JOURNAL OF AIRCRAFT
Vol. 36, No. 1, January–February 1999

Evolutionary Design of Controlled Structures

Brett P. Masters* and Edward F. Crawley†
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

An evolutionary controls/structures design method is developed. The basis of the method is an accurate model formulation for dynamic compensator optimization coupled with a genetic-algorithm-based updating of sensor/actuator placement and structural properties. One- and three-dimensional examples from the literature are used to validate the method. Frequency domain interpretation of the controlled structure systems provide physical insight as to how the objective is optimized. Several designs for a stellar interferometer precision structure are found for two different spectra of disturbance under assumed closed-loop optical conditions. Physical limitations in achieving performance are given in terms of average system transfer function gain and system phase loss.

Introduction

SPACE-BASED structures that are characterized and compensated to submicron levels have found utility in the stellar observation sciences. In such structures, control, other than that used in maintaining attitude, is typically introduced in the preliminary and detailed design stages, after the system is found to fail specifications passively. At this stage the structure is fixed in topology and member geometry, leaving the control designer to accept the given plant dynamics. In some cases the actuator/sensor design is also fixed, further limiting the achievable performance.

Consideration of structural control technology early in the design process of precision structures leads to possible benefits, but it also leads to numerous design variables and, subsequently, many potential cost functions, making the combined optimization problem very difficult. The payoff is that the actuators/sensors are simultaneously designed with the system, and can subsequently render the controls with greater influence over improving the performance. The drawback is that the combined optimization problem is plagued by large dimensions in both the controls and structural problems, and furthermore, it becomes combinatorial with the addition of discrete choices such as sensor/actuator location (distribution) and type, e.g., inertial or relative.

The objective of this paper is to provide a consistent method of designing a precision structure that leverages the use of control to meet performance requirements. In using the method, the necessity of considering both the sensor/actuator design and dynamic compensation in the system design process is shown.

In the past 20 years various investigators have tackled the dynamic controls—structures optimization problem, yielding a myriad of results and sparse implementation. Rao et al.,[1] Maghami,[2] Canfield et al.,[3] and Miller and Shim[4] provide approaches where controls and structures topologies are fixed and the control gains, control bandwidth, and member cross-sectional variables optimized with respect to either mass, motion error, or control effort costs. For the most part, the preceding items use nonlinear programming techniques and cover a variety of the control techniques [linear quadratic regulator (LQR), H_2/H_∞, and positive real techniques, respectively]. They do not consider actuator/sensor placement variation simultaneously with structural variation. Onoda and Haftka[5] consider both simultaneously, but the feedback is constant gains (LQR), and the actuation is coupled to rigid-body control while the objective is mass minimization. This will serve as a validation example.

Sepulveda et al.[6] provide a method that uses branch and bound techniques that places collocated actuators and sensors in a truss structure while optimizing the member cross sections. The control method is local position and rate output feedback. Using a branch and bound technique allows the integer placement problem to be solved using continuous problems at each branch. While Sepulveda et al.'s work does consider structural and actuator/sensor placement variations simultaneously, the optimization uses suboptimal (but robust) output feedback and uses an intermediate variable formulation for the model. The model formulation allows simple calculations of gradients but foregoes the accuracy required for control.

Open-loop structural topology optimization is performed and experimentally implemented by Keane,[7,8] with good narrowband disturbance rejection results. Here a genetic algorithm (GA) search is performed over structural topology (geometric location of truss joints) to minimize high-frequency narrowband energy levels in a structural member. Good performance (30-dB reduction in energy metric at a given frequency) is achieved by adjustment of the local dynamics (via length change) of the individual members. In these results it is unclear whether broadband performance is compromised.

Results from the literature are numerous and are difficult to compare because of various factors:

1) Dimensionality of the structure (beam vs truss vs even more complex in nature) and structural discretization (global vs local dynamics).

2) Sensor/actuator choice, differing in type, e.g., internal relative vs external inertial, and location (only sometimes variable).

3) Differing disturbances/performances.

4) Differing control techniques.

The question still remains: how does one design a structure to leverage control, i.e., what are the necessary considerations, besides choosing an optimizer, and how well will the design perform in practice? The answers require a uniform investigation, from a systems perspective, over typical precision structures and experimental verification. An investigation of an interferometer precision structure will be given here. Experiments have been performed[9] that validate the results presented in this paper.

Received May 19, 1997; revision received May 27, 1998; accepted for publication Aug. 9, 1998.

*Graduate Research Assistant, Space Systems Laboratory; currently Head of Engineering, Mide Technology Corporation, 247 Third Street, Cambridge, MA 02141. Member AIAA.

†Head, Space Systems Laboratory, Department of Aeronautics and Astronautics, and McVicar Faculty Fellow, 33-201 MIT. Fellow AIAA.

Development of the evolutionary design methodology will be given first. The method is then validated on one- and three-dimensional examples from the literature. Solutions to these examples are investigated in the frequency domain. Attention is then turned to a one-dimensional stellar interferometer design example. This application example will provide guidance and physical interpretation of what the optimal interferometric structure is, given the critical dimensions and performance emphasis, while simultaneously varying the sensor/actuator design and dynamic compensator. Application example cases will cover differing disturbance spectra and sensor suites. Insight into the optimized designs is given by system frequency domain interpretations in terms of pole-zero structure and average transfer function gain.

Method Formulation

A method was developed that can handle controlled structures with many degrees of freedom (DOF), topological variations, and dynamic compensation techniques. Figure 1 outlines the proposed algorithm. The method uses zeroth-order objective/fitness evaluations only, with a GA search outer loop.

In the next section, modeling and cost evaluation are first addressed, followed by a discussion of the design evolution. The discussion on modeling begins with structural models generated from component formulation, condensation, and synthesis. Eigen solutions are then found for the synthesized models and state-space models formed for control design. In the application example the proposed method for finding dynamic controllers is H_2 optimal design. The motion error part of the linear quadratic H_2 cost serves as the performance used for design ranking. Different control techniques and performances are used in the validation examples.

