JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 21, No. 3, May–June 1998

Fault Accommodation in Control of Flexible Systems

Peiman G. Maghami,* Dean W. Sparks Jr.,† and Kyong B. Lim‡
NASA Langley Research Center, Hampton, Virginia 23681-0001

New synthesis techniques for the design of fault accommodating controllers for flexible systems are developed. Three robust control design strategies, static dissipative, dynamic dissipative, and μ-synthesis, are used in the approach. The approach provides techniques for designing controllers that maximize, in some sense, the tolerance of the closed-loop system against faults in actuators and sensors, while guaranteeing performance robustness at a specified performance level, measured in terms of the proximity of the closed-loop poles to the imaginary axis (the degree of stability). For dissipative control designs, nonlinear programming is employed to synthesize the controllers, whereas in μ-synthesis, the traditional D–K iteration is used. To demonstrate the feasibility of the proposed techniques, they are applied to the control design of a structural model of a flexible laboratory test structure.

Introduction

CONTROL system design for flexible systems is challenging because of their special dynamic characteristics: a large number of structural modes within the controller bandwidth; low, closely spaced modal frequencies; very small inherent damping; and insufficient knowledge of the parameters. For a control design to be considered feasible, it must 1) be of reasonably low order, 2) satisfy the nominal performance specification, and 3) be robust to errors in the design model. A feasible control design must be robust to parametric and nonparametric uncertainties in the system model. In this study we focus on errors or failures in the control system hardware, such as sensors and/or actuators. It is desirable to design a controller that can accommodate, to a specified degree, the failure and/or degradation of the control system hardware.

Generally, instrument failures can be divided into two categories: hard failure, which indicates the total loss of the instrument, and soft failure, which indicates a partial loss of the instrument, resulting in input–output performance degradation. Hard failures have been traditionally dealt with through the introduction of redundancies, along with reliable fault detection systems. Soft failures have been traditionally accommodated through control designs that guarantee sufficient stability margins. In single-input/single-output systems, the stability margins are imposed through gain and phase margins. For multi-input/multi-output systems, stability margins can be attained through robust control theory. This paper investigates the use of robust control theory in developing a methodology for designing fault-accommodating controllers. Three robust control design techniques, namely, static dissipative, dynamic dissipative, and structured singular-value-based control designs, are considered. The paper considers soft instrument failure in the form of degradation of the sensitivity and output of the instrument(s), e.g., actuator force output has decreased for a given input voltage command.

Synthesis techniques are presented for each of the three control strategies to design controllers that provide a specified degree of stability while accommodating soft input–output failures. The term degree of stability refers to the distance from the pole of closed-loop system closest to the imaginary axis, which can be viewed as a performance measure for the system. The approach taken reduces the robust performance problem to a robust stabilization problem. Simple conditions for robust stabilization are established for both static and dynamic dissipative controllers. For dissipative control designs, nonlinear programming is employed to synthesize the controllers, whereas in μ synthesis, the traditional D–K iteration method is used. To demonstrate the feasibility of the proposed techniques, they are applied to the control design of a structural model of NASA's Phase-II Controls–Structure Interaction (CSI) Evolutionary Model test structure.

In the following sections, the presented fault accommodation theories have been formulated for the case of actuator failures. Sensor failures can be treated using similar approaches.

Plant Description

A typical model of a linear, time-invariant (LTI) flexible system may be represented by the following second-order dynamical equation:

$$M\ddot{\boldsymbol{x}} + D\dot{\boldsymbol{x}} + K\boldsymbol{x} = B\boldsymbol{u} \tag{1}$$

together with some set of measurement and performance output equations

$$\boldsymbol{y}_p = C_{mp}\boldsymbol{x}, \qquad \boldsymbol{y}_r = C_{mr}\dot{\boldsymbol{x}}, \qquad \boldsymbol{z}_p = C_{pp}\boldsymbol{x}, \qquad \boldsymbol{z}_r = C_{pr}\dot{\boldsymbol{x}} \tag{2}$$

where M is the positive definite mass matrix; D is the positive semidefinite damping matrix; K is the positive semidefinite stiffness matrix; B is the input influence matrix; C_{mp} and C_{mr} are the position/displacement and rate/velocity measurement output influence matrices, respectively; C_{pp} and C_{pr} are position and rate performance output influence matrices, respectively; $\boldsymbol{x}$ is a $k \times 1$ vector of coordinates, or degrees of freedom; $\boldsymbol{u}$ is a $m \times 1$ vector of inputs to the system; $\boldsymbol{y}_p$ and $\boldsymbol{y}_r$ are the position and rate measurement output vectors, respectively; and $\boldsymbol{z}_p$ and $\boldsymbol{z}_r$ are the position and rate performance output vectors, respectively. Usually, a finite element analysis is used to obtain these matrices. In most cases, the number of degrees of freedom, k, is quite large and, thus, impractical to work with for general design and analysis purposes. To make the problem more tractable, the vector $\boldsymbol{x}$ is transformed into modal coordinates using the transformation $\boldsymbol{x} = \Phi\boldsymbol{r}$, with $\boldsymbol{r}$ being a $n \times 1$ vector of modal amplitudes and $n \ll k$. The transformation matrix Φ contains n columns, which are the eigenvectors associated with the n modes of interest of the flexible system. The equations for the system, in transformed coordinates, are

$$M_r\ddot{\boldsymbol{r}} + D_r\dot{\boldsymbol{r}} + K_r\boldsymbol{r} = \Phi^T B\boldsymbol{u}$$

$$\boldsymbol{y}_p = C_{mp}\Phi\boldsymbol{r}, \qquad \boldsymbol{y}_r = C_{mr}\Phi\dot{\boldsymbol{r}} \tag{3}$$

$$\boldsymbol{z}_p = C_{pp}\Phi\boldsymbol{r}, \qquad \boldsymbol{z}_r = C_{pr}\Phi\dot{\boldsymbol{r}}$$

If normal modes are used, and their mode shapes have been normalized with respect to the mass matrix and modal damping is assumed, then $M_r = I_{n\times n}$, $K_r = \mathrm{diag}(\omega_1^2, \omega_2^2, \ldots, \omega_n^2)$, and $D_r =$

Received Dec. 5, 1996; revision received Nov. 7, 1997; accepted for publication Nov. 25, 1997. Copyright © 1998 by the American Institute of Aeronautics and Astronautics, Inc. No copyright is asserted in the United States under Title 17, U.S. Code. The U.S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner.

*Senior Research Engineer, Guidance and Control Branch. Senior Member AIAA.
†Aerospace Technologist, Guidance and Control Branch.
‡Research Engineer, Guidance and Control Branch. Member AIAA.

$\text{diag}(2\zeta_1\omega_1, 2\zeta_2\omega_2, \ldots, 2\zeta_n\omega_n)$, with ω_i and ζ_i being the open-loop natural frequency and damping ratio values, respectively, for the ith mode.

