ADAPTIVE FEEDFORWARD CONTROL FOR ACTIVELY ISOLATED SPACECRAFT PLATFORMS

Eric H. Anderson
CSA Engineering, Inc.
2850 West Bayshore Road, Palo Alto, CA 94303-3843 USA
Jonathan P. How
Dept. of Aeronautics and Astronautics
Stanford University, Stanford, CA 94305 USA

Abstract

Active vibration isolation systems are being considered to improve the performance of spacecraft instruments and sensors. Because of uncertainties inherent in on-orbit operation, adaptive control strategies and algorithms have relevance to these systems. In this paper, analysis of the algorithms, numerical simulation, and laboratory test data are used to evaluate adaptive feedforward control. Of particular interest are performance characteristics and limitations of the filtered-x LMS (FXLMS) algorithm and its finite impulse response (FIR) filter implementation. Combination feedback/feedforward control and the Augmented Error algorithm are two means investigated to extend the capabilities of FXLMS by desensitizing the algorithm to the specific dynamics of the plant. Several experiments were conducted on a laboratory testbed which serves as the prototype for a planned active vibration isolation flight demonstration.

Adaptive Vibration Isolation

The jitter requirements for certain spacecraft sensors and instruments are stringent, resulting in submicron motion specifications. In the presence of reaction wheels, motors, and other articulating devices, the level of vibration on the satellite may exceed requirements. Vibration isolation is one means of reducing the effect of motion at critical locations. A well-designed isolation system would begin with passive control and add active and finally adaptive control as needed to counter specific disturbance environments or meet particular instrument specifications.

Vibration isolation systems separate vibration sources from vibration-sensitive components. There are two basic types of isolation, classified by the component being isolated: 1) instrument or payload isolation and 2) base isolation. Both options have potential use on-board spacecraft. The immediate motivation for this paper is payload isolation.

Most isolators in use are compliant passive systems which support vibrating machines. Passive isolators effectively reduce transmission of high frequency energy, but are ineffective below the suspension frequencies of the isolated system. These mounts are also difficult to tailor to narrowband applications involving transmission of single or multiple tones. One extension is adaptive passive isolators — inherently passive devices adjusted periodically by an external control system. Examples include stiffening a mount to counter maneuvering loads in a vehicle, and then softening the mount during cruise. Adaptation in this case adds versatility, but the approach is still constrained by fundamental performance limitations associated with passive isolation.

Fully active isolation systems introduce fundamental advantages over their passive or adaptive passive counterparts. The transmissibility of the isolator can be tailored to selectively attenuate important inputs without the passive constraint relating isolation corner frequency and static deformation. Overall performance improvements can be significant. Active isolation is already common in terrestrial applications, but this paper is guided by the potential use on spacecraft, where the flight heritage is short and the number of applications is increasing. In one case, a three-axis active isolation system was demonstrated for a vibrating cryocooler [1]. Other active systems have been developed for microgravity isolation within modules of manned vehicles, on the Shuttle Orbiter [2] and the Space Station [3]. More recent systems [4] will isolate

*eric@csaenr.com; 415-494-7351; www.csaengineering.com;
howjo@sun-valley.stanford.edu
©1997 by Eric H. Anderson and Jonathan P. How

specific spacecraft instruments.

The disturbances on-board operational spacecraft include reaction wheels, thrusters, solar array drives, motors, solar arrays, cryocoolers, and other instruments. Their input may be large enough to disrupt operation of the more sensitive instruments and sensors onboard satellite imaging and communications systems. An active mount can isolate either the disturbance generator or the sensitive equipment.

Unfortunately, the benefits of active isolation are partially offset by the potential stability and performance difficulties associated with implementing the controllers in a space environment that is challenging to simulate accurately on the ground. Furthermore, the disturbance environment is often poorly characterized. However, many of these issues can be overcome using adaptive control to modify an active control structure, function, or form to improve performance based on a new understanding of the system dynamics or properties (See Table 1). Adaptive control makes possible the full performance benefits of an active isolation system.

The application of adaptive control to vibration suppression in general is increasing, driven by the availability of low cost digital signal processors (DSPs). Vibration isolation is a more tractable problem than general vibration control because of the potentially lower order dynamics and inherent limitations on the number of physical paths for energy flow.

Sommerfeldt [5] approached adaptive isolation from an active noise control perspective. His initial laboratory demonstration used an off-line FIR filter to model the secondary path dynamics and an on-line FIR filter for control. Performance was demonstrated for single tones, two tones, and limited broadband frequency regions. This was later replaced with on-line estimation of the FIR models of MIMO systems [5]. In the simplest cases, he demonstrated that use of a pure lag to model the secondary path was adequate. The primary emphasis of his work was on the interaction between the on-line estimation and FXLMS algorithms.

