Multiple Mode Vibration Suppression in Controlled Flexible Systems

James M. Hyde

MIT Artificial Intelligence Laboratory

This blank page was inserted to preserve pagination.

MULTIPLE MODE VIBRATION SUPPRESSION IN CONTROLLED FLEXIBLE SYSTEMS

by

James Matthew Hyde

B.S. Mechanical Engineering
Stanford University
June 1989

Submitted to the Department of Mechanical Engineering
In Partial Fulfillment of the Requirements
For the Degree of

Master of Science in
Mechanical Engineering

at the

Massachusetts Institute of Technology
May 1991

IV

MULTIPLE MODE VIBRATION SUPPRESSION IN CONTROLLED FLEXIBLE SYSTEMS

by

James Matthew Hyde

Submitted to the Department of Mechanical Engineering on
May 10, 1991 in partial fulfillment of the requirements for the
degree of Master of Science in Mechanical Engineering

Abstract

Spacecraft, space-borne robotic systems, and manufacturing equipment often
utilize lightweight materials and configurations that give rise to vibration
problems. Prior research has led to the development of input command pre-
shapers that can significantly reduce residual vibration. These shapers exhibit
marked insensitivity to errors in natural frequency estimates and can be
combined to minimize vibration at more than one natural frequency.

In this work we present a method for the development of multiple mode
input shapers which are simpler to implement and produce smaller system
response delays than previous multiple mode designs. The new technique
involves the direct solution of a group of simultaneous non-linear impulse
constraint equations.

An MIT/NASA experimental flexible structure, MACE, is employed as a test
article for the validation of the direct solution shaping technique. We
examine the results of shaper performance tests conducted on linear and non-
linear computer models of MACE. Vibration problems caused by system non-
linearities are identified and their eradication attempted.

The direct solution shapers are shown to be effective in suppressing multiple
mode vibration, even in the presence of mild kinematic and dynamic non-
linearities.

Thesis Supervisor: Warren P. Seering
                       Professor of Mechanical Engineering

3

Acknowledgements

There wasn't enough room on the title page to include all the names that truly belong there, so I am forced by tradition to thank everyone here on page five. I have to start with my thesis advisor, Professor Warren Seering. Warren's unique blend of steady guidance and casual freedom allowed me to pursue this research at my own pace, while keeping my motivation high. My lasting impression is that Warren has a solid grasp on the concept of balance, a trait I sincerely admire. .

Those with perhaps the most immediate impact on my work have been the members of Prof. Seering's research group in the MIT Artificial Intelligence Laboratory. I thank Brian Avery, Mike Caine, Ken Chang, Andrew Christian, Whit Rappole, Lukas Ruecker, Bill Singhose, Kamala Sundaram, Bruce Thompson, Tim Tuttle, and Erik Vaaler for their insightful advice, keen talent, and collective sense of humor. Erik Vaaler and Inaki Garabieta were particularly helpful in their willingness to share their vast pool of mechanical design knowledge.

The research group at the MIT Space Engineering Research Center has also been extremely helpful in bringing this work to fruition. I thank Prof. Ed Crawley, Dr. David Miller, Dr. Mathieu Mercadal, Paul Bauer, Mark Campbell, Roman Hachkowski, Carlos Padilla, Marco Quadrelli, Dan Rey, and Eric Saarma for their support and their "Aero/Astro" point of view.

I collectively thank the people who really run MIT, namely the secretarial, support and technical staffs who keep us all on track and reasonably sane.

My family has helped me in too many ways to describe here, but I can say that without the faith, guidance, and love of my father, mother, sister, and other relatives, I would not be writing this now. Special thanks also go to Angelina Beitia, whose support, understanding, and patience have been invaluable in my often harrowing academic lifestyle.

My roommates, Robert Borchers and Thomas Moyer, have shared with me the occasional broken window, scattered bike parts, and a pint or two. Their motivation, sincerity, and tolerance of myself leaves me in their debt.

Finally, I am grateful for the financial support I have enjoyed over the last two years, or, in more formal language...

Table of Contents

Table of Figures

Chapter 1: Introduction

1.1 Background

Vibration is a problem common to a large fraction of mechanical systems. The earliest reciprocating steam or water driven devices can claim a close bond to even the most advanced modern robotic systems, in that mechanism motion invariably induces troublesome vibration. Problems caused by vibration can range from life-threatening structural failures to expensive mass production assembly delays.

Reducing or even eliminating system vibration are goals worthy of concentrated effort, and in this thesis we will examine some approaches for achieving those goals. The immediate benefits of general vibration suppression are improved system performance, increased safety, greater reliability, and reduced operating costs. These in turn translate into the long-term advancement of industrial and scientific hardware, the effects of which can be viewed as global.

15

1.2 Motivation

One of the more promising areas for the implementation of vibration suppression techniques is robotics, broadly considered as the field of computer controlled electro-mechanical devices. We can distinguish two areas of closely related, yet somewhat divergent, fields of robotics: earth-based and space-based systems. Each area can benefit from the successful application of vibration suppression techniques, both in terms of increased performance and reduced cost.

In an earth-based mass production setting, just a few seconds lost in the assembly of an electronics board can add up to an annual financial loss of millions of dollars. Industrial robots increasingly utilize lighter structural elements to improve the speed of automated assembly. The combination of a lightweight structure with high performance requirements often leads to serious vibration problems, however, and the resulting increased settling times can undermine the original desired gains in maneuvering speed.

Another "industrial" example where vibration reduction techniques can afford great advantages is in the area of data storage, particularly relating to hard disk drive mechanisms. The head that floats over the disk, reading magnetic information, must be accurately positioned over a particular sector on the platen to ensure proper data readings. When the head servos to a different sector on the disk, the controller must pause after the move is completed, allowing the residual vibration to settle and ensuring an error free read. If the head was able to conclude a move and have no residual vibration

present, data reads could take place immediately, and the data access time would be decreased.

Space-based systems share many of the vibration problems of earth-based systems, especially the difficulties caused by the desire to decrease weight. The cost of transporting equipment to orbit has risen to thousands of dollars per pound, forcing designers of space borne robotic systems to trim the mass of their devices and operate under the strictest weight budgets. The low weights certainly facilitate launching, but high flexibility is a common result, giving rise to chronic vibration problems.

It is common, for instance, for operators of the Space Shuttle's Remote Manipulator System to spend 20 to 40 seconds waiting for oscillations to decay after maneuvering the arm. Manipulating a massive satellite acts to further degrade the maneuvering speed, and with shuttle operating costs at approximately $20,000 per minute, decreasing the maneuvering time of the arm through some means of vibration reduction is attractive, both financially and from a performance viewpoint.

Future space program robotic efforts also include the development of large, earth-monitoring satellites, as discussed in Chapter 4. Two systems mentioned in that chapter are the Earth Observing System (EOS) and the Geosynchronous Platform (GEOS). These systems are multi-body platforms, meaning that they feature many different instruments all mounted on a common structure. Again, to keep the system weight low, the structure is often designed to have a lightweight, and thus flexible, configuration. The vibratory interactions of multiple scanning payloads and antennae with each

other and the structure can potentially cause serious pointing and tracking accuracy problems. Reducing the vibration inherent in the spacecraft would allow the satellite a greater chance of successfully completing its mission.

We can conclude that partial or complete suppression of system vibration can improve spacecraft durability and performance, and would allow manufacturing systems to operate faster and thus more economically.

1.3 Previous Work

Many researchers in the past decade have examined the problems of vibration reduction, with varying success. This thesis stands firmly on the shoulders of those previous investigations, the more relevant of which we consider here.

1.3.1 Vibration Suppression

In this section we will examine some previous efforts aimed at general vibration reduction. These works follow no particular method, but are presented here as a survey of influential research.

Nurre, Ryan, Scofield, and Sims [20] wrote a general study of some of the problems caused by system vibration and possible approaches to curbing vibration. These researchers constrained their examination to vibration concerns relating to large space structures, but their paper outlines modelling and control techniques that are broadly applicable to any flexible, vibration plagued system.

Cannon and Schmitz [7] experimented with the non-colocated feedback control of a flexible beam. Through the use of highly accurate system models and optical tip position sensing they achieved significant vibration reduction and precise tip positioning in their planar test article.

Book, Maizza-Neto, and Whitney [4] examined three methods for controlling a two-link manipulator with flexible members. A joint-angle feedback control scheme incorporating transfer matrices and numerical techniques was shown to improve manipulator performance, causing only a slight increase in controller complexity. A second method, which fed back data on the manipulator beam flexible modes, improved performance but exhibited high sensitivity to parameter perturbations and required a significant increase in controller complexity.

Asada, Ma, and Tokumaru [1] employed inverse plant dynamics to develop feedforward vibration reducing manipulator trajectories. Virtual rigid link coordinates simplified the system vibration constraints, reducing the complexity of the inverse dynamics calculations. This technique demonstrated performance enhancement in simulations of a flexible two-link manipulator.

Turner and Junkins [29] studied the non-linear dynamics of a modelled spacecraft featuring a central hub and four flexible appendages. Through an application of optimal large angle reorientation principles and an "assumed modes" modelling of the appendage flexible behaviour, rapid maneuvers were achieved in simulation that suppressed the first four flexible modes of the structure.

Yurkovich, Pacheco, and Tzes [31] reduced vibration in the presence of unknown and/or varying payloads by employing on line system identification and controller tuning. By using frequency domain techniques to examine the system response following a sample input, enough information was collected to adjust the controller gain scheduling to compensate for vibration problems.

Wie and Liu [32] employed H_∞ controllers to reduce vibration while providing robustness to modelling errors. This technique displayed solid performance, but was relatively difficult to implement.

Crawley and de Luis [8] used distributed segmented piezo-electric actuators to actively control the dynamic vibration and shape characteristics of a flexible aluminum beam. They demonstrated the potential for piezo-electric materials to be used as imbedded actuators capable of suppressing the vibration of flexible structures.

1.3.2 Input Command Shaping

An attractive vibration reduction method not mentioned above is input command shaping. Using this method, commands can be fed through a shaper and into the flexible system, and ideally the resulting output will be vibration free. Shapers also usually reside completely outside of a given control system and are thus easily compatible with other vibration reduction schemes (see Figure 1.1). Smith [25] conducted early shaping investigations, using "positive-cast," or posicast control to create inputs that would act to suppress vibration.
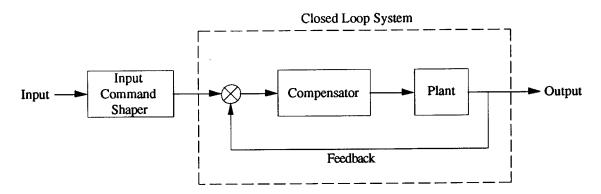
Figure 1.1: Shaper position in control system.

Meckl and Seering [16] examined the use of shaped force profiles to reduce vibration in manufacturing systems. Meckl created profiles by using a versine (1 − cosine) function to modify force commands. When integrated twice, these force profiles become input trajectories that reduce system vibration at a structure's first natural frequency.

Aspinwall [2] modelled the effects of shaping pulse input commands delivered to a free-free single flexible beam. By selecting the proper coefficients for a short Fourier series expression for the forcing function, the spectra of the system residual vibration was suppressed at desired points, namely at the system's natural frequencies.

Farrenkopf [10] established optimal open-loop slewing profiles for flexible spacecraft maneuvering between two quiescent attitude states. Studies of a simple, single structural mode dynamic spacecraft model revealed that the product of maneuver time and system natural frequency is a key parameter influencing structural excitation. For certain time-frequency products, calculus of variations methods can be used to form optimal profiles which reduce post-maneuver structural deformation.

A major problem with input command shaping techniques is that their success usually depends on solid prior knowledge of plant dynamics. Many attempts at input shaping have been criticized because the shapers exhibited significant dependence on precise system models.

Gupta, Lyons, Aubrun, and Margulies [12] examined the benefits of using frequency-shaping methods imbedded in the hierarchical control of large space structures. They developed approaches for controlling flexible systems with poorly damped modes, where only a small number of low-frequency modes were known accurately. By employing a hierarchical control structure of plant-dependent high-authority control laws, robust low-authority controllers, and passive actuator damping, they proposed to increase the parameter insensitivity of control systems that utilize shaping methods.

Swigert [27] employed torque shaping techniques to suppress the residual vibration of a two-link flexible structure. The shaped torque profiles were generated using terminal boundary conditions, and excited vibration during a particular maneuver, but pushed the vibration amplitude to zero at the conclusion of a slew. This shaping method is very sensitive to plant parameter variations if the torque input is minimized, but can be adjusted to be less sensitive through an iterative numerical derivation technique.

Singer [23] presented a computationally simple shaping algorithm that demonstrated strong insensitivity to modelling errors. The shapers were assembled from impulse sequences; to shape a command, one convolved it

with the impulse train. These shapers produced small delays in system response times, on the order of one period of a system's natural frequency.

Singhose [24] revised Singer's work to increase the shaper insensitivity to errors in parameter estimates. By allowing a small amount of residual vibration to occur when the system modelling is perfect, shapers can be designed which are insensitive to natural frequency estimate errors as large as ± thirty percent.

Tzes, Englehart, and Yurkovich [30] studied the effects of combining input shaping with a closed loop acceleration feedback controller. They conducted tests on a flexible beam and proved that each technique can complement the other, resulting in enhanced vibration reduction. This work supports the assertion that input command shaping can be used concurrently with other vibration suppression schemes.

Singer [23] originally assembled shapers designed to cancel single mode vibration and convolved these shapers together to handle multiple mode problems. Simpler multiple mode impulse trains can be assembled by directly solving a full set of constraint equations. These shapers are functionally equivalent to the original shapers, and exhibit savings in response time and implementation complexity.

1.4 Problem Statement

Our objective in this thesis is to develop a method of suppressing multiple mode vibration in flexible systems without invoking massive computational burdens or causing large system response delays. Since impulse sequence

based input shapers have been proven effective in the past, we would like the new method to be an extension of the established technique, exhibiting equivalent vibration reduction while avoiding some of the performance degradation caused by previous designs.

We will derive a direct solution multiple mode shaping method that can circumvent some of the problems caused by convolution based input shapers. We will examine the shaper effects on linear and non-linear simulations of a physical structure. Problems caused by the system non-linearities will be identified, and their eradication attempted.

The remainder of this thesis is divided into five chapters:

Chapter 2 examines fundamental single mode input command shaping, and develops the equations which serve as the theoretical foundation for much of the rest of the thesis. Simple experimental results are presented to confirm the vibration reducing effectiveness of the input shaping method.

Chapter 3 extends the fundamentals of Chapter 2 toward multiple mode problems. Two methods, convolution and direct solution, are examined in detail. We compare the relative strengths and weaknesses of the two multiple mode shaping approaches, and present a computational framework for the difficult generation of the direct solution sequences.

Chapter 4 introduces the Mid-deck Active Control Experiment (MACE), a test article which will be employed in validating the direct solution shaping

technique. We examine some of the motivational and developmental concerns of MACE.

Chapter 5 presents the linear and non-linear computer simulations of MACE, used to determine the direct solution shaper's effectiveness in suppressing vibration.

Chapter 6 concludes the thesis and provides suggestions for future work in the area of vibration suppression.

Chapter 1: Introduction

Chapter 2: Fundamental Input Command Shaping

2.1 Introduction

This chapter presents the basic dynamics which define input command shaping. This examination will begin with the simple concept of residual vibration in single mode systems, progress through the basic shaper equations and plant parameter insensitivity concerns, and finish with some examples demonstrating the effects of input command shaping on the response of physical systems.

Some of the material in this chapter is paraphrased from Singer [23]. This material is included here to provide a concise background on input command shaping, and to define the equations which will occupy a few of the later chapters. For a full derivation of the following equations, see Singer, but note that the equations presented here originate from time domain response relationships, not transformed frequency domain relationships, and therefore some of these equations will differ from those found in Singer.

2.2 Single Mode Shaping

A standard second order system, such as a car suspension, will oscillate in a decaying sinusoidal pattern after experiencing an impulse input, e.g. when you press your hand down quickly on a car hood and then let go. This oscillatory behaviour is shown in Figure 2.1. A second impulse, delivered at the proper time with the correct magnitude, will set up vibrations which are 180° out of phase with the motion caused by the first impulse, as shown in Figure 2.2. The response to the second impulse will tend to cancel the response to the first impulse, yielding a maneuver that is vibration free following the application of the second impulse, as shown in Figure 2.3.
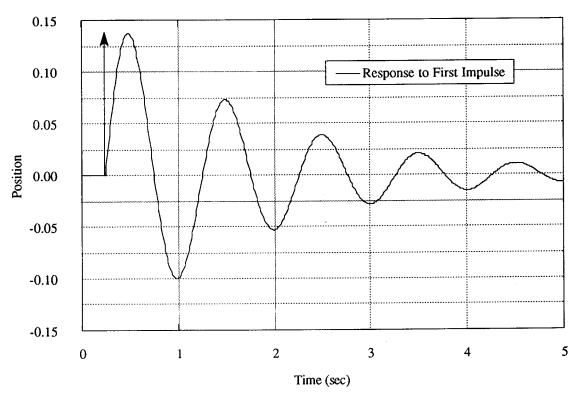
Figure 2.1: Second order system response to single impulse input.
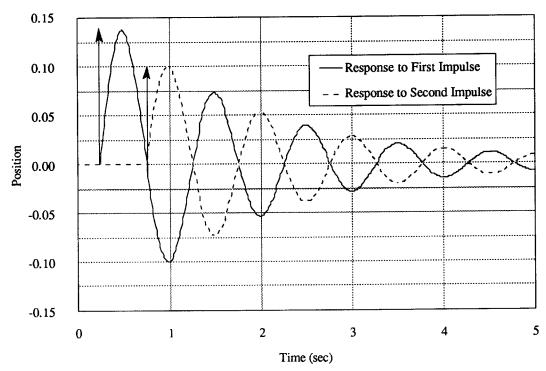
Figure 2.2: Individual second order system response to two impulses.
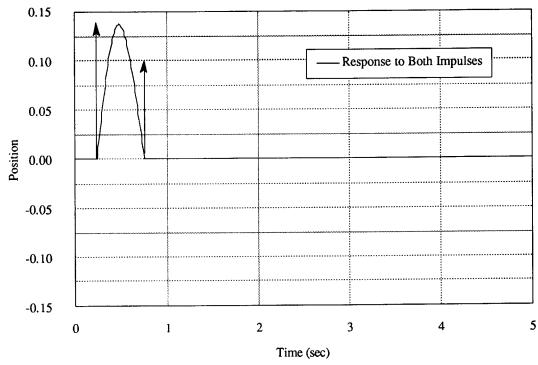
Figure 2.3: Combined second order system response to two impulses.