By defining a new vector $\boldsymbol{x}_s = [\boldsymbol{r}^T \ \dot{\boldsymbol{r}}^T]^T$, the second-order equations in Eq. (3) can be rewritten into first-order form as

$$\dot{\boldsymbol{x}}_s = A_s\boldsymbol{x}_s + B_s\boldsymbol{u} \tag{4}$$

$$A_s = \begin{bmatrix} 0 & I_{n\times n} \\ -K_r & -D_r \end{bmatrix}, \qquad B_s = \begin{bmatrix} 0 \\ \Phi^T B \end{bmatrix}$$

The measurement and performance output equations can then be written as

$$\boldsymbol{y} = \begin{bmatrix} \boldsymbol{y}_p \\ \boldsymbol{y}_r \end{bmatrix} = C\boldsymbol{x}_s, \qquad \boldsymbol{z} = \begin{bmatrix} \boldsymbol{z}_p \\ \boldsymbol{z}_r \end{bmatrix} = C_z\boldsymbol{x}_s \tag{5}$$

$$C = \begin{bmatrix} C_{mp}\Phi & 0 \\ 0 & C_{mr}\Phi \end{bmatrix}, \qquad C_z = \begin{bmatrix} C_{pp}\Phi & 0 \\ 0 & C_{pr}\Phi \end{bmatrix}$$

Here, $\boldsymbol{x}_s$ is the plant state vector, A_s is the plant state matrix, B_s is the control input influence matrix, and C and C_z denote the measurement output and performance output influence matrices, respectively. A thorough description of these matrices with normal modes is given in Ref. 1.

Problem Formulation

In this section, the development of the three control design strategies, static dissipative, dynamic dissipative, and μ synthesis, will be discussed.

Static Dissipative Controller

Dissipative controllers that utilize collocated and compatible actuators and sensors offer an attractive strategy for active control of flexible systems. For example, this means actuators such as torquers are used with attitude and attitude rate sensors and thrusters are used with linear position and linear velocity sensors. Note that complete sensor–actuator collocation and compatibility implies that in Eq. (3) $C_{mp} = C_{mr} = B^T$. Dissipative controllers utilize special passivity type input–output properties of the plant and offer robust stability in the presence of both nonparametric and parametric uncertainties.[1,2] The simplest controller of this type is the constant-gain static dissipative controller, which is given as follows:

$$\boldsymbol{u} = -G_p\boldsymbol{y}_p - G_r\boldsymbol{y}_r \tag{6}$$

where $\boldsymbol{y}_p$ and $\boldsymbol{y}_r$ are the measured position/attitude and rate and G_p and G_r are symmetric, nonnegative-definite gain matrices. The static dissipative controller can be used as a low-authority stabilizing controller in a hierarchical architecture, or as the primary controller.

The objective of this section is to develop a procedure for designing the gains of a static dissipative controller so as to provide robust performance against actuator or sensor failures described earlier. Robust performance means that a specified level of system performance is maintained under possible hardware failures. In this development, the time constants of the closed-loop system poles, i.e., the damping in the closed-loop poles, is used as a measure for performance robustness design. Here, for simplicity of presentation, we shall assume without any loss of generality that the static dissipative controller considered would provide damping enhancement only, i.e., the position/attitude gain matrix G_p is null. It is noted that if zero-frequency rigid modes, which are to be controlled by the dissipative controller, are present the position/attitude gain matrix cannot be null. Note that with collocated and compatible rate feedback alone $C = B_s^T$. Moreover, the controller is assumed to be decentralized (local feedback only), such that the rate gain matrix is diagonal. Although, decentralization may limit the capability of the controller, it would generally enhance its stability robustness and provide for easier and more practical implementation. Furthermore, because the controller is diagonal due to decentralization, diagonal input–output multiplicative uncertainties of the system are equivalent. Therefore, only input multiplicative uncertainties need to be considered.

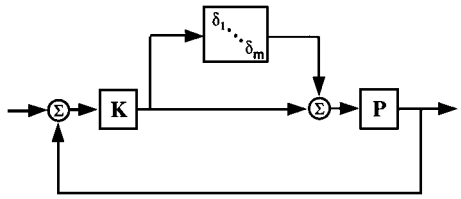

Fig. 1 Block diagram of the feedback system with multiplicative uncertainty in the input.

The problem is now to design a static dissipative controller such that the poles of the closed-loop system, represented by the eigenvalues of the closed-loop state matrix $A_s - B_sG_rB_s^T$, reside in a specified region in the left-half plane, i.e., to the left of a line that is parallel to the imaginary axis, intersecting the real axis at $\boldsymbol{x} = -\alpha$, and is designated as the $-\alpha$ line, for a defined uncertainty in the input to the plant. In other words, the system is to maintain an α degree of stability in the presence of the uncertainty. The block diagram of the system is shown in Fig. 1.

Here, P denotes the plant transfer function defined as $B_s^T(sI - A_s)^{-1}B_s$, Δ represents the uncertainty in the input, and K represents the controller G_r.

Lemma 1: Let P' denote the transfer function of the plant with its poles shifted to the right by α, defined as $B_s^T(sI - A_s - \alpha I)^{-1}B_s$. Then, the α-deg robust performance problem may be changed to a robust stabilization problem by requiring that the controller K stabilizes a modified plant P' for a defined uncertainty in the input to the plant.

Proof: Let K stabilize P' with respect to the uncertainty in the input Δ. This means that the poles of the modified closed-loop system are all stable (to the left of the imaginary axis) for all allowable uncertainty, indicating that the poles of the true closed-loop system would then be to the left of the $-\alpha$ line (because the poles of true system are the poles of the modified system shifted to the left of the imaginary axis by $-\alpha$).

The uncertainty Δ is assumed to have a diagonal representation, $\Delta = \text{diag}(\delta_1, \delta_2, \ldots, \delta_m)$, to allow for independent uncertainty in each input channel. However, because of the dissipative and decentralized nature of the controller one can use a simpler form of uncertainty as indicated by the following theorem.

Theorem 1: K stabilizes P' for all $\Delta = \text{diag}(\delta_1, \delta_2, \ldots, \delta_m)$, $|\delta_i| \le \delta_{\max}\ \forall i$, iff K stabilizes P' for all $\Delta = \delta_{\max}\bar{I}$, where $\bar{I}$ is any diagonal matrix with elements that equal either 1 or -1.

Necessity Proof: Obvious from the statement of theorem 1.

Sufficiency Proof: The closed-loop system equation for a distinct diagonal uncertainty block may be written as

$$A_{\text{cl}} = A_s' - B_s \begin{bmatrix} 1+\delta_1 & 0 & \ldots & 0 \\ 0 & 1+\delta_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 1+\delta_m \end{bmatrix} G_rB_s^T \tag{7}$$

with each $|\delta_i| \le \delta_{\max}$ and $A_s' = A_s + \alpha I$. Rewrite Eq. (7) as follows:

$$A_{\text{cl}} = \hat{A}_{\text{cl}} - B_s \begin{bmatrix} \eta_1 & 0 & \ldots & 0 \\ 0 & \eta_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \eta_m \end{bmatrix} G_rB_s^T \tag{8}$$

or

$$A_{\text{cl}} = \hat{A}_{\text{cl}} - B_s\hat{G}_rB_s^T \tag{9}$$

where $\hat{A}_{\text{cl}} = A_s' - (1 - \delta_{\max})B_sG_rB_s^T$ and $\eta_i = \delta_{\max} + \delta_i$, $\forall i$, with $0 \le \eta_i \le 2\delta_{\max}$, $\forall i$. The eigenvalues of $\hat{A}_{\text{cl}}$ are stable because by assumption K stabilizes P' for all $\Delta = \delta_{\max}\bar{I}$, which includes $\Delta = \delta_{\max}I$. Also, the gain matrix $\hat{G}_r$ is dissipative. Now, the sufficiency of the premise can be proved by contradiction. From the statement of the theorem, K stabilizes P' for all $\Delta = \delta_{\max}\bar{I}$, or by definition, for all $H \equiv \text{diag}(\eta_1, \eta_2, \ldots, \eta_m) = \delta_{\max}(I + \bar{I})$. However, assume that there exists a matrix $H_u \equiv \text{diag}(\eta_{u1}, \eta_{u2}, \ldots, \eta_{um})$,