For most structures of realistic dimension, a full model would be too large to compute during each evaluation step. To address this problem, the structure is first discretized into component models. Given each new set of structural parameters (attributes) component models are assembled using finite elements:

$$\begin{aligned} M_i\ddot{p}_i + C_i\dot{p}_i + K_i p_i &= \beta_{w_i} w_i + \beta_{u_i} u_i \\ y_i &= c_{yp_i} p_i + c_{y\dot{p}_i}\dot{p}_i \\ z_i &= c_{zp_i} p_i + c_{z\dot{p}_i}\dot{p}_i \end{aligned} \tag{1}$$

where p is the structural DOF, w are the disturbances, u is the actuation, z are the performances, and y are the measurements. Within the components, enough fidelity is included to capture variation in local dynamics with changes in grid point location, member cross section, and actuator/sensor location.

To reduce and synthesize component models into a manageable global model, component mode synthesis (CMS) is used.[10] Active locations within the components, i.e., those actuated, disturbed, sensed, and performing, are treated as interface DOF. Constraint modes are used for integral structural actuators such as piezoelectric struts in a truss, whereas attachment modes are better suited to inertial actuators. It is necessary to treat the active DOF with care to preserve the integrity of the system transfer functions, which in turn, will determine the ability of the system to be controlled. Treatment of active DOF as interface DOF preserves their static response, which is important for preserving the system zero locations that enable or inhibit control.

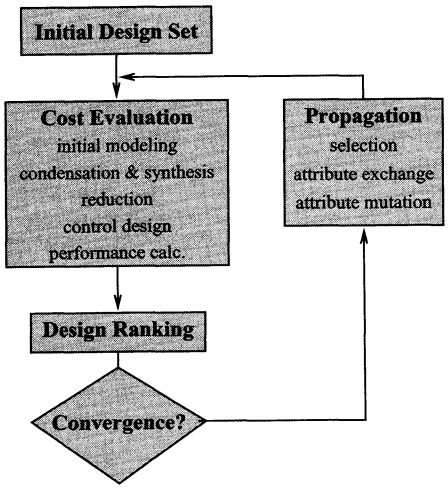

Fig. 1 Methods overview.

Constraint equations are used to combine the condensed components. For the most part, this is simply assembly of the interface DOF. If a component does not include active locations and if its internal dynamics are above the bandwidth of important performance, then quasistatic information is enough and internal DOF are not kept.

An eigen solution is now found for the CMS-synthesized high-fidelity model

$$(K_{\text{syn}} - \omega^2 M_{\text{syn}})\psi_d = 0 \tag{2}$$

with mass normalization

$$\psi_d^T M_{\text{syn}} \psi_d = I \tag{3}$$

Because this operation is performed for every model during each propagation cycle, it is necessary to keep the order of the synthesized system low without compromising the accuracy of the model. For example, local dynamics near active locations should be kept.

For control purposes, the model is now further reduced using only modal information. This is accomplished using the approximate balanced singular value reduction developed by Gregory.[11] In Gregory's work, two modes are decoupled for the purposes of balancing if their settling times are long compared to the time it takes for the motion of the modes to move in and out of phase. This will be true if

$$\frac{\max(\zeta_i, \zeta_j)\max(\omega_i, \omega_j)}{|\omega_i - \omega_j|} \cong \frac{T_{\text{beat}}}{T_{\text{settling}}} \ll 1 \tag{4}$$

Given that the preceding is approximately true, the ensuing relations are used for the approximate singular values:

$$\begin{aligned} \sigma^2_{zw_i} &\cong \frac{\beta_{w_i}\Sigma_{ww}\beta^T_{w_i}}{(4\zeta_i\omega_i)^2}\left\{\frac{c^T_{z_i}R_{zz}c_{z_i}}{\omega_i^2} + c^T_{\dot{z}_i}R_{zz}c_{\dot{z}_i}\right\} \\ \sigma_i^2 &= \alpha_{zw_i}\sigma^2_{zw_i} + \alpha_{zu_i}\sigma^2_{zu_i} + \alpha_{yw_i}\sigma^2_{yw_i} + \alpha_{yu_i}\sigma^2_{yu_i} \end{aligned} \tag{5}$$

where the subscript i refers to the mode number or corresponding row and column of the β and c matrices, respectively. The matrices R and Σ have been included to weight the disturbances w, controls u, performances z, and sensors y, relatively. Equation (5) is an example of computing the approximate singular values, $\sigma^2_{zw_i}$, for disturbance to performance. Equation (5) shows how they are combined in a weighted sum. Weighting of the inputs and outputs with respect to each other is included in the coefficients, α. Transfer function weights such as those on performances and disturbances are evaluated at every modal frequency and are included in the coefficients. Normalization of each of the I/O approximate singular values by the maximum singular value is also incorporated into the coefficients. The σ_i of Eq. (5) are ranked as the highest modes kept, or, a threshold may be set and the model order may naturally vary.

The assumption is that removing small residue modes will not degrade the ability to do control on a given plant

$$\begin{Bmatrix} z \\ y \end{Bmatrix} = \begin{bmatrix} G_{zw} & G_{zu} \\ G_{yw} & G_{yu} \end{bmatrix} \begin{Bmatrix} w \\ u \end{Bmatrix} \tag{6}$$

$$u = -Ky \tag{7}$$

where K is the dynamic compensation and G are system transfer functions. The explicit dependence on the Laplace variable s has been dropped. If there were a destabilizing mode with a small residue, the optimal dynamic compensator would, generally speaking, invert this pole, limiting the stability robustness of the controller design to this pole. This leads to the conclusion that this mode should be more disturbable and observable in the performance, i.e., have larger σ_{zw_i}. This is exactly what most robustification techniques will do for a given actuator/sensor suite.

After model reduction there is a necessary procedure that preserves the integrity of the plant model for the purposes of control. When accurately predicting the ability to control a plant, in closed loop it is necessary to preserve the system transmission zeros.[12] To do this, static modes are found for every disturbance and actuator, and appended in a state-space format.

Weights are incorporated in the system that emulate disturbance frequency content and possible disturbance/performance isolation. Further weights provide a loop-shaping capability, on u and y, that must reflect a realistic capability of the actuator/sensor topology. The system is next converted to real modal form for conditioning purposes.