To derive a mathematical description of this cancellation, we first note that the second order system response to an impulse input is described by:

$$y_i(t) = A_i\, e^{-\zeta\omega(t-t_i)} \sin\!\left((t-t_i)\,\omega\sqrt{1-\zeta^2}\,\right) \tag{2.1}$$

where $y_i(t)$ is the output, A_i is the impulse amplitude and t_i is the time at which the impulse occurs. The system's vibration frequency is ω, with damping ζ. If the system is linear, its total response to a series of N impulses can be expressed as a sum of the the responses to each impulse "i." The amplitude of the total response immediately following the delivery of the Nth impulse is given by:

$$Amp = \left[\left(\sum_{i=1}^{N} A_i\, e^{-\zeta\omega(t_N-t_i)} \sin\!\left(t_i\,\omega\sqrt{1-\zeta^2}\,\right)\right)^2 +\right.$$

$$\left.\left(\sum_{i=1}^{N} A_i\, e^{-\zeta\omega(t_N-t_i)} \cos\!\left(t_i\,\omega\sqrt{1-\zeta^2}\,\right)\right)^2\right]^{1/2} \tag{2.2}$$

A train of properly arranged impulses can suppress residual vibration by forcing *Amp* to equal zero. This can only happen when both the sine and cosine terms in equation (2.2) independently equal zero:

$$\sum_{i=1}^{N} A_i\, e^{-\zeta\omega t_i} \sin\!\left(t_i\,\omega\sqrt{1-\zeta^2}\,\right) = 0$$

$$\sum_{i=1}^{N} A_i\, e^{-\zeta\omega t_i} \cos\!\left(t_i\,\omega\sqrt{1-\zeta^2}\,\right) = 0 \tag{2.3}$$

To construct an impulse sequence that will act as a vibration reducing input shaper, we start by imposing two constraints:

$$t_1 = 0 \tag{2.4}$$

$$\sum_{i=1}^{N} A_i = 1 \tag{2.5}$$

The first is simply an origin specification, and the second is a normalization constraint. Normalizing a shaper's impulse magnitudes ensures that a shaped input will not exceed limitations imposed on the original input, such as actuator saturation or stress limits. We specify an arbitrary value for A_1, and with $N = 2$, we can use equations (2.3) to solve for the time and amplitude of the second impulse in a two-impulse shaper. Invoking equation (2.5) completes the shaper development. A two impulse shaper designed for $\omega = 1$ Hz and $\zeta = 0.1$ is shown in Figure 2.4.
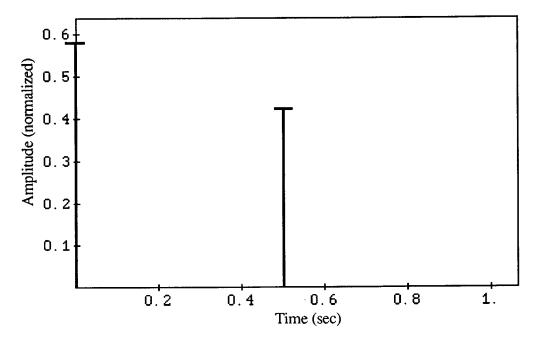
Figure 2.4: Two impulse sequence.

This impulse sequence will completely cancel residual vibration in a single mode system, as long as the natural frequency and damping ratio are perfectly known. For small errors in these plant parameter estimates, the post-maneuver residual vibration will grow rapidly. This behaviour is shown in Figure 2.5, the "insensitivity curve" for the impulse sequence of Figure 2.4. The horizontal axis on the graph is the actual plant frequency, and the vertical axis represents the normalized amplitude of the residual vibration. We create the insensitivity curve by plotting the residual vibration amplitude expression, equation (2.2), using ω as the independent variable. If the amplitude equals unity, the residual vibration will have the same peak to peak level as the response to an unshaped input. A residual vibration level of about 5% is generally considered acceptable.
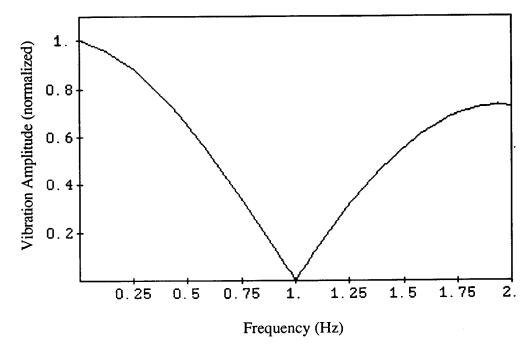
Figure 2.5: Insensitivity curve for two impulse sequence.

When the actual plant frequency is 1 Hz, the residual vibration amplitude is zero, as expected. When ω drifts away from 1 Hz, however, the amplitude curve rises steeply. By allowing for some uncertainty in our estimates of plant parameters, we can attempt to "flatten" the insensitivity curve, a topic of the next section.

2.3 Adding Insensitivity

The two impulse sequence of Figure 2.4 will completely suppress a single mode system's residual vibration, but only if the values of ω and ζ are known exactly. If there is an error in the estimate of one or the other of these parameters, the impulses in the sequence will be spaced incorrectly, and will

be assigned improper magnitudes. If the second impulse from Figure 2.3 is shifted slightly to the right, some of the residual vibration will remain in the system response, as shown in Figure 2.6.

The point to realize here is that the residual vibration is sinusoidal, regular, and thus probably repeatable. If we erred and caused the vibration to occur once, we can most likely do it again. And, if we create a duplicate of the residual vibration out of phase with the original residual vibration, the total post-maneuver vibration can be suppressed. Our flawed two impulse shaper was responsible for the improper cancellation in Figure 2.6, but if we copy the shaper and apply the copy immediately following the original shaper, we can produce the phase shift required to suppress the vibration caused by the first sequence. The duplicated two impulse shaper is equivalent to a single three impulse shaper, as shown in Figure 2.7.

The response to each individual two impulse sequence is shown in Figure 2.8. By adding the two traces, we get the total response shown in Figure 2.9. As the figures clearly illustrate, we have successfully employed a flawed two impulse shaper to suppress the system's residual vibration. The process of combining the two impulse shaper with itself to yield a three impulse shaper imparted some parameter insensitivity to our sequence. Since we can never perfectly know our system parameters, this measure of insensitivity will be crucial for proper vibration suppression.
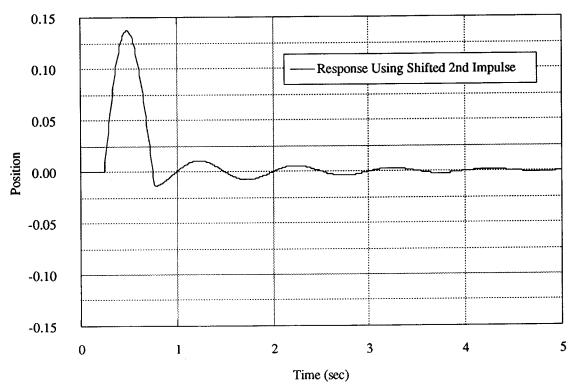
Figure 2.6:   System response using an improperly positioned second impulse.
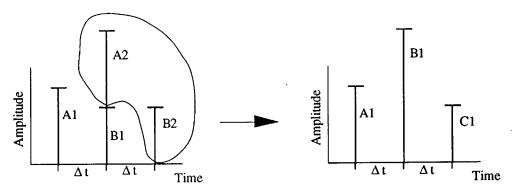
Figure 2.7:   Combining two two-impulse sequences into one three-impulse sequence.
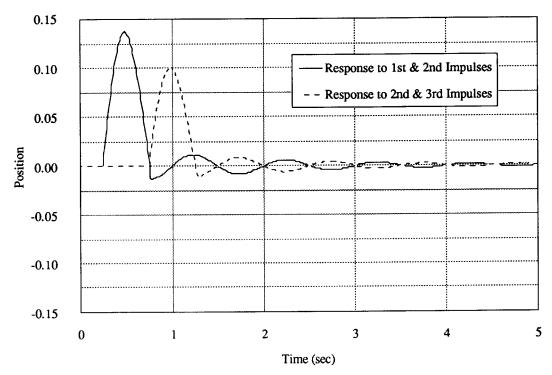
Figure 2.8: Repeating the effects of the incorrect two impulse shaper.
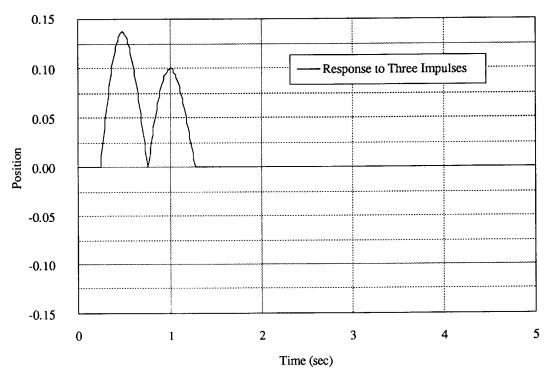
Figure 2.9: Response to three impulse shaper.

The above figures presented an intuitive approach for using three impulse sequences to deal with parameter insensitivity. To develop the equations defining the three impulse sequence, we start by differentiating equations (2.3) with respect to natural frequency. This generates two additional impulse constraints:

$$\sum_{i=1}^{N} A_i \, t_i \, e^{-\zeta \omega t_i} \sin\left(t_i \, \omega \sqrt{1-\zeta^2}\right) = 0$$

$$\sum_{i=1}^{N} A_i \, t_i \, e^{-\zeta \omega t_i} \cos\left(t_i \, \omega \sqrt{1-\zeta^2}\right) = 0 \tag{2.6}$$

Setting the partial derivative with respect to natural frequency equal to zero also sets the partial derivative with respect to damping ratio equal to zero, as shown in Singer [23]. These new constraints require the addition of a third impulse to our sequence; we have four equations, we need two unknown amplitudes and two unknown times. Given the constraint that t_3 be as small as possible, a closed form solution exists for the single mode three impulse shaper, making these sequences extremely simple to develop. The three impulse sequence, as we have shown, will also force the residual vibration to be low even if the system parameters are not precisely known.

The closed form solution for the single mode three impulse sequence is described by equations (2.7), where ω is the system natural frequency in radians, ζ is the damping ratio, and the "denom" term is used to normalize the impulse amplitudes, as dictated by equation (2.5). It is interesting to note that the impulse spacing depends on the values of ω and ζ, but the impulse amplitudes are only ζ dependent.

Single Mode Three Impulse Sequence: Closed Form Solution

$$K = e^{-\frac{\zeta\pi}{\sqrt{1-\zeta^2}}} \qquad \Delta T = \frac{\pi}{\omega\sqrt{1-\zeta^2}} \qquad denom = 1 + 2K + K^2$$

Impulse 1:   *time = 0*   *amplitude = 1 / denom*

Impulse 2:   *time = ΔT*   *amplitude = 2 K / denom*

Impulse 3:   *time = 2 ΔT*   *amplitude = K^2 / denom*   (2.7)

Figure 2.10 shows the three impulse equivalent of Figure 2.4. The parameter insensitivity is clearly demonstrated by Figure 2.11. Notice that the curve in Figure 2.11 is "flat" around the plant frequency, 1 Hz, and recall the steepness of the two impulse sequence insensitivity curve around the same frequency. Figure 2.11 illustrates that if our estimate of the natural frequency is slightly in error, the amplitude of the residual vibration will remain small.
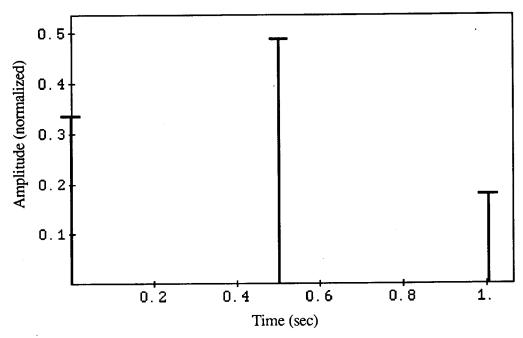
Figure 2.10: Single mode three impulse shaper.
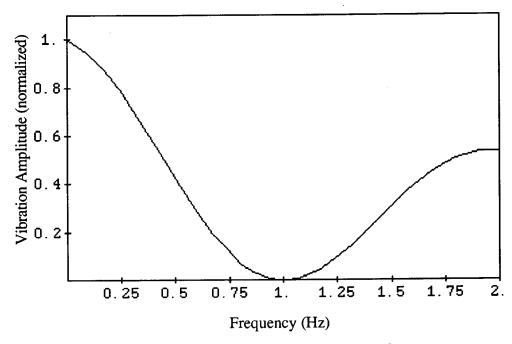
Figure 2.11: Insensitivity curve for three impulse sequence.

The penalty associated with the shaping method is a delay in the system rise time. By passing a step command through a three impulse sequence, for instance, the step will be delayed in reaching its full value for a time equal to the time of the final impulse in the shaper, a delay of about one period of the system's natural frequency. This delay will cause no instability problems, because the shaper remains outside of the control system, but it can act as an irritant in tele-operated systems. Minimizing the delay will be covered thoroughly in Chapter 3.

Note: the impulses in these sequences are constrained to having only positive amplitudes. If negative impulses are used, the time of the final impulse in the sequence can be decreased, but negative impulses tend to tax a system's actuators and introduce high stress levels. In the remaining chapters, all shapers will utilize impulses with positive amplitudes. For a full derivation of the above equations, see Singer [23].

2.4 Input Shaping Examples

In the following graphs, step inputs were fed through the shapers of Figures 2.4 and 2.10 and delivered to the $\omega = 1Hz$, $\zeta = 0.1$ system mentioned above. Figure 2.12 illustrates the effect of a two impulse shaper on the step input, and Figure 2.13 shows the corresponding response. Notice how shaping the step input results in a "staircased" command; we never deliver an impulse to the system. Figure 2.14 displays the step as modified by a three impulse shaper, with the responses shown in Figure 2.15. Figures 2.13 and 2.15 clearly illustrate the response delay caused by the shaper, and both graphs reveal the effectiveness of the shaper in reducing the system residual vibration.
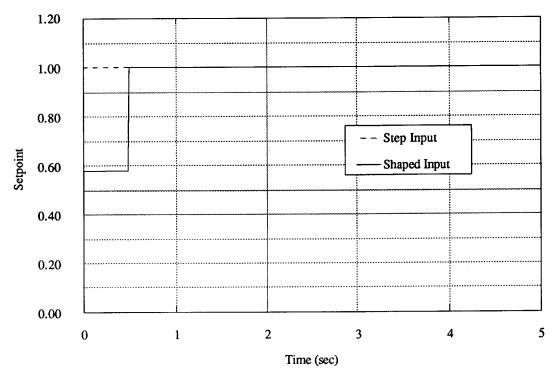
Figure 2.12: Step input as modified by two impulse shaper.
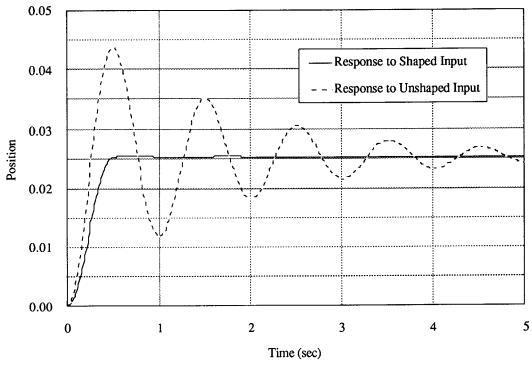
Figure 2.13: Responses to inputs of Figure 2.12.
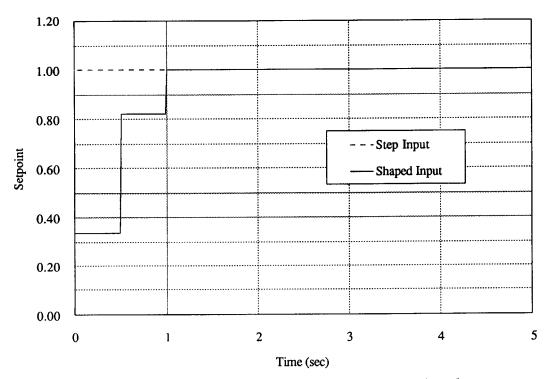
Figure 2.14: Step input as modified by three impulse shaper.
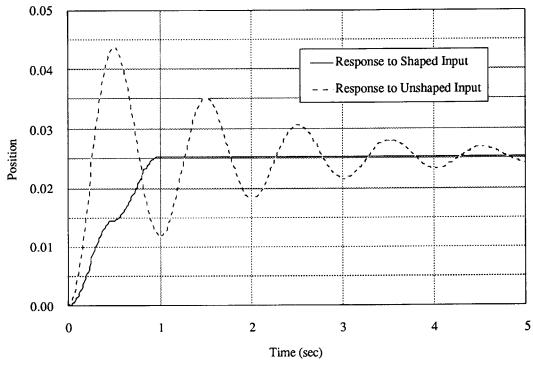
Figure 2.15: Responses to inputs of Figure 2.14.

2.5 Conclusion

In this chapter we developed a technique for using impulse sequences to suppress residual vibration in single mode flexible systems. Standard system inputs are convolved with an impulse sequence to generate an adjusted command that will produce vibration free motion. The "input shaping" method demonstrates significant insensitivity to variations in plant parameters and causes only slight delays in the system rise time, on the order of one period of the plant's natural frequency. The next step is to develop an approach for suppressing vibration in the large range of systems that oscillate at more than one frequency.

Chapter 3: Extending to **Multiple Mode** Problems

3.1 Introduction

In this chapter we move on to the more complicated problem of systems which vibrate at more than one frequency, or multiple mode systems. We showed in Chapter 2 that we can use impulse based input shaping to suppress vibration at a single mode, and in this chapter we'll examine two methods, convolution and direct solution, that can be used to extend the single mode shaping theory to deal with multiple mode problems.

3.2 Convolution

Convolution sequences utilize combinations of existing single mode shapers to suppress multiple mode vibration. The approach is to develop single mode shapers for all desired frequencies, and then convolve all of the shapers together into a single long impulse train. This is a relatively straightforward extension of the theory presented in Chapter 2.

Chapter 3: Extending to Multiple Mode Problems

3.1 Introduction

In this chapter we move on to the more complicated problem of systems which vibrate at more than one frequency, or multiple mode systems. We showed in Chapter 2 that we can use impulse based input shaping to suppress vibration at a single mode, and in this chapter we'll examine two methods, convolution and direct solution, that can be used to extend the single mode shaping theory to deal with multiple mode problems.

3.2 Convolution

Convolution sequences utilize combinations of existing single mode shapers to suppress multiple mode vibration. The approach is to develop single mode shapers for all desired frequencies, and then convolve all of the shapers together into a single long impulse train. This is a relatively straightforward extension of the theory presented in Chapter 2.

3.2.1 Building Convolved Sequences

In Chapter 2 we analyzed an impulse's tendency to cause vibration and then we positioned additional impulses to counteract that vibration. The resulting impulse train formed our input command shaper. In a multiple mode system, the first impulse in the train will cause vibration at more than one frequency. We can still place our additional impulses as in Chapter 2, and we will eliminate residual vibration at one frequency, but the system will be free to oscillate at frequencies not included in the shaper development. If we know what the other frequencies are, we can develop individual shapers to combat those modes, too, and we can convolve the sequences together so that they all work at the same time.

As an example, consider a system with modes at 1 Hz and 10 Hz, zero damping in both modes. Using equations (2.7), we can quickly generate individual impulse sequences tailored to eliminate vibration at each of the above frequencies. The single mode sequences for 1 Hz and 10 Hz are shown in Figures 3.1 and 3.2, respectively.
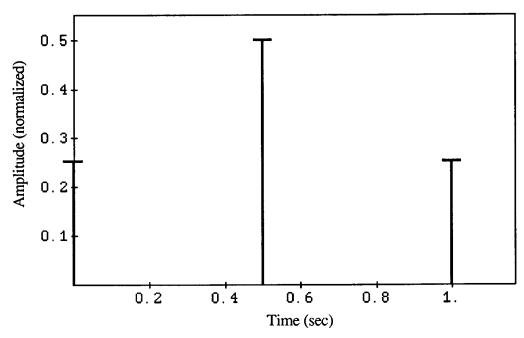
Figure 3.1: Three impulse sequence designed for 1 Hz.
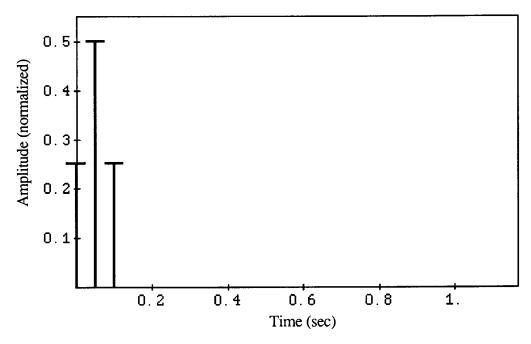
Figure 3.2: Three impulse sequence designed for 10 Hz.

To convolve these two sequences, we perform an outer sum of the impulse times in sequence #1 with the impulse times in sequence #2, and we perform an outer product of the impulse amplitudes. Another way to think of this is that every impulse in the 1 Hz sequence will touch off 10 Hz vibration. Each 1 Hz impulse is therefore immediately shaped for 10 Hz by the faster sequence. The result is a sequence with 3^m impulses, m being the number of modes suppressed by the shaper.

The product of convolving Figures 3.1 and 3.2 is shown in Figure 3.3, and the insensitivity curve for the convolved sequence is shown in Figure 3.4. We create a multiple mode insensitivity curve using the same method as in Chapter 2; we plot the residual vibration amplitude expression, equation (2.2), with ω varying as the independent variable.

