MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 1223   March 1990

Shaping Inputs to Reduce Vibration: A Vector Diagram Approach

William Singhose

Abstract

This paper describes a method for limiting vibration in flexible systems by shaping the system inputs. Unlike most previous attempts at input shaping, this method does not required an extensive system model or lengthy numerical computation; only knowledge of the system natural frequency and damping ratio are required. The effectiveness of this method when there are errors in the system model is explored and quantified. An algorithm is presented which, given an upper bound on acceptable residual vibration amplitude, determines a shaping strategy that is insensitive to errors in the estimated natural frequency. A procedure for shaping inputs to systems with input constraints is outlined. The shaping method is evaluated by dynamic simulations and hardware experiments.

Table of Contents

List of Figures

Chapter 1

Introduction

Vibration is a concern of virtually every engineering discipline. Mechanical engineers continually face the problem of vibration because mechanical systems vibrate when performance is pushed to the limit. The typical engineering solutions to vibration are to design "stiff" systems, add damping to flexible systems, or develop a good controller. Input shaping is another possibility for vibration control that can supplement the above methods. If inputs (velocity, torque, voltage, etc.) are shaped properly, a system will respond with a vibration-free movement [2, 5, 6, 7, 8, 9].

Vaaler and Seering [12] used an unorthodox geometry to build an assembly robot with high stiffness. They used a "four plus two" configuration to achieve the necessary six degrees of freedom. By giving the robot's "right hand" four degrees of freedom and the "left hand" two degrees of freedom, they avoided the flexibility that occurs in most six degree-of-freedom robots that are designed with their axes in series.

Plump, Hubbard, and Bailey [4] examined the use of piezoresistive polymer films to generate additional damping in a structure. Alberts, Hastings, Book, and Dickerson [10] used a thin layer viscoelastic material to obtain passive damping that enhanced system stability.

A great deal of work has been done in the area of feedback control of flexible systems. Cannon and Schmitz [11] examined feedback control with noncolocated endpoint position measurements for a one link flexible robot. Hollars and Cannon [3] compared four different control strategies for a two-link robot with elastic drives.

An early form of input shaping was the use of posicast control by O.J.M. Smith [9]. This technique breaks a step of a certain amplitude into two smaller steps, one of which is

delayed in time. The result is a reduced settling time for the system.

Farrenkopf [2] developed velocity shaping techniques for flexible spacecraft. He showed that control of decoupled modes could be added without exciting vibration. Swigert [5] demonstrated that torque shaping can be implemented on systems which modally decompose into second order harmonic oscillators. Utilizing terminal boundary conditions he generated shaped inputs that excited the vibration during movement, but reduced the vibration to zero when the movement ended. Swigert also showed that vibration from multiple modes could be eliminated by adaptive correction of the control law.

Singer and Seering [7] have shown that residual vibration can be significantly reduced for single mode systems by employing an input shaping method that uses a simple system model and requires very little computation. The system model consists only of the system's natural frequency and damping ratio. Constraints on the system inputs result in zero residual vibration if the system model is exact. When modeling errors exist, the shaped input function does not keep the system vibration at zero, but it does reduce it to a low level that is acceptable for many applications. Extending the method to multi-mode systems is straightforward [6].

The shaping method involves convolution of a desired input with a sequence of impulses to produce an input function that does not cause vibration. Selection of impulse amplitude and timing dictate how well the system performs. Figure 1-1 shows how impulse sequences can be convolved with system inputs to generate shaped inputs. Three-impulse sequences have been shown to yield particularly effective system inputs both in terms of vibration suppression and response time [6]. The shaping method is effective in reducing vibration in both open and closed loop systems.

My work concentrates on generating the impulse sequences to be used in the convolution that produces the vibration-reducing inputs. Most of the work in this text centers on vector diagrams, which are graphical representations of impulse sequences.

Vector diagrams are used to generate and evaluate the vibration-reducing impulse sequences. All sequences in this text will consist of three impulses, although the use of more impulses can be beneficial in some applications. By modifying the constraints used to produce the impulse sequences, a variety of sequences can be generated that give better performance than those reported previously.
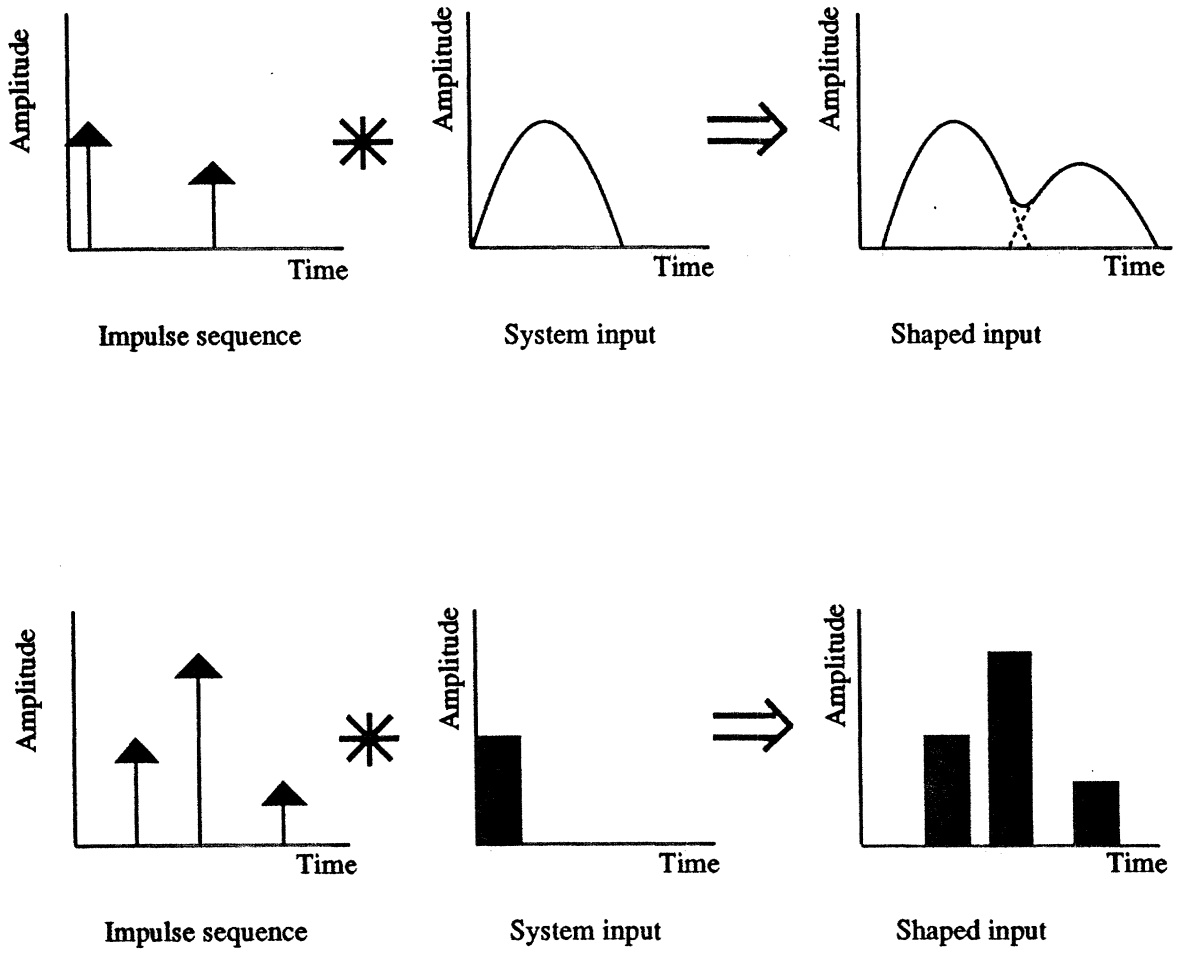
Figure 1-1: Convolving an impulse sequence with a system input produces a "shaped" input.

