An Integrated Approach to Structural Damping

by

Eric Russell Marsh

B.S., University of Illinois (1990)

M.S., Massachusetts Institute of Technology (1992)

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 1994

Author .....
Department of Mechanical Engineering
April 29, 1994

Certified by .....
Alexander H. Slocum
Associate Professor
Thesis Supervisor

Accepted by .....
Ain A. Sonin
Chairman, Departmental Committee on Graduate Students

An Integrated Approach to Structural Damping

by

Eric Russell Marsh

Submitted to the Department of Mechanical Engineering
on April 29, 1994 in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Mechanical Engineering

Abstract

Structural performance is often improved by developing designs with higher stiffness. The drawback to this approach is that increasing stiffness without addressing structural damping can result in poor dynamic performance. For example, alumina is now being used in place of cast iron as a structural material in some applications because of the superior stiffness (by a factor of two). However, the internal damping in alumina is much lower than that of cast iron. To achieve world class dynamic performance in mechanical structures (e.g., machine tools), the stiffness as well as damping of a design must be considered. For example, the machine tool industry is slowly making the transition from cast iron machine bases to steel weldments. Steel structures are stronger, lighter, and less expensive, but are prone to excessive vibration. Structural vibration damping will be critical to meet new challenges in manufacturing.

The fundamental contribution of this research is the development of a robust damping mechanism capable of energy dissipation over a wide range of frequencies and vibration amplitudes, and a cohesive theory that allows designers to readily predict performance. Viscoelastic constrained layer damping is a known method of damping thin, plate-like structures. This research makes it possible to use the same materials on the inside of a structure with robustness and design theory accuracy. The damping mechanism dissipates energy efficiently, and the proposed designs offer damping without compromising the stiffness of a structure. The mechanism is well modeled by the finite element method and an estimate of a structure's damping may be obtained without building and testing a prototype.

The result of this research is the development of a new method of providing damping in structures shown to give very high energy dissipation. This includes a mathematical model of the damping mechanism, an understanding of optimal implementation, and a number of case studies verifying the theory.

Thesis Supervisor: Alexander H. Slocum
Title: Associate Professor

Acknowledgments

This research was sponsored by a number of generous organizations: LeBlond-Makino Machine Tool Corp., makers of machining centers; Weldon Machine Tool, makers of cylindrical grinding machines; the Leaders for Manufacturing Program, which awarded support for research related to manufacturing; and the National Science Foundation.

A number of people made great contributions to this research effort. Mr. Fred Coté provided patience and wisdom during the experimental testing. My thesis committee, Prof. Alex Slocum, Prof. Carl Peterson, and Prof. J. Kim Vandiver, provided tremendous enthusiasm and offered innumerable suggestions and ideas that resulted in the wonderful successes this program has enjoyed. Many other MIT faculty provided timely advice along the way: Prof. Stephen Crandall, Prof. Mary Boyce, Prof. John Lienhard, Prof. Richard Lyon, and Prof. Ain Sonin.

Extra thanks goes to Prof. Earnest Rabinowicz who supplied endless technical advice and moral support.

I also thank Prof. Alex Slocum, who made this exciting research a reality. He provided an ideal research environment with eager industrial sponsors, a steady flow of new test equipment and materials, as well as boundless creative insight and energy.

Finally, I am grateful for the support of my wife Kristin who spent countless hours editing this dissertation. Kristin gave me the love, patience, and faith needed to complete this work.

Contents

Chapter 1: Modern Structural Damping Techniques

1.1 Contributions of the Dissertation

This dissertation presents the results of research performed in the area of structural damping. In the course of this research, a replicated in-place shear damping mechanism capable of dissipating energy in a variety of structural geometries was developed. Furthermore, a comprehensive analysis method was derived that allows the designer to work with the shear dampers in an efficient manner. The shear damping mechanism has been designed and built into full scale machine tool structures and providing very high damping for both bending and other vibration modes. These excellent results illustrate the tremendous impact that the shear damping mechanism will have on the machine tool and other industries.

MIT has filed a patent application for the shear damping mechanism which will allow the licensing of this technology to American industry.

1.2 Introduction

Structural vibration can result from energy sources such as floor noise, actuator activity, cutting forces, and the movement of various components. These sources often generate (relatively) wide band excitation that may excite one or more of the critical modes of vibration in a structure. Designers faced with these operating environments use materials such as cast iron for heavy structures and plastics or composites for light structures because of their favorable internal damping. Hollow structures filled with other materials such as hydraulic or polymer concrete may provide further vibration reduction. Isolation mounts, tuned mass dampers, and constrained layer treatments can be installed if the design still exhibits poor dynamic performance. While these approaches may be used to improve the performance of many devices, there are disadvantages to each that make them inappropriate for many applications. For example, tuned mass dampers may reduce vibration, but only in a fairly narrow frequency bandwidth. Furthermore, tuned masses can be physically large and difficult to build into an enclosure.

The damping mechanism developed in this dissertation offers a significant advantage over many competing damping methods: its effect on structural performance may be accurately estimated prior to prototype construction. The mechanism can be straightforwardly modeled in finite elements because of its linear, velocity proportional damping behavior.

The proposed damping mechanism shears a lossy material to dissipate energy, but unlike traditional constrained layer treatments, the proposed shear medium can be a viscous fluid *or* a viscoelastic solid. The mechanism provides damping over a wide range

of frequencies and bending amplitudes. The manufacturing of these damped structures is simpler than traditional constrained layer treatments because they are built directly into the inside of a structure with replicating epoxy.

For example, a thin layer of fluid placed between two smooth, simply-supported beams is sheared as the beams bend relative to each other (because one face will be in compression as the other is in tension). As the beams oscillate, energy is lost in shearing the fluid. This energy dissipation is modeled analytically and found to be proportional to the fluid viscosity and inversely proportional to the lossy layer thickness. Experimental study of the performance of damped slender beams verifies the relationship between loss factor and fluid layer thickness, but several new issues appear.

These issues include the behavior of the damping mechanism as the damping layer is made progressively thinner. Eventually, as the damping layer becomes very thin, the structure will not receive any damping because the layer has almost zero shear strain. In this case, the structure will have increased stiffness, but very little damping.

The optimal damping layer thickness for a particular structure lies between the two extremes of a thick layer with little damping and a very thin layer offering little damping. In practical applications, it is reasonable to expect that either a viscous fluid or a viscoelastic material is available that will provide the optimal damping. Viscous fluids are best used to damp low modulus structural materials, while viscoelastic materials are necessary to properly damp heavier metallic or ceramic structures.

1.3 Performance of Structural Damping Mechanisms

The motivation for structural damping research is the hope of finding a general strategy of reducing vibration in structures of arbitrary geometry and materials. While some industries, such as aerospace, have identified working solutions to specific vibration problems, many other industries are only now addressing the need for structural damping. The ever increasing structural stiffness of modern machine design dictates that the importance of damping will increase.

In order to provide the reader with an indication of the dynamic response of a steel structure, this section will demonstrate the behavior of a moderate aspect ratio steel beam subjected to various damping treatments. The beam is a square steel tube with an outer dimension of 10 cm and a wall thickness of 6 mm. The length of the tube is 0.8 m.

The dynamic response of the undamped beam is shown in Figure 1.1. Note that the modal peaks are very high, indicating very low damping (the modal loss factor η is around 0.002 for the various peaks). The dynamic response of the same beam filled with water (not shown) also exhibits modal loss factors of about 0.002 for the first several modes of vibration. The figure shows the dynamic response of a wet sand filled beam (the dry sand results are similar). Note that many modes shown in the undamped beam are not visible in the drive point mobility of the sand filled beam. The sand adds mass to

the structure, but only negligible stiffness. Friction between the grains of sand provides damping; modes closely coupled to the sand are most effectively damped.

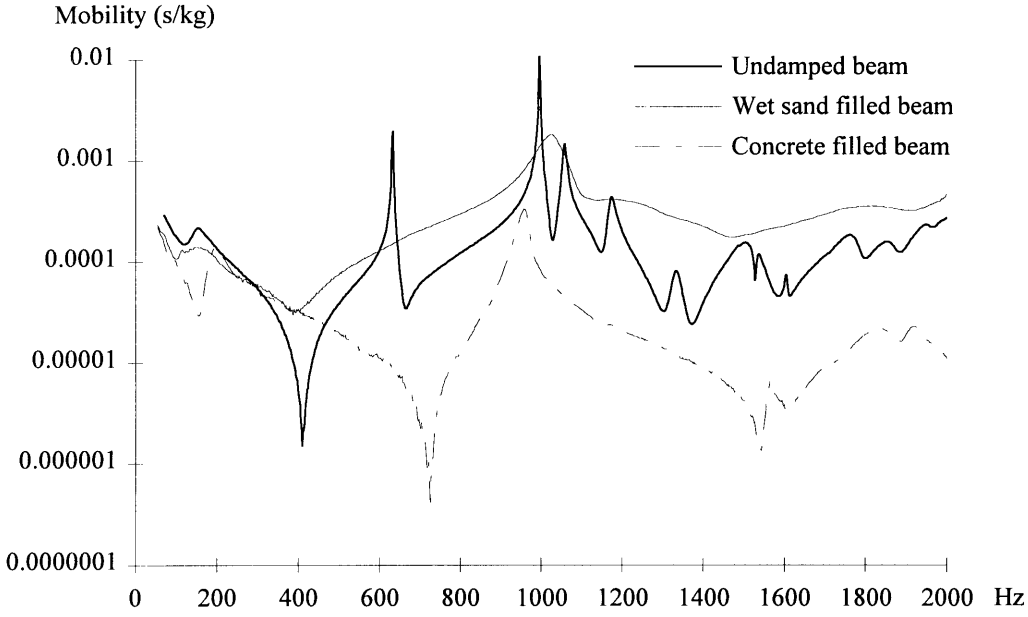

Figure 1.1 Dynamic response of an undamped and filled steel beams.

The concrete filled beam response shown in Figure 1.1 shows a substantial improvement over the undamped beam. Part of this improvement is due to the added stiffness of the concrete, which lowers the drive point mobility. The improved damping of the composite concrete and steel structure helps improve the response in the damping controlled regions near the modal peaks (loss factors around 0.02).

The frequency response functions shown in Figure 1.2 show the much larger improvement in dynamic performance of the shear damping mechanism developed for this dissertation. The undamped beam is again shown for reference. While both damped structures benefit from increased stiffness (the shear damped beam uses steel damping inserts), the figure shows how the damping is dramatically improved. The viscous fluid damped structure has modal loss factors of about 0.04, and the viscoelastically damped structure has modal loss factors around 0.10, depending on the vibration mode.

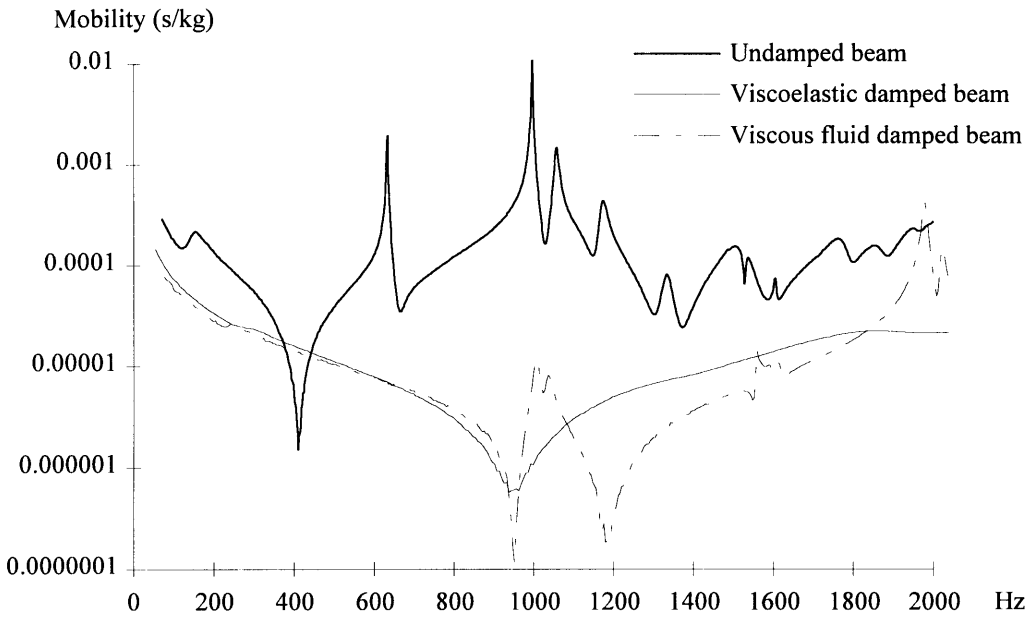

Figure 1.2 Dynamic response of undamped and damped steel beams.

Figure 1.3 shows the impulse responses of some of the same structures shown in Figure 1.2. These responses were drive point measurements taken from the same location on each of the four different structures.

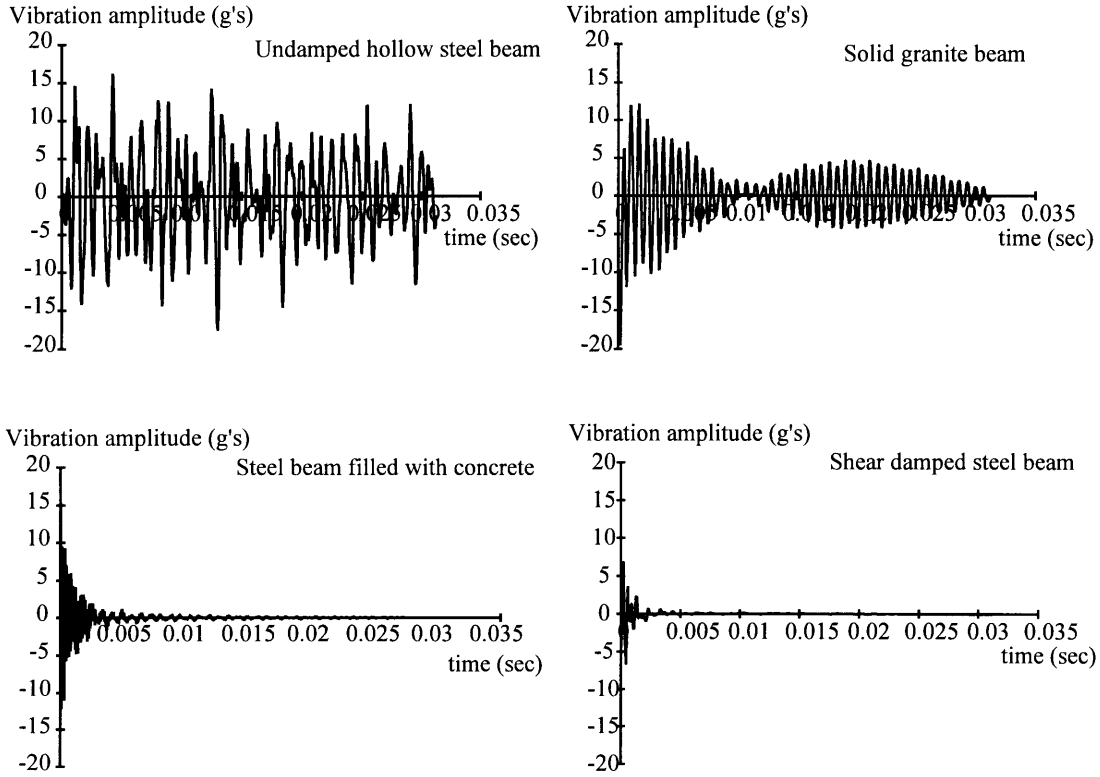

Figure 1.3 Time traces of structures shown in Figures 1.1 and 1.2.

The damping analysis presented in this dissertation will provide a complete set of equations for the design of structures such as the steel beam shown in this example. Methods of analytically investigating the damping in complex structures are traditionally based on the analysis of a three layer plate. While providing valuable insight into the damping problem, the three layer theory is not capable of handling arbitrary beam-like structures. This dissertation seeks to first identify the current state of the art in damping research and then extend it to more complicated applications with an improved damping mechanism and structural analysis technique.

1.4 Background

In 1959, two papers published in the *Journal of the Acoustical Society of America* sparked interest in the use of constrained laminates of lossy materials to reduce vibration in plates and beams. The first paper, by E. Kerwin, presented a closed-form analysis of the damping in a flat, three layer plate (plate, viscoelastic damping material, thin constraining layer). The second paper, by G. Kurtze, discussed an analysis method for estimating the damping that may be obtained by two parallel plates enclosing a viscous fluid. Both papers cited the two layer plate (a plate and a viscoelastic damping material layer) work done in the early 50's by H. Oberst and P. Liénard.

The seminal Kerwin paper has made a lasting impact on research on the constrained layer damping problem. The results of a closed-form analysis and experimental verification are presented in the paper, and virtually any work done in the field refers back to this article. The Kurtze paper, which presents an interesting approach to solving constrained fluid layer damping problems by analogy to electric circuit analysis, is of lesser importance today because: 1) we are no longer interested in solving differential equations with analog circuits, and 2) the fluid damping mechanism relies on the squeeze film effects between two parallel plates (with fluid layer thicknesses in the range of 0.5 to 5 cm!).

In 1959, the ASME Applied Mechanics Division published a book of papers given at the Structural Damping Colloquium. This book, edited by J. Ruzicka, contains a paper by D. Ross and E. Ungar and Kerwin presenting a more general analysis of the three layer beam damping problem (the constraining layer is no longer assumed to be thin with respect to the base plate). Since this publication, literally dozens of papers have re-visited the three layer damping analysis. These updates seek to expand the understanding and generality of the original Kerwin and Ross-Ungar-Kerwin work. Table 1.1 summarizes some of the important contributions.

Table 1.1 Summary of several important papers in constrained layer damping theory.

Author	Contribution to constrained layer theory
P. Liénard	Work on extensional damping: thin plate and damping layer, 1951.
H. Oberst	Work on extensional damping: thin plate and damping layer, 1952.
H. Plass	Three layer flexural damping: thin plate, damping layer, and thin constraining layer, 1957.
E. Kerwin	Three layer damping: arbitrary plate, damping layer, and thin constraining layer, 1959.
D. Ross, E. Ungar, & E. Kerwin	Three layer damping: arbitrary plate, damping layer, and arbitrary constraining layer, 1959.
R. DiTaranto & W. Blasingame	Five layer beam damping, 1965.
R. Plunkett and C. Lee	Use of discontinuous constraining layers, 1970.
M. Lalanne, M. Paulard, & P. Trumpette	Finite element modeling of viscoelastic damped structures, 1970.
B. Nakra	Review paper, 1975.
P. Torvik	Review paper, 1980.

Although the references cited in Table 1.1 all contain interesting details about the constrained layer damping problem, most of them use the fundamental approach taken by Ross, Ungar, and Kerwin, if not the simpler (and less general) analysis of Kerwin's own 1959 paper. The Ross, Ungar, and Kerwin analysis outlined in the following section explores this traditional analysis of constrained layer damping treatments.

1.5 RUK Constrained Layer Damping Theory

The following analysis is taken from *Damping of Plate Flexural Vibrations by Means of Viscoelastic Laminae*, presented at the 1959 ASME Structural Damping Colloquium. The analysis correctly predicts the stiffening of the undamped plate by the addition of the two damping laminates (one viscoelastic, one elastic). Once the full RUK analysis is summarized, two traditional beam configurations will be presented using simplifications of the general case.

A number of assumptions are made to make the problem solvable in closed-form:

1. The beam has simply-supported boundary conditions (this leads to purely sinusoidal mode shapes).

2. The beam is comprised of only three layers (other approximate techniques are available to estimate the damping in multi-layer configurations).

3. The viscoelastomer is modeled by a complex shear stiffness $G^* = G(1 + i\,\eta)$.

4. The elastic layers are maintained at constant spacing by the viscoelastic layer, even though the viscoelastic layer typically has a much lower elastic modulus.

The analysis is made possible by using a number of additional assumptions frequently made in beam vibrations work:

5. The beam has a wavelength sufficiently larger than its thickness.

6. The deflections of the structure are small enough such that the slope of the neutral axis is much less than unity.

The analysis begins by considering an element in the three layer plate. Figure 1.4 shows a close up detail of the three layers in bending. For our purposes, layer 1 is the base plate, layer 2 is the viscoelastic layer, and layer 3 is the constraining layer. As will be seen, any of the layers can have an arbitrary amount of damping; the analysis will provide an estimate of the damping in the composite plate.

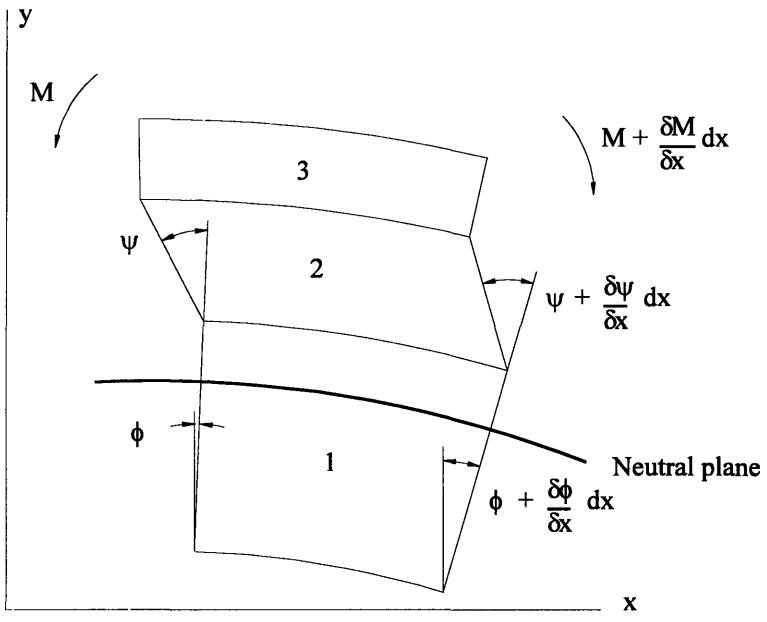

Figure 1.4 Section of the three layer beam model (after Ross, Ungar, and Kerwin, 1959).

The angle ϕ, as defined in the original RUK work, is the flexural angle of the base plate. The angle ψ is the shear strain in the middle layer. Notice that the angle ψ is defined in the opposite direction as the flexural angle. The neutral plane of the three layer

structure is displaced a distance D from the neutral plane of the base plate when the damping layers are applied. Figure 1.5 shows D and the notation of the other dimensions. Note that the thicknesses of the individual layers are H_i.

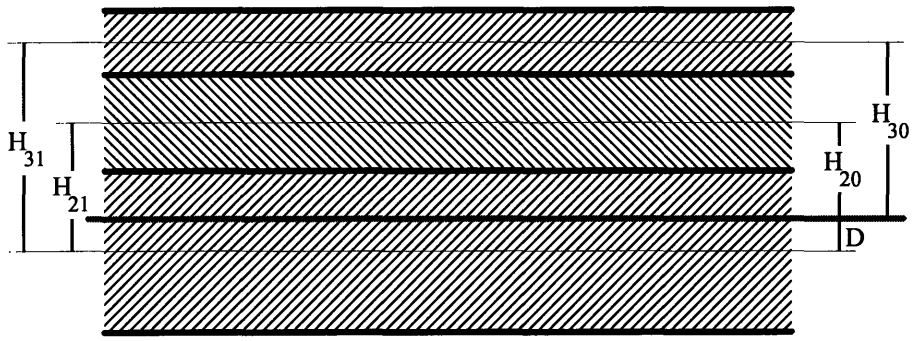

Figure 1.5 Dimensions of the three layer beam (after Ross, Ungar, and Kerwin, 1959).

The RUK analysis proceeds by considering the total moment acting on the three layer structure.

$$M = B\frac{\partial \phi}{\partial x} = \sum_{i=1}^{3} M_{ii} + \sum_{i=1}^{3} F_i H_{i0} \tag{1.1}$$

where:

$$\sum_{i=1}^{3} M_{ii} = K_1 \frac{H_1^2}{12}\frac{\partial \phi}{\partial x} + K_2 \frac{H_2^2}{12}\left(\frac{\partial \phi}{\partial x} - \frac{\partial \psi}{\partial x}\right) + K_3 \frac{H_3^2}{12}\frac{\partial \phi}{\partial x}$$

$$\sum_{i=1}^{3} F_i = K_1 H_{10}\frac{\partial \phi}{\partial x} + K_2\left(H_{20}\frac{\partial \phi}{\partial x} - \frac{H_2}{2}\frac{\partial \psi}{\partial x}\right) + K_3\left(H_{30}\frac{\partial \phi}{\partial x} - H_2\frac{\partial \psi}{\partial x}\right) \tag{1.2}$$

K is defined as the extensional stiffness per unit width of the plate ($K_i = E_i H_i$). The forces are obtained by integrating the stress over each layer in the beam element shown in Figures 1.4 and 1.5 (the stresses are continuous at each interface). The distance D is not known *a priori*, but may be solved for since we know that the total extensional force on the three layers must be zero. Setting the sum of the forces in the element to zero, we obtain an expression for D.

$$D = \frac{K_2 H_{21} + K_3 H_{31} - \left(\frac{K_2}{2} + K_3\right) H_2 \frac{\partial \psi}{\partial \phi}}{K_1 + K_2 + K_3} \tag{1.3}$$

The remaining unknown is the ratio between spatial derivatives $(\partial\psi/\partial x)/(\partial\phi/\partial x)$. Considering the shear stress acting on the middle layer, the strain may be written as a function of the shear modulus in the viscoelastic layer G_2 and the force on the upper face of the viscoelastic layer.

$$\psi = -\frac{1}{G_2}\frac{\partial F_3}{\partial x} \tag{1.4}$$

As a result of the sine wave mode shape assumption, the shear strain is related to its own second derivative by the bending wave number k_B.

$$\frac{\partial^2 \psi}{\partial x^2} = -k_B{}^2 \psi \tag{1.5}$$

The ratio $\partial\psi/\partial\phi$ is the same as the ratio between second derivatives because the three layers are undergoing vibration of the same wavelength. Using the constitutive relation for the shear strain, the previously stated expression for the force across the third layer, and the relationship between the shear strain and its second derivative, we can solve for the relationship $\partial\psi/\partial\phi$:

$$H_2 \frac{\partial \psi}{\partial \phi} = H_2 \frac{\left(\partial^2 \psi / \partial x^2\right)}{\left(\partial^2 \phi / \partial x^2\right)} = \frac{H_{31} - D}{1 + \dfrac{G_2}{k_B{}^2 K_3 H_2}} \tag{1.6}$$

The dimensionless quantity g (the shear parameter) gives an indication of the magnitude of the shear stress in the viscoelastic layer relative to the extensional stress in the beam.

$$g = \frac{G_2}{k_B{}^2 K_3 H_2} \tag{1.7}$$

We now have a complete closed-form expression for the flexural rigidity of the three layer beam. The equations have become too arithmetically involved to be substituted into one expression, so they are traditionally presented in the following format:

$$\begin{aligned} EI = {} & K_1 \frac{H_1^2}{12} + K_2 \frac{H_2^2}{12} + K_3 \frac{H_3^2}{12} + K_1 D^2 + K_2 (H_{21} - D)^2 + K_3 (H_{31} - D)^2 \\ & - K_2 \frac{H_2}{12}\left(\frac{H_{31} - D}{1+g}\right) - \left[\frac{K_2}{2}(H_{21} - D) + K_3 (H_{31} - D)\right]\left(\frac{H_{31} - D}{1+g}\right) \end{aligned} \tag{1.8}$$

where:

$$\begin{aligned} D &= \frac{K_2 (H_{21} - H_{31}/2) + g(K_2 H_{21} + K_3 H_{31})}{K_1 + K_2/2 + K_3 + g(K_1 + K_2 + K_3)} \\ H_{31} &= \frac{H_1 + H_3}{2} + H_2 \\ H_{21} &= \frac{H_1 + H_3}{2} \end{aligned} \tag{1.9}$$

These equations provide a framework for the design of three layer damping treatments. Often, the elastic modulus of the damping layer is so small that it can be assumed to be zero to reduce the complexity of the math.

1.6 Implementation of the RUK theory

The RUK equations are used by replacing the moduli of the three layers with complex moduli. In the most general case, all three layers may exhibit lossy behavior.

$$\begin{aligned} g &\rightarrow g(1+i\eta_\beta) \\ E_1 &\rightarrow E_1(1+i\eta_1) \\ E_2 &\rightarrow E_2(1+i\eta_2) \\ E_3 &\rightarrow E_3(1+i\eta_3) \\ E &\rightarrow E(1+i\eta) \end{aligned} \tag{1.10}$$

Using the substitutions listed in Equation (1.10), and the closed-form analysis results, the damping of a constrained layer damped structure of arbitrary geometry and material properties can be determined. However, the equations are too cumbersome to identify interesting trends without making some simplifying assumptions. Two cases will be considered: 1) the case of extensional (two layer) damping ($E_3 = H_3 = 0$), and 2) the case of a sandwich construction with a thin constraining layer.

1.6.1 Extensional damping ($E_3 = H_3 = 0$)

Extensional damping, which is the principal means of energy dissipation in a two layer structure (no constraining layer), was first investigated by Oberst and Liénard. This special case can be studied with the general three layer equations by noting that the shear parameter g goes to infinity as the height of the constraining layer goes to zero. This leads to a tremendous simplification of the equations. The result is an equation relating the bending stiffness of the composite two layer structure to the bending stiffness of the base plate.

$$\frac{EI}{E_1 I_1} = 1 + e_2 h_2^{\,3} + 3(1+h_2)^2 \frac{e_2 h_2}{1+e_2 h_2} \tag{1.11}$$

where:

$$\begin{aligned} e_2 &= \frac{E_2}{E_1} \\ h_2 &= \frac{H_2}{H_1} \end{aligned} \tag{1.12}$$

Replacing the viscoelastic modulus e_2 with the complex modulus $e_2(1 + i\,\eta_2)$ and the composite modulus E with $E(1 + i\,\eta)$, the dimensionless stiffness and damping can be found. The dimensionless quantities are given by a ratio of polynomials.

$$\frac{EI}{E_1 I_1} = \frac{1+4e_2h_2+6e_2h_2^{\,2}+4e_2h_2^{\,3}+e_2^{\,2}h_2^{\,4}}{1+e_2h_2}$$

$$\frac{\eta}{\eta_2} = \frac{e_2h_2\left(3+6h_2+4h_2^{\,2}+2e_2h_2^{\,3}+e_2^{\,2}h_2^{\,4}\right)}{(1+e_2h_2)\left(1+4e_2h_2+6e_2h_2^{\,2}+4e_2h_2^{\,3}+e_2^{\,2}h_2^{\,4}\right)} \tag{1.13}$$

The dimensionless quantities give important insight into the extensional damping problem. In the limit as h_2 goes to zero, the dimensionless stiffness approaches unity and the dimensionless damping approaches zero. In the case where h_2 goes to infinity, the dimensionless loss factor approaches unity and the dimensionless stiffness approaches infinity.

Considering an intermediate range of h_2, we can look at a typical viscoelastic material data sheet and investigate the performance of an extensional damping treatment. Using values from a Lord LD-400 material data sheet (taken at 200 Hz and 75 °C), the stiffness and loss factor of an extensionally damped plate may be plotted as a function of h_2 ($e_2 = 0.004$ and $\eta_2 = 0.6$).

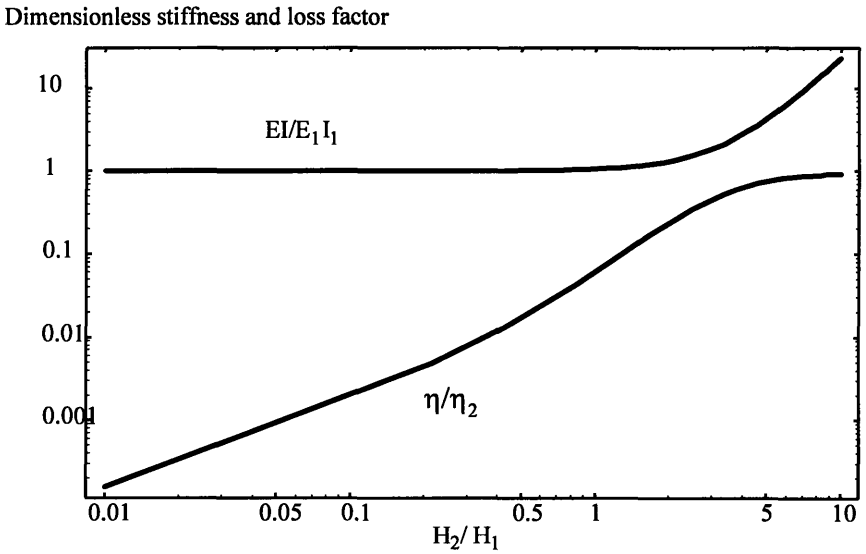

Figure 1.6 Dimensionless loss factor and bending stiffness of a 200 Hz beam and Lord LD-400 damping material.

As seen in Figure 1.6, the loss factor of the composite beam is low except when the viscoelastic layer is made very thick compared to the thickness of the base plate. In a real design application this approach is rarely acceptable. As will be shown in the next

section, constrained layer damping is almost always more effective than extensional (unconstrained) damping.

As a side note, the designer may notice that the simplified equations used to describe the performance of the extensional damping method are still rather cumbersome. Although the formulas are easily implemented with computer software, there is limited intuitive understanding to be obtained by inspection of the equations.

1.6.2 Damping in a Three Layer Treatment

To further illustrate the damping prediction capability of the RUK analysis, a second special case will be considered. In this case, the constraining layer is assumed to be thinner and of lower elastic modulus than the first and third layers ($K_3^2 << K_1^2$ and $E_2 = 0$). The results of this analysis correspond to the original Kerwin paper, where the constraining layer was assumed to be too thin to carry a bending moment.

Upon simplification, the loss factor can again be written as a dimensionless quantity using the loss factor of the viscoelastic material [Kerwin, 1959]:

$$\frac{\eta}{\eta_2} = \left(\frac{H_{31}}{1 + K_3 / K_1 \; g/(1+g)} \right)^2 \frac{K_3}{EI} \frac{g}{(1+g)^2} \tag{1.14}$$

where:

$$\frac{EI}{E_1 I_1} = 1 + 12 \frac{H_{31}^{\;2}}{H_1^2} \frac{g K_3 / K_1}{1 + g + g K_3 / K_1} \tag{1.15}$$

The shear parameter g involves two moduli. The first is the elastic modulus of the constraining layer; the second is the shear modulus of the viscoelastic material. The elastic modulus of the constraining layer is usually taken to have a loss factor of zero with negligible error. Therefore, the imaginary part of the shear parameter g is driven solely by the complex modulus of the shear in the viscoelastic layer.

A variety of manufacturers produce high loss viscoelastic materials and publish the complex modulus as a function of temperature and frequency. As the temperature of a viscoelastomer is lowered toward the material's glass transition temperature, the shear modulus tends to go up and the loss factor decreases. Increasing the shear rate has roughly the same effect; higher shear rates tend to stiffen the material and decrease the loss factor. Sample data sheets are given in Figures 1.7, 1.8, and 1.9 for the properties of 3M ScotchDamp ISD-112 [Nashif, Jones, and Henderson, 1985].

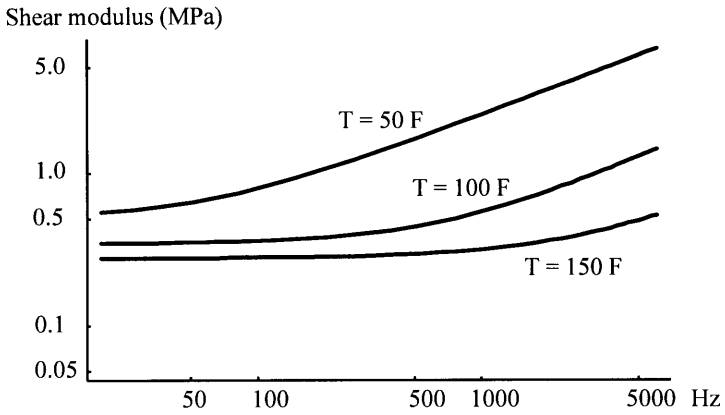

Figure 1.7 Shear modulus of 3M ISD-112 as a function of frequency and temperature.

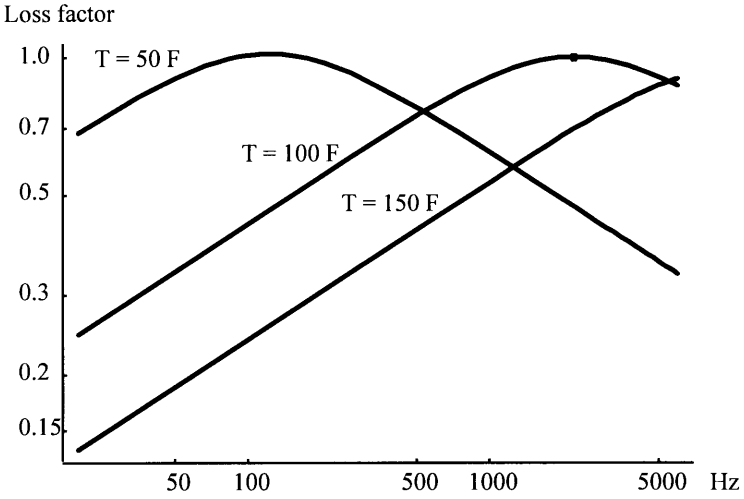

Figure 1.8 Loss factor of 3M ISD-112 as function of frequency and temperature.

The data shown in Figure 1.9 present the net damping capacity of the material, the imaginary part of the complex modulus. As before, this quantity is temperature and frequency dependent. Also plotted in Figure 1.9 is the damping capacity of a high viscosity silicone fluid, GE Silicone's Viscasil 600,000. While the fluid has lower damping capacity, it is relatively insensitive to temperature fluctuation.

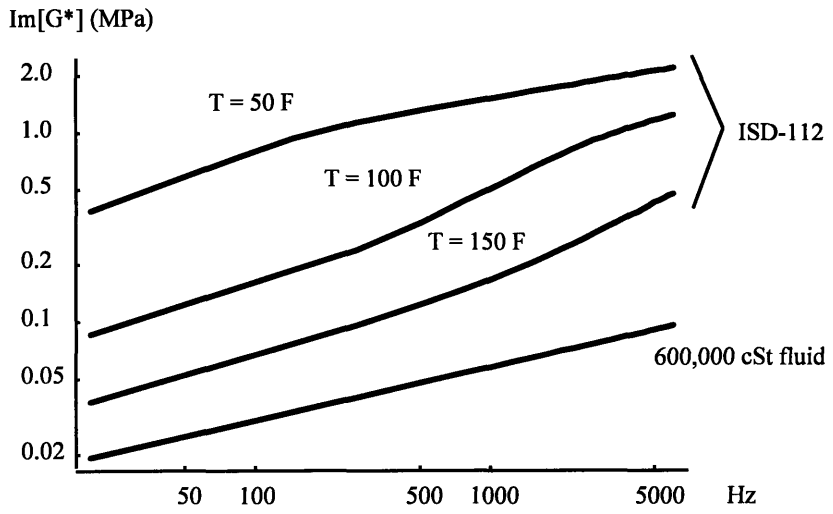

Figure 1.9 Imaginary part of the complex modulus of 3M ISD-112, Im[$G^* = G + i\,G\eta$].

In general, the designer will seek to maximize the amount of damping in a structure by considering the shear modulus and loss factor of a material and the desired thickness of the viscoelastic layer. The design formulas, while algebraically tedious, offer flexibility to allow the designer to select an optimum configuration.

For example, the natural frequency of a simply-supported steel beam 1 meter long and 4 cm thick is 93 Hz. Using the ScotchDamp data sheets shown above, we find that at room temperature, the loss factor of the viscoelastomer is 0.7 and the shear modulus is 0.45 MPa. Using the thin constraining layer design formulas, the loss factor of the damped plate may be plotted as a function of H_2 and H_3, as shown in Figure 1.10.

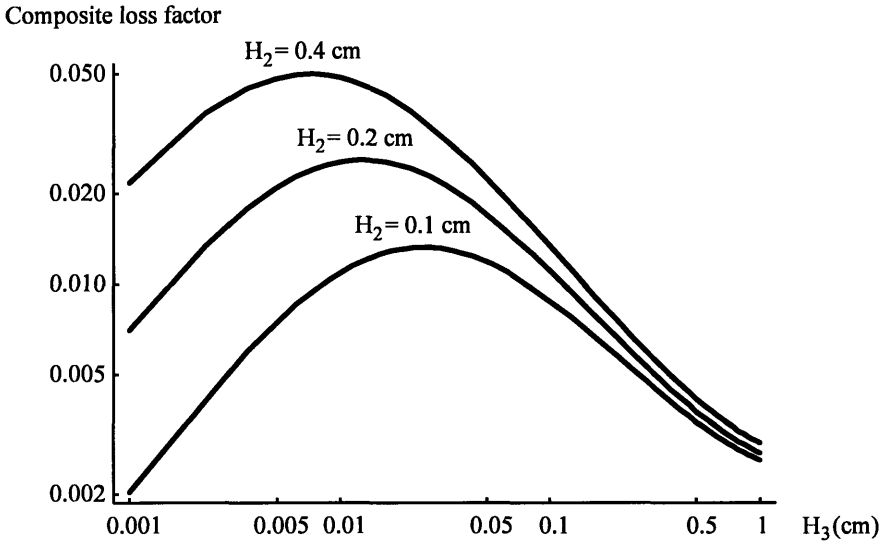

Figure 1.10 Loss factor of a 1 meter beam as a function of constraining layer thickness.

The loss factors predicted in Figure 1.10 show an important result of the RUK analysis: the damping is maximized by increasing the constraining (third) layer thickness. This effect is commonly referred to impedance matching. Simply stated, the effect of the constraining layer is optimized if its impedance (a measure of dynamic stiffness) is close to that of the base plate. In this example, we have assumed that the constraining layer is thin to simplify the equations, an assumption that does not allow us to investigate thicker constraining layers without appreciable loss of accuracy.

1.7 Comparison of RUK-Style Approach to the Proposed Analysis

As mentioned above, the shear damping mechanism developed for this dissertation uses the same viscoelastic materials to dissipate energy as the classic three layer damping treatment. The contribution of this dissertation is: 1) to recognize that the classic RUK constrained layer analysis is too specific to be used for typical structural damping problems, and 2) to provide a completely integrated approach to structural damping in complicated geometries. The three layer theory, while capable of rigorously modeling the interaction between the arbitrarily thick layers in a plate, does not provide any means of generalizing this capability to more complicated geometries.

The research outlined in this document considers the lessons learned from constrained layer theory and places the damping layers *inside* mechanical structures for improved robustness and ease of manufacture. An analysis is developed that can accurately predict how the geometric and material properties of the structure effect the total damping. Many case studies are presented highlighting the excellent performance of the shear damping mechanism.

The design engineering will appreciate the new analysis method because it accurately predicts the amount of damping without the complexity of the traditional constrained layer theory. This simplification is a result of considering the vibrating beam of arbitrary geometry: the interaction between the many components of the beam system cannot be obtained in a closed-form solution. This requires the use of a modal strain energy approach to the damping calculation. The resulting analysis accurately predicts the amount of damping in a structure for thick damping layers and offers an accurate empirical estimate of the optimal damping level (given geometric and material properties of the beam structure).

Figure 1.11 outlines the difference between the traditional constrained layer damping theory and the closed-form analysis developed for this dissertation.

	RUK Analysis	Proposed approach
Mode shape	pure sine wave	arbitrary mode shape
Structure geometry	3 layers of uniform width	arbitrary geometry
Optimal damping thickness	explicitly modeled	correction available

Figure 1.11 Comparison of classic RUK to the closed-form analysis developed herein.

An important note is that the finite element method can be used to accurately predict the damping of a structure regardless of the structural geometry or materials.

1.8 Conclusion

This chapter has discussed the motivation for investigating structural damping mechanisms. The need for damping in structures is ever-increasing, and research in this area has an opportunity to make a lasting impact on mechanical design. Dynamic responses of a steel beam were shown to highlight how traditional damping techniques compare to the shear damping mechanism. As shown, the shear damped structure shows a large improvement over other treatments without the typical high cost, weight, and thermal penalty.

The original damping work in constrained layer theory, which is now 40 years old, provides the foundation of this research. Recent work has shown that while multiple constrained layer treatments are one method of further increasing the damping in a plate, the effect of subsequent layers is much smaller because the viscoelastic layers added after the first are subject to increasingly small shear strains. The net effect is that multiple layer treatments are roughly equivalent to using a single, three layer system with the constraining layer thickness equal to the total thickness of the constraining layers used in a multi layer application [Nashif, Jones, and Henderson, 1985]. Examples of the RUK analysis have been provided so that the reader can understand the classic three layer damping theory.

In comparison, the shear damping mechanism of this dissertation is much more convenient to design and manufacture with the added advantage that any beam-like structure may be optimized with a simple closed-form solution. The remainder of this work will include a detailed discussion of the shear damping mechanism and the closed-form analysis. A finite element model of the damping mechanism is included for investigation of the performance of structures other than beams. The manufacturing issues of shear damped structures are also addressed and numerous analytical and experimental examples are presented to verify the theory.

1.9 References

DiTaranto, R. A., and W. Blasingame, *Composite Loss Factors of Selected Laminated Beams*, Journal Acoustical Society of America, Vol. 31, September 1959.

Inman, Scott, and Mitchell K. Enright, *Vibration Damped Apparatus*, U. S. Patent No. 4,706,788, Nov. 17, 1987.