In the application example, standard H_2 control design is implemented to minimize the linear quadratic cost

$$J = (\|G^{cl}_{zw}\|_2)^2 \tag{8}$$

where, for controller computation, z includes the estimator error and weighted control penalties as well as motion penalties, whereas w includes modeled disturbances and sensor noises. The closed-loop disturbance to performance transfer function can be expanded as

$$G^{cl}_{zw} = G^{ol}_{zw} - G_{zu}K(I + G_{yu}K)^{-1}G_{yw} \tag{9}$$

where, when considering only the motion error in z and the system disturbances in w, G_{zu} is the regulator transfer function matrix, G_{yw} is the estimator transfer function matrix, and G_{yu} is the actuators to sensors transfer function matrix. Equation (9) still holds when z includes estimation error and control penalties and w includes sensor noises, but now G_{zu} and G_{yw} have different meanings, depending on the interpretation of z and w. The solution for the dynamic compensator (K) is found from the solutions of the usual regulator and estimator Riccati equations.[13] These Riccati equations are solved using eigenvalue decomposition of the Hamiltonian matrices, which are perturbed if nondistinct eigenvalues exist.

The objective of Eq. (9) is clear. It is desired to make G^{cl}_{zw} small by shaping the dynamic compensator K to make the last term $G_{zu}K(I + G_{yu}K)^{-1}G_{yw}$ cancel G^{cl}_{zw}, within the constraints of stability. Roughly, if K is large, then the last term is $\sim G_{zu}G^{-1}_{yu}G_{yw} \sim G^{ol}_{zw}$. If, for example, we choose an analogy between measurements y and motion error performances z, then Eq. (9) reduces to

$$G^{cl}_{zw} = S_{yu}G^{ol}_{zw} \tag{10}$$

where S_{yu} is the sensitivity of the control loop $(I + G_{yu}K)^{-1}$. When, also, the actuator is collocated with y, the sensitivity S_{yu} can be shaped to achieve extremely good disturbance rejection (large K) over the controller bandwidth. Note, if u is made analogous to the disturbance w in Eq. (9), then a similar effect occurs showing that isolating the disturbance by directly acting on the disturbing forces is also a powerful performance enhancer in the controlled bandwidth.

The performance that is used to rank the designs in the application example is that part of J pertains to motion error only, J_{motion}. That is, G^{cl}_{zw} is computed, where z only includes motion errors, whereas w represents modeled disturbances and sensor noises. The result is scaled for the GA as

$$F_{\text{obj}} = M - 10 \log J_{\text{motion}} \tag{11}$$

where M ensures that F_{obj} is positive.

Even though the controller is only guaranteed stable on the model it is given, the resulting closed-loop performance is consistently limited by the loss of phase of the modeled system. In the bandwidth, unmodeled (truncated or poorly discretized) dynamics may still limit actual implementation, and it is this threat that requires the method be validated in design model experiments.[9]

Use of the H_2 control-design technique allows consideration of parameter robustness through noise modeling and unmodeled dynamics via frequency weights. At this stage in a design, it is unnecessary to deal with absolute guarantees of stability or performance robustness. For the application examples that follow, extra performances were added to the system that additionally penalize structural motion.

For the outer-loop optimization or the search, the GA appeals[14-16] because of its ability to handle discrete topological changes easily. In his experiences with optimizers and open-loop structural topology problems, Keane[17] finds the GA to work best. The key to the GA is effective encoding of the system properties in the design space, e.g., nodal locations, member cross sections, connectivity arrangements, sensor/actuator placement.

The method proceeds by a first-order propagation of a discrete sample space of systems. Initial sets of designs are usually chosen randomly, to populate the design space, and may require some projection to be feasible. Bounding the design space is often necessary as a multitude of options exist in topological variation. A simple representation of the design space is

$$\begin{gathered} \Omega = \{X^P, Y^Q, \ldots\} \\ \{x^1_j \quad x^2_j \quad \cdots \quad x^1_j\} \in X \\ \{y^1_j \quad y^2_j \quad \cdots \quad y^m_j\} \in Y \end{gathered} \tag{12}$$

where X and Y represent particular attribute types, e.g., nodal locations and member cross sections, with superscripts, P and Q, that represent the order of variation within the type. For example, for 50 structural members, each with a 3-bit representation for the cross section, a variation of order 8^{50} results. A particular element i of a type x for the jth design is denoted x^i_j. Propagation follows selection according to fitness (scaled objective) f_j, whereby the best designs are given a higher chance of proceeding. Attributes from the randomly selected designs are transposed to yield new designs in a process called crossover. A standard crossover operation for attributes of type x is

$$\begin{gathered} (c^1 \cdots c^l) \in \{0, 1\} \\ \begin{matrix} x^1_j \cdots x^l_j \\ x^1_k \cdots x^l_k \end{matrix} \rightarrow \begin{matrix} \hat{x}^1_n \cdots \hat{x}^1_n \\ \hat{x}^1_{n+1} \cdots \hat{x}^l_{n+1} \end{matrix} \\ \hat{x}^p_n = \begin{cases} x^p_j & \text{if} \quad c^p = 1 \\ x^p_k & \text{otherwise} \end{cases} \\ \hat{x}^p_{n+1} = \begin{cases} x^p_k & \text{if} \quad c^p = 1 \\ x^p_j & \text{otherwise} \end{cases} \end{gathered} \tag{13}$$

The crossover mask c is typically two uniform concatenated strings of ones and zeros, e.g., $c = (1, 1, 1, 1, \ldots, 1, 1, 0, 0, 0, \ldots, 0, 0, 0)$. In more complex GAs, rules are applied during the crossover to emulate biological functions such as dominance and learning by altering the regularity of c. In this method, crossover is performed in each individual chromosome. This crossover is not considered simple because it is information exchange in multiple chromosomes. Such information exchange may sever good strings of information, and thus, retard the method, or it may serve to compartmentalize the important information, and thus, accelerate the method. Before the new children designs are evaluated, a small amount of mutation introduces random allowable changes to the system attributes.

Calculated fitness values of the new designs [see Eq. (11)], generated from propagating attributes, are compared to the population from which they came and the best half of the total pool propagated. Identical designs are eliminated from the best half in favor of the next-best performers so that the next-generation designs maintain some diversity.

Constraints are usually dealt with using penalty functions and projection methods that find the closest design in the feasible design space.