Chapter 2

Vector Diagrams

A vector diagram is a graphical representation of an impulse sequence. Vector diagrams are graphs in polar coordinates (r-θ space). A vector diagram is created by setting r equal to the amplitude, A_i of the i^{th} impulse in a sequence and by setting $\theta_i = \omega \Delta T_i$, where ω (rad/sec) is an arbitrary frequency and ΔT_i is the time delay from an arbitrary time zero to the time when the i^{th} impulse occurs. Figure 2-1 shows a typical impulse sequence and its corresponding vector diagram.

Vector diagrams become useful tools for producing vibration-reducing impulse sequences when ω is set equal to the natural frequency of a system ($\omega = \omega_{sys}$) and the time of the first impulse is arbitrarily set to zero ($\Delta T_1 = 0$).[1] When a vector diagram is created in this manner, the resultant, R, from summing the vectors on the vector diagram has a special significance. R is proportional to the amplitude of residual vibration of a second order system of natural frequency ω_{sys} driven by a step convolved with the impulse sequence [6]. Because arbitrary inputs can be built as sums of steps, the amplitude of R is a measure of system response for arbitrary inputs. This result enables us to calculate residual vibration geometrically. The length of the resultant on the vector diagram is the amplitude of the vibration and the angle of the resultant is the phase of the vibration relative to the system response from the first impulse. Figure 2-2 compares a vector diagram representation of vibration with a time domain representation of vibration for an undamped system. On a vector diagram, vibration appears as a vector, whereas, in the time domain, vibration appears as a sinusoidal function.

---

[1]In this text, ω will refer to the modeling frequency and ω_{sys} will stand for the actual natural frequency of the system.
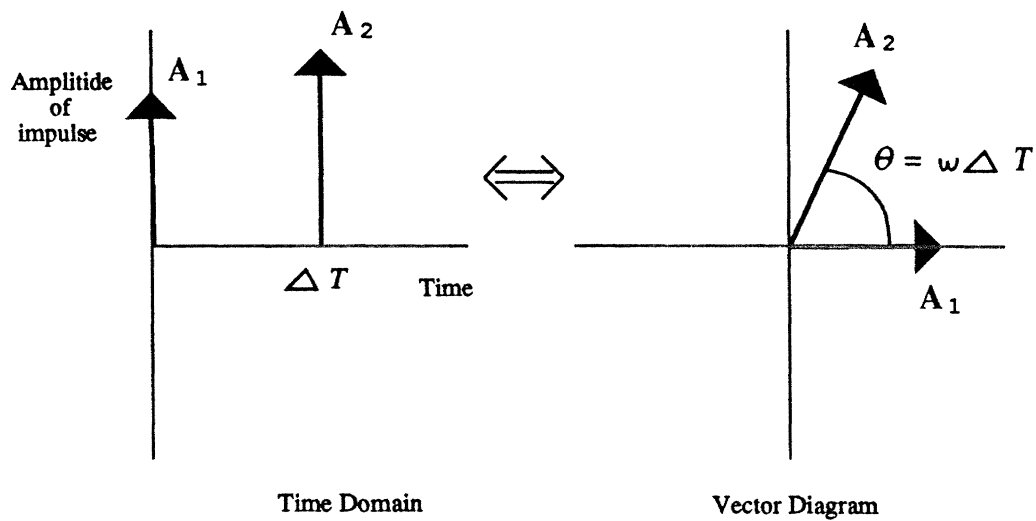
Figure 2-1: An impulse sequence and the corresponding vector diagram.

A_R = The amplitude of the vibration
after impulses A_1 and A_2
have been applied to the system.

Θ_R = The phase of the residual vibration.

vector diagram

response to A_1
response to A_2

Amplitude

A_1

A_2

Time

total response

Amplitude

A_1

A_2

Time

Figure 2-2: The vector diagram and time domain representations of the same vibration.

2.1 Canceling Vibration

If we place N vectors on a vector diagram such that the resultant equals zero, a second-order system of natural frequency ω given the corresponding time domain input will execute a vibration-free movement. We can use this fact to create vibration-reducing input functions directly from a vector diagram. We can place N arbitrary vectors on a vector diagram and then cancel the resultant of the first N vectors with an N+1st vector. When the vectors are converted into an impulse sequence, and the sequence is convolved with a step input,[2] the resulting shaped input will cause no residual vibration when applied to the system. And, if the impulse amplitudes are normalized so they sum to one, the system will come to rest at the point desired by the user.

The canceling vector, A_{n+1}, is given by the equations:

$$|A_{n+1}| = \sqrt{|R_x|^2 + |R_y|^2}$$
$$\theta_{n+1} = \pi + tan^{-1}[\frac{R_y}{R_x}] \tag{1}$$

where:

$$R_x = \Sigma A_i cos\theta_i$$
$$R_y = \Sigma A_i sin\theta_i \tag{2}$$

The above equations demonstrate an interesting fact; there are an infinite number of impulse sequences that will result in a vibration-free response. We can place N arbitrary vectors on a vector diagram and then use Eqs. 1 to find an N+1st vector that will cancel the N original vectors. When the N+1 vectors from the vector diagram are used in the convolution, a vibration-free input function is produced.

---

[2] Or, for that matter, with any desired input function.