Jones, David I. G., *Application of Damping Treatments*, Shock and Vibration Handbook, Third Edition, McGraw-Hill, New York, 1988.

Kurtze, Gunther, *Bending Wave Propagation in Multi-Layer Plates*, Journal Acoustical Society of America, Vol. 31, September 1959.

Lalanne, M., M. Paulard, and P. Trumpette, *Response of Thick Structures Damped by Viscoelastic Material with Application to Layered Beams and Plates*, Shock and Vibration Digest, Vol. 45, Part 5, June 1975.

Liénard, P., *Etude d'une Méthod de Mesure du Frottement Intérieur de Revêtements Plastiques Travaillant en Flexion*, La Recherche Aéronautique, Vol. 20, 1951.

Nakra, B. C., *Vibration Control with Viscoelastic Materials*, Shock and Vibration Digest, Vol. 8, No. 6, June 1975.

Nashif, Ahid D., David I. G. Jones, and John P. Henderson, Vibration Damping, John Wiley and Sons, New York, 1985.

Oberst, H., *Ueber die Dämpfung der Biegeschwingungen dünner Bleche durch fest haftende Beläge*, Acustica, Vol. 2, Akustische Beihefte No. 4, 1952.

Plass, H. J. Jr., *Damping of Vibrations in Elastic Rods and Sandwich Structures by Incorporation of Additional Visco-Elastic Material*, Proc. Third Midwestern Conf. Solid Mechanics, Univ. of Michigan, April 1957.

Plunkett R., and C. T. Lee, *Length Optimization for Constrained Viscoelastic Layer Damping*, Journal of the Acoustical Society of America, Vol. 48, No. 1, Part 2, 1970.

Ross, Donald, Eric Ungar, and E. M. Kerwin, *Damping of Plate Flexural Vibrations by Means of Viscoelastic Laminae*, Proc. Colloq. Structural Damping, ASME, 1959.

Torvik, P. J., *The Analysis and Design of Constrained Layer Damping Treatments*, Damping Applications for Vibration Control, AMD-Vol. 38, ASME, 1980.

Chapter 2: Development of Shear Damping Theory

2.1 Introduction

This section documents the development of a mathematical model of the shear damping mechanism in flexural wave vibration. The model will be derived for an Euler beam with arbitrary cross section, and many examples with simple geometries will be presented.

Several assumptions are necessary to arrive at a closed-form solution. The beam is assumed to be sufficiently long and slender to neglect the shear and rotary inertia effects included in the Timoshenko beam model. The resulting differential equation of motion may be solved closed-form given various boundary conditions such as pin-pin, free-free, and clamped-free (cantilever). The solution of the differential equation of motion includes the bending mode shapes and undamped natural frequencies. These results will be used in the development of the damping model. The sinusoidal mode shape of a simply-supported beam will be used in the presentation of the closed-form solution, but solutions for other beam boundary conditions are also summarized. The closed-form analysis may be applied to estimate the loss factor of virtually any beam-like structure.

Like the damping model used in the Ross-Ungar-Kerwin analysis, the shear damping mechanism uses a sandwich type construction to shear a lossy material. Viscous fluids with approximately Newtonian behavior and viscoelastic materials with complex shear moduli can be used in the analysis. Several manufacturers of highly viscous fluids market materials that adhere very closely to the Newtonian model. These fluids have been used in the experimental verification of the closed-form solution. Viscoelastic materials are available in a wide range of shear moduli, but are highly frequency and temperature sensitive. Furthermore, the complex shear modulus model for viscoelastic materials is only an approximation of the actual behavior. For this reason, many of the experiments in this dissertation are run with viscous fluids. In practice, the viscoelastic materials will usually be preferred because of their superior lossiness.

The derivation predicts the first bending mode loss factor with negligible error when compared to experimental results. The assumptions that have been made in the analysis are considered to be reasonable given the excellent agreement with experimental measurements.

Two vibrating beam geometries will be considered throughout the following analysis. The first is the general case of an arbitrary beam vibrating with pin-pin boundary conditions. Figure 2.1 shows a possible vibrating beam structure.

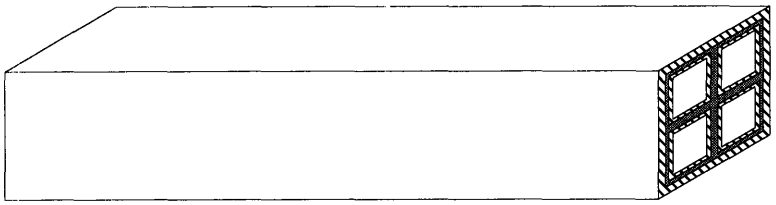

Figure 2.1 Arbitrary beam used in damping analysis.

The specific case of two beams vibrating with a shear layer between them is also used to simplify the geometry of the general analysis. Although the results of the specific case are applicable to just one beam configuration, the results show several critical trends that must be understood when designing structures with the shear damping mechanism. Figure 2.2 shows the beam geometry used in the specific case.

Figure 2.2 Simple beam geometry used in specific beam examples.

Subsequent chapters will explore several advanced topics that concern the shear damping mechanism, including:

1. Structural coupling between beam components when the damping layer is thin.

2. Optimal shear layer thickness.

3. Finite element modeling of the damping mechanism and correlation to experimental results.

4. Performance of the damping mechanism in complex beam and plate structures.

2.2 The Euler Beam Model

Variational calculus will be used to arrive at the equations of motion for an Euler beam. By neglecting the shear and rotary inertia effects, the derivation of the beam bending equation will result in the Euler beam equation. This assumes that the beam has uniform material properties and cross section, as well as a high aspect ratio.

This method requires the integral form of the potential energy and kinetic co-energy in the beam as well as the work done by forces acting on it. In Equation (2.1), the kinetic co-energy is represented by the integral T^* and the potential energy is V.

$$T^* = \int_0^L \frac{1}{2}\rho A\left(\frac{\partial y}{\partial t}\right)^2 dx \quad V = \int_0^L \frac{1}{2} EI\left(\frac{\partial^2 y}{\partial x^2}\right)^2 dx \quad \text{Work} = \int_0^L f\,\delta y\,dx \tag{2.1}$$

These energy terms are substituted into the Variational Indicator, which leads to the equation of motion for the dynamic system. The Variational Indicator may be written in integral form [Crandall, 1968]:

$$V.I.=\int_{t_1}^{t_2}\left(\delta T^{*}-\delta V+f\,\delta y\right)dt \tag{2.2}$$

The following equation is obtained upon substituting the energy integrals of the vibrating beam into the Variational Indicator.

$$V.I.=\int_{t_1}^{t_2}\left[\int_0^L \delta\left\{\frac{1}{2}\rho A\left(\frac{\partial y}{\partial t}\right)^2-\frac{1}{2}EI\left(\frac{\partial^2 y}{\partial x^2}\right)^2\right\}dx+\int_0^L f\delta y\,dx\right]dt \tag{2.3}$$

The variational operator δ works similarly to the differential operator with the additional property represented in Equation (2.4).

$$\delta\frac{\partial y}{\partial t}=\frac{\partial \delta y}{\partial t} \tag{2.4}$$

The following equation results from taking the variations of the energy terms in the indicator.

$$V.I.=\int_{t_1}^{t_2}\int_0^L\left\{\rho A\frac{\partial y}{\partial t}\frac{\partial \delta y}{\partial t}-EI\frac{\partial^2 y}{\partial x^2}\frac{\partial^2 \delta y}{\partial x^2}+f\,\delta y\right\}dx\,dt \tag{2.5}$$

Equation (2.5) can be rewritten using integration by parts.

$$V.I.=\int_{t_1}^{t_2}\int_0^L\left\{-\rho A\frac{\partial^2 y}{\partial t^2}-EI\frac{\partial^4 y}{\partial x^4}+f\right\}\delta y\,dx\,dt-\int_{t_1}^{t_2}EI\frac{\partial^2 y}{\partial x^2}\frac{\partial \delta y}{\partial x}dt\Big|_0^L+\int_{t_1}^{t_2}EI\frac{\partial^3 y}{\partial x^3}\delta y\,dt\Big|_0^L \tag{2.6}$$

Setting the left hand side of the equation to zero and considering only geometrically admissible motions of the system means that each term of Equation (2.6) must equal zero to satisfy the Variational Indicator. The first integrand is the equation of motion of the Euler beam:

$$\rho A\frac{\partial^2 y}{\partial t^2}+EI\frac{\partial^4 y}{\partial x^4}=f(x,t) \tag{2.7}$$

The second and third integrands are the boundary conditions of the beam.

$$0=EI\frac{\partial^2 y}{\partial x^2}\frac{\partial \delta y}{\partial x}\Big|_0^L \quad \text{and} \quad 0=EI\frac{\partial^3 y}{\partial x^3}\delta y\Big|_0^L \tag{2.8}$$

Now that the equation of motion and boundary condition equations have been found, specific beam boundary conditions can be considered.

2.2.1 Solution to the Pin-Pin Beam Equation of Motion

This analysis uses the Euler equation to solve for the transverse displacement of the beam in bending vibration as a function of time and position. Considering a simply-supported beam of known equivalent stiffness EI and linear density ρA, the equation of motion may be written as:

$$\rho A\frac{\partial^2 y}{\partial t^2}+EI\frac{\partial^4 y}{\partial x^4}=f\,(x,t) \quad \text{where } y(0)=y(L)=y''(0)=y''(L)=0 \tag{2.9}$$

Such a beam is shown in Figure 2.3.

Figure 2.3 Model of a simply-supported Euler beam.

The displacement $y(x, t)$ can be considered as a product of two functions, one displacement dependent and one time dependent. The solution to the eigenproblem posed by this equation may be obtained by setting the force equal to zero and treating the time dependent behavior as a function of the complex exponential $e^{i\omega t}$.

$$y(x,t) = Y(x)Y(t) = Y(x)e^{i\omega t} \tag{2.10}$$

$$-\omega^2\rho AY(x)+EIY(x)^{iv}=0 \tag{2.11}$$

The mode shape $Y(x)$ of the Euler beam takes the form of a fourth order differential solution. The wave number k_B represents the ratio of the beam parameters that result from solving the differential equation for the mode shape $Y(x)$ (k_B is the mode number dependent wave number).

$$Y(x) = A_1 \sin k_B x + A_2 \cos k_B x + A_3 \sinh k_B x + A_4 \cosh k_B x \tag{2.12}$$

where:

$$k_B{}^4 = \omega^2 \frac{EI}{\rho A} \tag{2.13}$$

The A_i's are a function of the beam boundary conditions. For a simply-supported beam, the boundary conditions at $x = 0$ may be used to quickly find that A_2 and A_4 are 0.

$$\begin{aligned} Y(0) &= 0 = A_2 + A_4 \\ Y''(0) &= 0 = -A_2 + A_4 \end{aligned} \tag{2.14}$$

The boundary conditions at $x = L$ may be used to find A_1 and A_3. Note that the sine function is periodic with π, but that the sinh function is not.

$$Y(L) = 0 = A_1 \sin\beta L + A_3 \sinh\beta L$$
$$Y''(L) = 0 = -A_1 \sin\beta L + A_3 \sinh\beta L \tag{2.15}$$

A_3 must be 0 to satisfy both equations, but sin $k_B L$ will satisfy the equations regardless of the value of A_1, provided that k_B is a multiple of π.

$$k_B L = n\pi \tag{2.16}$$

The natural frequencies of the simply-supported beam may now be determined by eliminating the wave number.

$$\omega_n = \left(\frac{n\pi}{L}\right)^2 \sqrt{\frac{EI}{\rho A}} \tag{2.17}$$

The mode shape of simply-supported bending vibration is simply a sine function (shown with an arbitrary scaling function A_1).

$$Y(x) = A_1 \sin k_B x \tag{2.18}$$

The magnitude of A_1 is arbitrary, but is traditionally scaled so that the modes are *normalized* to unit mass:

$$\int_0^L \rho A\, Y_m(x) Y_n(x) = \delta_{mn} \quad \begin{cases} \delta_{mn} = 1, m = n \\ \delta_{mn} = 0, m \neq n \end{cases} \tag{2.19}$$

Evaluation of Equation (2.19) gives the normalized mode shape of a simply-supported beam:

$$Y(x) = \sqrt{\frac{2}{\rho AL}} \sin k_B x \tag{2.20}$$

The mode shape and natural frequencies of other boundary conditions may also be determined using the same approach. The difficulty is that for other boundary conditions, the sinh and cosh functions do not drop out, leaving more complicated equations. Figure 2.4 summarizes the results for the first bending modes in pin-pin, free-free, and clamped-free beams. Blevins' Formulas for Natural Frequency and Mode Shape contains a longer list [Blevins, 1979].

	ω_1	$Y(x)$
clamped-free	$\frac{1.875^2}{L^2}\sqrt{\frac{EI}{\rho A}}$	$\cosh\frac{1.875x}{L}-\cos\frac{1.875x}{L}-0.734\left(\sinh\frac{1.875x}{L}-\sin\frac{1.875x}{L}\right)$
pin-pin	$\frac{\pi^2}{L^2}\sqrt{\frac{EI}{\rho A}}$	$\sin\frac{\pi x}{L}$
free-free	$\frac{4.73^2}{L^2}\sqrt{\frac{EI}{\rho A}}$	$\cosh\frac{4.73x}{L}+\cos\frac{4.73x}{L}-0.983\left(\sinh\frac{4.73x}{L}+\sin\frac{4.73x}{L}\right)$

Figure 2.4 Mode shape and natural frequency of common beam boundary conditions.

2.3 Closed-Form Solution to Damping Factor

A closed-form solution to the loss factor of a shear damped structure may be developed from an assumed mode shape of vibration. The mode shape of slender beams that were determined in the previous section will be used in this analysis.

Two possible shear materials have been evaluated in this research: viscous fluids and viscoelastic solids. Although these are decidedly different classes of materials with different mathematical models, both allow closed-form estimation of the damping in a structure (given a simple transformation in the model). The damping derivation shows that viscous fluid shear layers are fully developed and laminar (resulting in a linear fluid velocity profile). Viscoelastic materials are assumed to have a linear strain profile across their cross section. The analysis proceeds assuming that the damping layer is a viscous fluid. The transformation that allows consideration of viscoelastic materials is given at the end.

The second assumption is that the damping layer is thick enough that it does not dynamically stiffen the structure. As the layer becomes thinner, the damping medium will begin to act as a solid, and in the limiting case of zero thickness, the structure is integral with the shear members. In this case, there is no strain and therefore no damping in the layer, not infinite damping as predicted by the solution. A correction for this assumption will be presented in Chapter 3.

Figure 2.5 shows a sketch of how the shear mechanism works. As the beam vibrates, the damping material undergoes a periodic reversal. The shearing action dissipates energy at a rate proportional to the fluid viscosity and inversely proportional to the fluid film thickness. The following derivation will calculate the total energy in a vibrating beam as well as the energy dissipated per cycle in the shear layers to estimate the loss factor of an arbitrary beam.

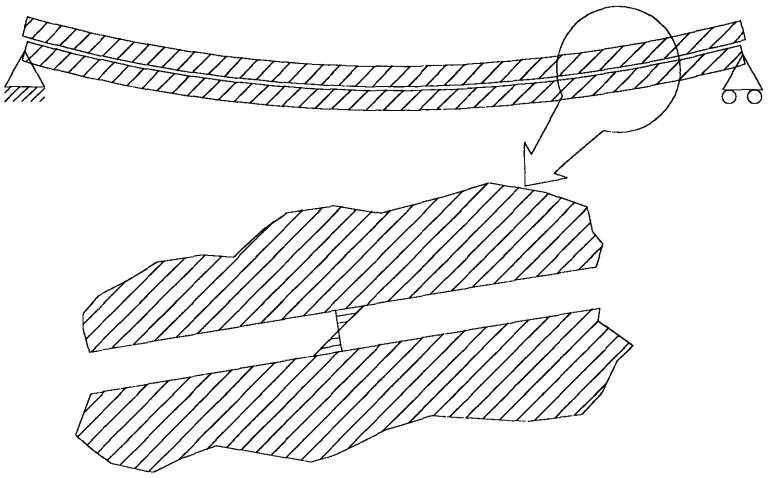

Figure 2.5 Schematic of the damping layer profile in a shear damped beam.

Figure 2.6 shows a sample shear damped structural member. The structure is an elastic element in bending vibration along its longest dimension. A structure of this sort often has a hollow cross section to save material and weight which allows the shear mechanism to be built into the structure.

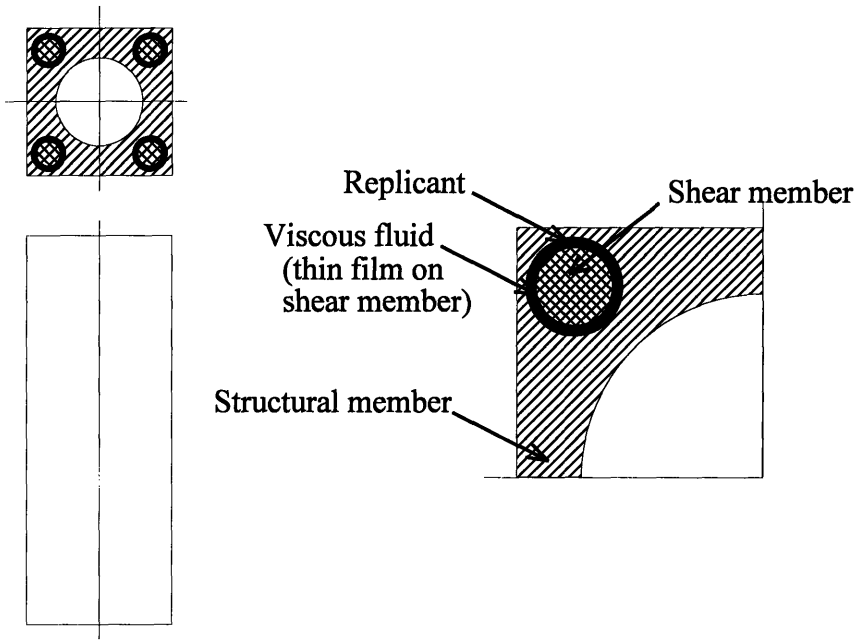

Figure 2.6 Possible shear damper configuration.

The displacement of the beam as a function of time and position (for the first bending mode) was shown for the simply-supported beam to be:

$$y(x,t) = Y \sin k_B x \, e^{i\omega t} \tag{2.21}$$

The modal amplitude Y is determined by the specific load case. In general, $y(t)$ can be written as a summation of many modes, each with its own amplitude. The value of the Y_i must be calculated to find the dynamic response of a beam to an arbitrary forcing function. In our case, we are only considering the first mode of vibration. The fluid damping loss factor derivation does not require that the Y be explicitly determined when calculating the loss factor of a single mode. This means that the damping loss factor may be conveniently estimated for any structure for which the mode shapes are known without a precise knowledge of the loading. For most loading cases this is a safe assumption because higher modes makes a relatively small contribution to the total displacement.

The damping factor analysis will proceed using the relationship between the total energy and the energy dissipated in one cycle of vibration. First, the total energy of the vibrating beam and then the energy dissipated in the damping layer will be derived. The loss factor is found from their ratio:

$$\eta = \frac{W_{diss}}{2\pi W_{total}} \tag{2.22}$$

The machine designer may be more comfortable with the quality factor Q, which gives the amplification at resonance. For a second order system, the quality factor is simply the inverse of the loss factor.

2.3.1 Nomenclature

$\delta(x)$ = longitudinal extension along beam

Y = amplitude of the first mode

ω = natural frequency of first mode

ρ_m = density of the m^{th} beam material

A_m = cross sectional area of the m^{th} beam member

E = modulus of beam material

L = length of beam

A_f = cross sectional area of fluid film

I = area moment of inertia of beam

c_i = distance from neutral axis of i^{th} shear member

p_i = cross sectional perimeter of i^{th} shear member

μ = dynamic viscosity of shear fluid

η = loss factor ($\eta = 2\zeta$, ζ = damping factor)

2.3.2 Calculation of Total Energy in the Beam System

At resonance, the total energy in the beam system may be calculated as the maximum amplitude of either the kinetic and potential energy.

$$W_{total} = \left| \frac{\rho A_{tot}}{2} \int_0^L \left(\frac{\partial y}{\partial t} \right)^2 dx \right|_t = \frac{\omega^2 Y^2}{4} \sum_m \rho_m A_m L \tag{2.23}$$

The summation is necessary because an arbitrary vibrating beam will have a total of m vibrating components. These components include the structure, the shear members, and any other vibrating element in the system.

2.3.3 Calculation of Energy Dissipated Per Cycle

The damping mechanism dissipates energy by shearing a Newtonian fluid. The fluid flow between the structural members is similar to the classic Couette flow problem. The solution is obtained by determining the longitudinal velocity of the shear members relative to the structural members and using this to calculate the dissipated energy.

The mode shape of a structure gives the transverse displacement of the structure as a function of time and position. The axial strain in the beam is related to the transverse displacement by:

$$\varepsilon_x = \frac{\partial \delta(x,t)}{\partial x} = -c \frac{\partial^2 y(x,t)}{\partial x^2} \tag{2.24}$$

This differential equation may be integrated to yield the longitudinal displacement $\delta(x)$ at a distance c from the bending neutral axis. The shear and structural members are assumed to be free floating so that the relative displacement is zero at the midpoint of the beam (at $x = L/2$). This configuration yields the least amount of damping in a given beam. Other designs are possible to increase the loss factor by fixing the displacement at one end.

Equation (2.25) shows the longitudinal displacement and its time derivative U at c.

$$\begin{aligned} \delta(x,t) &= -Yck_B \cos k_B x \, e^{i\omega t} \\ U(x,t) &= -iY\omega ck_B \cos k_B x \, e^{i\omega t} \end{aligned} \tag{2.25}$$

The net longitudinal deformation can be calculated by the superposition of the individual deflections between the core beam and the shear members. The longitudinal deflections at a structural/shear member interface will typically add because one face is in tension and the other is in compression.

The two-dimensional Navier Stokes equation in Cartesian coordinates will be used to find the velocity profile of the fluid between the shearing structural members using the longitudinal displacements as boundary conditions [Potter and Foss, 1982].

$$\frac{\partial u}{\partial t}+u\frac{\partial u}{\partial x}+v\frac{\partial u}{\partial y}=-\frac{1}{\rho}\frac{\partial p}{\partial x}+\nu\left[\frac{\partial^2 u}{\partial x^2}+\frac{\partial^2 u}{\partial y^2}\right] \tag{2.26}$$

The continuity equation for two-dimensional flow is also used in the derivation.

$$\frac{\partial u}{\partial x}+\frac{\partial v}{\partial y}=0 \tag{2.27}$$

A dimensional analysis of the Navier Stokes and continuity equations will allow insignificant terms to be identified. In the x-direction, the characteristic length may be taken to be the beam length L. The y-direction characteristic length is the fluid film thickness h. The characteristic velocity is U, the difference in velocity of the upper and lower boundary conditions. Substituting these parameters into the Navier Stokes and the continuity equations will highlight the most significant terms:

$$\frac{U}{L}+\frac{V}{h}=0 \tag{2.28}$$

$$\frac{U}{\tau}+U\frac{U}{L}+V\frac{U}{h}=-\frac{1}{\rho}\frac{\partial p}{\partial x}+\nu\left[\frac{U}{L^2}+\frac{U}{h^2}\right] \tag{2.29}$$

The dimensional analysis shows that the velocity gradient in the y-direction is critical (because the fluid film thickness is much smaller than the beam length). Furthermore, the left hand side terms may be neglected provided that three conditions are met (τ is the characteristic time constant: the period of vibration).

$$\frac{h}{L}<<1,\ \frac{h^2}{\tau\nu}<<1,\ \text{and}\ \frac{Uh^2}{L\nu}=\frac{d}{L}\frac{h^2}{\tau\nu}<<1 \tag{2.30}$$

The first condition is easily met by a typical shear damping design; the fluid thickness is on the order of microns and the length of the beam is on the order of meters. The second condition is also satisfied because the kinematic viscosity is on the order of one (SI units) for highly viscous silicone fluids. The time constant of vibration is small since most structures have high natural frequencies; however, the square of the fluid film thickness is much smaller, satisfying the condition.

The final condition results in the additional requirement that the amplitude of the longitudinal displacement be less than the length of the beam, which is clearly satisfied. In practice, the three dimensionless ratios are typically smaller than 1 part per million in a typical beam design. This allows a large simplification of the Navier Stokes equation.

$$\frac{1}{\mu}\frac{\partial p}{\partial x}=\frac{\partial^2 u}{\partial y^2} \tag{2.31}$$

The analysis will proceed by "flattening out" the vibrating beam into a one-dimensional flow field. This is accomplished by considering the upper and lower boundary conditions and ignoring their vertical component. This assumption is safely made because we have already shown that the velocity profile is highly dependent on the horizontal, not vertical component.

Because the beam is symmetric about its half-length $L/2$, there is no flow across the midpoint. Figure 2.7 shows the setup of the boundary conditions in the first half of the beam's length. The reference frame of the model is attached to the midpoint of the structural member.

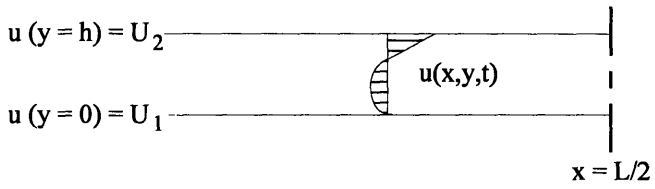

Figure 2.7 Model of flow in the shear layer.

The velocity profile may be found be double integrating the Navier Stokes equation and imposing the boundary conditions as shown in Figure 2.7. The velocity profile is symmetric about the midpoint so it is only necessary to consider one half of the beam.

$$u(x,y,t) = \frac{1}{2\mu}\frac{\partial p}{\partial x}\left(y^2 - yh\right) + \frac{(U_2 - U_1)y}{h} + U_1 \tag{2.32}$$

The mass conservation law will now be invoked to solve for the unknown pressure gradient $\partial p/\partial x$. Considering the "flattened" flow field, the most convenient control volume is a deformable volume that starts at a given location on the upper and lower surfaces and includes the beam midpoint at $L/2$. The fluid layer thickness is unchanging with time or position, as will be shown in the next chapter. Figure 2.8 shows the deformable control volume.

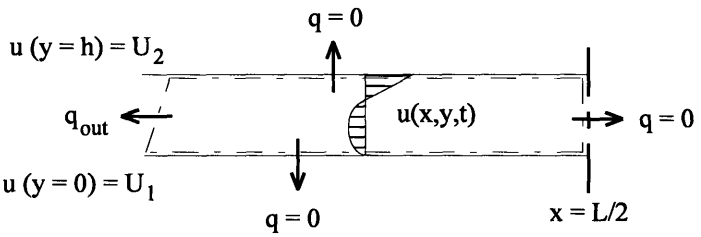

Figure 2.8 Deformable control volume used in fluid flow analysis.

Inspecting the two halves of a simply-supported beam, the flow across the midpoint of the beam is seen to be zero by symmetry. The flow across the top and bottom is also zero. Therefore, the net flow across the left hand side of the control volume must be

offset by the change in size of the control volume. Integrating the flow u over the height h of the fluid layer, the change in size of the control volume exactly equals the flow across the boundary (resulting in a pressure gradient of zero).

$$\frac{\partial p}{\partial x} = 0 \tag{2.33}$$

The fluid velocity profile may then be written as a function of the fluid film thickness and the differential velocity between the upper and lower surfaces:

$$u(x, y, t) = \frac{(U_2 - U_1)y}{h} + U_1 \tag{2.34}$$

where

$$\begin{aligned} U_1(x,t) &= i\omega\delta_1 = -Yc_1k_B \cos k_B x \; i\omega e^{i\omega t} \\ U_2(x,t) &= i\omega\delta_2 = Yc_2k_B \cos k_B x \; i\omega e^{i\omega t} \end{aligned} \tag{2.35}$$

The power dissipated in the fluid is given by the integral form of the viscous energy dissipation function [Özisik, 1985]:

$$\Pi_{diss} = \mu \int_V \left(\frac{\partial u}{\partial y}\right)^2 dV = \frac{\mu L A_f}{2h^2}\left((c_1 + c_2) Y k_B \omega e^{i\omega t}\right)^2 \tag{2.36}$$

The work dissipated per cycle can be determined by integrating the power dissipated over the period of oscillation:

$$W_{diss} = \Pi_{diss}\frac{\pi}{\omega} = \frac{\pi\mu L A_f \omega}{2h^2}\left((c_1 + c_2) Y k_B\right)^2 \tag{2.37}$$

2.3.4 Calculation of Damping Factor

The damping factor of the beam system may now be calculated using the total and dissipated energy. The contribution of the i damping members may be summed up to determine the total damping in a structural design.

$$\eta = \frac{W_{diss}}{2\pi W_{total}} = \frac{\mu k_B^{\,2}}{\omega \sum_m \rho_m A_m} \sum_i \frac{P_i c_i^{\,2}}{h_i} \tag{2.38}$$

This result shows several important trends with the shear damping mechanism. The most obvious is the direct relationship between the fluid viscosity and the damping factor. Another result is the inverse proportionality between the fluid film thickness and the amount of damping. From this closed-form solution, the damping appears to be infinite at zero fluid thickness. In practice, the fluid becomes so dynamically stiff at small thicknesses that the damping reaches an optimal limit. When the layer is made thinner

than this optimal limit, the damping decreases again as the relative displacement between the shearing members decreases. This phenomenon is documented in the next chapter.

2.3.5 The Damping Factor of a Simple Beam

Now that a closed-form solution is available for the damping ratio in a shear damped beam of arbitrary configuration, a simpler case will be considered to illustrate some important trends. Figure 2.9 shows the design of the double beam used in this exercise. The height of each beam is t, the width is b, and the length is L.

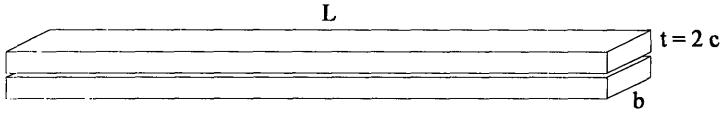

Figure 2.9 Simple beam geometry used to show damping trends.

Each beam will vibrate with the same amplitude and in phase with the other beam. The neutral axis of bending will be the centerline of each individual beam. Therefore, the quantity $c = t/2 + t/2 = t$. The first bending mode will be considered so $n = 1$. The perimeter of the viscous area is one width of the two beams b. Finally, the product of the density and the area (lineal density) is $2\rho bt$.

$$\eta = \frac{\mu k_B{}^2 t}{2\rho h \omega_n} \tag{2.39}$$

The natural frequency of a simply-supported beam is known and can be substituted into the equation for η to yield a simple expression given this specific geometry:

$$\omega_n = k_B{}^2 \sqrt{\frac{EI}{\rho A}} \tag{2.40}$$

$$\eta = \frac{\sqrt{3}\mu}{h\sqrt{E\rho}} \tag{2.41}$$

This result reveals several important issues to be considered when designing a simple fluid damped system:

1. The amount of damping is not independent on the length of the system.

2. The dominant terms are the fluid viscosity and fluid layer thickness. The more viscous the fluid and the thinner the layer, the higher the damping.

3. The damping factor also depends on the material properties of the structure.

An aluminum structure will be better damped than a steel structure of the same geometry (given a specific fluid layer film thickness). Similarly, a plastic structure will be better damped than a aluminum structure. This result is analogous to the loss factor in classic viscoelastic theory where the amount of damping is inversely proportional to the stiffness of the structure [Kerwin, 1959].

2.3.6 Damping in Beams of Other Boundary Conditions

The preceding section outlined the analysis for the loss factor of a simply-supported beam. This section will present the results for other beam boundary conditions, but will not step through the lengthy mathematics used to arrive at these quantities.

The damping equation can be written for a beam of any mode shape that has symmetric boundary conditions at the two ends. The mode shape $\phi(x)$ can be calculated directly, and also listed in the vibration literature for virtually any combination of beam boundaries.

$$\eta = \frac{\mu}{\omega_n \sum_m \rho_m A_m} \frac{\int_0^L \left(\left. \frac{\partial \varphi}{\partial x} \right|_x^{L/2} \right)^2 dx}{\int_0^L \varphi^2 dx} \sum_i \frac{P_i c_i^2}{h_i} \tag{2.42}$$

Using this equation and the mode shapes and natural frequencies of Figure 2.4, the loss factor for any mode of any symmetrically supported beam can be computed.

$$\omega_n = \frac{\lambda_1^2}{L^2} \sqrt{\frac{EI}{\rho A}} \tag{2.43}$$

$$\eta = \frac{K \mu k_B^2}{\omega_n \sum_m \rho_m A_m} \sum_i \frac{P_i c_i^2}{h_i} \tag{2.44}$$

Only the constant K changes from beam to beam, depending on the boundary conditions. Figure 2.10 shows K for several beam support conditions. Note that the factor K may be used without error in the general equation of the damping in a beam defined above.

Clamped-free $\delta = 0$ at $x = 0$	Pin-pin $\delta = 0$ at $x = L/2$	Free-free $\delta = 0$ at $x = L/2$
1.32	1.00	2.21

Figure 2.10 Damping constant K as a function of beam boundary conditions.

The way that the shear members are fixed relative to the structural members makes a significant impact on the amount of damping in the structure. If the beams can be fixed at one end, the damping will be three to five times greater than if the beams are free floating (depending on the boundary conditions). This is because the overall relative shear is much greater when the constraining layers are joined at one end, compared to when the layers are joined at their mid length.

2.4 Effects of Shear Rate on Apparent Viscosity

The apparent viscosity of the family of silicone fluids shows highly rate dependent behavior. At high rates of shear, the fluid viscosity rolls off by a significant amount. While the shear rate dependence of the silicone fluids is non-Newtonian, it does not invalidate the analysis because a given structure has a specific first bending mode frequency. When estimating the damping in a structure, the fluid viscosity must be corrected for the this frequency.

Figure 2.11 shows the roll-off for several moderate viscosity fluids. The roll-off trend shows that as the nominal viscosity increases, the break point frequency decreases. All of the fluids tend to converge on the same viscosity at very high shear rates.

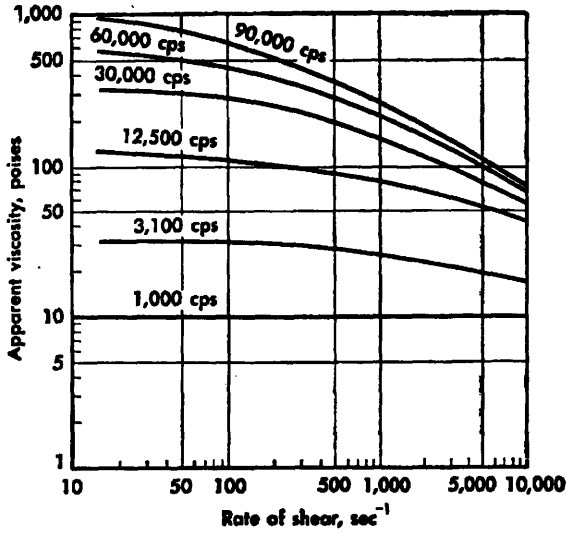

Figure 2.11 Apparent viscosity of silicone fluids vs. shear rate [MacGregor, 1954].

The viscosity vs. shear rate curves for the highly viscous silicone fluids used in this dissertation are not available from manufacturers such as GE Silicones and NuSil Corp. For this reason, a viscometer was designed and built to accurately measure the viscosity of the fluids used in the shear damper designs. Figure 2.12 shows the schematic of the viscometer system. A shaker, two accelerometers, and a dynamic signal analyzer were used to measure the behavior of the system at a range of frequencies. Two

accelerometers are used to obtain an average value of the motion to compensate for any rotation in the plate.

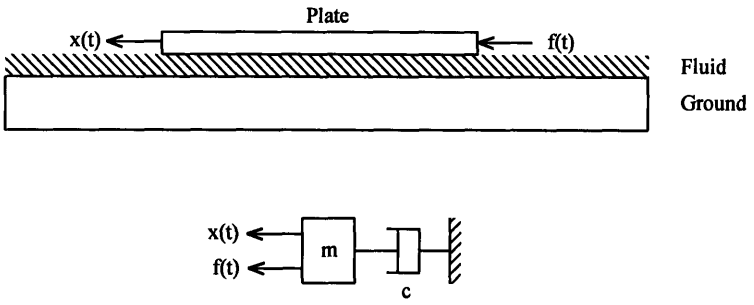

Figure 2.12 Viscometer design and mathematical model.

The nonlinear behavior of the viscous fluid requires that the viscometer data be considered one frequency at a time so that equilibrium can be reached with the effective viscosity at that frequency. For this reason, a swept sine input was used with a 50 cycle settling time and a 50 cycle measurement window. The signal analyzer used in these experiments automated this process to evaluate a large number of points. The results were used to calculate the effective viscosity at many different frequencies.

The experimentally measured amplitude of the transfer function x/f of the system can be used with the math model to deduce the correct viscosity at each frequency. The transfer function of the mathematical model may be obtained by considering the differential equations of motion of the viscometer system.

$$G = \frac{\ddot{x}(s)}{f(s)} = \frac{s}{ms + c} \tag{2.45}$$

The experimentally measured amplitude of the transfer function may be used in the following equation to estimate the damping factor.

$$R = \omega\sqrt{\frac{1}{G^2} - m^2} \tag{2.46}$$

The damping constant R can be found by considering the area of the plate A over which the fluid acts and the height h of the fluid film. The nominal viscosity is μ_{nom}. The multiplier γ gives the dimensionless shape of the viscosity profile.

$$R = \frac{A\gamma\mu_{nom}}{h} \tag{2.47}$$

The transfer function G is a complex number, therefore the effective magnitude of the argument in the square root of Equation (2.46) must be used to obtain the following estimate for the viscometer test results.

$$\mu = \frac{\omega h}{A}\sqrt{\frac{(1-m\,\mathrm{Re}(G))^2 + m^2\,\mathrm{Im}(G)^2}{\mathrm{Re}(G)^2 + \mathrm{Im}(G)^2}} \tag{2.48}$$

Figure 2.13 shows the viscosity of Dow Corning Viscasil 600,000 and NuSil Corp's 2,500,000 centiStoke fluid.

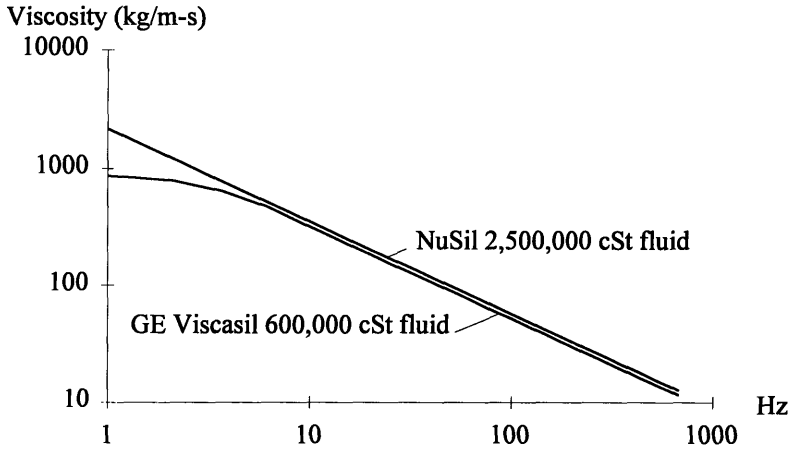

Figure 2.13 Viscosity of silicone fluids as a function of frequency.

The experimental data shown in Figure 2.13 were curve fit so that a viscosity estimate could easily be made for a given frequency (measured in Hz). Note that in the range of frequencies above 10 Hz, the two fluids are essentially the same, despite their large difference in nominal viscosity. This result is expected given the viscosity dependence of the frequency break point (the higher the nominal viscosity, the lower the break point frequency at which the effective viscosity rolls off).

$$\mu(f) = 1300 f[\mathrm{Hz}]^{-0.72} \tag{2.49}$$

2.5 Generalization of the Derivation to Viscoelastic Materials

The fluid flow profile in a sheared fluid was shown to have a linear profile. This profile leads to a shear stress in the fluid that is constant throughout the fluid, but dependent on the position x.

$$\tau = \mu \frac{\partial}{\partial t} u(x,y,t) = \mu i\omega \frac{\delta_2 - \delta_1}{h} \tag{2.50}$$

The velocity boundary conditions U_i are obtained directly from the mode shape of the beam and go with cos $k_B x$. The viscous fluid mechanism is therefore well modeled with velocity proportional damping. Considering a second order system, the viscous term may be easily included as the constant R.

$$m\ddot{x} + R\dot{x} + kx = f(t) \tag{2.51}$$

Viscoelastic materials are conveniently modeled as having complex modulus (stiffness) proportional to displacement, not velocity.

$$m\ddot{x} + k(1+i\eta)x = f(t) \tag{2.52}$$

A comparison between the two damping models may be obtained by Laplace transforming the differential equations of motion. The classic viscous damping model has the following roots in the Laplace domain:

$$s = -\zeta\omega_n \pm i\omega_n\sqrt{1-\zeta^2} \tag{2.53}$$

The structural damping roots are similar when the damping is small, but become increasingly different as the loss factor is made larger.

$$s = \pm i\omega_n\sqrt{1+i\eta} \tag{2.54}$$

When the damping is relatively low, we see that the two models yield very similar results given the approximation that the loss factor η is twice the damping factor ζ. Using this approach, the loss factor derivation can be easily modified to accommodate the viscoelastic materials using the structural damping formulation. The shear modulus of a viscoelastic material $G^* = G(1 + i\eta)$ is related to the shear stress in the layer by the constitutive relation for a solid:

$$\tau^* = G^*\gamma = G^*\frac{\delta_2 - \delta_1}{h} \tag{2.55}$$

The shear stresses in the viscous fluid and viscoelastic models may be equated to find the relationship between material properties. Noting that only the imaginary part of the complex modulus dissipates energy, the relationship may be simply written as follows:

$$\mu = \frac{G\eta}{\omega} \tag{2.56}$$

The loss factor for a shear damped beam may now be expressed as a function of the lossiness of the damping layer.

$$\eta = \frac{G^* k_B^{\,2}}{\omega^2 \sum_m \rho_m A_m} \sum_i \frac{P_i c_i^{\,2}}{h_i} \tag{2.57}$$

where:

$G^* = \mu\omega$ for viscous fluids

$G^* = G\eta$ for viscoelastic materials

2.6 Conclusion

This chapter outlined the analysis of the loss factor of a shear damped beam with arbitrary cross section and any number of shear damping sections. A simpler geometry was considered to show the importance of fluid film thickness and fluid viscosity. Other parameters, such as beam length, were shown to be unimportant (assuming the beam is sufficiently slender).

Through a dimensional analysis, the assumptions made of the fluid flow in the beam were shown to be valid. Experimental measurements show excellent agreement with the derivation provided in this chapter. The shear rate dependency of the silicone fluid was discussed to provide a complete understanding of the shear damping mechanism. Silicone fluids have desirable material properties including extremely high viscosity and only very slight temperature dependence. Viscoelastic materials have a greater modulus loss factor resulting in correspondingly higher damping for a given design. As will be shown, there is a practical limit to the amount of damping that may be built into a given design. Therefore the choice of using silicone fluids or viscoelastic materials for the damping medium will be dictated by the geometry and materials of the undamped structure.

2.7 References

Blevins, T. L., Formulas for Natural Frequency and Mode Shape, Van Nostrand Reinhold, New York, 1979.

Bridgman, P. W., Proceedings of the American Academy of Arts and Sciences, 77, 115, 1949.

Crandall, Stephen H., Dean C. Karnopp, Edward F. Kurtz, Jr., and David C. Pridmore-Brown, Dynamics of Mechanical and Electromechanical Systems, Krieger Publishing Company, Malabar, Florida, 1968.

MacGregor, Rob Roy, Silicones and Their Uses, McGraw-Hill Book Company, New York, 1954.

Meirovitch, Leonard, Elements of Vibration Analysis, McGraw-Hill, New York, 1975.

Özisik, M. Necati, Heat Transfer: A Basic Approach, McGraw-Hill, New York, 1985.

Potter, Merle, and John Foss, Fluid Mechanics, Great Lakes Press, Okemos, MI, 1982.

Chapter Three: Derivation of Dimensionless Beam Coupling Indicator

3.1 Introduction

The loss factor analysis of Chapter 2 made one assumption that requires further investigation: the damping material thickness was assumed sufficient to prevent coupling of the different components of the structure. This means that the shear force exerted by the damping material does not have a significant effect on the behavior of the bending beams. This is clearly not true for the limiting case of zero damping material thickness. As the layer thickness goes to zero, the adjacent structures begin to act as a homogenous structure. Correspondingly, the natural frequency of the structure will increase and the damping goes to zero (because there is no shear strain across the damping layer).

If the damping material is a viscous fluid, the layer thickness necessary to achieve structural coupling between layers is extremely thin for stiff materials such as aluminum, steel, and ceramic. Experiments show that physically realizable fluid film thicknesses are too large to generate substantial component coupling. Plastic beams, on the other hand, can be constructed that show experimentally component coupling. When using viscoelastic materials, the layer thickness is much more important. In general, the high shear modulus of viscoelastic damping materials means that the damping layer can readily couple with the vibrating elastic structure (coupling is inevitable when the damping layer thickness is near the optimal value, the goal is to avoid making the layer too thin).