The time of the final impulse in the convolved sequence of Figure 3.3, 1.1 seconds, is equivalent to the sum of the damped periods of the modes used in developing the sequences. This maximum time can be thought of as the shaper's "length," a term we'll use in the rest of the thesis. The length of a sequence is not the number of impulses in the sequence, it is the time of the sequence's final impulse, and concurrently, the delay in rise time caused by the shaper.
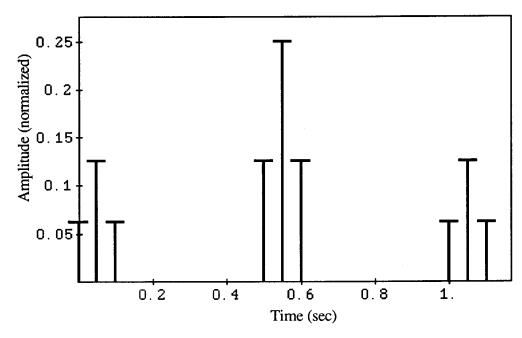
Figure 3.3: Convolved sequence designed for 1 Hz and 10 Hz, $\zeta = 0.0$.
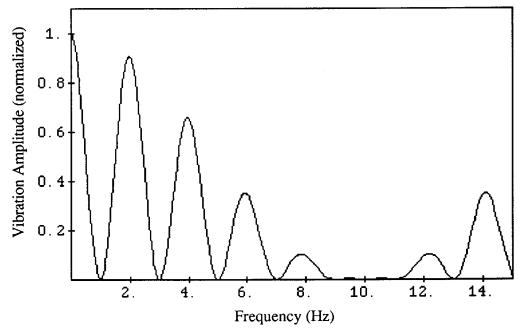
Figure 3.4: Insensitivity curve for convolved sequence of Figure 3.3.

As Figure 3.4 shows, vibration at 1 Hz and 10 Hz will have zero magnitude in the residual vibration. The insensitivity curve also illustrates an interesting feature of input shapers. If the system damping is zero, or very low, odd integer multiples of the frequencies used in developing the shaper will also be suppressed in the post-maneuver vibration. So in our example, modes at 3 Hz, 5 Hz, 7 Hz, etc. and 30 Hz, 50 Hz, etc. are also suppressed, incidentally providing excellent parameter insensitivity around 10 Hz.

This feature is caused by the periodic nature of the residual vibration, or any vibration, for that matter. When we defined our closed form solution for the single mode impulse shaper in equations (2.7), we specified that the time of the third impulse, t_3, be as small as possible. If we hadn't made this specification, we could have found an infinity of solutions with t_3 placed at increasing odd multiples of $2 \Delta T$ (the period of the mode's oscillation). So our 1 Hz shaper also cancels vibration at 3 Hz, by applying the impulse train over three cycles of the 3 Hz mode. The 1 Hz vibration is cancelled by applying the train over a single cycle, as originally intended. This suppression windfall disappears when the damping in the modes increases.

Figures 3.5 and 3.6 illustrate the effects of damping on the coincidental multiple mode suppression. Figure 3.5 shows a convolved sequence for 1 Hz and 10 Hz, with $\zeta = 0.1$ in both modes, and Figure 3.6 graphs the insensitivity curve for this sequence. We still cancel the necessary frequencies, at 1 Hz and 10 Hz, but the vibration amplitudes at the odd harmonic frequencies have drifted up from zero.

Recall from equations (2.7) that the presence of damping will affect the placement of the impulses in the shaper. When $\zeta = 0.0$ in all modes, the shaper for one frequency will also suppress vibration at all of the frequency's odd harmonics. When damping is introduced at the fundamental frequency, however, the spacing and magnitudes in the impulse train will be altered, and that shaper will not completely suppress vibration at the odd harmonics, where the damping might remain zero. This behaviour is characteristic of all shapers, including the direct solution sequences we'll examine in the later sections of this chapter.
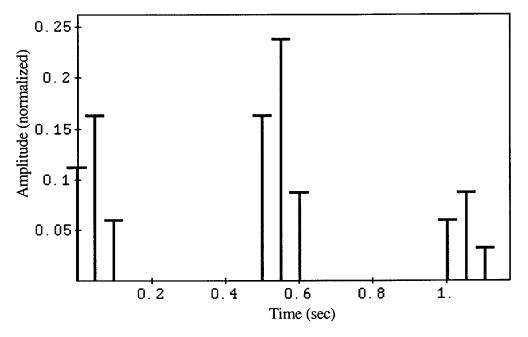
Figure 3.5: Convolved sequence for 1 Hz and 10 Hz, $\zeta = 0.1$.
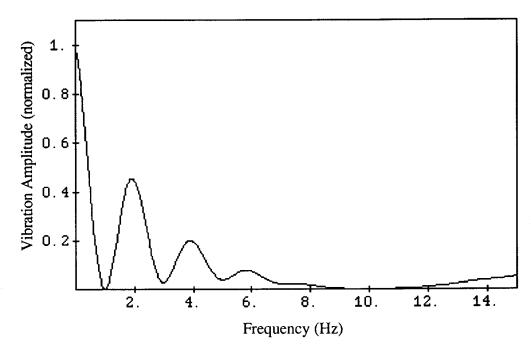
Figure 3.6: Insensitivity curve for sequence of Figure 3.5.

3.2.2 Implementing Convolved Sequences

Implementing any shaper in a digital control scheme involves modifying a setpoint by convolving it with the shaper's impulse sequence. In the computer code, the presence of one impulse requires one additional floating point multiplication operation and one additional floating point addition operation. In the previous section, we defined an expression for the number of impulses in a convolved sequence: # impulses = 3^m, where m equals the number of modes to be suppressed by the shaper. When the impulse population increases exponentially, using a convolved sequence to cancel even four or five modes will increase the computational burden and could force a decrease in servo rate.

Clever programming, however, can prevent the use of exponentially growing sequences. By utilizing successive ring buffers that handle the shaping for each mode, the setpoint can be passed along to each shaper in turn, allowing the controller to use only 3 * *m* impulses. This technique essentially convolves the single mode shapers on the fly, and can reduce the computer's workload.

3.2.3 Problems with Convolved Sequences

The convolved multiple mode sequences are easily generated, but their failings become clear when the cancellation of higher mode vibration is desired. We showed in section 3.2.2 that the number of impulses in a convolved shaper grows at a rate of three per mode, an improvement over the exponential growth expression of section 3.2.1, but still a potential cause of servo rate reduction.

Particularly painful is the application of convolved shapers to multiple mode problems involving several low frequencies. As an example, consider the Space Shuttle Remote Manipulator System (RMS) maneuvering a 7,000 pound satellite out of the Shuttle's payload bay. The first four frequencies of the manipulator arm / payload system, depending on orientation, can be estimated at about 0.20 Hz, 0.26 Hz, 0.45 Hz, and 0.59 Hz, with very low damping (around $\zeta = 0.01$). Recall that the length of a convolved shaper is equal to the sum of the damped periods of the shaper's modes. A shaper designed for cancellation of these four modes will be about 12.76 seconds long.

The Shuttle RMS is a tele-operated system, with an astronaut sending joint velocity commands to the arm to perform tasks. If we impinge a delay

of 12.76 seconds into that human controlled system, serious stability and hunting problems can result. Relatively precise maneuvers, such as docking the arm with a malfunctioning satellite, or positioning massive objects carefully within the payload bay, become extremely difficult. There are a variety of other space-borne and industrial systems with low frequencies that could experience similar problems caused by the convolved shapers.

So the major complaint of convolved shapers is their prohibitive length when used on multiple mode, low frequency systems. Another concern is their impulse population expression: we need three more impulses for every mode we add to our convolved shaper. An approach which addresses both of these problems is the direct solution technique, the subject of the next section.

3.3 Direct Solution

In the previous section we constructed convolved shapers by developing individual sequences for each mode, and then convolving them all together. In this section we'll pursue the idea of developing a single multiple mode shaper "all at once." This new approach will allow the generation of shapers which contain fewer impulses and cause smaller time delays than convolved sequences. We'll attempt to find such sequences through a "direct solution" of the original shaper constraint equations.

3.3.1 Re-formulating the Equations

To begin the search for direct solution sequences, we can rewrite the original single mode impulse sequence equations, (2.3) and (2.6), to include an arbitrary number of modes:

$$\sum_{i=1}^{N} A_i \, e^{-\zeta_j \omega_j t_i} \sin\left(t_i \, \omega_j \, \sqrt{1 - \zeta_j^2}\right) = 0$$

$$\sum_{i=1}^{N} A_i \, e^{-\zeta_j \omega_j t_i} \cos\left(t_i \, \omega_j \, \sqrt{1 - \zeta_j^2}\right) = 0 \tag{3.1a}$$

$$\sum_{i=1}^{N} A_i \, t_i \, e^{-\zeta_j \omega_j t_i} \sin\left(t_i \, \omega_j \, \sqrt{1 - \zeta_j^2}\right) = 0$$

$$\sum_{i=1}^{N} A_i \, t_i \, e^{-\zeta_j \omega_j t_i} \cos\left(t_i \, \omega_j \, \sqrt{1 - \zeta_j^2}\right) = 0 \tag{3.1b}$$

Repeating equations (3.1) for additional modes "*j*" generates a set of simultaneous non-linear impulse constraint equations. If we can find a solution to these equations, we'll derive a shaper that has the vibration reducing capabilities of the convolved shapers, and yet affords some advantages, mentioned in the next section.

3.3.2 Advantages of Direct Solution Sequences

An initial observation of equations (3.1) reveals that the solution of these equations will require only (2 * *m*) + 1 impulses, *m* being the number of cancelled modes. We can state that any sequence satisfying the above equations will automatically have fewer impulses than the equivalent

convolved sequence, which requires 3 * *m* impulses. The direct solution sequence's slowly increasing population leads to fewer impulses in higher mode shapers, reducing implementation time and computational burden.

We can also show that the direct solution sequences will always be shorter than the equivalent convolved sequences. On an intuitive level, this makes sense because we're defining one sequence from the outset which will suppress all desired modes, whereas with a convolved sequence, we must first develop a shaper for each mode and then stack all of the shapers together. We will not present a theoretical proof of the direct solution sequence's length savings, but we will clearly demonstrate this savings in examples throughout the remainder of the thesis. We will show that in some cases, a direct solution sequence can be as much as 20% shorter than the equivalent convolved sequence.

3.3.3 Problems with Direct Solution Sequences

The savings in length and impulse number that the direct solution sequences support are offset by an increase in sequence generation complexity. We mentioned that the single mode shaper equations have a closed form solution, equations (2.7). Convolved shapers achieve multiple mode vibration suppression through simple combinations of single mode sequences. The direct solution expressions, equations (3.1), have closed form solutions only for cases of two or three modes, and only if a certain structure is imposed on the resulting sequence. For higher mode systems, equations (3.1) require a strict set of constraints just to limit their infinite solution space, and no general closed form solution has been found.

3.4 Solving the Direct Solution Equations

The lack of a closed form solution makes matters difficult, but it is still certainly possible to develop a sequence that will satisfy equations (3.1). The key to solving the multiple mode equations thus far has been to employ a linear approximation. Equations (3.1) are non-linear only in terms of impulse time, so by constraining the impulse times to known values, we can linearize the equations and converge on a solution through optimization.

3.4.1 Formulating the Linear Problem

To linearize equations (3.1), we can pick a time for the sequence's final impulse, defining the sequence length, and then divide that length into a fine time mesh, as shown in Figure 3.7.
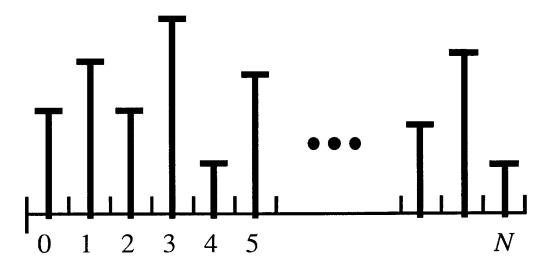
Figure 3.7: Defining the time mesh.

An impulse is placed at each time slot, with unknown amplitude but known time. With this knowledge of the impulse times, we can evaluate all of the non-linear terms in equations (3.1), forming constant coefficients for the impulse amplitudes. We now have a constant coefficient linear problem with impulse amplitudes as variables. We'll usually pick a fine time mesh to provide a good approximation to the actual solution, and we get a new variable with every time slot, so our linearized equations will be under-constrained. We can generate solutions to these equations through the use of linear programming with optimization.

To perform optimization, we need a matrix of constraints, and a cost function that we'll try to minimize. The constraints for our particular optimization problem are the equations we want to solve, the multiple mode equations (3.1). We'll add the normalization requirement of equation (2.5) to our constraint matrix to ensure that our final sequence will shape inputs properly.

Choosing a cost function can be a bit more complicated. What characteristics of our impulse sequence do we want to minimize? Several answers come to mind. We know that eventually we want to minimize the time of the final impulse, reducing the delay caused by the shaper. Also desirable is the minimization of the sequence's sensitivity to parameter variations. We can even think of using a cost function set equal to zero, so the optimization problem will be solved as soon as we find the first sequence that satisfies our constraint matrix.

It turns out that minimizing the sequence length makes the definition of the constraint matrix somewhat complex. Recall that we set up our time mesh by defining a length and then subdividing the length into slots. If we choose sequence length as a cost function, successfully reducing this cost function will mean that we'll have to re-define our time mesh, and therefore our constraint matrix will change throughout the problem. It is possible to make this work, but we'll show in section 3.4.2.b that we can reduce the sequence length using a simpler, yet perhaps less elegant method.

Minimizing the sequence's parameter insensitivity is also an attractive goal. We demonstrated in Chapter 2 that three-impulse sequences display greater parameter insensitivity than two-impulse sequences. Singer [23] showed that the parameter insensitivity increases still further when four-impulse sequences are used. Equations (3.1) are simply the three-impulse equations, written for multiple modes. We can increase the insensitivity of our three-impulse sequences by forcing them to behave more like four-impulse sequences.

To develop the equations defining the three-impulse sequence, we differentiated equations (2.3) with respect to natural frequency. Taking the second derivatives of equations (2.3) with respect to natural frequency forms the additional equations needed to develop a four-impulse sequence. We can use these "second derivative" equations, as written for multiple modes, in our cost function to impose some four-impulse sequence insensitivity on our direct solution sequence. The cost function formed by using these "second

derivative" equations is defined by equation (3.2). M is the number of modes in the problem, all other notation has already been defined.

$$\text{Cost} = \sum_{j=1}^{M} \left[\left(\sum_{i=1}^{N} A_i \, t_i^2 \, e^{-\zeta_j \omega_j t_i} \sin\left(t_i \, \omega_j \sqrt{1-\zeta_j^2}\right) \right)^2 + \left(\sum_{i=1}^{N} A_i \, t_i^2 \, e^{-\zeta_j \omega_j t_i} \cos\left(t_i \, \omega_j \sqrt{1-\zeta_j^2}\right) \right)^2 \right] \tag{3.2}$$

In the majority of the following optimization problems, the cost function will be defined by equation (3.2). Note that the summations over N are squared to eliminate any negative values. This action creates a non-linear cost function. One might argue that taking the absolute value of the summations over N would work just as well and produce a linear cost function, but the absolute value operator is also non-linear, around zero. Mixing linear constraints with non-linear cost functions is actually a common practice, and will not dramatically increase computation effort in the search for a solution to the linear problem.

We will also occasionally employ the "zero" cost function mentioned above. Using this cost function, our solution process stops when we find the first sequence that satisfies our constraints. This cost function will be useful when we want to get a quick, rough look at a non-optimal solution to the constraint equations.

With the constraint matrix formed by equations (3.1) and (2.5) and the cost function of equation (3.2), the linear problem becomes the classic mathematical exercise: "minimize the cost function subject to the stated constraints."

3.4.2 Solving the Linear Problem: GAMS

There are many techniques for solving linear optimization problems, and several existing routines and programs available that utilize those techniques. We chose the General Algebraic Modeling System (GAMS) [5], a comprehensive programming language that eases the process of problem definition, communication with solver routines, and solution presentation.

3.4.2.a Running GAMS

The GAMS system reads an input file designating problem data and various options, compiles the input file, sends the file to a specified solver routine, and writes an output file containing the results of the optimization. The output file details any solution problems such as infeasibility, local minima, etc. The actual technique employed by GAMS to perform linear optimization is a version of the popular primal simplex method.

To apply GAMS to our problem, we start by writing an input file that contains the modal information for our vibratory system. We add to the file our choice for t_N, the time of the final impulse, and Δt, the time separating the slots in our time mesh. This sets up an array of times containing $t_N / \Delta t$ entries. We define our array of variables, the impulse amplitudes A_i, to be the same length as the time array. We then specify equations (2.5) and (3.1) as

our constraints and equation (3.2) as our cost function. Alternatively we can specify a zero cost function. GAMS will automatically use the values in our time array to linearize the amplitude coefficients in the constraint matrix and the cost function. Finally we specify a solver routine and define what data we want to see in the output file. The next step is to simply call GAMS and specify the input file.

To continue the example we used for the convolved sequences in section 3.2.1, let's write an input file with the modal information: $\omega_1 = 1$ Hz, $\omega_2 = 10$ Hz, and $\zeta = 0.0$ in both modes. Let t_N equal the length of the old convolved sequence, 1.1 seconds, set $\Delta t = 0.005$ seconds, and use the zero cost function. We know that we'll be able to find a shorter direct solution sequence than 1.1 seconds, but let's see what happens with these parameters just for comparison. Using our input file, GAMS returns the sequence shown in Figure 3.8, with the parameter insensitivity shown in Figure 3.9.
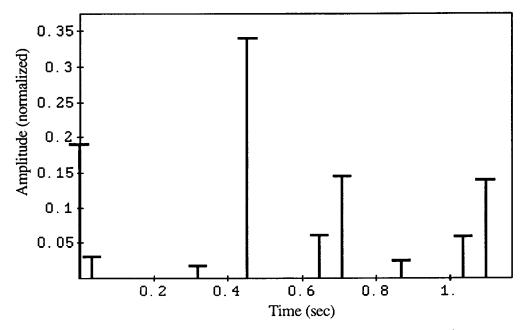
Figure 3.8:   Direct solution sequence for 1 Hz and 10 Hz, $\zeta = 0.0$, zero cost function.
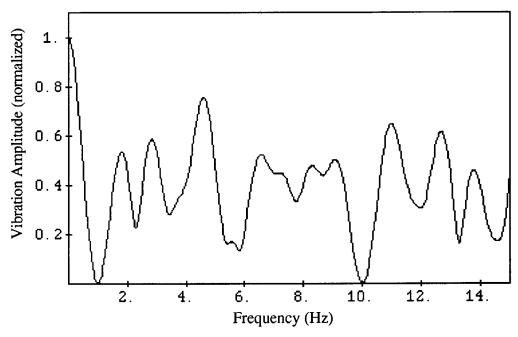
Figure 3.9: Insensitivity curve for sequence of Figure 3.8.

Our sequence certainly works; the insensitivity curve reveals zero vibration amplitudes at 1 Hz and 10 Hz, but a few interesting points should be mentioned here. First, recall that we specified a variable array of the amplitudes, A_i, with the same length as our time array, or t_N / Δt slots. As part of our linearization, we placed an impulse at each time slot, meaning our variable array must have 1.1 / 0.005 or 220 entries. The impulse sequence of Figure 3.8, therefore, actually has 220 impulses, but only 9 are shown. The remaining impulses are still there, but their amplitudes are zero.

The zero amplitude impulses result from the primal simplex method that GAMS uses to solve linear problems. If M is the number of modes to be cancelled, the GAMS constraint matrix consists of:

$$\text{rows: } r = (4 * M) + 1$$

$$\text{columns: } c = t_N / \Delta t$$

The primal simplex method dictates that at least $(c - r)$ variables will equal zero. The number of "zeroed" variables will be adjusted if the cost function is non-linear, because GAMS adds or removes "supervariables" to its constraint matrix to aid in the solution of non-linear problems. The main point is that the returned sequence will have about r impulses, not t_N / Δt. In our case, r = 9, and our returned sequence does, in fact, have 9 impulses. This number will be shifted slightly when we use the "second derivative" cost function.

