JOURNAL OF SPACECRAFT AND ROCKETS
Vol. 33, No. 1, January–February 1996

Microgravity Isolation System Design: A Case Study

R. David Hampton*
McNeese State University, Lake Charles, Louisiana 70609-1735
Carl R. Knospe†
University of Virginia, Charlottesville, Virginia 22903
and
Carlos M. Grodsinsky‡
NASA Lewis Research Center, Cleveland, Ohio 44135

Many acceleration-sensitive, microgravity science experiments will require active vibration isolation from manned orbiters on which they will be mounted. The isolation problem, especially in the case of a tethered payload, is a complex three-dimensional one that is best suited to modern control design methods. Extended H_2 synthesis is used to design an active isolator (i.e., controller) for a realistic single-input, multiple-output microgravity vibration isolation problem. Complex μ-analysis methods are used to analyze the isolation system with respect to sensor, actuator, and umbilical uncertainties. The design process employed and the insights gained are fully discussed. This design case study provides a practical approach for isolation problems of greater complexity. Issues addressed include a physically intuitive state–space description of the system, disturbance and noise filters, filters for frequency weighting, and uncertainty models. The controlled system satisfies all the performance specifications and is robust with respect to model uncertainties.

Nomenclature

c	= umbilical damping, lbf · s/ft
d	= orbiter inertial position, ft
f	= direct disturbance acting on payload, lbf
f_1	= mass-normalized direct disturbance, ft/s^2
k	= umbilical stiffness, lbf/ft
s	= Laplace variable, s^{-1}
t	= time, s
u	= control current, A
x	= payload inertial position, ft
x_1	= payload relative position, ft
x_2	= payload relative velocity, ft/s
x_3	= payload acceleration, ft/s^2
y	= plant output, ft/s^2 ($y = x_3$)
α	= proportionality constant, lbf/A
ζ	= damping factor
ω_h	= pseudostate-filter circular frequency, rad/s

Introduction

THE microgravity vibration-isolation problem has received considerable attention in recent years. It is anticipated that a number of materials-processes and fluid-physics experiments, planned for study in a weightless space environment, will experience unacceptably high background acceleration levels if not isolated.[1] The low-frequency disturbances of greatest concern are a natural accompaniment of space flight with large, flexible, unloaded structures and random, human-induced excitations. Passive isolation alone is incapable of providing the necessary isolation. The combined need, with many experiments, for human interaction and for umbilicals connecting orbiter with payload has resulted in a very difficult, three-dimensional active-isolation design problem.[1]

An earlier paper[2] by the authors introduced an extended H_2-synthesis framework, along with an associated general design philosophy, for developing a robust microgravity vibration-isolation controller. A subsequent paper provided an analysis framework and philosophy for evaluating a given isolation-controller candidate, with emphasis on the effective use of μ-analysis methods.[3] In the present work, extended H_2-synthesis and μ-analysis methods are applied to a realistic single-input, multiple-output (SIMO) microgravity vibration-isolation problem. The design process and results provide engineering insights useful in developing the designer's intuition for isolation problems of greater complexity.

Problem Statement

System Model

A one-dimensional isolation problem was chosen for study (see Fig. 1). The design objective was to develop a feedback controller for isolating a tethered experiment mass (payload) against low-frequency $10^{-3}g$ (stochastic) disturbances, without exceeding rattlespace constraints. The plant (i.e., tether plus payload) is subject to both direct and indirect disturbances. The direct disturbances are those that act directly upon the payload; for example, they could be caused by air currents, astronaut contact, the flow of fluids for lubrication or cooling, or rotating machinery mounted on the experiment platform. The indirect disturbances act upon the payload through the umbilical, and are caused by the vibratory motion of the experiment rack (or, equivalently, the orbiter, to which it is hard-mounted). Massless umbilicals, characterized by a stiffness and a damping, connect the orbiter and the payload. A

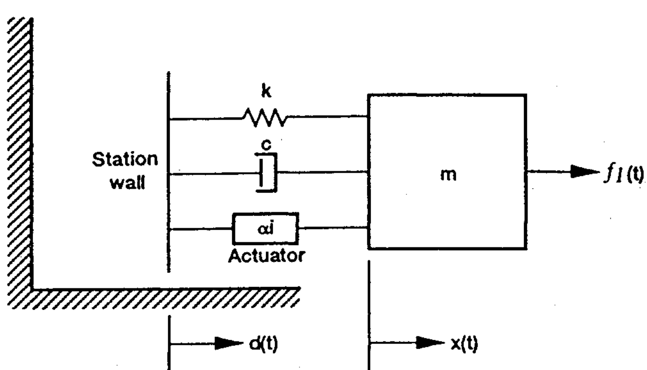

Fig. 1 Tethered mass *m* subject to disturbances *d* and f_1 and to actuator force αi.

