin and a test mass on the the y-axis of $\pm\Delta L$. At the same time, a test mass on the $x-$axis will change by $\mp\Delta L$ – and *vice versa*. If we consider a Michelson interferometer, positioned with its beam splitter at the origin and one of its arms (often referred to as the North arm) aligned with the $y-$axis, and its other arm aligned with the $x-$axis (the East arm), then the end mirrors of the arms will behave exactly as the test masses, resulting in a $4\Delta L$ path difference. (Unfortunately the Michelson interferometer is completely insensitive to waves which are × polarized with respect to the interferometer axes).

A Michelson interferometer (see figure 5.4(a)) is a device which allows very precise measurement of the change in path length of one arm with respect to the other. In a small-scale lab device, one observes the resulting interference fringes at point d. In an full-scale interferometric gravitational-wave detectors, the mirrors and beam-splitter are mounted on multi-stage pendulums to isolate them from seismic noise, and are enclose in vacuum systems. Electronic feedback allows the optical elements to be held relative to each other such that the interferometer is held on a dark fringe, i.e. no light is output at d (in the absence of a gravitational-wave disturbance) and all the light which enters the system is returned to the source, s. For much more detail than this

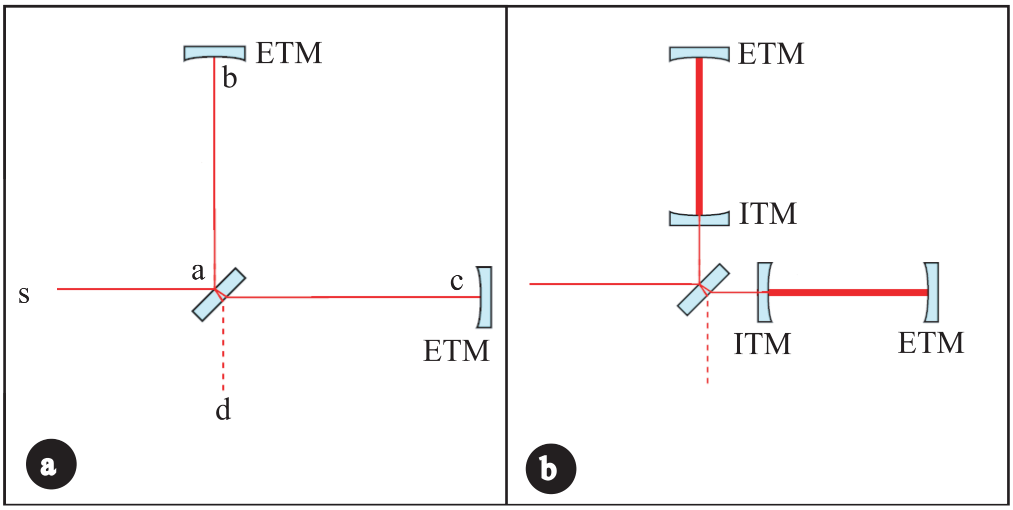

Figure 5.4: (a) A simple Michelson interferometer. (b) A Fabry-Perot Michelson interferometer

brief introduction can possibly convey about interferometer configurations and electronic locking procedures see [64].

The first generation of interferometric gravitational wave detectors are already in existence and have achieved peak sensitivities of $\sim 10^{-21}/\sqrt{\text{Hz}}$: no candidate gravitation wave signals have yet been detected [57]. Preparations for the construction of a second generation of detectors (for the most part, upgrades to existing instruments [65]) are well advanced, and a proposed third generation of instruments are being extensively investigated, both theoretically [66] and experimentally [67].

A simple Michelson configuration, even with arm lengths > 1km, is not sensitive enough to detect gravitational waves. Most existing interferometric detectors use a Michelson layout enhanced by the inclusion of either or both of signal recycling and power recycling [68], and Fabry-Perot cavities in the arms

[3]. For future interferometers, the high laser power (∼ 100 W) coupled with the extremely high finesse of the Fabry-Perot arm cavities (see figure 5.4(b)) means that the coating thermal noise of the Fabry-Perot mirrors will be the limiting factor for the sensitivity of the devices [32] in the region of a few hundred Hz. Our experiment investigates the possibility of replacing the conventional Fabry-Perot arm-cavities with instead a pair of diffractively-coupled cavities. To this end, we commissioned a single 10 m long diffractively-coupled cavity, in order to characterise its properties, and assess if fitness for use in a full-scale interferometric gravitational wave detector. The next two chapters will describe this process and its results, while the next section briefly considers the fundamental limits of sensitivity of interferometric gravitational wave detection.

5.4 Limits to the sensitivity of an interferometric gravitational wave detectors

Many different sources of noise combine to create the complicated noise floor of an interferometric gravitational wave detector. Some are environmental, some are fundamental to the physics of the detection process, and some are determined by the design of the detector. A study of some existing detectors can be instructive here, and we consider an early proposal for a LIGO detector [69].

One obvious sensitivity limit set by the individual detector design is the arm length, L. Since, for a simple Michelson,

$$\Delta L \simeq hL, \tag{5.10}$$

then, in general, a longer arm length means greater signal for the same strength of wave and thus a greater sensitivity. This also holds for a Fabry-Perot Michelson interferometer.

Seismic noise represents an unavoidable (for ground based interferometers) source of noise, and is often the limiting factor at low frequency. Isolation, using multistage suspensions, for the interferometer optics has become the standard technique for reducing this. Several different suspension designs have been tested, for instance, GEO600's monolithic silica triple suspension system or the VIRGO superattenuator [70]. Active cancelling of seismic noise (*via* hydraulic actuators) has also been carried out [71]. Figure 5.5 shows the seismic noise after reduction by suspension isolation; it also shows the noise due to thermal excitation of the suspension (dotted line – the spike just below 1 Hz is a pendulum mode of the suspension).

The high frequency limit is due to photon shot noise. This can be reduced by increasing the laser power and/or increasing the stored light power in the arm cavities. However, increasing the number of photons also increases the radiation pressure noise, which stems from the motion of the suspended mirrors caused by the back-reaction of reflected photons. Since these two factors so closely influence each other, the sum of the two is often referred to simply as quantum noise.

Increasing the laser power can also have implications for noise due to thermal effects in the reflective coatings of the suspended mirrors (by two methods: "coating brownian" noise, i.e. direct thermal excitation of the coating; and "coating thermo-optic" noise, which is due to the changes in refractive index with temperature) and the substrate, particularly highly transmissive optics such as the beam splitter and inner test mass of the arm cavities. Cryogenic cooling of the test masses has been suggested as a possible method to ameliorate this; another possibility is to circumvent the need for transmissive optics altogether, and move to all-reflective interferometry [72], or, as with our experiment, to use diffractively coupled-cavities.

Finally gravity gradient noise, represents the fundamental low frequency limit for ground based detectors [73]. Gravity gradient noise arises from the motion of nearby (i.e. earthbound) masses. The most ubiquitous of these are density fluctuations in the Earth's lithosphere and asthenosphere due to ambient seismic waves.

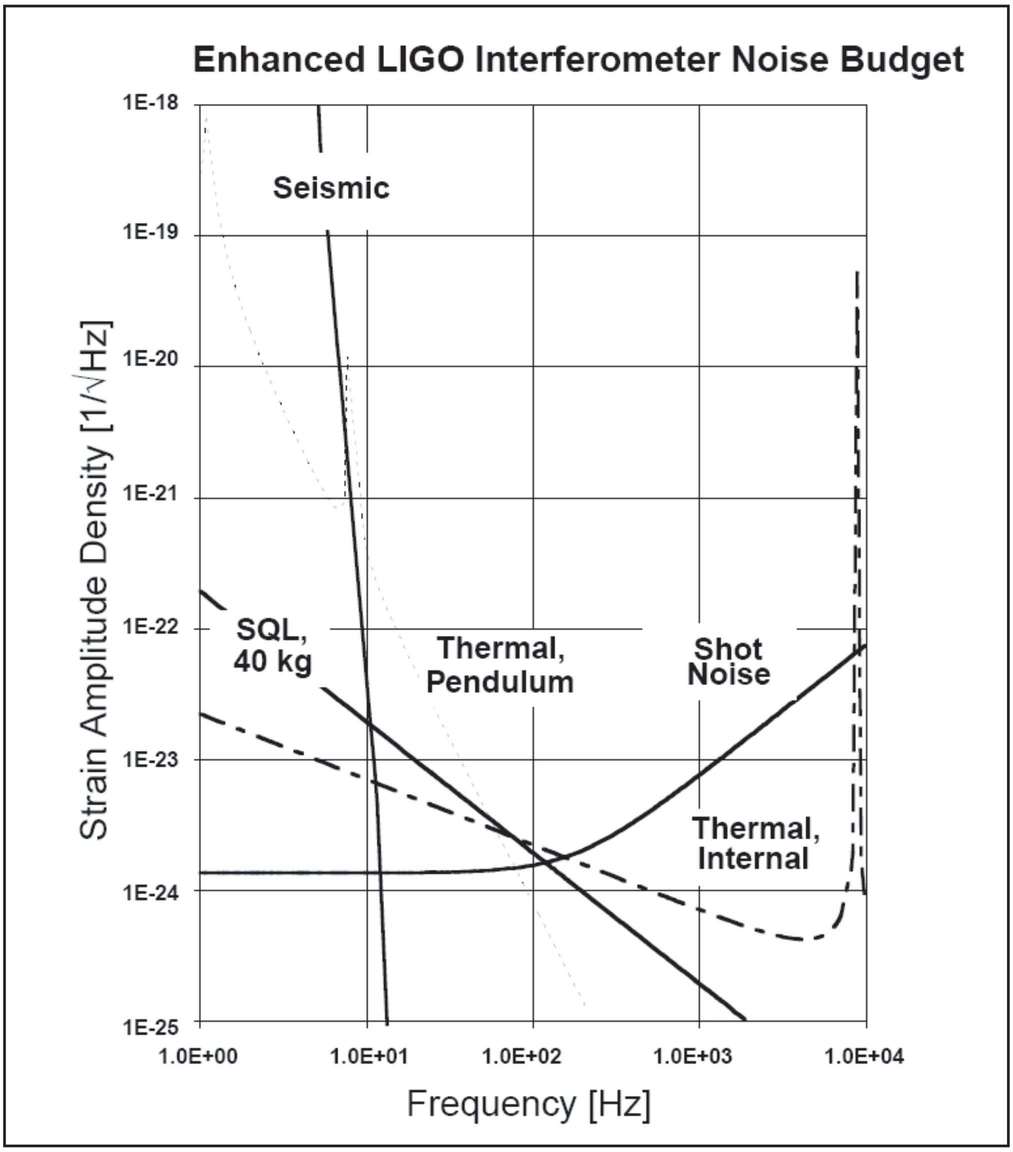

Figure 5.5: The dominant contributions to the strain spectral density of an Enhanced LIGO detector, from [69].

CHAPTER 6

Grating-coupled cavities for Gravitational Wave detectors

This section discuses grating coupled cavities and their possible advantages over conventional Fabry-Perot cavities (at least for use in Gravitational Wave detection), and also introduces the theoretical framework necessary to analyse their behaviour. It also details the specific problems of controlling and locking on-resonance cavities, the mirrors of which are suspended on pendulum mounts.

It should be noted that in this chapter, the term *cavity* means exclusively, a stable ($g_1 g_2 \leq 1$) optical resonator.

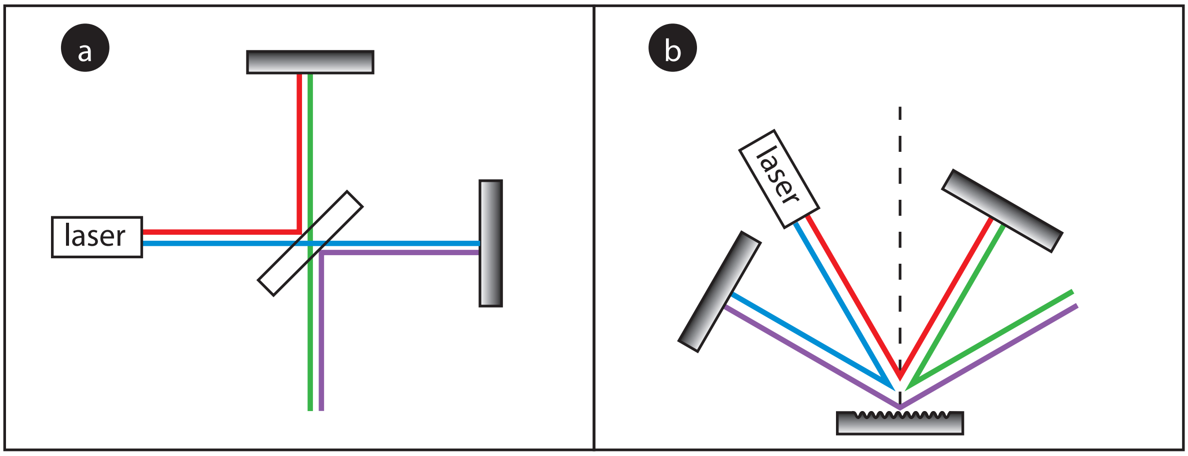

Figure 6.1: A comparison of two Michelson interferometers: (a) is the conventional geometry with a central beam-splitter, while (b) uses a four-port diffraction grating, but is topologically identical.

6.1 Why use grating coupled cavities?

With current designs, quantum shot noise is a strong limiting factor at high frequency [66]. Planned future detectors will have increased laser power, and one possible method of mitigating the increased coating noise [74, 75] is to move from transmissive optics to diffractive optics – for instance, to replace beam-splitters and inboard arm-cavity mirrors with diffractive couplers. The beam-splitter can be replaced by a four-port grating (see figure 6.1) and the inboard mirrors can be replaced by three-port gratings (see section 6.2).

6.2 Scattering Matrix formalism for grating-coupled cavities

An optical element [76] can be characterized by the number of *ports* which it has; that is, by the number of input/output directions. Thus a simple mirror (see figure 6.2(a)) is a two-port element which has a reflectivity, ρ, and a transmissivity, τ. We can relate the outputs of this mirror ($\tilde{b}_1$ and $\tilde{b}_2$) to the inputs ($\tilde{a}_1$ and $\tilde{a}_2$) using the scattering matrix formalism:

$$\begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = \begin{bmatrix} \rho & \tau \\ \tau & -\rho \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}. \tag{6.1}$$

The elements of the scattering matrix are the aforementioned reflectivity and transmissivity.

Without specifying a physical mechanism, we can imagine a general three-port element (see figure 6.2(b)). As with the simple mirror we can represent its input/output relations as follows:

$$\begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} = \begin{bmatrix} \rho_{11} & \tau_{12} & \tau_{13} \\ \tau_{21} & \rho_{22} & \tau_{23} \\ \tau_{31} & \tau_{32} & \rho_{33} \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix}. \tag{6.2}$$

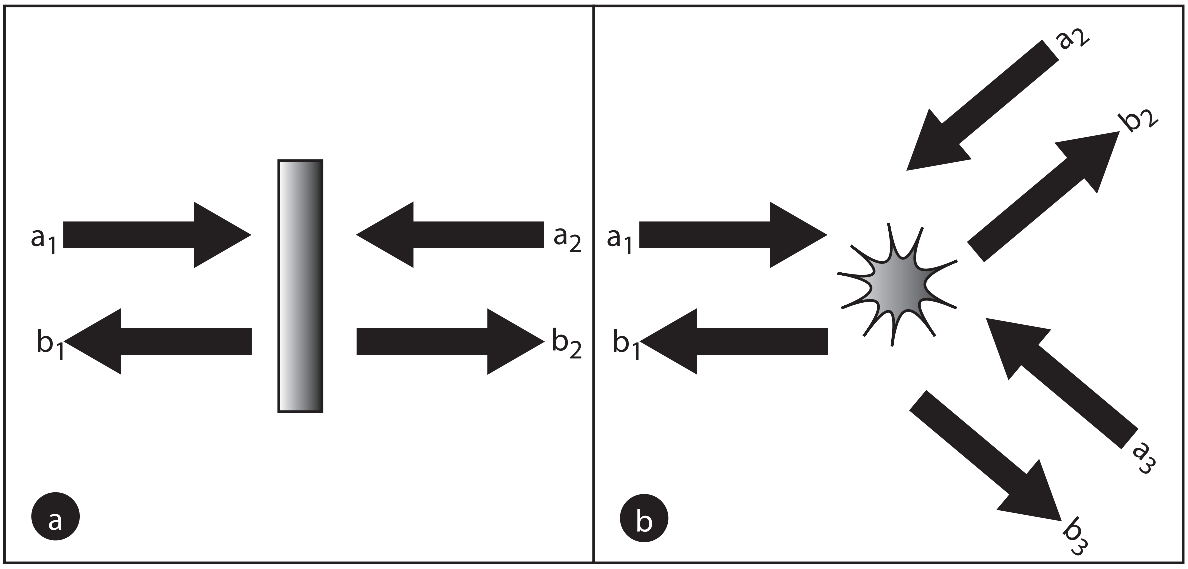

Figure 6.2: Reflection and transmission of optical waves from (a) a simple mirror, and (b) a general three-port optical element.