Using the "second derivative" cost function will improve the parameter insensitivity of Figure 3.9. Figure 3.10 displays the sequence output from GAMS when the zero cost function is traded for the second derivative expression, with the accompanying insensitivity curve shown in Figure 3.11.
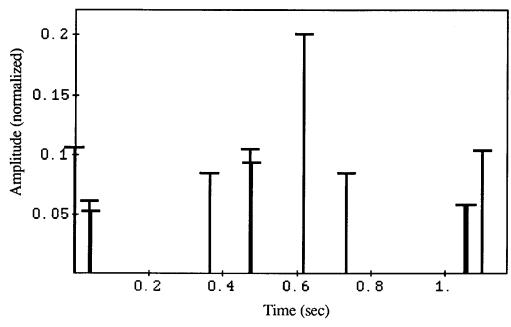
Figure 3.10: Direct solution sequence for 1 Hz and 10 Hz, $\zeta = 0.0$, "second derivative" cost function.
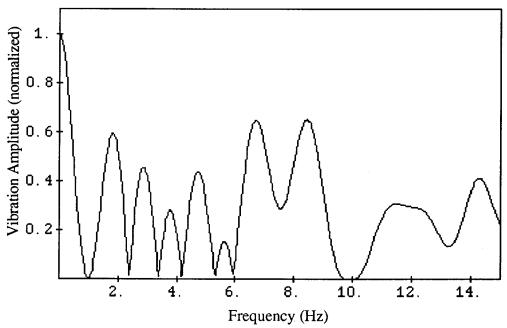
Figure 3.11: Insensitivity curve for the sequence of Figure 3.10.

In a comparison between Figure 3.11 and Figure 3.9, we can see that the insensitivity curve in Figure 3.11 is flatter around 10 Hz, reflecting the slight, but beneficial effects of using the second derivative cost function.

An interesting point brought up by Figure 3.10 is the "pairing" of impulses in GAMS sequences. In three cases in Figure 3.10, impulses were placed in adjacent time slots. This phenomena can be thought of as GAMS trying to place a single impulse in an imaginary time slot that exists between two defined slots. GAMS just splits the original impulse and places the two new impulses in the nearest time mesh locations available.

3.4.2.b Finding the Shortest Solution

Now that we've shown that linearizing the constraint equations can actually produce workable impulse sequences, the next step is to find the sequence that yields the greatest time savings. This can be accomplished by running GAMS multiple times, systematically reducing the final impulse time, t_N. GAMS will specify in the output file if the solution to the problem is infeasible, meaning that the sequence length has been reduced too far. A simple binary search algorithm was utilized to minimize the time spent searching through the solution space. This technique ensured that the final GAMS output was the shortest possible approximation.

While searching for the shortest sequence, the quickest method is to use the zero cost function. This expression is computationally less difficult for GAMS, and many iterations can be conducted quickly. Once the minimum time has been discovered, the problem can be run again using the

second derivative cost function to obtain the most insensitive sequence. The time mesh length may have to be increased by one or two slots to accomodate this new sequence.

The final, shortest, and most insensitive GAMS sequence for the 1 Hz & 10 Hz problem is shown in Figure 3.12. The insensitivity curve for this sequence is the subject of Figure 3.13. It is interesting to compare Figure 3.13 to Figure 3.4, the insensitivity curve of the equivalent convolved sequence. Our GAMS sequence does not perfectly suppress vibrations at the odd harmonics of the included frequencies. This mild performance degradation can be attributed to the the digitization of the impulse times in the GAMS sequence; an exact sequence can place impulses in continuous, not digital, locations. We should recover the odd harmonic suppression if we are able to convert the GAMS sequence into an exact, continuous sequence.

The length of the sequence in Figure 3.12 is 1.055 seconds, a savings over the convolved sequence length of about 4%. This length improvement is minor; greater time savings will be realized when we apply the direct solution method to higher mode, low frequency systems.

As a stronger example of the direct solution sequence's time savings potential, consider the Space Shuttle RMS case mentioned in section 3.2.3. A convolved shaper designed for the first four frequencies of that system, 0.20 Hz, 0.26 Hz, 0.45 Hz, and 0.59 Hz, zero damping in all modes, will have a length of 12.76 seconds. The equivalent direct solution sequence is shown in Figure 3.14. This sequence has a length of 10.70 seconds, a savings of 16%.
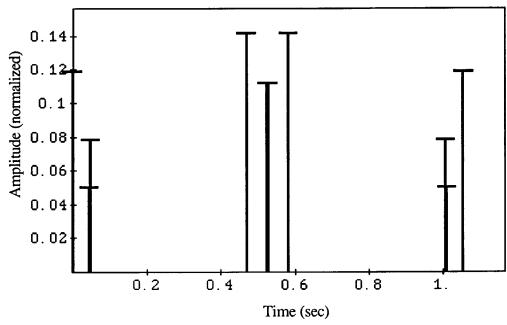
Figure 3.12: Minimal time impulse sequence for 1 Hz and 10 Hz, zero damping, second derivative cost function.
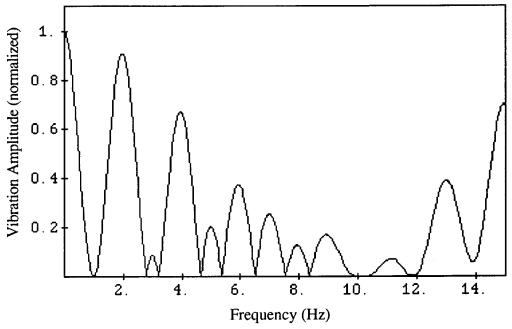
Figure 3.13: Insensitivity curve for sequence of Figure 3.12.
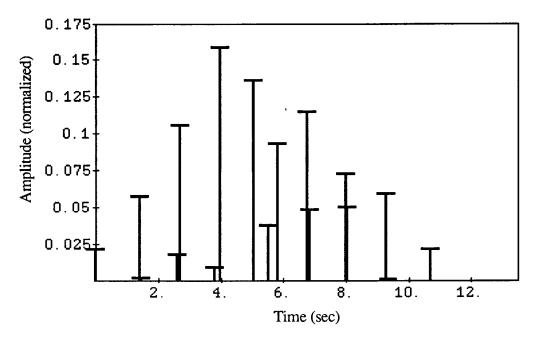
Figure 3.14: Direct solution GAMS sequence for modes at 0.20 Hz, 0.26 Hz, 0.45 Hz, and 0.59 Hz, zero damping.

3.4.2.c Using the GAMS Sequence

When any impulse shaper is implemented in a digitally controlled system, the impulse times must be rounded to their nearest value in servo counts. This rounding, depending on the servo rate, can cause the shaper to lose some of its effectiveness, as mentioned in Singer [23]. These errors are minimized if the impulse times are already synchronized with the controller servo rate. for instance, in the above GAMS sequences, we used a Δt of 0.005 seconds. If the system's controller were operating at 200 Hz, our impulse sequence would fit perfectly into the servo count structure.

In many cases, therefore, it is desirable to simply set the GAMS time mesh to correspond with the controller servo rate, and employ the optimized GAMS sequence as the input shaper. Occasionally, the number of impulses in the GAMS output will be high enough to slow the controller down and cause missed sampling problems. In these cases, the GAMS output can be adjusted slightly to minimize the impulses in the sequence. Recall from section 3.4.2.a that GAMS often returns sequences with "paired" impulses: impulses in neighboring time slots. We can "squash" two neighbors together by summing their magnitudes and pushing the new single impulses into one of the slots previously occupied by one of the neighbors.

Figure 3.15 shows the sequence of Figure 3.12 after it has been "squashed." The time mesh spacing is preserved; the neighboring impulses are simply combined. Figure 3.16 displays the insensitivity degradation caused by combining the neighboring impulses. A comparison to Figure 3.13 reveals this degradation as slight. We reduced the number of impulses from 11 to 7, and saving the computational burden of four impulses can possibly allow a digital controller to operate properly and not miss data samples.
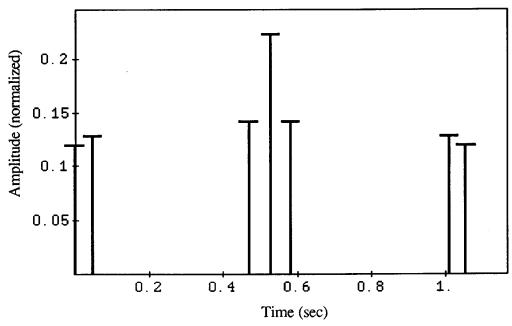
Figure 3.15: GAMS sequence of Figure 3.12 with neighboring impulses combined, a decrease of four impulses.
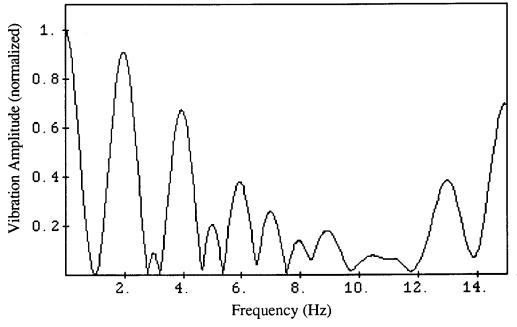
Figure 3.16: Insensitivity curve for GAMS sequence of Figure 3.15

3.4.3 Finding Exact Solutions

Now that we've developed short, insensitive sequences through a linear approximation, the next step is to try to convert those approximate sequences into exact, continuous-time sequences. We will accomplish this by considering the best GAMS results as initial guesses. These guesses will be used by a non-linear equation solver to attempt to find an exact sequence.

3.4.3.a Interpreting the Linear Problem Results

The GAMS output has to be interpreted to be able to function as an initial guess of the exact sequence. Recall that the exact solution to equations (3.1) requires only (2 * m) + 1 impulses, m being the number of modes. The raw GAMS sequences are returned with about 4 * m impulses, so we need some sort of algorithm that will reduce the number of the GAMS impulses while retaining the basic sequence structure.

The interpretation algorithm used on the GAMS sequence has two phases: first, neighboring impulses are "squashed" as mentioned in section 3.4.2.c, but instead of retaining the time mesh of the GAMS sequence, the new impulse is placed at a time halfway between the original neighbors. This "squashing" may or may not reduce the number of impulses far enough, it usually will not and phase two must be used.

The second phase in the algorithm searches for the closest non-adjacent neighbors, i.e. those impulses not squashed by phase one but still very close to each other. These impulse pairs are replaced, one pair at a time, by a single impulse. The amplitude of the new impulse equals the sum of the

original impulses, and the new impulse time is calculated as a weighted average of the times of the original neighbors. The weighting function for this average is simply the amplitude of each original impulse. A check is then made to determine if the sequence has the proper number of impulses, and if not, phase two is repeated.

This set of techniques will yield a sequence whose number of impulses matches that required by the non-linear multiple mode equations (3.1). As an example, consider the four-mode sequence shown in Figure 3.14. The resulting sequence, after employing the two-phase interpretation algorithm mentioned above, is shown in Figure 3.17.
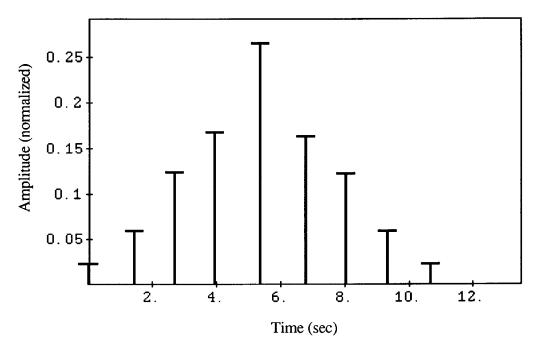
Figure 3.17: Interpreted sequence from Figure 3.14, to be used as initial guesses for non-linear equation solver.

3.4.3.b Using the Linear Results in a Non-Linear Solver

The availability of non-linear equation solving routines rivals that of linear approximation routines, with a variety of approaches utilizing several languages and formats. We chose Mathematica™ [33] as a non-linear equation solver, mainly because of its additional potential as a programming language that could envelop the entire computational side of the impulse sequence generation. A long Mathematica program can be used to interface with the user, write GAMS input files, analyze GAMS output files, and to attempt exact solutions.

To use the GAMS sequence as an initial guess for the Mathematica solver, we can hold constant the time and amplitude of the first ($i = 1$) impulse in equations (3.1), matching the first impulse from the interpreted GAMS output. The remaining times and amplitudes are allowed to vary, with initial guesses of their values provided by the reduced GAMS sequence.

Continuing our use of the Space Shuttle four-mode problem, we can feed the initial guesses of Figure 3.17 into the Mathematica solver routine. The resulting exact direct solution sequence is shown in Figure 3.18, with the corresponding insensitivity curve of Figure 3.19. The slight differences between the final sequence of Figure 3.18 and the guessed sequence of Figure 3.17 reveal that our original guesses were quite close to the actual solution. Our exact solution has virtually the same length as the GAMS sequence of Figure 3.14, retaining the length savings of 16% over the equivalent convolved sequence.
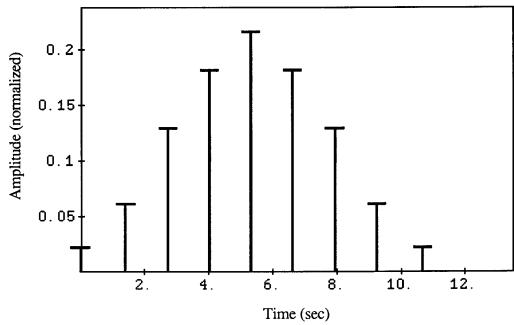
Figure 3.18: Exact direct solution sequence for problem of Figure 3.14; modes at 0.20, 0.26, 0.45, and 0.59 Hz, $\zeta = 0.0$.
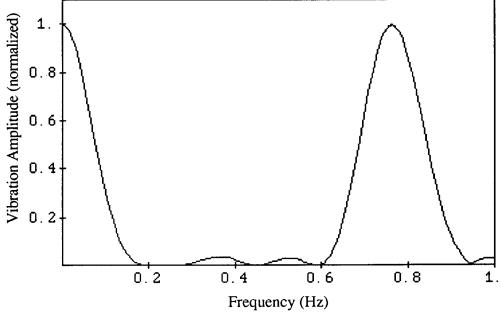
Figure 3.19: Insensitivity curve for sequence of Figure 3.18.

Mathematica employs a Newtonian gradient search algorithm to arrive at its solutions. This method works quite well, as long as our guesses are "sufficiently" close to the optimal solutions, meaning that our initial guesses reside on the same gradient as the final solution. As the non-linear equations are continuous and differentiable, adequate gradients are readily available, and points of singularity are usually easy to avoid.

GAMS can also be employed as a non-linear equation solver, and we can extend our Mathematica program to write GAMS input files designed for the exact solution of the constraint equations. Using GAMS instead of Mathematica produced little or no difference in ease of solution.

3.4.3.c Exact Solution Difficulties

The GAMS / Mathematica routine can successfully derive exact direct solution sequences for a variety of multiple mode problems. GAMS alone will produce sequences that, with minor modifications, can be perfectly meshed with digitally implemented controllers. The exact solution routine, however, is not entirely robust, and the two problems that regularly cause failure are modal spacing and number of modes.

The modal spacing characteristics that plague the generation of exact solutions are difficult to describe, and at the time of this writing, not entirely understood. Back in Section 3.2.1, we examined the input shaper's interesting ability of eliminating vibration not only at a designated frequency, but also for odd harmonics of that frequency. We explained that this phenomenon is a function of the vibration's periodic nature. An infinite number of shapers

can be found that will suppress vibration at a particular mode; we happen to be interested in the shortest possible sequence, that sequence with the smallest final impulse time.

The modal spacing problems that prevent the exact solutions from converging seem to be linked to the odd-harmonic suppression "feature." In a two mode case, for example, if the second mode is greater than three times the first, the GAMS / Mathematica routine will not converge on an exact solution. This difficulty appears to be governed by a ratio, R, that we have established through further experimentation and analysis. R is the theoretical length of the exact direct solution sequence designed for two modes:

$$R = 4 \, / \, (f_1 + f_2)$$

f_1 : frequency of mode 1

f_2 : frequency of mode 2

If f_2 is less than or equal to three times f_1, R will be greater than or equal to $1 \, / \, f_1$, the period of the lowest mode, and everything works fine. If f_2 is greater than three times f_1, however, R will be less than $1 \, / \, f_1$. When R drops below the first mode's period, our routine won't find a solution – to suppress a particular mode, our entire shaper must be at least as long as that mode's period. The solution algorithm will try to use negative impulses to force a shorter sequence, an action which we'd rather avoid for the reasons mentioned in Chapter 2.

The coupling between solution feasibility and modal spacing becomes unclear when additional modes are added to the system. We have so far been

unable to identify an R parameter for systems involving more than two modes. The clearest statement we can make at this point is that we realize modal spacing is a problem, and we have some insight into the nature of that problem, but we cannot predict exactly which spacings will cause the solution algorithm to fail for systems involving more than two modes. We are conducting further investigations into the difficulties caused by modal spacing, and hopefully future work will shed some more light on this area.

The solution algorithm's difficulty with high modal numbers is a bit more straightforward. To date, the routine has successfully generated exact impulse sequences for properly spaced groups of up to five modes. Linear GAMS sequences can be easily generated for many more modes, but the exact solutions will not converge. The complexity of the simultaneous solution process increases rapidly with additional modes, comparable to the rapidly increasing difficulty of finding roots of third or greater order polynomials. As we push past fifth-order modal problems, our current approach is not equal to the mathematically burdensome task of finding exact solutions.

One might think that if the GAMS time step were reduced far enough, the linear sequence would virtually emulate the exact sequence. This is sadly not the case, for the problem as posed to GAMS does not strictly bound the number of impulses; GAMS will use roughly 4 * M impulses, where the exact solution is constrained to using (2 * M) + 1 impulses, M being the number of modes. For finer and finer time meshes, therefore, GAMS will just continue to use its allotted number of impulses, and the linear sequence will not become a better approximation of the exact sequence.

Generating exact multiple mode shaping sequences can thus become a difficult, or even impossible computational problem. Fortunately, in many cases the shapers will be employed in digital control algorithms, where variations of the linear sequences, such as those output from GAMS, are more appropriate. Linear sequences can be generated for any number of modes, covering any modal spacing, provided the linear system time step can be reduced sufficiently. In situations where finding an exact solution is difficult or impossible, a linear sequence can be used to save a great deal of time and effort, at a cost of a mildly increased controller burden.

3.4.3.d New Approaches

Current, parallel research is aimed at finding closed form solutions for multiple mode problems. We can recognize that many of the exact sequences we found in this chapter, either by convolution or direct solution, have been symmetric, with an equal number of impulses ranged around a central pivot impulse. We can rewrite the impulse constraint equations to "build in" a symmetric grouping, which tends to reduce the complexity of the equations and can allow a direct solution to be expressed in a closed form.

As of this writing, the "symmetric re-formulation" technique has produced exact solutions for two- and three-mode problems, but modal spacing still causes difficulties and occasionally the new sequences are not as short as the equivalent GAMS linear solutions. We must also realize that some of the exact solutions, depending again on modal size and spacing, are not necessarily symmetric.

Other current work proposes to constrain the number of impulses that GAMS can use, bringing the raw GAMS output closer to the final exact sequence. This would improve the fidelity of the linear sequences and decrease the amount of interpretation required to use the GAMS sequences as guesses for non-linear equation solvers. Both of these new approaches show promise, and sections of future theses will be devoted to the increased capability of finding exact solutions.

3.5 Conclusion

In this chapter we have examined two methods for applying impulse sequence shapers to multiple mode problems: convolution and direct solution. Convolution involves combining several single mode shapers into an impulse train that will perform vibration suppression of multiple modes.

To use the direct solution method, we re-formulate the single mode shaper equations to include several modes, and solve those equations simultaneously, generating a single sequence that matches the performance of the convolved shaper. The direct solution sequences use fewer impulses than the convolved sequences, reducing the computational burden of implementing the input shaper. The new sequences are also shorter than convolved sequences, resulting in smaller system response delays. These advantages are offset by an increase in sequence generation complexity.

