HIGH PERFORMANCE VIBRATION ISOLATION FOR GRAVITATIONAL WAVE DETECTION

John Winterflood, B.App.Sc., P.G.Dip

This thesis is presented for the degree of
Doctor of Philosophy
of the University of Western Australia
Department of Physics
2001

LIGO-P020028-00-R

Abstract

This thesis describes research done in the endeavour to isolate better very sensitive measurement apparatus from the all pervasive seismic motion. The intended application is the detection of gravitational waves using laser interferometry with suspended optics, but the techniques researched are applicable to any situation in which seismic motion at frequencies above a fraction of a hertz is a problem. Very good isolation at frequencies above 10 Hz or so is relatively easy to achieve with a suspension chain of cascaded spring-mass isolators, and so an emphasis in this thesis is the push for better isolation and less residual motion at lower frequencies.

Remaining problems with suspension chains are first addressed. One problem in the isolators has been that the mass of the spring elements required for vertical isolation allow internal modes which bypass their isolation at undesirably low frequencies (a few tens of hertz). A remarkably simple and ingenious new structure for vertical isolation is presented which reduces this problem towards its ultimate minimum - and thereby minimises the number of stages required in the suspension chain.

Another problem in the suspension chain is the large residual motion of the isolated mass due to the normal mode resonances formed by its high Q-factor stages. A new method of strongly and passively damping these modes is presented which can allow normal mode Q-factors to be reduced well below 10. With such damping the residual motion is easily reduced by an order of magnitude.

Even with a well damped suspension chain, the residual motion is dominated by the motion at the normal mode resonances of the chain. A very effective method to reduce residual motion further is to add an ultra-low frequency (ULF) pre-isolator stage in front of the isolator chain. Such a device can reduce the seismic drive to the normal mode resonances and thus the residual motion, by two orders of magnitude. A selection of ULF structures for both vertical and horizontal isolation are briefly described. Then a few of the structures that seemed the most suitable and for which full sized isolation stages were built and measured are presented in detail.

All the methods presented to this point have been entirely passive, with active control only being required to compensate for temperature drifts, alignment etc. It is shown that there is an intrinsic limit to the isolation that can be obtained by such purely passive techniques and that a single stage of ULF pre-isolation together with a well damped chain is sufficient to reach this limit. This limit is due to the effect of seismic

tilt which will then dominate over translation and which can only be overcome by inertial sensing and active feedback control. The performance requirements of a tilt sensor and servo system are fully investigated and proposed as the final components required for a vibration isolation system of ultimate performance and minimal complexity.

Acknowledgments

First and foremost I must thank my wife Ming-Pei for her untiring labour and support in all areas of our life together while I undertook this research. I must also "blame" her for persuading me to start this PhD rather than the Master's I was contemplating!

Secondly I must thank my supervisors, primarily Prof David Blair for all the encouragement, appreciation, and freedom to pursue and contribute to this interesting area of research and for his untiring enthusiasm for the project. Also Dr Mark Notcutt who was a great guy to work with before he moved to greener pastures.

Thirdly I must thank the colleagues I worked alongside of through these years. Primarily Dr Ju Li, who's PhD research laid the foundation that I have built on, and who has been a real pleasure and encouragement to work beside. But also Liu, Zhao, Yang, Mitzuru, Fetah, and more recently David Coward and John Bongiovanni for their friendship.

I must also thank the workshop staff for the major effort that was put in to build some of the complex mechanical contraptions that I managed to dream up, and Peter and Mike for their electronics support.

During this research I spent several periods totalling almost a year with the VIRGO suspension group in Italy. I really appreciated all the help, friendships, and experiences of Italy during my times there. So many thanks to Adalberto, Riccardo, Rafaelli, Rosa, Giovanni, both Stephano's, and really too many others to mention all by name. A very special thanks however must go to Georgio Belletini at INFN, Pisa for the loan of the BMW, and having me to "skipper" his yacht on a Mediterranean sailing holiday!

Preface

This thesis contains mostly published papers but some chapters or sections that have not yet been condensed and formatted into papers. In addition, the order that the material is arranged is not the order in which the developments took place and in which the papers were written. To aid understanding the following is a brief summary of the published and unpublished content, and the order in which the main developments took place :

The earliest research was done on ultra-low frequency horizontal pre-isolation and this work resulted in the development of the Scott-Russel pre-isolator stage and the two papers forming chapter 4.

The next area of attention was ultra-low frequency *vertical* pre-isolation which resulted in the development of the Torsion-Crank pre-isolator and the paper forming the majority of chapter 5. Some significant graphical results were removed from the paper to shorten it in order to satisfy a referee and these results are included after the paper.

During this PhD research, the author had three visits to Italy and worked with the Virgo group in Pisa for a total of approximately 12 months. During one of the visits he had built and thoroughly investigated a small "wobbly table" inverted pendulum horizontal pre-isolator. This research aided in the development of Virgo's large 6 meter wobbly table horizontal pre-isolator. On a following visit, the author's main task became to adequately measure the transfer function of the large pre-isolator stage. This ended up being a relatively large effort which is described in the preface to chapter 6. The paper itself, although not actually written by this author, gives a good description of the structure measured, includes the author's results, how he measured them, the center of percussion theory on which the multiple measurements and adjustments were based, and some steps taken in consultation as a result. For this reason this paper is included verbatim as the majority of chapter 6 of this thesis.

Largely as a result of the major effort required to measure adequately and adjust a large pre-isolator such as Virgo's by shake-testing, the author saw the great benefit in having some means of shake-testing and/or adjusting the high performance pre-isolator stage in situ and in vacuum. This need turned into a tilt-rigid "servo-frame" which can be installed *prior* to a high performance pre-isolator and in fact ended up being a low-performance *pre*-pre-isolator. This possibility of applying double pre-isolation (with a servo-frame) prompted an investigation into the effects of tilt (which bypass pre-isolation) and what would be required in order to obtain maximum benefit from double pre-isolation. This research was published as a paper and is placed as chapter 7 because it represents a fairly good summing up of the best that passive isolation can achieve.

The self-damping concepts occurred at about the same time as the servo-frame ideas but did not get beyond modelling until the need arose for a new isolator chain for local research needs. At this time self-damping was ready to be tried and so the author incorporated it into the engineering design for the new isolators. This work has not yet been published (except for brief descriptions in conference proceedings and a web document) and forms the content of chapter 3.

With the need for new isolators it seemed worth trying to improve the vertical suspension in the same design effort. The Euler spring technique to improve vertical isolation was thus the last idea to be developed to satisfy this improvement but seemed so promising that it was completed ahead of the self-damping. It has been written up in the form of two papers most recently and both have been accepted for publication. These papers form the content of chapter 2.

The introduction (chapter 1) as always was written last. Topics originally envisioned as chapters, became briefly covered as sections in the introduction due to lack of time.

The detailed mathematics supporting various sections was separated from the main text for smooth reading and forms chapter 8 at the end. It was kept within the thesis rather than being relegated to an appendix as it is mostly original material.

The appendix only contains the author’s publication list and a paper on local control which is closely related to the thesis topic and is referenced several times within the thesis. This paper was written during the period of PhD research but the experimental and design work was done prior to this.

Table of Contents

1. Introduction

1.1 Vibration Isolation Basic Philosophy

The essence of vibration in this thesis is the undesirable motion produced in a body due to varying forces acting on it. While the term is usually applied to periodic motion in a frequency range from a few hertz through tens of hertz, here it is used of more random motion including frequencies well above and below this range. The primary source of these forces and this motion is from sound waves in the atmosphere and from the continuous background seismic activity in the surface of the earth on which our equipment needs to be mounted. This motion is of order microns and usually goes unnoticed until a very sensitive measurement is made or a strong earthquake occurs nearby.

An ideally isolated body is one which is in a state of free-fall such that there are no forces at all acting on it to disturb its free-fall in any way. For an earth based system this can be well approximated for frequencies above DC by operating in a high vacuum enclosure to isolate from sound waves, and by means of a sophisticated suspension system to isolate from seismic motion - which is the topic of this thesis. Earth's DC gravitational field must be counteracted by a constant force in the vertical direction, and the body can only be free to drift over a very limited distance before it must be acted upon by a variable force to keep it within the enclosure. This distance could be made some tens of centimetres but since seismic motion (in the absence of earthquakes) is only microns, this free drift distance may be limited to a millimetre or less.

In order to constrain the body to an approximate position within the enclosure, a restoring force needs to be applied when the mass moves away from its ideal position. Since the enclosure is continually moving under seismic motion, this means that a varying force will continually be applied to the body as a result of this motion. Reducing the vibrational motion of the body resulting from this force to a minimum is the primary purpose of the vibration isolation and suspension system.

One might imagine a force versus displacement scheme in which almost no force is applied to the body for small offsets from its ideal position, but which is strongly increased for large offsets. The problem with a nonlinear scheme like this is that it up-converts permissible low-frequency motion to unacceptable high-frequency motion.

This is obvious if we take the scheme to its limit by applying the nonlinearity of a brick wall - a slow drift comes to a sudden halt and reverses.

So we are limited to applying a restoring force which is linearly related to displacement and includes a damping effect to prevent oscillating forever. This describes a damped harmonic oscillator and suspension systems necessarily consist of such second order stages often realised with masses and springs or pendulums for horizontal motion. Active techniques may also be applied such as measuring acceleration inertially and applying forces to increase the apparent mass, or measuring the position and applying forces to change the apparent spring-rate, or even simulating entire mass and spring stages with electronic sensing and actuation. Multiple stages are needed in order to achieve the required isolation and they act as a multi-stage low pass filter preventing the higher frequency seismic motion from affecting the isolated body but constraining its position within the enclosure or framework at low frequencies.

1.2 Gravitational Wave Detection

1.2.1 Methods of Detection

The basic method of detecting gravitational waves is to suspend a large elastic body (resonant bar) or a widely separated pair of test masses (interferometer) in such a way that they cannot be acted upon by any well known and mundane forces. The length of the elastic body or the separation between the test masses is monitored by a very sensitive motion sensor to detect any unexpected motion. (There are some forces such as impacts from cosmic ray showers which cannot be isolated from but which are seldom enough to be detected by other means and vetoed). When all mundane disturbances have been prevented or vetoed, any inexplicable sudden motion must be caused by unexplored forces - gravitational waves from astronomical events being the only candidates expected.

In the case of resonant bar detectors (including the modern spherical “bars”) the passage of a gravitational wave causes a momentary stretching and relaxing of the elastic material of the bar, which causes a ringing or a change in the ringing at the resonant frequency of longitudinal stretch of the bar. The ringing motion is detected by a very sensitive motion sensor between one end of the bar and a small tuned spring-mass assembly attached to that end. For various reasons this frequency is usually in the region of several hundred Hz to 1 kHz and these detectors are usually only sensitive to components of the gravitational wave at frequencies close to this resonance. As a result of the high frequency and narrow band of detection sensitivity, these devices are relatively straightforward to isolate from seismic disturbance.

In contrast the laser interferometer detectors are a broadband detector and rely on the isolation systems to prevent the separation between individually isolated test masses several kilometres apart from being disturbed within this bandwidth - which is a much more challenging endeavour. In its simplest conception an interferometer detector uses three freely suspended test masses arranged as a very large Michelson interferometer as shown in figure 1.1. A passing gravitational wave of a suitable polarisation will cause space to be "warped" (or oscillating gravitational forces to appear) such that distances between freely suspended test masses change as indicated by the ring of blocks in the figure. The interferometer is ideally suited to detecting this type of motion and works as follows :-

A powerful laser beam is sent through a beam splitter mounted on a central test mass and travels down two perpendicular arms to be reflected back again from mirrors on test masses at the remote ends of the arms. The difference in path length travelled by the beams in the two arms determines their relative phase when they recombine at the beam splitter and this in turn determines how much of the recombined light exits towards the photo-detector and how much towards the laser. A change in path length difference of one half of a wavelength (i.e. 0.5 μm in several km) is sufficient to redirect all the light from the laser exit to the photo-detector exit and therein lies the sensitivity of the device.

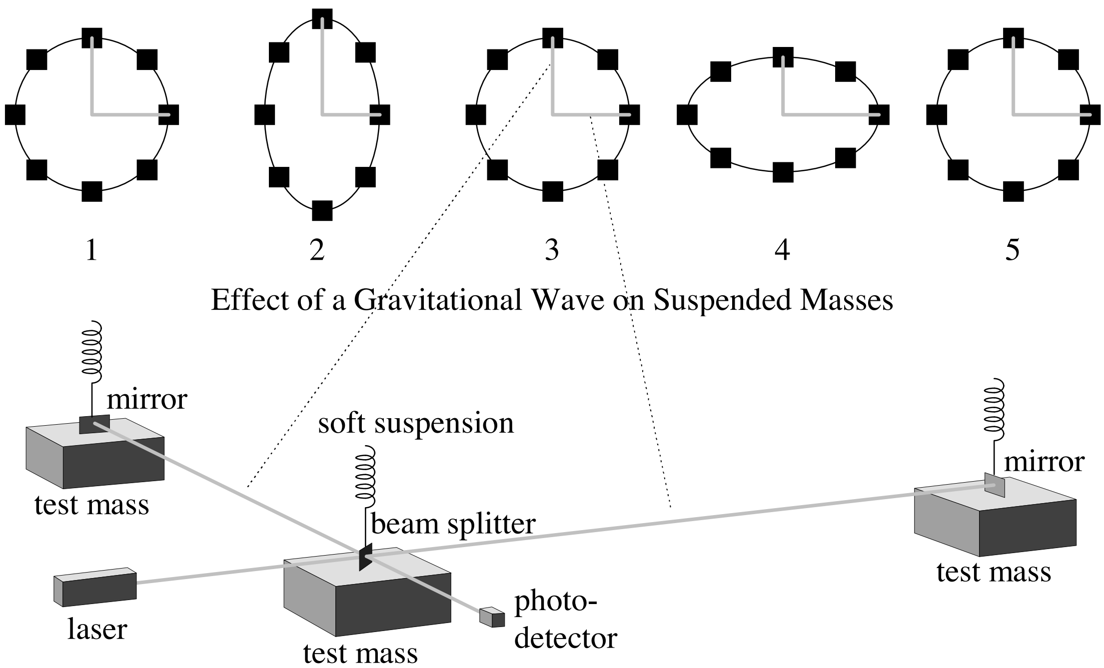

Figure 1.1 Application of a laser Interferometer to gravity wave detection

1.2.2 Sensitivity

A simple Michelson such as described is still not sensitive enough by many orders of magnitude and various schemes are used to increase the sensitivity - which is found to be determined by the light energy stored in the arms. The most common approach is to put additional mirrors that are slightly transparent just after the beam splitter so that the split light passes through (the back of) the extra mirrors and has to reflect many times between them and the end mirrors before it can escape from this cavity to recombine at the beam splitter. This increases the sensitivity of the basic instrument by the average number of times the light re-travels the arms. However with this arrangement the laser light only resonates in the arm cavities for round trip cavity lengths which are within a very small fraction (1/re-travel times) of an exact integer number of wavelengths. This poses a tough problem for the control system to initially align and obtain lock in the presence of even very small amounts of residual low-frequency seismic induced motion.

There are also additional schemes which may be employed to recycle and thus increase the laser and signal power which do not impact the suspension arrangements and so are beyond the scope of this simple introduction. The bottom line is that in order to stand a reasonable chance of detecting gravitational radiation from predicted sources within the lifetime of an experiment, motion sensitivities of order 10^{-20} m/√Hz are required. At this sensitivity many fundamental sources of noise dominate the spectrum,

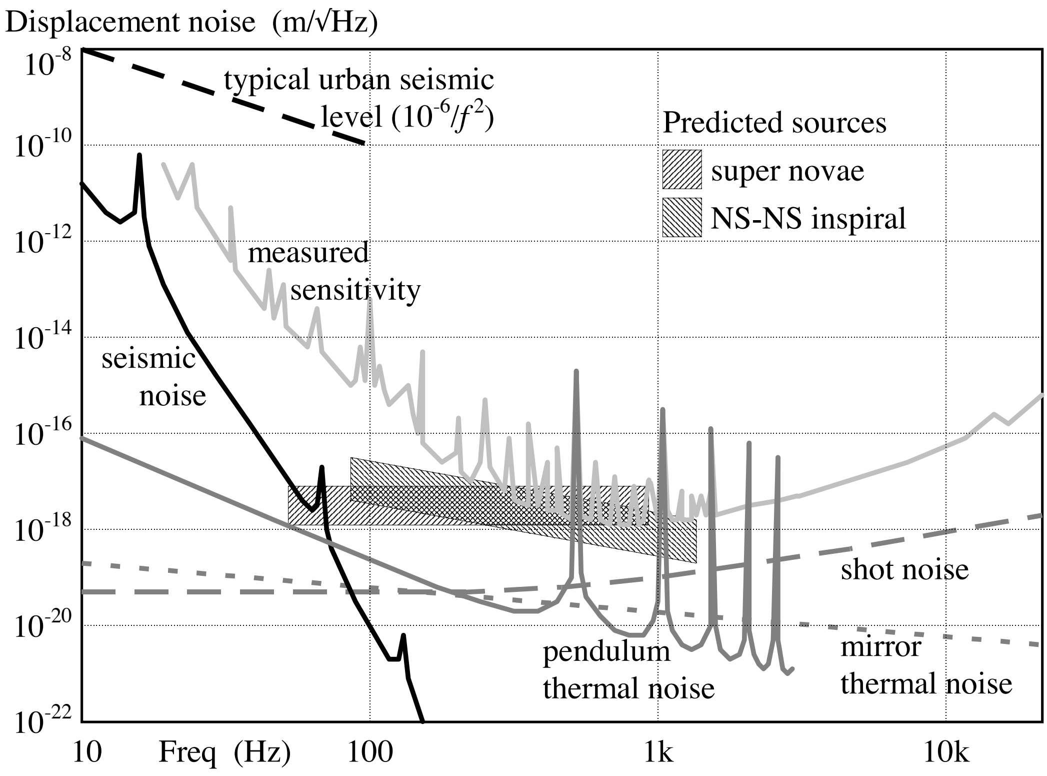

Figure 1.2 TAMA 300 noise floors and sensitivity attained in September 2000.

most of which are indicated in figure 1.2.

These curves are taken from web published TAMA 300 (Mitaka, Japan) interferometer data which is the only interferometer functional enough to intermittently collect data at the time of writing (September 2000) although there are four others under construction (LIGO-Hanford & LIGO-Livingston, USA; VIRGO Cascina, Italy; GEO600 Hannover, Germany) and one or two planned or waiting for funding (AIGO Gingin, Australia; LCGT Kamioka mine, Japan).

1.2.3 Seismic Noise

The smooth looking curves in figure 1.2 are theoretical levels which should be achieved when all the sub-systems of the interferometer are working at their best possible performance. It can be seen that the sensitivity of the instrument is bounded by the seismic wall on the left, final pendulum thermal noise below, and shot noise from the laser light at higher frequencies. Also superimposed are expected levels of gravitational wave signals from sources within our galaxy (binary neutron star in-spirals and super novae explosions). The measured sensitivity of the interferometer for September 2000 is also given and this indicates that there are still many sources of excess noise preventing the sensitivity from reaching these ultimate floor levels. The only noise source that this thesis is concerned with is the seismic noise which remains after vibration isolation and forms the low frequency wall on the left. A typical urban seismic level of $10^{-6}/f^{2}$ m/√Hz is also shown indicating that at a few tens of Hz the vibration isolation needs to reduce this seismic motion by approx 10 orders of magnitude (200 dB). This may seem difficult at first glance but is easily achieved at tens of Hz by cascading several isolation stages together. In fact with a simple 4 stage isolation chain of total height 2.5 m described in the next section, this level of isolation is achieved by about 16 Hz (see figure 1.5a). A major thrust of this thesis however is reducing the residual motion at frequencies well below those at which the chain isolates effectively.

Figure 1.3 shows the level of seismic noise measured in our Gingin laboratory on the cruciform concrete slab that the isolation systems are to be mounted on. These readings were measured with SM-6 geophones which have a 4.5 Hz natural frequency. They were 20 minute averages taken one after the other on a quiet evening in September 2000. The noise floor was measured with a 375 ohm resistor replacing the geophone and indicates that only the values above ~2 Hz are valid readings - the values below 2

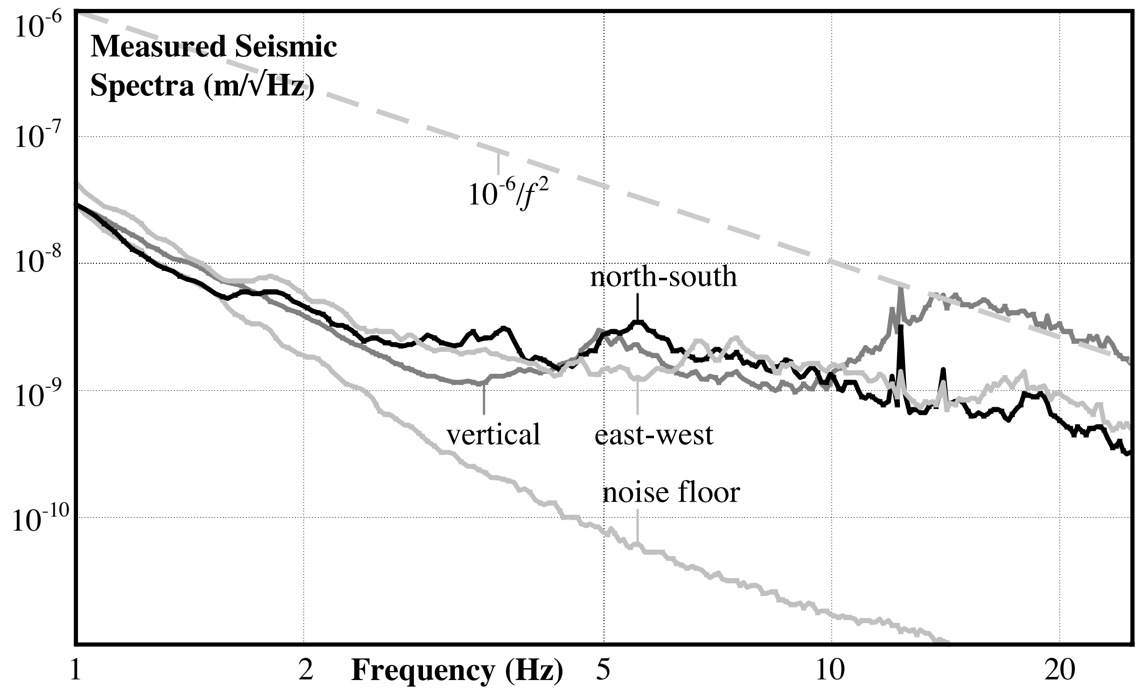

Figure 1.3 Seismic noise levels measured at the Gingin laboratory September 2000.

Hz are only Johnson noise in the resistance or input noise of the analyser. The standard seismic level of $10^{-6}/f^2$ m/√Hz which will be assumed throughout the rest of this thesis is included for reference. It is apparent that the site is almost two orders of magnitude quieter than this level at around 2 Hz so this provides a large margin for additional noise from personnel walking around, stormy days, etc.

1.2.4 Summary

1) The test masses in an interferometric gravitational wave detector require isolation from seismic motion to an unprecedented degree. Approximately 200 dB of isolation is required over the entire signal detection bandwidth from as low as is practical (a few hertz), through a few kilohertz.
2) The residual test mass motion below the detection bandwidth should be reduced as much as possible to ease control system design and lock acquisition.
3) Provision must be included to allow millimetres of slow, smooth translation of the isolation structure in order to maintain constant test mass separations against daily temperature expansion and contraction of the framework mounted on the earths crust.

Requirement (1) has been well met for frequencies above a few tens of Hz by the work of many previous researchers. The improvements presented in this thesis, of a novel vertical suspension technique, chain self-damping, ultra-low frequency pre-isolation, and tilt stabilisation, extend the detection band down to ~10 Hz while minimising the number of stages required and the total height of the suspension system.

These techniques meet requirement (2) by also reducing residual motion below that of a simple undamped suspension chain by five orders of magnitude – to the nanometre level – rendering control system design trivial by comparison.

Requirement (3) is inherently met by an ultra-low frequency horizontal pre-isolation stage, provided it is designed with millimetres of dynamic range. The structures presented in this thesis offer dual pre-isolation – a tilt-rigid 3-D translation stage with ±10 mm of motion in all directions, followed by a higher performance pre-isolation kept at its optimal operating point.

1.3 Standard Passive Suspension Chain Overview

1.3.1 Design Philosophy

The target of the isolation system in an interferometer gravitational wave detector is to suspend the mirrors and beam splitter in such a manner that the seismic disturbance in the detection band is either as low as possible, or lower than other fundamental sources of displacement noise at the same frequency. Clearly it is only motion in the one degree of freedom measured by the laser beam which counts, but since it is impossible in practice to avoid some degree of cross-coupling, all other degrees of freedom must be isolated to within two or three orders of magnitude of the main laser sensed degree of freedom.

The fundamental noise source which should be the noise floor at the low frequency end of the detection band is thermal noise (the random Brownian motion of the atoms) due to the finite temperature of the material that the structure is made of. At the lowest frequencies this shows up as a disturbance in the steady hang of the last pendulum even if that pendulum could be hung from a perfectly stationary support. The random thermal motion of the atoms making up the wire cause it to jiggle and wave around (to a very small degree). This effect may be reduced by lowering temperatures to cryogenic levels, but not without overcoming major engineering challenges. Nevertheless, this is being seriously investigated by some groups [Kuroda 1999, DeSalvo 2001].

The main method of minimising the random effect of the thermal noise is by the use of material with very high quality factor or low intrinsic loss. This is because just as the loss mechanism allows mechanical motion to be converted to thermal energy, so the same mechanism allows the reverse conversion of thermal energy to mechanical motion (fluctuation-dissipation theorem). Using high Q-factor materials allows the majority of the energy to appear in very narrow frequency bands (of the normal mode resonances) rather than being spread across the band in the “skirts” of the resonances. The normal mode resonances would preferably be placed outside of the detection band leaving it clear and at the residual skirt level. However this is not necessary because even if the resonances fall within the detection band, their very high Q-factor makes the normal mode motion extremely predictable and thus in principle removable.

The target of the vibration isolation system is to reduce the seismic disturbance reaching the test mass to a value below this thermal noise level at as low a frequency as possible. We should briefly consider all the six degrees of freedom to see which need the most attention. Generally in this thesis z is taken as the vertical axis and parallel to gravity, x is horizontal and typically parallel to the laser beam, while y is the remaining horizontal axis perpendicular to the beam. There is also rotation to be considered and these are called θ_x, θ_y and θ_z for rotation about each of the three translational axes.

1.3.2 Isolating All Six Degrees of Freedom

Rotational disturbance about a vertical axis (i.e. θ_z noise) is easiest to deal with in systems which have a single or near single wire suspension chain. This is because the torsional resonant frequencies and normal modes are extremely low and provide very good filtering very easily. This is evident from a torsion pendulum's use in many extremely sensitive experiments such as the Cavendish torsion balance. In systems with multiple widely spaced suspension wires, then the θ_z torsional modes and isolation simply become the differential case of the horizontal pendulum isolation. However rotational seismic motion (i.e. differential-mode horizontal) is much smaller at typical isolator stage lever-arm lengths than the common-mode horizontal motion.

Vertical disturbance (z) is also quite straight-forward to deal with using springs and masses although due to the gravitational potential, there is a relatively large amount of energy (mgh) which has to be stored and retrieved from the spring as the support vibrates up and down (if the suspended masses are not to follow the motion). This aspect is dealt with in great detail in the paper of section 2.2 and will not be repeated here. It is sufficient to note that in principle there is no real limit to how low the resonant frequency of a vertical stage may be made. With reasonable engineering it is straight forward to obtain lower vertical resonant frequencies than horizontal, and it is pointless to provide better isolation in the vertical than can be done in the horizontal.

Tilting disturbance (θ_x and θ_y) propagating down the suspension chain is harder to consider because it is inextricably coupled to the horizontal pendulum motion. However if one conceives of the suspension chain with the pendulum heights reduced to zero and the upper and lower wire bending attachment points replaced with (concentric) pivots of equivalent angular spring-rate, then it becomes obvious that these torsional spring-rates together with the stage's moment of inertia will give much lower resonant frequencies than the typical horizontal pendulums. If multiple widely spaced suspension wires are used, then the tilt suspension becomes the differential case of the vertical suspension. Again the tilting (differential-mode vertical) seismic motion is much smaller at typical isolator stage lever-arm lengths than the common-mode vertical motion. The height and gravity effects which couple the tilting to horizontal motion are dealt with by the horizontal filtering.

Disturbance from horizontal seismic motion forms the main limit to the isolation that can be obtained at the low frequency end of the spectrum. In order to provide horizontal isolation, multiple stages are inevitably pendulums effectively suspended from a single point on the stage above. Even if multiple widely spaced wires are used then each one must act as a pendulum of the same length working in parallel. Any other scheme produces a tilting moment on the stage above, which if the stage is soft to tilt (which it must be for isolating tilt) is unstable. It is possible to have one or even two stages which achieve lower than pendulum frequency, but they are only stable if their supporting stage is rigid to tilt (so that it cannot tilt appreciably when a moment is applied). Clearly this cannot be the case for the majority of the suspension chain as then there would be no tilt isolation. This is more fully described in the introduction to horizontal *pre*-isolation (section 1.5.1).

1.3.3 Cascading Isolation Stages

The fundamentals of basic vertical and horizontal isolation are adequately dealt with in the introduction of the paper of section 2.2 and will not be repeated here. However the conditions arising from cascading multiple isolation stages one after the other to achieve better isolation than a single stage, is reasonably complex and interesting and deserves some attention. One might naively imagine that hanging a second pendulum from a first would have the same effect on the vibration from the first, that the first had on the vibration from its mounting. This would be the case if the second pendulum has negligible mass in comparison to the first as then the first has no knowledge that the second pendulum is present and no reason to behave differently when it is added.

When the mass of the second pendulum is appreciable in comparison with the first, then it increases the tension in the suspension fibre of the first pendulum by a significant amount thereby increasing its restoring force for any given offset. This effectively makes its resonant frequency higher and makes it a less effective isolator stage. (One could observe the same effect by pulling on the pendulum mass with a long thin string in a constant direction (i.e. from infinity), or with a constant vertical magnetic field. Its mass is not altered but the spring-rate of its restoring force is increased.) For n cascaded pendulums of equal mass this effect makes the isolation worse at high frequencies by a factor of $n!$ (factorial) than it would be without this interaction (cf equations (8.7) & (8.8) in section 8.3).

The effect of this is illustrated in figure 1.4 for a 4 stage chain. The black trace is the transfer function of an equal height, equal mass pendulum chain with total height 2.5 m (and Q-factor of individual stages being 100). There are no resonances in this trace at the 0.63 Hz that each individual pendulum would show, but instead there is a distribution of resonant modes at other frequencies. The dashed trace shows what happens when the mass ratio between the stages approaches zero so that they do not affect each other. In this case the four resonant mode frequencies gather together and all approach the 0.63 Hz for a single (2.5/4) m pendulum. In the limit of the mass ratio becoming zero, they would all overlap and the total transfer function would simply be the transfer function of a single stage to the fourth power. The $n!$ difference between these two traces is also indicated in figure 1.4.

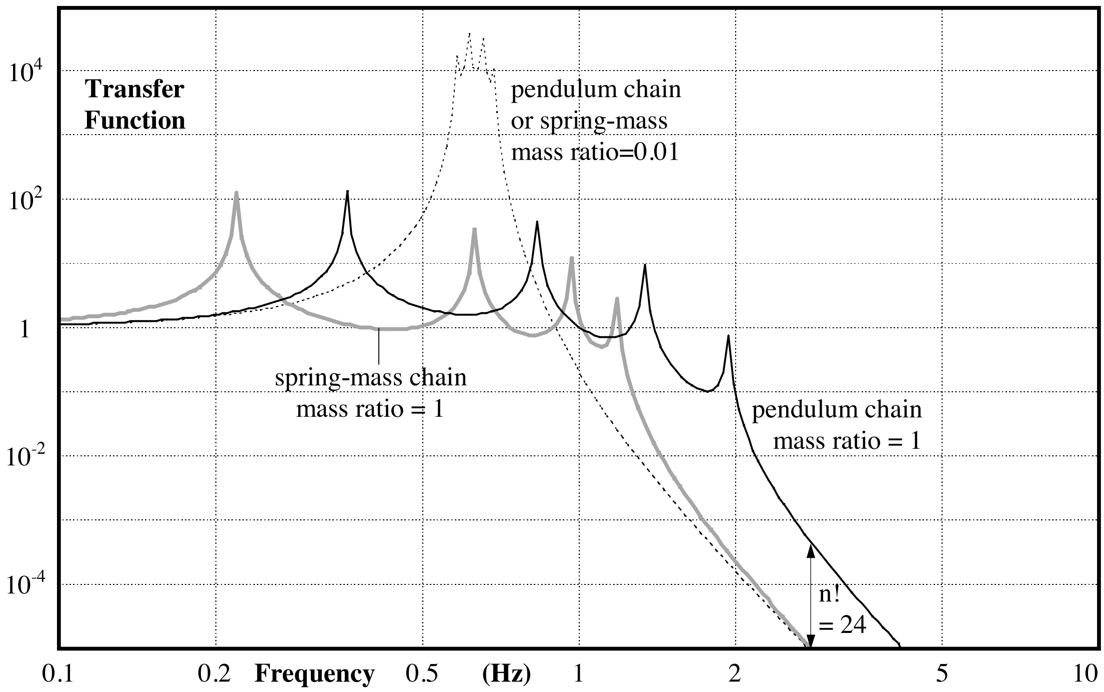

Figure 1.4 The effect of interactions and gravity between isolation stages (n = 4 stages, total height 2.5 m, equal height pendulums, spring-mass stages equal in resonant frequency to pendulums).

Another case to consider are cascaded isolator stages made of springs and masses such as might be used for vertical isolation. Supposing the resonant frequencies of the individual spring-mass stages are made equal to each other and equal to the previously discussed pendulum stages, then the transfer function obtained is the grey trace in figure 1.4. Once again there are no resonances at the 0.63 Hz of a single stage, but the interactions between neighbouring masses causes a distribution of resonances around the value for a single stage.

The reason for the different resonant frequencies is easiest to understand when there are only two masses. In this case there are only two resonant modes - one in which the masses move in unison, and one in which they move in contra-motion. Considering the mass at the end of the chain - the contra-motion mode is higher in frequency because the spring attaching this mass is being compressed and expanded at a higher rate than just the motion of this mass would normally cause. This makes the spring look stiffer and gives a higher frequency than it would be stand-alone. The other mode is lower in frequency for the opposite reason. For more than a few masses and springs of differing values it is difficult to guess the amplitudes and phase of motion for each mass, but they may be calculated using state space matrix methods (as is described in section 8.1). Often it is just the lowest frequency mode or the highest one that is of interest and these are straightforward. The lowest mode is where all masses move back and forth in unison, the lower stages moving further than the upper stages. The highest mode is where the motion alternates with each mass so that each is moving in the opposite direction to its neighbours at any given instant. If the mass of successive stages in a spring-mass system are made much smaller than the ones they are attached to, then once again this interaction disappears and the dashed trace of figure 1.4 is obtained.

If we consider hanging a spring-mass isolation stage in the presence of gravity, we find for a linear spring (which stretches in length proportional to the mass that it is loaded with), that its resonant frequency has the same relationship to its extension under load as the frequency of a pendulum has to its length - which is $(g/l)^{1/2}$. Now if we cascade such stages, it becomes apparent that the ones at the top will stretch a lot further than the ones at the bottom because they have to support more total mass. If we now strengthen the springs at the top in direct proportion to the amount of mass that they have to support so that they all stretch the same amount, then we have an identical situation to the pendulum case described in the second paragraph of this section where the spring-rates at the top are much stiffer than lower down. Indeed the transfer function for such an equal extension spring-mass chain is identical with that of a pendulum chain with length equal to the extension (being $n!$ worse at high frequency than before the springs were strengthened).

1.3.4 Maximising Low Frequency Bandwidth

Figure 1.5(a) shows the transfer functions obtained with a total height of 2.5 m available, and with 1, 2, 3, or 4 equal height pendulums or spring-mass stages cascaded from each other to take up that height. The number of resonances in each case is equal to the number of suspended masses (in a free fall situation one would count the earth as one of the masses and so it would be one less than the total number of interconnected masses).

It is apparent from figure 1.5(a) that as the number of stages is increased, the isolation at high frequencies improves dramatically. However the corner frequency of the highest resonance moves up with each increase in the number of stages so that the frequency at which the transfer function goes below a given threshold (the thermal noise level for instance) does not improve at a great rate after the first few stages. In fact if we add the condition that the total height of all the stages must stay below some limit, then this frequency reaches a minimum for some optimum number of stages after which is gets worse again (see derivation in section 8.3). This minimum occurs at 12 stages for the equal mass case and for an isolation of 200 dB (from section 1.2.4) but it exists for

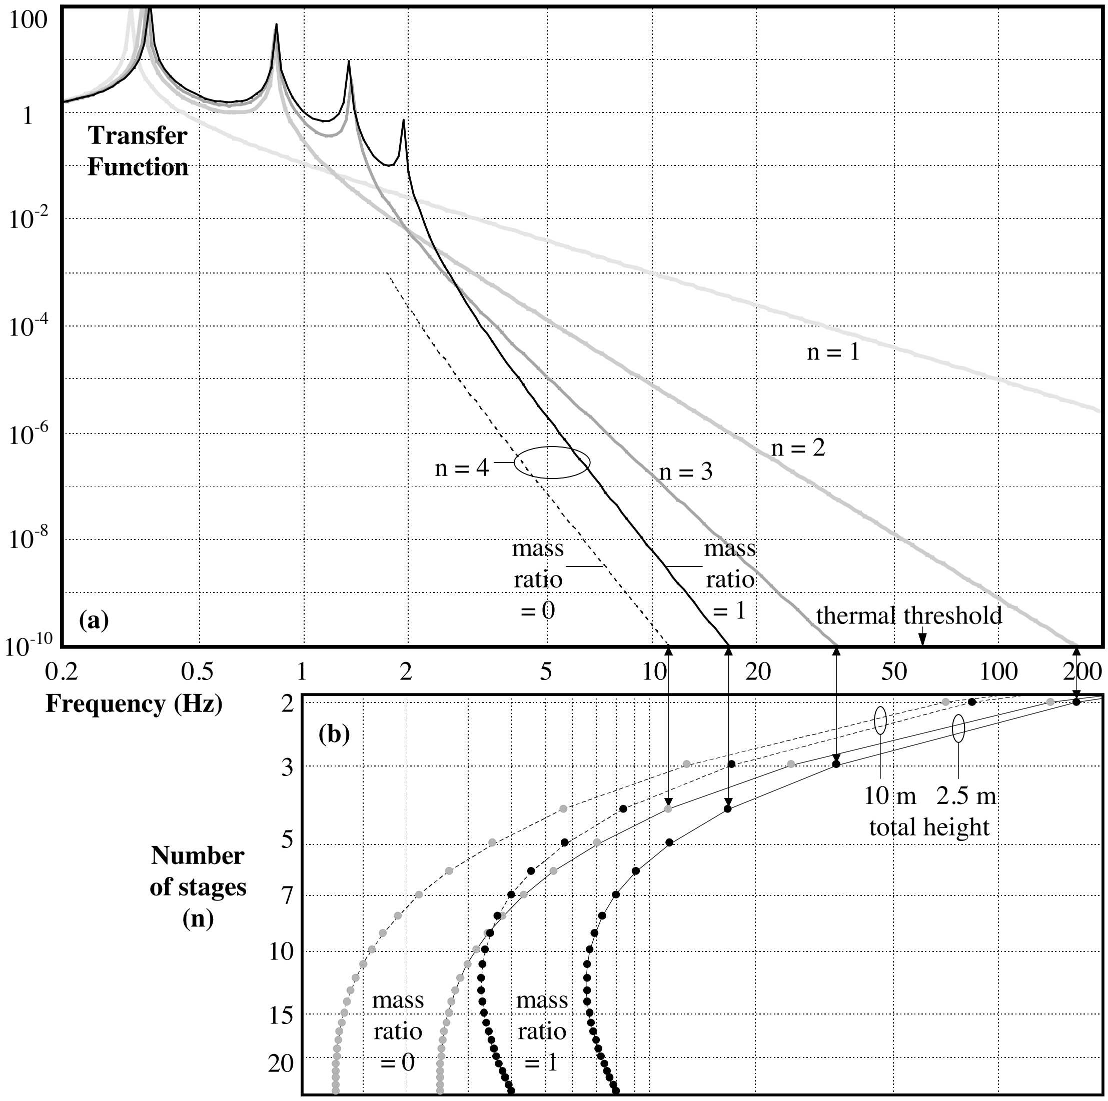

Figure 1.5 The effect of cascading many isolation stages.

all mass ratios and occurs (at 23 stages) for the impossible case of a mass ratio of zero - when the stages no longer interact.

This effect is shown in figure 1.5(b) which has two sets of points, one for a 2.5 m height limit (e.g. Aigo) and one for a 10 m limit (e.g. Virgo). It can be seen that the minimum occurs for 12 stages in both cases and the minimum reached goes as the square root of the available height as one might expect (6.6 Hz for 2.5 m and 3.3 Hz for 10 m). The minimum is very weak and practical issues of cost and complexity dictate that the actual number of stages used be typically considerably fewer than the position of this minimum might suggest. However it does illustrate the important point that this type of passive isolation is fundamentally limited by the height available for the chain and the fact that such a minimum exists seems interesting and initially surprising.

This optimum number of stages and minimum isolation frequency is plotted for a large range of isolation threshold requirements and several chain heights in figure 1.6. The minimums mentioned above appear on the 2.5 m and 10 m lines for the isolation requirement of 10^{-10} (i.e. 200 dB of isolation). However supposing one had a significantly reduced isolation requirement so that instead of 10^{-10} one only needed say 10^{-3}, then the graph indicates that 4 stages is the optimum number of equal mass stages to reach that isolation level while obtaining the lowest possible isolation frequency for the height. Supposing there was 1 m of height available, then one could expect to obtain 10^{-3} of isolation from a frequency of 3.5 Hz, (or if 10 m available then from 1.1 Hz).

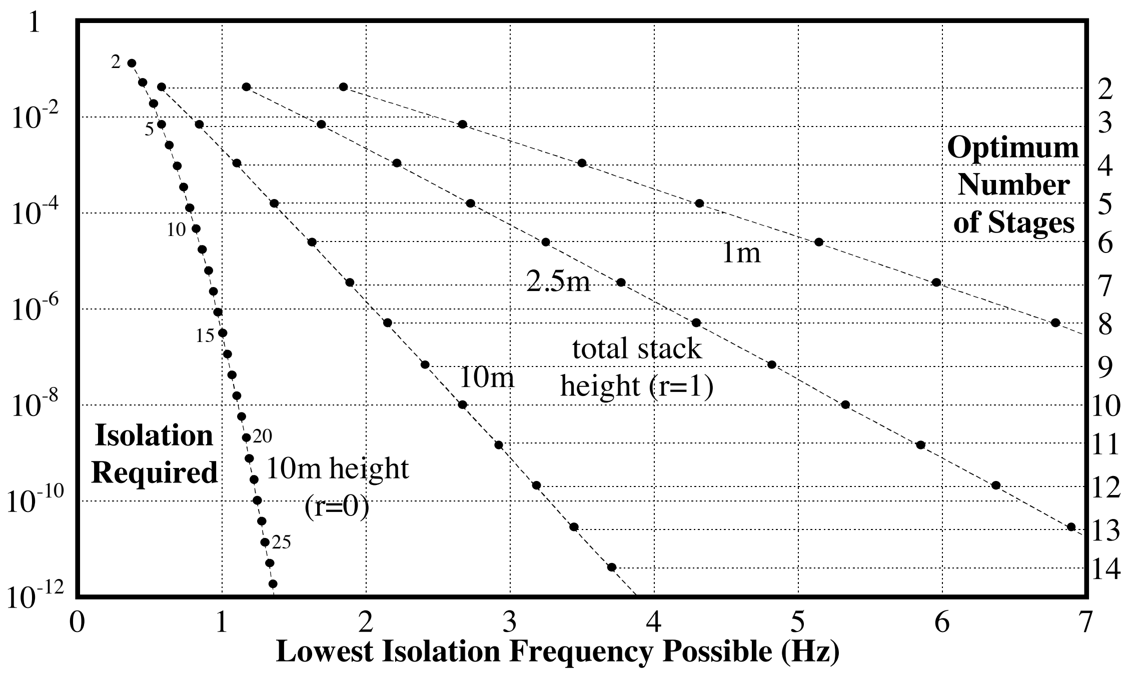

Figure 1.6 The optimum number of stages to obtain a frequency minimum for a range of isolation requirements.

The only way to do better than these solutions indicate with a given height constraint is firstly to use ultra-low frequency pre-isolation (see section 1.5) to reduce the horizontal feedthrough to the same level as the tilt noise effect at the top of the chain, and after that to use active isolation techniques. The lowest residual motion obtainable passively is given by the double integral of g times the tilt seismic motion (see 1.5.2.5 description of perfect pre-isolation) filtered by the optimum number of isolator stages as described above. To exceed this within a given height requires active isolation.

Active isolation allows masses to have more apparent inertia (in that they are harder to accelerate), while not weighing any more. This is achieved by sensing acceleration with respect to an inertial reference frame and feeding back some force to resist it. The mass appears to be increased by the gain of the feedback loop but does not increase the tension in the suspension wires. This allows the resonant frequency of a pendulum to be considerably lower than its height would normally dictate. Active isolation is not discussed in this thesis except for the case of tilt servoing in section 7.3 discussing double pre-isolation.

1.3.5 Summary

1) The purpose of the isolation chain is to prevent seismic motion from affecting the mass over the full detection bandwidth. Its performance at low frequencies directly determines the low limit of the detection bandwidth - which starts at about the frequency where the chain has reduced the seismic disturbance below the thermal noise level. Its performance below that level is immaterial.
2) The chain must be constructed in such a way as not to generate any excess noise from creep, friction, etc and to be compatible with a high vacuum.
3) Low frequency horizontal motion is fundamentally the most difficult to isolate and performance goes as the square root of the height available for cascading pendulums.
4) Without using height, a once-only advantage may be obtained by an ultra-low frequency pre-isolation stage (section 1.5). Adequate pre-isolation changes the primal seismic disturbance from translation to tilt and can provide an improvement of up to 2 orders of magnitude. Beyond this active isolation is required.
5) If the vertical springs are linear and the stages are designed to have equal displacements under load, then their performance is in principle identical to a pendulum chain of the length of their displacements.

6) It is sufficient to isolate the vertical to within 100 times worse than the horizontal. It is relatively easy to obtain low resonant frequencies for all other degrees of freedom.

1.4 Suspension Chain Damping

1.4.1 The Need for Damping

It is apparent from simply considering a well-constructed chain of pendulums that any disturbance given to the chain will take a very long time to die down. It may also be deduced from the transfer functions of the previous section that any seismic spectrum at the resonant frequencies of the normal modes of the chain will be resonantly enhanced or amplified far above their normal micron level. (The peak motion per root hertz amplifies as the Q-factor but the RMS motion only as the square root of the Q-factor because the motion becomes more narrow band). Since there will be an optical mirror at the end of the suspension chain, that has to be perfectly aligned with another at the distant end of a long vacuum tube and eventually fringe locked to a small fraction of a wavelength, this will obviously be a major source of problems.

There are two scenarios here to consider. Firstly there is the case of damping, and slowing or "cooling" the motion of the test mass mirror so that lock can be established. This may be done simply by acting on or near the test mass itself as it is at the end of the chain and will typically move the most. Secondly however, once lock has been achieved and forces are applied at the test mass essentially to fix its position in inertial space (or at least w.r.t. the remote test mass), then in this situation *no more local damping can be done by acting on the test mass*. It has become effectively rigidly connected to an infinite mass by the control system. In this condition there will be one less normal mode available to the chain and all of the normal modes will shift in frequency. (compare fig 3 with fig 4 in [Winterflood 1995]). The forces which now occur that the control system has to counteract are *only limited by the remaining Q-factor of the chain*. If there is no other damping being done then these forces can be orders of magnitude higher than they would be if the chain was damped. At worst the energy building up in the new undamped modes due to seismic excitation may overcome the locking force available and result in loss of lock. At best it will mean higher than necessary electronic and control system noise is injected along with the forces applied to keep the system locked. Damping of the chain, separately from the mirror control is clearly an area that needs attention for sensitive detection. Briefly we should look at what has been done in the field previously, and then consider more closely some aspects of some of the possible passive techniques available for simple vertical and horizontal mass suspensions.

1.4.2 Suspension Damping Approaches

Various techniques for applying damping to the high Q-factor modes have been used in the past and some are as follows :

1) A common method and one that we have used to damp the high Q-factor resonances in a pendulum style isolation chain, is to sense the low frequency motion with respect to the ground and actively apply low-pass filtered viscous forces to damp the sensed motion [Winterflood 1995]. It has proved difficult to avoid injecting excess noise using this method, and in principle this method *prevents* a reduction of the low frequency motion below the seismic level (due to seismic motion in the sensing). So it is totally inappropriate for use on pre-isolated systems where the disturbance at the top of the chain is already orders of magnitude below seismic. Besides it shouldn't be necessary to use a complex active control system to generate something as simple and stable as viscous friction!

2) A passive method that we [Saxey 1995] and others [Tsubono 1993] have used is to have a second parallel chain (or at least a suspended reaction mass) with different normal mode frequencies, and viscously couple the motion of the parallel systems together so that they damp each other (using magnetic eddy current coupling). However this method increases the suspended mass and the mechanical complexity of a system.

3) Another method might be to simply make the Q-factor of each stage very low by incorporating viscous damping. However normally applied viscous damping bypasses the isolation to a degree only providing $1/f$ isolation per stage (see next section 1.4.3) and so this method can double the number of stages (and hence complexity) required in order to achieve the same level of isolation. It is also far from easy to find high-vacuum compatible techniques for viscous damping. In section 1.4.4 we look at using viscous damping to simulate structural damping to retain the $1/f^2$ isolation per stage.

4) Various vacuum compatible methods of using non-vacuum compatible materials such as filling bellows with graphite loaded rubber [Plissi 1998] or making coil springs from tubes and filling the tubes with damping material (Ligo).

5) A high performance damping metal (alloy M2052) has recently become available which may be used for suspension wires and springs (lowest material Q-factor reputedly below 5 [Mio 2001]). This presumably will give structural damping performance.

6) Chapter 3 of this thesis presents a relatively new method of damping the Q-factors of these modes that we have termed “self-damping”. It is closely analogous to the dynamic vibration absorber commonly used to control vibration in industry but needs no additional reaction masses and instead couples between different degrees of freedom of the same mass. The standard vibration absorber has been considered for suspension chain damping before [Ju 1995], but not apparently for the horizontal pendulum regime where damping is most needed. Vibration absorber theory and how it relates to self-damping will be looked at in section 1.4.5.

1.4.3 Simple Viscous Damping

The most obvious method of reducing the high Q-factor modes of a spring-mass or pendulum suspension system is to add viscous damping as shown in figure 1.7. The dashpots shown have the characteristic that they generate a resisting frictional force proportional to the rate at which they are compressed or expanded. The spring-mass

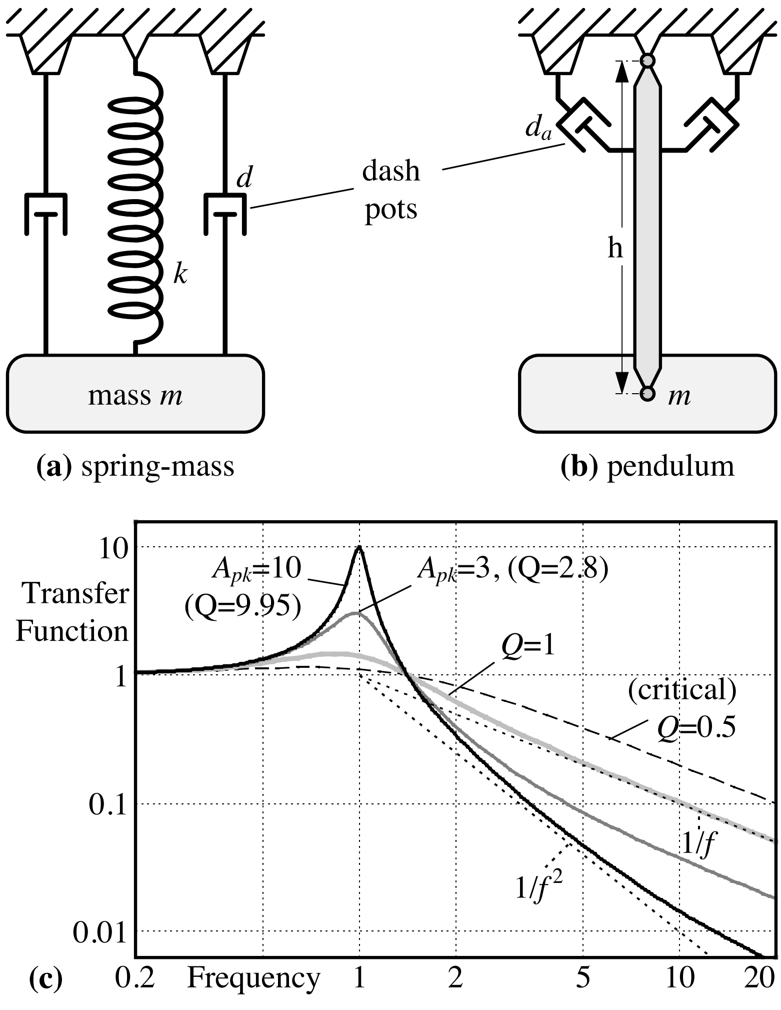

Figure 1.7 (a) Simple spring-mass oscillator with viscous damping, (b) Equivalent viscously damped pendulum. (c) Transfer functions of these systems with various damping.

linear damping coefficient has units of newtons/(meter/second) and the angular unit (for the pendulum pivot) has units of newton·meters/(radian/sec). One useful measure of the amount of damping present is to note its effect on a resonant oscillation of the system and express it as a fraction of the amount of damping which would make the oscillation die away in the minimum time (critically damped). This measure is referred to as the damping ratio. Another closely related measure which we will use in preference to the damping ratio in this thesis is the quality factor or Q-factor of the oscillation (which is defined as the peak energy present in the oscillation divided by the energy lost per radian). The two are simply related in that the Q-factor is the reciprocal of twice the damping factor. (An oscillation is critically damped when it has a Q-factor of 0.5). If the Q-factor of an oscillation is greater than a few units it may be approximated by counting the number of cycles required for the amplitude to fall to 1/3rd (actually $1/e$) and multiplying this value by 3 (actually π). The Q-factor is also approximately given by the amplitude enhancement of a simple resonant peak (actually $(1+Q^2)^{1/2}$ see table 1.1).

In figure 1.7(c) the isolation performance (or transfer function of the movement of the suspended mass m as a fraction of the movement of the support) for several different values of Q-factor are plotted . It can be seen that at low Q-factors with good damping of the peak, the isolation performance is significantly worse than that obtained with a high Q-factor. This is because the dashpots effectively bypass the spring or pendulum, applying a force directly to the mass and this force is proportional to velocity and thus to frequency. The dotted lines illustrate the general rule that the isolation performance changes from improving as $1/f^2$ to only $1/f$ at a frequency which is the Q-factor times the resonant frequency. The performance of these oscillators may be characterised as follows :

Physical Values		Units
m	suspended mass	(kg)
k	spring-rate of suspension spring	(N/m)
l	length extension of spring under load of m ($=k/(mg)$)	(m)
d	total dissipation coefficient of linear dashpots	(N/(m/s))
h	height or length of pendulum	(m)
I_a	pendulum moment of inertia about pivot ($=mh^2$)	(kg·m^2)
k_a	angular spring-rate of pendulum ($=mgh$)	(N·m/rad)
d_a	total dissipation coefficient of angular dashpots	(N·m/(rad/s))

Oscillation characteristics	Spring-Mass Suspension	Pendulum Suspension
Resonant frequency	$2\pi f_0 = \sqrt{k/m} = \sqrt{g/l}$	$2\pi f_0 = \sqrt{k_a/I_a} = \sqrt{g/h}$
Critically damped when	$d = 2\sqrt{k\,m}$	$d_a = 2\sqrt{k_a\,I_a} = 2mh\sqrt{g\,h}$
Quality factor of oscillation	$Q = \sqrt{k\,m}\big/d$	$Q = \sqrt{k_a\,I_a}\big/d_a = m\,h\sqrt{g\,h}\big/d_a$
Peak motion enhancement	$A_{pk} = \sqrt{1+Q^2}$	$A_{pk} = \sqrt{1+Q^2}$

Table 1.1 Characteristics of simple viscous damped oscillators.

It may be seen from the equations that there is an exact correspondence between the systems when the angular equivalents (I_a, k_a, d_a) of the linear values (m, k, d) are used.

It is easy to see how this damping technique could be applied to a vertical suspension chain. The pendulum version is almost as straightforward provided the pendulum links are made rigid rather than flexible fibres. With rigid links, the angular dashpots can be applied at each pivot in the chain.

1.4.4 Simulated Structural Damping

It has long been known that the type of damping most common in spring bending and structural flexing in nature is not proportional to velocity and frequency as is the viscous damping of the previous section but instead exhibits damping which is independent of frequency [Kimball 1927]. Its frictional force has the characteristic that for any given frequency, it is some constant fraction of the spring restoring force and leads the restoring force in phase by 90°. It is well modelled in the frequency domain by a complex spring-rate $k(1+i\phi)$, where the imaginary component is some approximately constant fraction ϕ of the restoring spring-rate k, and ϕ is commonly called the loss tangent. In the absence of other damping forces, the loss tangent ϕ has a simple reciprocal relationship to the Q-factor of an oscillation such that $Q=1/\phi$.

One remarkable aspect of this type of damping - if it could be applied to the systems of figure 1.7 - is that the isolation performance continues to improve with $1/f^2$ regardless of how low the Q-factor is made (within reason). This should allow very good damping of the resonant peaks to be achieved without sacrificing isolation performance. However it has been difficult to find materials with high structural losses and it is also difficult to enhance those losses to obtain damping of the same order as can easily be obtained by viscous techniques. (A high loss metal has become available recently [Mio 2001] and losses can be enhanced by engineering positive and negative spring-rates to null each other out in which case the losses sum).

It has been shown [Quinn 1992, Saulson 1994] that this damping performance can be duplicated by a conventional viscously damped structure consisting of one main spring in parallel with an array of spring-dashpot pairs (termed Maxwell units). The springs of the pairs have all the same small spring-rate, while their dashpot coefficients are related by a constant ratio (as a geometric series) so that they come into action at equal spaced frequencies on a logarithmic scale. This structure is capable of duplicating the frequency independent damping over as wide a frequency range as is desired simply by having enough Maxwell units. In our case we should like to have this damping performance only in the region of the resonant frequency of the oscillator and would like it to work as an ideal spring at high frequencies.

This is possible to achieve with just one Maxwell unit* and an equivalent structure is shown in figure 1.8(a). Instead of putting the Maxwell unit (spring in series with a dashpot) in parallel with the main spring, we have replaced the resulting structure with its equivalent consisting of two springs in series (together forming K) with a dashpot d in parallel with one of them. At frequencies near resonance, the dashpot should be strongly acting as it did for simple viscous damping, while at high frequencies the dashpot should effectively lock up with only the stiffer spring k remaining active (and giving $1/f^2$ isolation performance through high frequencies).

A direct pendulum equivalent of this spring-mass structure is shown in figure 1.8(b). At frequencies near resonance both sections H of the pendulum work together with the upper section working stiffly and providing damping, while at high frequencies the upper section of pendulum becomes locked to the support and only the shorter pendulum h remains active.

The additional complexity in these oscillators means that there are two parameters to be adjusted : (1) the viscous damping coefficient (d or d_a) which was present in the simpler oscillators and (2) the fraction of the spring or pendulum length μ which remains undamped by the dashpot. It is clear that if the dashpot damping is either zero

---

* Although being Maxwell's model of an anelastic solid, this particular damping structure has not been noticed elsewhere as a proposed damping method. This is probably because most damping applications do not care for $1/f^2$ isolation performance. However it is probably a close analogy (or a *dual*) of the springless "Lanchester damper" [Den Hartog 1956] which has a similar level of complexity but has an extra mass and dashpot instead of an extra spring and dashpot.

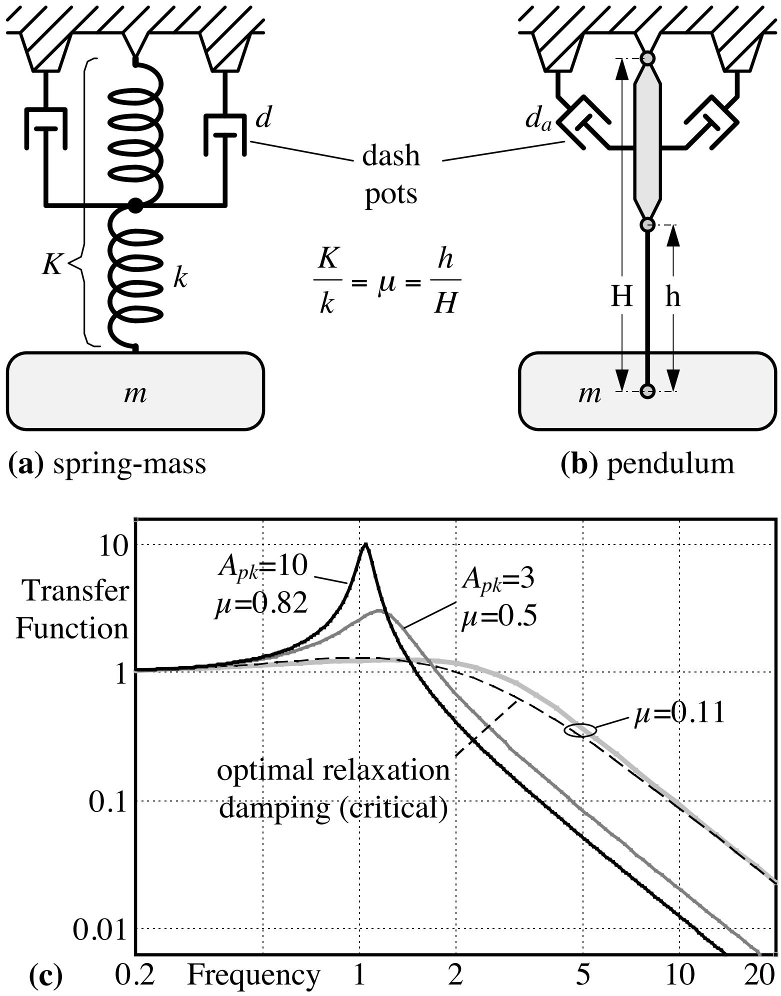

Figure 1.8 (a) Simple spring-mass oscillator with pseudo-structural damping, (b) Equivalently damped pendulum. (c) Transfer functions of these systems for constant K or H but various μ values and optimal damping.

or infinite, then no damping will be present and we simply have one or other resonant frequency corresponding to the two spring-rates (K and k) or pendulum lengths (H and h). There must be an optimum value somewhere between zero and infinity at which to set the damping for each fraction of spring or pendulum length bridged by the damper. Once this optimum value is found, then there remains only one usefully adjustable parameter - the undamped spring fraction μ - and some plots of the transfer function for different values of μ are shown in figure 1.8(c). It can be seen that they all fall off with $1/f^2$ above resonance although the resonant peak moves up in frequency as more of the spring or pendulum is bridged by the damper (the plots shown are for constant K or H).

There are actually at least two contenders for "optimum" damping in this case. One is the value that produces the lowest resonant peak*, and the other is the value that

---

* The math procedure used to derive the minimum peak value is common in texts on vibration [e.g. Den Hartog 1956] and is not shown here. It mainly consists in identifying a frequency at which the amplitude

produces the shortest relaxation time but which peaks slightly higher. An expressions for tuning for the lowest resonant peak height and the height obtained is given in table 1.3, but expressions for the shortest relaxation time were not obtained for this case. However it is obvious that a minimum relaxation time value must exist, and it is shown to be different from the lowest peak tuning by the numeric determination of a critical damping solution which is discussed below.

The performance of these oscillators may be characterised as follows :-

Physical Values		Units
m	suspended mass	(kg)
K	spring-rate of entire (both sections) spring	(N/m)
k	spring-rate of undamped section of spring	(N/m)
d	total dissipation coefficient of linear dashpots	(N/(m/s))
H	length of entire (both sections) pendulum	(m)
h	length of undamped section of spring	(m)
d_a	total dissipation coefficient of angular dashpots	(N·m/(rad/s))

Oscillation Characteristics	Spring-Mass Suspension	Pendulum Suspension
Resonant frequency with dashpot free	$2\pi F = \sqrt{\frac{K}{m}}$	$2\pi F = \sqrt{\frac{g}{H}}$
Resonant frequency with dashpot locked	$2\pi f = \sqrt{\frac{k}{m}}$	$2\pi f = \sqrt{\frac{g}{h}}$
Spring Ratio (dashpot locked / dashpot free)	$\mu = \frac{K}{k} = \frac{F^2}{f^2}$	$\mu = \frac{h}{H} = \frac{F^2}{f^2}$
Q-factor of oscillation with k = ∞ or h = 0	$Q = \frac{1}{d}\sqrt{\frac{m}{1/K - 1/k}}$	$Q = \frac{m\sqrt{g(H-h)^3}}{d_a}$

Table 1.2 Characteristics of pseudo-structurally damped oscillators.

---

of the transfer function is independent of the damping coefficient. This is achieved by equating the ratios between the corresponding frequency terms from the numerator and the denominator of the transfer function. This gives the lowest point through which all possible curves for varying damping must pass. If the slope of the amplitude is then set to zero at this point, the lowest amplitude damping solution is obtained.

Tuning and Performance	Lowest Resonant Peak	Minimum Relaxation Time
Optimum setting for Q-factor of oscillation (with k = ∞ or h = 0)	$Q_{opt} = \sqrt{\dfrac{2\mu(1-\mu)}{1+\mu}}$	Critical damping for μ=0.11, Q=0.5434
Peak transfer function motion enhancement	$A_{pk} = \dfrac{1+\mu}{1-\mu}$	

Table 1.3 Tuning and performance of pseudo-structurally damped oscillators.

It is interesting to note that this system can be made critically damped for values of μ less than approximately 0.11 (i.e. 89% of the spring bridged out by the dashpot). With μ=0.11 critical damping may be obtained with Q=0.5434. This is quite different from the minimum peak tuning (which gives Q=0.42) and gives a relaxation time of less than half of that given by the minimum peak tuning. The performance of the system with these two different types of optimum tuning are shown in figure 1.8(c). The dashed curve is the one which is optimised for shortest relaxation time and is critically damped, while the grey curve gives minimum peak height.

It is easy to see how this damping technique could be applied to a vertical suspension chain, but it is not quite so obvious how to apply it to a pendulum chain. However for intermediate stages of a pendulum chain, the rigid section of the pendulum link may be divided such that approximately half of its length is above each pivot and half is below (the pivot in the middle carries the mass). Then each pivot is made viscously stiff for damping while the flexible fibre sections of the links provide the good isolation at high frequency that we are trying to achieve with this damping technique.

1.4.5 Vibration Absorber and Self-Damping

The next increase in complexity is to add a second mass to the previous structure allowing it a second degree of freedom. There is a choice of whether to place the extra mass cascaded in series with the suspension chain or hung off to the side and in parallel with it. Since the purpose of the exercise is to achieve good damping by adding viscous coupling between two points, if the points are within the chain, then we will have used the height available for two stages but not achieved the $1/f^2$ isolation that we would like *per stage* due to the viscous coupling between a pair of them. Hanging a mass effectively “off to the side” or in parallel with the suspension chain allows the motion of the suspension chain mass to be damped while leaving the remaining height available for cascading further stages.

In the case of the spring-mass system shown in figure 1.9(a), there is physically little difference, except that the transfer function is taken from the suspension point to the first mass m_1, rather than to the second m_2. Additional stages would then be hung from the first mass and the second only serves to damp the motion of the first by being viscously coupled to it. However attempting to achieve the same effect in the horizontal pendulum regime immediately greets one with conceptual problems!

If the pendulums were suspended with two fibres so that they were forced to remain horizontal as they swung, then there would be no problem. A damping mass can be suspended from the main cascaded mass and viscously coupled to it with a rigid post reaching between the stages. Additional stages could be hung from the first mass rather than the second just as with the spring-mass system. In this case the analogy is exact and obvious. This approach also highlights the main difference between the vertical and horizontal systems. In the vertical, all the forces and motions lie on the same line of action with one another, whereas in the horizontal, the forces and motions of each pendulum have the offset of the pendulum height between their lines of action.

As a result such a system generates significant tilting moments due to the viscous force between stages. The target application for these techniques also needs to be soft in tilting in order to isolate that degree of freedom also. This approach will provide strong viscous coupling between these degrees of freedom which could be advantageous in that it could be used to damp both degrees of freedom to each other simultaneously. However it does not lend itself to the two degree of freedom analysis that we are attempting, and also the large majority of suspension chains are not of this multiple fibre variety.

An approach which could be applied to a single fibre pendulum would be to hang a second pendulum to coincide its mass at the same height as the first so that the horizontal forces are in line and can be viscously coupled without creating tilting moments. In order to make their frequencies different (so that they can be usefully coupled) they could have different upper suspension points (and thus different lengths). This would be easy to analyse but does not correspond to the vertical case, and would not be useful for more than one stage in a cascaded arrangement. If they are to be hung from the same height, then it is necessary to change the frequency of one of them by changing the tension in its fibre without a corresponding change in mass. This may be achieved by hanging more mass some distance below as a cascaded pendulum. Indeed one can imagine two complete chains hung concentrically "beside" each other, one being the main suspension chain and the other being solely a damping chain. By altering the mass distribution down the chain, and maybe an extra terminating mass on one of them, the frequency difference between each stage and its damper could probably be chosen to optimise damping effectiveness. Again however such a many degree of freedom system does not lend itself to this simple level of analysis.

The damping solution shown in figure 1.9(b), while possessing an identical analysis to the spring-mass system in (a), is really significantly different (and apparently new). The swinging of the dumbbell shaped mass as a pendulum and its rocking motion form a two degree of freedom system in which the two motions can be viscously coupled in a similar manner to which the two masses in the spring-mass system are coupled. It makes use of the height offset between the lines of action creating moments to achieve this coupling. The transfer function is taken between the suspension point and the lower

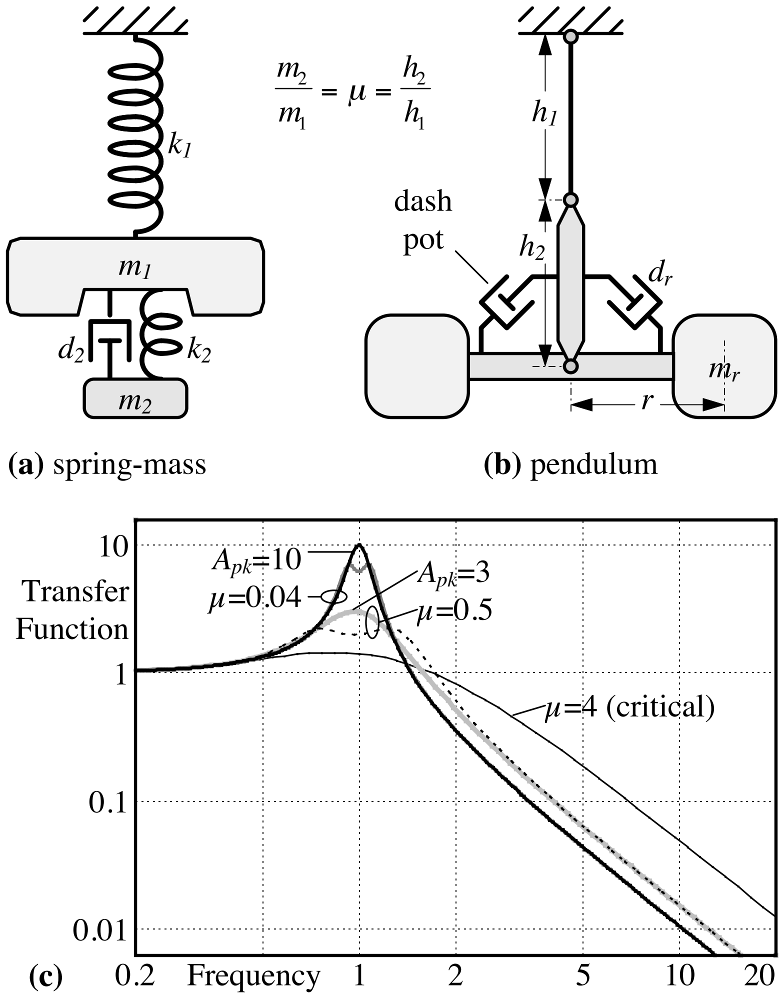

Figure 1.9 (a) Standard spring-mass vibration absorber with damping, (b) “Self-damped” pendulum of total mass m_r, at a radius of gyration r, total dashpot damping d_r segmented massless pendulum h_1+h_2. (c) Transfer functions of these systems with different types of optimal tuning (single peak = shortest settling time, double peak = minimum peak height)

pivot at the centre of the dumbbell shaped mass (hereafter called the *rocker*). Further stages would be suspended from this point for a cascaded system.

Comparing the spring-mass system in figure 1.9(a) with the pendulum in (b), we see that if the dashpot of the spring-mass system is locked, then it becomes a simple undamped single-mode oscillator with a resonant frequency f_S determined by the total mass m_1+m_2 together with spring-rate k_1 as detailed in table 1.5. The same is true of the pendulum system in (b) if its dashpots are *free*. It's single mode resonant frequency f_S is determined by the total pendulum height h_1+h_2 and is also given in table 1.5 (the rocker has no moments applied to it if its dashpots are inactive and it remains horizontal as it swings). Considering the opposite dashpot condition for both systems, the spring-mass system becomes a simple undamped dual mode oscillator with its dashpot free, and so does the pendulum with its dashpots locked (one swinging and one rocking mode). By varying the dashpot coupling, the two modes can be made to damp each other.

If we obtain the transfer function of each system (which is done in sections 8.2.5 and 8.2.7) and equate the coefficients of the various frequency terms we obtain relationships between the parameters of the systems required to make them dynamically equivalent to each other. There are many rearrangements of these relationships possible but the selection shown in table 1.4 seemed to be the most meaningful.

<u>Physical Values</u>		<u>Units</u>
m_1	main mass	(kg)
m_2	absorber mass	(kg)
k_1	main spring	(N/m)
k_2	absorber spring	(N/m)
d_2	dissipation coefficient of absorber dashpot	(N/(m/s))
l_1	length of main spring under load of m_1+m_2	(m)
l_2	length of absorber spring under load of m_2 only	(m)
h_1	length of upper flexible fibre pendulum section	(m)
h_2	length of lower rigid pendulum section	(m)
m_r	mass of entire rocker	(kg)
r	radius of gyration of rocker	(m)
d_r	total dissipation coefficient of rocker dashpots	(N·m/(rad/s))
I_r	moment of inertia of rocker (=$m_r r^2$)	(kg·m^2)
cop(h_2)	centre of *precession* (equivalent simple pendulum length) of compound pendulum h_2 & rocker with dashpot locked.	(m)

Spring-Mass Suspension	Equivalence Equation	Pendulum Suspension
Single mode resonant frequency with dashpot locked (d_s=∞)	$\frac{k_1}{m_1+m_2}=\frac{g}{h_1+h_2}$	Single mode resonant frequency with dashpot free (d_p=0)
Mass ratio	$\frac{m_2}{m_1}=\frac{h_2}{h_1}$	Pendulum length ratio
Standalone absorber resonant frequency (m_2 with k_2, dashpot free)	$\frac{k_2}{m_2}=\frac{g\,h_2}{r^2}=\frac{g}{cop(h_2)}$	Lower compound pendulum resonant frequency (h_2 pivot, dashpot locked)
Ratio of absorber damping coefficient to absorber spring-rate.	$\frac{d_2}{k_2}=\frac{m_r r^2}{d_r}=\frac{I_r}{d_r}$	Inverse ratio of angular damping coefficient to rocker moment of inertia

Table 1.4 Equivalence of Spring-Mass and Pendulum oscillators shown in figure 1.9(a) and (b).

It is clear from the correspondence in the table that as far as their transfer functions are concerned the two systems are *duals* of each other with spring-rates corresponding to inverse moments of inertia, and masses corresponding to pendulum lengths. The additional degree of freedom in these oscillators now means that there are three parameters to be adjusted : (1) the viscous damping coefficient (d_2 or d_r), (2) the mass ratio or pendulum length ratio μ , (both of which had parallels with the previous structure), and (3) the resonant frequency of the standalone absorber spring-mass.

Defining Parameters	Spring-Mass Suspension	Pendulum Suspension
Single Mode Resonant frequency	$2\pi f_S=\sqrt{\frac{k_1}{m_1+m_2}}$	$2\pi f_S=\sqrt{\frac{g}{h_1+h_2}}$
Mass or Pendulum length Ratio	$\mu=\frac{m_2}{m_1}$	$\mu=\frac{h_2}{h_1}$
Standalone absorber resonant frequency	$2\pi f_A=\sqrt{\frac{k_2}{m_2}}$	$2\pi f_A=\sqrt{\frac{g\,h_2}{r^2}}$
Q-factor of standalone absorber oscillation	$Q_A=\frac{\sqrt{k_2 m_2}}{d_2}$ (m_1 fixed, dashpot active)	$Q_A=\frac{d_r}{m_r r\sqrt{g\,h_2}}$ (h_2 pivot, dashpot active)

Table 1.5 Characteristics of oscillators with damped vibration absorbers.

The spring-mass system shown in figure 1.9(a) is well known to vibration engineers as a tuned "vibration absorber" with damping and is to be found with a full analysis in many standard texts on vibration [e.g. Den Hartog 1956]. The fact that the single stage self-damped pendulum in figure 1.9(b) is a dual of it means that the accumulated knowledge is also directly applicable to the pendulum version. There are well known formulas for adjusting both the resonant frequency of the absorber, and the damping coefficient in order to achieve minimum peaking of the transfer function. We have also derived a formula for the damping coefficient that provides minimum relaxation[*] time (the frequency tuning is the same) and these formulas are shown in table 1.6. In the case of the pendulum system an implication of the frequency tuning formula is that the radius of gyration r is determined by the pendulum lengths according to the equation given in the table. The equivalent spring system relation of spring lengths being related to mass ratio only holds if the springs are linear and would go to "zero-length" when released.

Tuning Parameters	Lowest Resonant Peak	Minimum Relaxation Time
Optimum setting for absorber standalone resonant frequency	$f_A = \dfrac{f_S}{\sqrt{1+\mu}} \quad \left(= \dfrac{f_1}{1+\mu} \right)$	
Frequency tuning implication for spring-mass system	$\dfrac{k_1}{k_2} = \dfrac{(1+u)^2}{u} \quad or \quad \dfrac{l_2}{l_1} = 1+u$	
Frequency tuning implication for pendulum system	$r = h_1(1+\mu)\sqrt{\mu} = \dfrac{h_2(1+u)}{\sqrt{\mu}}$	
Optimum setting for absorber standalone Quality factor[†]	$Q_A = \sqrt{\dfrac{2}{3}}\sqrt{1+\dfrac{1}{\mu}}$	$Q_A = \dfrac{1}{2}\sqrt{1+\dfrac{1}{\mu}}$

Table 1.6 Optimal tuning of damped vibration absorbers.

Having determined both the absorber resonant frequency and damping from the formulas, again leaves only one usefully adjustable parameter - the mass ratio or pendulum length ratio μ. Some plots of the transfer function for different values of μ are shown in figure 1.9(c). In this case we have plotted curves for both types of optimum tuning. It can be seen that the resonant peak does not shift up in frequency as

[*] The minimum relaxation time solution in this case was found by obtaining an expression for the roots of the denominator and making a pair of them degenerate by equating the surd expression to zero.

[†] The usual text book equivalent of this formula gives the damping ratio of the *absorber* mass at the resonant frequency of the *main* system resulting in a cubed $(1+\mu)$ term under the root. If the frequency of the absorber system is used instead, the cube disappears.

this type of damping is increased which results in a slightly better performance than the previous pseudo-structural damping technique. The isolation performance varies with frequency as $1/f^2$ as we would wish and as it did with the pseudo-structural case.

The performance expected from the two types of optimum tuning are shown in table 1.7 and it is interesting to note that with the minimum relaxation time tuning there is a very simple relationship between the mass ratio and the pole Q-factor obtained. From this relationship it is evident that critical damping is possible with optimal tuning when the ratio $\mu = 4$. The performance of a system with this ratio and tuning is also shown in figure 1.9(c). The width of the rocker mass hanging from the pendulum in this case needs to be a minimum of four times the height of the overall pendulum! Similarly for the critically damped vibration absorber, the absorber mass is 4 times the main mass and hangs below the main mass 5 times as far as the main mass hangs below its support (assuming linear springs). This means that the structure with its split mass is distributed over 6 times as much height as its stage height.

Performance Obtained	Lowest Peak Tuning	Minimum relaxation time
Peak transfer function motion enhancement	$A_{pk} = \sqrt{1 + \frac{2}{\mu}}$	$A_{pk} = \sqrt{1 + \frac{4}{\mu}}$
Frequency of Poles		$f_p = f_S$
Q-factor of Poles		$Q_p = \frac{1}{\sqrt{\mu}}$
Relaxation time		$\tau = \frac{Q_p}{\pi f_p}$

Table 1.7 Performance of oscillators with optimally tuned and damped vibration absorbers..

1.4.6 Self-Damping Other Degrees of Freedom

Despite the fact that the transfer functions of this vibration absorber and single stage self-damped pendulum can be made equivalent, there are other aspects which are quite different. Suppose we consider the critically damped version with the large absorber mass and the small main mass. If a second stage is suspended from the main mass, then it is clear that at high frequencies it will only see the main mass (i.e. 1/5 of the *stage* mass) appearing as a reaction mass. However the self-damped pendulum is quite different - at high frequency its entire stage mass appears as a reaction mass to subsequent stages. This must make a difference to the mode frequencies in a cascaded system where one generally desires that the mass ratios of preceding to following stages should be as high as possible (see section 1.3.3).

Since these damping stages are clearly quite different in some respects, it is interesting to consider whether the self-damping idea can be applied to the vertical spring-mass stage so that it would use another degree of freedom of the main mass for damping (rather than splitting the mass into parts and damping between them). This is indeed the case and it can be done quite readily - probably most usefully by coupling to the rotational θ_z degree of freedom. What is required is a means whereby vertical tension couples to rotational motion. Then some means of interposing a dashpot between the coupling and the suspended mass. The simplest to imagine might be a coil spring with a very coarse helix so that as it stretches it tends to uncoil and vice-versa. This coil spring is connected to a pivot with a vertical axis and in-built rotational dashpot. The other side of the pivot is connected to the suspended mass.

If the dashpot in the pivot is made free, then the system has a single vertical resonant mode where the uncoiling and recoiling of the spring freely spins the pivot back and forth. This is analogous to the self-damped pendulum with dashpot free. If the dashpot is locked, then the system has two modes - being in phase and opposite phase combinations of the vertical bouncing mode and rotational twisting mode. This is analogous to the self-damped pendulum with locked dashpot. If the dashpot is set to some ideal viscous coefficient then optimal damping between the two modes can be obtained. As with the self-damped pendulum and in contrast to the vibration absorber, the entire stage mass is kept rigidly together in a single lump, the entire stage mass appears as a reaction mass to subsequent stages, and no extra height is used for the structure other than the stage height.

One small flaw that could be picked with the analogy between these structures, is that the vertical system just described exerts torsional moments on the stage above as well as the vertical force normally expected. This would seem to be essentially different to the pendulum's single pivot point attachment. This is true but could be overcome by having a double system - two springs two pivots and two masses. One spring with a left hand helix and the other with a right hand helix so that the torsional moments cancel at the suspension point. Such a double system can be suspended from a single moment-free pivot point as can the self-damped pendulum. However the mass has now been split (although not in the same manner as the vibration isolator), adding a new differential rotational mode and the common θ_z rotational mode (which cannot exist with a moment-free pivot point suspension) is no longer cross-coupled or damped.

The author believes that the best suspension system would be a multi-fibre system in which the modes of all six degrees of freedom have similar resonant frequencies. With multiple pendulum fibres, θ_z resonant frequencies become similar to x and y, while with vertical suspension on those fibres, θ_x and θ_y tilting resonant frequencies become similar to z. To conceive a method of coupling between the vertical suspension and θ_z that was provided by the coarse helical spring, one might imagine a very large diameter helical spring in place but the end cut to only give a quarter turn. If the suspension fibre is attached to the end of this spring to use it as a cantilever, and similar segments of spring added for the remaining fibres, then θ_z to z coupling is in place.