Again, the vectors **a** and **b** represent the inputs and outputs of the element. The elements on the diagonal of the matrix are labeled ρ_{nn} because they represent the coupling of the input of a port with the output of the same port, i.e. a reflectivity. The off-diagonal elements are the couplings between the inputs and outputs of different ports and are thus labeled τ_{nm}.

Now let us consider a diffraction grating [77] where the grating period is d and all the transmitted orders are suppressed. From the grating equation,

$$\sin\alpha + \sin\beta_m = \frac{m\lambda}{d}, \tag{6.3}$$

we require simply that there be only two diffracted orders with $\beta_m < 90°$. With the further conditions that the light must be incident along a diffraction order – this requirement leads to the grating being in a so-called *Littrow* configuration [78] – we can represent a diffraction grating by just such an equation as 6.2. If we assume that the grating is symmetrical with respect to the grating normal [79], we can write the scattering matrix as

$$\mathbf{S}_{3p} = \begin{bmatrix} \eta_2 e^{i\theta_2} & \eta_1 e^{i\theta_1} & \eta_0 e^{i\theta_0} \\ \eta_1 e^{i\theta_1} & \rho_0 e^{i\theta_0} & \eta_1 e^{i\theta_1} \\ \eta_0 e^{i\theta_0} & \eta_1 e^{i\theta_1} & \eta_2 e^{i\theta_2} \end{bmatrix}, \tag{6.4}$$

where the $e^{i\theta_n}$ terms denote the phase shift associated with the n_{th} diffraction order.

In our experiment light is incident on the grating at the second-order Littrow angle(see figure 6.3), giving rise to reflected light (η_0), back-diffracted light (η_2), and the first-order diffracted light, η_1, which is coupled into the cavity (see figure 6.3(b)). Once inside the cavity, after being reflected from the End Test Mass (ETM) the light again encounters the grating, this time at normal incidence (see figure 6.3(c)). Most of the light is reflected back into the cavity, (ρ_0), but some is diffracted out of the cavity, η_1.

Assuming no loss at the grating surface, we have:

$$\eta_0^2 + \eta_1^2 + \eta_2^2 = 1, \tag{6.5}$$

and

$$\rho_0^2 + 2\eta_1^2 = 1, \tag{6.6}$$

where equation 6.5 refers to the situation in figure 6.3(b) and equation 6.6 refers to the situation in figure 6.3(c). So, the light coupled into the cavity depends entirely on the first order diffraction efficiency, η_0, while the light coupled out of the cavity depends only (if we consider for the moment the end mirror being perfectly reflective) on the same diffraction efficiency, but this time multiplied by a factor of 2 as there are two output diffracted orders. Clearly, if we want our cavity to have a high finesse, we should design our grating to maximise ρ_0 and minimise η_0 and η_1. Put simply, it should be more like a plane mirror than a diffraction grating.

Now we consider the entirety of the cavity. We start with the simplest case, two partially reflective mirrors aligned on the same axis and separated by a length, L. The mirror on the left, mirror 1, has a reflectivity ρ_1 and a transmissivity τ_1. We know from chapter 2 that, for the cavity to be resonant, we must have $L = \frac{m\lambda}{2}$. Assuming this to be the case, we start at $t = 0$ with light incident only from the left (see figure 6.4(a),) with field $\mathbf{E}_i$ (the red arrow),

$$\mathbf{E}_i = \mathbf{E}_0 e^{i\omega t}. \tag{6.7}$$

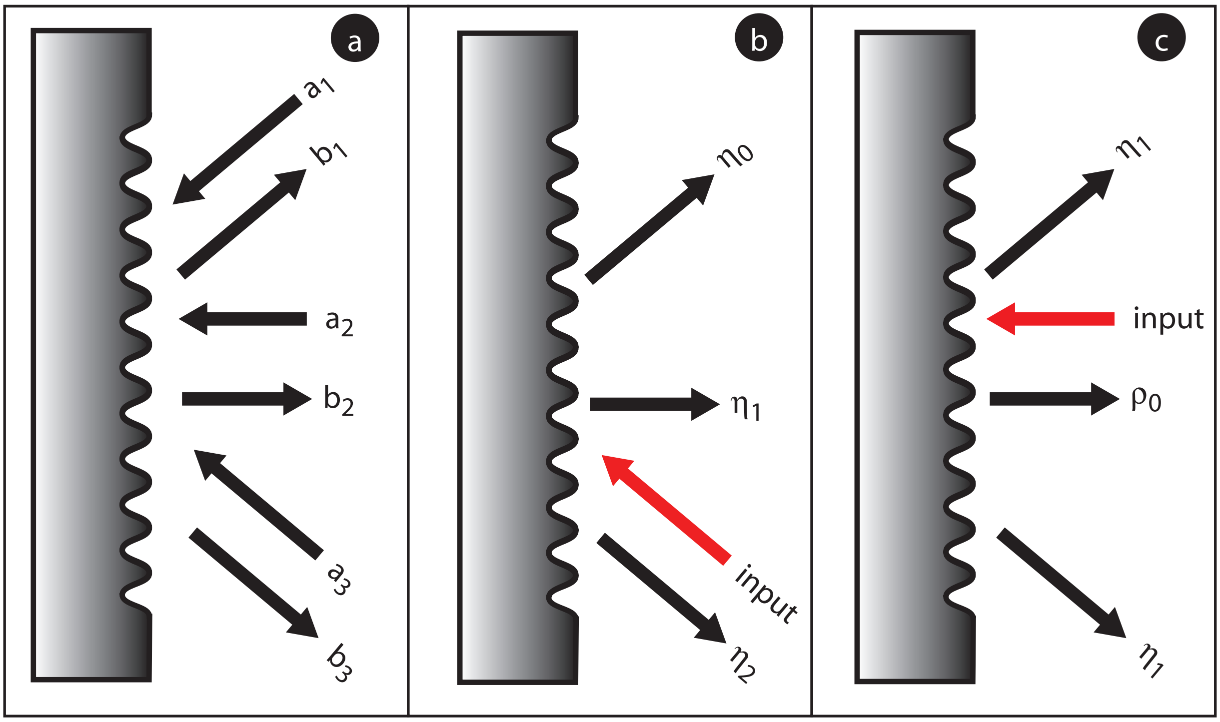

Figure 6.3: A three-port reflection grating. In (a) all the possible input and output ports are shown and labeled. In (b), the grating is arranged for cavity input with the port labeled η_1 coupling into the cavity. (c) shows the grating arranged for cavity output with both ports labeled η_1 coupling out of the cavity.

Since the incident field is constant, a mode quickly builds up in the cavity. This mode is represented by the two circulating fields, $\mathbf{E}_{c1}$ and $\mathbf{E}_{c2}$, which can be many times stronger that the incident field, and are given by

$$\mathbf{E}_{c1} = i\tau_1 \mathbf{E}_i - \rho_1 \mathbf{E}_{c2}, \tag{6.8}$$

and

$$\mathbf{E}_{c2} = -\rho_2 \mathbf{E}_{c1} e^{-i2\delta}, \tag{6.9}$$

where $\delta = 2\pi L/\lambda$ and ρ_n,τ_n are the amplitude reflection and transmission coefficients of the mirrors respectively. The term, δ is the phase delay due to the cavity length.

With these equations, we can now work out what the two outputs of the cavity are. $\mathbf{E}_t$ is simply the portion of the field $\mathbf{E}_{c1}$ which is transmitted through mirror 2, i.e.

$$\begin{aligned} \mathbf{E}_t &= \tau_2 \mathbf{E}_{c1} \\ &= \frac{\tau_1 \tau_2 e^{-i\delta}}{1 - \rho_1 \rho_2 e^{-i2\delta}} \mathbf{E}_i. \end{aligned} \tag{6.10}$$

Similarly, $\mathbf{E}_r$ is simply the portion of $\mathbf{E}_{c2}$ which is transmitted through mirror 1, plus the portion of $\mathbf{E}_i$ which is reflected from mirror 1,

$$\begin{aligned} \mathbf{E}_r &= \tau_1 \mathbf{E}_{c2} + \rho_1 \mathbf{E}_i \\ &= \rho_1 \mathbf{E}_i - \frac{\rho_2 \tau_1^2 \mathbf{E}_i (e^{-i2\delta} - \rho_1 \rho_2)}{1 + \rho_1^2 \rho_2^2 - 2\rho_1 \rho_2 \cos(2\delta)}. \end{aligned} \tag{6.11}$$

Now that we know all outputs of the cavity in terms of the incident field let us consider a simple example: a cavity with two identical mirrors ($\rho_1 = \rho_2 = \rho$

and $\tau_1 = \tau_2 = \tau$) separated by a distance $L = 1$m. Further, we specify that the mirrors are lossless ($\rho + \tau = 1$).

Firstly, from the resonance condition the factor $\delta = 2\pi L/\lambda$ becomes $\delta = 2\pi m$ (where m is an integer) since, on resonance, L/λ must be an integer. So, in equations 6.10 and 6.11 the the exponential and cosine terms are equal to 1, giving

$$\mathbf{E}_t = \frac{\tau^2}{1-\rho^2}\mathbf{E}_i, \tag{6.12}$$

and

$$\mathbf{E}_r = \rho\mathbf{E}_i - \frac{\rho\tau^2(1-\rho^2)}{(1-\rho^2)^2}\mathbf{E}_i. \tag{6.13}$$

Thus, after a period of time during which the power in the cavity builds up, the transmitted and reflected fields are simply equal to the incident field.

For this to be the case, what must be the circulating fields within the cavity? Looking at mirror 2 in isolation, we see that the field $\mathbf{E}_{c1}$ is incident upon it and that the field $\mathbf{E}_t$ is transmitted. And since we know that $\mathbf{E}_t = \mathbf{E}_i$ it is clear that,

$$\mathbf{E}_t = \tau \mathbf{E}_{c1}. \tag{6.14}$$

In contrast to this, the grating coupled cavity (figure 6.4(b)) has three output fields: $\mathbf{E}_{\eta 0}$, $\mathbf{E}_{\eta 2}$ and $\mathbf{E}_t$. The incident field is aligned with the η_2 diffraction port, i.e. the grating is in a second order Littrow mount. We can derive the input-output relations for this more complex cavity by using equation 6.4 - full details of the derivation can be found in [79]. Here we simply state that we find,

$$\mathbf{E}_{\eta 0} = \mathbf{E}_i \eta_0 + \mathbf{E}_i \eta_1^2 e^{[2i(\phi_1 + \phi)]}, \tag{6.15}$$

$$\mathbf{E}_{c1} = \mathbf{E}_i \eta_1 e^{(i\phi_1)d}, \tag{6.16}$$

$$\mathbf{E}_{\eta 2} = \mathbf{E}_i \eta_2 e^{i\phi_2} + \mathbf{E}_i \eta_1^2 e^{[2i(\phi_1 + \phi)]}, \tag{6.17}$$

$$\mathbf{E}_t = \mathbf{E}_{c1}(1 - \rho_1) i e^{i\phi}, \tag{6.18}$$

where d is the resonance factor ($d = 1/[1 - \rho_0 \rho_1 \exp(2i\phi)]$), and ϕ_n is the phase shift associated with the n^{th} diffraction order.

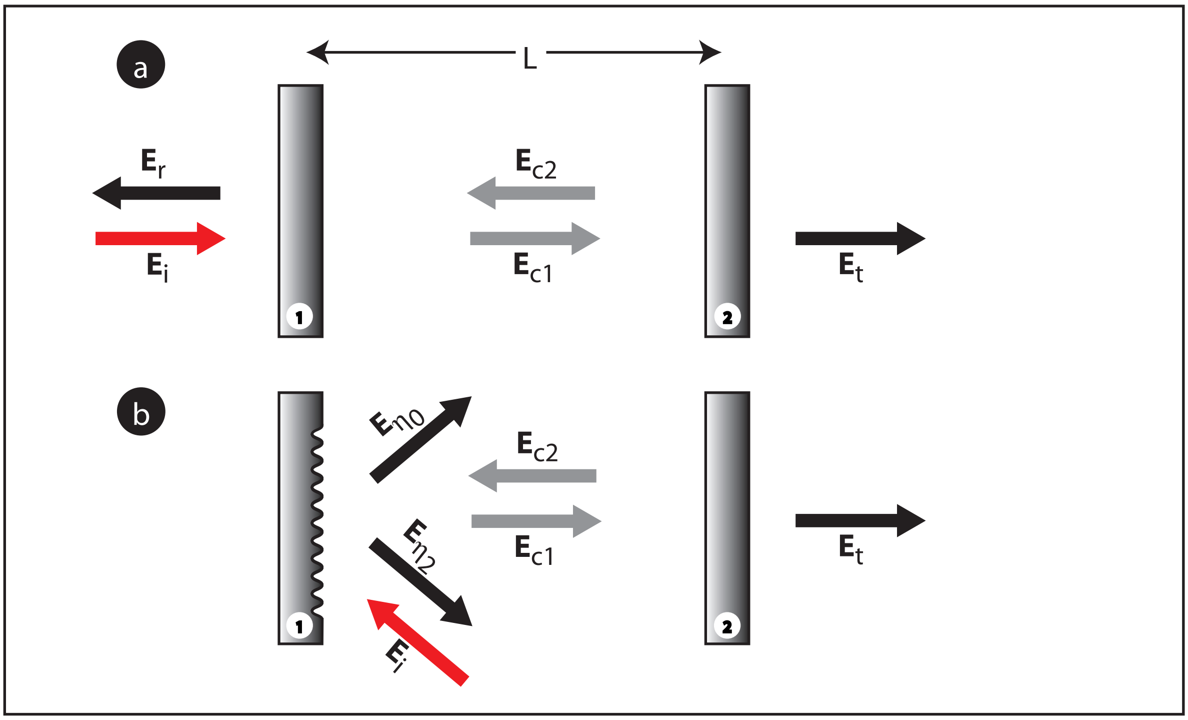

Figure 6.4: A two-port, conventional Fabry-Perot cavity (a) and a three-port, grating coupled cavity (b), both of length L. Red arrows represent incident field, grey arrows are fields circulating within the cavity and black arrows are fields which leave the cavity.

6.3 Length sensing and control for suspended cavities

So far we have talked about the expected properties of our diffractively-coupled cavity, but not about one of the major challenges of realising this experimentally. As we discussed in the previous chapter, for gravitational wave detection, the optical elements of the interferometer, including the arm cavities, must be suspended. This obviously removes our ability the rigidly define the length L of the arm cavities as we would be able to if our optical elements were, say, bolted to an optical bench.

One method of doing this for conventional Fabry-Perot cavities is the *Pound-Drever-Hall* [80] method. For this method, the light input into the cavity (which we call the *carrier*) is modulated to produce a set sidebands by either amplitude or phase modulating the carrier. Since our experiment involves only phase-modulation sidebands, we will discuss these.

In our experiments, phase sidebands are imposed on the carrier by passing it through an electro-optic modulator (EOM) [81], which is a birefringent $\mathrm{LiNbO_3}$ crystal sandwiched between two copper plates - effectively a capacitor (approximately 30 pF) with the birefringent crystal as the dielectric. The sinusoidal signal at RF frequencies (usually of the order of 10MHz) is applied to the copper plates and creates a pair of phase-modulated sidebands. We can represent the modulated field as

$$E_{mod} = E_0 e^{i(\omega t + m\cos(\omega_n t))}, \tag{6.19}$$

were E_0 in the unmodulated carrier field, ω is the angular frequency of the laser, t is the time, m is the modulation index and ω_m is the angular frequency of the modulation. We can expand this as,

$$E_{mod} = E_0 e^{i\omega t}[J_0(m) + J_1(m)e^{i\omega_n t} - J_1(m)e^{i\omega_m t} + \cdots], \tag{6.20}$$

where $J_i(m)$ are Bessel functions of the first kind, and the higher order terms have been left out. The modulation index is, then, an indication of how much of the carrier power is diverted into the sidebands.

Only light within the linewidth of cavity is resonant; the sideband fields are effectively reflected from the front of the cavity, when the cavity is resonant with the carrier (as long as a modulation frequency outside of the cavity linewidth is chosen). Therefore the light emerging from the cavity will interfere with the static phase reference of the sideband field resulting in a beat at the sideband modulation frequency. The beat The superposed carrier (after circulating in the cavity) and sideband field is detected and can be demodulated by mixing it with a local oscillator of the same frequency as the modulation (and also of the correct phase), resulting in a signal as shown in figure 6.5.