A rigorous closed-form solution that explicitly considers the stress distribution between components in a beam is not available because of the complexity of the problem. That is why the original Ross-Ungar-Kerwin analysis considered a special three layer geometry. The work done in this dissertation has approached the problem from another direction: find an expression for the damping in any beam geometry assuming minimal coupling between layers. Using this approach, analysis has shown that a figure of merit can be developed to evaluate the likelihood of dynamic coupling effects. This dimensionless quantity can be used as an indicator of when the coupling effects are significant.

The results shown in this chapter were made possible with finite element analysis. A finite element model of the shear damping mechanism was developed so that different beam geometries could be efficiently investigated. In fact, the high accuracy with which the finite element method can find the damping in a beam design is one of the many advantages of the shear damping concept; experimental and finite element results show very close agreement.

The chapter will begin with a thorough explanation and verification of a finite element model of the damping mechanism, followed by an investigation of the coupling effect found in beams with very thin damping layers. The case of a fluid damped structure will be used without loss of generality; the relationship between the fluid damping model and the viscoelastic model was shown in the previous chapter.

3.2 Finite Element Modeling of a Shear Damped Beam

The finite element method may be used to model the coupling effect and determine at what damping layer thickness it becomes important. For this discussion, the damping model is taken to be two fluid damped beams constraining a high viscosity fluid between two flat, smooth surfaces.

3.2.1 The Model

The finite element model of the shear damping mechanism is based on two-dimensional spring and dashpot elements, as shown in Figure 3.1.

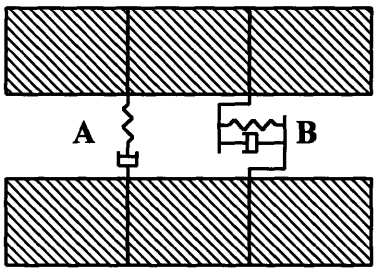

Figure 3.1 Finite elements used to model the shear damping mechanism.

The finite element model should be set up so that the nodes of each beam line up with each other in the vertical direction. This allows the use of two different types of spring/dashpot elements between adjacent structural members. The vertically acting elements in **A** model squeeze film and fluid compressibility effects (or vertical stiffness in the case of a viscoelastic damping material). The horizontally acting element **B** models the shear stiffness and damping effects. Note that *both* types of spring/dashpot elements (**A** and **B**) are used at each node pair along the structure/damping material/structure interface.

3.2.1.1 Squeeze Film/Compressibility Element (A)

Spring and dashpot elements in **A** model the effect of the squeeze film damping as well as the slight compressibility of the damping material. The spring is used to model the compressibility of the damping material and the dashpot models the squeeze film damping in thin fluid films. Note that the squeeze film effect only occurs in damper designs using viscous fluids.

The compressibility can be estimated by using the (bulk) modulus of the damping material. The family of silicone fluids used in the proposed damper designs are not measured directly by the manufacturer for bulk modulus; however, the literature contains the results of compressibility studies performed on silicone fluids. Experimental data indicate that silicone fluids have approximately 4.5% compression at 7100 psi [Bridgman, 1949]. This corresponds to a bulk modulus of 2.3 MN/m^2 which is similar to water or oil. Viscoelastic materials will have a modulus that can be used directly from manufacturer data sheets (the modulus is typically a small fraction of the modulus of steel or aluminum).

The damping associated with the squeeze film effect can be determined using the two-dimensional Navier Stokes equations to determine the fluid flow profile. The assumptions used in the derivation of the loss factor in Chapter 2 will again simplify the equations (fully developed flow). Figure 3.2 shows a cross sectional view of the squeeze film model. Note that the fluid flow profile is shown across the width of a shear member (there will be virtually zero flow into the page).

$$\frac{1}{\mu}\frac{\partial p}{\partial x} = \frac{\partial^2 v}{\partial y^2} \tag{3.1}$$

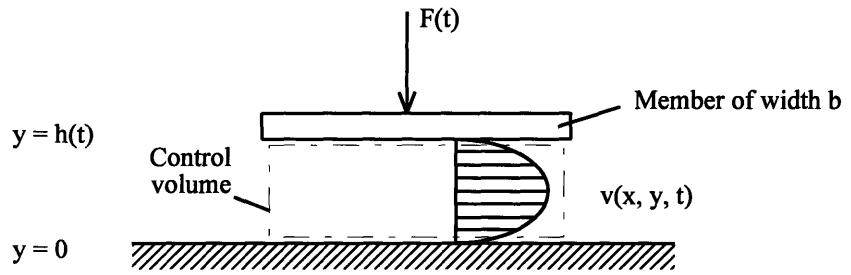

Figure 3.2 Model of squeeze film damping in the fluid layer.

The velocity profile may be found by double integrating the simplified Navier Stokes equations. The boundary conditions for the flow indicate that the velocity is zero at both structural surfaces.

$$v(y) = \frac{1}{2\mu}\frac{\partial p}{\partial x}\left(y^2 - yh\right) \tag{3.2}$$

A control volume drawn around the fluid between the two plates will relate the relative motion of the top plate to the flow out the sides. This relation allows calculation of the pressure gradient. Integration of the pressure gradient yields the equation for the pressure on the structural surfaces.

$$p(x) = \frac{6\mu}{h^3}\frac{dh}{dt}\left(\frac{b^2}{4} - x^2\right) \tag{3.3}$$

By balancing the applied force per unit length of the structure (into the page), an equation relating the force to upper plate velocity may be found. The width of the plate is b and the fluid viscosity is μ.

$$\frac{F}{L} = \mu\left(\frac{b}{h}\right)^3 \frac{dh}{dt} \tag{3.4}$$

The damping constant is simply the relationship between the total force and the relative velocity between the two plates.

$$R = \mu L\left(\frac{b}{h}\right)^3 \tag{3.5}$$

The viscous fluid damping mechanism has been shown to work best with relatively thin fluid films. At these layer thicknesses, the cubic relation between the damping and the fluid film thickness results in tremendous dynamic stiffness across the dashpot (resulting in essentially zero relative motion). As a result, no energy is dissipated by the squeeze film effect. Kurtze's 1959 paper documents the use of viscous fluids as a damping medium, but in his application the fluid films were so thick that relative motion could occur in the vertical direction (with appreciable squeeze film losses) [Kurtze, 1959].

The damping layer compressibility, which acts in series with the squeeze film damping, has a lower dynamic stiffness than the squeeze film. However, in the range of damping layer thicknesses that are suitable for efficient shear damping, the vertical compressibility is sufficiently stiff to prevent relative motion. The amount of stiffness in element **A** is based on the modulus of the layer and the area of the finite element.

$$k = B\frac{bL}{h} \tag{3.6}$$

The quantity bL is the area over which the element acts, B is the modulus of the layer material (the bulk modulus in the case of fluid damping), h is the fluid film thickness. The high stiffness of the spring and the dashpot means that, in general, the element **A** may be safely removed from the finite element model. A constraint may be applied in its place to keep the vertically adjacent nodes equally spaced in the vertical direction.

3.2.1.2 Viscous Damper Element (B)

The spring/dashpot element **B** is a horizontally oriented parallel spring and dashpot that models the shear damping effect. In the case of fluid damping the spring has zero stiffness because the fluid is assumed to be Newtonian. As shown in Chapter 2, the shear strain profile across the damping layer is linear for both viscous fluid and viscoelastic materials. The damping constant R and the spring constant k can be estimated by

considering the stress on the surface of the beam imparted by the sheared damping material.

$$R = bL\frac{\mathrm{Im}G^*}{\omega h}$$
$$k = bL\frac{\mathrm{Re}G^*}{h} \tag{3.7}$$

The area product bL is the area of a single constraining surface element. As with any finite element solution, using more elements typically gives higher accuracy. In practice, element types **A** and **B** are efficient enough that 20 of each distributed down the length of a two-dimensional beam gives good results.

The shear damping mechanism acts in every direction tangent to the surfaces of the structure. For example, a two degree of freedom beam problem will have one element modeling the fluid compressibility and one element modeling the shear damping between every facing node on the solid/fluid/solid interface. A plate will require one element for compressibility and two elements for the viscous damping between every pair of facing nodes.

3.2.2 Selection of Appropriate Element Types

The shear damping mechanism allows flexibility in the mesh design, but a few guidelines have been established. In the case of a two-dimensional structure, the natural element choices are high order beam elements and various 4, 8, or 9-node isoparametric elements [Bathe, 1982]. A fully three-dimensional structure typically requires 8 or more node brick type elements. If the three-dimensional structure is thin, shell elements may be used to model the structure.

Consideration of the shear mechanism introduces a constraint on the element type that may be used to model the damping in a structure. *The element must have nodes modeling thickness.* Shell and beam type elements require the calculation of quantities that include thickness; however, there is only a single element across the height of these element types. Planar or solid elements model the thickness of the structure and in doing so, model the longitudinal compression and tension of the material away from the neutral axis. This effect is shown in Figure 3.3.

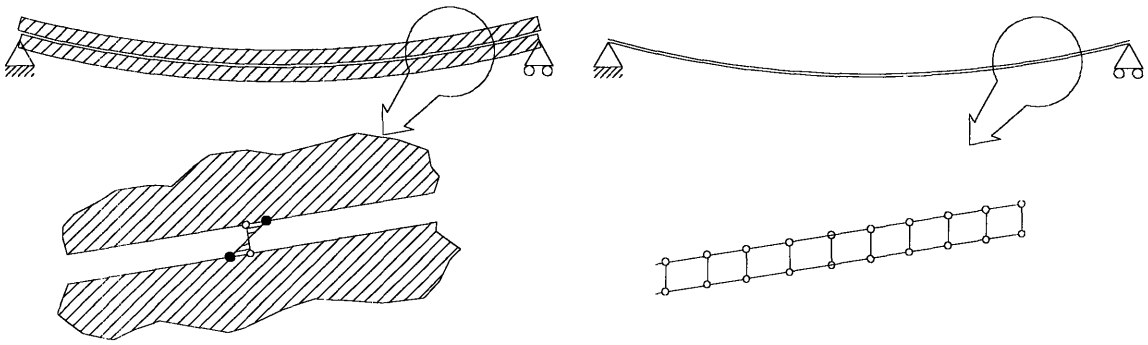

Figure 3.3 Contrast of element bending problem: isoparametric vs. beam element types.

In the case of the beam or shell element model, the nodes are assumed to be along the neutral axis so the nodes do not model longitudinal displacement when undergoing bending vibration. As a result, a model made with this approach would show zero damping because there is no relative motion across the horizontal damping element. Figure 3.4 shows elements that may be used to properly model the damping effect.

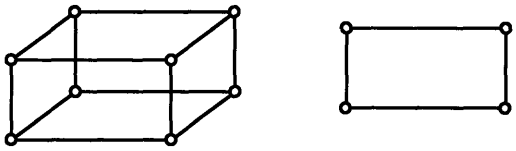

Figure 3.4 Suitable structural element choices for shear damping models.

Another issue that must be considered when meshing a shear damped structure is the number of elements across the thickness of the structure. The numerous finite element simulations made in preparation for this dissertation show the importance of having multiple elements across the thickness of the structure. As the damping layer thickness is made thinner, the shear stress that the fluid imparts to the elastic structure becomes important. This high level of shear cannot be accurately modeled with a single layer of elements across the thickness of the structure. Multiple layers of elements must be used to properly model the behavior of the structure (such as a flat plate or beam) when the shear stresses become high.

Trial finite element simulations did show that in the interest of saving computer time, reduced density meshes across the thickness of the structure can be used when the shear stresses are low (at damping layer thicknesses that provide minimal coupling between the adjacent structures). Figure 3.5 summarizes this phenomenon.

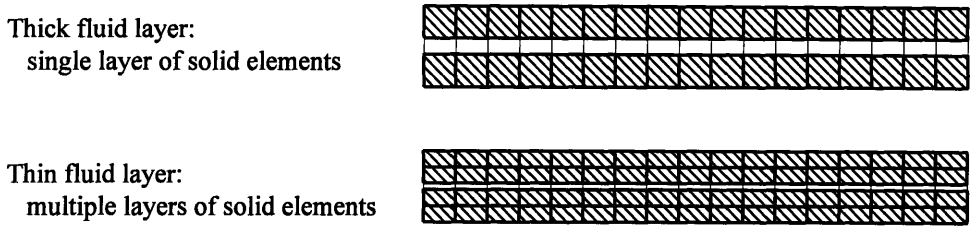

Figure 3.5 Minimum number of elements required to model a structure.

In general, the film thicknesses that result in coupling of structural components are not physically realizable with viscous fluids, therefore simplified finite element meshes may be safely used. This is an important feature of the damping method; the damping may be accurately modeled with minimal computational effort. Viscoelastic materials, as a result of their higher stiffness, may require finer mesh densities even when the damping layer is comparatively thick.

3.2.2 Model Verification

The finite element model of the squeeze film and shear damping phenomena was verified in a study of the dynamic performance of a damped structure. The structure has two identical beams constraining a thin fluid film. Each beam is 30 cm long, 2.5 cm wide and 1 cm thick. The material is steel and fluid film thicknesses between 25 and 250 microns were tested. Figure 3.6 shows a typical drive point comparison between experimental and finite element results. In this example, the viscosity is corrected for the first bending mode at 750 Hz.

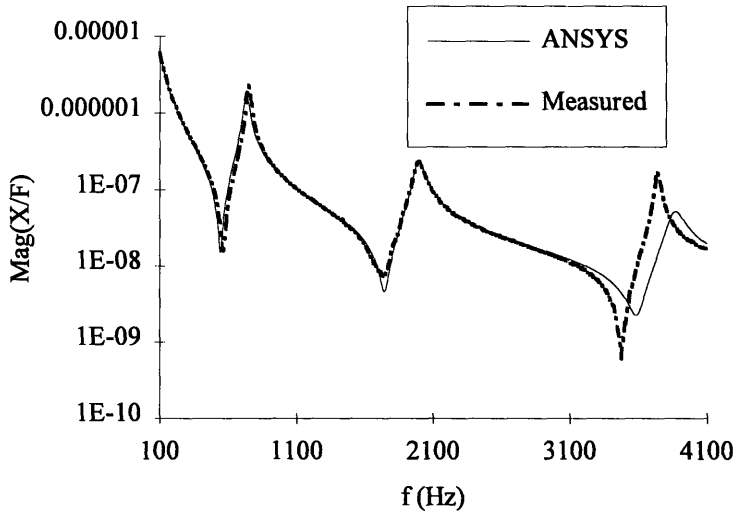

Figure 3.6 Measured and finite element drive point frequency response functions (μ corrected for 750 Hz).

The figure shows close agreement in the first and second bending modes. The third mode is not quite as well predicted by the finite element analysis. This is an inevitable result of using a viscosity corrected for a particular frequency (both viscous fluids and viscoelastomers will show this behavior). The viscosity correction factor at 3500 Hz is much lower than it is at 750 Hz, which is why the finite element results overestimate the damping.

Numerous plots such as Figure 3.6 were made in the process of correlating the finite element model to experimental data. Other beam geometries, materials, and boundary conditions were also tested. The results indicate that the finite element model accurately captures the dynamics of the shear damped beam. This model may be confidently used to explore alternative geometries to hasten the design process.

Figure 3.7 shows the first three bending mode shapes of a free-free beam obtained experimentally and by finite element analysis.

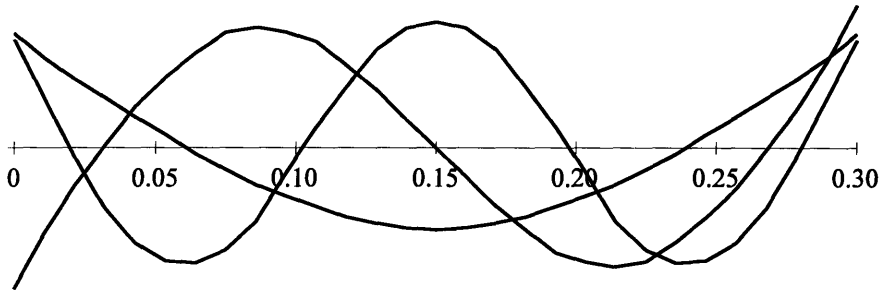

Figure 3.7 First, second, and third mode shapes of a free-free damped beam.

Note how the figure shows the node of the first mode at around 22 percent of the length of the beam (a classic result for free-free Euler beams). The plots shown in Figure 3.7 came from experimental data, explaining the slight jaggedness of the mode shape curves (30 points were tested along the length of the beam).

The close agreement of the modal parameters (natural frequency and damping), and the similarity of the predicted and experimental drive point frequency response functions verifies the validity of the two element types used to model the fluid damping mechanism. The next section will use this finite element model to investigate the behavior of the beams given very thin fluid film thicknesses (a case which is not well modeled by the closed-form solution).

3.3 Component Coupling from Thin Damping Layers

The results of finite element runs made with decreasing fluid film thickness show the point at which the closed-form solution ceases to accurately predict the damping in the beam. In the case of the double beam geometry, the assemblage begins to act as a homogenous structure as the damping layer is made thinner. Figure 3.8 shows results from a series of ANSYS runs made with a twin beam assemblage of steel strips. As the

damping thickness is made thinner and thinner, the natural frequency climbs until it reaches that of a solid beam.

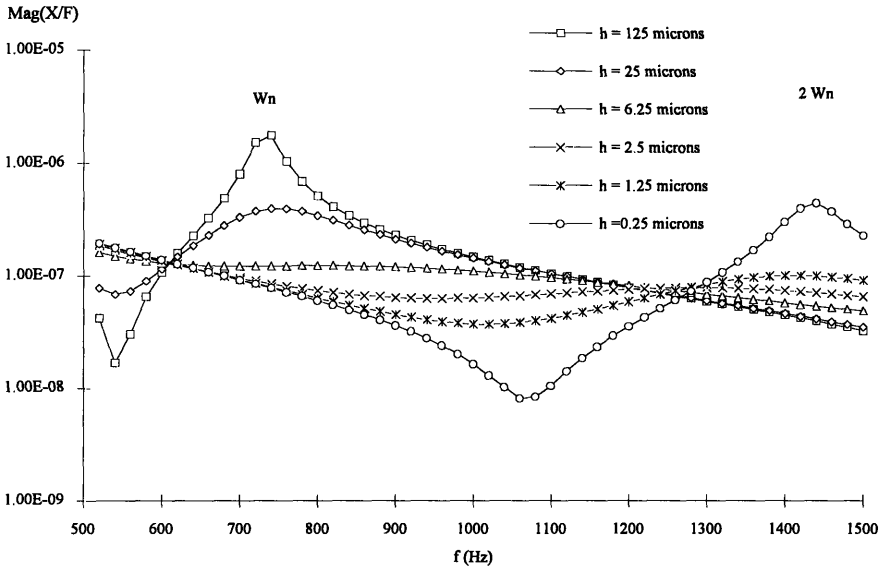

Figure 3.8 Finite element drive point frequency response function of a beam structure (μ corrected for 750 Hz).

Although the closed-form solution to this limiting case is not available, a dimensionless parameter has been derived that will predict the film thicknesses at which the shear stresses will become important.

3.3.1 Evidence of Coupling by Modal Analysis

The component coupling effect was first observed experimentally with a high density polyethylene beam. HDPE has a much lower density and elastic modulus so the beam showed coupling with relatively thick fluid films. The beams were also modeled in ANSYS to investigate the changes in mode shape of the beams as the fluid layer is made progressively thinner. Figure 3.9 shows ANSYS results for three different film thicknesses (experimentally verified results).

Fluid thickness (microns)	Natural Frequency (ω)	Damping Factor (η)
250	234 Hz	0.50
85	386 Hz	0.64
25	428 Hz	0.24

Figure 3.9 First mode results from HDPE beams simulated in ANSYS.

When the fluid layer is fairly thick, the amount of damping in the beam is very well predicted by the equations presented in Chapter 2. The mode shape of the beam is that of

the classic Euler solution, and the natural frequency of the assemblage is equal to that of the two identical (uncoupled) shear members. Figure 3.10 shows the x-direction shear strain obtained by finite element analysis with a thick fluid layer (no coupling between the two beams). Although the two beams are connected by the fluid layer, they act independently. The strain distribution in the beams is symmetric about each beams' neutral axis.

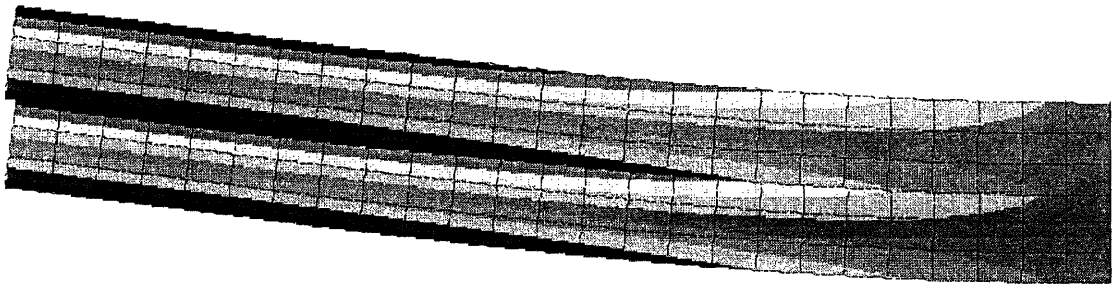

Figure 3.10 Horizontal strain on the left hand side of a free-free beam with a thick fluid layer (250 microns).

When the damping layer is made thinner, the mode shape of the beam deviates from the expected Euler solution, and the natural frequencies are in a transition between the uncoupled and coupled beam. Figure 3.11 shows the same beam with an intermediate thickness damping layer. The fluid layer is thin enough to result in some coupling between the two individual beams, and the mode shape is significantly different than the classic free-free beam solution (the ends are nearly straight).

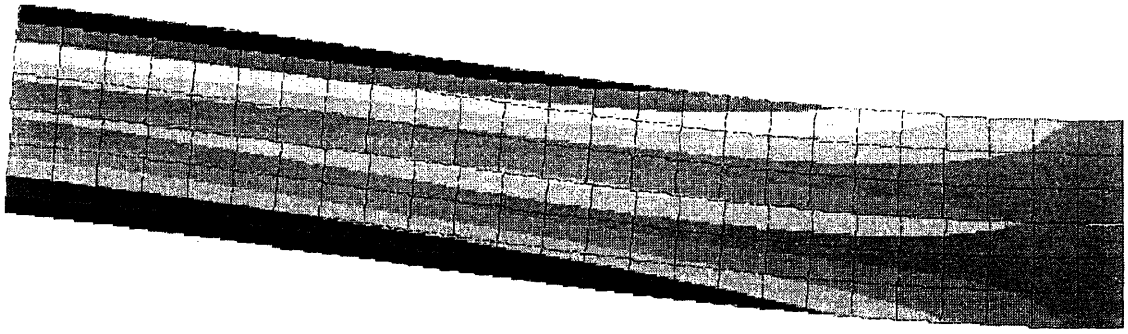

Figure 3.11 Horizontal strain on the left hand side of a free-free beam with a intermediate fluid layer (85 microns).

Figure 3.12 shows the structure with a thin damping layer. In this case, the two beams are bending as one, and the natural frequency of the assemblage has doubled. The strain distribution is that of a solid beam with virtually no strain in the region of the fluid film. There is very little damping in this case because the strain across the damping layer is so low.

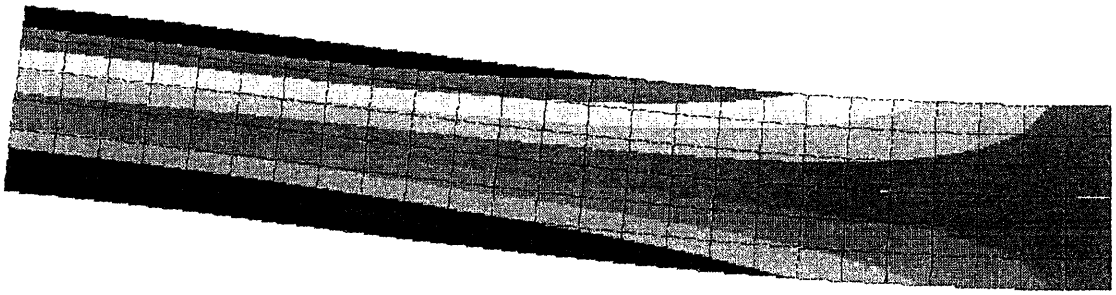

Figure 3.12 Horizontal strain on the left hand side of a free-free beam with a thin fluid layer (25 microns).

The mode shape is altered when the damping layer thickness is in the intermediate range. The mode shape of a simply-supported beam, for example, is changed from a pure sine wave to a sine wave with flattened out ends. If the fluid layer is very thin, the mode shape returns to a pure sine wave. The best way to visualize this mode shape distortion is by looking at the slope of the mode shapes of a beam of various film thicknesses, as shown in Figure 3.13 for the three film thicknesses (only the left hand side of the beam is shown).

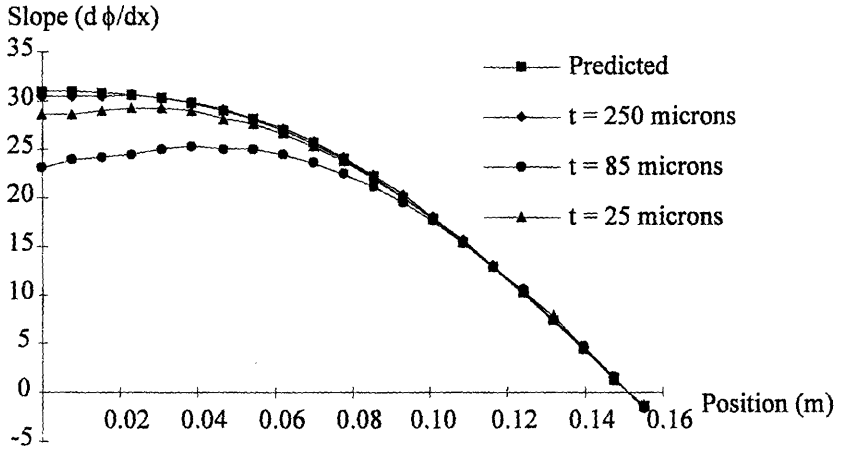

Figure 3.13 Stiffening of a beam resulting in an altered mode shape slope.

Figure 3.14 shows the changes in natural frequency and damping as a function of damping layer thickness in the parallel steel beams.

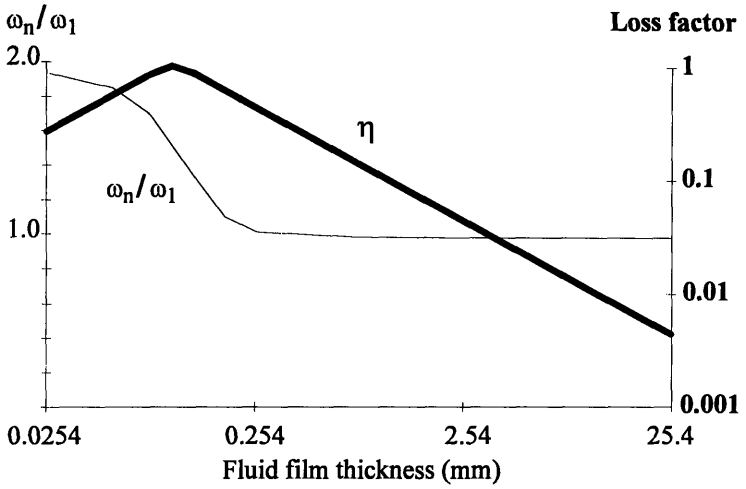

Figure 3.14 Natural frequency and damping factor as a function of fluid film thickness.

The finite element model was used to test many beam materials, damping layer materials, and geometries to gain further understanding in the trends in stiffening behavior. This approach was used because a closed-form solution to the stiffening effect is not available. The goal was therefore to find a dimensionless quantity that could be used as an indicator to determine if coupling effects are important in a given beam design.

3.3.3 Analysis of Structural Coupling Effect

The key to understanding the coupling phenomenon is the ratio between the shear strain in the damping layer and the spatial derivative of the mode shape. This ratio was derived in the RUK analysis presented in Chapter 1 [Ross, Ungar, and Kerwin, 1959].

$$H_2 \frac{\partial \psi}{\partial \phi} = \frac{H_{31} - D}{1 + \dfrac{G_2}{k_B{}^2 K_3 H_2}} \tag{1.6}$$

where:

$$g = \frac{G_2}{k_B{}^2 K_3 H_2} \tag{1.7}$$

The dimensionless parameter g is important because its magnitude controls the filter-like shape of the ratio.

Note that this equation was derived for the case of a three layer plate in the RUK analysis. We may adapt it to our arbitrary beam geometry by treating the first layer (denoted by subscript 1) as the core structural member of the beam, the third layer as a particular shear member (subscript 3), and the second layer as a damping material between them (subscript 2).

The damping layer thickness at which shear effects becomes important is the thickness at which the damping layer stops shearing and the shear member starts stretching. This layer thickness corresponds to the point of optimal damping in the structure. The ratio between the shear force at the wall (from the fluid) and the extensional stress in the shear member will help determine the critical damping layer thickness. For the purpose of example, the simply-supported beam mode shape is used.

$$y(x) = Y \sin k_B x \sin \omega t \tag{3.8}$$

Considering the first mode of vibration, the extensional stress in the shear member, σ_{ext}, is given by:

$$|\sigma_{ext}| = E_3 H_3 H_{30} \left| \frac{\partial^3 y}{\partial x^3} \right| = Y E_3 H_3 H_{30} k_B^{\ 3} \tag{3.9}$$

The shear stress τ that the damping layer imparts on the surface of the adjacent shear member is calculated from the relative displacement across the damping layer. The fluid velocity profile was calculated in Chapter 2 and is used again here.

$$|\tau_w| = \frac{G_2}{H_2} |\delta| = \frac{Y G_2 k_B H_{30}}{H_2} \tag{3.10}$$

The ratio of the shear stress at the wall and the extensional stress can be calculated to obtain a measure of when the shear effects become important:

$$\left\langle \frac{\tau_w}{\sigma_{max}} \right\rangle_{x,t} = \frac{G_2}{k_B^{\ 2} E_3 H_3 H_2} \tag{3.11}$$

Note that the spatially-averaged stresses from Equations 3.8 and 3.9 are used to find the ratio. This is because the bending stress goes to zero at the ends of the beam (which would make the stress ratio go to infinity). The distribution of the stresses along a beam is shown in Figure 3.15.

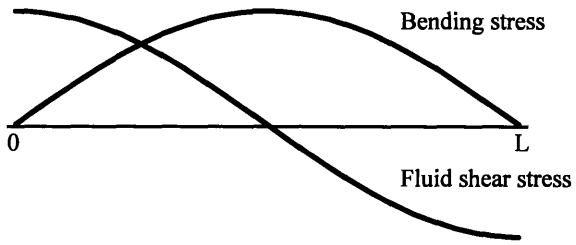

Figure 3.15 Shear and bending stresses along a simply-supported beam.

The ratio between the shear and extensional stress yields the same dimensionless parameter as the shear parameter g from the RUK analysis. The finite element method

used in the next section explores the validity of this dimensionless parameter as an indicator of coupling effects in a shear damped beam.

3.3.4 Dimensional Analysis of the Structural Coupling Problem

Figure 3.16 shows a sample beam that was used to develop a large database of finite element results to test the dimensionless parameter.

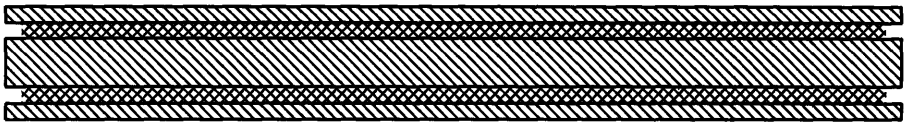

Figure 3.16 Schematic of a five layer shear damped beam (core beam, two shear members, and two damping layers).

Eight parameters effect the behavior of the five layer beam assemblage shown in Figure 3.16. These parameters will dictate both the natural frequency of the beam assemblage as well as the loss factor in the composite beam. The term d is an intermediate result (dubbed the effective radius of gyration) which is used to quantify the effective off-axis distance between the shear members and the structural (center) beam. For example, when the fluid layer is thick, the three separate beams are uncoupled and vibrate about their individual neutral axes ($d = 0.5\ [h_1 + h_3]$). When the fluid layer is very thin, the beams all bend as a unit so d approaches 0. Ultimately, determining the damping in a given beam configuration will depend on accurately estimating the distance d as a function of damping layer thickness.

	Units	Variable
E	kg/m-s^2	beam modulus
ρ	kg/m^3	beam density
μ	kg/m-s	fluid viscosity
L	m	length
h_1	m	inner beam thickness
h_2	m	fluid layer thickness
h_3	m	shear member thickness
λ	-	boundary condition constant
d	m	effective neutral axis term

$$d = f(E, \rho, \mu, L, h_1, h_2, h_3, \lambda) \tag{3.12}$$

Dimensional analysis of the eight parameters yields the following reduced set of independent dimensionless variables for the distance d.

$$d^* = \frac{d}{2(h_1 + h_3)} = f\left(\frac{\mu}{h_2\sqrt{\rho E}}, \frac{L}{h_2}, \frac{h_1}{h_2}, \frac{h_3}{h_2}, \lambda\right) \tag{3.13}$$

The distance d^* is a function of five dimensionless variables. Four complete sets of finite element runs were made to determine the relationship between d^* and the other variables. This testing was done in a Factorial-style matrix with one variable changed at a time. Figure 3.17 shows the table of dimensionless parameters used in the test matrix.

Set	$\frac{\mu}{h_2\sqrt{\rho E}}$	$\frac{L}{h_2}$	λ	$\frac{h_1}{h_2}$	$\frac{h_3}{h_2}$
One	c_1	c_2	4.730	ten trials	four trials
Two	c_1	c_2	π	ten trials	four trials
Three	c_1	1.333 c_2	4.730	ten trials	four trials
Four	0.5 c_1	c_2	4.730	ten trials	four trials

Figure 3.17 Test matrix used in the dimensional analysis.

Completing the test matrix required 160 finite element modal analyses in ANSYS. The natural frequency and damping of the first bending mode for each finite element run was recorded. The results were then tabulated creating a database of loss factors as a function of the five dimensionless parameters. A solution to the equation was then found by exploring the relationships between all the data points. After exhaustive investigation, several trends became apparent:

1. The effective radius of gyration d^* acts as a high pass filter with respect to the damping layer thickness. Figure 3.18 shows how the distance d^* corrects the loss factor for thin damping layers.

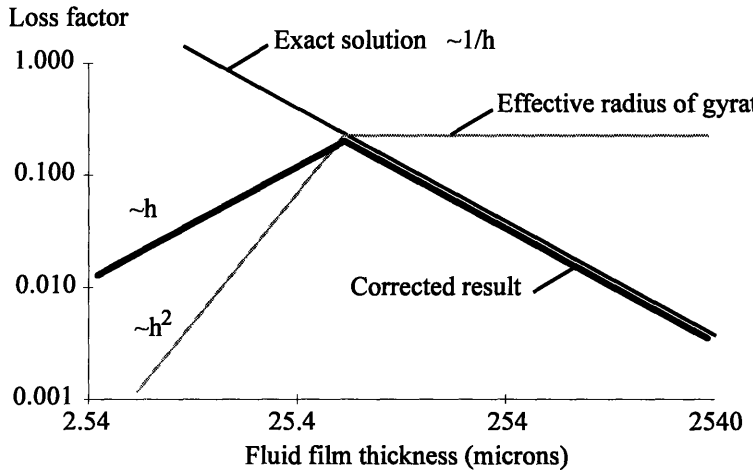

Figure 3.18 Relationship between effective radius of gyration d and damping factor.

2. The general trend of a high pass filter can be achieved by using an exponential function:

$$f(\cdot) = 1 - \exp^{-f'(\cdot)} \tag{3.14}$$

3. The function f' goes with the ratio between h_1 and h_3.

4. The function f' goes with the square of the boundary condition parameter λ.

5. The function f' does not depend on the length of the beam L.

The five observations were used to develop an expression for f' that satisfied all 160 finite element results, as well as countless simulations. This function takes the dimensionless quantity f' and multiplies it by the characteristic length 0.5 (h_1 + h_3), as shown below.

$$d = \frac{h_1 + h_3}{2}\left\{1 - \exp\left(\frac{-10}{g}\right)\right\} \tag{3.15}$$

$$\eta = \frac{KG^* k_B^{\,2}}{\omega_n^{\,2}\sum_m \rho_m A_m}\sum_i \frac{P_i d_i^2}{h_i} \tag{3.16}$$

where:

$$\begin{aligned} &K = 1 \text{ for pin} - \text{pin beams} \\ &K = 2.21 \text{ for free} - \text{free beams} \\ &K = 1.32 \text{ for cantilever beams} \end{aligned} \tag{3.17}$$

This equation is useful for any beam judged to have a damping layer thin enough to be near the theoretical maximum for a given beam geometry. Figure 3.19 shows a sample of the ANSYS results, and the corresponding empirical formula results.

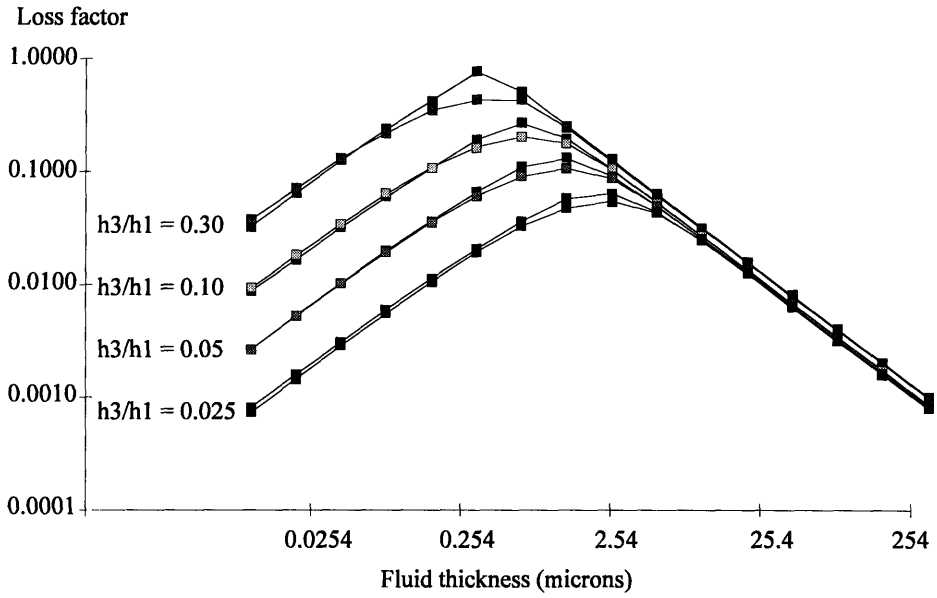

Figure 3.19 ANSYS and empirical results showing the effective radius of gyration correction factor.

The results shown in Figure 3.19 are typical of the corrected damping factors predicted by the equations above.

3.4 Maximum Loss Factor in Beams with Complex Geometries

The designer who understands the fundamentals of the shear damping mechanism will desire an equation estimating the maximum loss factor that can be obtained for an arbitrarily complex structure. Such an estimate was found by examining the huge database that was created during this research. While only an empirical estimate, the equation is not sensitive to boundary condition or beam length.

$$\eta = 0.4\left(\frac{EI_\infty}{EI_0} - 1\right) \tag{3.18}$$

EI_∞ is the bending stiffness of the beam if damping material is assumed to be infinitely stiff (i.e., the damped beam is solid). The quantity EI_0 is the bending stiffness of the beam if the damping material has no stiffness (i.e., the structural components are completely uncoupled). This equation may be readily applied to optimize a shear damper design.

3.5 Conclusion

Chapters 2 and 3 have presented the complete analysis of the shear damped beam in bending. Provided the shear layers are thick enough to prevent a high degree of coupling

between components, the damping in a beam of arbitrary cross section may be estimated by:

$$\eta = \frac{KG^* k_B^{\,2}}{\omega^2 \sum_m \rho_m A_m} \sum_i \frac{P_i c_i^{\,2}}{h_i} \qquad \text{(2.44) and (2.57)}$$

where:

$G^* = \mu\omega$ for viscous fluids

$G^* = G\eta$ for viscoelastic materials

$K = 1$ for pin – pin beams

$K = 2.21$ for free – free beams

$K = 1.32$ for cantilever beams

The maximum amount of damping that may be built into a structure can be predicted by a convenient empirical formula:

$$\eta = 0.4\left(\frac{EI_\infty}{EI_0} - 1\right) \qquad (3.18)$$

In the case of a five layer configuration (as well as other similar geometries), the damping in a shear damped structure may be accurately written as a function of the shear parameter g and the effective radius of gyration d:

$$\eta = \frac{KG^* k_B^{\,2}}{\omega_n^{\,2} \sum_m \rho_m A_m} \sum_i \frac{P_i d_i^{\,2}}{h_i} \qquad (3.16)$$

where:

$$g = \frac{G_2}{k_B^{\,2} E_3 H_3 H_2} \qquad (3.11)$$

$$d = \frac{h_1 + h_3}{2}\left\{1 - \exp\left(\frac{-10}{g}\right)\right\} \qquad (3.15)$$

3.6 References

Bathe, Klaus-Jurgen, Finite Element Procedures in Engineering Analysis, Prentice Hall, Englewood Cliffs, 1982.

Kurtze, Gunther, *Bending Wave Propagation in Multi-Layer Plates*, Journal Acoustical Society of America, Vol. 31, September 1959.

Ross, Donald, Eric Ungar, and E. M. Kerwin, *Damping of Plate Flexural Vibrations by Means of Viscoelastic Laminae*, Proc. Colloq. Structural Damping, ASME, 1959.

Ruzicka, Jerome, United States Patent 3,088,561, May 7,1963.

Chapter 4: Implementation of the Shear Damping Mechanism in Finite Elements

4.1 Introduction

This chapter documents the use of the closed-form theory and finite element model to estimate the damping in a variety of shear damped structures. Results from a large flat plate and a multilayer beam are presented with experimental verification.

Finite element modeling of a shear damped beam gives accurate results when used to estimate the loss factor of bending vibration modes. The finite element and experimentally measured modal loss factors are in close agreement, a result of the well modeled behavior of the damping mechanism. This ability to accurately predict the damping of a structure without building a prototype is a tremendous advantage over other damping treatments. Furthermore, the optimal structural design for stiffness and damping may be determined from a finite element approach to the problem.

4.2 Experimental and Finite Element Correlation of a Shear Damped Flat Plate

A flat plate with free boundary conditions cannot be modeled closed-form for lack of a analytical mode shape solution (note that only the simply-supported plate can be solved analytically). The finite element method allows calculation of its modal frequencies, damping, and mode shapes which are in close agreement with experimental results.

The example will proceed by first verifying the finite element model of the flat plate without added shear damping. Verification is carried out by comparing analytical and experimental modal analysis results from the plate. Once verified, the finite element model will be used to estimate the damping of two similar plates with a thin layer of viscous fluid between them. A second experimental modal analysis is performed of the damped structure to compare to finite element results.

4.2.1 Single Plate Finite Element/Experimental Analysis Correlation

The plate used in this case study is an acrylic sheet 41 cm wide, 51 cm long, and 5.4 mm thick. Acrylic was chosen because, as seen in Chapter 2, materials with low modulus and density give greater loss factors with thicker fluid films than stiffer materials such as aluminum and steel. Figure 4.1 summarizes the material properties used in the finite element model. Poisson's ratio of acrylic was not directly available, but a value of $\nu = 0.25$ was chosen based on data taken on similar materials [McClintock and Argon, 1966].

Plate width	41 cm
Plate length	51 cm
Plate thickness	5.4 mm
Plate density	1200 kg/m^3
Plate modulus	4.25 GPa
Plate Poisson's ratio	0.25
Nominal 8-node brick element size	2 × 2 × 0.5 cm

Figure 4.1 Single sheet plate configuration.

The following sections will show the close agreement between the finite element and experimental modal analysis results of the single acrylic sheet.

4.2.1.1 Finite Element Analysis

The acrylic plate was modeled with 8-node elastic solid elements. These elements offer three degrees of freedom at each node (displacements in the x, y, and z-directions). Meshes of various densities were tested to find a suitable model. A comparison of results obtained from shell elements was made as a final accuracy check. Figure 4.2 shows the mesh density of the converged model.

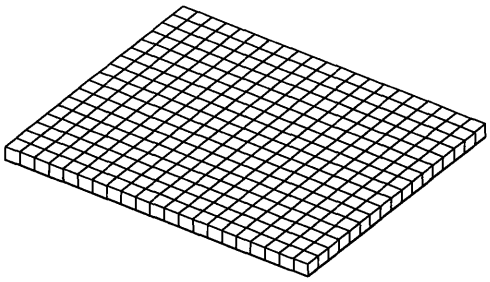

Figure 4.2 Finite element model of the acrylic sheet.

The mode shapes of the plate were obtained through the use of the complex eigensolver available in ANSYS. Figure 4.3 shows the five bending modes that occur below 100 Hz.