Operation of the GA can be understood by a growth equation for good schemes of attributes. Let a scheme of good attributes be H, where

$$H = \{z_j^p, z_j^{p+1}, \ldots, z_j^{p+n}\} \tag{14}$$

then let $m(H, j)$ be the number of this scheme present in generation j. Growth of this number can be represented as

$$m(H, j+1) \geq m(H, j)\frac{f(H)}{\bar{f}}\left[1 - p_c\frac{\delta(H)}{l-1} - \mathbb{O}(H)p_m\right] \tag{15}$$

where $f(H)$ is the average fitness of designs representing H at generation j; $\bar{f}$ is the average fitness of the entire generation; p_c is the probability of crossover; $\delta(H)$ is the string length of scheme H; l is the string length of the crossover mask; $\mathbb{O}(H)$ is the order or number of important attributes in H, e.g., if z_j^{p+4} and z_j^{p+5} do not affect the influence of H, then $\mathbb{O}(H) = n - 2$; and p_m is the probability of mutation. The equation shows the number m grows with the improvement to the average fitness, $f(H)/\bar{f}$, and deteriorates with finite probability of the scheme string being broken by a crossover operation, and finite probability of the important attributes within the scheme mutating. A conclusion is that short, low-order, above-average schema propagate exponentially and increase in number.

The growth equation shows that the GA is inefficient and converges slowly, like a power law, yet it is the inefficiency that allows a number of diverse solvable options to result when the design space is combinatorially hard. Typically, the GA will appear to converge quickly in the early generations. This is because the underlying population is generated from a random seed, the average fitness of which is easy to improve upon. It is the very foundation of random information that allows the final designs to surpass common solutions in performance.

Validation examples were sought to test the capability of the method against published results. The first example was published by Onoda and Haftka,[5] and represents a one-dimensional high-aspect ratio, low-model dimension structure with relatively simple control and useful design metric. The second example, published by Sepulveda,[6] tests the ability to deal with three-dimensional structures with simplified output feedback control. Results for these validation cases compared well and have been published by the authors in their conference paper.[18]

Application

Both of the validation example cases shown minimized structural mass, one with an appended control effort penalty function and motion error constraint, the other with constraints on motion error and control effort. An application was sought that exercised the ability of the method on a precision structure with true motion error objective. Structural mass and control effort are to be treated as weak constraints; i.e., mass is constrained through member sizing and control effort through control weighting in the H_2 controller design. The design example is a sparse aperture, space-based telescope.

Next-generation orbiting stellar observatories require high angular resolution to meet their objectives such as extra-solar planet detection, resolution of close binaries, imaging cores of galaxies, and direct measurement of parallax of extragalactic stars. Filled aperture telescopes much larger than the Hubble Space Telescope would be prohibitively expensive because of costs incurred during launch vehicle integration and manufacturing. As a result, space-based interferometers were conceptualized that used several discrete apertures for improved resolution at lower cost. In one concept the discrete apertures are structurally connected with active collecting optics at the spacecraft hub.

A line drawing of a structure to host such an interferometer is shown in Fig. 2 and will serve as a precision structures design example. The interferometer concept shown is the result of a systems design performed at MIT for NASA's Jet Propulsion Laboratory interferometry group.[19] The structure is a long truss boom mounted on a spacecraft bus hub. Collectors are located at the tips of the boom and are relatively massive compared with the structure, $m_{\text{col}} \approx 3m_s$. The hub houses attitude control actuators and sensors as well as collecting optics, $m_{\text{hub}} \approx 50m_s$.

The essence of operation of such an interferometer is the pairwise combination of light paths, from a common wave front, incident on the separated apertures. Star light from the combined legs must be held spatially correlated over a coherent integration time to yield fringes that provide intensity and spatial phase information (relative to a guide star). The coherent integration time is set roughly by the magnitude of the target star, and the spatial correlation is usually over a few

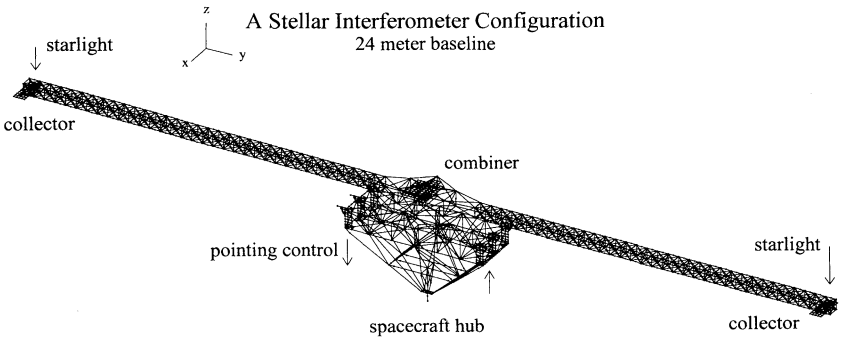

Fig. 2 Line drawing of an interferometric spacecraft.

wavelengths. Desired accuracy requires that each star light path be controlled to the same length within nanometer levels. Pathlengths are equilibrated using optical delay lines (ODL), of similar optics, that are in each light path. Under dynamic control action the ODLs induce pathlength to compensate for measured system pointing error.[20] When all of the optical loops are closed the performance is approximately the high-pass filtered, differential, tip displacement. In the design example, the performance is measured as the differential tip displacement and is weighted with a high-pass filter that represents the closed-loop optics S_{yu} of Eq. (10).

Dynamic disturbances that cause pathlength jitter may come from a variety of sources such as coarse pointing at the target sky, diurnal heating and cooling cycles, and reaction forces from optical compensation devices. For the design problem, disturbances will enter as torques about the sensitive axis acting on the hub.

These disturbances may be classified roughly as low frequency or high frequency in content. For example, cold gas attitude control thrusters generate fairly low-frequency steplike inputs, whereas reaction wheels generate high-frequency harmonics set by the wheel size and speed. Both low- and high-frequency disturbance cases are investigated as bounded disturbance weights, or filtered white noise, and are sketched in Fig. 3. The low-frequency shaping filter is a constant gain that rolls off at $1/\omega$ above a 1-Hz corner frequency. The gain is set so that the rms disturbance input is equivalent to that of a 0.1-Hz bandwidth attitude control system. The high-frequency disturbance rolls up like ω^2, corners to a constant gain at 20 Hz, and rolls off like $1/\omega^4$ beyond 50 Hz. For this filter the gain is set by expected energy from reaction wheel harmonics for a 600-rpm Hubble-sized wheel.

The structure shown in Fig. 2 was statically condensed into 12 elements a side. Fifteen different types of beam elements resulted that represented three different truss work topologies each with five different member designs. Nonstructural mass varied with the topology. The two centermost elements of the model structure were constrained to have very stiff properties. Fixed concentrated masses and inertias were added for the spacecraft hub and optical payload.