The error signal for a diffractive-coupled cavity is slightly more complicated since the light leaving the cavity has a different phase shift depending upon which port it exits from. These complications will be discussed in the next chapter.

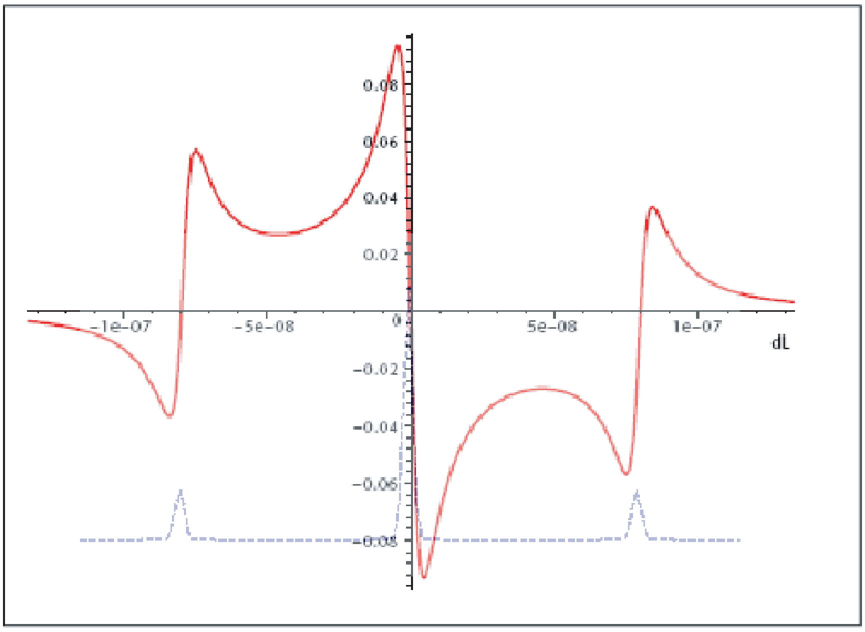

Figure 6.5: A typical (modelled) error signal for a Fabry-Perot cavity, which can be thought of as the derivative of the cavity transfer function (dashed line).

CHAPTER 7

The JIF interferometer: experimental demonstration of a grating-coupled suspended cavity

This chapter describes – more or less, chronologically – the commissioning of a diffractively-coupled cavity in the Glasgow 10m prototype. It then goes on to describe the experiments carried out to characterize this cavity – work which laid the foundation for the publication of [67] (and also [82]). This work was carried out in conjunction with B. Barr, M. Plissi, M. Edgar and S. Huttner.

As the requirements for circulating light power in gravitational wave detectors increase, thermal noise due to the test mass reflective coatings and substrate also increase. Numerous methods have been suggested to mitigate this, including the use of mesa beams [32], or coating-free corner-cube cavities [83]. The method which we choose to investigate here are diffractive reflection gratings. The worst thermal effects will occur at the beam-splitter and at the input couplers of the arm cavities; if these transmissive optics can be replaced by appropriate diffraction gratings then, perhaps, these thermal effects can be reduced to an acceptable level.

7.1 The Glasgow 10m prototype

The Glasgow 10m prototype is housed in the JIF facility – a class 100 clean laboratory, with class 10 areas [84] around the tanks housing the optics (see figure 7.1). The suspended optics are housed in 9 vacuum tanks connected by several tens of metres of beam tube. It is intended to be a test-bed for technologies intended to increase the sensitivity of interferometric gravitational wave detectors.

The optics are suspended on either 2 stage or 3 stage pendulums. In general the 2 stage optics are the beam alignment mirrors and the 3 stage optics are the cavity optics or beam splitters. From figure 7.2 we see that the beam comes from the optical bench (where the control sidebands are imposed upon it); through tank 4 (not shown) which houses the thermal noise experiment [85]; through tank 5 which usually houses the power-recycling mirror and into

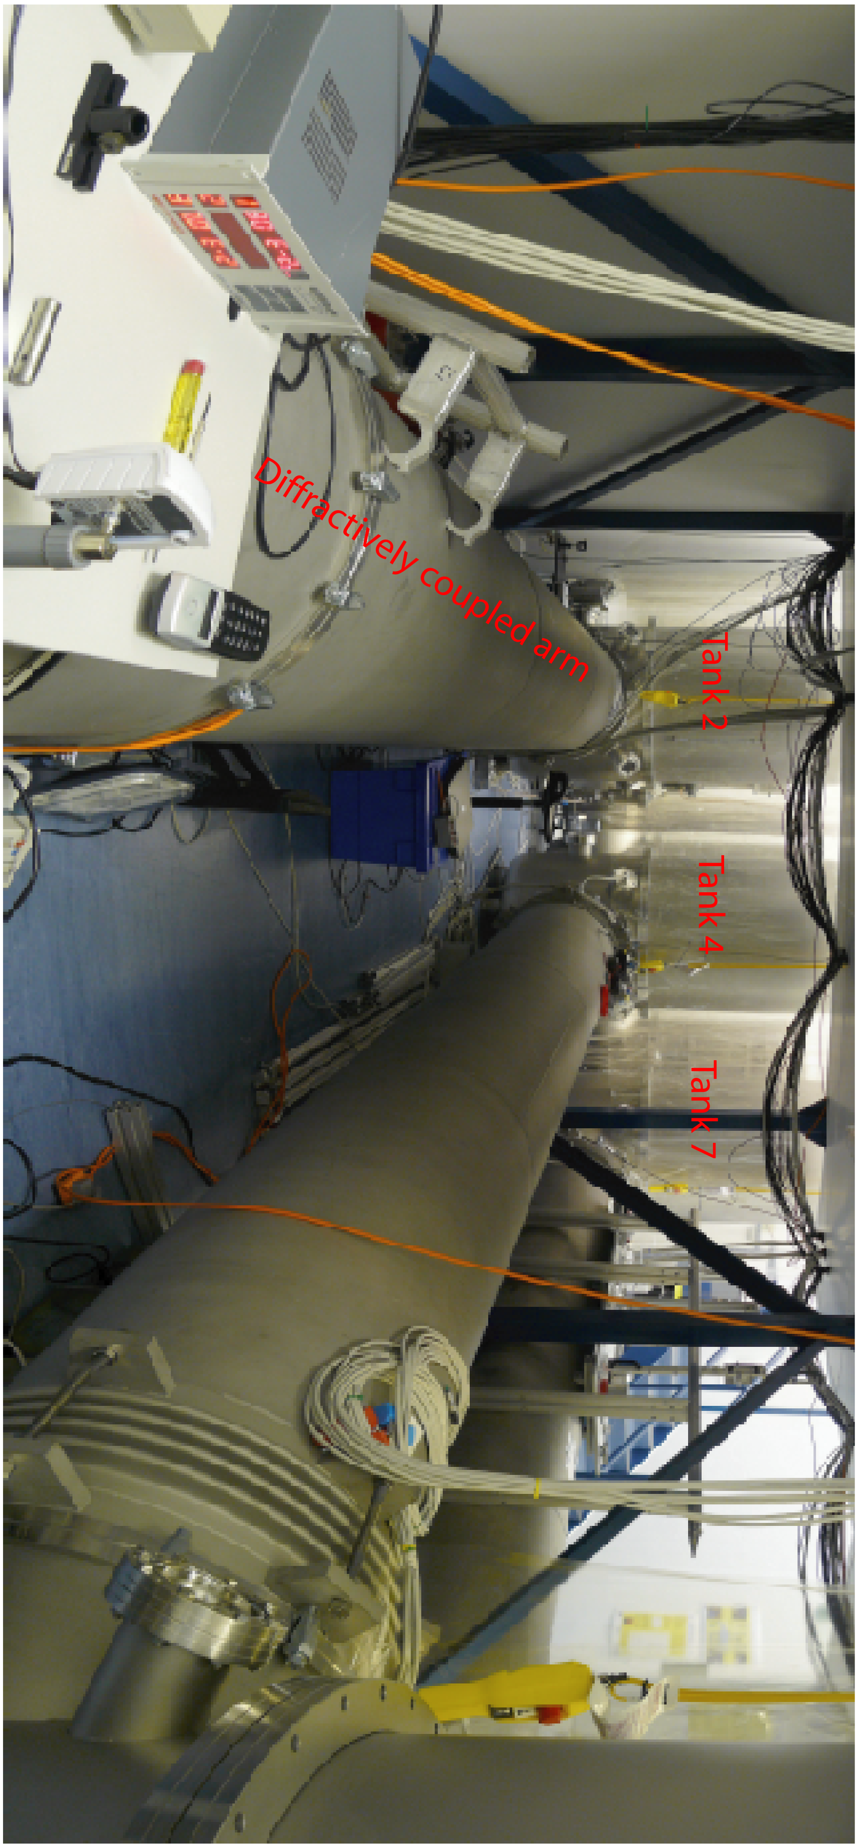

Figure 7.1: A photograph of the the JIF lab facility showing the 10m beam tubes and the ITM tanks.

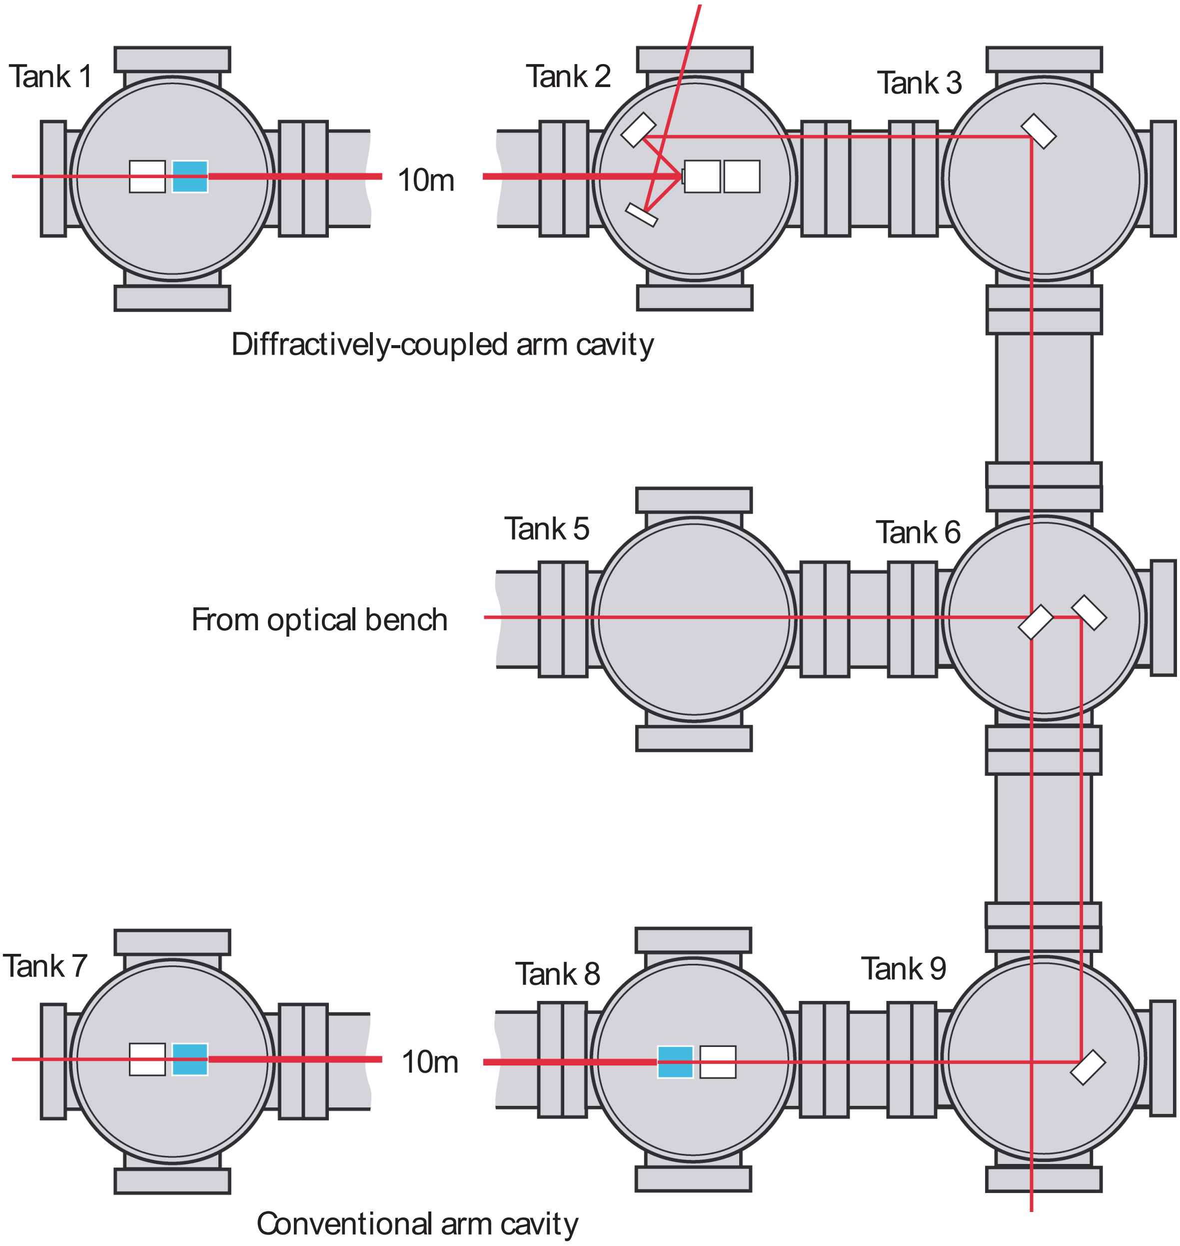

Figure 7.2: A plan view of the JIF lab facility. Tank 4 containing the thermal noise experiment [85] and the optical bench are not shown.

tank 6 were it encounters a $50 - 50$ beam-splitter. Half of the light power is diverted to our experiment in tanks 1,2 and 3, while the rest continues to the conventional Fabry-Perot cavity housed in tanks 7 and 8 (note that silica masses are shown in blue, while aluminium masses with 1” optics are shown in white).

Figure 7.3 shows the layout on the optical bench. Light from a 2 W Nd:YAG laser [86] passes through a Faraday isolator, and then the Electro-Optic Modulators (EOMs) which impose phase sidebands of 10 MHz and 15.24 MHz for control of the cavity (see section 7.4).

A more detailed plan of tanks 1 and 2 is shown in figure 7.4. Light from the beam splitter falls onto the first alignment mirror MA1 and is directed onto the second alignment mirror MAD. This, incidentally, was one of the trickiest procedures of the rough alignment: since the position of MAD is more or less fixed, due to lack of space in the tank, and the clearance between the reaction mass, M1Dr, and the supporting leg (not shown) is quite tight, the correct positioning of MA1 was critical. MAD reflects the light onto the diffractive optic, M1D, from where it is diffracted into the cavity formed between the almost-flat diffractive optic and the 15m focal-length silica mass, M1E.

Rough alignment of the beam bath was carried out by moving the mounting points of the steering-mirror suspensions and rotating the suspensions on their mounts. MA1, as the only suspension in tank 3, could be freely moved; MAD had a very small range of possible movement and rotation, since the interior of tank 2 was rather crowded.