Received Nov. 21, 1994; revision received Aug. 7, 1995; accepted for publication Aug. 29, 1995.

*Assistant Professor, Mechanical Engineering Department, P.O. Box 91735 (Drew Hall).

†Assistant Professor, Department of Mechanical, Aerospace, and Nuclear Engineering, Thornton Hall.

‡Senior Research Engineer, MS 500-216, 21000 Brookpark Road.

Lorentz actuator exerts a control force proportional to the applied control current, with proportionality constant α. Typical parameter values are assumed: mass 75 lbm, stiffness 1.544 lbf/ft, damping 0.01138 lbf · s/ft ($\zeta = 0.3\%$), and $\alpha = 2$ lbf/A.

Design Specifications

The final feedback controller should satisfy the following:

1) The payload should track perfectly the dc motion of the spacecraft, where almost no relative motion can be tolerated on account of rattlespace constraints.

2) Below 0.001 Hz, the payload vibration should track the orbiter vibration to within 10%, again to prevent collision of the payload with the walls of the experiment rack surrounding it.

3) Above 0.1 Hz, the payload acceleration should be 40 dB below the spacecraft acceleration, to provide adequate vibration isolation.

4) The loop gain of the system (plant and controller) should be less than 0.1 above 200 Hz, to avoid controller excitation of unmodeled modes at higher frequencies, where the system model is less accurate.

5) The system should remain stable and exhibit good performance for anticipated inaccuracies in the system model. (Note: This is too vague to be a specification in the strict sense of the term; it is more precisely a guideline for use as a point of comparison among competing controller candidates.)

Controller Design

Choice of States

The plant, shown in Fig. 1, can be modelled mathematically via a state–space description. A certain freedom exists in choosing the states used in such a model. The authors made the state choices so as to provide the greatest insight into the physics of the design problem and the requirements posed by the design specifications.

Since the overriding design objective was to reduce the acceleration of the payload (i.e., $\ddot{x}$), and since the H_2 problem is most fundamentally a weighted state-minimization problem, payload acceleration was an obvious state choice. With this selection, a heavier weighting of $\ddot{x}$ in the cost functional signals the H_2 machinery to attempt to increase effective system mass, a concept very familiar and physically intuitive to vibration engineers. Acceleration has the further advantage of being easily measurable in space.

A second logical choice of state, for space applications, is the payload relative position ($x - d$), a quantity that like $\ddot{x}$ is readily measurable. A heavier weighting of $x - d$ in the cost functional signals the H_2 machinery to attempt to increase the effective umbilical stiffness. Consequently this second state choice, like the first, yields a great deal of physical insight into the design problem. A further advantage is that specifications 1 and 2 can be expressed easily in terms of this state.

To complete the state–space description, the payload relative velocity ($\dot{x} - \dot{d}$) must also be included. This final state choice conveniently allows the design engineer to weight the effective umbilical damping.

With these three state choices, the frequency-weighting capabilities of H_2 synthesis also become relatively intuitive tools. For example, a heavier weighting of $x - d$ at low frequencies signals the H_2 machinery to bias its efforts toward a control that causes payload tracking of the orbiter in the low-frequency range, where rattlespace constraints would be most limiting.

Using these three states, the system of Fig. 1 can be written in state–space form as follows. The differential equation of motion for the system is

$$m\ddot{x} = -k(x - d) - c(\dot{x} - \dot{d}) - \alpha u + f_1 \tag{1}$$

Let capital letters indicate the Laplace transforms of the corresponding time-domain variables. Then using the definitions

$$x_1 := x - d \tag{2}$$

$$x_2 := \dot{x} - \dot{d} \tag{3}$$

and

$$f := f_1/m \tag{4}$$

one can rewrite Eq. (1) as a state–space model:

$$\begin{Bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{Bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ -k/m & -c/m & 0 \\ -\omega_h k/m & -\omega_h c/m & -\omega_h \end{bmatrix} \begin{Bmatrix} x_1 \\ x_2 \\ x_3 \end{Bmatrix} + \begin{bmatrix} 0 \\ -\alpha/m \\ -\omega_h \alpha/m \end{bmatrix} u + \begin{bmatrix} 0 & 0 \\ -1 & 1 \\ 0 & \omega_h \end{bmatrix} \begin{bmatrix} \ddot{d} \\ f \end{bmatrix} \tag{5}$$

$$y = [0 \quad 0 \quad 1] \begin{Bmatrix} x_1 \\ x_2 \\ x_3 \end{Bmatrix} \tag{6}$$

The dynamics of Eq. (1) are contained in the first two rows of Eq. (5). The state x_3 has been added so that the model will contain a state that acts like payload acceleration. For low frequencies (i.e., low relative to the pseudostate-filter frequency), the x_3 signal is equivalent to payload acceleration. Thus, the transfer functions of interest for vibration isolation are X_3/s^2D (that between orbiter acceleration and payload acceleration) and X_3/F (that between direct disturbance force and payload acceleration).