Structural actuators are included to improve pathlength compensation beyond that of the optical loops. The actuators are symmetrically placed local moment pairs that act differentially, and may be placed in elements 3–11 in the available 12 counted from the center. Note that a fixed impedance actuator was not chosen. The optimal designs will reflect the best actuator impedance.

Sensor suites used for low-frequency disturbances all use a hub-angle sensor plus a structural sensor. The hub sensor is nearly collocated with the disturbance. Structural sensor choices are differential tip motion (performance feedback), collocated differential angle, and collocated differential load. A series of runs was made for each structural sensor choice.

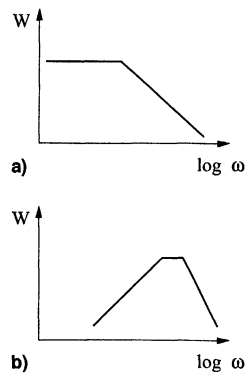

Fig. 3 Sketch of disturbance energy: a) broadband low frequency and b) narrowband high frequency.

Results from the performance and collocated load sensors will be presented. The collocated load sensor represents an interesting choice because the sensed signal depends on the stiffness of the local cross section.

In the case of high-frequency disturbances, the sensor suite is limited to be the hub-angle sensor and the collocated-load sensor. Here, the hub angle sensor is weighted so that it is barely used in the important bandwidth.

GA parameters were set for propagation of 30 designs over 80 generations. The design information strings were beam type and actuator location. Crossover occurred with probability $p_c = 0.8$, and mutation occurred with probability $p_m = 0.05$. Diversity was enforced during propagation so that no two copies of the same design proceeded. For each sensor choice and disturbance type, several runs were made using a purely random selection of initial designs. In each respective case, a final run (seeded) was made that used the best designs found in earlier runs along with random initial designs.

Low-Frequency Disturbance Cases

Results from several different runs for two different low-frequency disturbance cases are shown in Fig. 4 with the actuator location denoted by x. The nominal structure, with best actuator location, is shown for each case in the first column of the figure. The middle two columns are samples of GA runs for each case that started from a random initial population. The rightmost column is the final seeded GA run. Rows 1 and 2 are for the performance and collocated load structural sensors, respectively. Performance improvements over the nominal open loop are listed in dB in the bottom right-hand corner of each design. Control efforts are listed in the bottom right-hand corners of the nominal and seeded designs.

Several key characteristics are evident in low-frequency disturbance solutions of Fig. 4. All runs have tended to the maximum stiffness constraint. Local stiffness at each actuator location has been reduced, resulting in an active hinge. A particular characteristic of the performance structural sensor solutions is softening near the root.

Rows of GA designs show the perils of the design space. Here, on closer inspection, different designs show similar performance. Exploration of the sensitivity of the optimal designs to minor structural changes revealed that the design space was shallow and multimodal near these solutions. Sensitivity to local variations in the location of the active hinge was also found to be small.

For the performance sensor, the optimal actuator location is near the tips to collocate the actuator with the performance, yet removed from the tips to allow authority over them. The collocated load sensor apparently performs more poorly than the performance sensor in the relative rms performance numbers. This is a result of the load sensor impedance tradeoff. The optimal design is driving toward a hinge, whereas reducing the local stiffness reduces the overall sensor gain relative to its noise quality. Note that the load sensor designs in Fig. 4 are thicker at the hinge.

In general, the active hinging isolates and allows authority over the performance. Just how this is done can be seen by viewing the optimized system block transfer functions vs those of the nominal.

The nominal system transfer function matrix is shown in Fig. 5 [see Eq. (6)]. Modes that are visible in the response are antisymmetric modes of the system. Note how the plant transfer functions G_{zw} and G_{zu}, solid curves, are affected by the closed-loop optics weighting. Without the weighting, the low-frequency response in G_{zw} would roll down at $1/\omega^2$ and would dominate the motion error cost. Low-frequency disturbance weighting contributes to the rolloff seen in G_{zw} and G_{yw}. Dashed curves in G_{yw} and G_{yu} represent the hub sensor transfer functions. Control weights are plotted on the regulator transfer function, G_{zu}. They show that the control is rolled off near 100 Hz.

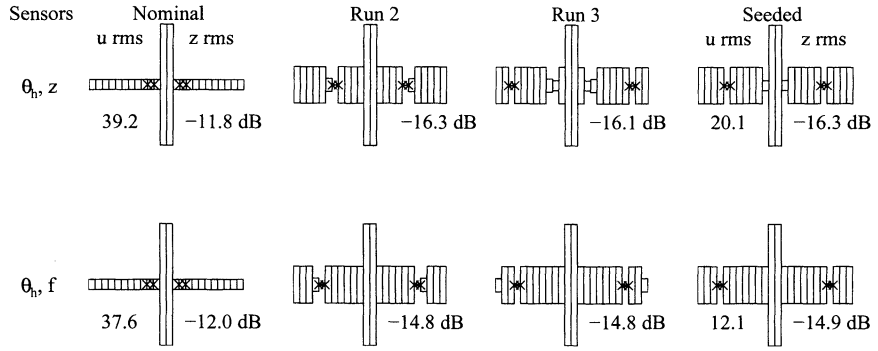

Fig. 4 Optimized structures for low-frequency disturbances; columns are differing runs; paired × mark actuator locations. Performances compared to open-loop nominal are shown in the bottom right-hand corner for each design. Root-mean-square control powers are compared with the bottom left-hand corners.