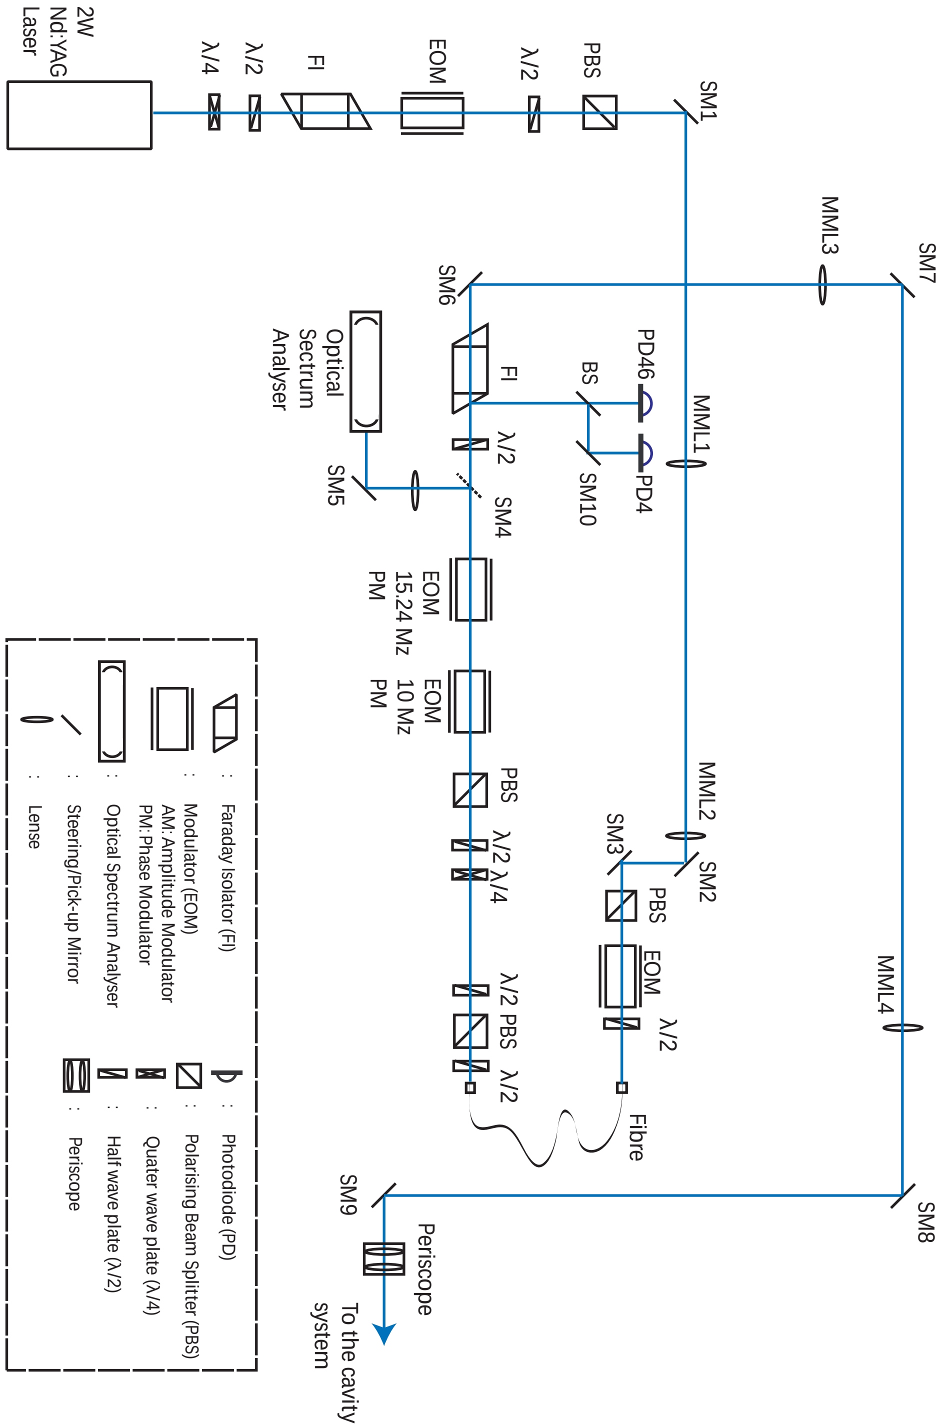

Figure 7.3: The optical bench: light from a 2 W Nd:YAG laser passes through a single-mode fibre and then through 2 EOMs which impose control sidebands upon it. Figure adapted from an original by S. H. Huttner.

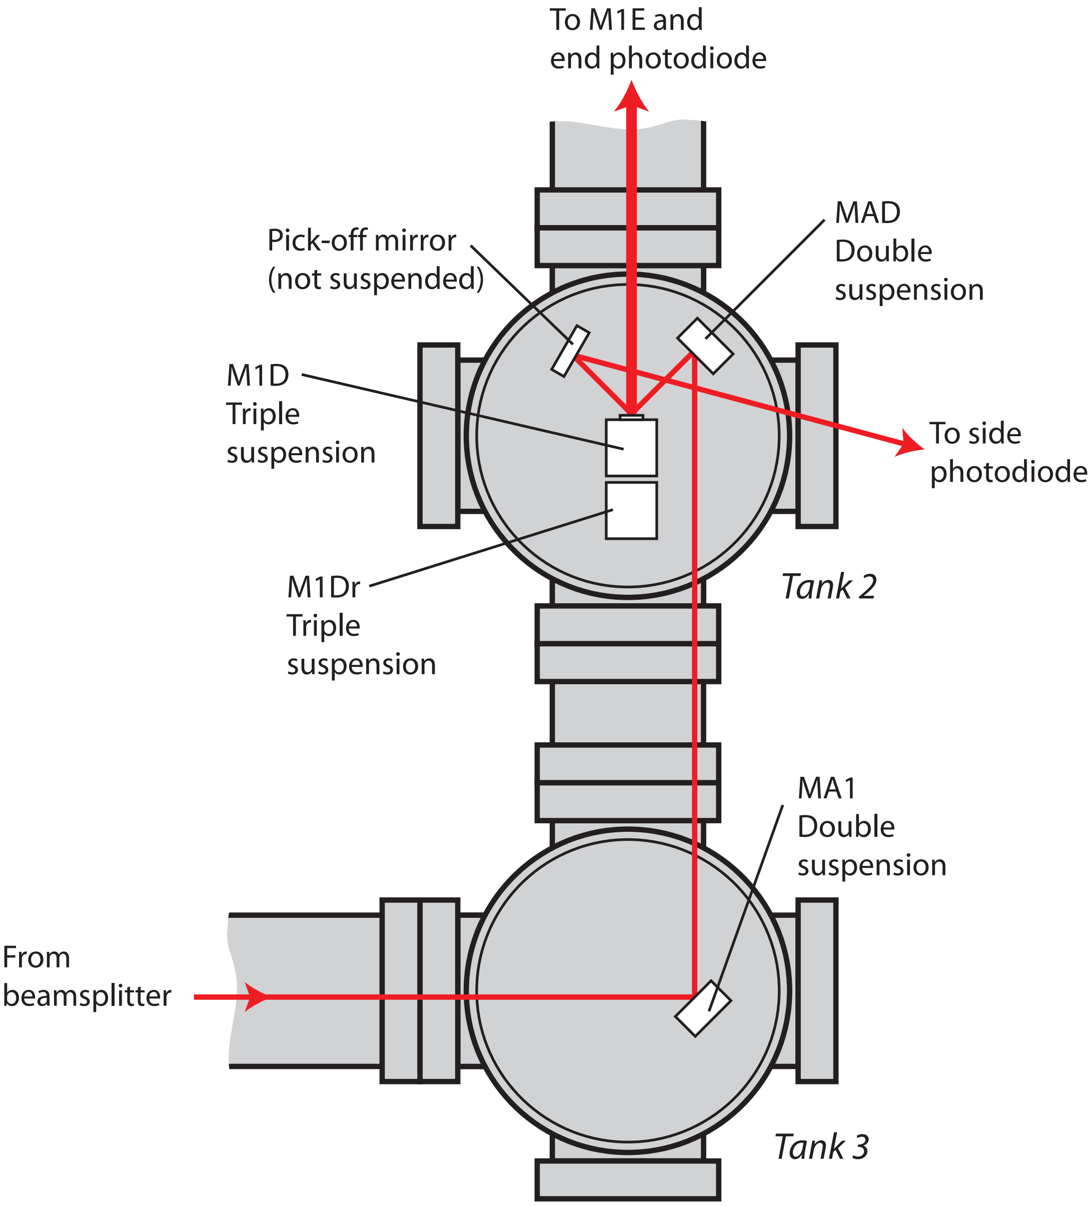

Figure 7.4: The interior of tank 2.

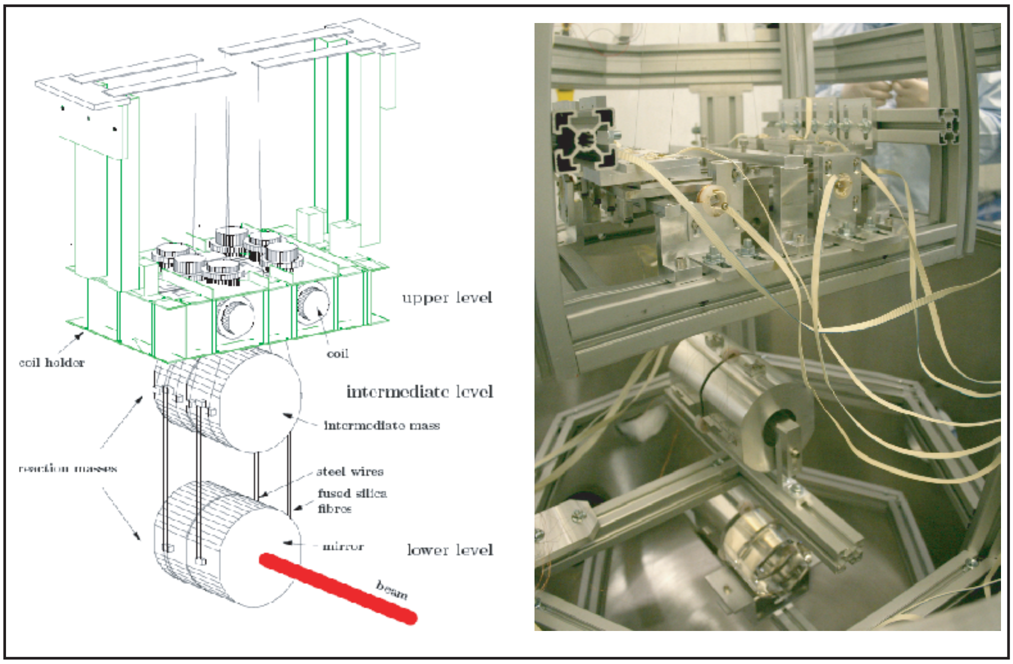

Figure 7.5: On the left is a diagram a triple pendulum – or rather, of two triple pendulums: a main optic pendulum and a reaction mass pendulum. Six shadow sensors on the upper stage of the pendulum give information about the orientation of the upper stage and hence (like a marionette puppet) the orientation of the optic in the lower stage. The photo on the right shows the rather more complicated reality: Bosch support frame, ribbon cables, and the various catchers in case of suspension breakage.

The triple suspensions of M1D and M1E were also roughly aligned using a variety of means. The attachment of the triple suspensions makes it more difficult to alter the rotation of the optic: instead of being directly attached to the Bosch framework (as the double pendulums are) they are suspended from a pair of cantilever blades (see figure 7.5), though this is not a procedure to be recommended . It is possible to gain a few degrees of rotation control by altering the alignment of the cantilever blades. Rough tilt control on both double- and triple-pendulums is controlled by adjustable masses on the upper stage of the suspension.

In order for the diffractive optic to function as a second order Littrow mount, the input light must encounter the diffractive at an angle of 47.2°. Furthermore (see figure 7.6), it is necessary that the light be propagating in the xz-plane, and that the grating itself be aligned parallel to the z-axis. The 3D grating equation tells us,

$$\vec{q} \times \vec{N} - \vec{p} \times \vec{N} = \frac{m\lambda}{d}\vec{G}, \tag{7.1}$$

where $\vec{p}$ is the unit vector of the incoming beam, $\vec{q}$ is the unit vector of the outgoing beam, $\vec{N}$ is the normal vector of the grating, $\vec{G}$ is the unit vector of the grating and d is the grating spacing. Thus proper alignment, in the terms of figure 7.6, has $\vec{N}$ parallel to the z-axis and $\vec{G}$ parallel (or antiparallel) to the y-axis.

Local control of suspended optics

Local control refers to the electronic alignment control – and damping – of the suspended optics. It involves a complex system of sensors, analog electronics and digital control software. Figure 7.7 shows a block diagram of the control and damping system for one pendulum, be it a double or a triple suspension.

The sensor and actuator for the pendulum is an integrated unit. The sensor is a shadow sensor (see figure 7.8) composed of and IR LED and photodiode

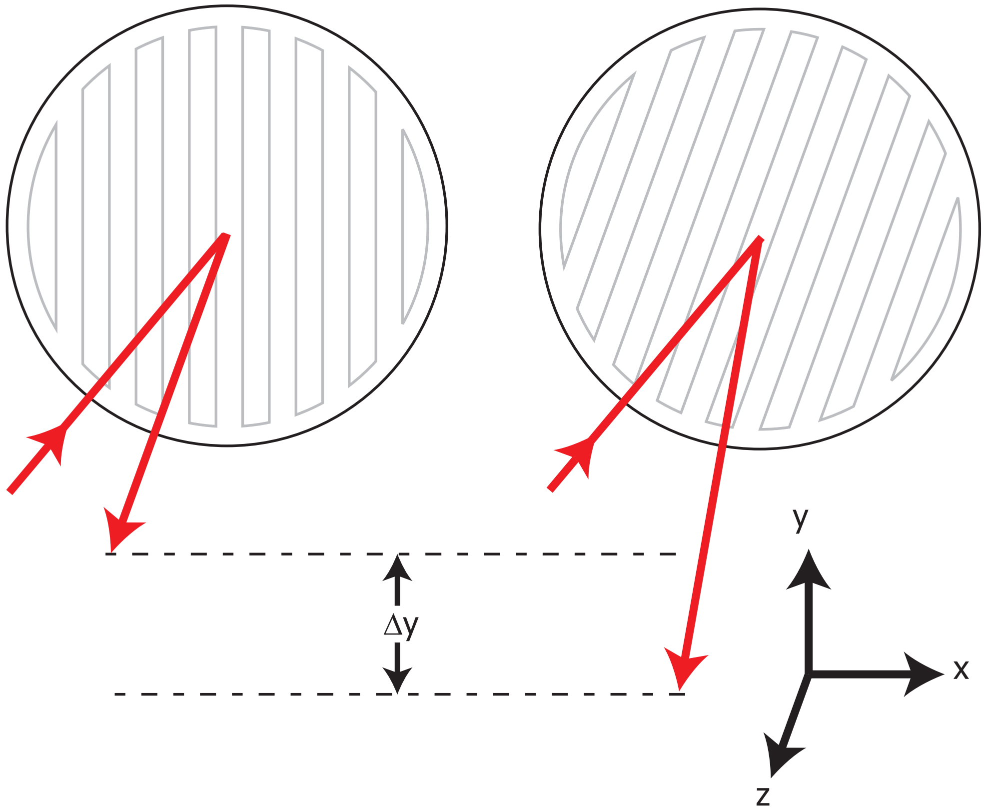

Figure 7.6: The effect of the grating roll degree-of-freedom on cavity alignment. The left figure shows the grating perfectly aligned, i.e. $\overrightarrow{g}$ parallel to the z-axis and the grating face is in the xy-plane: if the input light is in the xz plane at an angle of 47.2° to the z-axis, then the 1^{st} order diffracted beam propagates along the z-axis.

housed in a PTFE structure which also holds a copper coil. The magnet mounted on the top stage of the pendulum and the coil are the actuator. Each double suspension has four of these senor-actuator combinations and each triple suspension has six. These provide fine control of the rotation, tilt and roll of the end stage optic.

The MCP (a computer programme in LabVIEW designed to monitor and adjust the digital control boards for each suspension, see figure 7.7) keeps track

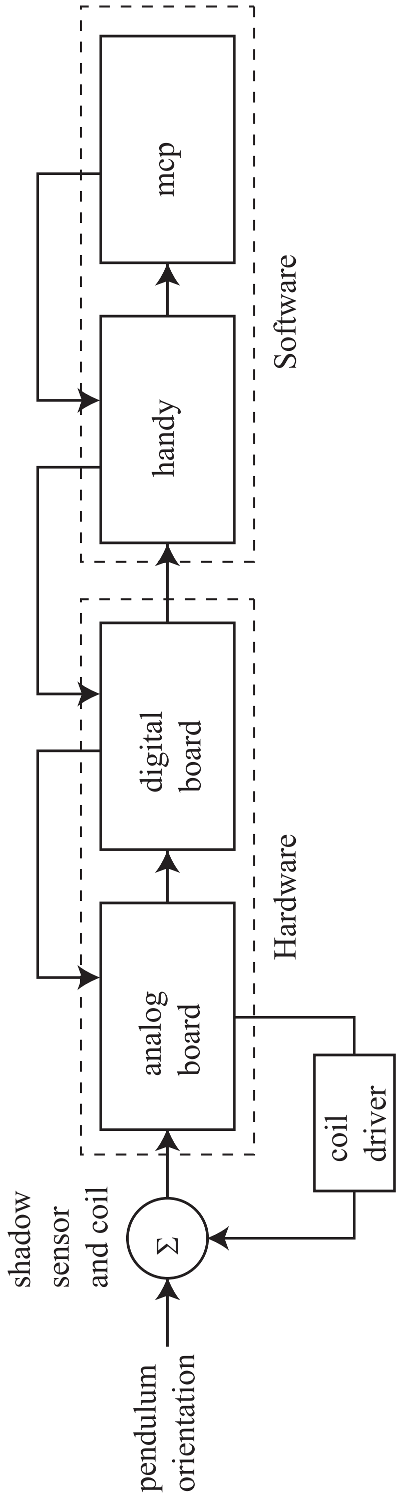

Figure 7.7: A block diagram showing the feedback path for damping movement of the pendulums, as well as the alignment .