The most important point however that this discussion has been leading to is now that the *one single large mass* of the stage is sufficient to self-damp *all degrees of freedom*! (The vibration absorber with its split mass could never do this). It is just a matter of mounting the mass in such a manner that it is free to rotate in all three degrees of freedom - such as it would be if it was simply mounted on a short piece of spring wire near its centre of mass. Eddy current damping can then be applied to its periphery in such a manner as to damp two rotational degrees of freedom with a single magnet. The details can be greatly improved - by replacing the cantilever suspension with the Euler spring technique described in chapter 2 for instance, but the concept seems to present a very elegant solution to multi-degree of freedom passive damping of cascaded isolation stages.

1.4.7 Summary

We have examined three damping approaches of increasing complexity and damping performance, the best of which is the well known resonant “vibration absorber” with damping. A new damping method similar to the vibration absorber has been presented which viscously cross-couples different degrees of freedom of the *same* mass rather than splitting the mass into two parts and only using one degree of freedom of each part. It is applicable to both vertical and horizontal isolation and can damp *all* 6 degrees of freedom simultaneously with the *same* single mass. It can also work suspended from a moment-free pivot point. This new method is identified as having single stage damping characteristics equivalent to the vibration absorber but with highly significant advantages over it when applied in a cascaded system.

We have identified an alternative “optimal” tuning which minimises relaxation time rather than transfer function peaking thereby allowing critical damping to be readily reached. Using this type of stage and tuning, it may be possible and practical to construct a critically damped suspension chain with: $1/f^2$ isolation performance per stage, $(g/h)^{1/2}$ stage resonant frequencies, and no wastage of available height for damping structures.

A practical model of a self-damped pendulum stage (not attempting to achieve critical damping) in a simple cascaded configuration is numerically assessed in section 3.3.

1.5 Ultra-Low Frequency Pre-Isolation Overview

1.5.1 Introduction

What has come to be termed a pre-isolator, is in principle just an isolation stage that has been designed to have a very low resonant frequency of suspension. However there are a couple of good reasons for the distinction between a *pre*-isolator and normal isolation stage. The main one being that structures that are designed to have very low resonant frequencies tend to be larger and more massive and consequently have lower internal resonances (which bypass their isolation) than normal isolator chain stages. As a result they generally do not provide useful isolation in the 10 Hz - 1 Khz detection band but are included mainly to reduce residual motion at frequencies below the detection band. The worst residual motion usually occurs near 0.5 Hz to 1 Hz (for 2.5 m total height) at the resonant frequencies of the lowest swinging mode(s) of the standard isolator chain. Good pre-isolation can easily reduce the seismic drive to these modes and thus the residual low frequency motion by two orders of magnitude. The pre-isolation structures also allow smooth translation of the entire chain and test-mass by macroscopic distances (millimetres) to counteract any thermal expansion or drift effects in distance between end masses or their alignment.

In the horizontal regime there is an additional difference between the two isolator types. Horizontal isolator stages are generally pendulums of an appropriate height (tens of cm's) which sets their resonant frequency $(g/h)^{1/2}$ and are attached by a single (ideally) perfect pivot to a fixed point or to the stage above. However in order to obtain a frequency lower than that given by $(g/h)^{1/2}$, it is necessary to attach the isolator stage in a different manner than a simple pivot. This is obvious when one considers what happens when the centre of mass of the load becomes displaced from being directly underneath the suspension point (as a result of vibration of the suspension point for instance). With a pivot suspension the force applied to the centre of mass of the load remains unavoidably in line with the moved suspension pivot point and so must contain a horizontal component which is the fraction of mg proportional to the tan of the offset angle. This horizontal force produces the normal pendulum acceleration and resonant frequency.

In order to apply less horizontal force than this, it is necessary that the restoring force is applied in a direction other than towards the previously offset suspension pivot point.

This requires an attachment other than a simple pivot from a single point[*]. In fact it is apparent that as the centre of mass of the load is horizontally offset, the suspension point must move horizontally with it to keep the net suspension force pointing almost vertically upwards. The desired effect is that the net force should point as if it comes from a very distant fixed point - like a ball rolling on a slightly concave surface with the suspension point at its center of curvature, or like the suspension point of a very long pendulum. A structure which provides a *net* force of this nature is obviously essential for a lower-than-pendulum-frequency horizontal pre-isolator.

A requirement that makes the design of isolators for gravitational wave detection more difficult than most other applications is that the mechanics have to be compatible with high-vacuum (i.e. no oily lubricants) and generate no stiction or other noise during motion. This rules out all sliding and rolling joints and reduces almost all motion to what can be done with flexing - which is ideally just intermolecular spacing stretching and compressing. This makes it more difficult to achieve the previously described motion (of a very long pendulum) in a reasonable amount of space and demands one or more of the following ingenious structures.

The fact that the structures are unavoidably quite massive, means that the dynamic effect of seismic *acceleration* must be considered and designed around. This is primarily a matter of making sure that attachment points are located at centres of percussion as defined in section 1.6 and as discussed in detail for just about every isolator covered in depth in this thesis.

---

[*] A nagging thought which occasionally troubles researchers in this field is whether there may not be some way to achieve lower than pendulum resonant frequencies in a single point cascadable manner. Barring the application of external forces, conservation of momentum dictates that the only way to allow some part of the suspended mass to accelerate slower than a simple pendulum, would be to have the remainder of the mass accelerate faster so that the *centre* of mass moves as the simple pendulum to keep the line of force through the pivot. However doing this requires more energy than a simple global acceleration, and a source of energy to provide this extra acceleration in an ongoing manner cannot be present in a *passive* system.

1.5.2 Horizontal Isolation Structures

1.5.2.1 General Characteristics

The horizontal isolation structures used to date may be divided into two main types based on the nature of the forces which are summed to provided the low net restoring force required for low frequency oscillation. These are illustrated in figure 1.10.

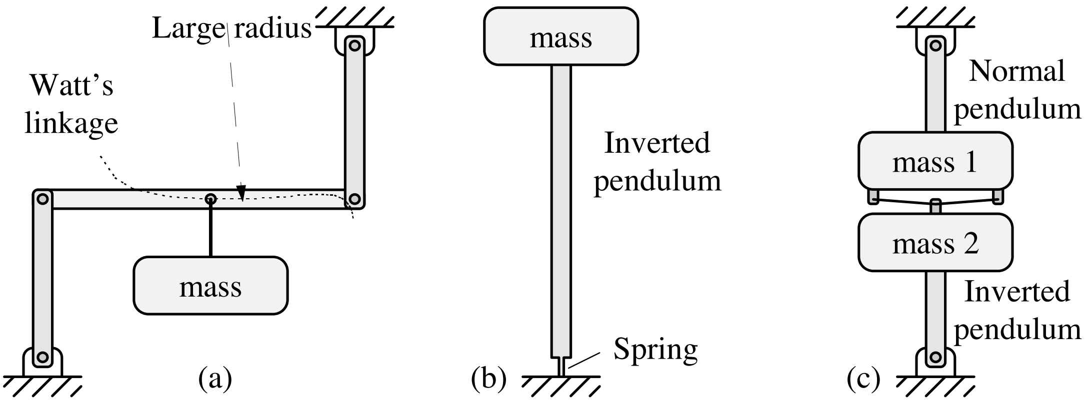

Figure 1.10 Three types of ultra-low frequency horizontal isolation structures :- (a) Geometrical, (b) Counter-Spring, (c) Counter-Weight.

The first type that may be termed "geometrical" work by causing the mass suspension point to move in an almost flat horizontal plane. The restoring force in this type could in principle be entirely gravitational, but in practice ends up being a combination of small spring-rates from flexes and strains, balanced with a very shallow gravitational hill or well. Their distinctive characteristic is that they try to keep stored energy as low as possible. All of the linkage based varieties are of this type. Figure 1.10(a) shows Watt's linkage with a mass suspended from an appropriate point which moves along a circle of very large radius (for small offsets at least).

The second type makes no attempt to keep stored energy low but excels in simplicity and relies on a strong counter-spring to null out the relatively deep gravitational potential hill or well that the suspension point moves in. The main examples of this type are the inverted pendulum and "wobbly table" devices of which the simplest version is illustrated in figure 1.10(b).

Considering the differences between these two types, suggests a third type illustrated in figure 1.10(c) in which a deep gravitational potential hill serves as a counter-weight to a deep potential well. To the author's knowledge, this type have only been discussed and never built but may offer the low energy storage of the geometrical type together with the simplicity of the counter-sprung type.

One main difference between gravitational forces and elastic forces from springs, is that the former are essentially lossless and the latter are not. When opposing forces are nulled out by balancing them against each other, only the lossless part is nulled and any lossy parts sum to appear as greatly enhanced or concentrated after the removal of the lossless. The result of this difference is that the second type (b) that store comparatively large amounts of energy in counter springs have much greater losses and lower Q-factors, and thus can never be tuned to as low a frequency as an equivalent type (a) or (c). It is possible that the lower Q-factor may be an advantage in providing damping for operation at a desired frequency which is not aiming to be as low as possible, but generally it is a disadvantage.

Experience shows that the damping encountered in this case is structural, and the Q-factor falls off with frequency squared. So the Q-factor of an isolator's resonance which is say 64 at 10 sec period will drop near unity by tuning down to ~80 sec. As an example of the difference between the performance of the geometric and that of a counter-spring type we could compare the Scott-Russel isolator which still had a Q-factor of ~20 at 150 sec (see sect 4.4), with a wobbly table counter-spring type of similar (~1 m) size and aluminium springs in which the Q-factor approached unity at only 40 sec. The Virgo wobbly table, being much larger, has a considerably shallower gravitational hill and less counter-spring requirement. It also used higher Q-factor material for the springs and achieved a Q-factor of ~33 at 33 sec period. (Thus one might expect it to be tunable down to ~190 sec - see figure 6.12). It should be noted that as the Q-factor approaches unity, so the device verges on instability [Saulson 1994] and one cannot expect to operate such a device at low Q-factors without a control system.

Another difference between gravitational and elastic forces is that the former are virtually constant whereas the latter usually have some temperature and time dependence. The result of spring-rate nulling is to strongly accentuate these effects by the amount of nulling taking place. For instance to tune a 1 metre inverted pendulum to have 30 sec resonant period requires a 99.6% cancellation of forces. At this level any temperature dependency will be amplified by a factor of 250. Saulson has shown that with only a factor of two or three times this nulling in steel flexures, time dependent effects bring on short term instability [Saulson 1994].

Another disadvantage with the counter-spring types is that they are very load dependent - they rely on the load to tune the frequency down. By comparison the tuning of the geometrical types can be made virtually independent of load. However due to the mass of the rigid structure making up the geometry (which mass typically does not move horizontally), it is usual for the geometrical structure to move in a shallow potential well, and for the mass to be suspended on a small potential hill to cancel with the well (similar to the counter-weight type (c)). This makes the resonant frequency of the geometrical types somewhat load dependent but to a much lesser degree than the counter-spring type. The important difference being that the height of the potential hill can be adjusted by the geometry, whereas the counter-spring type have no adjustment possible for different total masses except changing the spring - which is usually a major operation.

1.5.2.2 Geometrical Types with Minimal Energy Storage

Watt's linkage of figure 1.10(a) is also known as a folded pendulum and has been used extensively for low-frequency and sensitive experiments [Lindenblad 1967, Blair 1994, Liu 1997, Fan 1999, Bertolini 2001]. There are various others like it - such as the "swinging gate" type commonly used in seismometers - which give long period suspension, but are only capable of being applied to one degree of freedom at a time. There are some others shown in figure 1.11 which are all capable of being generalised to two degrees of freedom in the one device (with varying degrees of difficulty) and so are of considerably greater interest for isolation. The required length ratios of the various links are not included here but are obtainable from many texts such as [Beggs 1955].

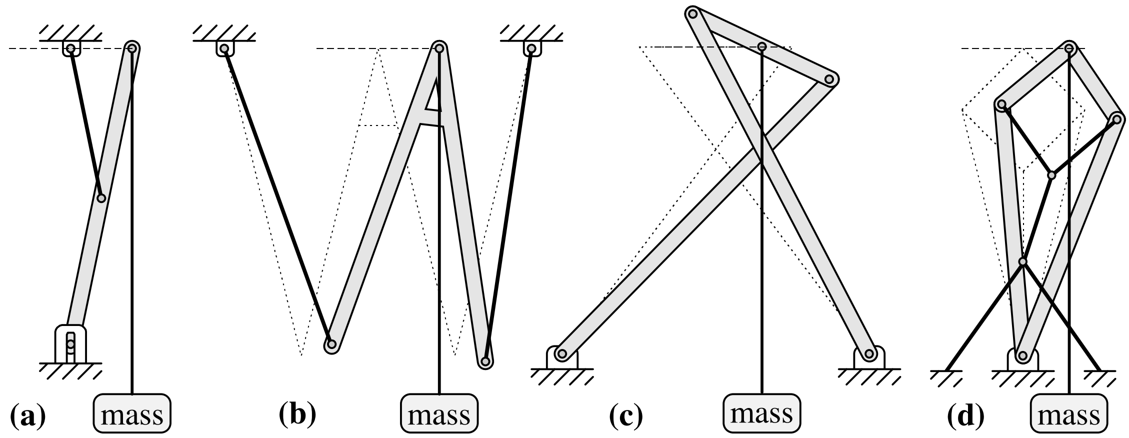

Figure 1.11 Various straight-line linkages :- (a) Scott-Russel, (b) Roberts Linkage, (c) Chebyshev Linkage, (d) Peaucellier Cell Linkage.

All of the linkages shown in figure 1.11 have been shown with the mass support point at the top as this is generally the most convenient for use as a pre-isolator so that the suspension chain can hang down from that point making double use of the height. They can all be inverted in which case some rigid links can be turned into fibres and vice-versa.

The Scott-Russel linkage in figure 1.11(a) seems the simplest but it needs a sliding joint of some sort. For large motion this can be provided by another approximately straight-line linkage, but for small motion can rely on the inherent flexibility of materials. We have used this linkage extensively and chapter 4 reports fully on it. Lindenblad built an inverted version of it [Lindenblad 1967] which seems considerably more difficult than the versions we have built. Once the Scott-Russel has been rearranged to allow the mass to hang freely down the centre, it becomes approximately as complex as the Roberts linkage which is naturally clear in the centre.

The Roberts linkage in figure 1.11(b) is probably the simplest linkage to apply to 2-D pre-isolation. It is not immediately obvious that it can be exactly generalised to 2-D, but one can imagine a copy of the device turned 90° about a vertical axis and the rigid centre section joined to the first copy. Then looking North-South one can see the East-West W shape in full profile but the North-South W shape is a vertical line. As the suspension point is moved fully East across one's view the N-S becomes a W shape of half the width of the East-West one. Since both end points of this half width W shape remain on the top horizontal line joining the suspension points, it follows that the half-width W shape must work in the same manner as the one seen in full profile.

We have built ultra-low frequency seismometers using this linkage (using 3 wires instead of 4) and it works very well indeed although as always one has to make allowance for flexibility and stretch in the materials which adds an anti-restoring force (see similar effect in Scott-Russel in 4.2.3 & 4.2.4). The structure does not seem to be well known but has been incorporated as the horizontal isolation component for vibration isolation legs for optical tables for instance [Nelson 1998](it appears that this patent only applies to the specially arranged combination of this linkage with a pneumatic vertical cushion).

The Chebyshev linkage in figure 1.11(c) is not nearly so convenient as the Roberts since the links have to cross over and when generalised to 2-D this problem gets worse. Also when generalised to 2-D in the orientation shown it becomes unstable in the torsional θ_z degree of freedom. This is readily corrected by inverting it, which wastes height but removes the need to keep the links away from the centre to avoid the suspended payload. However the links still need to pass close to each other without touching and this seems to require that they be made of rigid members which can be dog-legged to miss each other. The closest case of this linkage being used is the X-Pendulum [Barton 1996] which only achieved one degree of freedom in a single stage.

Wires were used as the links and the length ratios were not the classical Chebyshev with the result that the point giving long radius circular motion was a long way from the crossed fibres. Two degrees of freedom were achieved by cascading two stages in an ingenious height saving manner [Barton 1999].

The Peaucellier Cell in figure 1.11(d) is the most remarkable of the linkages shown here in that it achieves *exact* straight line motion using only pivots (whereas the others are only approximate or use a sliding joint). What is also remarkable and initially surprising is that it can be immediately and perfectly generalised to 2-D isolation in a single stage. The two points which have to be held fixed are the bottom of the kite shape and the bottom of the Y shaped fibres (which is only prevented from left-right motion by the tethers in the diagram), and the bottom of the kite shape has to be able to pivot on two axes. After some thought it becomes clear that the geometry is the same regardless of turning the kite shape about its long axis - and thus it must work identically in the other direction also. The author built an inverted version of this linkage approximately 1 m in length in which the outer kite shape was wires, and the inner Y shape was aluminium tubes with wide hinges on the ends. It worked very well indeed and it was initially a surprise that it worked so well in the unintended direction! However despite its beauty, it remains considerably more complex than the first two in this set and is unfavourable for this reason. (There is also an alternate arrangement of the Peaucellier linkage that looks totally different but works identically).

Once generalised to 2-D isolation, these linkages may be applied in multiple copies at each corner of a platform for instance to provide a “wobbly table” isolation stage similar to the inverse pendulum type to be described next but without the high energy storage. A diagram of such a scheme was presented for the Scott-Russel as figure 4.5.

1.5.2.3 Counter Sprung Types - Inverse Pendulums

Due to the inherent simplicity which makes this suspension technique attractive, it does not come in many different varieties. The normal requirement that the center be free for the suspended payload, means that inverse pendulums for isolation purposes are almost always built into “wobbly tables” as shown in figure 1.12 [Pinoli 1993, Holloway 1997, Losurdo 1999]. Due to the devices accentuation of the imperfections of the material, the single leg version shown in figure 1.10(b) has also been used to probe the anelastic properties of springs [Saulson 1994].

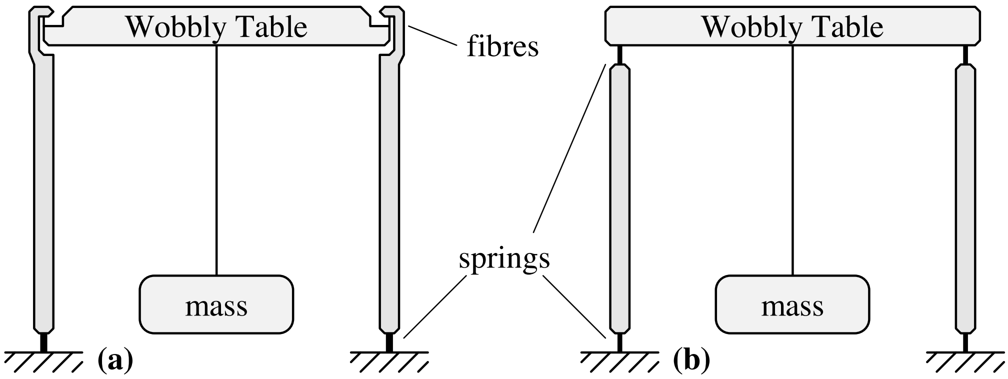

Figure 1.12 (a) Fibre suspended wobbly table, (b) Double flexure wobbly table.

The wobbly table pre-isolator shown in figure 1.12(a) is suspended by thin fibres at the top of the legs (it could just as well be springs at the top and fibres at the bottom). The purpose of this is to make the torsional θ_z degree of freedom just as soft as the translational - in effect achieving 3 axes of isolation in one stage. (In fact it seems quite difficult to obtain 2 axes without the third being included - even when it is not desired.) This fibre suspended version of the wobbly table has the disadvantage that the load on the table top needs to be well centred or else the leg or legs that are bearing most load become unstable and flop over to one side in a twisting action with the table top turning about the remaining leg(s) which remain upright

The wobbly table shown in figure 1.12(b) does not suffer from this effect and the table top can be loaded up by piling weights in almost any position without it collapsing in this twisting mode. However this table has a relatively high torsional θ_z resonant frequency - much higher than the translational. This is generally not a problem since the θ_z torsional seismic motion is much smaller at the radius of the legs than the translational and so does not need as much filtering. Also if a single fibre suspension chain is used then the θ_z filtering down the chain is vastly better than any other degree of freedom. It is important to adjust the suspension point of the chain accurately to the center of mass of the table top so that horizontal cross-coupling is avoided. This can be done by a nulling adjustment while exciting the θ_z normal mode - after which the weight distribution on the table top must remain balanced about this suspension point.

There are other considerations with these devices but they are well covered in the paper on Virgo's wobbly table pre-isolator in chapter 6.

1.5.2.4 Counter Weighted Types

It seems to the author that this device type has considerable benefits to offer which have not been fully appreciated to date. It might be guessed from the simplistic version shown in figure 1.10(c) that the two pendulums need to be the same length but this is not so for small displacements. If one pendulum is made a lot shorter than the other one then it needs much less weight hanging on it than the other one. This is illustrated in figure 1.13(a) in which the normal counter-weight pendulum is 1/4 of the weight and 1/4 of the length of the spring-less wobbly table inverse pendulum which supports the suspension chain.

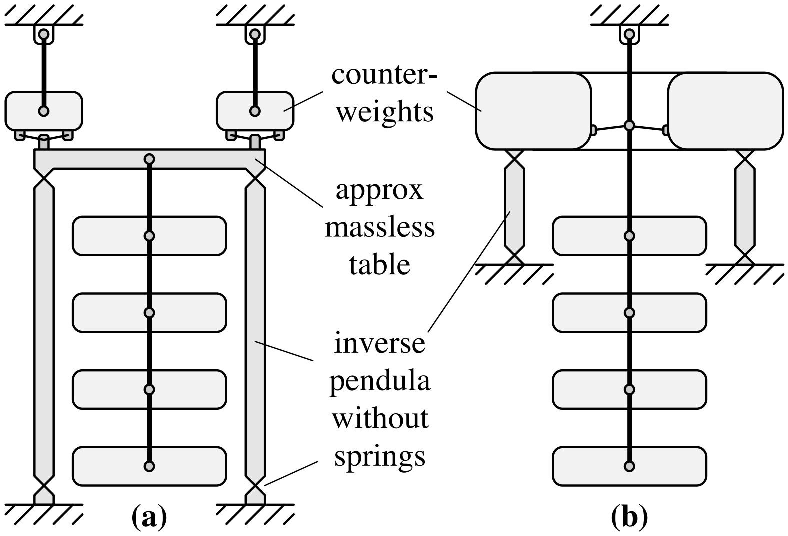

Figure 1.13 Counter-weighted pre-isolators :- (a) Suspension chain inverse, counter-weights normal, (b) Counter-weights inverse, suspension chain normal.

Figure 1.13(b) shows the roles reversed where the suspension chain is suspended from a normal pendulum and an approximately equal mass counter-weight is added to it as a springless inverse pendulum to tune its resonant frequency down. In all cases a pre-isolator stage needs a reasonable amount of mass in order to avoid a high frequency mode of interaction with the following stage. It seems possible to make use of this mass by the methods shown to provide a nulling gravitational spring to reduce the resonant frequency of the pre-isolator. (The wobbly table in figure 1.13(b) as shown is actually unstable in the θ_z degree of freedom. This can be overcome but the change required did not display the concept so simply).

1.5.2.5 Tilt Coupling in Horizontal Pre-Isolators

These ultra-low frequency pre-isolation structures all work by providing a rigid constraint against the very large gravitational force of the total suspended mass while also providing extremely soft horizontal compliance (e.g. a very long pendulum or a ball rolling on a slightly concave plate). This requires the structure to be maintained in very accurate alignment with the gravitational field. However the earth's surface is subject to a certain amount of seismic tilt noise which perturbs any fixed alignment. This tilt noise multiplied by g then appears directly as horizontal acceleration in the mass of the pre-isolator stage, which in turn integrates up to become the equivalent of translational horizontal seismic motion.

As an example, consider a mass which can slide around frictionlessly on a perfectly flat surface. This represents a perfect passive horizontal isolator with an infinitely low resonant frequency - perfectly isolating all translational horizontal motion of the surface. If the surface is subject also to tilt noise, then the perfectly isolated mass will be accelerated around randomly in a similar manner as if it was being disturbed by translational motion. Once horizontal isolation has reduced the translational feedthrough down to this level (and it seems not very hard to reach with a good pre-isolator), then there is no more improvement to be made without sensing and counteracting tilt actively. This whole subject is the main topic of chapter 7.

1.5.3 Vertical Isolation Structures

As mentioned previously, there is nothing special about a vertical *pre*-isolator in comparison with an isolator except that it is to have a very low resonant frequency and it is acceptable if it no longer acts as a filter at the normal detection band frequencies. In order to achieve a very low resonant frequency without an enormous spring length extension it is essential to have a nonlinear force-displacement characteristic. So any method of engineering a nonlinear spring with a spring-rate which can be adjusted and maintained to be almost zero for a reasonable range of motion around the operating point is likely to be a good candidate for a pre-isolator. Figure 1.14 shows a selection of one well known type and three newer types for discussion.

The time honoured method of achieving a very low vertical frequency oscillator for gravimeters and vertical seismometers is shown in figure 1.14(a). A method for winding the “zero-length” spring necessary for its best operation at least was reported by LaCoste [LaCoste 1934] and the characteristic arrangement with the diagonal spring bears his name. A more complex arrangement again using diagonal springs that achieves straight line motion instead of the circular motion of the figure was reported considerably later [LaCoste 1983]. Contrary to popular belief, the angle of the spring does not have to be 45° but can be any angle whatsoever. The vertical offset between the spring ends simply determines the approximately constant lifting force, while the horizontal offset serves to reduce the vertical spring-rate close to zero. It can be conceived of as an inverse pendulum on its side with the horizontal offset of the spring

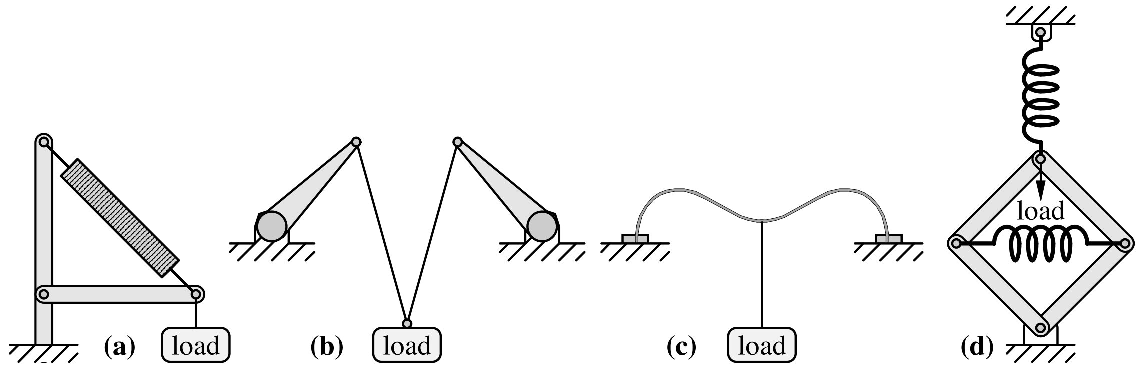

Figure 1.14 Low frequency vertical oscillators :- (a) LaCoste, (b) Torsion-Crank, (c) Bow-Spring, (d) Rhombus negative spring.

providing the longitudinal compressive force tending it towards instability. Assuming a zero-length spring (in which the force is directly proportional to its length without any offset), then if the fixed end of the spring is attached anywhere vertically in line with the pivot as shown, then the vertical spring-rate is zero (i.e. vertical force on the load is constant). If the fixed end of the spring is moved to the left of this point then the vertical spring-rate becomes negative and the device is unstable, and vice-versa if the fixed end is moved to the right. So the vertical offset sets the approximately constant value for the lifting force, while any horizontal attachment offset (each side of the pivot) tunes the spring-rate and thus resonant frequency around zero.

Figure 1.14(b) shows a torsion-crank suspension which gets away from the coil springs of the LaCoste and uses torsion rod springs instead. It has a nonlinear force-displacement characteristic which can be flattened to give low resonant frequencies over a reasonable operating range, but nowhere near as large a range as either the LaCoste (a) or the Rhombus (d). It is fully covered in chapter 5 of this thesis.

Figure 1.14(c) shows a spring technique that the author was investigating for use in the filters of Aigo's suspension chain, but a large and carefully tuned version could also be applied as pre-isolation. The spring itself is intended to be initially flat and is stressed into the position shown which is seen edge on. The spring material needs to be quite thin in order to bend to the sort of radii required, but can be as wide (into the page) as is required to support the load. The principle is that if the ends are clamped at a suitable distance apart (the diagram shows the correct spacing for a vertical launching angle), then a flat region is obtained in the force-displacement characteristic which provides for low frequency resonance. If the ends are moved closer together, then the region becomes negative sloped and unstable and vice-versa. By being able to alter this spacing with a screw adjustment, the resonant frequency of the stage can be adjusted. The ends of the spring do not need to start off vertical as shown but this launching angle is a parameter which may be chosen to maximise working range or used as an alternative tuning control. Before the author finished applying this technique, the Euler technique of chapter 2 was discovered and this bow-spring technique was discarded in favour of the Euler technique. Another group however has implemented vertical suspension using a very similar technique to this in which the (plan view) shape of the strip is not rectangular, but sculpted and possibly pre-bent near the middle. They achieved excellent results [Ando 2001] with their device that they term a "monolithic geometric anti-spring" (MGAS).

Figure 1.14(d) shows a remarkable structure which the author accidentally discovered while mathematically simulating various arrangements of levers and springs. If the vertical spring (which for this case needs to be equal to the horizontal spring and both of them "zero length") is connected across the rhombus instead of above it, then we have the interesting situation of neutral equilibrium where the rhombus can be squashed to any shape and the springs will not produce any net force to change the shape. This is because the energy stored in a spring is proportional to the square of its length, and the sum of the squares of the diagonals of a parallelogram is constant. Thus if the energy is constant, then there is no force ($F=dE/dx$). From this fact, and the fact that we know that the force being generated by the vertical spring acting alone is directly proportional to its length, we can deduce that the vertical force generated by the horizontal spring acting through the rhombus linkage must be the exact negative of the vertical spring (i.e. a perfectly constant negative spring-rate over the entire length of travel!). So if we reverse the direction of action of the vertical spring by attaching it to the upper support as shown, the negative spring-rate of the rhombus spring must cancel exactly the positive spring-rate of the vertical spring to produce a constant force over the entire travel. Knowing that the rhombus spring simply has a constant negative spring-rate, it is obvious that the force from the two springs do not have to be equal at all - just their spring-rates need to be equal. The vertical spring can be very large and support a lot of weight, while the rhombus negative spring can be quite small and only contribute a small amount of force. As long as the spring-rates are matched the force will be constant and the suspended load will have a very low resonant frequency.

As was pointed out some time later [Holloway 1996], the structure has an even more remarkable property - if the suspension point is allowed to be offset a small amount horizontally, then it generates a horizontal restoring force as low as the vertical restoring force. Also since the device is cylindrically symmetric it must work just as well in both horizontal directions. So by making the spring-rate of the vertical spring slightly larger than that of the rhombus negative spring, this device provides ultra-low frequency, 3-D isolation in a single stage. One can imagine an isolator using three such structures, one attached to each corner of a triangular platform to support its weight and that of the isolated payload.

The author attempted to build a demonstration model of a simple rhombus negative spring implementation very early on in this research but failed to complete it for various reasons. Since then however the Euler spring technique of chapter 2 has appeared and this may usefully expand the possible realisations of this rhombus spring because Euler springs can be applied from the other direction as compression springs, and can easily be adjusted to be equivalent to zero length extension springs.

There are various other ultra-low frequency vertical suspension techniques not mentioned here - such as the magnetic anti-springs of the Virgo group [Beccaria 1997, DeSalvo 1999] and a "geometric anti-spring" [DeSalvo 2001] - but some will be mentioned elsewhere in this thesis. The vertical pre-isolation technique in greatest favour of this author at the time of writing would be the Euler spring techniques of chapter 2 and some methods of spring-rate reduction are described in that chapter.

1.6 Center of Percussion

1.6.1 Definition

The terms “center of percussion” and “center of percussion tuning” are used in several places in this thesis and the concept does not seem to be a well known one. In addition it seems surprisingly difficult to find a standard text that even has the term in the index. Since this is the case it seems useful to briefly define and discuss the concept as the author has intended it to be understood.

For any rigid body undergoing translation and rotation in a plane, a point may be found which is momentarily stationary with respect to the plane. For instance on a yoyo rolling down a stationary string, this point is the point at which the string leaves the coil on the yoyo.

If a rigid body is *accelerated* by a force which is not in line with its center of mass, then it undergoes rotational as well as translational acceleration. In this case there is similarly a point which may be found which is momentarily not undergoing acceleration. This point is commonly called the “center of percussion”. It is the “sweet spot” on a bat or racket where there is no apparent recoil from a sudden acceleration or force impulse applied at another point.

This point is always paired with a second “force impulse point” on the opposite side of the centre of mass where the force is being applied. If the force is applied instead to the center of percussion, then the previous force impulse point becomes the new center of percussion. If a force is applied to any other point which is not one of the pair, then a different centre of percussion has to be found for the different force application point.

In suspension work, this concept finds a lot of usefulness because high frequency seismic motion behaves in the same manner as a force impulse. Thus when high frequency seismic motion is applied to a typical rigid member of a suspension structure, there will be a point on that member which is almost unaffected and scarcely moves in response to high frequency seismic motion. In fact the point only rotates, and if further structures are attached to this point by a pivot, then almost no motion is coupled.

1.6.2 Equations

Following are some equations for making calculations about this effect.

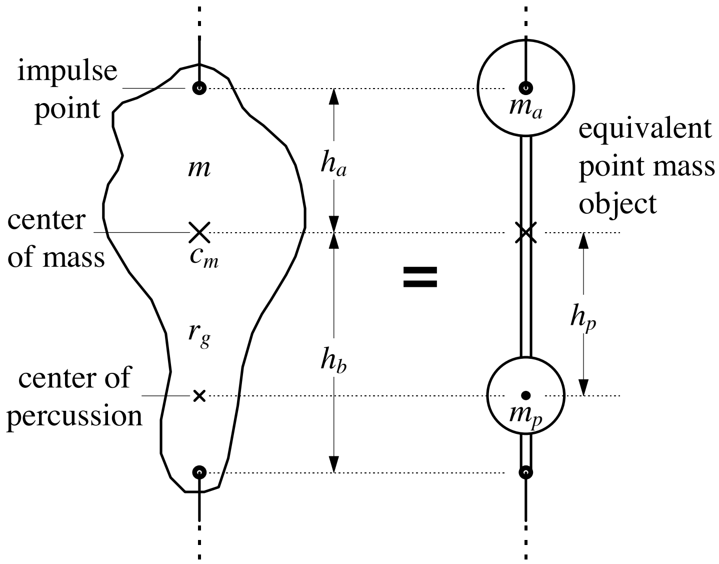

Figure 1.15 Point mass equivalent of a rigid body and its center of percussion.

Any massive body constrained to planar motion, can be completely characterised by its mass m, the position of its centre of mass c_m, and by its radius of gyration r_g. Thus for in-line attachment points on a typical pendulum suspension we can expect an upper attachment point some distance h_a above the centre of mass c_m of the standalone body. At some other distance h_b below the centre of mass there is usually a lower attachment point for cascaded pendulum stages. The distributed mass of the body itself can be replaced with an exact equivalent consisting of two point masses m_a and m_p located along a rigid massless rod joining the attachment points. One of the point masses m_a can be placed at the upper attachment point, and the other m_p must lie somewhere below the centre of mass such that the total mass is unaltered :

$$m = m_a + m_p$$

the moment of inertia (or radius of gyration) is unaltered :

$$m r_g^2 = m_a h_a^2 + m_p h_p^2$$

and the centre of mass is unaltered :

$$m_a h_a = m_p h_p$$

Solving these three equations simultaneously gives the relationships :

$$m_a = \frac{m r_g^2}{h_a^2 + r_g^2}$$

$$m_p = \frac{m h_a^2}{h_a^2 + r_g^2}$$

$$h_p = \frac{r_g^{\ 2}}{h_a}$$

It is apparent that if a horizontal impulse force is applied at the position of the upper mass, then the lower mass will have no reason to translate but will only rotate. Likewise if a horizontal impulse force is applied at the position of the lower mass, then the upper mass will not translate but will only rotate. So the positions of the two point masses at h_a and h_p from the center of mass, are the force impulse point and centre of percussion pair in this case. The last equation gives the useful relationship that their product is always equal to the radius of gyration squared (or their geometric mean is equal to the radius of gyration).

To avoid high frequency seismic motion from appearing at the lower attachment point h_b as a result of motion applied at the upper attachment point h_a, h_b should be made equal to h_p which is the center of percussion for h_a.

1.6.3 Adjustment Procedure

In practice this positioning often needs to be very accurate and is best achieved by applying a high frequency shaking signal at h_a and adjusting for a null at h_b. The amplitude of the feedthrough indicates the distance away from optimum, while the phase (polarity) indicates whether it is above or below.

An incorrect adjustment for centre of percussion can be recognised in a transfer function as a flat floor level below which the isolation cannot go. Where this joins with the $1/f^{2n}$ fall-off after resonance, there is either a gentle asymptotic curve to meet the floor, or a notch. This effect is shown in figure 6.6 and is fully described in that section.

2. Euler Spring Vertical Isolation

2.1 Preface

The idea of using Euler buckling as a low spring-rate suspension first occurred to the author while reading the chapter on Elasticity in Feynman's lectures on Physics [Feynman 1964] where he describes the critical Euler force and the buckling phenomenon. He mentions that "a beam will not necessarily collapse completely when the force exceeds the Euler force" because the force is larger than the simplified expression obtained. Unfortunately he gives no clue as to what the character of the force-displacement might be aside from this suggestion that it might be stable.

An early (and incorrect) attempt by the author at deriving what the force-displacement relationship might be like, suggested that it started out with zero gradient before curving gently upwards (like a parabola of x^2+1 starting at unity). Excited by this possibility, the author tried to dream up methods of attaching crossed-over wires to opposite ends of straightened clock springs so that when the wires are stretched the clock springs are compressed to buckle. Unfortunately all easily dreamt up schemes of compressing something long and thin, by crossing over wires to apply tension to opposite ends of it, seemed to either be unstable, or require too much complexity. So the idea remained in limbo for many months.

The next major stimulus resulted from an over dinner discussion at a Tama conference between the author, Warren Johnson and Ron Drever. One of them said "wouldn't it be nice if you only had to worry about the energy going in and out due to the vibration instead of all the strain required to suspend against gravity." Another noted that that was exactly how an active system normally worked requiring only very small AC forces to be generated to counteract the vibration (after having the DC gravitational component balanced with a large spring of some sort). At this the author pointed out that Euler column buckling was a passive system which was almost as good - needing no pre-strain to support the mass, and providing almost zero spring-rate just post buckling. He was wrong about the "almost zero" of course and on being asked where he learnt that from, was unable to give a satisfactory answer, and also was unable to suggest a stable structure to utilise the effect. The discussion moved on but not before that particular idea had been well and truly removed from the back-burner in the author's mind.

This impetus occurred at just about the time that there was a strong need to complete a design for a new suspension chain. The author had been very unhappy with the previous curved cantilever springs having seen the bad internal-mode performance of apparently optimally designed and beautifully engineered cantilever springs of Virgo. In addition the author had strongly desired to come up with some vertical spring-rate reduction scheme to apply to the isolation stages in the chain to improve their isolation and lower the Q-factor of the vertical modes to go with the self-damped horizontal system. The necessity on the one hand of coming up with something tolerable, together with the reminder that the Euler buckling should be looked at again, resulted in some rapid designs, great measurement results, an application for a patent on the idea, and eventually the following papers. Terran Barber's name is on the second paper as he set up the experiment and measured the force displacement data as part of his honours project work. The rest however is entirely the author's own work.

2.2 Seminal Euler-Spring Paper

High Performance Vibration Isolation Using Springs in Euler Column Buckling Mode.

J. Winterflood, D. G. Blair

Department of Physics, University of Western Australia, Perth 6009, Australia.

A new vertical suspension technique utilising the remarkable properties of column springs in Euler buckling mode allows the mass of suspension springs to be reduced towards their ultimate minimum, greatly increasing the resonant frequency of internal modes, and allowing near ideal vibration isolation to substantially higher frequencies than achievable conventionally.

2.2.1 Introduction

Vibration isolation consists in suspending a system in such a manner that motion of the support results in minimum motion of the isolated system. This is achieved by the mass of the isolated system being kept in place by soft linear restoring forces, thereby forming a harmonic oscillator with respect to the support. The softer the restoring force, the lower the resonant frequency, and the better the isolation obtained.

The stringent demands of gravitational wave detection has provided the motivation for the development of greatly improved vibration isolators [Ju 2000b]. In this paper we present a new use of elastic springs which allows vertical vibration isolation to achieve performance comparable to the horizontal vibration isolation of a simple pendulum. To explain the significance and benefits of the new technique it is necessary to discuss the intrinsic difficulties of conventional vertical isolator approaches.

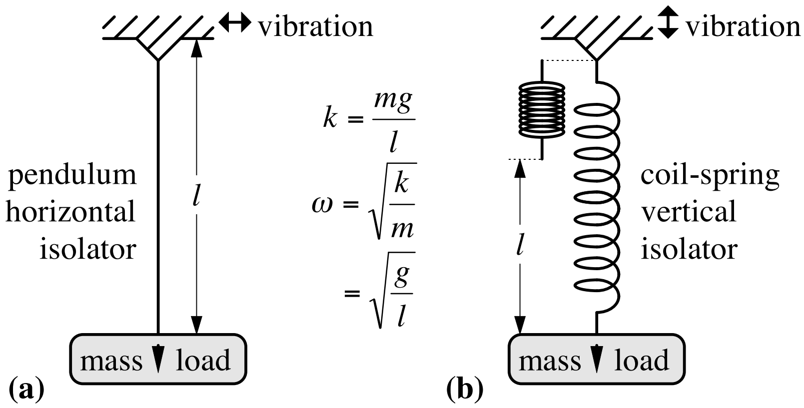

Figure 2.1 Typical horizontal and vertical isolation structures.

The simplest structure for isolating from *horizontal* motion is the pendulum as shown in figure 2.1(a), together with its isolation performance being the dashed curve in figure 2.2. A 25 cm pendulum has a resonant frequency of approx 1 Hz. Below resonance the isolated mass moves with the same amplitude as the support - and thus the transfer function has a magnitude of 1.0. At resonance the pendulum amplifies vibration by a factor determined by its Q-factor (typically several hundred) although this motion can normally be damped to greatly reduce this amplification. Above resonance the transfer function has a slope of -2 (on a log-log plot) so that isolation improves with frequency squared. Thus at 10 Hz for instance, the transfer function is approximately 0.01 or 1% (meaning that for 10 mm motion of the support, the isolated mass only moves 0.1 mm). If the support shakes at higher frequencies, or if the pendulum is made longer thereby lowering the resonant frequency, then the isolation is better still.

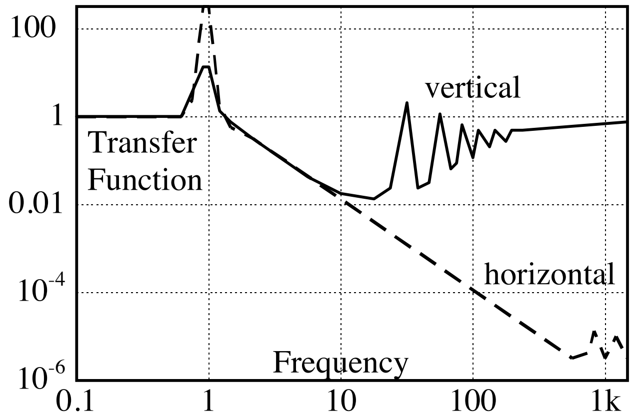

Figure 2.2 Typical transfer functions of horizontal and vertical structures.

In our laboratory we have demonstrated synthetic pendulum structures with very long periods that have achieved a transfer function with near ideal performance over almost 3 decades of frequency comparable to the dashed transfer function in figure 2.2 but at much lower frequency [Liu 1995, Winterflood 1999]. For 1 mm applied vibration this corresponds to just 10 nm amplitude at the optimum isolation frequency.

Vertical vibration isolation presents more problems than horizontal isolation because the vibrational motion (between the support and the isolated mass) in the presence of gravity requires the dynamic storage of significant amounts of energy (*mgh*) to momentarily absorb this motion. This energy storage is often provided in the form of mechanical springs. One of the simplest structures for isolating from vertical motion is the coil spring suspension shown in figure 2.1(b). In order for a *linear* spring (i.e. displacement $\propto$ force) such as this to achieve the same resonant frequency as a pendulum, it must be stretched under load (beyond its relaxed length) by the same length as the equivalent pendulum. Thus if a relaxed 10 cm long coil spring is loaded with mass so that it becomes 35 cm long, then the mass will oscillate with a frequency of 1 Hz.

Most high performance vertical vibration isolators operate in a regime where the amplitude of vibration is very small compared with their extension under load. Under vibration only this small amount of dynamic energy is exchanged in and out of the spring while a large amount of energy remains statically in the spring due to its initial extension under load. The key point is that the total energy stored in the spring is much larger than the dynamic energy storage required in operation, as figure 2.3 illustrates.

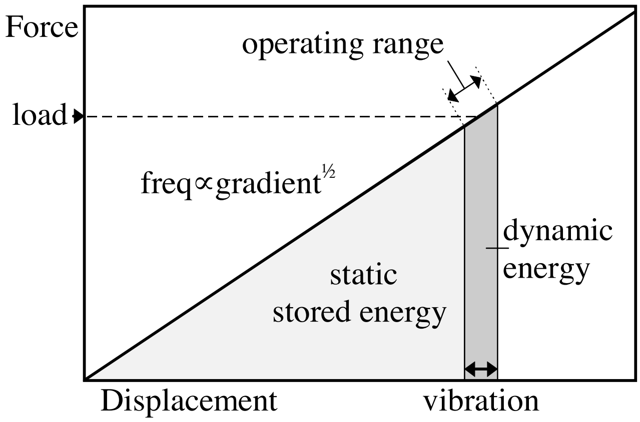

Figure 2.3 Force - displacement plot for a simple spring vertical isolator.

The large static energy storage necessitates a large mass of elastic material - proportional to the total static and dynamic energy to be stored. (In a pendulum on the other hand there is in principle no static energy and the dynamic energy is stored as gravitational potential energy - not requiring elastic material with mass) The disadvantage of the large mass requirement is that it supports low frequency *internal* resonant modes of the spring (e.g. "surging" in coil springs). These resonances have a large effective mass which strongly couple vibration between the support and the isolated mass at these internal resonant frequencies. Because the internal mode density increases with frequency, performance is strongly degraded as is illustrated in figure 2.2.

The same effect also occurs in the pendulum, but because the suspension fibre can be made so much less massive than the spring, its first internal (violin string mode) resonant frequency is much higher. In addition it couples much less energy due to the large ratio of the suspended mass to the fibre mass (the coupling depends also on its Q-factor).

2.2.2 Advanced Vertical Isolation

There are three main areas to consider in order to alleviate the spring mass problem inherent in vertical vibration isolators:- (1) ensure that the entire volume or mass of spring material is usefully storing energy by being stressed to its limit, (2) redistribute the mass of the spring to minimise its velocity and thus the kinetic energy of any internal mode motion, and (3) reduce the static energy and mass of the spring while keeping a low resonant frequency by producing a nonlinear force vs displacement relationship.

The first of these areas only offers a small gain (a maximum of 2 for torsion and 3 for bending) and is rarely considered. A simple example would be to wind a coil spring from a tube rather than from solid material. This removes the central mass which is scarcely stressed and does not store much energy. This would be difficult but not inconceivable to apply to bending (one can imagine filling a central cavity with some lightweight and almost incompressible liquid!).

However when a spring is only stressed in one polarity, then in principle it is possible to preset initial stresses within the material so that the entire volume is stressed to its limit at some maximum deflection. For material loaded in torsion this is commonly called "scragging" (from "to wring the neck of"!). Consider a torsion bar which is plastically twisted well beyond its initial yield point and then released. The internal layers will not be able to relax with the amount of untwist that relaxes the surface layer, and so will strain the surface layers in the reverse direction until the torque from the internal layers is balanced. It is apparent that in order to twist such a bar again to its yield point, it must be twisted considerably further than before - initially to undo the reverse surface strain, and then the same as before to reach yield. Since the spring-rate is unaltered, such a torsion bar (which may be formed into a coil spring) can store considerably more energy for one direction of twist than it could before. The equivalent approach for flat springs is to bend them well past yield in the loaded direction. An example of this is curved cantilever blades and we believe this technique can also result in reduced creep at high stress levels [Ju 2000a].

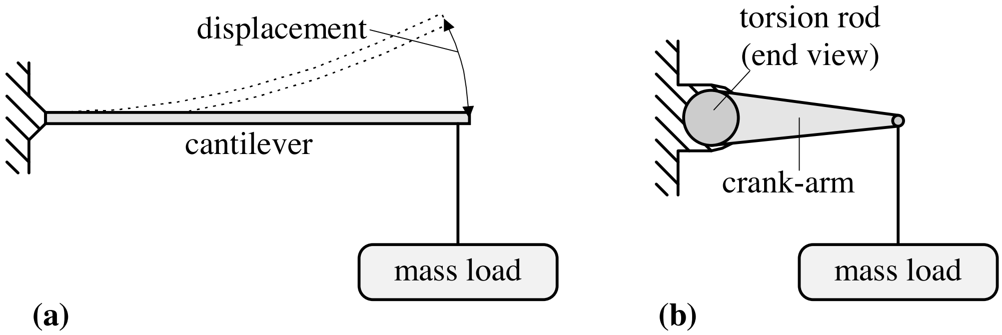

Figure 2.4 Concentrating the energy storage and thus mass to low velocity regions.

Examples of better distribution of the mass are shown in figure 2.4. The cantilever blade [Ju 1994b] in figure 2.4(a) with constant stress over its surfaces is in the shape of a triangle (top view) with the fixed attachment being the wide base of the triangle and the suspending tip being the apex of the triangle. This puts most of the mass near the fixed attachment where motion is minimum, thus achieving higher internal mode frequencies. Probably the best that can be achieved is the torsion rod suspension [Blair 1993] shown in figure 2.4(b). Here the crank-arm can be made very rigid (to have very high internal mode resonances) and the spring material is all located very close to the fixed centre of rotation.

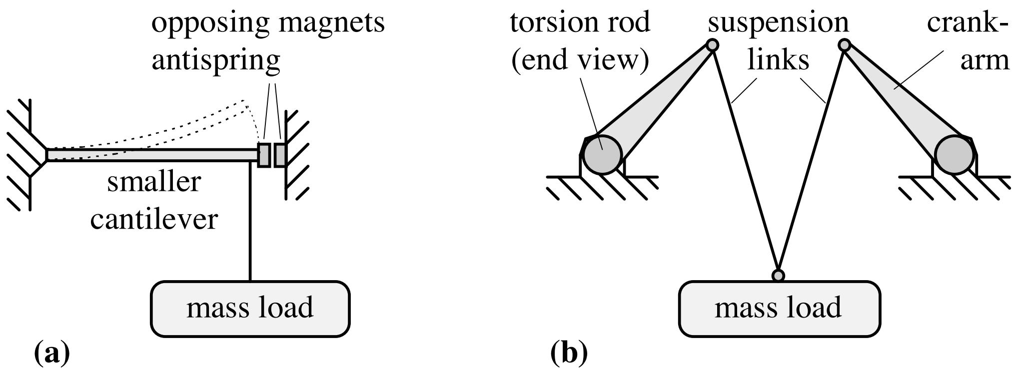

Figure 2.5 Anti-spring and nonlinear techniques allow much smaller spring mass to achieve low resonant frequencies.

Examples of reducing the static energy and mass stored in the spring using nonlinear spring techniques are shown in figure 2.5. The cantilever in figure 2.5a is fitted with magnets which strongly repel and try to drive the mass away from the normal operating position [Beccaria 1997]. This is an anti-restoring force or “anti-spring” which when added to the normal restoring force of the spring produces a region of reduced gradient on the force displacement curve as is illustrated in Figure 2.6. The main resonant frequency of an isolation system (eg: 1 Hz in figure 2.2) is determined by the (square root of the) spring-rate which is the gradient of the force-displacement characteristic. By operating in the flattened region of the curve, a much lower resonant frequency is obtained than would normally be the case for the static displacement and energy storage employed. The torsion-crank suspension in figure 2.5b achieves a similarly flattened

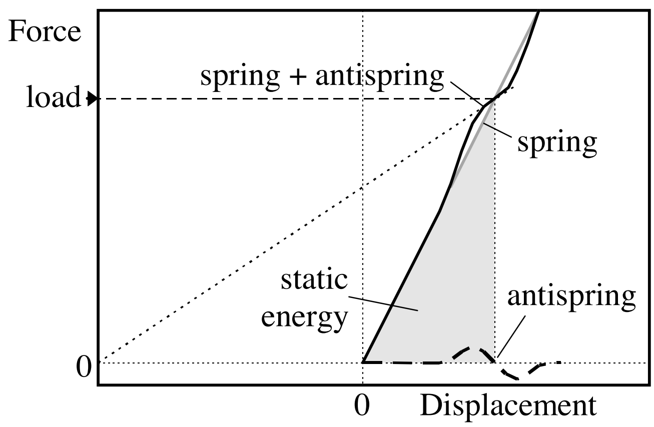

Figure 2.6 Force - displacement diagram for typical anti-spring system.

region in its force-displacement characteristic by its geometry and has been fully described previously [Winterflood 1998].

Since the mass of spring material used must be proportional to the total (static + dynamic) energy stored in the spring, it is clear from figure 2.6 that for a given resonant frequency, the spring mass used in such a nonlinear system can be greatly reduced from its linear equivalent. It also becomes apparent that the best possible situation would be obtained if the static energy could be reduced to zero so that only dynamic energy storage need contribute to the spring mass. Remarkably this ideal is readily obtained by the very simple spring arrangement described hereafter.

2.2.3 Euler Buckling Spring

It is well known in engineering that a column of elastic material will support a load with virtually no deflection until at some critical value of load (dependent on its modulus and not on its yield strength) it suddenly starts to buckle. This is the sort of wall-shaped nonlinearity (figure 2.8a) that is required in the force-displacement characteristic to provide zero static energy. Figure 2.7 shows two cases of buckling columns - figure 2.7a having pinned ends (a pivot point at each end) and figure 2.7b having rigidly fixed or clamped ends.

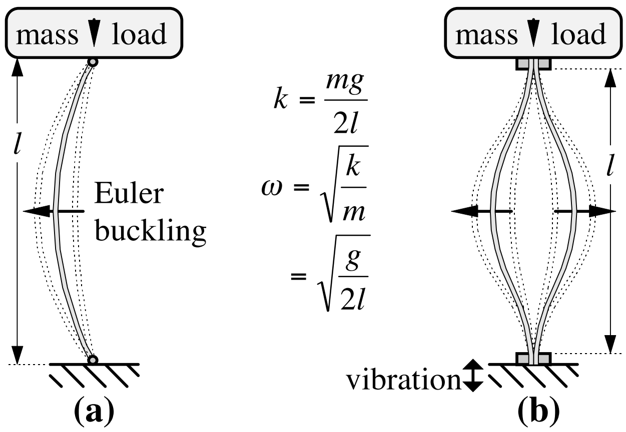

Figure 2.7 Spring blades in Euler column buckling mode.

The well known Euler column formula giving the critical load at buckling for the pin-ended case is $P_{cr}=\pi^2 EI/l^2$. Here E is modulus of elasticity, I is the area moment of inertia and l is the length of the column [Feynman 1964]. The critical load for the clamped end case is also well known and may be obtained by comparison :- It is apparent that the pin-ended column is the same curve as the centre section of the clamped end column between the points of inflection, which in this case is half the length. It follows that the critical load for a single clamped spring will be four times as great as for the same spring with pinned ends.

What is not well known about Euler column buckling is the relationship of force vs displacement beyond the initial onset of buckling at critical load (as buckling is generally to be avoided!). The shape assumed by an inextensional elastically buckling spring is called an *elastica* and exact analysis of large deflections in such a spring involves the use of elliptic integrals. The application of elliptic integrals to the elastica is well covered in advanced texts [Reddick 1947] and used by ourselves in a related analysis of nonlinear spring-rate reduction [Winterflood 2001b]. If the degree of buckling of a clamped spring as in figure 2.7b is specified by the maximum angle α_0 it attains with respect to its unbent centre line, then the force F (normalised to the critical load) and displacement x (normalised to the unbent spring length) are given by :-

$$F = 4\,K(k)^2 / \pi^2$$

$$x = 2\left(1 - E(k)/K(k)\right) \tag{2.1}$$

where $k=\sin(\alpha_0/2)$ is called the modulus and $K(k)$, $E(k)$ are complete elliptic integrals of the first and second kinds. These equations are plotted in figure 2.8b and show that the force vs displacement characteristic is remarkably well behaved with a low spring-rate and this remains reasonably constant even up to very large amounts of buckling.

Figure 2.8a illustrates that if the working range is designed to start just at the onset of buckling, then all of the energy stored by this type of spring is the dynamic energy exchanged in and out while operating within its working range. It also becomes apparent that the mass of spring required is directly proportional to the working range that can be accepted. For example suppose we can accept a working range of 0.5 mm at a resonant frequency of 1 Hz (normally requiring 25 cm static extension in a linear system), then if made of the same elastic material the Euler spring may be 1/250th the mass of its linear equivalent. One would expect the internal resonant modes of the Euler spring to be roughly proportional to the square root of this mass ratio (i.e. $\sqrt{250} \approx$ 16 times higher than the linear case!) and the effect of the coupling at the internal resonances to be reduced by this mass ratio (i.e. 1/250th). In addition we will show (section 2.2.6) that spring-rate reduction can also be achieved, to allow, in principle, even greater relative improvement.

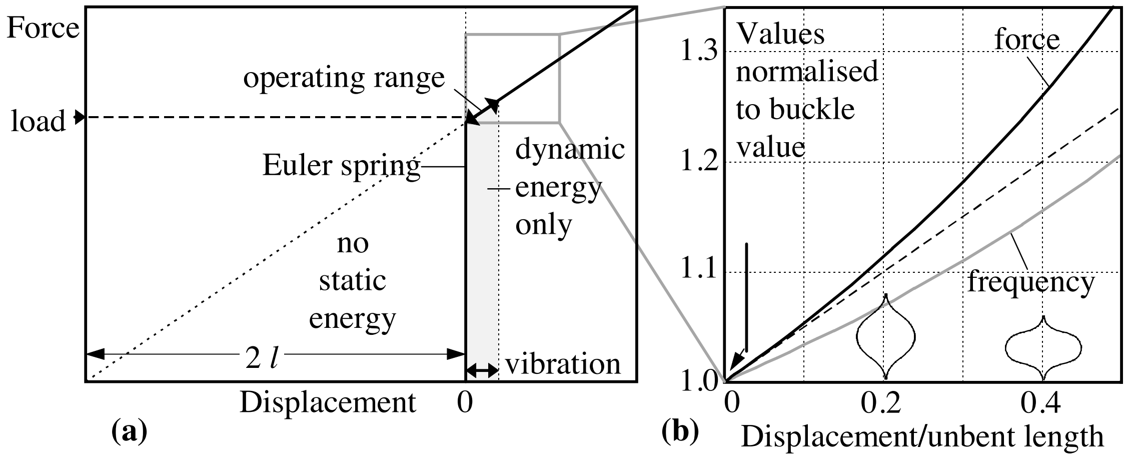

Figure 2.8 Force - displacement characteristic of elastic buckling in Euler springs.

2.2.4 Resonant Frequency

The ratio of the derivatives of equations (2.1) yield the spring-rate $k_s=\delta F/\delta x$ and this together with the mass $m=P_{cr}/g$ required to critically load the spring determine the resonant frequency $\omega=(k_s/m)^{½}$ of the system. Remarkably when an Euler spring as shown in figure 2.7 is critically loaded so that it just starts to buckle, then the vertical resonant frequency of the mass-spring system depends only on the length of the spring (and g = accel due to gravity). This might be expected from familiarity with the fact that a linear spring in vertical suspension and a pendulum share the same equation for resonant frequency - $\omega=(g/l)^{½}$ - with l being the extension of the linear spring under load or the length of the pendulum as the case demands. However for the Euler spring an extra factor of 2 is involved - the expression becomes $\omega_e=(g/2l)^{½}$ - so that the suspended mass moves as though it was suspended by a linear spring which had been extended by an amount *twice* the length of the Euler spring.

Figure 2.8a shows this initial spring-rate slope in the context of the critical load, while the grey rectangle is expanded in figure 2.8b to show in detail how the spring-rate and consequent resonant frequency vary with displacement. It should be noted that the displacement axis in figure 2.8b covers a very large travel range compressing the Euler springs to half their initial length (see spring shapes indicated), whereas a useful displacement (for minimising static energy etc) would typically compress by less than 1% of spring length. It is apparent that in theory the spring-rate remains almost constant for such small displacements.

The resonant frequency relative to that at the start of buckling is also plotted and it can be seen that it only increases by 20% for a 50% compression of spring length.

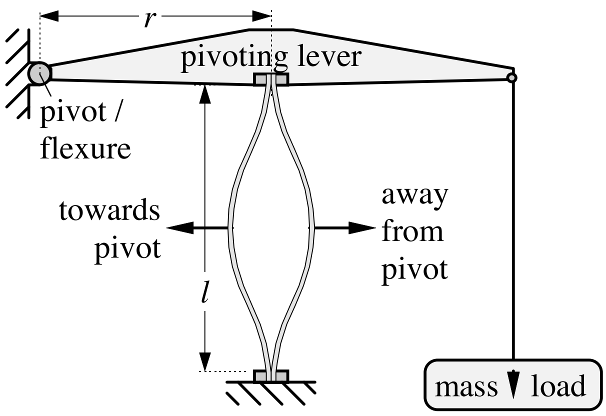

Figure 2.9 Simple structure for mounting and constraining Euler springs.

2.2.5 Supporting Structures

In order to be useful as a suspension device the Euler springs need to be constrained within some structure to limit motion to the desired longitudinal compression. A simple mechanism to provide this constraint is the cantilever or pivoted lever arm shown in figure 2.9. The motion constraint in this case is the arc of a circle rather than a straight line, but it is approximately linear for the small working range required.

The slightly nonlinear motion also provides some advantages. One is that the load can be supported or suspended at a different distance along the lever to that at which the spring blade is clamped - allowing a mechanical advantage ratio to be used to match an available spring to a particular load. This lever ratio may be adjustable to allow a continuous trade-off between supported mass and working range (mass × range = constant fixed by spring blade energy storage capacity). Another advantage if a bell-crank is employed (figure 2.10a) is that the spring blade can be mounted at any desired orientation around the pivot to make use of available space. However the main advantage is the ease with which spring-rate reduction techniques may be applied to rotational motion.

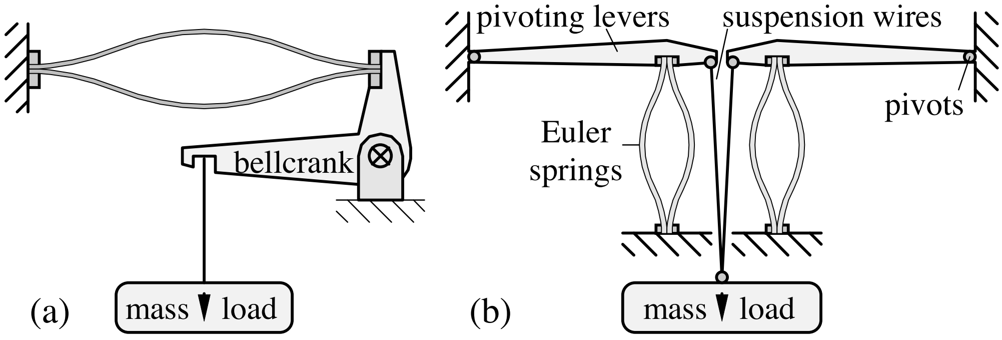

Figure 2.10 Additional structures for application of Euler springs.

2.2.6 Spring-Rate Reduction

As a flat spring blade starts to buckle, the bulk of the blade can be offset in either one of two directions. If it is mounted in the pivoted support structure shown in figure 2.9, then the effect of offset in one direction is markedly different from the effect of offset in the other direction. If the offset occurs towards the pivot then a very low (and even negative = unstable) spring-rate is obtained (curve a in figure 2.11). If the offset occurs away from the pivot then a much higher spring-rate is obtained (figure 2.11d). If a pair of matched spring blades are employed with one going in each direction, then the spring-rate (figure 2.11c) is graphically indistinguishable from what it would have been if it had been constrained to move linearly rather than in the rotating support structure actually employed.

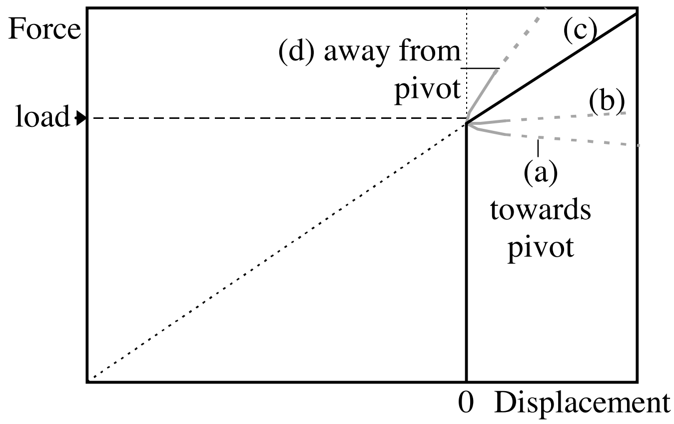

Figure 2.11 Force - displacement characteristic of individual springs.

It follows that by choosing an appropriate ratio between the bending stiffness of the blade(s) moving towards the pivot to those moving away, a suitable mix of figure 2.11a and figure 2.11d can be obtained giving a much reduced spring-rate figure 2.11b. This allows much lower resonant frequencies and also results in reduced Q-factor giving better damping. However a full analysis (to be published separately [Winterflood 2001b]) shows that all the curves apart from figure 2.11c are significantly nonlinear and it is not so simple to reduce the spring-rate over a reasonable operating range with this method.

Figure 2.12 shows two spring-rate reduction techniques which only use the linear case of pairs of Euler springs - one deflecting in each direction. Figure 2.12a has the load attached to what is in essence an inverse pendulum of height h positioned on top of the Euler springs. This structure has merit in that only very small forces are placed on the pivot allowing a lightweight structure and a very thin flexure to serve as the pivot. Taking the most simplistic view of a vertical force f acting in one suspension wire and from the Euler springs on one side, then this inverse pendulum produces an anti-rotational spring-rate of *fh* in the pivot. The normal rotational spring-rate due to the Euler springs is r^2 times the linear rate i.e. $r^2 f/(2l)$. Equating these two expressions gives an approximation of the dimensions required to achieve spring-rate cancellation :- $h/r = \frac{1}{2}\, r/l$.

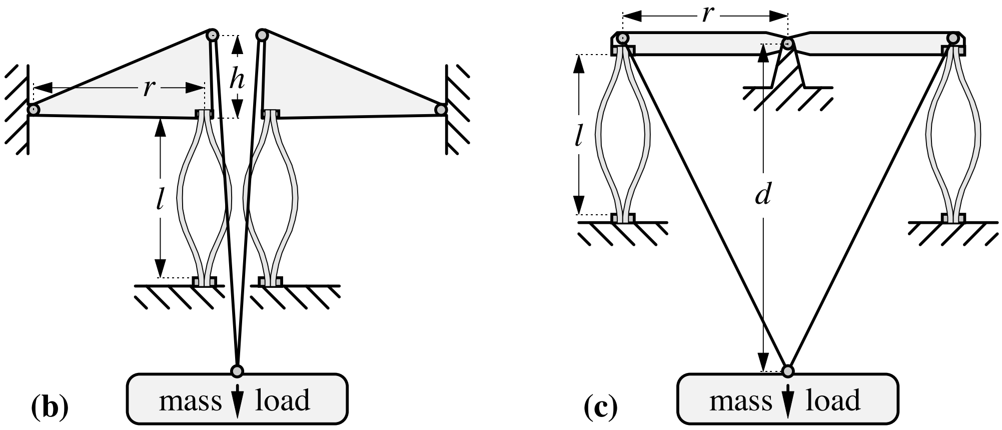

Figure 2.12 Spring-rate nulling in structures using pairs of Euler springs.

Figure 2.12b places one quarter of the total lifting force on the pivot and it is this horizontal force acting on the horizontal arm that produces the anti-spring effect. In this case the horizontal force is fr/h giving an anti-rotational spring-rate of fr^2/h. Equating this expression to the rotational spring-rate due to the Euler springs of $r^2 f/(2l)$ gives dimensions for spring-rate cancellation of simply $d=2l$.

Techniques such as these which cancel almost constant spring-rates against each other, offer easily obtained low spring-rates with large operating range and low Q-factors. The structures that we have investigated (the two in figure 2.10) indicate that the predicted spring-rates near the buckling knee can only be achieved if great care is taken to fix the boundary conditions. Imperfect clamping rigidity and small offset angles causes the sharp corner discontinuity of figure 2.8a and figure 2.11 to be rounded into a gradual asymptotic approach to the theoretical spring-rate gradient at larger displacements. This can cause the spring-rate to be increased by an order of magnitude for small buckling displacements, giving a factor of 3 higher frequency than expected.

2.2.7 A High Performance Vertical Vibration Isolator

Initially a bell-crank version (very similar to figure 2.10a) was built which allowed fine adjustment of various parameters such as the clamping angles for each end of the Euler springs. This allowed many aspects of the theory to be verified [Winterflood 2001b] but the unbalanced structure with its strong coupling to tilt did not give a clean transfer function. So following this prototype, a higher performance balanced version similar to figure 2.10b and shown in figure 2.13 was built and tested.

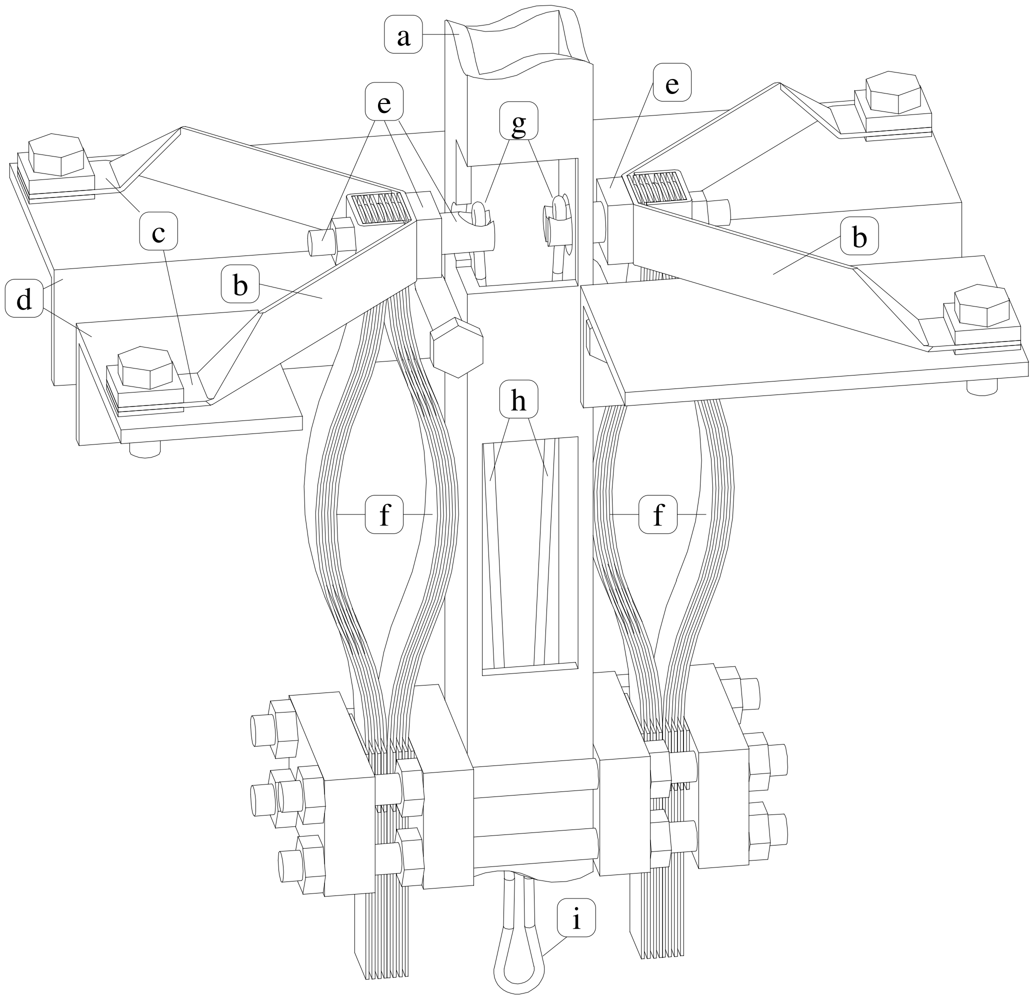

Figure 2.13 High performance vertical vibration isolator.

This particular implementation was designed as a single vertical isolation stage in a cascaded chain of stages hung vertically one below the other. The square tube (a) at the centre is the mounting platform for all the parts comprising the isolation stage (only the suspension components are shown and sections are cut away for viewing). The pivoting arms (b) are a “wishbone” shape for maximum rigidity and minimum mass. They are made of thin sheet and are intended to flex in the flat sections (c) near the clamp which can be thinned for this purpose. They are mounted on pieces of angle (d) clamped on to the central tube. There is a pair of special bolts (e) which pass through a hole in the ends of the spring blades (f) (made of feeler gauge stock) clamping a whole set of blades together and to the wishbones with spacers between each blade. The lower ends of the spring blades are clamped rigidly together with spacers to the central tube. The special bolts (e) clamping the upper ends of the blades pass through large clearance holes in the central tube to limit their motion to a safe operating range and contain a tapered socket to engage firmly with the folded over ends (g) of a loop of suspension wire (h). The suspended mass or following stages are hung by a pin through the centre of this loop of suspension wire (i).

2.2.8 Measured Results

The measurements shown here were for the structure of figure 2.13 but with a total of only two Euler spring blades (one each side bending outwards). Having blades only bending in the outward direction is normally unstable (figure 2.11 curve a) but can be stabilised by the contributing stiffness of the flexing wishbones (figure 2.13c) and bending wire (between g&h). In this case we made the wishbone flexing section (c) quite thick and made it adjustable in length using slots in the wishbone clamps. The spring blades were 0.8 mm thick strips of feeler gauge with a length between clamps of 126 mm and the system was loaded with a mass of 32 kg. It was designed to have a working range of ±1 mm (i.e. 2 mm total from unbuckled to motion limit).

A number of swept sine shake tests were done in air and a typical result is shown in black in figure 2.14. Here it can be seen that the first troublesome internal mode occurs around 400 Hz. Evidence mentioned below indicates that this mode is the mass of the clamping bolts (e) resonating with the wire (h) and that the first blade internal mode is the small one appearing at almost 500 Hz.

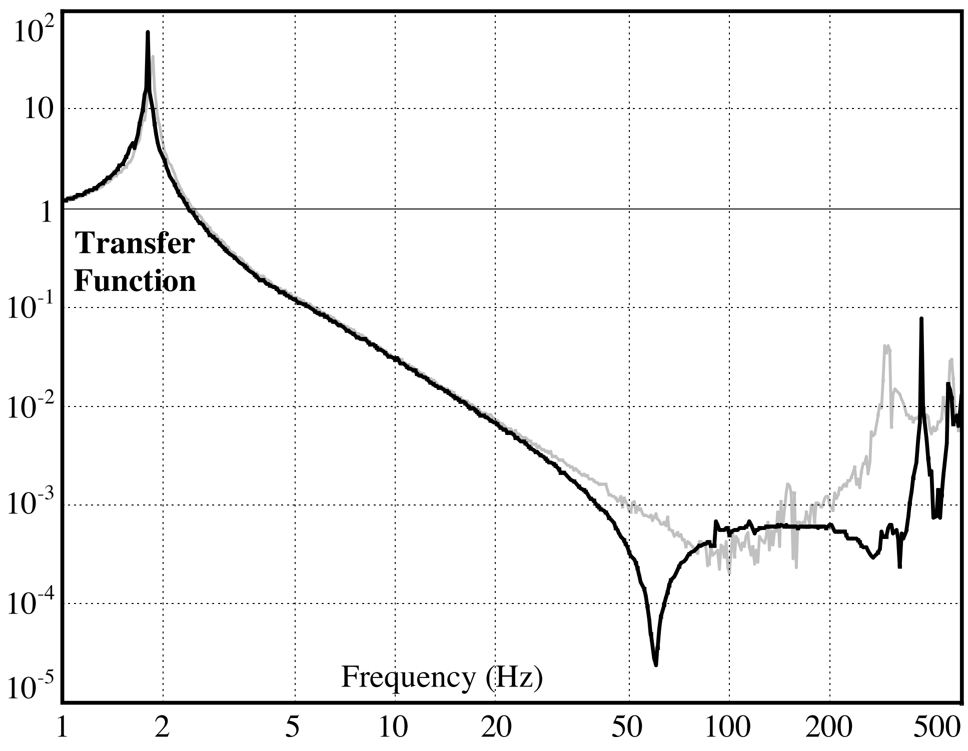

Figure 2.14 Transfer functions of Euler sprung vibration isolation stage.

There is a clear notch at 60 Hz and above that an isolation floor of -65 to -70 dB. The notch is strongly suggestive of dynamic inertia (or centre of percussion) effects. To test this we loaded the wishbones (b) at about their mid point with small blocks of lead (of order 60 g each side) with the resulting transfer function shown in grey. The notch at 60 Hz disappeared as expected and a rounded shape falling to -70 dB around 100 Hz was obtained confirming this suggestion. In addition the first internal mode resonance moved from 400 Hz down to 300 Hz indicating its dependence on the wishbone and clamping bolt mass, while the small mode near 500 Hz remained unaltered. This indicated that the 400 Hz mode was due to the wire and clamping bolts, rather than the Euler spring blades. Normally with very flexible joints at (c), this internal mode would not couple significantly, but in this particular case we had deliberately made the joints (c) very stiff to stabilise the negative spring-rate of the outward bending blades. The reason we were particularly interested in this configuration was because with the design of figure 2.13 it allows the outward bending blades (f) to be clamped very close to the tube (a) and wire (g), minimising forces in the wishbone (b).

2.2.9 Dynamic Effects

Normally with massive suspension structures, significant consideration must be given to the dynamic effect of the structure's inertia (commonly called centre of percussion tuning) [Winterflood 1999]. With the minimal spring mass to suspended mass ratio of this technique, it should be possible to neglect this effect. The transfer function in figure 2.14 shows that some improvement should be gained by appropriate counter-balancing (counter-weights between the 0 g and 60 g values used in figure 2.14 should give a lower transfer function floor level). However ideal counter-balancing is not possible and can only be approximated - due to the fact that the lateral motion of the spring mass is not proportional to compression but varies as its square root. This does not seem to be a significant disadvantage.

2.2.10 Spring Blade Variations

It is apparent that for a given working range, spring material yield stress, and mass to be supported, there are different selections of spring blades which may be applied. For example instead of a single pair of thick blades, two pairs of thinner and shorter blades may be used instead. It is of interest to know whether any advantage is to be gained by using more thin short blades rather than fewer thicker and longer blades.

As an example we may suppose that the thickness of a given pair of blades (e.g. feeler gauge strip 100 mm long, 0.64 mm thick giving a 1 mm working range) are halved so that they may be bent to half the radius at the same stress level. To retain the same working range, the length of the thinner blades may only be reduced by a factor of 0.63. This increases the system resonant frequency by a factor of 1.26. The internal mode resonances scale as thickness/length^2 so in this case they will also increase by a factor of $0.5/0.63^2 = 1.26$. The suspension force decreases and therefore the number of blades must be increased by a factor of 3.17. A rule of thumb might be that if we halve the thickness of the blades, we can shorten them to 2/3 the length but must triple their width or number and we gain an increase in resonant frequency of internal modes of 25% but the ratio of internal mode to system resonant frequency remains unaltered.

It may have been noted, that the preceding discussions have only fully addressed the last of the three areas mentioned in section 2.2.2 - that of reducing the static energy and mass of the spring by producing an almost ideal wall-shaped nonlinear force-displacement relation. There remains the possibility of improving performance further by (1) presetting some initial stresses within the springs, and (2a) moving some of the spring mass away from the centre of the Euler springs towards the ends where velocity of motion is less, or (2b) moving the spring mass closer to the centre of rotation as in a torsion based variant.

Presetting initial stresses may be done by starting with tightly rolled spring material, then unrolling it and bowing it somewhat the opposite way so that when released it lies flat. It must then be loaded in the bowed direction and this treatment should allow considerable extra deflection before yielding. This should allow a reduction in the mass of the springs by a factor of almost 2. We have not tried this, and neither have we explored any aspects of (2a) non-uniform Euler springs. However the gains available from (2b) moving the spring mass closer to a centre of rotation are obvious from a lever example :-

A very effective approach to obtain higher internal modes is to reduce the working range of the springs while retaining the required working motion of the stage with a lever as mentioned and illustrated in section 2.2.5. Halving the length, working range, and thickness of a blade also halves the force. If the working range is restored with a 2:1 lever, then this further halves the force requiring 4 times the number of blades to support the same load. However in this case the internal resonant frequencies are increased by a factor of 2 while the system resonant frequency remains unaltered. This impressive gain is the result of moving the spring mass closer to the centre of rotation of the lever.

2.2.11 Conclusion

Vertical isolation using mechanical springs (as opposed to compressed gas or magnetic suspension) has been plagued by the problem of internal modes bypassing the isolation at relatively low frequencies. The new technique presented here, which is most applicable for suspending constant loads under conditions of very small vibration, is capable of providing orders of magnitude better performance than previous approaches. Some remaining inefficiencies have been identified which may allow even better performance to be obtained from spring material of a given yield strength.

2.2.12 Acknowledgment

We thank Bram Slagmollen of the Australian National University gravitational wave group for assistance with the transfer function measurements. This work was supported by the Australian Research Council, and is part of the research program of the Australian Consortium for Interferometric Gravitational Astronomy.

2.3 Euler Spring Mathematical Analysis Paper

Mathematical Analysis of an Euler Spring Vibration Isolator.

J. Winterflood, T.A. Barber, D. G. Blair

Department of Physics, University of Western Australia, Perth 6907, Australia.

A new suspension technique uses flat springs in Euler buckling mode. This numerical analysis of the basic lever constrained configuration identifies the effect of the main geometrical parameters, and suitable values for obtaining the most useful force-displacement characteristic. The theory is shown to agree well with experimental measurements.

2.3.1 Introduction

A new suspension technique was presented recently [Winterflood 2001a] which minimises the mass of spring material required for suspension against a constant force such as gravity. This has major advantages in reducing the isolation bypassing effect of internal resonances, and the technique also avoids the large pre-loading strain which must generally be established. The technique primarily involves using the post-critical elastic buckling properties of a longitudinally compressed column or spring. Elastic buckling is often analysed in the literature, but almost always in order for unstable conditions to be avoided - very rarely in order to make use of its unusual characteristics. An excellent source of theory together with extensive experimental comparison of the general post-buckling behaviour of various structures is given by Britvec [Britvec 1973]. This paper focuses on one particular structure which we have developed for vibration isolation and analyses the effect of various configuration parameters on the final force-displacement and frequency-displacement relationships obtained.

A typical vibration isolation structure employing Euler springs is shown in figure 2.15 in which the suspended mass is supported by spring elements loaded in compression so that they just start to buckle elastically. A pivoted lever is a simple way

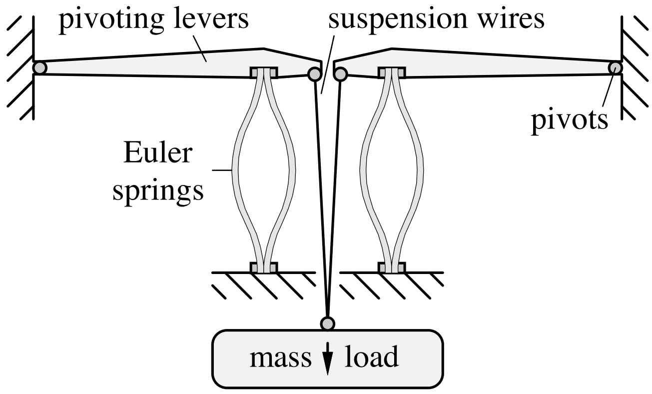

Figure 2.15 The balanced form of the geometry to be analysed.

to keep the motion planar and prevent lateral instability. Two are used in figure 2.15 to form a balanced arrangement to avoid cross-coupling from vertical to other modes of motion. A result of using the pivoted lever mounting is that one end of the clamped flexures is constrained to rotate through a certain angle as the load displaces and the springs compress. This paper analyses the effect that this rotation has on the resulting force-displacement characteristic and how it changes with the main parameters such as lever arm radius and spring launching angles.