2.2 The Effects of Damping

When damping is considered, the vector diagram must be modified in two ways. First, we must use the damped natural frequency of the system to plot the vector diagram. This corresponds to using:

$$\theta = \sqrt{1-\xi^2}\ \omega\Delta T \tag{3}$$

when plotting the vector diagram. Second, the amplitudes of the vectors must be scaled to account for damping. As time progresses, the amplitude of the canceling vector decreases. For example, if we give a system an impulse with amplitude, A_1 at time zero, the impulse that will cancel the system's vibration is located π (180^o) out of phase with the first impulse, but it has a smaller amplitude. Figure 2-3 demonstrates this result in the time domain. The amplitude of the second impulse is [6]:

$$A_2 = A_1 e^{-\xi\omega t} = A_1 e^{-\xi\theta} \tag{4}$$

The effective amplitude of a vector, A_i at ΔT_i is the amplitude of a vector at time zero whose vibration it could cancel. Written in equation form this means that the effective amplitude, $|A_{eff}|$ of a vector A is:

$$|A_{eff}| = \frac{|A|}{e^{-\xi\theta}} \tag{5}$$

This scaling effect of damping can be represented on a vector diagram by superimposing the spiral, $Ae^{-\xi\theta}$. Any vector whose tip lies on the spiral has the effective amplitude of a vector A at time zero. See Figure 2-4. When we attempt to cancel N vectors with an N+1st vector on a vector diagram we must assign each of the N vectors an effective amplitude before using Eqs. 1 to solve for the resultant. When we include the effects of damping, the equations describing the N+1st canceling vector are:

$$|A_{n+1}| = (e^{-\xi\theta}n+1)\sqrt{|R_x|^2 + |R_y|^2}$$
$$\theta_{n+1} = \pi + tan^{-1}[\frac{R_y}{R_x}] \tag{6}$$

where $\theta = \sqrt{1-\xi^2}\ \omega\Delta T$, and R_x and R_y are given by:

$$R_x = \Sigma A_{i_{eff}} cos\theta_i$$
$$R_y = \Sigma A_{i_{eff}} sin\theta_i \qquad\qquad\qquad\qquad (7)$$
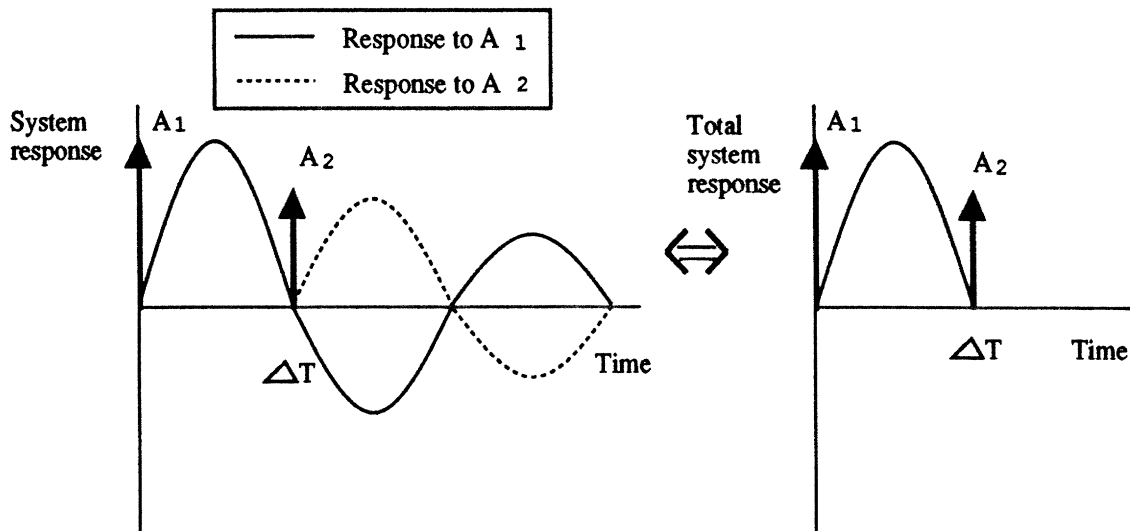

Figure 2-3: The scaling effect of damping on the "canceling" impulse. A_2 is smaller than A_1 because the vibration has been partially damped out after ΔT.

vector diagram

Figure 2-4: The damping spiral, $Ae^{-\xi\theta}$ superimposed on a vector diagram. Each of the vectors shown have the same effective amplitudes.

Chapter 3

Insensitivity

It is possible to create an infinity of vibration-reducing input functions. The "best" would seem to be the one that worked most effectively on a real system; they are all vibration-free when the system model is exact. When the system model is not exact, some residual vibration will occur when the system is moved. A plot of the vibration *vs.* error in estimated natural frequency for a three-impulse sequence developed by Singer [6] is shown in Figure 3-1 along with the corresponding vector diagram. This impulse sequence produces a system response that is fairly insensitive to errors or changes in the system parameters. That is, there is relatively little vibration in the system even when the resonant frequency estimate is off by 20% as shown.

3.1 Effects of Modeling Errors on the Vector Diagram

Figure 3-1 can be obtained directly from a vector diagram if we analyze how a modeling error changes the diagram. When the natural frequency of a system differs from the assumed natural frequency, the error can be represented on a vector diagram by shifting the vectors through an angle ϕ [6]. If ω_{sys} is the actual natural frequency of the system and ω is the modeling frequency, then the error in frequency is $\omega - \omega_{sys}$. The angle through which the vectors are shifted, ϕ, is related to the frequency error by the equation:

$$\phi_i = (\omega - \omega_{sys})\Delta T_i \tag{8}$$

The error in modeling causes a resultant to be formed on the vector diagram; the vectors no longer satisfy Eqs. 6. The resultant that is formed represents the vibration that is caused by the error in frequency.

vector diagram

Figure 3-1: Vector diagram of a three-impulse sequence and its insensitivity curve.

Given that modeling errors cause a resultant, R_{err}, on a vector diagram, we can compare the insensitivity of different input functions by plotting the amplitude of R_{err} vs. the error in frequency. For now, the errors in the damping ratio are ignored because the errors in natural frequency have been shown to be far more important [6].

If we plot an "insensitivity curve" like the one shown in Figure 3-1, we can determine how much vibration will result from a given error in estimated frequency. To make an insensitivity curve, we must develop an expression for the resultant as a function of the error in frequency ($\omega - \omega_{sys}$).

If we subtract the angle due to the error, ϕ from the original angle, θ, then the i^{th} vector on the vector diagram has a total angle of:

$$\theta_{i_{total}} = \theta_i - \phi_i \tag{9}$$

Given this, the amplitude of the resultant is:

$$|R_{err}| = \sqrt{|R_{xerr}|^2 + |R_{yerr}|^2} \tag{10}$$

where:

$$R_{xerr} = \Sigma A_{i_{eff}} cos(\theta_i - \phi_i)$$
$$R_{yerr} = \Sigma A_{i_{eff}} sin(\theta_i - \phi_i) \tag{11}$$

and:

$$A_{i_{eff}} = \frac{A_i}{e^{-\xi(\theta_i - \phi_i)}}$$

$$\phi_i = \frac{\omega - \omega_{sys}}{\omega}\theta_i \tag{12}$$

Eq. 10 is the expression that we were seeking. It gives the amplitude of the resultant as a function of the error in frequency.

3.2 Defining Insensitivity

To compare impulse sequences and determine which is the "best" for decreasing vibration in the presence of modeling errors, we need a formal criterion. Therefore, the insensitivity of a sequence will be defined as the width of the insensitivity curve at a given level of residual vibration. If the acceptable level of vibration is 5%, then we can draw a horizontal line across the insensitivity curve at 5%. The distance between the points of intersection is the insensitivity. For example, the insensitivity of the impulse sequence shown in Figure 3-1 is 0.286, because it causes less than 5% of the step vibration from $(\omega/\omega_{sys})_{lower}= 0.857$ to $(\omega/\omega_{sys})_{upper}= 1.143$.