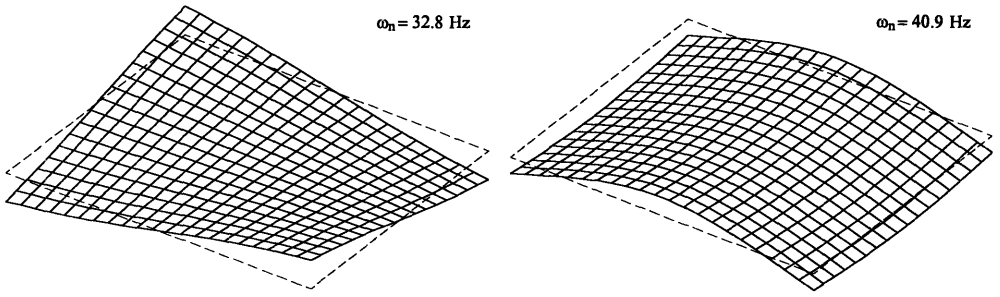

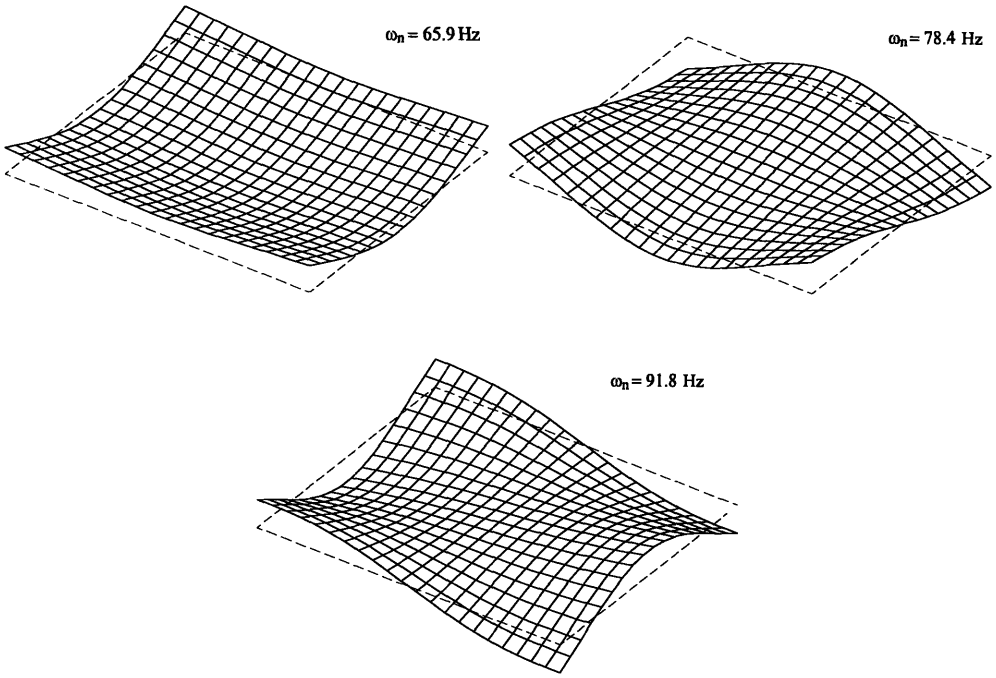

Figure 4.3 First five mode shapes of a flat plate (obtained from ANSYS).

Once this step was completed, an experimental modal analysis of the single acrylic plate was performed.

4.2.1.2 Experimental Modal Analysis

The acrylic plate experimental analysis was performed with a force shaker, a low mass force transducer, and a very low mass accelerometer. The use of low weight transducers is important during the testing of structures because mass loading effects can distort the measurements. Measurement point locations such as the corners of the plate undergo a large deflection, a result of the relatively low stiffness associated with these regions. The mass of an accelerometer can alter the natural frequency of the plate when placed in these low stiffness regions. As a result, the collection of frequency responses for the plate would have discrepancies as the accelerometer is roved around the plate. The use of low weight transducers minimizes this effect at the expense of some low frequency resolution. The first mode of the acrylic plate is around 30 Hz, safely within the region that can be accurately measured with low weight piezoelectric transducers.

Figure 4.4 shows the plate, its compliant supports (which provide a close approximation of free boundaries), and the location of the test points. The solid dot indicates the location of the shaker, which was left constant throughout the modal testing. A single accelerometer was roved around the plate at the locations indicated by the hollow dots.

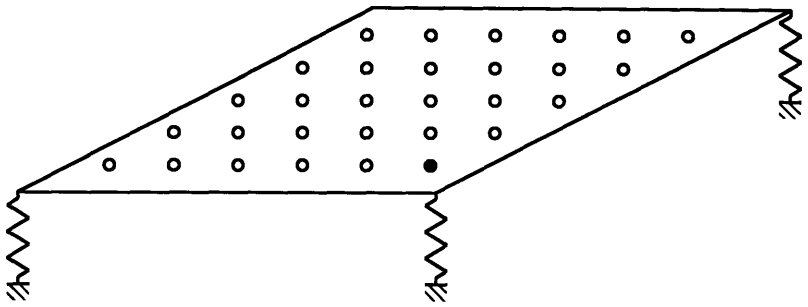

Figure 4.4 Schematic of experimental modal analysis measurement points.

The data collection proceeded by optimizing the setup parameters on a signal analyzer to give frequency response functions with very high coherence (greater than 0.99). A zoom frequency range was used to avoid the low frequencies where low transducer output leads to poor coherence. Data filtering (e.g., Hanning or flat top windows which smooth the data) was not required because the compliant plate supports effectively isolated the plate from external noise sources. Figure 4.5 shows the setup of the analyzer used to collect the response functions.

Frequency span	15 to 115 Hz
Frequency resolution	0.125 Hz (8 second sample time)
Excitation waveform	Burst random (90 percent duration)
Windowing	Uniform (none)
Number of averages	25
Curve fitting technique	Polyreference (frequency domain)

Figure 4.5 Data collection and modal analysis parameters.

The data collected in the experimental modal analysis are summarized in Figure 4.6. Although the first two modes of vibration are of primary interest, higher modes are included for completeness. The natural frequencies of the modes given by the ANSYS simulations are also listed for comparison. In general, the experimental and finite element data are in close agreement.

Mode no.	Experimental nat. freq. (Hz)	ANSYS nat. freq. (Hz)	No. of 1/2 wavelengths on short side	No. of 1/2 wavelengths on long side
1	33.1	32.8	1	1
2	41.5	40.9	0	2
3	66.0	65.9	2	0
4	78.2	78.4	1	2
5	91.4	91.8	2	1

Figure 4.6 Summary of the first five plate bending modes.

Figure 4.7 shows the first five bending modes of the flat plate measured experimentally and processed in the STARModal analysis software. The 30 test locations were not centered exactly on the plate, so the mode shapes are not quite symmetrical on the wireframe model. The data are of excellent quality, and show the correct second order bending of the plate in modes 2 and 3.

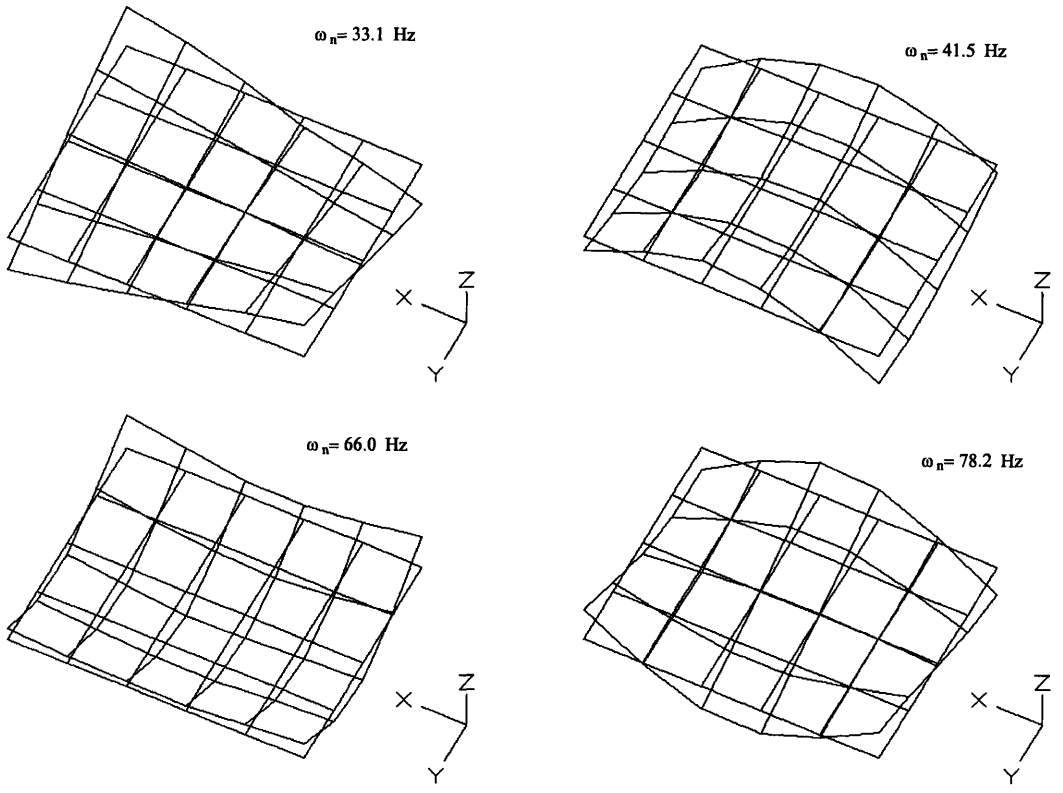

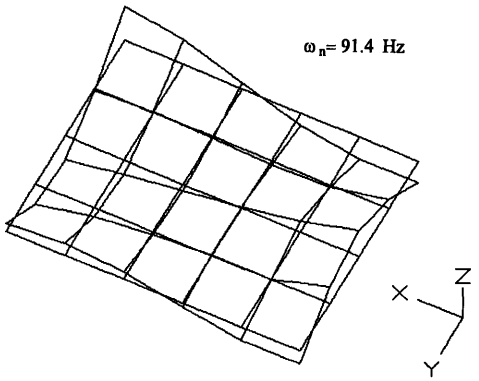

Figure 4.7 First five bending modes of the plate from experimental modal analysis.

Based on the close agreement of the ANSYS results and experimental modal data, the model is considered to be accurate and the viscous fluid damped plate may now be considered.

4.2.2 Viscous Fluid Damped (Double Plate) Analysis

The analysis of the fluid damped plate will proceed by first considering the finite element model and then the experimental results. As will be shown, the finite element model is very effective at predicting the amount of damping available in a fluid damped plate. The plate configuration used for the comparison is the same as the single plate tests, except now two plates are used with a thin fluid layer between them.

Thin plates require large numbers of brick elements to keep the element aspect ratio reasonably low. The wavefront of such a model can become large enough that a PC cannot quickly solve the problem because of time and hardware constraints. In order to simulate the performance of a thin plate assemblage with very thin fluid film thicknesses, a finer mesh is required. In practice, the fluid damping mechanism does not introduce high shear stresses for realizable film thicknesses. As a result, the mesh density requirements are greatly reduced and a PC can easily compute the amount of damping that may be obtained from a shear damped plate.

4.2.2.1 Finite Element Analysis

The finite element analysis of the fluid damped plate was made with a model containing the parameters listed in Figure 4.8. Note the correction factor that is needed to account for the slight decrease in the apparent viscosity of the fluid at higher shear rates. The plate modes encountered in this study were in the range of 30 to 100 Hz. The apparent viscosity at 50 Hz is 80 percent of the nominal viscosity.

Plate width	41 cm
Plate length	51 cm
Plate thickness	0.5 mm
Plate density	1200 kg/m^3
Plate modulus	4.25 GPa
Plate Poisson's ratio	0.25
Fluid thickness	500 microns
Fluid viscosity	100 kg/m-s
Fluid frequency dependence correction factor	80 %
Nominal 8-node brick element size	2 ×2 × 0.5 cm

Figure 4.8 Fluid damped plate geometry and materials.

The finite element simulations of the fluid damped plate show two important results. The first is that the mode shapes are essentially identical between the single plate and the fluid damped, double plate. The second result is that the amount of damping available from a plate design shows the same linear behavior as the fluid damped beams. As expected, if the fluid film thickness is halved, then the amount of damping is doubled. The loss factors obtained from the finite element analysis will be shown in the next section for comparison with the experimental measurements.

4.2.2.2 Experimental Modal Analysis of the Damped Plate

The shear damped, double plate was tested experimentally using the same modal analysis setup as the single plate testing. Data were taken on both plates to investigate the compliance of the fluid layer in the transverse (vertical) direction. Figure 4.9 shows the setup used in the data collection and modal analysis.

Frequency span	15 to 115 Hz
Frequency resolution	0.125 Hz
Excitation waveform	Burst random (90 percent duration)
Windowing	Uniform (none)
Number of averages	25
Curve fit	Polyreference (frequency domain)

Figure 4.9 Data collection and modal analysis parameters.

Figure 4.10 shows the results of the experimental tests alongside the results from the finite elements analysis.

Mode no.	ANSYS nat. freq. (Hz)	ANSYS damping	Experimental nat. freq. (Hz)	Experimental damping
1	32.8	10.2%	35	11%
2	41.0	15.4%	45	15%
3	66.2	14.1%	66	10%
4	78.7	10.3%	81	10%

Figure 4.10 Summary of finite element and experimental damped plate analysis.

As can be seen in Figure 4.10, the finite element results are in reasonable agreement with the experimental measurements. The slight error in the higher modes is due to the viscosity correction used in the finite element model. Therefore, the reported experimental natural frequencies and modal damping factors are less accurate for modes away from the corrected frequency.

In general, the modal damping is very high (first mode damping factor of 0.1, loss factor of 0.2). Additional measurements were made with other film thicknesses in an effort to achieve even higher levels of damping. Figure 4.11 shows a sample frequency response from these tests. In this case, all of the modal damping factors exceed 0.15 (a loss factor of 0.3). The mode numbers are shown for reference with the modes documented in the previous sections (note that the modal frequencies have increased).

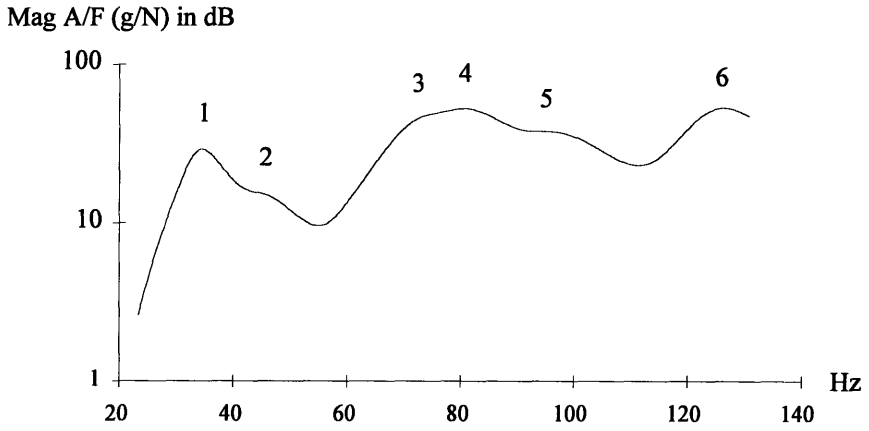

Figure 4.11 Drive point frequency response function of a well damped plate.

4.3 Experimental and Finite Element Correlation of Fluid Damping in a Slender Beam

A one meter beam was studied to find the optimal viscous fluid shear damper design. Two design questions were investigated during the optimization: 1) how does the low stiffness epoxy effect the amount of damping that may be obtained, and 2) what is the importance of the shear member thickness. The finite element method was used to

efficiently collect a large database of information, the results of which are shown on the following pages. Figure 4.12 shows the geometry of the beam studied in this optimization procedure.

Figure 4.12 Schematic of the one meter beam optimization parameters.

The composite beam studied in this example closely models an actual structure with viscous fluid damping. As shown in the schematic, there is a large core section and two shear members. In the real implementation, low stiffness epoxy layers are used to replicate the shear members into the main structure. The analyses performed in this optimization includes the lower stiffness of the epoxy layers to obtain results closely simulating actual structures.

The seven layer beam was optimized by first developing a very large database of dynamic performance results in ANSYS and then drawing conclusions from the results. Figure 4.13 shows the parameters investigated in the optimization.

Beam length	1 m
Beam width	75 mm
Boundary conditions	free
Nominal natural frequency	500 Hz
Nominal fluid viscosity	15 kg/m-s
Core height	75 mm
Fluid heights tested	octaves from 2.5 to 5000 microns
Shear member heights tested	3, 6, and 12 mm
Epoxy heights tested	0, 3, and 6 mm
Beam materials tested	aluminum, steel
Shear member materials tested	aluminum, steel

Figure 4.13 Seven layer beam parameters and variables.

A five variable test matrix was set up for the optimization problem: beam material, shear member material, fluid layer thickness, shear member height, and epoxy layer thickness. 400 finite element runs were performed (one for each combination of variables) and the natural frequency and damping factor of the first bending mode was recorded for each. The following four figures show the loss factor as a function of the other four variables.

Figure 4.14 shows the dynamic performance of an aluminum beam damped with steel shear members. This is the most logical combination of materials for achieving the maximum damping because the impedances of the beam core and shear members are the most favorable.

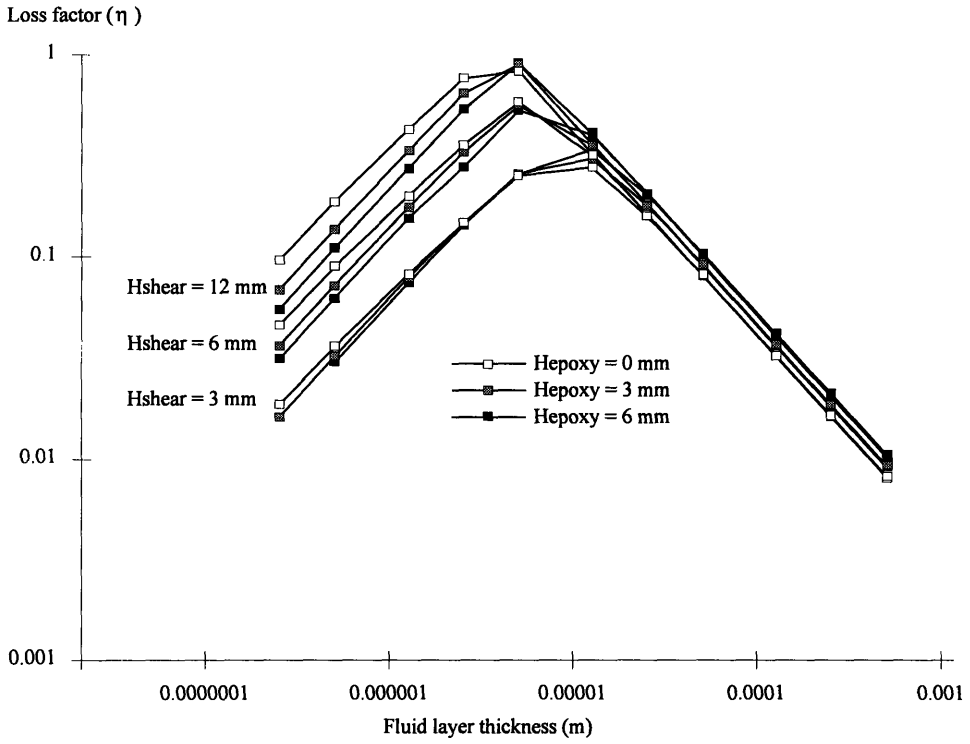

Figure 4.14 Loss factor as a function of fluid film thickness for an aluminum beam with steel shear members.

As shown in Figure 4.14, the damping in beams with fluid layer thicknesses above the critical thickness are not strongly affected by the shear member thickness. The performance of beams with thinner films shows greater dependence on the shear member height. Also shown is the minimal importance of the epoxy layer thickness, even below the critical film thickness.

Figure 4.15 shows the same plot made for a steel beam damped with steel shear members. The trends are similar to those shown in Figure 4.14, but the maximum damping levels are slightly lower.

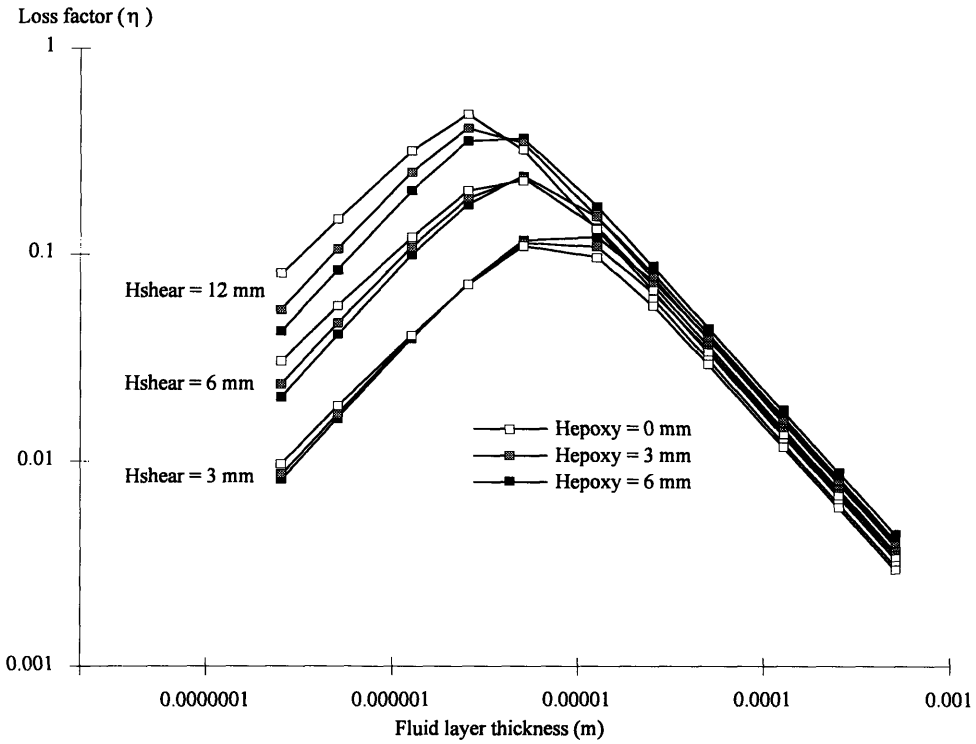

Figure 4.15 Loss factor as a function of fluid film thickness for a steel beam with steel shear members.

Figure 4.16 shows the damping results of a steel beam damped with aluminum shear members. This combination is the least favorable from the standpoint of impedance matching. Note that the maximum loss factor is less than 0.2 and that the critical film thicknesses occur at much higher levels than in the aluminum beam with steel shear members.

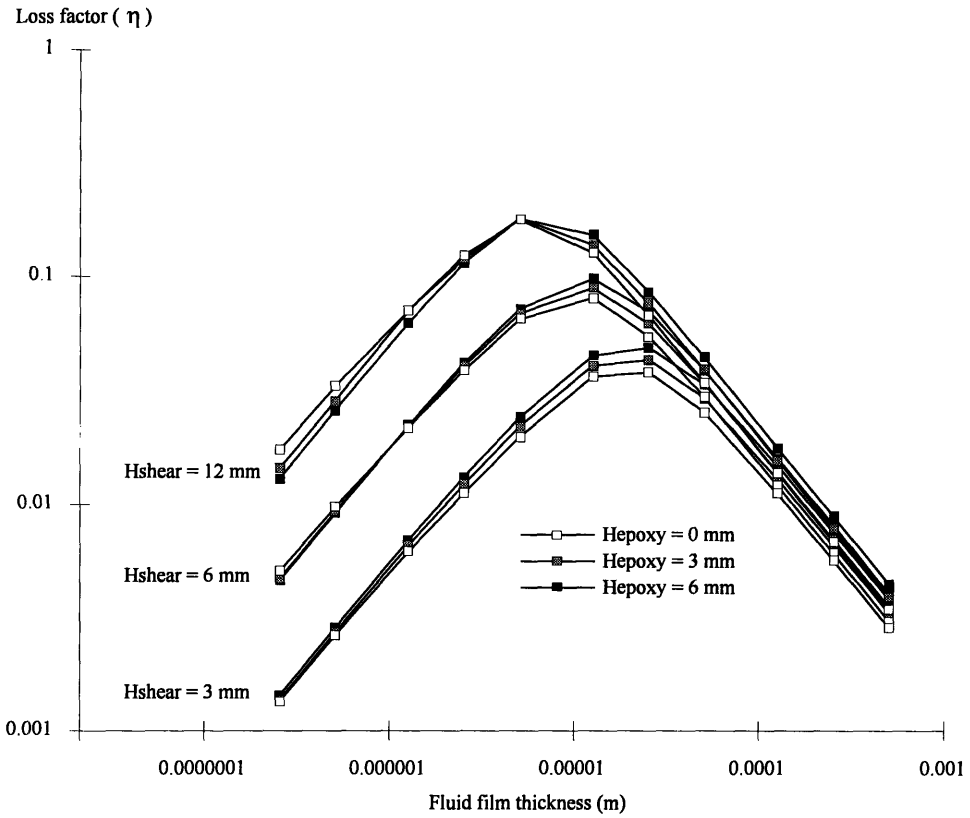

Figure 4.16 Loss factor as a function of fluid film thickness for a steel beam with aluminum shear members.

Figure 4.17 shows the final beam/shear member material combination (an aluminum core with aluminum shear members). Note that the all-aluminum beam performs very similarly to the all-steel beam in Figure 4.15.

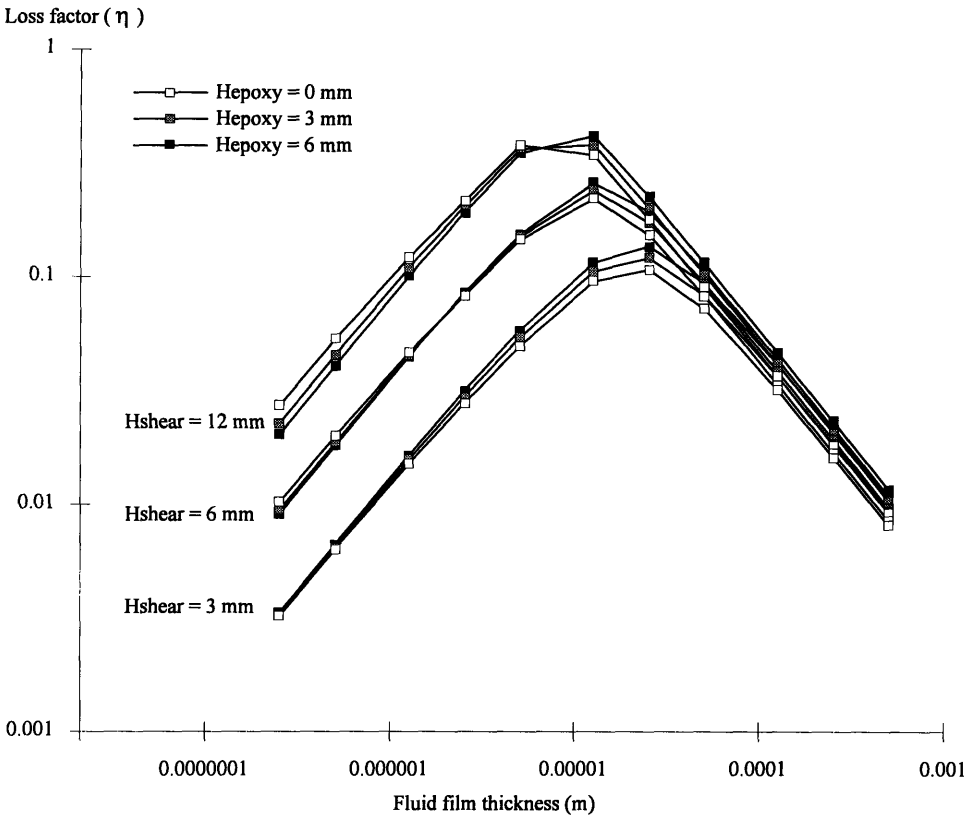

Figure 4.17 Loss factor as a function of fluid film thickness for an aluminum beam with aluminum shear members.

Investigation of the database revealed some important results concerning the performance of a multilayer beam:

1. The amount of damping that can be built into a typical structure is not significantly affected by thin layers of low stiffness epoxy. In general, the epoxy layer thickness is not important on either side of the critical fluid film thickness because the epoxy has a stiffness higher than the damping material.

2. The thickness of the shear members is important below the critical damping layer thickness. The thicker the shear member, the thinner the fluid layer must be to reach the critical thickness. Larger shear members are beneficial above the critical thickness because they offer a larger effective radius of gyration, below critical they resist elongation better than thin shear members.

3. The critical film thickness is typically far below the limit of the current manufacturing technology (for viscous fluids). For this reason, the behavior of a physically realizable beam will be adequately predicted by the general damping formula, which is valid for a structure of any number of layers.

Figure 4.18 shows the finite element results of a seven layer, fluid damped beam with an epoxy layer. Also plotted are the values obtained by the five layer beam analysis developed in Chapter 3.

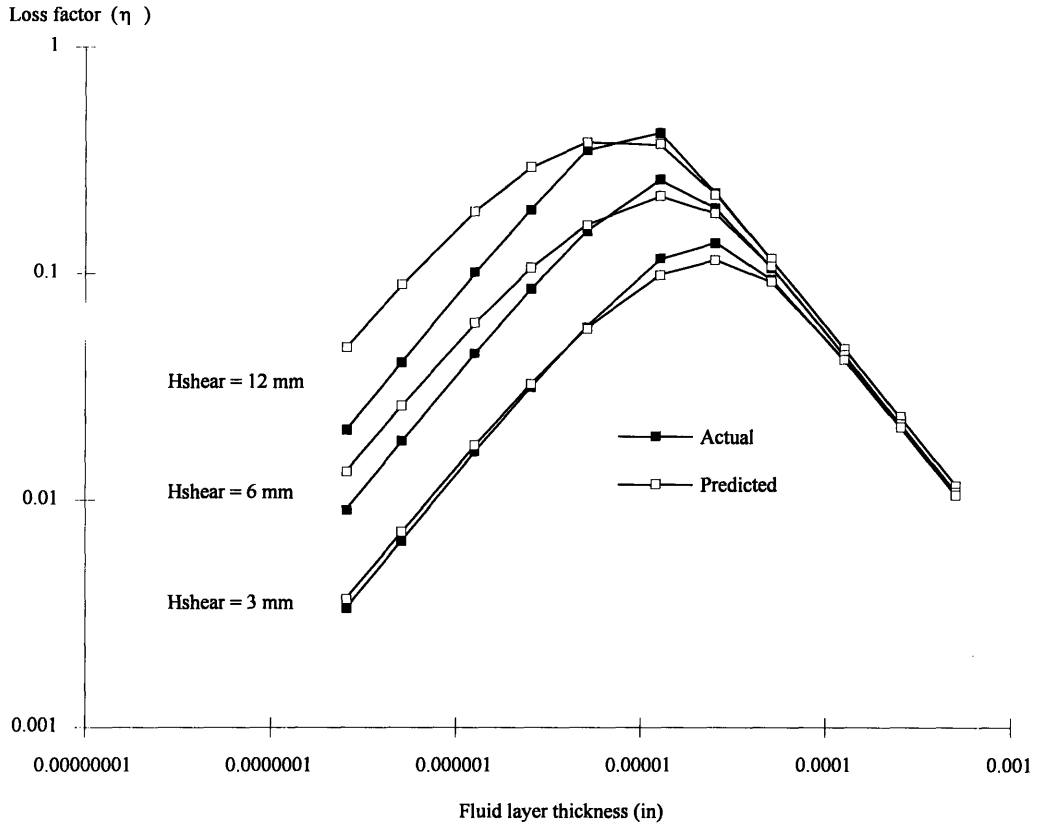

Figure 4.18 Predicted and actual loss factors as a function of the fluid thickness (an aluminum beam with aluminum shear members).

The five layer beam equation, when applied to the seven layer beams used in this study, works reasonably well. The five layer equation accurately predicts the critical thickness, and exactly matches the damping values for film thicknesses above critical. The damping in beams with fluid films thinner than the critical thickness are less accurately predicted, but the overall trend is still apparent.

4.3.1 Investigation of Quadruple Shear Member Design

A quadruple strip shear damper design was studied to investigate the effects of multiple damping layers in a structure. The one meter long geometry was kept, including the 75 mm core beam height. Figure 4.19 shows the two-dimensional model of the quadruple strip shear damper design.

Figure 4.19 Schematic of the quadruple shear member beam.

The study was performed in the same manner as the seven layer shear damping investigation outlined in the previous section; however, the epoxy layer thickness was kept at a constant 6 mm. This simplification was made because the results from the double strip damper designs indicate that the epoxy layer thickness is less important than other factors. The shear member thickness was also fixed at 6 mm.

As before, the beams were simulated in finite elements with both aluminum and steel cores. As expected, the aluminum core beams showed more damping than the steel core beams (by approximately a factor of three, the ratio between E_{steel}/E_{alum}). Figure 4.20 shows the results for aluminum core beams with double and quadruple aluminum damping strips.

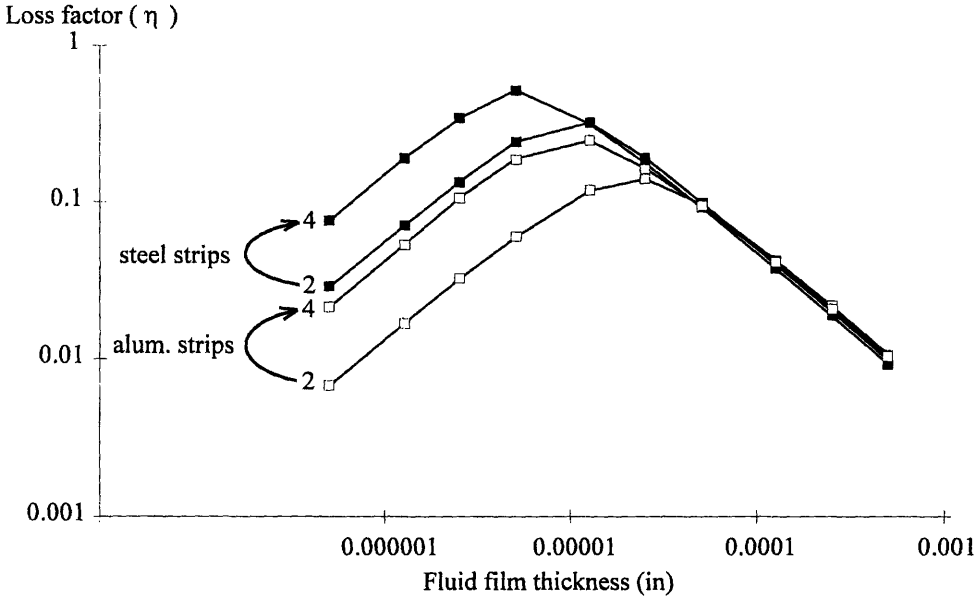

Figure 4.20 Damping factor of aluminum beam with 2 and 4 damping strips.

The interesting result of the quadruple damping member testing is that there is no benefit to having multiple layers of damping shear members above the critical damping layer thicknesses. Current manufacturing technology allows production of fluid films in the neighborhood of 25 microns thick, which is above critical for most materials. Viscoelastic materials are available that allow a particular design to reach or exceed the optimal damping layer thickness. As a result, using more than one damping layer only

adds damping in some viscoelastic designs. This simplifies the design of fluid damped structures.

4.4 Conclusion

The shear damping theory that has been presented in this dissertation has provided a closed-form equation for the loss factor that may be obtained in a given structure. The theory developed in Chapter 2 accurately predicts the amount of damping in a beam for *any* combination of beam materials and laminates. This theory is valid whenever the shear layer thickness is above the critical thickness (which is often only 10 microns for fluid damping applications). Chapter 3 provides an investigation of the stiffening effects for a five layer beam. This work generalizes the damping calculation for any damping layer thickness; however, the beam is assumed to be of a single material and constructed of five components of equal width (one main beam, two damping layers, and two symmetric shear members). This theory, while based on several specific assumptions, is useful because it provides an understanding of the coupling in a beam with very thin damping layers. Furthermore, the theory works very well even for beams that do not exactly fit the five layer model. For example, the seven layer beam shown in Figure 4.12 can be reasonably well predicted by the structural coupling model.

The results presented in this chapter illustrate the relative ease with which the shear damping mechanism can be modeled using finite elements. The close agreement between the finite element and analytical results further validate the analytical model of the damping mechanism. More importantly, the damping available in viscous fluid damped structures has been shown to be fully predictable without the burden of building and measuring a prototype.

4.5 References

McClintock, Frank, and Ali Argon, Mechanical Behavior of Materials, Addison-Wesley Publishing Co., Reading, MA, 1966.

4.6 Sample of ANSYS Finite Element Simulations

The ANSYS code that was used to make the plate and beam models has been included in this section so that the analysis technique is readily available to the designer.

4.6.1 Single Plate Model

```
W1=.61                                  ! Plate width
W2=.51                                  ! Plate depth
T1=.0054                                ! Brick thickness
D1=.02                                  ! Element size

/PREP7                                  ! Enter the preprocessor
/VIEW,,-1,.25,.9
ET,1,45
```

```
MP,EX,1,4.5E9                          ! Define material properties
MP,DENS,1,1190
MP,NUXY,1,.25
K,1,0,0                                ! Generate the plate volume
K,2,W1
K,3,W1,,W2
K,4,,,W2
K,5,0,T1
L,1,5
A,1,2,3,4
VDRAG,1,,,,,,1
ESIZE,D1
VMESH,ALL                              ! Mesh the volume
SAVE
FINISH

/SOLU                                  ! Enter the solver
ANTYPE,MODAL
MODOPT,SUBSP,10,5
SAVE
SOLVE
FINISH

/SOLU                                  ! Expand the mode shapes
EXPASS,ON
MXPAND,10
SOLVE
FINISH
```

4.6.2 Damped Double Plate ANSYS Model

```
W1=20.2*.0254                          ! Plate width
W2=16.2*.0254                          ! Plate depth
T1=.0054                               ! Plate thickness
GAP=.010*.0254                         ! Fluid thickness
D1=.06                                 ! Element size
FAC=.20                                ! Damping factor
EL1=NINT(W1/D1)                        ! Number of elements along side 1
EL2=NINT(W2/D1)                        ! Number of elements along side 2
NNUM=(EL1+1)*(EL2+1)                   ! Number of nodes in one plate surface
C1=FAC*600*(W1*W2/(EL1*EL2))/GAP       ! Damping coefficient

/PREP7                                 ! Enter the preprocessor
/VIEW,,-1,.25,.9
ET,1,45
ET,2,14,,1                             ! Damper x-dir element
ET,3,14,,3                             ! Damper z-dir element
ET,4,14,,2                             ! Compressibility spring element
R,2,0,C1
R,3,0,C1
R,4,1E5
MP,EX,1,4.5E9                          ! Define material properties
MP,DENS,1,1190
MP,NUXY,1,.25
K,1,0,0                                ! Generate the plate volume
KGEN,EL1+1,1,,,W1/EL1
KGEN,EL2+1,ALL,,,,,W2/EL2
KGEN,2,ALL,,,,GAP
```

```
*DO,INCR,1,NNUM                          ! Generate the x-direction dampers
L,INCR,NNUM+INCR
*ENDDO
TYPE,2 $REAL,2
LESIZE,ALL,GAP
LMESH,ALL

LGEN,3,1,NNUM,1,0,,,,1                   ! Generate the z-direction dampers
*DO,INC2,1,NNUM
L,2*NNUM+(2*INC2-1),2*NNUM+(2*INC2)
*ENDDO
LSEL,S,LINE,,NNUM+1,2*NNUM
LESIZE,ALL,GAP
TYPE,3 $REAL,3
LMESH,ALL

*DO,INC3,1,NNUM                          ! Generate the y-direction springs
L,4*NNUM+(2*INC3-1),4*NNUM+(2*INC3)
*ENDDO
LSEL,S,LINE,,2*NNUM+1,3*NNUM
LESIZE,ALL,GAP
TYPE,4 $REAL,4
LMESH,ALL

A,1,EL1+1,NNUM,NNUM-EL1                  ! Mesh the two plates
A,1+NNUM,EL1+1+NNUM,NNUM+NNUM,NNUM-EL1+NNUM
LSEL,S,LINE,,3*NNUM+1,3*NNUM+7,2
LESIZE,ALL,,,EL1
LSEL,S,LINE,,3*NNUM+2,3*NNUM+8,2
LESIZE,ALL,,,EL2
K,10000,0,-T1
K,10001,0,GAP+T1
L,1,10000
L,1+NNUM,10001
LSEL,S,LINE,,3*NNUM+9,3*NNUM+10,1
LESIZE,ALL,T1
ALLSEL
VDRAG,1,,,,,,3*NNUM+8+1
VDRAG,2,,,,,,3*NNUM+8+2
TYPE,1
VMESH,ALL

NUMMRG,NODE
wsort,all
SAVE
FINISH

/SOLU                                    ! Enter the solver
ANTYPE,MODAL
MODOPT,DAMP,16,5
SAVE
SOLVE
FINISH

/SOLU                                    ! Expand the mode shapes
EXPASS,ON
MXPAND,16
SOLVE
FINISH
```

4.6.3 Replicated Beam ANSYS Model

```
H1=3*.0254                              ! Beam height
L1=1                                    ! Beam length
GAP=.001*.0254                          ! Fluid thickness
HMOG=.25*.0254                          ! Replicant layer thickness
HCON=.25*.0254                          ! Shear member thickness
B1=3*.0254                              ! Beam width
BULK=2.2E9                              ! Fluid bulk modulus
LNUM=20                                 ! Number of elements along beam length
LNOD=LNUM+1                             ! Number of nodes along beam length
C1=15*B1*L1/(GAP*LNUM)                  ! Value of fluid damping dashpot
K1=BULK*B1*L1/(GAP*LNUM)                ! Value of compressibility of spring

/PREP7
ET,1,42,,,3                             ! Beam element model
ET,2,14,,1                              ! Fluid damping dashpot element
ET,3,14,,2                              ! Compressibility spring element
ET,4,14,,,2                             ! Support springs
R,1,B1
R,2,0,C1
R,3,K1
R,4,250
MP,EX,1, 70E9                           ! Define material properties
MP,DENS,1,2700
MP,NUXY,1,.3
MP,EX,2,.69E9
MP,DENS,2,1200
MP,NUXY,2,.3
MP,EX,3,205E9
MP,DENS,3,7850
MP,NUXY,3,.25
K,1,0,0                                 ! Generate the beam geometry
K,2,0,HCON
K,3,0,HCON+GAP
K,4,0,HCON+GAP+HMOG
K,5,0,HCON+GAP+HMOG+H1
K,6,0,HCON+GAP+HMOG+H1+HMOG
K,7,0,HCON+GAP+HMOG+H1+HMOG+GAP
K,8,0,HCON+GAP+HMOG+H1+HMOG+GAP+HCON
KGEN,2,1,8,1,L1
N,1,0,HCON                              ! Generate the dashpots dampers
N,2,0,HCON+GAP
NGEN,LNOD,2,1,2,1,L1/LNUM
N,2*LNOD+1,0,HCON+GAP+HMOG+H1+HMOG
N,2*LNOD+2,0,HCON+GAP+HMOG+H1+HMOG+GAP
NGEN,LNOD,2,2*LNOD+1,2*LNOD+2,1,L1/LNUM
N,4*LNOD+1,0,5*H1
N,4*LNOD+2,L1,5*H1
TYPE,2
REAL,2
E,1,2
EGEN,2*LNOD,2,1
TYPE,3
REAL,3
E,1,2
EGEN,2*LNOD,2,2*LNOD+1
TYPE,4
REAL,4
```

```
E,2*LNOD+2,4*LNOD+1
E,4*LNOD,4*LNOD+2
A,1,9,10,2
A,4,12,13,5
A,7,15,16,8
A,3,11,12,4
A,5,13,14,6
LESIZE,ALL,L1/LNUM
TYPE,1
REAL,1
MAT,1
AMESH,2
TYPE,1
REAL,1
MAT,3
AMESH,1,3,2                          ! Mesh the two plates
TYPE,1
REAL,1
MAT,2
AMESH,4,5,1
D,4*LNOD+1,UX,0,,4*LNOD+2,1,UY
NUMMRG,NODE,1E-7
WSORT,all
/pbc,all,1
EPLOT
SAVE
FINISH

/SOLU                                ! Enter the solver
ANTYPE,MODAL
MODOPT,DAMP,6,200
SOLVE
FINISH

/SOLU                                ! Expand the mode shapes
EXPASS,ON
MXPAND,6
SOLVE
FINISH
/CLEAR
```

Chapter 5: Use of the Shear Damping Mechanism to Control Boring Bar Chatter

5.1 Introduction

The shear damping mechanism was used to address the problem of boring bar chatter. A critical issue in boring bar design is the amount of overhang that may be safely used without inducing chatter in the tool. The overhang of a boring bar, the ratio between the length of the bar extending beyond the fixture and the bar diameter, is the dimensionless parameter used to decide which boring bar design is the most appropriate for a given task. Steel shank bars may be used with overhang ratios up to about 3:1 or 4:1 without chatter in a material such as medium steel. Much costlier tungsten carbide boring bars may be used with overhang ratios of 8:1. The high cost of the carbide bars has led to a great deal of work with passive damping treatments in steel shank bars. Historically, these innovations have met with varying success, usually at a significant added expense [Alev, 1969; Peter and VanHerck, 1969; New and Au, 1980; Rao, Rao, and Rao, 1988; Rivin and Kang, 1989].

In practice, chatter results from instability in boring bars performing boring/profiling operations. Compliant fixturing, incorrect machine spindle speed, improper feed rate, etc., can all result in excessive chatter and unacceptable surface finishes. A boring bar capable of precision turning at higher overhang ratios would be useful in one pass machining of parts such as journal bearing lands in engine blocks and deep profiles in long cylinders.