Choice of Input Disturbances for the Synthesis Problem

As noted previously, there are two kinds of disturbances to consider, viz., the indirect and the direct. The former type, transmitted via the umbilical, is generally considered to be of greater concern for most types of experiments. However, for some payloads (e.g., those having either moving mechanical parts or flowing liquids, or those requiring direct human intervention), direct disturbances may be significant. A control that attenuates both disturbance types is thus to be preferred.

There is another important reason to include direct disturbances in the design problem: their inclusion can significantly improve the stability robustness of the resulting system. State feedback for this problem corresponds to changes in one or more of the following: payload mass, umbilical stiffness, or umbilical damping. All of these changes correspond to passive isolation strategies. Thus, it is useful to examine how the transfer functions of interest are affected by changes in these parameters. (See Ref. 4 for a more extended treatment of these parametric issues from a classical perspective.) It can be shown readily from Eq. (1) that the respective acceleration-reduction transfer functions are as follows:

$$\frac{s^2X}{s^2D} = \frac{cs + k}{ms^2 + cs + k} \tag{7}$$

and

$$\frac{s^2X}{F} = \frac{s^2}{ms^2 + cs + k} \tag{8}$$

The payload acceleration because of indirect disturbances [see Eq. (7)] can be reduced either by lowering the effective stiffness of the system or by raising its effective mass. The former method lowers the system's stability robustness to variations in umbilical stiffness; the latter suffers no corresponding penalty. Further, whereas reducing the effective umbilical stiffness adversely affects system transmissibility to direct disturbances [see Eq. (8)], raising the effective system mass lowers the transmissibility to both disturbance types. Consequently the latter approach is preferable; this is especially true in view of the likely modeling inaccuracies in umbilical stiffness. The extended H_2 synthesis machinery can be biased to seek an increased-mass control solution, by using disturbance-accommodation techniques with a direct-disturbance model. If only an indirect disturbance is included in the model, extended H_2 synthesis has no reason to prefer increasing mass over lowering stiffness. If the latter is less costly, in terms of the performance index, the synthesized controller may indeed have good performance (as was found in the authors' experience), but it will have an unacceptably low stability robustness to variations in umbilical stiffness.

An increased-mass solution can be obtained as follows. Beginning with the state–space formulation of the problem [Eq. (6)], incorporate the disturbances into the extended H_2 synthesis problem using standard disturbance-accommodation methods.[2] If the power of the direct-disturbance model, f, is made large relative to that of the indirect-disturbance model, $\ddot{d}$, in some frequency range (again, see Ref. 2), then the extended H_2 synthesis machinery will be biased to seek a control solution that adds effective mass to the system for those frequencies. Note that it is not important for these disturbance models actually to match the physical disturbances. The disturbance models are included in the system model simply to provide the designer with a degree of control over the type of approach used by the synthesis machinery. To exploit this capability with respect to robustness concerns, and to attenuate the direct disturbances, which are significant with some payloads, direct disturbances were included in the problem formulation.

Choice of Design Filters for the Synthesis Problem

The ability to use state and control frequency-weighting filters, and to assign disturbance-accommodation filters, adds a great degree of flexibility to the problem formulation. But it also requires that the designer now make some filter choices. Each filter adds to the problem's complexity, so the simplest approach is to use no filters; i.e., to use a basic H_2 synthesis approach, with no extensions. Without extensions, however, H_2 synthesis chooses the optimal controller feedback gains under the erroneous assumption of a perfect plant model. The synthesis machinery simply seeks a stabilizing feedback control that minimizes the quadratic performance index, paying no particular attention to the method by which such minimization is achieved. There is no fundamental reason, for example, for basic H_2 synthesis to prefer a greater-mass solution over one that merely reduces stiffness. In view of the inevitable modeling inaccuracies, the consequence is generally a controller that lacks stability or performance robustness to plant parameter uncertainties. In the present case, all design attempts without the use of filters failed to produce a suitable control; it was necessary to employ the H_2-synthesis extensions.