Spanos *et al.* [6] demonstrated isolation on a six-strut laboratory hexapod system and are participating in a flight demonstration system [4]. Beginning from a feedback background, and then moving toward FXLMS algorithms, their approach makes use of accurate identified models of the inverse plant dynamics. Haynes *et al.* [7] used an adaptive FIR feedforward controller to minimize the acceleration error while simultaneously running a second FIR filter to model the plant dynamics. Relatively high order FIR filters were used to achieve an average of 15 dB attenuation between 50-250 Hz. Other promising recent work has applied a neural framework to vibration suppression

Table 1: Applications for adaptive active isolation.

Physical Effect	Rate of Variation	Examples
Time-varying dist.	low-high	Reaction wheels, solar arrays
Base/component modes	low	Flexible satellite
Inertia changes	low-med	New instrument, moving stage
Changes in reqd. perf	low-high	Launch/on-orbit pointing
Subsystem failure	low-high	Transducer, spring
Nonlinearities	low	Aging, friction

problems [8].

This paper supports the separate development and testing of a planned passive/active/adaptive isolation space flight experiment. While the Satellite Ultraquiet Isolation Technology Experiment (SUITE) development continues, we seek to address more broadly applicable areas of adaptive vibration isolation. The work here parallels some of the developments of Sommerfeldt and Spanos, but we also explore both feedforward and feedback control, slewing disturbances, and the potential role and benefits of new adaptive feedforward control algorithms. Adaptive feedback and adaptive feedforward are the two broadest classes of control with potential application to the problem. Adaptive feedback is the more general formulation, but typically depends more on an accurate system model, and places greater computational demands on the physical control system. The main advantages cited for a feedforward system are the straightforward physical implementation and enhanced stability properties. We discuss several connections between the feedforward vibration isolation architecture and the adaptive signal processing framework.

The remainder of the paper focuses on adaptive feedforward control, beginning with a discussion of LMS-based algorithms. The limitations of finite impulse response (FIR) filters for control of lightly damped systems are reviewed. The hardware testbed used for experimental demonstration is described. Finally, results are presented and implications of these results for application to spacecraft vibration isolation are assessed.

Adaptive Feedforward Algorithms

The Least Mean Square algorithm has a long heritage in the fields of adaptive signal processing and control, and thus need not be developed in detail here. However, there are several features of this algorithm that

are of particular interest because they strongly influence the resulting compensator and the implementation, and these are discussed below.

The standard form of the LMS algorithm is written as a feedforward compensator H with finite impulse response (FIR) filter coefficients h_i (see Figure 1). The objective is then to select N filter weights to generate a control signal u to cancel (through S) the disturbance created by the input x through the primary path P . Thus, for the SISO case, the control signal at time k is generated from the filter

$$u_k = \sum_{i=0}^N h_i(k) x_{k-i} = \mathbf{h}_k^T \mathbf{x}_k \quad (1)$$

which includes the vector of weights $\mathbf{h}_k = [h_0(k) \cdots h_N(k)]^T$ and a vector of filter inputs $\mathbf{x}_k = [x_k \cdots x_{k-N}]^T$. Typically, $\mathbf{x}_k$ is obtained from the input x through a tap delay line.

As discussed in detail in Widrow and Stearns [9], the filter weights $\mathbf{h}_k$ are selected to minimize the instantaneous mean-square value of the measured error signal e_k . The optimal weight selection is then given by the update equation

$$\mathbf{h}_{k+1} = \mathbf{h}_k - \mu e_k \mathbf{x}_k. \quad (2)$$

It is well known that the convergence parameter μ significantly impacts the performance of this adaptive system, as it determines the bandwidth of the adaptation process. Note that the resulting compensator (Eqs. 1 and 2) is a linear function of the error signal, but a nonlinear function of the disturbance x , which complicates the stability and performance robustness analysis [10, 11].

A variation on this problem, called filtered-x (FXLMS), uses an estimate of the secondary path $\hat{S}$ to generate a filtered version of the reference signal $\tilde{x} = \hat{S}x$ that is used to drive the filter update equations [12]. With the assumption that the filter update process is relatively slow, the FXLMS algorithm has the same control output given in Eq. 1, but the update equation for the filter weights is of the form

$$\mathbf{h}_{k+1} = \mathbf{h}_k - \mu e_k \tilde{\mathbf{x}}_k; \quad \tilde{\mathbf{x}}_k = [\tilde{x}_k \cdots \tilde{x}_{k-N}]^T. \quad (3)$$

Note that these algorithms were originally developed as feedforward systems, with the architecture shown in Fig. 1(a). The interpretation in this case is that the measured reference signal x is fed forward, through the adaptive filter in Eqs. 1, 3 and the secondary path S to generate a signal that cancels the output of the primary path P . However, with the error signal e used to update the (FX)LMS weights, there is an implicit feedback loop in the overall system, the loop gain of which is determined by μ in Eq. 3. This feedback loop is made explicit in an equivalent formulation shown

in Fig. 1(b) [13, 14]. This approach consists of a relatively simple feedback compensator G_c , the coefficients of which can be adaptively tuned using a frequency following circuit [13]. The feedback approach is relatively easy to implement for harmonic disturbances because simple closed form expressions are available for G_c .