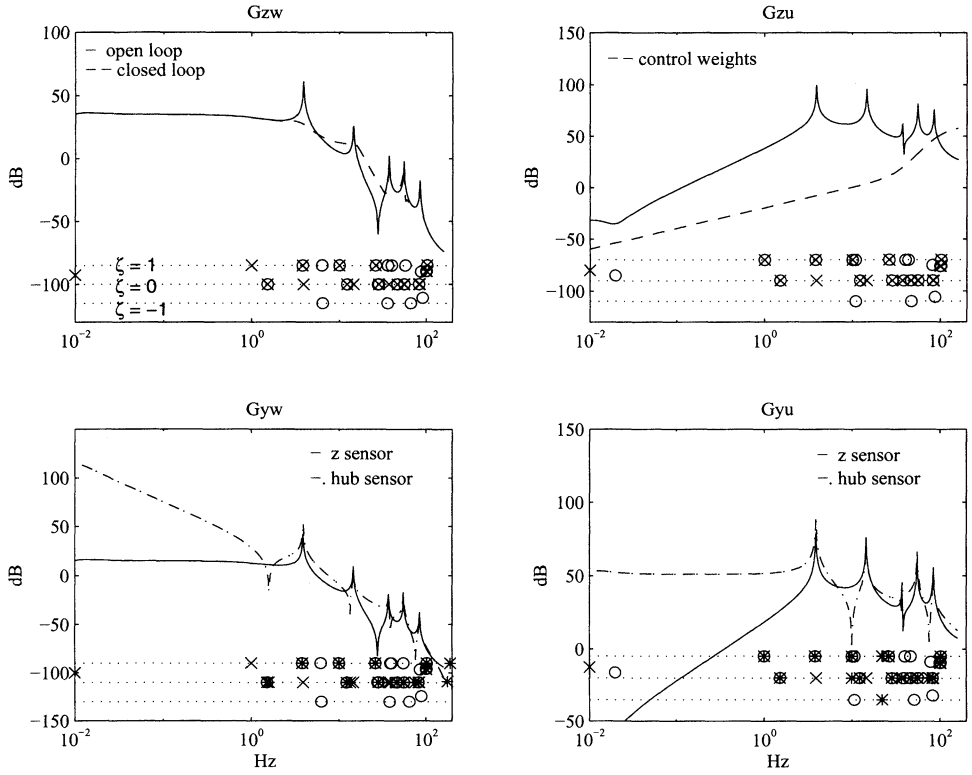

Fig. 5 Nominal structure block transfer functions for low-frequency disturbances using the performance sensor. ×, open-loop system poles; ○, zeros to the structural sensor z; and ⋆, zeroes to the hub sensor.

On each transfer function a pole-zero frequency vs damping plot, $p/z - \zeta$, is provided where the x are system poles, the o are zeros to the structural sensor, and the * are zeros to the hub sensor. On the $p/z - \zeta$ plots, the lower most dotted line, $\zeta = -1$, represents real right-half s-plane (RHP) axis, and the upper most dotted line, $\zeta = 1$, represents real left-half s-plane (LHP) axis. Poles and zeros not on the $\zeta = 1$ or $\zeta = -1$ lines are duplicate; i.e., they at least occur in conjugate pairs.

The slightly stiffened rigid-body mode is evident at 0.01 Hz. Stable weighting poles can be seen on the $\zeta = 1$ line. The disturbance weighting pole is at 1 Hz. Performance weighting poles appear to be canceled by zeros. These poles are duplicated because the performance is used as a sensor. Symmetric mode pole-zero cancellations are depicted on the $\zeta = 0$ line with increasing high-frequency density.

As is characteristic in beam problems, real minimum–nonminimum phase zero pairs occur in transfer functions with noncollocated inputs and outputs. The pairs become less damped and gradually offset in frequency as their frequency increases. This later effect is a result of modeling discretization. The nonminimum phase zero of these pairs is known to limit controller performance.[12] A limit on performance, more particular to structural systems, is that the pair of zeros raise the average transfer function gain for no phase increase. This impedes compensator rolloff because the plant continues to lose phase with the occurrence of modes. In a sense, structural delay from noncollocated inputs to outputs is manifested in the occurrence of modes.

Closed-loop G_{zw} is plotted using a dashed line in Fig. 5. The first antisymmetric mode is damped heavily and the second is only slightly damped. Modes from the second one onward contribute little to the total closed-loop performance. The control design can be thought of as separate regulator and estimator stages with the combined compensator acting on the plant as seen through the actuators and sensors, G_{yu}. In the nominal design, the regulator is clearly limited by the nonminimum phase zero at 10 Hz. The estimator primarily uses the hub sensor that is nearly collocated with the disturbance. There are no * (zeros with respect to the hub sensor) below the $\zeta = 0$ line in the bandwidth shown. The performance sensor is mixed in near the zeros of the hub sensor transfer function. Even though u is near the root, it is not collocated with the hub

sensor beyond the second antisymmetric mode, as shown by the zero at 25 Hz in the G_{yu} loop. Because the control is focused on the first antisymmetric mode, it is really the regulator that limits performance when combined with G_{yu} loop nonminimum phase behavior.

Block transfer functions for the optimized structure shown in the right-most column of the first row in Fig. 4 are depicted in Fig. 6. Comparing the G_{zw} transfer functions of the nominal and the optimized structures shows softening of the first and second modes and stiffening of the higher modes. The hinging has improved the average gain in the G_{zu} transfer function through the softening. Locating the actuator near the tip yields good pole-zero structure out to near 100 Hz (alternating poles and zeros in G_{zu} until the real zero pair near 100 Hz). Moving the hinged actuator toward the tip does not compromise estimator phase, because again, the hub sensor is primarily used. Collocation to the secondary performance sensor is worse. The important nonminimum phase behavior for the optimized system is now in the G_{yu} transfer function. The improved authority over the first two modes can be seen by comparing the closed-loop disturbance to performance plots.

The optimized design shows mass-dominated transfer functions, i.e., zeros occurring close to the left of the poles in G_{zu} and G_{yu}, suggesting that the trend is to structurally isolate the tip payloads. This isolation-type solution is further aided by the high-pass nature of the performance weights and the low-pass nature of the disturbance weights.

Results from optimizing the system are limited. Improvement over the nominal system is just 4 dB in the closed loop. The reason for this is that performance is dominated by the system rigid-body mass, as can be seen in the low-frequency region of the G_{zw} transfer functions. This is predominantly the hub inertia and collector mass times the baseline. Optimizing the structural system can only have limited effect on this response. However, it is possible to demand higher gain from the closed-loop optical systems at these low frequencies (if the guide star is bright enough) that further attenuates the response. The effect would further emphasize the GA design results.

High-Frequency Disturbance Case

Under high-frequency disturbances, the search finds designs that are structurally discontinuous (Fig. 7). In this case, the collocated load sensor is predominantly used. Emphasis on the higher modes has resulted in a complicated structural filter producing a design with actuator/sensor pair at the midspan, reduced cross-sectional inboard, and stiffened cross-sectional outboard. The best nominal actuator location is also at the midspan; see the left-most design in Fig. 7. Again, the cross section at the actuator/sensor location is soft, but not so thin as to overly reduce the sensor transfer function gains.