$0 \le \eta_i \le 2\delta_{\max}$, $\forall i$ such that the closed-loop system is unstable. It is known from the behavior of LTI systems with decentralized dissipative controllers that, as the gain elements are increased, the poles of the closed-loop system approach the transmission zeros from the left-hand side.[3–5] For example, keeping $\eta_{u2}, \eta_{u3}, \ldots, \eta_{um}$ the same, as η_1 is increased from 0 to $2\delta_{\max}$, the closed-loop poles move from the open-loop poles toward the transmission zeros of the system (A_s', B_{s1}, B_{s1}^T). Either the closed-loop poles are stable at $\eta_1 = 0$, or they are not. If at $\eta_1 = 0$, all poles are stable, then because the closed-loop system is assumed to be unstable at $\eta_1 = \eta_{u1}$, and the loci approach the transmission zeros from the left side, it is obvious that the closed-loop system would continue to be unstable as η_1 approaches $2\delta_{\max}$. This indicates that if the closed-loop system is unstable at $\eta_1 = \eta_{u1}$, it will also be unstable at either $\eta_1 = 0$ or $\eta_1 = 2\delta_{\max}$, or both $\eta_1 = 0$ and $\eta_1 = 2\delta_{\max}$. Following the same argument for the other channels, it can be concluded that if the closed-loop system is unstable for $\eta_{u2}, \eta_{u3}, \ldots, \eta_{um}$, it will be unstable as well for some Δ in $\delta_{\max}\bar{I}$. This violates the initial assumption that K stabilizes P' for all $\Delta = \delta_{\max}\bar{I}$ and, thus, theorem 1 is proved.

Therefore, the robust performance problem is redefined in terms of a robust stability problem under a restricted diagonal uncertainty structure in the input. The problem is to compute $\delta_{\max}$ that satisfies the conditions of theorem 1.

The approach taken for synthesizing a static dissipative controller is optimization based. Here, nonlinear programming is used to synthesize the elements of the rate gain matrix to provide maximum allowance for uncertainty in the inputs. One possible optimization scenario may be to use a repeated diagonal uncertainty structure and optimize the elements of the rate gain matrix to maximize a bound on $\delta_{\max}$, for a prescribed degree of stability α, and subject to the conditions of theorem 1. Now, given a plant $P'(s)$ and a controller K, a bound on the uncertainty (repeated diagonal uncertainty structure) is established in the following.

Theorem 2: Given the robust stabilization problem, a bound on the uncertainty $\Delta = \delta I$: $|\delta| \le \delta_{\max}$ may be established as follows:

$$\delta_{\max} = \left\{ \sup_{\omega} \rho([I + KP']^{-1}KP') \right\}^{-1} \ge \left\{ \sup_{\omega} \bar{\sigma}([I + KP']^{-1}KP') \right\}^{-1} \tag{10}$$

In this equation, ρ represents the spectral radius and $\bar{\sigma}$ denotes the largest singular value.

Proof (given in a corollary of Doyle[6]): The functional form used here for $\delta_{\max}$ is the H-infinity norm bound given in Eq. (10). The reason for using the H-infinity norm bound rather than the supremum of spectral radius is that the former is computationally tractable (with commercial software) whereas the latter is not. Therefore, the optimization problem is posed as follows.

Minimize the objective function

$$\sup_{\omega} \bar{\sigma}([I + KP']^{-1}KP') \tag{11}$$

over the diagonal elements of the rate gain matrix G_r, which satisfy the conditions set in theorem 1, K stabilizes P' for all $\Delta = \delta_{\max}\bar{I}$, which may be defined as

$$Re\{\lambda_i(A_{\text{cl}}')\} < 0, \qquad \forall i \text{ and } \forall \Delta = \delta_{\max}\bar{I} \tag{12}$$

as well as side constraints on the elements of the gain matrix.

The number of constraints associated with the conditions of theorem 1 grow exponentially with the number of input uncertainties. Alternatively, one may optimize the Eq. (10) objective function without side constraints to ensure the conditions of theorem 1, although the omission of the side constraints may result in the conditions of theorem 1 not being satisfied by the final design. However, experience has shown that the minimum spectral norm solution generally provides an optimal solution that typically satisfies, or almost satisfies, the conditions of theorem 1. Therefore, the recommended synthesis procedure is to first find an optimal design, then check the conditions of theorem 1. If these conditions are not met, then one has three choices. First, accept a reduced α degree of stability for the designed $\delta_{\max}$, which is readily available from the check of these conditions. Second, reduce the $\delta_{\max}$ until the conditions of theorem 1 are satisfied. Last, restart the design optimization procedure with the conditions included as side constraints.

Dynamic Dissipative Controller

The constant-gain static dissipative controller, though attractive in its simplicity, may result in limited performance due to its inability to provide signal shaping. To achieve better performance, another type of controller, the dynamic dissipative controller, is considered. Like static dissipative controllers, dynamic dissipative controllers that use collocated and compatible sensor/actuator pairs are robust against both parametric and nonparametric model uncertainties. Unlike the static dissipative controller, however, the dynamic dissipative controller has the advantage of having the freedom of controller phase shaping, as well as gain shaping, thus enhancing its ability to affect the overall system dynamics as desired.

The dynamic dissipative controller takes the following form:

$$\dot{\boldsymbol{x}}_c = A_c\,\boldsymbol{x}_c + B_c\,\boldsymbol{y}, \qquad \boldsymbol{u} = -C_c\,\boldsymbol{x}_c \tag{13}$$

where A_c, B_c, and C_c are the controller's state, input influence, and output influence matrices, respectively. The vector $\boldsymbol{x}_c$ is the controller state vector, $\boldsymbol{y}$ is the vector of measurements fed into the controller, and $\boldsymbol{u}$ is the resulting controller output vector. To ensure that the controller is dissipative, the matrices in the controller equation must satisfy the following[7]:

$$A_c^T R + RA_c = -Q, \qquad C_c = B_c^T R \tag{14}$$

where R is positive definite, Q is positive semidefinite, A_c is strictly stable, the pair $(A_c, Q^{1/2})$ is controllable, and the pair (A_c, C_c) is observable.