2.3.2 Geometrical Model

In order to analyse the performance of the structure shown in figure 2.15, we consider just the right hand lever and only show the flexure which is deflecting towards the pivot. Figure 2.16 shows this essential geometry drawn so that the flexure is approximately horizontal. Lengths have been normalised by taking the length of the flexure to be unity so that all lengths (such as the radius R of the lever) are given as a ratio to the flexure length. Lever pivoting angles are measured with respect to a reference position where the flexure clamping points would form a right angle with the pivoting centre. The main benefits of this suspension technique are obtained when operated at small spring buckling deflections. For this reason the pivoting angles considered only vary very slightly from this square position. The length D is the length of a straight line between the clamping points on the flexure and this line is used as a reference for various angular measurements. In particular when the flexure is stretched straight so that it lies along this line (i.e. D=1 at maximum clockwise pivot position θ_0) then the launching angles α_F and α_P (typically 0 deg) that the ends of the flexure are clamped at are measured with

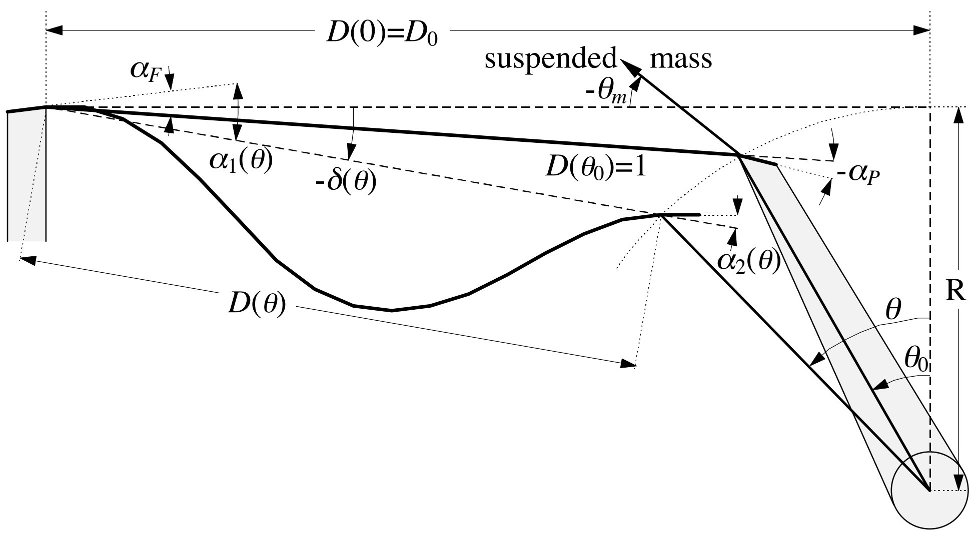

Figure 2.16 Variables for analysis of the geometry.

respect to this line.

When the flexure is buckled to form a curve, it can be deflected either towards the pivot (downward as shown) or away from it. With more than one flexure (e.g. figure 2.15) deflection can occur in both directions and the force contribution from each direction can be infinitely varied from one, to both, to the other (e.g. by varying their thicknesses). A parameter B_P is used to represent the fraction of the force being contributed from the flexure bending towards the pivot.

The main parameters which may be configured in order to obtain desirable system performance are identified in figure 2.16 and are summarised as follows :-

R	Lever radius (normalised to the flexure length).
α_F	Flex clamp angle at fixed end (w.r.t. strip at θ_0)
α_P	Flex clamp angle at pivot end (w.r.t. strip at θ_0)
B_P	Force fraction from strip deflecting towards pivot.

Another parameter which may be varied is the attachment angle of the suspended mass θ_m. If the mass is made to pull in a significantly different direction than along the line of the flexure, some sinusoidal nonlinearity can be obtained. However this parameter, together with a few other possible variables (for instance θ_0) typically have to be pushed to unlikely values in order to contribute a significant effect. The variables are included in the mathematics for completeness but they are left at zero for the plots and their effect is not explicitly analysed. (The angle of the suspended mass θ_m required to null the basic spring-rate is indicated in figure 2.12b of the earlier paper [Winterflood 2001a].)

The geometry of the pivoting lever and end clamps determine three boundary conditions applied to the flexures which dictate the two possible stable states that the flexure can take on. These boundary conditions are :-

D	Distance between clamps as ratio to flexure length
α_1	Left hand clamp angle (w.r.t. straight joining line)
α_2	Right hand clamp angle (w.r.t. straight joining line)

Given the previous set of configuration parameters, these three flexure boundary conditions can be found from :-

$$D = R\sqrt{\left(D_0/R - \sin\theta\right)^2 + \left(1 - \cos\theta\right)^2}$$

$$\alpha_1 = \alpha_F - \left(\delta(\theta) - \delta(\theta_0)\right)$$

$$\alpha_2 = \alpha_P - \left(\theta - \theta_0\right) + \left(\delta(\theta) - \delta(\theta_0)\right)$$

where $\delta(\theta) = \tan^{-1}\left[(\cos\theta - 1)/(D_0/R - \sin\theta)\right]$ (2.1)

and $D_0/R = \sqrt{1/R^2 - (1-\cos\theta_0)^2} + \sin\theta_0$

2.3.3 Euler Strip Analysis

Consider the thin flat strip of spring material, which in an unstressed state is flat, but is clamped at both ends and has a longitudinal compressive force applied so that it is kept in a buckled state with two points of inflection along its length. The shape that such a strip assumes once clear of the clamps is called an *elastica* and the parametric equations describing it make use of elliptic integrals of the first $F(k,\phi)$ and second $E(k,\phi)$ kind where we follow [Reddick 1947] and use k as the modulus and ϕ as the amplitude.

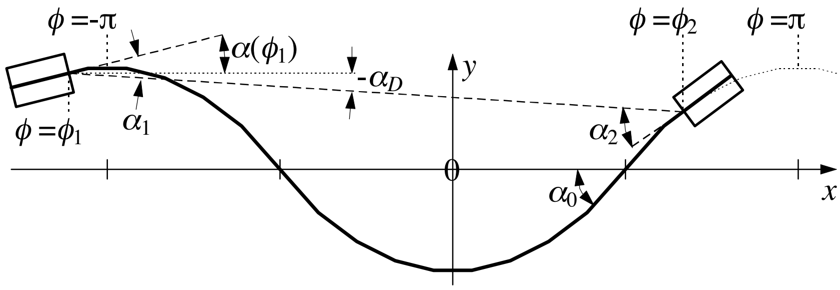

Figure 2.17 Flexure strip boundary conditions.

Figure 2.17 shows this curve positioned on its x-axis which passes through the points of inflection on the strip, and a y-axis through a point of maximum curvature. ϕ becomes the parametric variable which defines points along the curve and x and y coordinates are given by the equations :-

$$x = c\left(2E(k,\phi) - F(k,\phi)\right)$$
$$y = B\,2ck\cos(\phi) \quad (2.2)$$

where B=-1 in equation (2.2) gives the downward bending solution and B=+1 would give an upward arching solution. The curve in figure 2.17 has ϕ ranging from $\phi_1 \approx$-(10π/9) on the left to $\phi_2 \approx$ (2π/3) on the right. The value of c is determined by the flexure's modulus of elasticity E, its moment of inertia I and the compressive force F_x applied along its x-axis :-

$$c = \sqrt{EI/F_x} \quad (2.3)$$

The modulus k indicates the curvedness of the flexure and is related to the maximum angle α_0 occurring between the elastica and the x-axis. Since α_0 is easier to visualise, we use it to specify the curvedness and k is found from it by :-

$$k = \sin(\alpha_0/2)$$

The angle $\alpha(\phi)$ at any point along the elastica is given by :-

$$\alpha(\phi) = 2\sin^{-1}\left(k\sin\phi\right)$$

The arc length L along the flexure between ϕ_1 and ϕ_2 is :-

$$L = 1 = c\left(F(k,\phi_2) - F(k,\phi_1)\right)$$

Using these equations, it is easiest to generate a particular elastica given values for :-

α_0	maximum angle of flexure w.r.t. x-axis of elliptic
ϕ_1	left end termination of incomplete elliptic
ϕ_2	right end termination of incomplete elliptic
B	Bending direction (+1 = up, -1 = down as figure 2.17)

Given this set of shape defining values, the flexure boundary conditions (D,α_1,α_2) can be found by the following calculation steps :-

$$\begin{aligned}
k &= \sin(\alpha_0/2) \\
c &= 1/\left(F(k,\phi_2) - F(k,\phi_1)\right) \qquad (2.4) \\
\Delta x &= 2c\left(E(k,\phi_2) - E(k,\phi_1)\right) - 1 \\
\Delta y &= 2Bck\left(\cos\phi_2 - \cos\phi_1\right) \\
D &= \sqrt{\Delta x^2 + \Delta y^2} \\
\alpha_D &= \tan^{-1}(\Delta y/\Delta x) \\
\alpha_1 &= -\left(\alpha(\phi_1) - \alpha_D\right) \\
\alpha_2 &= -\left(\alpha(\phi_2) - \alpha_D\right)
\end{aligned}$$

where $\quad \alpha(\phi) = 2B\sin^{-1}\left(k\sin\phi\right)$

The inverse problem of finding (α_0,ϕ_1,ϕ_2) given (D,α_1,α_2) is solved numerically by an iterative approach such as by using FindRoot[eqns,vars] in *Mathematica*® [Wolfram 1988]. If B_P is non-integral, then two solutions need to be found - one with B=1 for the flexure bending upwards and one with B=-1 for the flexure bending down. These two solutions are used in proportion to determine the overall force-displacement characteristic.

2.3.4 Obtaining Force vs Displacement

As a by-product of finding values for (α_0,ϕ_1,ϕ_2) that solves for a given set of (D,α_1,α_2), we also obtained a value for c (eqn (2.4)) from which we can obtain a value for the force using equation (2.3). This longitudinal force appears acting along the curve's x-axis without torque at the inflection points of the flexure (placing a perfect pivot at an inflection point would have no effect). The torque on the lever is this force multiplied by the perpendicular offset to the lever pivot. This torque (added to any torque from the pivot spring-rate) can be translated to a force applied in the direction of the suspended mass giving finally the force vs displacement characteristic.

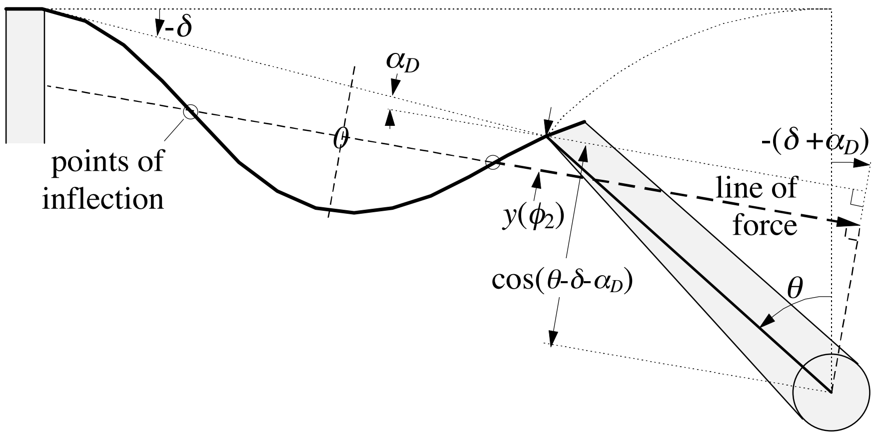

Figure 2.18 Variables for torque and force determination.

If we normalise the force obtainable from equation (2.3)

$$F_x = EI/c^2$$

to the Euler force of a straight clamped flexure of length L(=1) as it starts to buckle

$$F_e = 4\pi^2 EI/L^2 \tag{2.5}$$

we obtain the normalised force :-

$$F_n = F_x/F_e = 1/\left(4\pi^2 c^2\right)$$

If this force is applied to the lever along the x-axis of the elastica as in figure 2.18 we obtain a torque T_e on the lever of :-

$$T_e = F_n\left(R\cos(\theta - \delta - \alpha_D) - y(\phi_2)\right)$$

where δ is found from (2.1) and $y(\phi_2)$ from (2.2).

The angular spring-rate of pivoting k_p (which includes suspension point pivoting) provides an additional torque $k_p(\theta\text{-}\theta_k)$ dependent on the angle θ away from neutral position θ_k. The sum of these two torques appears as a force F_m in the direction of the suspended mass (θ_m in figure 2.16) :-

$$F_m = \left(T_e + k_p\left(\theta - \theta_k\right)\right) \big/ \left(R\cos(\theta - \theta_m)\right) \tag{2.6}$$

It is apparent that the suspension point pivoting (i.e. bending of a suspension wire as the lever rotates) will contribute a much stiffer angular spring-rate, than the lever pivot which is almost unloaded and can be made very thin and flexible. An expression may be obtained for the normalised angular spring-rate k_n from a long round suspension wire strong enough to support the mass m, as a ratio to the angular spring-rate $k_e = r^2mg/(2l)$ of an Euler spring length l acting at the same suspension radius r :-

$$k_n = k_p / k_e = l\sqrt{Emg/\pi} \big/ \left(r^2\sigma\right) \tag{2.7}$$

where E = wire modulus and σ = max wire stress. Using typical values for spring steel wire (loaded to 800 MPa) and length = radius = 100 mm, this ratio is approx 3% for a suspended mass of 10 kg or 10% for 100 kg. This poorly defined value is sufficiently small to be neglected in comparison with the Euler spring-rate contribution for qualitative purposes. The force value from (2.6) is plotted against mass displacement given by :-

$$x_m = R\left(\sin(\theta - \theta_m) - \sin(\theta_0 - \theta_m)\right)$$

to obtain the following plots. In all of these plots the x-axis displacement and radius R is given as a ratio to the spring blade length, and the force or frequency is given as relative to the value for a perfectly clamped blade under parallel buckling (i.e. eqn (2.5) for force or $\omega = (g/(2l))^{½}$ for frequency).

2.3.5 The Effect of Bending Direction and Lever Radius

Figure 2.19 indicates how the force-displacement characteristic varies with bending direction and lever radius. When the lever radius is infinite (i.e. parallel straight-line compression), then the central, almost straight dashed line is obtained. Its gradient (or spring-rate) has a normalised value of 1/2.

As the lever radius reduces, the curves for each direction of bending (towards or away from the pivot) separate as negative and positive roots of a square root function (to first order). These curves take the form of approximate parabolas of increasing focal length straddled about the basic Euler spring-rate for parallel compression (the dashed line in the centre). The upper part of the curve corresponds to bending away from the pivot (giving a stiffer spring-rate) while the lower part of the curve corresponds to bending towards the pivot (often giving a negative spring-rate). Varying the proportion of force contributed by springs bending in each direction allow any intermediate curve

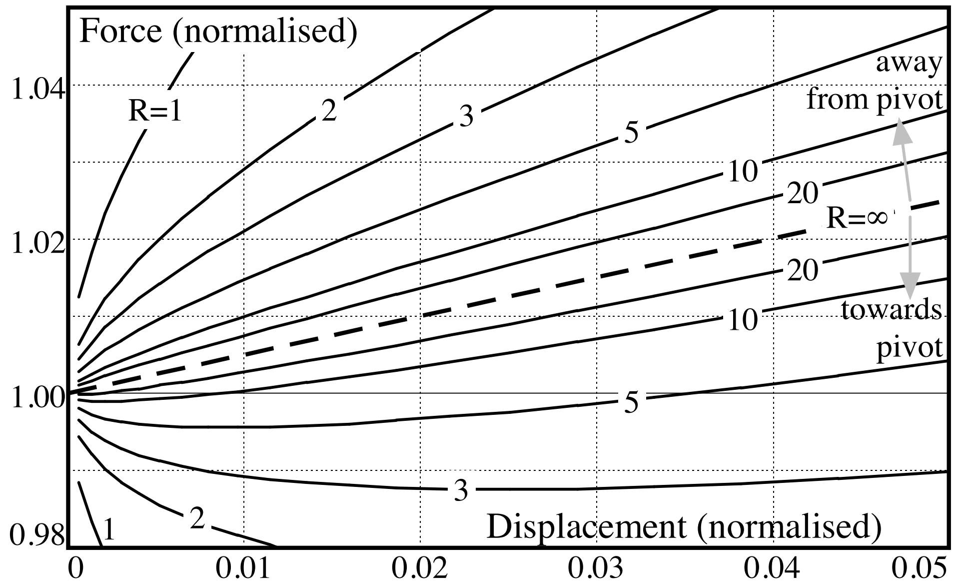

Figure 2.19 Force-displacement dependence on direction and radius.

to be obtained (i.e. the curves belonging to levers of greater radii). If equal springs are bending in both directions then a force-displacement relationship graphically indistinguishable from the infinite radius one is obtained.

The significance of these parabolic shaped curves for the isolation technique is that there is no curve that can be selected by radius and spring direction ratio that will give a significantly lower spring-rate than the dashed central line over a reasonable displacement range starting from buckling. Any curve that promises a reasonable range is significantly unstable close to the buckling point. Also if it is desired to only use springs bending towards the pivot (as we tried in [Winterflood 2001a]), then a rather large radius (between 5 and 10 times the spring length) would normally be required.

2.3.6 The Effect of Non-Zero Flexure Clamping Angles.

Figure 2.20 indicates how the force-displacement characteristic varies with non-zero flexure clamping angles. The lever radius used for all curves is unity (lever radius = spring blade length) and sets of curves for other radii are not shown but are very similar in general appearance. The dashed line is the infinite-radius or equal-springs-in-each-direction line. The central parabolic curve shown in black is for zero clamping angles and is the same as the R=1 curve of figure 2.19 (with different vertical and horizontal scaling).

The grey curves indicate the force-displacement relationship for various non-zero clamping angles. It does not seem to matter which end of the spring is given an offset angle - they both have almost identical effect. In fact if one end is offset to an angle in one direction and the other in the opposite direction (to form a slight "S" when

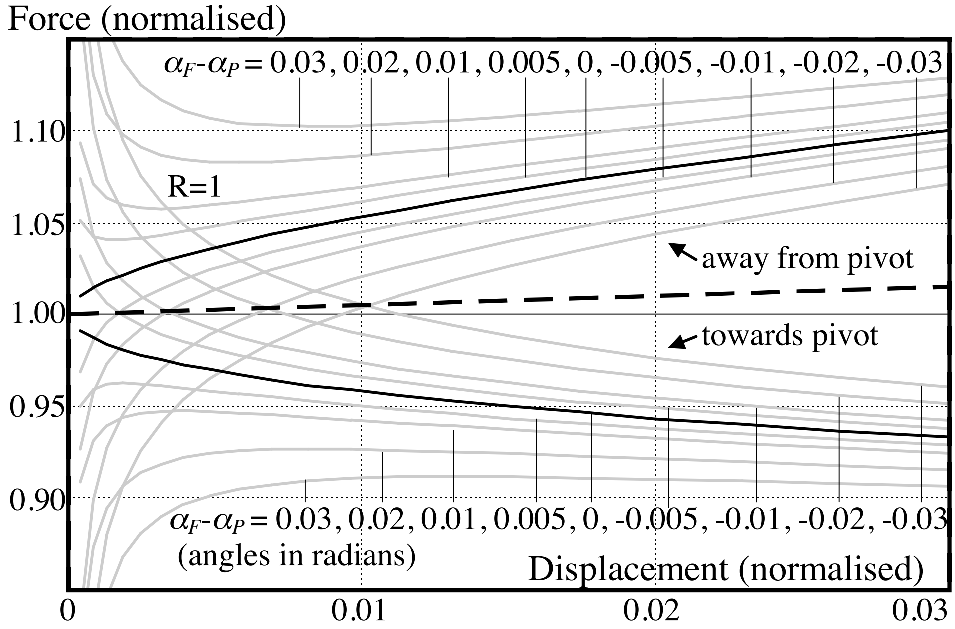

Figure 2.20 Force-displacement dependence on clamping angle.

uncompressed), then a curve graphically indistinguishable from the zero angle clamping curve is obtained. For this reason the curves on the plot have been labelled with the difference in clamping angle α_F-α_P because this is the value that counts.

Although not shown, the zero-offset black curves reach unity at almost zero displacement before falling sharply to zero below that. The convex-up grey curves reach zero by gently falling away, while the others reach zero somehow by passing unstably through infinity (in fact the force is quite limited by the compressibility of the neutral axis which is neglected in this mathematical model). It may be seen that clamping angle offsets which launch the blade towards its main deflection direction “rounds” the left side of the curve, while a launch in the opposite direction to its main deflection “sharpens” it into a rise to infinity (so that the blade wants to “snap-through” to the opposite deflection as it becomes uncompressed).

What is significant in these curves is that if the springs are chosen to bend predominantly towards the pivot (to provide a reduced spring-rate), then the normally unstable section of the curve near zero compression can be ameliorated with a non-zero clamping angle to round the left side of the curve. In fact by choosing the best radius (or in/out bending ratio) together with some clamping angle, it is possible to obtain a spring-rate as low as desired which is also a minimum (force-displacement point of inflection) at the operating point - giving stability above and below. This optimisation will be explored in section 2.3.8.

2.3.7 *Experimental Verification*

A bell-crank based isolator stage shown in concept in figure 2.21a, was built as shown in figure 2.21b to allow verification of the suspension technique. The suspension arm of the crank was 140 mm long and the sprung arm 70 mm giving a 2:1 lever ratio. The crank was pivoted with a flexural pivot (Lucas “Free-Flex” 6016-600) to allow friction free pivoting and the clamps for the spring blades were made of brass and are shown in more detail in figure 2.21c. The clamps were attached with a single bolt allowing clamp angle adjustment and the angle was set with a screw adjustment protruding vertically from the clamping block. Horizontal screws acting on the block allow fine setting of initial buckling, and a motion limit adjustment prevents overstressing.

The clamps were fitted with 4 spring blades of unmodified feeler gauge stock 0.635 mm thick (0.025 in) and 12.7 mm wide and clamped tightly with bolts between blades 1 and 2 and between blades 3 and 4. All of the measurements mentioned in this report were made with a free blade length between clamps of 100 mm and with a total loading of approximately 45 kg with small amounts of mass being added or removed to obtain force vs displacement measurements. The suspension wire (most readily to hand) was greatly oversized and should contribute an additional spring-rate of 52% of the Euler rate, and the flex pivot another 3.5%. Initially the resonant frequency and Q-factor were

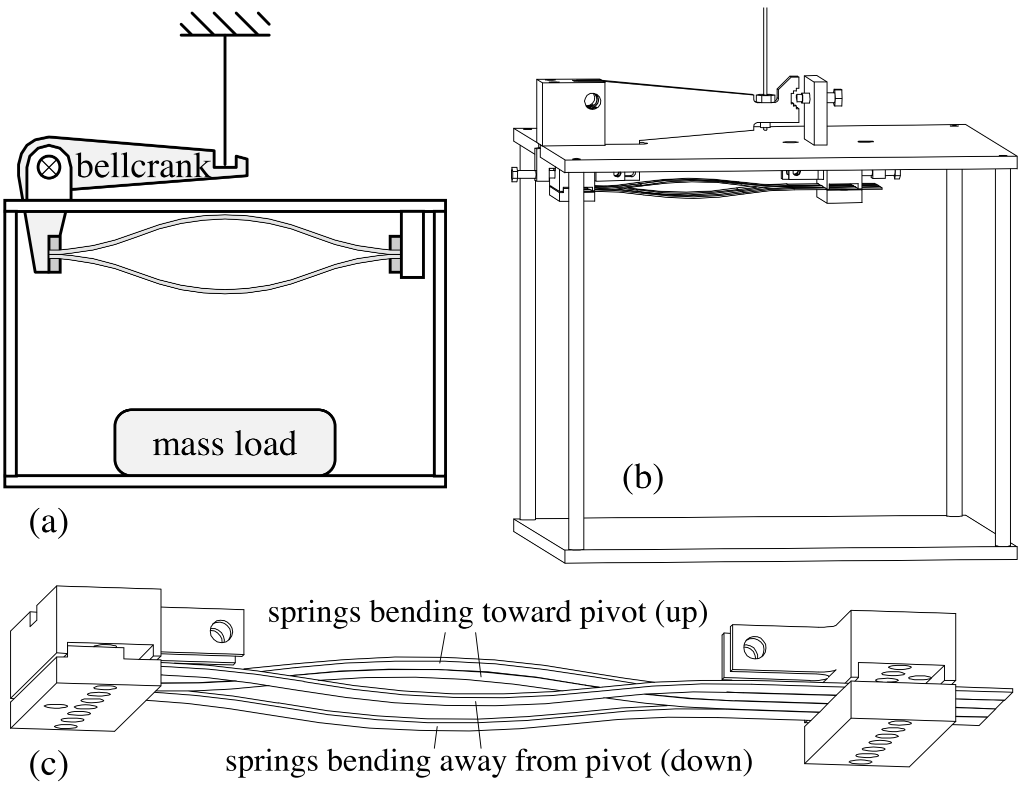

Figure 2.21 (a) Schematic and (b) Drawing of prototype isolator for testing the Euler spring technique. (c) Detail of spring clamping arrangement.

measured from a ring-down with two springs up and two down. In theory, taking into account the 2:1 lever ratio and including the additional suspension spring-rates, a resonant frequency of 0.98 Hz should be obtained. However a frequency of 1.77 Hz was measured with a Q-factor of 50. This Q-factor suggests that if a spring-rate reduction technique was applied, it should be possible to reduce the resonant frequency down to approx 0.25 Hz before the Q-factor approaches unity (Q-factor typically varies as frequency squared in this situation).

We also investigated the characteristics of the structure by applying various masses and noting displacements with various spring up/down settings and clamping angles, and measured the resonant frequency at each setting. Two of the 5 data sets are shown. Figure 2.22 shows the measurements made with clamping angles approximately zero, and figure 2.23 shows the measurements made with α_P = -0.023 rad. For each set a particular mass was applied, and displacement readings taken for the various spring up/down combinations that were stable. The dots at the centre of the short line segments are the force-displacement coordinates while the slope of the line through the dot indicates the gradient found from the measured resonant frequency at that point.

The continuous lines through the points were obtained by the numerical routines outlined above and lines were plotted for all five possible spring up/down settings even though some were unstable and could not be measured by simply applying mass. The theoretical lines had some parameters that were not exactly known which were adjusted for best fit. Initially these only included the zero displacement and Euler critical load for each graph. However it was not possible to obtain a reasonable fit between the zero

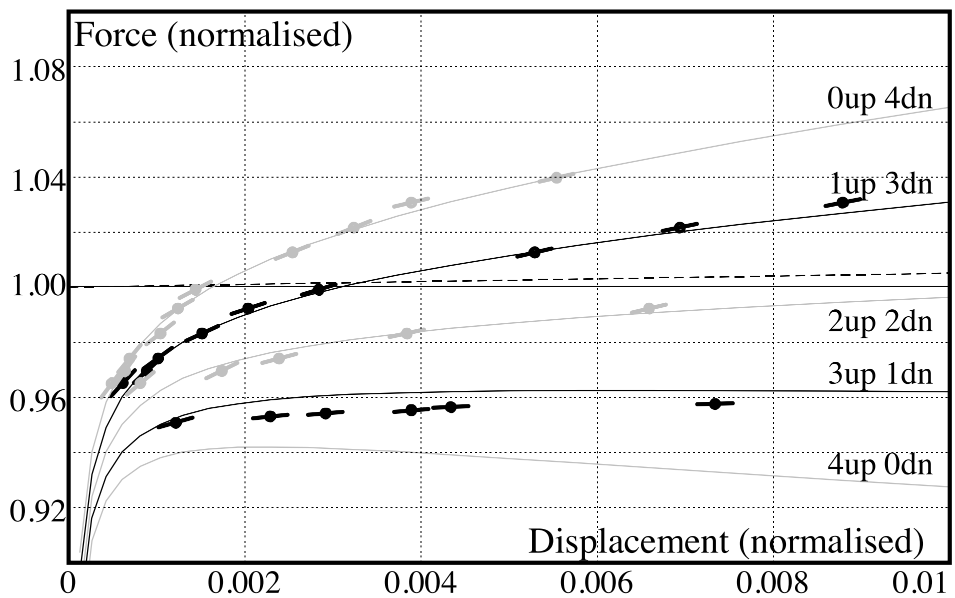

Figure 2.22 Comparison of measured force-displacement with theory.

offset angle data of figure 2.22 and the sideways parabolas of figure 2.19 particularly when the gradients were taken into account. This is obvious by comparing the gradients of the 2up/2down data against the dashed line which is the approximate force-displacement relation expected for this case.

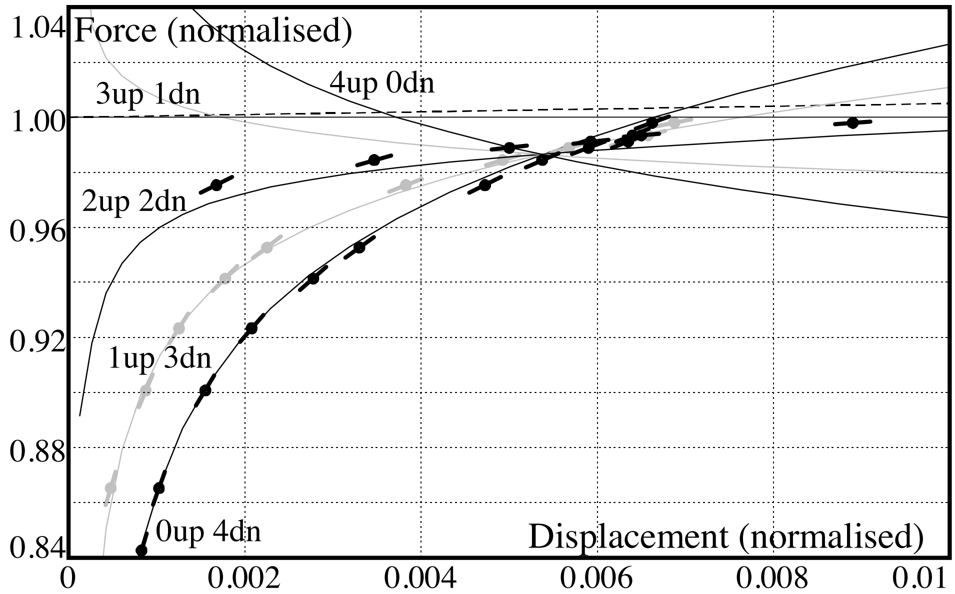

Figure 2.23 Comparison of measured force-displacement with theory.

Our explanation for this obvious departure from theory is that the clamping is non-ideal in that even though the clamping jaws of all blades are at the same angle (being side-by-side between the same jaws) yet it is possible that there is some "slip and stick" effect at the launching edges which allows springs bending in one direction to be launched effectively at a small angle in that direction while blades bending in the opposite direction are launched at a small angle in the opposite direction. Adding a single additional global parameter to account for a small "slip and stick" angle allowed quite good fits to all curves to be obtained. The global fitted value for this angle occurring at both ends of all four springs was 0.0039 rad. This small error (in relation to the angles shown in figure 2.20) results in a spring-rate more than an order of magnitude higher than expected for the 2up/2down case. It was also apparent that this effect had slowly increased with time and usage. Initially the resonant frequency was measured at 1.77, followed by 2 Hz when measuring some transfer functions, and finally 3 Hz when measuring the effects of clamping angles - and this drift seems to have occurred without the spring clamping being loosened or changed.

We are confident from Britvec's analysis [Britvec 1973] (2-16) that the non-extensible theory used above is applicable and the theoretical curves obtained are correct. We note that he also found significant departures from theory when measuring fixed-end columns (e.g. figs 5-4.3 & 5-4.5) whereas his measurements on pin-ended columns (section 3-4) were more in agreement with theory. However due to better clamp design, his departures were small enough to be put down to “imperfections”.

2.3.8 Optimising Spring-Rate Cancellation Range

From an examination of figure 2.20 and its discussion in section 2.3.6 it is evident that there exists a unique solution for the lever arm radius and clamping angle offset that will provide a turning point minimum in the spring-rate at some operating displacement. Of particular interest is this solution for the simplest case which only uses blade spring(s) bending towards the pivot. These solutions are plotted in figure 2.24.

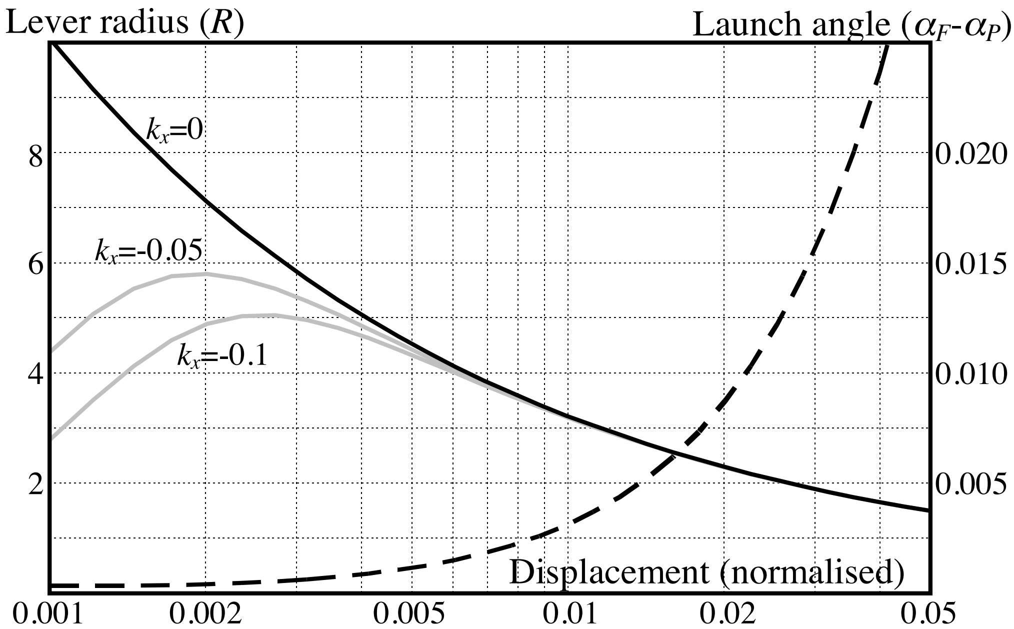

Figure 2.24 Turning point solutions giving radius and launching angles dependent on operating displacement.

The solid lines give the lever radius (as a ratio to blade length) required to obtain a minimum resonant frequency at the operating displacement shown on the x-axis, while the dashed line gives the launch angle required. Since these angles are small and the required value is complicated by the exact boundary conditions such as the difference between clamping and launching angle mentioned in the previous section, this angle needs to be made adjustable. The black lines are the solutions to obtain a zero vertical spring-rate from the combined effects of lever radius and launch angle, while the grey lines are solutions for slightly negative spring-rates to give values of -5% and -10% of the balanced Euler spring-rate $mg/(2l)$. Some small negative value of spring-rate may be desired in order to compensate for the finite spring-rate contribution of pivot and suspension wire bending (see discussion of eqn (2.7)).

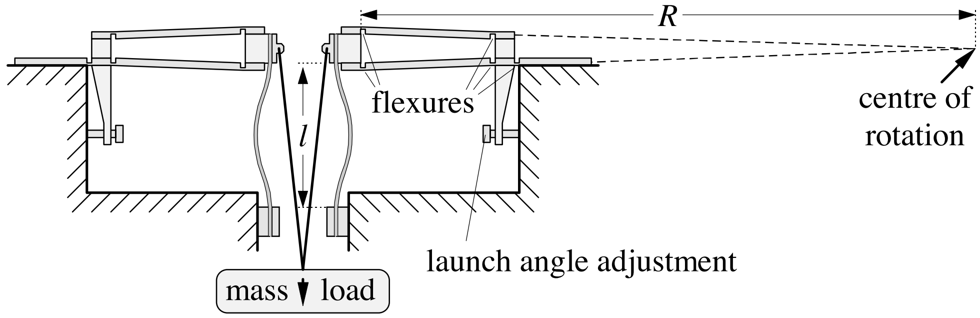

Figure 2.25 Schematic of a reduced spring-rate isolator.

It is apparent from figure 2.24 that the lever radius required for typical displacements in the 0.005 to 0.01 range are quite large and impractical for the simple lever structure shown in figure 2.15. However using a trapezoidal shaped linkage structure for the lever as shown in figure 2.25 allows a very long radius to be obtained from short structural elements. The radius is readily preset by choice of the height difference between the front and back blocks of the trapezoid. The launching angle of one end of the spring blades is then finely screw-adjusted to obtain the desired low resonant frequency.

The frequency tuning curves for one particular operating point (Euler spring compressed by 0.65% of its length) have been calculated and are shown in figure 2.26. The black curve indicates how the resonant frequency varies as a function of displacement for the zero spring-rate solution calculated in figure 2.24. The grey curves indicate how this curve is approached as the launching angle approaches the correct value for zero frequency, ε being the angle away from this correct value. Bearing in mind that the limited Q-factor prevents usefully reducing the resonant frequency below

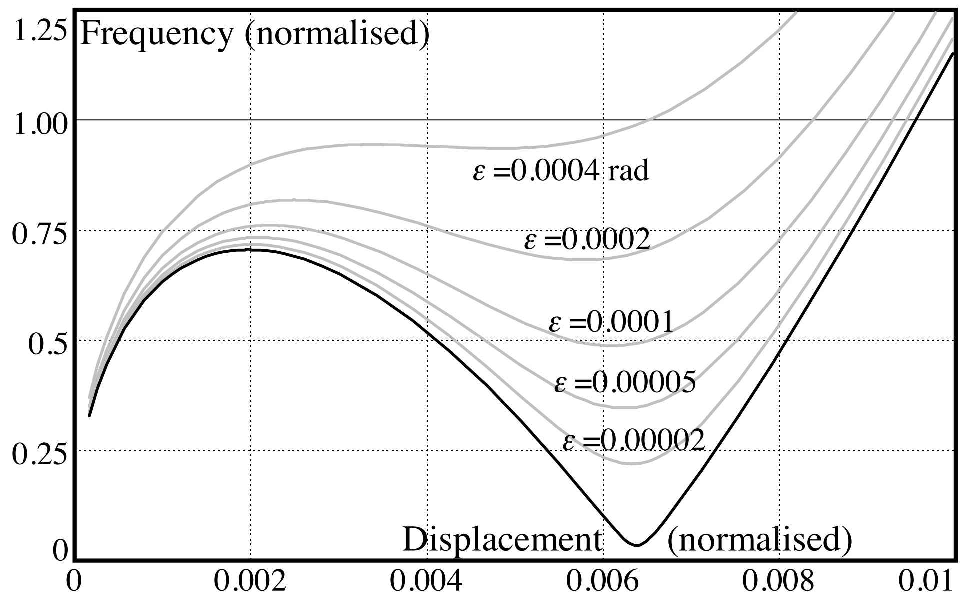

Figure 2.26 Frequency tuning for an operating displacement of 0.0065.

about 0.25 of the unreduced Euler value, it seems that a useable but quite limited working range can be obtained with the resonant frequency only rising to 0.75 of the Euler value.

2.3.9 Conclusion

The effect of the geometrical parameters on the force-displacement characteristic of a simple lever constrained Euler spring mechanism have been explored by theory and experiment. Apart from the case with equal springs bending towards and away from the pivot (which is almost identical to parallel compression and has an almost straight-line characteristic), the nature of the nonlinearities obtained are quite unsuited to spring-rate reduction.

A technique to use the radius and clamp angle effects together has been presented which allows a reduced spring-rate characteristic to be obtained together with a simplified (single-direction) spring arrangement, albeit with a limited working range. Given that the technique is most usefully applied to conditions of constant load and small vibrations, and that spring mass can be reduced (and performance increased) in direct proportion to a reduction in working range [Winterflood 2001a], this limited range is likely to be acceptable.

Theoretical spring-rates near the buckling knee are not easily obtained without careful clamp design or provision for launch angle adjustment. We might suggest making the clamp jaws from material which is at least as hard as the springs themselves and which has slightly protruding jaw edges to stress the material almost to yield along the launching edges when clamped tight.

2.3.10 Acknowledgment

This work was supported by the Australian Research Council, and is part of the research program of the Australian Consortium for Interferometric Gravitational Astronomy.

2.4 Postscript

These two papers on the Euler spring technique were both accepted for publication by Physics Letters A in November 2001, but the referee was concerned that they may be a bit long for a letters journal but that he would also like to see the papers published as they were. The editor agreed that they were rather lengthy and pointed out that the number of figures (15 and 12 respectively) was excessive for a letters journal and needed to be reduced. The author has corresponded with the editor agreeing that some figures from the introductory section in the first paper are superfluous and textual description or references are sufficient, but also pointing out that all of the figures in the second paper are essential to its textual content. The present status is that the modifications to the first paper will be made as soon as the author has sufficient time to attend to them, but the textual content will remain virtually unaltered.

A presentation of the Euler suspension technique was made at the 4th Amaldi conference on gravitational waves and a considerably abbreviated version (6 page limit) of the first paper has subsequently been accepted as a special issue article for publication in Classical and Quantum Gravity [Winterflood 2001c].

Having worked significantly now with this technique, the author has become very enamoured by it indeed. One thing which never fails to please is the fact that such a large mass can be supported by such small, light, and easy to obtain strips of spring steel. The second thing is that there is no struggle with powerful springs trying to get them clamped in position and strained up to support a heavy mass or additional forces applied in order to produce spring-rate nulling. The structure is assembled with ease and doesn't move at all as load is added until its full load is reached. At that point it just starts to bounce gently. Truly a pleasure to work with compared to some other techniques!

3. Passively Self-Damped Chain

3.1 Preface

After struggling to obtain adequate damping of a high Q-factor suspension chain, the author longed to come up with a new design in which the masses just hung "dead" after any disturbance without complex and troublesome motion-sensing and feedback control loops. After all why should one have to *artificially generate friction* with a complex control system when it is so prevalent in nature and never gives stability problems!

Initially one might think that the problem is just that one has tried to obtain too much isolation with too few stages necessitating 2nd order stages with high Q-factors to obtain $1/f^{2}$ isolation per stage above cutoff. The obvious solution being to simply have twice as many first order stages instead. However the problem with such a simple suggestion is always the difficulty of obtaining a suitably frictional joint (for pendulums) that is vacuum compatible and not subject to noisy forms of friction such as stiction.

It was realised early on that if a suitably lossy joint could be created, it would be quite easy to simulate structural damping with it (which falls off with $1/f^{2}$ even when strongly damped) by putting it in series with a lossless joint which is strongly sprung (i.e. by joining the lever arm of the lossy joint with a short pendulum wire - see section 1.4.4). With this modification only the original number of stages are required and not twice the number.

With the dream of obtaining a really "dead" chain the author gave serious consideration to a 2-D gimbal type pivot with wire cut slots (similar to figure 3.8(a) but much smaller) and some sort of putty filling the slots and sealed in with metal bellows. To find a suitable material he built a simple apparatus and measured the viscosity of "silly putty". This unusual compound and a few others like it offer strong viscosity at low rates of shear and yet are quite elastic with low loss at higher frequencies.

Another scheme that was given some thought was the idea of hanging a complete second chain with different pendulum tuning beside (actually concentric with) the first to act as a damping reaction mass using Eddy current coupling between them (some simple versions of this have been done previously - see 1.4.2). At some stage during this brainstorming it was realised that when one considers all 6 degrees of freedom, one actually already has the equivalent of 6 single-degree-of-freedom chains all hanging concentric with each other and all typically tuned to different frequencies. So the problem is not how to fit in another chain, but how to viscously cross-couple between some of the 6 chains already present !

Once this was realised, very rapid conceptual progress was made. It was quickly realised that x displacement and θ_y rotation form a perfect pair for cross-coupling, as do y and θ_x. The remaining z and θ_z are also a good pair and the author came up with a neat but more complex scheme to cross-couple them also (see 1.4.6). To do this properly involved discarding the single wire suspension and replacing it with several widely spaced wires in order to get θ_z rotational resonances high enough to be able to usefully couple to them and use them for damping.

After the pleasure of solving the problem, reality returns and something workable has to be built. At this point it was realised that there are better ways to solve the z damping problem and θ_z is always so soft and slow it has scarcely been a problem in any case (settling time after a disturbance can be a problem but normal running is not). If low frequency vertical suspension is strived for by any low spring-rate technique then the vertical Q-factors automatically become very low and no extra damping should be necessary. For these reasons the idea of cross-damping the z modes with the θ_z fell by the wayside. In hindsight this was partially because the Euler suspension with its natural lever arm constraint had not yet been thought of, and the bow-spring technique (see figure 1.14c) that was being investigated at the time did not lend itself to multiple wire suspension.

The next problem was finding an efficient way to analyse this type of structure to be able to play with parameters manually or automatically and iteratively adjust them to obtain the best damping. The state-space methods which had been used up to this time were quite inadequate and too slow for generating transfer functions for a many staged system with many algebraic parameters left in. This was the incentive for coming to grips with and mastering the transfer matrix methods (see section 8.2). Using this method transfer functions are generated almost instantly regardless of number of stages and even if all parameters are algebraic.

In order to determine how many and what sort of magnets and thickness of copper plate would provide the required viscous coupling, experiments were done with the help of a summer scholarship student Joseph Andersen to measure eddy current damping as a function of copper plate thickness, air gap, etc.

Eventually the author completed a set of engineering drawings and had a set of prototypes built. Just hanging one up with a mass below it and giving it a push was sufficient to show that the damping was working even if not very effectively. Although transfer functions of such a single stage arrangement have been measured by other students, to date no careful adjustment to obtain and verify optimum damping has been attempted.

The author initially termed this new damping technique "cross-coupled mode damping" or "cross-mode damping" for short. The more descriptive and better name of "self-damping" was suggested by his supervisor. Other than the name, the following chapter is entirely the work of the author.

3.2 Self-Damping Concept

3.2.1 Introduction and Motivation

The vibration isolation systems that are typically applied in gravitational wave detectors make use of all metal springs and masses in order to be compatible with high-vacuum requirements. This usually means that all the vibrational modes of the structure have low energy dissipation and high Q-factors. This is useful at high frequencies (>10 Hz) since it provides 2nd order filtering action per stage (isolation improves as $1/f^2$ per stage), but is very troublesome at the natural resonant frequencies of the pendulums or springs and masses forming the isolation chain (~1 Hz). These high Q-factor normal modes need to be damped so that the isolation chain doesn't continue to swing and bounce for a long time after being excited, and also so that the broad-band seismic motion which is being isolated at high frequencies is not resonantly amplified at the normal mode frequencies.

This damping has been achieved by various methods in the past (see section 1.4.2), all of which add significant mechanical and/or electronic complexity to the system. The method presented here is relatively minimal in its addition of complexity (only adding one gimbal joint between each stage's main mass and its vertical suspension spring) and uses simple vacuum-compatible passive eddy current damping.

The basic approach is to viscously cross-couple existing modes on the same isolator stage which have different degrees of freedom. Thus the rocking (or tilting) modes which are inevitably coupled to a small extent to the horizontal translational modes, may be deliberately engineered to be strongly and viscously coupled and to have mode frequencies that differ sufficiently to optimise damping. This overcomes two problems at once - damping both the main horizontal modes as well as the rocking modes.

3.2.2 Concept Development

The progressive development of this scheme is illustrated in the following figures. The normal pendulum link in figure 3.1(a) is made rigid in figure 3.1(b) and the large mass that is normally suspended at the end of the pendulum link is made into a dumbbell shape to emphasise its large rotational inertia. It is flexibly attached to the rigid pendulum by a soft spring pivot so that it is not forced to tilt as the pendulum swings, but can rock or tilt relatively independently at another frequency. Now to apply self-damping an angular dashpot is added between these two motions to obtain optimal damping one against the other. This structure can use x-y pivots and a doughnut shaped mass to operate in both horizontal degrees of freedom simultaneously, damping x pendulum motion against θ_y rocking and y motion against θ_x rocking. Vertical suspension is retained as telescopically extensible pendulum rods at this stage in the concept although this is far from practical.

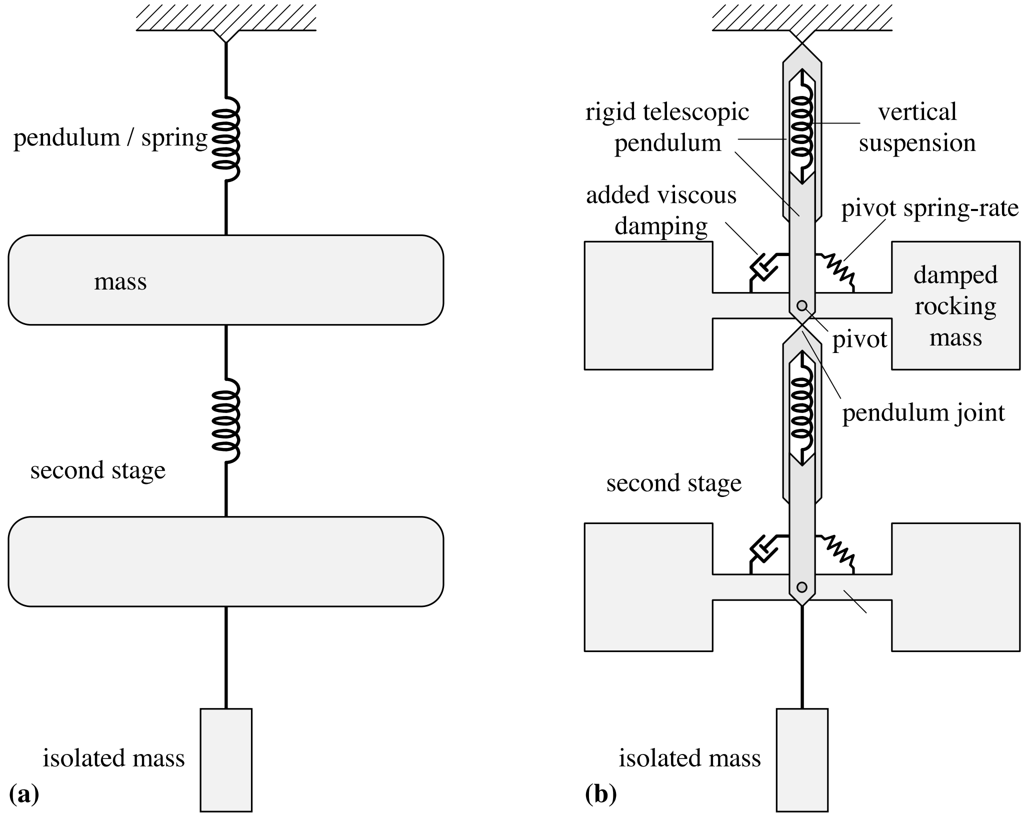

Figure 3.1 (a) Schematic of a typical double stage isolation chain, (b) a horizontally self-damped chain.

3.2.3 Practical Considerations

A problem arises in trying to convert the simple concept of the rigid extensible pendulum rod in figure 3.1(b) into a practical structure. The requirement of the pendulum being rigid but extensible demands a structure which produces constrained linear motion. Such a structure is unavoidably massive in comparison to a simple suspension wire. This means great care must be taken to ensure that the pivoting attachment point supporting subsequent stages is located at a centre of percussion (CoP) to avoid dynamic coupling at higher frequencies (see section 1.6 for CoP description). The fact that the stage is extensible (comprising a vertical suspension element) adds to this problem because the attachment to the subsequent stage needs to *remain* at the CoP regardless of extension. This is not impossible but is a significant engineering challenge and demands a complex structure if the vertical suspension range is to be significant.

A solution to these problems is shown in figure 3.2(a). It consists of breaking the isolation stage into two sections - one being a simple pendulum section (with almost massless suspension fibre and hence no problems with CoP tuning), and the second being a shorter, massive vertical suspension section in which the horizontal CoP effect may be neglected (thanks to the simple pendulum section) and the vertical CoP given maximum consideration if necessary. The damped rocking mass is pivoted on this vertical suspension section and the angle between them is viscously damped by magnetic eddy current damping at their extremities as illustrated in figure 3.2(b).

Most methods of vertical suspension can be incorporated since size and mass is not an issue. We have been very pleased with the performance and simplicity offered by our Euler spring technique [Winterflood 2001a] and a vertical suspension section incorporating this technique is also shown in figure 3.2(b).

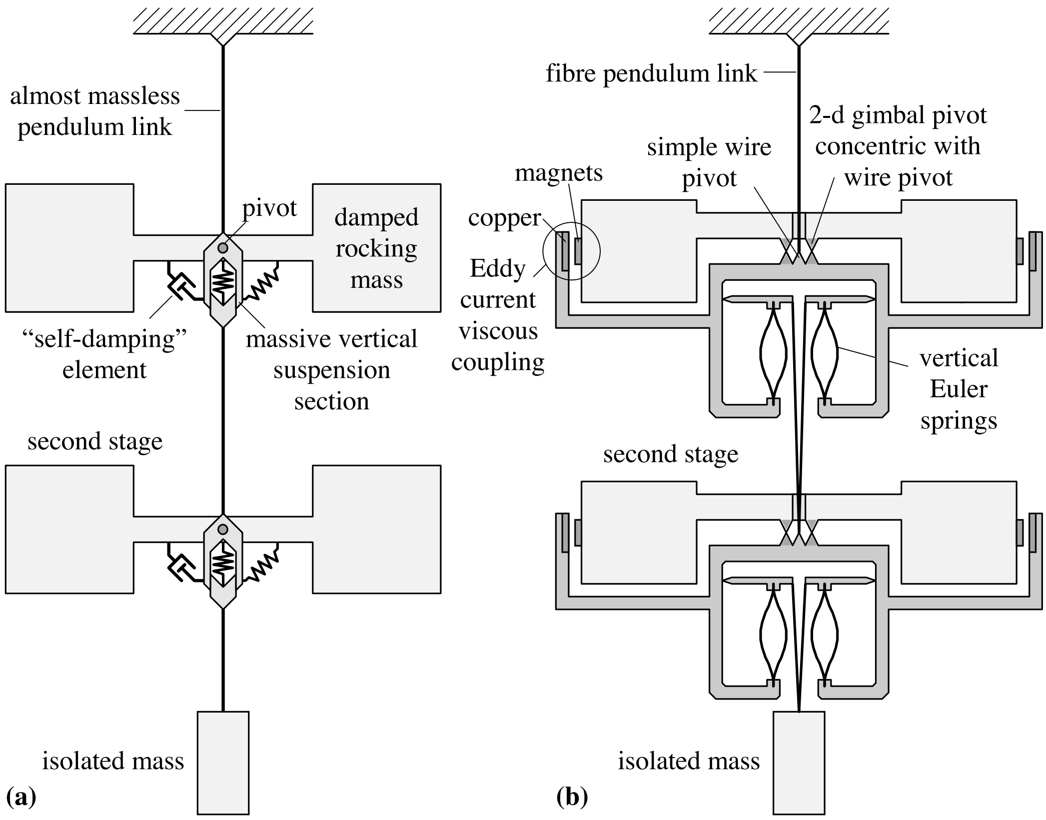

Figure 3.2 Schematic of a more practical self-damped chain, alongside a possible realisation.

3.2.4 System Arrangements

There are various re-arrangements that this basic approach may take and some of these are illustrated in figure 3.3. The most obvious is that the vertical suspension section (the telescopic pendulum sections with in-built springs) and the horizontal section (the massless fibre with massive rocker) may be reversed in order (by turning the stages up-side-down for instance) as is illustrated in figure 3.3(b). This seems less ideal because the massless fibre no longer terminates at the centre of the main large rocker mass forming a simple pendulum section which isolates the vertical suspension structure with untuned CoP below it. However a simple numerical check of what happens to the transfer function when the rocker mass and vertical suspension structure are interchanged indicates that it changes very little indeed. The normal mode frequencies change due to the changes in mass location and their related tensions in the suspension chain, but with the expected damping of these modes, there is very little to choose between them. They are both just as effective at higher frequencies.

One useful consideration is the final mass requirement at the lower end of the chain. For gravitational wave detection, the horizontal isolation is far more critical than the vertical (by a factor of order 10^3) and so the final isolation structure is inevitably a very low-loss horizontal pendulum, and the test mass itself is made from very low-loss material in a compact shape to keep internal modes frequencies as high as possible. It makes sense to also avoid low frequency structural resonances on the mass that the final test mass is suspended from. This requirement suggests that the vertical suspension sections are best located below the rockers on each stage, so that the final two stages need not have any vertical suspension structures mounted on them. It becomes apparent that this arrangement provides no vertical suspension before the first stage mass. It seems best to add a separate vertical suspension section in here so that the horizontal isolation has only one extra stage of isolation instead of two. This arrangement is illustrated schematically in figure 3.3(a).

Another consideration is the attachment at the top of the chain. The simplest

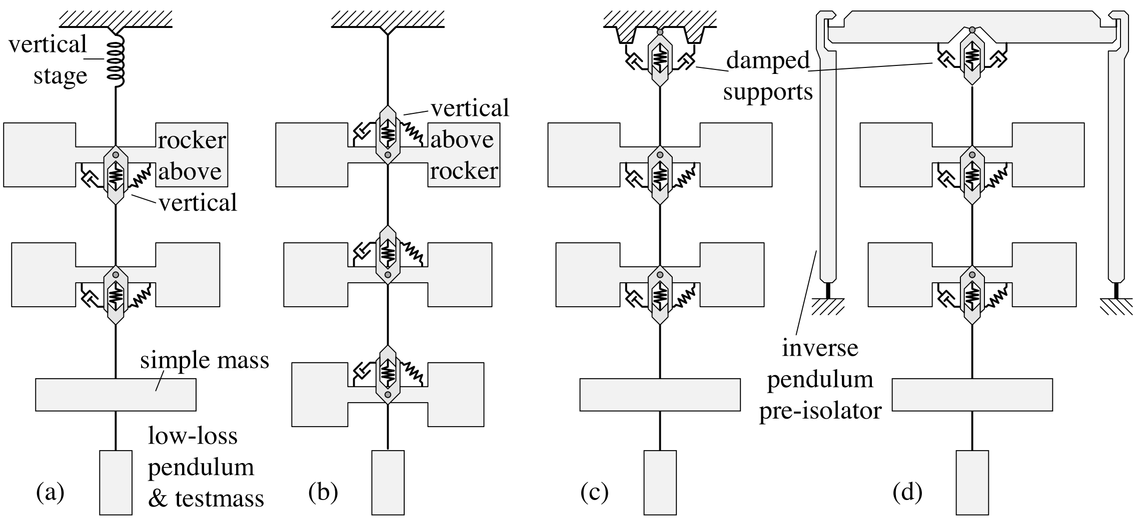

Figure 3.3 Arrangements self-damping stages into system isolation chains.

possible attachment is a simple support (infinitely flexible pivot) which we have assumed for this investigation as it is the most universal and makes the least assumptions about any supporting structure above it. However there are various supporting structures which provide a tilt-rigid attachment to the earth permitting additional torsional viscous damping at the attachment point. These include a simple rigid stand - figure 3.3(c), and some tilt-rigid ultra-low-frequency pre-isolators - figure 3.3(d). We have not investigated these as the damping is very effective when just suspended from a simple support.

3.3 Mathematical Modelling

3.3.1 Model of a Self-Damped Stage

An initial investigation was made assuming that the mass of the vertical suspension section was zero, and the results obtained were very promising. However in practice, the mass (and rotational inertia) of the vertical suspension section will be significant and may actually be some reasonable fraction (say 5-20%) of the mass of the entire isolation stage. Taking the parameters associated with this into account, the full set of parameters used for this more accurate analysis are listed in figure 3.4.

Figure 3.4 shows a schematic of a self-damping stage drawn to work in the single horizontal direction in the plane of the page. The majority of the mass m_r is contained in the large rocker (light grey) and the mass is deliberately positioned at a significant radius of gyration r_r to provide the rocker with a reasonable moment of inertia ($m_r\ r_r^2$). The centre of mass of the rocker is treated as the main reference point of the isolation stage in this analysis.

The rocker is connected by a soft pivot to the remainder of the stage (darker grey)

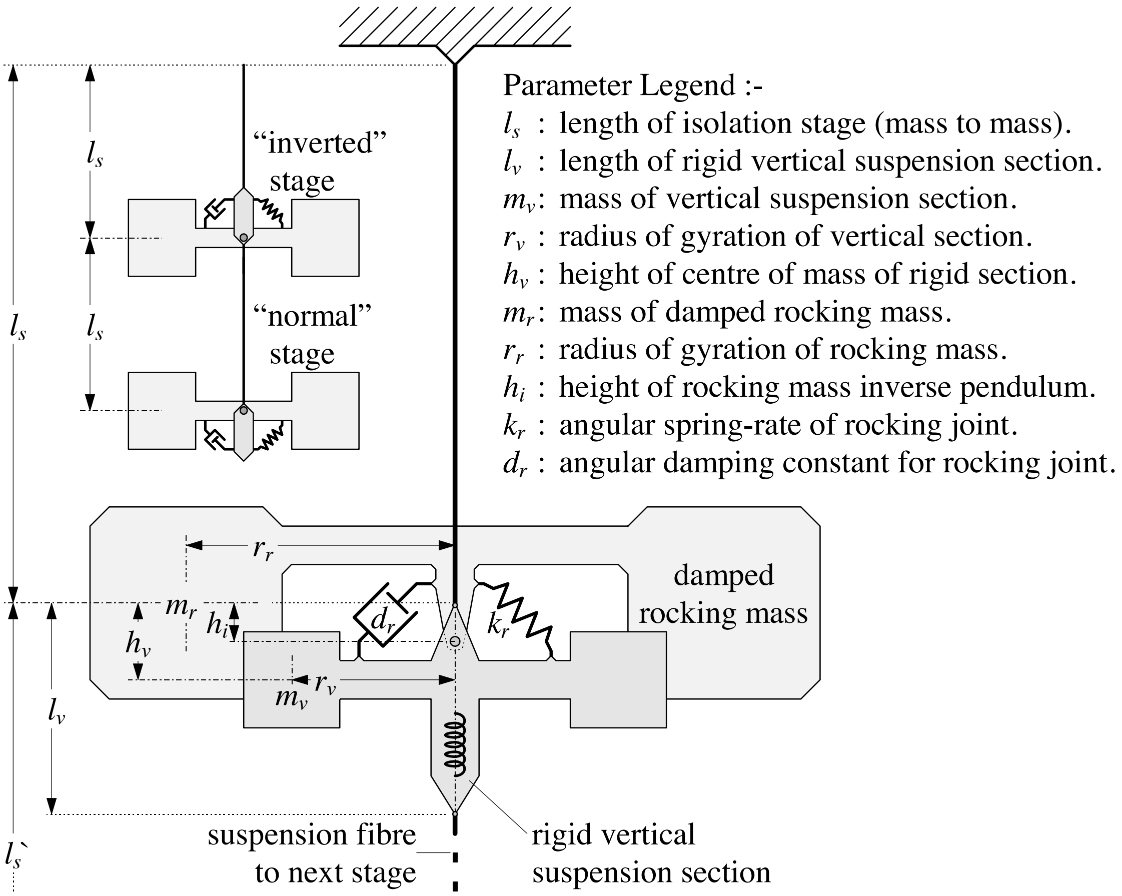

Figure 3.4 Schematic of isolation stage for mathematical analysis.

which joins the upper and lower pendulum suspension fibres. This rigid section between the suspension fibres contains the vertical suspension mechanism. The vertical suspension is not shown in the diagram because it has very little influence upon the horizontal characteristics except that its mass is included in the mass of this rigid vertical section m_v and its moment of inertia is bound up in the value of r_v.

The axis of the rocking pivot should coincide approximately with the centre of rotation of a joint between the rigid section and the suspension fibre in order to minimise motion coupling through the rigid section by see-saw pivoting against the large rocker mass. (This is obvious if one considers making the rocker mass infinite. If the rocker pivot is then located at either of the pendulum fibre attachment points, no motion can be coupled through. If the rocker pivot is located near the centre then forces applied at the top couple directly to the bottom as the rigid section acts as a see-saw.) This means that there are only two ideal positions to place the rocker pivot on the rigid section - the arrangements labelled in the diagram as “normal” and “inverted”. In fact the ideal positions are centres of percussion (CoP - see section 1.6) with respect to each other, on the combination of rigid section and rocker mass, but because the rocker contains the majority of the mass and appears as a point mass due to the pivot, one CoP will always be very close to the pivot.

The rocker pivot has a finite angular spring-rate denoted by k_r and a deliberately large component of viscosity denoted by d_r which is obtained by eddy current coupling at the rocker’s extremities as suggested in figure 3.2(b). The effective spring-rate of the pivot may be reduced by increasing the height h_i of the centre of mass of the rocker with respect to the pivot to make a small inverse pendulum, or the effective spring-rate may be increased by making h_i negative. This height adjustment allows its rocking frequency to be tuned in order to optimise damping (otherwise its spring-rate is determined by manufacture and cannot be easily adjusted).

This isolation stage structure is intended to work as part of a cascaded system with additional stages weighing of order its own mass suspended below it. The tension produced by this suspended mass strongly affects the working of this structure (and the stability of any inverted pendulum content) and so it must be analysed with this loading present. Since this is the case, the angular spring-rate of the joints to the pendulum fibres can be neglected in comparison with the gravitational spring-rate due to the suspended mass below and tension in the chain. In addition the deliberate damping due to the viscously damped rocker totally dominates over any loss that these joints may contribute and so they may be modelled as ideal and perfectly flexible joints.

The system is analysed by means of transfer matrices and is described in section 8.2. The chain of suspension elements is broken up into a cascade of very simple two-port (or four-port) mechanical networks for which the matrix equations are easily determined, then these matrices are multiplied together to obtain the transfer matrix of the entire system. The transfer function is readily obtained from the matrix expressions and can be plotted, or the Q-factors of the poles printed out and parameters adjusted manually or automatically to obtain a minimum.

3.3.2 Simplified Tilt Coupled System

Other than the mass suspended below the stage (giving the tension in the suspension fibres) there are at least six or seven unconstrained parameters per stage which may be adjusted to obtain a well damped result. In order to gain some insight into the effect of each parameter it is necessary to start by analysing a simple system and adjusting one main parameter at a time to observe its effect. For this purpose we consider variations on a simple double pendulum system. Two pendulums of equal mass (performance is independent of actual mass) and of height 25 cm are suspended one below the other. Before including self-damping, we assume pendulum Q-factors of 1000 to limit the heights of the peaks to something reasonable. We look at the case where the pendulums are simple (i.e. no moment of inertia or rocking modes), and then at cases where the pendulum has significant moment of inertia and is given increasing amounts of coupling as shown in the insets of figure 3.5.

It can be seen from the transfer function curves in figure 3.5 that the simple pendulum system has normal modes at 0.76 and 1.84 Hz. When the upper mass is endowed with moment of inertia and allowed to couple to the horizontal motion by spacing its suspension points slightly apart (i.e. forming a short rigid section in the pendulum chain), then an additional parasitic mode appears as a closely spaced pole-zero pair. This is the rocking mode of the pendulum mass. The resonant frequency of this mode may be tuned over a wide range by altering the spacing of the suspension points as shown in figure 3.5 (which alters the angular spring-rate), or by altering the moment of inertia (by changing the radius of gyration). The rocking resonance can be placed above, or below, or made to coincide with either of the two normal modes that form the filter chain.

3.3.3 Vibration Absorber Analogy

This situation is closely analogous to a tuned resonant absorber in a simple two mass system which is routinely applied in all sorts of vibration situations from ships rolling in swell with periods of many seconds, to hundreds of hertz on motor crankshafts and springy structures [Den Hartog 1956]. When a massive structure is subject to a vibrating force, the technique is to attach an extra, usually much smaller mass by means of a spring so that it resonates at the troublesome frequency. If the excitation is narrow-band and at a fixed frequency, then an *undamped* absorber may be used and tuned to the same frequency as the excitation. Due to its high Q-factor it oscillates at much greater amplitude and in anti-phase to the main body almost nulling the force on, and motion of, the main body. This effect produces the notch in the main body response that may be noted on each of the curves in figure 3.5 that couple rocking.

If the excitation is broader band and the vibration problem is due to normal modes of the structure (as is the case here), then tuning such an undamped absorber to null the structure resonance, merely splits it into two resonances, one each side of the original structure resonance. The two new resonances have just as high a Q-factor and would be just as troublesome as the single original resonance - as may be seen in the dashed curve in figure 3.5. However if some damping is applied to the motion of the vibration absorber with respect to the main mass, then the two peaks with intervening notch can

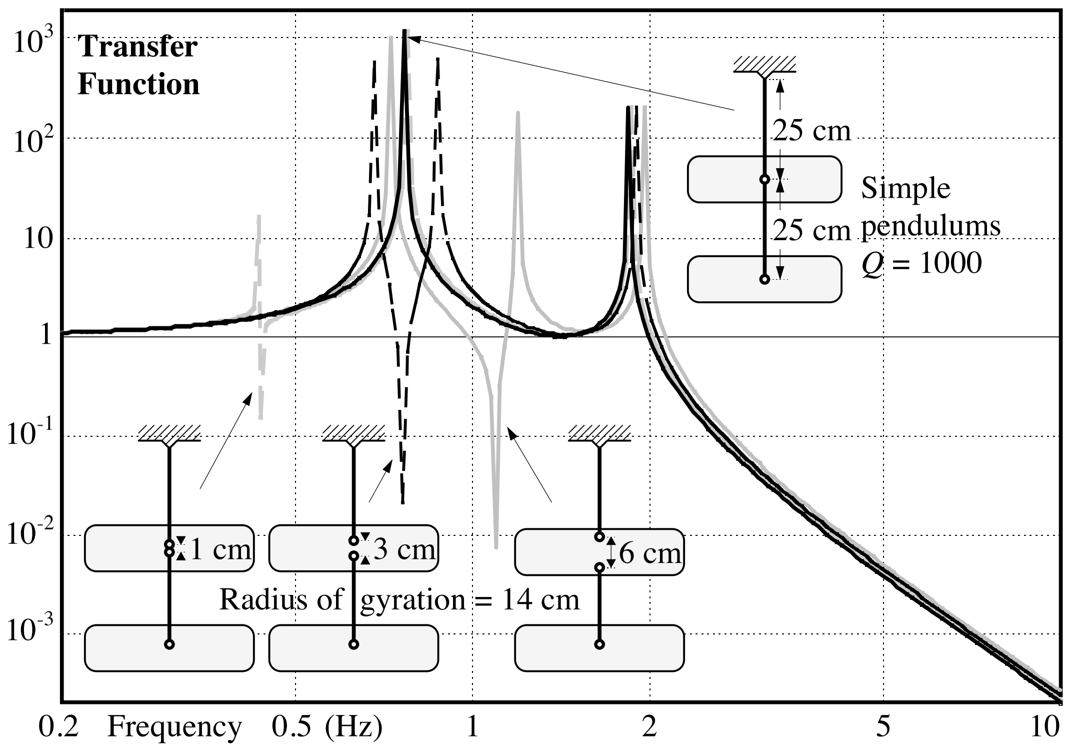

Figure 3.5 Effect of upper pendulum moment of inertia.

be smoothed into each other to produce a much flatter response. If too much damping is applied, then the vibration absorber mass will become locked to the main mass and will resonate with a single high resonant peak again, only at a slightly lower frequency (due to the extra absorber mass). So the challenge in this application is to be able to tune the available parameters in such a way as to optimally damp the overall system.

This raises the question of *what is optimal* ?. The usual criterion applied is that the peak amplitude of the transfer function across the whole frequency band should be a minimum [Den Hartog 1956]. But the transfer function usually minimised is the motion resulting from force applied to the main mass (dynamic deflection / static deflection) and so the optimisation obtained is best for a flat force spectrum. Puksand noted that this situation rarely applies to variable speed vibrating machinery in which the force increases with frequency squared and he obtained an alternate expression for optimum tuning in this case [Puksand 1975]. Our case is different again - the seismic displacement and thus force being applied through the main spring decreases with frequency squared. Indeed we may not even wish to minimise the residual motion. It only makes sense to minimise motion while the interferometer is trying to obtain lock, but once locked it makes more sense to minimise the force spectrum required to maintain lock - which may require quite different conditions as the normal modes change significantly when the end of the suspension chain becomes effectively fixed instead of free (see notes in section 1.4.1). In section 1.4.5 we provide expressions for both minimum transfer function peaking and minimum relaxation time tuning for the simplest case of a standalone self-damped stage. When such a device is cascaded into a chain of isolators it becomes almost impossible to obtain such an expression and in any case we do not see much value in a complex expression when it is so easy to use computer code to determine and simulate various solutions numerically and interactively. In this investigation we have aimed at reducing the Q-factor of the most troublesome normal mode of a simple multi-mode system as low as possible just to see what can be done and how to go about it. In engineering a real system we should rather consider it in the locked stage and reduce the Q-factor of the modes that then appear - as they are what directly determine the locking force requirement.

3.3.4 Self-Damped System Analysis

In our experience it has always been the case that the lowest horizontal swinging normal mode is the most troublesome and the hardest to damp effectively. For this reason we only consider tuning the vibration absorber to damp the 0.76 Hz mode and let the 1.84 Hz mode do what it will. In the numeric solutions which follow we have chosen to adjust the resonant absorber to the same frequency as the *resulting* filter mode as this produces the lowest Q-factors and thus the shortest relaxation time. This is the optimum choice if the criterion is that the system should settle down most rapidly after a disturbance - i.e. minimum relaxation time. The filter system ends up with two optimally low Q-factor poles sitting at the same frequency which together produce a higher peak (amplification factor) than if they were separated slightly with higher individual Q-factors (which is the standard [Den Hartog 1956] tuning). In fact applying this optimal tuning in a standard damped vibration absorber situation shows that absorber frequency tuning would be identical but our damping is $(8/3)^{½}\approx 1.6$ times as strong (see 1.4.5 for optimal tuning). The resulting difference between the two methods of tuning may be seen in figure 3.6. In practice we are sure that we will be pleased if we can just manage to adjust the absorber somewhere close!

The parameters which are intended to be used for adjustment are the viscous damping coefficient d_r (adjusted with number of magnets and their spacing from copper plates), and the effective spring-rate of the pivot k_r (by manufacture and with fine adjustment by varying the height of the rocker inverse pendulum h_i). By varying these two parameters only (k_r and d_r), it is always possible to obtain the optimum damping mentioned in the preceding paragraph. For each variation of the system, we tabulate the values of these two parameters required to obtain optimum damping, together with frequencies and Q-factors obtained for the double resonant peak near 0.763 Hz and the second filter mode near 1.842 Hz.

In practice it is quite difficult to measure the spring-rate k_r and damping coefficient d_r, whereas if the rigid section is held fixed then the resonant frequency and Q-factor of the rocker can typically be measured *in situ*. Also these latter values are mass independent, such that if the system is scaled up or down in mass then these values remain invariant. For this reason these values have been tabulated also. Values of k_r and d_r are still required in order to easily determine the practicality of manufacture. The k_r and d_r values have been based on a total system mass of 200 kg (i.e. 100 kg shared between the rocker and rigid section and 100 kg lower pendulum mass). These values scale in direct proportion to mass, such that a system with half the mass requires half the values for k_r and d_r.

Table Variables Summary		Units
k_r	angular spring-rate of rocking pivot	(N·m/rad)
d_r	dissipation coefficient of angular dashpot	(N·m/(rad/s))
f_r	resonant frequency of rocker with rigid section clamped	(Hz)
Q_r	Q-factor of rocker resonance with rigid section clamped	-
f_1	resonant frequency of low frequency filter normal mode (~0.76 Hz)	(Hz)
Q_1	Q-factor of low frequency normal mode (~0.76 Hz)	-
f_2	resonant frequency of high frequency filter normal mode (~1.84 Hz)	(Hz)
Q_2	Q-factor of high frequency normal mode (~1.84 Hz)	-
l_v	length of rigid vertical suspension section of pendulum	(m)
r_r	radius of gyration of rocker mass	(m)
μ	mass ratio (absorber/main) of vibration absorber achieving same Q-factor	-
h_i	height of inverse pendulum : rocker mass offset above pivot	(m)
r_v	radius of gyration of rigid vertical suspension section mass	(m)
h_v	height of center of mass of rigid vertical suspension section	(m)
f_v	resonant frequency of mode associated with vertical section	(Hz)
Q_v	Q-factor of mode associated with vertical section	-

3.3.5 Parameter Characterisations

3.3.5.1 Length of Rigid Section

We have already seen how the length of the rigid section in the compound double pendulums of figure 3.5 changes the resonant frequency of the rocking mode by affecting its spring-rate and allowing it to be tuned to match a filter mode. When the rotational inertia of the mass is separated from the rigid link by a spring pivot, then the two spring-rates are in series and the combination has a reduced spring-rate.

Considering first the limiting case where the pivot has zero spring-rate, but is given a significant amount of viscous damping. The spring-rate due to the rigid link length in *series* with the viscous damping of the pivot, is the equivalent of an alternate spring-rate in *parallel* with a corresponding alternate viscous damping coefficient. So even with zero spring-rate in the pivot we can expect to obtain an optimally tuned resonant system just by varying the rigid link length and damping coefficient. There are two cases - one with the rigid link above the rocker pivot (l_v negative), and one with it below (l_v positive). These solutions are shown in table 3.1. These are the minimum values of l_v for which it is still possible to obtain optimal damping. As can be seen, if the rocker mass adds to the tension in the rigid link by being at the lower end of it, then much less length is needed to achieve the required spring-rate for resonance at 0.76 Hz. The Q-factor of this resonance is very low but the amount of damping required is quite high and would need a significant mass of magnets to obtain.

l_v (cm)	k_r	d_r	f_1	Q_1	f_2	Q_2
-2.2	0	21.8	0.763	4.69	1.89	84.0
4.3	0	15.3	0.764	3.33	1.95	29.5

Table 3.1 Minimum rigid section length l_v for spring-rate k_r=0 (r_r=14 cm, m_r=100 kg, l_s=25 cm, m_v=r_v=h_v=h_i=0, d_r optimally adjusted).

If we increase the length of the rigid section and adjust both k_r and d_r for optimal damping (include the previous solutions) then we obtain the values shown in table 3.2.

l_v (cm)	k_r	d_r	f_r	Q_r	f_1	Q_1	f_2	Q_2
-20	43.6	4.34	0.751	2.13	0.746	4.64	1.88	15.1
-10	47.6	5.41	0.784	1.79	0.747	4.59	1.89	14.8
-7.0	51.1	6.62	0.813	1.51	0.747	4.53	1.89	15.2
-2.2	0	21.8	0	-	0.763	4.69	1.89	84.0
4.3	0	15.3	0	-	0.764	3.33	1.95	29.5
7.0	39.7	12.4	0.717	0.714	0.744	2.90	1.99	9.93
10	43.0	9.22	0.746	0.995	0.735	3.03	1.97	5.92
20	40.5	6.37	0.723	1.40	0.732	3.21	1.89	5.52

Table 3.2 Optimum tuning for varying rigid section length l_v (l_s=25 cm, r_r=14 cm, m_r=100 kg, m_v=r_v=h_v=h_i=0, k_r & d_r optimally adjusted).

The solutions with the rigid section below the pivot (positive l_v) obtain the lowest Q-factors for both filter normal modes, but require significantly higher damping coefficients. The Q-factor seems to reach a very shallow minimum around l_v = 7 cm and this value has been selected as the value to fix l_v at while varying other parameters. It is interesting that for some cases the clamped rocker resonant frequency f_r is set considerably higher than the filter mode frequency and other cases it is considerably lower. This is true whether one looks at the tabulated resultant filter mode or considers the simple mass mode frequency of 0.763 Hz. The reason for this is not understood at this time. We are sure that the optimum settings tabulated are unique and that there is no discontinuity between neighbouring solutions as we found far more solutions within and beyond the range than we have shown to conserve space.

3.3.5.2 Rocker Moment of Inertia

The rocker moment of inertia ($m_r\, r_r^2$) is related to the vibration absorber mass (dual is inverse of absorber spring-rate - see 1.4.5). If it is increased, we expect to have to increase rocker pivoting spring-rate somehow and also the damping coefficient in order to maintain optimum tuning. We also expect to obtain better damping with higher moments of inertia. For interest in this section we have also tabulated a value for μ which is the ratio of the "absorber mass" to "main mass" required to achieve the same Q-factor with the same type of tuning in an "equivalent" vibration absorber system. It is apparent that because the main mass in this case is a modal mass instead of the simple mass on a spring that μ normally applies to, the equivalence is not necessarily valid. The expression relating the minimum Q-factor achievable from a given mass ratio μ for this type of optimum tuning is $Q_{min}=1/\mu^{½}$ (from table 1.7)

There are a couple of interesting cases. Firstly we look at what minimum length of rigid section is required to still be able to obtain optimum tuning with various settings for moment of inertia (which is set by r_r). These values are shown in table 3.3.

r_r (cm)	l_v (cm)	d_r	f_1	Q_1	μ	f_2	Q_2
7	1.13	7.73	0.763	6.62	0.023	1.87	238
10	2.26	11.0	0.763	4.64	0.046	1.89	81.5
14	4.27	15.3	0.764	3.33	0.090	1.95	29.5
20	8.09	21.7	0.767	2.35	0.18	2.06	9.85
28	13.8	29.7	0.776	1.69	0.35	2.31	3.25

Table 3.3 Minimum rigid section length l_v for spring-rate k_r=0 and varying rocker moment of inertia $r_r^2m_r$ (l_s=25 cm, m_r=100 kg, $m_v=r_v=h_v=h_i=0$, d_r optimally adjusted).

This table shows that really very low Q-factor resonances for both filter modes can be achieved with reasonably large moments of inertia. As expected the damping and rigid section length must be increased and this length clearly sets a limit as the entire pendulum stage is only 25 cm high and we were keen to have a reasonable length of thin fibre taking up some of that height. It may also be noted that the equivalent mass ratio μ determined simply from the Q-factor Q_1 varies with moment of inertia (or r_r^2) as one might expect. In order to obtain critical damping with this optimisation, the mass ratio μ would need to reach 4 (i.e. absorber mass is 4 times main mass) which extrapolates to a value for r_r of almost 1 metre. (This is twice what one might expect from the discussion of the equations in table 1.6. This reason for this discrepancy is not known).

The second case we look at is to suppose that we leave the rigid section length at 7 cm and vary the moment of inertia from low values up to its maximum. These values are shown in table 3.4.

r_r (cm)	k_r	d_r	f_r	Q_r	f_1	Q_1	μ	f_2	Q_2
5	6.03	0.296	0.782	4.14	0.759	9.24	0.012	1.85	113
7	12.3	0.921	0.797	2.66	0.755	6.56	0.023	1.87	41.5
10	26.1	3.44	0.813	1.48	0.747	4.47	0.050	1.90	14.8
14	39.7	12.4	0.717	0.714	0.744	2.90	0.12	1.99	9.93
18.4	0	20.0	0	-	0.766	2.55	0.15	2.03	12.8

Table 3.4 Optimum tuning for varying rocker moment of inertia $r_r^2 m_r$ (l_s=25 cm, l_v=7 cm, m_r=100 kg, m_v=r_v=h_v=h_i=0, k_r & d_r optimally adjusted).

Here again it may be seen that the mass ratio μ (determined from the Q-factor Q_1) varies with moment of inertia. It is interesting that the Q-factor of the second filter mode reaches a minimum somewhere before maximum moment of inertia.

3.3.5.3 Rocker Inverse Pendulum

Since we are intending to adjust the height between the centre of mass of the rocker and its pivot in order to adjust its resonant frequency (to obtain the f_r values tabulated), it is important to check that it operates as might be expected. Table 3.5 lists optimum settings for various heights of the rocker inverse pendulum adjustment.

h_i (cm)	k_r	d_r	f_r	Q_r	f_1	Q_1	f_2	Q_2
1.7	65.4	18.5	0.794	0.529	0.773	4.85	2.06	29.5
1.5	71.6	19.7	0.857	0.536	0.756	2.70	2.06	27.3
1	59.9	17.1	0.805	0.581	0.752	2.75	2.04	19.1
0	39.7	12.4	0.717	0.714	0.744	2.90	1.99	9.93
-1	23.7	8.45	0.658	0.959	0.733	3.16	1.89	5.68
-2	11.7	5.26	0.636	1.49	0.721	3.73	1.68	4.29
-2.5	7.29	3.92	0.641	2.01	0.715	4.27	1.55	4.67
-3	3.89	2.75	0.656	2.94	0.710	5.21	1.44	6.09
-4	0.298	0.851	0.715	10.3	0.703	12.6	1.26	22.1
-4.5	0.173	0.084	0.757	111	0.702	112	1.19	284

Table 3.5 Inverse pendulum h_i effect on tuning (l_s=25 cm, l_v=7 cm, r_r=14 cm, m_r=100 kg, m_v=r_v=h_v=0, k_r & d_r optimally adjusted).

The values of h_i in the first column are positive for inverse pendulum (i.e. rocker mass above pivot) and negative for normal pendulum (rocker mass below pivot). The important thing is that this value should be able to compensate for likely values of k_r. There is a range between $-2.5<h_i<1.5$ where this is so with reasonable linearity. Manufactured values of k_r from 7 to 70 N·m/rad can be accommodated by simply adjusting h_i within this range to obtain tuning. The tuning coefficient obtained around zero is ~18 N·m/rad per cm of height change. This can be compared with the actual angular spring-rate of a standalone pendulum which is mgh or 9.8 N·m/rad per cm. The *in situ* value is almost twice this value because of the non-rigidity of the suspension chain mounting which adds an anti-restoring force.

The clamped rocker resonant frequency f_r required does not change much within this adjustment range, but perhaps surprisingly the Q-factor of the clamped rocker resonance Q_r does by almost a factor of 4, and the closely related damping, changes by a factor of 5. The reason for this is not well understood at this time, but must be related to the fact that the restoring force due to pendulum length has the gravitational vertical as its absolute reference, while the restoring force due to k_r depends on the relative angle between the rocker and the rigid section which has the resonant filter mode affecting its complex stiffness.