This effort was undertaken fairly early on in the shear damping research; the boring bar problem was an ambitious project, but the program has resulted in some important progress. At present, boring bars have been made with first bending mode loss factors of 0.3 in free-free vibration.

5.2 Background

The vibration mode which plagues deep boring operations is not the more common forced vibration that results from sources like spindle imbalance, but rather self-excited vibration. The forces which result in self-excited vibration originate from the cutting process itself, not an external source. This type of regenerative instability can occur when the tool encounters a small imperfection in the work piece. When the workpiece rotates 360 degrees in the lathe, the tool may not feed all the way past the imperfection so the tool will encounter what amounts to a slightly larger imperfection. This instability will continue until the tool leaves large chatter marks on the turned work (nonlinearity in the boring bar will eventually limit its maximum displacement).

This type of chatter represents a type of unstable closed-loop feedback between the work and the tool. An error is sent to the tool by the work and is fed back to the work by the tool. There are several known solutions to reduce the instability problem.

1. The most common method is to adjust the spindle speed so that natural frequencies of the boring bar are avoided. This will increase the dynamic stiffness of the boring bar, thereby reducing the tendency to chatter.

2. Adjusting the machine feed rate may also reduce chatter. In some cases, the boring bar can be effectively preloaded by increasing the feed rate of the machine axes, and actually reduce chatter. In other cases, the feed can be decreased to result in a smaller cutting force which can also decrease chatter.

3. The tool geometry can be adjusted to reduce the cutting forces, possibly at the expense of surface finish.

As the list indicates, tool chatter is a difficult problem to analyze and in many cases, there are no convenient solutions to real world problems. The motivation for applying the shear damping mechanism was to increase the dynamic stability of the boring bar so that tool chatter could be reduced for a variety of machining conditions.

5.3 Development of the Shear Damped Boring Bar

The shear damping concept has been shown to be an effective means of damping a structure undergoing bending vibration. Because a chattering boring bar shows bending vibration (as well as some torsional vibration), the damping mechanism was implemented in a number of boring bar designs to find the optimal shear damped configuration.

The finite element solutions outlined in the previous sections were unavailable at the time of the testing (the appropriate three-dimensional meshes were too large for PC finite element codes). Therefore, numerical analyses and some engineering judgment were used to develop a number of potential designs. The boring bars selected for this study are 35 cm long, 3.8 cm in diameter, and made of low carbon steel.

5.3.1 Design of the Shear Damped Boring Bar

The design constraints were determined by considering the static strength of the boring bar. Because the damping mechanism is placed on the inside of the hollow shanks, the maximum allowable core size is calculated. Figure 5.1 shows the natural frequency and stiffness of a steel shank as a function of the ratio of wall thickness to outside diameter.

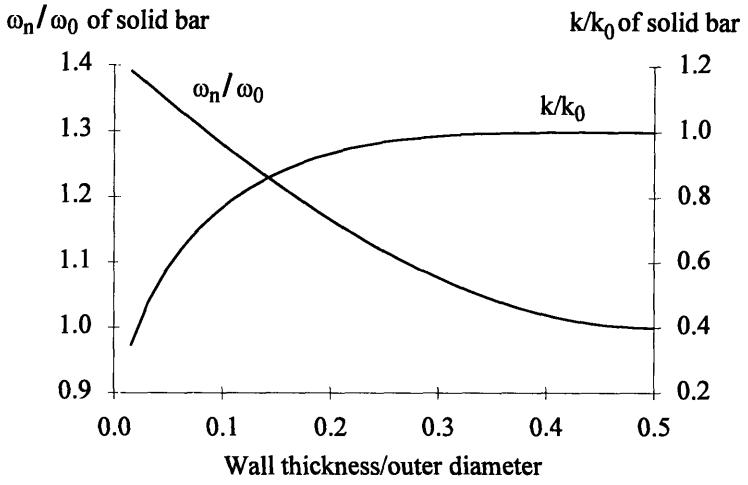

Figure 5.1 Dynamic performance of a steel boring bar shank as a function of wall thickness.

A convenient choice for the inside diameter of a 3.8 cm diameter steel shank is 2.54 cm. This corresponds to a wall thickness to outer diameter ratio of 0.167. The stiffness of the hollow shank is only 10 percent less than the stiffness of a solid bar and a reasonable amount of space is left inside the boring bar shank for the shear damping components. The strength to weight ratio of the hollow bar is greater than the solid bar, so the natural frequency is about 20 percent higher.

Many shear tube configurations were cast into the boring bars to find the best configuration. Shear members were designed and cast into the bars with replicating material filling the annulus between the shear members and the steel shank. Figure 5.2 shows three boring bar designs that were manufactured and tested in the lathe. Designs (a), (b), and (c) were tested with viscous fluid layers and designs (a) and (c) were tested with viscoelastic damping layers.

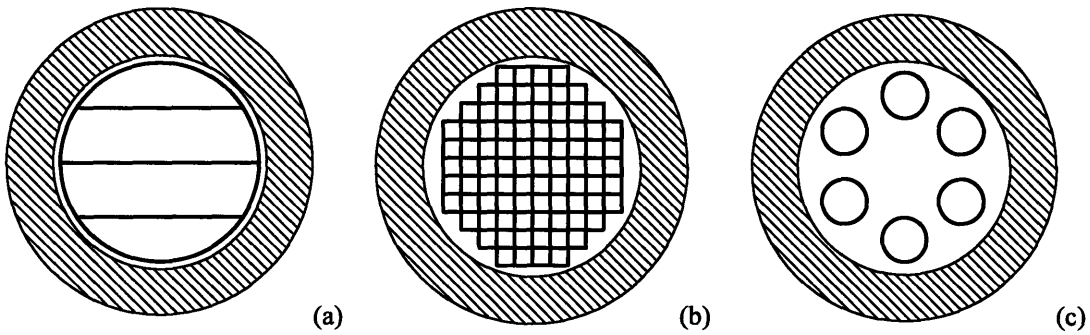

Figure 5.2 Boring bar designs using the viscous shear damping concept.

The quality of each boring bar design was assessed by two methods. The first method was a measurement of the damping in free-free bending vibration. The second

method was to run cutting tests in a NC lathe and examine the surface finish of the cut metal samples.

The free-free bending vibration performance indicated the amount of damping obtained from a particular shear member configuration. This testing was performed using experimental modal analysis. These tests ignored the effects of boring bar/lathe interaction and served as a good preliminary indication of the damping of each design. The performance of each bar was measured using an impulse hammer and an accelerometer.

The damped boring bars showed much better free vibration characteristics than the solid boring bars. Figure 5.3 shows a time response of a solid bar, as obtained from the manufacturer. As shown, the bar has very light damping (η = 0.0004). Figure 5.3 also shows a typical damped boring bar design, the second bar shown in Figure 5.2(a) (η = 0.3). The loss factor in the damped boring bar is almost 1000 times greater than the undamped bar.

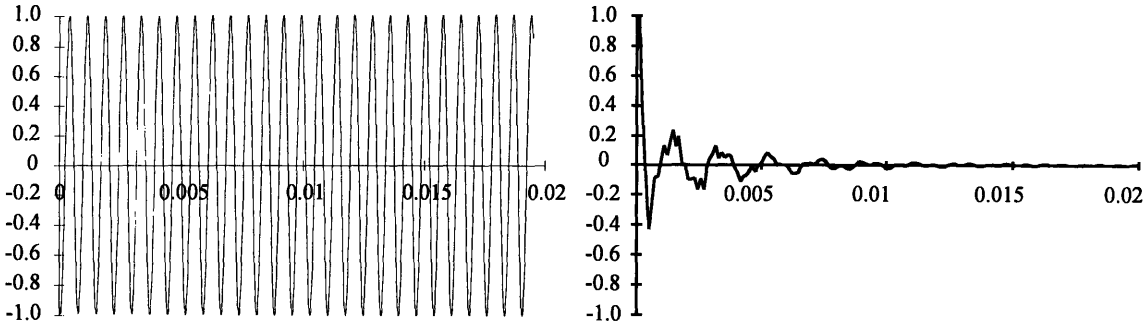

Figure 5.3 Vibration time history of a damped and undamped (solid) boring bar (in seconds).

The second part of the evaluative testing, the lathe cutting tests, were performed on hot-rolled 4140, 10 cm round stock. Two sets of tests were made on the 4140 steel samples for each bar: a heavy roughing cut and a light finishing cut. Figure 5.4 outlines the parameters of the cutting tests.

	Roughing cut	Finishing cut
Depth of cut	2.54 mm	0.64 mm
Spindle speed	350 fpm	350 fpm
Feed per revolution	125 microns	12.5 microns
Insert type	Valenite general purpose uncoated carbide (VC5)	Kennametal finishing ceramic /metal binder (KT175)

Figure 5.4 Lathe test cutting parameters.

All of the cutting tests were run with an overhang ratio of 8:1, regardless of the boring bar that was being tested (solid steel, damped steel, or tungsten carbide). In general, the carbide shank boring bars offer the best performance; however, in some cases the shear damped boring bars offer very good performance.

Figure 5.5 summarizes some typical results from cutting tests. The solid boring bar chattered in heavy cutting and produced unacceptable surface finishes. The shear damped designs sometimes gave excellent surface finishes under heavy cutting conditions (although the shear damped designs also showed heavy chatter in other trials). In finishing cut testing, the shear damped designs out-performed the solid bars (although the difference in performance is less pronounced). The surface finishes obtained in the finishing cut tests were not as good as the heavy cut tests.

	Roughing cut	Finishing cut
Solid bar	fair/poor	fair/poor
Cluster design Fig 5.2(b)	excellent/poor	good/poor
Slice design Fig 5.2(a)	excellent/poor	good/poor

Figure 5.5 Typical lathe test results.

Figure 5.5 shows two ratings for the shear damped boring bar designs (excellent/poor and good/poor). This is because the boring bars gave surface finishes that varied as the cutting inserts were worn in. During the cutting testing, it was found that the carbide inserts worked best if they were worn so that the cutting edge took on a honed surface. Both the solid and shear damped boring bars produced significantly better surfaces once the inserts had been worn in. Valenite and Kennametal confirmed that this is a common observation.

Neither the shear damped nor the solid steel boring bars were capable of repeatably producing chatter-free cuts with fresh inserts. When the inserts were worn in, both designs had an increased probability of producing good surface finish cuts (the difference being that even with a honed tool, the solid bars were difficult to tune for chatter-free performance). The shear damped boring bars had enough added stability to give much more predictable performance *when the tool was worn in.*

The result of this testing is that under certain conditions, the shear damped boring bars were found to produce excellent quality surface finishes, even in very heavy cuts (2.5 mm). This is a substantial improvement over conventional solid shank boring bars because the window of lathe operating conditions giving satisfactory cutting performance is larger. The difficulty is that the insert wear that is necessary to achieve these good results is not easily quantified, and somewhat tricky to achieve.

5.3.2 Evaluative Testing Using 5 cm Boring Bars

The results of the first round of lathe cutting tests suggests that the shear damped boring bar designs help increase the operating window in which cuts can be made without chatter. To further investigate the feasibility of using the damping mechanism in boring bars, a second round of cutting tests was performed with 5 cm diameter boring bars (loaned to the project by Kennametal).

The cutting tests with the 5 cm boring bars included a comparison of shear damped boring bars to commercially available bars such as tungsten carbide designs with impact dampers built into the shank. Figure 5.6 summarizes the size and construction of the three bars tested.

	Bar 1	Bar 2	Bar 3
Shank construction material	solid steel shank	tungsten carbide on steel core	hollow steel with shear damping components
Length	40 cm	50 cm	50 cm
Approx. Cost	$350	$3600	$700
Overhang ratio	3:1 to 4:1	6:1 to 8:1	???

Figure 5.6 Construction and geometry of 5 cm boring bars.

Figure 5.7 shows the shear damper design used in the modified boring bar. Note that this design is adapted from the 3.8 cm boring bar design that gave the best results in the first round vibration and cutting tests.

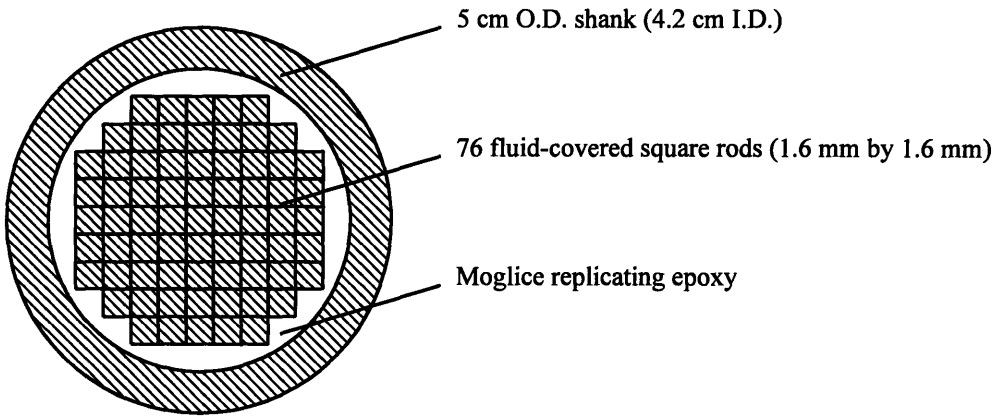

Figure 5.7 Cross section of 5 cm damped boring bar design.

A 3 × 3 × 3 cutting test matrix was constructed to evaluate the three boring bars' performance (this matrix is identical to the test procedure used by Kennametal's tooling engineers). The test matrix variables are cutting speed (in feet per minute), tool feed rate (in inches per revolution), and depth of cut (inches). The inserts used in the cutting tests

were identical and changed for each new boring bar. This means that each insert face was used for less than three minutes of actual cutting (the inserts were not used long enough to reach the honed state that was found to be desirable in the earlier testing).

Every effort was made to keep the cutting conditions consistent between tests with each bar. The boring bar designs featured an interchangeable head design so the same tool head was used with each bar. The repeatability of the cutting results was also investigated by completely disassembling the boring bar and insert holder and then reinstalling and repeating the cutting matrix with a fresh insert. The repeatability of the cutting tests was very good. All tests were run with coolant. The cutting material was hot rolled 4140 steel with a nominal diameter of 10 cm. The overhang ratio for the testing was 6.5:1.

An informal evaluation scheme was developed to quantify the quality of each cut surface. Figure 5.8 shows the scheme used in the testing. In general, any part with a surface finish rating of an A or B would be acceptable for many applications. Grades of D or better would be required for parts that would be finished with a final grinding operation. Machined parts receiving a grade of F generally had very poor surface finishes.

	Grade
Excellent, mirror-like finish	A
Very good, near mirror-like finish	B
Good, smooth finish with light scratches	C
Poor, fairly rough finish or smooth with light chatter	D
Chatter, unacceptable finish	F

Figure 5.8 Evaluation scheme used in round two cutting tests.

Figure 5.9 shows the test matrix results of the solid 5 cm boring bar. Note that most of the surface finishes received a D rating, indicating that the machined part was fairly rough.

Speed (fpm)	315			450			540		
Feed (ipr)	0.002	0.004	0.008	0.002	0.004	0.008	0.002	0.004	0.008
DOC: 0.020"	D+	D	D	D	D+	C	D	C	B
DOC: 0.030"	D	D	D	D	D+	C	D	C	B
DOC: 0.050"	D	D	D	D	C-	C	D	C	B
DOC: 0.100"	D+	D	C-	C	F	B+	B	F	B

Figure 5.9 Results of cutting tests with 6.5:1 overhang ratio (solid boring bar).

Figure 5.10 shows the results of the cutting test matrix with the shear damped boring bar. On the average, the results are very similar to the solid boring bar. These results verify the conclusions of the testing performed on the 3.8 cm boring bars; the fresh inserts do not work well with either the solid or the shear damped boring bar.

Speed (fpm)	315			450			540		
Feed (ipr)	0.002	0.004	0.008	0.002	0.004	0.008	0.002	0.004	0.008
DOC: 0.020"	D	D	D	D	D-	D	D	D	C-
DOC: 0.030"	D	D	D	D	D-	D+	D	D	C-
DOC: 0.050"	C-	D+	D+	C	D-	C-	C+	D-	C
DOC: 0.100"	D-	F	C	B	F	F	A	D	F

Figure 5.10 Results of cutting tests with 6.5:1 overhang ratio (shear damped boring bar).

Figure 5.11 shows the results of the cutting tests run on the tungsten carbide boring bar. With the exception of the tests run at the slowest spindle speed, at least one axis feed rate gave very good results for each depth of cut.

Speed (fpm)	315			450			540		
Feed (ipr)	0.002	0.004	0.008	0.002	0.004	0.008	0.002	0.004	0.008
DOC: 0.020"	D	D	D+	A-	C	C+	A	B+	B-
DOC: 0.030"	D	C-	C	C+	C+	B-	A	A	B
DOC: 0.050"	D-	C-	C+	A-	A-	B+	A	A	A-
DOC: 0.100"	A	B+	B-	B	A	B+	B	A	A-

Figure 5.11 Results of cutting tests with 6.5:1 overhang ratio (carbide boring bar).

This round of testing confirmed the findings of the earlier testing: the fresh cutting inserts gave poor results for the solid shank and shear damped boring bars. The carbide bars performed much better, but also had to be run within a certain operating window (that may not be known *a priori*).

5.3.3 Further Testing Using 5 cm Boring Bars

As a final investigation of the performance of the shear damped boring bars, the overhang ratio was extended to 8.5:1 and the shear damped boring bar and the tungsten carbide boring bar were re-run through the 3 × 3 × 3 cutting test matrix. This time, the shear damped boring bar was run with a fresh insert, as well as a "honed" insert for comparison. Figure 5.12 shows the shear damped boring bar results using fresh inserts.

Speed (fpm)	315			450			540		
Feed (ipr)	0.002	0.004	0.008	0.002	0.004	0.008	0.002	0.004	0.008
DOC: 0.020"	F	F	F	F	F	F	F	F	F
DOC: 0.030"	F	F	F	F	F	F	F	F	F
DOC: 0.050"	F	F	F	F	F	F	F	F	F
DOC: 0.100"	F	F	F	F	F	F	F	F	F

Figure 5.12 Results of tests with 8.5:1 overhang ratio (shear damped, new inserts).

Clearly the shear damped boring bar performed unacceptably when subjected to extreme cutting conditions. Figure 5.13 shows the results of the same boring bar run with the honed inserts.

Speed (fpm)	315			450			540		
Feed (ipr)	0.002	0.004	0.008	0.002	0.004	0.008	0.002	0.004	0.008
DOC: 0.020"	D	D	D	D	D	D+	D	D	C
DOC: 0.030"	D	D	D	D+	D	D+	D+	D	C
DOC: 0.050"	D	D+	D+	D	C-	C	B	B+	B
DOC: 0.100"	B	C+	C+	A	B+	B	B	B	A-

Figure 5.13 Results of tests with 8.5:1 overhang ratio (shear damped, worn in inserts).

With the very high overhang ratio, the honed inserts still gave some chatter, but in comparison to Figure 5.14, the shear damped boring bar with honed inserts showed performance similar to that of the carbide boring bar.

Speed (fpm)	315			450			540		
Feed (ipr)	0.002	0.004	0.008	0.002	0.004	0.008	0.002	0.004	0.008
DOC: 0.020"	B-	D	D+	B-	A	B	C	B	B
DOC: 0.030"	B-	B-	C	B-	A	B	F	B	A-
DOC: 0.050"	C	B	C+	B	A	B+	C	A	A-
DOC: 0.100"	D	B+	B	D	F	B+	F	A	A-

Figure 5.14 Results of tests with 8.5:1 overhang ratio (carbide boring bar).

The conclusion to be drawn from the extensive testing of the lathe boring bars (much of which is not documented here for brevity) is that the shear damped boring bar shows promise, but could not be made to make reliably good cuts without paying special attention to the cutting inserts. When the inserts were worn in, the performance of the shear damped boring bar approached that of the carbide bar. This promising result is important given the extreme expense of carbide tools.

5.4 Conclusion

The shear damping mechanism has been shown to effectively reduce vibration in boring bars and structural members such as laminated beams. However, the damped boring bars did not make a decisive improvement in the turning capabilities of the NC lathe. Several factors have come to light which explain this outcome. The first is that undamped boring bars, despite their inherently low internal damping, have reasonably high damping when fixtured in the lathe. This fact became apparent when a modal analysis was performed on the undamped boring bar side by side with the damped bar in the lathe turret.

The two bars (shear damped and solid) have virtually indistinguishable dynamic performance once fixtured in the lathe, as shown in Figure 5.15. The plotted quantity is the accelerance at the toolpoint of the boring bar as a function of frequency. The excitation was also applied at the toolpoint, making the plot a drive point measurement.

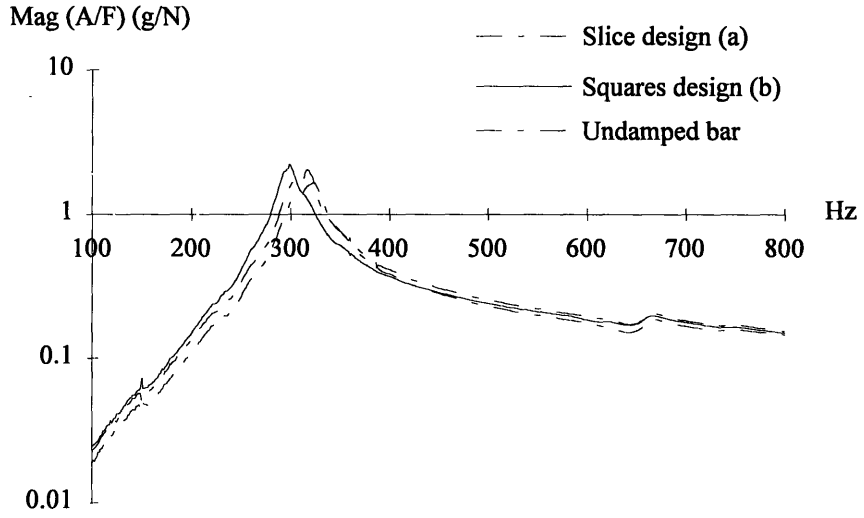

Figure 5.15 Frequency response of boring bars in NC lathe.

The results indicate that the total damping available in the fixtured boring bars is dominated by damping other than in the bar itself. In the case of the undamped bar, the damping of the bar is low, but the lathe damping is reasonably high, therefore the lathe damping dominates and the low damping of the bar is not apparent. Similarly, the damped boring bar, while offering a reasonable amount of damping, does not exceed the damping in the machine, so again the lathe damping dominates.

5.4.1 Next Generation Boring Bar

A second generation boring bar was developed near the close of this research. This boring bar benefited from a much improved understanding of the damping process and a better knowledge of potential damping materials.

Unlike the round one and two boring bars, which were made with silicone fluid damping material, the next generation boring bar was made with a thin viscoelastic layer in the configuration shown in Figure 5.16. As shown previously, the viscoelastic material is a much more effective energy dissipater.

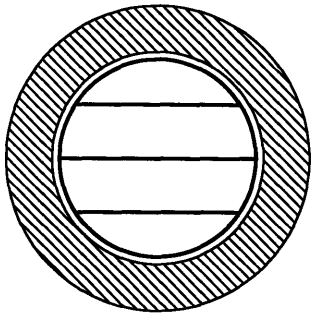

Figure 5.16 Boring bar design using viscoelastic damping.

When constructed with a 125 micron thick layer of 3M ScotchDamp ISD-112, this boring bar design gave a first bending mode loss factor of 0.25. However, lathe testing indicated an important issue in the design of this bar: although the free-free boundary condition testing of the bar gave an exceptionally high loss factor, the cantilever mode was not nearly as high. This is a result of the way the damping shear members are left free-floating in the steel shank. The shear members need to be attached at the end of the boring bar to maximize the shearing of the damping material down the length of the bar.

At the time of writing, a new boring bar with provisions to attach the shear members at the clamped end of the cantilever is being built.

5.5 References

Alev, Ali R., *Design and Devices for Chatter Free Boring Bars*, American Society of Tool and Manufacturing Engineers Conference, March 1969.

Bernett, Frank, *Tuned Array Vibration Absorber*, U. S. Patent No. 4,924,976, May 15, 1990.

Goodman, Lawrence E., *Material and Slip Damping*, Shock and Vibration Handbook, Third Edition, McGraw-Hill, New York, 1988.

New, R. W., and Y. H. J. Au, *"Chatter-Proof" Overhang Boring Bars-Stability Criteria and Design Procedure for a New Type of Damped Boring Bar*, Journal of Mechanical Design, Vol. 102, July, 1980.

Peters, J., and P. Vanherck, *Theory and Practice of Fluid Dampers in Machine Tools*, Proc. of 10th Int. Machine Tool Design and Research Conference, 1969.

Rao, P. N., U. R. K. Rao, and J. S. Rao, *Towards Improved Design of Boring Bars Part 1 and 2: Dynamic Cutting Force Model with Continuous System Analysis for the*

Boring Bar Performance, Int. Journal Machine Tools Manufacturing, Vol. 28. No. 1, 1988.

Rivin, Eugene, and Hongling Kang, *Improvement of Machining Conditions for Slender Parts by Tuned Dynamic Stiffness of Tool*, Int. Journal Machine Tools Manufacturing, Vol. 29. No. 3, 1988.

Chapter 6: Manufacturing with the Shear Damping Mechanism

6.1 Introduction

This chapter outlines several issues that must be addressed when manufacturing structures with the shear damping mechanism. These observations are the result of designing and building literally dozens of structures of various geometry during the course of this research:

1. If a viscous fluid is used as the damping medium, the layer should be of uniform thickness over the shear member. This is necessary to achieve the expected amount of damping. Several methods of applying a uniform fluid layer have been developed.

2. The replicating epoxy must be pourable into the structure. Vacuum and positive pressure assisted methods of forcing the epoxy into the structure were developed.

3. Sharp edges on the shear members must be protected so that the viscous fluid is not displaced during curing of the epoxy. Edge guards were developed to keep the fluid distributed over the shear members.

These manufacturing techniques will be discussed in the following chapter so that the shear damping mechanism may be fully understood by the reader.

6.2 Achieving a Uniform Fluid Coating on Shear Members

The fluid layer must have a smooth uniform thickness and coat all of the shear members prior to casting into a structure. The viscous behavior of the silicone family of fluids makes it easy to achieve a smooth layer given sufficient time, but the thickness of the film is harder to control. Over a period of days, a highly viscous film of fluid will flow down a vertically oriented shear member (before it is cast into the structure) so that a sufficiently thin layer is obtained; however, this process may be too slow for practical structural construction.

The Navier Stokes equations can be used to estimate the time to reach a given film thickness as a function of the fluid viscosity and the geometry of the shear member. In the case of beams that are being built up in laminates, the beams can be pre-assembled and pressed together to squeeze out excess fluid. This speeds up the process

considerably, and a solution is available to estimate the time required. Of course, a laminated beam may be checked with a micrometer to determine the exact fluid thickness.

During the development of the fluid damping mechanism, a third method of generating smooth, thin fluid films was invented. Provided that the shear member is of uniform cross section along its length, a tight fitting sleeve can be used to apply a thin film of fluid. The sleeves may be weighted so that they travel relatively quickly down the length. For example, a meter long shear member can be coated with a 125 micron fluid layer in about 30 minutes with a 2 kg sleeve. Many half meter shear members were also coated in times ranging from 5 to 15 minutes with various fluid film thicknesses and sleeve weights.

This method results in the most uniform fluid layers as well as the quickest application times. It is also the messiest method of applying the viscous fluid. At the time of the writing of this dissertation, applicator sleeves were being used exclusively for applying silicone fluids to shear members. Figure 6.1 shows a round shear member and a cylindrical applicator sleeve.

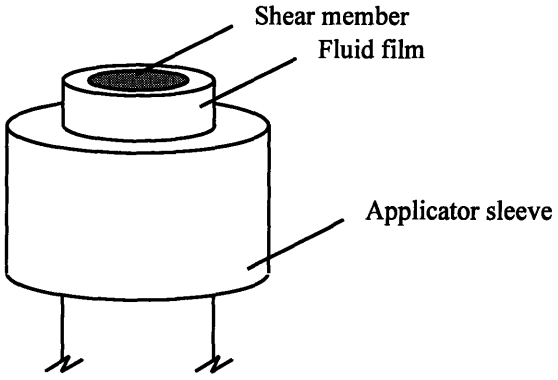

Figure 6.1 Fluid applicator on a round shear member.

The time to travel the length of a shear member can be found by considering the two forces acting on the applicator sleeve. The downward force is due to the weight of the sleeve; the upward force is a result of the shearing viscous fluid.

$$\frac{\mu A v}{h} = mg \tag{6.1}$$

The velocity is essentially uniform down the length of the shear member and can be approximated as the length of the tube divided by the total time. The solution of this equation for time yields a function of the viscosity and the dimensions of the shear member.

$$t = \frac{\mu A L}{mgh} \tag{6.2}$$

The quantity *mg* can be augmented to speed up the application process. If too much weight is added, the viscous fluid tends to tear, but this condition is easily avoided.

The applicators that were made during this research had small pockets that were milled to contain the excess fluid. The pocketed applicators were found to be self-centering on the shear members, yielding fluid films of constant thicknesses.

6.3 Casting Shear Members into Structures

A variety of structures were built in the course of this research. One of the goals of building so many structures was to ensure that virtually any design can be accommodated with the fluid damping mechanism. The use of replicating epoxy is a well known procedure in the literature, but a few items specific to the shear damping technique should be mentioned. The first is that epoxy has a limited work time (typically less than 30 minutes) so the structure should be prepared in advance with access holes and pouring funnels. Furthermore, an estimate should be made of how much epoxy is needed to fill the structure.

Several two component epoxies were used in this research. All had fairly thick consistencies, so thin crevices were difficult to fill. Gaps of about 6 mm were easily filled with epoxy, but gaps less than 4 mm did not fill properly. Long, horizontal structures do not fill under the force of gravity, but vacuum assisted pours can be made (even through annuli of 6 mm). Figure 6.2 shows a horizontal machine base structure that was filled under vacuum assist.

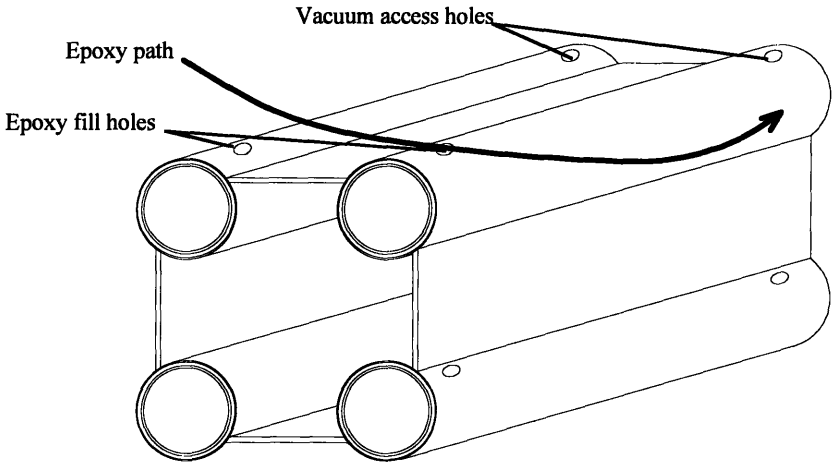

Figure 6.2 Epoxy filling technique for a horizontal structure.

The structure shown in Figure 6.2 was over a meter long, and a vacuum was applied to quickly fill it with epoxy. Larger structures are readily castable with similar or larger pumps (the largest structure tested was three meters long and was filled with a grout pump).

In the course of designing and building the fluid damped sample structures, a number of epoxy/silicone fluid combinations were evaluated. As shown in Chapter 2, silicone fluids of a wide range of viscosities converge to a similar reduced effective viscosity at high shear rates. Silicone fluids donated from GE Silicones of viscosities from 30,000 to 600,000 centiStokes were used. NuSil Corp. makes a 2,500,000 centiStokes fluid that was evaluated, but this costly specialty fluid is indistinguishable from the GE Viscasil 600,000 silicone above shear rates of 10 Hz.

A variety of manufacturers make viscoelastic materials that offer self-adhesive coatings for easy installation. Products from 3M and Soundcoat were used in this research with excellent results.

Chockfast Orange, made by ITW Philadelphia Resins, was found to be a cost effective epoxy to replicate shear tubes into the test structures. Of the epoxies tested, the Chockfast Orange is the least expensive and easiest to pour (due to its relatively low viscosity when freshly mixed).

6.4 Maintaining Fluid Layer Integrity During Epoxy Curing

Once a uniform fluid layer is applied to the shear members and the structure has been prepared with epoxy access holes, the shear members must be installed and cast such that the fluid layer remains intact. It was found during experimentation that the epoxy layers are distorted by the curing epoxy during the replicating process. The epoxy was forced away from sharp corners so that the edges of rectangular shear members were uncoated. Figure 6.3 shows a schematic of the silicone fluid before and after the epoxy cures.

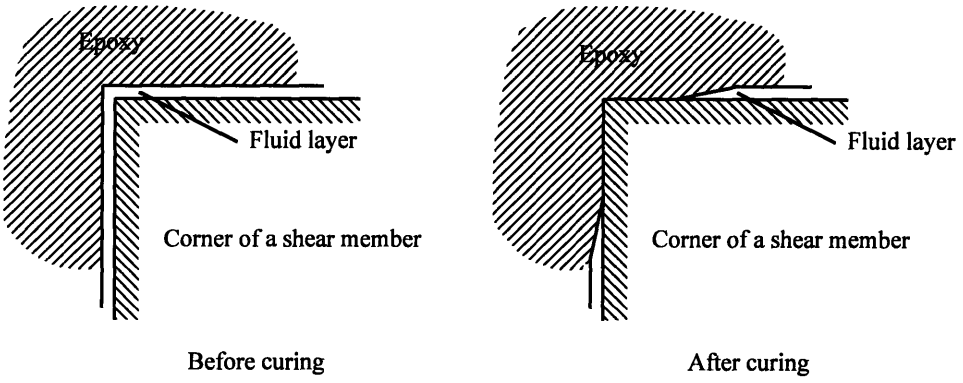

Figure 6.3 Schematics of fluid layers during the casting process.

The epoxies used to replicate the shear members inside structures have a small amount of shrinkage during the curing process. Furthermore, the shear members can be nicked during installation. Both the shrinkage and the nicks create spots on the shear members that stick to the epoxy because there is no silicone to act as mold release. Consequently, the shear members have to be broken free from the structure once the epoxy has cured. These forces were found to exceed 50 kN for small structures.

Figure 6.4 shows Instron data from small 20 cm long sample castings that were made with 50 × 6 mm steel strips cast into replicating epoxy (200 micron layer of silicone fluid).

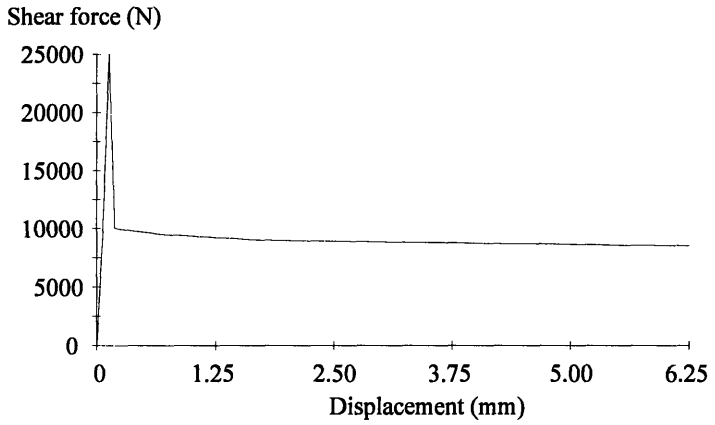

Figure 6.4 Breaking force of an epoxy replicated shear member.

The force that is expected for the shear member shown in Figure 6.4 is 100 N. Clearly, the shear member was not completely covered by the silicone fluid because the measured force was 25 kN. Note the very high peak at the beginning of the time trace corresponding to the breaking free of the shear member inside the structure.

The solution to this problem is to protect the edges of the shear members with plastic covers. The plastic is placed onto the shear members after they have been coated with the proper thickness of silicone fluid. Figure 6.5 shows how plastic edge guards may be used to protect the silicone fluid on a 5 cm wide steel shear member.

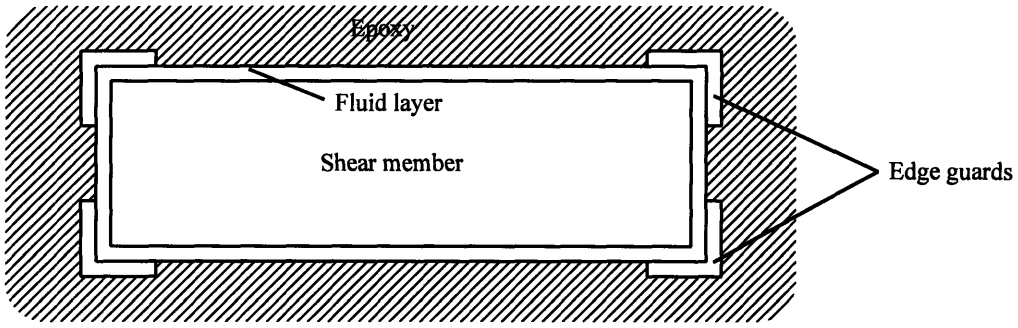

Figure 6.5 Edge guard covered shear member.

Once the edge guards are in place, the structure may be cast as before with the replicating epoxy. Not only do the edge guards protect the sharp corners of the shear members from being displaced, but the fluid films are less likely to scrape against the sides upon insertion into the structure. Figure 6.6 shows the Instron force history of a shear member protected with edge guards.

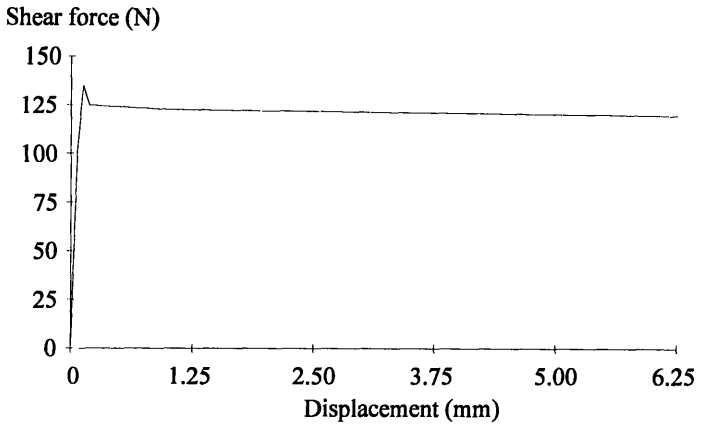

Figure 6.6 Shear force of an edge guard protected shear member.

Figure 6.6 shows the virtual elimination of the breaking free spike as shown in the time trace without the plastic edge guards. More importantly, the shear members are properly floating in the replicating epoxy and move with the expected force.

6.5 Conclusion

The manufacturing of shear damped structures has been investigated to ensure that the shear damping technique can be easily built into real world structures. The approach to casting the shear tubes into various structures has been tested with excellent results; these structures have shown the tremendous improvement in dynamic performance that the analytical model predicts.

The final test of the manufacturing technique was the production of a full scale machine tool base featuring the damping mechanism. This test was performed in cooperation with a sponsor of the shear damping research on a cylindrical grinding machine. The test results showed that the casting process is workable and that the damping of the first bending mode was very high (a loss factor higher than 0.1). This represents two orders of damping improvement over the undamped grinder base.

Chapter 7: Conclusion

This dissertation has provided a complete development of the shear damping mechanism including an analytical derivation for its performance in damping beam-like structures and a thorough investigation of the hardware implementation. A finite element model has been provided that allows other structural configurations to be analyzed and optimized. This novel means of incorporating damping into a structure has been shown to be very effective at reducing vibrations in structures.

The dynamic performance of a shear damped structure is much higher than equivalent untreated structures. Several machine tool scale structures have been built and tested that offer modal loss factors between 0.1 and 0.3, depending on the constraints of added weight and available space for shear members. Other key results have been observed in the damping mechanism:

1. Energy dissipation over a range of frequencies and vibration amplitudes.

2. Cohesive theory that allows designers to readily predict performance.

3. Manufacturability and placement on the inside of a structure.

4. Damping without compromising structural stiffness.

Future work with the shear damping mechanism will include completion of the boring bar stability problem, as well as adaptation to other structural damping problems. At the time of writing, a patent is being sought by MIT so that the technology may be licensed to American companies seeking improvement in the structural performance of their designs.

Appendix A: An Introduction to Experimental Modal Analysis

A.1 Introduction

The dynamic performance of a mechanical system is often dictated by a small number of modes of vibration. These modes may be modeled with an assemblage of lumped parameter elements. The resulting coupled differential equations of motion are traditionally linearized and written with three parameter matrices: mass, damping, and stiffness. A transformation of the equations characterizes the system with three new matrices: modal frequency, damping, and mode shape. Once a mathematical model has been determined analytically for a system, experimental modal analysis may be used to check the model on the physical hardware. Modal analysis is valid for linear, time invariant systems.

Experimental modal analysis begins by collecting data at discrete points on the test article. The data are obtained through the use of sensors that measure the response of the system to a known input. For example, an impact hammer can be used to provide a measurable force input and accelerometers will measure the output response. Data are collected in the form of time histories recording the impact and the resulting free vibration. The time histories can then be Fourier transformed to find the frequency response functions relating the output to the input. Figure A.1 shows a sample frequency response function with three modal peaks. The different peaks in the response represent vibration modes, each with its own modal frequency and damping factor.

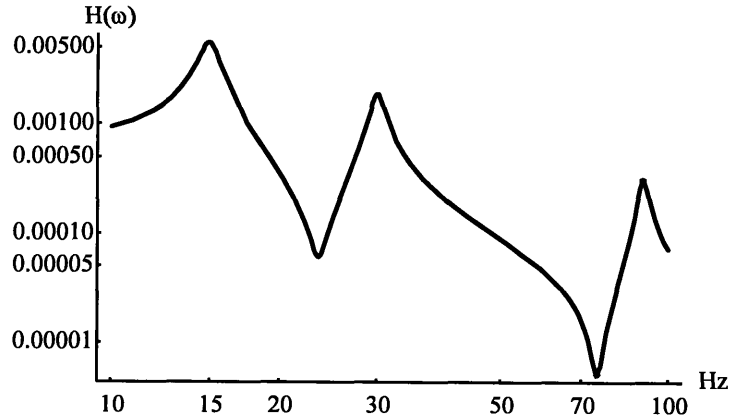

Figure A.1 A typical frequency response measurement.

A variety of data reduction techniques are available to solve the modal analysis problem. The most straightforward use frequency domain data to identify the modes of vibration. Within the broad range of frequency domain system identification techniques, the literature describes many curve fitting algorithms which may be classified as either single or multiple degree of freedom methods. Single degree of freedom methods may be conveniently applied whenever the modes are well separated. The modal peaks shown in Figure A.1 show fairly wide spacing (each mode may be easily distinguished from neighboring modes). Multiple degree of freedom methods are best used when the test article has closely spaced or well damped modes. Multiple degree of freedom methods are more complicated and are usually implemented on a computer.

Using single degree of freedom curve fitting techniques, each modal peak shown in Figure A.1 can be independently cast into the form of a SDOF system. The superposition of the three SDOF systems will be the combined response shown in the figure. The combined response may be written as:

$$H_{ij}(\omega) = \sum_{k=1}^{n} \frac{\phi_{ik}\phi_{jk}}{m_k\left(\omega_k^{\,2} - \omega^2 + j2\zeta_k\omega\omega_k\right)} \tag{A.1}$$

where:

$H_{ij}(\omega)$ = matrix of FRF's
ϕ = mode shape vector
ω_k = modal frequency
ζ_k = modal damping factor

The modal parameters ω_k and ζ_k can be obtained from a single FRF, but the mode shapes ϕ are obtained by fitting all the frequency response curves taken for a test article. Note that $n = 3$ for the FRF shown in Figure A.1.

The following sections will show the underlying theory of the modal analysis problem and the fundamentals of modal testing. The results from an experimental modal analysis are provided to illustrate the method.

A.2 Dynamics of a Single Degree of Freedom System

A single degree of freedom system is a mathematical idealization of a single mode of vibration. However, SDOF methods may be applied to structures with many vibration modes that are spaced sufficiently far apart in frequency that each mode may be modeled independently. For this reason, a study of the dynamics of a single degree of freedom system is important. The traditional SDOF model has a mass, spring, and dashpot as shown in Figure A.2.

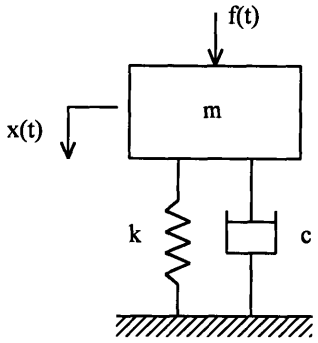

Figure A.2 Single degree of freedom mass-spring-dashpot model.

When the mass is vibrating, energy is passed between the mass (kinetic energy) and the spring (potential energy). Figure A.3 shows the kinetic and potential energy of the system under steady state sinusoidal excitation by the force $f(t)$.