Like the static dissipative controller, the dynamic dissipative controller is assumed to be decentralized, i.e., having local feedback connections only. Under this assumption, the matrices A_c, B_c, and Q can be written in block diagonal form as follows:

$$A_c = \begin{bmatrix} A_{c1} & 0 & \ldots & 0 \\ 0 & A_{c2} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & A_{cm} \end{bmatrix}, \qquad B_c = \begin{bmatrix} B_{c1} & 0 & \ldots & 0 \\ 0 & B_{c2} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & B_{cm} \end{bmatrix} \tag{15}$$

$$Q = \begin{bmatrix} Q_1 & 0 & \ldots & 0 \\ 0 & Q_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & Q_m \end{bmatrix}$$

with m the number of sensor/actuator pairs used in the controller. In turn, for the ith sensor/actuator pair, the individual nth order controller parameters can be written as

$$A_{ci} = \begin{bmatrix} 0 & 1 & \ldots & 0 \\ 0 & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ -a_1 & -a_2 & \ldots & -a_n \end{bmatrix}, \qquad B_{ci} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix} \tag{16}$$

$$Q_i = \begin{bmatrix} q_1 & 0 & \ldots & 0 \\ 0 & q_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & q_n \end{bmatrix}$$

It should be noted that the order of the controller for each sensor/actuator pair need not be the same. The a and q terms are the design parameters for the dynamic dissipative controller. The a terms control the locations of the controller poles and zeros, whereas the q terms set the gain levels. Instead of computing the matrix C_c directly, the matrix Q is obtained first. Once Q is known, Eq. (14) can be used to solve for the corresponding C_c matrix.

Referring again to the block diagram in Fig. 1, the problem now is to design a dynamic dissipative controller K, such that the eigenvalues of the closed-loop system, represented by the eigenvalues of the closed-loop system matrix

$$A_{\text{cl}} = \begin{bmatrix} A_s & -B_s C_c \\ B_c C & A_c \end{bmatrix} \tag{17}$$

reside to the left of a specified $-\alpha$ line in the left-half of the complex plane, for a given uncertainty in the input to the plant P. Again, the plant transfer function is denoted by $B_s^T (sI - A_s)^{-1} B_s$.

Lemma 2: The α-deg robust performance problem may be changed to a robust stabilization problem by requiring that the controller K', obtained by α shifting the controller K to the right, stabilizes a modified plant P', again defined as $B_s^T (sI - A_s - \alpha I)^{-1} B_s$, for a defined uncertainty in the input to the plant. The significance of the prime notation associated with the controller will be explained shortly.

Proof: Let K' represent a dynamic dissipative controller that stabilizes P' with respect to the uncertainty in the input Δ. This means that the poles of the modified closed-loop system are all stable (to the left of the imaginary axis) for all allowable uncertainty, indicating that the poles of the true closed-loop system would then be left of the $-\alpha$ line (because the poles of true system are the poles of the modified system shifted to the left of the imaginary axis by $-\alpha$).

Once again, the uncertainty Δ is assumed to have a diagonal representation, $\Delta = \text{diag}(\delta_1, \delta_2, \ldots, \delta_m)$, to allow for independent uncertainty in each input channel. In turn, because of the dissipative and decentralized nature of the controller, one can use a much simpler form of uncertainty as indicated by the following theorem.

Theorem 3: K' stabilizes P' for all $\Delta = \text{diag}(\delta_1, \delta_2, \ldots, \delta_m)$, $|\delta_i| \le \delta_{\max} \forall i$, iff K' stabilizes P' for all $\Delta = \delta_{\max} \bar{I}$, where $\bar{I}$ is any diagonal matrix with elements that equal either 1 or -1.

Necessity Proof: Obvious from the statement of theorem 3.

Sufficiency Proof: Let K' be a decentralized, dynamic dissipative controller represented by the system (A_c', B_c, C_c) so that $A_c'^T R + R A_c' = -Q'$ and $C_c = B_c^T R$ for some $Q' \ge 0$ and $R > 0$. Because K' stabilizes P' for all $\Delta = \delta_{\max} \bar{I}$, then K' will stabilize P' for $\Delta = -\delta_{\max} I$. Now, for any distinct uncertainty block, the controller output matrix can be written as

$$\begin{aligned} C_c &= [I + \text{diag}(\delta_1, \delta_2, \ldots, \delta_m)] B_c^T R \\ &= [(1 - \delta_{\max}) I + \text{diag}(\eta_1, \eta_2, \ldots, \eta_m)] B_c^T R \\ &\equiv [(1 - \delta_{\max}) I + H] B_c^T R \end{aligned} \tag{18}$$

with $\eta_i = \delta_{\max} + \delta_i, \forall i$ and $0 \le \eta_i \le 2\delta_{\max}$. It is observed from the preceding controller output matrix equation that the feedback loop has been decomposed into a stabilizing dissipative feedback term, with $C_c^1 = [(1 - \delta_{\max}) I] B_c^T R$, which stabilizes P', along with a feedback term $C_c^2 = H B_c^T R$. Now, define the matrix $\Gamma = [\Gamma_1, \Gamma_2, \ldots, \Gamma_m]$, where $\Gamma_i = \eta_i I_i$ and I_i is an identity matrix the size of A_{ci}', i.e., the size of Γ equals the size of A_c'. Because K' is decentralized, A_c', B_c, C_c, Q', and R are all block diagonal, such that

$$A_c'^T R\Gamma + \Gamma R A_c' = -\Gamma Q', \qquad C_c^2 = B_c^T R\Gamma \tag{19}$$

Now, because $\Gamma \ge 0$, $\Gamma Q' \ge 0$, and $R\Gamma > 0$ (except for the trivial case of any $\Gamma_i = 0$), therefore, the second feedback term C_c^2 is dissipative as well. From here, the same proof by contradiction arguments used to prove theorem 1, i.e., the static dissipative case, apply.

Just as in the static dissipative controller design strategy, the robust performance problem is redefined in terms of a robust stability problem under a restricted diagonal uncertainty structure in the input. For the dynamic dissipative case, the problem is to compute $\delta_{\max}$ that satisfies the conditions set in theorem 3.

The approach taken for synthesizing a dynamic dissipative controller is optimization based, just as in the synthesis of the static dissipative controller. Again, nonlinear programming is used to synthesize the controller matrices to provide maximum allowance for uncertainty in the inputs. Using a repeated diagonal uncertainty structure, the elements of the matrices A_c' and Q', which define the dynamic dissipative controller K', are used as design variables and are optimized to maximize a bound on $\delta_{\max}$, for a prescribed degree of stability α, subject to the conditions of theorem 3. The optimization problem is posed as follows.

Minimize the objective function

$$\sup_{\omega} \bar{\sigma}([I + K'P']^{-1} K'P') \tag{20}$$

over the elements of A_c' and Q', more specifically, the parameters a' and q' [see Eq. (16)], that meet the conditions of theorem 3, K' stabilizes P' for all $\Delta = \delta_{\max} \bar{I}$, which may be defined as

$$Re\{\lambda_i(A_{\text{cl}}')\} < 0, \qquad \forall i \text{ and } \forall \Delta = \delta_{\max} \bar{I} \tag{21}$$

as well as side constraints on the elements of the matrices A_c' and Q'. In dealing with the side constraints, one can follow the discussion presented for the static dissipative case.

Note that the preceding fault accommodation formulations refer to the dynamic dissipative controller K'. This is not the real dynamic dissipative controller that is desired because K' was designed for the artificially shifted plant P', not the actual plant P.

Lemma 3: The real controller K can be computed from K' by reverse shifting the eigenvalues of K' by α, i.e., the same amount that the eigenvalues of the plant P were moved to create P', but in the reverse direction. That is, $A_c = A_c' - \alpha I$, where A_c is the state matrix of K and A_c' is the state matrix of K'. The resulting controller K will also be a dynamic dissipative controller provided that the poles of K are to the left of the $-\alpha$ line.