Now that we have a precise definition for insensitivity, we can compare various impulse sequences quantitatively.

3.3 Increasing Insensitivity by Changing Vectors

The three impulses in the function shown in Figure 3-1 are in the ratio 1:2:1 and are located on the vector diagram at 0, π, and 2π respectively. We know that we can arbitrarily place two vectors on the vector diagram and cancel them with a third, so we can vary the amplitudes and angles of the first two vectors and then cancel the vibration they cause with a third vector. By the definition of a vector diagram, the angle of the first vector is always zero and its amplitude is one. Any change from the value of one will simply scale the second and third vectors accordingly.

If we modify the impulse sequence in Figure 3-1 by placing the second vector at an angle less than π (keeping the amplitude fixed at 2), the insensitivity curve changes in an interesting way. It gets wider and shifts to the right. Figure 3-2 shows the insensitivity curve when the second vector is at 154°. The insensitivity for this input function is 0.408 (0.93 to 1.338), a 43% improvement over the impulse sequence of Figure 3-1. A drawback

is that the insensitivity curve is skewed to the right, i.e., it is more insensitive to errors that are higher in frequency than the modeling frequency. This may be a desired property of an input function if the system being moved increases its natural frequency during some part of its operation [1]. However, for most applications it is desirable to have equal insensitivity on either side of of the modeling frequency $((\omega/\omega_{sys}) = 1)$.

We can shift the insensitivity curve of Figure 3-2 back to the left by choosing a new modeling frequency that is in the center of the skewed insensitivity curve. The modeling frequency is not shifted to get it closer to the actual natural frequency of the system, rather it is shifted to obtain an insensitivity curve that has equal insensitivity on either side of the modeling frequency. When the modeling frequency is shifted to the center of the skewed curve for the above example $(\theta_2 = 154^o)$, the new modeling frequency is: $\omega_{new} = 1.134\omega_{old}$. Adjusting so that the new modeling frequency is at one causes the insensitivity curve to shift to the left and shrink. The shrinkage occurs because the new modeling frequency is larger than the original, so when (ω/ω_{sys}) is recalculated using the larger ω, the difference between $(\omega/\omega_{sys})_{lower}$ and $(\omega/\omega_{sys})_{upper}$ decreases. In the above example, the "true" insensitivity at 5% vibration is 0.36; smaller than the "skewed" insensitivity of 0.408, but still much larger than the insensitivity of the 1:2:1 impulse sequence shown in Figure 3-1.

Figure 3-2: Insensitivity curve when the 2^{nd} vector is placed at 154 degrees and the third vector exactly cancels the first two vectors.

3.4 Increasing Insensitivity by Relaxing the Zero Vibration Constraint

In the cases discussed previously, it was assumed that the residual vibration should be zero when the estimated natural frequency exactly matched the system natural frequency. As we shall see, the relaxation of this constraint can improve insensitivity. We introduce an error at the modeling frequency if we do not exactly cancel the first and second vectors with the third vector, i.e., do not use the exact solution given by Eqs. 6. For most systems, it would seem desirable to have an insensitivity curve which is symmetric about the modeling frequency. So, to generate an error at the modeling frequency and maintain symmetrical insensitivity, we should change the amplitudes of the vectors, but always place them at 0, π, and 2π.

The largest insensitivity that has been discovered for a three-impulse sequence occurs when the error at the modeling frequency exactly matches the vibration limit, V_{lim},[3] and the insensitivity curve falls off to zero on both sides of the modeling frequency. See Figure 3-3. This "hump" in the insensitivity curve widens the curve and, therefore, increases insensitivity.

Using the above conditions, we can derive the three-impulse sequence that yields the maximum known insensitivity for a given vibration limit. The insensitivity curve should be symmetrical about the modeling frequency, so when vectors are shifted by a modeling error or a shift in system natural frequency, the angle of the third vector, Φ_3, is always twice the angle of the second vector, Φ_2. In equation form:

$$\Phi_3 = 2\Phi_2. \tag{13}$$

When the resultant at the modeling frequency is set equal to the vibration limit, V_{lim}, we have:

$$|A_1| - |A_2| + |A_3| = V_{lim}(|A_1| + |A_2| + |A_3|) \tag{14}$$

---

[3]V_{lim} for the above examples was 5%; that is, the allowable residual is 5% of the residual which would have resulted had the system been given a step input.

A negative value is assigned to $|A_2|$ on the left side of Eq. 14 because A_2 points in the opposite direction of A_1 and A_3 on the vector diagram. We have arbitrarily set $|A_1|$ equal to one, so we can rearrange Eq. 14 to get an expression for $|A_2|$:

$$|A_2| = \frac{(1-V_{lim})(1+|A_3|)}{(1+V_{lim})} \qquad (15)$$

Because the insensitivity curve falls to zero on both sides of the modeling frequency, the resultant on the vector diagram equals zero when vector A_2 is at some angle, Φ_2, and vector A_3 is at some angle, Φ_3. This corresponds to:

$$0 = 1 + |A_2|cos\Phi_2 + |A_3|cos\Phi_3 \qquad (16)$$
$$0 = |A_2|sin\Phi_2 + |A_3|sin\Phi_3 \qquad (17)$$

Eqs. 13, 15, 16, and 17 are four equations with five unknowns, $(A_2, \Phi_2, A_3, \Phi_3,$ and $V_{lim})$. If an upper bound on allowable vibration is known, the sequence that yields the maximum insensitivity for the given V_{lim} can be determined. Putting Eq. 13 into Eq. 17 and reducing gives:

$$|A_2| = -2|A_3|cos\Phi_2 \qquad (18)$$

Combining Eq. 15 with Eq. 18 we get $cos\Phi_2$ in terms of $|A_3|$:

$$cos\Phi_2 = -\frac{(1-V_{lim})(1+|A_3|)}{2|A_3|(1+V_{lim})} \qquad (19)$$

If we put Eqs. 13 and 15 into Eq. 16, we get:

$$1 + [\frac{(1-V_{lim})(1+|A_3|)}{(1+V_{lim})}] \, cos\phi_2 + 2|A_3|cos^2\Phi_2 - |A_3| = 0 \qquad (20)$$

Putting Eq. 19 into Eq. 20 and solving for $|A_3|$, gives:

$$|A_3| = 1 \qquad (21)$$

We can use the result of Eq. 21 to solve Eq. 15 for $|A_2|$. When this is done, the result is:

$$|A_2| = \frac{2(1-V_{lim})}{(1+V_{lim})} \qquad (22)$$

Therefore, the three-vector combination that yields the largest known insensitivity for a given vibration limit is:

$$|A_1| = 1, \qquad\qquad \theta_1 = 0$$

$$|A_2| = \frac{2(1-V_{lim})}{(1+V_{lim})}, \qquad\qquad \theta_2 = \pi \qquad\qquad (23)$$

$$|A_3| = 1, \qquad\qquad \theta_3 = 2\pi$$

Converting the above vectors into the time domain we find that the three-impulse sequence that yields the largest known insensitivity for a given vibration limit is:

$$A_1 = 1, \qquad\qquad \Delta T_1 = 0$$

$$A_2 = \frac{2(1-V_{lim})}{(1+V_{lim})}, \qquad\qquad \Delta T_2 = \frac{\pi}{\omega} \qquad\qquad (24)$$

$$A_3 = 1, \qquad\qquad \Delta T_3 = \frac{2\pi}{\omega}$$

Figure 3-3 shows the insensitivity curve for the above impulse sequence when the vibration limit is set to 5%. For larger vibration limits, insensitivity increases significantly as Figure 3-4 demonstrates. The insensitivity increases from 0.398 to 0.56 when V_{lim} is changed from 5% to 10%.