It is possible to gain much further insight into these adaptive algorithms by analyzing the case of a stationary harmonic disturbance at frequency ω_0 . For this case, several authors have shown that the LMS and FXLMS controllers have the LTI form [14, 11]

$$G_{\text{LMS}}(s) = \mu_c A_d^2 \frac{s}{s^2 + \omega_0^2} \quad (4)$$

$$G_{\text{FXLMS}}(s) = \mu_c A_d^2 \hat{M} \frac{\cos \hat{\phi}_0 s + \sin \hat{\phi}_0 \omega_0}{s^2 + \omega_0^2} \quad (5)$$

where $\hat{S}(j\omega_0) = \hat{M}e^{j\hat{\phi}_0}$ models the secondary path dynamics at the disturbance frequency, μ_c is the convergence parameter for the continuous system, and A_d is the amplitude of the harmonic signal. Note that the form of the controllers in Eq. 4 shows that μ determines the overall compensator gain. Also note that these compensators have infinite gain (theoretically) at the disturbance frequency ω_0 , independent of the value of $\mu > 0$. This indicates that μ should not influence the depth of the notch filter introduced to cancel the harmonic disturbance. However, since μ controls the loop gain, it will influence the width (bandwidth) of the notch, which is determined by the frequency range about ω_0 for which the loop gain is greater than one. Thus the feedback interpretation of the role of μ is consistent with the feedforward one.

The FXLMS formulation has proven to be very useful for the acoustic control problem since it is often difficult to accurately model the system gain and phase, even at one frequency point. However, several issues arise when using these approaches to adaptively control isolation systems:

1. If additional modeling information is available, how could this be included in the formulation?
2. Can the algorithm be reformulated to avoid some of the approximations in the FXLMS approach?
3. Can the adaptive controller design be cast in a more formal setting to enable additional extensions?

The following outlines recent developments in the field that address these points by providing an overview of the *Augmented Error* algorithm and a comment on the $\mathcal{H}_\infty$ formulation of LMS.

To compare the augmented error algorithm with FXLMS and LMS, we first rewrite these algorithms

Figure 1: Terminology for the various interpretations of the adaptive controllers.

in a consistent form. Following the derivation of Bayard [11], the problem is to cancel a disturbance signal $y(t)$ that is assumed to be related to the measured signal $x(t)$ through an optimal, fixed set of N weights w^o , i.e. $y(t) = w^{oT}x(t)$, where $x(t) \in \mathcal{R}^N$ is a *regression* vector taken from the measured data. The objective is to tune the weights of the approximate filter $w(t)$ to cancel the disturbance using the output from the adaptive controller G_{AE} given by $\hat{y}(t) = w(t)^T x(t)$. Define $\phi(t) = w^o - w(t)$, and write the error as

$$e = S(s)[y - \hat{y}] = S(s)[\phi^T x] \quad (6)$$

where $[\cdot]$ implies that the filtering operation is applied to each element of the vector. The block diagram for this formulation is shown in Fig. 1(c). Note that $y(t)$ enters as an input disturbance to the system dynamics (which actually corresponds to the secondary path). By redefining the weights w^o it is possible to rewrite this problem in the more standard form with an output disturbance.

As discussed, the standard assumption with the FXLMS algorithm is that the weights change slowly so that the order of the filtering and adaptation can be interchanged yielding

$$e = S(s)[\phi^T x] \approx \phi^T \tilde{x} \quad (7)$$

where $\tilde{x} = S(s)[x]$. The associated stabilizing weight update algorithm is then the same as before

$$\dot{w} = \mu_c e \tilde{x} \quad ; \quad \hat{y} = w^T x. \quad (8)$$

Monopoli [15] introduced a technique which includes extra terms in the error, and as a result does not need to invoke this assumption. The technique, called *Augmented Error*, is well known in the design of adaptive feedback controllers [16], but, until recently, has not been applied to the adaptive feedforward case [11]. As the name suggests, this approach hinges on adding extra terms to the error $e(t)$ used in the LMS and FXLMS algorithms to develop the augmented error

$$e_a = e + \phi^T S(s)[x] - S(s)[\phi^T x] = \phi^T \tilde{x} \quad (9)$$

$$= e + S(s)[\hat{y}] - \tilde{y} \quad (10)$$

where e is measured, $\hat{y}$ is defined above, and $\tilde{y} = w^T \tilde{x}$. The last term in Eq. 9 cancels the contribution of the error $e(t)$, leaving a form for the augmented error that is similar to the FXLMS error in Eq. 7. However, Eq. 10 shows that the augmented error can be computed without knowledge of the optimal weights, and thus can be determined in real-time. Furthermore, with this augmented error, a stable tuning algorithm is (see Refs. [16, 11])