The system block transfer functions for the best design under high-frequency disturbances are shown in Fig. 8. Disturbance emphasis, seen in the average gain of G_{zw}, has resulted in tailoring of the dynamics to achieve improved regulator collocation; see the pole zero spacing in G_{zu}. There is good authority over the emphasized region, shown in the gain and the pole zero spacing in G_{yu}. Poor estimator collocation, real RHP zero pairs in the load G_{yw}, has resulted from the very flexible

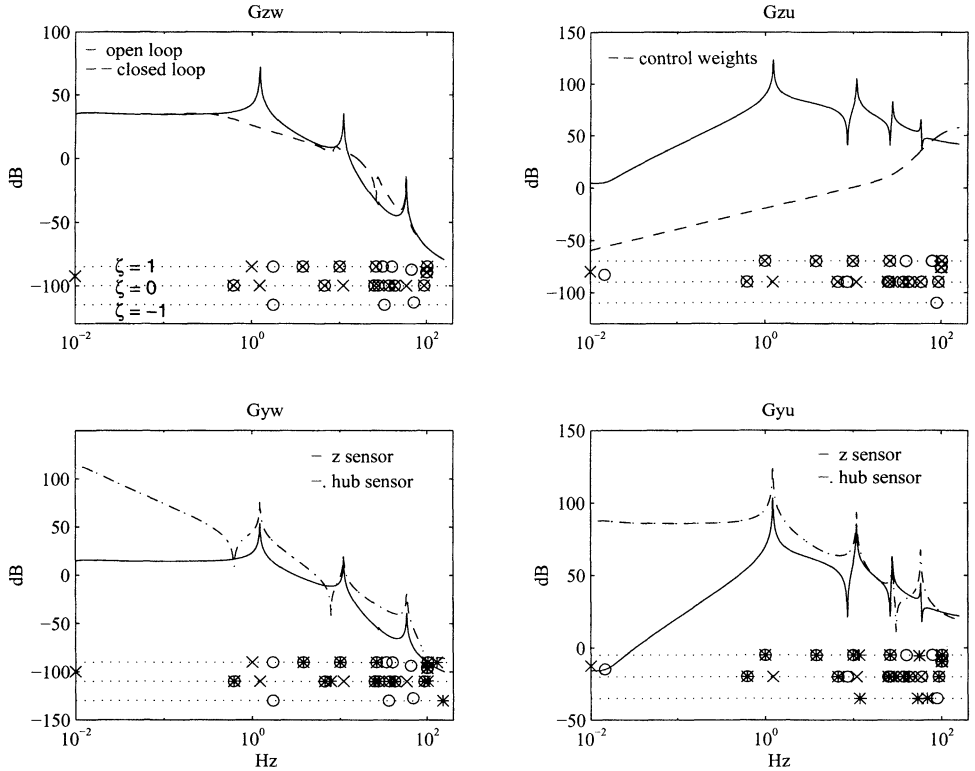

Fig. 6 Optimized structure block transfer functions for low-frequency disturbances using the performance sensor.

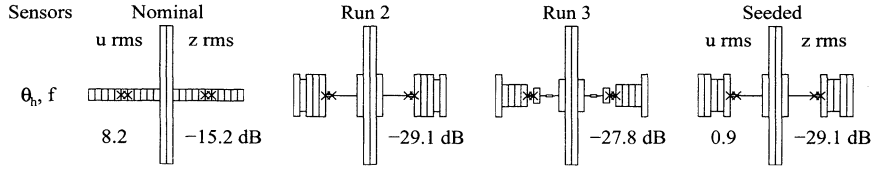

Fig. 7 Optimized structures for high-frequency disturbances using the structural load sensor.

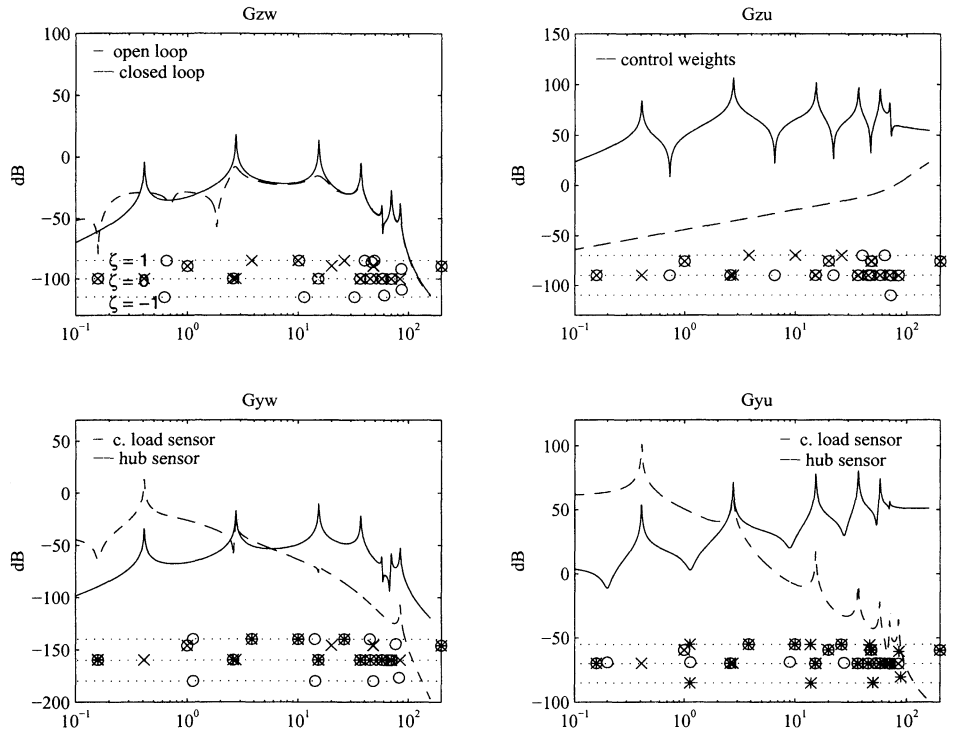

Fig. 8 Optimization structure block transfer functions for high-frequency disturbances using the collocated load structural sensor.

inboard portion of the structure and limits the control to damping of the first, second, and third modes.