In general high damping values d_r such as are required for the positive inverse pendulum cases in the table are quite costly in terms of magnets. These results are an incentive to reduce manufactured k_r values as low as possible so that normal pendulum (i.e. -ve) values of h_i can be applied for tuning, greatly reducing the requirements on magnetic damping.

3.3.5.4 Rigid Section Mass and Moment of Inertia

So far we have been considering an ideal rigid section which has no mass or moment of inertia. It should be possible to make it as light as a few per cent of the total mass of the stage, leaving the great majority of the mass to the rocker. This is the condition for which we expect the most effective system damping.

Adding mass to the rigid section adds at least two more variables - mass m_v and radius of gyration r_v giving it moment of inertia (h_v considered in next section). With two extra variables, we thought it may be possible to step one (we tried stepping the rigid section mass from 4 to 50% of stage mass) and tune the other (the moment of inertia) by minimising the 2nd filter mode Q-factor (in addition to obtaining optimum damping of the 1st filter mode as before). However there was no clear preference for a particular moment of inertia to choose at each mass percentage. Failing to narrow down the variables in this manner, we chose to fix the rigid section mass at 10% and vary its moment of inertia from small to large. This data is shown in table 3.6.

r_v (cm)	k_r	d_r	f_r	Q_r	f_1	Q_1	f_2	Q_2	f_v	Q_v
5	39.2	10.3	0.750	0.806	0.741	3.09	1.98	9.37		
7	38.8	10.3	0.746	0.801	0.741	3.06	1.98	9.15		
10	37.8	10.3	0.737	0.791	0.741	3.01	1.99	8.73		
14	36.0	10.3	0.719	0.774	0.740	2.92	2.00	8.10		
20	32.3	10.1	0.681	0.745	0.738	2.74	2.03	7.56	2.12	0.514
28	26.0	9.42	0.611	0.719	0.736	2.40	2.09	10.1	1.33	0.686

Table 3.6 Rigid section moment of inertia $r_v^2 m_v$ effect on tuning (l_s=25 cm, l_v=7 cm, r_r=14 cm, m_r=90 kg, m_v=10 kg, h_v=h_i=0, k_r & d_r optimally adjusted).

Adding this rigid section mass produces another degree of freedom in the system and we can expect that with low damping another parasitic resonance will appear (similar to changing from a simple to a compound pendulum in figure 3.5). With the optimal damping aimed at the 0.74 Hz filter mode, this new resonance (labelled f_v and Q_v) is overdamped and doesn't appear as a mode until its frequency gets quite low (approaching that of the 2nd filter mode) at a radius of gyration of 20 cm. At this setting the 2nd filter mode Q-factor Q_2 reaches a minimum and as r_v is increased further the frequency Q_v gets lower and so doesn't interact so well.

It should be noted however that this variation in Q_2 is a very small effect and there is really very little to choose between these solutions. It is just satisfying to see that moving some of the stage mass from the rocker to the rigid section does not significantly affect the previously obtained results.

3.3.5.5 Height of Centre of Mass of Rigid Section

Once the rigid section is endowed with mass and moment of inertia, there remains one last parameter which is the height h_v at which the center of this mass appears along the rigid section. This height is measured from the end with the rocker pivot and is positive in the direction of the fibre joint at its other end. Negative values are possible if the majority of the mass is placed on the opposite side of the rocker pivot to its main length. Table 3.7 shows the effects of varying the height of the mass of the rigid section.

h_v (cm)	k_r	d_r	f_r	Q_r	f_1	Q_1	f_2	Q_2	f_v	Q_v
-10	26.9	10.8	0.621	0.636	0.746	2.56	2.02	4.62		
-5	33.5	10.8	0.694	0.714	0.742	2.8	1.99	6.29		
0	37.8	10.3	0.737	0.791	0.741	3.01	1.99	8.73		
5	40.2	9.79	0.760	0.861	0.740	3.18	1.99	12.2		
10	41.3	9.28	0.770	0.919	0.739	3.29	2.00	17.2		
15	41.4	8.85	0.771	0.966	0.738	3.36	2.01	24.9	3.09	0.627

Table 3.7 Rigid section center of mass height h_v effect on tuning (l_s=25 cm, l_v=7 cm, r_r=14 cm, m_r=90 kg, m_v=10 kg, r_v=10 cm, h_i=0, k_r & d_r optimally adjusted).

It is clear from the tabulated values that for reasonable positioning of the center of mass of the rigid section, the performance of the system is not adversely affected. It is interesting to note that if the mass is positioned a long way down (+ve h_v) by extending its length beyond the lower fibre attachment, that the resonance of the rigid section again appears as a parasitic mode in a similar fashion to when the rigid section was given a large moment of inertia. However it is not obvious why the Q-factor of the 2nd filter mode is not reduced as this mode approaches it (and passes it - the analysis was continued well beyond the range of the table data shown).

3.3.5.6 Sensitivity to Tuning Tolerance

The final effect to check is that reasonable damping is still obtained when the tuning parameters are significantly away from the optimum values. In this case we try altering the damping coefficient and spring-rate up and down by 10% and 25%, and make an approximately equivalent adjustment to the rocker inverse pendulum height (±4 mm and ±1 cm). In these cases the 0.74 Hz modes which were tuned to be degenerate, split and have different frequencies and Q-factors as listed in table 3.8.

	k_r	d_r	f_r	Q_r	f_{1a}	Q_{1a}	f_{1b}	Q_{1b}	f_2	Q_2
optimal	37.8	10.3	0.737	0.791	0.741	3.01	degenerate		1.99	8.73
d_r -25%	37.8	8.26	0.737	0.989	0.646	2.75	0.789	5.25	1.96	9.04
d_r -10%	37.8	9.39	0.737	0.870	0.680	2.62	0.778	4.20	1.97	8.78
d_r+10%	37.8	11.4	0.737	0.719	0.720	4.60	0.797	2.10	2.01	8.80
d_r+25%	37.8	12.9	0.737	0.633	0.715	5.88	0.861	1.90	2.01	9.08
k_r -25%	30.3	10.3	0.659	0.707	0.734	1.52	0.746	5.48	1.99	8.17
k_r -10%	34.4	10.3	0.703	0.754	0.733	1.95	0.745	4.50	1.99	8.47
k_r+10%	41.6	10.3	0.773	0.829	0.695	3.60	0.794	3.14	1.99	9.02
k_r+25%	47.3	10.3	0.824	0.884	0.679	4.46	0.823	3.46	1.99	9.47
h_i = -1 cm	37.8	10.3	0.818	0.878	0.693	9.14	0.894	2.78	1.95	5.69
h_i = -4 mm	37.8	10.3	0.771	0.827	0.698	5.27	0.829	2.58	1.97	7.31
h_i =+4 mm	37.8	10.3	0.702	0.753	0.671	1.74	0.767	5.30	2.01	10.5
h_i =+1 cm	37.8	10.3	0.645	0.692	0.591	1.22	0.773	7.56	2.02	14.2

Table 3.8 Sensitivity to tuning tolerances (l_s=25 cm, l_v=7 cm, r_r=14 cm, m_r=90 kg, m_v=10 kg, r_v=10 cm, k_r, d_r & h_i offset as detailed).

From this data it can be observed that when the tuning is away from optimal and the 0.74 Hz degenerate modes separate, that the Q-factor of one of the modes becomes a bit lower while its partner becomes significantly higher. However even for large setting errors of 25% in damping (i.e. 1 extra magnet in 4!) and spring-rate, the system remains very well damped; the Q-factor and thus relaxation time rarely exceeding twice that obtained with optimum settings. It will be demonstrated in the next section that this is no worse than using the conventional optimal tuning.

3.3.6 Damped Transfer Functions

Figure 3.6 shows the transfer function of an undamped double pendulum with Q-factor of 1000 (black curve) together with a couple of typical damping solutions (grey and dashed curves) obtained with the self-damping structure and parameters shown in the inset. The self-damping diagram in the inset is drawn approximately to scale with the areas of the lumped masses of the rocker and rigid section proportional to their mass and positioned at their radii of gyration. These are the parameters that have been central to most of the preceding tables of values.

The dark grey curve in figure 3.6 is the optimally adjusted system that has also been central to most of the preceding tables of values (it is at the top of table 3.8 for instance). The low frequency filter mode and the absorber mode are degenerate at a frequency of 0.74 Hz and a Q-factor of 3, while the second filter mode is at 2 Hz with a Q-factor of almost 9.

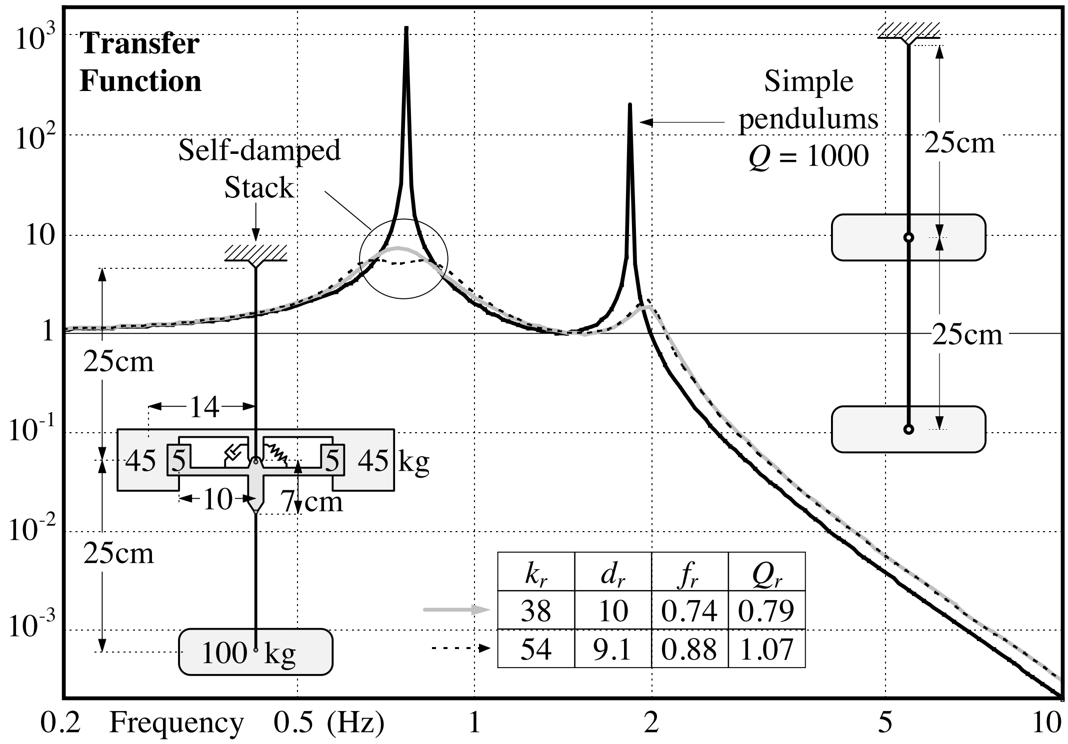

Figure 3.6 Performance of a self-damped chain compared with a simple pendulum chain.

If instead of the shortest relaxation time tuning, the minimum peak height tuning of [Den Hartog 1956] is employed, then the dashed curve is obtained. This has the resonant peaks split to 0.66 Hz and 0.84 Hz with Q-factors of 5.5 and 4.6 respectively while the Q-factor of the second mode rises to greater than 10. It is clear from these values (5.5/0.66 *vs* 3/0.74 = 2.1) that this adjustment has more than twice the relaxation time ($\propto Q/f$) of the other even though the curve looks lower and the resonances better suppressed.

3.4 Engineering Design

3.4.1 Central Tube and Suspension Wires

Due to the need to locate a fibre attachment point at the same height as the rocker pivot (see discussion in 3.3.1), it is necessary to make one of these flexible joints as a gimbal ring (or universal coupling made with flexures). Locating the second flexible joint in the space in the middle of the ring then allows the two joints to be concentric and located at the same height. (In fact this is only really the case for a single fibre suspension which was the intended arrangement at the time the gimballed rocker pivot was designed. For the multiple wire Euler suspension now used it would be possible to pivot the rocker on a single central flexure, and still allow concentric suspension with wires equally spaced around it.)

Another feature which was thought highly desirable at the time was to provide some intrinsic motion limit and safety catching arrangement in case a suspension wire should break. The method thought of for this was to provide a solid structural member which extended from the top to the bottom which could then be interlinked with the stages above and below as a very strong chain. Of course the “links” of the chain would be loose with plenty of clearance and arranged to never touch except in the case of excessive displacement such as might be caused by an earthquake or a wire breakage (or clumsy manual demonstrations of the suspension!).

For this reason it was decided to make the gimbal ring rather larger than need

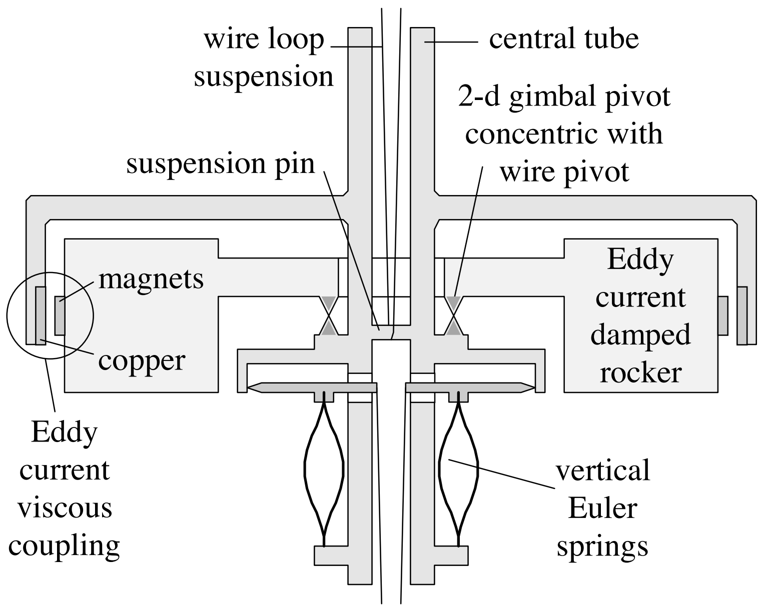

Figure 3.7 Schematic showing basic mechanical arrangement of self-damping stage.

otherwise have been, in order to fit a solid tube through it to extend from the top to the bottom and be interlocked with the central tubes of neighbouring stages. This central tube then forms a chassis on which to mount all the suspension structures and a short section of it forms the “rigid section” of the preceding discussions and analyses. This basic arrangement can be seen in .figure 3.7.

The balanced Euler vertical isolator (see section 2.2.7) naturally gives a dual wire suspension. This has some useful aspects. The angular flexing spring-rate is decreased by $1/\sqrt{2}$ of the value for a single wire, and the maximum swing angle before overstressing is increased by a factor of $\sqrt{2}$. (The diameter of 2 wires can be reduced by $\sqrt{2}$ to still support the same weight as 1 wire, and the spring-rate of an infinite length flexure varies as its diameter squared and the square root of its loading.) Also the attachment can be quite convenient at one end - a loop is present which may be simply taken around a pin and clamped in some suitable fashion.

3.4.2 Gimbal Pivot and Eddy Current Dampers

The gimbal flexure pivot was made from aluminium machining rod (which has a higher yield strength than most aluminium alloys) and the flexures were cut with electric discharge machining (EDM) wire cutting. Its basic form can be seen in figure 3.8 inset (b). The rocker mass was made from 100 mm thick aluminium plate and was intended to end up weighing approx 40 kg (with magnet assemblies etc). The gimbal flexure thickness and length to support this mass were 0.2 and 0.4 mm respectively and are loaded in compression. If additional mass is added to the rocker (see the next section) then these sizes are increased by $\sqrt{2}$ for double the mass and twice these values for 4 times the mass (160 kg total). Their short and thick (2:1) aspect ratio was designed for high frequency horizontal resonance (“S” bending of the flexure) but it also ensures that they will never buckle in compression. The thin (~0.2 mm) wire cuts at the extremity of the gimbal ring (∅100 mm) provide simple protection from over stressing and limit the angular extent of the rocking to ±4 mrad or ±1.2 mm at the damping combs (approx 300 mm lever arm).

The approximate angular spring-rates calculated for the three flexure sizes supporting the different rocker masses are 6 N·m/rad for 40 kg rocker, 12 N·m/rad for 80 kg rocker, and 24 N·m/rad for 160 kg rocker. These values are approximately one third of the typical values required for optimal damping from scaling the 90 kg rocker configuration analysed (spring-rate scales proportionally to mass). It would require the centre of mass of the rocker to be lowered almost two centimetres with respect to the pivot in order to obtain optimal tuning (see section 3.3.5.3). This in turn allows almost a 50% reduction in the damping coefficient, but optimum tuning after this lowering doesn't give such a well damped system as the higher spring-rate would have. It is a simple matter to stiffen the flexure, but usually trickier to make it much softer. In this case there is scope for reducing the pivoting spring-rate by maybe an order of magnitude by choosing a more appropriate material than aluminium, but it seems unnecessary.

The rotational viscous damping appearing mathematically in the pivot is actually

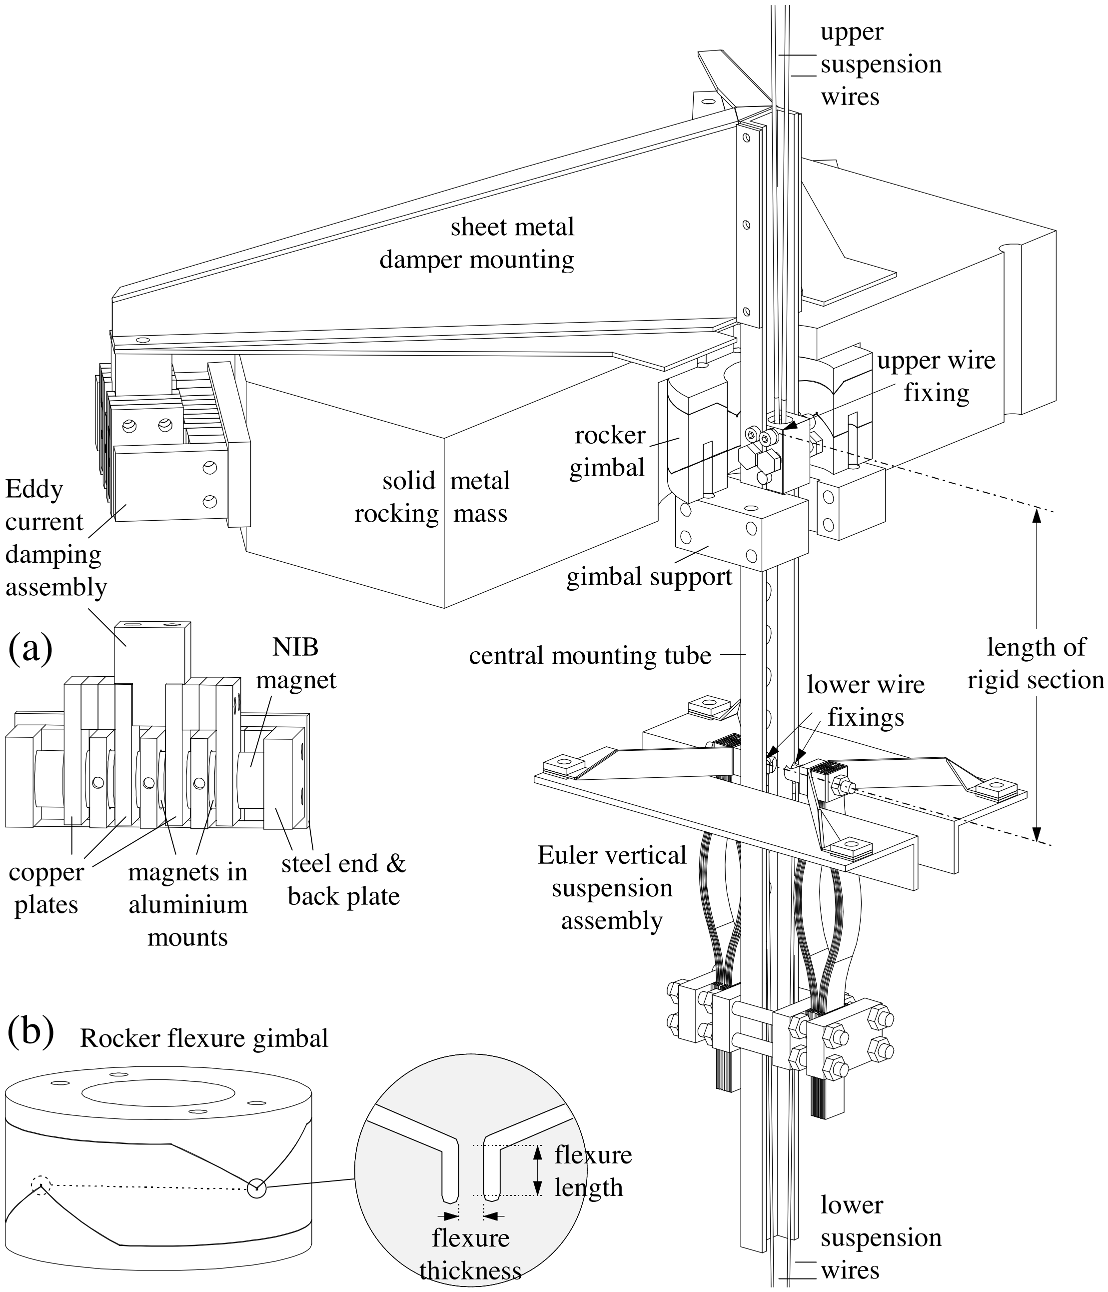

Figure 3.8 Cut away drawing showing engineering details of a self-damping isolator stage. Inset (a) shows interleaved magnets and copper combs of viscous damper. Inset (b) shows detail of 2-D rocker flexure gimbal.

applied at the extremity of the rocker mass. This is done in order to minimise the number of magnets required to produce the Eddy current damping (it reduces as the length of the lever arm squared). The NIB magnets and 6 mm copper plates are interleaved as shown in figure 3.8 inset (a) in order to make use of both poles of the magnets (halving the number which would otherwise be required). The magnets between two plates are gripped in an aluminium mounting frame while the ones on the ends hold themselves to the steel end plates. This allows the end magnets to be slid sideways to reduce the coupling to the copper plates. Further reduction is achieved by removing magnets or plates. In addition the entire assembly of steel end plates, spacers and magnets is bolted together with non-magnetic bolts and holds itself to the steel backplates attached to the corners of the aluminium rocker. This allows these assemblies to be easily slid around to adjust for equal gaps between magnets and plates and even to fine balance the rocker.

We have measured the linear damping coefficient of ∅23×15 mm (27 g) NIB magnets moving near a 6 mm copper sheet to be approx 4 N/(m/s) per magnet. Using both faces of such a magnet with a 300 mm lever arm should give an angular damping coefficient of ($2\times0.3^2\times4$) ~0.72 N·m/(rad/s). This means that the maximum angular damping coefficient we could obtain with a fully fitted set of magnets and copper combs (mounted on both opposite corners of the rocker) would be eight times this or ~5.8 N·m/(rad/s). If we take the value of 10 N·m/(rad/s) required for the optimum damping of the 90 kg rocker shown in figure 3.6 and scale it down for a rocker mass of 40 kg then we obtain 4.6 N·m/(rad/s). Hence this maximum damping of 5.8 N·m/(rad/s) is sufficient to optimally damp a 40 or 50 kg system but not much more. Analysis not presented here suggests that in a multi-stage system, the higher stages do not need nearly as high a damping coefficient as the lowest and lightest one does.

It is interesting to note that if the magnets and plates could be positioned at the extremities of the vertical z axis instead of the horizontal x and y axes, then only half the number of magnets would be required because they would all work on both rocking degrees of freedom instead of only one.

3.4.3 Cascading Self-Damped Stages and Extra Masses

For higher stages of an isolation chain slightly better performance is obtained if the masses are increased towards the top of the chain while keeping the stage heights equal (see section 8.3). For this reason the rocker was designed to have additional weights

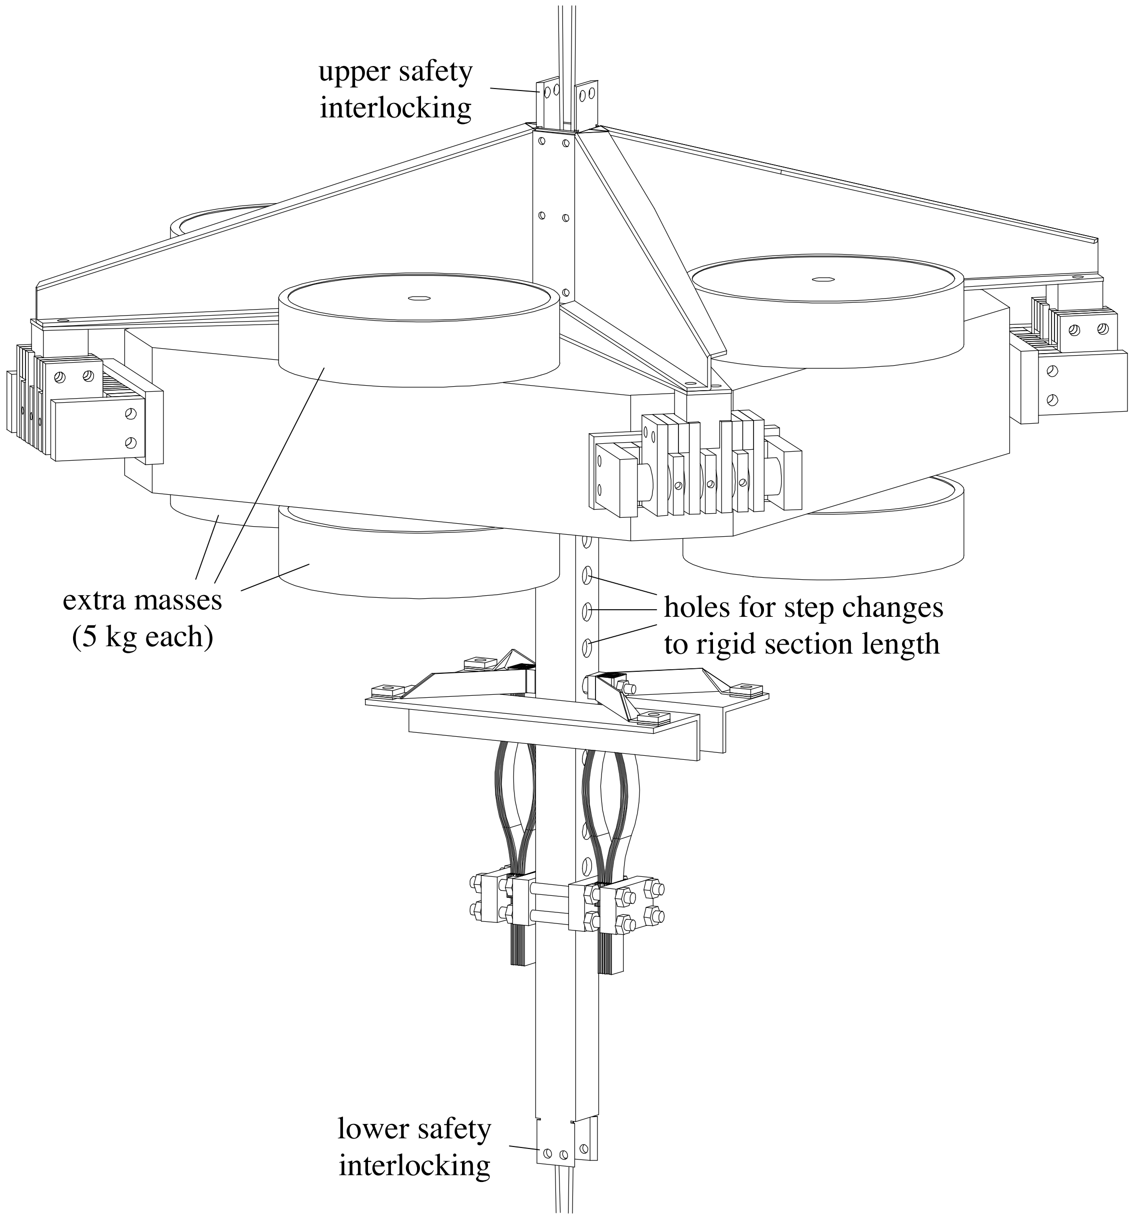

Figure 3.9 Overall view of a self-damped stage.

bolted on to the rocker plate as is shown in figure 3.9 . The weights shown are 5 kg each adding an extra 40 kg to the rocker making it 80 kg total. Alternatively triple thickness masses of 15 kg each may be bolted on to make a total of 160 kg.

Having stages which halve in mass down the chain and terminated by a mass which is equal to the last stage make a system in which each stage sees a mass hanging below it equal to its own mass. Each stage in a system such as this is working in quite similar conditions to each other and is also quite similar to the system analysed in depth in the previous section 3.3 (apart from more and differing mode frequencies). This is the case because the system as analysed was independent of actual mass and just depended on the mass ratios.

3.5 Conclusion

Unfortunately the technique of self-damping is still in quite an early stage of development. In writing this report it becomes apparent that we have only scratched the surface of investigating the possibilities that this technique has to offer in effective damping of low frequency pendulum type suspension systems. Judging by the dearth of ideas for such apparently simple low-frequency damping compared to the wealth of literature on every other conceivable form of vibration damping, one can only guess that there has been very little need for high performance suspension systems consisting of multiple low frequency pendulum stages prior to the advent of gravitational wave research.

One criticism that could be aimed at these results is that there is very little experimental evidence to indicate that they are correct. However the theory is very simple, and the author has re-written the analysis code almost from scratch approximately four times in order to improve the efficiency of machine optimisation of the damping on multiple stage systems etc. (one rewrite for instance was to specify the chain stage parameters in terms of frequencies and Q-factors rather than spring-rates and damping coefficients. This allowed the automatic optimisation progress to be tracked more intelligently by seeing what normal modes were being tuned for, and also prevented spring-rates from being adjusted by the automatic process to negative values!) At least twice there were discrepancies between the results obtained from the old and new code which required delving into the details to find and correct the error. As a result of the extensive use and experience with the routines, the author has a high degree of confidence that the results reported here are correct. Completed isolator stages of this design are currently being measured.

The author has analysed various 4 and 5 staged systems with both manual and automatic tuning (of very many parameters!), sometimes for lowest Q-factors of normal modes, and sometimes for lowest integrated residual motion when driven with a standard seismic spectrum. Apart from the (many) final results that the model indicates should be possible to obtain with a system similar to what we happen to have built, there is not a lot of general knowledge that can be gained from this rather specific modelling. For this reason the author has chosen to concentrate on this simplest possible system in this report in an attempt to gain understanding of the dynamic processes at work. As mentioned several times in the report there are unexpected effects that are not understood and need more investigation even with this simple system.

3.5.1 Future Directions

Major project scheduling dictated that the design presented in section 3.4 be built and used. However the system briefly described in section 1.4.6 - using multiple wire suspension to make all degrees of freedom have similar modal frequencies, and damping the vertical z translation against the θ_z rotation - remains a very attractive approach.

The possibility of designing a multiple staged isolation chain which achieves critical damping for all degrees of freedom (as mentioned in 1.4.7), while probably being unuseful and impractical, for some reason (maybe after the author being annoyed for so long with continually bouncing and swinging pendulums) seems a worthy challenge. Certainly a pendulum chain that could be pulled to one side and released and not overshoot on the swing-back would be impressive!

There may be a considerably simpler engineering solution to the present method of viscous Eddy current damping - in the form of fluid filled chamber of some sort. One can imagine for instance a spherical tank filled with mercury for mass and filled also with coarse steel wool or some fibre matrix to viscously resist movement of the fluid through it in all three rotational degrees of freedom. Tanks like this suspended with multiple pendulum wires terminating at Euler sprung levers (arranged to couple z to θ_z) would seem to be a paragon of simplicity. This would represent springless, untuned, viscous-only damping (i.e. a “Lanchester” damper).

Supposing instead of a spherical tank, a tank was constructed as three doughnut shaped tubes intersecting at right angles (as gyroscope gimbals). Such a design could be made from 12 standard vacuum tube 90 deg bends for instance, joined with 6 standard vacuum cross tubes. With this arrangement, fine gauze or suchlike could be clamped in some of the joints between the tubes to provide viscous damping, and thin sheets of rubber could be put across others to provide resonant tuning!

A possible problem with fluid filled tanks which could render such designs useless for gravitational wave isolation would be if temperature gradients set up flows in the fluid.

4. Scott-Russel Horizontal Pre-Isolator

4.1 Preface

The idea of the Scott-Russel isolator occurred to the author very early on in suspension work. It started out as a niggling thought that somewhere he had seen a linkage that produced almost straight line motion - which is the locus of motion needed for a center of mass to stay at a constant gravitational potential. The author found the motion in a book called "How things work" as one or two of the loci produced by a beam and crank mechanism. The particular arrangement in the diagram was not useful for an isolator but it had sowed a seed and before long the idea became more developed and a library search was done to see what other ideas were in the literature. There the author discovered many delightful mechanisms and before long had selected the very few symmetrical ones which might be applicable to suspension.

At the time the Scott-Russel linkage seemed to be the only linkage which was sufficiently symmetrical to be generalised into two dimensions and become truly rotationally symmetric about a central axis. There were a couple of other very interesting ones :- The Robert's and Chebyshev's linkages (see section 1.5.2.2) looked perfect for 1-D work and in the author's opinion should have long replaced the dreadfully asymmetric folded pendulum (Watt's linkage)[Blair 1994]. Peaucellier's cell linkage was truly a marvel in that it produced perfect straight line motion or very long radius circular motion only using pivots. However the Scott-Russel was simplest and most obviously applicable to two dimensions and so a proof of principle model of an ultra-low frequency pendulum using the Scott-Russel linkage was built (as per figure 4.2). The performance of this simple structure was carefully measured and these results formed the basis for the first paper following.

The next step was to design a full sized version suitable for incorporation into the existing stands and pendulum stacks that were being used for isolating the 8 meter interferometer at UWA. In order to measure the performance of the device, it was necessary to build a large "cradle" which could be used to suspend the pre-isolator, together with a full test load of several hundred kilograms for shake testing. The cradle was of a triangular design to be hung from three steel rods from the ceiling to allow soft translation of the pre-isolation structure in both horizontal directions without any tilting. The author had no 3-D drafting package for these designs and they were done with a combination of *Mathematica*® [Wolfram 1988] and Microsoft Word Picture. As a result of having to calculate so many double angled cuts the author swore to never design another 3-D structure with triangular symmetry (which is why a following design for the servo-frame pre-isolator was made all rectangular)!

The performance obtained from the pre-isolator was very pleasing indeed and the data trace of one particularly low resonant frequency - still with reasonable Q-factor - which was not considered important enough for the paper is included in the postscript. One can't help noting with these two papers how with so little design and construction work, the first could be written, and after so much design and construction work, only one more could be written!

Giovanni Losurdo who the author worked with in Italy visited Perth and helped with some setting up of the shaker drive of the cradle and measuring of some initial transfer functions. In fact the measurements presented in the paper required a lot more care in setting up, centre of percussion adjustment, shielding from air currents, and so forth and were taken by the author some time after Giovanni returned to Italy. But his visit and assistance is the reason that his name was included on the second paper.

4.2 Seminal Scott-Russel Pre-Isolator Paper

A long-period conical pendulum for vibration isolation.

J. Winterflood, D. G. Blair

Department of Physics, University of Western Australia, Crawley, Australia.

A practical design for a long-period conical pendulum based on the Scott-Russel linkage is presented together with proof-of-concept measurements on a simple prototype showing good stability for periods in excess of 20 seconds. The design is intended as a vibration isolation component for gravitational wave detection.

4.2.1 Introduction

The sensitivity of terrestrial gravitational wave detectors is limited at low frequencies by the degree of seismic isolation obtainable. Efforts to improve this isolation have resulted in several proposed very-low frequency passive isolators. The two most promising devices, the folded pendulum [Liu 1995] and the X-pendulum [Kanda 1994], suspend a test mass from a linkage arrangement such that its motion mimics that of a very long simple pendulum. Unfortunately both devices are most easily implemented as one dimensional isolators. Two one-dimensional isolator stages may be cascaded to obtain isolation in both horizontal directions (x and y) but there are significant constructional problems with this approach. The main problems arise from the requirement that both stages must be rigidly aligned with gravity within their operating plane and the rigidity required increases as the square of the pendulum period. Also the first cascaded stage (say x) provides no isolation in the y direction, which means that the y stage must be very well balanced and orthogonally aligned to the x stage to avoid coupling to the y motion which will bypass the first stage.

The design presented here avoids these problems by using a linkage arrangement which mimics the motion of a very long *conical* pendulum achieving x-y isolation in a single stage. It is based on the Scott-Russel linkage [Beggs 1955] which is then generalised to cylindrical symmetry to provide the required large radius spherical motion.

4.2.2 Geometry of Motion

Figure 4.1 shows a schematic of the essential linkage geometry. It consists of a normal pendulum of length r joined near the mid-point of a beam of length $a+b$. The normal pendulum is under tension and supports the entire weight of the structure and suspended mass. The top section of the beam supports the suspended mass under compression (and bending). The lower end of the beam is merely constrained to move in a vertical line directly under the main support.

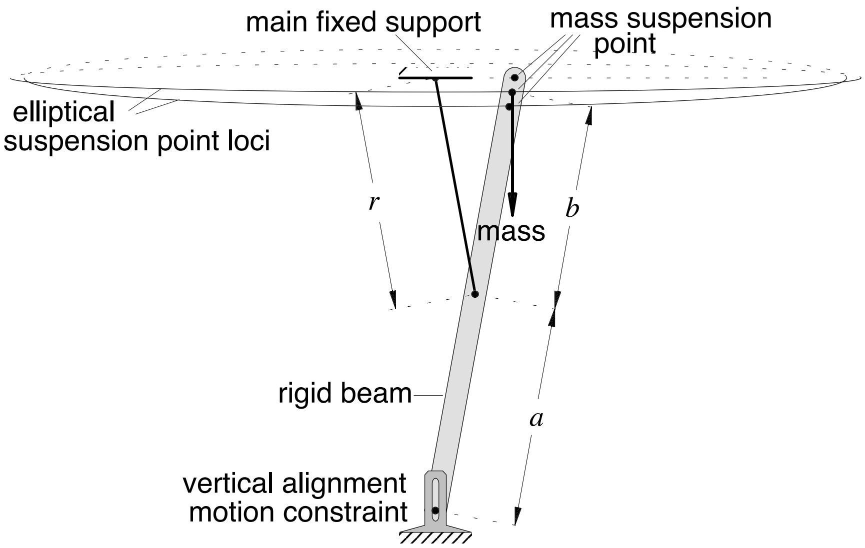

Figure 4.1 Geometrical schematic. Pendulum r is under tension and carries the weight of the structure. The lower end of the rigid beam is constrained to move in a vertical line and a is typically made equal to r. The suspension point b may be chosen to give straight-line or large-radius curved motion.

If length a is made equal to r then the upper section of the beam b may be considered as an inverted pendulum which is constrained to follow the same angle of deflection as the normal pendulum a. If b is also equal to r then the effect of the normal and inverted pendulums cancel and the suspension point follows a straight horizontal line. If a slightly lower suspension point is chosen, then the suspension point follows an exact elliptical path as shown which may in principle be set to any large radius for small displacements. The linkage may be given cylindrical symmetry to produce spherical motion. There are some spatial conflicts, but these are readily overcome with a little ingenuity.

Letting $A=a/r$ and $B=(a+b)/a$, the full equation of motion of the suspension point in a single plane $(x\text{-}z)$ with the fixed support at $(x=0,z=0)$ is

$$z = (B-1)\sqrt{A^2r^2 - x^2/B^2} - \sqrt{r^2 - x^2/B^2} \tag{4.1}$$

If $a=r$ (A=1) then this equation may be simplified to

$$\frac{x^2}{B^2r^2} + \frac{z^2}{(B-2)^2r^2} = 1$$

which is the equation of an ellipse centred about the fixed support with turning points at $\pm Br$ horizontally and $\pm(B-2)r$ vertically.

For this application $a\approx r$ and $b\approx r$, so we may replace $a \rightarrow r\text{-}a'$ and $b \rightarrow r\text{-}b'$ in equation (4.1) where a' and b' now represent the difference in length of a and b from the straight-line arrangement of $a=b=r$. The restoring force produced by gravity g acting on a mass m suspended by this geometry is then given by

$$F = -m\,g\left(\frac{b' - 2a'}{4r^2}x + \frac{b' - 4a'}{32r^4}x^3 + \cdots\right)$$

By choosing the relationship $a'=b'/4$, the x^3 term vanishes and the motion becomes closer to parabolic rather than elliptical. However this improvement is probably insignificant for the small displacements expected from seismic disturbances. Assuming the simplest arrangement $a=r$ then the resonant frequency in Hz is given by

$$f = \frac{\sqrt{g\,b'}}{4\pi\,r}$$

For a beam length of 1 metre ($r=0.5$ m) and a resonant period of 60 seconds, the suspension point is ~1 mm below the straight-line or quasi-stable point ($b'=1.1$ mm).

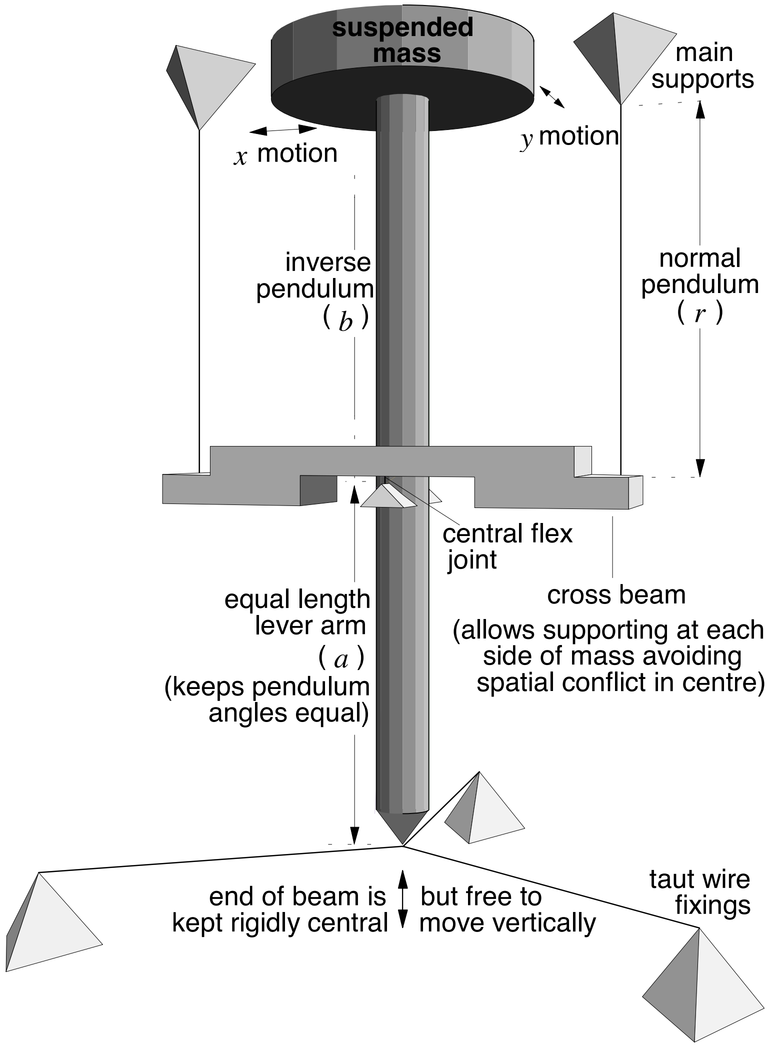

Figure 4.2 Prototype arrangement. The vertical wires carry the entire weight of the structure as a normal pendulum. The vertical beam was a threaded rod and the suspended mass a threaded disk for ease of height adjustment.

4.2.3 *Measurements*

A simple prototype as shown in figure 4.2 was constructed in order to prove the concept. The beam was actually a threaded rod approx 500 mm long and the suspended mass was a threaded disc. This allowed the "suspension point" to be adjusted up and down easily. With the cross beam arrangement the pendulum lengths (r) in the sideways direction (x) and the forward-and-back direction (y) can be independently set. Length r_y is the average length of the support wires at each side, while length r_x may be different if the central flex joint is not arranged to be at exactly the same level as the lower end of the support wires. This is the reason for the dog-leg shaped cross beam. These lengths were about 1 mm different for the measurements taken and this shows up as different resonant frequencies for the x and y directions. The beam was rather thin (6 mm threaded rod) for the mass supported but this allowed the effect of beam flexibility to be explored. The mass was initially adjusted for maximum repeatable period and then moved down in small amounts measuring the resonant frequency for each direction at each position. The process was repeated with a much heavier mass for comparison and the results are shown in figure 4.3. The amplitude of the motion used was between 1 cm and 3 cm p-p.

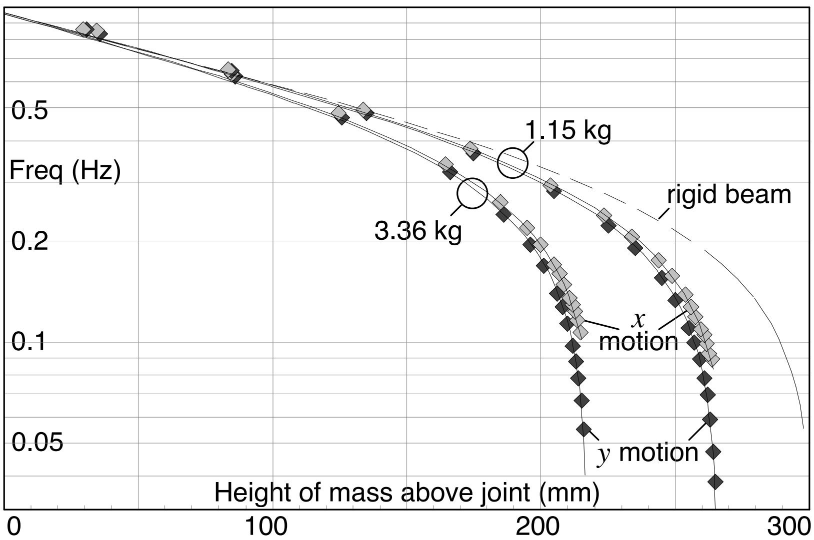

Figure 4.3 Plots of frequency vs height. Points show the frequency of two different masses in the x and y directions as a function of height. The lines are curves fitted from theory.

The large difference in the frequency vs height curves for the two different masses was due to the flexibility of the beam. The difference between the x and y directions was partly due to the 1 mm joint offset mentioned and partly due to differing beam flexibility from the flex joint construction in the centre of the beam.

The Q-factor of each resonance was measured by counting the number of cycles required for a ratio drop in amplitude. These results are shown in figure 4.4.

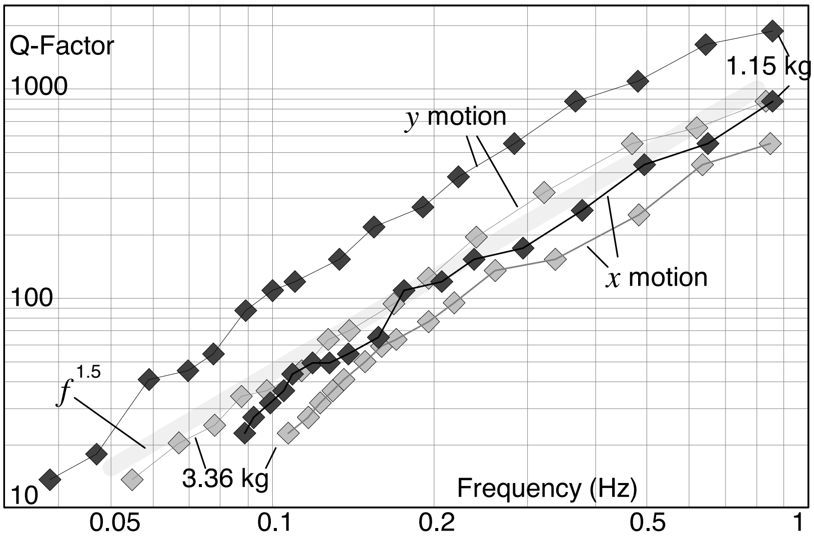

Figure 4.4 Plots of Q-factor vs frequency. Points show the Q-factor of each resonance as a function of frequency. The joining lines are only to aid the eye in grouping the points. The slopes show a close $f^{1.5}$ dependence.

The poor Q-factor in the x direction by comparison with the y direction was due to the central flex joint which is only active for x motion. This joint was rather lossy - the wire just passing around a slight V slot (not visible in drawing) whereas the main supporting wires had crimped aluminium joints. The difference in Q-factor between the heavy and light masses varies oppositely to what might be expected, but this is not surprising in the light of other experimenters experience and is discussed below.

4.2.4 Resonant Frequency

Ideally the only force acting on the mass would be the restoring force due to the geometry of the motion and its alignment with gravity. Deriving this from equation (4.1) with $a \approx r$ and to first order in x gives a spring constant *($k=-dF/dx$)* of

$$k_{geom} = mg\frac{a(ab + ar - 2br)}{(a+b)^2 r^2} \tag{4.2}$$

The effect of most of the non-ideal aspects of a physical implementation (beam, frame, and mounting flexibility, taut wire stretch, etc) is to produce an anti-restoring force proportional to the offset. It may be modelled as a spring constant with opposite polarity to that of the restoring force. The dominant component of this spring constant for this prototype was the beam flexibility resulting from its area moment of inertia I_A and Young's modulus E_Y. Its value is given by

$$k_{bend} = -(mg)^2 \frac{2b^2}{3(a+b)E_Y I_A} \tag{4.3}$$

The effect of the finite stiffness of the flexible joints may be modelled as a positive spring constant k_{flex} which simply adds to the other two. For an inertial mass of m_i the resonant frequency of this combination of spring constants is then

$$f = \frac{1}{2\pi}\sqrt{\frac{k_{geom} + k_{bend} + k_{flex}}{m_i}} \tag{4.4}$$

The extra mass of the beam m_b and the cross-member m_c need to be added to the value of the suspended mass m_s before appearing as m in equations (4.2) and (4.3). The effective height b in these equations is also less than the measured height b_m :-

$$m = m_s + m_b + m_c$$

$$b = b_m m_s / (m_s + m_b + m_c)$$

The effective inertial mass is not simply the gravitational mass m because although the centre of mass may move in an almost straight line, the beam and suspended mass rotate significantly. The effective inertial mass of an object with moment of inertia I which is constrained to rotate as it moves, about a radius of length $(a+b)$ is

$$m_{eff} = m + I/(a+b)^2$$

Applying this to both the suspended mass (a flat disc of radius r_m) and the vertical beam (a long thin rod) gives an effective inertial mass m_i for equation (4.4) of

$$m_i = m_s\left(1 + \frac{r_m}{4(a+b)^2}\right) + \frac{4m_b}{3} + m_c$$

This mathematical model is sufficient to fit the measured points quite well as shown in figure 3. Most parameters were known but some which were not easy to determine exactly were obtained by optimising the curve fit. The known parameters were :-

r_y = 258 mm, r_x = 257 mm, a_y = 260 mm, a_x = 261 mm,

m_{s1}= 1.15 kg, m_{s2}= 3.36 kg, m_b = 94 g, m_c = 34 g

r_{s1} = 9 cm, r_{s2} = 12 cm.

The fitted parameters were :-

b_m to b scale coincidence offset for m_{s1}, m_{s2}, x, y,

$E_Y I_A(x)$ = 8.67 N·m^2, (calc 6.3 for 5 mm, 13 for 6 mm),

$E_Y I_A(y)/E_Y I_A(x)$ ratio = 1.04, k_{flex} = 0.25 N/m.

4.2.5 Q-factor

The Q-factor measurements were more approximate but show some interesting details. The overall trend as indicated by the faint wide band is that the Q-factor varies with frequency to a power of approx 1.5. We shall consider some models to explain this.

If the damping force was proportional to velocity (viscous) then we would have the classical equation for damped harmonic motion $m x = -k x - d x$ for a mass m with restoring force factor k and dissipation factor d. Solving this gives a resonant frequency ω_0 of $(k/m)^{1/2}$ with a Q-factor of $(k\,m)^{1/2}/d$. Equating k in the two equations as the parameter being varied by adjusting the height of the mass gives $Q = \omega_0 m/d$ indicating that this model predicts that the Q-factor would vary proportionally to resonant frequency and mass.

It has been found that there are damping mechanisms which are frequency independent (i.e. no x term) over many decades of frequency. The effect is often termed structural damping and has been the subject of considerable investigation recently [Saulson 1994, Quinn 1995]. Two main areas have been noted [Quinn 1995] one being the intrinsic loss of the material and the other being stick and slip losses at clamped joints. The first is independent of amplitude whereas the second is strongly amplitude dependent. The damping effect is usually modelled in the frequency domain by a restoring force factor which has a small complex component. This gives an equation of motion of the form $m x = -k\left(1 + i\phi\right)x$ where ϕ is termed the loss angle and is a characteristic of the stressed material. Solving this gives a resonant frequency of $(k/m)^{1/2}$ with a Q-factor of $1/\phi$. In our case only the material spring constants (i.e. k_{bend} and k_{flex}) which we will collectively call k_m can have this loss component ϕ_m. The main adjustable geometrical spring constant k_{geom} which we will now call k_g is due to gravity which is conservative. Substituting the sum of k_m (with ϕ_m) and k_g into the equation of motion and solving gives $\omega_0^2 = \left(k_g + k_m\right)/m$ and $Q = 1/\phi_m \times \left(k_g + k_m\right)/k_m$. Equating the $\left(k_g + k_m\right)$ terms gives $Q = \omega_0^2 m/\left(k_m \phi_m\right)$. So from this model we would expect the Q-factor to vary as ω_0^2 provided k_m stays reasonably constant. The flexure component of k_m should not change with k_g adjustment and it is expected that this is the main contributor to the losses since the flexing is very localised with high stresses at joints whereas the beam bend is evenly distributed over a monolithic rod.

So it seems that there is a significant component of both viscous and structural damping present which gives the intermediate power relationship. There is also a distinct concave downward tendency and this suggests that structural damping becomes more dominant at the lower resonant frequencies and vice-versa. We expect that air friction and turbulence become significant at the higher frequencies.

With these relationships indicating that the Q-factor should vary directly with the mass for either damping mechanism, it may seem odd that an almost inverse relationship was observed. However the experimenters in [Quinn 1995] showed that stick and slip losses in clamps are strongly amplitude and stress dependent (significant amplitude dependence was observed in our case) and this dominates any intrinsic material loss. This means that the loss angle ϕ_m can increase greatly with additional loading. Also they found that even carefully designed kinematic weight mountings produced significant losses which disappeared after screwing them down tightly. In our case the heavier mass was measured first and when the lighter mass was measured it was found to inconsistently produce high losses due to movement between it and the threaded rod. So a lock nut was added to prevent this - possibly resulting in a considerable improvement over the heavier mass which was not tightly clamped. We thank a referee for drawing our attention to ref [Quinn 1995].

4.2.6 A Practical Implementation

The linkage structure shown in figure 4.2 is not suitable for suspending further structures as the rigid beam gets in the way. If it were inverted and the pendulum wires made into rigid rods with torsionally rigid flex hinges then it would become more suitable. However the entire height of this low-frequency stage would then add to the height of any structure suspended below it which could be excessive. A better approach seems to be to provide three linkages, each of which support one corner of a triangular platform which then provides a level surface as well as being horizontally isolated to low frequencies. In this way the height of the low-frequency isolating stage can be made to overlap with the height of any suspended structure ,the overall height being greatly reduced.

Beam rigidity is a prime requirement with this design. Large diameter tubular section is most suitable and rather than weakening it with holes in the middle for a cross-beam, the main support pendulum wire can be made concentric within the tube. This also ensures that the x and y pendulum lengths are identical. Figure 4.5 shows two

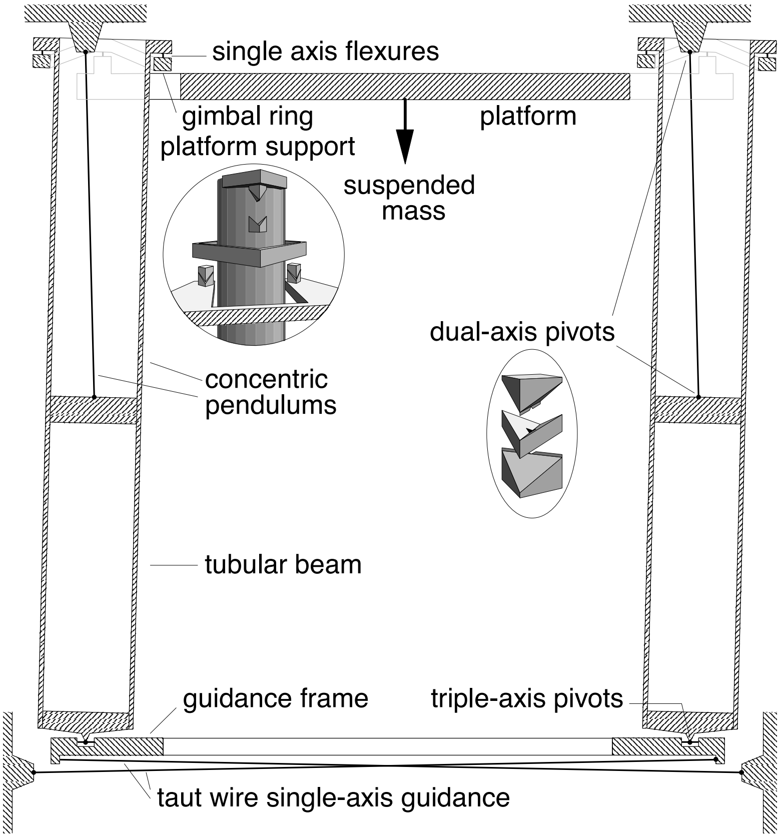

Figure 4.5 Schematic of a practical implementation. Three linkages (only two shown) can support a triangular platform providing a level, horizontally isolated surface. Concentric pendulums ensure equal *x* and *y* performance.

such tubular beams in cross-section with a platform suspended between them. The main support is at the top centre of the tube and the weight is taken by the thin central wire under tension. The lower end of the beam is connected by a flexible joint and a frame to the lower end of the other beams. This frame is constrained to only move vertically in some manner - taut wire guidance being simple and adequate. The upper end of the beam can move horizontally a few centimetres with the main support acting as a limit to this motion. The triangular platform is not easily supported centrally by the beam as this space is used by the main support. However a gimbal arrangement made with flexure pivots may be used to overcome this difficulty. An exploded diagram of a simple arrangement is shown as an inset together with a simple design for a dual axis flexible pivot. The flexures shown may be made from a monolithic block by spark erosion wire cutting. The lower flexible joints need to be able to cope with some twist about the vertical axis for the torsional motion which becomes possible with a platform supporting design. However these lower joints do not carry a significant load.

As with any vibration isolation system incorporating a rigid (i.e. massive) pendulum, care must be taken to suspend the isolated stage from the centre of percussion of the rigid pendulum. This is easily accomplished in this design by extending the beam tubes above the platform by an appropriate amount and adding some weight to their top ends. This has not been shown in the figure for the sake of simplicity.

4.2.7 Conclusion

From the measurements made it is evident that this Scott-Russel linkage performs as expected and can readily be generalised to simulate large-radius spherical motion. It is also evident that with some ingenuity it can be applied as a low-frequency horizontal isolation stage achieving *x-y* isolation in a single stage. We consider that the application of this concept to isolation for gravitational wave detection is worthy of further investigation.

4.3 Full-Size Scott-Russel Measurement Results Paper

Initial results from a long-period conical pendulum vibration isolator with application for gravitational wave detection.

J. Winterflood, G. Losurdo†, D. G. Blair

Department of Physics, University of Western Australia, Crawley, Australia.

†INFN, Sezione di Pisa and Scuala Normale Superiore, Pisa, Italy.

Initial performance measurements of a Scott-Russel pre-isolator stage with 400 Kg load capability are presented. Isolation better than 90 dB was obtained around 2 Hz. A Q-factor of ~500 was obtained with ~30 sec resonant period setting and a longest resonant period of 153 sec was obtained with Q-factor >3.

4.3.1 Introduction

Significant advantages are to be obtained by cascading an ultra-low frequency (ULF) pre-isolation stage in front of the standard multistage isolator chains used in the suspension of gravitational wave detector test masses. The internal modes of a ULF structure will typically bypass its isolation above a few tens of hertz [Blair 1993]. However the multistage isolator following the ULF stage provides such enormous isolation at these frequencies that this is no disadvantage. The advantages provided by a ULF stage include a large reduction in the seismic motion driving the multistage isolator normal modes and an extension to the low frequency end of the detection band. The seismic induced normal mode motion is so reduced that damping is no longer required and the residual low-frequency disturbing motion which needs to be counteracted by the fringe locking control loop is vastly reduced. This provides many control system advantages including easier locking with reduced loop gain and reduced force actuation with consequent reduction in noise injection. In addition, cascading a ULF pre-isolator allows the entire multistage isolator and test mass assembly to be smoothly translated useful distances with low power coil-magnet type actuators. For these and other reasons several designs for ULF horizontal pre-isolators have appeared in the literature recently [Liu 1997, Losurdo 1999, Barton 1996]. We have recently presented a pre-isolator design based on the Scott-Russel linkage [Winterflood 1996], which we believe offers the best features including 2-dimensional isolation in a single stage, and primarily *gravitational* restoring force. This allows a much higher Q-factor as well as insensitivity to temperature and mass loading in contrast with an inverse pendulum for instance. Having a high Q-factor at some already low resonant frequency, allows the resonance to be adjusted lower still (improving isolation with f^{-2}) before the Q-factor falls to a point where it ceases to provide further isolation effect. A full sized pre-isolator based on this concept has now been built and tested. Performance measurements are presented here.

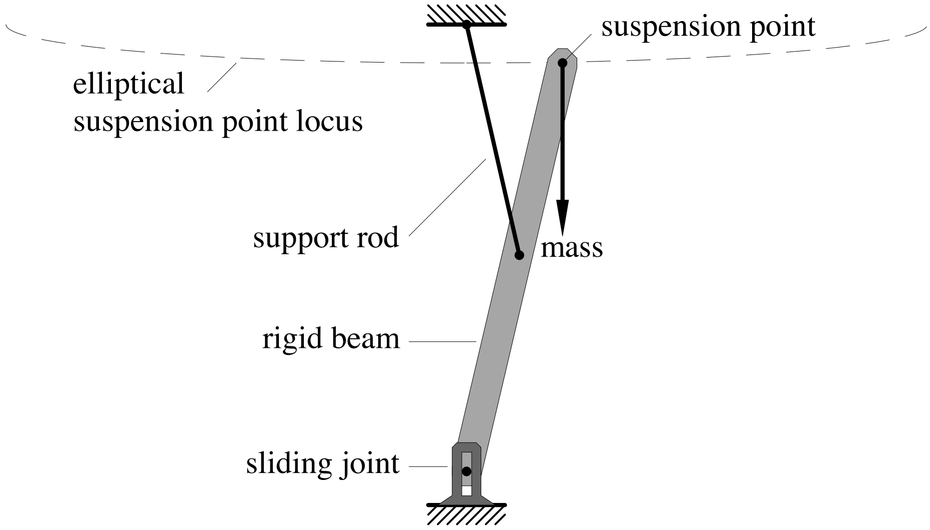

Figure 4.6 Scott-Russel linkage basic geometry.

4.3.2 Pre-isolator Realisation

A schematic of the Scott-Russel linkage and how it is applied to simulate a long radius pendulum is illustrated in figure 4.6. This linkage is given cylindrical symmetry with 2-D pivots at each end of the support rods, at the suspension point and at the sliding joint. This produces near spherical motion of the suspension point - simulating a long-radius conical pendulum with the suspended mass swinging in a very shallow potential well. The topology may then be totally rearranged while keeping the same basic geometry to give the inverted U-beam realisation shown in figure 4.7 (as well as the tubular coaxial design illustrated in figure 4.11). This inverted U-beam realisation was employed to fit over our existing isolator frames and around our existing multistage isolators. Gimbal rings form the central and lower 2-D flexible joints of the beam and the suspension point is at the centre of the top crossbar. If the maximum horizontal translation required at the suspension point is of the order of 1 cm with a rigid beam length of 1 m, the maximum vertical range required from the sliding joint is of the order of 0.2 mm. This is readily provided by the inherent flexibility in “rigid” structures such as the lower gimbal ring and a deliberate “sliding” joint is not required.

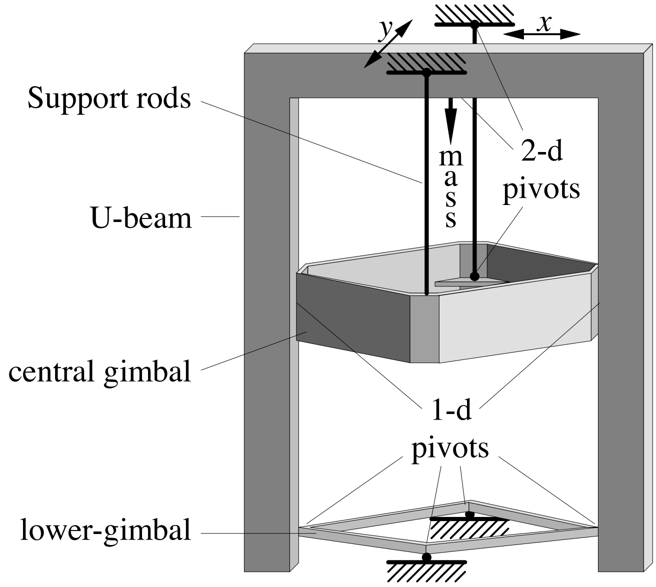

Figure 4.7 Pre-isolator realisation as upside-down U-beam.

This realisation of the Scott-Russel pre-isolator allows our unmodified multistage isolator chain and test-mass to be attached to the suspension point at the middle of the top crossbar and to hang freely down the centre. Also with this arrangement the height of the pre-isolator overlaps the height of the multistage isolator chain and so the cascaded structure is no taller than the original isolator chain without pre-isolator.

4.3.3 Design Considerations

In order to obtain very low resonant frequencies it is necessary to aim for as high a Q-factor as possible. This is necessary because the normal dissipative mechanism (structural damping) causes the Q-factor to fall off with the resonant frequency squared. This means that the system quickly becomes overdamped and ceases to behave as an isolator as the resonant frequency is lowered. Ideally the only energy stored would be due to the gravitational potential-well produced by the geometry. In reality the losses in flexures and the limited rigidity of the structure (including the mounting frame and foundations) limit the Q-factor obtainable. To assess the significance of these non-ideal effects, we will compare the unwanted spring constants with the gravitational spring constant of a fully loaded (400 Kg) 40 sec pendulum (400 m pendulum height and 10 N/m or 1 Kg restoring force per metre offset).

For the pivots in this design we employed commercial flex-pivots (Bendix Free-Flex® type from Lucas Aerospace) to give well defined load rating, centres of rotation, and friction-free low-spring-constant pivoting. One flex-pivot provides a 1-d pivot point directly but the 2-d pivot points require three flex-pivots (arranged as two lighter ones spaced along one axis and a third heavier one directly between the two and perpendicular to them). The final spring constant contributed by all the flex-pivots (23 total) to the motion is approximately +1.5 times our gravitational reference spring constant. This ends up being the dominant material spring constant in the motion.

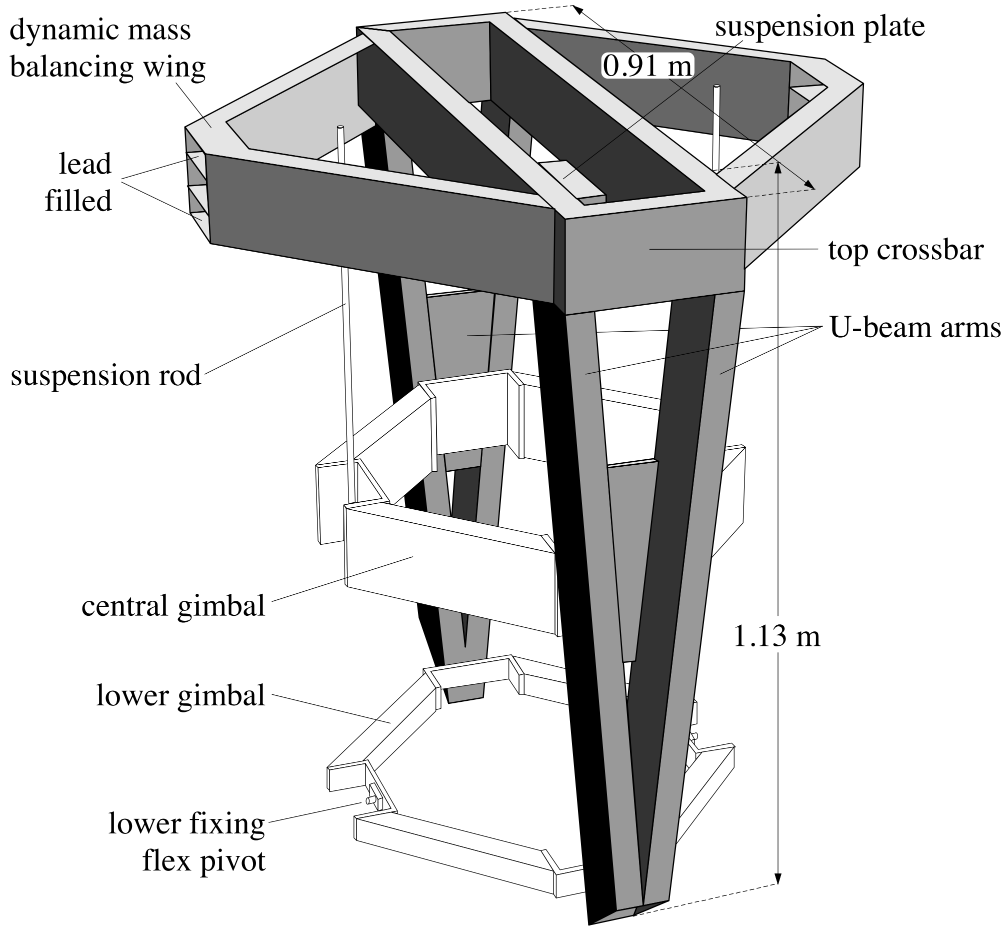

Figure 4.8 U-Beam construction showing dynamic mass balancing.

Figure 4.8 shows the final engineering design of the inverted U-beam structure. It was made primarily of hollow aluminium section (125×40×3 mm wall) giving it good rigidity. The contribution of the flexibility of the beam in the x and y directions are slightly different but the worst case (for x motion) is calculated to be approximately -0.25 of our gravitational reference (this effect produces an anti-restoring force and so subtracts from the previous spring constants but its dissipatory part is additive). Likewise the effect of the limited rigidity of the frame that it will finally be mounted on, or the cradle and the stretchiness of its suspension rods, will also subtract from the spring constant but all these have been calculated to amount to typically less than 10% of our gravitational reference.

All these material spring constant effects serve to reduce the Q-factor of the final motion, but as can be seen from these ratios, there is no reason why the Q-factor at our reference setting (400 Kg and 40 sec) should not be of the same order as the Q-factor of the flex-pivots (i.e. ~10^3).

4.3.4 Dynamic Effects

In order to obtain good isolation at higher frequencies (where the transfer function should be falling off as $1/f^2$), it is again necessary to consider the effects of non-ideal construction materials. The primary effects to be considered here are reaction forces produced by acceleration of massive elements, and normal mode resonances produced by limited rigidity in conjunction with massive elements.

When the supporting structure is accelerated suddenly in a horizontal direction (by high frequency seismic motion), then the U-beam is forced to rotate about an instantaneous rotation centre or “centre of percussion” (COP). In order to prevent this high frequency motion from affecting the suspended mass, the mass must be attached very near to the COP. The COP point depends on the distance from the centre of mass at which the accelerating force is applied and on the radius of gyration (moment of inertia / mass)$^{1/2}$ of the beam for the axis under consideration. Unfortunately the simple U-beam (figure 4.7) has a much larger moment of inertia about an axis parallel to the y direction than it does about one parallel to x, so that the COP point is quite different depending on the direction of excitation. For this reason the balancing “wings” shown in figure 4.8 were added. The hollow ends of the wings were filled with about 6 Kg of lead on each side in order to make the moment of inertia and COP identical for excitation in the x and y direction. For excitation applied at the lower gimbal, the COP for both horizontal directions can be made to coincide at the suspension point which is a few centimetres above the suspension plate shown in figure 4.8.

The normal mode resonances of the structure were measured to determine what upper frequency the structure would continue to provide good isolation at. The lowest mode affecting the structure is at 16 Hz. This is a torsional mode in which the top of the U-beam twists about a vertical axis against the lower gimbal. This mode does not bypass horizontal isolation and its effect may be readily nulled by adjusting the mass distribution on the top crossbar or by centring the suspension point. Vertical modes at 38 and 39 Hz are due to the lead-loaded wings “flapping” and the central gimbal support rods stretching. Neither of these modes couple to or bypass horizontal isolation. The first modes bypassing the horizontal isolation occur at 50 Hz (x) and 52 Hz (y). These are shear type modes in which the top of the U-beam moves horizontally in the x or y

direction. The legs act as cantilevers with a pivot constraint where they are attached to the lower gimbal.

4.3.5 Measurement Arrangement

The pre-isolator performance was tested by suspending it from a “cradle” structure to allow horizontal translation to be applied without tilting. The cradle was driven with an electromagnetic shaker (capable of 100 N force) fixed to the floor. The horizontal motion of the cradle, and of the top suspension point on the pre-isolator were monitored with shadow sensors giving a position sensitive measurement. The shadow sensors had a linear range of about ±1 mm and were mounted to the floor and ceiling respectively for minimum cross-coupling. A 1.3 metre pendulum with a dummy load of 160 Kg was hung from the suspension point. This pendulum loading (rather than directly attaching mass to the pre-isolator) closely simulates our intended application but produces what may be a confusing pole-zero pair in the transfer function (TF) where the pre-isolator mass and load pendulum resonate in contra-motion. Normally with such a cascaded system one would be interested in measuring the overall TF, in which case the pendulum load would just add a pole, but this would make the high frequency motion extremely small and difficult to measure.

4.3.6 Adjustment

Adjustment of the resonant period can be achieved by several means but the most convenient was found to be a double adjustment in which the suspension rods are say shortened by a small amount and then the position where the central gimbal pivots on the U-beam is raised by the same amount. This does not alter the suspension height or the lower gimbal height. This adjustment changes the x and y resonant periods simultaneously. If the x and y periods differ from each other, then the position where the suspension rods pivot on the central gimbal can be adjusted to bring them to match. Due to the finite weight of the U-beam, the tuning is not entirely geometric and may also be altered finely by varying the mass loading. This is convenient for making small adjustments once a suitably long period has been obtained.

The x and y centres of percussion are adjusted by adding or removing mass from the balancing wings (for y) and/or ends of the U-beam crossbar (for x). High frequency (a few Hz) excitation needs to be applied and the real part of the TF may be checked to determine how much mass to add or remove. If mass is added at the suspension point height (intended COP) then the x and y adjustments do not interact significantly and only a couple of iterations (adjust x, adjust y,..) are required to achieve a very good null.

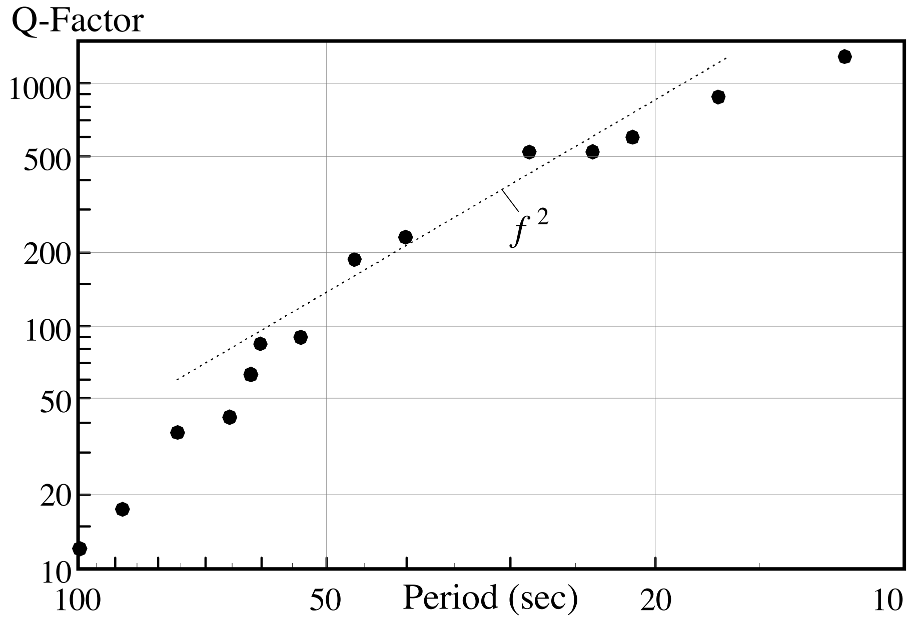

Figure 4.9 Q-factor measured at various resonant frequencies.

4.3.7 Q-Factor and Transfer Function

The Q-factor of the long period resonance in the y direction was measured from ring-down envelopes as a function of the resonant period setting. These results are shown in figure 4.9. As expected the Q-factor was several hundred for periods up to 40 sec even with only 40% of load (giving a larger material spring constant contribution than for full load). At very long periods it fell off with decreasing frequency slightly faster than the expected f^2 for structural damping but remained well over unity for periods right up to the longest obtained - a period of 153 sec (~6.5 mHz - equivalent to a 5.8 km pendulum!).

Various TFs were measured and two with a resonant period setting of approx 60 sec but slightly different COP balancing are shown in figure 4.10. The dotted curve shows the fitted theoretical expectation for perfect COP balancing. The notch at 0.8 Hz and peak at 1 Hz are the pole-zero pair previously mentioned that are due to the pendulum loading while measuring only the pre-isolator TF. The response falls off as $1/f^2$ as expected for structural damping until about 1 or 2 Hz. (With viscous damping we would expect a change in slope to $1/f$ at a frequency of the Q-factor times the resonant frequency, and no deep notches would be measurable since the 90 deg phase shifted motion from viscous coupling cannot be nulled by COP balancing.)

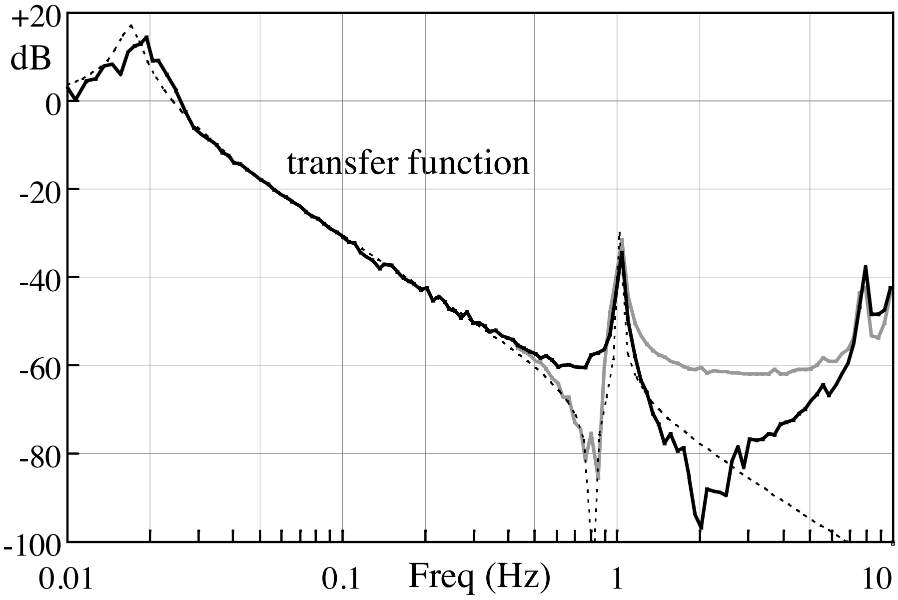

Figure 4.10 Transfer functions: black - COP nulled at ~2 Hz, grey - nulled at ~0.9 Hz by adding 40 gm to each wing, dotted - fitted theoretical.

Above 1-2 Hz the isolation is limited due to a combination of the COP effect and acoustic coupling. With acoustic coupling present, it is only possible to null the residual motion at one particular frequency and at nearby frequencies the isolation is limited to about 60 dB. First a null was obtained for excitation at 2 Hz and the TF obtained with this balance is shown in black. In this case the acoustic coupling hides the notch at 0.8 Hz and also causes the TF to drift up to about 50 dB at 10 Hz instead of an expected 80 dB. When 40 gm mass was added to each wing in order to move the null to about 0.9 Hz, the notch at 0.8 Hz was revealed but the isolation above 1 Hz was limited to 60 dB. A careful and complex arrangement of rigid baffles would help prevent acoustic coupling during testing but operation as intended in a vacuum is the best solution. The small peak at about 7.5 Hz is a torsional mode of the cradle assembly and shaker drive.

4.3.8 Conclusion

We have constructed a full-sized horizontal pre-isolator based on the Scott-Russel linkage and obtained excellent isolation performance and extremely low resonant frequencies with reasonable Q-factors. These measurements suggest that with drift cancelling feedback added, it should be possible to operate this system at resonant periods of up to 150 sec.

However at this level of isolation, cross-coupling from other modes of motion (vertical and tilt) into the horizontal will easily exceed that from direct horizontal feedthrough unless careful steps are taken to prevent it. We have reported [Winterflood

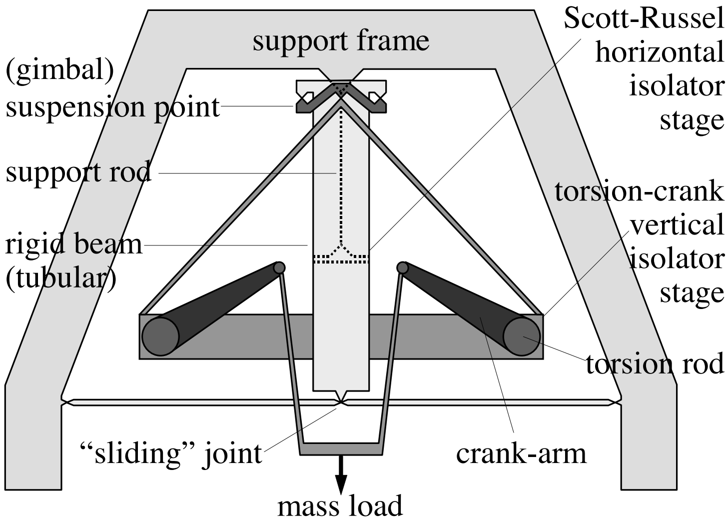

Figure 4.11 Coaxial Scott-Russel pre-isolator cascaded with a ULF vertical pre-isolation stage.

1998] on a suitable method of achieving very low frequency vertical pre-isolation and a conceptual diagram of the scheme is shown in figure 4.11 together with an alternative Scott-Russel pre-isolator realisation. Seismic tilt feedthrough can only be prevented by an active system which senses tilt and counteracts it - and is a necessity if any usefulness is to be obtained from horizontal pre-isolation with resonant periods longer than 30 or 40 seconds.

If combined with very low frequency vertical pre-isolation and an active seismic tilt cancelling system, this Scott-Russel isolator would give a very desirable horizontal isolation of 75 dB at our 0.5 Hz worst normal mode stack resonance and considerably better isolation above that. Alternatively its centers of percussion could be adjusted to give a deep null in the TF at 0.5 Hz while still providing 75 dB isolation for frequencies and stack resonances above 0.5 Hz.

With the centres of percussion normally adjusted, the suspension point well centred to null torsional coupling, and operated in a vacuum, we expect the high degree (~90 dB) of isolation obtained to continue up through 10 Hz or so as the first troublesome internal modes start at 50 Hz.

4.3.9 Acknowledgments

We thank Graeme Warburton of the Physics Workshop for construction of the pre-isolator and test cradle and Miss Seval Hackett for assistance with the Q-factor measurements.. This work was supported by the Australian Research Council.

4.4 Postscript

Much, much later, after all was built and tested, in following up some other references, the author discovered a 2-D Scott-Russel pendulum (as well as the 1-D Watt's linkage) that had been built by a retiree Nils Lindenblad and described in an Amateur Scientist article in Scientific American [Lindenblad 1967]. The most remarkable thing about the Scott-Russel design presented there was that it was inverted to the design presented here - which really makes it very hard to engineer into a 2-D structure. One can't help being surprised at the insight that was sufficient to find and apply the idea, but that stopped short of applying it in the simplest and apparently most obvious manner. He was aware that it was the straight-line aspect of the linkages that made them useful for this work and knew them by their correct name - unlike many who have written about the "folded pendulum"!

The >150 second period (and its ~6 km equivalent height!) obtained using the large Scott-Russel pre-isolator is often mentioned but the evidence for it was not put into the Physics Letters A paper (although it did appear in one or two conference proceedings). For the sake of completeness the trace is reproduced here as figure 4.12.