When the driven system has viscous damping the amplitudes in Eqs. 24 must be replaced by effective amplitudes and the damped natural frequency must be used. The sequence then becomes:

$$A_1 = 1, \qquad\qquad \Delta T_1 = 0$$

$$A_2 = \frac{2(1-V_{lim})}{(1+V_{lim})}e^{-\pi\xi}, \qquad\qquad \Delta T_2 = \frac{\pi}{\sqrt{1-\xi^2}\,\omega} \qquad\qquad (25)$$

$$A_3 = e^{-2\pi\xi}, \qquad\qquad \Delta T_3 = \frac{2\pi}{\sqrt{1-\xi^2}\,\omega}$$

where ω is the modeling frequency and ξ is the approximate damping ratio of the system.

Figure 3-3: Insensitivity curve for vibration limit of 5%.

Figure 3-4: Insensitivity curves for vibration limits of 5% and 10%.

Chapter 4

Input Shaping

The following steps should be taken when shaping a system input:

1. Determine the system's natural frequency and damping ratio. (The insensitivity provided by this method allows for substantial errors in the measurement of these two parameters).

2. Determine the impulse sequence for the system by using Eqs. 25.

3. Normalize the amplitudes of the impulses so they sum to one.

4. Convolve the normalized impulse sequence with any desired input.

5. Use the result of the convolution as the input to the system.

The shaped input causes the system to arrive at the position it would have if the desired input had been used. However, when the movement is finished, the system's vibration will be much less than if the desired input had been used. In addition, the vibration will be less than the established vibration limit, provided the system model is within the interval shown on the insensitivity curve. The cost of the improved performance of the system is a time penalty equal to the period of the system's natural frequency. (The shaped input takes one period of the natural frequency longer to execute than the unshaped input). For most systems, however, the time penalty is not a real penalty because the vibration caused by an unshaped input usually takes more than one period to decay. A system can actually be moved faster with shaped inputs because at the end of the input sequence, the system is *at* the desired position rather than oscillating *about* the desired position.

4.1 Digital Systems

The impulse sequence given by Eqs. 25 was derived based on the assumption that the amplitude of the input can be changed at any time. This is not possible with digital systems because the inputs can be changed only at discrete time increments. The impulse sequence of Eqs. 25 can, however, be transformed into an equivalent sequence where the impulses occur only at the discrete time steps. The transformation is straightforward and is discussed in detail in reference [6]. Basically, the transformation works by replacing an impulse that is not located at a discrete time step with two impulses that are located at the digital time steps on either side of the impulse to be replaced. The two impulses are chosen so as to be equivalent to the impulse they replace.

4.2 Constant Amplitude Inputs

When a system is driven by constant amplitude inputs; for example, when the space shuttle is moved by its stabilizing thrusters; the shaped input cannot be implemented in its original form. (The result of the convolution cannot be used because it requires varying the amplitude of the input). Although we cannot get all of the vibration-reducing ability of this input-shaping method into a control system that uses constant amplitude inputs (pulses), there is an approximate method that works well when the shortest possible driving pulse is small compared to the period of the system's natural frequency. When this is the case, the result of the convolution[4] can be modified by converting the amplitude of the three pulses into the width of the three pulses. The pulse with the smallest amplitude should be assigned the shortest width possible. The widths of the other pulses are then adjusted according to the height to width ratio used to convert the smallest pulse.

---

[4]The convolution of the impulse sequence with the shortest possible driving pulse will produce three pulses of unequal amplitudes separated by some amount of time. The greater the separation between the pulses, the better the approximate method will work.

Once the three constant-amplitude pulses have been determined, their centers should be placed at the time spacings given by Eqs. 25. When this is done, one half of the first pulse will exist in negative time, a simple time shift will yield a three-pulse sequence that moves the system with much less vibration than would occur if a single pulse had been used to achieve the same movement.

The approximate method for constant amplitude inputs can be summarized as follows (see Figure 4-1):

1. Convolve the impulse sequence of Eqs. 25 with a pulse equal to the shortest possible driving pulse.

2. Modify the result of the convolution by converting the amplitude of each pulse into a proportional width.

3. Center the three constant amplitude pulses around the times given by Eqs. 25.

4. Multiply the width of the three pulses by a proportionality constant to achieve the length of movement desired.

5. Shift the pulses to place the first pulse at zero time.

The above method is a simple and straightforward process for dealing with the restriction of constant amplitude inputs. The method significantly reduces vibration when the input pulses can be made much shorter than the period of the system's natural frequency. The next chapter gives the results of dynamic simulations that demonstrate the above method's vibration reducing ability.

Figure 4-1: Approximate shaping method for systems with constant amplitude inputs.

Chapter 5

Experimental Results

The performance of the impulse sequence given by Eqs. 25 was verified by dynamic simulations and with hardware experiments.

5.1 Simulation of Single Mode Control

The ability of the shaping method to reduce vibration was shown by implementing it on the simple two-mass lumped parameter model shown in Figure 5-1. In the first series of tests, the following values were used for the parameters:

$$M_1 = M_2 = 1$$

$$k = 19.74$$

$$b = 0.$$

Figure 5-1: Simple lumped parameter model used for simulating a single-mode system

The system was moved by giving a step in the force, F(t), on M_1. The amplitude of the step was one and it was applied from 0.5 sec. to 0.7 sec. The step input and the corresponding system response is shown in Figure 5-2.

The following five steps were then taken to shape the step input so the residual vibration of the system would be reduced to five percent of the vibration that resulted from the step input.

1. The system's natural frequency was determined to be 2π. [5]

2. The values of $V_{lim} = 0.05$ and $\omega = 2\pi$ were put into Eqs. 24, giving an impulse sequence of:

$$A_1 = 1, \qquad\qquad\qquad \Delta T_1 = 0$$

$$A_2 = 1.81, \qquad\qquad\qquad \Delta T_2 = 0.5 \qquad\qquad (26)$$

$$A_3 = 1, \qquad\qquad\qquad \Delta T_3 = 1$$

3. The amplitudes of the impulses were normalized to give: $A_1 = 0.262$, $A_2 = 0.475$, $A_3 = 0.262$.

4. The normalized sequence was convolved with the step input shown in Figure 5-2.

5. The result of the convolution was used as an input to the system.

The shaped input and the corresponding system response are shown in Figure 5-3. By comparing Figures 5-2 and 5-3 we can see that the shaped input reduced the residual vibration just as predicted. Figure 5-3 also shows the time penalty that is incurred when the input is shaped. The time penalty is equal to one period of the natural frequency (in this case, 1 sec). The step input ends at 0.7 sec., but the shaped input ends at 1.7 sec.