$$\dot{w} = \mu_c e_a \tilde{x} \quad ; \quad \hat{y} = w^T x \quad (11)$$

Bayard [11] presents a very detailed analysis of this algorithm for a variety of potential disturbance environments. For a stationary harmonic disturbance at ω_0 , the equivalent LTI controller is

$$G_{AE}(s) = \frac{G_{FXLMS}(s)}{1 + |\hat{S}(j\omega_0)|^2 G_{LMS}(s) - \hat{S}(s)G_{FXLMS}} \quad (12)$$

where as before, $\hat{S}(s)$ is a model of the secondary path. The expression for the controller in Eq. 12 was chosen to emphasize the close links between AE and the LMS and FXLMS algorithms, and in particular, the very similar properties that they share. For example, AE compensators have the same robustness properties at ω_0 as the FXLMS designs (i.e. can tolerate phase errors as large as $\pm 90^\circ$) [11].

As discussed, the compensator poles of the FXLMS and LMS controllers are at the disturbance frequency, and effectively result in infinite gain at this frequency. Eq. 4 shows that LMS includes no further information about the secondary path dynamics. However, the FXLMS algorithm also includes information about the phase of the system (secondary path) at the frequency of the applied disturbance. The phase information determines the locations of a real zero in the compensator G_{FXLMS} , which will have a large impact on the compensator and loop phase. There are an infinite number of secondary paths that could have the same magnitude and phase at the given disturbance frequency, and we would use the same FXLMS controller for each of them. Clearly the stability of the compensator on these systems for a given value

of μ would be determined by the loop gain and phase at other frequencies as well, but these are not taken into account in the design. The analysis of the AE algorithm is more complex than LMS and FXLMS, as illustrated by the fact that the location of the poles of AE compensator are not obvious from Eq. 12. In particular, the AE algorithm includes the entire plant model in the controller term $\hat{S}G_{\text{FXLMS}}$. This part of the control is used to cancel SG_{FXLMS} terms that appear in the characteristic equation for the system in Figure 1(c), thereby phase stabilizing the system and guaranteeing stable closed-loop performance independent of the value of μ [11]. Of course, the impact of modeling errors in $\hat{S}$ is a key issue, and for systems with resonant modes, this will require study beyond the multiplicative error model used in Ref. [11].

The development of a formal framework for the adaptive controllers to enable further extensions will be addressed in more detail in subsequent work. It is interesting to note however that Hassibi *et al.* have recently shown that the LMS algorithm is $\mathcal{H}_\infty$ optimal, in the sense that it minimizes the maximum energy gain from the disturbances to the predicted errors [17]. This result is important because it shows a close connection between adaptive control and modern robust estimation/control. The approach provides the opportunity for many further developments such as adaptive approaches based on optimal $\mathcal{H}_2/\mathcal{H}_\infty$ estimators and formulations of FXLMS-like estimators that do not require any assumptions about the adaptation rate.

Limitations of FIR Filters for Isolation

This section reviews an important limitation of LMS-type algorithms implemented with FIR filters. The limitation involves light damping in the physical system under control. For certain physical isolation architectures it is possible for lightly damped modes to contribute to the primary (or secondary) path response. In many cases, the attractive aspects of FIR filters, and careful selection of actuators and sensors, can mitigate the effects discussed here.

The general issue is illustrated well by considering a single-channel LMS algorithm. Complete cancellation occurs if the weights converge to the primary path times the inverse of the secondary path, $P(z)S(z)^{-1}$, which is a single transfer function in this case. The performance of the adaptive feedforward algorithm is strongly related to the ability of the weights to model the transfer function $P(z)S(z)^{-1}$, and this ability is affected by the filter architecture chosen for the weights, specifically whether finite impulse response (FIR) filters or infinite impulse response (IIR) filters are used.

Each filter can be expressed as a finite difference

equation, or equivalently, a ratio of polynomials in z . The general expression for the z -transform of an N tap FIR filter is

$$H(z) = b_0 + b_1z^{-1} + \dots + b_Nz^{-N}, \quad (13)$$

and considering $H(z)$ as the transfer function to the input-output relationship of two variables,

$$H(z) = \frac{Y(z)}{X(z)} \quad (14)$$

we can rewrite the output of an FIR filter at each sample, n , as

$$y(n) = b_0x(n) + b_1x(n-1) + \dots + b_Nx(n-N). \quad (15)$$

In contrast, the general expression for an IIR filter is

$$H(z) = \frac{b_0 + b_1z^{-1} + \dots + b_Nz^{-N}}{1 + a_1z^{-1} + a_2z^{-2} + \dots + a_Mz^{-M}}, \quad (16)$$

while the filter output at each sample, n , is

$$y(n) = b_0x(n) + b_1x(n-1) + \dots + b_Nx(n-N) - a_1y(n-1) - \dots - a_My(n-M). \quad (17)$$

The primary difference between the two filter types is that an IIR filter inherently contains feedback due to the delay terms in the denominator of the transfer function. Comparing equation 15 and equation 17, it is clear that an FIR filter is simply an IIR filter with all of the feedback terms set equal to zero. This lack of feedback places constraints on the ability of the FIR filter to match the magnitude and phase of an arbitrary transfer function.