The first layer of complexity added was disturbance-accommodation filtering of the direct-disturbance model. This filtering is necessary if one is to affect the synthesis of the feedback gain matrix K by the relative weightings of the two disturbance types. Without such filtering the direct and indirect disturbances are modeled simply as white noise (specifically, zero-mean white Gaussian), and the relative weightings can affect only the synthesis of the observer gain matrix L (Ref. 2). When a low-pass filter was added to the direct-disturbance model, so that the direct disturbances were now both large at low frequencies (as before) and also represented by their pseudostates in the performance index (the purpose of adding the filter), the synthesized controller led to much greater system robustness to umbilical stiffness variations. This was because now the synthesis machinery sought a feedback gain matrix K that could reject these large, low-frequency direct disturbances. The performance, however, was still short of the design specifications. In particular, the controller did not meet the higher-frequency requirement of specification 4, viz., controller turnoff above 200 Hz. This requirement could be satisfied only by using frequency weighting.

To require the controller to turn off at higher frequencies, high-pass filtering was used to penalize high-frequency control. (Note that there are limits on the type of control filtering allowable, since the control filter must have a high-frequency asymptote with zero slope if a solution is to exist.[3,5]) This high-pass control filtering, however, is inadequate in itself to demand the required high-frequency controller rolloff. It is necessary also that all state weightings roll off, so that at high frequencies the states will make negligible demands for control effort. Otherwise, performance-index minimization will not permit complete rolloff of the control. Similarly, it is advisable to have disturbance models with negligible power at the higher frequencies, so that the state responses to those disturbances will not make unnecessary control demands. It was found that a low-pass filter in the direct-disturbance model, along with a small, frequency-independent weighting in the indirect-disturbance model, could provide the desired results (i.e., controller turnoff) with minimum added complexity. The improvement in the higher-frequency performance, however, was accompanied by a degradation in the performance at the lower frequencies. In particular, with flat frequency weightings on the states for the lower frequencies (i.e., with zero dc slopes) the H_2-synthesis machinery could not be compelled to move the closed-loop system poles down to the 0.001-Hz region. Consequently, the disturbance rejection in the intermediate range (approximately 0.001–0.1 Hz) was insufficient. Additional frequency weighting in and below the intermediate-frequency region was necessary to surmount this obstacle.

The open-loop Bode α plot depicting the system transmissibility to orbiter acceleration (Fig. 2) indicates unit transmissibility up to the system natural frequency at about 0.1 Hz. If this system knee, or corner frequency, can be moved down by two orders of magnitude to about 0.001 Hz, the corresponding controlled system will satisfy the first two specifications, viz., perfect tracking of the orbiter motion at dc and unit transmissibility below 0.001 Hz. It is well known that the natural frequency of a spring–mass–damper system can be reduced either by raising system effective mass, or by lowering the effective stiffness, or both. And it has already been noted that, for robustness reasons, the first of these three is the preferred approach.

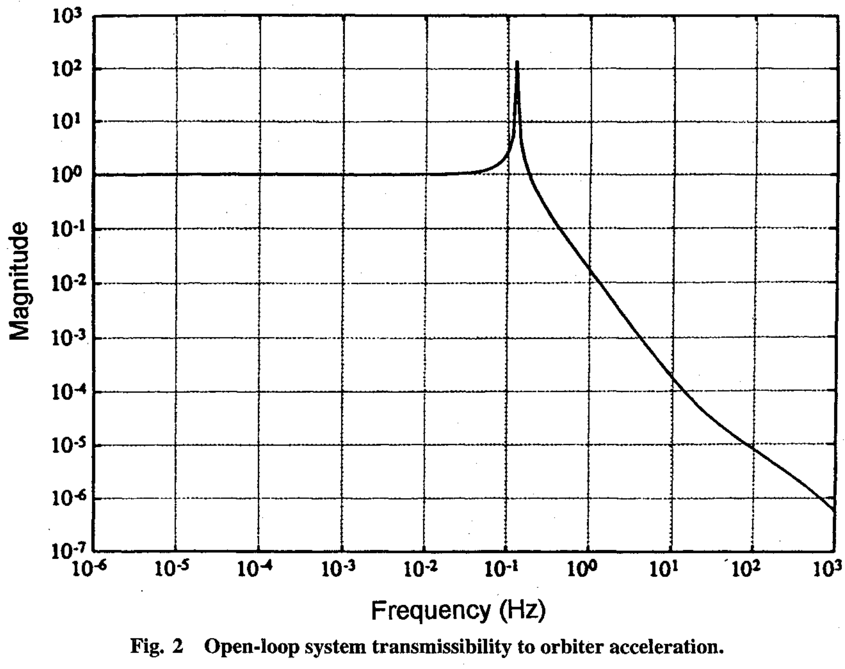

Fig. 2 Open-loop system transmissibility to orbiter acceleration.