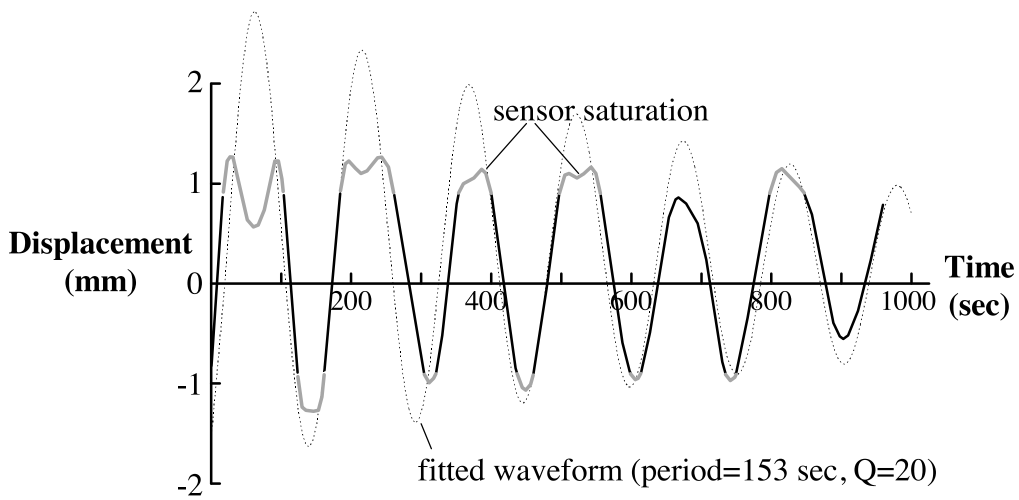

Figure 4.12 Trace of a very low resonant frequency setting.

Some time after completing this work, the author had more opportunities to design and construct ultra-low frequency conical pendulums for a seismometer project. Wanting to try something different, he used the Robert's linkage (see section 1.5.2.2) to make a dual degree of freedom pendulum instead of the Scott-Russel and was quite pleased with its simplicity and how well it performed. For the particular space that was available around the existing suspension chains and stands in UWA's 8 meter interferometer, it still seems that this Scott-Russel design is the most ideally suited. But in a different situation, such as is presently the case of needing a pre-isolator to cascade within the servo frame pre-isolator, the author would likely prefer the Robert's linkage for its slightly improved mounting simplicity.

5. Torsion-Crank Vertical Pre-Isolator

5.1 Preface

The idea for this device sprang from the growing need for an ultra-low frequency vertical isolation stage to complement the horizontal pre-isolators which by that time had produced excellent performance. The magnetic anti-spring technique of the VIRGO group had produced reasonable performance at the time, but had significant engineering and other difficulties associated with it. The single spring LaCoste pendulum was also well known but unfavourable due to its coil springs and asymmetry. So the author again searched the literature for alternative vertical techniques but really found very little indeed. It was clear that some nonlinear geometry was required, and a torsion rod spring element seemed highly favoured for various reasons. As mentioned in the paper, the torsion-crank implementation seems to be the simplest possible symmetrical arrangement of this type, which can provide a point of inflection in the force-displacement relationship.

5.2 Seminal Torsion-Crank Pre-Isolator Paper

A long-period vertical vibration isolator for gravitational wave detection.

J. Winterflood, D. G. Blair

Department of Physics, University of Western Australia, Crawley, Australia.

A novel design for a long-period passive vertical suspension system based on a torsion crank linkage is presented together with proof-of-concept measurements. A working range of centimetres and periods in excess of 20 sec were obtained in agreement with theory. The design is cascadable with a horizontal isolator for 3-D isolation.

5.2.1 Introduction

The sensitivity of terrestrial gravitational wave detectors is limited at low frequencies by the degree of seismic isolation obtainable and by gravity gradient noise. Efforts to improve low frequency seismic isolation of test masses have taken two main approaches. One is to use sensitive low-frequency accelerometer sensing and feedback to actively suppress seismic noise [Flaminio 1994,Stebbins 1994]. The other approach is passive isolation. Both methods require a test mass to be suspended with as low a resonant frequency as possible (either the small test mass within the accelerometer or the entire large test mass assembly). Suspending the entire assembly has some advantages in that low resonant frequencies are easiest to obtain using large dimensions and heavy masses. Also the entire structure may then be readily translated useful distances with low power actuation. For these reasons several designs for ultra-low frequency (ULF) horizontal pre-isolators have appeared in the literature recently [Winterflood 1996,DeSalvo 1996,Barton 1996].

While interferometer based gravitational wave detectors are primarily affected by horizontal motion, they are also affected by vertical seismic noise - through mechanical cross-coupling and also due to the local vertical not being exactly perpendicular to the long baseline laser beam. For this reason vertical isolation is also essential – and presents quite different problems from the horizontal. ULF horizontal isolators must be carefully aligned to cancel the effect of gravity and may be designed so that minimal energy is stored in elastic elements. Vertical isolators however have to support the entire weight against the force of gravity, and large energy storage in elastic elements is almost unavoidable.

The resonant frequency of a mass suspended by a linear elastic element can be given as $(g/l)^{1/2}$ where g is the acceleration due to gravity and l is the length extension under load. In order to obtain a low resonant frequency from a passive system without requiring extravagant vertical extension or pre-stress, it is possible to arrange a nonlinear force-displacement characteristic with a flat region near the operating point. If this is done then the resonant frequency is better expressed as $(k/m)^{1/2}$ where $k = \partial F/\partial y$ (F = force, y = vertical displacement) is the effective spring constant at the operating point and m is the suspended mass. Significant effort has been applied to obtaining low vertical resonant frequencies in seismometer and gravimeter designs. One well known approach is the LaCoste pendulum which makes use of a helical tension spring acting at 45 deg on a horizontal pendulum [LaCoste 1934]. In order to obtain good cancellation of the nonlinear effects, LaCoste innovated the use of a "zero-length" spring and obtained very constant period with significant displacement (i.e. 37±1 sec for ±9 deg pendulum deflection). Another technique which does not provide nearly such good cancellation is to sum a magnetic anti-spring with a mechanical spring. The Benioff variable-reluctance seismometer utilises this principle and while not aimed at long periods, obtains a spring constant cancellation of ~90% with an operating range of a few millimetres [Benioff 1955]. The VIRGO group have investigated this technique thoroughly on their super-attenuators (SA) using blade springs and permanent magnets [Braccini 1993 Beccaria 1997]. A standard SA stage uses blade springs 354 mm long and the resonant period remains between 5 and 10 seconds for an operating range of ±1.3 mm. In a special prototype SA with slightly more than double length blade springs and magnet array height they obtained an operating range of ±5.2 mm for the same resonant period limits [Cella 1997].

These results show that it is difficult to obtain good spring constant neutralisation over a large operating range. Another disadvantage of the magnetic anti-springs is the large temperature coefficient of the ferrite magnets (ferrite was chosen for minimum eddy current losses). In this letter we present a theoretical analysis and experimental confirmation of a new ULF vertical stage - the torsion crank linkage (TCL). It uses only elastic material to produce the spring constant and modifies it with geometry to obtain the required long periods. We avoid the use of helical tension springs - in particular the problem of manufacturing powerful zero-length springs. Our geometry does not produce as unvarying a spring constant as the more ideal LaCoste geometry, but its operating range appears to be almost fifty times larger than a standard VIRGO SA stage for comparable parameters.

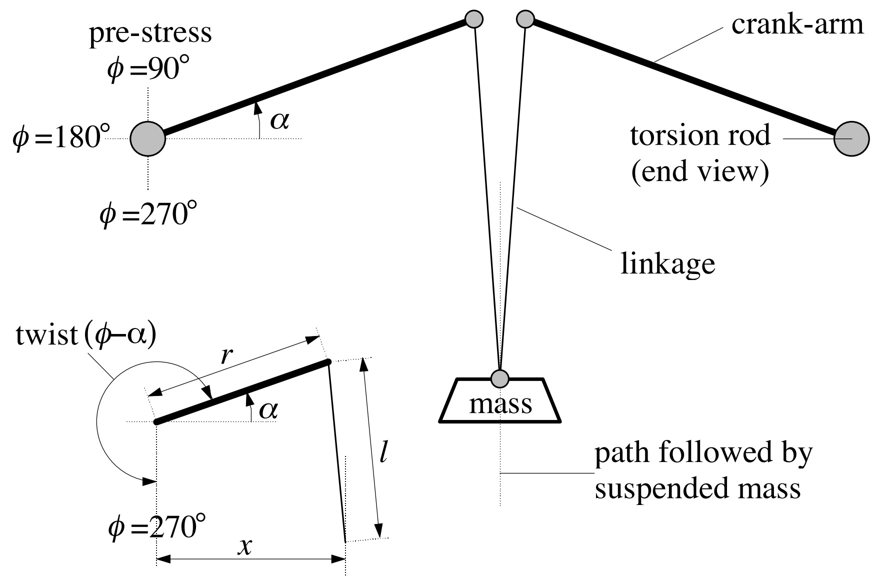

Figure 5.1 Geometry schematic. A balanced pair of torsion-sprung crank-arms connected to a single mass provides a nonlinear single axis suspension.

5.2.2 Design of Torsion Crank Linkage

The TCL makes use of the nonlinearity produced by a torsion sprung crank-arm connected to a suspension link, the loaded end of which is constrained to move in a vertical line. Figure 5.1 shows the concept of the isolator. The simplest way of achieving the vertical motion constraint is to use a pair of cranks in a symmetrical

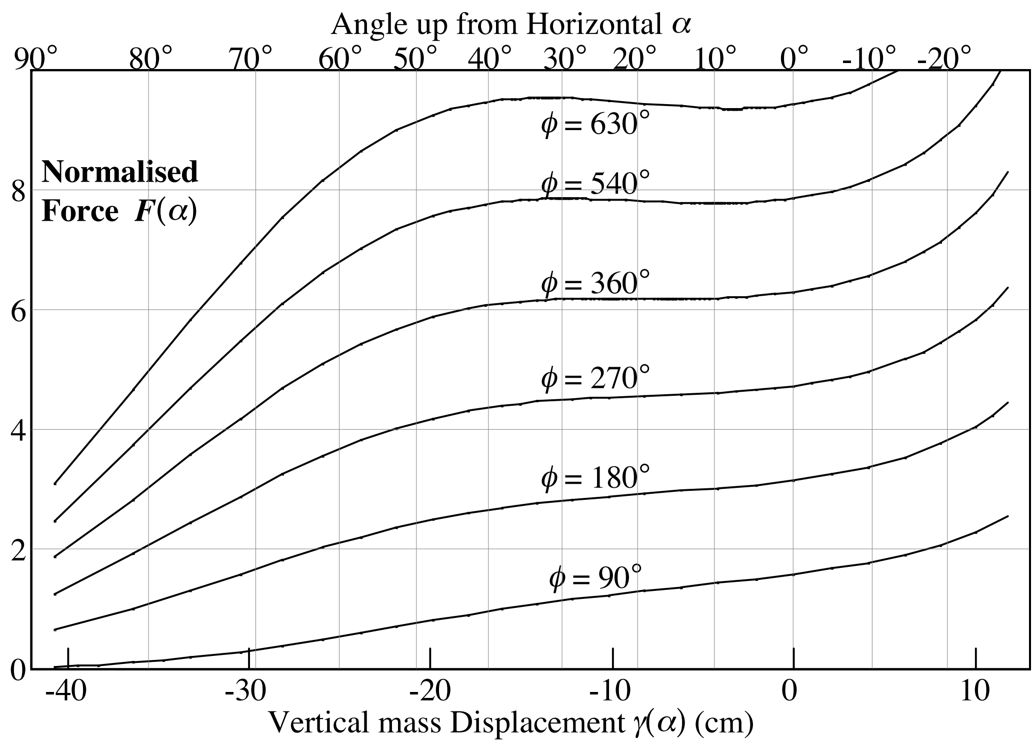

Figure 5.13 Plot of force vs height. These curves are for equal length crank, link, and offset. The nonlinearity in the characteristic increases with pre-stress and the slope changes from positive to negative for angles near 20 deg.

arrangement as shown. The torsion-rods are pre-stressed to provide an upward acting torque on the crank-arms and the mass is supported by links with flexible joints at each end.

If we normalise the geometry variations by fixing the crank-arm length to a constant r, then there are three other parameters which may be changed to adjust the force-displacement characteristic :- (1) the length l of the supporting link, (2) the offset distance x from crank-arm centre to the vertical path followed by the end of the link, and (3) the amount of pre-stress ϕ in the torsion spring. We define this pre-stress angle ϕ as the angle through which the crank-arm must be turned (in the loading direction) from being unstressed to being horizontal.

In figure 5.13 the force-displacement characteristic is shown for the case where the suspension link and offset are the same length as the crank-arm ($l = x = r$) and the pre-stress is varied in steps of 90 deg. It is apparent that as the pre-stress is increased, so the nonlinearity in $F(y)$ becomes more pronounced until $\partial F/\partial y$ becomes zero and then negative (in the region where the crank-arm angle α is about 20 deg). This is the desired effect. We wish to operate with the spring constant $\partial F/\partial y$ almost zero but remaining slightly positive, thus giving a very long vertical resonant period for a suspended mass. Typical resonant period vs vertical position curves for the same

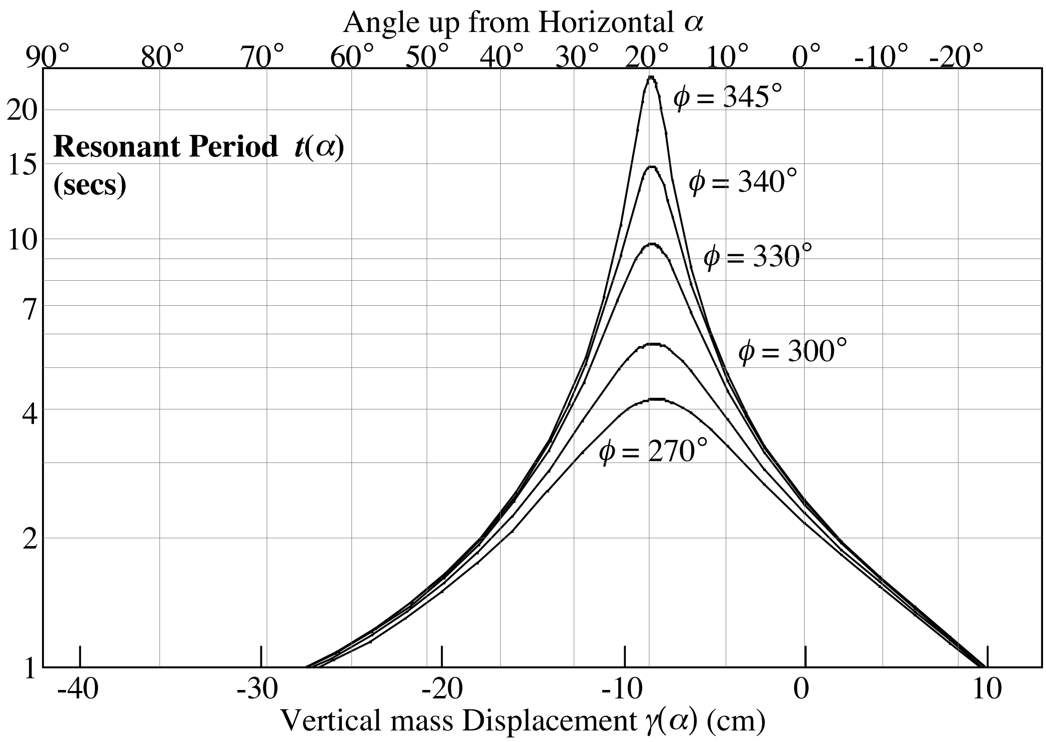

Figure 5.2 Resonant period vs height. As the pre-stress is increased towards 345 deg, the force-displacement slope approaches zero and the resonant period increases rapidly.

geometry ($l = x = r = 25$ cm) are shown in figure 5.2. This figure shows that as the pre-stress is increased towards 345 deg, the resonant period increases rapidly - as the force-displacement characteristic approaches zero slope.

We believe this is the simplest torsion based geometry capable of providing an almost flat $\partial F/\partial y$ together with inflection at the same point ($\partial^2 F/\partial y^2 = 0$). (The geometry can be arranged to also give $\partial^3 F/\partial y^3 = 0$ at the same point, but we show that this is not particularly advantageous for our application.) This design has advantages over other schemes in compactness, ease of manufacture and adjustment, and application to arbitrarily large masses.

5.2.3 Geometry Analysis

Consider a torsion lever arm (length r) at an angle α (anticlockwise and unwinding the torsion-rod) from horizontal, and with pre-stress to horizontal ϕ as shown in figure 5.1. A link of length l is joined to it and the lower end constrained to move in a vertical line at a distance x from the torsion-rod centre. The energy E stored in the torsion-rod as a function of angle will be

$$E(\alpha) = \frac{1}{2} k_a (\phi - \alpha)^2$$

where k_a is the angular spring constant of the torsion-rod. If we neglect other energy storage mechanisms such as the finite flexibility of the joints etc, then the force acting in the vertical direction as a function of vertical position is given by the rate of change of this energy with position $\partial E/\partial y = \partial E/\partial \alpha \times \partial \alpha/\partial y$.

$$F(\alpha) = \frac{k_a(\phi - \alpha)}{y'(\alpha)} \tag{5.1}$$

It is the nonlinearity of the derivative $y'(\alpha)$ that allows the force $F(\alpha)$ to remain almost constant with displacement in the region of a particular operating angle. The relationship between the vertical position y of the mass and the angle α is given by the geometry :-

$$y(\alpha) = \sqrt{l^2 - (x - r\cos\alpha)^2} - r\sin\alpha + const \tag{5.2}$$

Using $t = (m/k)^{1/2}$ where k is the effective vertical spring constant $\partial F/\partial y = \partial F/\partial \alpha \times \partial \alpha/\partial y$ and $m = F(\alpha)/g$, if follows that

$$t(\alpha) = 2\pi \sqrt{\frac{F(\alpha)\, y'(\alpha)}{g\, F'(\alpha)}} \tag{5.3}$$

Since both $F(\alpha)$ and $F'(\alpha)$ are proportional to the torsional spring constant k_a this expression for the period $t(\alpha)$ is independent of this spring constant and only depends on the geometry function $y(\alpha)$ and the gravitational acceleration g.

In order to obtain a constant resonant period for small vertical motion, it is necessary to choose an operating point where the resonant period is at a turning point – i.e. $\partial t/\partial y = \partial t/\partial \alpha \times \partial \alpha/\partial y = 0$. Since $\partial \alpha/\partial y$ is slowly varying and non-zero near the operating point, we may simply solve for $\partial t/\partial \alpha = 0$ deriving it from equation (5.3). Keeping only the variable terms from the numerator gives the turning point condition

$$F'(\alpha)\,y'(\alpha) - F(\alpha)\,F''(\alpha)\,y'(\alpha) + F(\alpha)\,F'(\alpha)\,y''(\alpha) = 0 \tag{5.4}$$

Equations (5.3) and (5.4) become rather large once the substitutions from (5.1) and (5.2) are made due to the multiple derivatives present. However they are readily solved numerically.

The solution of equation (5.3) to obtain a particular period (we chose 20 sec) together with equation (5.4) with fixed r (we chose 25 cm), is sufficient to fix only one of the three parameters l, x and ϕ. Thus there is a two-dimensional space of solutions which needs to be searched for an optimum.

5.2.4 Optimising the Geometry

We have searched the space of solutions of equations (5.3) and (5.4) in an attempt to find an optimum geometry for which the resonant frequency is minimised over the largest possible vertical range. We determined that the special geometry giving a line of solutions with $\partial^3 F/\partial y^3 = 0$ gives rise to a false optimum because it is associated with a horizontal instability over part of the operating region. (The symmetric opposing cranks become horizontally unstable over approximately the upper half of the vertical range, reducing the useful range to about the same as that of a simple maximum. Horizontal motion occurs when both crank-arms turn the same way allowing one arm to rise and the other to fall). This horizontal instability can be overcome by adding constraints to limit motion to vertical translation only. The VIRGO magnetic anti-spring solves a similar problem by the use of horizontal centring wires, at the expense of significant complexity.

Other than this special line of solutions requiring horizontal stabilisation we find that there is no strongly optimum geometry. Over an area spanned by x (horizontal offset) varying from 10 cm to 40 cm and l (link length) varying from 10 cm to 50 cm using a crank-arm radius of 25 cm, approximately one third of the area gives roughly equally acceptable solutions (another third of the area is not so useful, and the remaining area has no realisable solutions). The tuning curves shown in figure 5.2 are typical of the useful solutions. The geometry is so flexible that it can be chosen depending on the available space for mounting a practical structure and the amount of prestress that can be applied to a particular torsion element. However for dynamic balancing reasons (see sect 5.2.6) there is some advantage in choosing a geometry which has an operating angle close to 90 deg between the crank-arm and link. Detailed graphical results of the parameter space analysis will be published elsewhere.

5.2.5 Functional Test

We used a pair of high-tensile steel torsion-rods 2.3 mm in diameter and 600 mm long, operated at a wind-up angle of just over 3/4 of a revolution. The crank-arm length r was 25 cm and the pivots were knife edges. The link length l was 29 cm and we suspended an adjustable mass of the order of a kilogram from the system. Adjustment is simple. If the resonant period is too short, then more mass is added and compensated by increased pre-stress to maintain a constant operating point (increasing ϕ in figure 5.13 and figure 5.2). If the motion becomes bistable with vertical instability in between two stable points, then both pre-stress and mass are reduced. In this way the system is readily tuned for long periods.

The fine wires forming the links were terminated by winding them around a 4 mm diameter rod at the ends of the pivot arms and at the mass. This avoided the significant energy storage of strong flexures and the anelastic bending problems of thin flexures for the large flex angles under test. However this termination geometry is slightly inconsistent with the mathematical analysis (e.g. l is slightly dependent on α due to the 2 mm winding radius). Regardless of this, the general agreement between theory and experiment was remarkable, as shown in figure 5.3. As predicted, the resonant period remained greater than 10 sec over a vertical range of almost ±2 cm. Scaling the crank-arm up to match the 354 mm length of the VIRGO SA blade springs and choosing the same period limits (5 to 10 sec) would give a vertical range of almost ±6 cm. Periods greater than 20 sec were readily and repeatably achievable with a Q-factor of about 2.

5.2.6 A Practical Implementation

Figure 5.4 shows a full scale conceptual implementation of a TCL vertical isolation stage, cascaded after a Scott-Russel horizontal isolator stage [Winterflood 1996]. A passive ULF horizontal stage always requires rigid alignment with gravity and so needs to precede the vertical stage. The diagram shows a single beam Scott-Russel stage in the centre with torsion cylinders placed on either side of it to make good use of the space.

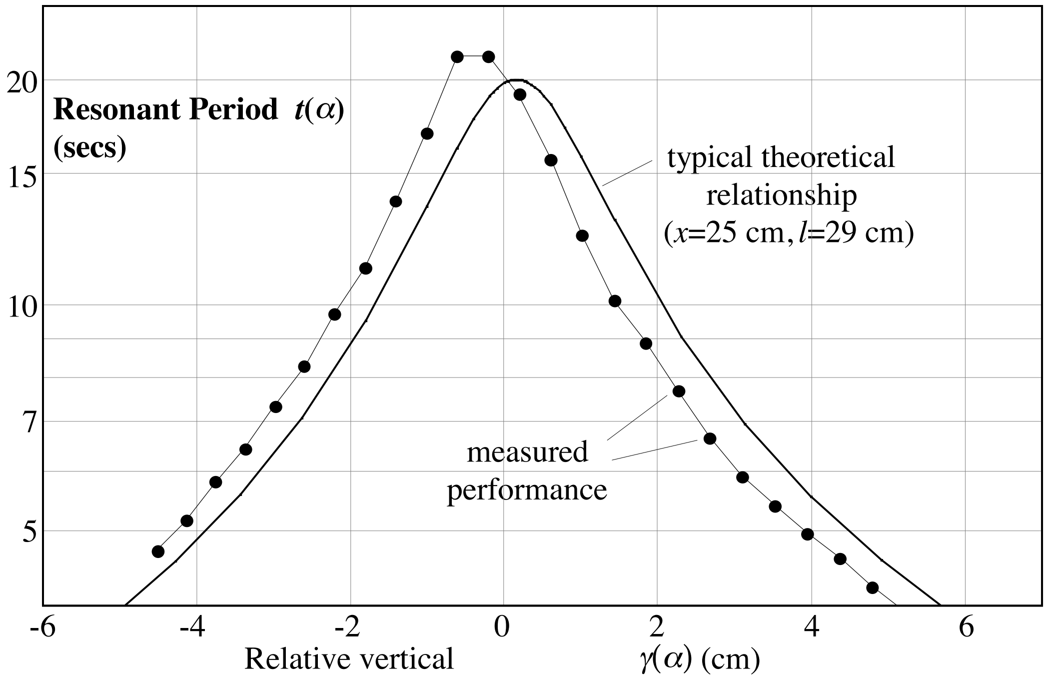

Figure 5.3 Measured resonant period tuning curve. Points show the resonant period measured as a function of vertical position (lower axis) or angle (upper axis). The curve is a theoretical resonant period curve assuming simple pivots (as opposed to the wrapped flexures used).

The torsion cylinder can be made from a coil spring or multiple concentric coil springs (used in torsion mode rather than tension). Alternatively many thin torsion-rods may be summed together in a cylindrical arrangement to allow support of arbitrary loads in a compact structure. It should be noted for practical purposes that in order to obtain a whole revolution of wind-up angle in a cylindrical rod without exceeding the yield point in a typical spring steel (e.g. SAE1095 with torsional elastic limit 0.69 GPa) a ratio of length to diameter in excess of 360:1 is needed. In addition a helical instability occurs for wind-up angles greater than about 7/8th of a turn unless the rod (or coil spring) is held in tension [Wahl 1963].

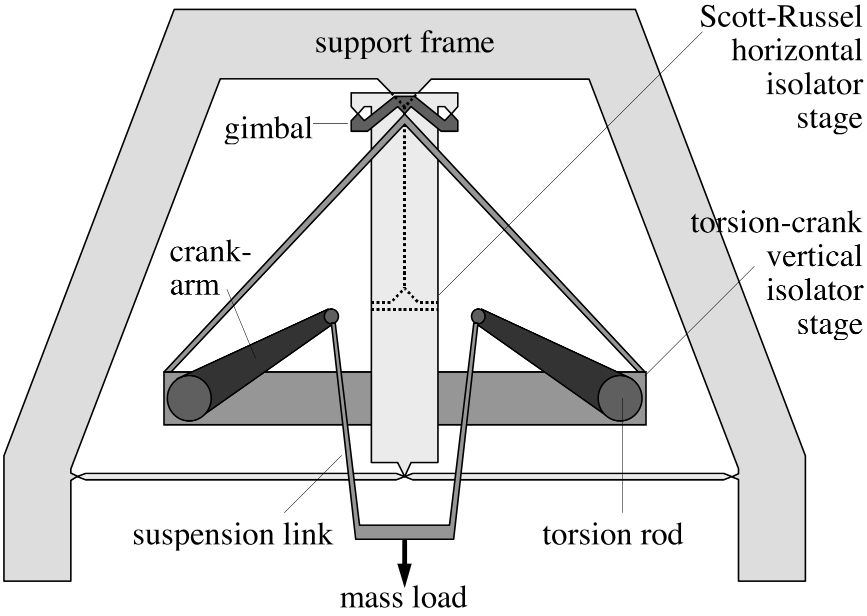

Figure 5.4 Schematic of a practical implementation. A torsion crank vertical isolator with dual torsion cylinders is shown suspended from a single beam Scott-Russel horizontal isolator. This provides full three-dimensional isolation to very low frequencies.

In order to achieve good isolation at frequencies a decade or so above the resonant frequency the crank-arm must be counter-weighted so that no vertical acceleration (to first order) appears on the load as a reaction to sudden vertical accelerations applied at the suspension point. In order to achieve this it is necessary to ensure that the instantaneous rotation centre of the crank-arm structure in response to vertical acceleration applied at the torsion-rod centre, occurs at the point where the link crosses an imaginary horizontal line through the torsion-rod centre. If the crank-arm is considered as a rigid massive beam, then it is necessary to simply extend one end of the beam with mass beyond the torsion-rod / link attachment points, to place its centre of percussion to line up (perpendicular to the beam) with the required instantaneous rotation centre. This dynamic motion null will remain effective over a larger vertical range for geometries which have close to 90 deg operating angle between the crank-arm and link.

Temperature compensation can be built into a design of this type by arranging that the differential expansion rates of say aluminium and steel, cause additional torque to be applied to counteract the normal reduction in modulus of elasticity with temperature. However it may be easier to control the temperature and even to use temperature as the means of level adjustment control.

5.2.7 Conclusion

We have introduced and demonstrated a torsion crank suspension system suitable for gravitational wave detectors. The TCL has a large vertical operating range and can provide exceptional vertical seismic isolation at very low frequencies. Cascaded with a Scott-Russel ultra-low frequency two-dimensional horizontal stage, this structure offers full 3-dimensional isolation to very low frequencies. Such isolation is expected to dramatically improve interferometer performance by greatly reducing the control forces required to maintain dark fringe locking.

5.2.8 Acknowledgment

This work was supported by the Australian Research Council.

5.3 Graphical Solutions of the Parameter Space

5.3.1 Introduction

The torsion-crank vertical suspension structure presented in the preceding paper 5.2, after choosing the scale of the device by fixing the crank-arm length r, had three other parameters which could be adjusted to alter the force-displacement characteristic :- (1) the length l of the supporting link, (2) the offset distance x from crank-arm centre to the vertical path followed by the end of the link, and (3) the amount of pre-stress ϕ in the torsion spring. Adding the condition that the gradient of the force-displacement characteristic should approach zero at some point (eqn. (5.3)) which should also be a point of inflection (eqn. (5.4)), reduces these variable parameters to only two. For the intuitive graphical display of the following figures, we chose to span this solution space using l and x as the undetermined parameters and the equations were solved for the remaining parameter ϕ. The corresponding values for the operating angle (α_0), the operating twist in the torsion rod (ϕ-α), and operating range (with 10 < period < 20 sec) were plotted as contour graphs on a grid spanned by $10<x<40$ cm and $10<l<50$ cm.

Figure 5.5(a) shows how the solution space area relates to the torsion-crank structure. If the top left corner of a graphical area is placed at the appropriate position relative to the centre of rotation of the left crank arm, then the l parameter is the height that the suspension link would hang vertically down from a horizontal crank, and the x parameter is the offset of the vertical path constraint. “X” marks the spot on the graphs where the solution values may be read for this l and x.

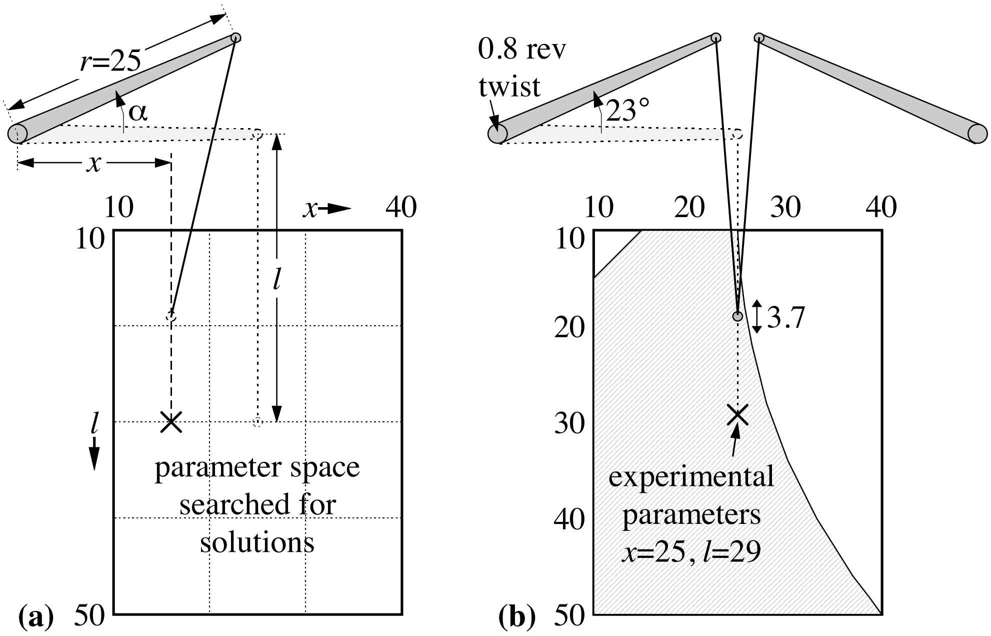

Figure 5.5 (a) Indicates the area of parameter space searched for solutions. It is bounded by 10<*x*<40 cm and 10<*l*<50 cm. (b) **X** marks the point in parameter space of the experimental measurements of the preceding paper, the diagram of cranks and links above it being a scale representation.

Figure 5.5(b) shows a scale representation of the experiment described in the preceding paper 5.2, and the "X" indicates where its parameters of x=25 cm and l=29 cm would place it in relation to the parameter space searched. The angle of 23 deg indicated is the operating angle α , the twist of 0.8 revolutions is the amount of twist in the torsion rod from the point at which it is fixed to the point of attachment to the crank. Finally the range indicated by the double headed arrow just beside the join in the suspension wires indicates the operating range for which the period remains between 10 and 20 seconds. These three values may be read off the following contour graphs for any set of l and x parameters for which a solution was found. The area for which solutions were found is indicated by the cross-hatching. There are two areas with no solutions. The structures specified by the large area on the right do not have solutions to the equations, while those specified by the small area at the top left cannot be constructed. (It should be noted that the contour graphs are rotated 90 deg from these to give a more normal aspect ratio).

5.3.2 Solutions at Extremities of Parameter Area

Figure 5.6 shows scale representations of the solutions at all four corners of the parameter area searched. To the left of centre the cranks cross each other at low elevations, while to the right the linkages form up-side-down "W" shapes. Also the geometries to the left have quite acute operating angles between the crank-arm and link, while those on the right approach 90 deg. Structures with this angle acute have much larger forces acting radially on the torsion support and also in the wire links if they are at shallow angles such as is the case in (a). However they require much less wind-up angle in the torsion spring and so are much easier to realise with a reasonable length torsion rod. It should be noted that the angle between the crank-arm and the link must always be slightly acute. It is this angle increasing towards 90 deg as the load descends which increases the mechanical advantage and cancels most of the torsional spring-rate.

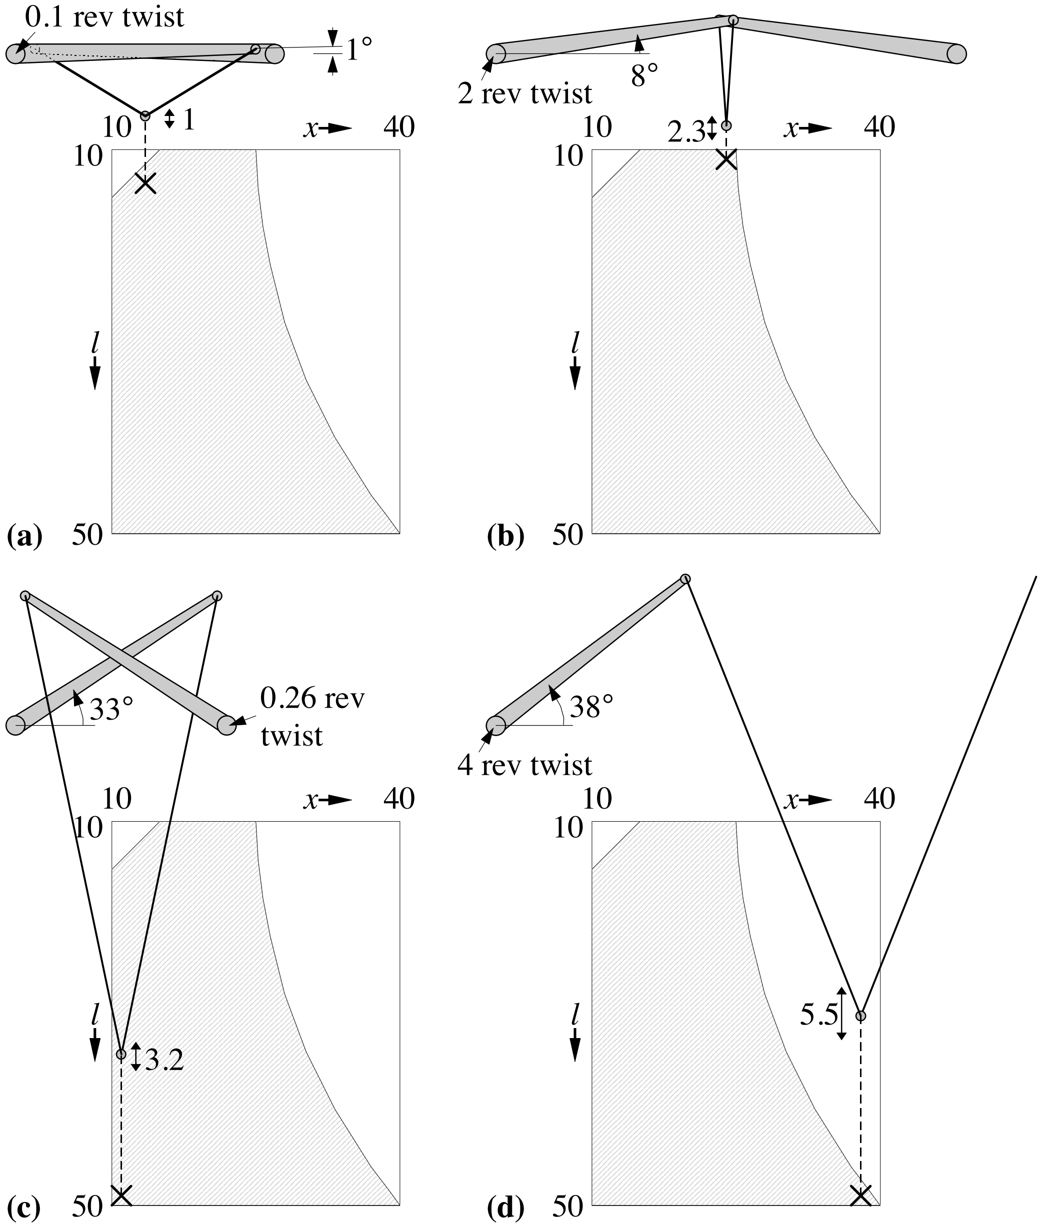

Figure 5.6 Solutions at extremities (marked by X) of searched parameter area. (Axes units are cm).

5.3.3 Operating Angle α_0

Stepping through values of x and l we obtained ϕ and α_0 by solving equations (5.3) and (5.4) simultaneously. Figure 5.7 shows the value of α for which the period peaks at 20 seconds ($t(\alpha_0)$ = 20 sec). All large peak-period values occur at about the same angle as is evident from figure 5.2.

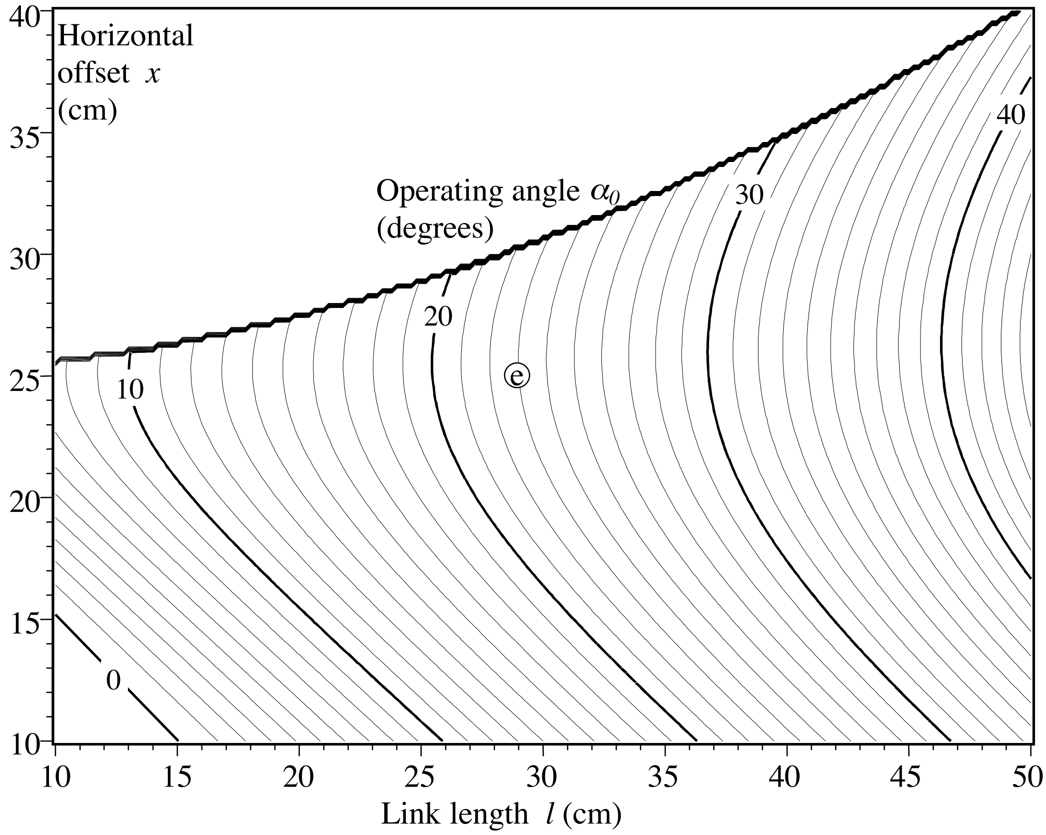

Figure 5.7 **Operating angle α_0.** A 20 sec period peak is accessible for the range of offset distances and link lengths shown. This period occurs at an operating angle α_0 indicated by the contour curves. The values used for the experiment are marked by (e).

5.3.4 Operating Twist (ϕ-α_0)

Figure 5.8 shows the operating twist in the torsion rod (ϕ-α_0 measured in revolutions) required to achieve this 20 sec period. This is generally the main determining parameter of the design. This value also does not vary much for changes in peak resonant period – figure 5.2 indicates that a change of only 5 deg (0.014 rev) changes the peak period from about 15 to 25 sec. The operating twist required to achieve this 20 sec peak-period escalates rapidly to infinity just beyond the 10 rev contour, and the area above this and below zero have no physical solution.

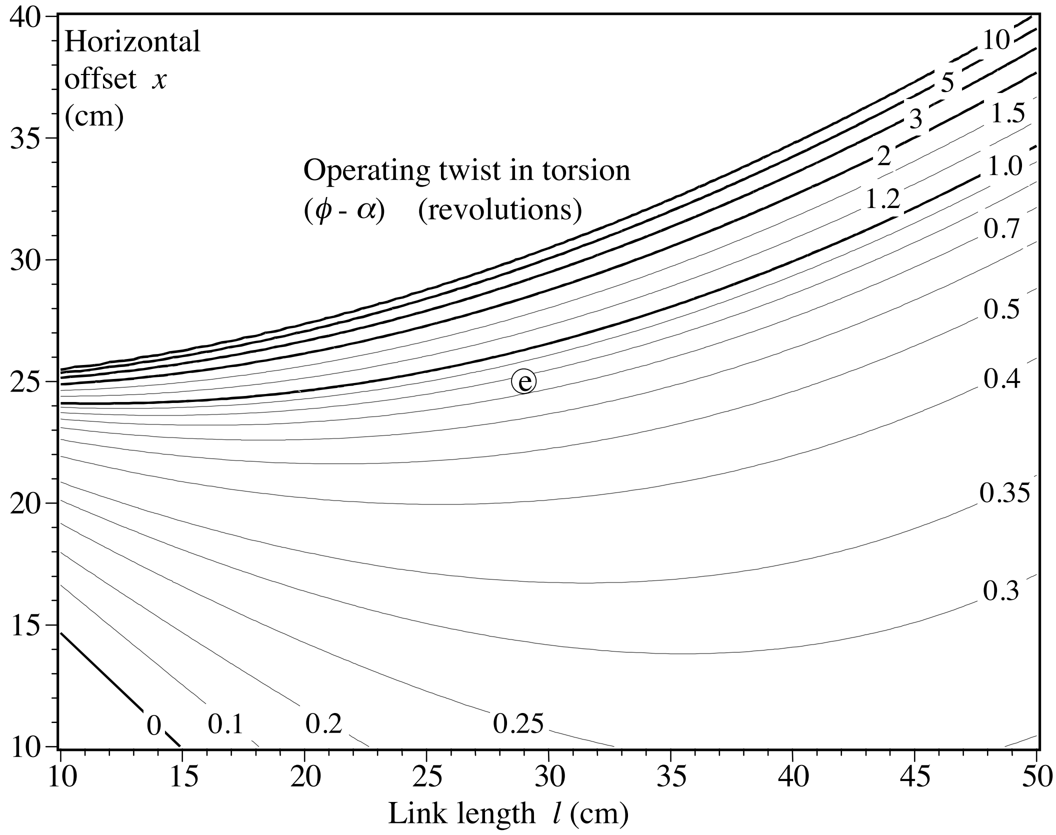

Figure 5.8 **Operating twist.** The total torsional twist in the torsion rod required to achieve a 20 sec period-peak for a range of offset distances and link lengths. The values used for the experiment are marked by (e).

5.3.5 Vertical Range

Figure 5.9 shows the expected vertical operating range for each setting ($y(\alpha_1)$–$y(\alpha_2)$ such that $t(\alpha_1) = t(\alpha_2) = 10$ sec). Longer values of l and wider spacing between the torsion rods (greater x) give larger operating ranges. However the height of the resultant structure increases faster than does the operating range. So the optimum geometry depends very much on the available space for mounting a practical device and the amount of twist that can be given to an available energy storage element. The flexibility of the system is remarkable however in the wide range of geometries and torsion elements that can be used to provide long period resonance and large vertical operating range.

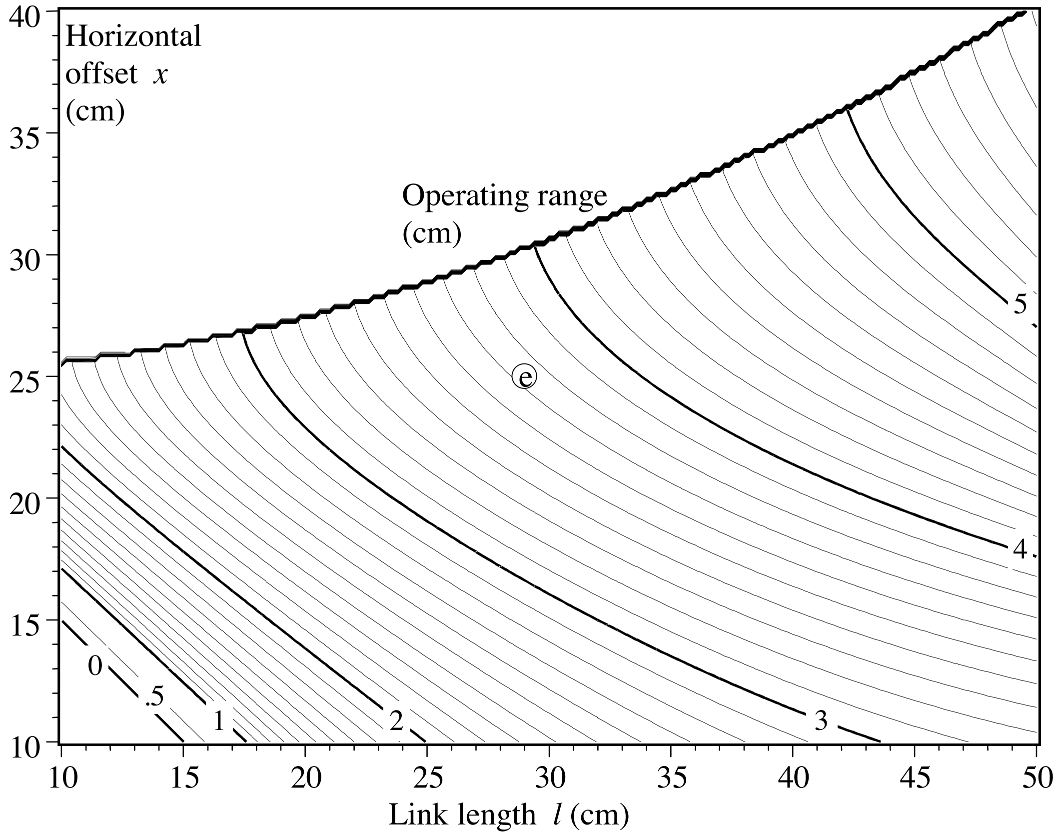

Figure 5.9 **Vertical range.** The vertical position range over which the resonant period remains between 10 and 20 sec for a range of offset distances and link lengths. The values used for the experiment are marked by (e).

5.3.6 Conclusion

The torsion-crank vertical suspension has several degrees of freedom available to its design parameters. In this report the effect of varying these parameters has been investigated numerically to obtain contour plots of operating conditions and expected performance. These plots allow a designer to quickly choose suitable operating parameters for a particular application.

5.4 Postscript

This device was intended as a single attachment point device (i.e. not tilt rigid) able to give high performance (requiring a high Q-factor to allow ultra-low frequencies to be obtained). While a great idea, and one that served as a seed [DeSalvo 1998] for similar ideas such as the so called “geometric anti-springs” [Bertolini 1999], we currently intend to replace it with an extra large and low frequency version of the Euler spring stage. The low measured Q-factor of Euler stages is a caution, but should be solvable by careful engineering of the boundary clamping condition, (maybe cutting from monolithic block), and by use in a vacuum. The small dynamic range of the Euler technique is not a disadvantage in the case of double pre-isolation (i.e. preceding it with the servo-frame of section 7.2) because the servo-frame provides any macroscopic translation required, leaving only seismic level motion required of the Euler stage. In addition the major advantage of the Euler technique is that very lightweight springs are used which require no large strains associated with initial loading.

6. Inverse Pendulum Horizontal Pre-Isolators

6.1 Preface

During this PhD research, the author had the opportunity to work with the Virgo group in Pisa several times. During the first visit the author worked on a high-tech active damping system intended to be applied to the second superattenuator in the Virgo suspension chain. Unfortunately there is not the time to include a report on this interesting project in this thesis but it was as a result of this project later being abandoned that the inverse pendulum project at Virgo was started.

The active damping approach was abandoned around the time that the author was present on a second visit to the Virgo group. The hope at that time was still to be able to actively damp filter #2 of the suspension chain. Instead of using 6 accelerometers mounted *on* the filter itself (with the troubles of tilt appearing as horizontal acceleration), the idea was to suspend a proof mass (of some kilograms) with an ultra-low frequency suspension and then servo the filter to follow the proof mass. So initially the wobbly table was to be quite small to provide the horizontal suspension for this proof mass. The vertical suspension was to be a LaCoste horizontal pendulum with a short spring and long lever arm. The author drew up and had made a small triangular prototype table with 60 cm legs which he then spent a lot of time testing thoroughly.

It has always been the author's contention that ultra-low frequencies are easiest to achieve with big dimensions and large masses. So when the idea of suspending the entire Virgo suspension chain from a very large wobbly table was discussed, it had the author's enthusiastic support. At the time there was ongoing development of a system to allow the chain suspension point to be servoed around by some centimetres while rolling on very large ball bearings. The fact that a large inverse pendulum stage would also provide this facility in a perfectly noise free manner, made the decision to switch to it a lot easier.

On his next visit to the Virgo group, the large inverse pendulum table had been built and was working with some load masses suspended from it. It became the author's job together with Giovanni Losurdo to measure the transfer function of the table to make sure that it was working as expected. We removed the load masses to avoid the extra resonances produced by the suspended masses and put the masses on the table top (on nuts to give a firm seating). Early on we tried to obtain a transfer function by relying on natural seismic motion as the shaking signal and left it accumulating a measurement all night long. For some reason we could not get a good result in this manner (probably it was the low amplitude of the seismic in relation to air disturbances). The transfer functions we measured in this manner did not fall off as $1/f^2$ but were more like $1/f^{1.5}$ or even $1/f$.

The author had always obtained good results from swept sine measurements owing to the high signal to noise ratio and he suggested that we try to shake test the entire 2 tonne structure - base, stand and inverted pendulum. Someone (Riccardo DeSalvo?) had foreseen that some rollers for moving the structure might be desired at some stage and some *x-y* roller plates had already been manufactured for the job. However the problem of how to shake the structure had not be considered.

No suppliers of suitable shakers (at reasonable cost) could be found and so we decided to build a pneumatic one using a double acting piston that was available, together with proportional solenoid valves and position sensor from Radio Spares, and the digital signal processor (DSP) controller that the Virgo group designed and built for their own use. This shaker setup is shown in figure 6.10 (which the author drew).

The high pressure air supply from the compressor (varying from about 7.5 bar up to 11 bar) was first regulated to a constant value (as high as possible or 7.5 bar) and this supplied via a proportional valve to one side of the piston. A second proportional valve was also fitted to allow this side of the piston to be exhausted to the room. This allowed a variable pressure from 0 to about 7.5 bar to be applied to this “pushing” side of the piston. The other side of the piston was fed with a middle bias pressure of approx 4 bar by a second regulator. This gave the required “pull” action for low applied pressures and “push” action for high applied pressures.

In order to obtain the best possible response speed from the system (i.e. for shaking at 10 Hz or so), the variable pressure end of the piston was arranged to have the minimum possible volume. This possibly made the control system harder to stabilise as the response parameters vary strongly within a short distance, but was necessary for high speed performance. The motion of the push rod was limited by a (piston mounted) plate and (casing mounted) rods and nuts, to stay within a few millimetres of this end limit.

The DSP controlled the solenoid valves in a proportional manner with a specially modified power driver module. (The module was modified from the standard +/- 24 volt version so as to only produce a positive voltage +2.5 V to +42.5 V to prevent the control polarity from inverting.) One of the problems found with the system was that the pressure between the solenoid valves readily saturated to the high or low limit. This immediately causes nonlinearity in the control loop and with high gain causes overshoot of the integral action and oscillation. To prevent this a pressure sensor was fitted to allow the control system to be stabilised more readily.

The shaker system was marginally underpowered and it was found that frequency sweeping was best done in a direction from high amplitude (usually low frequency) to low amplitude (usually high frequency). This was because the balls rolling in the V-grooves built up little humps at the extremities of travel which become too hard for the shaker to push past if increasing amplitude is required.

Even with such a good strong signal, we were still not able to measure the expected transfer function and suspected air coupling between the inverse pendulum and a low frequency resonance of the tall surrounding framework. Initially the author damped this resonant motion at the top of the frame with a large damper consisting of one set of aluminium plates fixed to the frame and an interleaved set connected with a push-rod to the wall with oil between them for viscous friction. This helped a little but still the transfer function did not match theory.

The final solution was to unbolt the entire fixed framework and support it from the ceiling with the crane. This left just the base and inverse pendulum table to be shaken. This finally produced excellent measurements and we were able to explore the variation in transfer function with the critical centre-of-percussion adjustment. It was just as well that we did because the correct adjustment was not possible within the range of the available adjustment. It needed the available adjustment to be set to maximum and another 2.7 kg of lead added to each of the three legs (see figure 6.13).

The paper reporting these results (following in section 6.2) was written by Giovanni Losurdo as he had a lot more access to other drawings, measurement results, and information after this author had returned to Western Australia. The reviewers comment is worthy of note :

“The article describes the application of a very simple concept - the inverted pendulum - in a very extraordinary application - test mass suspension for gravitational wave detection. The parameters are unprecedented: height (> 6 metres), suspended mass (~ 1 tonne), and resonant frequency (~30 mHz), and the details of methods used, problems encountered and results obtained are of great interest and should be published.”

6.2 VIRGO's 6 Meter Inverse Pendulum Pre-Isolator

An inverted pendulum preisolator stage for the VIRGO suspension system.

G. Losurdo

Istituto Nazionale de Fisica Nucleare, Sez. Di Pisa, Via Livornese 1291, 56010 S. Piero a Grado (Pisa), Italy and Scuola Normale Superiore, 56100 Pisa, Italy

M. Bernardini, S. Braccini, C. Bradaschia, C. Casciano, V. Dattilo, R. De Salvo, A. Di Virgilio, F. Frasconi, A. Gaddi, A. Gennai, A. Giazotto, H. B. Pan, F. Paoletti, A. Pasqualetti, R. Passaquieti, D. Passuello, R. Taddei, and Z. Zhang

Istituto Nazionale de Fisica Nucleare, Sez. Di Pisa, Via Livornese 1291, 56010 S. Piero a Grado (Pisa), Italy

G. Cella, E. Cuoco, E. D'Ambrosio, F. Fidecaro, S. Gaggero, P. La Penna, S. Mancini, R. Poggiani, and A. Vicere`

Istituto Nazionale di Fisica Nucleare, Sez. Di Pisa and Dipartimento di Fisica, Universita` degli Studi di Pisa, 56100 Pisa, Italy

M. Mazzoni and R. Stanga

Dipartimento di Astronomia, Universita` di Firenze, 50100 Firenze, Italy

L. Holloway

Department of Physics, University of Illinois, Urbana, Illinois 61801

J. Winterflood

Department of Physics, University of Western Australia, Nedlands 6807, WA, Australia

The design of a new preisolator stage for the VIRGO superattenuator is presented. The device is essentially a 6 m high inverted pendulum with horizontal resonant frequency of 30 mHz. An isolation of 65 dB at 1 Hz has been achieved. Very low forces are needed to move the whole superattenuator acting on the inverted pendulum. For this reason, the system is a suitable platform for the active control of the mirror suspension.

6.2.1 Introduction

The sensitivity of interferometric antennas for gravitational wave detection [Abramovici 1992, Bradaschia 1990, Hough 1994, Kuroda 1997, McClelland 1994] is limited at low frequencies by seismic noise. Various suspension systems to isolate the mirrors from ground vibrations have been studied and developed by the different groups involved in this field. The attenuation performance of a suspension system and, consequently, the low frequency threshold of the detection range is determined by the resonances of the suspension itself. Several devices to improve the low frequency performance are being studied worldwide [Barton 1996, Winterflood 1996, Winterflood 1998]. In this article we present the design and prototype testing results of the horizontal element of the ultralow frequency seismic preisolator stage (PIS) for the VIRGO superattenuator (SA) chain suspension system [DeSalvo 1996, Holloway 1997, Losurdo 1998].

The device producing ultralow frequency horizontal oscillations is a 6 m high, three-legged inverted pendulum (IP). On the top of this structure a special filter (called *filter 0*) conceived as the vertical component of the preisolator [DeSalvo 1999] is connected at three points to form a wobbly table moving in the horizontal plane. This design, if well aligned, achieves nearly complete decoupling of the vertical from the horizontal degrees of freedom (dof). The PIS has been specified and designed to have resonance frequencies of 30 mHz or lower.

Additionally, it must be free of resonances and other mechanisms that may hinder its performance as a seismic excitation barrier between ground and the SA chain resonances, which are spread between 0.1 and 2 Hz, both in the horizontal and vertical directions. A full scale prototype of the IP has been built and tested in Pisa. The IP tests described in this paper have been performed replacing the filter 0 (F0) by a disk of equal diameter and weight loaded with the mass corresponding to the SA chain weight. This article reports on the passive performance of the IP. A forthcoming article will report on the tests of position control, active damping and inertial seismic noise suppression in the PIS.

6.2.2 VIRGO Superattenuator

A well designed and reliable isolation system must be effective in six dof because, in a real mechanical system, angular and translational dof are unavoidably coupled.

In the horizontal plane (in the following we refer to an xyz frame, where x is the axis parallel to the laser beam, y is the horizontal axis perpendicular to x , and z is the vertical axis) the SA is essentially a multistage pendulum composed of a cascade of units, hereafter called *filters*, as sketched in figure 6.1. Each filter is connected to the previous and the following ones by two steel wires 1.15 m long, with a pendulum frequency of 0.45 Hz. The whole chain, from the suspension point to the center of the mirror, is about 9 m high, with the lowest normal mode frequency $f \approx 0.2$ Hz.

To generate low tilt frequencies, the filters are designed with a large moment of inertia and short lever arms between the point from which a filter is supported by its suspension wire and the point to which the wire supporting the next filter is attached. Also the two wires are connected as close as possible to the filter center of mass. This design leads to tilt mode frequencies below 1 Hz. Very low angular frequencies about the vertical axis are naturally obtained thanks to the low angular spring constant of the long, small diameter steel suspension wires.

Vertical attenuation is obtained by providing each filter with a set of high tensile strength steel, cantilever, triangular ‘‘blade’’ springs pointing to the center of the filter [Braccini 1996] from which the load of the following stage is suspended. With this design, the system acts as a chain of serially coupled oscillators also in the vertical direction. The lowest vertical resonance achieved with a single stage in this way is 1.4 Hz (to be compared with the 0.45 Hz horizontal pendulum frequency). A magnetic antispring system [Beccaria 1997] is used to further ‘‘soften’’ the springs and to tune the resonant frequency to 0.4 Hz. The SA chain is purely passive. No active elements are used between the PIS and the locking actuators of the mirror and the preceding

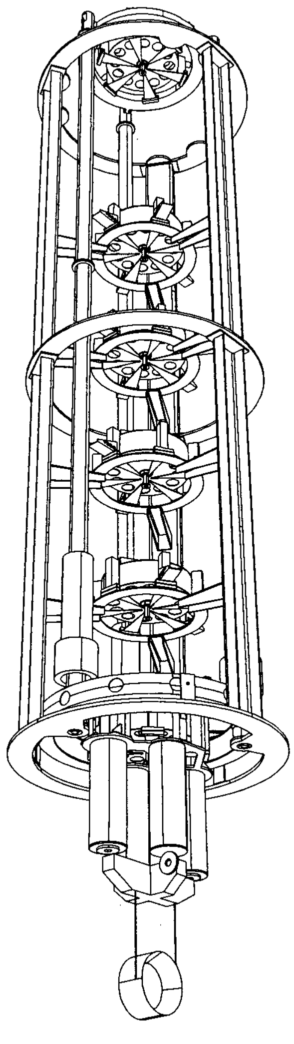

Figure 6.1 The superattenuator with the 6 m inverted pendulum preisolator stage. From the top to the bottom one can see the filter 0, four standard filters, the steering filter (with its “legs”), the marionetta and the mirror.

intermediate mass. Figure 6.2 shows the calculated SA chain transfer function, including the effect of the resonances of the cantilever springs. The composite resonances of the five filters can be observed to spread between 0.2 and 2 Hz. The overall attenuation at 10 Hz is expected to be 10^{-15}. The effect of the PIS will be discussed in section 0.

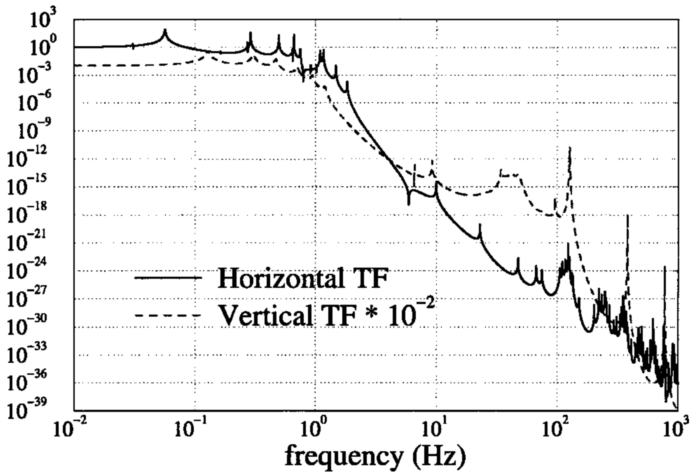

Figure 6.2 The SA transfer function magnitude, obtained assuming a vertical-horizontal coupling of 0.01.

6.2.3 The Preisolator Stage

The SA chain is suspended from the top stage which is designed to fulfill three main functions.

(i) To provide the SA with a suspension point positioning system. During locked beam operation the suspended mirror position (with respect to local ground) may drift substantially (~0.1 mm) over the time scale of days due to local thermal effects, tidal strains, etc. The amplitude of the slow drifts is far beyond the interferometer locking system dynamic range. Additionally, during out of lock operations the mirror suspension point may have to be moved by larger distances to optimise the interferometer working conditions. Each SA is then provided with a mechanical positioning system, able to set the SA chain suspension point in the operating position within a mechanical range of ±20 mm in the horizontal plane and ±2 mm vertically, and to control it dynamically within a few hundred microns. These movements must all be obtained minimising mechanical noise production and with small power consumption to avoid thermal disturbances in the antispring systems. Therefore, soft springs connecting the top table to a motor driven slide clamped on the external frame are used to set the SA in the right working position and to control the position at frequency f << 1 mHz, while coil–magnet actuators are used for the active control at higher frequency.

(ii) To introduce a very low frequency (about 30 mHz) filtering stage, limiting the amount of seismic energy feeding into the SA chain resonances in the 0.2–2 Hz range. The effect on the suspended mirror residual motion of the purely passive filtering generated by an ultralow frequency filter is illustrated in figure 6.3.

(iii) To provide a soft suspension stage at the top of the chain that allows low power active damping of the residual resonant mode excitation and seismic noise depression by means of position and inertial sensors coupled to electromagnetic actuators.

To fulfill all of the above requirements in the horizontal plane a device conceived on the principle of the inverted pendulum has been built and tested.

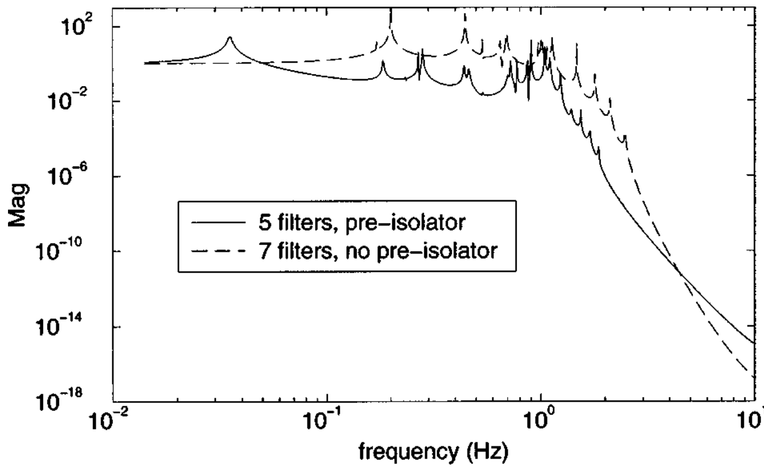

Figure 6.3 Calculated SA transfer function: (a) former SA with seven filters and no preisolator, (b) new SA with five filters and the IP preisolator stage.

6.2.4 The Inverted Pendulum Concept

An inverted pendulum is a simple device with rich dynamics (see for instance the works of Saulson *et al.*[Saulson 1994], Pinoli *et al.*[Pinoli 1993]). To understand the basic dynamics of an inverted pendulum it is useful to consider a simple ideal model schematised in figure 6.4. The model consists of a load of mass M supported by a massless rigid rod of length l, hinged by a flex joint fixed to ground. The hinge has an equivalent spring constant k for forces applied to the top of the rod. When the mass is displaced from the vertical position by a distance x ($x << l$) the total force acting on it is:

$$F = F_{el} + F_{grav} = (k - M\,g\,/\,l)x + O(x^3) \approx \tilde{k}x \tag{6.1}$$

When $\tilde{k} > 0$, the system is an oscillator resonating at frequency:

$$f_0 = \frac{1}{2\pi}\sqrt{\frac{k}{M} - \frac{g}{l}} \tag{6.2}$$

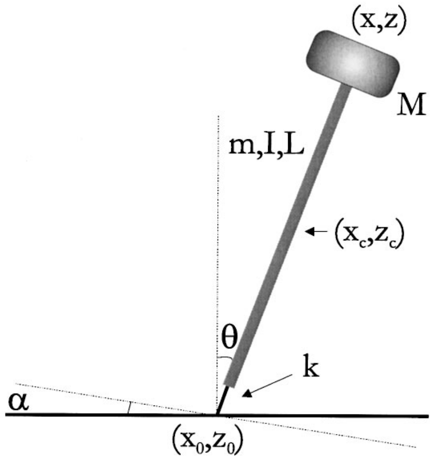

Figure 6.4 Schematic model of the inverted pendulum.

By properly tuning the spring stiffness and the suspended load one can obtain an arbitrarily small resonant frequency. When the load exceeds the critical value ($M > k\ l/g$), gravity dominates over the elastic forces ($\tilde{k} < 0$) and the system collapses.

An IP is an ideal device as a horizontal preisolator stage for several reasons:

(i) It acts as a vibration isolator in the horizontal plane and can be tuned to arbitrarily low resonant frequencies within contained dimensions (the nominal 30 mHz tuning frequency would be equivalent to a 280 m long pendulum).

(ii) Its geometry, extending upwards from the ground supports, provides gain of vertical height for the subsequent suspended pendula without building tall rigid structures.

(iii) The force required to move the IP a distance x at frequencies lower than the resonant frequency is:

$$F \approx M\omega_0^2\, x \tag{6.3}$$

Assuming a mass $M = 1$ tonne, at the nominal frequency of 30 mHz only 0.36 N are needed to translate the suspension point by 1 cm.

(iv) The IP is a convenient platform for applying active controls (resonance damping, inertial seismic suppression, position control). Because it is virtually orthogonal to the vertical and tilt dof, the horizontal dof can be treated separately. This greatly simplifies the feedback loops complexity.

The IP (and F0) are the first elements of the SA chain directly connected to ground and therefore are the best locations to take advantage of the still strong and detectable seismic signals. One can apply feedback active compensation without being limited by the sensor sensitivity. Other concepts of ultralow frequency horizontal oscillators have been considered. The IP concept was retained for its simplicity, reliability and its geometrical advantage in the VIRGO towers.

6.2.5 Design of the Preisolator Stage

A simple minded inverted pendulum could not be used as a SA chain base because its rigid arm would occupy the central axis from which the SA chain is supposed to hang. The problem is solved by using three legs linked together at the top by a rigid structure and hooking the load from its center. This geometry introduces a soft angular oscillation around the vertical axis, exploitable for fine angular positioning of the mirror. In a complete PIS the linking rigid structure would be the F0, which, in the tests reported here was replaced by a dummy filter (a rigid disk of similar mass and equal dimensions). The three IP legs are made of 6.18 m long, 130 mm outer diameter, 125 mm inner diameter aluminium tubes. The large diameter of the legs is chosen to provide rigidity and high resonant frequencies of the leg vibration modes (the first one is at 9 Hz). Thin (2.5 mm) walls minimise mass and moment of inertia which can become a bypass route transferring seismic noise to the suspended load. The legs are connected at the bottom to a flexible joint, a steel rod of length and diameter adapted to generate the required elastic constant k at the top of the legs. A length $l = 216$ mm and a diameter $d = 27$ mm have been chosen for a load M of about 1 tonne. The joint is made of Maraging steel (Marval 18), chosen because of its high mechanical stability and its low internal dissipation factor [Beccaria 1998]. IP tests performed with elastic joints manufactured with AISI40 steel have shown mechanical quality factors less than half those made with maraging. Due to mechanical imperfections, the legs top end may not coincide perfectly with the corresponding suspension points of the F0 (this happens, for instance, if one of the legs is bent and then not perfectly vertical). Therefore, linking the top end of the legs of the IP to the disk introduces an offset force which, if not compensated, displaces the IP from the nominal zero position. Such offsets are counter-balanced by means of three parasitic springs. One end of each spring is connected to the top of the leg and the other to a motor driven slide clamped on the external frame. A displacement sensor is used to monitor the initial offset and to provide error signals to the motors. This system does not feed seismic noise to the IP top disk since the elastic constant of the parasitic springs is much smaller than the one of the IP flexible joints.

To minimise the angular rigidity of the system, the top end of each leg is provided with a rigid loop from which the F0 hangs by means of maraging wires 3 mm in diameter and 31 mm long. With this solution a negligible elastic constant is added to the transverse and torsional dof. The three connection wires constitute a high frequency pendulum which is effectively mounted in series with the IP one. Keeping the wires as short as 31 mm pushes this additional frequency above 3 Hz, far enough from the main resonant frequency that it does not generate any interference problem.

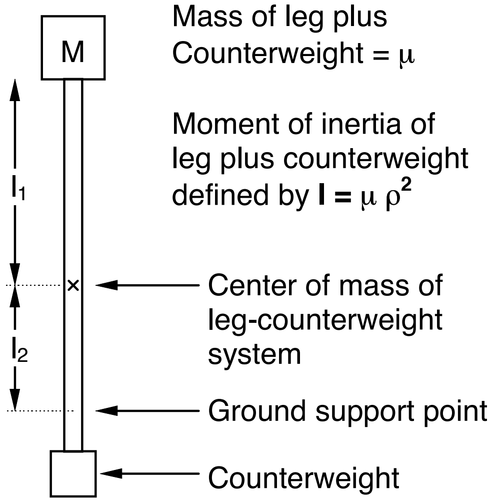

Figure 6.5 Schematic diagram defining the parameters of the moment of inertia effect.

The tops of the legs are made of titanium to provide rigidity while keeping the leg mass to a minimum. Although the legs have been built as light as possible, they still have a substantial moment of inertia. The fact that the supporting legs of the inverted pendulum have a finite moment of inertia affects the transfer function at high frequencies. Referring to figure 6.5, let the mass of the supported weight be M, the mass and length of the supporting legs be μ and L, and the moment of inertia be defined by $I = \mu\rho^2$ (where ρ is the *radius of gyration*). Then the transfer function of the displacement of the table x, with respect to ground x_0, is

$$x = x_0 \frac{\omega^2 - \omega_z^{\ 2}}{\omega^2 - \omega_p^{\ 2}} \tag{6.4}$$

where the pole frequency is given by

$$\omega_p^{\ 2} = \left[\frac{k}{M} - \frac{g}{l}\left(1 + \frac{\mu}{M}\frac{l_1}{L}\right)\right]\left[1 + O(\mu / m)\right] \tag{6.5}$$

and the position of the zero is given by

$$\omega_z^{\ 2} = \omega_p^{\ 2}\left(\frac{L^2}{l_1 l_2 - \rho^2}\right)\frac{M}{\mu}\left[1 + O(\mu / m)\right] \tag{6.6}$$

Here, l_1 and l_2 are the distances from the center of mass of the support leg to the seismic driving point and the table suspension point. The condition $\rho^2 = l_1\, l_2$ moves the frequency of the zero to infinity and, as a consequence, an impulse applied at l_1 is transmitted minimally to l_2. This effect is well known to batsmen in cricket and baseball.

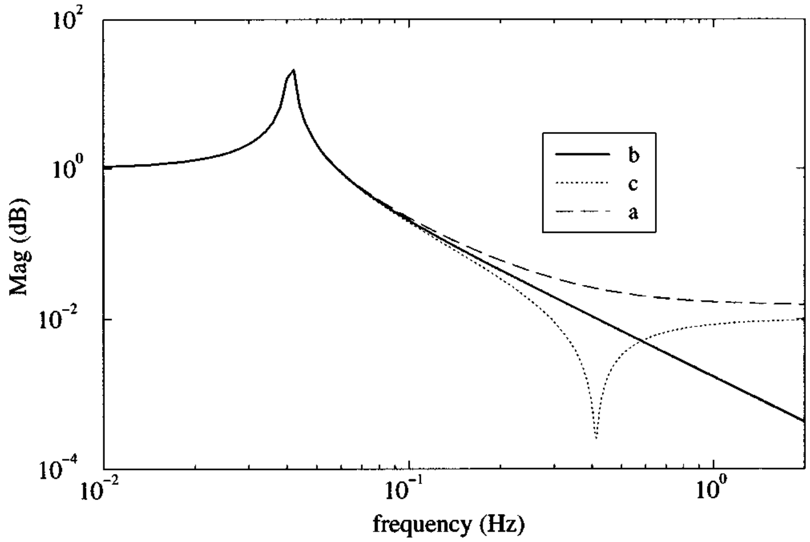

Figure 6.6 Calculated transfer function for various values of ρ^2 : (a) $\rho^2 > l_1 l_2$, (b) $\rho^2 = l_1 l_2$, (c) $\rho^2 < l_1 l_2$.

The value of ρ^2 can be adjusted by affixing a counterweight below the support point. In figure 6.6 the calculated effects of various values of ρ^2 on the transfer function are shown. To fix this problem counterweights can be provided on leg extensions so that either the flexible joint or the wire connecting the leg top to the F0 be placed at a percussion point of the leg. The percussion point is defined as the point around which a leg will rotate if pushed with an impulsive transversal force applied to its farthest end. The counterweight reduces the seismic energy transmission caused by the inertia of the legs and the unwanted saturation in attenuation properties. It was mechanically simpler to fix the counterweight below the flex joint.

The mechanical problem of bringing the leg percussion point to coincide with the hinging point was solved by mounting each flex joint on top of an 0.8 m long rigid column; a counterweight is attached to a bell-shaped, 0.9 m long skirt bolted to the bottom end of the leg. This skirt wraps the joint and its support column as shown in figure 6.7. The required counterweight mass is about 25 kg. The mass or the position of the counterweight must be fine tuned to minimise the excitation transfer. The three supporting columns were solidly attached on a 2 m diameter rigid ring (first resonance at 150 Hz) which is itself supported by three tunable feet. In order to allow for the necessary tunability each foot was built in the form of a stack of twin cup springs, very rigid in all dof except in the vertical direction where an elastic constant of 10^6 N/m is allowed for (figure 6.8). The foot supports the weight of the bottom ring over a

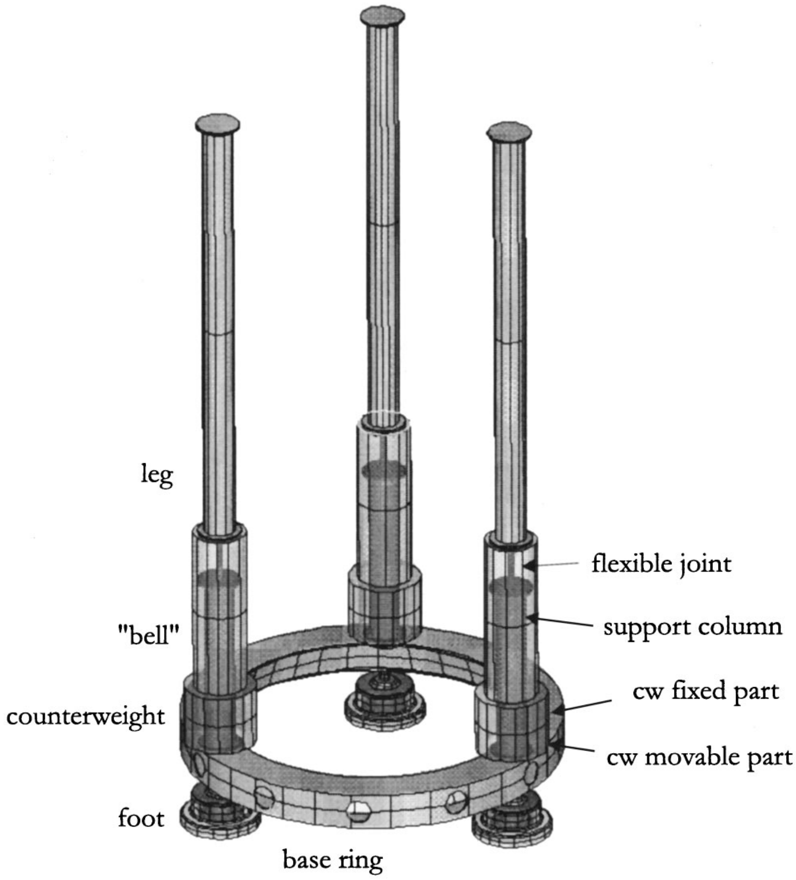

Figure 6.7 View of the bottom part of the inverted pendulum. From the top to the bottom one can see part of the legs, the light bell-shaped structure, the counterweight, the ring, and the feet. In transparency the support columns and the flexible joints are visible.

hemispherical stud screwed on the foot central body. The horizontality of the base ring can be adjusted by turning the fine thread screws of the studs. A piezoelectric actuator, mounted coaxially below each foot, allows dynamical control of the base ring tilt.

The horizontality of F0 is important: to ensure that the effective length of each leg can be adjusted with two counter-threaded nuts moving the wire attachment points with respect to the F0 support arms. The parallelism of the legs will be guaranteed only if the three top hinging points of the F0 are positioned on a triangle identical to that on which

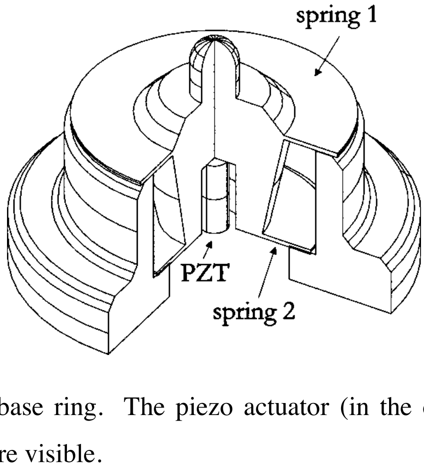

Figure 6.8 Foot supporting the base ring. The piezo actuator (in the center) and the twin cup springs providing the vertical elasticity are visible.

the legs elastic joints are located. In this case the F0 will remain horizontal during the IP motion. In the case of converging or diverging legs the F0 would move along a curved surface. This movement would affect the performance of the horizontal accelerometers mounted on the F0 body since they are also sensitive to tilt. A dimensional tolerance of 0.1 mm was placed on the positioning of the columns on the bottom rings. This precision is achieved by using the F0 body as a column assembly jig.

6.2.6 Experimental Setup

The entire IP setup was mounted on special roller skates to allow for two-dimensional horizontal motion. The dummy filter was loaded with lead blocks to tune the horizontal resonant frequency at 30 mHz. This value was easily achieved preserving mechanical stability. The required mass was less than predicted. Indeed, the elasticity of the legs, the support columns, the base and the ground had been neglected in the calculations causing an overestimate of the stiffness. In the future the joints stiffness will be over-dimensioned with a safety factor, and the resonant frequency will be tuned to the desired value adding ballast weights.

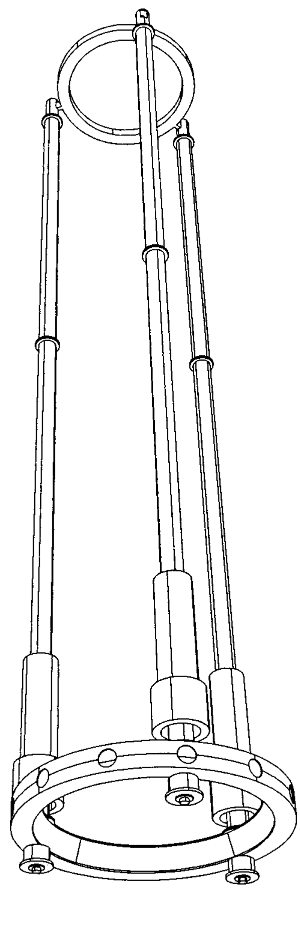

Figure 6.9 IP prototype used for the tests. Accelerometers were clamped on the top table and on the base ring. The whole system was layed on a steel base which could slide over rolling plates, not shown in this drawing.

The tests of the IP were performed with the SA chain weight placed on the top table instead of hanging from it. This allowed us to better characterise the IP: hanging weights would have introduced their own resonances on the transfer function measurement, thus making it more difficult to distinguish the IP properties to be measured. To measure the IP performance, two horizontal accelerometers [Braccini 1995] were used: one clamped in the center of the top table, and the other clamped on the IP base ring, parallel to the first. The accelerometers work in the band 0-400 Hz, with a sensitivity better than 10^{-9} (m/s^2)/√Hz below 2 Hz. The IP transfer function is measured as the ratio of the top to the base acceleration spectra.