By changing the value of V_{lim} to zero and repeating steps 2-5, the residual vibration can be reduced to zero as shown in Figure 5-4. The sequence that is produced by setting V_{lim} equal to zero is the same as Singer's three-impulse sequence [6].

---

[5]

$$\omega = \sqrt{\frac{2k}{M}}$$

F(t) (N)

1.0

0.5

0.5   1.0   1.5   Time (sec)

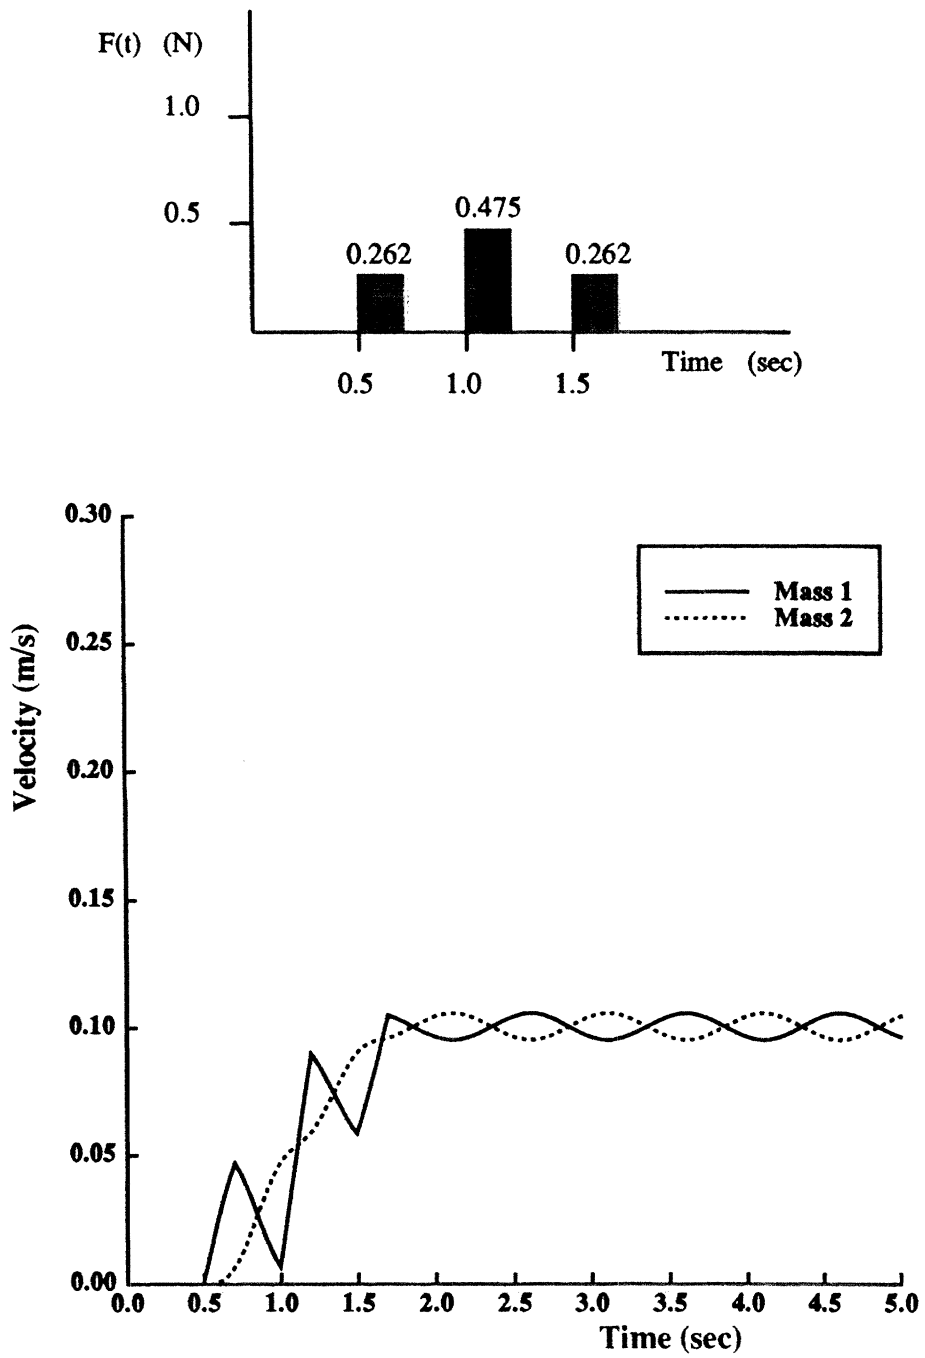Figure 5-2: A step input and the corresponding response of
the single-mode system.

Figure 5-3: The shaped input when $V_{lim} = 0.05$ and the corresponding response of the single-mode system.

Figure 5-4: The shaped input when $V_{lim} = 0$ and the corresponding response of the single-mode system.

5.1.1 Single Mode With Damping

Damping was added to the dynamic simulation by setting the damper of the lumped parameter model equal to 0.5. Figure 5-5 shows the response of the damped system to a step input. The damped natural frequency and damping ratio were calculated and then put into Eqs. 25 to produce the vibration reducing impulse sequence. V_{lim} was set to zero and the impulse sequence was convolved with the step input. Figure 5-6 shows the shaped input and the corresponding system response.

5.2 Energy Consumption of Shaped Inputs

In addition to reduced settling time, input shaping has another benefit: it saves energy. Qualitatively, it can be reasoned that because it takes energy to excite vibration, the shaped inputs require less energy than unshaped inputs. We can calculate the energy savings by integrating the force times velocity with respect to time. When this is done for the movement of the damped system above, we find that the step input requires approximately 0.018 joules, but the shaped input uses only about 0.010 joules. The shaped input requires less energy because it decreases the energy loss in the damper.[6] A plot of the energy used by the inputs *vs.* time is shown in Figure 5-7.

---

[6]The average velocity difference across the damper is less when the input is shaped.

Figure 5-5: A step input and the corresponding response of the damped single-mode system.