In general, FIR filters are simple to implement since they require only knowledge of the past and present inputs. Stability of the weights is achieved when there is no feedback between the filter output and the filter input. In contrast, the lack of input-output feedback in the filter structure makes it difficult to model the poles of a transfer function. In general, a large number of taps is required to model a dynamic system with lightly-damped poles.

Using IIR filters in the weight updates for the LMS algorithm introduces advantages and disadvantages. Poles and zeros can be matched exactly with an IIR filter because of the input-output feedback; a low-order IIR filter can achieve the same performance as an FIR filter that contains a large number of taps. The primary disadvantage of an IIR implementation is that the filter is not guaranteed to be stable, which requires that the poles of filter be monitored during the on-line adaptation. Several pole monitoring techniques exist, but they require intensive computations and therefore

Table 2: Comparison of FIR and IIR filters for adaptive feedforward design

	Advantages	Disadvantages
FIR	<ul style="list-style-type: none"> • simple implementation • stability guaranteed 	<ul style="list-style-type: none"> • require large number of taps to model lightly-damped poles
IIR	<ul style="list-style-type: none"> • poles and zeros can be matched exactly • can achieve performance equal to FIR design with lower order filters 	<ul style="list-style-type: none"> • stability not guaranteed • adaptation may converge to a local minimum • convergence of adaptation algorithm could be slow compared to FIR design

increase the convergence time of the algorithm. Furthermore, due to the feedback in the filter, the algorithm is no longer guaranteed to converge to the global minima during the adaptation.

A study was conducted to understand the convergence properties of FIR filter designs for plants with dynamics similar to those of vibration isolation systems. Of specific interest are the ability of FIR filters to model the resonant dynamics associated with the suspension modes of the isolation system, and the effect of damping on the convergence of the FIR filter weights.

For the investigation, the magnitude and phase responses of the secondary and primary paths were assumed to be linear combinations of second-order filters with natural frequencies at 50 Hz and 80 Hz. The damping ratios of the filter were varied, and for each damping ratio, a constrained optimization was performed to minimize the error between the magnitude and phase of $P(z)S(z)^{-1}$ and $H(z)$, where $H(z)$ is the transfer function of the FIR filter. The constrained optimization was performed for an FIR filter with 4, 8, 12, $\dots$, 32 taps for damping ratios of $\zeta = 0.01$, 0.03, 0.05, and 0.10. The target frequency range for the constrained optimization was 30 Hz to 95 Hz. No weight was placed on the frequency response outside this band.

The optimization result for $\zeta = 0.01$ and a 32-tap FIR filter is shown in Figure 2. The dashed line in both plots is the magnitude and phase response of $P(z)S(z)^{-1}$, and the rapid changes in magnitude and phase of $P(z)S(z)^{-1}$ as a function of frequency are due to the light modal damping. As the plot indicates, even a 32-tap FIR filter does not accurately model the magnitude and phase of $P(z)S(z)^{-1}$ in the 50 to 60 Hz region. Outside the target band, the FIR filter is a very poor representation of the dynamics because the optimization did not attempt to shape the magnitude and phase of $H(z)$ below 30 Hz or above 95 Hz.

The optimization result for a damping ratio of $\zeta = 0.10$ and a 32-tap filter is shown in Figure 3. As with the previous result, the frequency response of the FIR filter does not even attempt to model the magnitude

Figure 2: Optimal 32 tap FIR filter design for $\zeta = 0.01$: frequency response of $P(z)S(z)^{-1}$ (dashed), frequency response of the optimal 32-tap filter (solid).

and phase of $P(z)S(z)^{-1}$ below 30 Hz and above 95 Hz, resulting in a poor representation of the dynamics at those frequencies. In contrast to the previous result, though, the FIR filter accurately models the dynamics of $P(z)S(z)^{-1}$ within the target band. The increased damping ratio smooths out the magnitude and phase response of $P(z)S(z)^{-1}$, making it easier for the FIR filter to represent the dynamics between 30 and 95 Hz.