Chapter 4: The MACE Test Article

4.1 Introduction

In the previous chapters we examined input command shaping as a method to suppress residual vibration in flexible systems. Now that we have a clear, theoretical framework for the development and implementation of direct solution impulse sequence shapers, we need a physical system upon which we can test, and hopefully validate, the shaping technique. In this chapter we will discuss the development of an experimental flexible structure, the Middeck Active Control Experiment (MACE). We will use this experiment as a test article in later chapters.

4.2 Motivation

Future plans in the U.S. Space Program call for the development of large, earth-orbiting satellites that will monitor various global trends relating to weather and long term temperature fluctuations. These proposed satellites will be "multi-body" platforms, meaning that each spacecraft will carry an

array of independent sensors, antennae, and instrumentation, all mounted on a common structure. Two space platform programs in various stages of development are the Earth Observing System (EOS) [19], and the Geostationary Platform (GEOS) [11]. Figure 4.1 is a schematic of the early EOS-A configuration, and Figure 4.2 is a drawing of the GEOS platform.

9-4072
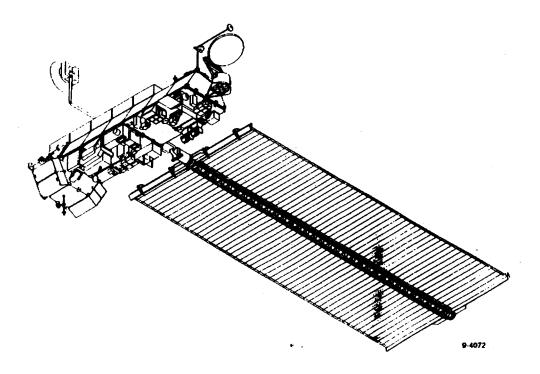
Figure 4.1: EOS-A Spacecraft Configuration
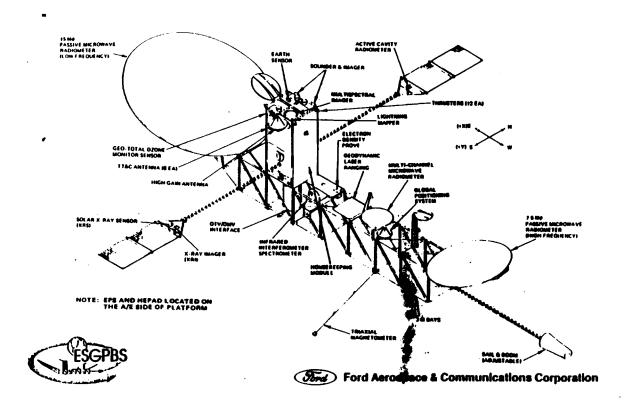
Figure 4.2: GEOS Platform

One problem common to such space systems is structural vibration. The several sensors and instruments all rest on a common platform. To facilitate launching, the platform is designed to be as lightweight as possible, while still maintaining strength and rigidity. These structures will not pass for bedrock, however, and a moving antenna will transmit vibration through the structure, affecting all of the other on-board devices. The problems caused by several scanning antennae and optical systems, all moving at once, are easily imagined.

There are several approaches for dealing with this system vibration. Materials research continues to produce lighter, stronger structures, and

heavy-lift rockets are a viable, if expensive, option for orbiting stiffer platforms. The payloads can also be isolated from the common bus, utilizing magnetic bearing technology to filter out small vibrations that can disrupt radar imagery or communications. This concept can add weight, because each payload requires an individual isolation mechanism.

A different method is to coordinate the control systems of all the slewing objects on the satellite, and coupled with vibration sensing and accurate system models, this global control network can act to circumvent vibration problems. Each payload has full knowledge of the vibration in the system, and can undertake measures to counteract that vibration, under the instruction of the "overlord" controller.

Another technique, pursued in the MIT Space Engineering Research Center (SERC) by Crawley and de Luis [8], among others, is the use of active elements embedded within the platform. The platform designs are often truss structures, formed of many linking struts and nodes. An active strut is one which can shrink or expand to exert forces on the surrounding truss. When the active element is linked to vibration sensors and a controller, it can be used to dampen the vibration in the truss. These "smart" struts can work in various ways; a popular approach is to use piezoelectric material that will expand or contract under an applied voltage.

The input shaping techniques can be used to supplement any of the above ideas. By shaping slewing commands, vibration can be prevented from exciting the structure in the first place, easing the load on the isolation platforms, global controllers, or active strut systems.

To experiment with micro-gravity spacecraft vibration suppression strategies, particularly the active piezoelectric strut options, personnel at SERC have been developing a joint MIT/NASA experimental test article, the Mid-deck Active Control Experiment (MACE) [9,18]. MACE will be tested on the Space Shuttle's mid-deck, a shirtsleeve environment less hazardous than the Shuttle's payload bay.

4.3 The MACE Project

MACE is composed of three major subsystems: structure, attitude control unit, and pointing systems. The structure is a flexible segmented tube, or bus, about 1.5 meters long. The bus is divided into four polycarbonate tubular struts, with each strut threaded into two cubic aluminum nodes. Each face of any particular node has fittings for attaching a strut or other instruments.

The attitude control unit is a set of three-axis torque wheels mounted to the node at the center of the bus. A three-axis rate gyroscope package is attached on the opposite side of the center node to provide inertial attitude information. When the gyros indicate that the structure is rotating, the wheels will be accelerated or decelerated, preventing the structure from spinning out of control during experiments.

High precision two-axis pointing mechanisms are mounted on either end of the bus. These gimbals will slew payloads to conduct pointing and tracking tests. Each payload contains a two-axis rate gyro package, giving the gimbal mechanism access to the payload's inertial position and allowing the payloads to be pointed independent of bus orientation. Each gimbal can point

its payload roughly within a 120° included cone, that is, each payload can move ± 60° away from vertical, in two directions.

To conduct active control experiments, a plain strut can be easily swapped for a segment fitted with piezoelectric actuators. The replacement strut can have various cross sections, giving freedom to the type of active control desired for a particular experiment. The total weight of MACE is projected at about 60 pounds, and a simple system schematic is shown in Figure 4.3.
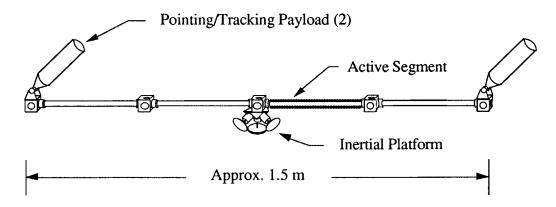
Figure 4.3: The MACE flexible test article.

The main research intent behind all of this hardware is to be able to point or track with one of the payloads while the other payload is also moving. Any slewing gimbal will transmit vibration through the bus, affecting the performance of the opposite payload. With this test article, we can examine various methods for suppressing the bus vibration in order to increase the overall pointing and tracking performance of the individual gimbals.

The MACE project is divided into two separate physical systems, a Ground Test Article (GTA), to remain at MIT, and a Flight Test Article (FTA), scheduled for a Space Shuttle mid-deck flight in 1993. The following two sections, 4.3.1 and 4.3.2, provide a rough description of the two test articles, for general reference. The writer has also performed a significant amount of detailed survey and design work for both test articles; refer to the Appendix for a concise description of this activity.

4.3.1 Ground Test Article

The Ground Test Article will remain in the MIT SERC lab throughout the life of the MACE experiment. To emulate a flexible spacecraft, the GTA will be suspended in a tall bay by an active suspension system. The active suspension allows the experiment to swing laterally, in two directions, and also employs a sophisticated pneumatic weight matching system, providing about four inches of essentially zero friction vertical travel. MACE will be hung with the heavy payloads pointing downward, opposite from the schematic in Figure 4.3, to keep the system naturally stable. Suspension wires will be attached at three points; to the undersides of the central and endpoint nodes.

The gimbals on the GTA are designed to have pointing accuracy of about one arc minute, or $0.01667°$. Some of the EOS and GEOS pointing mechanisms will require accuracy of about one arc second, or $0.00028°$, but the relatively low accuracy GTA gimbals should provide insight into much of the vibration caused performance degradation, and gimbal cost rises

exponentially with accuracy. To keep the GTA practically affordable, the one arc minute gimbals have been deemed adequate.

Electric signals for the sensors and actuators will be sent to the bus via an umbilical cord, attached to the center node. All of the subsystems on the structure will be managed by a real-time computer, which will implement various control strategies.

Through a thorough study of the GTA, we can analyze many of the vibration problems facing high performance components on actual space systems. The tests conducted on the GTA will hopefully lead to possible methods for dealing with those problems. The ground article will also be useful as a support mechanism for the FTA while tests are being conducted on orbit. Any major unanticipated problems experienced by the FTA could be quickly communicated to SERC and simulated with the ground article. A rapid study might generate corrections that could be relayed to the Shuttle, allowing experiments to continue.

As of this writing, the MACE GTA is still under construction. The bus and attitude control systems are complete, and one of the gimbals has been fabricated. The second gimbal may be designed simply as a disturbance source, requiring less engineering effort than the first gimbal, but this is still under consideration. The completed article should be available for testing in the summer of 1991.

4.3.2 Flight Test Article

The MACE Flight Test Article is needed because some of the problems facing orbiting space structures simply cannot be simulated on Earth. We can use our complicated suspension system to provide relatively frictionless travel in three dimensions, and we could even suspend the article in a vacuum chamber to eliminate damping effects caused by the atmosphere. The one inescapable phenomenon on Earth, however, is gravity. Even when the GTA is suspended actively, gravity will affect the vibratory behaviour of the structure, compressing or extending the struts and gimbal mounts and shifting the modal characteristics of torsional or in-plane bending.

To provide a full examination of vibration problems, therefore, we need to examine a structure residing in a micro-gravity field. In 1993, when the FTA is slated for a Space Shuttle flight, we will hopefully gain comprehensive insight into the complex behaviour of multiple mode vibration of flexible space structures. We will be able to apply controllers designed on Earth, and detect flaws in our modelling of micro- and one-G effects. Proposed methods of vibration reduction will also be tested, and validated, if our simulations were accurate.

The FTA differs from the GTA in several respects. First, where the GTA could be assembled at leisure and left hanging in the laboratory, perhaps indefinitely, the FTA must be launched disassembled and stowed in a mid-deck compartment. The bus segments will have to be hand-assembled by a weightless astronaut once the Shuttle reaches orbit, and disassembled for re-entry. To ease this process, the struts and nodes are designed for easy

coupling; the connection is secured by a simple threaded collar. The gimbals and attitude control system could be pre-fastened to their respective nodes before launch. The FTA struts will also provide a simultaneous mechanical and electrical connection, with about 50 conductors running through the center of the strut. The wiring on the GTA, however, will probably be external to the struts.

The gimbals on the FTA will be designed to be more accurate than those on the GTA, to provide a closer approximation to the systems proposed for the larger satellites. Proposed FTA gimbal designs call for accuracy of around 5 arc seconds. The flight gimbal actuators will be less powerful than the ground systems, simply because the FTA gimbals will not need to overcome gravity in order to slew their payloads.

Another significant difference between the two articles is in their controlling computers. The real-time computer governing the control of the GTA has a relatively straightforward architecture. In space, unfortunately, there is no atmosphere to provide shielding from cosmic rays and other types of radiation. The FTA computer, therefore, will require special shielding and perhaps a modified architecture, and all of the FTA wiring will have to be carefully shielded to avoid problems caused by random background radiation.

The increased gimbal accuracy of the FTA, combined with the special electronics concerns caused by spaceflight, boost the cost of the flight hardware to several times that of the GTA. The rigorous safety guidelines imposed by NASA only add to the difficulty of producing an effective flight experiment. All of the trouble is necessary, though, to provide an accurate examination of

the effects of micro-gravity environments on the vibration characteristics of flexible structures, and on the methods utilized to suppress that vibration.

4.4 Conclusion

The Mid-deck Active Control Experiment (MACE) is a flexible structure under development at the MIT Space Engineering Research Center (SERC). The research for MACE is prompted by NASA proposals for large, earth orbiting satellites that will be used in the future to monitor long term environmental trends such as weather patterns and global warming.

The MACE structure is composed of a 1.5 meter segmented tubular bus with articulated pointing payloads located at the bus endpoints and an attitude control system situated at the bus midpoint. The experiment is divided into two phases: the first phase utilizes a Ground Test Article (GTA) which will be actively suspended in a high bay at MIT to emulate a free-floating spacecraft. The second phase involves a Flight Test Article (FTA), which will extend the study of the GTA during a Space Shuttle mid-deck flight in 1993.

By studying the vibration characteristics of this flexible structure, researchers can gain insight into methods for suppressing vibration, and thus improving the performance of tracking and pointing devices flown on space-borne flexible systems. The input shaping techniques developed in previous chapters are possible methods for minimizing the system vibration. Some of the input shaping techniques have been tested on physical structures in the past, but the MACE test article is a prime candidate for additional validation

of the shaping methods, and for initial tests with the direct solution sequence of Chapter 3.

Chapter 5: Simulations and Results

5.1 Introduction

In the previous chapters we have examined methods for developing multiple mode input command shapers, and we have considered a suitable test article (MACE) for validation of the new shaping techniques. The remaining task is to conduct experiments using the test article and analyze the results. While the MACE structure has been under construction, personnel at the MIT Space Engineering Research Center (SERC) have developed several computer models of the test article. In this chapter we will apply various input shapers in tests of linear and non-linear computer simulations of MACE.

Note: the development of these models is not examined thoroughly here; the writer employed existing SERC simulations, and performed no significant modifications of the modelled plant in those simulations. For a detailed study of the development of MACE models, consult Miller [17] and Padilla [22].

5.2 Linear Simulations
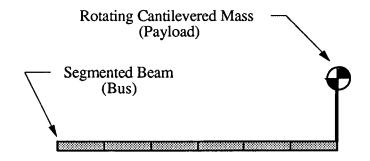
The linear tests were performed in MATLAB™ [28] using a simple finite element model of MACE. These tests established some basic knowledge about the behaviour of MACE in response to shaped and unshaped slewing commands.

5.2.1 The Linear Model

The linear model features only the bus and one of the endpoint gimballed payloads. The bus is treated as a uniform Bernoulli-Euler beam, divided into six finite elements. Each finite element has two degrees of freedom, one translational, and one rotational. The modelled payload is considered to be a point mass mounted at the end of a rigid member.

Torques can be commanded at the payload axis of rotation, simulating a gimbal, and the bus nodal translations and rotations are available as an output state vector. The model is constrained to two-dimensional motion, so the linear experiments examined only planar slews. Our version of the linear model also runs open loop, so any control system torques, including attitude controller inputs, are ignored. Gravity, of course, is also ignored.

Figure 5.1 shows a schematic of the MATLAB model. This linear model is conceptually a bit simpler than the actual MACE structure, but we can use tests conducted on this model to gain an understanding of the basic response characteristics of MACE. Consult Miller [17] for a detailed description of the background theory and specifications for this model.

94
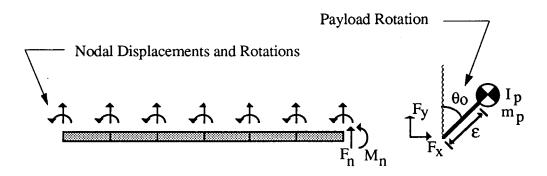
Figure 5.1: MATLAB linear finite element model of MACE.

The finite element model can be converted to a state-space form, which allows the system eigenvalues to be easily identified and simplifies the MATLAB simulation. The model has twelve modes of flexible vibration, and the first five linear model natural frequencies are:

Mode	Natural Frequency (Hz)
1	9.8
2	31.7
3	66.2
4	114.1
5	175.9

These frequencies will be used to develop input shapers that we will employ to suppress the model's post-slew residual vibration.

5.2.2 Experiments

We can use the eigenvalues from the above table to generate input shapers, as outlined in Chapter 3. Figure 5.2 shows the GAMS / Mathematica direct solution input shaper designed for the first three modes. The solution algorithm failed when higher modes were included (see Chapter 3, Section 3.4.3.c) so we pushed on with our three-mode shaper. The final impulse time of this sequence is 0.118 seconds, a response time savings over the equivalent convolved sequence of about 20%.
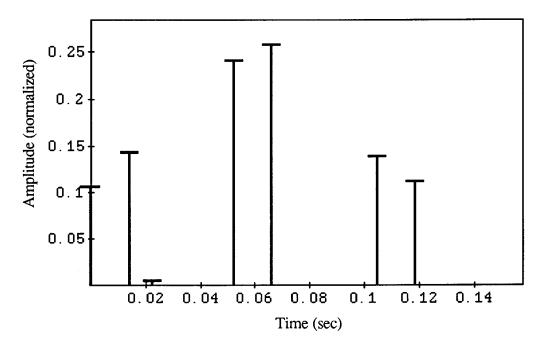
Figure 5.2:   Direct solution input shaper designed for 9.8 Hz, 31.7 Hz, and 66.1 Hz, ζ = 0.0 in all modes.

The next step is to define an input command to the payload gimbal. Recall that in this case, MACE is modeled as a free-floating article running open loop, so we need to use some sort of input that will not cause the experiment to go spinning wildly out of control. If we accelerate the payload up to a steady angular rate, pause to let the payload slew through some angle, and then decelerate the payload along a profile that mirrors the original acceleration, we should achieve bounded motion. The torque input corresponding to this maneuver is a pair of delayed square pulses.

A candidate torque input and the equivalent shaped input are shown in Figure 5.3. To employ the shaper of Figure 5.2 in the MATLAB code, we rounded the sequence impulse times to their nearest equivalent in simulation time steps. Note how convolving a step input with the input command shaper results in a "staircased" input; we never actually deliver an impulse to the physical system. This input command will cause the payload to rotate counter-clockwise, and the bus will therefore rotate clockwise in order to keep the system center of mass stationary.

To study the system vibration, we can examine the displacement history of the left-most bus finite element node, that node on the opposite end of the bus from the payload. Figure 5.4 shows the nodal displacement response history, and Figures 5.5 and 5.6 are detailed views of the unshaped and shaped responses, respectively.

The model had a system of twelve modes of vibration, and only the first three were used in forming the input shaper. It is clear from the figures, however, that cancelling these three modes was sufficient to suppress the

majority of the structure's vibration. The vibratory response to the unshaped input had a peak-to-peak value of roughly 1.2 mm, but in a space-based pointing application, vibration of this magnitude is a serious concern.
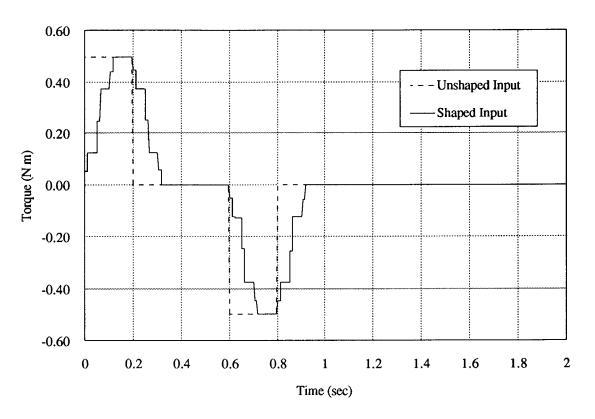
Figure 5.3: System inputs adjusted by the input shaper.
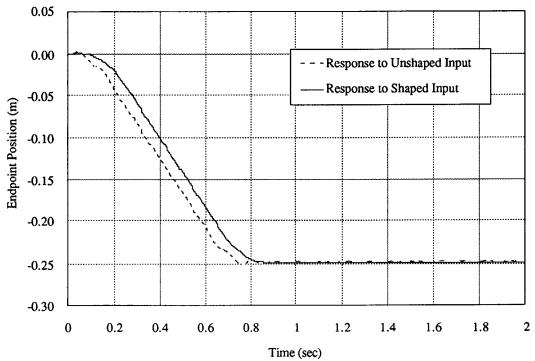
Figure 5.4: Bus endpoint translation due to inputs.
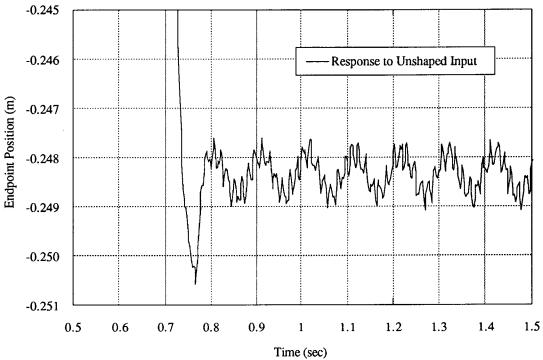
Figure 5.5: Response to unshaped input (detail).
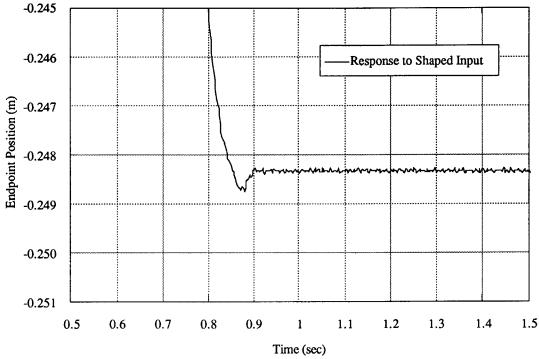
Figure 5.6: Response to shaped input (detail).