There are two means by which the designer can call for a greater-mass solution from extended H_2 synthesis. One method is to incorporate a large direct disturbance into the system model, as previously noted. (It turns out that this greater-mass solution need be requested only in the general region surrounding the open-loop- and closed-loop-system corner frequencies.) A second method is to place a high penalty on payload relative position at intermediate and lower frequencies (i.e., in the general vicinity of the open- and desired closed-loop-system knees, and down to dc, respectively), so that the design machinery will tend to reject a lower-stiffness solution as too costly. If at the same time the designer attaches a high cost to intermediate-frequency acceleration, he can increase the attractiveness of a greater-mass solution.

For a single-input, single-output (SISO) system, the disturbance-accommodation and state-frequency-weighting approaches to design can always be used to produce equivalent controllers (i.e., having identical pole–zero patterns).[6] For a SIMO or multiple-input, multiple-output (MIMO) system, however, although there is still a duality relationship, the systems in general cannot be made equivalent. This means that controllers designed respectively by the two methods will be related (by duality) but not identical in performance. For the present SIMO problem, it was found that the frequency-weighting approach led to the most robust controller.

In the very low frequency range, as one approaches dc, it is important to note (and perhaps not immediately obvious) that no penalty should be applied to either acceleration or control. A high penalty on acceleration would call for increased effective mass in that range, which would militate against the desired unit transmissibility for indirect disturbances. The payload must track the orbiter at dc, so the system effective stiffness at these very low frequencies must be high, and the effective mass low. Low-frequency control should not be penalized either, since the H_2-synthesis machinery should not have unnecessary constraints placed on it in determining the optimal control solution. Since the open-loop system already has the desired transmissibilities in the dc region, a nonzero control cost at dc places unnecessary (and, as it turns out in practice, apparently debilitating) restrictions on the control that can be used. These additional low-frequency considerations complete a logical design strategy for weighting the states and control.

In summary, the competing demands across the entire frequency range call for the following state- and control-frequency-weighting design filters: an integrating filter to weight relative position, a bandpass filter to weight payload acceleration, either a bandpass or a low-pass filter on relative velocity (the latter is simpler and was the shape ultimately chosen), and a high-pass control filter. Such a set of filter choices calls for a greater-mass solution in the vicinity of the open-loop- and closed-loop-system corner frequencies, and a greater-stiffness solution in the region below. A low-pass filter on a high-power direct-disturbance model, and an unfiltered, very low-power, indirect-disturbance model, can also be used to help drive the synthesis machinery to seek a greater-mass solution.

Choice of Relative Noise Levels for the Synthesis Problem

Several other options are available for influencing the design through the H_2 machinery. These include 1) incorporating control noise into the system model, 2) modeling sensor noise power spectra via the addition of output disturbance-accommodation filters, and 3) adjusting the covariances of the various input (i.e., process), output (i.e., sensor), and control noise signals. In most realistic design problems the existence of a solution to the optimal observer-gain problem requires that the system model have noise in all sensor channels.[5] Since for the present problem an observer is required for state reconstruction (relative velocity is here considered unmeasurable, or at least unmeasured), the omission of sensor noise is not an option. On the other hand, control noise is optional, as is output disturbance accommodation. For the sake of maintaining controller simplicity, it was decided to use these extensions to H_2 synthesis only if necessary to achieve the desired system robustness; they were ultimately found not to be needed.

Choosing intelligently the design process and sensor noise levels (i.e., those levels to be used in the model) requires considering the relative importance to be ascribed to each by the observer, in its task of state reconstruction. Since the dominating system uncertainties were considered to lie in the umbilical model, it was decided to weight the direct input disturbance much more heavily than the indirect input disturbance, even though both were assumed to be uncolored. It was anticipated that this would tend to make the observer more robust to umbilical-stiffness modeling errors. As to the relative size of the two sensor noise levels, it was noted that the observer must use the control signal (assumed to be noise-free) and the two sensor measurements (both noise-contaminated) in performing its state reconstruction. The relative displacement signal was modeled as being contaminated with a much higher noise level than the acceleration signal. The purpose of including this high noise contamination was to bias the extended H_2-synthesis machinery to place much greater confidence in the acceleration signal, and therefore to give it preeminence in its state reconstructions. It was anticipated (correctly) that the result would be improved observation of the acceleration state. This accuracy in acceleration reconstruction is desirable, since the optimal control is fundamentally a smart form of acceleration feedback.