Another indicator of the beneficial effects of damping is shown in Figure 4. In this figure, the normalized mean-square error of the single-channel control system is plotted as a function of filter length for systems with increasing modal damping. The normalized mean-square error is defined for this case as

$$E(\zeta, N) = \left[\frac{\int_{60\pi}^{190\pi} (|P(z, \zeta)| - |S(z, \zeta)| |H_N(z)|)^2 d\omega}{\int_{60\pi}^{190\pi} |P(z, \zeta)|^2 d\omega} \right]^{\frac{1}{2}}, \quad (18)$$

which is equivalent to the root mean square response for an input with a unity power spectra between 30 and 95 Hz and zero power spectra elsewhere. The mean-square error is parameterized by the damping ratio, ζ , and the filter length, N . As Figure 4 illustrates, increased damping increases the maximum amount of error attenuation achieved by the FIR filter and decreases the filter length required to obtain an achievable mean-square error. For example, a 32-tap filter only achieves a thirtyfold decrease in the open-loop mean-square error when the damping ratio is 0.01,

Figure 3: Optimal 32 tap FIR filter design for $\zeta = 0.10$: frequency response of $P(z)S(z)^{-1}$ (dashed), frequency response of the optimal 32-tap filter (solid).

whereas the same order filter attenuates the open-loop mean-square error by a factor of 250 for a damping ratio of 0.10. Similarly, a 24-tap filter is required to achieve almost a fiftyfold reduction in the mean-square error for a damping ratio of 0.03, whereas the same error attenuation can be achieved with a 4-tap filter if the damping ratio is increased to 0.10.

Primary results of the numerical investigation are:

- Accurately modeling sharp changes in the magnitude and phase of $P(z)S(z)^{-1}$ requires an FIR filter with a large number of taps.
- Increasing the modal damping in the primary and secondary path transfer functions increases the amount of achievable attenuation using FIR filters as the weights in the adaptation. Similarly, increased damping enables equivalent attenuation with a lower number of filter taps.

Combination of feedback with feedforward is therefore expected to be beneficial, with the presence of low-order feedback enabling use of substantially fewer weights in the FIR filters, and largest benefit expected for broadband performance objectives.

Adaptive Feedforward Implementation

The adaptive feedforward control approaches can be implemented using both analog and digital methods.

Figure 4: Normalized mean-squared error (Equation 18) as a function of filter length and damping ratio is reduced for greater damping, ζ .

Figure 5: Measured $P(z)S(z)^{-1}$ (solid) and two reconstructed FXLMS filters for 79 and 95 Hz disturbances (dashed lines) on three-axis testbed

Figure 6: Measured phase difference between the fourth order quadrature filters used for AE algorithm implementation

Collins [18] provides a detailed analysis of the advantages and issues associated with an analog implementation. However, for a development system (and potential space flight experiment), there are significant advantages, such as increased flexibility and greater broadband performance, in a digital implementation.

The ability of FXLMS to model $P(z)S(z)^{-1}$ was considered in hardware for a tonal disturbance. This property was verified experimentally on a simple three-axis isolation testbed for which feedforward control was implemented in the frequency range of the suspension modes. A 16-tap FXLMS filter executed at 3500 Hz was allowed to converge for the same value of μ for two different disturbances, sinusoidal inputs at 95 and 79 Hz. The filter was then reconstructed and compared in each case to the previously measured $P(z)S(z)^{-1}$. Figure 5 plots the measured phase data, showing the modes of $S(z)$ at 82 and 105 Hz, and the two different reconstructed filters. In both cases, the filter matches the magnitude and phase of $P(z)S(z)^{-1}$ at the appropriate disturbance frequency. Outside the respective narrow frequency ranges, the reconstructed filter is dissimilar to the measured $P(z)S(z)^{-1}$.

Numerous authors have investigated implementation using various length FIR or IIR filters. It is well known that only two weights are required for a harmonic disturbance, but more weights are needed for broadband disturbance rejection. FIR filters were used in this work. The FXLMS was implemented here using 4 weights in a standard tap delay filter. However, the AE algorithm was implemented using an alternate structure employing only two weights, with a quadrature filter developed using the networks from Bedrosian [19]. A digital implementation of the two

Figure 7: Stewart platform microprecision testbed used for vibration isolation experiments

fourth order quadrature filters provides an accurate 90° phase difference (less than 0.5 degree error over 1.5 decades in frequency), which is useful for slewing disturbances (Figure 6).

Experimental Results

This section describes the laboratory experimental system and then summarizes the results of several experiments with the adaptive feedforward algorithms.

Six-axis Payload Isolation Testbed

Figure 7 shows the primary test platform which serves as a prototype for a planned space flight system called SUITE (Satellite Ultraquiet Isolation Technology Experiment). Each member in the six-strut Stewart platform uses a combined passive-active isolator arranged mechanically in series. The stiffness of the damped passive stage can be varied to produce the desired suspension frequencies for a particular payload mass and inertia. The suspension mode frequencies ranged from 25–80 Hz with the mass (6 Kg) and inertia used. The active stage in each laboratory strut uses a piezoelectric actuator (50 μ m stroke) and a velocity sensor that can resolve 50 nm of motion at 10 Hz.