5.2.3 Conclusions

These MATLAB results are somewhat predictable. The input shapers are defined by equations that predict the response of linear systems, and the MATLAB model is also linear. Cancelling the vibrations of the MATLAB model, therefore, serves mainly as a confirmation of the proper solution of the constraint equations (3.1), and allows for concrete visualization of what the MACE system experiences when the input shapers are employed.

The next step in this program is to employ more accurate models of MACE. The actual structure will probably exhibit a variety of non-linear phenomenon, which our MATLAB model cannot emulate. We do have

access to non-linear MACE models, however, for the test article has also been simulated using the DISCOS program.

5.3 Non-Linear Simulations

No physical flexible system behaves in exact concordance with linear system theory. Inevitably, there is some factor involved in a physical experiment that will cause some amount of non-linear behaviour. Phenomena such as kinematic motion, non-linear dynamics, joint hysteresis, and friction will disrupt the accuracy with which linear equations can describe the motion of actual structures. It is possible to model many of the above non-linear phenomena, however, and doing so can allow a computer model to closely approximate the behaviour of an actual physical test article.

Having recognized the fallibility of their linear MATLAB model, SERC personnel also developed a MACE model using DISCOS (Dynamic Interaction Simulation of Controls and Structure), a program capable of describing non-linear behaviour [3]. This model can provide a more accurate picture of the MACE system behaviour, and can also be used to develop an understanding of the effects of actual system non-linearities on the input shaping method.

5.3.1   DISCOS

DISCOS is a large program written in FORTRAN that was originally developed to analyze spacecraft dynamics. The program can simulate several interacting bodies, all connected by specified hinges. The user can define sensors at desired points on the modelled bodies and hinges, to obtain time

101

histories of position, acceleration, etc. Initial body orientations and input torque or force commands can also be injected into a model. DISCOS uses numerical integration techniques to describe the non-linear motion of each body in the system.

We will not examine the DISCOS program in detail here; for more information consult the work by Bodley, et.al. [3], or for a concise overview of some of DISCOS' capabilities, see Sundaram [26].

5.3.2 The Non-Linear Model

The non-linear model expands the description of the MACE test article by adding the second payload and incorporating attitude and gimbal pointing proportional + derivative control systems. Three-dimensional motion can also be considered.

To retain similarity between the various MACE models, the DISCOS model was constrained to two-dimensional motion, and its left-most payload was clamped in a vertical orientation. One might argue that by constraining the non-linear model we lose some of our ability to accurately predict the motion of the physical MACE structure, but we want to develop a base level of understanding before we launch into examinations of the behaviour of more complex 3-D models.

The body and sensor layouts for the non-linear model are shown in Figure 5.7. Note that the bus sensors are not placed at the endpoints or midpoint of the bus; this arrangement will be helpful when we identify the

natural frequencies of the structure. For an account of the development of this model and its exact specifications, consult the work of Padilla [22].
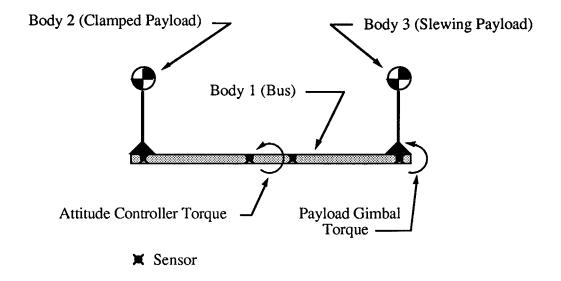
Figure 5.7: DISCOS non-linear flexible model of MACE.

Our next step in the non-linear study is to determine the natural frequencies of the model. We have two cases here, open loop and closed loop, which we will examine independently. To find the system natural frequencies, we can inject white noise into the attitude controller and the unclamped gimbal axis, and average the Fourier transforms of the responses at sensor points on the MACE structure.

For the open loop case, we set all the controller gains to zero, and with the payloads standing in the orientation shown in Figure 5.7, we averaged the spectra from several 30-second runs. The first four non-rigid frequencies were identified as:

Non-Linear Model Open Loop Frequencies

Mode	Natural Frequency (Hz)
1	1.65
2	5.42
3	6.90
4	19.08

Most of our DISCOS tests would be conducted using the model's built in controllers, so we wanted to obtain an accurate picture of the closed loop frequencies. We conducted two closed loop frequency identification tests; one with the payload canted inboard by 60°, and the second with the payload tilted out by 60°. These angles covered the designed slewing range of the payload gimbal, and allowed us to get an idea of how the frequencies might shift during a slew. We determined that a 6-second run was sufficient to identify the model frequencies, and by averaging twenty runs we found the first four non-rigid modes for the inboard and outboard closed loop cases:

Non-Linear Model Closed Loop (60° Inboard) Frequencies

Mode	Natural Frequency (Hz)
1	1.88
2	13.40
3	14.20
4	15.90

Non-Linear Model Closed Loop (60° Outboard) Frequencies

Mode	Natural Frequency (Hz)
1	2.18
2	14.25
3	15.25
4	15.90

As expected, the frequencies shift slightly when the payload orientation changes. SERC personnel estimated that the damping in each mode would be about 1%, or $\zeta = 0.01$, and higher modes were unreadable from the frequency spectra. We should note that with non-linear systems, the concept of a mode becomes somewhat unclear. Only linear systems can truly have modes and the corresponding eigenvalues and eigenvectors. With a non-linear system, we can only hope to identify the significant natural frequencies, as we have done above. We will use the term "mode" when referring to these frequencies to imply some continuity between the linear and non-linear experiments.

A second, very important note: in the closed loop system, there is actually another mode below the first. Strong damping by the control system prevents this mode from appearing on the Fourier spectrum, but it is evident when a step slewing command, for instance, is delivered to the payload gimbal. The payload will slew around to its commanded position, with a peak overshoot time of about 4.2 seconds, corresponding to a natural frequency of about 0.15 Hz. This slow mode will contribute little to the

residual vibration, and so it will not be included in the development of input shapers designed for the closed loop system.

We can anticipate two major types of non-linear behaviour from our DISCOS model: frequency shifts brought on by kinematic motion and centripetal acceleration effects caused by the slewing payloads. Other possible non-linearities, such as gimbal friction, saturation limits, or hysteresis are not included in the model. Estimates from SERC predict that if the lowest system frequency (0.15 Hz for the closed loop case) is less than one-tenth the frequency of the fundamental flexible mode, payload centripetal acceleration effects will be minor. Quantifying the exact effects of these problems is difficult, and as yet we have employed no metric for describing the actual severity of the system non-linearities.

We can predict, however, that the kinematic non-linearities will be more severe for large angle slews, and the centripetal acceleration effects will increase as the payload angular velocity increases. Our objective is to identify the supplemental residual vibration contributed by non-linear behaviour, and to suppress as much as possible all of the residual vibration in the modelled non-linear system.

5.3.3 Open Loop Tests

We conducted open loop experiments to provide some basis for comparison with the MATLAB tests conducted earlier, and to serve as a base case before controllers were added. We fed the identified open loop frequencies into the GAMS routine, generating the sequence shown in Figure 5.8. We did not

attempt an exact solution in this case, we merely set the GAMS time step equal to the DISCOS simulation time step, and thus our GAMS shaping sequence is tuned to the DISCOS "servo rate." This sequence affords a response time savings of 25% over the equivalent convolved sequence.
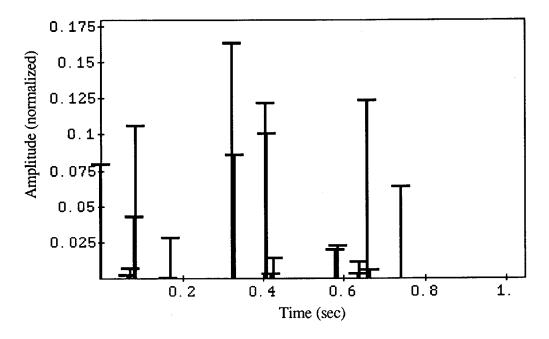
Figure 5.8:   Shaping sequence for 1.65 Hz, 5.42 Hz, 6.90 Hz, and 19.08 Hz, $\zeta = 0.01$ in all modes.
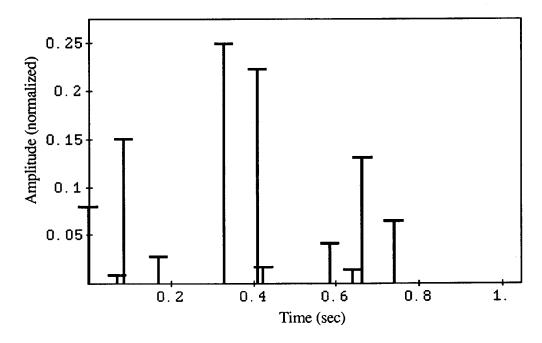
Figure 5.9: Modified sequence from Figure 5.8.

We can "clean up" the sequence of Figure 5.8 by combining neighboring impulses while retaining the time mesh spacing, yielding the sequence of Figure 5.9. A study of the new sequence's insensitivity curve reveals it to be nearly as effective at suppressing vibration as the raw GAMS output (see Chapters 2 and 3 for descriptions of insensitivity curves and methods for adjusting the GAMS sequences). Impulse population is not a significant concern in this case, it will only slow down the simulation, so by using the sequence of Figure 5.9 we're just saving the trouble of typing in a few more impulse specifications. If we were implementing the sequence on a physical controller, however, the sequence of Figure 5.9 would decrease the computational burden on the controller, possibly making room for a higher servo rate.

We wrote a few subroutines to implement the shaper in the DISCOS source code, and set the controller gains to zero to achieve open loop running. To retain similarity with the MATLAB tests, square wave torque inputs were delivered to the unclamped payload gimbal axis, producing about a 40° slew. We examined the resulting vibration occurring at the opposite end of the bus.

Figures 5.10 and 5.11 display the vertical motion of the left-most sensor, in response to unshaped and shaped torque inputs, respectively. These graphs are centered around zero because the DISCOS sensors are measuring relative, not absolute motion. The vibrations of the response to the unshaped inputs are small, in the sub-millimeter range, but again, many feasible applications can be disrupted by vibrations of this magnitude.
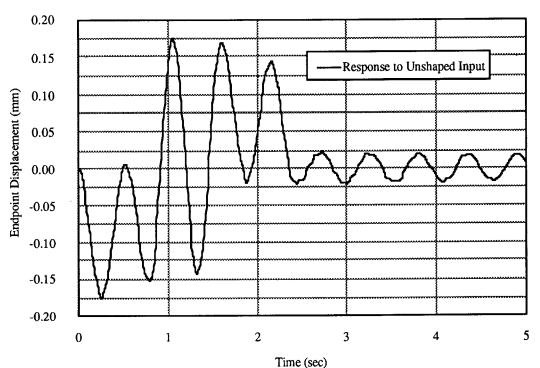
Figure 5.10: Bus endpoint response to unshaped input.
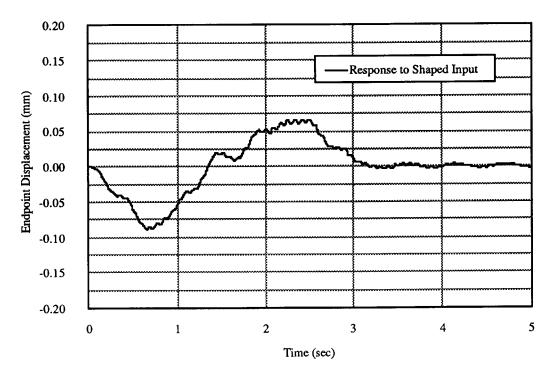
Figure 5.11: Bus endpoint response to shaped input.

The shapers again performed well, greatly reducing the amount of residual vibration, but we still have some leftover oscillation. A close-up inspection of the plot in Figure 5.11 reveals that the system is still vibrating at about 1.9 Hz and 12.5 Hz. We would not observe this phenomenon with a linear model. Clearly, these are new frequencies that the structure adopts in its post-slew configuration. These new frequencies were not included in the shaper development and they differ relatively significantly from the original frequencies, so they were not removed by the input shaper. We had identified our open loop frequencies with the payloads standing vertical; the frequencies apparently shifted as a result of the slewing maneuver.

This kinematic frequency shift is a non-linear phenomenon, one of the problems we anticipated in the previous section. Since most of the actual

MACE experiments will be run closed loop, we will leave a detailed investigation of this frequency shift to the next section.

In our open loop tests we have demonstrated a rough correspondence with the linear MATLAB experiments, and we recognize that the input shapers can be effective in suppressing vibration in non-linear systems, but non-linearities can cause some shaper performance degradation. We will continue our examination of system non-linearities and begin a thorough study of the associated shaping problems by using our closed loop model.

5.3.4 Closed Loop Tests

We conducted a broad range of closed loop slewing maneuvers to span the performance envelope of the MACE gimbal and to hopefully induce varying non-linear motion. To shift from open loop to closed loop running, we simply brought the controller gains up from zero to their specified values in Padilla [22]. The simulations in this section feature symmetric slews, with the payload starting at an outboard angle (rotated clockwise, or negative Θ, from vertical), and moving to the equivalent inboard angle.

We will examine two test cases, slews from negative 20° to positive 20° and from negative 60° to positive 60°. We know that the system frequencies will vary as the payload moves, but as an initial experiment we will employ a shaper derived for the 60° outboard frequencies as defined in Section 5.3.2. Similar to the open loop tests of the previous section, the shaper can be tuned to the DISCOS "servo rate" by setting the GAMS time step equal to the DISCOS simulation time step. The raw GAMS shaper for the four identified

frequencies is shown in Figure 5.12, with the equivalent "cleaned up" shaper shown in Figure 5.13. An 17% savings in response time is gained by using the sequence in Figure 5.13 instead of the equivalent convolved shaper.
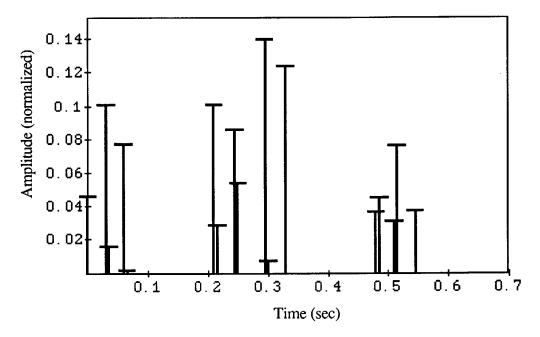
Figure 5.12: Shaping sequence for 2.18 Hz, 14.25 Hz, 15.25 Hz, and 15.90 Hz, $\zeta = 0.01$ in all modes.
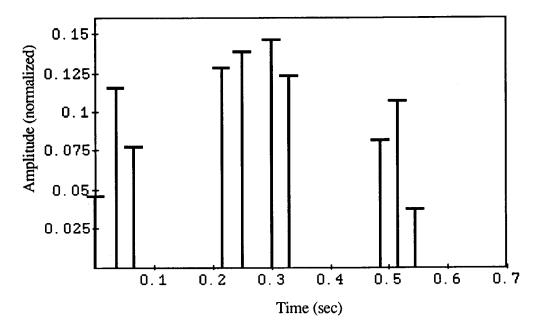
Figure 5.13: Modified sequence of Figure 5.12.

The absolute inertial payload angle history in response to shaped and unshaped 40° slew commands is shown in Figure 5.14. This graph clearly shows the slight delay caused by the shaper, similar to the plot in Figure 5.4. The opposite endpoint vertical displacement caused by the unshaped command is detailed in Figure 5.15, and Figure 5.16 shows the response to the shaped command.

Once again, the shaper effectively inhibited most of the residual vibration, and this example illustrates the shaper's ability to work with existing control systems. There is still some leftover vibration, however; the plot in Figure 5.16 clearly shows oscillations at about 1.8 Hz. Recall that we used the 60° outboard frequencies to develop our input shaper. We know

that the actual system frequencies during the 40° slew will differ from the frequencies included in the shaper, and apparently the shifts were significant enough to degrade the shaper performance. The payload centripetal acceleration could also act to increase the residual vibration, but we know from Section 5.3.2 that the centripetal acceleration effects will be small. In any case, the problems caused by system non-linearities in this slewing case are minor, and we'll examine a larger slew case to get a better view of the non-linear problems.
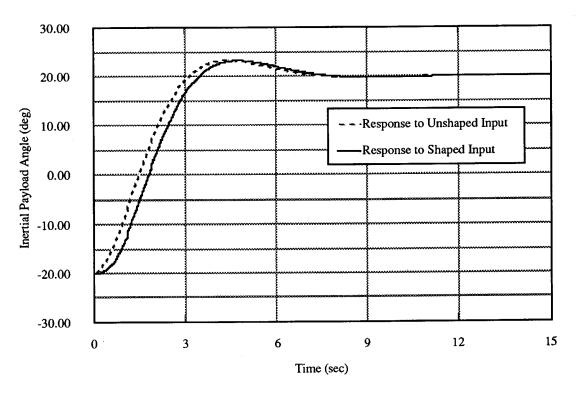
Figure 5.14: Closed loop 40° payload slew.
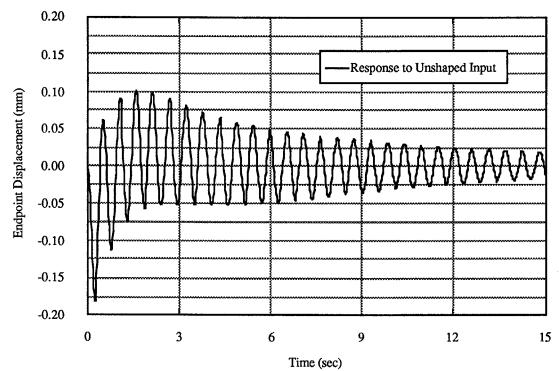
Figure 5.15: Endpoint response to unshaped 40° slew.
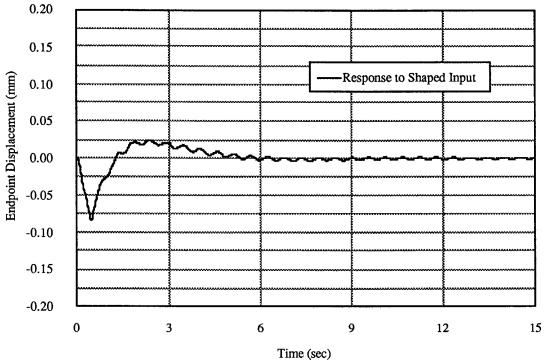
Figure 5.16: Endpoint response to shaped 40° slew.