Proof: Let the dynamic dissipative controller K' be defined by the following equations: $\dot{\boldsymbol{x}}_c = A_c' \boldsymbol{x}_c + B_c \boldsymbol{y}$ and $\boldsymbol{u} = -C_c \boldsymbol{x}_c$. Because K' is a dynamic dissipative controller, the following conditions (positive realness conditions) must be true: $A_c'^T R + R A_c' = -Q'$, $R > 0$, and $Q' \ge 0$. By adding $-2\alpha R$ to both sides, the preceding Lyapunov equation can be rewritten as $A_c'^T R + R A_c' - 2\alpha R = -(Q' + 2\alpha R)$, or as $(A_c' - \alpha I)^T R + R(A_c' - \alpha I) = -(Q' + 2\alpha R)$, and finally, as $A_c^T R + R A_c = -Q$, with $A_c = A_c' - \alpha I$ and $Q = Q' + 2\alpha R$. Because $\alpha > 0$, and $2\alpha R > 0$, then $Q > 0$. Now, if A_c' has its eigenvalues to the left of the $-\alpha$ line, i.e., A_c is strictly stable, (A_c, B_c, C_c) also satisfies the positive realness conditions and, therefore, the real controller K will also be a dynamic dissipative controller. Note that neither B_c nor C_c change during the reverse shift operation. The reverse shift operation is unnecessary for the static dissipative controller because that controller has no dynamics associated with it.

μ Controller

The robust performance of interest in this study is to guarantee that real components of all closed-loop eigenvalues remain to the left of $-\alpha$ under predefined uncertainties.

Consider two LTI feedback systems, $G \equiv (A_s, B_s, C_s, 0)$, $K \equiv (a, b, c, d)$ and $G' \equiv (A_s + \alpha I, B_s, C_s, 0)$, $K' \equiv (a + \alpha I, b, c, d)$. From the closed-loop system matrices, it is clear that the eigenvalues Λ of the closed-loop system (G, K) are related to the eigenvalues Λ' of the closed-loop system (G', K') by

$$\Lambda = \Lambda' - \alpha I \tag{22}$$

Let $\Delta \in B_\Delta$, where B_Δ denotes the set of bounded and structured uncertainty. Let P_{22}' and K' be α-shifted systems corresponding to P_{22} and K. Define the augmented plants for linear fractional transformations (LFTs) partitioned as

$$P \equiv \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix}, \qquad P' \equiv \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22}' \end{bmatrix} \tag{23}$$

and the lower LFT

$$\begin{aligned} F_l(P', K') &= P_{11}' + P_{12}' K' (I - P_{22}' K')^{-1} P_{21}' \\ &= P_{11} + P_{12} K' (I - P_{22}' K')^{-1} P_{21} \end{aligned} \tag{24}$$

The following summarizes the α degree of stability robustness with respect to structured uncertainty.

Theorem 4: The eigenvalues of the set of perturbed closed-loop system (P, K, Δ), $\forall \Delta \in B_\Delta$ remains on the left of $-\alpha$ iff

$\sup_\omega \mu(F_l(P', K')_{11}) \leq 1$, where the subscript 11 denotes the upper-left partition of ( ), which in this case corresponds to the uncertainty.

Proof: From Doyle,[6] the closed-loop system (P', K', Δ) will be robustly stable $\forall \Delta \in B_\Delta$ iff $\sup_\omega \mu(F_l(P', K')_{11}) \leq 1$. Then from Eq. (22), (P, K, Δ) will be robustly stable with degree α.

The design implication is clear from theorem 4. The performance index, $\sup_\omega \mu(F_l(P', K')_{11})$, can be minimized with respect to the given uncertainty structure, to attain or get close to α degree of robust stability. The actual minimization can be carried out using D–K iteration via μ tools.[8]

The design strategy is summarized in the following steps:

1) First α shift P_{22} in Eq. (23) to obtain the shifted plant P'_{22}.

2) Next, use μ synthesis to design a controller K' to minimize the $\sup_\omega \mu(F_l(P', K')_{11})$, which will maximize the uncertainty allowed.

3) Last, reverse α shift K' to obtain the implementing controller K.

Numerical Examples

To demonstrate the feasibility of the proposed control synthesis techniques, they were applied to the control design of the Phase-II CSI Evolutionary Model (CEM), a testbed at NASA Langley Research Center for ground-based experimental studies of multipayload space platforms. The structure, shown in Fig. 2, consists of a 620-in.-long aluminum main truss with several appendages. Three two-axis gimbals, each implemented with a laser source and a high-precision scoring system, are attached to the main bus to simulate science payloads on a space platform. Eight bidirectional thrusters are used for platform control and/or disturbance generation, whereas eight collocated accelerometers provide acceleration measurements for feedback control. These acceleration signals are integrated with the aid of washout filters to provide collocated velocity measurements. A 20th-order state-space model of the structure was used in the analysis.

Static Dissipative Controller

A static dissipative controller was synthesized using the nonlinear programming approach described earlier. Using the eight diagonal elements of the rate gain matrix as design variables, and the desired degree of stability α taken as 0.30, the elements of the rate gain matrix are optimized to maximize the fault tolerance of the closed-loop system. Here, the optimization problem presented for the static dissipative controller case used Eq. (11), i.e., the optimal gains were computed by minimizing the H-infinity norm bound of the closed-loop system, as formed by the nominal α-shifted plant P' and the designed controller K, to maximize $\delta_{\max}$. Constraints were placed to ensure that the set of eigenvalues of the closed-loop system were stable. However, closed-loop stability checks based on theorem 1 were not done in the controller synthesis procedure. The automated design synthesis (ADS) software[9] was used to carry out the optimization. The interior penalty function method of ADS was used to solve the nonlinear programming problem. The initial values of the gain elements were set at 1.0, with the lower bound and upper bound values set at 0.001 and 1000, respectively. The initial value of the objective function (H-infinity norm of the closed-loop transfer function at the input) was 5.51. The optimization decreased the objective function to a value of 2.97 (a 46% reduction), thereby increasing the upper bound constraints on the uncertainty norm by almost a factor of two. The true measure of uncertainty allowance, however, is almost inversely proportional to the maximum value of the spectral radius over the frequency spectrum [see Eq. (10)], which turned out to be 2.71 for the optimal design. This provides for an input multiplicative uncertainty bound of 0.36 (equal to 1/2.71). It should be noted that the optimal controller provides robust performance, in the form of closed-loop time constants of less than 33 s, for diagonal multiplicative uncertainty in the input of less than or equal 36%. The optimal rate gain matrix was $G_r = \text{diag}(0.86, 1.93, 1.74, 5.44, 2.99, 1.25, 0.46, 0.77)$. To further demonstrate these results, a 1000 run Monte Carlo simulation was performed wherein the signal in each input channel was randomly and independently modified by as much as $\pm 36\%$, using MATLAB's random number generating function with normal distribution. Figure 3 shows the locations of two poles of the closed-loop system closest to the imaginary axis, as computed from the Monte Carlo simulation. As expected, the loci of these poles stayed to the left of the -0.3 line.