Choice of Uncertainty-Block Types for the Analysis Problem

Complex-μ analysis methods were used to find guarantees on system stability robustness. A multiplicative-input uncertainty block provided a measure of the allowable in-channel phase or gain variation from the controller output to the associated plant input. The primary source of such variations was expected to be the actuator, whether Lorentz or magnetic. The multiplicative-input uncertainty-block weighting function[3] was expressed in terms of the maximum phase variations expected (or allowable), as a function of frequency, at the control input. The weighting function could just as easily have been expressed in terms of expected maxima in the gain variation. Identical information about the MIMO phase and gain margins results from the two formulations of the anticipated in-channel variations.

At the plant output, both structured and unstructured multiplicative uncertainty blocks were used to determine stability-related guarantees on allowable sensor modeling errors. Using an unstructured uncertainty block, stability could be guaranteed only for very small coupling between measurements; the resulting guarantees were negligible. For example, even for the final controller design selected, the MIMO phase-margin guarantee was only 0.000046 deg if unstructured uncertainty blocks were used. This indicates that it may be very important that these measured quantities (relative position and payload acceleration) not be directly dependent on each other. The authors do not expect this to be a problem for an actual active microgravity isolation system. The structured uncertainty test (which implies no cross coupling between sensor channels) yielded guarantees on stability for variations in the sensors which were quite large. (For the final controller design, the MIMO phase-margin guarantee was found to be 50.3 deg.)

The effects of system modeling inaccuracies at higher frequencies were not directly evaluated by complex-μ analysis methods. The high-frequency system modes were handled by forcing controller turnoff by 200 Hz, using the frequency-weighting-design filters described in the previous section.

The use of multiplicative-input and -output uncertainty blocks alone was found to be insufficient to guarantee system stability robustness to umbilical parameter uncertainties; a feedback uncertainty model and an associated analysis framework were developed for this purpose.[3] Like the classical root-locus method, this analysis tool provides guarantees of stability for single-parameter variations from the nominal. But it can also provide stability guarantees for combinations of real stiffness, damping, and mass variations within a continuous region of real values.

Design of the Optimal Controller

A logical design strategy to use extended H_2 synthesis for the specified design problem has been presented. The implementation of this strategy, for determining a practical controller design, involved iteration between synthesis and analysis. The former was to develop a controller candidate; the latter, to evaluate its suitability. A step-up, step-down procedure was followed, with layers of complexity

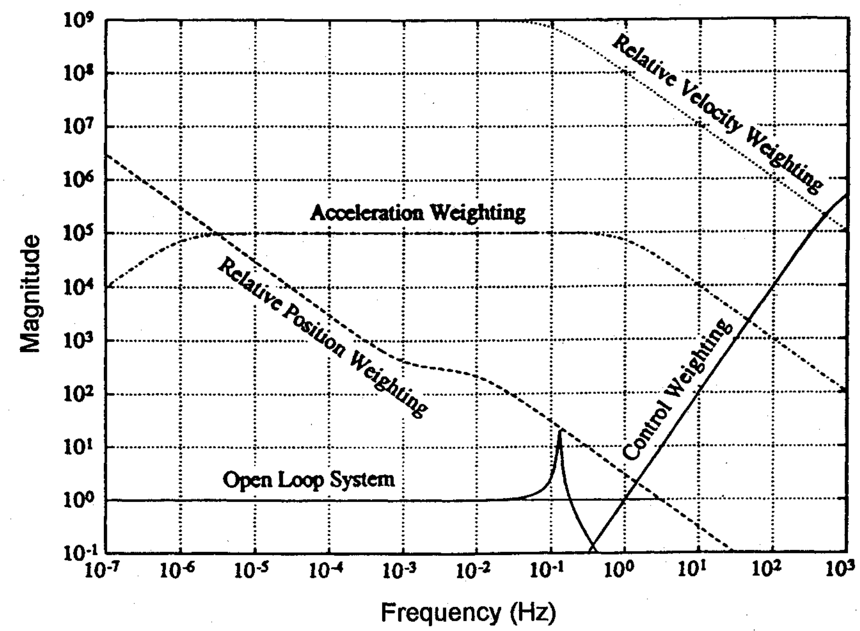

Fig. 3 Bode α plots of the weighted frequency-weighting and disturbance-accommodation filters.