The IP prototype was mounted in a noisy environment: neither vacuum system nor thermo-stabilisation were present. The motion of the top table was affected by air currents overwhelming the attenuated seismic signal. In order to raise the attenuated signal above the noise level and measure a meaningful transfer function, an ‘‘artificial seism’’ had to be provided, with a large amplitude. An excitation system (shaker) was designed and built for this purpose. The shaker was made by a compressed air double

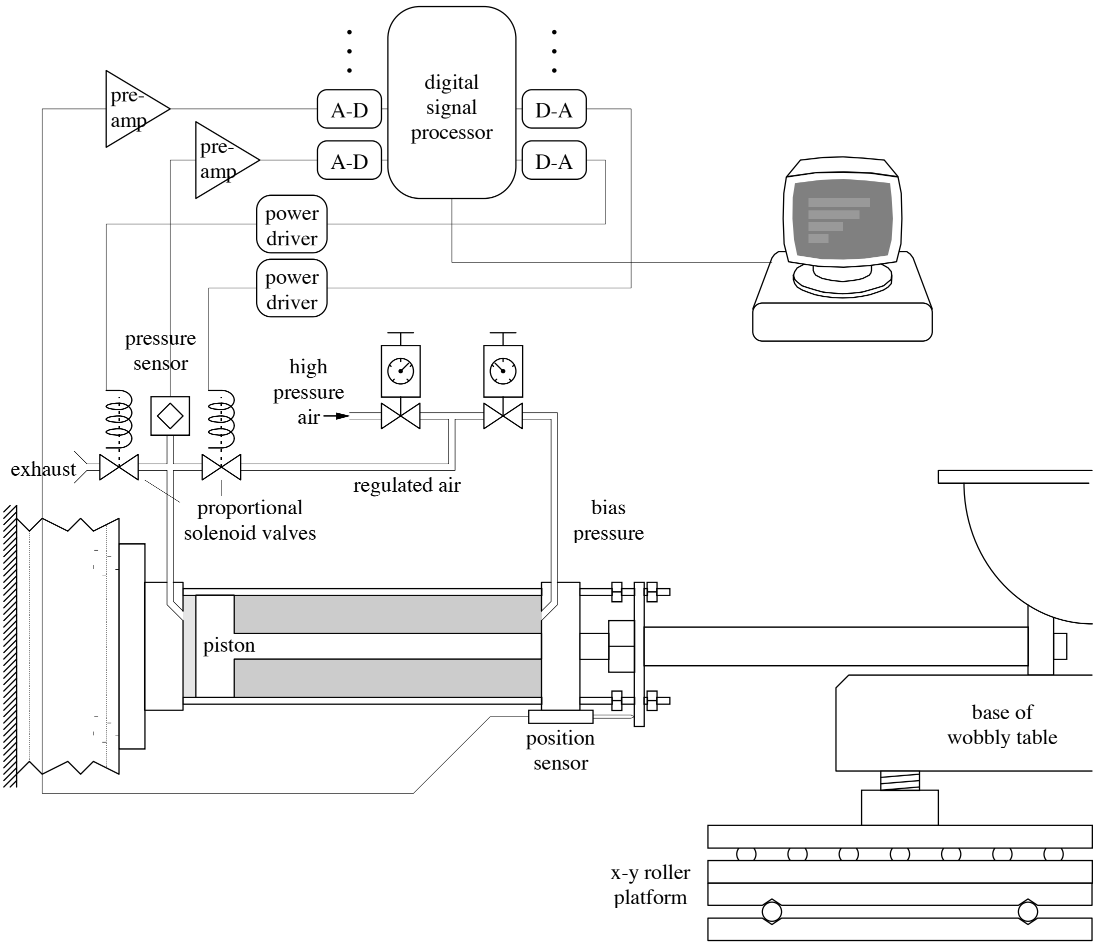

Figure 6.10 The pneumatic shaker used for IP excitation with the control setup: the shaker is mounted on the wall and connected to the IP base.

action piston, shown in figure 6.10. The motion of the piston was sensed by a linear potentiometer and feedback controlled through a programmable digital signal processor (DSP). The DSP was calibrated and programmed to slave the piston movement to an external white noise or to a fixed frequency sinusoidal signal. The body of the piston was solidly attached to the laboratory wall and its stem was linked to the base of the inverted pendulum setup. The roller skates allowed for low friction movement in one or two directions. The active control piston attached to the almost 2 tonne system presented a useful band pass extending from dc to about 20 Hz.

The experimental setup is sketched in figure 6.9 and figure 6.10.

6.2.7 Resonant Frequencies and Quality Factors

In figure 6.11 the measured horizontal resonant frequency is plotted as a function of the load. The curve is fit by the function:

$$f_0 = \frac{1}{2\pi}\sqrt{\frac{k - (M/3 + m/2)g/l}{M/3 + m/2}} \tag{6.7}$$

where M is the load (dummy table + added load) and m is the leg weight. Equation (6.7) has been derived treating the legs as uniform beams and neglecting the contribution of the bell-counterweight device to the resonant frequency. From the fit one finds the value of the joint stiffness: $k = 566 \pm 10$ N/m. A finite element analysis of the system was performed and the predicted stiffness was $k_{th} = 570$ N/m, in good agreement with the measured value.

Some measurements of the IP resonance quality factor were also performed and are

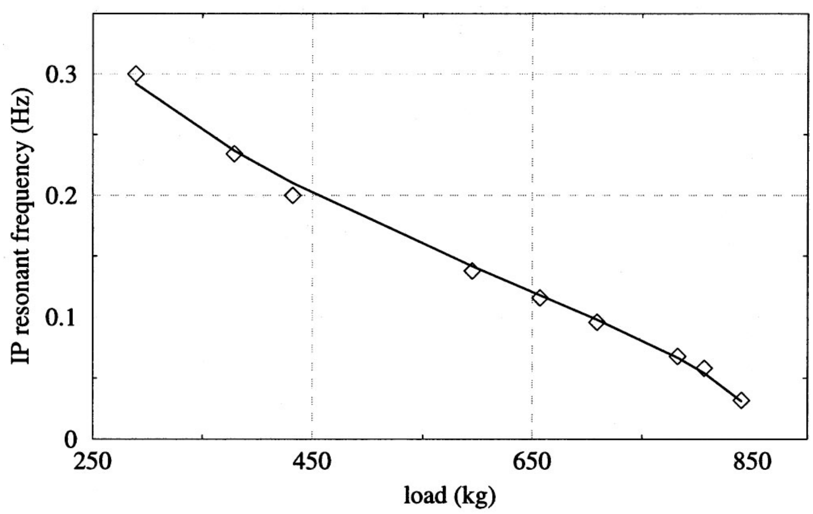

Figure 6.11 IP horizontal frequency measured as a function of the added load (the zero of the x axis corresponds to the weight of the dummy filter: $M_0 = 296$ kg). The fitting function has been obtained from a simple model of IP with homogeneous leg. The linear stiffness turns out to be 566±10 N/m.

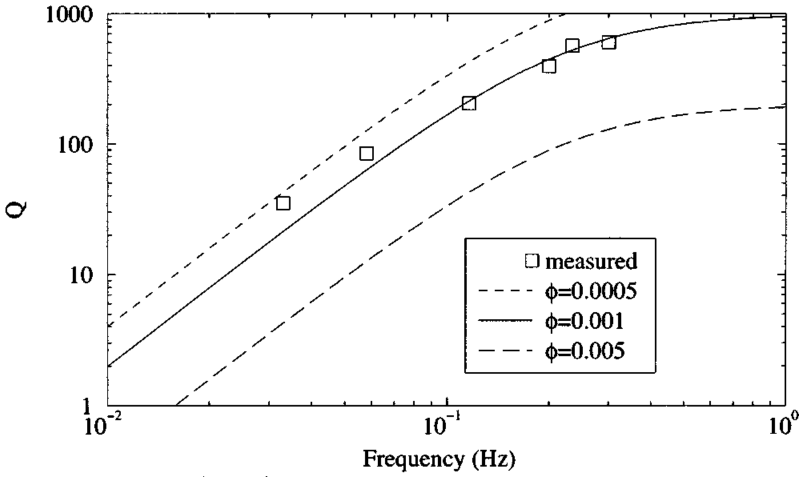

Figure 6.12 IP quality factor as a function of the resonant frequency. The data are compared with the theoretical predictions for structural damping for different values of the loss factor: $\phi = 0.0005$, 0.001, and 0.005.

shown in figure 6.12. The experimental points are compared with the theoretical predictions for structural dissipation. When the dissipative term is included, the IP effective stiffness is written:

$$\tilde{k} = k(1 + i\phi) - M\,g\,/\,l \tag{6.8}$$

where ϕ is the loss factor, and the IP quality factor is written:

$$Q = \frac{1}{\phi}\frac{\omega_0^{\,2}}{\omega_0^{\,2} + g\,/\,l} \tag{6.9}$$

The data are compatible with a loss factor $\phi \sim 10^{-3}$. This number must be considered an upper limit for the flexible joint loss factor: in fact, the overall elasticity is contributed by the aluminium leg itself, which degrades the Q one would measure if the maraging steel joints were the only dissipative elements.

6.2.8 Measurement of the IP Transfer Function

The IP transfer function has been measured using a swept sine excitation: the shaker was driven with a sinusoidal signal whose frequency was increased step by step after each measurement cycle. The excitation amplitude was frequency dependent in order not to exceed safe IP oscillation amplitudes. The measured transfer function is shown in figure 6.13. The results can be summarised as follows:

(i) The main resonant frequency was tuned at 30 mHz, and showed a quality factor Q of 33±2.

(ii) Figure 6.13 shows five measurements performed with different counterweight loads or positions. The dip associated with the center of percussion effect moves when the counterweight mass and position are changed. The five curves correspond to progressive changes from the minimal counterweight configuration (without the movable counterweight) to the maximal one (movable counterweight 3 cm down, 5.7 kg per leg added). When the movable counterweight is taken away the IP transfer function would flatten out as per the curve (a) of figure 6.6. In fact, a dip at 5 Hz is present due to the interaction with the leg internal resonance.

(iii) An attenuation of 65 dB is achieved, with peak attenuation depending on the counterweight configuration.

(iv) The five measurements were performed for 10 days without noticeable changes of the 30 mHz resonant frequency, which is a proof of the system stability.

The resonances at about 9 and 15 Hz are associated with the main transversal beam modes of the leg-counterweight systems. These resonances limit the attenuation performance of the IP at high frequency. In principle, it would be useful to raise these resonant frequencies as much as possible, using larger diameter legs and/or a material stiffer than aluminium (like titanium or beryllium). From the point of view of the overall SA performance these resonances are not very important because they come in a

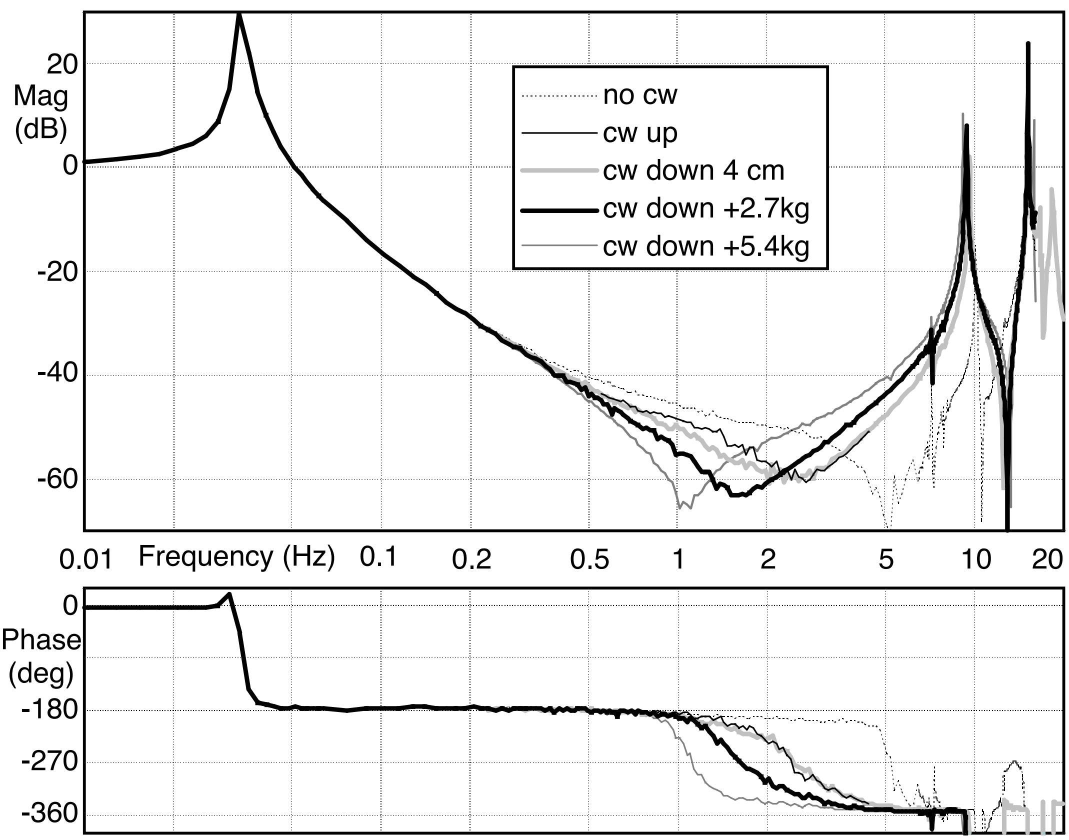

Figure 6.13 IP transfer function: five measurement are displayed with different counterweight configurations.

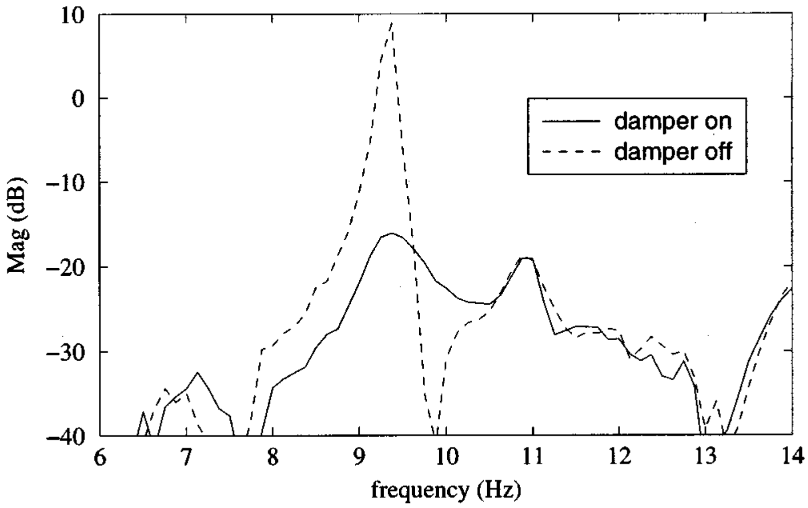

Figure 6.15 Effect of the damper on the legs resonance: a peak damping of 24 dB has been obtained.

frequency band where the attenuation properties of the standard filters have already kicked in.

As the IP is only required to filter in the region of the filters' resonances (0.2-2 Hz), the legs resonances were simply passively damped using the experience achieved with the blade springs of the standard filters [Braccini 1996]. The damper is sketched in figure 6.14. It is composed of a viton threaded rod 12 mm in diameter and 50 mm in length clamped to the leg corresponding to the resonance maximum excursion point. The viton rod carries a 0.3 kg block at a distance which can be adjusted to match the damper resonant frequency to the leg's. The effect of this device is shown in figure 6.15. A reduction of 24 dB has been achieved on the peak amplitude. Nevertheless, even without damping, the leg resonances do not affect the SA performance: the calculations show that the residual seismic motion at the leg resonant frequency will be largely dominated by thermal noise.

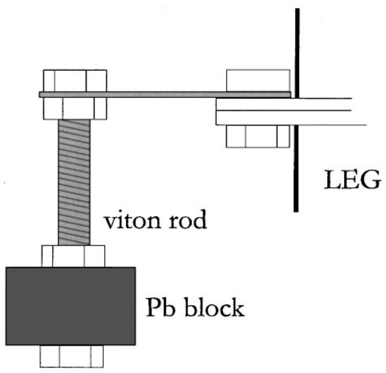

Figure 6.14 The damper for the IP legs.

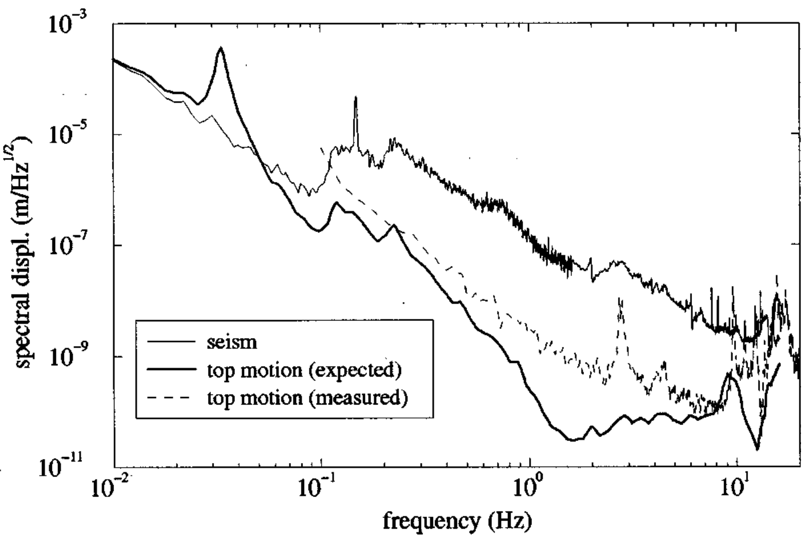

Figure 6.16 The measured seismic spectral density measured on floor is compared with the residual noise expected on the top of the IP according to the transfer function of Fig. 13 and with the measured residual noise (which is contributed by the acoustic noise).

6.2.9 RMS IP Residual Movement

It is possible to estimate the rms resonant motion of the IP at its fundamental resonance. The seismic activity at 30 mHz (measured on a Sunday to avoid human activity contributions) is $\tilde{x}_0(f_0) \approx 10^{-7}$ m/√Hz. The intensity of seismic noise on a working day is usually larger by a factor of 2-3. In figure 6.16 the seismic noise measured on the ground is compared with the spectral noise *expected* after the IP. The expected spectrum is calculated by multiplying the ground noise times the measured IP transfer function of figure 6.13. A spectrum measured directly on the IP top is also shown: the *measured* noise is larger than the *expected* noise due to the effect of acoustic noise. A useful estimate of the residual rms motion of the IP when operated under vacuum can be obtained only from the *expected* spectrum $\tilde{x}_{\exp}(\nu)$:

$$x_{rms}(f) = \sqrt{\int_f^{\infty} \tilde{x}_{\exp}^2(\nu) d\nu} \qquad (6.10)$$

The residual rms motion as a function of the starting integration frequency is plotted in figure 6.17. The expected rms at 5 mHz is about 70 mm, while the rms at 20 mHz is about 20 mm. This result corresponds to the completely passive IP performance and can be greatly improved with active damping and/or inertial seismic motion active suppression.

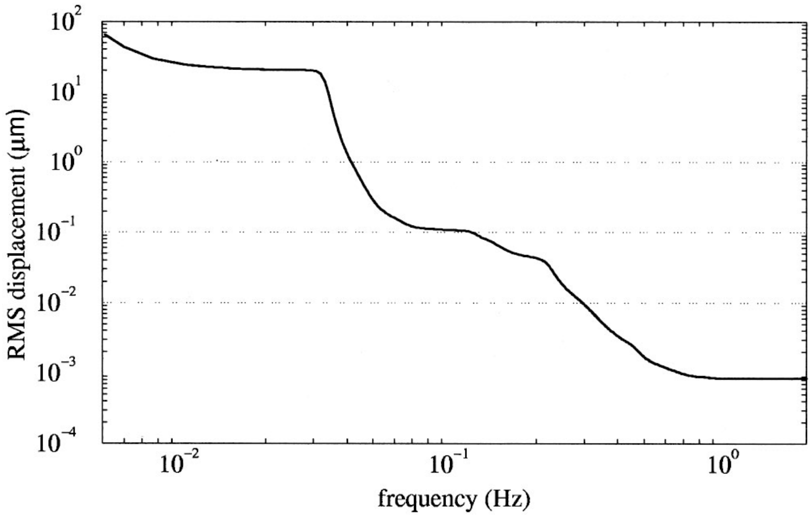

Figure 6.17 The rms motion of the IP top is calculated from the expected residual noise (see figure 6.16) as a function of the frequency using Eq. (6.10).

6.2.10 Further Developments

We have measured the inverted pendulum performance as a horizontal plane filter in the range from 30 mHz to a few Hz. Its performance as the horizontal component of the VIRGO preisolator stage provides a large reduction of the rms residual motion of the mirror, thus reducing required dynamic range of the locking actuators to acceptable levels. Ongoing tests are aimed at performing low frequency control of the IP position, to demonstrate its capabilities as a positioning stage for the superattenuator chain as well as the feasibility of active damping and inertial seismic noise suppression. Complete active suppression of the IP resonant peaks has been achieved and will be presented in a forthcoming article.

7. Double Pre-Isolation and Tilt Suppression

7.1 Preface

In order to adequately measure the performance of the large Scott-Russel pre-isolator (chapter 4), the author had to design and build a special test rig (called a cradle) capable of suspending the pre-isolator and its full payload for shake-testing. This was reasonably successful and achieved a very good level of isolation and text book performance at low frequencies (<2 Hz). However at slightly higher frequencies (>5 Hz) the isolation which should have stayed low was spoilt by what we believe to be acoustic coupling between the shaking cradle and the pre-isolator framework. As a result it becomes impossible to set the center of percussion correctly.

The author encountered the same problem when trying to measure the transfer function of Virgo's large inverse pendulum (chapter 6). In that case it was proved that the acoustic coupling was to blame by disconnecting all of the shaking framework from the base and suspending it with the crane so that the inverse pendulum was actually freestanding in still air on the shaking platform and this finally allowed sensitive detection of the center of percussion notch.

In addition to the acoustic coupling problems, there is also the problem of suitably suspending the pre-isolator stage so that it can be shaken without tilting.

With this experience in mind, and with the impending task of designing new pre-isolators for the 80 meter interferometer to be installed at the Aigo site in Gingin, the author was determined from the beginning to design in some effective and convenient way of being able to shake-test the pre-isolators. Ideally it should be possible to shake-test them in all three directions, without any tilting, in a vacuum, and preferably *in situ* - i.e. after final installation so that they don't need to be unmounted and moved again after adjustment and tuning. One can imagine that such a facility would have other uses also - such as to allow testing of the entire seismic isolation system and the various control systems with deliberately engineered disturbance signals. It also provides an ideal point at which to apply active inertial damping in the future

Aigo's vacuum tank design was already complete at this time and quite unsuitable for adding a structure on the outside to achieve the translational shaking. The only remaining option was to build a tilt-rigid servo frame to go inside the vacuum tank and from which the whole pre-isolator and isolation chain could be suspended. In order not to have to apply large forces to do the shaking, it would be preferable if the servo frame was designed as a fairly soft isolation stage itself.

At this point it becomes apparent that we are really considering cascading the pre-isolated suspension system with a *second*, albeit low-performance and power-actuated, pre-isolator. A *pre*-pre-isolator if you will. The question then arises, whether there is any benefit in trying to obtain good isolation performance out of it instead of just using it as an actuation frame. This is the main topic that the paper in this chapter deals with and the answer has to do with the level of seismic tilt noise present and how well it can be measured and actively servoed out. Sensitive measurement of tilt in the required frequency range is a new field and so some significant space is given to discussing the design requirements of a tilt sensor and control system that is feasible and would make double pre-isolation useful.

All of the theoretical calculations and predictions in the paper are the author's. All experimental work and measurements (on the tilt sensor and readout) were carried out by Zhou, the second author on the paper, who wrote a paper himself [Zhou 2000] detailing his measurements of the walk-off sensor. Apart from designing and building the shadow (or aperture) sensor and electronics for the initial tilt experiment, and apart from grinding the flat on the mirror for the walk-off sensor, this author has only been involved in these experiments with the tilt sensor and readout in a supervisory role. More recently another researcher (Y. Cheng) has been working with us and has built the 2-D version of the tilt sensor fashioned after the one shown in figure 7.11 and made measurements with it [Cheng 2001]. But his contribution occurred after this paper was published.

Regardless of isolation benefit or not, having the facility of a power actuated servo frame immediately in front of the suspension system is clearly highly desirable and the author searched to find the simplest and most reliable design with which to achieve this aim. Section 7.2 of this chapter gives a brief overview and some drawings of the author's design that was finally built and investigated.

The author was aware that three-legged and triangular structures for both the inverse pendulum and the LaCoste straight line stage were ideal, but having had bad experience manufacturing triangular based structures before, he chose to make it easy to manufacture even if over-constrained. In addition, the interlinking cubes leave a much larger central volume for the isolation chain, than would interlinking triangular prisms.

In order to obtain top performance from it, it will presumably be necessary to shake-test it also in order to adjust its center of percussion tuning. We have tried to do this but the system has very bad acoustic coupling from the large surface areas in close proximity and it wasn't possible to even get a change in phase from one extreme adjustment to the other. We have not yet tried in vacuum.

Some of these concepts have been presented at a few conferences and in conference papers [Winterflood 2000a] but no paper yet exists dedicated to this fully 3-D, tilt-rigid, ultra-low frequency pre-isolator.

7.2 Tilt-Rigid 3-D Servo-Frame Pre-Isolator

7.2.1 Introduction

Ultra-low frequency pre-isolation has been shown to greatly reduce the amount of residual motion affecting the test mass at frequencies around the normal mode resonances of the suspension chain. We have built a high-performance horizontal pre-isolator [Winterflood 1999] and intend to cascade it with a high-performance vertical pre-isolator [Winterflood 1998]. In order to achieve good isolation with a device of this type, it needs to be shake-tested and to have its centre of percussion finely adjusted - ideally in vacuum and preferably in situ where it will finally be operated long-term.

The primary requirement of such a shake-test facility, it that it should be capable of soft and smooth translational motion while remaining quite rigid to tilting moments. The requirement that it should operate in vacuum and in situ argues strongly against any lubricated, or sliding or rolling type of bearings for the motion - really only leaving flexing of elastic material to allow movement. This is the same requirement as any type of isolator or pre-isolator for gravitational wave detection and so we can consider any suitable suspension techniques that are known to the field.

7.2.2 AIGO 1 Metre Inverse Pendulum Pre-Isolator

Several methods of tilt-rigid horizontal suspension have appeared in the literature and consist of multiple matched horizontal suspension structures each providing a single isolated suspension point. For instance two X-pendulums give 1-D tilt-rigid isolation [Barton 1996] or three Scott-Russel suspensions can support a triangular platform in a tilt-rigid manner [Winterflood 1995] giving 2-D suspension (actually 3 degrees of freedom :- x, y, & θ_z) with the high performance and insensitivity to mass loading

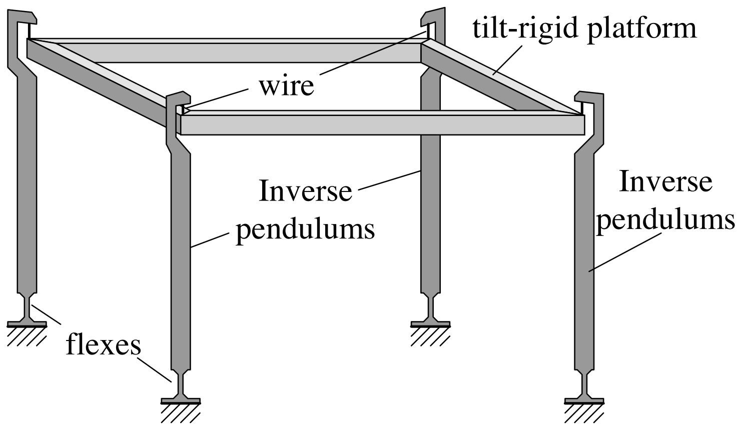

Figure 7.1 Schematic of an inverse pendulum style tilt-rigid horizontal pre-isolator.

provided by a geometry based design. However for the basic requirements of what is primarily a shake-test frame, the simplicity and strong tilt-rigidity of multiple inverse pendulums supporting a platform as a “wobbly table” [Losurdo 1999] illustrated in figure 7.1 is difficult to beat (even if it is only tuneable to low frequency with one particular mass loading).

The construction of the horizontal stage is illustrated in figure 7.2. The horizontally isolated frame is a cube shape approximately one meter on a side made of heavy gauge aluminium angle and supported at each corner with hollow aluminium inverse pendulums. The device was made with four sides and legs rather than the more ideal three sided structure for simplified manufacturing and increased central volume (see figure 7.7). The over-constrained nature of four legs as opposed to three does not prove to be a problem in practice (most tables have four legs rather than three!) as long as some adjustment is provided to ensure that each leg bears the same load. This adjustment can be made with low load so that under full load the inherent elasticity of the material takes up any final dimensional imperfections. Particular attention was given to making the frame as rigid as possible against cascaded loads with the large diagonal struts on each side. The design allows for +/-12 mm of motion in all horizontal directions before meeting a motion limit.

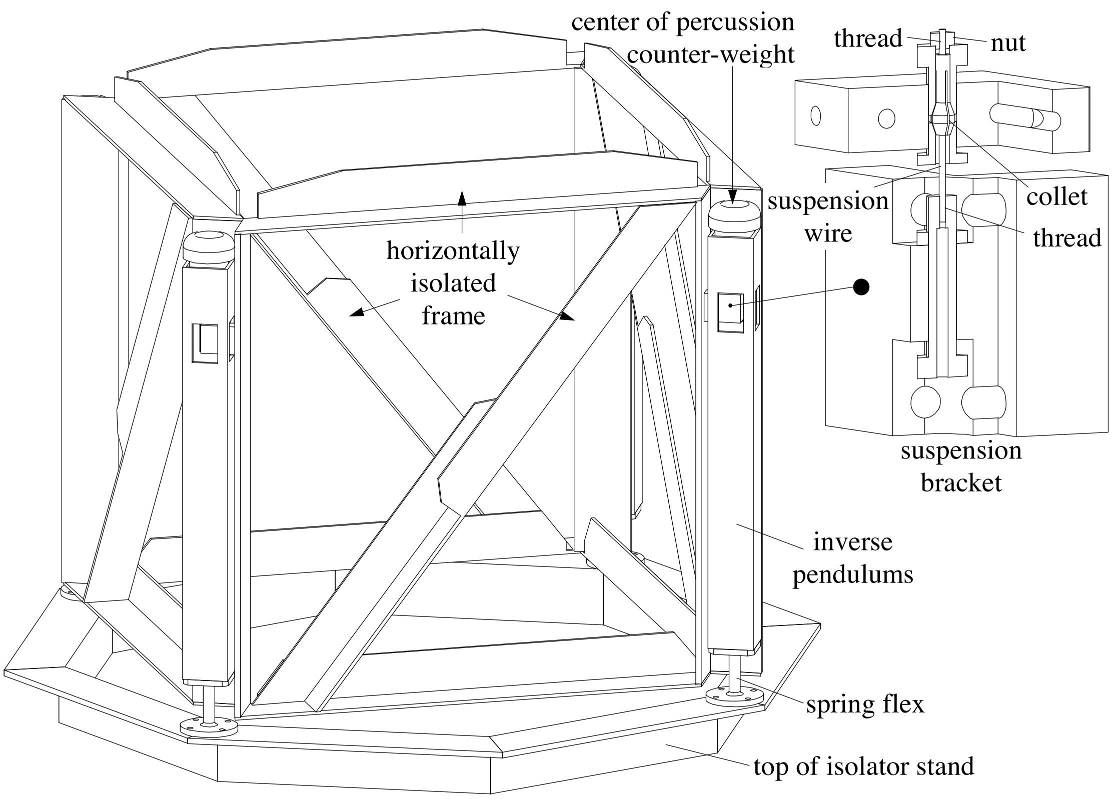

Figure 7.2 Horizontally isolated frame. Inset shows cutaway suspension bracket.

For the prototype we chose to build the wire suspended variety (see figure 1.12) and went to considerable trouble with the design to try to have well defined boundary conditions for the wire suspension - supporting its longitudinal tension by fine screw thread cut on the wire itself and a nut, but also gripping it after loading in a 4-jaw collet, as shown in the inset of figure 7.2.

Counterbalancing for center of percussion tuning is provided as the large stainless steel knobs in the tops of the legs (rather than in a ring at the bottom as Virgo has). This makes mounting at the lower end simpler and also makes it easier to change the spring flexures for an alternate loading. The flexures were turned from readily available aluminium machining rod so that they are easy to remake to a new size and replace.

The lower end of the flexures were not bolted directly onto a supporting surface as is suggested in the figure, but are seated on a special plate which has screws to allow fine adjustment of both the radial and the circumferential angle that the lower end of the flexure is clamped at. This is crucial to allow balancing and tuning :-

Initially the four plates are set accurately level with a precision spirit level. After mounting and suspending the frame so that all wires are carrying approximately the same load, the stage is loaded up with mass until its frequency starts to lower and any imbalance becomes evident. The angles of the plates may be adjusted in the following manner. To correct x and y offsets, opposite pairs of circumferential angles are adjusted in the same global direction by the same amount. To correct θ_z rotations, all four circumferential angles are adjusted by the same small amount differentially (i.e. in the same direction around the circumference). To tune the θ_z rotational resonant frequency down with respect to the translational resonances, all four *radial* angles are adjusted to tilt the legs more towards the centre, and vice-versa to tune θ_z up in frequency. (The x and y resonant frequencies cannot be adjusted with respect to each other except by sanding the round flexures slightly oval.) Care should be taken to always adjust all four angles or at least pairs of angles by the same amount whether differentially or collectively. This should avoid a build-up of large balanced stresses which have no external effect and so cannot be detected.

The four leg mounting plates have one additional screw adjustment, which allows each one to be translated towards or away from the clockwise adjacent leg. This allows the parallelism between pairs of adjacent legs to be adjusted. Non-parallel legs can be detected by a change in tilt angle with translation using a precision spirit level. These adjustments allow any sensed correlation between translation and tilt to be nulled out.

7.2.3 AIGO 1 Metre LaCoste Straight-Line Pre-Isolator

Tilt-rigid vertical suspension can be obtained from a linear guidance structure which rigidly constrains all degrees of freedom except vertical translation, and sprung to provide support against gravity with a low spring-rate. We chose the Wilmore suspension figure 7.3(b), turned it to work vertically and added springs in the LaCoste manner [LaCoste 1934] shown in figure 7.3(a) for low spring-rate. This is achieved by attaching diagonal coil springs between two sides of the shearing parallelograms formed by the pivot arms and frames to give the structure in figure 7.3(c). (Later the author found out that a three sided version of this structure is in the literature and called a LaCoste straight-line suspension [LaCoste 1983]!). The low internal mode frequency of coil springs (~20 Hz) is not a problem for this application as isolation operation is only required up to about 3 Hz. Figure 7.4 shows how this structure was implemented.

The horizontal pivot arms were made of three separate pieces - a pivot block at each end with a heavy gauge tube joining them (the bolts holding them together are not shown in the drawings). This was done both for convenience of manufacture, and to allow some degree of self-alignment and adjustment during assembly. Tightening the joints after assembly prevents dimensional imperfections from producing large stresses as a result of having four sides and eight pivot arms, instead of the kinematically ideal solution of three sides with only five pivot arms. Four of the eight arms were fitted with counter-weights sufficient in mass to provide center of percussion tuning for all eight.

The pivots at each end of the arms were made from a monolithic block of aluminium machined into a crossed flexure type of pivot as shown in the inset of figure 7.5. Monolithic structures such as this give lower losses (higher Q-factors) than the equivalent materials in a clamped situation and obtaining low loss is of prime

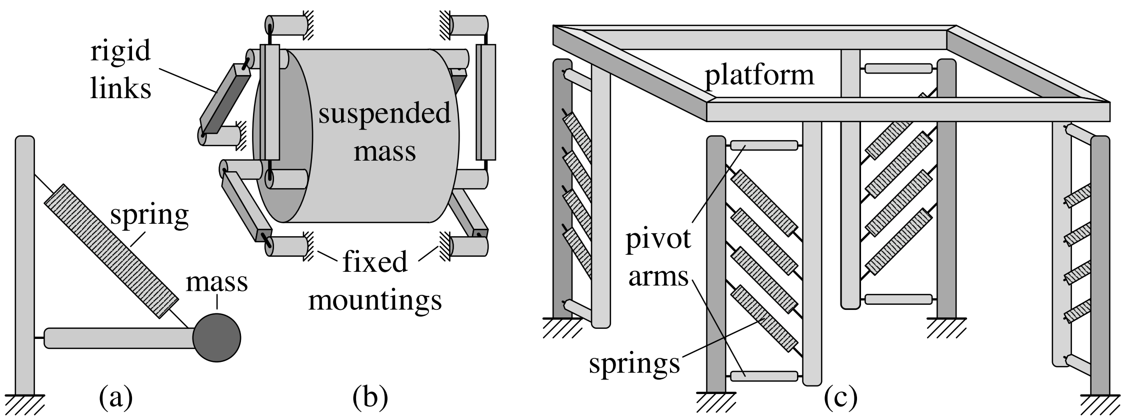

Figure 7.3 (a) Classical LaCoste Pendulum, (b) Wilmore horizontal suspension, (c) combined to make a tilt-rigid vertical stage.

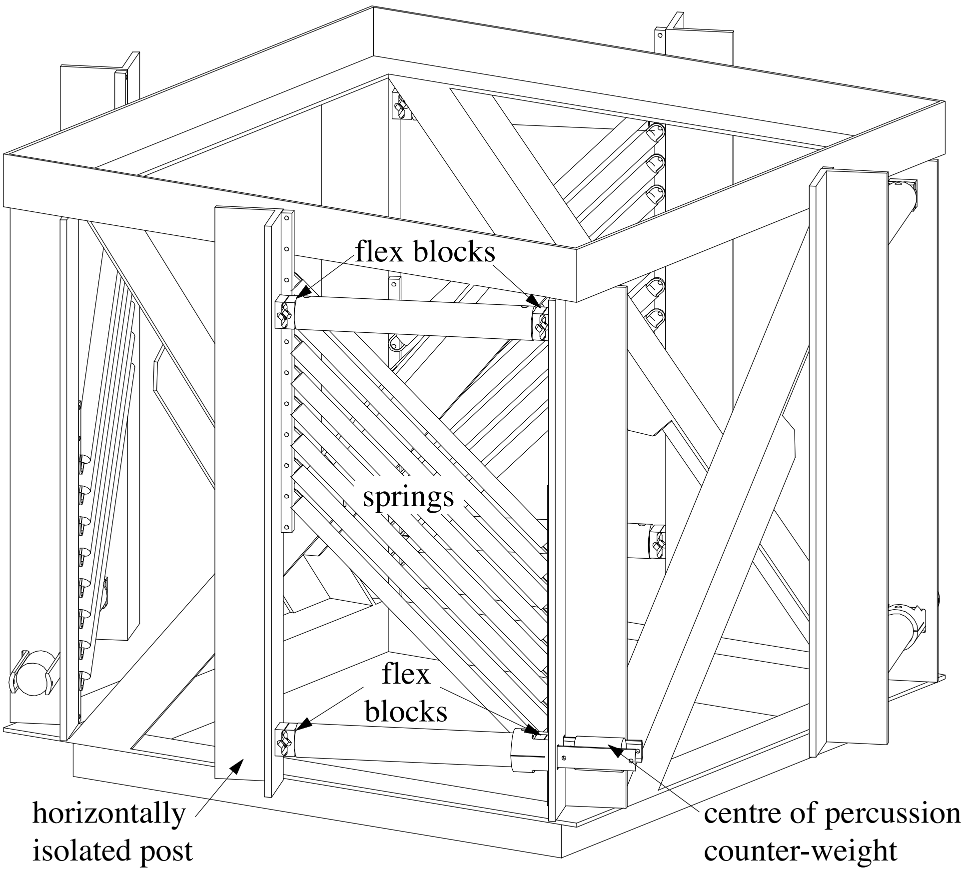

Figure 7.4 Vertically isolated frame (with outer layer of springs removed to show horizontal arms, flex blocks and counter-weighting).

importance for this type of energy storage intensive application which is aimed at achieving low resonant frequencies.

We used springs with a high degree of pre-tension to give a good approximation to the “zero-length” requirement of the LaCoste geometry[LaCoste 1934]. Each spring only provides about 10 Kg of supporting force - making it easy for even a child to add or remove springs to support different loads. Adjustments smaller than a whole spring are made with mass ballast. In the hope of maximising the final Q-factor, significant thought was given to finding a simple method of clamping the ends of the springs which was low loss and would provide a well defined end-of-spring flexing point. Unfortunately the sheer number of springs ruled out manufacturing any special flexible end terminations and an acceptable solution seemed to be simply to clamp both spring ends tightly between heavy gauge spring steel coned washers.

The separation between the pivot joints on the horizontal arms is ideally made the same as the separation between the flexing points on the rigidly clamped springs so that adding and subtracting springs does not alter the vertical spring-rate. (If this separation is kept fixed then varying the angle at which the springs are mounted - usually but not necessarily ~45 deg - also does not alter the spring-rate, but just varies the static vertical lifting force). The ~2 mm thickness of the spring wire would mean that the end flexing point with wire clamped is some centimetres away from the clamping point and depends strongly on tension. However if the first coil of the spring is placed very near to where the wire is clamped, almost all of the flexing occurs in this first coil which can be at a fairly well defined position and is certainly constant enough for empirical determination and use. (The effective flexing point can be found empirically by first nulling out the flexure-block spring-rates and then seeing how the resonant frequency changes as more load and springs are applied. If the flexing points are in line with the pivots, then no change in resonant frequency should occur. If the horizontal separation between spring flexing points is less than the separation between flexure pivot points then the resonant frequency will increase and vice-versa).

To null out the significant spring-rate of the flexure pivots in the horizontal arms, several additional springs on each side shown in figure 7.5 are stretched horizontally to a much greater width than the separation between the pivoting points. This inverse pendulum effect produces a negative spring-rate and making the width adjustable allows the overall vertical spring-rate to be tuned to obtain very low resonant frequencies.

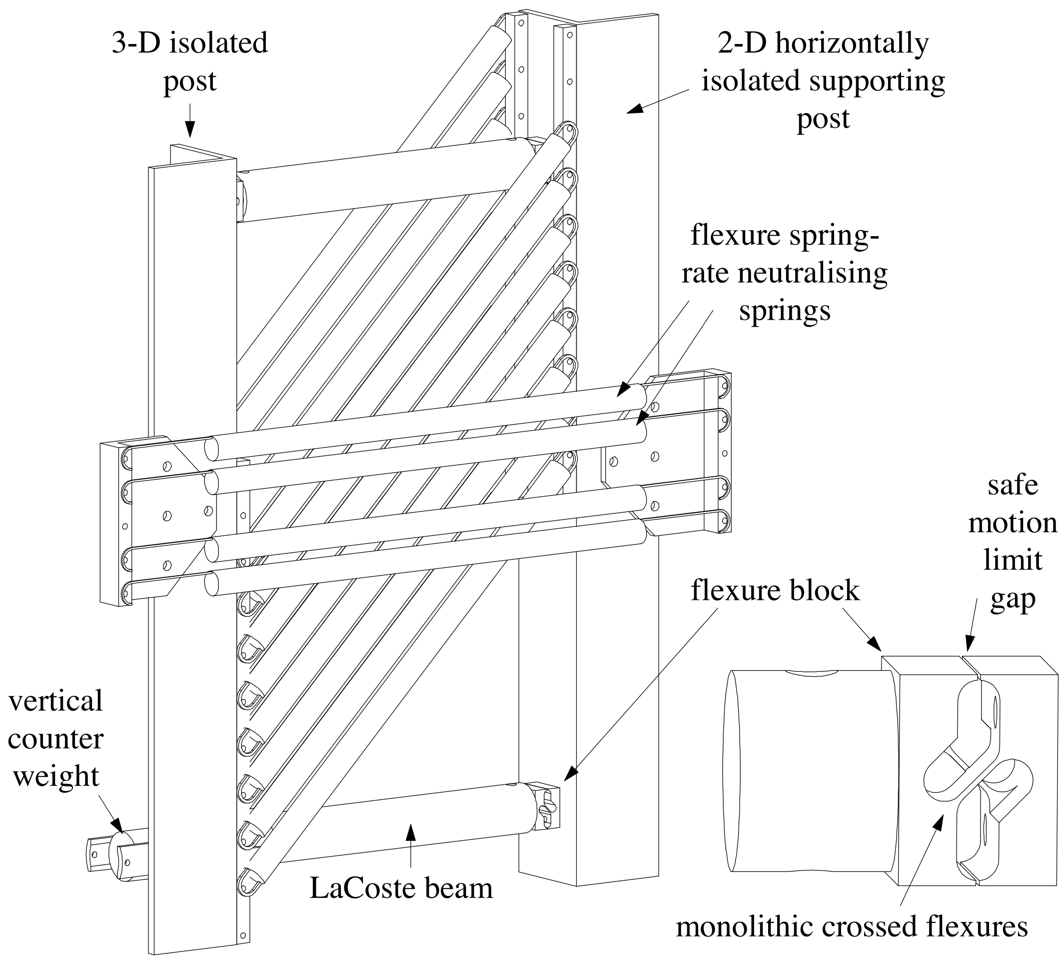

Figure 7.5 Detail of LaCoste sprung parallelogram with neutralising springs and flexure block detail.

7.2.4 *The Cascaded Structure*

The cascaded structure is shown in figure 7.6. The two cubic frames interweave with each other and a little forethought was required to choose which joints could be welded and which should be bolted to allow assembly after manufacture!

Figure 7.7 give a view looking down from the top through the centre of the pre-isolator to show the large clear central volume that is available for cascading the intended second pre-isolation stage and/or the normal isolator chain.

Table 7.1 summarises the characteristics and the performance that we measured with this device.

7.2.5 *Conclusion*

The device met its design objectives and performed better than what had been assumed for such a device in the following paper. This prototype was made with spring flexures at the lower ends of the legs and wire links at the top as shown in the figures. This provides soft suspension and isolation in the θ_z rotation as well as horizontal translation. The author believes this is not the best choice for this application and future "production" models will have identical spring flexures at both ends of each leg. This

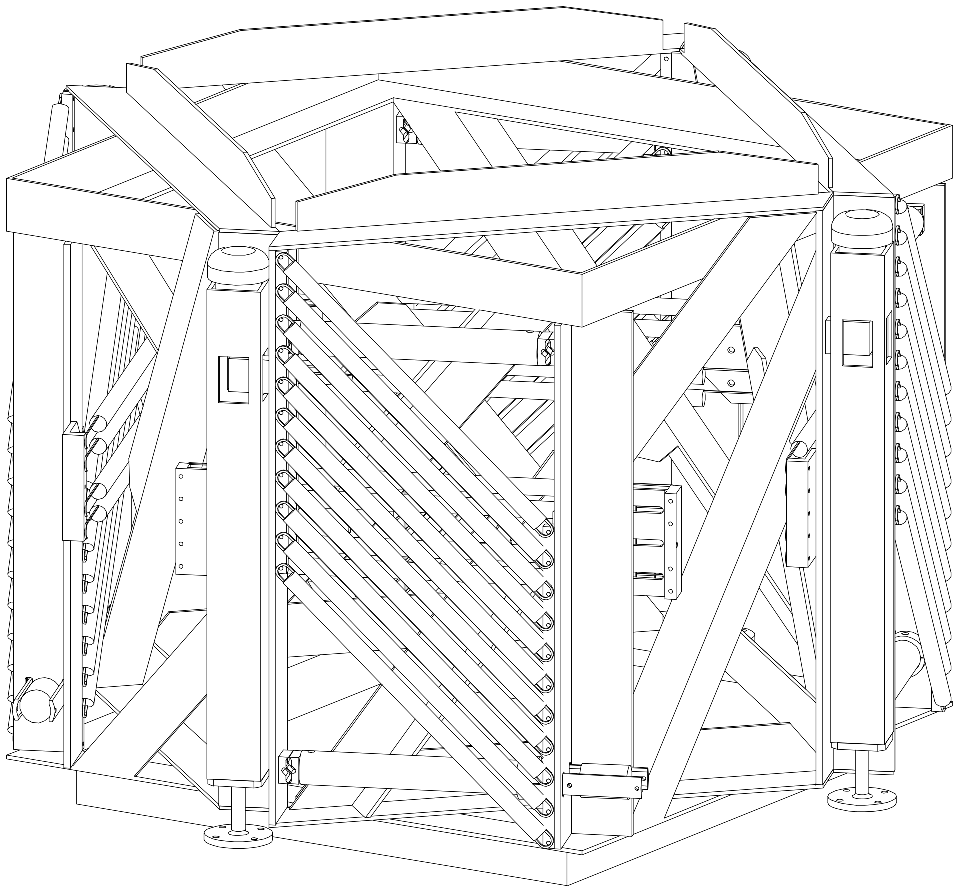

Figure 7.6 Complete 3-D tilt-rigid pre-isolator (horizontal and vertical stages intertwined).

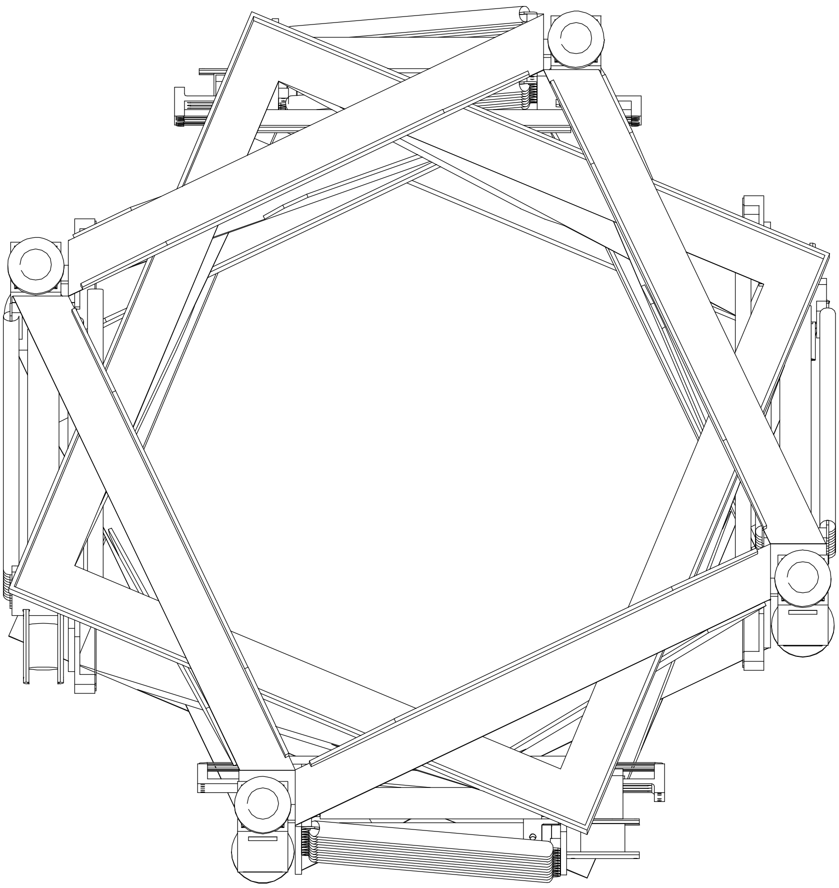

Figure 7.7 Top view of 3-D tilt-rigid pre-isolator showing central free useable area.

makes for much simpler and more robust system and also means any weight positioning on the platform is much less critical (see section 1.5.2.3)

General	
Tank diameter requirement	1.5 m
Central useable area diameter	0.85 m
Maximum design loading	1000 kg
Target isolation frequency range	0.2 - 3 Hz
Vertical Measurements	
Max vertical resonant period	25 sec
Diagonal spring internal mode	17 Hz
Horizontal Measurements	
Max horizontal resonant period	40 sec
Leg bending internal mode	80 Hz

Table 7.1 Parameters and measurements of 3-D tilt-rigid pre-isolator.

7.3 Tilt Suppression for Ultra-Low Residual Motion Paper

Tilt suppression for ultra low residual motion vibration isolation in gravitational wave detection.

J. Winterflood, Z. B. Zhou[†], L. Ju, D. G. Blair

Department of Physics, University of Western Australia, Perth 6907, Australia.

[†]*Department of Physics, Huazhong University of Science and Technology, Wuhan, 430074, P.R. China.*

It is shown that a combination of ultra low frequency vibration isolators combined with high sensitivity tilt sensing and feedback allows RMS vibration levels in laser interferometer gravitational wave detector test masses to be reduced from ~100 μm to ~1 nm greatly simplifying operation and reducing control system noise injection.

7.3.1 Introduction

Gravitational wave (GW) detection offers a new spectrum with which to explore the universe. The technology of GW detection requires the reduction of all sources of vibration towards or below limits set by quantum measurement theory. Isolation from seismic noise is critical – not only in the GW detection band, but also at frequencies below this because even low frequency disturbances make locking of high finesse cavities very difficult and call for large servo control forces to maintain lock which can themselves inject noise (from the limited dynamic range of electronics and through seismic coupling of imperfect actuators). It is therefore now widely recognised that a reduction of residual motion is of critical importance to the achievement of the high sensitivity goals of advanced GW detectors.

For the simplest control scheme, the dynamic range of the electronic control system (~10^{10}) should be matched to the dynamic range of the forces to be applied to fringe lock the test mass, where the largest force is the force required to maintain lock in the presence of low-frequency (~0.5 Hz) residual motion, and the smallest is the force noise that produces motion not exceeding the instrument noise level (~10^{-22} m/√Hz - advanced LIGO) at the low end of the detection band (~20 Hz). If an applied force produces 10^{-22} m of motion at 20 Hz, then 10^{10} times this force can be expected to produce only ~10^{-9} m of motion at 0.5 Hz. Thus the residual motion at the test mass needs to be reduced to this level if the locking forces are to be applied directly to the test mass at this sensitivity.

The use of mechanical or electrical filters to reduce the control system noise reaching the test mass allows this requirement to be eased somewhat. For instance, if the test mass is suspended as a 1 Hz pendulum from a previous stage and the locking forces applied to this previous stage, then the noise at 20 Hz reaching the test mass will be reduced by a factor of 20^2 bringing the residual motion requirement to a more attainable value of $\sim 5\times10^{-7}$ m. However adding filter stages between the fringe locking drive and the test mass (within the control loop), can make control considerably more difficult or even impossible. These arguments suggest a target residual motion of order 1 nm. Arguments presented later in this paper based on assumed seismic levels indicate that residual motion less than 1 μm RMS above 0.2 Hz cannot be achieved without tilt control.

In following sections, we point out various problems to be overcome in order to reduce the residual motion to the nanometre level. It becomes apparent that a key element in this progression is an adequate high frequency (0.1 - 5 Hz) seismic tilt sensor. This is an area that has remained largely unexplored in seismology but has recently begun to be addressed by the GW community. The design requirements of this tilt sensor are the main topic of this paper and we present some preliminary experimental measurements.

7.3.2 Seismic Isolation

The size of the control forces required to fringe lock the test mass is determined by the residual horizontal low frequency motion which feeds through an isolation stack below its intended cutoff frequency (a few hertz). To enable us to express our results in linear displacement units, we shall assume a standard seismic noise spectrum (the dotted curve in figure 7.8) of $10^{-6}/f^2$ m/√Hz above 1 Hz levelling out below the microseismic peak at about 0.2 Hz. This spectrum is typical of an urban environment but the isolated sites at which GW detectors are located may be two or even three orders of magnitude quieter around 1 Hz (although much the same at microseismic frequencies). A quieter site reduces all the absolute levels (i.e. the vertical scale on the graphs) by the seismic factor but leaves the relative relationships unchanged.

Figure 7.8 indicates the residual spectrum and RMS motion of a typical test mass as a result of applying various isolation structures. The curves on the left show the spectral densities in m/√Hz, while the ones on the right have been integrated to give the RMS motion in metres for all frequencies above the value shown on the *x*-axis. All of these curves consist of the isolator transfer function multiplied by the standard seismic spectrum. (The isolator transfer functions were calculated by constructing a transfer matrix for each component such as a stiff pivot, pendulum link, mass etc and multiplying them together to obtain the complete system response. Where necessary matrices with 4 variables were used: force, displacement, moment, angle. Sub-systems constructed with 4 variables were converted to 2 variables – typically force and displacement – at connections with simplifying assumptions such as perfectly flexible pendulum joints.)

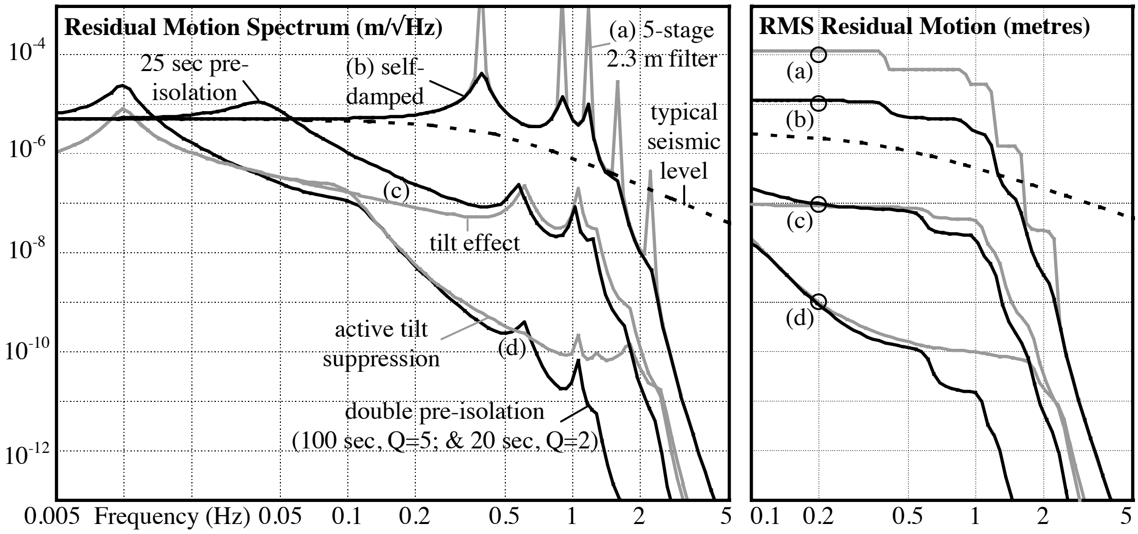

Figure 7.8 Successive reduction of residual motion - (a) high Q-factor internal modes resonantly amplify the seismic motion at the normal mode frequencies, (b) these modes can be passively damped for a factor of 10 reduction in residual motion, (c) ULF pre-isolation can provide an additional factor of 100 reduction, (d) dual ULF pre-isolation and active tilt suppression can gain another factor of 100.

Excellent vibration isolation in the tens of hertz to kilohertz range has been demonstrated using all-metal isolator stacks [Blair 1993, Beccaria 1997] consisting of cascaded pendulum stages for horizontal isolation, and incorporating springs for vertical isolation. With reasonable engineering the horizontal component remains dominant. Curve (a) shows the response of a typical 5 stage, 2.3 m high, all-metal isolation stack (Consisting of 4 cascaded pendulums 0.5 m and 100 kg each, Q-factors ~1000, finishing with 20 kg test mass on 0.25 m, Q ~10^6 pendulum). It cuts off steeply above a few hertz, but it amplifies the seismic noise at low frequency due to the high Q-factor normal mode resonances. The residual motion spectrum is dominated by these internal resonant modes and the RMS residual motion is almost two orders of magnitude above the seismic level.

This amplified motion may be reduced by damping the Q-factors of these modes. Direct active damping using sensors and actuators mounted to a seismically affected frame is limited in performance due to seismic noise pickup and injection. However using suspended reaction masses, or the new passive technique we have called "self-damping" [Winterflood 2000a] can allow normal mode Q-factors to be reduced below

10. With such damping the residual motion is reduced to the level of curve (b) – an order of magnitude improvement, but still above seismic and far above the nanometre level.

7.3.3 Ultra-Low Frequency Pre-Isolation

A very effective method to further reduce residual motion is to add an ultra-low frequency (ULF) pre-isolator stage in front of the isolator stack. Such a device can reduce the seismic drive to the normal mode resonances and thus the residual motion, by two orders of magnitude. It also provides a structure before the isolation stack to which very low frequency drift correction forces may be applied (suspension point servoing) - avoiding the noise injection problems associated with applying large DC forces near to the test mass. For this reason only the residual motion above approx 0.2 Hz is considered in this discussion as motion below this frequency can be easily counteracted above the stack at the pre-isolator level. (This scheme also invalidates any attempt to damp the stack resonances by frame mounted sensing because the frame has two orders higher motion than the resonances to be damped).

With this incentive, various new ULF pre-isolators (1-D, 2-D and vertical) have been introduced and tested, which have passive isolation performance equivalent to springs or pendulums up to 6 km in length [Winterflood 1999]. These include the inverse pendulum [Losurdo 1999], the X-pendulum [Barton 1996], the folded pendulum [Liu 1995], the Scott-Russel linkage [Winterflood 1999, Winterflood 1996], and the Torsion Crank suspension [Winterflood 1998]. Adding a pre-isolation stage (of 150 kg, 25 sec period and Q-factor ~2) to the previously described isolation stack gives the black curve (c) when only horizontal translational motion (i.e. no tilting) is considered. This reduces the residual RMS motion (integrated above 0.2 Hz) down to the 1 micron level. It will be seen that this is about the best that can be achieved without resorting to an active tilt sensing and control scheme.

There are a couple of reasons for the distinction between a *pre*-isolator and normal isolation stage. The main one being that structures that are designed to have very low resonant frequencies tend to be larger and more massive and consequently have lower internal resonances (which bypass their isolation) than normal isolator stack stages. As a result they generally do not provide useful isolation in the 10 Hz - 1 Khz detection band but are included mainly to reduce residual motion below the detection band. In the horizontal regime there is an additional difference between the two isolator types.

Normal isolator stages are generally pendulums of an appropriate height (tens of cm's) which sets their resonant frequency $(g/h)^{1/2}$ and are attached to a single (ideally) perfect pivot on the stage above. However in a pre-isolator stage the connection to the previous stage is multi-point allowing it to transfer a tilt moment thereby obtaining a resonant frequency much lower than that expected for a pendulum of comparable height (e.g. 150 sec [Winterflood 1999]).

7.3.4 Effect of Tilt on Pre-Isolation

Horizontal ULF pre-isolation structures work by providing a rigid constraint against the very large gravitational force of the total suspended mass, while also providing extremely soft horizontal compliance. This requires the structures to be maintained in accurate alignment with the gravitational field - i.e. rigidly fixed against tilting. However seismic tilts perturb the alignment and produce horizontal accelerations directly.

For a suspension system without horizontal pre-isolation the effect of seismic tilt is far smaller than seismic translation and can almost always be neglected. However when a ULF horizontal pre-isolator is employed, tilt induces significant translational horizontal motion. As the period of the pre-isolator is increased, the residual *translational* motion is reduced but the tilt coupling remains at the base level for an object sliding on a flat friction-less surface. Once the translational motion is smaller than this base level, there is no improvement to be gained by further horizontal isolation.

The tilt coupled translation for the above ULF pre-isolated system is shown by the grey curve (c). It is of the same approximate magnitude as the horizontal translational signal (black curve) in the 0.5 Hz - 1 Hz region. If the pre-isolation is improved (e.g. by increasing its resonant period or adding a second stage), then the horizontal signal (black curve) reduces but the tilt coupling (grey curve) remains at the same level and becomes the dominant source of residual motion. (The total residual motion is the RMS sum of these two curves)

In order to further improve isolation performance and reduce residual motion, it is necessary to counteract the tilt effect. The only effective means of achieving this is by active tilt control. A tilt sensor to sense tilt in two degrees of freedom (θ_x and θ_y) is required, and a servo system and actuators to counteract the measured tilt. The control loop is required to provide gain over a range starting near the pre-isolator resonant frequency (~50 mHz), peaking at ~1 Hz and falling off as quickly as possible above that.

We discuss the design requirements of this tilt sensor and associated control system in the following sections.

With tilt stabilisation present, it becomes advantageous to improve the ULF pre-isolation performance again. The best approach here seems to be to add a second stage of ULF pre-isolation (rather than try for impossibly low resonant frequencies). With tilt stabilisation in place and with a second ULF pre-isolation stage included, it should be possible to achieve the residual motion indicated by the lowest curves (d) in figure 7.8. In order to cascade a second ULF pre-isolator from a first, the first must be made rigid to tilt. Our cascaded ULF pre-isolation design will be the topic of a future letter.

7.3.5 Tilt Suppression Requirement

A tilt suppression requirement could be derived from a residual motion or servo force specification at the test mass. However a more practical approach is to require sufficient tilt suppression to allow full benefit to be obtained from a second ULF pre-isolation stage. This specification has the advantage that the level required is independent of the actual seismic level at a particular site. We will use the properties of our existing pre-isolator in deriving this specification. That is, the tilt contribution is required to be at the level of the translational horizontal coupling of a dual pre-isolator.

The translational horizontal motion also contains a component which is mechanically cross-coupled from vertical motion - making it also dependent on vertical pre-isolation performance. This cross-coupling is usually assumed to be at the 1% level set by fabrication accuracy of mechanical structures, but provided it can be measured and tuned, it can be made as low as 10^{-4} - 10^{-5} at a particular frequency. This has been demonstrated in centre of percussion tuning of high performance isolators achieving 100 dB isolation [Winterflood 1999, Liu 1995].

We have recently built and tested a 3-D ULF pre-isolator which is rigid to tilt [Winterflood 2000a] so that a second pre-isolator can be cascaded after it. It is constructed as cascaded horizontal and vertical sections. For the horizontal isolation we used a platform on 1 m inverse pendulum legs. This tilt-rigid structure supports a LaCoste straight-line suspension [LaCoste 1983] for vertical isolation. We obtained vertical resonant periods in excess of 25 sec with Q-factor>1, and horizontal resonant periods of ~40 sec. For a second stage, we assume performance based on our previous work - a full sized Scott-Russel pendulum, in which we obtained a period of ~150 sec.

For the purposes of this analysis we will aim at conservative vertical and horizontal periods of 15 sec and 20 sec respectively for a first tilt-rigid pre-isolation stage, and 20 sec and 100 sec for a second higher performance stage. These parameters imply that at the top of the stack the horizontal seismic noise will be about 2% of the vertical. This ratio is sufficient to be able to ignore vertical seismic noise in the following discussion. Hence we assume two cascaded horizontal pre-isolators with the following parameters: p_1=20 sec, m_1=150 Kg, p_2=100 sec, m_2=500 Kg. Cascading these give the mode periods apparent in figure 7.8(d) of 100 sec and 8 sec. This pair produces a seismic attenuation in the 0.5 Hz frequency region of $\sim 2\times10^{-5}$. This sets the tilt suppression requirement in this frequency region to this value or below.

7.3.6 Seismic Tilt Expected

The seismic tilt spectrum in the frequency range of interest is rarely measured and to our knowledge has not been correlated with the vertical and horizontal spectrum amplitudes. However making the reasonable assumption that the seismic motion travels as waves in the surface of the earth, suggests that the tilt spectrum amplitude will be approx $2\pi/\lambda$ times the vertical spectrum amplitude where λ is the wavelength of the waves. The phase velocity, which determines λ varies greatly with frequency and terrain composition (which is rarely homogeneous). For instance at microseismic peak frequencies (50 - 200 mHz) in eastern Montana this velocity was measured at ~3 km/s [Haubrich 1969], but for shorter wavelengths where the bulk of the strain exists in a much softer surface layer such as sand we expect the velocity to be much less. A velocity given for shear waves in sand below the water table [Press 1966] is 500 m/s and we have taken this value as a likely and conservative figure for the sand-plains of Western Australia in the frequency region of most concern which is around 1 Hz. (The velocity above the water table is given as 400 m/s and Raleigh waves travel somewhat slower than pure shear waves).

This figure finds some support in a measurement made of the ground tilt seismic spectrum [Luiten 1997] at the VIRGO laboratory near Pisa. Unfortunately the vertical and tilt seismic spectra were not measured simultaneously, but assuming the usual $10^{-6}/f^2$ m/√Hz above 3 Hz, this measurement suggested a phase velocity of 700-800 m/s. We have assumed that the horizontal and vertical seismic noise have similar spectra (with the level shown by the dotted curve in figure 7.8), and are related to tilt using a constant seismic wave velocity of 500 m/s. This relationship directly produces the expected seismic tilt spectrum shown by the dotted curve in figure 7.10. The expected higher velocities at long wavelengths will make the low frequency end of the spectrum fall off much quicker and higher velocities in general will reduce the tilt sensor sensitivity and loop gain requirement (by reducing the area between grey curves (c) and (d) in figure 7.8).

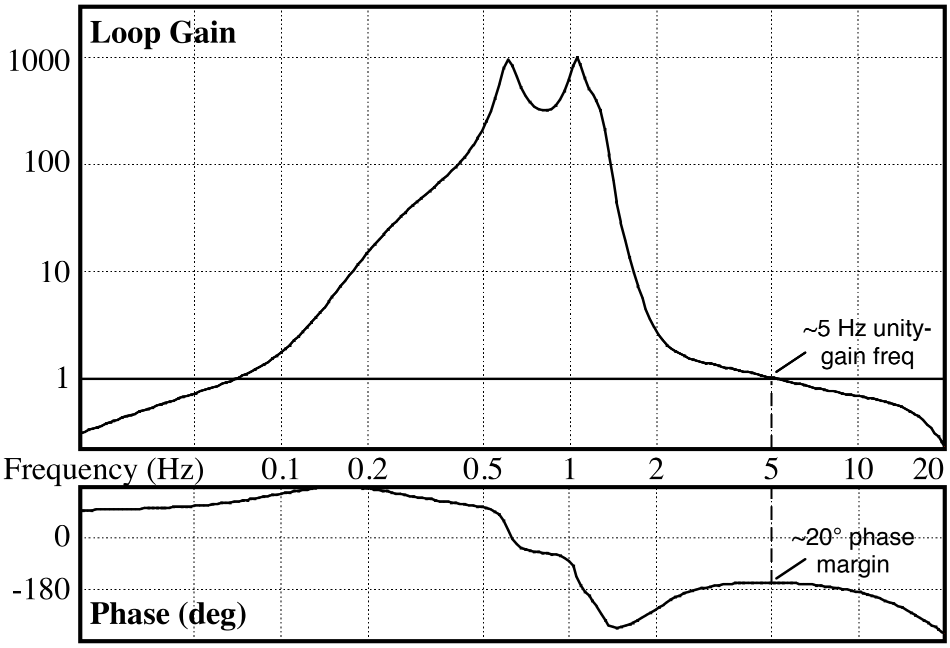

Figure 7.9 Loop gain requirement for the tilt suppression in figure 7.8(d).

7.3.7 Tilt Cancellation Servo Control Loop

With these assumptions and tilt suppression requirement, a control system may be designed to counteract tilt (e.g. by means of piezo stack actuators) somewhere prior to the second pre-isolation stage. The tilt must be sensed after the actuation (through as rigid a connection as is possible) and a high loop gain provided to null the tilt as well as possible. The loop gain may be tailored to only provide high suppression (suppression ≈ 1 / loop gain) at the necessary frequencies and a suitable loop gain is shown in figure 7.9. It uses 6 complex poles and 2 complex zeros and cuts off sharply above 20 Hz. It was manually tailored to minimise the residual motion under free-mass condition (prior to fringe lock), and also to minimise the high frequency performance required of the servo system. This loop gain was used to obtain the residual motion curves shown in figure 7.8(d).

7.3.8 Tilt Sensitivity and Noise Requirements

Clearly the sensitivity (noise floor) required from a tilt sensor applied in this system must be such that tilt readout noise does not contribute, through the servo system, significant additional translation noise to the test mass (over and above residual seismic feedthrough and final pendulum thermal motion). This sensitivity requirement is the black curve shown in figure 7.10. At high frequencies (>5 Hz) the thermal noise in the final pendulum dominates. At medium frequencies the seismic tilt and translation feeding through the isolation dominates. At low frequencies (<0.2 Hz) a drift control system which is not considered here will be used to centre the entire stack using a low frequency position error signal and applying its control through the tilt actuators. Its effect is not shown in any of the figures.

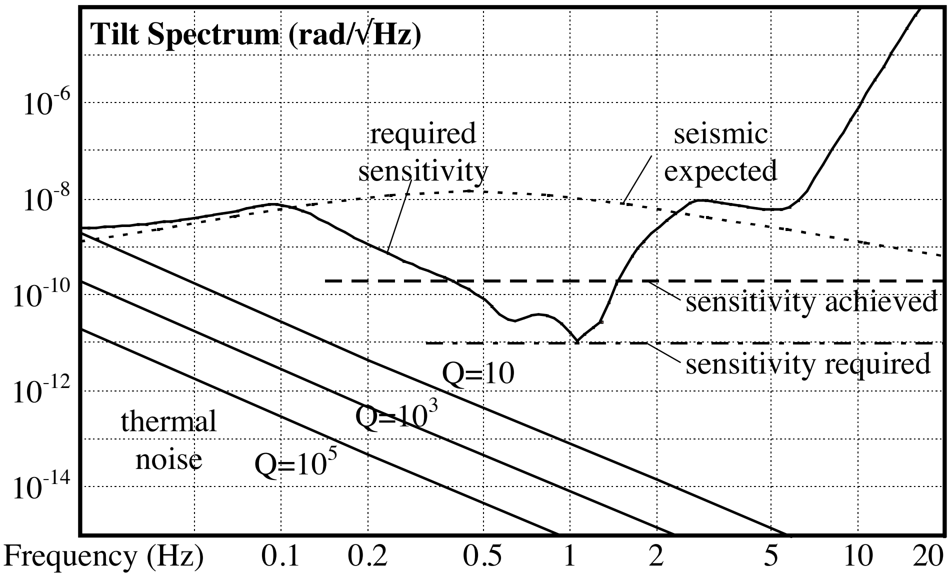

Figure 7.10 (a) tilt sensor sensitivity requirement. (b) thermal noise floor. (c) shot noise floor achieved.

In figure 7.10 the required sensitivity is compared with the Brownian motion thermal noise predicted from a 15 kg mass with 8 cm radius of gyration pivoted with a resonant period of 100 sec and various Q-factors. It is evident that for the frequency range of interest, thermal noise should not be a problem.

Also shown in figure 7.10 is the readout noise level that we have achieved with a relatively simple shadow sensor (actually LED and 3 mm aperture with balanced photo-diodes). This is clearly inadequate but we expect that the “walk-off” angle sensor described in following paragraphs will achieve the 10^{-11} rad/√Hz required.

Other sources of tilt sensing noise are cross-coupled translational motion appearing as a tilt signal, and temperature fluctuation gradients. Using a conductive metal test mass in a vacuum with an isothermal shield, we do not expect temperature fluctuation noise to contribute to the noise floor. Cross-coupled noise is covered in a following section.

7.3.9 Recent Angular Accelerometer / Tilt-meter Research

Seismic tilt in this high frequency range of 0.1 to 5 Hz has seen very little investigation until quite recently. In 1989 Speake and Newell [Speake 1990] reported on a high frequency (0.1-10 Hz) tiltmeter and a measurement of the seismic tilt in their laboratory. Their tiltmeter employed a dumbbell shaped test mass of approx 4 kg and length ~20 cm, suspended very lightly and as close as possible to its centre of mass using a commercial (Bendix) flex-pivot. The resonant period used was ~10 sec, Q-factor ~100, and the relative position between the frame and test mass was measured with a balanced shadow sensor arrangement. A damping torque was applied by differentiating this position signal and applying a small force using magnet and coil actuation. Low frequency stability was also obtained by integrating this signal and applying it to drive the dumbbell towards the electrical zero for times longer than ~50 sec. The seismic tilt measurement showed a level of the order of 1 nrad/√Hz and was approximately flat with frequency. However this level was only a factor of 2 above the noise floor of the instrument.

In 1992 Usher and Mat-Isa [Usher 1992] reported on a similar but smaller device using a test mass of approx 0.3 kg and length ~10 cm, again suspended with flex-pivots. The free resonant period obtained was approx 1.4 sec but they applied force feedback to obtain an angular acceleration signal and the closed loop resonant frequency was about 10 Hz (with a loop gain of ~100). They used a differential capacitive transducer for sensing the offset (error signal) but obtained a very poor signal to noise ratio of approximately unity. This was apparently due to the thermal noise (Brownian motion) of their test mass - exacerbated by the low frequency of interest which was a narrow band centred on 0.05 Hz. A factor of 4 improvement would be expected at 0.1 Hz.

In 1997 the VIRGO group [Luiten 1997] reported on a similar device and used it to obtain a measurement of the tilt acceleration in the laboratory above 1 Hz. They used a test mass of approx 2 kg and ~24 cm long suspended initially with beryllium copper crossed blade springs. Their resonant frequency could be tuned between 0.3 and 3 Hz and also used the force feedback technique to obtain angular acceleration. They used a differential inductor transducer (LVDT) for sensing the offset and the closed loop resonant frequency was of the order of 10 Hz. This allowed measuring the seismic tilt acceleration spectrum up to 10 Hz and above that the tilt displacement spectrum was obtained with the feedback loop open. The signal to noise ratio obtained was more than an order of magnitude above 2 Hz, but there was unexplained high noise below 1 Hz.

7.3.10 New Tilt Sensor Design

The tilt sensor we propose in figure 7.11 incorporates several novel features which allow wide band sensitivity extending to DC, low sensitivity to translational motion in the band of interest, and the ability to sense both θ_x and θ_y tilts with a single test mass (allowing simple central mounting of a single sensor unit).

Precision electric discharge machining (EDM) allows very thin and wide (<50μm thick × cm's wide) flexures to be made from a monolithic block of low loss metallic material. With an arrangement of cuts similar to the inset in figure 7.11 it is possible to make *x-y* flexures with intersecting axes from a monolithic block. Suspending a disk shaped test mass very near to its centre of mass with such a flexure allows sensing both θ_x and θ_y tilts. Utilising two of these *x-y* flexures with a pendulum link between provides continued sensitivity down to DC (thanks to the hang of the pendulum) and can be made surprisingly insensitive to translational motion. Sensitivity to DC is not a necessity for this application but it provides a means to measure earth tides and other low frequency tilts for correlation in noise studies or for calibration.

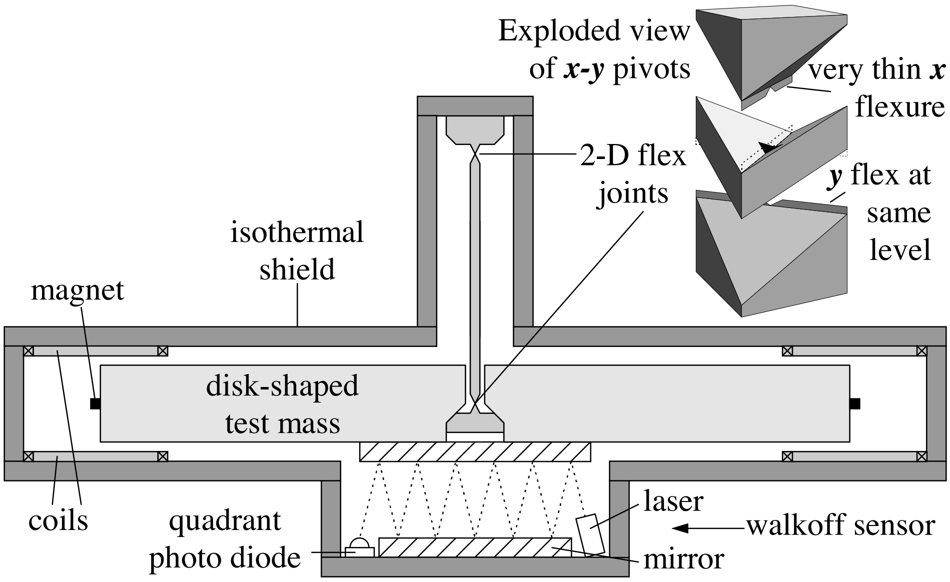

Figure 7.11 Dual axis tilt sensor.

7.3.11 Translational Motion Rejection

Knowing both the translational seismic motion expected and the tilt sensitivity required, allows us to plot the level of translational motion rejection required of the tilt-sensor (sensitivity required / motion expected). This is shown in figure 7.12. The lower broken curve is the rejection required if the tilt-sensor is mounted before the first stage of pre-isolation and the higher curve is for mounting after it. The latter mounting has a much more relaxed rejection requirement because its motion noise is reduced as a result of the pre-isolation stage.

It follows from the equivalence principle, that if a device is made sensitive to the gravitational field in order to provide a signal down to DC, then it must also be sensitive to horizontal acceleration. For this reason other workers have not attempted to obtain DC tilt sensitivity in order to obtain the best immunity from translational cross-coupling. However we find from analysis that we can obtain very good translational motion rejection over the main frequency band of interest while retaining DC tilt sensitivity below it. This can be done by tuning the spring-rates of the pivots at the top and bottom of the pendulum so that the angle of the test mass mirror under translational excitation remains constant with respect to the fixed mirror. The typical performance of such tuning is the solid curve in figure 7.12. Between the low frequency tilting resonance and the high frequency pendulum resonance, the cross-coupling can be made very low and almost constant. It seems that this is quite a wide tolerance adjustment (of order 1% in spring constant) and the level of rejection then obtained depends primarily on the Q-factor (1/loss tangent) of the flexure material.

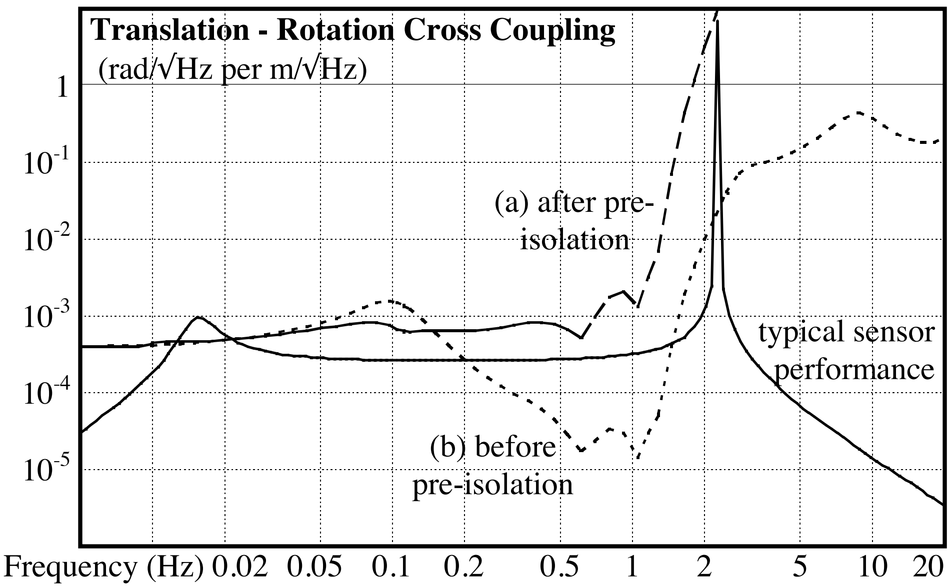

Figure 7.12 Required translational motion rejection if sensor is mounted (a) after or (b) before first stage of pre-isolation. Solid curve is a typical sensor performance available for particular design parameters.

The best material Q-factor we have measured for a thin metal flexure membrane (50 μm thick by 2 cm wide) is ~10^5 and was obtained from niobium [Baker 2000]. However for this application a value of 10^3 seems to be plenty and is readily obtainable from most metals (e.g. steel) if cut from a monolithic block. The particular curve shown was calculated for a 15 kg disk with 8 cm radius of gyration, suspended by a pendulum 5 cm long with flexures made from material with a (hysteretic damping) Q-factor of 10^3.

This length was chosen to place the pendulum resonance above 2 Hz where its effect is readily kept below the required level. Also being so short, both flexures may be cut in a single EDM operation.

The flexure parameters for this design were not particularly challenging although analysis shows that reducing the effective Q-factor of the lower flexure (the low frequency resonance) gives a much flatter sensitivity function at low frequency (shown in figure 7.13). This is achieved by making the lower flexure much stiffer than need be and then offsetting this with some inverse pendulum effect by attaching it at a point significantly below the centre of mass of the disk. The upper flexure was nominated to be 50μm thick and 2 mm wide to support the load at 50% yield (Rockwell 49 tool steel with extrapolated yield strength 2.8 GPa) giving an angular spring-rate of 0.026 N·m/rad, while the lower flexure was nominated to have ten times this spring-rate but attached 1.68 mm below the centre of mass of the disk to give a 96% spring-rate cancellation. With these values the resonant rocking period of the disk with pendulum held rigid would be 20 sec, but with the pendulum free to tilt, the interaction gives the mode period of 65 sec appearing in the graphs. The actual period of the low frequency rocking resonance is not important and it is intended that adjusting screws providing a small centre of mass height adjustment be used to adjust the effective spring-rate of the lower pivot for the best cross-coupling rejection.

A remarkable feature of most of these pendulum suspended arrangements is that the resonances do not appear in the sensitivity function at all. Figure 7.13 shows the sensitivity function for the arrangement described and it is apparent that there is no visible effect from the high Q-factor resonance just above 2 Hz. The poles and zeroes are present in the transfer function but they cancel almost perfectly (the plotting routine used solves for all roots and plots a point at each root). In fact it was found that for many areas of parameter space the low frequency resonance also disappears entirely from the sensitivity function. Unfortunately it doesn't seem possible to obtain this low frequency resonance disappearance at the same time as the spring-rate tuning required to minimise the translational cross-coupling. However in the case of low frequencies, we intend to servo the disk to the electrical zero, the force output becoming the tilt signal at these frequencies. This method has the useful effect of flattening out any sensitivity variations by the gain in the feedback. The disappearance of the 2 Hz resonance is convenient in that its effect doesn't need to be considered when determining the stability of the main tilt cancellation servo control loop.

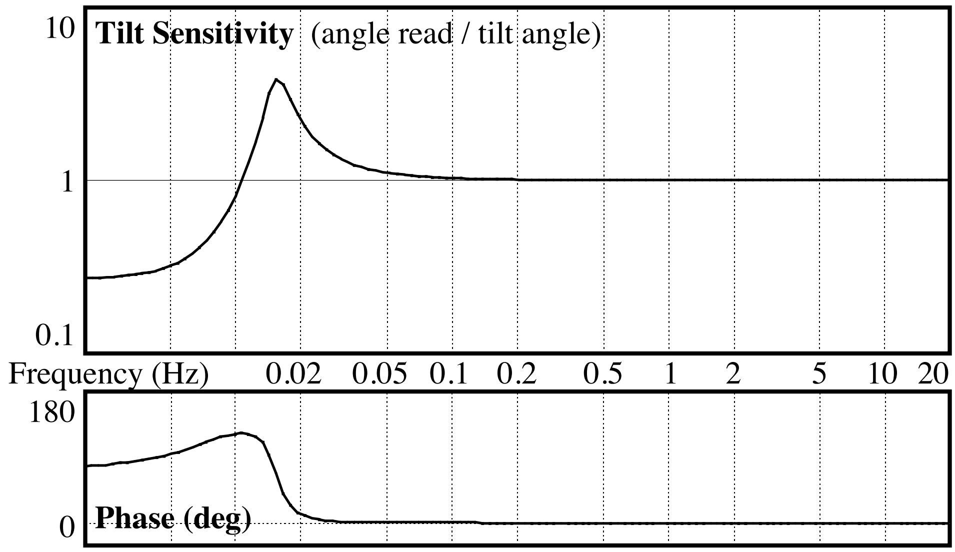

Figure 7.13 Transfer function of angle sensed from seismic tilt.