To evaluate the effect of the α level on the uncertainty that can be tolerated, several static dissipative controllers were synthesized for various α levels, using the nonlinear programming technique described earlier. Figure 4 shows the result of this trade study. Here, it is observed that, as expected, performance and robustness to uncertainty follow an inverse relationship. The higher the performance levels are, the smaller the uncertainties tolerated, and the smaller the performance levels are, the larger the uncertainties allowed. The

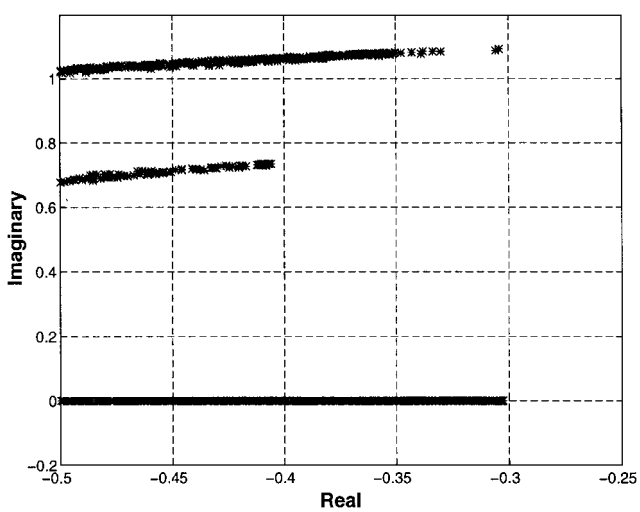

Fig. 3 Loci of the closed-loop poles closest to the imaginary axis, static dissipative controller.

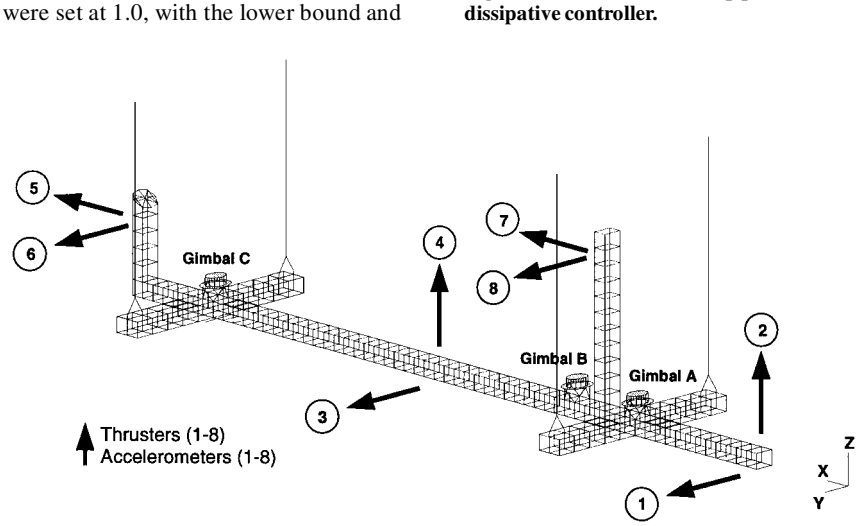

Fig. 2 Schematic of the Phase-II CEM.

Table 1 Maximum spectral radius values for the static dissipative controllers

Degree of stability, α	Maximum spectral radius	δ_{max}
0.05	1.0357	0.9655
0.1	1.1539	0.8666
0.2	1.6025	0.6240
0.3	2.7116	0.3687
0.4	15.0675	0.0664

Table 2 Uncertainty allowance trade study

Degree of stability, α	Static dissipative controller uncertainty allowance	Dynamic dissipative controller uncertainty allowance
0.05	0.961	0.960
0.1	0.857	0.858
0.2	0.607	0.611
0.3	0.336	0.337
0.4	0.060	0.087

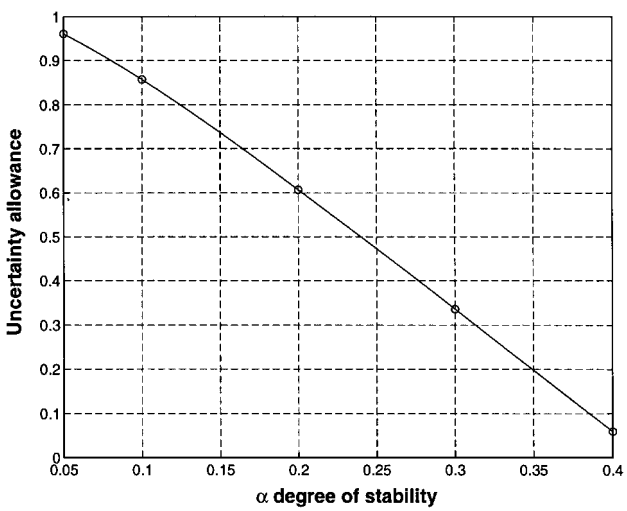

Fig. 4 Tradeoff between performance and uncertainty, static dissipative controller.

level of performance is bounded on the right by the location of the finite zeros of the flexible plant because as the gain matrix is increased a number of closed-loop poles (equal to the number of finite zeros of the system) approach the finite zeros. On the left, the tradeoff curve continues to rise as the performance level α decreases. If the performance level α is the only driving factor in the controller design, then as the α level drops below the performance level of the open-loop system, the uncertainty level that can be tolerated goes beyond 1.0 (allowing 100% or more variation in the inputs), and the control gain would approach zero (because no control is necessary to obtain desired performance).

To verify these results with respect to theorem 1, the maximum spectral radius values for the closed-loop systems formed with the described optimal static dissipative controllers were computed to determine the less conservative maximum input uncertainty levels (see Table 1). See Table 2 to compare the maximum input certainty levels for these same controllers as computed from the H-infinity norm bound. Using the maximum input uncertainty levels from Table 1, the stability of the closed-loop system, formed with each controller, was checked by computing $Re\{\lambda(A_s - B_s(I + \delta_{max}\bar{I})G_r B_s^T)\}$, for all possible permutations of $\bar{I}$, i.e., the stability conditions of theorem 1. The condition for degree of stability set in theorem 1 was considered violated if any eigenvalue lay to the right of the specified $-\alpha$ line for the given controller. This is equivalent to the condition set in theorem 1, where the eigenvalues of the closed-loop system formed by the α-shifted plant, A_s', and controller must lie to the left of the $j\omega$ axis. A violation would indicate that the particular controller could not guarantee the desired closed-loop α degree of stability for all possible distinct input uncertainties up to the specified δ_{max} level. For a controller with m inputs, there would be $2^m - 2$ permutations of the matrix $\bar{I}$, excluding the permutations $\bar{I} = I$ and $\bar{I} = -I$. In the example, with 8 inputs, 254 permutations of $\bar{I}$ and, thus, 254 sets of closed-loop eigenvalues were computed for each static dissipative controller. The results of these computations show that for those controllers corresponding to α ranging from 0.05 to 0.3, the closed-loop eigenvalues lay to the left of the $-\alpha$ lines for all permutations of $\bar{I}$, thus satisfying the requirement set in theorem 1. This supports the effectiveness of the optimization problem posed for the synthesis of a static dissipative controller without including the side constraints of theorem 1. However, for the $\alpha = 0.4$ controller, there were several permutations of $\bar{I}$ where some of the closed-loop eigenvalues lay to the right of the -0.4 line; thus, that specific static dissipative controller would not be able to guarantee a closed-loop system degree of stability of 0.4 for all possible distinct input uncertainties of up to 0.0664. Further computations revealed that the maximum allowable input uncertainty level, to guarantee an $\alpha = 0.4$ degree of stability, would be 0.0663. Conversely, if a maximum uncertainty allowance of 0.0664 was truly desired, then the achievable guaranteed degree of stability would be a slightly lower value of 0.39998. Again, this shows the effectiveness of this synthesis procedure and how well the spectral radius bound in Eq. (10) approximates the allowed uncertainty tolerance. If the design does not satisfy the conditions of theorem 1, one can easily determine what performance level it can provide.