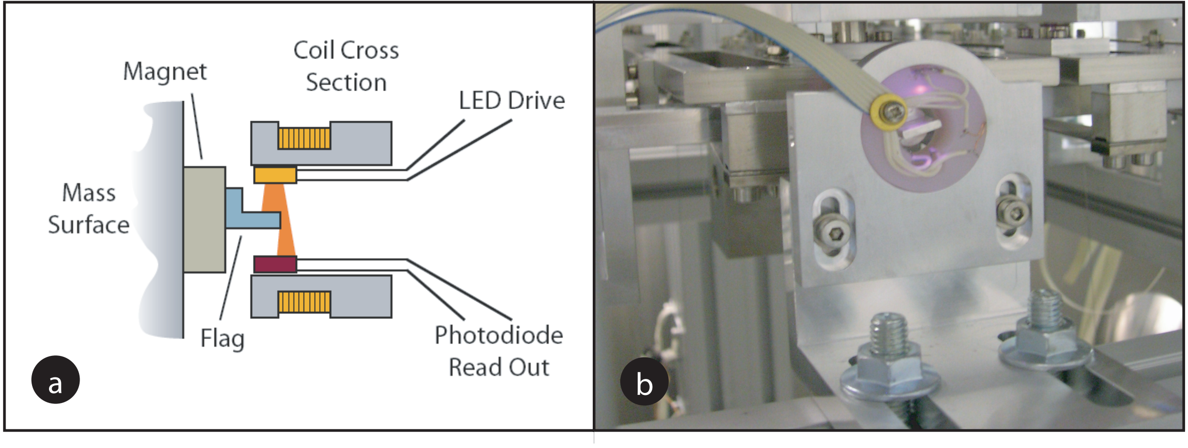

Figure 7.8: A diagram (a) of shadow sensor and magnet-coil actuator used for damping and steering of suspended optics. A photograph (b) of one of the six shadow sensor and magnet-coil actuator used on the top stage of the three-stage pendulums.

of, and archives, the setting for each of the 7 pendulums in the diffractive cavity experiment (and, also, the pendulums for the three-mirror coupled-cavity experiment – but we are not concerned with these here). Figure 7.9 shows an example of the look-up table for one double-stage pendulum. The module number is set here in software and must be matched correctly to the hardware module by correctly setting two hexadecimal digital potentiometers (i.e. 256 possible settings).

One of the author's tasks during the commissioning of the diffractively coupled cavity was to rewrite the local control software to allow for the control of more than the 14 suspensions it had initially been designed for.

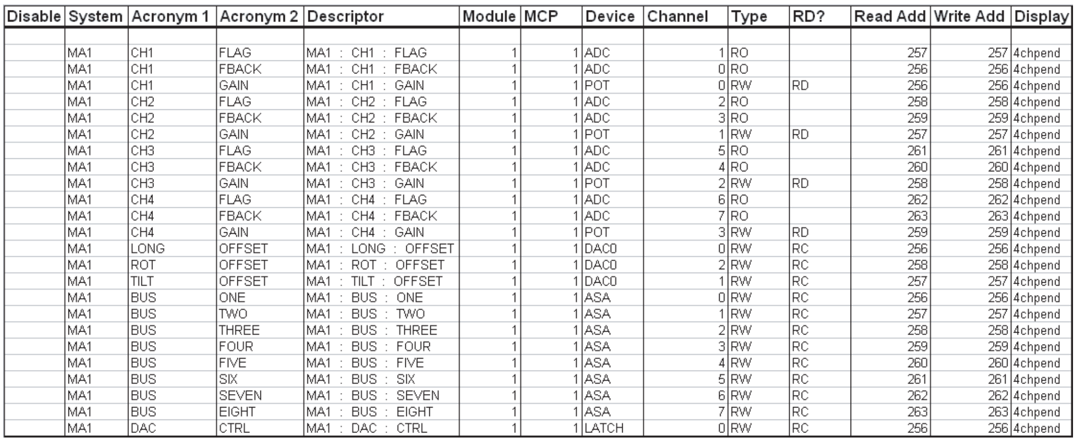

Disable	System	Acronym 1	Acronym 2	Descriptor	Module	MCP	Device	Channel	Type	RD?	Read Add	Write Add	Display
	MA1	CH1	FLAG	MA1 : CH1 : FLAG	1	1	ADC	1	RO		257	257	4chpend
	MA1	CH1	FBACK	MA1 : CH1 : FBACK	1	1	ADC	0	RO		256	256	4chpend
	MA1	CH1	GAIN	MA1 : CH1 : GAIN	1	1	POT	0	RW	RD	256	256	4chpend
	MA1	CH2	FLAG	MA1 : CH2 : FLAG	1	1	ADC	2	RO		258	258	4chpend
	MA1	CH2	FBACK	MA1 : CH2 : FBACK	1	1	ADC	3	RO		259	259	4chpend
	MA1	CH2	GAIN	MA1 : CH2 : GAIN	1	1	POT	1	RW	RD	257	257	4chpend
	MA1	CH3	FLAG	MA1 : CH3 : FLAG	1	1	ADC	5	RO		261	261	4chpend
	MA1	CH3	FBACK	MA1 : CH3 : FBACK	1	1	ADC	4	RO		260	260	4chpend
	MA1	CH3	GAIN	MA1 : CH3 : GAIN	1	1	POT	2	RW	RD	258	258	4chpend
	MA1	CH4	FLAG	MA1 : CH4 : FLAG	1	1	ADC	6	RO		262	262	4chpend
	MA1	CH4	FBACK	MA1 : CH4 : FBACK	1	1	ADC	7	RO		263	263	4chpend
	MA1	CH4	GAIN	MA1 : CH4 : GAIN	1	1	POT	3	RW	RD	259	259	4chpend
	MA1	LONG	OFFSET	MA1 : LONG : OFFSET	1	1	DAC0	0	RW	RC	256	256	4chpend
	MA1	ROT	OFFSET	MA1 : ROT : OFFSET	1	1	DAC0	2	RW	RC	258	258	4chpend
	MA1	TILT	OFFSET	MA1 : TILT : OFFSET	1	1	DAC0	1	RW	RC	257	257	4chpend
	MA1	BUS	ONE	MA1 : BUS : ONE	1	1	ASA	0	RW	RC	256	256	4chpend
	MA1	BUS	TWO	MA1 : BUS : TWO	1	1	ASA	1	RW	RC	257	257	4chpend
	MA1	BUS	THREE	MA1 : BUS : THREE	1	1	ASA	2	RW	RC	258	258	4chpend
	MA1	BUS	FOUR	MA1 : BUS : FOUR	1	1	ASA	3	RW	RC	259	259	4chpend
	MA1	BUS	FIVE	MA1 : BUS : FIVE	1	1	ASA	4	RW	RC	260	260	4chpend
	MA1	BUS	SIX	MA1 : BUS : SIX	1	1	ASA	5	RW	RC	261	261	4chpend
	MA1	BUS	SEVEN	MA1 : BUS : SEVEN	1	1	ASA	6	RW	RC	262	262	4chpend
	MA1	BUS	EIGHT	MA1 : BUS : EIGHT	1	1	ASA	7	RW	RC	263	263	4chpend
	MA1	DAC	CTRL	MA1 : DAC : CTRL	1	1	LATCH	0	RW	RC	256	256	4chpend

Figure 7.9: An example of the look-up table for a 4-channel pendulum. Each of the 7 pendulums in the diffractive-cavity chain have a similar look-up table.

7.2 Systematic alignment technique

After the first, rather haphazard, alignment of the diffractively-coupled cavity, a systematic technique was developed to align the cavity; this was important since the cavity needed to be realigned after almost every change made in tank2 (the following section will show that there were a considerable number of these).

To a very high degree, the diffractive optic can be considered to be simply a plane mirror ($> 99\%$) – this gives a clue as to the best method of alignment. Recall that the cavity-throughput beam exits tank 1 through a window and is detected on the throughput photodiode. With tanks 1 and 2 open to allow visual inspection of the cavity optics, we replace the throughput photodiode with a HeNe laser (see figure 7.10). Laser light (helpfully, visible) is shone into the cavity from the conventional-mirror end. First we confirm that the HeNe is positioned correctly by making sure that the visible spot on both M1E and M1D is central; then, using the coarse and fine alignment techniques outlined previously, we move M1D until the return spot from M1D (visible on the face of M1E) coincides with the input spot. Then, looking at the return spot from M1E on the face of M1D (i.e. 30m downstream from the input spot), we move M1E until the return spot coincides with the input spot. This may need to be repeated with spots further and further downstream from the input spot.

The triple pendulums are damped by the local controls (at the digitally preset gain) but are receiving no global cavity control: therefore the optics of both M1E and M1D are moving very small distances ($< 1\,\mathrm{mm}$ if not disturbed) at the pendulum frequency. When the cavity is well enough aligned the optics will swing through a resonance condition (i.e. $D = m\lambda/2$) and out again – the visual effect of this is that the laser spot on the cavity optics will flash brightly (as the cavity field builds up and dies down). At this point the cavity is aligned (although not locked), but only as a conventional Fabry-Perot cavity.

The next stage of the procedure is to align the grating's roll degree-of-freedom, and the input pointing defined by the two steering mirrors, MA1 and MAD. For this stage we use the beam from the Nd:YAG laser. By changing the rotation and tilt of MA1 and MAD we align the input spot onto the centre of M1D: fine control is usually sufficient to accomplish this. We then look for the 1^{st} order diffracted spot from M1D, on the face of M1E. Since the cavity is aligned, if the diffracted spot is not central it is due either to incorrect input pointing of the beam from MAD (i.e. it is not in the xz-plane, but at an angle to it) or to the grating not being aligned with the z axis. We can correct the input pointing by adjusting MA1 and MAD so that the 0^{th} order (reflected) beam spot on the pick of mirror is at the same height (measured from the lower Bosch frame with a meter stick) as the beam spot on MAD, while maintaining the input spot central on M1D. Then a combination of fine roll and tilt control of M1D should bring the 1^{st} order diffracted spot to the centre of M1E. We then examine the return spot from M1E on the face of M1D. An fluorescent IR viewing card, improved by the subtraction of a small hole (slightly bigger than the beam) is the ideal tool for such a task: the diffracted beam is allowed to pass though the hole, and the reflected is observed on the face of the card. Very fine rotation and tilt of M1E will bring the reflected spot to the centre of M1D.

In the final stage of the alignment procedure, a CCD camera connected to a monitor is positioned at the window of tank 1. As the two pendulums swing at their natural frequency, a quick succession of TEM modes (and superpositions of TEM modes) can be seen on the monitor. If the cavity is poorly aligned, modes with a central dark spot – such as TEM11 and 33 – will predominate; or if the cavity is well aligned in one axis, but not the other higher order modes such as TEM16,26 will appear. With a great deal of patience, tiny changes can be made to the alignment of all four mirrors, so that modes with a central bright spot (TEM00, 22, etc.) predominate. At this point the cavity is ready for global controls to be applied.

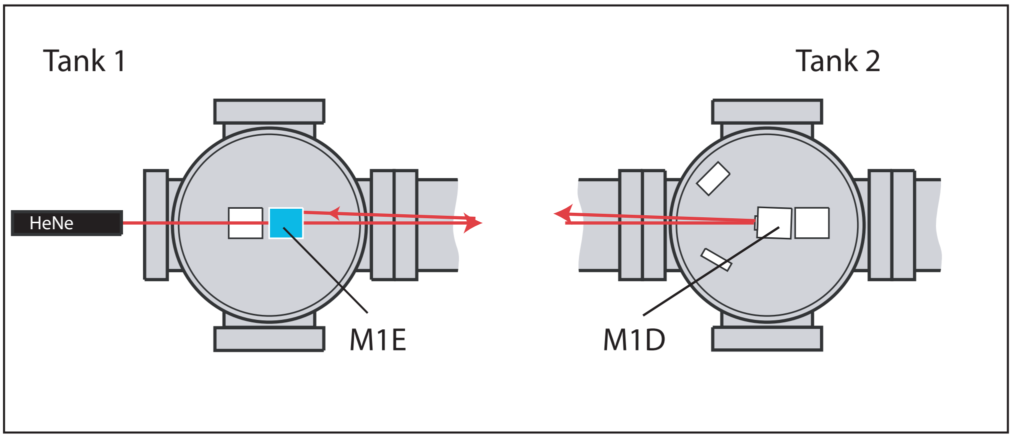

Figure 7.10: Alignment procedure for a grating coupled cavity. With both tanks open to allow visual inspection of the cavity optics, light from a HeNe laser is shone into the cavity from the conventional-mirror end.

7.3 The laser bench: mode matching and control signals

With the rough alignment complete, it was necessary to make some changes on the setup of the laser bench. During the preliminary stages of our experiment it had been left unchanged from the previous experiment (characterizing a 3-mirror coupled cavity [87]). In this experiment, the three cavity mirrors were – traveling from the laser bench to the end of the cavity – as follows: the PRM (power recycling mirror), with a focal length of 15m; the ITM (inner test mass), a plane mirror; and the ETM (end test mass), with a focal length of 15m. Thus the mode matching requirement from the laser bench into the cavity was that the beam should have a wavefront curvature of -15m, approximately 7m from the optical bench (i.e. where it encounters the PRM).

Our requirements were slightly more complex: the cavity input-optic is flat (neglecting it's diffractive properties for the moment), so we require the beam to have a flat wavefront where it meets this optic, some 11.5m from the optical bench, i.e. we must arrange that the beam has a waist at this point. Furthermore, because the beam meets the diffractive optic at an angle of 47.2°, we require that the beam be elliptical (when measured perpendicular to the beam) in order that the projection of the beam into the cavity be circular. In fact we require that the $x-$waist be 0.68 times smaller than the $y-$waist.

This requires the use of either one or two cylindrical lenses to create an appropriately sized waist at the correct position from the laser bench. We can theoretically work out where to place the cylindrical lenses (using, for instance, JamMt [88] or by hand, using the $ABCD$-matrices) if we know where the waist of the current beam is. We measure the beam profile using a Melles-Griot Super BeamAlyzer – a 7-blade knife-edge beam profiler. We decided that it was impractical to measure the beam inside the vacuum tanks. Instead, by diverting the beam from the optical table before it entered the vacuum system (see figure 7.11), we could reflect the beam back onto the optical bench from a mirror mounted on a camera tripod, 8m distance away, and measure the beam profile at various points on the bench.

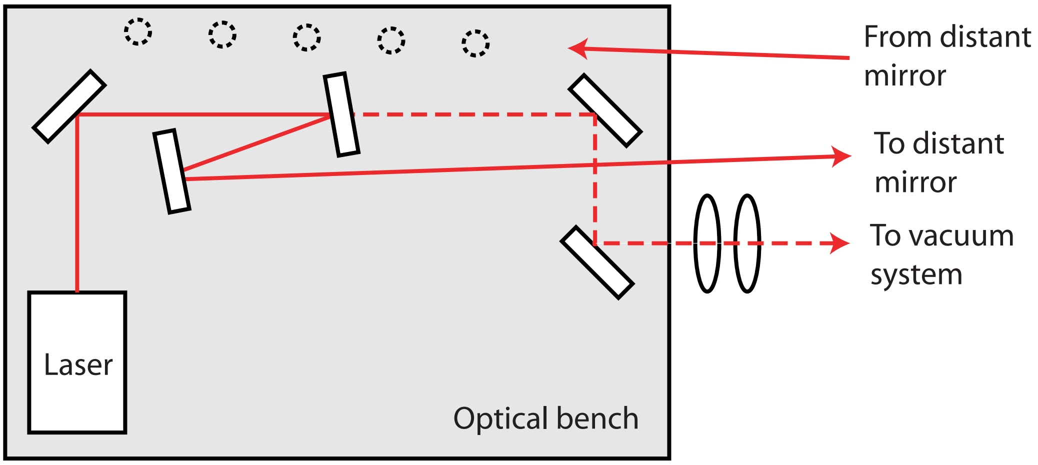

Figure 7.11: Measuring beam profiles.

7.4 Cavity control

In order for the diffractively-coupled cavity to act as an optical resonator, the two components of the cavity – the suspended diffractive coupler and end mirror – must be held at a distance of $D = m\lambda/2$. The same actuators used to control the alignment can be used to control the length of the cavity.