Major improvements over the closed-loop nominal system are realized because the overall average transfer function gain has been reduced by 10–15 dB. The control merely damps the system to this average gain. At high frequencies the average gain is dominated by the flexible properties of the system. Evidently, structural system optimization has greater impact here, and the assumed optical control has less.

Conclusions

This paper provides a method by which to design precision structures for control. The method is somewhat inefficient in that it uses first-order stochastic propagation of a parallel model space (GA); however, the method is not specialized to any one modeling or control technique and allows the simultaneous update of topological variables. In the proposed method a consistent effort is made via condensation, model reduction, and static mode correction to provide a good representative model for control. These discrete operations are complemented by utilizing an outer-loop GA that searches over structural and topological variables. Successive refinements and variations of the outer-loop variables can be pursued. The authors have exercised the method on complicated three-dimensional truss problems with variations in member connectivities, nodal locations, and member geometries. Along with discrete variations in structural control variables, the method allows the freedom to use differing inner-loop control optimization techniques (this has also been explored by the authors).

Results are presented in the frequency domain for both the open and closed loop. Appropriate transfer functions (with appended weights and filters) allow easy inspection of model fidelity and control action. Pole-zero-ζ plots enable the observation of systematic performance limits. The major limit to the closed-loop control is noncollocation in either regulator, estimator, actuator-to-sensor transfer functions. Noncollocation, or delay, in structural systems manifests itself as the average transfer function gain staying high while the phase decreases as a result of the occurrence of system modes.

A one-dimensional interferometer design is used as an application of the method. The results for this example show that, for low-frequency disturbances that roll off and performance weights that act as high pass, the best controls/structures design globally stiffens and locally isolates the performance. Structural isolation achieves good regulation gain and modal residues necessary for effective control. The price is delays from the actuators to structural sensors. For structural moment actuators, low substrate stiffness results, yielding good authority over the performance and high overall gain in the structural sensor transfer functions.

Under high-frequency disturbances that are bandpass and performance weights that act as high pass, the best controls/structures interferometer designs involve complicated structural filtering. In this case, several modes have been softened to lower frequencies to take advantage of the bandpass nature of the disturbances. The actuator/sensor design is well coupled to the softened modes. This enables improved closed-loop performance through reduced average gain in the critical bandwidth.

References

[1]Rao, S., Venayya, V., and Khot, N., "Game Theory Approach for the Integrated Design of Structures and Controls," *AIAA Journal*, Vol. 26, No. 4, 1988, pp. 463–469.

[2]Maghami, P., Gupta, S., Elliot, K., and Joshi, S., "Experimental Validation of an Integrated Controls-Structures Design Methodology for a Class of Flexible Space Structures," NASA TP 3462.

[3]Canfield, R., Grandhi, R., and Venkayya, V., "Optimum Design of Structures with Multiple Constraints," *AIAA Journal*, Vol. 26, No. 1, 1988, pp. 78–85.

[4]Miller, D., and Shim, J., "Combined Structural and Control Optimization for Flexible Systems Using Gradient Based Searches," AIAA Paper 86-0178, Jan. 1986.

[5]Onoda, J., and Haftka, R. T., "An Approach to Structure/Control Simultaneous Optimization for Large Flexible Spacecraft," *AIAA Journal*, Vol. 25, No. 8, 1987.

[6]Sepulveda, A. R., Jin, I. M., and Schmit, L. A., "Optimal Placement of Active Elements in Control Augmented Structural Synthesis," *AIAA Journal*, Vol. 31, No.10, 1993.

[7]Keane, A. J., "Passive Vibration Control Via Unusual Geometries: The Application of Genetic Algorithm Optimization to Structural De-

sign," *Journal of Sound and Vibration*, Vol. 185, No. 3, 1995, pp. 441–453.

[8]Keane, A. J., "Passive Vibration Control Via Unusual Geometries: Experiments on Model Aerospace Structures," *Journal of Sound and Vibration*, Vol. 192, No. 1, 1996.

[9]Masters, B. P., and Crawley, E. F., "An Experimental Investigation of Optimized Precision Optical Controlled-Structures," *Proceedings of the SPIE Society for Photonics and Optical Instrumentation Engineering* (San Diego, CA), 1997.

[10]Craig, R. R., and Bampton, M. C., "Coupling of Substructures for Dynamic Analyses," *AIAA Journal*, Vol. 6, No. 7, 1968.

[11]Gregory, C. Z., "Reduction of Large Flexible Spacecraft Models Using Internal Balancing Theory," *Journal of Guidance, Control, and Dynamics*, Vol. 7, No. 6, 1984.

[12]Freudenberg, J. S., and Looze, D. P., "Right Half Plane Poles and Zeros, and Design Tradeoffs in Feedback Systems," *IEEE Transactions on Automatic Control*, Vol. 30, 1985.

[13]Kwakernaak, R., and Sivan, R., *Linear Optimal Control Systems*, Wiley, New York, 1972.

[14]Goldberg, D. E., *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison–Wesley, Reading, MA, 1989.

[15]Ackley, D. H., "An Empirical Study of Bit Vector Function Optimization," *Genetic Algorithms and Simulated Annealing*, edited by L. Davis, Morgan Kaufmann, Los Altos, CA, 1987, pp. 170–204.

[16]Baker, J. E., "Adaptive Selection Methods for Genetic Algorithms," *Proceedings of the International Conference on Genetic Algorithms and Their Applications*, edited by J. J. Grefenstette, Pittsburgh, PA, 1985, pp. 101–111.

[17]Keane, A. J., "Experiences with Optimizers in Structural Design," *Adaptive Computing in Engineering Design and Control*, Sept. 1994.

[18]Masters, B. P., and Crawley, E. F., "Evolutionary Design of Controlled Structures Systems," AIAA Paper 97-1263, April 1997.

[19]Shao, M., Colavita, M. M., Hines, B., and Staelin, D., "The Mark III Stellar Interferometer," *Journal of Astronomy and Astrophysics*, Vol. 193, 1988, pp. 357–371.

[20]O'Neal, M., and Spanos, J., "Optical Pathlength Control in the Nanometer Regime on the JPL Phase B Interferometer Testbed," *Proceedings of the SPIE Conference on Active and Adaptive Optical Systems* (San Diego, CA), 1991.