Incorporating both the maximum spectral radius method of calculating δ_{max} (instead of the H-infinity norm bound method) and permutation calculations would make the controller design synthesis procedure either too computationally intensive or intractable, particularly if large numbers of inputs and/or system states were involved. The current design synthesis procedure gives good results, which can be checked and adjusted as needed off line.

Dynamic Dissipative Controller

A dynamic dissipative controller was synthesized, using the same nonlinear programming approach used to create the static dissipative controller. For the dynamic dissipative case, the optimization problem involved minimizing Eq. (20). Like the static dissipative controller design example, rate information was used as the feedback signals. For the desired closed-loop α degree of stability of 0.3, the a and q parameters were optimized to maximize the input uncertainty allowance to the plant. Again, the ADS software package was used to perform the optimization. For each of the eight sensor/actuator pairs, a second-order controller was synthesized, making the total order of the complete dynamic dissipative controller equal to 16. As initial conditions, the controller for each sensor/actuator pair was set as

$$G_{ri}\frac{100(s+50)}{(s+100)(s+50)} \tag{25}$$

where G_{ri} was the optimal gain associated with the ith sensor/actuator pair of the static dissipative controller described in the earlier numerical example. The initial pole and zero locations were selected to make the dynamic dissipative controller behave like the constant-gain static dissipative controller in the frequency bandwidth of the plant model (34 rad/s). The ADS program was free to adjust these pole and zero locations, as well as the gains, during the optimization. The optimization decreased the objective function to 2.96, which was slightly lower than the static dissipative controller. The uncertainty allowance value, as computed by the maximum of the spectral radius over the frequency spectrum of interest, was 2.70, which corresponds to a multiplicative uncertainty bound of 0.37. The pole and zero locations of the optimal dynamic dissipative controller changed very slightly from those of the initial values. The largest change occurred in the controller portion associated with the second sensor/actuator pair: the optimal pole locations were at -50.019 and -99.981; the optimal zero location was at -49.973; and the optimal dc gain was 1.931, where the initial, i.e., the optimal gain for the second sensor/actuator pair of the static dissipative controller, dc gain was 1.933.

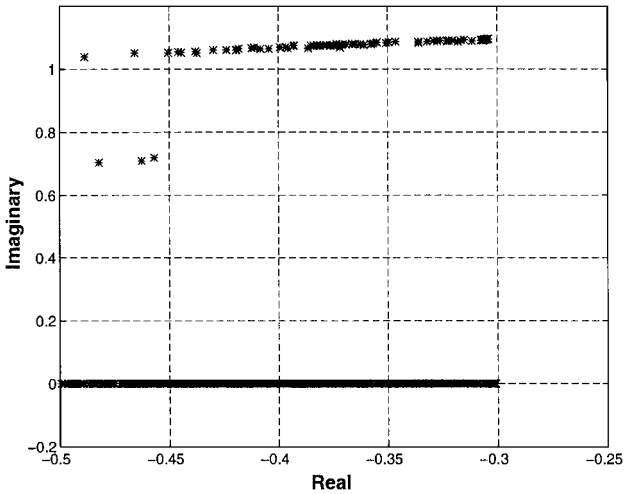

Fig. 5 Loci of the closed-loop poles closest to imaginary axis, dynamic dissipative controller.

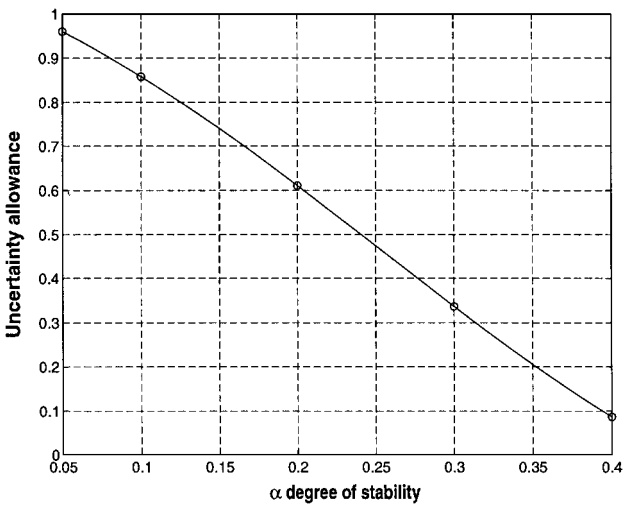

Fig. 6 Tradeoff between performance and uncertainty, dynamic dissipative controller.

To verify these results, a 1000 run Monte Carlo simulation was performed wherein the signal in each input channel was randomly and independently modified by as much as ±37%. Figure 5 shows the locations of two poles of the closed-loop system closest to the imaginary axis for the simulation. As expected, the loci of these poles stayed to the left of the −0.3 line.

To evaluate the effect of the α level on the uncertainty that can be tolerated, several dynamic dissipative controllers, again using the static dissipative controller optimal gains as initial conditions, were synthesized for various α levels using the nonlinear programming technique described earlier. Figure 6 shows the result of this trade study. Here, as was observed in the static dissipative controller designs, performance and robustness to uncertainty follow an inverse relationship. The higher the performance levels are, the smaller the uncertainties tolerated, and the smaller the performance levels are, the larger the uncertainties allowed. When comparing Figs. 4 and 6, it can be observed that the dynamic dissipative controllers were able to achieve, if only marginally, better performance than the static dissipative controllers. This was attributed to the increased design freedom of controller zero and pole placements to better affect the system's overall dynamics. Other dynamic dissipative controller designs, with higher required degree of stability α, achieved more noticeable improvements over the static dissipative controller. Table 2 lists the uncertainty allowance values, as computed from the H-infinity norm bound, for several degrees of stability α, taken from the respective static and dynamic dissipative controller trade studies.

These dynamic dissipative controllers were tested with respect to the stability requirement set in theorem 3. As in the case with

Table 3 Maximum spectral radius values for the dynamic dissipative controllers

Degree of stability, α	Maximum spectral radius	$\delta_{\max}$
0.05	1.0376	0.9638
0.1	1.1525	0.8667
0.2	1.5980	0.6258
0.3	2.7027	0.3700
0.4	10.4592	0.0956

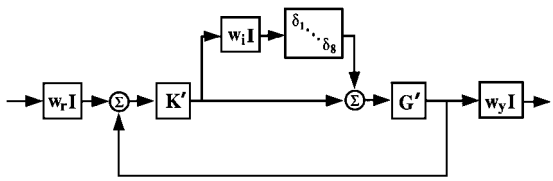

Fig. 7 Block diagram for the robust stabilization problem.