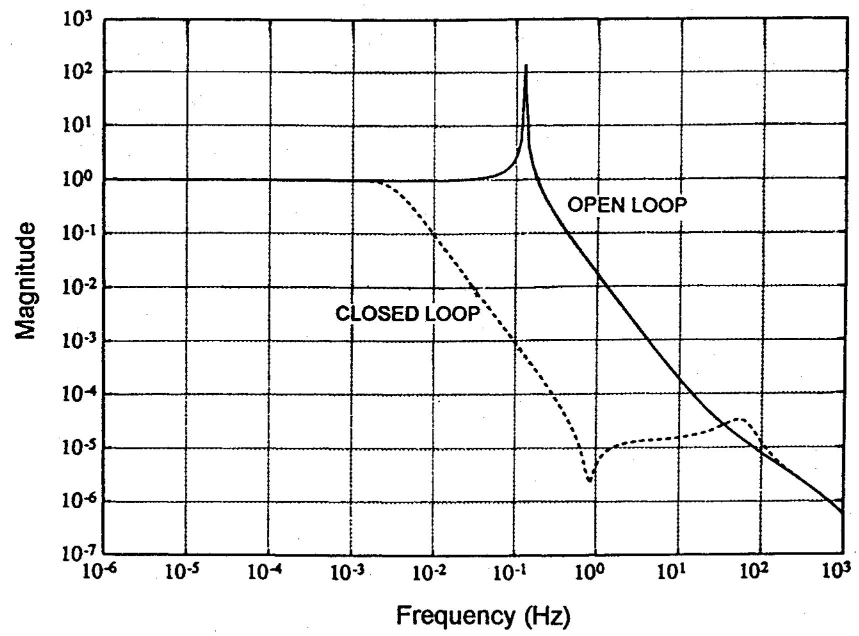

Fig. 4 Open- and closed-loop system transmissibilities to indirect accelerations.

added progressively as the need was determined. After the final design was developed, it was reduced in size using balance-and-truncate (Moore's method) and modal truncation.

In performing these iterations, the authors used an educated trial-and-error approach, informed by the logic presented above. Various trials were needed 1) to select from among the various reasonable alternatives in design-filter shapes, 2) to tune the actual pole and zero locations for the respective design filters, 3) to determine suitable relative weightings among the various frequency-weighted states and the control, 4) to choose suitable relative power levels among the various disturbance-input models, and 5) to reduce the controller size without unacceptably degrading its performance. Computer programs were written in MATLAB to accomplish the necessary synthesis and analysis tasks.

Results

The final design used the weighted filter shapes shown in Fig. 3. The result was reduced to fourth order and then connected to the nominal plant; the closed-loop transmissibility curves for the controlled system are shown in Figs. 4 and 5. The performance of this nominal system met all specifications. Structured singular values of the system seen by multiplicative-input and -output complex uncertainties were used to determine guarantees, respectively, on allowable input phase and gain margins (phase margin: [−51, +51 deg]; gain margin: [0.3118, 7.2060]) and on output MIMO phase and gain margins (phase margin: [−34, +34 deg]; gain margin: [0.6343, 2.3610]). If these in-channel margins are not exceeded, the controlled system is guaranteed not to go unstable. Recall that these guarantees are conservative. A feedback complex-uncertainty Δ block was also used, to determine stability guarantees for uncertainties in real parameters. It was found that for damping essentially unknown ($\pm 10^4$%) and for mass known to within ± 10%, stability could be guaranteed for stiffness known only to within ± 101%. Real parametric studies indicated that the closed-loop system performance remains acceptable (for the various combinations of parametric uncertainties examined) with mass, stiffness, and damping varied within these ranges.