Pre-existing servo electronics from a previous experiment [89, 87] were used, with very minor modifications, to lock the diffractively-coupled cavity, using feedback to both the end test mass of the diffractively-coupled cavity, M1E, and to the laser frequency. The laser frequency feedback was through two paths: fast (give some idea of speed) frequency changes were made using the laser's built in piezoelectric actuator; slower frequency changes were made by changing the temperature (and therefore the length) of the laser crystal.

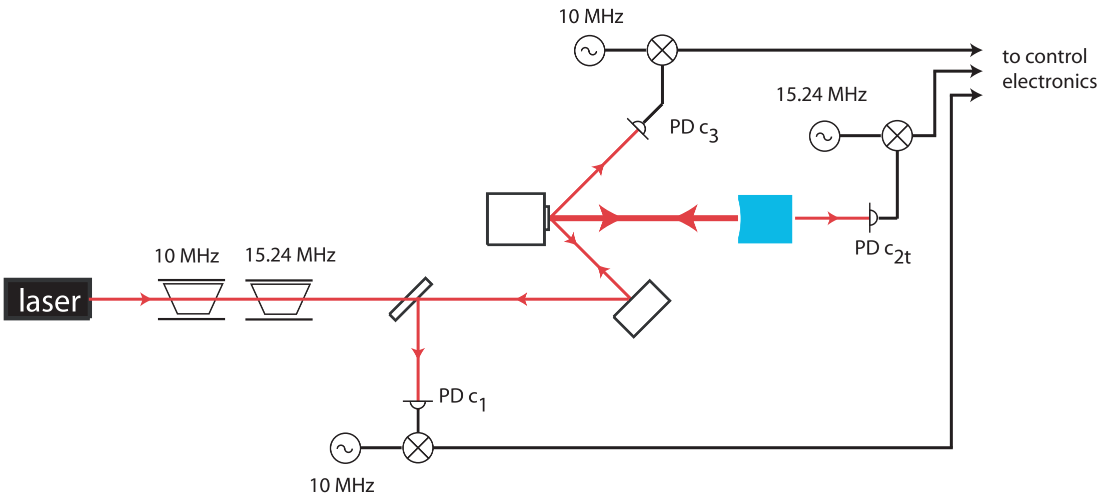

Figure 7.12: The cavity control scheme for our suspended, diffractively-coupled cavity. EOMs impose 10 MHz and 15.24 MHz phase sidebands on the 1064 nm carrier. Three photodiodes read out the demodulated signals and send them to the control electronics.

Phase sidebands of 10 MHz and 15.24 MHz were imposed upon the carrier (see figure 7.12). The value of 10 MHz was chosen for historical reasons, but the value of 15.24 MHz was chosen to be close to the free spectral range (FSR) of the diffractively-coupled cavity. The FSR was 15.27 MHz which ensure that the higher frequency sidebands would be partially resonant and thus large enough to result in a demodulated signal in transmission, measured at the photodiode PD c_{2t}. The 10 MHz sidebands could resulted in a demodulated signal at PD c_1 and PD c_3, i.e., the back-reflected and forward-reflected fields.

Using the coupling relations for a three-port grating developed in the previous chapter, and the parameters for the grating measured by the Albert-Einstein-Institute in Hannover (plus the known reflectivity of the end mirror), we could compare our measure power output of the cavity with what we would theoretically expect – this agreed to within 12%. The measured cavity finesse of

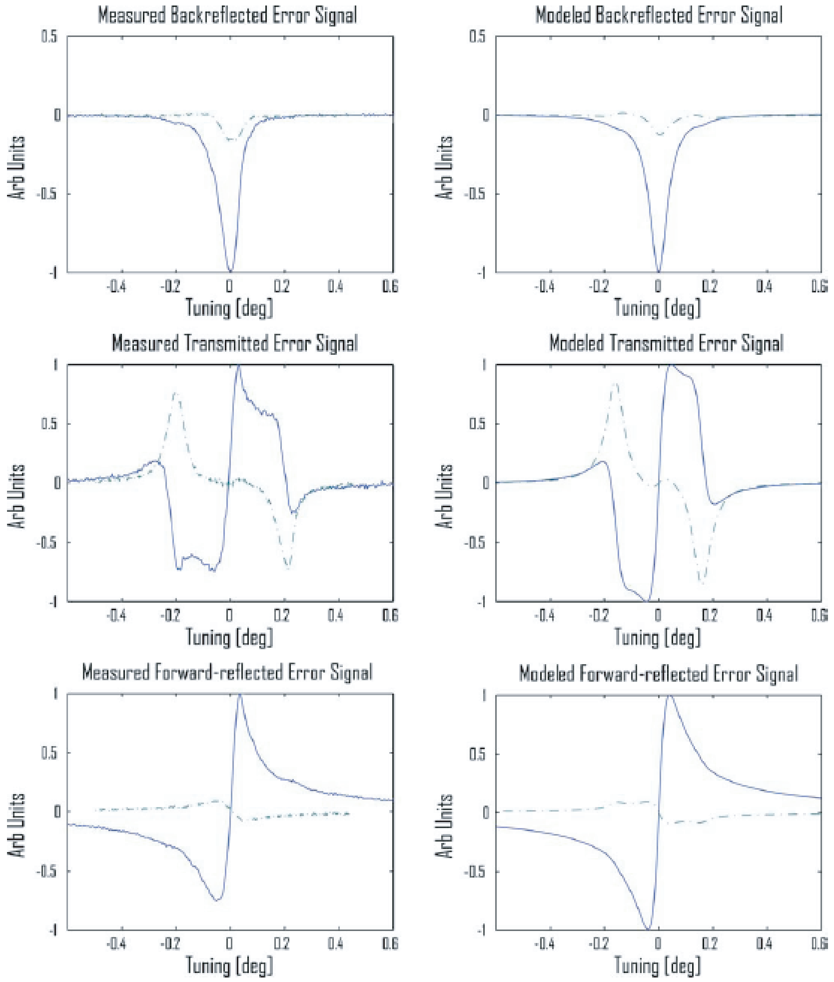

Figure 7.13: Normalized error signals from the transmitted, back-reflected and forward-reflected ports of the diffractive cavity compared with modelled error signals. The in-phase demodulation measurements are shown as dashed lines and the quadrature phase demodulation as solid lines.

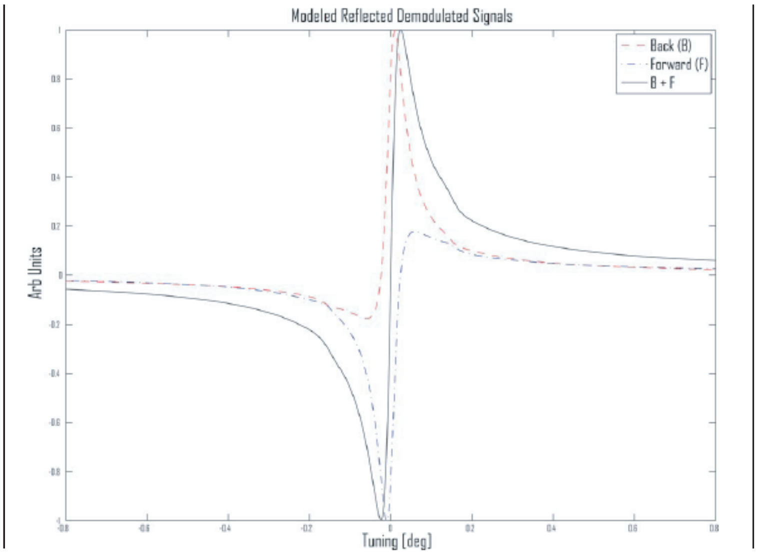

Figure 7.14: A simulated reconstruction of a traditional PDH error signal by combining back- and forward-reflected error signals for a lossless grating.

1107 ± 51 was also in close agreement with the calculated value of 1177 ± 27.

After my work on the experiment ceased, Barr and Edgar continued to characterise the cavity. Their aim was to compare the demodulated error signals obtained from all three photodiodes with those predicted by modelling the cavity. In order to do this, the cavity length was swept through one FSR. The maximum and minimum signals, obtained by mixing the local oscillator 'in-phase' and in 'quadrature-phase' (i.e. 90° away from the minimum, are shown in figure 7.13 compared with the modelled signals. The discrepancies and asymmetries in the measured values can be explained by the finite rate of sweep of the cavity length, when compared with the quasistatic model.

Previous work on a table-top grating-coupled cavity (by Hallam *et al.* [90]) had shown an asymmetry in the behaviour of the reflected carrier power levels,

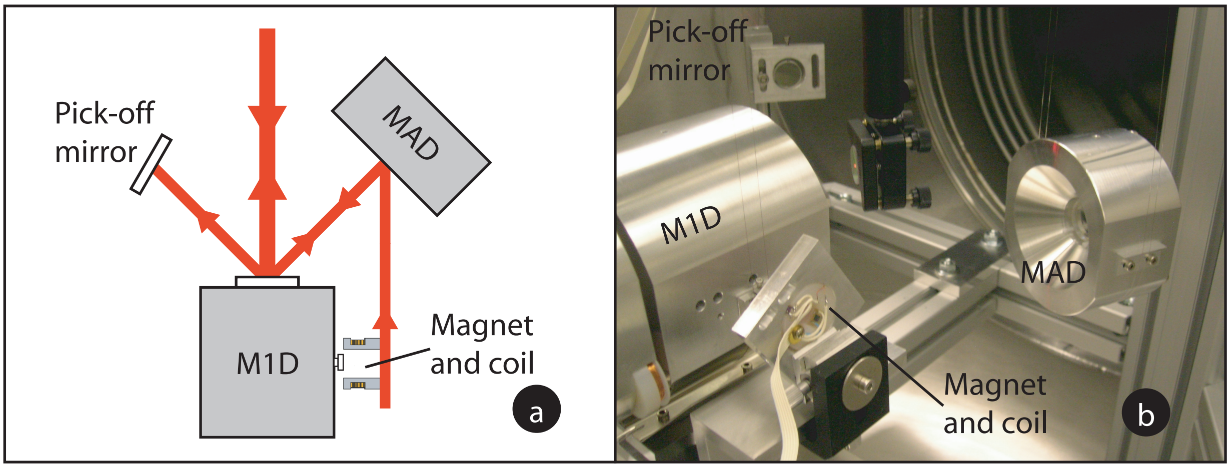

Figure 7.15: The suspended mass containing the diffractive optic, inside tank 2.

determined by η_0 and η_2. Barr and Edgar discovered that a careful choice of demodulation phase would allow the reconstruction of a traditional PDH error signal by combining the two reflected signals (see figure 7.14).

7.5 Side-motion coupling

As the previous chapter explained, the motion of the grating-coupler from side-to-side results in a frequency shift to the ±1 coupled orders. To investigate this effect, we implemented a modification to the control of M1D that allowed the grating to be displaced in the x-axis. Visible in figure 7.15 is the magnet and coil structure which drives the end-stage mass of M1D laterally. The coil was mounted on a 2-stage micrometer - this allowed the coil to be moved so that the force on the magnet pushes normal to the side of the end stage mass. The magnet was clamped, and then glued, onto the break-off of the end-stage mass.

In order to examine the movement of the end-stage mass (and thus of the diffraction grating) we decided to use a laser Doppler vibrometer (Polytec OFV-505 sensor head and OFV-5000 controller [91]). The vibrometer is a heterodyne interferometer, which compares the backscattered light from the target with a reference beam [92, 93]. The vibrometer was mounted in tank 2 and a mirror (or some cases two mirrors) positioned to direct the vibrometer beam onto the end-stage mass at one of four measurement points (see figure 7.16). Since only one point on the end-stage mass could be measured at a time the procedure was as quite laborious. The mirror, or mirrors, were arranged to direct the vibrometer's laser spot to the chosen measurement point on the end-stage mass and the vibrometer calibrated for that path-length. Often, at this point the setup was tested by applying a swept-sine signal to the side-coil (while the cavity was locked) and checking the resulting transfer function using a FFT-based dynamic signal analyser. Although the diffractive arm could not be evacuated (the cable between the vibrometer sensor head and controller compromised the vacuum seal of tank 2), we found that closing tank 2 gave more consistent results from the swept-sine measurements. It seems reasonable to believe that closing the tank gave M1D and MAD some protection from air-currents arising from the air conditioning system.

The first thing we looked for was the the signals from measuring at points 1 and 2 were balanced. This was achieved by taking successive measurements and then adjusting the micrometer stages on which the coil was mounted to improve the balance between the transfer functions measured at these points.

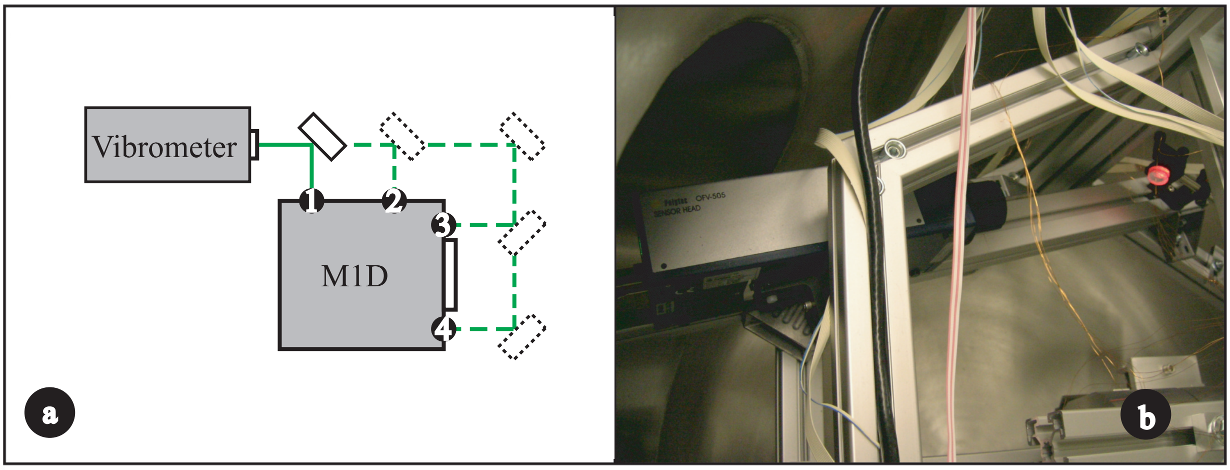

Figure 7.16: Taking vibrometer measurements inside tank 2. In (a) the four measurement and various mirror arrangements needed are shown. (b) shows a photograph of the vibrometer mounted inside the tank, measuring position 1

After a few iterations, the transfer function shown in figure 7.17 was measured. This shows that, although the end-stage mass might not be moving parallel to the x-axis (i.e. it might have some z component), points 1 and 2 remained parallel to the z-axis throughout the movement.

Results from points 3 and 4 were initially less successful. Examination of the transfer functions gained from these points seemed to show that significant back-action from the side-coil being transmitted through the Bosch frame to the mirror used to steer the vibrometer spot onto the front face of M1D. This could be improved by mounting the mirror in an inverted position from the upper Bosch frame (as can be seen in figure 7.15(b)), but not sufficiently for us to have confidence in the measurements.

Despite extensive optimisation of the experimental set-up, these problems per-

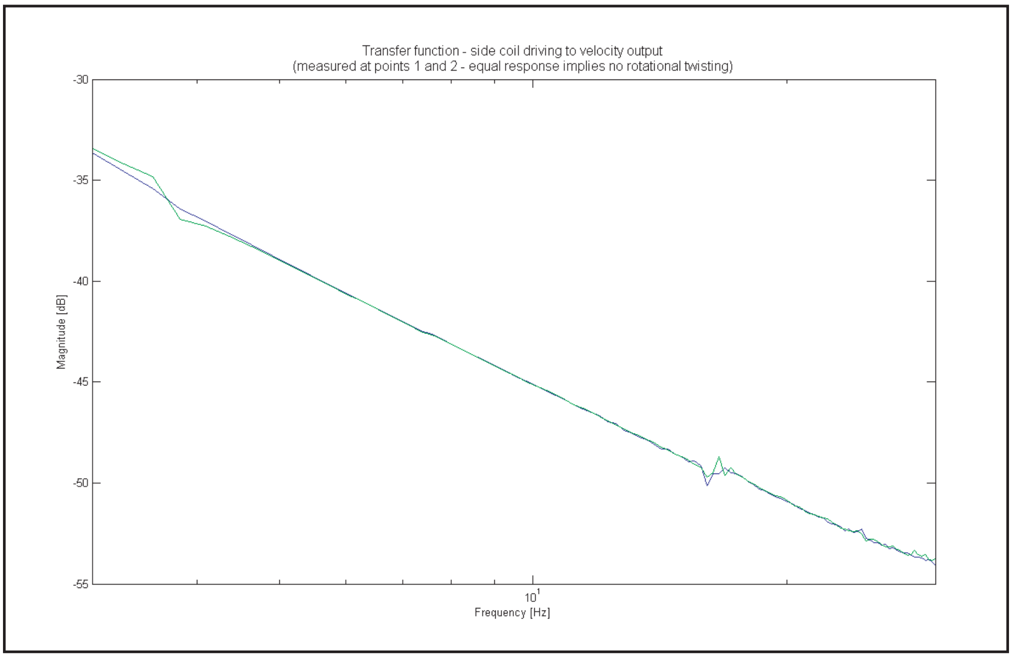

Figure 7.17: Transfer function of points 1 and 2 in response to a signal driven through the side coil. This shows that the end-stage mass of the M1D suspension moves in the x direction without twisting, but gives no information on whether it moves in the z direction as well.

sisted. Only after abandoning this approach, and rebuilding the coil mount (using an alternating rubber/lead stack) as well as a modified strategy for driving M1D, did Barr and Edgar succeed in measuring a total transfer function. This confirmed the expected $1/f$ response of the demodulated signal to the sideways motion [90].