F(t) (N)

1.0

0.4923

0.5

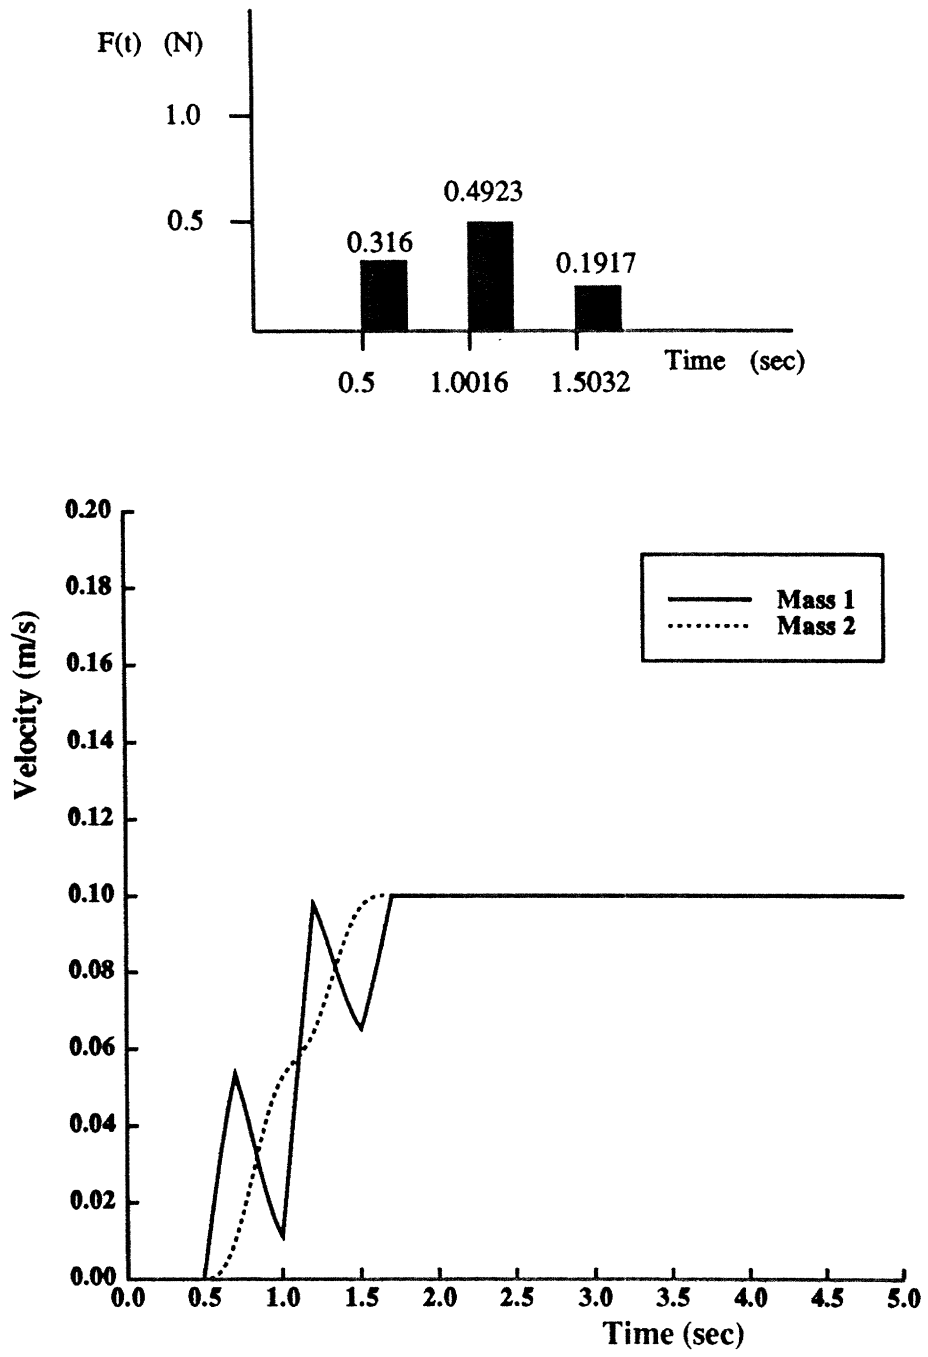0.316

0.1917

0.5 1.0016 1.5032

Time (sec)

Figure 5-6: The shaped input when $V_{lim} = 0$ and the corresponding response of the damped single-mode system.

Figure 5-7: The energy savings of shaped inputs.

5.3 Simulation of Constant Amplitude Input Control

When a system can only be driven by constant amplitude inputs, the shaped input that results from the convolution cannot be used. (This was discussed in the previous chapter.) The approximate method proposed in the last chapter was tested on a computer simulation developed to mimic the rotational dynamics of the space shuttle deploying the Hubble space telescope. The dynamic model is shown in Figure 5-8. The model consists of two unequal rotational inertias connected by a torsional spring. The large inertia (the shuttle) is acted on by a small constant torque (a gravity gradient) and it can be rotated in the opposite direction by a large constant amplitude torque (a stabilizing jet). The values for the system parameters were obtained from the DRS and SDAP simulations that Draper Laboratory uses to verify and plan shuttle missions.

Figure 5-8: Rotational inertia dynamic model.

The control action of the dynamic model was designed to copy the actions of the shuttle's digital autopilot (SDAP). Basically, the simulation attempts to maintain the large rotational inertia within a given angular region. When the inertia is rotated out of the

region by the gravity torque, a constant amplitude torque is applied in the opposite direction to drive the inertia back into the desired region. This would correspond to the firing of a stabilizing jet to reorient the shuttle. The reorienting torque causes large rotational oscillations in the system. Figure 5-9 shows the typical dynamic response of the model to a reorienting input. By 280 sec. the gravity torque has caused the large inertia to rotate out of the desired angular region. A large reorienting torque is applied to drive the inertia back into the desired region.

The procedure for shaping constant amplitude inputs that was outlined in the previous chapter was used to develop a shaped input for the above simulation. The system response to the shaped input is shown in Figure 5-10. A comparison between Figure 5-9 and Figure 5-10 reveals that the shaped input significantly reduces the oscillations in the system. The results shown in Figure 5-10 are impressive because the reorienting torque could be turned on and off at any time. However, when the system has a large digital time step the approximate shaping method is not as effective because the width of the input pulses cannot be set arbitrarily. Developing a method for dealing with constant amplitude inputs and large digital time steps is an area of ongoing research.

Inertias (N*m2)=8953000.000000 56470.000000
Rotational spring = 10000.000000
Rate limit (deg/sec)=   0.200000
Jet force (N)= -3860.000000
Moment arm for jet (m)=   3.800000
Gravity torque (N*m)=   20.000000

Figure 5-9: Response of rotational inertia model to unshaped input.

Inertias (N*m2)=8953000.000000 56470.000000
Rotational spring = 10000.000000
Rate limit (deg/sec)=   0.200000
Jet force (N)= -3860.000000
Moment arm for jet (m)=   3.800000
Gravity torque (N*m)=   20.000000

Figure 5-10: Response of rotational inertia model to shaped input.

5.4 Hardware Experiments

The hardware verification of the impulse sequence was performed on the assembly robot described in reference [12]. The experimental setup is shown in Figure 5-11. A steel beam with a mass at one end was attached to the turntable of the robot. The table was driven in the theta direction by a dc motor and its position was determined by an optical encoder.