The isolation system is controlled with a dedicated electronic support system including a Texas Instruments TMS320C31 processor which receives downloaded control code from a PC. The Stewart platform was mounted on a generic satellite model (mass 40 Kg). This structure was in turn connected to a rigid plate through soft rubber mounts. During the tests, a

Figure 8: Closed-loop feedforward control performance on 2 struts of the Stewart platform (TITO)

50 lbf capacity VTS shaker driving a reaction mass was mounted vertically on one corner of the satellite model to generate a disturbance input to the base satellite. The signal input to the shaker drive amplifier was used as the reference quantity in the control.

MIMO Feedforward

The particular strut design of the hardware configuration shown in Figure 7 allows a simple extension from SISO to MIMO control because of minimal interaction between the control loops on each leg. Figure 8 shows typical results with feedforward (FF) control on two legs that are attached at adjacent ports on the top plate. The FF control was independently designed and implemented for each leg. The plot shows excellent reduction (> 50 dB) in the transmission of a stationary harmonic disturbance at 100 Hz from the base to the top plate through these two legs. These results suggest that the FF control should easily extend to the full 6 leg MIMO control without requiring the additional complication of successive loop closure.

FXLMS and Augmented Error

As discussed, various algorithms exist to design adaptive feedforward controllers, and both FXLMS and the augmented error methods have been tested on the isolation mount. The FXLMS approach is well known and widely used within the active noise control field, and thus has a significant experimental heritage. On the other hand, the AE approach is much less well known, and there are limited experimental results within this field to support the available analysis. Figure 9 compares the performance achieved for initial SISO experiments using both the FXLMS and AE ap-

Figure 9: Feedforward control with FXLMS and AE

proaches for intermediate values of μ (approximately one half the values that resulted in closed-loop instability for the FXLMS). These graphs plot the ratio of the closed-loop to open-loop system response in the presence of background noise plus a strong stationary harmonic disturbance at 100 Hz. The cancellation of the disturbance is clearly shown by the sharp notch in the plots. As indicated, both approaches achieve greater than 40 dB performance improvement. However, much further work is required to demonstrate the theoretical advantages of AE over FXLMS that were discussed previously. In this implementation, the added plant information in the AE controller did not allow increased gain compared to FXLMS.

Feedback Control

Both feedforward and feedback control have been tested on the isolation mount in Figure 7. Typical feedback results shown in Figure 10 which demonstrate a broadband 20-30 dB reduction in the disturbance transmission (10-100 Hz) for a single strut. Of course, the feedforward and feedback can be combined to provide both broadband and tunable narrowband control. This was tested using a stationary harmonic disturbance at 100 Hz, where the feedback provides essentially no attenuation, but does modify the secondary path transfer function to an extent that the FXLMS must incorporate a modified plant estimate. These experimental results also demonstrate very good reduction of the narrowband disturbance (≈ 40 dB), but there was some additional amplification in the system response near 100 Hz, which indicated a potential interaction between the feedforward and feedback controllers. Further investigation is required to identify the cause of this phenomenon.

Figure 10: Single strut response with feedback control

Slewed Harmonic Disturbances

Figure 11 compares the open and closed-loop response of the system to a harmonic disturbance that is slewed at various rates from 80 to 100 Hz. This test is representative of the disturbance types that are expected on a spacecraft as a result of an imbalance in a reaction wheel that spins up in response to a commanded input. The curves show the square root of the power spectra computed from 20 seconds of measured time data, with the open-loop response for comparison. Case A shows the FXLMS response to a slow slew (80→100 Hz in 10 sec) with a relatively small value of μ , and illustrates that the controller tracks the change in the frequency in the input signal, thereby providing some attenuation from the OL response. However, the tracking is not fast enough to achieve the same levels of attenuation achieved for the initially stationary disturbance at 80 Hz. In comparison, case B shows that the FXLMS response with a larger μ (increased by a factor of 10) can adapt to the change in the signal frequency fast enough that it provides much better attenuation. As shown, the attenuation is particularly good in this case for the first half of the slew. Cases C and D investigate the system performance for even faster slews of 5 sec and 1 sec, respectively. Since the slews are faster, the power at a given frequency is lower, which should lower the sensor response. However, for a given value of μ we would expect the tracking performance to degrade, which should increase the system response. In case D, the very high slew rate causes the second effect to dominate. Note that when the value of μ was increased by another factor of four, the closed-loop system was destabilized. These results, which show the excellent *dynamic* performance of the FXLMS controllers, complement the stationary results shown previously.

Figure 11: Response to slews in the harmonic disturbance from 80→100 Hz

Summary/Conclusions

This paper was motivated by a planned space flight experiment which will benefit from adaptive control. Adaptive control will be implemented on-orbit to complement a planned set of feedback algorithms. The paper considers single and multiple tones for SISO or two-input two-output control. MIMO and broadband control problems will be considered in future work.