When the payload gimbal travels between its extreme maneuvering limits, negative 60° and positive 60°, the kinematic frequency shifts should be greater than those experienced in the 40° slew. The inertial payload angle history for the 120° slew is shown in Figure 5.17. Note the strong similarity between this graph and the 40° trace of Figure 5.14. The only difference is the beginning and ending points of the slew; thanks to the MACE control systems, the same amount of time is required to complete both maneuvers. The 120° slew, therefore, requires greater acceleration and higher velocity, which should boost any payload centripetal acceleration non-linear effects. The delay between the shaped and unshaped responses is about the same as in Figure 5.14; we used the same input shaper for both the 40° and 120° cases.
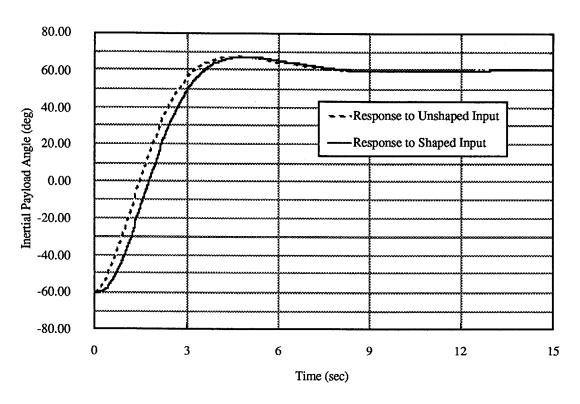
Figure 5.17: Closed loop 120° payload slew.

The graph of the unshaped response to the 120° slew is shown in Figure 5.18; note the strong resemblance to the equivalent curve for the 40° case, Figure 5.14. The major difference is that the maximum peak to peak displacement per cycle in the large slew case measures about 0.50 mm, significantly greater than Figure 5.14's 0.15 mm displacements. These larger amplitude vibrations are a reflection of the greater energy required to complete the 120° slew in the same time as the 40° slew.

The response to the shaped 120° slew (Figure 5.19), correspondingly, looks nearly identical to Figure 5.16, the difference again being only in vertical scale. Where the graph in Figure 5.16 peaks at about 0.025 mm and flattens out at about 6 seconds, the shaped response to the 120° slew peaks at 0.09 mm and dies out also at about 6 seconds. A closer look at Figure 5.19, however, reveals significant residual vibration, more serious than the leftover oscillations from the 40° slew (see Figure 5.20).
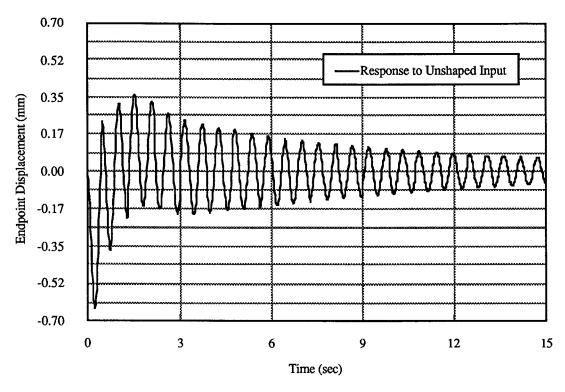
Figure 5.18: Endpoint response to unshaped 120° slew.
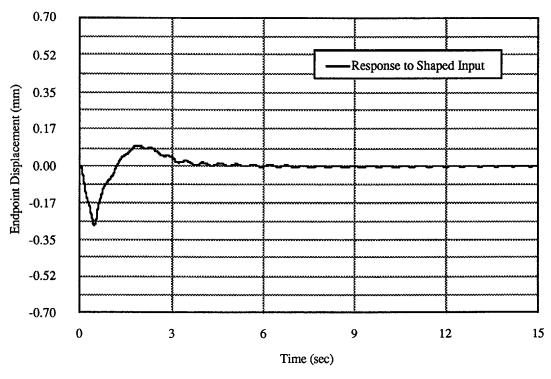
Figure 5.19: Endpoint response to shaped 120° slew.

118
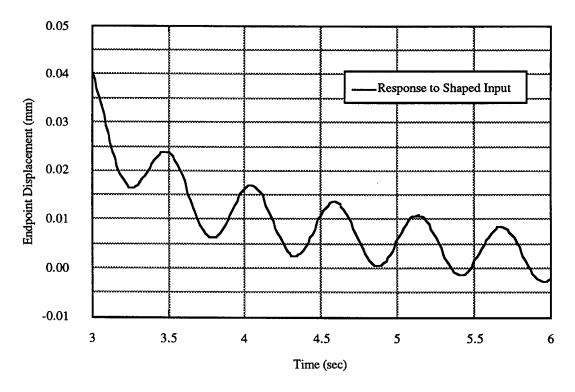
Figure 5.20: Detail of residual vibration from Figure 5.19.

The shaper suppressed most of the residual vibration, but a roughly 1.83 Hz leftover oscillation is clearly evident in Figure 5.20. Recall that we used the 60° outboard frequencies, with a fundamental of 2.18 Hz, to develop our shaping sequence. Apparently, the system frequencies shifted far enough during the course of the slew to disrupt the shaper's performance. This fast, large angle maneuver would also cause stronger payload centripetal accelerations, invoking dynamic effects that our technique is not prepared to manage. The challenge now is to determine if it is possible to modify our original shaper to compensate for some of these non-linear problems.

5.3.5 Improving the Non-Linear Response

There are countless modifications of our original shaper that we might use to improve the response curve of Figure 5.19. We'll look at a few in this section that can help; keep in mind that there probably is no one sequence that will always be the most effective. A particular shaper's performance will depend on slew size and speed.

The first modification we can try is to use the 60° inboard frequencies instead of the 60° outboard values. When we deliver our step input to the system, several frequencies are excited. The lowest, well damped frequency is regulated by our control system, as mentioned earlier. We also cause vibration at four other system frequencies, those mapped at the 60° outboard configuration. By the time the payload slews over to the 60° inboard position, the earlier frequencies have shifted, reflecting the kinematic motion of the MACE structure. If we shape for the frequencies at the final, rather than the initial position, we will be anticipating this frequency shift, and we should observe stronger vibration suppression.

The input shaper designed for the 60° inboard frequencies is shown in Figure 5.21. This sequence was the result of a simple modification of the raw GAMS output from the inboard four-mode problem. As with the previous sequences of this section, the neighboring impulses in the GAMS output were combined, yielding a "cleaner" sequence.
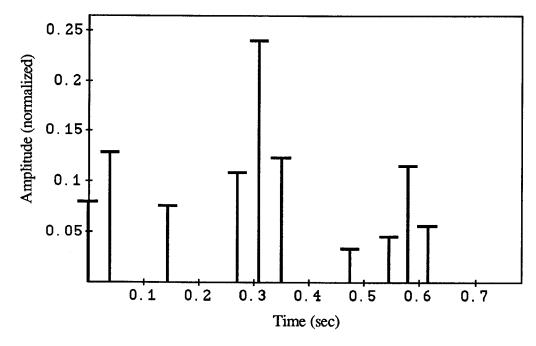
Figure 5.21: Shaping sequence for 1.88 Hz, 13.40 Hz, 14.20 Hz, and 15.90 Hz, $\zeta = 0.01$ in all modes.

When we use this shaper to adjust our 120° slew command, the resulting endpoint displacement is as shown in Figure 5.22, with a detailed plot shown in Figure 5.23. A comparison with Figure 5.20 reveals that we have managed to decrease the peak-to-peak value of the residual vibration, but the roughly 1.83 Hz oscillation is still present. The lesson here is that it helps to build the shaper using the frequencies present at a slew endpoint, not those active at the initial configuration.
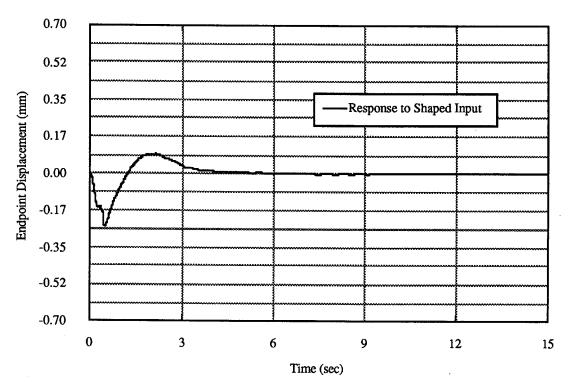
Figure 5.22: Endpoint response to shaped 120° slew.
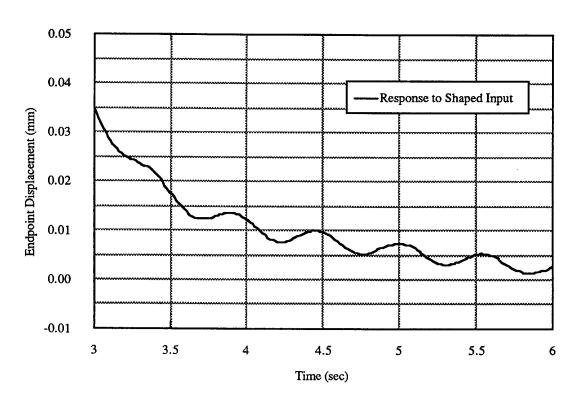
Figure 5.23: Detail of residual vibration from Figure 5.19.

122

We have now used a shaper designed for the frequencies that seem most appropriate, and still we observe the 1.83 Hz residual vibration. The fundamental mode in our shaper is 1.88 Hz, which should be close enough to eliminate the remaining oscillations. We might surmise that the shaper's parameter insensitivity is undermined by the non-linear behaviour, but in tests where the shaper used 1.83 Hz as a fundamental, the vibration was not decreased significantly.

We can try to force our sequence to be more insensitive to the varying parameters by adding a frequency or two to the shaper. The residual motion of Figure 5.23 oscillates at a frequency lower than the shaper's included fundamental; let's try adding a frequency to our shaper that is also lower than the current fundamental. Figure 5.24 shows a shaper designed for the original four frequencies, plus a fifth at 1.50 Hz.

The insensitivity curve for the shaper of Figure 5.24 is shown in Figure 5.25. The flattened trace at 1.50 Hz and 1.88 Hz is clearly evident; the curve doesn't quite touch zero simply because we're using a modified GAMS sequence; this in itself will cause no appreciable increase in the residual vibration. The associated system response curve for this sequence is shown in Figure 5.26. A comparison with Figure 5.23 reveals that adding the extra frequency didn't exactly help, the residual vibration in Figure 5.26 is if anything greater than the vibration shown in Figure 5.23.
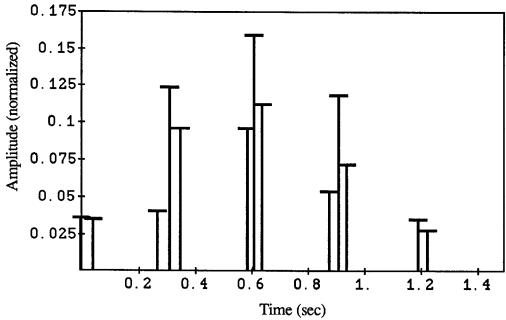
Figure 5.24: Shaping sequence for 1.50 Hz, 1.88 Hz, 13.40 Hz, 14.20 Hz, and 15.90 Hz, ζ = 0.01 in all modes.
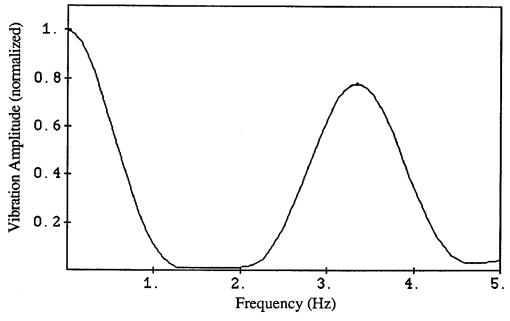
Figure 5.25: Insensitivity curve for shaper of Figure 5.24.
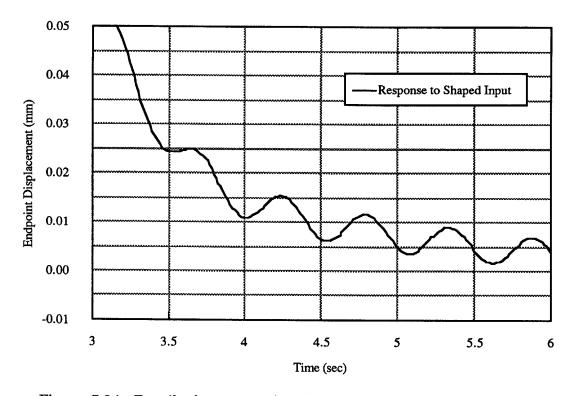
Figure 5.26: Detail of response found using sequence of Figure 5.24.

We can extend this approach and attempt to improve the response by adding yet another frequency, this one greater than the original fundamental. Figure 5.27 shows the resulting sequence when frequencies of 1.5 and 2.0 Hz are added to the original four 60° inboard frequencies. The corresponding insensitivity curve is shown in Figure 5.28. We have now completely flattened the curve in the region surrounding the original 1.88 Hz fundamental. When we use the new sequence, the system responds as detailed in Figure 5.29.
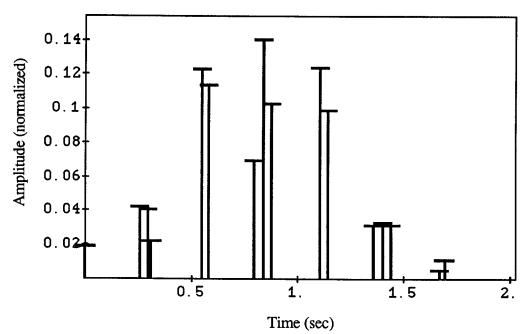
Figure 5.27: Shaping sequence for 1.5 Hz, 1.88 Hz, 2.0 Hz, 13.40 Hz, 14.20 Hz, and 15.90 Hz, ζ = 0.01 in all modes.
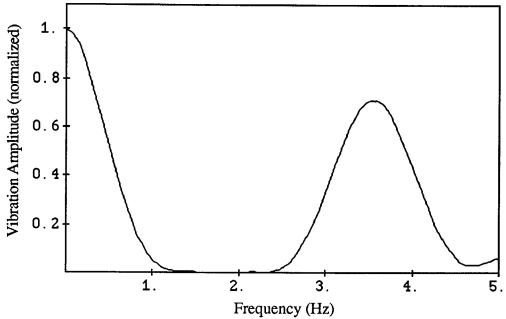
Figure 5.28: Insensitivity curve for sequence of Figure 5.27.

126
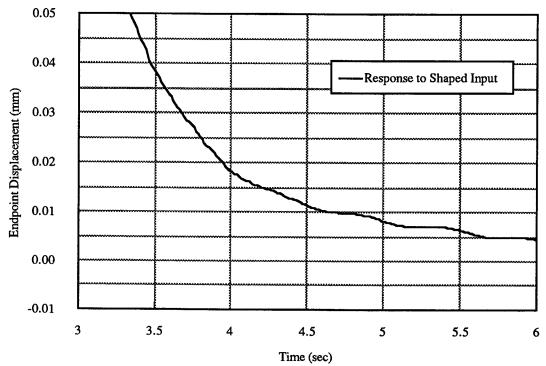
Figure 5.29: Detail of response found using sequence of Figure 5.27.

This last modification causes significant improvement over the response detailed in Figure 5.20. Bounding the shaper's fundamental frequency to increase the parameter insensitivity is thus a viable means to suppress residual vibration even in systems that experience notable frequency shifts. The penalties associated with adding new frequencies to the shaper are an increase in controller burden due to the impulse population increase, and an increased delay in the system response thanks to the longer shaper length. We have only to compare Figure 5.29 to the equivalent section of the unshaped response curve of Figure 5.18, however, to realize the overall benefits resulting from shaping the system input.

We have not examined the specific problems caused by payload centripetal acceleration. This non-linear effect will certainly increase with faster slew rates, and some current work at SERC is dedicated to identifying the exact magnitude of the centripetal acceleration terms in the equations that govern the MACE response. We will leave the task of a detailed centripetal acceleration examination to future writers.

5.4 Conclusion

In this chapter we have conducted tests of input shaper effectiveness on linear and non-linear models of MACE, the test article described in Chapter 4. We have shown that in linear systems, shapers can completely suppress all modelled modes included in the shaper development. Simulations run on non-linear models indicate that shapers can effectively suppress identified vibration frequencies, but system non-linearities tend to disrupt the shaper's general performance.

Shaper performance degradation caused by kinematic frequency shifts can be reduced by boosting the shaper's inherent insensitivity to plant parameter variations. Adding frequencies to the shaper which bound the aberrant oscillation will allow the shaper to effectively suppress even those vibrations which tend to shift in frequency as the system moves.

produce vibration free motion. The "input shaping" method demonstrates insensitivity to variations in plant parameters, but causes small delays in system response time.

In Chapter 3 we extended the single mode input shaping approach to multiple mode problems. We examined two multiple mode methods: convolution and direct solution. Convolution works by combining several single mode shapers into an impulse train that will perform simultaneous vibration suppression of several modes.

The direct solution method involves a re-formulation of the single mode shaper equations to include several modes. These equations are then solved simultaneously to generate a single sequence that matches the performance of the convolved shaper. The direct solution sequences are computationally less intensive to implement than convolved sequences, and they are also shorter, causing smaller system response delays. These advantages are offset by an increase in sequence generation complexity.

The Mid-deck Active Control Experiment (MACE) was the focus of Chapter 4. MACE is an experimental flexible structure under development at the MIT Space Engineering Research Center (SERC). The research for MACE is prompted by NASA proposals for large, earth orbiting satllites that will be used to study long term environmental trends including weather patterns and global warming problems.

The MACE structure features a 1.5 meter segmented tubular beam, or "bus" with 2 degree of freedom pointing payloads located at the bus end-
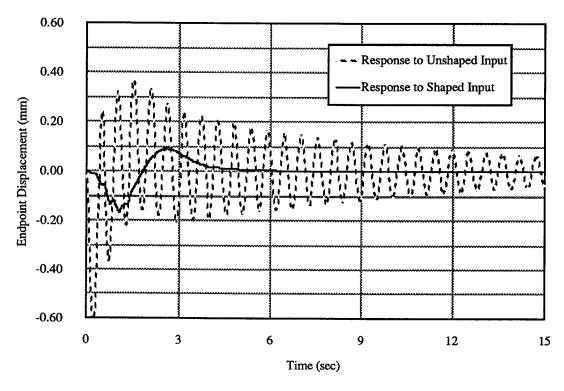
Figure 6.1: MACE endpoint response to shaped and unshaped 120° slew.

6.2 Future Work

A crucial step that must be carried out in the near future is the hardware verification of the conclusions of Chapter 5 through tests on the MACE test article. The results we gained from the non-linear simulations are valuable and show great promise for the shaping technique, but even the most careful simulation will suffer from some inaccuracies, and applying the shapers to physical systems is unquestionably a critical activity. One of the writer's regrets in this work is that sufficient time was not available for physical testing.

A second important endeavour is the improvement of the direct solution generation algorithms. We have shown that the direct solution sequences can afford advantages over the convolution multiple mode approach, but the complexity of the direct solution sequences drastically increases the computational effort required for their creation. To allow direct solution sequences to be generated rapidly, possibly even in real-time, adaptive applications, this computational burden must be lessened.

Closed form solutions to the direct solution equations are currently being sought, as mentioned at the end of Chapter 3. These approaches, if successful, could usher in expanded possibilities for the input shaping technique. Using a closed form solution, sequences could be generated with great speed, and in non-linear environments where the plant frequencies might be shifting, new shapers could be created as the system moves to account for such plant variations.

A final interesting step is to investigate the best means for blending input shaping with standard vibration reduction techniques. Many approaches for shaping movement trajectories or commanded velocity profiles can only benefit from the supplemental employment of the shaping technique. Future studies might be directed toward discovering an optimal combination of input shaping with the different vibration reduction techniques, those both currently available and newly emerging.

[7]   Cannon, R. H., and Schmitz, E.,
        "Initial Experiments on the End-Point Control of a Flexible One-Link
        Robot," The International Journal of Robotics Research, Vol. 3, No.3,
        Autumn 1984.

[8]   Crawley, E. F., and de Luis, J.,
        "Use of Piezo-Ceramics as Distributed Actuators in Large Space
        Structures," AIAA Paper No. 85-0626, Structures, Structural Dynamics
        and Materials Conference, Orlando, Florida, April 1986.