It is apparent from the cross-coupling curve in figure 7.12, that both of the resonances are coupled strongly to tilt. This means that they can be sensed and actuated on by the control loop and magnet-coil actuators and can be actively damped if required. This aspect has not been investigated however.

7.3.12 New Angle Sensor

The readout of angle may be achieved by various means but we propose a laser beam walk-off sensor. A laser beam is bounced from a mirrored surface onto a quadrant photo-diode thereby obtaining both angles with a single readout device. The laser beam displacement produced from a given change in test-mass angle can be greatly magnified by mounting another mirror on the frame almost parallel to the test mass mirror surface (as shown in figure 7.11), and bouncing the laser beam back and forth between them multiple times before landing on the detector. We have termed this angle sensing technique a “walk-off” sensor. The spot displacement produced for a given mirror angle change increases as the number of bounces squared, but due to finite beam divergence the sensitivity gain obtained will normally only increase as the number of bounces. A report [Zhou 2000] on our initial experimental investigation of the performance of such a device will appear shortly.

The number of bounces obtainable from a given mirror size can be increased fourfold over the simple parallel arrangement in figure 7.11 by presetting the fixed mirror to a small offset angle so that the laser traverses its width twice as illustrated in figure 7.14(a). If in addition one mirror is made slightly concave so that the diverging beam is re-focussed back at the detector as in figure 7.14(b), then sensitivity can be greatly increased (to go as the number of bounces squared) and the signal becomes insensitive to laser beam pointing fluctuations.

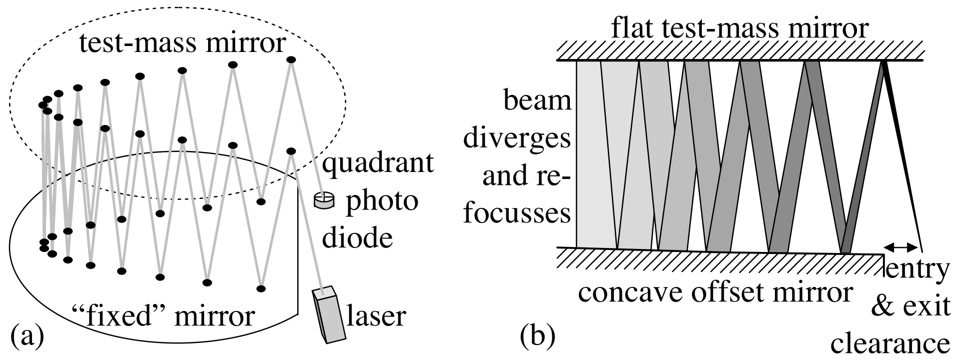

Figure 7.14 (a) Walk-off sensor laser path, (b) Laser focussing profile.

The sensitivity (and inversely dynamic range) of the device is adjustable in steps by changing the preset angle of the fixed mirror (and aim of laser) so that greater or fewer beam bounces occur. A smooth variation is also possible by changing the spacing between mirrors to obtain a better or worse focussed spot at the photo-diode.

If a flat mirror is used on the test mass, then it is clear that the device is insensitive to test mass motion parallel to this flat surface (i.e. x, y and θ_z). If the beam is detected at the same x-y position that it is sourced from the laser (e.g. with a beam splitter or by sourcing and detecting at the same position on opposite mirrors) then the sensor is also insensitive (when centred) to spacing between the mirrors (z motion). Since the z dimension is quite rigid, it is probably sufficient to simply mount laser and photo-diode offset slightly from each other. However if the mirror surface is not perpendicular to the pendulum it is suspended by, then horizontal motion in the direction of this offset will appear as variation in the spacing between the mirrors and will be sensed as cross coupling. This perpendicularity can be adjusted coarsely by mass redistribution and finely by a servo force applied to the test mass.

7.3.13 Experimental Test

A one-dimension tilt balance was constructed as shown in figure 7.15. A brass bar 38 cm long and 8 cm diameter, weighing 13.6 kg, was suspended near its centre by a metallic glass flexure (Goodfellow Co66/Si16/B12/Fe4/Mo2) 5.5 cm long. The thickness of the flexure was 25 μm, and its width was 20 mm giving an angular spring-rate under tension of 0.023 N·m/rad. Suspending this bar with this flexure with the centre of flexing at the centre of mass would give a period of 17 sec. It was balanced and tuned with the adjusting screws to a give a period of ~20 sec and the Q-factor was measured at ~190. (This suggests that the maximum period obtainable by further raising of the centre of mass would be ~270 sec before becoming critically damped).

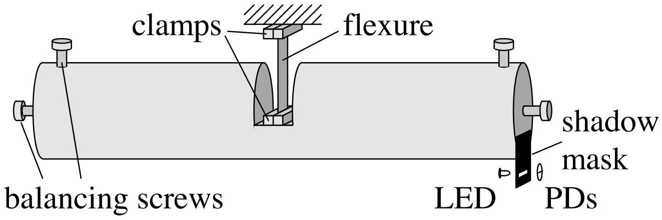

Figure 7.15 Experimental single axis tilt sensor using shadow sensor.

A simple 3 mm illuminated slot and split photo-diode shadow sensor with a sensitivity of 4×10^{-11} m/√Hz was used to detect the motion of the balance at one end as shown in figure 7.15. The whole system was put in a container but without vacuum. The bar is zero balanced with the end screws and its resonant period is adjusted coarsely by adding small lead plates on the top of the bar and finely with the top screws. The tilt noise measured by the shadow sensor is the solid curve shown in figure 7.16 (also indicated in figure 7.10), and it showed that the tilt sensor's sensitivity was $\sim2\times10^{-10}$ rad/√Hz limited by the shadow sensor noise. The peaks at about 2.1 Hz and 6.4 Hz were the horizontal mode of the tilt sensor and that of the support frame respectively. Even though this sensor was twice as sensitive as Speake and Newell's [Speake 1990], it seems that the seismic tilt in our basement laboratory is not measurable at this level.

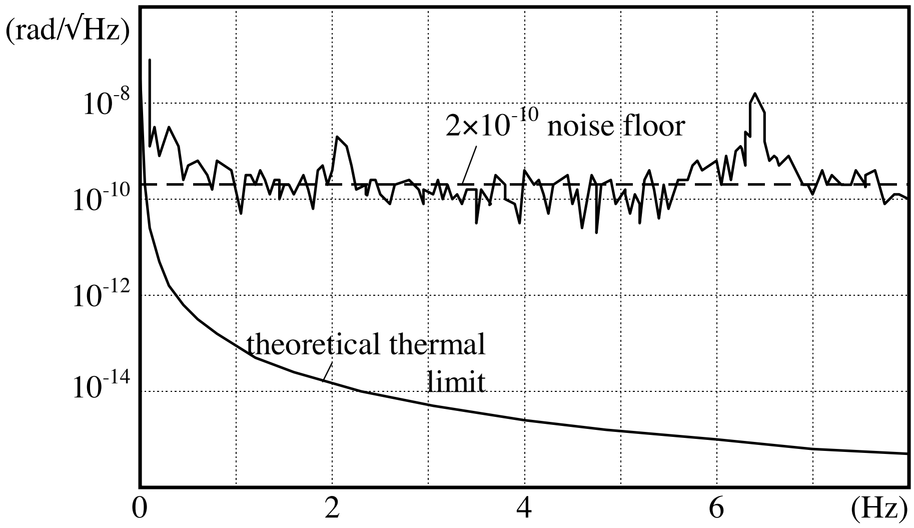

Figure 7.16 Seismic tilt measurement and predicted noise levels.

7.3.14 Conclusion

We have shown that the use of double stage ULF pre-isolation can result in the substantial reduction in residual motion of test masses provided it is combined with effective tilt control. This greatly reduces the size of the actuation forces required to control an interferometer. This can be quantified for our assumed seismic level as shown in figure 7.17. Forces are reduced from ~0.1 N for a simple all metal pendulum stack, to ~1μN for a damped, double pre-isolated and tilt stabilised system.

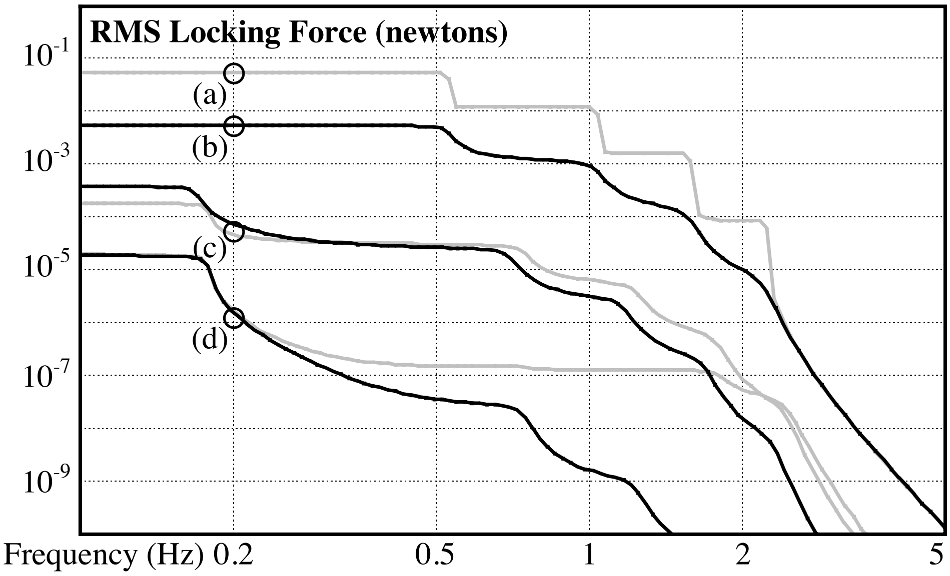

Figure 7.17 RMS Force required to hold test mass fixed - (a) no pre-isolation or damping, (b) normal modes damped, (c) single pre-isolation (black = translation, grey = tilt), (d) dual pre-isolation with active tilt suppression (black = translation, grey = tilt).

This system also reduces residual motion to the nanometre level at which it should take of order 100 seconds to drift through a 1 micron laser fringe prior to acquiring lock. Depending on the low frequency sensitivity required of the overall instrument, this system may allow direct actuation on the test mass by frame mounted actuators. We believe this system to be a minimally complex and highly practical solution to achieve the high sensitivity goals of advanced GW detectors.

7.3.15 Acknowledgment

This work was supported by the Australian Research Council.

7.4 Conclusion

While we expect this double passive pre-isolation with active tilt control to provide more than adequate seismic isolation, this structure provides an excellent upgrade path to ultimate active isolation. With the second stage of high-performance ultra-low-frequency 3-D suspension in place, it becomes possible to use the position of this last suspension point (with respect to the tilt-rigid frame) to feedback to the tilt-rigid frame to create effectively a "force-free" suspension (to some low-frequency limit). This is the same technique as applied in the "super spring" [Faller 1979], but in this case would be applied in all three dimensions.

Another way of conceiving this active vibration control is that the second stage of pre-isolation together with the entire suspended mass, becomes a gigantic 3-D accelerometer mounted on the tilt-rigid frame. Feedback is applied to the tilt-rigid frame to null any acceleration sensed, thereby actively stabilising this frame. This now provides a high-performance pre-isolator, a multistage isolation stack, and the test mass, all suspended from an actively stabilised frame. It seems very difficult to improve on such an ultimate structure!

8. Analysis Methods and Examples

8.1 Analysis of Isolation Systems Using State-Space Matrices

Consider the vertical and horizontal isolation chains depicted in figure 8.1. The vertical chain has a variable number of springs in parallel, the number suggesting the strength of the spring required. In each case however the spring-rate k_n is that of the whole set working together. The dissipation rate d_n is similarly the value for one whole set of springs and represents the viscous force per unit velocity (of expansion or contraction). For the horizontal stages the spring-rates are given by the equations in the figure - being simply the tension in the fibre above the mass divided by the fibre length. (The same equations hold for linear springs if h_n is the extension under system load). The analysis below for horizontal and vertical regimes using the spring-rates k_n found in this fashion are identical. The parameter d_n may be determined by each stage's individual oscillation frequency ω_n (rad/sec) and Q-factor Q_n using the simple relation $d_n = \omega_n m_n / Q_n$.

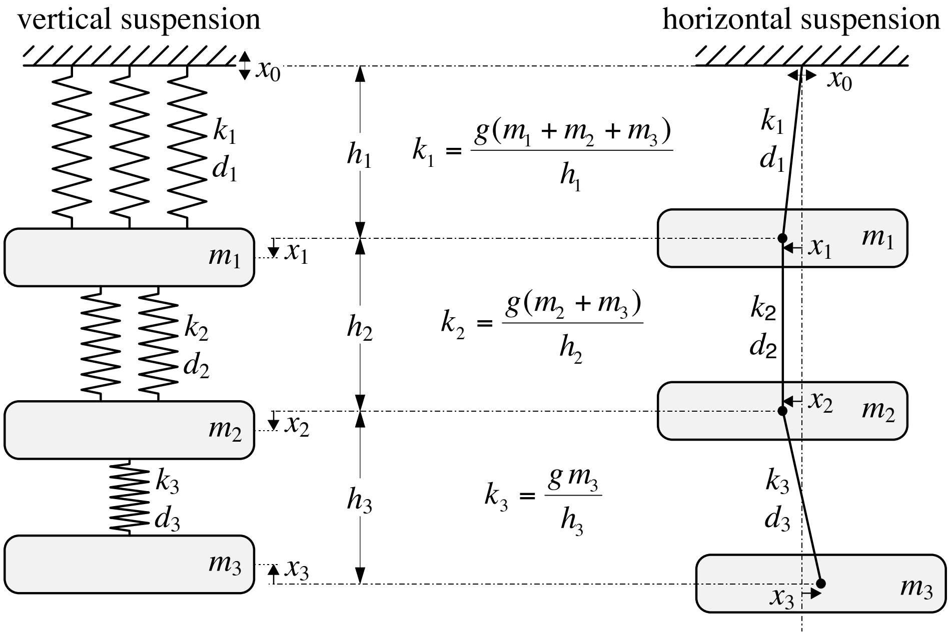

Figure 8.1 Parameters for analysis of vertical and horizontal isolation chains.

In order to use the state-space method of system analysis we may define the state variables as being the position and velocity of each of the suspended masses :- $x_1, x_2, \ldots, x_n$, and $v_1, v_2, \ldots, v_n$. The equation of motion for each mass may then be written as the mass times its acceleration = the sum of the forces acting on the mass :

$$
\begin{aligned}
m_1 v_1 &= -k_1(x_1 - x_0) - k_2(x_1 - x_2) - d_1(v_1 - v_0) - d_2(v_1 - v_2) \\
m_2 v_2 &= -k_2(x_2 - x_1) - k_3(x_2 - x_3) - d_2(v_2 - v_1) - d_3(v_2 - v_3) \\
m_3 v_3 &= -k_3(x_3 - x_2) \qquad\qquad\qquad\quad - d_3(v_3 - v_2)
\end{aligned}
$$

The full set of first order differential equations for this system then may be written as :

$$
\begin{array}{lcl:l:l:l:l:l:l}
\dot{v}_1 & = & -\dfrac{d_1+d_2}{m_1}v_1 & +\dfrac{d_2}{m_1}v_2 & \cdot & -\dfrac{k_1+k_2}{m_1}x_1 & +\dfrac{k_2}{m_1}x_2 & \cdot & +\dfrac{k_1+sd_1}{m_1}x_0 \\
\dot{v}_2 & = & \dfrac{d_2}{m_2}v_1 & -\dfrac{d_2+d_3}{m_2}v_2 & +\dfrac{d_3}{m_2}v_3 & +\dfrac{k_2}{m_2}x_1 & -\dfrac{k_2+k_3}{m_2}x_2 & +\dfrac{k_3}{m_2}x_3 & \\
\dot{v}_3 & = & \cdot & +\dfrac{d_3}{m_3}v_2 & -\dfrac{d_3}{m_3}v_3 & \cdot & +\dfrac{k_3}{m_3}x_2 & -\dfrac{k_3}{m_3}x_3 & \\
\dot{x}_1 & = & v_1 & \cdot & \cdot & \cdot & \cdot & \cdot & \\
\dot{x}_2 & = & \cdot & v_2 & \cdot & \cdot & \cdot & \cdot & \\
\dot{x}_3 & = & \cdot & \cdot & v_3 & \cdot & \cdot & \cdot &
\end{array}
$$

Where we have replaced d_1v_0 with sd_1x_0 (anticipating the transformation to the complex frequency domain). These may be rewritten in matrix form as :

$$
\begin{bmatrix} \dot{v}_1 \\ \dot{v}_2 \\ \dot{v}_3 \\ \hdashline \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix}
=
\left[\begin{array}{ccc:ccc}
-\dfrac{d_1+d_2}{m_1} & \dfrac{d_2}{m_1} & 0 & -\dfrac{k_1+k_2}{m_1} & \dfrac{k_2}{m_1} & 0 \\
\dfrac{d_2}{m_2} & -\dfrac{d_2+d_3}{m_2} & \dfrac{d_3}{m_2} & \dfrac{k_2}{m_2} & -\dfrac{k_2+k_3}{m_2} & \dfrac{k_3}{m_2} \\
0 & \dfrac{d_3}{m_3} & -\dfrac{d_3}{m_3} & 0 & \dfrac{k_3}{m_3} & -\dfrac{k_3}{m_3} \\
\hdashline
1 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0
\end{array}\right]
\begin{bmatrix} v_1 \\ v_2 \\ v_3 \\ \hdashline x_1 \\ x_2 \\ x_3 \end{bmatrix}
+
\begin{bmatrix} \dfrac{k_1+sd_1}{m_1} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} x_0
$$

This has the general form :-

$$\dot{\mathbf{x}} = \mathbf{A}\,\mathbf{x} + \mathbf{B}\,u$$

The solution in the complex frequency domain (where **I** = identity matrix) is :-

$$\mathbf{X}(s) = \left(s\mathbf{I} - \mathbf{A}\right)^{-1}\mathbf{x}_{(t=0)} + \left(s\mathbf{I} - \mathbf{A}\right)^{-1}\mathbf{B}\,U(s)$$

The output/input transfer function (ignoring initial conditions) is :-

$$\frac{\mathbf{X}(s)}{U(s)} = \left(s\mathbf{I} - \mathbf{A}\right)^{-1}\mathbf{B}$$

This gives an array of all the transfer functions from the input U(s) to each of the state variables. The second half of this array (to the x_n positions) gives the transfer function from X_0 to each of the n masses - but usually only the last one (X_3/X_0) is of interest.

In order to find the motion amplitudes (and phases) for a particular mode, it is necessary to find the mode frequency (from the roots of the denominator of a transfer function) and then evaluate the array of n transfer functions at that frequency to find the relative normal mode motions of the n individual masses for that normal mode.

8.2 Analysis of Isolation Systems Using Transfer Matrices

8.2.1 Introduction

The state-space analysis described in the previous section is powerful and universal. However it generally involves significantly more computation than is necessary to solve the problem of interest. For instance the way it is usually applied, it finds the transfer function between every possible combination of two state variables in a system. Typically each mass has 2 state variables associated with it, unless rotation has to be taken into account as well as translation, in which case each mass has 4 state variables. Usually only the transfer function between one variable at the start and one at the end of a chain is of interest. In addition it is not such an obvious and simple procedure to determine the matrices themselves given a mechanical schematic. For larger systems, the inverting of large matrices rapidly becomes time consuming even when the values are all numeric and consequently the motivation to find an alternative approach is quite strong.

The transfer matrix approach overcomes all these deficiencies to a large degree. For straightforward mechanical arrangements, it is a simple matter to determine the required matrices. All intermediate nodes are eliminated by simple matrix multiplication so that the largest matrix operations being handled are 2x2 matrices for single degree of freedom systems, or 4x4 matrices to handle rotation and translation.

The origins of the method have not been investigated, but it seems likely that it was developed, or at least became widespread for the analysis of electrical networks (interconnected two-port networks). Since the differential equations governing the relationships between voltage and current in electrical systems, and motion and force in mechanical systems are analogous to each other, it is possible to use these same techniques to analyse complex mechanical systems. Various analogous and dual relationships may be used, and the concepts of electrical impedance and admittance, can be made direct analogues of mechanical compliance and stiffness. Also the main mechanical components (springs, dashpots, masses and levers) may be found to have a direct electrical equivalent (inductors, resistors, capacitors and transformers). These analogy ideas are fully explored in several texts ([Harman 1962] for instance) but making heavy use of such analogies are only useful where a person is more familiar with the analogous system than with the system under consideration. With the simple concepts introduced here, electrical analogies may not be helpful, and we have chosen to use the variables of force F and position x rather than say force and velocity which have a closer correspondence to electrical networks. Also it goes without saying that all work is done in the complex frequency domain using the usual variable s to denote the complex frequency.

In vibration isolation structures, there are six degrees of freedom to be considered. Usually x and y are considered to be in the horizontal plane and z is considered to be the vertical. The angular degrees of freedom are usually considered to be rotation about these three axes giving θ_x, θ_y and θ_z. Some degrees of freedom (such as z the vertical) are typically quite uncoupled from other degrees, whereas others (such as x and θ_y) are almost unavoidably strongly coupled and need to be treated as a pair. Thus vertical isolation (z) typically only requires a one-dimensional analysis - a cascaded system of masses on springs has only motion in a vertical line along the z axis. The same is typically true of θ_z torsional isolation - again the only motion present being the angle about the z axis. The one-dimensional treatment is adequate for these degrees of freedom.

A one-dimensional treatment is also sufficient for horizontal isolation (x and y), if significant care has been taken to ensure that the stages are close approximations to simple pendulums. However when it is desired to include the effects of rocking modes (θ_x and θ_y), rotational inertia, finite mass connecting links, finite flexibility joints, and joints not at centres of percussion, then a two-dimensional treatment is required. In this case the x-θ_y and y-θ_x systems are usually sufficiently uncoupled to be treated as separate pairs.

8.2.2 One Dimensional Overview

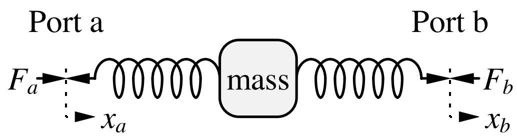

Figure 8.2 A simple “two-port” mechanical network.

Consider the simple two-port mechanical network depicted in figure 8.2. Because the connection between port-a and port-b is not rigid, displacements at port-a are different from displacements at port-b. Also because of the mass between the springs which interacts with the inertial reference frame, forces applied at port-a do not immediately appear at port-b. However only two of the four variables shown are independent and the other two can be derived from them. As a result there are 6 possible pairs of independent variables and associated sets of 4 matrix parameters. We are most interested in the set called the transmission (also called ABCD or "chain") parameters used to characterise the two-port network. Using this parameter set the matrix equation may be written in general form as follows :-

$$\begin{bmatrix} F_a \\ x_a \end{bmatrix} = \begin{bmatrix} \left.\dfrac{F_a}{F_b}\right|_{x_b=0} & \left.\dfrac{F_a}{x_b}\right|_{F_b=0} \\ \left.\dfrac{x_a}{F_b}\right|_{x_b=0} & \left.\dfrac{x_a}{x_b}\right|_{F_b=0} \end{bmatrix} \begin{bmatrix} F_b \\ x_b \end{bmatrix}$$

Thus the first term is the ratio of the force at port-*a* to the force at port-*b* with no displacement present at port-*b* (i.e. hold x_b fixed). It can be found by changing any one of the three unfixed variables to find the ratio between the two of interest. The lower right term is the ratio of the position at one end to the position at the other end with no force applied to the right hand end. If the left hand end is the attached end of an isolating structure, and the right hand end as the isolated mass, then this term gives the (reciprocal of the) transfer function that we are usually interested in.

The power of the technique is that when such networks are cascaded, the transfer matrix of the cascade is simply the dot product of the transfer matrices of the individual networks. Suppose we separated the network of figure 8.2 into a cascade of simpler networks with only a single element in each network as shown in figure 8.3.

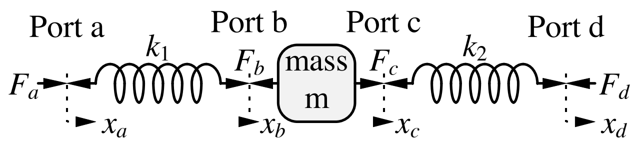

Figure 8.3 The two-port network broken up into a sequence of simpler networks.

In this case the individual transfer matrices are so simple that they can be written (with a little experience) by inspection as follows :-

$$\begin{bmatrix} F_a \\ x_a \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1/k_1 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & ms^2 \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 \\ 1/k_2 & 1 \end{bmatrix} \begin{bmatrix} F_d \\ x_d \end{bmatrix}$$

Multiplying the matrices out is a very simple matter for a computer :-

$$\begin{bmatrix} F_a \\ x_a \end{bmatrix} = \begin{bmatrix} 1 + \dfrac{ms^2}{k_2} & ms^2 \\ \dfrac{1}{k_1} + \dfrac{1 + ms^2/k_1}{k_2} & 1 + \dfrac{ms^2}{k_1} \end{bmatrix} \begin{bmatrix} F_d \\ x_d \end{bmatrix}$$

Which gives a transfer function of :-

$$\frac{x_d}{x_a} = \frac{1}{1 + ms^2/k_1}$$

This is the essence of the technique. With this overview, we can step back and cover the topic a little more methodically.

8.2.3 Basic Elements

The basic elements consist of masses, springs (or pendulum links for the horizontal regime) and dashpots, and the main property of each element is its dynamic stiffness (equal to force applied divided by the resulting displacement). The inverse of this value may be called receptance or dynamic flexibility, or compliance. This author has chosen to use the symbol Y for the dynamic stiffness (as these stiffnesses combine in parallel and series in the same manner as admittances - with symbol Y - in electrical networks)

The elements differ significantly in how force can be applied to them. A mass is a one-terminal device - the “second” terminal being its intrinsic connection to the inertial reference frame or ground. Springs, pendulum links, and dashpots are two-terminal devices and can be connected in series and/or parallel with each other to give a more complex two-terminal device. The two-terminal devices may be turned into a one-terminal device if the second terminal is fixed (connected to ground or to an effectively infinite mass).

The spring-rate of a spring can have some structural damping which is adequately modelled by the imaginary component of a complex spring-rate value. Viscous damping is put in separately using a dashpot. Since the restoring force of a pendulum link proper is purely gravitational, it can have no loss. In order to model a lossy pendulum system without having to include the rotational degree of freedom (which is where the losses actually occur in the angular spring-rates), provision is made as shown for hysteretic damping in the pendulum link also. This is achieved by making the tension due to gravity a lossy force. This is not physical of course but the result is graphically indistinguishable from having rotational damping in the pendulum joints. The basic elements from which transfer matrices may then be constructed include :-

Element	Stiffness Y	Parameters
mass	$m\,s^2$	m = mass
spring	$k\,(1+i/Q)$	k = spring-rate, Q = structural Q-factor
dashpot	$d\,s$	d = damping coefficient: force per unit velocity
link	$T\,(1+i/Q)/h$	T = tension in link, h = length of link, Q = structural Q

Table 8.1 Basic mechanical elements useable in mechanical networks.

8.2.4 Combining Elements into Matrices

The two-terminal elements may be combined in series and/or parallel in order to construct a more complex two terminal element. The stiffness values are added for a parallel connection and the compliance values (reciprocals) are added for a series connection (which needs to be reciprocated again to regain a stiffness value). One-terminal elements (mass) can also be combined in parallel and series with the two-terminal elements as long as all the "second" terminals of the masses are to have a common connection to ground, thereby resulting in a more complex one-terminal element. In order to construct a network with *more* than two terminals - i.e. a two-*port* network with an input terminal, an output terminal, and a ground connection, the result needs to be specified by a *matrix* as overviewed in section 8.2.2.

There is only one way of building a one-terminal element into a two-port network, and that is to have the displacement common to both ports ($x_a = x_b$), the network only modifying the applied forces. There is also only one way of building a two-terminal element into a two-port network, and that is to have the force common at both ports ($F_a = F_b$), the network only modifying the displacements. (The two-terminal element may be turned into a one-terminal element and then applied in the former manner however). The two networks that can be built from one-terminal and two-terminal elements may be called respectively a "Gnd" matrix and a "Via" matrix as follows :-

If Y is a one-terminal stiffness $\qquad Gnd(Y) = \begin{bmatrix} 1 & Y \\ 0 & 1 \end{bmatrix}$

If Y is a two-terminal stiffness $\qquad Via(Y) = \begin{bmatrix} 1 & 0 \\ 1/Y & 1 \end{bmatrix}$

When these two port networks are chained together, port-*b* of the leading network being connected to port-*a* of the following network, the matrices are simply multiplied to obtain the resulting two-port network.

There is one additional method which is quite useful to know, and that is how to turn a two-port network into a one-terminal stiffness element (when the second terminal is left unconnected for instance - i.e. no force is applied to it). This would be useful if we wanted to calculate the transfer function between the attached end of a suspension chain and some mid-point mass half way down. In this case the transfer function between the attached end and the mid-point is easy to find, but that mid-point needs to be *loaded* with the dynamic stiffness of the remainder of the chain in order to obtain the correct function. For this we would like to be able to use transfer matrices to calculate the matrix of the remainder of the chain and then turn that matrix into a simple one-terminal stiffness to add in parallel with, or to include in a Gnd matrix just before or after the mid-point mass. This stiffness is simply the ratio of force applied F_a divided by resulting displacement x_a with no force applied at the far end i.e. F_b=0. Simply re-arranging the two matrix equations gives the required value. The dynamic stiffness presented at port-*a* Ya, and in a similar fashion the transfer function x_b/x_a TF, are thus obtained from the transfer matrix as follows :-

$$Ya\left(\begin{bmatrix} A & B \\ C & D \end{bmatrix}\right) = \frac{B}{D} \tag{8.1}$$

$$TF\left(\begin{bmatrix} A & B \\ C & D \end{bmatrix}\right) = \frac{1}{D} \tag{8.2}$$

This is a sufficient basis to be able to construct most one-dimensional mechanical networks commonly encountered.

8.2.5 One Dimensional Example - Vibration Absorber

Consider the vibration absorber of section 1.4.5 duplicated in figure 8.4. We wish to find the transfer function from the attachment point at x_a to the main mass at x_b. This is a case for which we would like to find the stiffness from the matrix for the section from x_b to x_c.

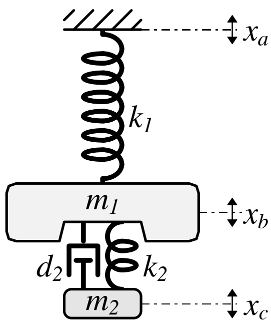

Figure 8.4 Single spring-mass stage with vibration absorber.

Firstly finding the transfer matrix T_{bc} from x_b to x_c :-

$$T_{bc} = Via\left(d_2\,s + k_2\right)\cdot Gnd\left(m_2\,s^2\right)$$

$$= \begin{bmatrix} 1 & 0 \\ 1/\left(d_2\,s + k_2\right) & 1 \end{bmatrix}\cdot\begin{bmatrix} 1 & m_2\,s^2 \\ 0 & 1 \end{bmatrix}$$

$$= \left[\begin{array}{c|c} 1 & m_2\,s^2 \\ \hline 1/\left(d_2\,s + k_2\right) & 1 + m_2\,s^2/\left(d_2\,s + k_2\right) \end{array}\right]$$

The dynamic stiffness Y_{bc} for the section from b to c is found from (8.1) :-

$$Y_{bc} = \frac{m_2\,s^2}{1 + m_2\,s^2/\left(d_2\,s + k_2\right)}$$

Then finding the transfer matrix T_{ab} from x_a to x_b :-

$$T_{ab} = Via\left(k_1\right)\cdot Gnd\left(m_1\,s^2\right)\cdot Gnd\left(Y_{bc}\right)$$

$$= \begin{bmatrix} 1 & 0 \\ 1/k_1 & 1 \end{bmatrix}\cdot\begin{bmatrix} 1 & m_1\,s^2 \\ 0 & 1 \end{bmatrix}\cdot\begin{bmatrix} 1 & Y_{bc} \\ 0 & 1 \end{bmatrix}$$

$$= \left[\begin{array}{c|c} 1 & m_1\,s^2 + Y_{bc} \\ \hline 1/k_1 & 1 + \left(m_1\,s^2 + Y_{bc}\right)/k_1 \end{array}\right]$$

If it was desired to cascade several of these damped absorbers in a suspension chain, it is just a matter of multiplying the appropriate number of this matrix for a single stage together. In this case we only require the transfer function of a standalone stage TF_{ab} which is the reciprocal of the lower right matrix element as given by (8.2) :-

$$TF_{ab} = \frac{k_1}{k_1 + m_1\,s^2 + Y_{bc}}$$

$$= \frac{k_1}{k_1 + m_1\,s^2 + \dfrac{m_2\,s^2}{1 + \dfrac{m_2\,s^2}{d_2\,s + k_2}}}$$

8.2.6 Two Dimensional Analysis

In a typical horizontal isolator (such as cascaded pendulums) the motion is transverse to the direction of support. Unless the attachment point(s) are perfect pivots at the centre of mass, then the motion of each mass will consist of rotation as well as horizontal translation. To take rotational motion and inertia of the masses into account in an analysis it is necessary to include the rotational variables of moment M and angle , in addition to the linear variables of force F and displacement x, for each network to be connected in a way which is capable of acting in both degrees of freedom. Figure 8.5 shows these variables for a single network. The boxes shown are both the same mechanical system - one box for each degree of freedom. They are shown vertically as if suspended with the line of action perpendicular to the line of suspension in order to correspond closely with the intended application of analysing pendulum systems and to assign an unambiguous polarity to the variables. The a and b port labelling is convenient for "above" and "below" the network.

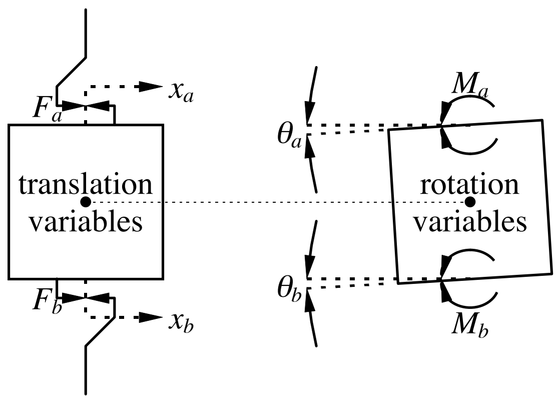

Figure 8.5 Transfer matrix variables for a network incorporating moment and angle.

Each network will now be characterised by a 4×4 matrix instead of the 2×2 that was used for the single degree of freedom case. The matrix equation now has the form :-

$$\begin{bmatrix} F_a \\ x_a \\ \hline M_a \\ \theta_a \end{bmatrix} = \left[\begin{array}{cc|cc} F_a/F_b & F_a/x_b & F_a/M_b & F_a/\theta_b \\ x_a/F_b & x_a/x_b & x_a/M_b & x_a/\theta_b \\ \hline M_a/F_b & M_a/x_b & M_a/M_b & M_a/\theta_b \\ \theta_a/F_b & \theta_a/x_b & \theta_a/M_b & \theta_a/\theta_b \end{array}\right] \begin{bmatrix} F_b \\ x_b \\ \hline M_b \\ \theta_b \end{bmatrix} \tag{8.3}$$

Where each ratio in the 4×4 matrix is measured while keeping the three remaining independent variables (on the right hand side) fixed at zero.

The matrix may be divided up into four quadrants as shown, in which case the upper left quadrant is the old 2×2 translational matrix that has been used so far, while the lower right is its equivalent but for rotational forces and motion. The other quadrants allow for cross-coupling between these two degrees of freedom.

The mechanical elements that may be used in the rotational quadrant are the rotational analogues of the first three elements in table 8.1 namely :-

Linear Element	Rotational Equivalent	Angular Stiffness Y_a	Units
mass	moment of inertia	$m\, r^2\, s^2$	mass × (r = radius of gyration)^2
spring	angular spring	$k_a\,(1+i/Q)$	Newton meters/radian
dashpot	angular dashpot	$d_a\, s$	Newton meters/(radian/sec)
link	-		

Table 8.2 Rotational equivalents of basic linear mechanical elements.

The two networks that can be built from one-terminal and two-terminal elements now are most useful if they allow both translational and rotational elements to be included together :-

One-terminal stiffnesses
$$Gnd(Y,Y_a)=\begin{bmatrix}1 & Y & 0 & 0\\ 0 & 1 & 0 & 0\\ 0 & 0 & 1 & Y_a\\ 0 & 0 & 0 & 1\end{bmatrix}$$

Two-terminal stiffnesses
$$Via(Y,Y_a)=\begin{bmatrix}1 & 0 & 0 & 0\\ 1/Y & 1 & 0 & 0\\ 0 & 0 & 1 & 0\\ 0 & 0 & 1/Y_a & 1\end{bmatrix}$$

There are many useful building block networks which sometimes are not easily made using these two basic networks and which are easy enough to write the equations for and turn into matrix form. Some useful building blocks which will be required by a following example are a rigid link of length (or height) h and under tension from total mass suspended below the link of T :-

$$RigidLink(h,T)=\begin{bmatrix}1 & 0 & 0 & 0\\ 0 & 1 & 0 & -h\\ h & 0 & 1 & Th\\ 0 & 0 & 0 & 1\end{bmatrix}$$

A pivot with angular spring-rate k_a and angular damping-rate d_a :-

$$Pivot(k_a, d_a) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1/(k_a + d_a s) & 1 \end{bmatrix}$$

Having obtained the transfer matrix for a section of mechanical network which includes the rotation variables, it is almost always required to drop those variables in order to continue cascading that section to other sections by means of simple pivots. This may be done by rearranging the equations in the matrix so that the variables which are to be dropped are given the correct terminating conditions. The commonest case is of the section being joined by an ideally perfectly flexible pivot at both ends. In this case there is a negligible moment applied at the connection, and so the matrix equations should be rearranged so that the moments going to zero are on right hand side (i.e. both M_a and M_b on the right). That sets half of the terms in the matrix to zero, but there are still four (non-independent) equations. The next step is simply to choose from the variables on the left hand side, which two are to be continued with and discard the equations for the others. Applying this method to a general 4×4 arranged as equation (8.3) for the case of moment-free ends gives :-

$$MomentFreeEnds\left(\begin{bmatrix} A & B & c & D \\ E & F & g & H \\ I & J & k & L \\ m & n & o & p \end{bmatrix}\right) = \begin{bmatrix} \dfrac{A - DI}{L} & \dfrac{B - DJ}{L} \\ \dfrac{E - HI}{L} & \dfrac{F - HJ}{L} \end{bmatrix} \tag{8.4}$$

Where the terms from the matrix that are ignored are shown in lower case.

8.2.7 Two Dimensional Example - Self-damped Pendulum

Consider the standalone self-damped pendulum of section 1.4.5 duplicated in figure 8.6. We wish for the transfer function from the attachment point at x_a to the lower pivot at x_c. This is a case for which we need to use 4×4 matrices to find the transfer matrix from x_b to x_c, then convert this to a 2×2 matrix in order to cascade it with the simple fibre link transfer matrix from x_a to x_b, giving finally the 2×2 matrix x_a to x_c and the transfer function.

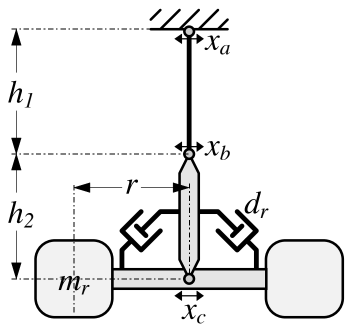

Figure 8.6 Single self-damped pendulum stage.

Firstly finding the 4×4 transfer matrix $T4_{bc}$ from x_b to x_c :-

$$T4_{bc} = RigidLink\left(h_2, m_r g\right) \cdot Pivot\left(0, d_r\right) \cdot Gnd(m_r, r^2 m_r)$$

$$= \left[\begin{array}{cc|cc} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & -h_2 \\ \hline h_2 & 0 & 1 & m_r g h_2 \\ 0 & 0 & 0 & 1 \end{array}\right] \cdot \left[\begin{array}{cc|cc} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 1 & 0 \\ 0 & 0 & 1/\left(d_r s\right) & 1 \end{array}\right] \cdot \left[\begin{array}{cc|cc} 1 & m_2 s^2 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & r^2 m_r s^2 \\ 0 & 0 & 0 & 1 \end{array}\right]$$

$$= \left[\begin{array}{c|c|c|c} 1 & m_r s^2 & 0 & 0 \\ \hline 0 & 1 & -h_2/\left(d_r s\right) & -h_2 - h_2 r^2 m_r s/d_r \\ \hline h_2 & h_2 m_r s^2 & 1 + m_r g h_2/\left(d_r s\right) & m_r g h_2 + r^2 m_r s^2\left(1 + m_r g h_2/\left(d_r s\right)\right) \\ \hline 0 & 0 & 1/\left(d_r s\right) & 1 + r^2 m_r s/d_r \end{array}\right]$$

The 2×2 transfer matrix T_{bc} for this section with perfect pivots at top and bottom is found using equation (8.4) :-

$$T_{bc} = \left[\begin{array}{c|c} 1 & m_r s^2 \\ \hline \dfrac{h_2\left(h_2 + h_2 r^2 m_r s/d_r\right)}{m_r g h_2 + r^2 m_r s^2\left(1 + m_r g h_2/\left(d_r s\right)\right)} & 1 + \dfrac{h_2 m_r s^2\left(h_2 + h_2 r^2 m_r s/d_r\right)}{m_r g h_2 + r^2 m_r s^2\left(1 + m_r g h_2/\left(d_r s\right)\right)} \end{array}\right]$$

This 2×2 transfer matrix T_{bc} is simply preceded in cascade with the transfer matrix for a simple pendulum link to give the transfer matrix for the entire stage T_{ac} :-

$$T_{ac} = Via\left(Link\left(h_1, m_r g\right)\right) \cdot T_{bc}$$

$$= \begin{bmatrix} 1 & 0 \\ h_1/\left(m_r g\right) & 1 \end{bmatrix} \cdot T_{bc}$$

$$T_{ac} = \left[\begin{array}{c:c} 1 & m_r\, s^2 \\ \hdashline \dfrac{h_1}{m_r\, g} + \dfrac{h_2\left(h_2 + h_2\, r^2 m_r\, s/d_r\right)}{m_r\, g\, h_2 + r^2 m_r\, s^2\left(1 + m_r\, g\, h_2/\left(d_r\, s\right)\right)} & 1 + \dfrac{h_1\, s^2}{g} + \dfrac{h_2\, m_r\, s^2\left(h_2 + h_2\, r^2 m_r\, s/d_r\right)}{m_r\, g\, h_2 + r^2 m_r\, s^2\left(1 + m_r\, g\, h_2/\left(d_r\, s\right)\right)} \end{array} \right]$$

The desired transfer function TF_{ac} is the reciprocal of the lower right matrix element :-

$$TF_{ac} = \frac{1}{1 + \dfrac{h_1\, s^2}{g} + \dfrac{h_2\, m_r\, s^2\left(h_2 + h_2\, r^2 m_r\, s/d_r\right)}{m_r\, g\, h_2 + r^2 m_r\, s^2\left(1 + m_r\, g\, h_2/\left(d_r\, s\right)\right)}}$$

It should be noted that unlike the previous example, multiple cases of this transfer matrix cannot be simply cascaded because there are terms present that depended on the tension in the supporting links - which depends on the mass of stages cascaded after it. If provision is made to update the tension producing mass depending on the masses cascaded, then cascading can be done.

The author uses *Mathematica*® [Wolfram 1988] routines that automatically accumulate the masses and calculate the tension up the chain based on the masses of the devices cascaded. This has proved to be very worth while, saving a lot of labour and doubtless avoiding many mistakes.

8.3 Derivation of Expression for Multistage Pendulum

The transfer matrix equation of motion relating the force and position above (F_a,x_a) to the force and position below (F_b,x_b) a single pendulum stage of mass m, resonant frequency ω_0, and no damping is :-

$$\begin{bmatrix} F_a \\ x_a \end{bmatrix} = \begin{bmatrix} 1 & m s^2 \\ \dfrac{1}{\omega_0^{\,2}\,\Sigma m} & 1 + \dfrac{m s^2}{\omega_0^{\,2}\,\Sigma m} \end{bmatrix} \begin{bmatrix} F_b \\ x_b \end{bmatrix}$$

where Σm stands for the sum of the masses producing tension in this pendulum's fibre (which in the case of a single stage is simply m), and s is the complex frequency. The x_a/x_b transfer function of this system in the case of the lower end being free (F_b=0) is seen by inspection to be the lower right term of the 2×2 matrix. (The more usual x_b/x_a isolation transfer function is simply its reciprocal.)

If several stages of varying mass m_1,m_2,··· are cascaded one above the other, then the transfer matrix equation relating force and position at the top (F_a,x_a) to those at the very bottom (F_b,x_b) becomes :

$$\begin{bmatrix} F_a \\ x_a \end{bmatrix} = \begin{bmatrix} 1 & m_1 s^2 \\ \dfrac{1}{\omega_1^{\,2}\Sigma m_1} & 1 + \dfrac{m_1 s^2}{\omega_1^{\,2}\,\Sigma m_1} \end{bmatrix} \begin{bmatrix} 1 & m_2 s^2 \\ \dfrac{1}{\omega_2^{\,2}\Sigma m_2} & 1 + \dfrac{m_2 s^2}{\omega_2^{\,2}\,\Sigma m_2} \end{bmatrix} \quad \begin{bmatrix} F_b \\ x_b \end{bmatrix}$$

In order to obtain the lower-right-hand (x_a/x_b) term by multiplying out two adjacent 2×2 transfer matrices, the product of the lower left and upper right term is summed with the product of the two lower right hand terms obtaining from the first two matrices :

$$\frac{m_2 s^2}{\omega_1^{\,2}\Sigma m_1} + \left(1 + \frac{m_1 s^2}{\omega_1^{\,2}\,\Sigma m_1}\right)\left(1 + \frac{m_2 s^2}{\omega_2^{\,2}\,\Sigma m_2}\right) = \frac{m_1\, m_2}{\Sigma m_1\, \Sigma m_2}\,\frac{s^4}{\omega_1^{\,2}\,\omega_2^{\,2}} + O\!\left[s^2\right]$$

If we are only interested in the systems performance at high frequencies, then we may discard all but the highest order in s. Repeating this process to multiply out all of the matrices and only keeping the term of highest order in s and then taking the reciprocal, the expression obtained for the transfer function at high frequency is simply :-

$$\frac{x_b}{x_a} \approx \frac{\Sigma m_1\, \Sigma m_2 \quad \Sigma m_n}{m_1\, m_2 \quad m_n}\,\frac{\omega_1^{\,2}\,\omega_2^{\,2} \quad \omega_n^{\,2}}{s^{2n}} \qquad \left(s >> \tilde{\omega}\right) \tag{8.5}$$

Replacing the resonant frequency of each stage with the expression containing its height $\omega^2 = g/h$ we obtain

$$\frac{x_b}{x_a} \approx \frac{\Sigma m_1 \, \Sigma m_2 \quad \Sigma m_n}{m_1 \, m_2 \quad m_n} \frac{g^n}{h_1 \, h_2 \quad h_n \, s^{2n}} \qquad \left(s >> \tilde{\omega}\right)$$

If we wish to minimise this expression to obtain the best isolation, and if the total height available (being the sum of the heights $h_1 \ldots h_\mathrm{n}$) is fixed, then the product of the h's is maximum when the total height is evenly divided between the pendulums so that they are all equal. (If a+b=constant, then a×b is maximum when a=b). This is the normal arrangement and using ω_0 now for the common resonant frequency for all the stages allows us to simplify equation (8.5) to

$$\frac{x_b}{x_a} \approx \frac{\Sigma m_1 \, \Sigma m_2 \quad \Sigma m_n}{m_1 \, m_2 \quad m_n} \frac{{\omega_0}^{2n}}{s^{2n}} \qquad \left(s >> \omega_0\right) \tag{8.6}$$

This expression may be further simplified for suspension chain arrangements which have simple relationships between the masses of the various stages. For instance one of the commonest arrangements is to have n stages of equal height and equal mass. In this case the high frequency transfer function becomes :-

$$\frac{x_b}{x_a} \approx n! \left(\frac{{\omega_0}^2}{s^2}\right)^n \qquad \left(s >> \omega_0\right) \tag{8.7}$$

Another possible arrangement is to have n stages of equal height but with masses changing towards the lower end by some mass ratio r between adjacent stages. If this value r is less than 1 then each term $\Sigma m_i / m_i$ is reduced because the initial m_i in the sum Σm_i dominates the sum as all the masses below it are getting smaller in size. It is useful to consider the best possible isolation which is obtained by making the ratio very small r→0 such that $\Sigma m_i / m_i \to 1$. This is equivalent to cascading filters which are buffered so that subsequent stages do not load or influence preceding stages to any degree. In this case the high frequency transfer function becomes very simply :-

$$\frac{x_b}{x_a} \approx \left(\frac{{\omega_0}^2}{s^2}\right)^n \qquad \left(s >> \omega_0\right) \tag{8.8}$$

It is clear that in order to approach this best possible transfer function, in a real situation where the total mass M_{total} ($=\Sigma m_1$) is limited, and the final test mass at the end of the chain m_n is finite, the ratio r should be minimised. Its value may be determined by solving the expression for the sum of a geometric series (which is not easily inverted to yield r on the left hand side) :-

$$M_{total} = m_n \frac{r\,(r^n - 1)}{r - 1} \tag{8.9}$$

There is probably a reason why having a *constant* ratio as opposed to some other decreasing function between adjacent masses actually optimises the isolation but it is not immediately obvious. It is obvious however that some sort of decreasing mass function towards the bottom of the chain must optimise the isolation and the ratio scheme used in section 3.4.3 was the result of this knowledge.

For the case of an isolation chain having stages with equal mass, equation (8.7) indicates that the transfer function approaches $n!(f_0/f)^{2n}$. Since for a given total height h the height of each individual stage will be h/n and resonant frequency $(n\, g/h)^{½}$. This means that an n stage isolator chain will provide a required thermal noise isolation A_{th} when $A_{th} = n!\,(n\, g/h)^n/f^{2n}$. Given the total available height h and the isolation required to reach thermal noise (which is approximately constant as the seismic noise has a similar spectrum to the thermal noise), gives a simple relationship between the number of stages and the frequency at which the isolation gets below the threshold. This relationship is shown in figure figure 1.5 for a height of 2.5 m and 10 m and reveals the interesting fact that for a given isolation threshold, there is an optimum number of stages to reach that threshold and still obtain the lowest cutoff frequency.

References

Abramovici 1992 A. Abramovici *et al.*, *LIGO: the Laser Interferometer Gravitational-Wave Observatory*, Science **256** (1992) 325

Ando 2001 M. Ando, *et al.*, *Report to NSF on the TAMA-SAS Passive Seismic Attenuation Chain Test for the Tokyo University 3 meter Interferometer*, Ligo technical report: LIGO-T010006-00-R.

Bachor 1989 H-A Bachor *et al.*, *The AIGO Proposal*, (A.N.U Canberra, 1989).

Baker 2000 M. Baker, L. Ju, I. Loo, D.G. Blair, *The Q-factor of Flexure Membranes*, submitted, but finally published as [Ju 2000c].

Barton 1996 M.A. Barton, N. Kanda, K. Kuroda, *A low-frequency vibration isolation table using multiple crossed-wire suspensions*, Rev. Sci. Instrum. **67** (1996) 3994-3999.

Barton 1999 M.A. Barton, T. Uchiyama; K. Kuroda; N. Kanda; H. Ishizuka, *Two-dimensional X pendulum vibration isolation table*, Rev. Sci. Instrum. **70** (1999) 2150-2154.

Beccaria 1997 M. Beccaria, *et al.*, *Extending the VIRGO gravitational wave detection band down to a few Hz: metal blade springs and magnetic antisprings*, Nucl. Instrum. Methods Phys. Res A **394** (1997) 397

Beccaria 1998 M. Beccaria *et al.*, *The creep problem in the VIRGO suspensions: a possible solution using Maraging steel*, Nucl. Instrum. Methods Phys. Res. A **404** (1998) 455.

Beggs 1955 see eg. J.S. Beggs, “Mechanism”, McGraw-Hill (1955).

Benioff 1955 H. Benioff, Advances in Geophys. 2 (1955) 219.

Bertolini 1999 A. Bertolini; G. Cella; R. DeSalvo; V. Sannibale, *Seismic noise filters, vertical resonance frequency reduction with geometric anti-springs: a feasibility study*, Nucl. Instrum. Methods Phys. Res. A **435** (1999) 475-83.

Bertolini 2001 A monolithic folded-pendulum accelerometer development is the subject of the doctoral thesis of Alessandro Bertolini of the University of Pisa.

Billing 1979 H. Billing, K. Maischberger, A. Rüdiger, R. Schilling, L. Schnupp, and W. Winkler, *An argon laser interferometer for the detection of* *gravitational radiation*, J. Phys E: Sci. Instrum., **12** (1979) 1043-1050

Blair 1991 *The detection of gravitational waves*, ed. D.G. Blair, Cambrige University Press, Cambridge, (1991)

Blair 1993 D.G. Blair, L. Ju, and H. Peng, *Vibration isolation for gravitational wave detection*, Class. Quantum Grav. **10** (1993) 2407-2418.

Blair 1994 D. G. Blair, J. F. Liu, L. Ju, *Performance of an ultra low-frequency folded pendulum*, Phys. Lett. A. **193** (1994) 223-6.

Boston 1994 PC-30 available from Boston Technology (1994).

Braccini 1993 S. Braccini *et al.*, *An improvement in the VIRGO Super Attenuator for interferometric detection of gravitational waves: the use of a magnetic antispring*, Rev. Sci. Instrum. **64** (1993) 310.

Braccini 1995 S. Braccini *et al.*, *Low noise wideband accelerometer using an inductive displacement sensor*, Rev. Sci. Instrum. **66** (1995) 2672.

Braccini 1996 S. Braccini *et al.*, *Seismic vibrations mechanical filters for the gravitational waves detector VIRGO*, Rev. Sci. Instrum. **67**, (1996) 2899.

Bradaschia 1990 C. Bradaschia, *et al.*, *The Virgo Project: A Wide Band Antenna for Gravitational Wave Detection*, Nucl. Instrum. Methods Phys. Res. A **289** (1990) 518-25.

Britvec 1973 S.J. Britvec, “The stability of elastic systems”, New York: Pergamon Press Inc, 1973.

Cella 1997 G. Cella *et al.* (incl. J. Winterflood), *Filter Zero Prototype Performances*. VIRGO:VIR-TRE-PIS-4600-125 (1997) later published as [DeSalvo 1999]

Cheng 2001 Y. Cheng, J. Winterflood, L. Ju, D.G. Blair, *Tilt sensor and servo control system for gravitational wave detection*, Class. Quantum Grav., accepted.

DeSalvo 1996 R. DeSalvo A. Gaddi, G. Gennaro, L. Holloway, G. Losurdo, and J. Winterflood, *A proposal for a Pre-Isolator Stage for the VIRGO SuperAttenuator* VIRGO:NTS-096/034 (1996).

DeSalvo 1998 R. DeSalvo, LIGO Project, Caltech, Private communication.

DeSalvo 1999	R. DeSalvo *et al*. (incl. J. Winterflood), *Performances of an ultralow frequency vertical pre-isolator for the VIRGO seismic attenuation chains*, Nucl. Instrum. Methods A **420** (1999) 316-35.
DeSalvo 2001	R. DeSalvo, "Pathfinding Towards Cryo-LIGO", seminar presentation at 4th Amaldi conference, July 2001 in Perth, Australia.
Faller 1979	J.E. Faller; R.L. Rinker, *Super spring mechanical isolation*, Dimensions NBS **64** (1979) 25-6
Feynman 1964	R.P. Feynman, R.B. Leighton, M. Sands, "The Feynman Lectures on Physics", Vol II, Addison-Wesley, 1964, Sec. 38-5.
Flaminio 1994	R. Flaminio, PhD Thesis, Physics Dept., Univ. of Pisa, Italy, 1994.
Giazotto 1989	A. Giazotto *et al*., *The VIRGO Project*, (INFN, Pisa, 1989).
Giazotto 1991	A. Giazotto et al, in "Gravitational Astronomy", pp 127, Ed D.E. McClelland & H.A Bachor, World Scientific (1991), Singapore.
Harman 1962	W.W. Harman, "Electrical and mechanical networks : an introduction to their analysis", New York : McGraw-Hill, 1962.
Den Hartog 1956	J.P. Den Hartog, "Mechanical Vibrations", McGraw-Hill, 3rd Ed., 1956.
Fan 1999	Fan Shuhua, Cai Yunxia, Wu Shuchao, Luo Jun, Hsu Houtse, *Response of a folded pendulum to tilt tides*, Phys. Lett. A **256** (1999) 132-40
Hafitz 1994	Aron Hafitz, MIT (private communication).
Haubrich 1969	R.A. Haubrich, K. McCamy, Rev. Geophys. **7** (1969) 539
Hello 1991	P. Hello, and J.Y. Vinet, *Modelling of the Virgo Interferometer*, in "Gravitational Astronomy", Ed D.E. McClelland & H-A Bachor, World Scientific (1991), Singapore.
Holloway 1996	L. Holloway, INFN, Pisa, Italy, private discussions.
Holloway 1997	L. Holloway, *The Superattenuator: a mirror suspension system for reducing seismic induced noise in the VIRGO gravitational wave interferometer*, Nucl. Phys. B, Proc. Suppl. **54B** (1997) 176.
Hough 1994	J. Hough *et al*., GEO600 Proposal (unpublished, 1994).
Ju 1994a	L. Ju, DG Blair, *Compound Pendulum Test Mass Systems for Laser Interferometer Gravitational Wave Detectors*, Meas. Sci .Technol. **5** (1994) 1053-1060.

Ju 1994b L. Ju, DG Blair, *Low Resonant Frequency Cantilever Spring Vibration Isolator for Gravitational Wave Detectors*, Rev. Sci. Inst., **65** (1994) 3482-3488.

Ju 1995 L. Ju, *Vibration isolation and test mass suspension systems for gravitational wave detectors*, PhD thesis, Dept of Physics, UWA, 1995.

Ju 2000a L. Ju, D.G Blair, J. Winterflood, *Long-term length stability and search for excess noise in multi-stage cantilever spring vibration isolators.*, Physics Letters A **266** (2000) 219-227.

Ju 2000b L. Ju, D.G. Blair, C. Zhao, *Detection of gravitational waves*, Reports on Progress in Physics, **63** (2000) 1317-427.

Ju 2000c L. Ju, M. Baker, D.G. Blair, *High-Q niobium membrane flexure pendulum*, Phys. Lett. A, **280** (2001) 182-184

Kanda 1994 N. Kanda, M.A. Barton, and K. Kuroda, *Transfer function of a crossed wire pendulum isolation system*, Rev. Sci. Instrum. **65** (1994) 3780-3.

Kawamura 1994 S. Kawamura, M.E. Zucker, *Mirror Orientation Noise in a Fabry-Perot Interferometer Gravitational Wave Detector*, Applied Optics. **33** (1994) 3912-18

Ke 1994 Ke W, Blair D.G, Notcutt M, Ju L, *Normal Mode Suppression in All Metal Cantilever Vibration Isolators*, Phys. Lett. A **197** (1994) 275-81

Kimball 1927 A.L. Kimball, D.E. Lovell, Phys. Rev. **30** (1927) 948

Kuroda 1997 K. Kuroda *et al.*,in *Proceedings of the 1st International Conference on GW: Sources and Detector*s, edited by F. Fidecaro and I. Ciufolini (World Scientific, Singapore, 1997).

Kuroda 1999 Kuroda K, *et al.*, *Large-scale cryogenic gravitational wave telescope*, Int. J. Mod. Phys. D **8** (1999) 557

Labs 1994 Glasgow, Munich, LIGO (MIT and Caltech), ISAS, U. Tokyo, Tokyo Nat. Observatory, VIRGO (Pisa, Orsay).

LaCoste 1934 L.J.B. LaCoste, Physics **5** (1934) 178.

LaCoste 1983 L. LaCoste, *LaCoste and Romberg straight-line gravity meter*, Geophysics, **48** (1983) 606-10

Lindenblad 1967	N.E. Lindenblad, *Little pendulums that oscillate like big ones*, Sci. Am. 216 (1967) 124-6.
Liu 1995	J. Liu, J. Winterflood, D.G. Blair, *Transfer function of an ultralow frequency vibration isolation system*, Rev. Sci. Inst. **66** (1995) 3216
Liu 1997	J. Liu, L. Ju and D.G. Blair, *Vibration isolation performance of an ultra-low frequency folded pendulum resonator*, Phys. Lett. A **228** (1997) 243.
Losurdo 1998	G. Losurdo, Ph.D. thesis, Scuola Normale Superiore, Pisa, 1998.
Losurdo 1999	G. Losurdo, *et al*. (incl. J. Winterflood), *An inverted pendulum preisolator stage for the VIRGO suspension system*, paper 6.2, Rev. Sci. Instrum. **70** (1999) 2507-15
Luiten 1997	A.N. Luiten, *et al*., *Ground tilt seismic spectrum measured with a new high sensitivity rotational accelerometer*, Rev. Sci. Instrum. **68** (1997) 1889
McClelland 1994	D. McClelland, D. G. Blair, and R. J. Sandeman, in *Proceedings of the 7th Marcel Grossmann Meeting on General Relativity, Part B, Stanford University, July 1994*, edited by R. T. Jantzen, G. M. Keiser, and R. Ruffini (World Scientific, Singapore, 1996), pp.1415–1417.
Meystre 1983	see e.g. articles on gravitational wave detection in: *Quantum Optics, Experimental Gravitation, and Measurement Theory*, eds. P. Meystre and M. O. Scully, NATO ASI Series B94, Plenum, New York, (1983)
Mio 2001	N. Mio, *Application of high-performance damping metal to gravitational wave detectors*, Paper presented at the 4th Edoardo Amaldi Conference on Gravitational Waves (Amaldi4), Perth, July8-14 2001.
MPQ 1989	Proposal for a Joint German-British Interferometric Gravitational Wave Detector, September 1989 (MPQ report 147).
Nelson 1998	P.G. Nelson, Technical Manufacturing Corporation, *Suspended low-frequency horizontal pendulum isolator for vibration isolation systems*, US Patent No. 682855.

Notcutt 1994 M. Notcutt et al, UWA Report titled: *A Test System for Local Control Actuators in Laser Interferometer Gravitational Wave Detectors*. (1994).

Pinoli 1993 M. Pinoli, D.G. Blair and L. Ju, *Tests on a low-frequency inverted pendulum system*, Meas. Sci. Technol. **4** (1993) 995-9.

Plissi 1998 Plissi-MV; Strain-KA; Torrie-CI; Robertson-NA; Killbourn-S; Rowan-S; Twyford-SM; Ward-H; Skeldon-KD; Hough-J, *Aspects of the suspension system for GEO 600*, Rev. Sci. Instrum. **69** (1998) 3055-61.

Press 1966 F. Press, in Handbook of Physical Constants, rev. Ed., Memoir 97, Geological Soc. of America, 1966.

Puksand 1975 H. Puksand, *Optimal conditions for dynamic vibration absorbers for variable speed systems with rotating or reciprocating unbalance*, Int. J. Mech. Eng. Educ. **3** (1975) 145-152.

Quinn 1992 T.J. Quinn; C.C. Speake and L.M. Brown, *Materials problems in the construction of long-period pendulums*, Philosophical Magazine A, **65** (1992) 261-76

Quinn 1995 T.J. Quinn, C.C. Speake, R.S. Davis and W. Tew, *Stress-dependent damping in Cu-Be torsion and flexure suspensions at stresses up to 1.1 GPa*, Phys. Lett. A **197** (1995) 197.

Reddick 1947 H.W. Reddick, F.H. Miller, "Advanced mathematics for engineers", 2nd Ed. Wiley, New York, 1947, Art. 33.

Rüdiger 1994 Computer program written by A. Rüdiger, communicated by R. Schilling, Max-Planck-Institut für Quantenoptik.

Saulson 1994 P.R. Saulson, R.T. Stebbins, F.D. Dumont and S.E. Mock, *The inverted pendulum as a probe of anelasticity*, Rev. Sci. Instrum. 65 (1994), 182-91.

Saulson 1997a P.R. Saulson, *If light waves are stretched by gravitational waves, how can we use light as a ruler to detect gravitational waves?*, Am. J. Phys. **65**, (1997) 501-5

Saulson 1997b P.R. Saulson, *How an interferometer extracts and amplifies power from a gravitational wave*, Class. Quantum Grav. **14**, (1997) 2435-54

Saxey 1995 D. Saxey, *Strategies for optimising damping of vibration isolators for a laser interferometer gravitational wave detector*, BSc. Honours thesis, Dept of Physics, UWA, 1995.

Shoemaker 1988 D. Shoemaker, R. Schilling, L. Schnupp, W. Winkler, K. Maischberger, A. Rüdiger, *Noise behaviour of the Garching 30-meter prototype gravitational wave detector*, Phys. Rev. D, **38**, 423-432 (1988).

Speake 1990 C.C. Speake, D.B. Newell, *The design and application of a novel high-frequency tiltmeter*, Rev. Sci. Instrum. **61** (1990) 1500

Stebbins 1994 R. Stebbins *et al.*, *Low-frequency, active vibration isolation system*, Proc. SPIE 2264 (1994) 27-37.

Stephens 1991 M. Stephens, P. Saulson, and J. Kovalik, *A double pendulum vibration isolation system for a laser interferometric gravitational wave antenna*, Rev. Sci. Instrum. **62** (1991) 924-932

Thorne 1987 K.S. Thorne in *300 Years of Gravitation*, Eds S. W. Hawking and W. Israel, Cambridge University Press (1987).

Tsubono 1991 K. Tsubono *et al.*, *Gravitational Astronomy*, pp 75, Ed D.E. McClelland & H-A Bachor, World Scientific (1991), Singapore.

Tsubono 1993 K. Tsubono, A. Araya, K. Kawabe, S. Moriwaki, and N. Mio, *Triple-pendulum vibration isolation system for a laser interferometer*, Rev. Sci. Instrum. **64**, 2237-2240 (1993)

Usher 1992 M.J. Usher, A.R. Mat Isa, *An instrument for discriminating between acceleration and ground tilt in horizontal component seismometers*, Meas. Sci. Technol. **3** (1992) 574

Veitch 1993 P. J. Veitch, N.A. Robertson, C.A. Cantley and J. Hough, *Active control of a balanced two-stage pendulum vibration isolation system and its application to laser interferometric gravity wave detectors*, Rev. Sci. Instrum., **64** (1993) 1330

Vogt 1989 R.E. Vogt *et al.*, *LIGO Proposal*, (Caltech, Pasadena, 1989).

Wahl 1963 A.M. Wahl, "Mechanical springs" (McGraw-Hill, New York, 1963) 2nd Ed. p136.

Wolfram 1988 *Mathematica* is a registered trade mark of Wolfram Research Inc. (http://www.wolfram.com).

Winterflood 1995 Winterflood J, Blair D.G, Schilling R, Notcutt M, *A Position Control System for Suspended Masses in Laser Interferometer Gravitational Wave Detectors*, paper 9.1, Rev. Sci. Inst. **66** p2763-76 (1995)

Winterflood 1996 J. Winterflood, D.G. Blair, *A long-period conical pendulum for vibration isolation*, paper 4.2, Phys. Lett. A, **222** (1996) 141

Winterflood 1998 J. Winterflood, D.G. Blair, *A long-period vertical vibration isolator for gravitational wave detection*, paper 5.2, Phys. Lett. A, **243** (1998) 1

Winterflood 1999 J. Winterflood, G. Losurdo, D.G. Blair, *Initial results from a long-period conical pendulum vibration isolator with application for gravitational wave detection*, paper 4.3, Phys. Lett. A, **263** (1999) 9

Winterflood 2000a J. Winterflood, Z.B. Zhou, D.G. Blair, *Ultra-low Residual Motion Suspension System*, in Proceedings of the 2nd TAMA International Workshop on Gravitational Wave Detection held at National Olympics Memorial Youth Center, Tokyo, Japan on October 19-22, 1999, Gravitational Wave Detection II, eds. S. Kawamura and N. Mio (Universal Academy Press, Tokyo, 2000), 301-310

Winterflood 2000b J. Winterflood, Z.B. Zhou, L. Ju, D.G. Blair, *Tilt Suppression for ultra-low residual motion vibration isolation in gravitational wave detection*, paper 7.3, Physics Letters A **277** (2000) 143-155.

Winterflood 2001a J. Winterflood, D.G. Blair, *High performance vibration isolation using springs in Euler column buckling mode*, paper 2.2, Phys. Lett. A., accepted.

Winterflood 2001b J. Winterflood, T.A. Barber, D.G. Blair, *Mathematical analysis of an Euler spring vibration isolator*, paper 2.3, Phys. Lett. A., accepted.

Winterflood 2001c J. Winterflood, T. Barber, D.G. Blair, *Using Euler Buckling Springs for Vibration Isolation*, Class. Quantum Grav., accepted.

Zhou 2000 Z.B. Zhou, J. Winterflood, L. Ju, D.G. Blair, *Investigation of a laser walk-off angle sensor and its application to tilt measurement in gravitational wave detection*, Physics Letters A **280** (2001) 197-203.

9. Appendices

9.1 Local Control System Paper

A Position Control System for Suspended Masses in Laser Interferometer Gravitational Wave Detectors

J. Winterflood, D. G. Blair, R. Schilling* and M. Notcutt

Department of Physics, University of Western Australia, Crawley, Australia.

**Max-Planck-Institut für Quantenoptik, Garching, Germany.*

We review and present an analysis of the major issues in the design of test mass control systems for laser interferometer gravitational wave detectors. Based on this analysis, we present a design for a computer controlled modular analog servo control system which is well suited to the control of a large number of degrees of freedom in long baseline instruments. The system has been tested on an interferometer using multistage cantilever spring isolators. The system enables simple monitoring, testing, and display of many channels simultaneously, while retaining the advantages of analog PID control electronics.

9.1.1 Introduction

A laser interferometer gravitational wave detector (LIGWD) consists primarily of a very large Michelson interferometer with a baseline of several kilometres. The sensitivity of the basic interferometer must be increased by a multi-reflection scheme - normally by forming optical cavities in the arms [Meystre 1983, Blair 1991]. Massive, low-loss, vibration isolated mirrors are the test masses. For operation of the interferometer these masses, and other optical components, must have their positions controlled to within a very small fraction of a wavelength. The displacement sensitivity goal is $\sim 10^{-20}$ m/√Hz in the audio frequency band from tens of hertz to several kilohertz [Hough 1994, Vogt 1989, Giazotto 1989]. At this sensitivity, gravitational wave signals from neutron star coalescence and supernovae are expected to be detectable [Thorne 1987].

The test masses in a LIGWD are suspended from a multi-stage vibration isolator which isolates each test mass from seismic noise and allows approximately free motion above the isolator cut-off frequency. To achieve operation of the interferometer, a feedback system must control the test mass position to high precision. Firstly a local control system is necessary to suppress large scale motions and align the mirrors to the point where interference fringes may be observed. Secondly a global control system is required to control the interferometer arm lengths by locking onto an interference fringe. These two systems are often intimately connected and may share signals, electronics, and transducers.

Control systems have been developed by most of the laboratories developing technology for LIGWDs [Labs 1994]. Unfortunately design details have not generally been published. There is a wide choice of technology available and a wide range of control schemes. For example, the servo systems may be entirely analog, or entirely digital, or as we present here, something in-between. A variety of transducers may be used for both sensing and force feedback. Feedback forces may be applied with reference to the seismically grounded frame which supports the isolator structure, or via suspended and vibration isolated reaction masses. Motion sensing may also be referenced to the frame, or it may be inertial (using accelerometers), or it may also be referenced to a suspended mass. Ultimately of course, the motion of the test masses along the beam directions are referenced to each other via the laser interferometer.

The purpose of this paper is twofold. One is to present a flexible and modular electronic solution to the design of a local control system suitable for large scale observatories. The second is to present an analysis and overview of the local control system design problem in general which we believe has been inadequately addressed in the literature. In section 9.1.2 of this paper we analyse the control system requirements for a test mass suspended by a typical high-Q multistage isolator in the presence of seismic noise. We consider approaches to overcome the problem of noise injection by the control system, and estimate the level of electronic noise that such a system injects. We then consider the scale of the servo system required, which is set by the total number of degrees of freedom needed to be controlled for various interferometer configurations. Based on this analysis, we go on in section 9.1.3 to describe the design and implementation of a multi-channel, computer interfaced, analog control system. The system is very modular and extendable and allows remote operation from a computer terminal.

9.1.2 Principal Design Issues

One of the main issues when considering the design of a vibration isolator control system is the lack of an ideal reference from which both to sense and to apply forces. In order to remove signal band seismic noise from the reference, another isolator may be provided to serve as a filtered reference and/or as a reaction mass. A reaction mass design has been reported by Veitch et al [Veitch 1993]. A nested design which nulls the error signal between the test mass and reaction mass by servoing the reaction mass has been reported by Stephens et al [Stephens 1991]. Here we consider the simpler approach of using a frame attached to the local seismically noisy earth, as our reference, and attempting to avoid noise injection by electronic filtering and position independent actuators.

9.1.2.1 Isolation Stack Control Requirements

Isolation stacks consist of cascaded mechanical low-pass filter stages. Traditionally these have been made with alternating masses and damped springs usually made of rubber. However this inherent damping introduces thermal noise, requires twice the number of stages, and rubber has poor vacuum performance and poor dimensional stability. If alternative low-loss elastic elements are used, low frequency seismic vibrations can be resonantly enhanced by the normal modes of the stack. Some means of Q-factor reduction or damping of these resonances is required in this case. Various stand-alone passive or active schemes have been used [Billing 1979, Blair 1993, Tsubono 1993], but here we achieve this damping as an integrated function of the local control system.

Low-loss isolator stages may be readily realised with masses and springs for vertical isolation, and pendulums for horizontal isolation. The resonant frequency of a single stage is determined by its pendulum length for horizontal motion or by its spring constant (i.e. extension length) for vertical motion. If these lengths are equal, then the resonant frequencies are also equal. Similarly if stages are cascaded with ideal springs (zero unextended length or their mechanical equivalent for small displacements) then the vertical transfer function is the same as the horizontal transfer function for the

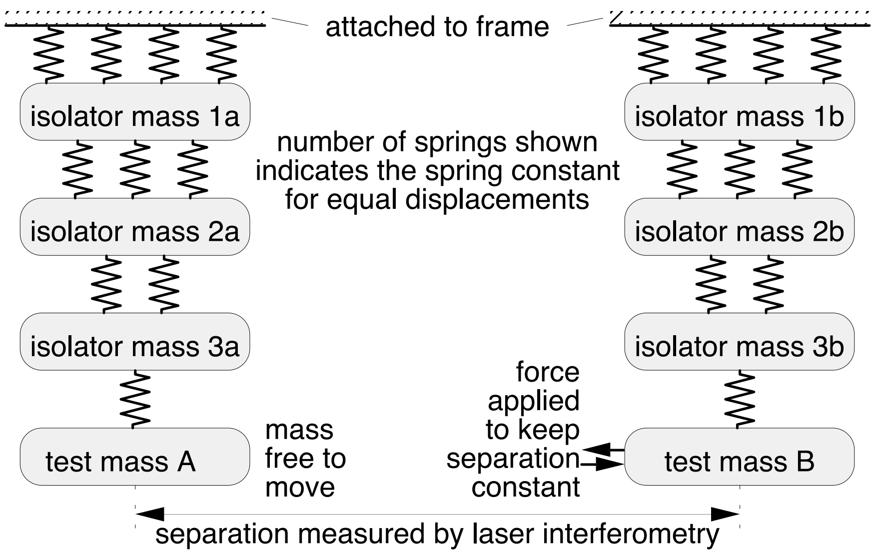

Figure 9.1 Test-mass isolator stacks. The distance between the test masses is held constant to within a small fraction of the laser wavelength by forces applied to one mass by a servo loop. They behave as a single mass joined together by a rigid rod with the servo force being a measure of the tension or compression in the rigid rod.

overall stack. While multi-stage vibration isolators may be constructed using stages with widely differing resonant frequencies, it is sufficient to consider stages of equal mass and deflection under gravity. The equal deflection arrangement is the natural solution to use if the main constraint is the yield strength of the elastic elements.

For interferometry, the separation between a pair of test masses must be fringe locked to within a very small fraction of a wavelength. This is achieved as shown in figure 9.1 by applying forces to one of the test masses. In order to determine the forces required, it is useful to consider three frequency bands: (a) near DC, (b) the normal mode band (e.g. 0.1 to 10 Hz), and (c) the signal detection band. Within the signal detection band all external forces acting on the test mass should be so low that no servo control is required. In the near DC band, forces are required to counteract thermal drifts, earth tides, and low frequency seismic tilt noise. However these forces will be small because of the softness of the isolation suspension. In the normal mode band the resonances of the stack will couple seismic noise between the frame (which supports the isolator) and the test masses. This coupling is directly related to the Q-factors of the normal modes and for an all metal isolator, this coupling is quite strong and determines one of the main parameters of the control system design - the maximum force required to maintain position lock.

To estimate this force we consider the transfer function for differential motion

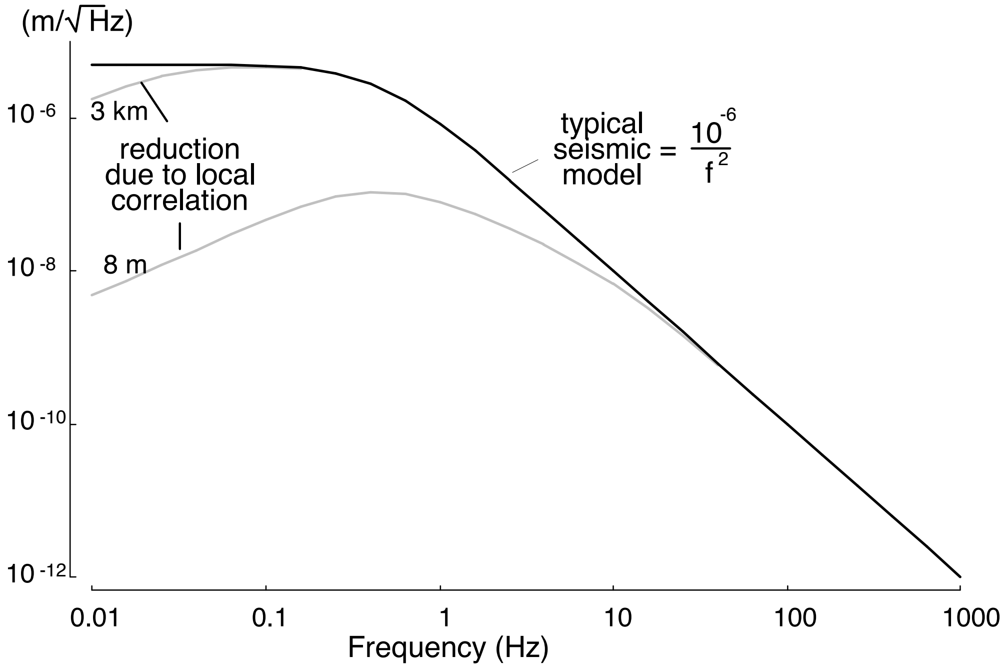

Figure 9.2 Seismic noise model. Seismic noise may be conservatively modelled as $10^{-6}/f^2$ m/√Hz. For small scale prototypes local correlation reduces the differential motion. This effect will not be noticeable for large scale instruments.

between two isolator stacks a distance d apart. Assuming that the seismic noise is in the form of a random field of shear waves (or Love / Rayleigh waves) with velocity v, we may estimate the reduction in the differential seismic motion due to its being correlated for nearby locations. For wavelengths which are large compared to d, the differential displacement is reduced by a factor $2\pi(d/\text{wavelength})$ and this reduction effect disappears for frequencies above which d represents about half a wavelength. Assuming a seismic noise spectrum given by $10^{-6}/f^2$ m/√Hz (f in Hz) which flattens off below the microseismic peak around 0.2 Hz, we obtain the relative seismic noise spectra shown in figure 9.2. For a laboratory scale instrument of ~10 m, there is a break point in the seismic spectrum around 10 Hz. Below this frequency, correlation creates a substantial reduction in the observed relative motion. However for a full scale instrument this break point occurs around 0.03 Hz (these frequencies are for a seismic wave velocity of 500 m/s in wet sand [Press 1966]). Thus large scale instruments place additional demands on the control system at low frequencies.

The suspension system in figure 9.1 may be analysed as a 10 element second order system (consisting of 2 frame mounts and 8 masses). The masses, damping coefficients, and spring constants are defined by sparse matrices of size 10×10. With the servo loop closed, the position of the two test masses are equated, and a transfer function may be found as the ratio between the motion of one of the frame mounts (the other being kept

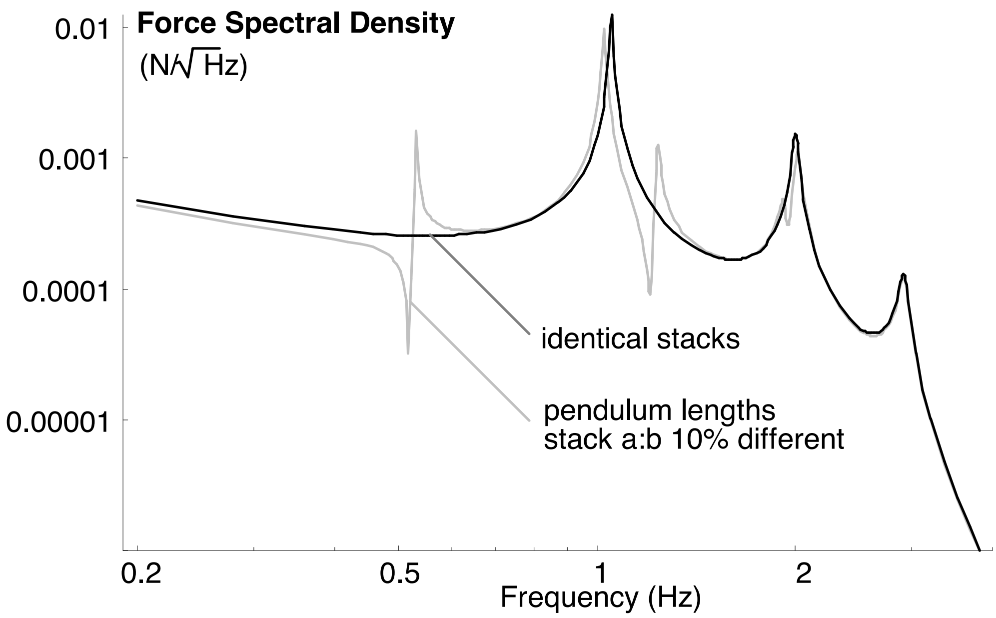

Figure 9.3 Spectrum of force required to lock a pair of isolators together. Masses are 100 kg, spring / pendulum length = 25 cm, Q = 100. The force value was obtained by assuming seismic motion of figure 9.2 applied to both stacks with the reduction factor for locality correlation. Additional modes become visible when the pendulum lengths of the two stacks differ by a few percent.

stationary), and the force on the test masses. Applying the seismic model in figure 9.2 to a system with 100 kg masses, 25 cm pendulums, and Q-factors of 100, gives the force spectrum shown in figure 9.3.

It may be seen from figure 9.3 that the required force peaks at frequencies corresponding to only three of the normal modes of the system. The other four modes do not appear as they are asymmetric modes which do not produce differential motion of the test masses. If the model stack parameters are deliberately unbalanced by say 10%, then as shown in figure 9.3, the asymmetric modes do appear. However the force required to counteract the imbalance is still small in comparison to that required to counteract the differential modes. The RMS force requirement is found by integrating the spectral density and is usually dominated by the largest peak. (If this is plotted as $(\text{force})^2$ against frequency on linear scales the dominant peak becomes very obvious). Provided there is a dominant peak with a Q-factor greater than 10 or so, the RMS force may be estimated from the formula $a_o\sqrt{(\pi f_0 Q / 2)}$, where a_o is the value at the base of the peak (highest side dominates), f_0 is the frequency of the peak in Hz, and Q is given very approximately by the ratio of peak value to base value. In figure 9.3 the base of the peak may be estimated as 3×10^{-4} N/√Hz, the peak as 1.5×10^{-2} N/√Hz, and f_0 as 1 Hz. These figures yield an RMS force value of approx 0.0026 N or about 1/4 g weight. The RMS value so obtained is the long term mean and the instantaneous force required will vary randomly above and below this value.