When the table was given a step in position, large oscillations were induced in the system. Figure 5-12 shows a typical system response to a step input. Data from ten step responses were recorded and the vibration amplitudes were averaged to get a baseline value for the vibration caused by a step input. The amplitude of the vibration was determined by finding the highest value in the data and then subtracting the lowest following value.

The system parameters were determined by examining the data from a step response. The natural frequency was found to be 2.8 Hertz and the damping ratio was approximated as 0.05. The step input was then shaped by the impulse sequence given by Eqs. 25 with the vibration limit set to 5%. Figure 5-13 shows the system response to the shaped input. The amplitude of the vibration in Figure 5-13 is only 5% of the baseline value for a step input. The residual vibration of the system can be virtually eliminated by setting the vibration limit to zero. See Figure 5-14.

5.4.1 Insensitivity Curves

Insensitivity curves for the impulse sequence were experimentally determined by purposely introducing errors in the system model. The experimentally determined natural frequency was chosen as the "exact" frequency (ω_{sys} = 2.8 Hz). Impulse sequences were then derived for frequencies ranging from $0.6\omega_{sys}$ to $1.8\omega_{sys}$ (1.8 Hz - 5.04 Hz). The value for the damping ratio was kept at 0.05 for all cases. Each impulse sequence was used

Figure 5-11: The experimental setup: Robot turntable and steel beam with a large mass attached to the end.

**Figure 5-12:** Typical system response to a step in position.

**Figure 5-13:** System response to a shaped step input when $V_{lim} = 0.05$.

Figure 5-14: System response to a shaped step input when $V_{lim} = 0$.

to shape the step command and the shaped input was applied to the system. The amplitude of the resulting vibration was recorded. The amplitude value was then divided by the baseline value to get the percentage of the step vibration caused by the shaped input. By plotting the percentage of the step vibration versus the normalized modeling error (ω/ω_{sys}), insensitivity curves were obtained.

Figure 5-15 shows the experimentally determined insensitivity curve when the vibration limit was set to 5%. Figure 5-15 shows the experimental curve has the same general shape as the theoretically determined insensitivity curve (Figure 3-3). The curve has a non-zero value when the system model is correct and it slopes down toward zero on either side of the modeling frequency. The experimentally determined insensitivity for V_{lim} = 0.05 is approximately 0.42. (The vibration is less than 5% from 0.82 to 1.24 on the insensitivity curve). The theoretical insensitivity when V_{lim} equals 0.05 is 0.398.

The vibration limit was set to 10% and the above tests were repeated. Figure 5-16 shows the resulting insensitivity curve. The insensitivity increased to approximately 0.67 (The vibration was less than 10% from 0.75 to 1.42 on the insensitivity curve). The theoretical insensitivity when V_{lim} equals 0.10 is 0.56.

The roughness of the insensitivity curves is largely due to the resolution of the optical encoder. The vibrations resulting from the shaped inputs are only a few ticks on the encoder. For example, a vibration that is 5% of the baseline value for a step is only three encoder spacings. Another limitation on the accuracy of the experimental tests is the servo rate (1 kHz) of the robot controller. The shaped input that results from the convolution does not usually change at exactly the same time that the servo loop is called. As a result, the change in input must occur during the next call of the servo loop. This error causes the experimental value for the vibration amplitude to differ slightly from the theoretical value.

**Figure 5-15:** Experimentally determined insensitivity curve for $V_{lim} = 0.05$

Figure 5-16: Experimentally determined insensitivity curve for $V_{lim} = 0.10$

Chapter 6

Discussion

It has been shown that a variety of impulse sequences can be convolved with a desired system input to create a shaped input that moves a system without causing residual vibration. By plotting the residual vibration of the system *vs.* the error in estimated natural frequency, we can determine how insensitive the impulse sequence is to shifts or errors in the natural frequency of the system. Some of the impulse sequences cause a skewed insensitivity, so the input function is more insensitive to errors in one direction. When an upper bound on the acceptable level of residual vibration is known, the three-impulse sequence that gives the maximum known insensitivity can be determined.

When the system inputs are constrained to one amplitude, an approximate method can be used to shape the input by varying the width of the input instead of the amplitude. The approximate method works well when the constant amplitude input can be turned on or off at any time. When the constant amplitude constraint is combined with a large digital time step, the shaping method becomes difficult to implement.

Computer simulations and hardware experiments verified the vibration-reducing ability of the impulse sequences. The hardware experiments showed that the impulse sequences were more insensitive than the theory predicts. The theory is based on the assumption that the system is a second-order damped harmonic oscillator. It is possible that nonlinear affects, such as stiction, caused the impulse sequences to be more insensitive in practice than in theory. The limitations of the hardware, such as the optical encoder spacing and the motor servo rate, may have also contributed to the high value for the experimentally determined insensitivity.

References

[1]   Christian, Andrew D.
       Design and Implementation of a Flexible Robot.
       Master's thesis, Massachusetts Institute of Technology, August, 1989.

[2]   Farrenkopf, R. L.
       Optimal Open-Loop Maneuver Profiles for Flexible Spacecraft.
       Journal of Guidance and Control (Vol. 2, No. 6), Nov.-Dec., 1979.

[3]   Hollars, M. G.; Cannon, R. H.
       Experiments on the End-Point Control of a Two-Link Robot With Elastic Drives.
       AIAA Paper 86-1977 , 1986.

[4]   Plump, J. M.; Hubbard, J. E.; Bailey, T.
       Nonlinear Control of a Distributed System: Simulation and Experimental Results.
       Journal of Dynamic Systems, Measurement, and Control , June, 1987.

[5]   Swigert, C. J.
       Shaped Torque Techniques.
       Journal of Guidance and Control (Vol. 3, No. 5), Sep.-Oct., 1980.

[6]   Singer, Neil C.
       Residual Vibration Reduction in Computer Controlled Machines.
       PhD thesis, Massachusetts Institute of Technology, February, 1989.

[7]   Singer, Neil C.; Seering, Warren P.
       Design and Comparison of Command Shaping Methods for Controlling Residual
            Vibration.
       IEEE International Conference on Robotics and Automation (Scottsdale, Arizona;
            May 14-19), 1989.

[8]   Singer, Neil C.; Seering, Warren P.
       Preshaping Command Inputs to Reduce System Vibration.
       Journal of Dynamic Systems, Measurement, and Control , June, 1990.

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

[10]   Alberts, T. E.; Hastings, G. G.; Book, W. J.; and Dickerson, S. L.
       Experiments in Optimal Control of a Flexible Arm with Passive Damping.
       Fifth VPI&SU/AIAA Symposium on Dynamics and Control of Large Structures
            (Blacksburg, VA; June 12), 1985.

[11]   Cannon, Robert H.; Schmitz, Eric.
       Initial Experiments on the End-Point Control of a Flexible One Link Robot.
       The International Journal of Robotics Research (Vol. 3, No. 3), Fall, 1984.

[12]   Vaaler, Erik; Seering, Warren P.
       Design of a Cartesian Robot.
       Presented at the winter annual meeting of the ASME , 1986.