CHAPTER 8

Conclusions

8.1 Simulations of novel optical resonators

Although there was little in the way of novel physics in the ray-traced simulations of resonator interiors, these studies did give an excellent insight into the imaging properties of resonators. Further work on the properties of ray-optical metamaterials has been carried out by members of the Optics group [26].

Our Fox-Li simulations of unstable resonators revealed some surprising never-before-seen properties of a very well studied system. It was unfortunate that simulations of spherically aberrated resonators could not be fully trusted due to an, as yet, unexplained pathology in the simulation technique.

8.2 Experimental realisation of a suspended diffractively-coupled cavity

This experiment successfully demonstrated that it was possible to use a conventional PDH scheme to control a suspended, diffractively coupled cavity. Further, it demonstrated that the traditional PDH error signal could be reconstructed by combining both reflected error signals. And finally, it gained a valuable insight into how translational, rotational and vibrational movement of the diffractive element couples into the obtained error signals [82]. Though the analysis is not yet complete, it suggests that the phase noise generated be the motion of a suspended grating-coupler would be greater than the expected decrease in coating brownian noise, making this an unsuitable technique for full-scale interferometric gravitational wave detectors.

APPENDIX A

WaveTrace - Numerical beam propagation

A.1 WaveTrace

This appendix describes WaveTrace: a library of LabVIEW VIs (see figure A.1) first developed by Johannes Courtial [94], and since built upon by Courtial and others.

WaveTrace is a based around a 2D Fast Fourier Transform beam propagation algorithm (described in the next section). It allows for an arbitrary beam to be propagated from one plane to another, and for that beam to interact with

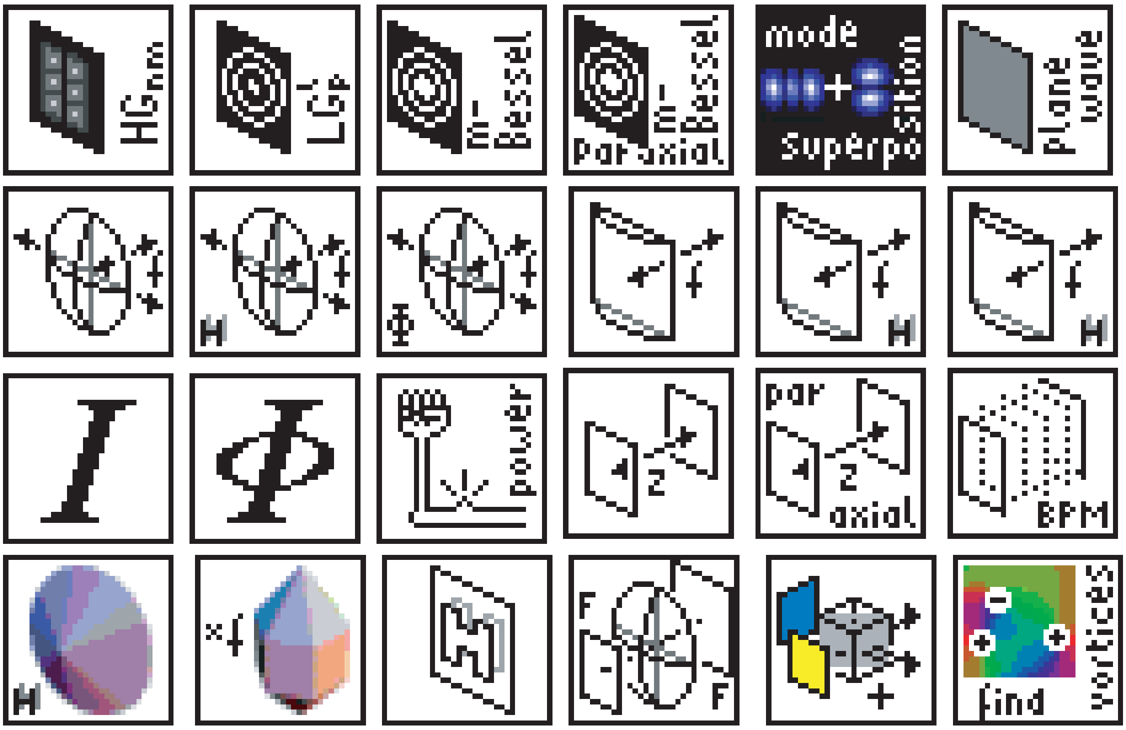

Figure A.1: WaveTrace v2.0 – a library of LabVIEW VIs for simulating and propagating arbitrary beams. Shown here are a few of the available VIs for creating various types of beams and lenses, measuring their properties and various means of propagating them.

phase holograms represent lenses, prisms or indeed any imaginable diffractive optical elements.

A.2 Beam Propagation Algorithm

If we take an initial electric field $E_{z_0}(x, y)$ in the plane $z = z_0$, then to propagate this to a plane a distance Δz away, we proceed as follows. First the 2D Fourier transform of the initial field is calculated,

$$e_{z_0}(k_x, k_y) = FT[E_{z0(x,y)}(k_x, k_y)], \tag{A.1}$$

which corresponds to a plane wave decomposition of the form

$$e_{z_0}(k_x, k_y)\exp(-i(k_x x + k_y y)). \tag{A.2}$$

That is, the superposition of the 3D plane waves of which equation A.2 gives the 2D coss sections) creates the field $E(x, y, z)$. The k_x and k_y components of each plane wave are the same in every z-plane, and the k_z component is given by

$$k = \sqrt{k_x^2 + k_y^2 + k_z^2} = \frac{2\pi}{\lambda} \tag{A.3}$$

By propagating to a plane a distance Δz from the initial plane, the cross section through each plane wave changes phase by $-k_z \Delta z$, thus in the plane $z = z_0 + \Delta z$ we have,

$$e_{z_0+\Delta z}(k_x, k_y) = e_{z_0}(k_x, k_y)\exp(-i\sqrt{k^2 + k_x^2 + k_y^2}\Delta z). \tag{A.4}$$

The cross-sectional electric field in this plane, $E_{z_0+\Delta z}(x, y)$, is then given by the inverse Fourier transform of equation A.4.

APPENDIX B

POV-Ray code listing: visualization of a canonical optical resonator

A sample of the POV-Ray code used to produce the figures in chapter 3 is included for reference.

```

// Simple symmetric resonator g=0.2
// 17.5.06


global_settings { assumed_gamma 2.2 max_trace_level 50 }

#include ”shapes.inc”
```

```
#include ”colors.inc”
#include ”textures.inc”
#include ”skies.inc”
#include ”metals.inc”
#include ”woods.inc”
#include ”golds.inc”
#include ”glass.inc”


// Choose quality of focal blur - affects rendering time

#declare FB_Quality_Off     =  0;
#declare FB_Quality_Fast    =  1;
#declare FB_Quality_Default =  2;
#declare FB_Quality_High    =  3;

#declare FB_Quality= FB_Quality_Fast;



// Macros
//=============================================================================

// g parameter macro:
//        Takes input g-parameter, G
//        Returns corresponding focal length, F
//        Mirror separation, L
#macro gParameter(G,L)

#switch(G)

```

```
  // deal with infinite focal length, which occurs for g=1
  #case(1)
    #local part1 = 1000;
    #local part2 = L;
  #break
  #else
    #local part1 = 0.5*L;
    #local part2 = 1/(1-G);
  #end
  part1*part2
#end



// Spherical mirror macro
// parameters:
//  F = focal length of mirror
//  A = radius of aperture of mirror
// returns a spherical mirror



#declare MirrorSurface = texture {
    pigment { BrightGold }
    finish {
        ambient .1
        diffuse .1
        specular 1
        roughness .001
        metallic
        reflection {
```

```
                1
                metallic
            }
        }
}



#macro SphericalMirror(F, A)
    #local R = 2*F; // radius of curvature of mirror
    #local Z = sqrt(R*R - A*A);
    #if (R>0)
        difference {
            cylinder {
                <0,0,0.001>, <0,0,-(R-Z)>, A
                pigment {Black}
            }
            sphere {
                <0, 0, -R>, R
                texture {MirrorSurface}
            }
        }
    #else
        intersection {
            cylinder {
                <0,0,-0.001>, <0,0,-R-Z+0.05*A>, A
                pigment {Black}
            }
            sphere {
                <0, 0, -R>, R
                texture {MirrorSurface}
```

```
            }
        }
    #end
#end

// End of macros
//==========================================================================



//Resonator parameters

#local G=0.2;                      // resonator g-parameter
#local A=1;                        // radius of mirrors
#local L=4;                        // length of resonator
#local M=L/2;


#local F=gParameter(G,L);          // focal length of mirrors



// Camera position 1 - inside the resonator

camera {
  location  <0, 1, 0.1>
  direction 2*z
  right     x*image_width/image_height
  look_at   <0, 1 , L>

// Camera position 2 - side view
```

```

//camera {
//   location   <7, 1, -L/2>
//   direction  1.5*z
//   right      x*image_width/image_height
//   look_at    <0, 1, L/2>

// Focal blur

#if(FB_Quality != FB_Quality_Off)
        aperture 0.7                  // Set depth of field
        focal_point <0, 1 , 7.7>      // Set focal point
#end

#switch(FB_Quality) #case(FB_Quality_Off)
    aperture 0
          #debug "\nNo focal blur used...\n"
#break
#case (FB_Quality_Fast)
    blur_samples 7
    confidence 0.5              // default is 0.9
    variance 1/64               // default is 1/128 (0.0078125)
          #debug "\nFast focal blur used...\n"
#break
#case(FB_Quality_Default)
    blur_samples 19
    confidence 0.90             // default is 0.9
    variance 1/128              // default is 1/128 (0.0078125)
          #debug "\nDefault focal blur used...\n"
#break
#case(FB_Quality_High)
```

```
   blur_samples 37
   confidence 0.975           // default is 0.9
   variance 1/255             // default is 1/128 (0.0078125)
        #debug "\nHigh Quality focal blur used...\n"
#break
#else
        #debug "\nNo focal blur used...\n"
#end }


// Make the sky and ground

background { color rgb<0.2, 0.4, 0.8>  } light_source { <30, 30,
-30> color rgb 1 } plane {
  y, 0
  pigment { checker color Brown, color Yellow scale 1 }
}

sky_sphere {
   pigment {
     gradient y
     color_map {
       [0.0 rgb <0.6,0.7,1.0>]
       [0.7 rgb <0.0,0.1,0.8>]
     }
   }
}



//Declare objects
```

```

#declare mirror_1=SphericalMirror(F,A) #declare
mirror_2=SphericalMirror(F,A) #declare P_object= text {
    ttf "timrom.ttf" "P", 0.01, 0
    pigment { Col_Glass_Green }
    finish { reflection 0 specular 0 }
    scale 2
  }


//Place objects in scene

object{mirror_1
  translate 1*y
  translate L*z
}

object{mirror_2
  rotate 180*y
  translate 1*y
}

object{P_object
        translate 0.1*y
        translate -0.6*x
        translate 1.5*z
}

//==================================================================
```

Bibliography

[1] C. Fabry and A. Perot. Theorie et applications d'une nouvelle methode de spectroscopie interferentielle. *Ann. Chim. Phys.*, 7:115, 1899.

[2] A. A. Michelson and E. W. Morley. On the relative motion of the Earth and the luminiferous ether. *Am. J. Sci.*, 34:333, 1887.

[3] R. W. P. Drever, J. Hough, A. J. Munley, S. A. Lee, R. Spero, S. E. Whitcomb, H. Ward, G. M. Ford, M. Hereld, N. A. Robertson, I. Kerr, J. R. Pugh, G. P. Newton, B. Meers, E. D. Brooks, and Y. Gursel. Optical cavity laser interferometry for gravitational wave detection. *Laser Spectroscopy*, 5:33, 1981.

[4] A. G. Fox and T. J. Li. *Bell. Syst. tech. J.*, 41:453, 1966.

[5] A. E. Siegman. A canonical formulation for analyzing multielement unstable resonators. *IEEE J. Quantum Elect.*, QE-12:35, 1976.

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

[7] A. Forrester, M. Lönnqvist, M. J. Padgett, and J. Courtial. Why are the eigenmodes of stable laser resonators structurally stable? *Opt. Lett.*, 27:1869, 2002.

[8] M. Bondarescu and K. S. Thorne. A new family of light beams and mirror shapes of future ligo interferometers. *arXiv.*, gr-qc/0409083, 2004.

[9] S. Sridhar. Experimental observation of scarred eigenfunctions of chaotic microwave cavities. *Phys. Rev. Lett.*, 67:785, 1991.

[10] G. Toraldo Di Francia. Flat-roof resonators. *Appl. Opt.*, 4:1267, 1965.

[11] A. E. Siegman. *An Introduction to Lasers and Masers.* McGraw-Hill, 1971.

[12] G. Toraldo Di Francia. Optical resonators. In S. M. Kay and A. Maitland, editors, *Quantum Optics*, volume 41, page 453, 1966.

[13] H. Kogelnik and T. Li. Laser beams and resonators. *Appl. Opt.*, 5:1550, 1966.

[14] J. Nelson, J. Courtial, and G. Whyte. Photorealistic visualization of imaging in cononical optical resonators. *Am. J. Phys.*, 76:991, 2008.

[15] J. Courtial and J. Nelson. Ray-optical negative refraction and pseudoscopic imaging with Dove-prism arrays. *New J. Phys.*, 10:023028, 2008.

[16] POV-Team. Pov-ray hall of fame, accessed on 6/3/2006.

[17] D Sweet, E Ott, and J A Yorke. Topology in chaotic scattering. *Nature*, 399:6734, 1999.

[18] POV-Team. Pov-ray hall of fame, accessed on 6/3/2006.

[19] A. E. Siegman. *Lasers*, page 902. University Science Books, Mill Valley, California, 1986.

[20] J. Courtial and M. J. Padgett. Monitor-outside-a-monitor effect and self-similar fractal structure in the eigenmodes of unstable optical resonators. *Phys. Rev. Lett.*, 85:5320–5323, 2000.

[21] POV-Team. Pov-ray reference (for pov-ray version 3.5.1), 2006.

[22] A. E. Siegman. *Lasers*, page 599ff. University Science Books, Mill Valley, California, 1986.

[23] A. E. Siegman. *Lasers*, chapter 19. University Science Books, Mill Valley, California, 1986.

[24] A. E. Siegman. *Lasers*, chapter 15.3. University Science Books, Mill Valley, California, 1986.

[25] Calum M. G. Watterson, Miles J. Padgett, and Johannes Courtial. Classic-fractal eigenmodes of unstable canonical resonators. *Opt. Commun.*, 223:17–23, 2003.

[26] A C Hamilton and J Courtial. Metamaterials for light rays: ray optics without wave-optical analog in the ray-optics limit. *New J. Phys.*, 11:013042, 2009.

[27] S Lansel, M A Porter, and L A Bunimovich. One-particle and few-particle billiards. *arXiv*, 2005.

[28] C. A. Kruelle, A. Kittel, J. Peinke, R. Richter, and R. P. Huebener. Chaotic billiards seen as mirror cabintets. *Physica D*, 102:227, 1997.

[29] L A Bunimovich. The ergodic properties of certain billiards. *Funkt. Anal. Prilozh.*, 8:73, 1974.

[30] W Dultz. The bust of the tyrant: an optical illusion. *Appl. Opt.*, 23:200, 1984.

[31] A. G. Fox and T. Li. Resonant modes in a maser interferometer. *Bell. Syst. tech. J.*, 40:453, 1961.