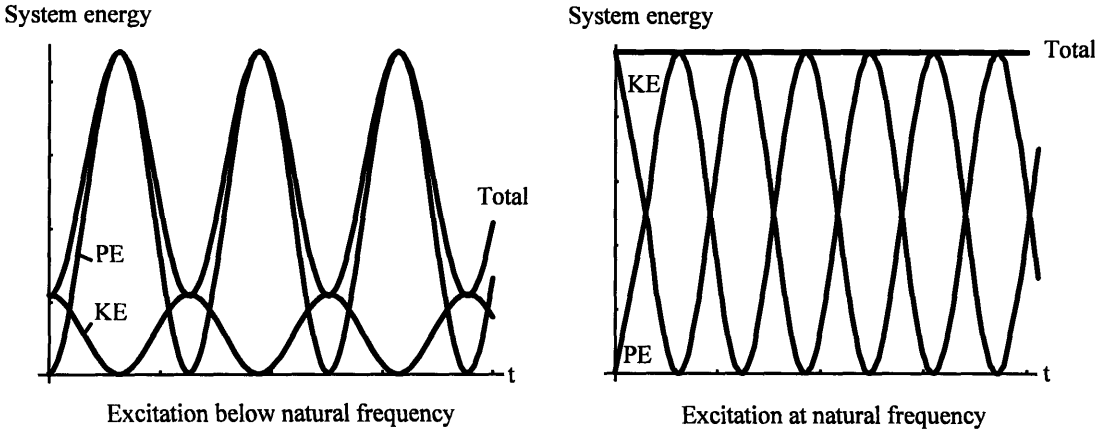

Figure A.3 Energy in a second order system under pure tone (sinusoidal) excitation.

The potential and kinetic energy in a mass-spring-dashpot system have equal magnitudes if the sinusoid excitation is at the natural frequency of the system. When the excitation is below the natural frequency, the potential energy will be greater; the kinetic energy will dominate above the natural frequency (the two sides of a modal peak are referred to as the stiffness and mass controlled regions for this reason).

The dashpot dissipates energy at a rate typically assumed to be proportional to velocity (viscous damping). Although there are other models of damping such as hysteritic and friction damping, viscous damping is often assumed because it is most conveniently cast into a workable analysis problem. Many mechanical structures have light enough damping to assume viscous damping without introducing large errors.

The equation of motion of the SDOF system may be obtained by a force balance acting on the mass. The spring restoring force and the damping force oppose the forcing function $f(t)$ to yield the classic equation of motion for a second order system.

$$m\ddot{x} + c\dot{x} + kx = f(t) \tag{A.2}$$

The under-damped, free vibration solution of this equation may be obtained by setting the force equal to zero and determining the response of the mass to initial conditions on x and $\dot{x}$.

$$\begin{aligned} x(t) &= e^{-\frac{c}{2m}t}\left(A\sin\sqrt{\tfrac{c^2}{m^2}-\tfrac{4k}{m}}t + B\cos\sqrt{\tfrac{c^2}{m^2}-\tfrac{4k}{m}}t\right) \\ x(t) &= e^{-\zeta\omega_n t}\left(A\sin\omega_n\sqrt{1-\zeta^2}t + B\cos\omega_n\sqrt{1-\zeta^2}t\right) \end{aligned} \tag{A.3}$$

where:

$$\begin{aligned} \zeta &= \text{damping factor} = \frac{c}{2\sqrt{km}} \approx \frac{1}{2Q} \\ \omega_n &= \text{natural frequency} = \sqrt{\frac{k}{m}} \end{aligned} \tag{A.4}$$

The smaller ζ is, the longer the settling time. Figure A.4 shows a typical under-damped impulse response of a single degree of freedom system. Virtually all mechanical systems exhibit damping factors much less than unity. A typical welded structure may have modal damping factors on the order of $\zeta = 0.001$. A bolted structure may have damping factors closer to $\zeta = 0.01$.

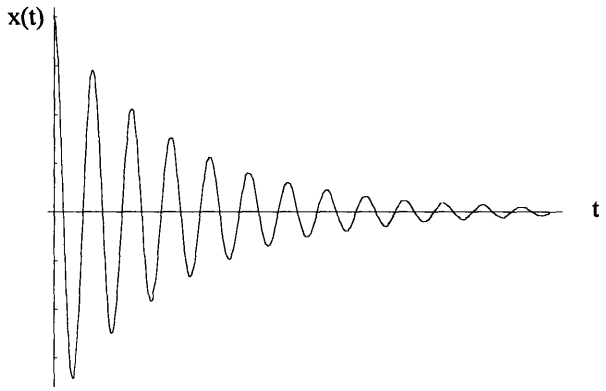

Figure A.4 Under-damped impulse response of a single degree of freedom system.

The steady state solution to the equation of motion, given a sinusoidal forcing function of amplitude F and frequency ω can also be calculated:

$$x(t) = \frac{F}{k}\frac{(\omega^2/\omega_n{}^2 - 1)\sin\omega t + jc\omega/k\cos\omega t}{(-\omega^2/\omega_n{}^2 + 1)^2 - (jc\omega/k)^2} \tag{A.5}$$

The Laplace transform of the time domain equation of motion may be easily written if the initial conditions of $x(t)$ are assumed to be zero. The ratio of output displacement to

input force is called the transfer function of the system. When the Laplace variable s is evaluated along the $j\omega$-axis, the transfer function is called a frequency response function.

$$H(\omega) = \frac{X(\omega)}{F(\omega)} = \frac{1}{m\left(-\omega^2 + 2j\zeta\omega\omega_n + \omega_n^2\right)} \tag{A.6}$$

The damped natural frequency and damping factor of this system is obtained by setting the denominator of the frequency response function equal to zero and solving for s. An under-damped single degree of freedom system has one complex pole pair. The real part (damping rate) of the complex pair controls the rate of the exponential decay of the impulse response envelope. The imaginary part (damped natural frequency) controls the frequency of the sinusoidal decay. Figure A.5 shows a complex pair plotted on the real and imaginary axes of the s-plane.

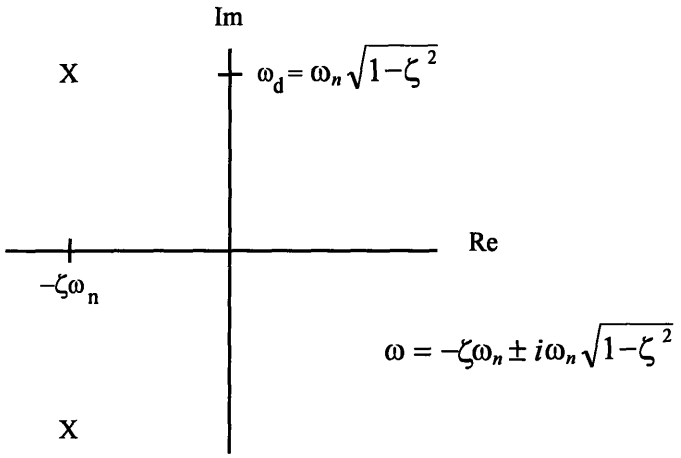

Figure A.5 Roots of a second order system plotted on the s-plane.

A.2.1 Design for Dynamic Response of a Second Order System

Examination of the second order system response offers valuable insight into the dynamic behavior of mechanical systems. An understanding of the behavior of second order systems is crucial when designing because the trends observed in changing the mass, stiffness, and damping of a second order system are valid for more complex structures.

The mass-spring-dashpot system shown above will be used again in the following discussion. Plots of the frequency response function magnitude will be presented as the system parameters are varied to provide insight into their effects on the system.

A.2.1.1 Effects of Changing Mass on the System

Intuitively, decreasing the mass of the system will enhance the ability of the system to respond quickly to command signals. The tradeoff is that the system will lose the ability to attenuate high frequency noise and vibration. This trend is shown in Figure A.6.

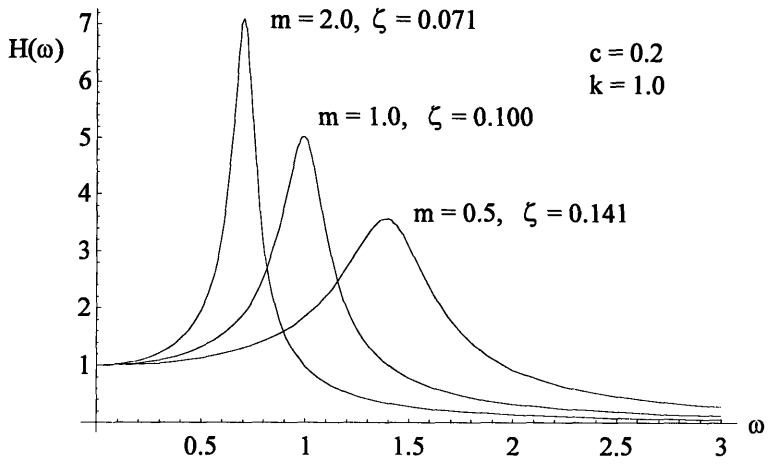

Figure A.6 Second order system response as a function of frequency and mass.

As shown in the figure, the system with less mass offers a higher natural frequency and damping factor. This means that higher frequency controller signals can be used without increasing the power requirements of the actuators. However, the low mass system has the greatest amplitude at frequencies above the peak frequency (the mass controlled region). This means that the system will be less able to attenuate high frequency noise and vibration. In machine design, high frequency vibration has a limited effect on the positioning accuracy of the system but can effect the acoustic quality of the design.

A.2.1.2 Effects of Changing Stiffness on the System

Figure A.7 shows the trends in second order system response as the stiffness is increased. Raising the stiffness increases the natural frequency and reduces vibration displacements for a given force input. At high frequencies, the compromise of decreased noise attenuation is not as dramatic as is the case with lowering the system mass (because high frequency behavior is dictated by the mass, not the stiffness of the system).

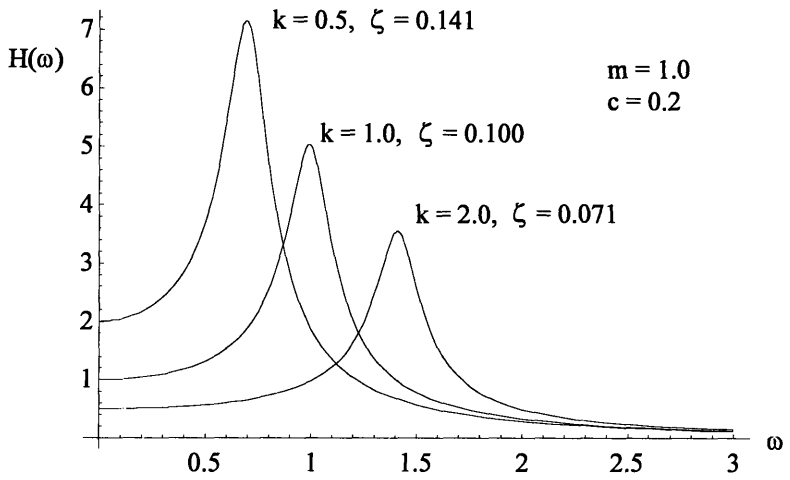

Figure A.7 Second order system response as a function of frequency and stiffness.

Figure A.7 suggests that raising the stiffness of a system is always desirable, but acoustical noise may be worsened by adding stiffness (because increasing the stiffness reduces the damping factor). Acoustical radiation can therefore become significant in the panels that enclose a machine. In general, structural components are built as stiffly as possible, whereas components likely to radiate noise are intentionally made less stiff (e.g., thin plate-like covers and panels).

A.2.1.3 Effects of Changing Damping on the System

Increasing the system damping makes a dramatic improvement in the system response near its natural frequency (the damping controlled region). Figure A.8 shows how increasing the damping lowers the amplification at resonance of the system. Although a damping factor of 0.2 is difficult to obtain in practice, the plot shows the dramatic improvement available by doubling the system damping.

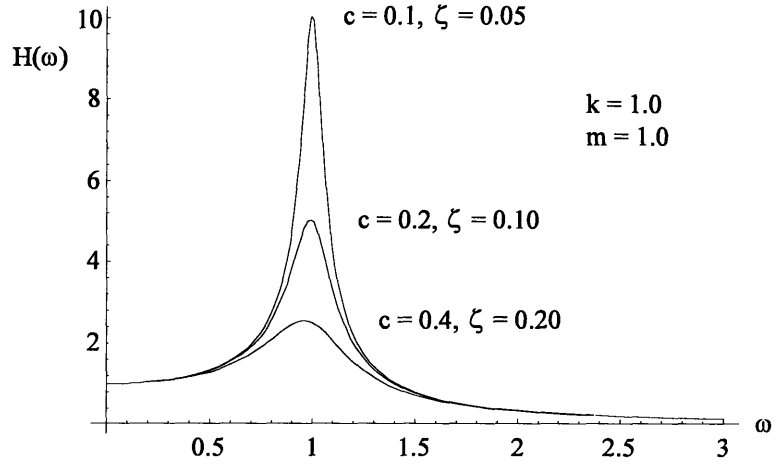

Figure A.8 Second order system response as a function of frequency and damping.

In summary, the stiffness of most structures should be maximized to improve accuracy and the mass should be minimized to reduce controller effort and improve the frequency response and damping factor (ζ). Damping must be present to attenuate vibration and compensate for the added stiffness.

A.3 Dynamics of a Multiple Degree of Freedom System

Like the single degree of freedom representation, a multiple degree of freedom system is used as a mathematical model of a continuous system. The techniques used to analyze a single degree of freedom system can still be used in the MDOF case; however, the scalar masses, damping values, and stiffnesses in the equation of motion become matrices. A sample MDOF system is shown in Figure A.9. The frequency response of this system is shown in Figure A.10, along with an illustration of how the individual modes are superimposed to equal the total frequency response.

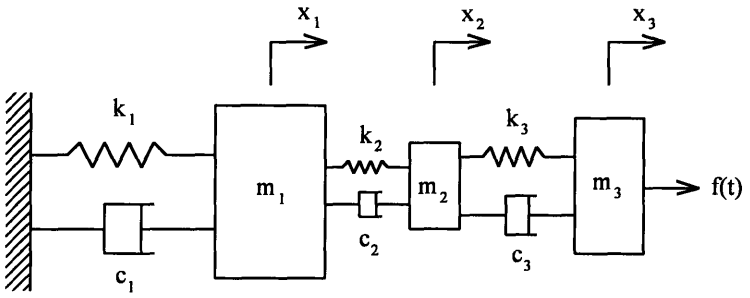

Figure A.9 Sample three degree of freedom system.

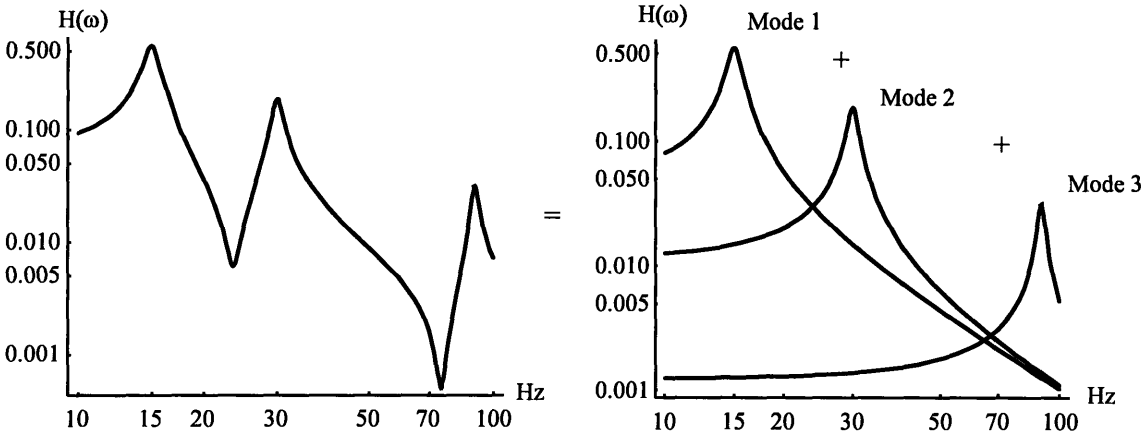

Figure A.10 Components of a MDOF frequency response.

A MDOF model is often a lumped parameter representation of a physical system. These lumped parameters consist of an assemblage of masses, dashpots, and springs. Using force balance, energy balance, or other techniques, the relationships between the

lumped parameters may be cast into the form of a set of coupled, linear differential equations of motion.

$$\mathbf{M}\ddot{x} + \mathbf{C}\dot{x} + \mathbf{K}x = f(t) \tag{A.7}$$

These time domain equations may be Fourier transformed to yield the frequency domain representation of the equations of motion:

$$\left(-\omega^2\mathbf{M} + j\omega\mathbf{C} + \mathbf{K}\right)X = F \tag{A.8}$$

The eigenvalues and eigenvectors may now be calculated to find the three modal parameter matrices (natural frequency, damping, and mode shape) of the system. This is done by solving the eigenproblem $-\omega^2\mathbf{M} + j\omega\mathbf{C} + \mathbf{K} = 0$. In the general case, the eigensolution will contain complex numbers. The eigenvectors of the system give the mode shapes of the different modes of vibration. The eigenvalues give the natural frequency and damping factor of each mode.

The damping matrix is often assumed to be proportional, meaning that the damping matrix is a linear combination of the mass and stiffness matrices ($\mathbf{C} = \alpha\mathbf{M} + \beta\mathbf{K}$). If $\mathbf{C}$ is proportional, then the equations of motion can be completely decoupled by pre- and post-multiplying by the appropriate eigenvector matrix Φ. Note that Φ is traditionally scaled so that the modal masses are unity ($\mathbf{I} = \Phi^T\mathbf{M}\Phi$).

$$\left(-\omega^2\mathbf{I} + j\omega\left[2\zeta\omega_n\right] + \left[\omega_n{}^2\right]\right)\mathrm{X} = \Phi^T\mathrm{F}\Phi \tag{A.9}$$

Because the left hand side square matrices are decoupled, the frequency response function $\mathbf{H}$ can be rewritten in its more familiar form.

$$\mathbf{H}(\omega) = \Phi\left[\omega_n{}^2 + j2\zeta\omega_n\omega - \omega^2\right]^{-1}\Phi^T \tag{A.10}$$

This matrix equation may be rewritten as a summation for each column of $\mathbf{H}$.

$$\mathbf{H}_{ij}(\omega) = \sum_{k=1}^{n} \frac{\phi_{ik}\phi_{jk}}{\omega_k{}^2 - \omega^2 + j2\zeta_k\omega\omega_k} \tag{A.11}$$

A.3.1 Example: Two Degree of Freedom System - Analytical Modal Analysis

A two degree of freedom system will be investigated to illustrate the fundamentals of analytical modal analysis. This example has two mass-spring-dashpot systems in series and will therefore have a pair of modal frequencies, damping factors, and mode shape vectors.

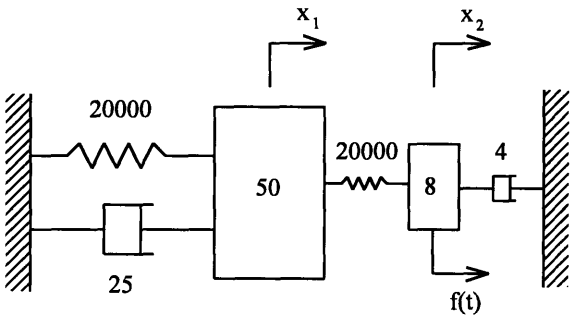

Figure A.11 Sample two degree of freedom system model.

The equation of motion of this system is:

$$\begin{bmatrix} m_1 & 0 \\ 0 & m_2 \end{bmatrix}\begin{Bmatrix} \ddot{x}_1 \\ \ddot{x}_2 \end{Bmatrix}+\begin{bmatrix} c_1 & 0 \\ 0 & c_2 \end{bmatrix}\begin{Bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{Bmatrix}+\begin{bmatrix} k_1+k_2 & -k_2 \\ -k_2 & k_2 \end{bmatrix}\begin{Bmatrix} x_1 \\ x_2 \end{Bmatrix}=\begin{Bmatrix} 0 \\ f \end{Bmatrix} \tag{A.12}$$

The Fourier transform of this equation may be written:

$$\begin{bmatrix} -\omega^2 m_1 + j\omega c_1 + k_1 + k_2 & -k_2 \\ -k_2 & -\omega^2 m_2 + j\omega c_2 + k_2 \end{bmatrix}\begin{Bmatrix} X_1 \\ X_2 \end{Bmatrix}=\begin{Bmatrix} 0 \\ F \end{Bmatrix} \tag{A.13}$$

The next step is to solve the eigenproblem of the free vibration. To do this, the forcing function is set to zero and the eigenvectors and eigenvalues of the 2 by 2 matrix are calculated. Since the problem was set up with proportional damping, the eigenvectors will be real-valued even though the roots of the determinate are complex. Because the eigenproblem becomes very complicated as the system order increases, the results are best calculated numerically.

$$[\omega_d]=\begin{bmatrix} 18.4 & 0 \\ 0 & 54.4 \end{bmatrix},\ [\zeta\omega_n]=\begin{bmatrix} 0.25 & 0 \\ 0 & 0.25 \end{bmatrix},\ \text{and}\ [\Phi]=\begin{bmatrix} 0.908 & -0.168 \\ 1.049 & 0.908 \end{bmatrix} \tag{A.14}$$

Figure A.12 shows the two mode shapes by plotting the location of the deformed system (in bold) over the location of the undeformed mass. Note that the first mode shows the two masses moving in phase at approximately the same amplitude. The second mode shows the two masses moving out of phase with the right hand mass moving much farther than the left hand mass. The deflections shown in the plot are unscaled; only the relative motion and phase between the two masses is given by the mode shapes.

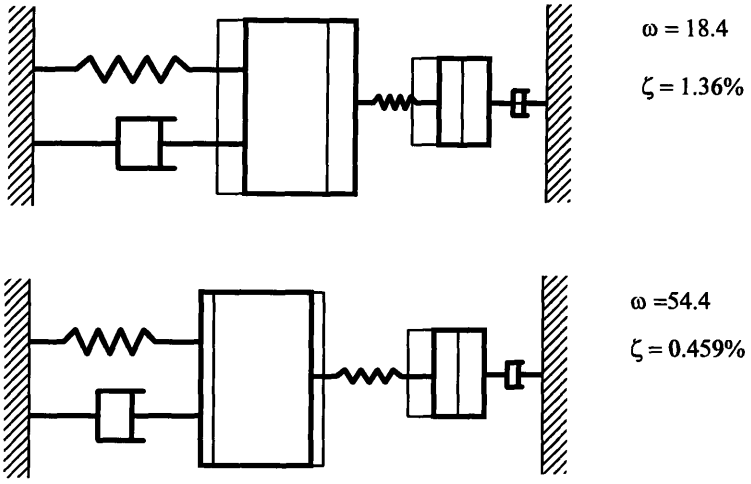

Figure A.12 Mode shapes of the two DOF system.

This example has shown the modal analysis of a lumped parameter model. If the same problem was solved experimentally the method would be the same, except experimental measurements be taken to determine **H**. To demonstrate this technique, the same problem will now be worked from experimental data assuming known frequency response functions for $X_1(\omega)/F(\omega)$ and $X_2(\omega)/F(\omega)$.

A.3.2 Example: Two Degree of Freedom System - Experimental Modal Analysis

The frequency response functions $X_1(\omega)/F(\omega)$ and $X_2(\omega)/F(\omega)$ may be obtained using a variety of experimental methods. One of the most common techniques is to use an accelerometer and a force exciter such as an electrodynamic shaker or an impulse hammer. The test article in this example requires only two frequency response functions (one for each degree of freedom; this system requires a measurement taken on each mass). The first FRF is obtained by exciting the first mass and measuring its response (a drive point measurement since the input and output are measured at the same location). The second FRF is obtained by again exciting the first mass but measuring the response of the second mass (a cross point measurement). Most portable signal analyzers have at least two channels (one for the input force transducer and one for the output response). A signal analyzer would typically calculate the two FRF's with 800 lines of resolution. Figure A.13 shows the two FRF's taken for this system.

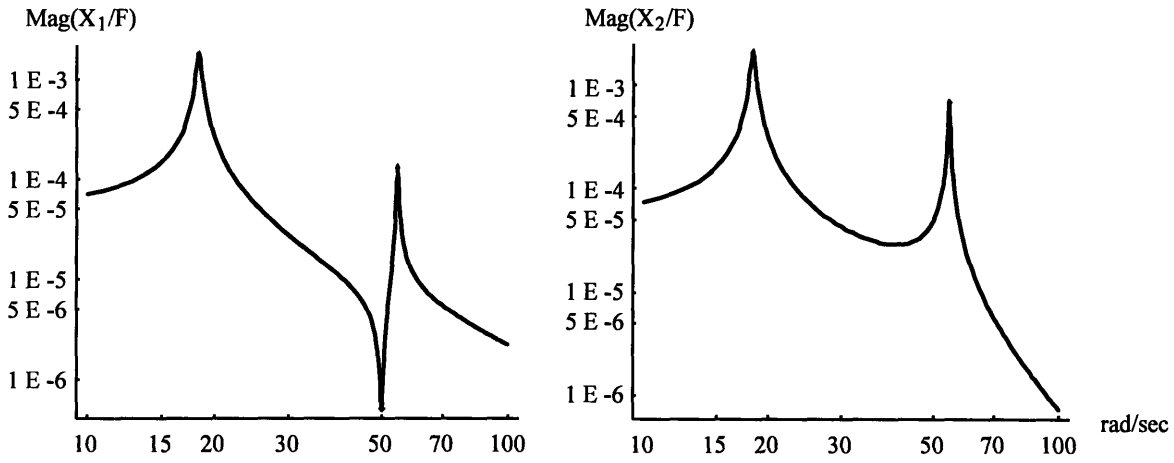

Figure A.13 Drive and cross point frequency responses of the example system.

The narrow peaks indicate that both modes have relatively low damping. The modes are also well separated so that a single degree of freedom analysis will give good results. While the magnitude plots in Figure A.13 show some of the important characteristics of the modes, a plot of the imaginary, or quadrature, part of the response shows even better separation of the two modes. The effects of coupling between the two modes are negligible in the quadrature plots shown in Figure A.14.

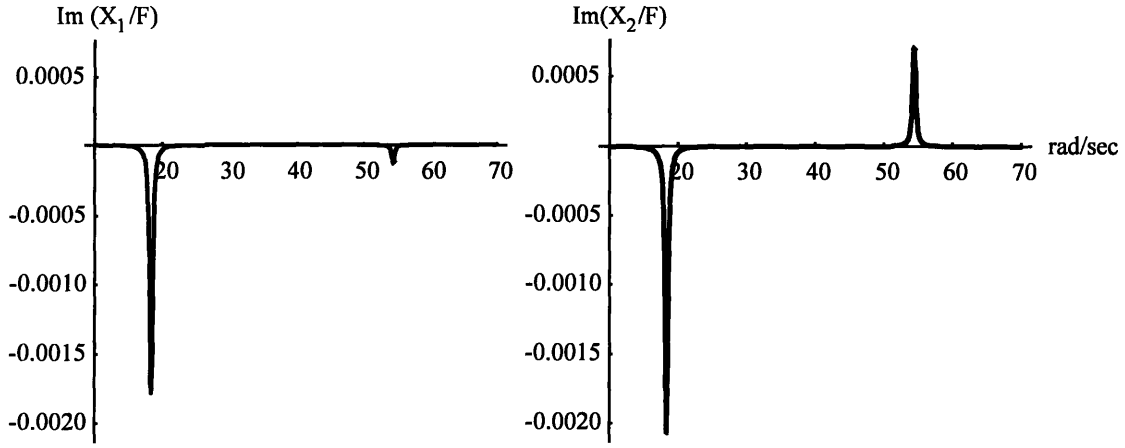

Figure A.14 The quadrature response of the drive and cross point FRF's.

The experimental modal analysis for lightly damped modes allows the modal natural frequencies, damping factors and mode shapes to be estimated using only the magnitude and imaginary response plots. A good estimate of the natural frequencies may be obtained by simply locating the peaks in the quadrature response. Using Figure A.14, the two peaks have natural frequencies around 18 and 54 rad/sec (2.9 and 8.6 Hz).

The damping factors may be estimated using the half power bandwidth of the frequency responses. The half power bandwidth relates the damping factor to the width

of a modal peak at $2^{-1/2}$ (70.7%) of the amplitude of the peak. Using Figure A.13, the first modal peak has an amplitude of 0.0018 m/N and the second peak has an amplitude of 0.00013 m/N. The frequency bandwidth at an amplitude of $0.0018 \times 2^{-1/2} = 0.00127$ and $0.00013 \times 2^{-1/2} = 0.0000919$ for the first and second modes are needed. These measurements can be read directly off the drive point frequency response plot. The formula for the half power bandwidth calculation for a force excited system is given by:

$$\frac{\Delta\omega}{\omega_n} = 2\zeta \tag{A.15}$$

The first mode frequency bandwidth is approximately 0.6 rad/sec; the second frequency bandwidth is also around 0.6 rad/sec. Using the natural frequencies that were obtained for each mode using the quadrature plots, this corresponds to a damping factor of 0.016 for the first mode and 0.0055 for the second mode.

The final modal parameters to be estimated from the experimental data are the mode shapes. A simple technique called quadrature peak picking is used in this example. Peak picking uses the quadrature plots of each measurement location and tracks the change in amplitude of the peaks. In the drive point data shown in Figure A.14, the first peak has an amplitude of –0.0018 m/N. The second peak has an amplitude of –0.00013 m/N. In the cross point measurement, the two peaks have an amplitude of –0.0020 m/N and 0.00070 m/N. These measurements are the unscaled mode shapes of the MDOF system:

$$[\Phi] = \begin{bmatrix} -0.0018 & -0.00013 \\ -0.0020 & 0.00070 \end{bmatrix} \tag{A.16}$$

Now that all of the modal parameters have been estimated a comparison can be made between the exact and experimental values. The results are summarized in Figure A.15.

	Experimental	Analytical
Natural frequency - first mode	18 rad/sec	18.4 rad/sec
Natural frequency - second mode	54 rad/sec	54.4 rad/sec
Damping factor - first mode	1.6 %	1.4 %
Damping factor - second mode	0.55 %	0.46 %
Mode shape - first mode	{0.88, 1.0}	{0.865, 1.00}
Mode shape - second mode	{-0.18, 1.0}	{-0.185, 1.00}

Figure A.15 Comparison of experimental and analytical modal parameters.

As shown in the table, the experimental modal estimates agree closely with the analytically determined values from the lumped parameter model. The modal damping factors would show better accuracy if the modes had more damping; the half power bandwidth method offers low resolution in poorly damped systems. Other means of estimating damping in poorly damped systems are available (e.g., circle fitting).

A.4 Experimental Data Collection

Quality modal test data result from careful selection of the most appropriate instrumentation and setup. Many times a little trial and error is the best way to approach a new test problem. If results are obtained two different ways the modal analyst is assured that a satisfactory approach is being taken.

This section is divided into two parts. The first part documents the instrumentation that is used in modal testing of mechanical structures, namely, the FFT signal analyzer, a force exciter (either an impulse hammer or shaker), and the transducers that measure input and output response (usually piezoelectric gages). The second part of this section describes the techniques that are used to transform the raw data from the sensors into usable frequency response functions. Although some modal tests require more sophisticated instrumentation such as non-contact sensors and multichannel data acquisition systems, many modal surveys can be handled with simple 2 channel equipment.

A.4.1 Instrumentation

The data used in an experimental modal survey are the input and output time histories to a system. The test article is treated as a black box with unknown modal parameters and the input and output time histories are used to identify them.

There are two common ways of generating input/output data for a structure. The first is to use an impact hammer at various points on the test article and measure the acceleration response at a fixed reference point. The second is to use a shaker attached at a fixed location and rove an accelerometer over the test article. Figure A.16 shows the test equipment setup for an impulse hammer modal test.

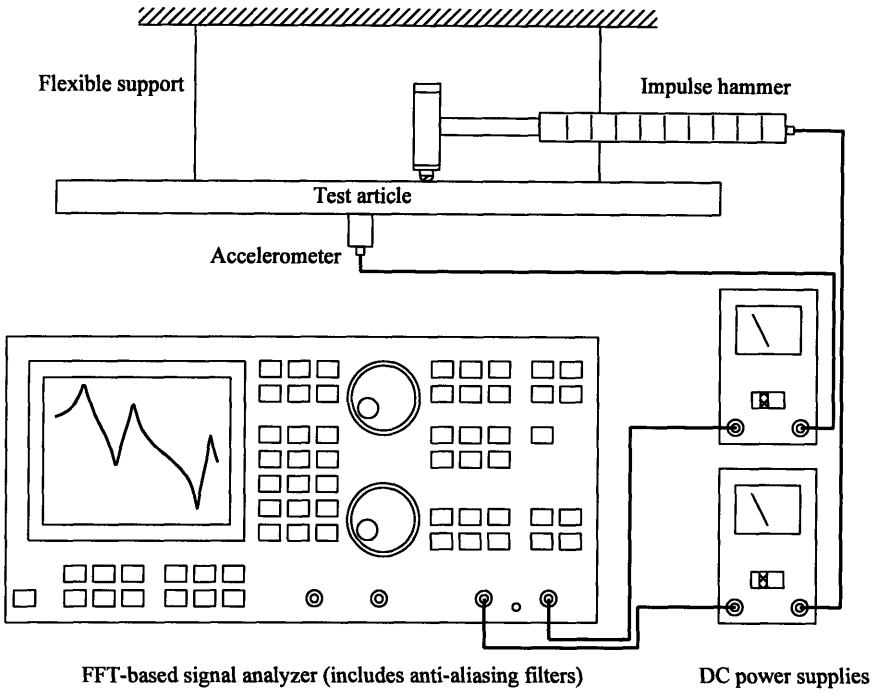

Figure A.16 A typical test article and modal test equipment.

Compliant support of the test article (simulating free boundaries) is often preferred because the article is isolated from background vibration. Free boundaries are also the easiest to simulate in finite elements for model verification. This method is shown in Figure A.16. Although the test article is connected to ground, very compliant supports such as bungy cords or rubber bands effectively decouple the structure from ground interaction. A structure truly floating in space would have rigid body modes at 0 Hz, but rubber bands make the rigid body modes occur at slightly higher frequencies. However, if the supports are compliant enough, the rigid body modes of the test article will be so low in frequency that they will not affect the structural response (the rule of thumb is that the rigid body modes of the test article should be less than one-third to one-tenth of the frequency of the first structural mode).

In some cases, the test article is so large that it cannot be supported with free boundary conditions. For example, large structures will often be tested with a connection to ground. In this case, it is important to investigate the behavior of the structure on its elastic base to see which degrees of freedom are constrained by the ground. An understanding of these boundary conditions is important, especially if the experimental results are going to be correlated to finite element results.

Once the test article has been suspended for testing, the next choice is the type of excitation that will be used as an input to the structure. Hammer testing is very popular because it is relatively simple to perform and the most economical. While shakers are more costly to purchase and use, they are the best way to input different types of forcing functions that maximize instrumentation accuracy and resolution. In practice, impulse

hammers are typically used in field testing because of the impracticality of setting up a heavy shaker for temporary use, but shakers are used in lab tests whenever possible.

An impulse hammer inputs a Dirac delta-like impulse to a structure once per sample period. Because a hammer blow closely approximates a delta function in the time domain, it has uniform frequency content over a wide bandwidth inversely proportional to the impact duration. The tradeoff is that a typical sample time may be 4 seconds, but the hammer blow might last 12 milliseconds. During the rest of the sample, only noise is recorded by the signal analyzer on the input channel. This noise can corrupt the data, especially if the noise has a large amplitude. The solution is to use a softer impact tip that increases the duration of the impact. Figure A.17 shows the frequency spectra of a hammer blow with three different impact tips.

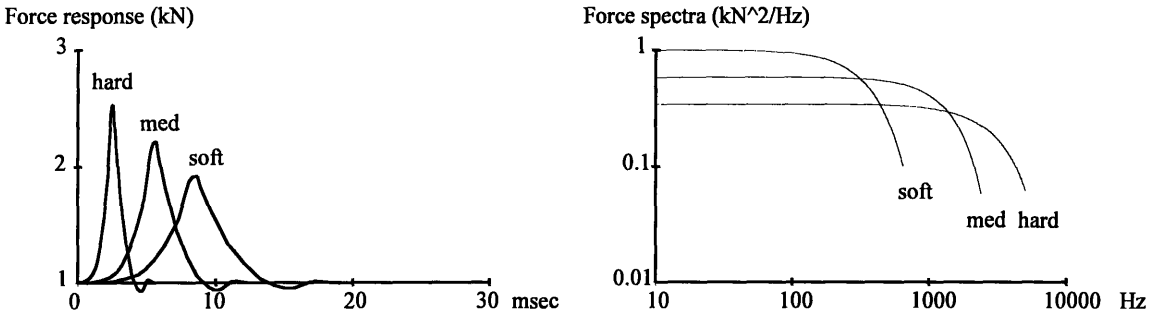

Figure A.17 Force spectra for different hammer tip hardnesses.

Unfortunately, longer impacts have a corresponding drop in the spectral bandwidth, as shown. The best practice is to use the softest tip that still contains enough high frequency content to excite all the modes under investigation. The correct hammer tip for a given test article can be found by experimentation.

An electrodynamic shaker allows many different types of forcing functions to be input to the test structure. The best choice of forcing function depends on the article, the frequency range of interest, and other factors that vary from test to test. One of the biggest concerns is that of minimizing the leakage in a measurement. Leakage is the result of Fourier transforming non-periodic components of a time history. If a wave is sampled an integer number of times during a data record, then the Fourier transform of this sequence has minimal leakage (there is always some leakage due to the finite length sampling period). If a wave is sampled a non-integer number of times during a data record, then the Fourier transform will show artificially widened frequency bandwidth. This is because the Fourier transform assumes that the signal repeats itself forever and is sampled an integer number of times. Figure A.18 shows two sine waves and their Fourier transforms. The upper sine wave has been sampled an integer number of times; the lower sine wave has been sampled with an extra half wavelength at the end. Note that the lower frequency spectrum shows a wide bandwidth even though the waveform is a pure

sinusoid (the kink at the end of the time series is caused by the assumption that the signal repeats itself, introducing wide band content into the spectrum).

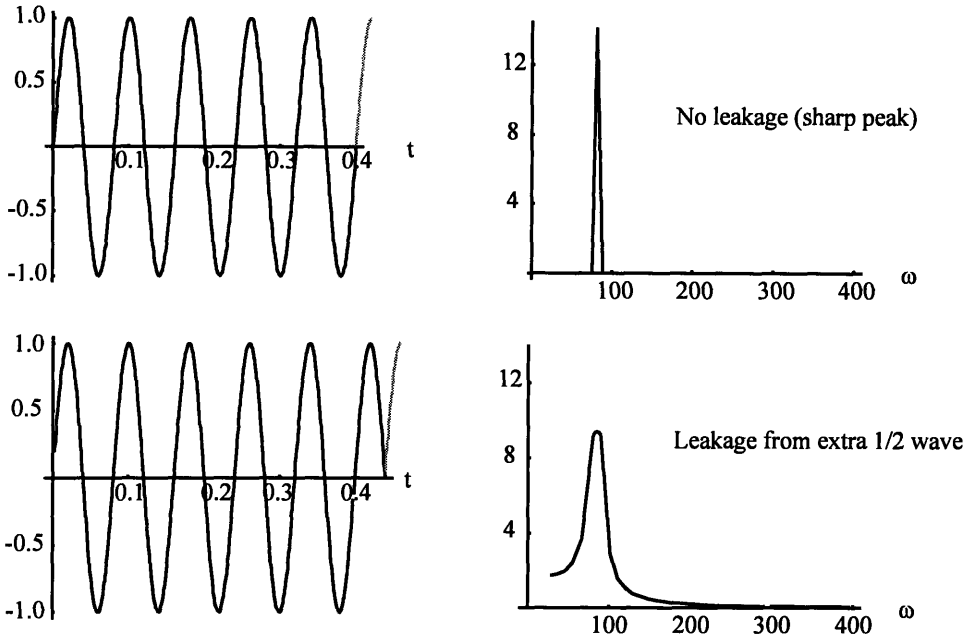

Figure A.18 Effects of leakage on a measurement.

There are three ways of minimizing the leakage in a measurement. The first is to carefully construct an excitation waveform that only contains components that will be sampled an integer number of times during the time record. This excitation waveform is called pseudo random.

An appropriate time window can also minimize leakage effects. Some time windows allow non-periodic components in the sampled waveform by filtering out the beginning and the end of the data record. This filtered record is Fourier transformed, not the raw time data. Both the input and output data must be filtered by the same window to minimize the error introduced by this technique. The final result will be much improved when compared to unfiltered data, but the windowing process introduces artificial smoothing (i.e., damping) to the measurements. This error may or may not be acceptable depending on the particular test. Figure A.19 shows an example of a filtered time series.

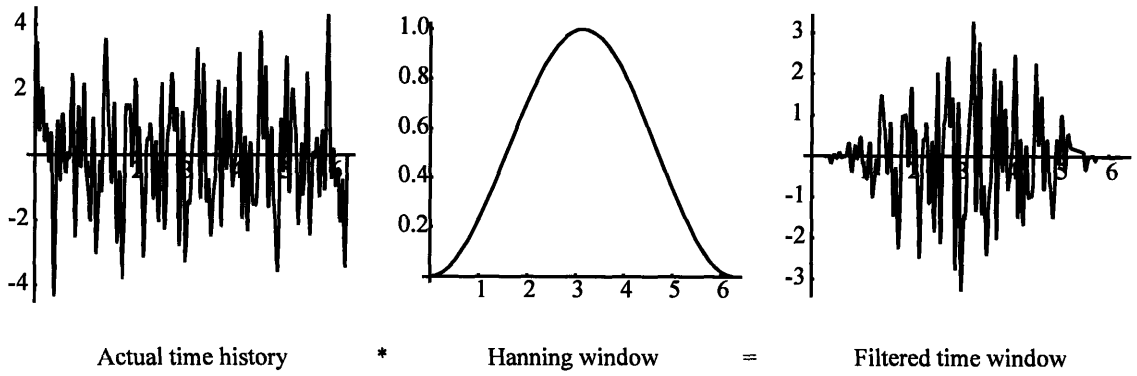

Figure A.19 Sample filtered time series with a Hanning window to remove non-periodic content.

Another way to avoid leakage is to use an excitation waveform that results in zero response at the beginning and end of the sample period. This ensures that the Fourier transform of the data will be valid and that no artificial damping will be included. In general, the burst or impact-type excitation waveforms minimize leakage because the response dies out by the end of the sample period. Figure A.20 shows a table of several common forcing functions.

	Steady sine	Swept sine	Burst sine (chirp)	True random	Periodic random	Burst random	Impact
Leakage	poor	poor	good	poor	good	good	good
Signal to noise ratio	good	good	good	fair	fair	fair	poor
Characterizes nonlinearity	yes	yes	yes	no	no	no	no

Figure A.20 A summary of forcing functions.

Burst random is one of the most commonly used excitation functions. It cannot be used to characterize nonlinearities because it is not repeatable, but burst random is a decent general purpose means of exciting structures over a wide range of frequencies. The burst sine waveform is useful when nonlinearities are encountered because it can be repeated with the exact same pattern with varying amplitude. A nonlinear structure will respond differently to an input of varying amplitude. For example, a hardening spring will have a higher natural frequency at higher force amplitude, and a softening spring will have reduced natural frequencies at higher force amplitudes. These trends can only be located with a repeatable waveform like burst or swept sine.

The selection of transducers used in a modal test has been greatly simplified in recent years by the proliferation of piezoelectric transducers with onboard pre-amplification. Many modal surveys of mechanical systems may be carried out with these

transducers. The appropriate sensor for any given modal test is chosen by considering the size of the test article, the magnitude of the response, and the frequency range of interest.

Accelerometers use a seismic mass and a piezoelectric crystal stack to generate a voltage proportional to the acceleration of the mass. The larger the seismic mass, the higher the resolution of the sensor, but the total weight of the accelerometer can result in undesirable mass loading of the structure. High mass accelerometers will interact with the structure and distort the measurements (particularly at high frequencies). The tradeoff in accelerometer selection is therefore between resolution and mass loading effects. In many cases, structures requiring the highest resolution and signal output are often relatively heavy so larger accelerometers are safe to use.

Force transducers must also be selected with the same tradeoff between output resolution and mass loading. The effects of mass loading may be reduced by moving the force input location to one with a higher effective mass (i.e., closer to a node).

A.4.2 Overview of Data Reduction

Figure A.21 shows a typical test article supported by compliant supports. An impulse hammer is shown at the point where energy will be input into the system. A force transducer will measure the input by sending a voltage to an analog anti-aliasing filter. The accelerometer mounted on the beam will send a voltage proportional to the acceleration of the beam at its mounting point.