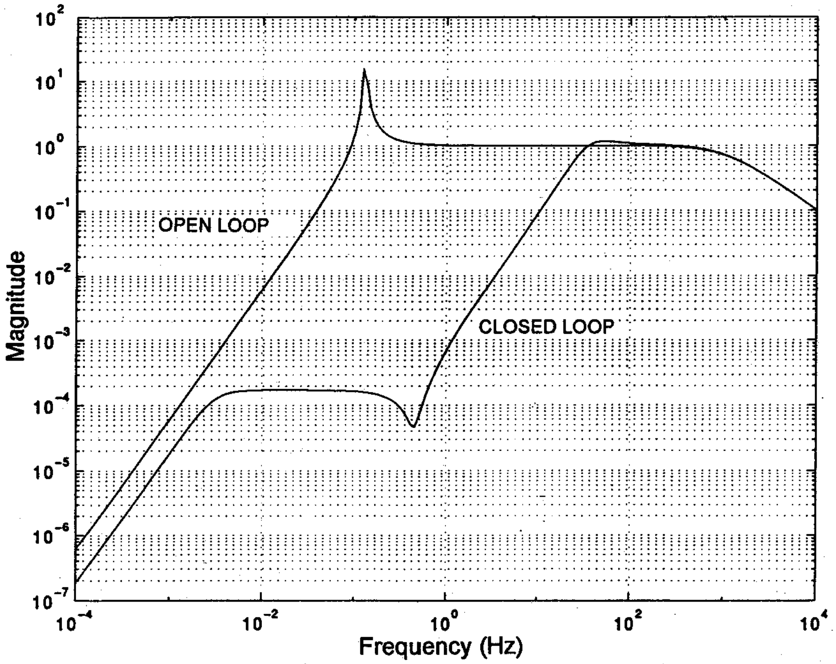

Fig. 5 Open- and closed-loop system transmissibilities to direct accelerations.

Concluding Remarks

This paper has presented the application of extended H_2-synthesis–μ-analysis techniques to a one-dimensional isolation problem. The requirement of the problem was to design a feedback controller to isolate a tethered mass against low-frequency $10^{-3}\, g$ (stochastic) disturbances applied through the umbilical, without exceeding rattlespace constraints. The umbilical was modeled as a stiffness and damping. It was assumed that acceleration and relative position were measurable system outputs, and that a linear (Lorentz) actuator provided the input force to be optimized.

The one-dimensional problem was selected to provide a design paradigm for developing engineering insights into the treatment of isolation problems of greater complexity. Extended H_2 synthesis is, of course, overkill for a problem of this size (although the solution is by no means trivial even by classical means); but the modern control techniques investigated permit ready application to three-dimensional isolation problems, where classical methods bog down.

The system states were chosen to be the relative position, relative velocity, and acceleration of the payload. These choices allowed the design to proceed along intuitive lines, using the extended H_2-synthesis and μ-analysis frameworks that exist in the literature.[2,3] It was found that these design methods yield excellent results for the one-dimensional microgravity vibration-isolation problem; it is not necessary to have a good umbilical model to design a good isolation system.

Frequency weighting in the higher frequencies, of each of the states and of the control, was needed to reduce the controller bandwidth enough to preclude exciting the higher system modes. Frequency weighting was also necessary at lower frequencies to provide the appropriate loop shaping for meeting the other design requirements. When only an indirect disturbance (i.e., one acting through the umbilical) was included in the plant model, it was found that the synthesis machinery could not be required to produce a robust solution. This robustness problem was correctable by using disturbance accommodation to include a direct disturbance in the plant model. Adding such a disturbance effectively biased the extended H_2-synthesis machinery to prefer a greater mass rather than a lower-stiffness solution. Alternatively, it was found that appropriately frequency-weighting the states and the control could accomplish the same end; the best results were obtained using the latter approach. Although the controller found by the design process was of a greater dimension than desired because of the incorporation of frequency-weighting pseudostates, it was readily reducible to a very practical size (fourth order), using standard techniques.

Acknowledgments

The authors are grateful to NASA Lewis Research Center and the Commonwealth of Virginia's Center for Innovative Technology for their partial funding of this work.

References

[1]Nelson, E. S., "An Examination of Anticipated g-Jitter on Space Station and Its Effects on Materials Processes," NASA TM-103775, April 1991.

[2]Hampton, R. D., Knospe, C. R., and Grodsinsky, C. M., "Microgravity Isolation System Design: A Modern Control Synthesis Framework," *Journal of Spacecraft and Rockets*, Vol. 33, No. 1, 1996, pp. 101–109.

[3]Hampton, R. D., Knospe, C. R., and Grodsinsky, C. M., "Microgravity Isolation System Design: A Modern Control Analysis Framework," *Journal of Spacecraft and Rockets*, Vol. 33, No. 1, 1996, pp. 110–119.

[4]Knospe, C. R., Hampton, R. D., and Allaire, P. E., "Control Issues of Microgravity Vibration Isolation," *Acta Astronautica*, Vol. 25, No. 11, 1991, pp. 687–697.

[5]Hampton, R. D., "Controller Design for Microgravity Vibration Isolation Systems," Ph.D. Dissertation, Univ. of Virginia, Charlottesville, VA, Jan. 1993.

[6]Sievers, L. A., and Von Flotow, A. H., "Comparison of Two LQG-Based Methods for Disturbance Rejection," *Proceedings of the 28th IEEE Conference on Decision and Control*, Vol. 1, Inst. of Electrical and Electronics Engineers, Piscataway, NJ, 1989, pp. 483–485.

E. A. Thornton
Associate Editor