Once the differential positions of the test masses are locked with respect to each other, they are still free to move with common mode motion (all in unison). In a delay line interferometer the system could in principle also sustain a "breathing" mode where both end mirrors move relative to the beam splitter. Although these motions need not affect the signal detection operation, the control force must end up reacting against the frame in order to maintain position lock. Even a reaction mass must be damped to the frame at its normal mode frequencies. This means that the damping force transducer arrangement must be capable of coping with the full range of these motions. For simple frame mounted force transducers, this motion must be kept to a minimum because of the practical difficulties of making a position independent force transducer with a large range (as discussed further in section 9.1.2.4). If high Q isolation stages are used, then considerable enhancement of seismic motion results at the normal mode frequencies of the isolation stack. Figure 9.4 shows that, for Q=100, the lowest normal mode is enhanced by more than a factor of 100. So a secondary task of the local control system

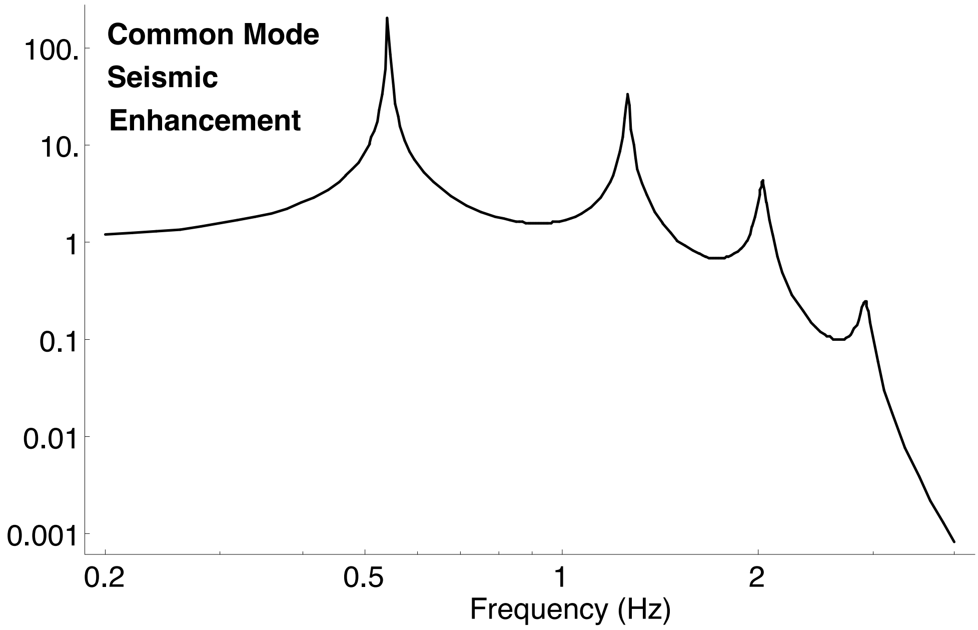

Figure 9.4 Common mode seismic enhancement of a 4-stage isolator. Masses are 100 kg, spring / pendulum lengths = 25 cm, Q = 100. Strongly damping the test masses with respect to the frame around the normal mode frequencies shown here brings the common mode motion back down toward the seismic level. This facilitates initial locking and minimises the position independent range required of the force transducers.

is to reduce these common mode motions as much as possible to minimise the position independent range required by the force transducer. Since a lower frequency reaction mass cannot be made available (if it was we should use it as the main isolator), it is not possible to do more than strongly damp this common mode motion with respect to the frame. The forces required to damp this common mode motion may be provided by the same force transducers that lock the test masses together. (Damping a test mass clearly requires less force than locking its position to a reference under conditions of similar disturbance and transfer function). The resulting common mode motion is of a similar magnitude to seismic noise and is continuously present during normal operation. This means that the force transducers of our example must be capable of continuously providing position independent control forces of the order of 10^{-2} N in the presence of seismic motion.

In order to be able to damp the common mode motion (with respect to the frame), the motion must be measured by sensors mounted on the frame, or inertially by an accelerometer. So here we have the main input and output requirements of the system for position control :-

a) Differential position from laser interference (DC through signal frequencies).

b) Frame referenced or inertially referenced position (around stack normal mode

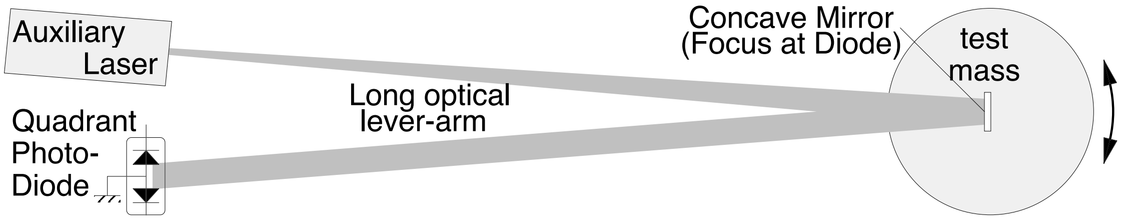

Figure 9.5 Mirror orientation measurement. Orientation may be measured by shining an auxiliary laser beam on a special mirror which focuses it back to a quadrant photo-diode located near the laser. This prevents laser beam pointing noise from appearing as rotation. The seismic motion of the laser / detector appears as the beam moving with twice the amplitude across the detector. With a long-lever-arm this is minimised.

frequencies).

c) Force output to be applied with respect to the frame.

9.1.2.2 Mirror Orientation Control

The test-mass mirrors need their orientation to be controlled to high precision, particularly if optical cavities are being formed in the arms. An optical lever such as that shown in figure 9.5, can use the full length of the interferometer for orientation sensing. If the laser and photo-diode are mounted on the frame, then only a very small noise signal will be produced by their transverse seismic motion. The magnitude of this orientation noise is 2 × (seismic noise / lever length) and is therefore of the order of $10^{-9}/f^2$ rad/√Hz for a full size instrument. This is around two orders of magnitude smaller than the orientational positioning accuracy required for a Fabry Perot cavity [Hello 1991] (tilt $\theta < 0.2\ \mu$rad), but perhaps only one order smaller than that required in a power recycling configuration [Hafitz 1994]. High frequency noise injected in the signal sensing axis as a result of low frequency misalignment noise has been thoroughly investigated by the LIGO team [Kawamura 1994].

A much more significant source of orientation noise could be the beam alignment jitter of laser sources. In small HeNe lasers we found alignment variations of ±10 μrad over periods of seconds. Diode lasers are much more stable. The worst instability in a 670 nm laser diode occurred on millisecond time scales with amplitudes of ±0.5 μrad. The angular seismic vibration of the laser mounting can also be a significant noise source. However both intrinsic and seismic alignment noise may be overcome by using a concave mirror on the test mass which focuses the laser light back at the detector. If a focal length of half the separation is used, then an image of the laser source is formed at the detector which is insensitive to beam alignment noise of the laser source. At the same time the position of the image is fully sensitive to the orientation of the mirror itself.

Orientation control adds another input to the local control system. It must have negligible effect at frequencies within the signal band but must be in full operation at any torsional resonant frequencies in the isolation stack and on down to DC in order to maintain alignment.

9.1.2.3 Avoiding Seismic Noise Injection

Seismic noise can be injected into the servo loop both in sensing and in force actuation. If motion sensing (for damping) and orientation sensing are done with respect to the frame (which vibrates with seismic motion), it is impossible to avoid seismic noise in the sensing. This sensing noise has components in the signal band which must be prevented from driving the test mass and appearing in the signal output. This is readily achieved by electronic filtering of the signal.

The amount of filtering required for the normal mode damping may be estimated as follows: Supposing the motion sensing is designed to handle seismic motion enhanced by a normal mode Q-factor of 100. From figure 9.4 this will typically occur around 0.5 Hz. Supposing also that full immunity from seismic noise is required above 100 Hz. At this frequency the amplitude of the seismic noise will be less than 10^{-6} of the normal mode sense signal (100 from Q-factor, $4{\times}10^4$ from seismic f^{-2} dependency). A 4-pole low-pass filter with a cut-off of 10 Hz will reduce this further by a factor of 10^{-4} which brings it below the electronic noise floor (10^{10} is the best signal to noise ratio achievable with standard operational amplifiers:- 10 V signal / 1 nV noise). Since damping requires the derivative of the signal to be taken which adds a +1 slope to the spectrum (in log-log representation), an additional pole is required by the filter. In any case an odd ordered filter is preferred as this allows the last pole to be constructed as a passive R-C stage which will filter out noise contributed by the op-amps of the preceding stages.

The amount of filtering required for the orientation control may be estimated in a similar manner to the normal mode damping. The main differences are that the normal mode frequencies are typically lower and full proportional-integral-derivative (PID) control is required rather than simple derivative only (damping) control. Both of these differences tend to ease the filtering demand. It is not unreasonable to apply the same 5-pole 10 Hz filter to both normal mode damping and orientation control. Analysis shows that a 5-pole 10 Hz filter applied to a derivative damped resonant mode, still leaves enough phase margin to be stable for normal mode frequencies up to about 3 Hz. However clearly the choice of filter cut-off frequency is dependent on the isolator characteristics used.

9.1.2.4 Position Independent Force Transducer

The second area of potential noise injection is the actuator force path. If an isolated reaction mass is used as a reference for force application as mentioned in section 9.1.2, then the problem is largely avoided. We consider here the simpler approach of applying the force directly with respect to the seismically affected frame. In this case the discussion must centre around the design of a force transducer which can apply a constant force independent of seismic motion in its mounting.

Magnetic force transducers consisting of a frame mounted coil and a permanent magnet are the most natural, and most widely used choice for LIGWD control systems [Shoemaker 1988]. The translational force acting on a point magnetic dipole m in a non-uniform field B is given by $\nabla(m \cdot B)$. This indicates that provided the field gradient $\partial B/\partial x$ in the direction of the required force is constant, and the dipole of the permanent magnet m is parallel to this force (so that the transverse gradients which are bound to exist have no effect), then the force developed between the two is constant regardless of position change (which in this case arises from seismic motion).

This situation is obtained to first order by positioning the coil such that the dipole of the permanent magnet (attached to the test mass) is in line with the coil's axis and at a position of $r/2$ from the coil where r is the radius of the coil. At this position $\partial^2 B/\partial x^2$ vanishes and the force is approximately constant for small displacements. This simple approximation can be considerably improved by placing additional coils to linearise the field [Giazotto 1991], or separated magnetic poles with differing fields at each pole to linearise the interaction [Rüdiger 1994]. The field gradients due to a large and small coil with opposing currents are shown in figure 9.6a. If the smaller coil has half the radius r of the large one and $1/16^{th}$ of the turns, and is wired to oppose the field of the large coil, then the second derivative of the gradient becomes zero at $r/2$ from the large coil. This combined gradient is flatter than that of a single coil and resembles a constant value for a wider displacement range. The field gradients due to two matched coils with opposing currents are shown in figure 9.6b. If the coils are positioned $\sqrt{3}\ r$ apart forming a Maxwell pair, then the second derivative of the gradient is again zero at the centre of the coils. This provides an exceptionally constant gradient if the topography is not a problem. Expanding the three arrangements in figure 9.6c, over a typical operating area, gives some indication of how superior the Maxwell pair is in providing a constant force for small displacements.

We have tested the effectiveness of a large/small coil scheme [Notcutt 1994]. Our conclusion which also applies to a single coil is that a point dipole is not an accurate enough approximation for practical sized magnets. The calculations need to take into account the volume of the magnetic material and integrate over the volume. A consideration which has not yet been fully addressed is how accurately the permanent

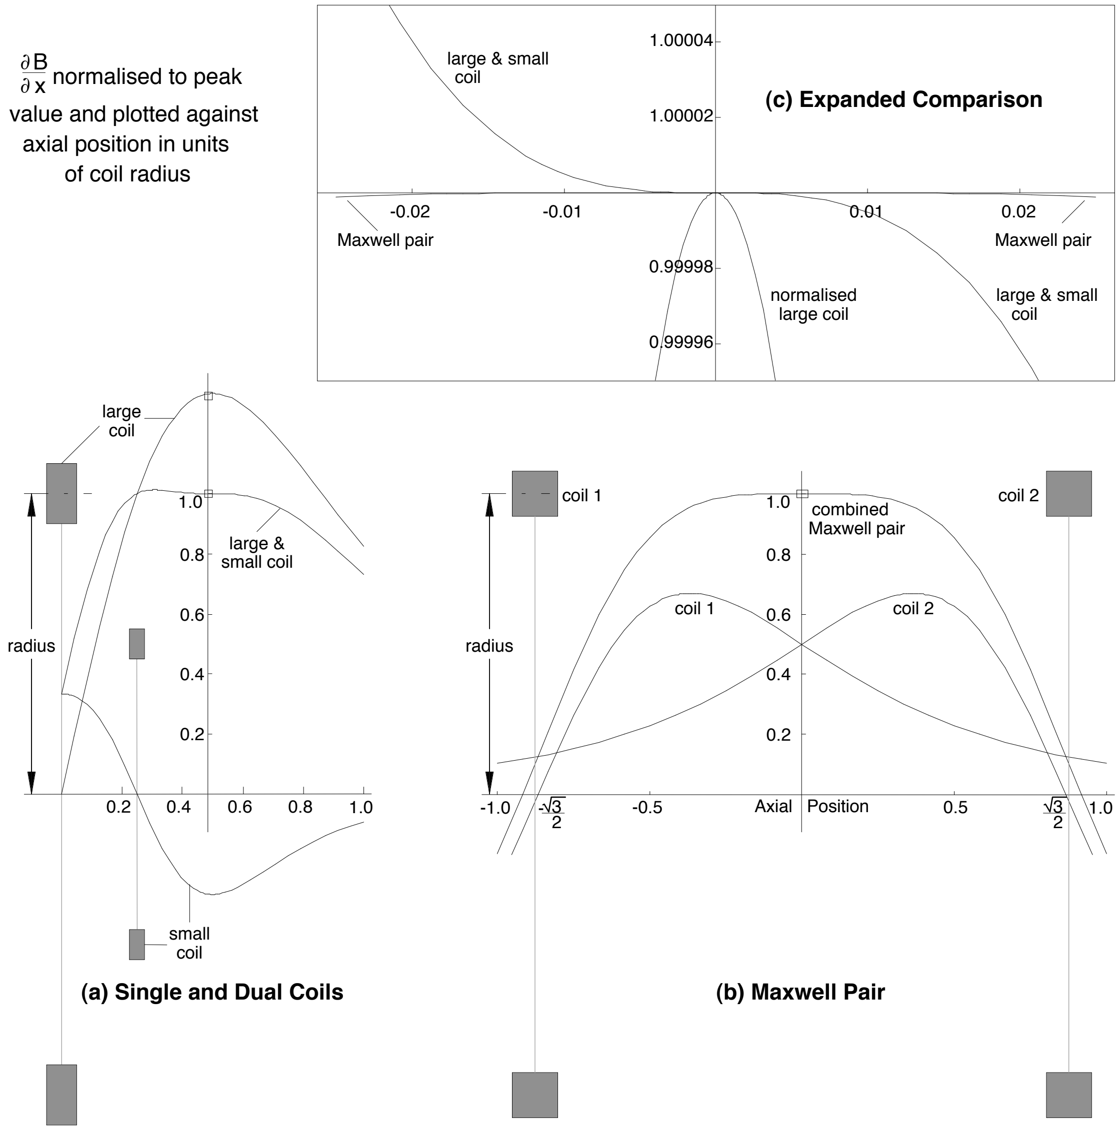

Figure 9.6 Field gradient of various coil arrangements. The force acting on a magnetic dipole is proportional to the field gradient and this is plotted for axial positions of various coil arrangements. Positioning the dipole at a maximum allows this force to be position independent to first order. If a second coil is added, position independence to second order may be obtained. Expanding the three arrangements over a typical operating area gives some indication of relative performance in providing a constant force for small displacements.

magnet needs to be aligned with the coil axis in order to avoid noise injection from transverse motion in the non-constant transverse gradient, and how to achieve this alignment accuracy.

With these coil-magnet arrangements, the current in the coil must be held proportional to the required force - i.e. kept constant for a constant force. This means that it must be driven with a *constant current* type of drive circuit providing as high a dynamic source impedance as possible. If this is not done then the seismic motion between the magnet and coil will produce an EMF in the coil which will result in a current change through the drive impedance, and this will react against the magnet in the same manner as an induced eddy current in a conductor. Even the stray capacitance in a few metres of cable connecting the coil to the current source driver is sufficient to allow seismic induced current to produce an unacceptable acceleration coupling between the coil mounting and the magnet. To avoid this the coil current driver should be located close to the drive coil.

Consideration must also be given to other modes of coupling between the seismically affected frame and the test mass in the presence of the magnetic fields due to this type of force transducer. There must be no seismically affected ferromagnetic parts anywhere near to magnets mounted on the test mass (or nearby stages). Likewise there must be no seismically affected conducting surfaces sufficiently close that induced eddy currents produce unacceptable velocity coupling. The same is true of the field produced by the coil. Thus any ferromagnetic or conducting surfaces on the test mass must be considered as possible routes to noise input. Ideally even the permanent magnets themselves should be made from non-conducting material such as ferrite. Their ferromagnetic property however is intentional as they are located by design in a position where the field gradient from the coil is constant. The effect of this type of ferromagnetic or eddy current coupling is difficult to estimate in most cases and we have not attempted to do so at this stage.

9.1.2.5 Unavoidable Noise Injection Level

Assuming the seismic affected inputs (damping and orientation) have been sufficiently filtered, and the force transducer design provides adequate position independence, the primary remaining noise source due to the local control system is the electronic noise in the final stages of the coil current driver. This is due to the fact mentioned in section 9.1.2.3 - that the best dynamic range (or signal to noise ratio) of electronic amplifiers is about 10^{10}. Thus if a peak force of say 10^{-2} N is required to counteract seismic motion, then the same hardware will be generating random forces of the order of 10^{-12} N due to the electronic noise in the drive amplifier. For an active current source output this seems unavoidable. It can not be passively filtered as any filtering components would tend to spoil the very high impedance drive required to avoid eddy current type seismic coupling. The options are therefore either (a) to ensure the peak force requirement is at a level where a factor of 10^{-10} takes the displacement noise produced in the test mass below the signal detection sensitivity required, or (b) to achieve the same effect by providing another stage of isolation between the mass affected by this noise (which now becomes the control stage), and the test mass proper [MPQ 1989].

We may estimate the effect of this noise force from the analysis in section 9.1.2.1 which yielded an RMS force requirement of 2.6 mN. Since this force is only the mean of a random distribution, our system must be capable of values greater than this to ensure that it is rarely exceeded. Choosing a peak force capability of say 10 mN means that the electronic hardware required to generate this maximum force will also generate a white noise force of 10^{-12} N. This will result in displacement noise of the lowest mass of $10^{-12}/(4\pi^2 f^2 m)$ which for m=100 kg, and f=100 Hz, gives a value of 2.5×10^{-20} m/√Hz. This is on the borderline of acceptability for gravity wave detector target sensitivities. Thus there is strong motivation for implementing at least one of the above options. In section 9.1.3.3 we show a compound pendulum which separates the test mass from the control stage. We are also researching a low frequency pre-isolator [Blair 1994] (of order 10^{-1} Hz to 10^{-2} Hz) which would allow a return to direct actuation on the test mass (as long as there were no additional difficulties in attaching a small magnet to it).

9.1.2.6 Servo Control Loops Required

The interferometer control system can be clearly broken down into three main types of control: (1) interference fringe locked loops, (2) orientation control loops, and (3) normal mode damping loops. In this section we asses the total number of loops required in each category, to estimate the magnitude of the total control system required.

1) Interference Fringe Locked Loops

These loops lock the relative separation between suspended components by laser interference and operate from DC up to the signal detection band. In a basic interferometer there would only be one of these control loops. The separation between the central mass (beam splitter) and the mirror at the end of one arm is taken as the reference length and the other arm is locked to it. A multi-reflection cavity system (for increased sensitivity) requires two more global control loops - one for each extra mirror, and power recycling (forming a cavity on the incoming laser port of the beam splitter) requires a further one. Depending on their suspension arrangements, each of these optical elements may need another set of local controls.

2) Orientation Control Loops

The orientation control loops use the long optical lever for sensing and must operate from DC to beyond the suspension system torsional modes (which may be much lower than its translational modes). Seven orientation control loops are required for a basic suspended Michelson interferometer. Each end test mass has two orientational degrees of freedom (pan and tilt), while the central mass with beam splitter has an additional degree of freedom associated with the aiming of the incoming laser beam (assuming optic fibre fed and terminated on central mass). If Fabry-Perot cavities are used in the arms another four control loops are required (two for each extra mirror). However power recycling only requires one additional loop, since the laser aiming control is then moved to the recycling mirror.

3) Normal Mode Damping Loops

The isolation stack normal mode damping uses frame referenced position sensing (or possibly inertial sensing) to provide damping over the decade or so of normal mode resonances. The derivative of the position (i.e. velocity) is required for damping, so these loops are AC coupled but need to be rolled off steeply above the normal modes with the filtering mentioned in section 9.1.2.3. Since each suspended mass needs to have its motion damped in five or six degrees of freedom, a full scale interferometer needs a very large number of these damping loops. Note that until damping has been achieved, it is generally impossible to implement the loops discussed above. For instance the laser beam for the long-lever-arm orientation sensing will not even fall on the quadrant photo-diode until the orientation is sufficiently stable, while fringe locking requires unacceptable loop gain if the test mass velocity is too large. So damping is required in duplication to initially lock these loops, although it may be turned off once fringe locking is achieved.

The table below shows the number of loops required for the various interferometer configurations assuming a maximum of six damping loops for each suspended mass, and allowing one degree of freedom (rotation about the laser axis) to be neglected for mirror masses except for beam aiming with the beam splitter. Some isolator designs may not allow adequate coupling to all of the normal modes at the physical location of the control mass. In this case additional damping loops may be required at an anti-node location on another isolator mass, or alternatively a more sophisticated transfer function implemented instead of simple velocity damping.

Type of System	Simple Interferometer	Fabry-Perot	Power Recycling
Number of Masses	3	5	6
Fringe Locked Loops	1	3	4
Orientation Loops	7	11	12
Normal Mode Damping	16-18	26-30	31-36

Table 9.1 This shows the number of control loops required for some typical experimental interferometer configurations.

9.1.2.7 Computerisation Requirements

A full scale gravity wave interferometer extends over several kilometres, and all the mirror local controls are intimately involved in obtaining cavity resonance and interference. For this reason it is essential that all positioning signals, set-point adjustments, etc, are available at all the useful locations for setting up, trouble-shooting, experimenting, etc. This is most easily done by interfacing the local control system to a computer and using standard computer networking and screen displays for this remote monitoring and adjustment. It is also necessary to bring the control loops into operation in a sequence :- First damping must be implemented. Then orientation control must be achieved. Only at this stage can the interference fringe locked loops be implemented. At this point some of the damping loops may be opened. A computer allows this to be automatically sequenced.

Due to the complexity and precision of the complete instrument, it is also very useful to be able to monitor the operation of each sub-system. A computer with access to all the control loop information, allows the performance of all the sub-systems to be monitored, analysed, and data logged. It also allows an alarm to be raised when any sub-system performance degrades from optimum. Thus we believe that extensive computerisation is essential. This could involve a fully digital control system, as under development by the VIRGO group. The alternative, that we present below consists of digitally controlled analog loops.

9.1.3 Control System Implementation

The control system design is necessarily quite dependent on the suspension arrangement, on the choice of frame or inertial referencing, and on the type of sensors and actuators. The implementation described here is applied to all-metal cantilever-spring isolators which have relatively high Q normal modes. Most servo loops are used for damping. We used frame referenced optical shadow position sensors with integrated magnetic actuator coils (following the example of the Garching group). This gives straightforward control loop design and means that every actuator output has a corresponding sensor input. We use photo-diode sensing for both damping (shadow sensors) and accurate aligning (quadrant photo-diodes). Many aspects of the design inevitably reflect these specific system parameters.

9.1.3.1 Transducer I/O versus Control Axes Oriented Design

There are two possible approaches to the control system design: one involves careful mechanical design in which degrees of freedom are carefully decoupled and isolated. The second allows a flexible mechanical design, and puts the complication into the electronic configuration. In the first case six degrees of freedom are controlled using a separate servo loop for each degree of freedom. In the second, sensor-actuators are used in combination to act on an individual degree of freedom through appropriate summing and differencing. For example, three transducers can be arranged in an upright triangle to control longitudinal translation and pan / tilt orientation:- (1) the longitudinal signal is applied in common to all three actuators for translation; (2) the pan signal is applied in opposition to the left and right actuators for left-right steering; (3) the tilt signal is applied to the top actuator, and applied with opposite polarity and half the amplitude to the lower two actuators for vertical tilting. This approach also allows unwanted coupling into other modes or degrees of freedom to be electrically adjusted to a minimum.

It is sometimes difficult to act on a particular degree of freedom with a single transducer. For example, if the line through the centre of mass is inaccessible, a pair of sensor-actuators on each side of the mass can be used in combination to control the centre of mass (see left-right damping in figure 9.15). In this case one axis signal (translation) is the sum of the two input signals, while the other axis signal (rotation) is the difference of the two. A flexible design needs to allow for likely transducer arrangements. It must be able to take sums or differences of input signals (to extract a single axis signal). It must then be able to apply PID processing to each axis signal, before recombining the axis signals to create the separate actuator channel outputs. A dual sensor-actuator arrangement is shown in figure 9.7 which may easily be extended to the triple sensor-actuator configuration mentioned above.

From figure 9.7 it is clear that a major portion of the electronic circuitry is required for sensor-actuator interfacing. Thus a modular design can easily be created in which standard servo channels are designed around the sensor-actuator interfacing circuitry, with configurable interconnections to a standard control processing block. We now go on to describe the implementation of this structure.

9.1.3.2 Servo Channel Configurations and Computer Control

As discussed above, we require control loops in three general configurations:

1) Single point sensing and derivative damping.

2) Dual point (or more) sensing and derivative damping - sum and difference channels.

3) DC coupled quadrant photo-diode and long-lever-arm PID control.

These three configurations are similar enough to be implemented with identical circuitry which can be configured for either application. Figure 9.8 shows a block diagram of just such a universal channel design. The differences between the configurations are provided by a few jumpers and resistor value changes. All three

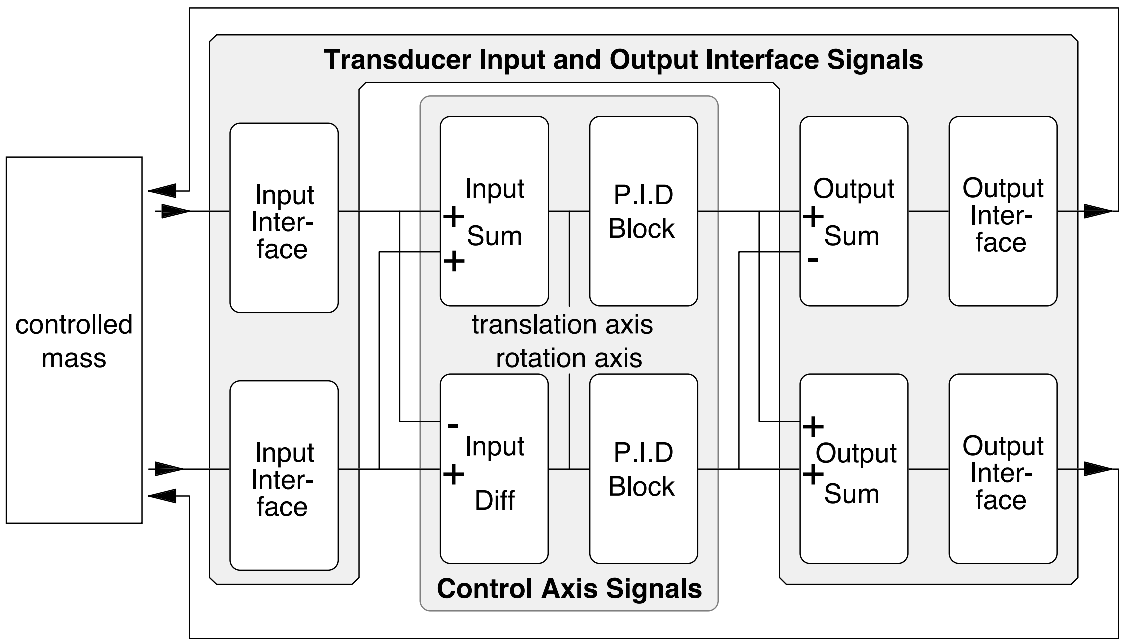

Figure 9.7 Transducer I/O versus axis control. A modular channel is designed around the sensor-actuator interfacing. A single axis signal is derived from an arithmetic combination of sensor inputs, processed by the PID control block, and recombined again to give the appropriate actuator output signal.

configurations are illustrated later in figure 9.14 which shows several interconnected channels.

For reasons of cost and simplicity we have chosen to use conventional high quality analog devices for the majority of the circuitry rather than considering a digitally based solution. However as shown in figure 9.8, computer interfacing is included for the reasons discussed in section 9.1.2.7. It is worth noting that a D/A converter has greater resolution and precision than most multi-turn potentiometers, so it is an economical and practical alternative to the front panel controls usually used in PID systems. Similarly A/D converters are a stable, low-noise means of adequately monitoring the control channels. The main signals to be monitored are obviously the controlled axes. It is also important to know whether the sensor inputs and outputs are within their dynamic range.

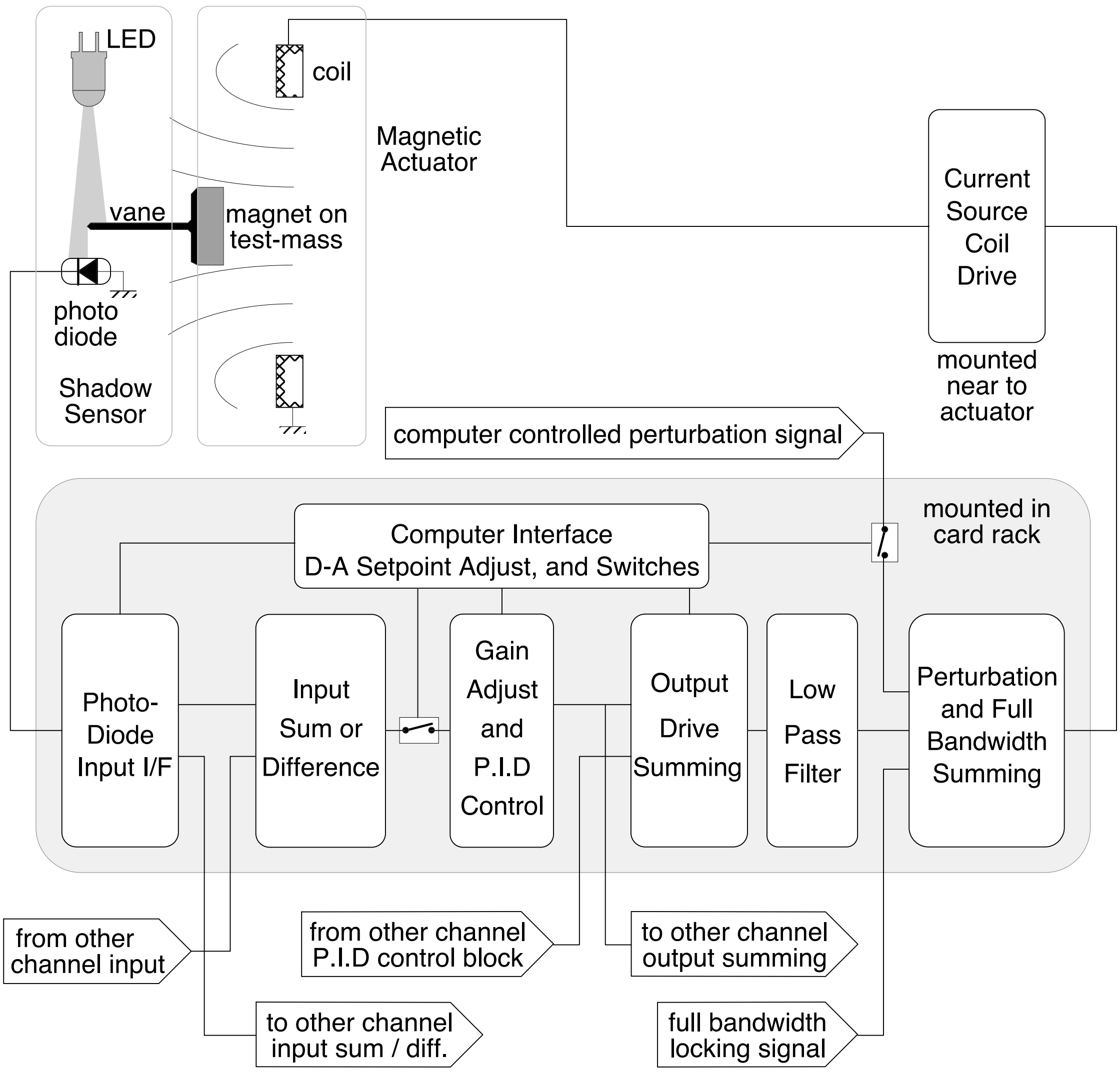

Figure 9.8 Single sensor-actuator servo channel. Two of these channels are fitted on a double eurocard sized PCB allowing most signal summing and differencing to be done on board. Each channel includes its own D/A converter and several computer controlled switches. Signals may be routed to and from channels on other boards via backplane connections as required.

Thus these two items are also monitored for each channel. The main output control adjustment required during normal operation, is to allow a positive or negative offset to be added to a controlled axis. This allows initial alignment to be achieved and serves many testing and experimental purposes. A computerised PID gain adjustment is not provided since precision gain control is not required, and once a loop is set up, it does not normally need altering.

Figure 9.9 shows how we have interfaced the control system to an IBM compatible PC using a "PC-30" [Boston 1994] analog I/O board. This board provides sixteen analog inputs and two analog outputs, all with 12 bit resolution, and 24 digital I/O lines. The digital lines are used to address select a particular channel, turn electronic switches ON and OFF, and store a value to the D-A converter. Three analog inputs are used to monitor the signals of interest from the selected channel and one of the analog outputs is used to source the perturbation signal.

An interface board is used to connect these PC-30 signals to the backplane wiring of the rack used to house the modular channel boards. Complete isolation between the PC-30 and the local control system is provided on this board by opto-isolators for the 24 digital lines. Analog isolators are used for the four analog lines. This effectively prevents the computer's electronic noise from contaminating the low-noise analog servo channels.

Three latched electronic switches are provided on each channel. The switch allowing the loop to be opened and closed is used to change from damping to orientation control and for open loop testing. The second switch allows the integral action of the orientation PID control to be turned off or reset. It is not used in simple damping configurations. The third switch allows the computer generated perturbation output to be switched so that it is added into the channel's output signal.

The universal control channel also contains three electronic switches which cannot be latched but are turned on automatically whenever a particular channel is addressed (or selected). These switches allow the three signals of interest to be applied to the analog bus lines so that their levels can be measured and monitored by the computer :-

1) The first signal of interest (ES1) is the channel signal itself. If the channel is configured for dual point sensing, then this signal is the summed or differenced position signal derived from the inputs of both channels. If it is configured for orientation control, then this signal is the sensed orientation.

2) The second signal of interest (ES2) is the quiescent value of one input and thus indicates how well the signal is falling within the dynamic range of the input. If the channel is configured for damping, then it indicates the quiescent level of current

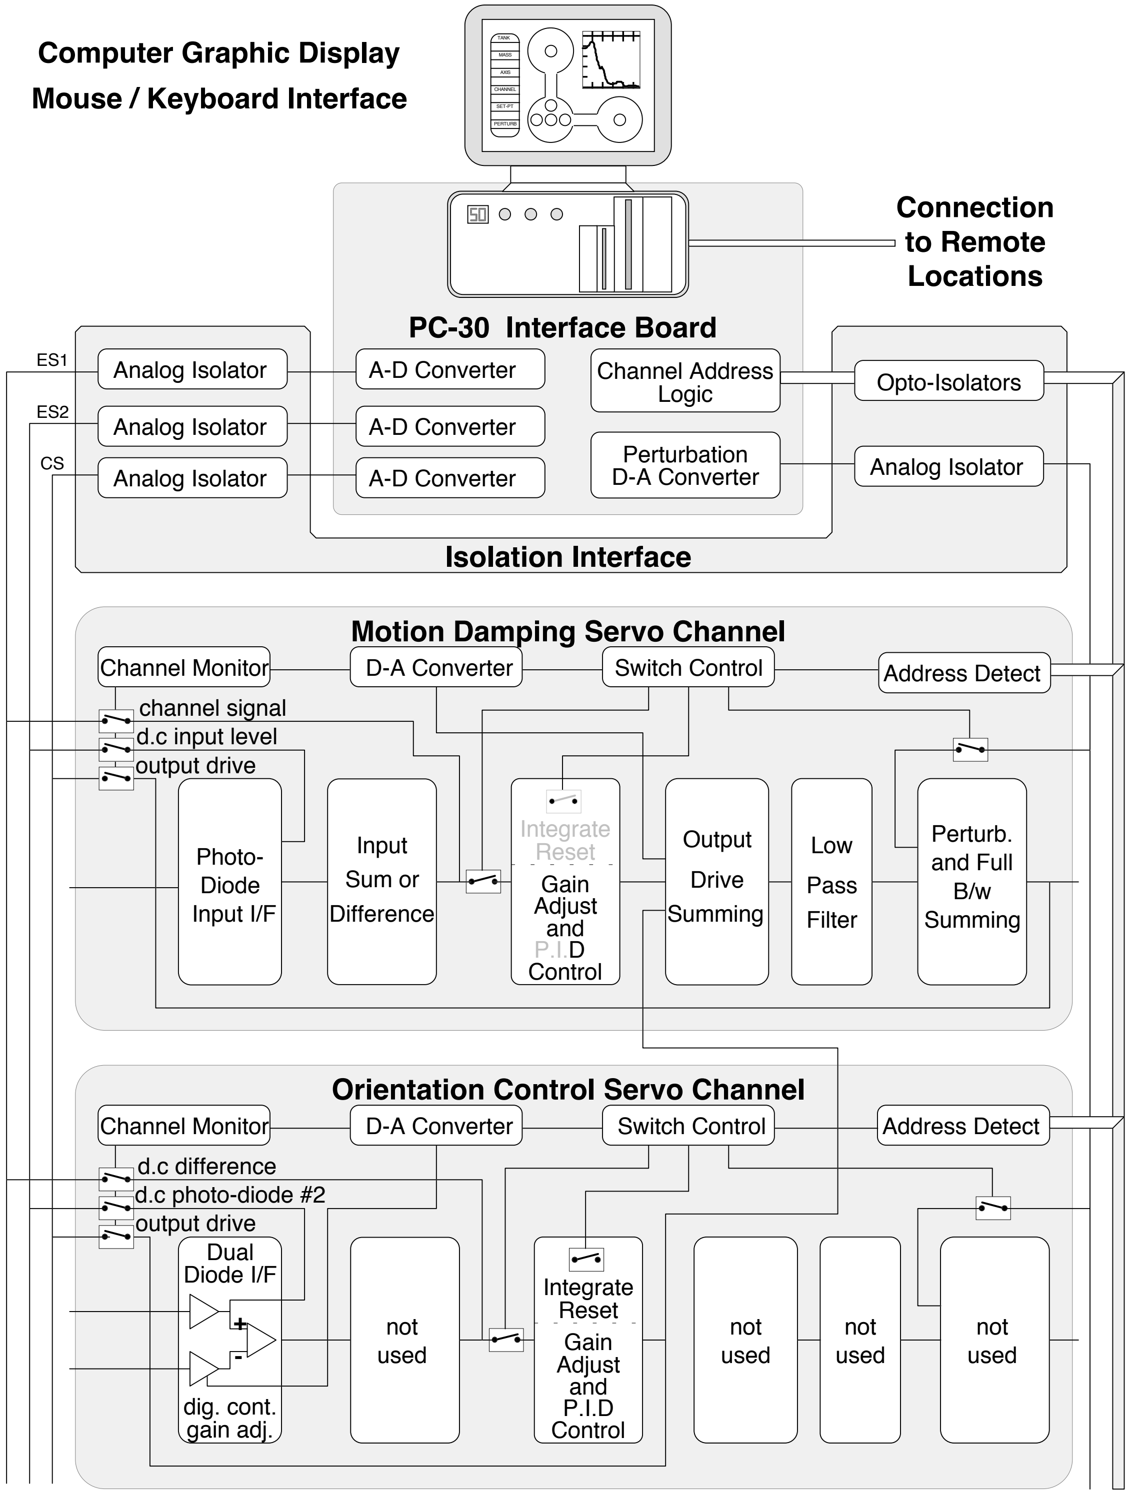

Figure 9.9 Computer interfacing to servo channels. The servo channels are interfaced using a commercial analog I/O board. This is isolated to keep computer noise out of the servo electronics, and is arranged to provide control and monitoring of each channel individually. Three signals of interest in a selected channel may be monitored, an analog adjustment made, control switches latched, and a perturbation signal added into the output for measurement and testing.

through the photo-diode (i.e. whereabouts in its total range that the shadow sensor is working). If the channel is configured for orientation control then it indicates what level of current is flowing through one side of the split photo-diode.

3) The other signal of interest (CS) is the output control signal being applied to the coil. This is the processed (derivative or PID) signal plus all other signals (perturbation, DC offset, etc) that may have been added into it. This allows the computer to check that the dynamic range of the coil drive is not being exceeded.

Each servo channel is fitted with a D/A converter which allows an adjustment to be made to that channel. The nature of the adjustment depends on the configuration of the channel :-

1) The single point sensing and damping configuration has the D-A converter output wired to the output drive summing so that a bipolar DC offset can be added into the otherwise AC coupled feedback loop. This allows an adjustable DC bias current to be caused to flow in the coil which may be used for manual testing or for initially aligning the test mass against a mechanical offset so that the auxiliary laser falls on the quadrant diode. This then allows the orientation control loop (which is DC coupled to provide any necessary offset) to be closed.

2) The dual point sensing and damping configuration has much the same connection and usage as the single point except that the bipolar offset is fed to both channels of the pair in the same way that the damping control signal is fed to both. That is if the inputs are summed then the bipolar offset is applied to both outputs. If the inputs are differenced then the bipolar offset is fed positive to one output and negative to the other. By this arrangement, the offset affects the same degree of freedom that the particular channel of the pair is damping (i.e. separated rotation or translation).

3) The DC coupled quadrant photo-diode configuration makes use of an entirely different arrangement. In this case the D-A converter is wired as an adjustable gain element in one of the two photo-diode signal paths. This allows the *ratio* of the currents from the split photo-diode sections to be adjusted as a set-point. This means that the long-lever-arm laser can be caused to fall on an area which is adjustably offset from the centre of the split photo-diode, and also means that fluctuations in laser beam intensity or photo-diode sensitivity will not affect the alignment as it is the current *ratio* that is controlled.

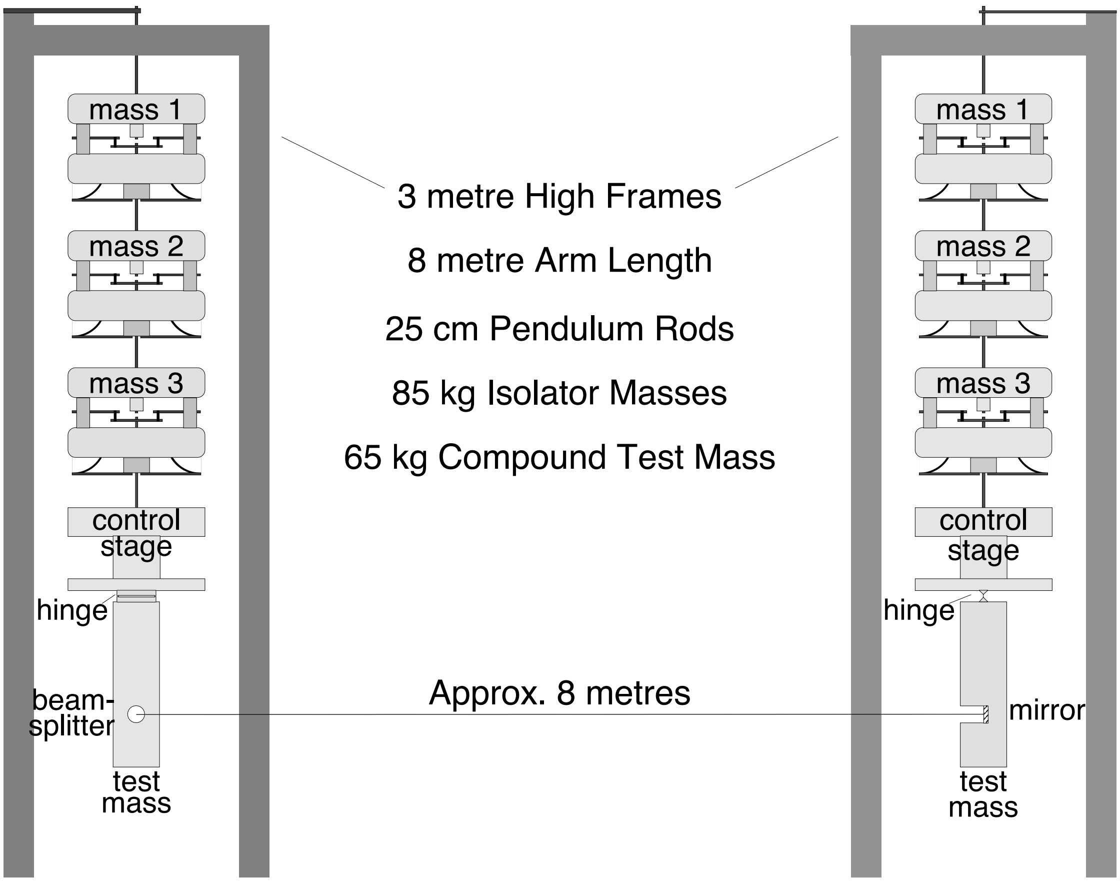

Figure 9.10 UWA Suspension and Test Mass Arrangement. We are currently working with a 4-stage stack as analysed, and have separated the control stage from the test mass proper along the main sensing axis by a high-Q hinge joint.

9.1.3.3 UWA Suspension and Test Mass Arrangement

At the University of Western Australia we have built a laser interferometer test facility, consisting of an 8 m arm length Michelson interferometer. The control system described here has been extensively tested on this system. The test mass suspension arrangements are as shown in figure 9.10.

Each test mass has 4 isolator stages. The test masses are compound pendulum stages as shown in figure 9.11. The isolator masses are approx 85 kg and are suspended by soft cantilever springs near to their centre of mass. The last stage consists of a control mass of 50 kg and a 15 kg compound pendulum test mass. The horizontal normal modes of the system are between 0.5 Hz and 2.5 Hz, while the vertical normal modes are between 0.9 Hz and 5 Hz. The test mass is suspended from the control stage by a low-spring-constant hinge. As reported elsewhere [Ju 1994a], this overcomes problems of detection band violin modes in suspension wires and in addition should provide a higher Q-factor for this last stage. The beam splitter or mirrors are mounted at the centre of percussion as this provides equivalent isolation to a simple pendulum - high frequency horizontal vibration at the hinge appears as tilting motion only and this does not affect the laser interference to first order.

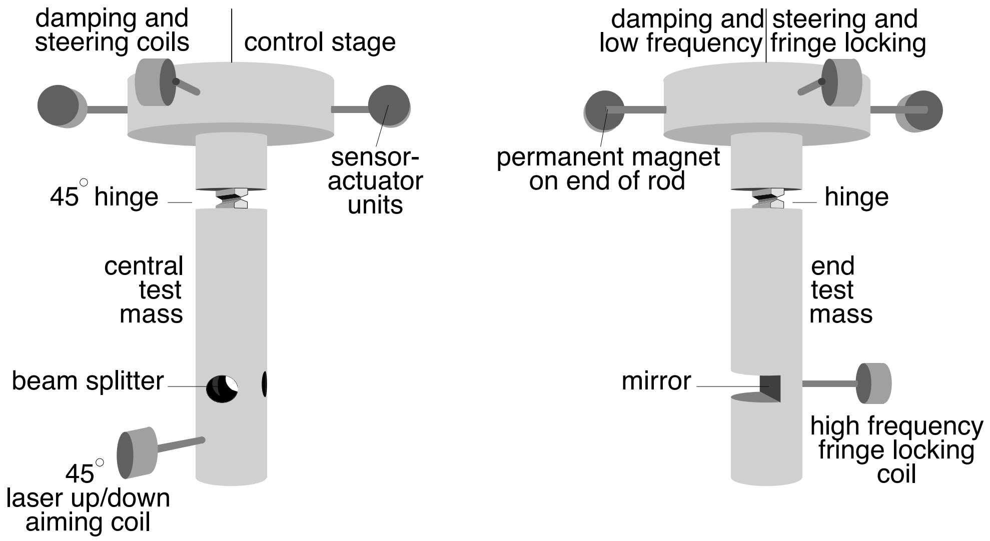

Figure 9.11 UWA Test mass and sensor-actuator configuration. The central test mass houses the beam splitter and the coil at the bottom is used to aim the laser up or down towards each of the end mirrors. The coils on the control stage are used for steering and damping. The mirror test mass only hinges in the laser direction and the high frequency part of the fringe locking signal is applied to its coil which acts in line with the centre of percussion. The low frequency part of the fringe locking signal is applied to the control stage, and the up down steering is redirected here also after attempting to offset the test mass (which is locked).

The compound hinges on the mirror test masses allow them to swing forward and back along the axis of the laser beam. The local control for these is divided between acting on the control stage and the test mass. Damping, left-right orientation, and low frequency fringe locking control is applied to the control stage; up-down orientation and high frequency fringe locking control is applied to the test mass as shown in figure 9.11.

The central test mass which houses the beam splitter is also used to provide the alignment control for the incoming laser beam. Its hinge is oriented at 45 deg to the laser paths to allow the mass to be swung forward and back along either of the two laser beam axes. The local control for this unit is also divided between acting on the control stage and the test mass although there is no fringe locking force applied to it. Damping and left-right orientation control is applied to the isolator mass; separate up-down orientation for each of the two beams is applied to the test mass.

9.1.3.4 Sensor-Actuator Units

The position sensing and magnetic actuator is constructed as a single unit as shown in figure 9.12. The coil assembly is attached to the frame and the permanent magnet and shadow vane are attached to the isolated mass. Motion is detected by the variation in light falling on the photo-diode as the vane moves, and current in the coils generates a largely position independent force on the magnet.

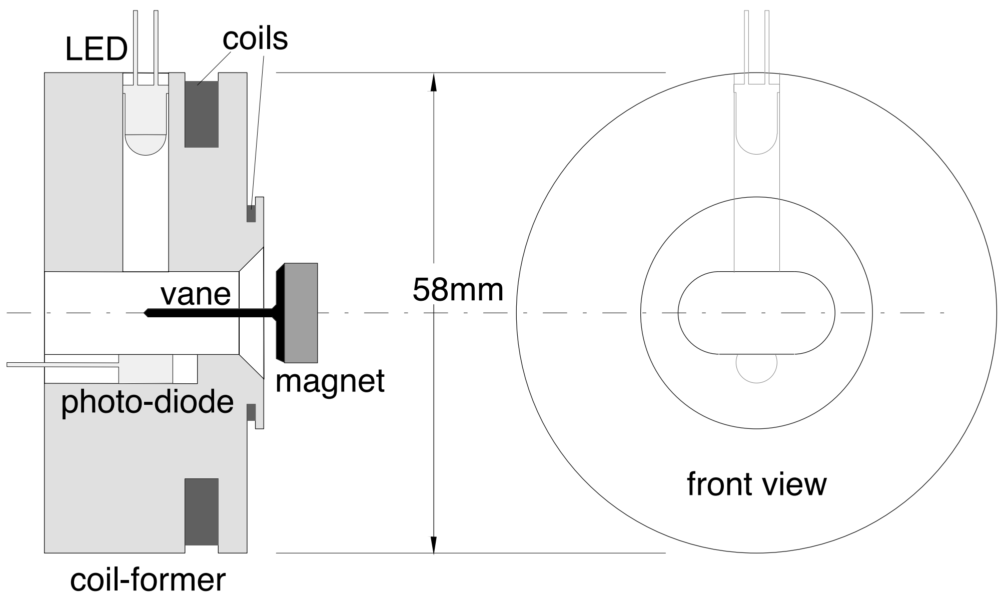

Figure 9.12 UWA sensor-actuator design. The coil-assembly is attached to the frame and the magnet is mounted on the end of a non-conductive pod attached to the isolated masses. The shadow vane is arranged to be half covering the photo-diode when the magnet is at a distance of half the radius from the large coil. Motion of the vane is detected from the variation in light falling on the photo-diode and current in the coils generates a position independent force on the magnet.

We used a narrow-angle infra-red LED as it gave the best current transfer ratio as well as having non-magnetic leads (the visible LEDs often have ferro-magnetic leads). The most suitable photo-diode we found had very slightly magnetic leads so we cut the leads off short and used copper wire in their place. The diode had a rectangular active surface for linearity, with an area of 2.6×2.6 mm giving about 2.5 mm dynamic range. The unshadowed photo-diode current was approx 0.45 mA for 50 mA through the LED. The cross-coupled signal from motion of the shadow vane parallel with the light beam is significant but not a problem at this stage. It could be improved with a lens or simply silvering the sides of the tube.

The shadow vane is arranged to be half covering the photo-diode when the magnet is at a distance of $r/2$ from the large coil (where r is the radius of the large coil) as discussed in section 9.1.2.4. The main coil is wound with 800 turns and the small linearising coil 50 turns. The resistance of the windings is approximately 100Ω and this allows a maximum current of around 100 mA to be passed through it. The permanent magnet used was measured as having a dipole moment of about 0.4 $A \cdot m^2$ which should produce a force of around 0.03 N with 100 mA flowing through the coils.

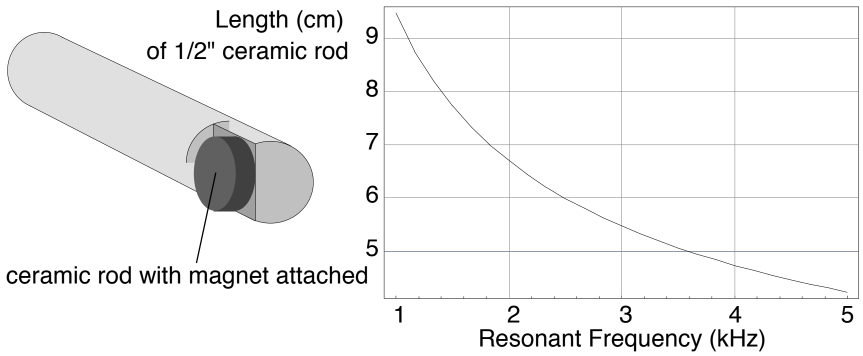

Figure 9.13 Rod mounted magnets. The permanent magnets are mounted on the end of an insulating rod to position the coil away from the metal of the control stage test mass. The natural resonance of the rod must not fall in the signal bandwidth (100 Hz to 1 kHz) and so it must be kept shorter than about 8 cm.

Since our isolated masses are all metallic (and the control stage is ferro-magnetic !), we decided to mount the permanent magnets on the end of insulating rods as shown in figure 9.13, to keep the ferromagnetic material and conducting surfaces as far away from the coils as possible. We chose a machinable ceramic for our initial experiment and determined that it was necessary to keep it shorter than about 8 cm in order to keep its lowest resonant frequency above about 1.5 kHz. We have not determined whether this length provides sufficient distance to avoid magnetic interaction, but the pod arrangement is advantageous for use with a Maxwell pair arrangement (figure 9.6).

9.1.3.5 Sensor-Actuator Positioning and Servo Loops

As illustrated in figure 9.14, a sensor-actuator mounted at each side of the control stage provides for forward-back translation (summed signals), and left-right rotation (differenced signals). Once the auxiliary laser falls steadily on the quadrant photo-diode, the left-right orientation channel signal takes over from the difference damping signal providing full PID orientation control

A sensor-actuator positioned to act on the test mass in line with the main sensing direction provides for up-down orientation control. A force applied at this point will act to move the test mass forward or back, simultaneously swinging the light beam up or down. This is the case for the mirror in the reference arm and the beam-splitter. However once fringe locking is achieved, the fringe locking system prevents the mirror in the locked arm from actually moving by generating control forces at the test mass and control stage which result in the control stage moving instead. This achieves the required up-down aiming effect. (Also shown in figure 9.14)

Damping of the left-right translational motion is stand-alone. It is only a single degree of freedom but as the most appropriate site for applying a damping force is already occupied (by the sensor-actuators on each side of the control stage), it is necessary to position a pair sensor-actuators - one at the front and one at the back - to be able to apply the force in line with the centre of mass. Both sensor-actuators could then be driven from a single channel by wiring both photo-detectors in parallel to the input and wiring both coils in series to the output. However this provides no means to adjust for an imbalance in the strengths of the two magnets. So a separate channel is used for each sensor-actuator to provide individual interfacing and the signals summed

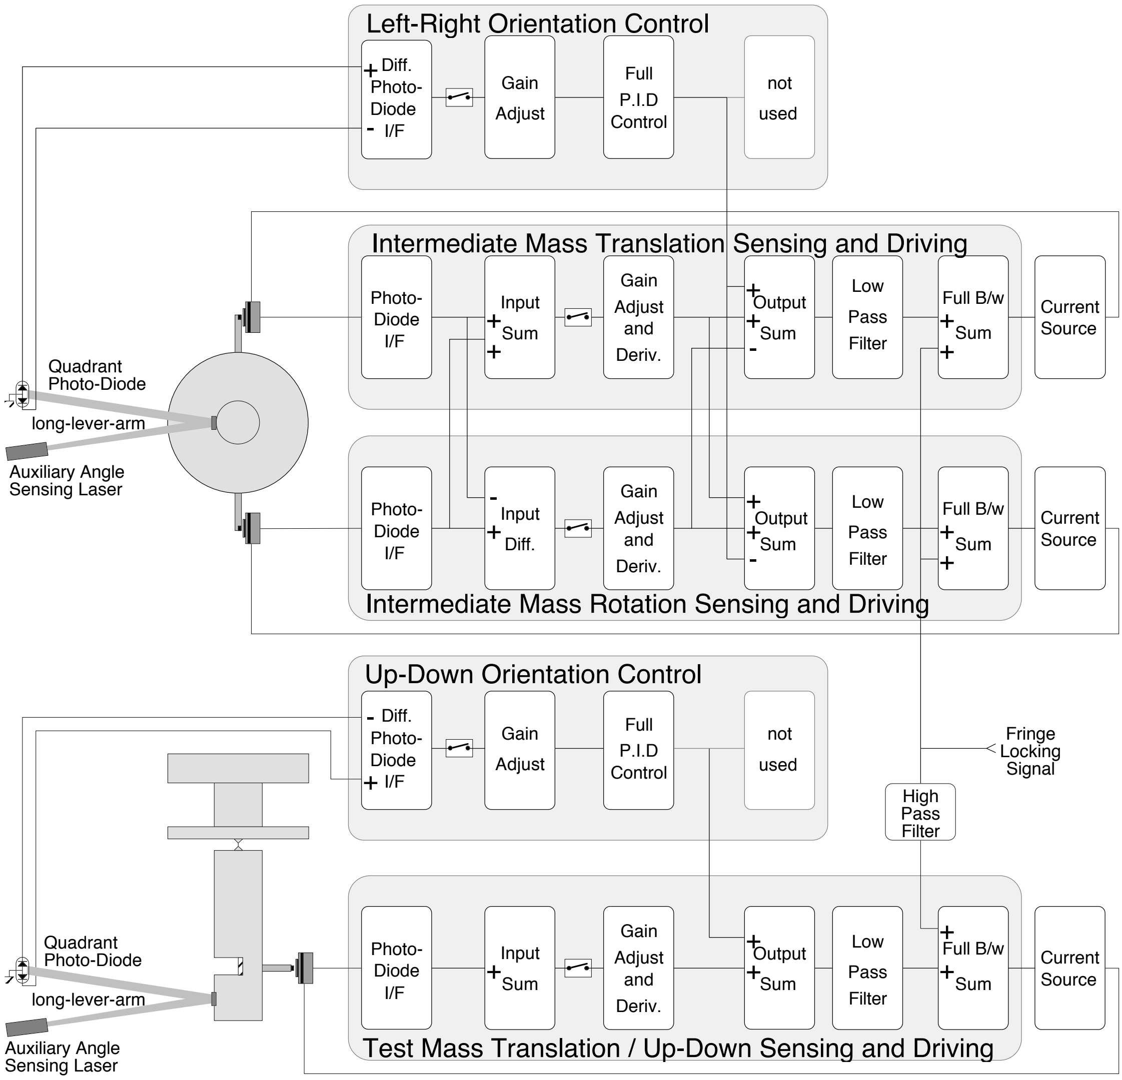

Figure 9.14 Interconnected servo loops for local control of test masses. Only one of the mirror masses incorporates the fringe locking signal shown, the other serves as the reference and is free to move. The beam-splitter servo loop arrangement is also without fringe locking and the up-down orientation control and driving channel is repeated for the perpendicular arm. Once fringe lock is obtained, the up-down drive on the locked channel will be appear as a low frequency error signal and be transferred to the control stage instead, achieving the desired up-down effect.

electronically. In this case the difference signal is unused but the coil force may be balanced by an adjustment provided on the current source modules. This is illustrated in figure 9.15.

The sequence of operation is that initially the system acts to damp all motion of the suspended masses (at the stack normal mode frequencies) in order to cool them down to the seismic level. Once the long-lever-arm laser beams fall steadily on the quadrant photo-diodes, the orientation control loops are closed and the damping signals for these orientation axes turned off (since they now only act to disturb the orientation as the damping sensing will still be picking up low frequency seismic noise). Having obtained the required mirror alignment, an interference signal should be available for fringe

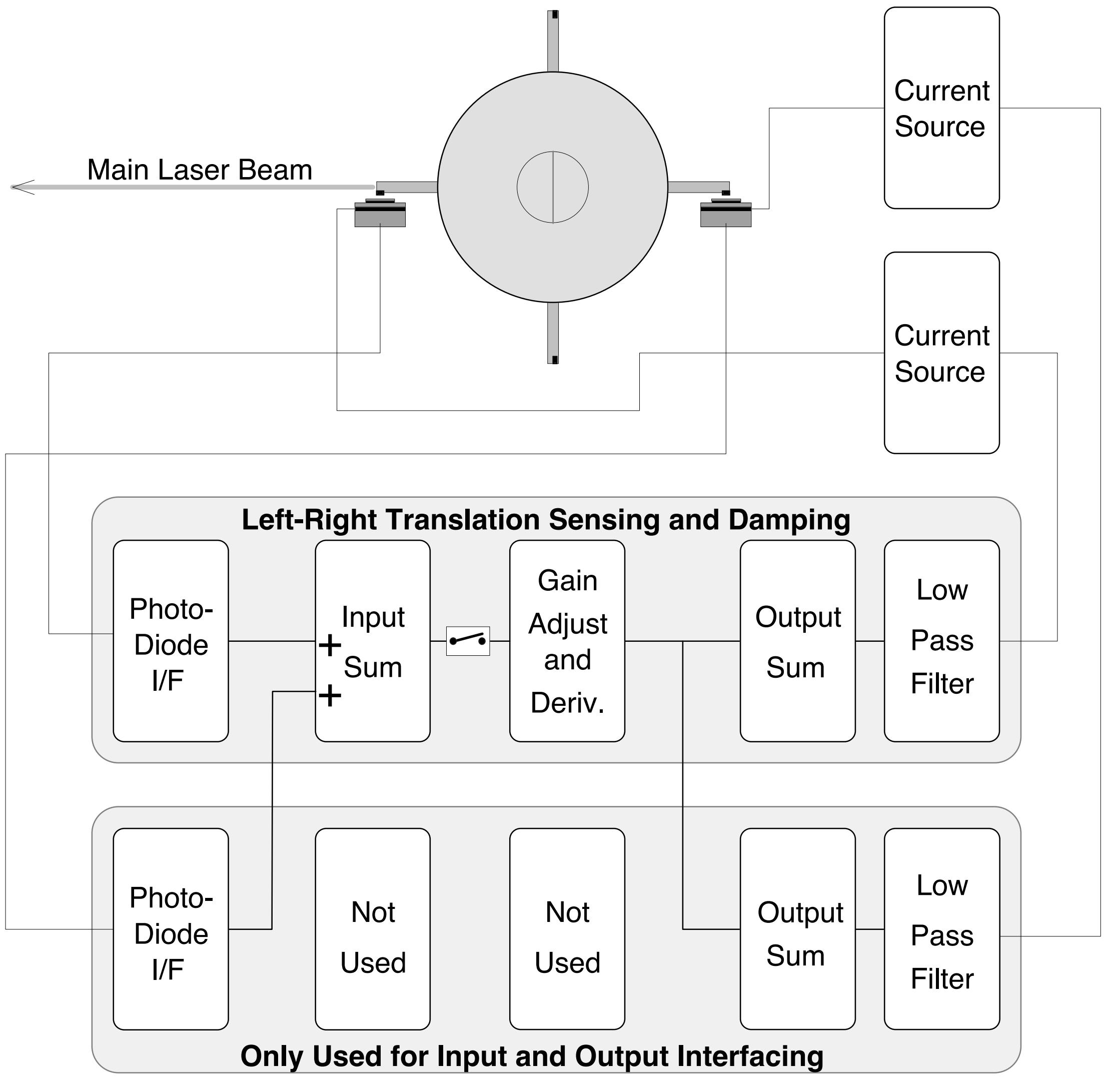

Figure 9.15 Servo loop for damping of left-right translation. Since it is difficult to position a single sensor-actuator to act in line with the centre of our test mass, one is positioned each side to act in unison. With this arrangement it would be possible to wire both photo-detectors in parallel to a single input, and wire both coils in series to a single output thus only using a single channel. However this does not allow balance adjustments to be made.

locking the interferometer arms. If resonant cavities are formed in the arms then it is very likely that the interference signal will still be too transient to produce lock without some additional circuitry which can derive a velocity signal from the nature of its transience to provide further viscous damping. Once fringe lock is obtained, the damping loops are left operational as their function is to damp common mode motion. They will attempt to disturb the lock, but only at unimportant stack resonance frequencies and the fringe locked loop transfers the force to the remote mass nullifying any differential effect.

9.1.4 Performance Measurements

The open loop gain of a pair of channels has been carefully measured in closed loop configuration. The channels were connected to a pair of sensor-actuators positioned one each side of the control stage such that forward and back translation produced a common mode signal and rotation produced a differential mode signal. A spectrum analyser was then connected up with its signal sourcing facility injecting into the closed loop, and the open loop response of both the translational and the rotational loops were measured (using the analyser's maths function capability). The gain was increased until instability appeared and then backed off a little and the measurement taken averaging over many hours. The resulting plots showed that the rotational open loop gain could be set quite high - reaching a peak of over 100 at a rotational mode of 0.3 Hz giving very good damping. However the translational gain shown in figure 9.16 could only be pushed up to a peak of just over 10 at the 0.5 Hz normal mode (ref figure 9.4 for normal modes). The resultant damping was acceptable but was not as good as hoped for. The reason for the instability which occurred at about 6 Hz was found to be due to a cantilever resonance of the entire 3 m high steel frame (ref figure 9.10) supporting the stack and to which the sensor-actuators are mounted. (The fixed ends of the cantilever are the welded corners at the top). This resonance is excited by the reaction force from the active damping of the isolator stack. Extra mass was placed on the top of the frame and the resonance was seen to move down in frequency. When diagonal struts were clamped to the steel frame, the frequency was seen to move upwards slightly but not sufficiently to allow extra gain to be applied. A possible solution would be to reduce the Q-factor of the resonance by attaching diagonal straps of some vacuum compatible low Q material (i.e. lead) or a passive spring/magnet eddy current damper [Blair 1993]. This would squash down the resonance peak allowing more gain to be applied to whatever extent the peak can be lowered. Despite this resonance limiting the translational loop gains which can be achieved, the normal modes are still well damped and the system is quite useable without attention to this problem.

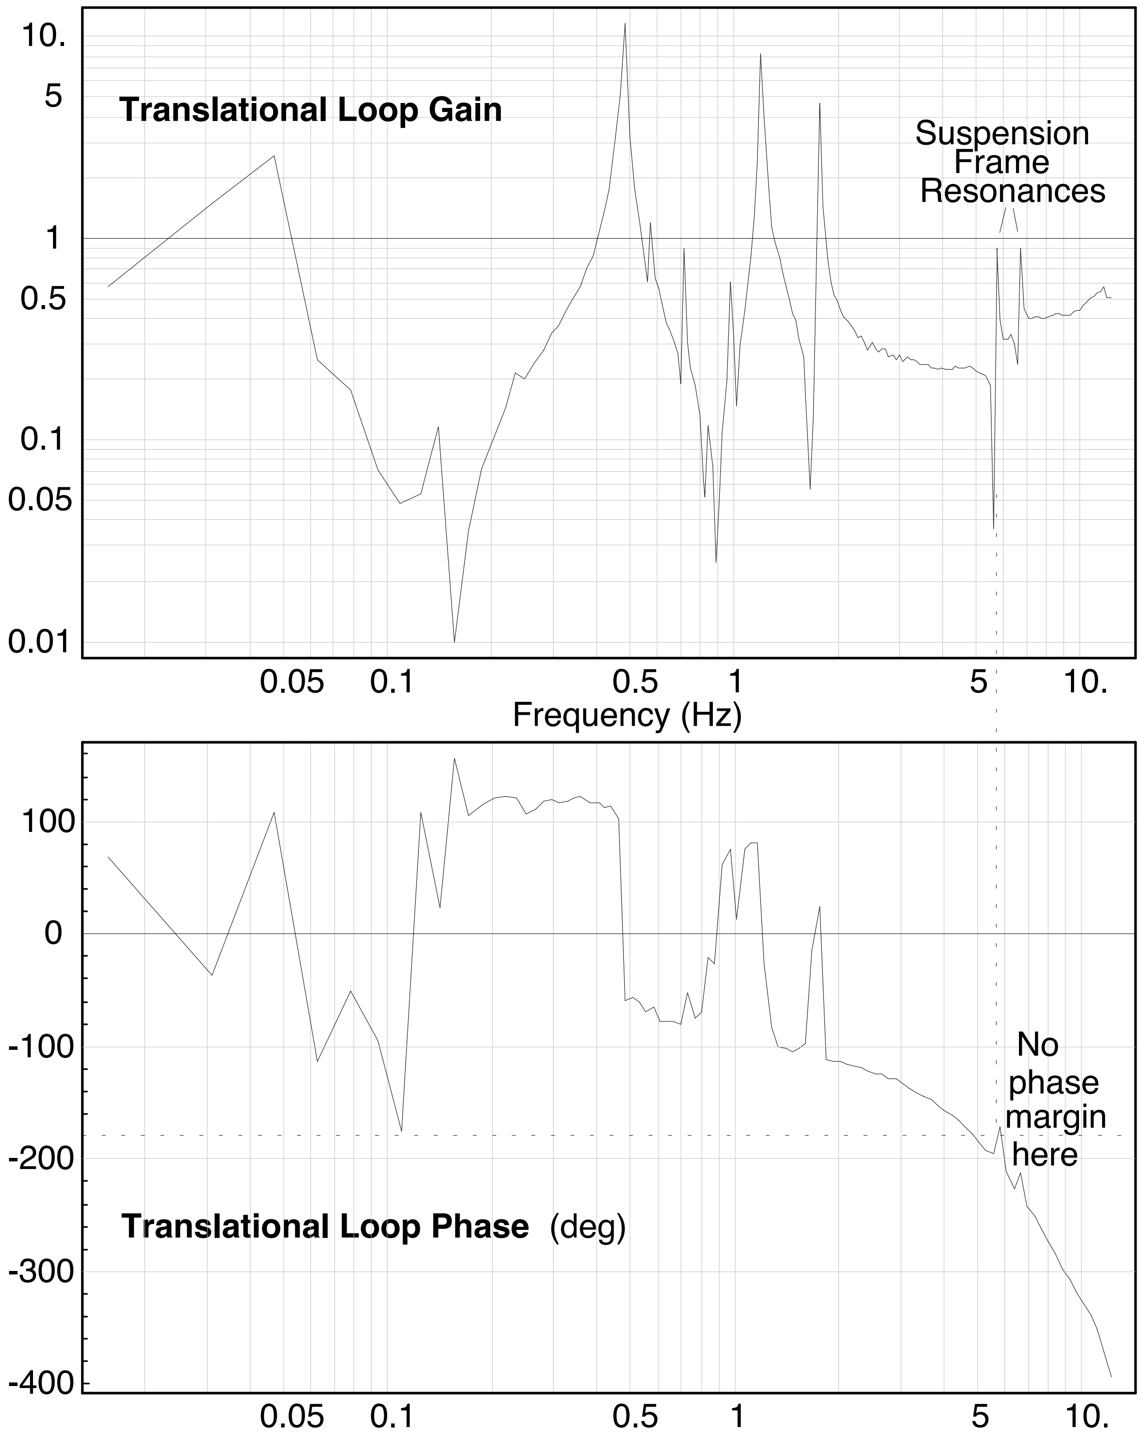

Figure 9.16 Translational loop gain and phase. The in-situ loop gain and phase was measured for the horizontal translational damping channel with the above result. The gain was limited by a structural resonance that produced instability at about 6 Hz.

9.1.5 Conclusion

We have presented an analysis of the control system requirements for a laser interferometer gravitational wave detector, with emphasis on systems without strongly internally damped isolation stacks. The system that has been described has been tested with damping active on three stacks controlling a total of 13 degrees of freedom (4+5+4). The long-lever-arm PID control has also been tested controlling a total of three degrees of freedom, and fringes are readily obtained. The normal modes of the vibration isolators described can all be strongly suppressed, typically by more than 20 dB and in some cases by more than 30 dB. Details of the isolator performance and normal mode suppression will be published elsewhere [Ke 1994]. Results have shown the importance of avoiding frame resonances in the normal mode frequency band, emphasising the need for very massive frame structure with low Q-factors and moderately high normal mode frequencies. As the isolator normal mode frequencies are reduced, the requirement on the frame rigidity is also reduced.

The servo system has shown itself to be very versatile, and the computer control facility including analog display of error signals and adjustment of analog offsets is very convenient to operate.

Further details of circuits including PC board layouts etc are available from the authors.

9.1.6 Acknowledgments

We would like to thank Peter Smith and Ian Russel of the Electronics Workshop for their contribution in terms of PC software development and circuit board manufacture. This project was supported by the Australian Research Council.

9.2 Publication List

J. Winterflood, D.G. Blair, *High performance vibration isolation using springs in Euler column buckling mode*, **PAPER 2.2**, **Phys. Lett. A, accepted**.

J. Winterflood, T.A. Barber, D.G. Blair, *Mathematical analysis of an Euler spring vibration isolator*, **PAPER 2.3**, **Phys. Lett. A, accepted**.

J. Winterflood, T. Barber, D.G. Blair, *Using Euler Buckling Springs for Vibration Isolation*, **Class. Quantum Grav., 19 (2002) 1639-1645**.

H. Hoffmann, J. Winterflood, Y. Cheng, D.G. Blair, *Cross-Correlation Studies with Seismic Noise*, **Class. Quantum Grav. 19 (2002) 1709-1716**.

Y. Cheng, J. Winterflood, L. Ju, D.G. Blair, *Tilt sensor and servo control system for gravitational wave detection*, **Class. Quantum Grav. 19 (2002) 1723-1729**.

D. Paget, J. Winterflood, Ju Li, D.G Blair, *Improved Technique for Measuring High Pendulum Q-factors*, **Meas. Sci .Technol. 13 (2002) 218-221**.

D.E. McClelland, S.M. Scott, M.B. Gray, D.A. Shaddock, B.J. Slagmolen, A. Searle, D.G. Blair, L. Ju, J. Winterflood, F. Benabid, M. Baker, J. Munch, P.J. Veitch, M.W. Hamilton, M. Ostermeyer, D. Mudge, D. Ottaway and C. Hollitt, *Second-generation laser interferometry for gravitational wave detection: ACIGA progress*, **Class. Quantum Grav. 18 (2001) 4121-4126**.

Z.B. Zhou, J. Winterflood, L. Ju, D.G. Blair, *Investigation of a laser walk-off angle sensor and its application to tilt measurement in gravitational wave detection*, **Physics Letters A 280 (2001) 197-203**.

J. Winterflood, Z.B. Zhou, L. Ju, D.G. Blair, *Tilt Suppression for ultra-low residual motion vibration isolation in gravitational wave detection*, **PAPER 7.3**, **Physics Letters A 277 (2000) 143-155**.

L. Ju, D.G Blair, J. Winterflood, *Long-term length stability and search for excess noise in multi-stage cantilever spring vibration isolators*, **Physics Letters A 266 (2000) 219-227**.

J. Winterflood, Z.B. Zhou, D.G. Blair, *Ultra-low Residual Motion Suspension System*, in **Proceedings** of the 2nd TAMA International Workshop on Gravitational Wave Detection held at National Olympics Memorial Youth Center, Tokyo, Japan on October

19-22, 1999, Gravitational Wave Detection II, (Universal Academy Press, Tokyo), p301-310, (2000).

J. Winterflood, Z.B. Zhou, and D.G. Blair, *Reducing Low-Frequency Residual Motion in Vibration Isolation to the Nanometre Level*, in **Proceedings** of the Third Edoardo Amaldi Conference, Pasadena, California, 12-16 July 1999, Gravitational Waves, Ed. Sydney Meshkov, (AIP, New York), p325-321, (2000).

J. Winterflood and D.G. Blair *Ultra-Low Frequency Pre-Isolation*, in **Proceedings** of the 8th Marcel Grossmann Meeting on General Relativity, Jerusalem, 1997, Ed: T. Piran, Series Ed: R. Ruffini, World Scientific Publ. Co. Part B pp 1071-1073, (1999).

J. Winterflood, G. Losurdo, D.G. Blair, *Initial results from a long-period conical pendulum vibration isolator with application for gravitational wave detection*, **PAPER 4.3**, **Physics Letters A. 263** p9-14 (Nov 1999).

G. Losurdo, M. Bernardini, S. Braccini, C. Bradaschia, C. Casciano, G. Cella, E. Cuoco, V. Dattilo, R. De-Salvo, E. D'Ambrosio, A. Di-Virgilio, F. Fidecaro, F. Frasconi, A. Gaddi, S. Gaggero, A. Gennai, A. Giazotto, L. Holloway, P. La-Penna, S. Mancini, M. Mazzoni, H.B. Pan, A. Pasqualetti, R. Passaquieti, D. Passuello, F. Paoletti, R. Poggiani, R. Stanga, A. Vicere, R. Taddei, J. Winterflood, Z. Zhang, *An inverted pendulum preisolator stage for the VIRGO suspension system*, **PAPER 6.2**, **Rev. Sci. Inst. 70** p2507-15 (1999).

R. DeSalvo, A. DiVirgilio, A. Errico, G. Feng, I. Ferrante, F. Frasconi, A. Gaddi, A. Gennai, G. Gennaro, A. Giazotto, L. Holloway, G. Losurdo, M. Maggiore, S. Mancini, M. Mazzoni, F. Palla, P.H. Bao, F. Paoletti, A. Pasqualetti, R. Passaquieti, D. Passuello, R. Poggiani, P. Popolizio, F. Raffaelli, D. Righetti, R. Ruberti, V. Rubino, P. Ruggi, R. Stanga, R. Taddei, R. Valentini, A. Vicere, J. Winterflood, Z. Zhang, M. Bernardini, A. Bozzi, S. Braccini, C. Bradaschia, C. Casciano, G.C. Cella, A. Ciampa, G. Curci, E. D'Ambrosio, V. Dattilo, G.C. DeCarolis, A. Delapierre, *Performances of an Ultra Low Frequency Vertical PreIsolator for the VIRGO Seismic Attenuation chains*, **Nucl Instrum & Methods in Phys. Res. A**, **420** p316-35 (1999).

J. Winterflood, D.G. Blair, *Ultra-low Frequency Pre-Isolation*, in **Proceedings** of the 8th M. Grossmann Meeting on General Relativity, Jerusalem Israel, Jun 1997, PartB, (World Scientific, Singapore), p1071-3, (pub May 1999).

J. Winterflood, D.G. Blair, *Ultra-low Frequency Pre-Isolation in Three Dimensions*, in **Proceedings** of the 2nd E. Amaldi Conf. on Gravitational Waves, Geveva, Switzerland, Jul 1997, (World Scientific, Singapore), p485-9, (pub Feb 1999).

M. Bernardini, S. Braccini, C. Bradaschia, C. Casciano, G. Cella, A. Ciampa, E. Cuoco, G. Curci, V. Dattilo, R. De Salvo, A. Di Virgilio, G. Feng, I. Ferrante, F. Fidecaro, F. Frasconi, A. Gaddi, A.Gennai, A. Giazotto, L. Holloway, P. LaPenna, G. Losurdo, S. Mancini, M. Mazzoni, F. Palla, H. B. Pan, A. Pasqualetti, R. Passaquieti, D. Passuello, R. Poggiani, P. Popolizio, F. Raffaelli, S. Rapisarda, R. Stanga, A. Vicere, J. Winterflood, Z. Zhang, *Active Control Hierarchy in VIRGO Superattenuator: The Role of the Inverted Pendulum* in **Proceedings** of the 2nd E. Amaldi Conf. on Gravitational Waves, Geveva, Switzerland, Jul 1997, (World Scientific, Singapore), p334-8, (pub Feb 1999).

J. Winterflood, D. G. Blair, *A long-period vertical vibration isolator for gravitational wave detection*, **PAPER 5.2**, **Physics Letters A. 243** p1-6 (1998).

G. Losurdo, J. Winterflood, *Inverted Pendulum Transfer Function*, VIRGO **report** VIR-TRE-PIS-4600-123, September 1997.

R. DeSalvo, A. Gaddi, G. Gennaro, L. Holloway, G. Losurdo, J.Winterflood, *Pre-Isolator Stage for VIRGO*, VIRGO **report** NTS-096/034 May 1996.

DG Blair, M Notcutt, L. Ju, ME Tobar, YM Yang, E Wang, C Taylor, J Winterflood, JF Liu, M Taniwaki, CN Zhao, D McClelland, C Harb, J Sandeman, D Mudge, D Ottaway, MW Hamilton, J Munch, PJ Veitch, *Progress in the Development of Technology for Advanced Laser Interferometric Gravitational Wave Detectors*, **Proceedings** of the TAMA Workshop on Gravitational Wave Detection, Saitama, Japan, Dec 1996. Ed: K. Tsubono, M.-K. Fujimoto, K. Kuroda. Universical Acad. Press (Tokyo) pp 75-94 (1997).

G. Losurdo, A. Gaddi, A. Gennai, D. Passuello, R. De Salvo, L. Holloway, J. Winterflood, R. Woode, *Velocity damping of a prototype Superattenuator pre-isolator stage*, VIRGO **report** VIR-NOT-PIS-4600-105, October 1996.

C.N. Zhao, J. Winterflood, M. Taniwaki, M. Notcutt, J. Liu, L. Ju and D.G. Blair, *The development of technology for high performance laser interferometer gravitational wave detectors*, in **Proceedings** of the Int. Conf. on Gravitational Waves: Sources and

Detectors, Cascina, Italy, March 1996, Ed: I. Ciufolini & F. Fidecaro, World Scientific, pp.211-216 (1997).

C.N. Zhao, J. Winterflood, M. Taniwaki, M. Notcutt, J. Liu, L. Ju and D.G. Blair, *The development of technology for high performance laser interferometer gravitational wave detectors*, in **Proceedings** of the First Australasian Conf. on General Relativity and Gravitation, Univ. of Adelaide, February 1996. Ed: D.L. Wiltshire, Inst. For Theoretical Physics, University of Adelaide, pp.7-12 (1996).

M.G. Zadnik, J. Winterflood, A.J. Wiliams, K.J. Wellington, B. Vaile, J. Tarter, R. Norris, G. Heiligman D.G. Blair and P. Backus, *Interstellar Communication Channel Search of 49 Stars Closer than 11.5 pc*. in **Proceedings** of the Capri Symposium for BioAstronomy, July, 1996.

J. Winterflood, *Low-Frequency Pre-Isolation for VIRGO*. **Seminar transparencies** for seminar during the joint VIRGO conference at Pisa in March 1996.

J. Winterflood, D.G. Blair, *A long-period conical pendulum for vibration isolation*, **PAPER 4.2**, **Physics Letters A**, **222** p141-147 (1996).

J. Winterflood, D.G. Blair, R. Schilling, M. Notcutt, *A Position Control System for Suspended Masses in Laser Interferometer Gravitational Wave Detectors*, **PAPER 9.1**, **Rev. Sci. Inst. 66** p2763-76 (1995)

J. Liu, J. Winterflood, D.G. Blair, *Transfer Function of an Ultra Low Frequency Vibration Isolation System*, **Rev. Sci. Inst., 66** p3216-18 (1995)

A. Giazotto, L. Ju, R. Poggiani, J. Winterflood, *Analysis of the Interaction of the Marionetta Coils / Magnets with the Vacuum Baffle Plate*, VIRGO **report**, January 1995.