[32] E. D'Ambrosio and K. S. Thorne. Nonspherical mirrors to reduce thermoelastic noise in advanced gravitational wave interferometers. *Phys. Rev. D*, 67:102004, 2003.

[33] P. E. Jackson. *The physics and technology of laser resonators*, chapter Numerical Resonator Calculations, page 106. Adam Hilger, 1989.

[34] A. E. Siegman. *Lasers.* University Science Books, Mill Valley, California, 1986.

[35] L. A. Vaĭnshteĭn. Open resonators for lasers. *Sov. Phys. JETP*, 17:709, 1963.

[36] S. Zaidi and D. L. MacFarlane. Mode evolution in a droplet. *Opt. Lett.*, 17:562, 1992.

[37] J. R. Meyer-Arendt. *Introduction to Classical and Modern Optics.* Prentice-Hall, 1972.

[38] M. G. Tarallo, J. Miller, J. Agresti, E. D'Ambrosio, R. DeSalvo, D. Forest, B. Lagrange, J. M. Mackowski, C. Michel, J. L. Montorio, N. Morgado, L. Pinard, A. Remillieux, B. Simoni, and P. Willems. Generation of a flat-

top laser beam for gravitational wave detectors by means of a nonspherical fabry-perot resonator. *Appl. Opt.*, 46:6648, 2007.

[39] M. Born and E. Wolf. *Principles of Optics*, chapter 9, pages 523–527. Cambridge University Press, 1999.

[40] P. V. P. Yupapin and K. T. V. Grattan. The stability properties of an axicon resonator. *Meas. Sci. Technol.*, 2:686, 1991.

[41] M. Endo. Doughnut like beam generation by a w-axicon resonator with variable geometry. *Jpn. J. Appl. Phys.*, 46:593, 2007.

[42] G. P. Karman and J. P. Woerdman. Fractal structure of eigenmodes of unstable-cavity lasers. *Opt. Lett.*, 23:1909–1911, 1998.

[43] G. P. Karman, M. W. Beijersbergen, A. van Duijl, D. Bouwmeester, and J. P. Woerdman. Airy pattern reorganization and subwavelength structure in a focus. *J. Opt. Soc. Am. A*, 15:884–899, 1998.

[44] G. P. Karman, G. S. McDonald, G. H. C. New, and J. P. Woerdman. Fractal modes in unstable resonators. *Nature*, 402:138, 1999.

[45] J. Nelson and J. Courtial. Geometrical volume imaging in laser resonators and 3d fractal laser modes. Photon06, Optics and Photonics Division: Advanced Imaging, September 2006.

[46] B. B. Mandelbrot. *The Fractal Geometry of Nature.* Freeman, San Francisco, New York, 1982.

[47] H.-O. Peitgen, H. Jürgens, and D. Saupe. *Chaos and Fractals: new frontiers of science*, page 20. Springer, New York, 1992.

[48] J. Courtial, J. Leach, and M. J. Padgett. Fractals in pixellated video feedback. *Nature*, 414:864, 2001.

[49] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie (The foundation of the the general theory of relativity). *Annalen Der Physick*, 7:50, 1916.

[50] K. S. Thorne. Gravtational radiation. In S. W. Hawking and W. Israel, editors, *300 Years of Gravitation*, chapter 9, page 330. Cambridge University Press, 1987.

[51] J. J. Sakurai. *Modern Quantum Mechanics.* Addison-Wesley, 1993.

[52] R. A. Mould. *Basic Relativity.* Springer-Verlag, 1994.

[53] R. Maté. When Worlds Collide. *Paramount Pictures*, 1951.

[54] F A E Pirani. Invariant formulation of gravitational radiation theory. *Phys. Rev.*, 105:1089, 1957.

[55] J H Taylor and L A Fowler. Measurements of general relativistic effects in the binary pulsar psr $1913+16$. *Nature*, 277:437, 1979.

[56] Michele Maggiore. Gravitational wave experiments and early universe cosmology. *Physics Reports*, 331(6):283 – 367, 2000.

[57] B. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. Amin, D. P. Anderson, S. B. Anderson, W. G. Anderson, M. A. Arain, M. Araya, H. Armandula, P. Armor, Y. Aso, S. Aston, P. Aufmuth, C. Aulbert, S. Babak, S. Ballmer, H. Bantilan, B. C. Barish, C. Barker, D. Barker, B. Barr, P. Barriga, and M. A. Barton. Einstein@home search for periodic gravitational waves in ligo s4 data. *Phys. Rev. D*, 79:022001, 2009.

[58] J. Weber. Anisotropy and polarization in the gravitational-raditation experiments. *Phys. Rev. Lett.*, 22:1320, 1969.

[59] S. E. Whitcomb. Ground-based gravitational-wave detection: now and future. *Class. Quantum Grav.*, 25:114013, 2008.

[60] F. B. Estabrook and H. D. Wahlquist. Response of doppler spacecraft tracking to gravitational radiaton. *General Relativity and Gravitation*, 6:439, 1975.

[61] J. W. Armstrong. Low-frequency gravitational wave searches using spacecraft doppler tracking. *Living Reviews in Relativity*, 9, 2006.

[62] B. Abbott, R. Abbott, R. Adhikari, A. Ageev, J. Agresti, B. Allen, J. Allen, R. Amin, S. B. Anderson, W. G. Anderson, M. Araya, H. Armandula, M. Ashley, F. Asiri, P. Aufmuth, C. Aulbert, S. Babak, R. Balasubramanian, S. Ballmer, B. C. Barish, C. Barker, D. Barker, M. Barnes, B. Barr, M. A. Barton, K. Bayer, and R. Beausoleil. First all-sky upper limits from ligo on the strength of periodic gravitational waves using the hough transform. *Phys. Rev. D*, 72:102004, 2005.

[63] W. H. Steel. *Interferometry.* Cambridge University Press, 1967.

[64] B. W. Barr. *Experimental investigations into advanced configurations and optical techniques for laser interferometric gravitational wave detectors.* PhD thesis, University of Glasgow, 2003.

[65] M Punturo, M Abernathy, F Acernese, B Allen, N Andersson, K Arun, F Barone, B Barr, M Barsuglia, M Beker, N Beveridge, S Birindelli, S Bose, L Bosi, S Braccini, C Bradaschia, T Bulik, E Calloni, G Cella, E Chassande Mottin, S Chelkowski, A Chincarini, J Clark, E Coccia,

C Colacino, J Colas, A Cumming, L Cunningham, E Cuoco, S Danilishin, K Danzmann, G De Luca, R De Salvo, T Dent, R Derosa, L Di Fiore, A Di Virgilio, M Doets, V Fafone, P Falferi, R Flaminio, J Franc, F Frasconi, A Freise, P Fulda, J Gair, G Gemme, A Gennai, A Giazotto, K Glampedakis, M Granata, H Grote, G Guidi, G Hammond, M Hannam, J Harms, D Heinert, M Hendry, I Heng, E Hennes, S Hild, J Hough, S Husa, S Huttner, G Jones, F Khalili, K Kokeyama, K Kokkotas, B Krishnan, M Lorenzini, H Lück, E Majorana, I Mandel, V Mandic, I Martin, C Michel, Y Minenkov, N Morgado, S Mosca, B Mours, H Müller-Ebhardt, P Murray, R Nawrodt, J Nelson, R Oshaughnessy, C D Ott, C Palomba, A Paoli, G Parguez, A Pasqualetti, R Passaquieti, D Passuello, L Pinard, R Poggiani, P Popolizio, M Prato, P Puppo, D Rabeling, P Rapagnani, J Read, T Regimbau, H Rehbein, S Reid, L Rezzolla, F Ricci, F Richard, A Rocchi, S Rowan, A Rdiger, B Sassolas, B Sathyaprakash, R Schnabel, C Schwarz, P Seidel, A Sintes, K Somiya, F Speirits, K Strain, S Strigin, P Sutton, S Tarabrin, J van den Brand, C van Leewen, M van Veggel, C van den Broeck, A Vecchio, J Veitch, F Vetrano, A Vicere, S Vyatchanin, B Willke, G Woan, P Wolfango, and K Yamamoto. The third generation of gravitational wave observatories and their science reach. *Classical and Quantum Gravity*, 27(8):084007, 2010.

[66] S. Hild, S. Chelkowski, and A. Freise. Pushing towards the et sensitivity using 'conventional' technology. *arXiv.org*, 2008.

[67] M P Edgar, B W Barr, J Nelson, M V Plissi, K A Strain, O Burmeister, M Britzger, K Danzmann, R Schnabel, T Clausnitzer, F Brückner, E-B Kley, and A Tünnermann. Experimental demonstration of a suspended,

diffractively coupled fabry-perot cavity. *Class. Quantum Grav.*, 27:084029, 2010.

[68] K. A. Strain and B. J. Meers. Experimental demonstration of dual recycling for interferometric gravitational-wave detectors. *Phys. Rev. Lett.*, 66:1391, 1991.

[69] A Lazzarini and R Weiss. Ligo science requirements document (srd). Technical report, California Institute of Technology, 1996.

[70] F. Acernese, F. Antonucci, S. Aoudia, K.G. Arun, P. Astone, G. Ballardin, F. Barone, M. Barsuglia, Th.S. Bauer, M.G. Beker, S. Bigotta, S. Birindelli, M. Bitossi, M.A. Bizouard, M. Blom, C. Boccara, F. Bondu, L. Bonelli, L. Bosi, S. Braccini, C. Bradaschia, A. Brillet, V. Brisson, R. Budzynski, T. Bulik, H.J. Bulten, D. Buskulic, G. Cagnoli, E. Calloni, E. Campagna, B. Canuel, F. Carbognani, F. Cavalier, R. Cavalieri, G. Cella, E. Cesarini, E. Chassande-Mottin, A. Chincarini, F. Cleva, E. Coccia, C.N. Colacino, J. Colas, A. Colla, M. Colombini, C. Corda, A. Corsi, J.-P. Coulon, E. Cuoco, S. D'Antonio, A. Dari, V. Dattilo, M. Davier, R. Day, R. De Rosa, M. Del Prete, L. Di Fiore, A. Di Lieto, M. Di Paolo Emilio, A. Di Virgilio, A. Dietz, M. Drago, V. Fafone, I. Ferrante, F. Fidecaro, I. Fiori, R. Flaminio, J.-D. Fournier, J. Franc, S. Frasca, F. Frasconi, A. Freise, L. Gammaitoni, F. Garufi, G. Gemme, E. Genin, A. Gennai, A. Giazotto, M. Granata, C. Greverie, G. Guidi, H. Heitmann, P. Hello, S. Hild, D. Huet, P. Jaranowski, I. Kowalska, A. Krlak, P. La Penna, N. Leroy, N. Letendre, T.G.F. Li, M. Lorenzini, V. Loriette, G. Losurdo, J.-M. Mackowski, E. Majorana, N. Man, M. Mantovani, F. Marchesoni, F. Marion, J. Marque, F. Martelli, A. Masserot,

F. Menzinger, C. Michel, L. Milano, Y. Minenkov, M. Mohan, J. Moreau, N. Morgado, A. Morgia, S. Mosca, V. Moscatelli, B. Mours, I. Neri, F. Nocera, G. Pagliaroli, C. Palomba, F. Paoletti, S. Pardi, M. Parisi, A. Pasqualetti, R. Passaquieti, D. Passuello, G. Persichetti, M. Pichot, F. Piergiovanni, M. Pietka, L. Pinard, R. Poggiani, M. Prato, G.A. Prodi, M. Punturo, P. Puppo, O. Rabaste, D.S. Rabeling, P. Rapagnani, V. Re, T. Regimbau, F. Ricci, F. Robinet, A. Rocchi, L. Rolland, R. Romano, D. Rosinska, P. Ruggi, F. Salemi, B. Sassolas, D. Sentenac, R. Sturani, B. Swinkels, A. Toncelli, M. Tonelli, E. Tournefier, F. Travasso, J. Trummer, G. Vajente, J.F.J. van den Brand, S. van der Putten, M. Vavoulidis, G. Vedovato, D. Verkindt, F. Vetrano, A. Vicer, J.-Y. Vinet, H. Vocca, M. Was, and M. Yvert. Measurements of superattenuator seismic isolation by virgo interferometer. *Astroparticle Physics*, 33(3):182 – 189, 2010.

[71] C Hardham, B Abbott, R Abbott, G Allen, R Bork, C Campbell, K Carter, D Coyne, D DeBra, T Evans, J Faludi, A Ganguli, J Giaime, M Hammond, W Hua, J Kern, J LaCour, B Lantz, M Macinnis, K Mailand, K Mason, R Mittleman, J Nichol, J Niekerk, B O'Reilly, D Ottaway, H Overmier, C Parameswariah, J Phinney, B Rankin, N A Robertson, D Sellers, P Sarin, D H Shoemaker, O Spjeld, G Traylor, S Wen, R Wooley, and M Zucker. Quiet hydraulic actuators for ligo. In *4th IFAC Symposion on Mechatronic Systems*, 2006.

[72] K.-X. Sun and R. L. Byer. All-reflective Michelson, Sagnac, and Fabry-Perot interferometers based on grating beam splitters. *Opt. Lett.*, 23:567, 1998.

[73] S A Hughes and K S Thorne. Seismic gravity-gradient noise in interfero-

metric gravitational-wave detectors. *Phys. Rev. D*, 58:122002, 2002.

[74] W. Winkler, K. Danzmann, A. Rüdiger, and R. Schilling. Heating by optical absorption and the performance of interferometric gravitation-wave detectors. *Phys. Rev. A*, 44:7022, 1991.

[75] W. Winkler, A. Rüdiger, R. Schilling, K. A. Strain, and K. Danzmann. Birefringence-induced losses in interferometers. *Opt. Comm.*, 112:245, 1994.

[76] A. E. Siegman. *Lasers, Chapter 11.* Oxford University Press, 1986.

[77] A. Freise, A. Bunkowski, and R. Schnabel. Phase and alignment noise in grating interferometers. *New J. Phys.*, 9:433, 2007.

[78] Daniel J. Schroeder. *Astronomical Optics.* Academic Press, 1987.

[79] A. Bunkowski, O. Burmeister, K. Danzmann, and R. Schnabel. Input-output relations for a three-port grating coupled fabry-perot cavity. *Opt. Lett.*, 30:1183, 2005.

[80] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Appl. Phys. B*, 31:97, 1983.

[81] C. C. Davis. *Lasers and Electro-optics*, chapter Chapter 19, pages 472–507. Cambridge University Press, 1996.

[82] B. Barr, M. P. Edgar, J. Nelson, M. V. Plissi, S. H. Huttner, K. A. Strain, O. Burmeister, M. Britzger, R. Schnabel, K. Danzmann, J. Hallam, A. Freise, T. Clausnitzer, F. Brückner, E. B. Kley, and A. Tünnermann.

Translational effects in diffractively coupled optical cavities. In preparation.

[83] V. B. Braginsky and S. P. Vyatchanin. Corner reflectors and quantum-non-demolition measurements in gravitational wave antennae. *Phys. Lett. A*, 324:345, 2004.

[84] Federal standard airborne particulate cleanliness classes in cleanrooms and clean zones, 1992.

[85] J. R. Taylor. *Interferometric experiments towards advanced gravitational wave detectors.* PhD thesis, University of Glasgow, 2008.

[86] Innolight. *Ultrastable, CW, single-frequency lasers.*

[87] S. H. Huttner, B. W. Barr, M. V. Plissi, J. R. Taylor, B. Sorazu, and K. A. Strain. Novel sensing and control schemes for a three-mirror coupled cavity. *Class. Quantum Grav.*, 24:3825, 2007.

[88] N Lastzka and A Thuering. *JamMt - Just another mode matching tool*, 0.22 edition, 2004.

[89] B W Barr, S H Huttner, J R Taylor, B Sorazu, M V Plissi, and K A Strain. Optical modulation techniques for length sensing and control of optical cavities. *Appl. Opt.*, 46:7739, 2007.

[90] J. Hallam, S. Chelkowski, A. Freise, S. Hild, B. Barr, K. A. Strain, O. Burmeister, and R. Schnabel. Coupling of lateral grating displacement to the output ports of a diffractive fabry-perot cavity. *J. Opt. A: Pure Appl. Opt.*, 11:985592, 2009.

[91] Polytec. *OFV-5000 Vibrometer Controller*, 2008.

[92] Y. F. Zastrogin. Optical noncontacting doppler methods of mechanical vibration measurement. *Meas. Techniques*, 16:365, 1973.

[93] C. M. Sutton. Accelerometer calibration by dynamic position measurement using heterodyne laser interferometry. *Metrologia*, 27:133, 1990.

[94] J Courtial. *Angular momentum of light, self-imaging beams, and fractal resonator modes.* PhD thesis, University of St. Andrews, 1999.