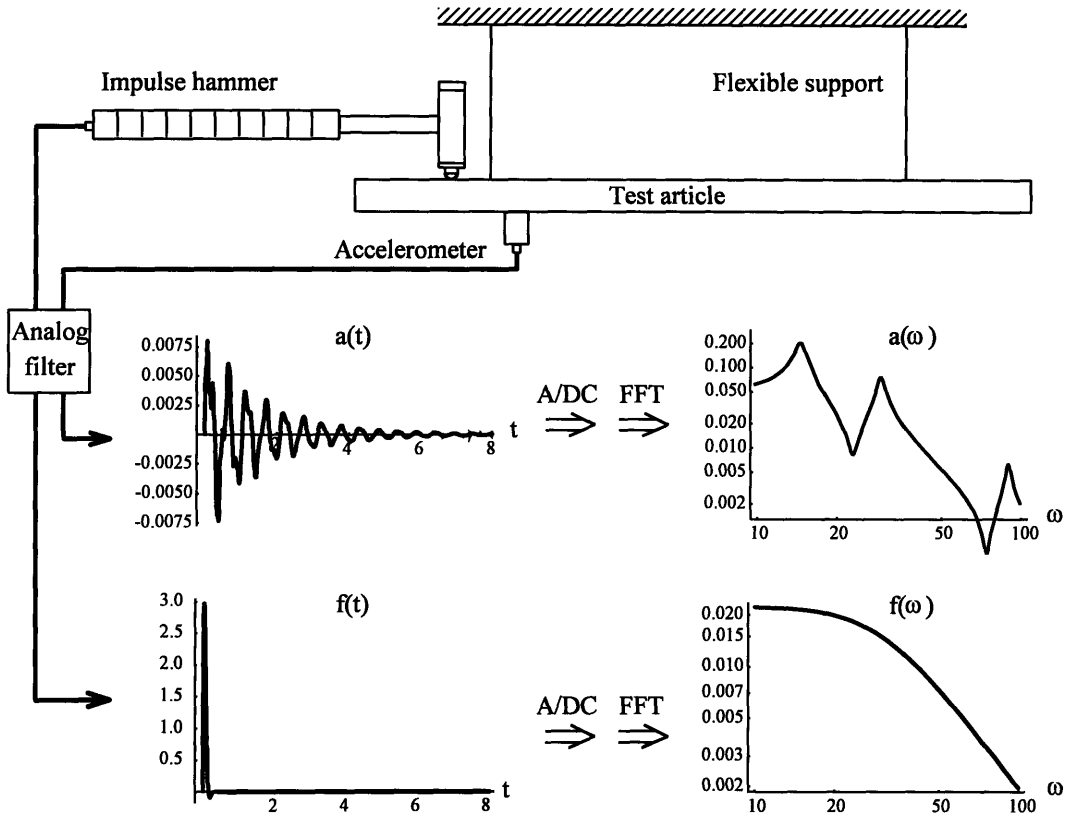

Figure A.21 Flow of data from sensors, to anti-aliasing filters, analog to digital converters, and finally a discrete Fourier transform calculator.

Once the spectra of the input and output sensors are found, the frequency response function can be calculated. Figure A.22 shows the calculated response function from the data collected in Figure A.21.

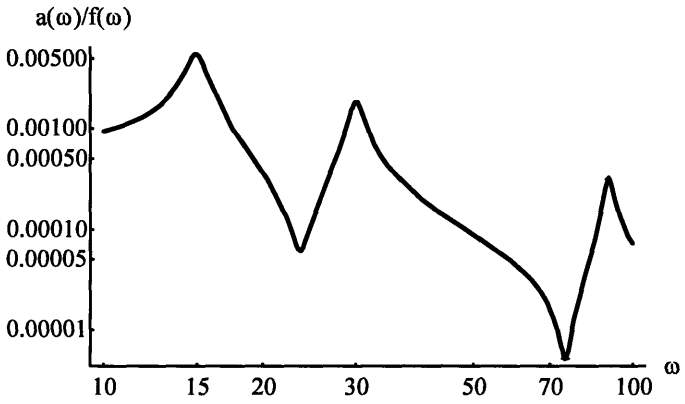

Figure A.22 Calculated frequency response function.

Other algorithms are often used to calculate the frequency response function. The most common frequency response calculation is $H_1(\omega)$. By using the cross power

spectrum and the auto spectrum of the input, the effect of noise in the output is minimized. $H_2(\omega)$ can be used instead of $H_1(\omega)$ if the input noise is considered greater than the output noise.

$$H_1(\omega) = \frac{S_{in-out}}{S_{in-in}} = \frac{\text{cross power spectrum}}{\text{input auto power spectrum}} \tag{A.17}$$

$$H_2(\omega) = \frac{S_{out-out}}{S_{in-out}} = \frac{\text{output auto power spectrum}}{\text{cross power spectrum}} \tag{A.18}$$

A current research topic in experimental modal analysis is the investigation of other methods of computing the most accurate approximation of the correct frequency response function.

A.4.3 Data Collection

The measurement process can begin once a test article has been properly supported and the correct instrumentation has been assembled. The first step is to locate a point on the structure where the vibration amplitude is expected to be large for the modes of interest. For example, if a sensor were to be located at the midpoint of a simply supported beam, all even numbered modes of vibration would be completely unmeasured. Figure A.23 illustrates this point. If the first four modes were the most important, the right hand location would be an excellent location for a sensor.

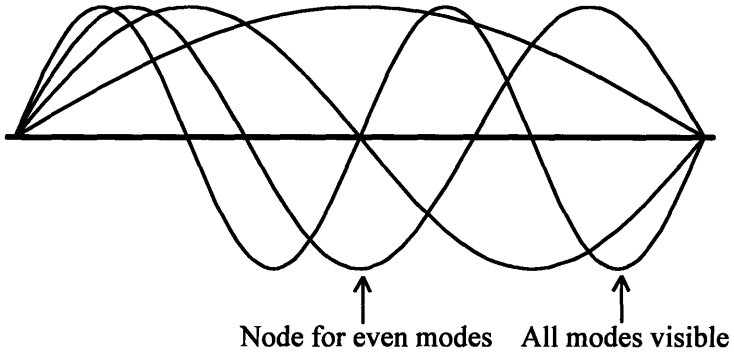

Figure A.23 The mode shapes of a simply supported beam.

Now that the drive point has been located on the structure showing all the modes of interest, a grid can be set out marking the locations on the structure where output data will be collected. This grid will include the drive point measurement, as well as many cross point measurements. In practice, the number of cross point measurements required to adequately describe the vibration of a complex structure can get quite large; complicated structures might easily require 100 or more points. In the case of the simply supported beam, only a handful are needed. This is because the vibration is essentially one dimensional in the frequency range of interest (the longitudinal vibration occurs at

much higher frequencies). Therefore, data need only be collected in one direction. If the first four modes of vibration are of interest, then 10 or 12 points will be sufficient to avoid spatial aliasing. Figure A.24 shows a possible test point location grid for the simply supported beam.

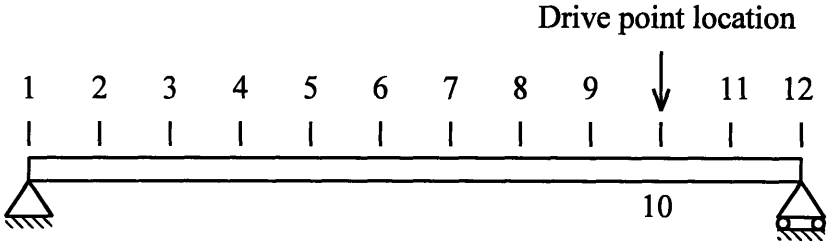

Figure A.24 Measurement locations on a simply supported beam.

With the test location grid mapped out and the correct instrumentation in place, a few sample drive point measurements can be taken with the signal analyzer. There are four different quantities that must be checked to ensure that the collected data are valid. The first is the frequency response function. The FRF will show peaks for each of the modes within the frequency range of interest. Note: the frequency range is a setup parameter for the signal analyzer. The analyzer must be configured to look at a certain frequency bandwidth. The process of locating the correct frequency bandwidth is often a matter of trial and error on a new test article.

Once the correct frequency bandwidth has been established, each of the modes of interest should appear as distinct peaks. If one or more modes is barely recognizable in the drive point frequency response, than a different drive point location may need to be selected. On complicated structures, the location of a good drive point location may take some investigation.

The coherence function is the next indication of the quality of the test results. If the coherence plot is shown alongside the frequency response plot, then the peaks in the frequency response should coincide with a coherence very close to unity. The coherence function is defined for a single input-single output system by (the S's indicate auto and cross power spectra of the input x and output y):

$$\eta_{yx}^2 = \frac{\left|S_{yx}(\omega)S_{xy}(\omega)\right|}{\left|S_{xx}(\omega)S_{yy}(\omega)\right|} \tag{A.19}$$

The coherence function should be as close to unity as possible. This indicates that the system output is fully correlated to the input. A coherence less than unity indicates the presence of noise in the system. In practice, the coherence should be greater than 0.85 for a measurement to be considered usable. In many cases, the coherence can be consistently 0.99 or better, indicating that the data are probably very good. Figure A.25 shows an acceptable frequency response measurement and corresponding coherence.

Note that there are two distinct drop-offs in the coherence. The first is a single spike, the second is in the area of a node.

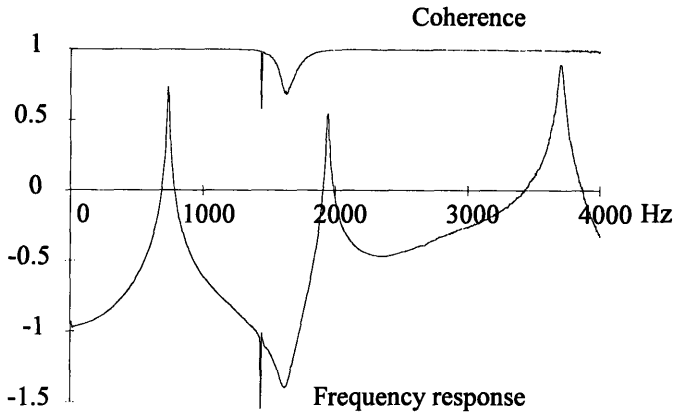

Figure A.25 Sample frequency response function and coherence.

Low coherence can indicate several problems in the data collection setup. If the coherence is low at modal peaks, then leakage is probably effecting the measurement. Low coherence at nodes (locations of zero output for any amount of input) is acceptable because there is no response to measure (resulting in a low signal to noise ratio). When using a shaker as the excitation force, it can be difficult to input power to the structure at low frequencies resulting in low coherence at low frequencies. For example, if a shaker is being used to excite a structure from 0 to 400 Hz, the frequency range of 0 to 20 Hz will typically show poor coherence because the shaker is physically unable to inject as much energy at these lower frequencies. The only way to capture vibration modes at very low frequencies is to use a reduced bandwidth. Piezoelectric sensors tend to perform poorly at very low frequencies because their output is proportional to the shear *rate* in the crystals. If an accelerometer consistently shows poor coherence in a low frequency range with important vibration modes, then an accelerometer with a larger seismic mass may be required to provide better resolution.

During the initial setup of a modal test, the time histories from the input and output transducer should be monitored to confirm that the signal analyzer is triggered at the correct time and that no spurious noise is entering the system. Assuming that the triggering is correct, the most likely problem that may arise, especially in impact hammer testing, is that noise may be compromising the quality of the data. Figure A.26 shows an impact hammer sample with a noisy accelerometer time history.

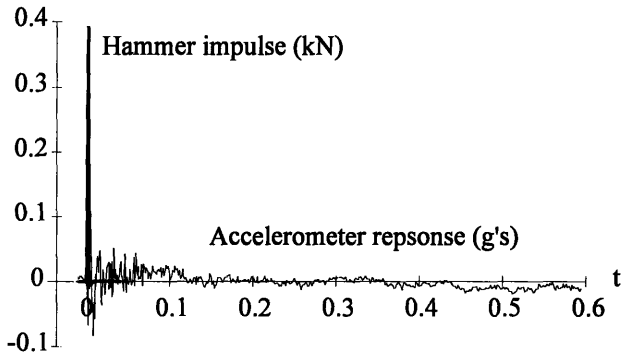

Figure A.26 Impact hammer and accelerometer data showing excessive noise.

Time windowing was discussed as a technique that may be employed to reduce the noise problem shown in Figure A.26. In fact, time windows may be used to reduce noise and filter the data to remove any content near the beginning and end of the time sample period. The pitfall of using windows is that they introduce artificial damping to the measurement. Any filtering process performed in the time domain will *smooth* the Fourier transformed data in the frequency domain so the filtered data will show more damping than is actually present. Figure A.27 shows a number of common filters and their frequency spectra.

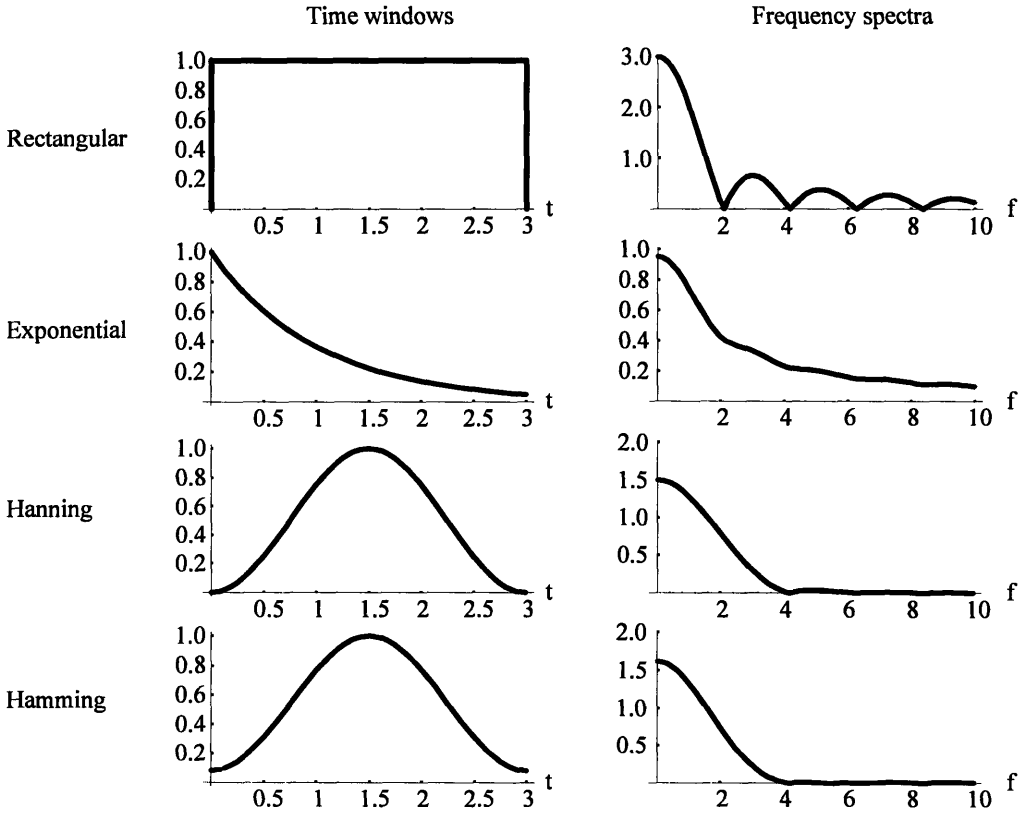

Figure A.27 Common time domain filters and their spectra used to smooth data.

Averaging is commonly used to improve the quality of frequency response function measurements. A typical modal test may use data averaged 10 times at each test location. This is one of the advantages of using a shaker for modal testing because it is easy to set up the signal analyzer to use as many averages as desired. In general, the more averages the better, but diminishing returns are realized when excessive averages are taken. From random vibration theory, the ratio of the standard deviation of a measurement to the mean is given by:

$$\frac{\sigma}{m} = \frac{1}{\sqrt{N}} \frac{1}{\sqrt{BT}} \tag{A.20}$$

This suggests that the number of averages N has little effect on the quality of data after a certain point. Note that the product BT (bandwidth × sample time) is a constant for a given analyzer and cannot be used to improve the quality of the data.

Using the techniques described in this section, a modal survey can be carried out on a variety of test articles. Data will be collected at each of the test locations on the structure. On the simply supported beam example, there were 12 points evenly distributed over the beam. For each of these locations, a number of averaged frequency response functions will be computed and saved on a computer (or printed out for manual computations). These FRF's are then curve fit to estimate the modal parameters of natural frequency, damping factor, and mode shape.

A.5 Case Study: Semiconductor Cassette Handling Robot

A modal analysis of a wafer cassette handling robot was conducted to demonstrate experimental modal analysis. The results of the modal test characterized the robot's dynamic performance and uncovered an important fault in the system. The fault is a low frequency (6 Hz), poorly damped rolling mode that originates from the cleanroom floor/robot structure interface. This rolling mode is a result of compliance in the floor, a weakness in the design of the clean room which profoundly effects the performance of the robot and other floor-mounted equipment.

A.5.1 Modal Test Setup

The survey began with some preliminary measurements being taken to optimize the instrumentation setup and data filtering parameters. The location of the drive point measurement was also selected using the pretest measurements. The drive point location on the test structure was chosen that clearly showed all the critical modes. The tip of the end effector was a natural choice because the deflection of the robot is largest at this point. Although the largest deflections of the first mode occurred in the axial direction, several modes were more visible when the accelerometer was placed in the vertical orientation.

Once the location of the drive point measurement was established, the locations of the other test points were chosen. The test locations were placed at the eight corners of the box-shaped welded steel structure, at four locations on the horizontal track, and several more locations on the robot arm. Figure A.28 shows the test structure and the location of the vibration measurement points. Note that the modal tests were run with the arm fully extended. This was done to simulate a worst case (least stiff) orientation of the robot arm.

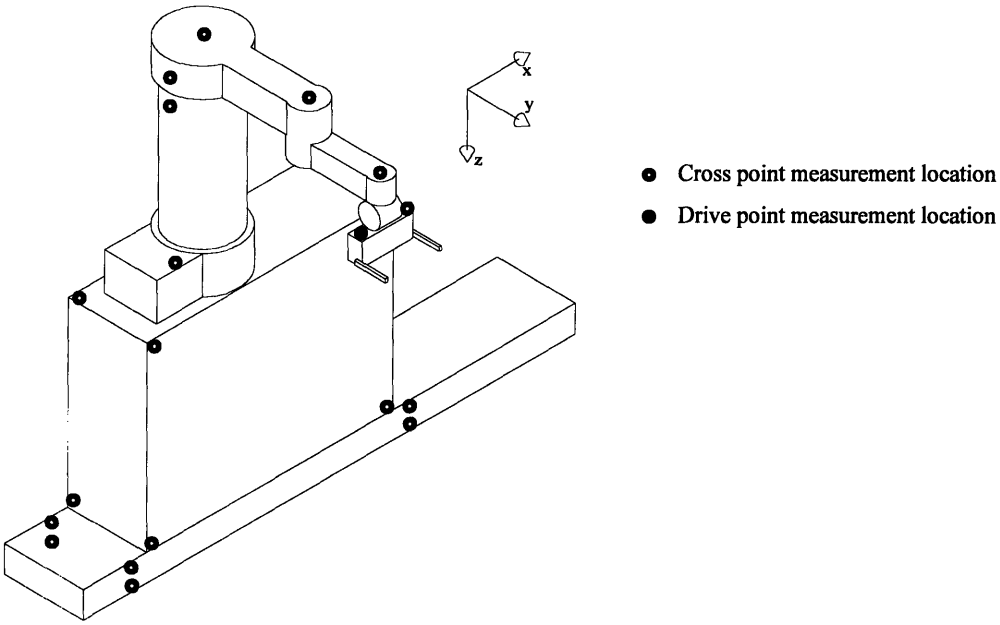

Figure A.28 Robot structure and measurement points.

During the pretest measurements, several checks were made to verify that valid data were being collected by the signal analyzer. The coherence of the drive point frequency response function was monitored, along with a few cross point measurements to make sure that ambient noise in the test site would not effect the results. The coherence was very good (approximately equal to unity) at all test locations. There was a slight drop off in quality of the measurements taken at the base of the track due to the very low impedance at these locations; however, this is a common occurrence at node-like locations.

Significant background noise was present during vibration testing. Vibration from people and machinery in the clean room had a noticeable effect on the measured data. The solution was to minimize movement during testing and to monitor the coherence function to reject bad samples from the averaged measurements.

The drive point measurement is shown in Figure A.29. The first six modal peaks have been identified as important. These are the modes of interest to the accuracy of the robot. Figure A.30 shows the same frequency response function, this time plotted as a

function of displacement per unit force input. This plot was obtained by double integrating the accelerance FRF.

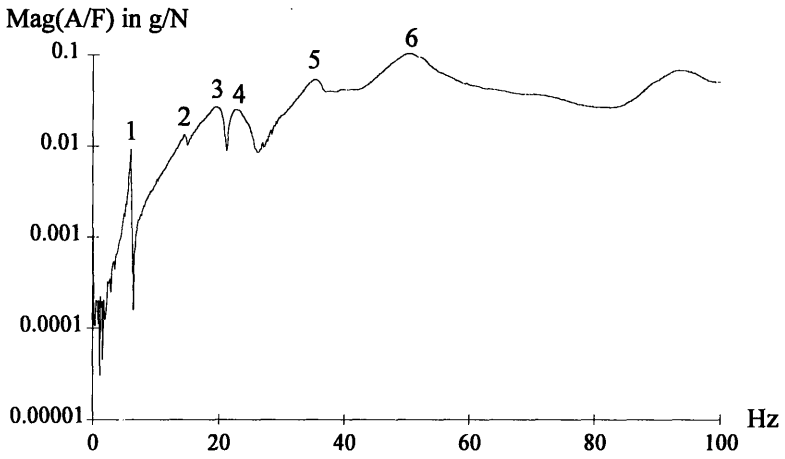

Figure A.29 Drive point accelerance of the robot arm.

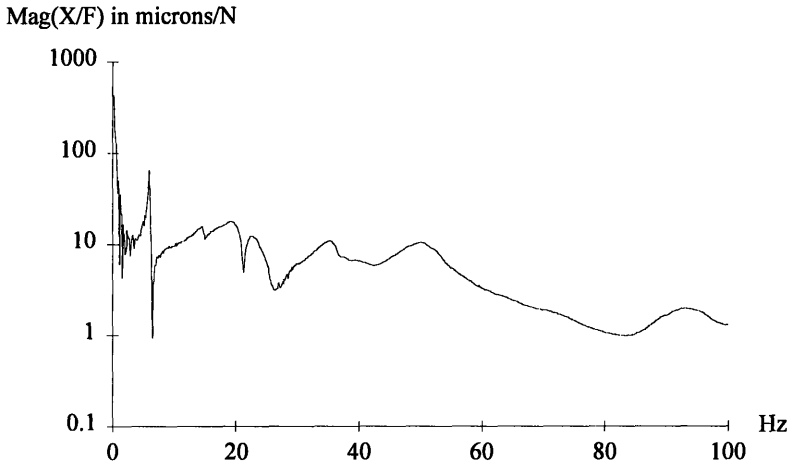

Figure A.30 Drive point compliance of the robot arm.

The 6 Hz rolling mode has a much larger amplitude than any of the structural modes of the robot. This speaks highly for the stiffness of the robot structure. Furthermore, the modes of the robot structure have well rounded peaks indicating a high degree of damping. The modal analysis results will show that most of the modes appear to originate from compliance of recirculating roller bearings which offer moderate damping.

The following sections detail the instrumentation and equipment setup during the modal testing.

A.5.2.1 Time Windowing and Averaging

Time windowing was not used in this modal survey. The frequency range of interest was narrow (100 Hz) so the resolution of the critical modes is high. Because the sampling

time for a digital Fourier transform is inversely proportional to the frequency bandwidth, the sample time was long (8 seconds). The free vibration decay time of the hammer impulse was much less than 8 seconds so the traditional exponential windowing used in impact testing was not required.

Measurement averaging was used in the modal survey to reduce the effects of imprecise hammer impacts and other variables. Although power spectrum quantity averaging does not improve the signal to noise ratio, it is the best averaging scheme when precise triggering is not available. Eight averages were taken at each data point location. In some cases, the frequency response functions from two complete sets of eight averages were compared at a single test measurement location to ensure that the results were repeatable.

A.5.2.2 Test Equipment and Configuration

The force excitation was provided with a three pound impact hammer and a soft impact tip. The soft tip was chosen to lengthen the duration of the impact, thus improving the time resolution of the force measurement. The force spectrum was checked to make sure that the soft tip had sufficient high frequency content to excite the vibration modes of interest.

An accelerometer with a relatively large seismic mass was selected to measure the response of the robot. A large seismic mass was desirable in this test case because the first vibration mode occurred at only 6 Hz. A smaller accelerometer, while capable of measuring very high frequencies, would not be able to accurately resolve the accelerations at 6 Hz. The large accelerometer mass, which is undesirable for testing lightweight structures, was not a problem for the heavy robot structure.

The piezoelectric sensors used in the accelerometer and the force transducer are powered by battery operated, constant current signal conditioners. The output voltage of the conditioners is measured by the signal analyzer and discretely sampled into digital data. Battery powered signal conditioning is preferred over AC power because of the reduced noise introduced into the measurement.

A Hewlett Packard 3562A signal analyzer was used to collect, anti-alias filter, Fourier transform, and average the raw transducer time histories. The HP analyzer provided visual displays of the raw and transformed data so that the quality could be checked as the modal survey took place. Figure A.31 summarizes the setup of the equipment and the analyzer configuration.

Frequency range	100 Hz
Sample time	8 seconds
Pre-triggering	0.25 seconds
Excitation	Roving 3 lb impulse hammer
Response	Low frequency accelerometer - fixed on end effector
Excitation window	Uniform
Response window	Uniform
Number of averages	8

Figure A.31 Test equipment and analyzer configuration.

A.5.2 Modal results

The results obtained in the modal survey confirm the existence of the six modes that are visible in the drive point frequency response measurement (Figure A.29). SMS Star Modal™ software was used to reduce the collection of frequency response functions to a set of vibration modes with calculated natural frequencies and damping factors. Star Modal helped identify the dominant modes of the drive point measurement and then curve fit the appropriate sections of each response function to calculate the natural frequency, damping factor, and amplitude of each mode. The results are summarized in Figure A.32.

Mode	Nat. Freq. (Hz)	Damping (%)
1	6.00	1.36
2	15.6	2.33
3	20.5	6.14
4	22.7	5.23
5	35.9	4.31
6	50.7	6.75

Figure A.32 Summary of the first six modes.

The validity of the six modes identified during the data analysis can be checked by the modal assurance criterion (MAC). The MAC is a square matrix of dimension equal to that of the number of curve fit modes. The MAC matrix shows the orthogonality of the identified experimental mode shapes. Ideally, all modes should be mutually orthogonal to each other so the off diagonal terms should ideally be 0's. The main diagonal of the MAC matrix should be 1's (because each mode coincides with itself). Figure A.33 shows the MAC matrix for the identified modes. Overall the modes show good orthogonality except for the first and second modes. This resulted from the small magnitude of the

second mode acceleration which made it difficult to separate out from the experimental noise floor.

	Mode 1	Mode 2	Mode 3	Mode 4	Mode 5	Mode 6
Mode 1	1.00	0.57	0.05	0.17	0.01	0.01
Mode 2	0.57	1.00	0.07	0.08	0.00	0.00
Mode 3	0.05	0.07	1.00	0.23	0.02	0.02
Mode 4	0.17	0.08	0.23	1.00	0.11	0.21
Mode 5	0.01	0.00	0.02	0.11	1.00	0.44
Mode 6	0.01	0.00	0.02	0.21	0.44	1.00

Figure A.33 The modal assurance criterion matrix for the identified modes.

The following sections summarize the modal survey test results and show how each mode is manifested in the vibration of the robot structure. Figure A.34 shows a schematic of the cassette robot as well as a wireframe of the structure. The wireframe representation is used in the following sections to illustrate each of the six important mode shapes.

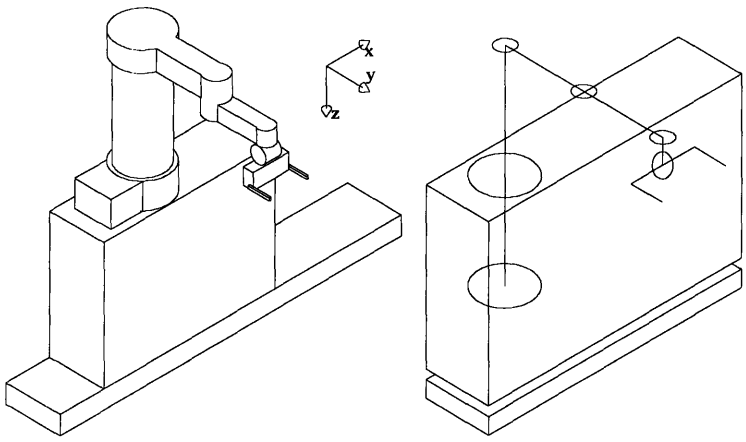

Figure A.34 Schematic and wireframe representation of the robot structure.

A.5.3.1 Mode 1

Figure A.35 shows the entire robot and track pivoting about the track/ground connection. The rolling motion of the entire structure is clearly undesirable because of its effects on the positioning accuracy of the robot end effector. However, the mode is not due to compliance in the robot structure, but rather a compliance in the floor.

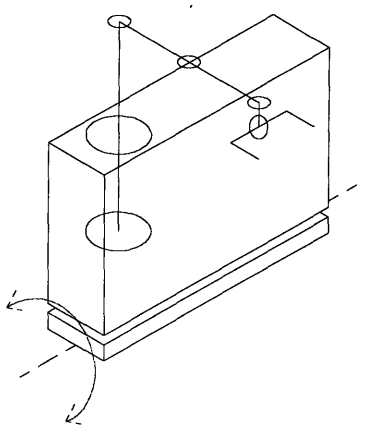

Figure A.35 First mode shape.

The grate-like tiles in the clean room floor were supported by three foot columns spaced two feet apart. This support system has the advantage of allowing easy access below the clean room for machine installation, maintenance, and cleaning. The disadvantage is that this floor design is very compliant. An operator walking around the robot can deform the floor enough to effect the positioning accuracy of the robot by several thousandths of an inch.

The 6 Hz mode, which has the highest peak shown in Figure A.29, may be eliminated by a more robust floor design.

A.5.3.2 Mode 2

The second mode shown in the drive point measurement occurs at 15.6 Hz with modest damping (ζ = 0.02). This mode has a very small amplitude, as shown in Figure A.29. Although this mode is difficult to resolve in the experimental data, the modal analysis software was successfully used to animate the wire frame image of the robot structure. The mode shows some rolling movement in the welded steel frame that supports the robot arm, but most of the motion occurs in rolling of the robot arm and support cylinder. This motion is shown in Figure A.36.

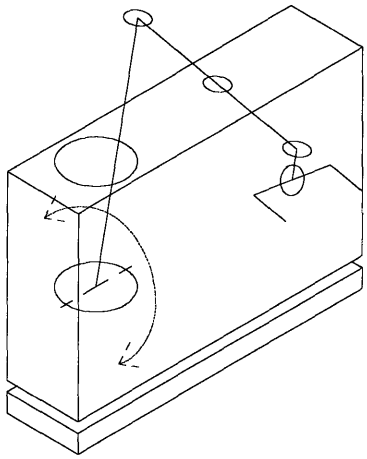

Figure A.36 Second mode shape.

This small amplitude mode is a result of the large modal mass that the cantilevered robot arm effectively places on the bearings supporting the vertical degree of freedom.

A.5.3.3 Modes 3 and 4

The third and fourth modes of the robot structure vibration show similar behavior. Both modes occur at very close frequencies (21 and 22 Hz), and show moderate damping ($\zeta = 0.06$ and 0.05). The effective source of compliance, the first axis bearing, is shown in Figure A.37.

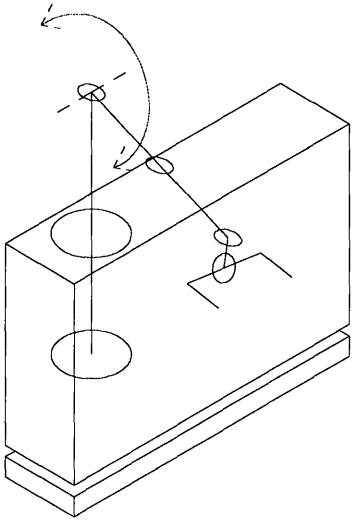

Figure A.37 Compliance of robot arm bending in third and fourth mode shape.

The source of compliance appears to be the bearings at the first rotational degree of freedom. It should be mentioned that the frequencies of these modes are fairly high, indicating that the robot is actually a stiff structure. The bearings will always be more compliant than the arm links because the links themselves have a massive cross section.

A.5.3.4 Mode 5

The fifth mode identified in the modal survey occurred at 36 Hz with moderate damping ($\zeta = 0.04$). This mode, shown in Figure A.38, is clearly a rocking of the robot arm on the linear guides supporting the vertical axis. The rocking occurs about an axis parallel to the orientation of the arm.

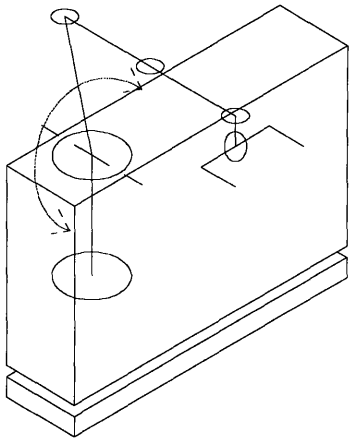

Figure A.38 Fifth mode shape.

Like modes three and four, this mode is reasonable well damped and occurs at a high enough frequency that the performance of the robot is not compromised. Additional stiffness in the vertical axis bearings would increase the natural frequency of this mode, further reducing its impact on the robot performance.

A.5.3.5 Mode 6

The sixth mode occurred at around 50 Hz with fairly high damping ($\zeta = 0.07$). The animated mode shapes output by the modal analysis software indicated that this mode resulted from compliance in the linear guides supporting the welded steel frame. This compliance is judged to be fairly small considering the weight of the robot.

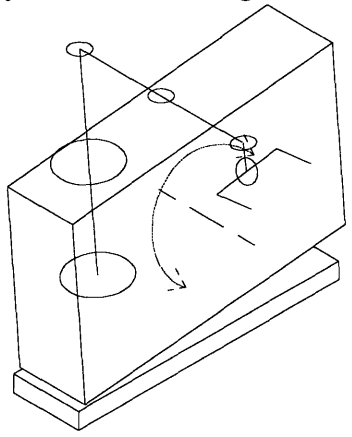

Figure A.39 Sixth mode shape.

The magnitude of the sixth mode of vibration could be further reduced by using a stiffer bearing design. This design change is not warranted given the performance of the robot in its current application.

A.6 Conclusion

Modal analysis is an excellent way to experimentally assess the dynamic performance of a mechanical system. The calculated modal matrices (frequency, damping, and mode shape) provide important information that the designer can use to efficiently identify inadequacies in a design and upgrade the weak components.

The system identification problem that must be solved when performing a modal survey can become difficult when closely spaced or highly damped modes occur in a structure. The quadrature peak picking method shown in this appendix cannot be used in these cases. The literature describes a vast number of more sophisticated techniques, many of which are implemented on commercially available modal analysis software. The analyst must remember that the identification techniques are highly sensitive to any distortion or noise in the data set. The experimentalist must therefore be prepared to invest the effort and time necessary to obtain good data.

Appendix B: Analytic Hierarchy Process

B.1 Introduction

Several tools have become available in recent years to address the need for an efficient method of stepping through complex decision making problems. These tools include the Analytic Hierarchy Process, house of quality, the Pugh Process, and the Kepner-Tregoe method [Saaty, 1980; Hauser and Clausing, 1988; Pugh, 1981; Kepner and Tregoe, 1981]. The Pugh Process and the Kepner-Tregoe method offer good results for a modest investment of time and effort. The AHP offers higher accuracy at the expense of added effort. The AHP also provides a more comprehensive framework for group decision making.

The Pugh Process and the Kepner-Tregoe method may be considered subsets of the AHP; the AHP will give the same results as the Pugh Process if the evaluation criteria are all assigned equal weight. The Kepner-Tregoe method offers a less rigorous (but quicker) approach to assigning a weighting vector to the evaluation criteria. Although the AHP generates more accurate results than Kepner-Tregoe or Pugh, these two methods allow quick, first order estimates (provided the user understands these limitations). The spreadsheet template developed in this appendix will make the AHP a more suitable alternative to Kepner-Tregoe and Pugh by providing a pre-formatted framework that is easily completed to resolve decision making problems.

Saaty developed the AHP to systematically tackle complex decision making problems in a wide variety of applications. This appendix presents an adaptation of the AHP to help engineering design teams select the most appropriate concept from many possible choices and identify desirable features that should be incorporated into the final design.

To achieve this result, the evaluation criteria are assembled into a hierarchical framework. This hierarchy represents the natural decomposition of the decision making problem from high level, broad issues to lower level, detail issues. The relative importance of the criteria is considered one level at a time. The decision solutions are then rated on their ability to satisfy the criteria independently of their relative importance. The interrelationships of the criteria are effectively decoupled by considering the weighting and the solution evaluations separately. An example of this decoupling is evident in the ubiquitous cost criterion. In one respect, cost is the deciding factor in almost every decision. However, the AHP can be used to break down the total cost into discernible components such as manufacturing cost, maintenance cost, and the cost associated with producing a design with compromised quality (such a product may be less expensive now, but not in the long run). This separation of criteria is analogous to the decoupling sought by axiomatic design in that each criterion's effect on the decision

outcome may now be considered individually [Suh, 1985; Steward, 1981]. This technique of considering the criteria and design concepts separately allows both novice and expert to systematically handle complex decision making problems.

The AHP is particularly useful in fields such as design concept selection because a wide range of goals must be satisfied by a single concept. The proposed adaptation is suited to engineering design because:

- Evaluation criteria are logically and mathematically decoupled.
- The framework focuses discussions on key design issues.
- The best decision is identified and approved upon completion of the exercise.

The AHP guides the design team to a consensus because of the structured decision making framework. The inherent mathematical and logical rigor of the AHP framework encourages efficient group discussion and reduces the opportunity for one team member to dictate the decision.

This paper will present the AHP, its adaptation to engineering design decision making, and a spreadsheet template that packages the process into a user friendly tool. An introduction to other decision making tools is presented first to provide a thorough background.

B.2 House of Quality

The house of quality is a design tool that may be used to ensure that customer requirements will be addressed by a product design [Hauser and Clausing, 1988]. The house of quality provides a framework for locating and recording interrelationships between all the customer requirements and mapping them onto engineering characteristics, customer perceptions, and evaluation of competitive products. The result of a house of quality study is target values for each customer requirement. The relative importance of the customer requirements may then be assessed to determine which target values must be met so that the product will have the best chance of success. Other targets relating to less important customer requirements may be ignored.

Upon completion of the study, the entire design team must use the house of quality results to begin making design decisions; decision making is not handled by the process itself. A method such as the Pugh Process is used so that the questions raised during the house of quality study can be resolved.

B.3 Pugh Concept Selection Process

The Pugh Process and the AHP share several similar features, in fact, the Pugh Process is a specific case of the more general AHP [Pugh, 1981]. Both methods ask the user to itemize the key requirements (or evaluation criteria). The different concepts are then evaluated on their ability to satisfy each of these requirements.

The difference is that the Pugh Process makes no distinction between the relative importance of evaluation criteria. For example, product cost is always rated equally important to other criteria such as ergonomics. This is not an oversight of the inventors; the assumption is that the design team will be able to efficiently settle on a concept that satisfies most or all of the criteria.

The Pugh Process uses one pre-selected concept as a datum and applies a plus, same, or minus rating to each concept for each criterion. The number of pluses and minuses are totaled for each concept. The process is then repeated using the first round winner as the new datum. The most favorable concept from the second round of comparisons is chosen as the final winner. Using this technique, the strength of each of the design concepts will be well defined and in many cases, the design team may be able to use this information to develop an improved concept that incorporates the strengths of many concepts.

B.4 Kepner and Tregoe Decision Making Process

The Kepner and Tregoe process is another popular method of handling difficult decision making problems [Kepner and Tregoe, 1981]. The mathematical rigor of the method exceeds that of the Pugh Process, but is not as involved as the AHP.

The Kepner and Tregoe method proceeds like the others in that a list of decision concerns is developed along with a list of possible solutions. This list of concerns is then broken down into "musts" and "wants." The musts are listed first and any solution that fails to satisfy each must is dropped from consideration. The wants are assigned a weight on a scale from 1 to 10 (10 being the most important).

The solutions are then rated on their ability to satisfy the wants of the decision makers. Like the AHP, the weight assigned to the wants is multiplied by the evaluation of the solutions to satisfy each want. In this fashion, a total score is assigned to each solution. Close calls are resolved by other means of comparison of the nearest choices.

The Kepner and Tregoe method will yield a fairly accurate solution with a reasonable effort. The method is implementable on a spreadsheet for the purposes of carrying out sensitivity analyses.

B.5 The Analytic Hierarchy Process

The foundation of the AHP is a hierarchy of matrices that record the decision making criteria and their relative importance. By breaking down the complex problem into a hierarchy, the fundamental relationships between criteria can be examined without having

to make broad judgments about the overall system. These matrices may be recorded by pencil and paper in the simplest analyses, but are more easily manipulated using spreadsheet programs. This ensures that the decision making process is as dynamic as the design process itself.

The Analytic Hierarchy Process involves three distinct steps:

- Step One (setup): Decision making criteria are determined, often by brainstorming or past experience. Hierarchical relationships are drawn between the criteria and are then represented in matrix form.

- Step Two (weighting): The matrices are filled with criteria comparisons. The comparisons allow calculation of the criteria weighting vector.

- Step Three (evaluation): Different problem solutions are evaluated on their ability to satisfy the various criteria. The final solution ratings are then calculated using the ratings determined in this step and the weighting vector calculated in Step Two.

B.5.1 AHP: Step One

The first task of Step One is to decide on the problem statement. This statement becomes *Level One* of the hierarchy and will be broken down into a short list of key concerns. These become the *Level Two* criteria and may include broad issues such as cost, reliability, ergonomics, etc. Additional levels are used to further decompose the *Level Two* criteria. For example, cost may be broken down into the *Level Three* criteria: labor, overhead, and variable cost.

When beginning a new problem, the user develops criteria by brainstorming different design concerns that must be satisfied by the solution. Past experience is extremely helpful because old criteria are readily modified for a new problem. The criteria are then organized into the various levels of the hierarchy. *While the criteria on a particular level are not necessarily equally important, they must be on the same order of magnitude.* Figure B.1 shows an example of a typical hierarchy.

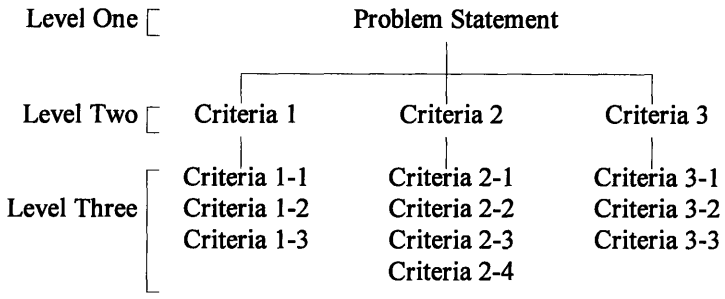

Figure B.1 Example of a criteria hierarchy.

Once the hierarchy has been completed, matrices are constructed with the criteria labels on each axis. For example, the *Level Two* matrix may have the form shown in Figure B.2.

Level Two	Cost	Ergonomics	Reliability
Cost			
Ergonomics			
Reliability			

Figure B.2 Example of a *Level Two* matrix with three common design criteria.

The levels below each of the *Level Two* criteria (such as cost) are also represented in matrix form. Figure B.3 shows an example of a possible cost matrix.

Cost	Labor	Overhead	Variable
Labor			
Overhead			
Variable			

Figure B.3 Example of a *Level Three* matrix with cost sub-criteria.

Once the hierarchy of the criteria is determined and all of the necessary matrices have been constructed, the user may proceed to Step Two.

B.5.2 AHP: Step Two

The matrices shown in the preceding section must now be filled with comparisons of the relative importance of the criteria on the two axes. The comparisons are used to calculate the weighting vector that gives the relative importance of all the criteria. The weighting vector is calculated from comparisons made between two criteria at a time. This is how the AHP breaks complex decision making problems into manageable parts. Figure B.4 shows the scale (1 to 9) proposed by Saaty for indicating the relative importance between criteria [Saaty, 1980].

1: Both criteria of equal importance	
3: Left hand weakly more important than top	1/3: Top weakly more important than left hand
5: Left hand moderately more important than top	1/5: Top moderately more important than left hand
7: Left hand strongly more important than top	1/7: Top strongly more important than left hand
9: Left hand absolutely more important than top	1/9: Top absolutely more important than left hand

Figure B.4 Comparison scale used to complete the weighting matrices.

A criterion along the left hand vertical axis is always compared with respect to a criterion along the top horizontal axis. For example, if a criterion on the left hand is more important than another criterion on the top axis, then a number between 1 and 9 is used to record this relationship. If the top axis criterion is more important then the reciprocal is used.

In Figure B.5, cost is of equal importance with respect to itself so a 1 is put in the upper left hand corner. Cost is decided to be weakly more important than ergonomics, so a 3 is recorded in the matrix as shown.

Level Two	Cost	Ergonomics	Reliability
Cost	**1.000**	**3.000**	
Ergonomics			
Reliability			

Figure B.5 Example of the first comparison made in a weighting matrix.

In Figure B.6, reliability is considered weakly more important than cost. This means that a number less than unity must be used. In this case, the number is the reciprocal of 3, or one-third.

Level Two	Cost	Ergonomics	Reliability
Cost	**1.000**	**3.000**	**0.333**
Ergonomics			
Reliability			

Figure B.6 Example of the second comparison made in a weighting matrix.