[9]   Crawley, E. F., de Luis, J., and Miller, D. W.,
        "Mid-deck Active Control Experiment (MACE): Phase A Final Report,"
        MIT Space Engineering Research Center Report No. 16-90-R, June 1989.

[10]   Farrenkopf, R. L.,
        "Optimal Open-Loop Maneuver Profiles for Flexible Spacecraft,"
        Journal of Guidance and Control, Vol. 2, No. 6, November-December
        1979.

[11]   Ford Aerospace Corp.,
        "Geostationary Platform Bus Study Final Report - Volume II," Report
        No. WDL-TR10860, Comprehensive Report NAS 8-36104, September
        1987.

[12]   Gupta, N. K., Lyons, M. G., Aubrun, J. N., and Margulies, G.,
        "Frequency-Shaping methods in Large Space Structures Control,"
        AIAA Paper No. 81-1783, AIAA Guidance and Control Conference,
        Albuquerque, New Mexico, August 1981.

[13]   Hollars, M. G., and Cannon, R. H.,
        "Initial Experiments on the End-Point Control of a Two-Link
        Manipulator with Flexible Tendons," ASME Winter Annual Meeting,
        November 1985.

[14]   Hyde, J. M. and Seering, W. P.,
        "Using Input Command Pre-Shaping to Suppress Multiple Mode
        Vibration," IEEE Conference on Robotics and Automation, April 1991.

[15]   Laskin, R. A., and Sirlin, S. W.,
        "Future Payload Isolation and Pointing System Technology," Journal of
        Guidance, Control, and Dynamics, Vol. 9, No. 4, July-August 1986.

[16]   Meckl, P. H. and Seering, W. P.,
        "Controlling Velocity-Limited Systems to Reduce Residual Vibration,"
        IEEE Conference on Robotics and Automation, April 1988.

[17] **Miller, D. W.,**
"Mid-deck Active Control Experiment (MACE): Preliminary Linear Control Analysis," MIT Space Engineering Research Center Document No. MACE-1-120, November 8, 1990.

[18] **MIT Space Engineering Research Center,**
"Mid-deck Active Control Experiment (MACE): Engineering Model Design Document," Version 1.0, March 16, 1990.

[19] **National Aeronautics and Space Administration,**
"Earth Observatory Satellite: A Mission to Planet Earth," NASA Headquarters Publication No. 75-K, February 1990.

[20] **Nurre, G. S., Ryan, R. S., Scofield, H. N., and Sims, J. L.,**
"Dynamics and Control of Large Space Structures," Journal of Guidance, Control, and Dynamics, Vol. 7, No. 5, September-October 1984.

[21] **Ogata, K.,**
Modern Control Engineering, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1970.

[22] **Padilla, C. E.,**
"Mid-deck Active Control Experiment (MACE): Nonlinear Modelling, Simulation, and Preliminary Control of the Baseline Test Article, Final Report," MIT Space Engineering Research Center, August 28, 1990.

[23] **Singer, N. C.,**
"Residual Vibration Reduction in Computer Controlled Machines," MIT Artificial Intelligence Lab Technical Report #1030, February 1989.

[24] **Singhose, W. E.,**
"Shaping Inputs to Reduce Residual Vibration: A Vector Diagram Approach," MIT Bachelor of Science Thesis, February 1990.

[25] **Smith, O. J. M.,**
Feedback Control Systems, McGraw-Hill Book Company, Inc., New York, N.Y., 1958.

[26] **Sundaram, K. J.,**
"The Use of Joint Torque Feedback for Force Control," MIT Master of Science Thesis, May 1990.

[27] **Swigert, C. J.,**
"Shaped Torque Techniques," Journal of Guidance and Control, Vol. 3, No. 5, September-October 1980.

[28] **The MathWorks, Inc.,**
PRO-MATLAB User's Guide, South Natick, Massachusetts, 1989.

[29] **Turner, J. D., and Junkins, J. L.,**
"Optimal Large-Angle Single-Axis Rotational Maneuvers of Flexible Spacecraft," Journal of Guidance and Control, Vol. 3, No. 6, November-December 1980.

[30] **Tzes, A. P., Englehart, M. J., and Yurkovich, S.,**
"Input Preshaping With Frequency Domain Information For Flexible-Link Manipulator Control," AIAA Guidance, Navigation, and Control Conference, August 1989.

[31] **Yurkovich, S., Pacheco, F., and Tzes, A.,**
"On-Line Frequency Domain Information for Control of a Flexible-Link robot with Varying Payload," Ohio State University Control Research Lab Report #CRL-1036-Su88-P, 1989.

[32] **Wie, B., and Liu, Q.,**
"Feedforward/Feedback Control Synthesis for Performance and Robustness," AIAA Guidance, Navigation, and Control Conference, 1990.

[33] **Wolfram, S.,**
Mathematica: A System for Doing Mathematics by Computer, Addison-Wesley, Redwood City, CA, 1988.

Appendix: MACE Test Article Development

A.1 Introduction

This appendix describes design and vendor survey activity conducted by the writer in support of the MACE test article development. Tasks relating to the MACE Ground Test Article (GTA) included design of the attitude control unit and specification of the corresponding purchased hardware. MACE Flight Test Article (FTA) activity included vendor surveys and initial price quote acquisition for the FTA attitude control unit and gimbal systems.

A.2 GTA Attitude Control Unit Development

The Attitude Control Unit is composed of five elements: inertial wheels, motors and supporting electronics, the node/motor interface, and a protective housing. The cluster can exert torques of 40 oz. - in. around three orthogonal axes for up to five seconds, and wheel rotation rates can vary from zero to 6000 RPM, with about ± 60 RPM accuracy.

The inertial wheels attach to the motor shafts using a split hub and clamping collar arrangement. The wheels were designed to allow for safe operation at 12,000 RPM, twice the normal maximum speed. Running at 6000 RPM, the disks have a safety factor of approximately 2.5. Each wheel is fabricated from aluminum, measures 6.75 inches in diameter, and weighs about 2.2 lbs. Figure A.1 shows the engineering drawing of an individual inertial wheel.

The motors that drive the inertial wheels are Aerotech 1017 servo units. The inertial wheels are clamped to the motor encoder shaft, not the standard motor output shaft. This unconventional arrangement is used because the encoder shaft is smooth, and has a 0.375 in. diameter, while the output shaft has two milled flats and has only a 0.250 in. diameter. Allowing the inertial wheel to grip the larger, fully round encoder shaft generates a more stable configuration. The motor bearings are fully capable of supporting the wheel, even though the load is attached to the "wrong" end of the shaft. The unorthodox mounting arrangement allows the motor/inertial wheel subsystem to be simply bolted to the node/motor interface using the standard motor face plate hole pattern.

The motors are supplied with an Aerotech DS8020 pre-matched 3-axis amplifier unit, which can easily power each motor to supply 40 oz. in. of torque from 0 to 6000 RPM. The motors weigh 1.5 lbs, and to retain this low weight, tachometers are omitted from the system. The back-EMF signal from the motor is used instead to sense and control angular velocity. The bandwidth capability of this control scheme has yet to be determined.

The node/motor interface connects the motor/inertial wheel assemblies to the standard MACE node. The interface's geometry defines the three orthogonal axis configuration, and allows for the torque wheel unit's center of mass to hang directly underneath the node. The block is manufactured from aluminum and weighs about 1.5 lbs. Engineering drawings of the interface are shown in Figures A.2 through A.5.

The protective housing will be constructed of thin plexiglas sheeting and is intended only to prevent human interaction with the wheels, not to contain fragments in the event of a wheel burst. Overload protection is provided by the wheels themselves and the limited RPM capabilities of the motors. Various configurations are still under consideration for the housing so an engineering drawing is not shown. The housing should weigh about 3 pounds, giving an overall attitude control unit weight of approximately 15.6 pounds.

Figure A.1: Inertial wheel engineering drawing.

NOTES:

1) MAT'L: 2024-T4 ALUMINUM.

2) HOLES 'A': DRILL 9/32" DIA THRU AND C-BORE 0.422 DIA (0.30 DEEP).

3) HOLE 'B': DRILL AND REAM TO 0.500 DIA. THIS HOLE SERVES AS THE GEOMETRICAL CENTER FOR THE OVERALL PART.

4) SEE SHEET #4 FOR HOLE PATTERNS ON SLOPING FACES.

TOP VIEW

2.742"

2.670"

120.00°

2.121" DIA B.C.

3.000"

60.00°

3.250"

6.250"

TOLERANCES		Space Engineering Research Center
X.X ±0.020		Massachusetts Institute of Technology
X.XX ±0.010		Cambridge, MA
X.XXX ±0.001		
ANGLES ± 0.06°		
(5 MIN OF ARC)	TITLE	NODE INTERFACE (SHEET #1)
Unless otherwise specified		
DATE	DRAWN	J.M. HYDE
8/8/90	CHECKED	

| REV | SCALE | QUANTITY | DRAWING NO. |
| A | 1:2 | 1 | 003 |

Figure A.2: Node/motor interface engineering drawing, sheet #1.

143

NOTES: 1) MAT'L: 2024-T4 ALUMINUM

FRONT VIEW

2.087"

0.250"

TOLERANCES		
X.X ±0.020		
X.XX ±0.010		
X.XXX ±0.001		
ANGLES ± 0.08°		
(5 MIN OF ARC)		
Unless otherwise specified		

Space Engineering Research Center
Massachusetts Institute of Technology
Cambridge, MA

TITLE NODE INTERFACE (SHEET #2)

DRAWN	J.M. HYDE
CHECKED	

REV	SCALE	QUANTITY	DRAWING NO.	DATE
A	1:2	1	004	8/6/90
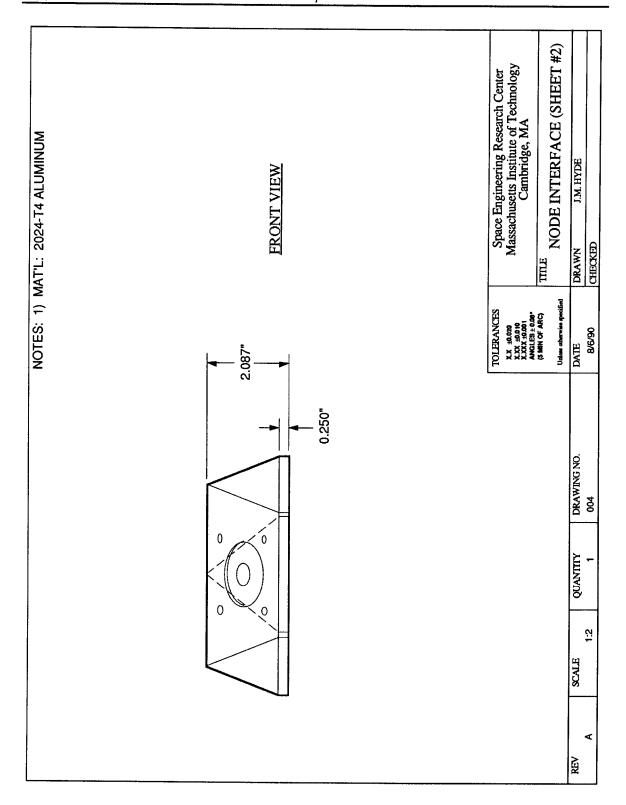
Figure A.3: Node/motor interface engineering drawing, sheet #2.

144

NOTES: 1) MAT'L: 2024-T4 ALUMINUM

SIDE VIEW

125.26°

TOLERANCES		Space Engineering Research Center Massachusetts Institute of Technology Cambridge, MA	
X.X ±0.020 X.XX ±0.010 X.XXX ±0.001 ANGLES ±0.08° (5 MIN OF ARC)			
Unless otherwise specified		TITLE	NODE INTERFACE (SHEET #3)
DATE	8/6/90	DRAWN	J.M. HYDE
		CHECKED	

REV	SCALE		QUANTITY	DRAWING NO.
A	1:2		1	005
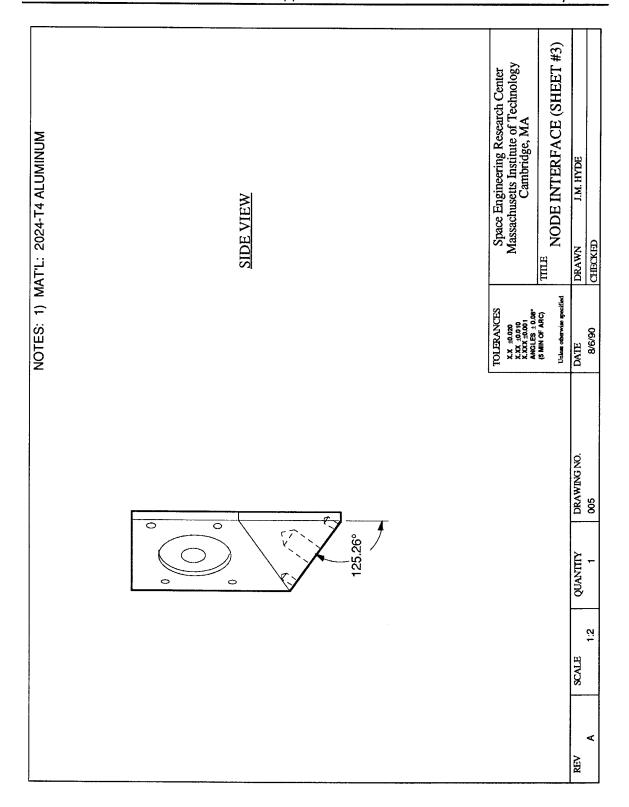
Figure A.4: Node/motor interface engineering drawing, sheet #3.

145

NOTES: 1) MAT'L: 2024-T4 ALUMINUM.
2) HOLES ARE DRILLED ORTHOGONAL TO THEIR RESPECTIVE FACES ON THE NODE INTERFACE (SEE SHEET #1).
3) THIS DRAWING DEPICTS THE SLOPING FACES AS VIEWED ORTHOGONAL TO THEIR SURFACES.

DRILL AND TAP FOR 10-32 THREADS (0.20" DEEP) 4 PL

C-BORE 1.505" DIA (0.125" DEEP)

CENTERLINE OF FACE

DRILL 0.50" DIA (0.675" DEEP - AS MEASURED FROM BOTTOM OF C-BORE)

0.930"
0.930"
0.930"
0.930"

Space Engineering Research Center Massachusetts Institute of Technology Cambridge, MA		
TITLE		
NODE INTERFACE (SHEET #4)		

TOLERANCES X.X ±0.020 X.XX ±0.010 X.XXX ±0.001 ANGLES ±0.05° Unless otherwise specified		DRAWN	J.M. HYDE
		CHECKED	
DATE 8/6/90			
DRAWING NO. 006			
QUANTITY 3			
SCALE 1:1			
REV A			

Figure A.5: Node/motor interface engineering drawing, sheet #4.

146

A.3 FTA Vendor Surveys

The following sections are paraphrased from reports delivered at the MACE Critical Design Review on February 20, 1990. Specifications and vendor information is listed for both the FTA Attitude Control Unit and the FTA Gimbal Systems.

A.3.1 FTA Attitude Control Unit

The attitude control unit provides MACE with three axis attitude control and can also be used to suppress low frequency vibration. The limited vibration suppression can be achieved by rapidly accelerating and decelerating the inertial wheels, providing an external torque on the structure that can be used to counteract vibration. The unit will interface to a standard MACE node, and will normally reside under the central node on the bus. The physical requirements for the torque wheels are:

Torque:	40 oz-in (5 seconds max duration)
Resolution:	+/- 0.5 oz-in
Bandwidth:	1-100 Hz
Operation:	Unidirectional

In addition to these requirements, low weight is desired, on the order of 15 lbs total. Some sort of protective housing must also be featured, to prevent astronaut interaction with the wheels. The following attitude control unit manufacturers have been identified and contacted:

147

Bendix: Teterboro, NJ
Draper Labs: Cambridge, MA
Honeywell: Glendale, AZ
JPL: Pasadena, CA
Versatron: Healdsburg, CA

Draper has been asked to identify an internal group which would be interested in supplying the components, with no coherent response thus far. The other potential vendors received an informal request for quote on February 6, 1990.

JPL wishes only to be involved as advisors, and Bendix, although they provided some details of their proposed system, failed to deliver a price quote. Honeywell claimed to be able to deliver a unit meeting our specifications for $100,000 per axis. Versatron communicated the lowest price estimate, placing the cost of three torque wheels at $100,000 total.

A.3.2 FTA Gimbal Systems

The gimbal systems serve as the core of the two-axis articulated pointing mechanisms residing on either end of MACE's bus. The gimbals will interface to a standard MACE node, and point a five pound payload comprised of dead weight and a rate gyroscope package. These pointing systems will serve both as disturbance sources, and measurements of MACE's ability to dampen vibration. The physical requirements for the gimbals are (per axis):

Slew Range:	+/- 60°
Slew Speed:	50°/sec
Acceleration:	125°/sec/sec
Torque:	40 oz-in (normal operation)
	560 oz-in (desired max. stall)
Resolution:	30 arcsec (pointing <u>and</u> tracking accuracy)
Bandwidth:	1-100 Hz

Low unit weight is also desired, with each gimbal system weighing on the order of 15 lbs. The following gimbal system manufacturers have been identified and contacted:

Ball Aerospace: Boulder, CO
Contraves: Pittsburgh, PA
Honeywell: Glendale, AZ
JPL: Pasadena, CA
Versatron: Healdsburg, CA

Each potential vendor received an informal request for quote on January 23, 1990. These letters were FAX'ed only after a series of conversations with company representatives and MIT personnel had clarified our objectives and requirements for the gimbals.

To date, all of the vendors have replied, and JPL has indicated that they wish only to be involved in an advisory capacity. Honeywell did not mail a response, but their rough verbal quote is defined below. The rough

estimates for non-recurring engineering, production, and delivery of three

gimbal systems (two for MACE plus a spare) are:

Vendor	Price	Lead Time
Ball Aerospace	$900,000	14 mos.
Contraves	5,000,000	24 mos.
Honeywell	750,000	-unspecified-
Versatron	320,000	-unspecified-

A.4 Conclusion

The MACE project is an ongoing MIT SERC / NASA endeavour, enjoying the

technical support of the Lockheed Missles and Space Company and the Jet

Propulsion Laboratory. For more information on MACE and the design and

development of MACE test articles, consult the MACE Phase A Final Report

[9] and the MACE Engineering Model Design Document [18].

REPORT DOCUMENTATION PAGE

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE May 1991	3. REPORT TYPE AND DATES COVERED technical report

4. TITLE AND SUBTITLE

Multiple Mode Vibration Suppression in Controlled Flexible Systems

5. FUNDING NUMBERS

N00014-86-K-0685
N00014-86-K-0124

6. AUTHOR(S)

James M. Hyde

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Artificial Intelligence Laboratory
545 Technology Square
Cambridge, Massachusetts   02139

8. PERFORMING ORGANIZATION REPORT NUMBER

AI-TR 1295

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Naval Research
Information Systems
Arlington, Virginia   22217

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

AD-A259444

11. SUPPLEMENTARY NOTES

None

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Distribution of this document is unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Spacecraft, space-borne robotic systems, and manufacturing equipment often utilize lightweight materials and configurations that give rise to vibration problems. Prior research has led to the development of input command pre-shapers that can significantly reduce residual vibration. These shapers exhibit marked insensitivity to errors in natural frequency estimates and can be combined to minimize vibration at more than one natural frequency.

In this work we present a method for the development of multiple mode input shapers which are simpler to implement and produce smaller system response delays than previous multiple mode designs. The new technique involves the direct solution of a group of simultaneous non-linear impulse constraint equations.

(continued on back)

14. SUBJECT TERMS   (key words)

vibration   robotics
oscillation   flexible manipulators

15. NUMBER OF PAGES

150

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED

An MIT/NASA experimental flexible structure, MACE, is employed as a test article for the validation of the direct solution shaping technique. We examine the results of shaper performance tests conducted on linear and non-linear computer models of MACE. Vibration problems caused by system non-linearities are identified and their eradication attempted.

The direct solution shapers are shown to be effective in suppressing multiple mode vibration, even in the presence of mild kinematic and dynamic non-linearities.

Scanning Agent Identification Target