the static dissipative controllers in the preceding section, the maximum spectral radii of the closed-loop systems formed by these dynamic dissipative controllers were computed. These values, along with their corresponding maximum input uncertainty allowances $\delta_{\max}$ are listed in Table 3. Using the maximum input uncertainty levels from Table 3, the degree of stability of the closed-loop system, formed with each controller, was checked by computing the real parts of the eigenvalues of the closed-loop system $[A_s \;\; -B_s(I + \delta_{\max}\bar{I})C_c;\; B_cB_s^T \;\; A_c]$, for all possible permutations of $\bar{I}$, where (A_c, B_c, C_c) are the dynamic dissipative controller matrices. If any eigenvalue lay to the right of the specified $-\alpha$ line for the controller in question, then the condition given in theorem 3 was considered violated, i.e., that particular controller could not guarantee the desired closed-loop α degree of stability for all possible distinct input uncertainties up to the specified $\delta_{\max}$ level. Again, this is equivalent to the requirement given in theorem 3, where the eigenvalues of the closed-loop system $[A_s' \;\; -B_s(I + \delta_{\max}\bar{I})C_c;\; B_cB_s^T \;\; A_c']$, i.e., with the shifted plant and controller, must lie to the left of the $j\omega$ axis. For each controller, 254 sets of closed-loop eigenvalues were computed. The results of these computations show that for those controllers corresponding to α ranging from 0.05 to 0.3, the closed-loop system eigenvalues lay to the left of the $-\alpha$ lines for all permutations of $\bar{I}$, thus satisfying the requirement set in theorem 3. However, for the $\alpha = 0.4$ controller, there were several permutations of $\bar{I}$ where some of the eigenvalues lay to the right of the −0.4 line; thus, that specific dynamic dissipative controller would not have been able to guarantee a closed-loop system degree of stability of 0.4 for all possible distinct input uncertainties of up to 0.0956. Further computations revealed that the maximum allowable input uncertainty level, to guarantee an $\alpha = 0.4$ degree of stability, would be 0.091. Conversely, if a maximum allowable uncertainty allowance of 0.0956 was truly desired, then the achievable guaranteed degree of stability would be 0.3977.

Similar to the static dissipative controller synthesis procedure, the dynamic dissipative controller synthesis procedure shows its effectiveness for providing fault accommodation. If the design does not satisfy the conditions of theorem 3, one can easily determine what performance level it can provide.

μ Controller

Actuator failures for any combination or number are modeled as independent scalar multiplicative uncertainties in all channels at the input to the CEM structure. This results in a diagonal multiplicative uncertainty at the input. Although not strictly necessary for the sole purpose of optimizing stability robustness, weighted external inputs and outputs were introduced to formulate the problem in terms of a performance robustness problem for a general LFT. For this study, very small values for the external input–output weights were chosen to emphasize only the μ of the robust stability block. The D–K iteration method was used to design μ controllers for all cases using only constant scales. Figure 7 shows the α shifted closed-loop

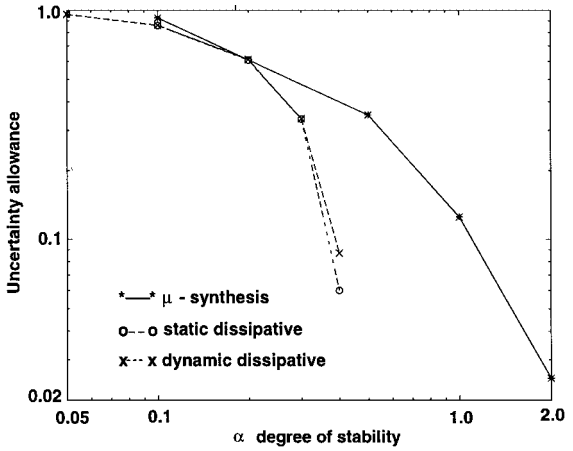

Fig. 8 Tradeoff between performance and uncertainty, μ synthesis and static and dynamic dissipative cases.

system used for controller design. The magnitude of the diagonal uncertainties, $\delta_1, \ldots, \delta_8$, were assumed to be bounded by unity.

Figure 8 shows the tradeoff between the desired degree of stability, α, and the corresponding maximum size of the uncertainty tolerated. For a given α value, the uncertainty weight was increased until the $\sup_\omega \mu[F_l(P', K')_{11}]$ approached unity. As in the dissipative case, the general trend is that as α increases linearly, Δ decreases exponentially. Although not shown in Fig. 8, an arbitrary level of uncertainty can be accommodated at $\alpha = 0$. This corresponds to the special condition where the given structural plant is stable and the controller basically opens the loop to guarantee robust stability. Note that because the α shift does not change the plant eigenvectors, the controllability and observability of the α shifted nominal structural plant will not change.

Although the inputs and outputs of the structure were physically collocated, the input and output responses are not diagonally dominant. Therefore, it is not surprising that the controllers obtained had large, but not dominant, diagonal components. In addition, the diagonal nature of the uncertainty did not significantly increase robustness over the corresponding unstructured case. It is also noted that μ controllers are generally not positive real; hence, these controllers, although they are robust to particular uncertainty, are not automatically robust to parametric uncertainty and spillover problem. However, additional uncertainty models can be introduced to robustify the μ controller. Simulation was also done to verify α degree of stability for the set of perturbed systems.

Concluding Remarks

This paper presents the development of novel methodologies that allow direct synthesis of controllers that guarantee robust stability with a prescribed degree. This is useful for fault accommodation because robust stability and performance are of primary concern. For soft failures in the form of large degradation in the actuator/sensor signals, both dissipative- and H-infinity-based controllers can significantly mitigate significant (up to 20 or 30%) errors and still guarantee various degrees of stability robustness. The μ-based compensator consistently provided larger uncertainty tolerance (for a desired degree of stability) than the static dissipative controller. The dynamic dissipative controller was second best in terms of providing robust performance in the presence of failures, particularly for large desired degrees of stability. Nonetheless, the constant diagonal nature of collocated velocity feedback provides good performance when compared to the much more complicated μ or dynamic dissipative compensators. The simplicity in the static dissipative controller and its guaranteed stability under hard failures with mild assumptions is certainly attractive.

References

[1] Joshi, S. M., *Control of Large Flexible Space Structures*, Springer–Verlag, Berlin, 1989.

[2] Joshi, S. M., and Maghami, P. G., "Dissipative Compensators for Flexible Spacecraft Control," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 28, July 1992, pp. 768–774.

[3] Kwakernaak, H., "Asymptotic Root Loci of Multivariable Linear Optimal Regulators," *IEEE Transactions on Automatic Control*, Vol. 21, June 1976, pp. 378–382.

[4] Kouvaritakis, B., "The Optimal Root Loci of Linear Multivariable Systems," *International Journal of Control*, Vol. 28, July 1978, pp. 33–62.

[5] Williams, T., "Transmission-Zero Bounds for Large Space Structures, with Applications," *Journal of Guidance, Control, and Dynamics*, Vol. 12, No. 1, 1989, pp. 33–38.

[6] Doyle, J. C., "Analysis of Feedback Systems with Structured Uncertainties," *IEEE Proceedings*, Vol. 129, Pt. D, No. 6, 1982, pp. 242–250.

[7] Lefschetz, S., *Stability of Nonlinear Control Systems*, Academic, New York, 1965.

[8] Balas, G. J., Doyle, J. C., Glover, K., Packard, A., and Smith, R., "μ-Analysis and Synthesis Toolbox," Mathworks, Natick, MA, Nov. 1995.

[9] Vanderplaats, G. N., "Ads—A Fortran Program for Automated Design Synthesis—Version 1.10," NASA CR 177985, 1985.