A consistent matrix formulation allows the remainder of the matrix to be completed given the information in the top row. Since the relationship is known between cost and ergonomics and between cost and reliability, the relationship between ergonomics and reliability can be found. In this case, the matrix entry for ergonomics versus reliability would be 1/9. Similarly, the rest of the matrix can be computed using the formula $a_{jk} = a_{1k} / a_{1j}$, as shown in Figure B.7.

Level Two	Cost	Ergonomics	Reliability
Cost	**1.000**	**3.000**	**0.333**
Ergonomics	0.333	1.000	0.111
Reliability	3.000	9.000	1.000

Figure B.7 Completed *Level Two* weighting matrix.

At this point, the reader may wonder why the entire matrix is displayed if only the first row is needed to complete the entire matrix and therefore determine the weighting vector. The other rows provide a built-in visual check of the first row comparisons. For example, if there is a question as to the proper value of the comparison between cost and ergonomics, the other rows may be visually checked to see if these results are reasonable. Saaty suggested that users fill in all the comparisons above the main diagonal (the entries below will always be the reciprocal of those above). Once the entire upper half of the matrix is filled in, the entries below the first row should be close to estimated results by $a_{jk} = a_{1k} / a_{1j}$. This check leads to a rating called the consistency ratio, a measure of how close the manually filled-in upper diagonal is to a consistent matrix.

In the design concept selection adaptation, we insist on consistency. This simplification yields two important results. First, the intuitiveness of the method is assured. For example, if we know that $y = 2x$ and $z = 4x$, then we also know that $z = 2y$. A fundamental tenet of machine design is determinism, which states that a design should be based on an analytical relationship between components [Slocum, 1992; Evans, 1989; Bryan, 1984]. Without a deterministic method of completing the weighting matrices, the matrix entries cannot be visually scanned for accuracy. Furthermore, the mathematics of the method are simplified because the weighting vector (which is taken to be the principal eigenvector of the matrix) is simply the first column of the matrix normalized so that the sum of its entries equals unity.

The weighting vector for our example is shown in Figure B.8. The first column (1.000, 0.333, 3.000) is normalized so that the sum of its entries is 1.0. Therefore, each entry in the weighting vector gives the percentage of the total weight applied to each criterion (this weight will be carried down to subsequent levels in the hierarchy).

Level Two	Cost	Ergonomics	Reliability	Weighting Vector
Cost	**1.000**	**3.000**	**0.333**	***23.1%***
Ergonomics	0.333	1.000	0.111	***7.7%***
Reliability	3.000	9.000	1.000	***69.2%***

Figure B.8 Weighting vector calculation made from criteria comparisons.

This comparison process is repeated for all the matrices to be used in the analysis. The weighting vectors of the lower matrices will be normalized so that their total weight

will equal that of the previous level criterion. For example, the cost sub-criteria (labor, overhead, and variable) will be given a total weight of 23.1%, as shown in Figure B.9.

Cost	Labor	Overhead	Variable	Weight. Vec.	Norm. Vec.
Labor	**1.000**	**3.000**	**5.000**	65.4%	15.1%
Overhead	0.333	1.000	1.667	21.6%	5.0%
Variable	0.200	0.600	1.000	13.1%	3.0%
				100.0%	**23.1%**

Figure B.9 Normalized weighting vector for the *Level Three* cost matrix.

B.5.3 AHP: Step Three

The final step of the AHP is to evaluate how the potential decision solutions satisfy the criteria generated in Step One. The evaluations are made for the most detailed level of each branch of the hierarchy.

The evaluations that are recorded with a user defined numerical scale (e.g., a scale from 1-10). The scale is unimportant; the results are normalized so the evaluations do not weight the results. Continuing with the example, if there are three possible solutions to our problem (A, B, and C), the evaluations would be represented, as shown in Figure B.10 (sample user-supplied judgments are shown in italics):

		Concept Evaluations			Normalized Results		
Cost	Norm. Vec.	A	B	C	A	B	C
Labor	15.1%	*1*	*4*	*6*	0.014	0.055	0.082
Overhead	5.0%	*5*	*7*	*2*	0.018	0.025	0.007
Variable	3.0%	*40*	*60*	*50*	0.008	0.012	0.010
	23.1%				**4.0 %**	**9.1 %**	**10.0 %**

Figure B.10 Weighting vector, concept evaluations, and final results for the cost matrix.

The final rating for the cost criterion is shown in boldface type. The normalized results in Figure B.10 are calculated by the following scheme:

- Task 1: Evaluate problem solutions using any preferred scale (such as 1-10).
- Task 2: Normalize each entry by the sum of each row.
- Task 3: Multiply the normalized rows by the normalized weighting vectors.
- Task 4: Add the columns of normalized scores to obtain the final rating.

Normalization to unity assures that the final scores will equal the percentage of the total weight assigned to the matrix. For example, the cost matrix was assigned 23.1% of

the total decision weight in the *Level Two* matrix; the weight assigned to the three sub-criteria in the *Level Three* cost matrix also equals 23.1%. When the other *Level Three* matrices are calculated (reliability and ergonomics), the total rating score will be 100%.

The result of this series of operations is a weighted rating for each solution. The highest rated solution will best meet the problem criteria.

B.6 The Relation Between the AHP and Axiomatic Design

The design principles identified by Suh state that an acceptable design solution must satisfy two axioms [Suh, 1985]. The first, *The Independence Axiom*, requires that the functional requirements of a design must remain independent of each other. This means that there should exist an uncoupled design matrix that relates the design parameters to the customer requirements. If these two design domains cannot be represented in an uncoupled form, then the designer may not be able to solve the design problem (i.e., as a mathematician would solve a set of linear equations).

The second axiom, *The Information Axiom*, states that the final design concept must have the highest probability of satisfying the customer requirements. These two axioms have been used in developing the proposed adaptation of the AHP. As a result, two design formulation and concept selection corollaries are evident:

> *Corollary One*: A weighting vector may be used to rate the relative importance of the uncoupled selection criteria. These uncoupled criteria result from satisfaction of The Independence Axiom.
>
> *Corollary Two*: The best decision may be made by assessing the rank in which the alternatives satisfy The Information Axiom using the results of the weighting vector of Corollary One.

Given a need for a design decision and understanding of the factors affecting the decision, the design team may develop a list of criteria. If these criteria may be arranged into an uncoupled hierarchy (satisfying The Independence Axiom), Corollary One states that a unique weighting vector may be calculated for the hierarchy given judgments on the relative importance between criteria. The result of Corollary One is that a weighting vector may be obtained for the complete hierarchy without the need for considering all criteria as a whole. Instead, individual comparisons are made that break the difficult problem into smaller decisions.

Corollary Two states that given a set of decision choices and information about them, the choices may be rated on their ability to satisfy the individual criteria. This rating will automatically satisfy The Information Axiom by maximizing the probability of satisfying the criteria. The result of satisfying Corollaries One and Two is the optimal

solution to the decision problem. This solution will have satisfied the fundamental tenets of axiomatic design.

B.7 Example Application of the Analytic Hierarchy Process

This example comes from experience with a machine tool manufacturer interested in designing a horizontal spindle jig boring machine. Ten design concepts were evaluated using the AHP, two of which will be presented here. The first concept features z-axis bearings that move on slideways mounted on the machine base. This configuration is called a *constant stiffness design* because the bearings, which are mounted on the tool head, are always a fixed distance from the tool point. The second concept is called a *maximum stiffness design* because the bearings are mounted on the machine base and the rails are on the tool head. As a result, if the tool head is fully retracted, the tool point acts on a short cantilever beam. If the tool head is extended, the cantilever is much longer. The cantilever length is always the same in the constant stiffness design, as shown in Figure B.11.

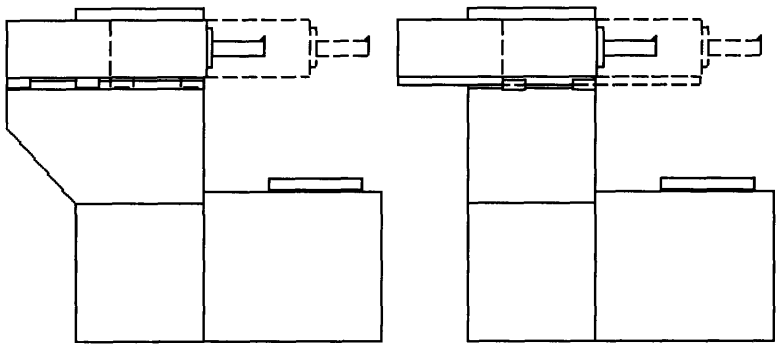

Figure B.11 Constant and maximum stiffness jig borer design concepts.

Level One of the hierarchy, the problem statement, is to select a jig boring machine concept that best satisfies the tenets of precision machine design. The next task, brainstorming the hierarchy of criteria, was performed by an entire design team of engineering, manufacturing, marketing, sales, service, and managerial personnel. The result of these Step One activities is shown in Figure B.12.

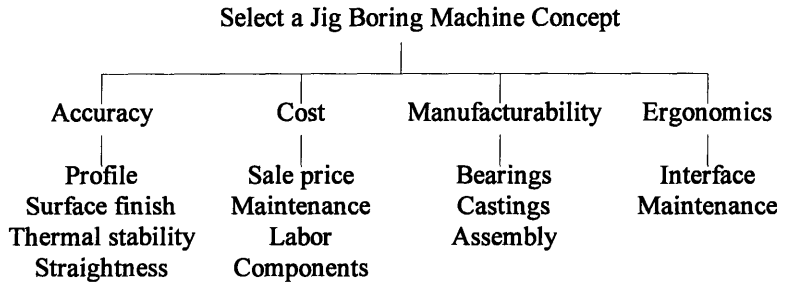

Figure B.12 Jig borer's hierarchy of concept selection criteria.

After the evaluation criteria and hierarchy were established, individual weighting matrices were constructed. Then, the first row of each matrix was filled in using the AHP rating scheme. These Step Two activities were completed by the entire team to ensure that the comparisons given in the matrices benefited from everyone's experience.

The results of the *Level Two* comparisons (accuracy, cost, manufacturability, and ergonomics) are shown in Figure B.13.

Level Two	Accu.	Cost	Manu.	Ergo.	Weighting Vector
Accuracy	**1.00**	**3.00**	**5.00**	**5.00**	***57.7%***
Cost	0.33	1.00	1.67	1.67	***19.2%***
Manufacturability	0.20	0.60	1.00	1.00	***11.5%***
Ergonomics	0.20	0.60	1.00	1.00	***11.5%***

Figure B.13 Jig borer's *Level Two* comparisons and weighting vector.

For this example, the design team felt that the accuracy of the machine was the most important design quality. The cost was also a considerable concern, but not as important as accuracy. Note that the sum of the weighting vector entries is 100%.

The *Level Two* criteria shown above are broken down into smaller parts in four *Level Three* matrices. The first criterion, accuracy, is broken down into *Level Three* criteria as shown in Figure B.14.

Accuracy	Profile	Surf.	Therm.	Straight.	Weighting Vector
Profile	**1.00**	**1.00**	**1.00**	**1.00**	***14.4%***
Surface finish	1.00	1.00	1.00	1.00	***14.4%***
Thermal stability	1.00	1.00	1.00	1.00	***14.4%***
Straightness	1.00	1.00	1.00	1.00	***14.4%***

Figure B.14 Accuracy matrix comparisons and weighting vector.

The design team felt that all four accuracy criteria were equally important. Note that the entries in the weighting vector are the same, but their sum is not equal to unity. The total weight of these *Level Three* criteria must equal the weight assigned to accuracy in *Level Two*. In this case, accuracy accounts for 57.7% of the total decision. This calculation is made by normalizing each *Level Three* weighting vector and multiplying the entries by the appropriate *Level Two* weight (57.7%).

The next *Level Two* criteria is cost. The design team determined that the labor cost was the most important component. The cost of maintaining the machine was rated much less important. The sales price and machine components cost were moderately important considerations, as shown in Figure B.15.

Cost	Sale	Maint.	Labor	Comp.	Weighting Vector
Sale price	**1.00**	**3.00**	**0.50**	**1.00**	***4.4%***
Maintenance	0.33	1.00	0.17	0.33	***1.5%***
Labor to produce	2.00	6.00	1.00	2.00	***8.9%***
Components	1.00	3.00	0.50	1.00	***4.4%***

Figure B.15 Cost matrix comparisons and weighting vector.

Within the *Level Two* criterion of manufacturability, the ease of assembly was considered the most important, as shown in Figure B.16.

Manufacturability	Bearings	Castings	Assembly	Weighting Vector
Bearings	**1.00**	**5.00**	**0.50**	***3.6%***
Castings	0.20	1.00	0.10	***0.7%***
Assembly	2.00	10.00	1.00	***7.2%***

Figure B.16 Manufacturability matrix comparisons and weighting vector.

Figure B.17 shows the *Level Three* breakdown of the elements of ergonomics. The design team found that the interface was very important to the overall design. In fact, the total weight given to the interface (9.6%) approaches that of the various accuracy criteria (14% each).

Ergonomics	Interface	Maintenance	Weighting Vector
Interface	**1.000**	**5.000**	***9.6%***
Maintenance	0.200	1.000	***1.9%***

Figure B.17 Ergonomics matrix comparisons and weighting vector.

Step Two of the AHP was complete once the last matrix in the hierarchy was filled in and the weighting vector was calculated. The design team then studied the different design concepts and decided how each would satisfy the criteria. The two design concepts were similar except for differences in the way the z-axis was mounted. The AHP provides an excellent method of making this tough choice.

Figure B.18 shows the design evaluation scores given to the constant and maximum stiffness configurations. The boldface columns are ratings based on a scale of 1 to 10. In the second pair of columns, the evaluations are normalized and multiplied by the corresponding weighting vector entry from the Step Two matrices.

Accuracy	Constant stiffness	Maximum stiffness	Normal. Const. K	Normal. Max. K
Profile	**5.00**	**3.67**	0.08	0.06
Surface finish	**6.33**	**5.00**	0.08	0.06
Thermal stability	**5.00**	**5.00**	0.07	0.07

Straightness	**5.67**	**3.67**	0.09	0.06
Cost				
Sale price	**5.00**	**6.00**	0.02	0.02
Maintenance	**5.00**	**5.00**	0.01	0.01
Labor	**2.00**	**5.00**	0.03	0.06
Components	**5.00**	**3.00**	0.03	0.02
Manufacturability				
Bearings	**3.67**	**5.00**	0.02	0.02
Castings	**6.67**	**5.33**	0.00	0.00
Assembly	**3.67**	**5.00**	0.03	0.04
Ergonomics				
Interface	**5.00**	**5.00**	0.05	0.05
Maintenance	**5.00**	**2.00**	0.01	0.01
		Rating	***0.52***	***0.48***

Figure B.18 Jig borer concept evaluations.

The final rating was found to be 52% for the constant stiffness design and 48% for the maximum stiffness design. Because the difference between ratings was fairly small, the design team performed additional analyses to investigate the validity of the evaluation scores given to critical items. In this case, accuracy was the critical issue, so the design team performed an error budget analysis to determine the effects of the different z-axis configurations on the accuracy of the machine [Slocum, 1992]. The results, summarized in Figure B.19, show that while the maximum stiffness design has smaller deflection, the deflection changes more over its stroke. The constant stiffness design has relatively high deflection but it remains constant regardless of the toolhead position. The AHP helped uncover this fact by guiding the team to systematically consider the merits of each design and identify areas to be investigated on the next level of detail.

	Constant K design	Maximum K design
Tool to work error (retracted)	1.4 μm	0.6 μm
Tool to work error (extended)	1.5 μm	0.9 μm
Change in deflection	0.1 μm	0.3 μm

Figure B.19 Error budget summary of machine accuracy of the constant and maximum stiffness designs.

Now *what-if* questions can be asked to further ensure that the proper choice is made. A sensitivity analysis may be performed to see how variations in comparisons affected the outcome. Small variations in the final rating indicate that the AHP has already determined the optimal choice. Larger variations may indicate that more effort should be spent on verifying the choices made in Steps Two and Three.

The jig borer example shows how the AHP provides a decision making framework to consider design concepts. Like the Pugh chart, the AHP allows designers to identify parts of different designs that are superior and might be combined to create a new super-concept. Beyond the Pugh Process, the AHP identifies which aspects are worth combining and which are not.

B.8 Calculations Used in the Analytic Hierarchy Process

A spreadsheet implementation of the AHP can be conveniently adjusted to handle different problems with varying numbers of hierarchical levels and criteria. Due to the consistent matrix formulation, the calculations are straightforward. This section documents the calculations made in the AHP and how to implement them on a spreadsheet.

B.8.1 Calculation of the Step Two Weighting Vector

The first calculation is the determination of the principal eigenvector of each weighting matrix. The *Level Two* matrix used in the jig boring machine example will illustrate this calculation. Figure B.20 shows three arbitrary choices made for the relationships between variables along the first row.

Level Two	Accu.	Cost	Manu.	Ergo.
Accuracy	**1.00**	**x**	**y**	**z**
Cost				
Manufacturability				
Ergonomics				

Figure B.20 Arbitrary first row comparisons.

Once the first row has been determined, the rest of the matrix can be completed using the formula $a_{jk} = a_{1k} / a_{1j}$. As shown in Figure B.21, the main diagonal of all the weighting matrices are 1's (because each criterion is of equal importance to itself).

Level Two	Accu.	Cost	Manu.	Ergo.
Accuracy	**1.00**	**x**	**y**	**z**
Cost	1/x	1.00	y/x	z/x
Manufacturability	1/y	x/y	1.00	z/y
Ergonomics	1/z	x/z	y/z	1.00

Figure B.21 Completion of matrix by consistent formulation.

Now that all the matrix entries have been determined, the eigenvector corresponding to the largest eigenvalue must be calculated (the principal eigenvector). The largest eigenvalue is simply the dimension of the matrix; in this example the eigenvalue is four.

All of the other eigenvalues are zero, a result of the rank one formulation. *The normalized weighting vector is the principal eigenvector.* A spreadsheet representation of this calculation is shown in Figure B.22.

	A	B	C	D	E	F	G
1	**Level Two**	Accu.	Cost	Manu.	Ergo.		Weighting Vector
2	Accuracy	**1.00**	**x**	**y**	**z**		**B2/SUM(B2:B5)**
3	Cost	1/C2	1.00	D2/C2	E2/C2		**B3/SUM(B2:B5)**
4	Manufacturability	1/D2	C2/D2	1.00	E2/D2		**B4/SUM(B2:B5)**
5	Ergonomics	1/E2	C2/E2	D2/E2	1.00		**B5/SUM(B2:B5)**

Figure B.22 Spreadsheet adaptation of weighting vector calculations.

Levels of hierarchical decomposition below *Level Two* require an additional calculation. The weight given to each item in the vector is adjusted by the weight given on the previous level. For example, the *Level Three* accuracy matrix shown in the jig borer example is:

Accuracy	Profile	Surf.	Therm.	Straight.	Weighting Vector	Norm. Vector
Profile	**1.00**	**1.00**	**1.00**	**1.00**	***0.25***	***0.14***
Surface finish	1.00	1.00	1.00	1.00	***0.25***	***0.14***
Thermal stability	1.00	1.00	1.00	1.00	***0.25***	***0.14***
Straightness	1.00	1.00	1.00	1.00	***0.25***	***0.14***

Figure B.23 Normalized vector that results from weighting vector and *Level Two* results.

Matrices in *Level Three* and beyond are calculated in a similar fashion. First the principal eigenvector is calculated and its magnitude set to unity. Then each entry in the vector is multiplied by the weight assigned to the matrix from the previous level. For example, if a criterion on *Level n* receives x percent of the total weight, then all sub-criteria (for that criterion) on *Level* $n + 1$ will receive a *total* of x percent. As a result of the normalizing scheme, the total weight of all the end branches will be equal to unity, regardless of what level each branch may be on.

B.8.2 Calculation of the Step Three Design Concept Ratings

Once the weighting vector is calculated, the design concepts may be evaluated and recorded on a spreadsheet. In the jig boring machine example, sample evaluations have been made to illustrate the process.

Accuracy	Weight. Vec.	Const. K	Max. K
Profile	0.14	**5.00**	**3.67**
Surface finish	0.14	**6.33**	**5.00**
Thermal stability	0.14	**5.00**	**5.00**
Straightness	0.14	**5.67**	**3.67**

Figure B.24 Accuracy matrix weighting vector and concept evaluations.

The evaluation for each of the accuracy criteria is shown in bold type. The evaluations are normalized and then multiplied by the appropriate weighting vector entry. In this case, the weighting vector is 14.4% for each accuracy criterion, as determined by the Step Two calculation outlined above. The spreadsheet calculations are shown in Figure B.25. Sample calculations and totals for these criteria are shown in Figure B.26. Note that the sum of the two accuracy columns is 57.7%. When all the criteria are considered, this total will be exactly unity.

	A	B	C	D	E	F
1					Normalized	Normalized
2	**Accuracy**	Weight. Vec.	Const. K	Max. K	Const. K	Max. K
3	Profile	0.14	5.00	3.67	**B3*C3/(C3+D3)**	**B3*D3/(C3+D3)**
4	Surface finish	0.14	6.33	5.00	**B4*C4/(C4+D4)**	**B4*D4/(C4+D4)**
5	Thermal stability	0.14	5.00	5.00	**B5*C5/(C5+D5)**	**B5*D5/(C5+D5)**
6	Straightness	0.14	5.67	3.67	**B6*C6/(C6+D6)**	**B6*D6/(C6+D6)**

Figure B.25 Spreadsheet formulas for normalized concept evaluations.

	A	B	C	D	E	F
1					Normalized	Normalized
2	**Accuracy**	Weight. Vec.	Const. K	Max. K	Const. K	Max. K
3	Profile	0.14	5.00	3.67	0.08	0.06
4	Surface finish	0.14	6.33	5.00	0.08	0.06
5	Thermal stability	0.14	5.00	5.00	0.07	0.07
6	Straightness	0.14	5.67	3.67	0.09	0.06

Figure B.26 Example evaluations for the accuracy matrix of the two design concepts.

The total ratings for a design concept selection are found by repeating this process for all the criteria. The sums of each normalized column can then be taken to determine the best design concept.

B.9 The Eigenvalue Problem Posed by the Analytic Hierarchy Process

The Analytic Hierarchy Process proposed by Saaty has been discussed and an example has been given which illustrates the power of the method. This section investigates some

of the numeric estimations available for finding the principal eigenvector of a matrix. Several approximate methods are given in the literature [Saaty, 1981; Wilkinson, 1965]. The reciprocal matrix used in the consistent formulation of the AHP offers a uniquely simple means of accurately calculating the principal eigenvector.

An arbitrary consistent matrix **A** will be used to demonstrate the proposed solution technique.

$$\mathbf{A} = \begin{bmatrix} 1 & x & y & z \\ 1/x & 1 & y/x & z/x \\ 1/y & x/y & 1 & z/y \\ 1/z & x/z & y/z & 1 \end{bmatrix} \tag{B.1}$$

The best way to solve the eigenvector problem is to examine the special properties of **A**. The eigenvector problem is written as:

$$\mathbf{A}\phi = \lambda\phi. \tag{B.2}$$

A is a rank one matrix and can be written as the product of two vectors.

$$\mathbf{A} = \begin{bmatrix} 1 & x & y & z \\ 1/x & 1 & y/x & z/x \\ 1/y & x/y & 1 & z/y \\ 1/z & x/z & y/z & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1/x \\ 1/y \\ 1/z \end{bmatrix} \begin{bmatrix} 1 & x & y & z \end{bmatrix} \tag{B.3}$$

A rank one matrix has only one non-zero eigenvalue. From the above equation, the eigenvector corresponding to this eigenvalue can be obtained by inspection.

$$\begin{bmatrix} 1 \\ 1/x \\ 1/y \\ 1/z \end{bmatrix} = \phi \quad \text{and} \quad \lambda = 4 \tag{B.4}$$

The only non-zero eigenvalue is simply the dimension of the matrix and its eigenvector is the reciprocal of the entries in the first row. This is the eigenvector shown in previous sections during the explanation of the AHP calculation methods. The advantage of the consistent formulation is that the eigenvector can be obtained by inspection of the **A** matrix. The only computation necessary is the normalization of the eigenvector.

B.9.1 Approximate Solutions to Inconsistent Weighting Matrices

This section outlines several approximation techniques for use when the weighting matrices are inconsistently formulated. These techniques attempt to find the eigenvector associated with the largest eigenvalue (the principal eigenvector). Here we present an overview of several of the most popular approximation methods and evaluate their suitability to spreadsheet adaptation. These methods will be presented in reverse order of

accuracy; the first technique is the worst guess for arbitrary matrices, but it is the easiest to perform. To illustrate these methods, the following *consistent* matrix will be considered. Note that all the methods arrive at the exact answer when applied to a consistently formulated weighting matrix.

$$\mathbf{A} = \begin{bmatrix} 1 & x & y & z \\ 1/x & 1 & y/x & z/x \\ 1/y & x/y & 1 & z/y \\ 1/z & x/z & y/z & 1 \end{bmatrix} \tag{B.5}$$

The matrix **A** will be used as an example of the calculations required for each technique. The symbolic representation of the entries a_{ij} will be useful in visualizing the operations performed.

B.9.2 Approximation One

The simplest method simply takes the sum of the elements in each row and normalizes them so that the total of the sums adds up to one. This method is shown below:

Row1	Row2	Row3	Row4	Sum
1	x	y	z	$1+x+y+z$
$1/x$	1	y/x	z/x	$(1+x+y+z)/x$
$1/y$	x/y	1	z/y	$(1+x+y+z)/y$
$1/z$	x/z	y/z	1	$(1+x+y+z)/z$

(B.6)

This method gives the exact result for the consistent matrix formulation.

B.9.3 Approximation Two

The second technique gives a better approximation to the principal eigenvector of an arbitrary matrix. In this technique, the reciprocal of the sum of each column is normalized to find the principal eigenvector. Shown below are the columns of matrix A and the sum of the entries in each column.

*Row*1	*Row*2	*Row*3	*Row*4
1	x	y	z
$1/x$	1	y/x	z/x
$1/y$	x/y	1	z/y
$1/z$	x/z	y/z	1
$1+1/x+1/y+1/z$	$x(1+1/x+1/y+1/z)$	$y(1+1/x+1/y+1/z)$	$z(1+1/x+1/y+1/z)$

(B.7)

The reciprocal of each column sum is now taken and normalized to find the resulting eigenvector. The result is the same as that found by Approximation One.

B.9.4 Approximation Three

The most accurate approximation to the principal eigenvector uses the geometric mean of each row. The n elements in each row of the matrix are multiplied and then taken to the n^{th} root. Using our example matrix **A** again,

Row1	Row2	Row3	Row4	Geo.Mean
1	x	y	z	$\sqrt[4]{xyz}$
$1/x$	1	y/x	z/x	$\sqrt[4]{\frac{xyz}{x^4}}$
$1/y$	x/y	1	z/y	$\sqrt[4]{\frac{xyz}{y^4}}$
$1/z$	x/z	y/z	1	$\sqrt[4]{\frac{xyz}{z^4}}$

(B.8)

Normalizing the geometric mean gives the exact result if the matrix is consistent.

B.10 Accuracy and Robustness of the Consistent Formulation

Saaty's AHP does not require that the weighting matrices be constructed with a consistent formulation. However, by requiring consistency, the entries and the accuracy of the entire weighting matrix are completely determined by the first row. While simplifying the math, this constraint decreases the resolution of the process. This section investigates the tradeoffs that might accompany this simplification. Numerical studies have been performed to investigate the sensitivity of the weighting matrix to the first row entries.

B.10.1 Accuracy of Consistent Formulation

The consistent formulation has been used to calculate the principal eigenvector for two reasons. The consistent formulation ensures transitivity (for example, if $x = y$ and $x = z$, then $y = z$). Without transitivity, the comparisons lose their intuitive basis. Because the accuracy of the AHP results rely on these subjective judgments, the comparisons must be as intuitive as possible.

An argument for inconsistency is that the user has a better chance of arriving at the correct result if more judgments are made. A matrix that is constructed with the user providing all the upper triangular entries requires $n(n-1)/2$ comparisons. A matrix constructed by the consistent formulation requires only $n-1$ comparisons. The ratio between the two, $n/2$, indicates the amount of information lost in the consistent formulation.

For the decrease in resolution, the user gains: 1) an intuitive feel of the relationships between criteria, and 2) speed in completing the analysis. The consistent method is a compromise between achieving maximum accuracy and presenting the process in a user-friendly format. The consistent formulation allows the novice to quickly learn the process, and the spreadsheet adaptation is formatted to allow the expert to scan the

comparisons that would be made manually in the inconsistent formulation. Visual scanning of the complete matrix reclaims the resolution lost with the consistent formulation.

Entries below the first row with values outside the 1/9 to 9 range should warn the user of a potential problem with the comparisons assigned in the first row. Two possible causes may lead to this numerical warning. The first is that erroneous values were entered in the first row that led to problems elsewhere in the matrix. In this case, the user should use the warning as a reminder to rethink some of the first row evaluations. The second source of over-range values is that the items being compared in the matrix are not actually on the same order of magnitude in importance. In this case, the user must redefine the hierarchy so that the criteria on each level are of the same relative importance. In either case, the entries in the matrix must fall within the range of 1/9 to 9.

B.10.2 Numerical Study of the Accuracy of the Consistent Formulation

The accuracy of the consistent AHP is dependent on the ability of the first row judgments to capture the true relationships of the criteria. A numerical analysis was performed to study this capability. Sample four by four weighting matrices were constructed by randomly generating the first row of the matrix and calculating the rest of the matrix by the consistent formulation. The upper triangular entries in the second through fourth rows were then multiplied by a random error ranging from $1/(1+\varepsilon) \le \varepsilon_i \le 1+\varepsilon$. The lower triangular portion of the matrix was the inverse of the perturbed upper triangular portion as shown.

$$\mathbf{A} = \begin{bmatrix} 1 & x & y & z \\ 1/x & 1 & (\varepsilon_1 y)/x & (\varepsilon_2 z)/x \\ 1/y & x/(\varepsilon_1 y) & 1 & (\varepsilon_3 z)/y \\ 1/z & x/(\varepsilon_2 z) & y/(\varepsilon_3 z) & 1 \end{bmatrix} \tag{B.9}$$

The exact principal eigenvector was then calculated and compared to the eigenvector given by the consistent solution. Using these matrices, the maximum percent error between the consistent and the exact principal eigenvector was calculated for a large sample of matrices. The random perturbation simulated how the user judgments might fluctuate from the actual relationships. This provides insight into how the consistent formulation may introduce some error into the weighting vector.

ε	0.0	0.2	0.4	0.6	0.8	1.0
Error	0.00%	5.84%	10.50%	13.80%	17.80%	23.50%

Figure B.27 Error of consistent approximation made on inconsistent matrices.

Figure B.27 shows that the inconsistent formulation is only moderately useful in fine-tuning the first row comparisons. The data suggest that even random perturbations

that vary from halving to doubling the values of the non first row entries only effect the results by about 24%. This is a worst case estimate of the error between the inconsistent and consistent formulation. Careful scanning of the automatically calculated results by the consistent formulation will result in much smaller errors. These results, combined with the deterministic philosophy of mechanical design dictate that the consistent formulation be used in this adaptation of the AHP.

B.10.3 Restrictions on the Position of Criteria Within a Matrix

Depending on the construction of the weighting matrices, some of the first row comparisons may range from 1/9 to 9. The consistent formulation would then dictate that the matrix look, as shown in Figure B.28 (given a sample first row of 1, 1/9, 3, and 9):

	A	B	C	D	
A	**1.000**	**0.111**	**3.000**	**9.000**	**9.6%**
B	9.000	1.000	27.000	81.000	**86.2%**
C	0.333	0.037	1.000	3.000	**3.2%**
D	0.111	0.012	0.333	1.000	**1.1%**
					100.0%

Figure B.28 Sample criteria weighting matrix violating AHP requirement that all criteria be of the same order of importance.

While the first row numbers fall within the AHP scale (1/9 to 9), the other rows have entries outside this range. In fact, this over-range will occur any time the first row judgments contain entries greater than 3 and less than 1/3. Therefore, this matrix is in violation of the AHP requirement that criteria compared on a single level all be of the same order of importance. In the above matrix, the B and D criteria are almost two orders of magnitude apart in importance. This indicates that either the judgments are incorrect, or that criteria C and D actually belong on a sub-level.

If the user believes that all the criteria compared in this matrix are of the same order of magnitude in importance, then a simple solution is available. Out of range entries can be avoided by designing the matrix so that the criteria of highest importance occupies the first position in the matrix. In the example above, criterion B is the most important, so it would be relocated to the first position. The matrix would then look like Figure B.29. Note that the judgments have been adjusted to fit within the appropriate range.

	B	A	C	D	
B	**1.000**	**5.000**	**7.000**	**9.000**	**68.8%**
A	0.200	1.000	1.400	1.800	**13.8%**
C	0.143	0.714	1.000	1.286	**9.8%**
D	0.111	0.556	0.778	1.000	**7.6%**
					100.0%

Figure B.29 Possible reordering of criteria to improve weighting results.

The weighting vector now satisfies the order of magnitude requirement. By reordering the entries of the example matrix, the comparisons were forced to satisfy this constraint. If the new matrix does not accurately reflect the user's opinion of the relative weights throughout the matrix, than the items of lesser importance (C and D) must be demoted to a sub-level to meet the AHP requirement.

In summary, the consistent formulation presents an intuitive approach to completing the weighting matrices. The proposed spreadsheet adaptation will be user-friendly and accurate. The consistent formulation does require that the user estimate the most important criteria in advance of completing the weighting matrices. This is done for several reasons including:

- The first row entries will all be greater than or equal to 1, eliminating a potential source of confusion (do I use 3 or 1/3?).

- The remaining rows will all contain entries within the 1/9 to 9 range. The upper triangular entries will be between 1 and 9. This automatically ensures satisfaction of the AHP requirement that entries in a particular matrix be of the same order of magnitude in importance.

- The most important criterion is often readily identified in advance. Even if the second most important criterion is mistakenly put in the first position, the results will be better than if a midrange criteria is used. A criteria of midrange importance in the first position means that the first row could contain entries such as 1/5 and 5 which leads to over-range row entries.

B.11 Case Study: Consumer Product Design Concept Selection

As an example of the adaptation of the AHP for design concept selection, results from a case study are presented.

This case study involves a consumer goods company in the process of developing a new electronics product. Two design consulting firms were retained to develop design concepts in parallel with internal efforts to produce the most satisfactory design. At the

end of several months of concept development, a decision was required as to which concept should be taken to market.

Engineers and managers from the consumer goods company and the two design firms were briefed on the use of the AHP. The three groups were also provided with a spreadsheet template and instructions. The consumer goods company generated criteria for the AHP analysis based on their experience with similar products and the specific goals for the new product.

The three design teams were able to experiment with the AHP spreadsheet to become comfortable with its use, as well as to apply it to their own designs prior to the final decision making sessions held by the company. During this time, the author and users made suggestions on how to refine the criteria. This led to enhanced understanding of the design problem.

After the three design teams were familiar with the AHP, the company managers and engineers assembled in groups and completed the weighting matrix evaluations. The groups included: engineering and project management; electrical, mechanical, and reliability engineers; and designers. A method of obtaining an average weighting vector from the entire design team was developed using a transformation of the data. Half of the AHP scale is linear between 1 and 9 and the other half is decidedly nonlinear (1/9 to 1). For this reason, the evaluations were linearized by the following equations:

$$\text{Linear weight} = \frac{\text{AHP weight} - 1}{-1/81 + \text{AHP weight}} \qquad \text{for AHP weight} < 1$$

$$\text{Linear weight} = \text{AHP weight} \qquad \text{for AHP weight} \geq 1 \tag{B.10}$$

These linearization equations transform the original AHP scale of 1/9 to 9 into a linear scale of -9 to 9. The mean and standard deviation are then calculated and transformed back into the AHP scale. Two important results were found during this study:

1. The design team converged on good results using the proposed adaptation of the AHP.

2. The importance of having well-defined weighting criteria cannot be overstated.

There was a direct correlation between the standard deviation of the group weighting evaluations and the ambiguity surrounding the individual criteria. For example, the results for *product reliability* were excellent. Another criterion, *invention*, was not as clearly understood. Invention was supposed to be a measure of the innovativeness of the solution. In one sense, a product with a new invention would be desirable because of

patent protection, but a product requiring the development of new technologies may take too long to take to market. This ambiguity was reflected in the final results.

In conclusion, this case study illustrates the power of the proposed AHP adaptation. The process was used as an efficient framework for making a design concept evaluation.

B.12 Spreadsheet Adaptation

The hierarchical structure of the AHP makes it ideal for computer spreadsheet adaptation. Math calculations are transparent to the user, and more importantly, the spreadsheet can quickly generate all of the necessary results given just a few comparisons provided by the user. Once the analysis is complete, the spreadsheet is a tool to perform sensitivity analyses and answer *what-if* questions.

B.12.1 Spreadsheet Adaptation

The upper left hand corner of the spreadsheet, `template.xls`, will be visible when loaded into a popular spreadsheet program. The spreadsheet window should look something like Figure B.30.

1. Develop criteria for concept evaluation and arrange in hierarchical form
2. Replace shaded axes labels with criteria names
3. Fill in top row of matrices (shaded) using the following scale:

1: both criteria of equal importance	0: no criteria for this position
3: left weakly more important than top	1/3: top weakly more important than left
5: left moderately more important than top	1/5: top moderately more important than left
7: left strongly more important than top	1/7: top strongly more important than left
9: left absolutely more important than top	1/9: top absolutely more important than left

4. Choose design concepts and enter names in results box (shaded)
5. Evaluate concepts ability to satisfy criteria (shaded)
6. Check final rating results!

Final Results:				
Des 1	Des 2	Des 3	Des 4	Des 5
20%	20%	20%	20%	20%

Level Two

	A	B	C	D	E	F	G	H	
A	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	13%
B	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	13%
C	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	13%
D	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	13%
E	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	13%
F	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	13%
G	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	13%
H	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	13%
									1.00

Level Three Matrix 1

	A1	B1	C1	D1	E1	F1	G1	H1	
A1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	2%
B1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	2%
C1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	2%

User evaluations here:

Level Three Matrix 1	Des 1	Des 2	Des 3	Des 4	Des 5	Des 1	Des 2
A1	1	1	1	1	1	0.00	0.00
B1	1	1	1	1	1	0.00	0.00
C1	1	1	1	1	1	0.00	0.00

Figure B.30 Landscape view of `template.xls`.

In the upper left corner of Figure B.30 is a summary of the rules for using the spreadsheet. The final AHP results are shown in a box below the rules. Three sets of columns are shown at the bottom of the figure. The first column is formatted to contain the matrices that will be used to weight the user's criteria. The second column records the evaluation scores the user provides. The final column is used by the spreadsheet to calculate the final results. The entries (cells) that the user must supply are shaded. For example, the design choice names in the final results box must be filled-in for each problem.

Selection of a design concept for the jig borer machine will demonstrate the capabilities of the template. A step by step demonstration of the selection process and its final results are provided in the following sections.

B.12.1.1 Step One: Criteria Selection and Hierarchy Formation

For the jig borer design concept selection, the same hierarchical structure will be used as before. The spreadsheet template can handle a more detailed analyses, but this example shows the key concepts. The spreadsheet template is designed to handle as many as eight *Level Two* criteria, with eight *Level Three* criteria for each.

The *Level Two* matrix and the eight blanks on each axis may be seen by scrolling down the spreadsheet. Step Two requires the replacement of the generic axes labels with the criteria names used in the analysis. The jig borer hierarchy indicates that the *Level Two* criteria are accuracy, cost, manufacturability, and ergonomics. These criteria labels, as well as the *Level Three* labels are entered at this stage. The *Level Two* matrix will then look something like Figure B.31 (note that the top row entries are automatically abbreviated to fit the column width).

Level Two									
	Accu.	Cost	Manu.	Ergo.					
Accuracy	**1.00**	1.00	1.00	1.00	1.00	1.00	1.00	1.00	**12.5%**
Cost	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	**12.5%**
Manufacturability	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	**12.5%**
Ergonomics	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	**12.5%**
	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	**12.5%**
	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	**12.5%**
	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	**12.5%**
	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	**12.5%**
									100.0%

Figure B.31 Partially initialized *Level Two* matrix.

This labeling process is repeated for all matrices in the AHP analysis. Note that each matrix is pre-formatted to hold eight criteria. The template allows for smaller numbers of

criteria to be used by simply typing a 0 (zero) in the unwanted cells on the shaded first row. These zeros will not be seen in the spreadsheet window, as shown in Figure B.32.

Level Two									
	Accu.	Cost	Manu.	Ergo.					
Accuracy	**1.00**	**1.00**	**1.00**	**1.00**					**25.0%**
Cost	1.00	1.00	1.00	1.00					**25.0%**
Manufacturability	1.00	1.00	1.00	1.00					**25.0%**
Ergonomics	1.00	1.00	1.00	1.00					**25.0%**
									100.0%

Figure B.32 Completely initialized *Level Two* matrix.

All the remaining matrices are initialized in a similar manner. Note that the weighting vector automatically changed from 8 × 12.5% to 4 × 25.0%. In both cases the total weight is 100%.

B.12.1.2 Step Two: Criteria Comparison

After labeling all the axes, comparisons are made that will determine the weighting vector. Only the shaded numbers in the first row need to be supplied by the user. The remaining rows of the matrix will be computed automatically and serve as a check to ensure that the first row numbers are correct. For the jig borer example, the appropriate *Level Two* weights are as shown in Figure B.33.

Level Two									
	Accu.	Cost	Manu.	Ergo.					
Accuracy	**1.00**	**3.00**	**5.00**	**5.00**					**57.7%**
Cost	0.33	1.00	1.67	1.67					**19.2%**
Manufacturability	0.20	0.60	1.00	1.00					**11.5%**
Ergonomics	0.20	0.60	1.00	1.00					**11.5%**
									100.0%

Figure B.33 *Level Two* matrix with criteria comparisons.

The other matrices will require completion by the same procedure. The spreadsheet automatically calculates the weighting vector as the user enters the first row of data. This vector is shown in bold face and presented as a percentage of the total weight.

B.12.1.3 Step Three: Concept Evaluation

After determining the weighting vector, the design concepts may be compared on an individual basis. The spreadsheet will use the evaluation scores and the weighting vector to calculate a total rating for each of the design concepts.

The spreadsheet is set up to handle five different design concepts. The user must make evaluations for each of the *Level Three* criteria. The default evaluation score for each design concept is 1. The user replaces these ratings with the appropriate scores. Figure B.34 shows a portion of this design concept evaluation section of the spreadsheet.

Accuracy	Des. 1	Des. 2	Des. 3	Des. 4	Des. 5	Des. 1	Des. 2	Des. 3	Des. 4	Des. 5
Profile	1	1	1	1	1	0.013	0.013	0.013	0.013	0.013
Surface finish	1	1	1	1	1	0.013	0.013	0.013	0.013	0.013
Thermal stability	1	1	1	1	1	0.013	0.013	0.013	0.013	0.013
Straightness	1	1	1	1	1	0.013	0.013	0.013	0.013	0.013
	1	1	1	1	1					
	1	1	1	1	1					
	1	1	1	1	1					
	1	1	1	1	1					

Figure B.34 Uninitialized sample concept evaluation matrix.

Using the evaluations listed in a previous section, the results box shows the best design concept based on the weighting matrix. Since only two design concepts were considered, the third, fourth, and fifth design concept have scores of zero.

Final Results:					
Design 1	Design 2	Design 3	Design 4	Design 5	
51.6%	**48.4%**	**0.0%**	**0.0%**	**0.0%**	100.0%

Figure B.35 Final results from the AHP study.

This final result is the conclusion of the process. Using the spreadsheet, a sensitivity analysis may be conveniently performed to check the results.

B.13 Conclusion

This appendix has reviewed the Analytic Hierarchy Process and provided examples of its application to design concept selection problems. The key to this application is group participation. A project team can achieve excellent results in a reasonable amount of time using the AHP framework. Furthermore, by uniting key people early on in the design process, the final concept is approved automatically at the close of the AHP discussions.

An important outcome of group participation is that it forces the team to constantly make comparisons between designs that will guide the selection of the superior concept.

The AHP elucidates the strengths of the design concepts being considered so that combination concepts may be developed.

The spreadsheet adaptation is presented in a conveniently packaged tool for practicing engineers to use in decision making. The structured framework of the AHP template allows the methodical breakdown of difficult problems into manageable parts, the results of which are combined transparently to identify the best solution.

B.14 References

Bryan, James B., The Power of Deterministic Thinking in Machine Tool Accuracy, Lawrence Livermore National Laboratory, Livermore, CA, 1984.

Evans, Chris, Precision Engineering: An Evolutionary View, Cranfield Press, Bedford, 1989.

Hauser, John R. and Don Clausing, The House of Quality, Harvard Business Review, May-June 1988.

Kepner, Charles H., and Benjamin B. Tregoe, The New Rational Manager, Princeton Research Press, Princeton, New Jersey, 1981.

Pugh, Stuart, Concept Selection: A Method that Works, Proceedings of the International Conference on Engineering Design (ICED), Rome, 1981.

Saaty, Thomas, The Analytic Hierarchy Process, McGraw-Hill, New York, 1980.

Slocum, Alexander H., Precision Machine Design, Prentice Hall, Englewood Cliffs, 1992.

Steward, Donald V., The Design Structure System: A Method for Managing the Design of Complex Systems, IEEE Transactions on Engineering Management, August 1981.

Suh, Nam P., Axiomatic Design, Oxford University Press, New York, 1985.