The need for adaptive isolation has been motivated and key adaptive feedforward algorithms have been reviewed. One class of limitations of FIR filter representations was discussed as a means of motivating design of systems free of lightly damped modes. Basic FXLMS control was extremely effective at rejecting single fixed tone and rapidly slewing inputs. The Augmented Error (AE) algorithm has potential to expand the utility of FXLMS controllers to systems with non-trivial dynamics. In the present study, AE was implemented and reasonable performance was achieved. Combining feedback and feedforward isolation, another extension to standard FXLMS, demonstrated 40 dB performance improvement in one case. There remain open issues in the control interaction which may be investigated further by considering feedforward controllers from the feedback perspective.

Acknowledgements

The authors thank Donald Leo for his contributions to the analysis and discussion of FIR filters, and Mark Holcomb, Scott Pendleton, and Michael Evert for developing the testbed.

References

- [1] Glaser R., Garba J., and Obal M. STRV-1B Cryocooler Vibration Suppression (1995) In *AIAA Structures, Structural Dynamics, and Materials Conference*, April 1995 AIAA 95-1122.
- [2] Edberg D., Boucher R., Schenck D., Nurre G., Whorton M., Kim Y., and Alhorn D. Results of the STABLE Microgravity Vibration Isolation Flight Experiment (1996) In *19th American Astronautical Soc. Guidance and Control Conference*, February 1996 AAS 96-071.
- [3] Bushnell G., Jacot D., Anderson T., and Becraft M. Active Rack Isolation System Development for the International Space Station (1997) In *AIAA Structures, Structural Dynamics, and Materials Conf.*, April 1997 AIAA 97-1203.
- [4] Sullivan J., Cobb R., Rahman Z., and Spanos J. Closed Loop Performance of a Vibration Isolation and Suppression System (1997) In *American Control Conference*, June 1997 I-97115D.
- [5] Sommerfeldt S. (1989) *Adaptive Vibration Control of Vibration Isolation Mounts Using an LMS-based Control Algorithm* PhD thesis, The Pennsylvania State University.
- [6] Spanos J. and Rahman Z. Narrow-band Control Experiments in Active Vibration Isolation (1994) In *SPIE Proc., Vibration Monitoring and Control*, July 1994 SPIE 2264-01.
- [7] Geng Z. and Haynes L. (1994) Six Degree of Freedom Active Vibration Control Using the Stewart Platform. *IEEE Transactions on Control Systems Technology*. pages 45–53.
- [8] Hyland D. and Harris L. A Multiple-Input, Multiple-Output Neural Architecture for Suppression of a Disturbance of Multiple Tones (1996) In *19th American Astronautical Soc. Guidance and Control Conference*, February 1996 AAS 96-067.
- [9] Widrow B. and Stearns S. (1985) *Adaptive Signal Processing* Prentice-Hall, Englewood Cliffs, New Jersey.
- [10] MacMartin D. A Feedback Perspective of the LMS Feedforward Algorithm (1994) In *American Control Conference*, pages 1632–1636, 1994.
- [11] Bayard D. (1996) A General Theory of Linear Time Invariant Adaptive Feedforward Systems with Harmonic Regressors JPL Internal Report D-13556, Jet Propulsion Laboratory, Pasadena, CA.
- [12] Morgan, D. (1980) An Analysis of Multiple Correlation Cancellation Loop with a filter in the Auxiliary Path. *IEEE Trans. on Acoustics, Speech, and Signal Proc.* ASSP-28(4), 454–467.
- [13] Scribner K. Active narrow band Vibration Isolation of Machinery Noise From Resonant Substructures (1990) Master's thesis, Department of Aeronautics and Astronautics, MIT, Cambridge, MA.
- [14] Sievers L. and von Flotow A. (1992) Comparison and Extensions of Control Methods for Narrow-Band Disturbance Rejection. *IEEE Trans. on Signal Proc.* 40, 2377–2391.
- [15] Monopoli, R. (1974) Model Reference Adaptive Control with an Augmented Error. *IEEE Trans. on Auto. Control.* AC-19(5), 474–484.
- [16] Narendra K. and Annaswamy A. (1989) *Stable Adaptive Systems* Prentice Hall, Englewood Cliffs, NJ.
- [17] Hassibi B., Sayed A., and Kailath T. (1995) $\mathcal{H}_\infty$ Optimality of the LMS Algorithm. *IEEE Trans. on Acoustics, Speech, and Signal Proc.* AC-44(2), 267–280.
- [18] Collins S. (1994) *Multi-Axis Analog Adaptive Feedforward Cancellation of Cryocooler Vibration* PhD thesis, Dept. of Aero. and Astro. Mass. Inst. of Tech., Cambridge, MA MIT SERC Report #8-94.
- [19] Bedrosian, S. (1960) Normalized Design of 90° Phase-Difference Networks. *IRE Trans. on Circuit Theory*. pages 128–136.