JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 19, No. 4, July–August 1996

Derivation of Zero-Gravity Structural Control Models from Analysis and Ground Experimentation

Roger M. Glaese* and David W. Miller†
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

The Middeck Active Control Experiment (MACE) is a Space Shuttle flight experiment that flew on STS-67 in March 1995, to demonstrate high-authority active structural control in microgravity conditions. Because no experimental on-orbit data were available before flight, the model used to derive these controllers was obtained from a combination of analysis and open- and closed-loop ground testing. A modeling approach to obtain this 0-*g* model is presented along with its application to the MACE test article. This approach starts with a 1-*g* model, which includes gravity and suspension effects. This 1-*g* model is improved using ground test results in a process called updating. Once updated, it is verified through a combination of structural perturbation and closed-loop control. Finally, the gravity and suspension effects are removed to arrive at a prediction of the 0-*g* behavior of the structure. Flight data are used to demonstrate the effectiveness of this approach.

Introduction

MANY future spacecraft will require active structural control to meet stringent performance specifications. It is highly desirable that the controllers for these vehicles be guaranteed, prior to flight, to be stable and perform well on orbit. However, prior to launch, these controllers can be analyzed and tested only on the ground. Unfortunately, these ground tests are affected by gravity and the suspension system, which couple with the structural dynamics. Thus, it is very difficult to use ground tests to validate the controller that is selected for on-orbit operations. These are the issues that are addressed with the Middeck Active Control Experiment (MACE).

The modeling goal of MACE is the prediction, with confidence, of the on-orbit behavior of the structure. Prior to launch, there are no on-orbit experimental data that can be used to derive measurement-based models[1] and resulting controllers. There is an alternative approach, the finite element analytical modeling technique. However, the accuracy of the 0-*g* finite element model (FEM) is unknown prior to flight. Our approach is to add gravity and suspension effects to the 0-*g* model to form a 1-*g* FEM.[2] This 1-*g* model is compared with ground test data to determine the accuracy of the FEM. Physical parameters of this model, such as masses and material properties, are updated to match 1-*g* experimental data.[3,4] This 1-*g* model, as well as the hardware, is perturbed to determine whether modeling accuracy is preserved, thereby indicating whether the updates were made appropriately. Once this 1-*g* model has sufficient accuracy, the direct gravity and suspension effects are removed from the model. The result is a 0-*g* model that represents the best prediction of the on-orbit behavior. This paper traces this development for the MACE test article.

Although the MACE test article is a relatively simple structure, this modeling approach is applicable to much more complicated structures that can be tested on the ground. If, on the other hand, the entire structure cannot be tested as a whole, the approach is equally well suited for the testing of components of the structure. In this case, however, extreme care must be exercised when dealing with the boundary conditions of the components, both in the model and the experiment. Using this modeling approach, high-performance controllers can be designed and evaluated prior to flight. Even with this detailed modeling approach, it is still possible that, because of remaining modeling inaccuracies and nonlinearities, the desired on-orbit performance levels are not met by preprogrammed controllers. For this reason, it is desirable that provisions for on-orbit system identification and subsequent control redesign are retained in the spacecraft. This capability was a fundamental aspect of the MACE flight operations, a discussion of which is beyond the scope of this paper but can be found in Campbell et al.[5]

Received July 31, 1995; revision received April 9, 1996; accepted for publication April 10, 1996.

*Graduate Research Assistant, Space Engineering Research Center. Student Member AIAA.

†Principal Research Scientist, Space Engineering Research Center. Member AIAA.

General Modeling Approach

The modeling objective of MACE is to develop an I/O model with sufficient accuracy to design controllers for implementation on the test article in 0 *g*. The MACE approach consists of both open- and closed-loop testing on the ground and in space, using controllers formulated from both finite element[6] and measurement models.[7] The modeling effort described here emphasizes the FEM because an accurate 0-*g* version is required for control design prior to flight.

The first step in the MACE approach is to obtain an initial 1-*g* model, which includes the suspension and gravity effects. The FEM is used to generate 1-*g* frequencies and mode shapes that are combined with models of the sensors and actuators, including electronic gains and models of all signal conditioning and servo controls. A block diagram showing this is given in Fig. 1. Through previous experience with MODE[8] and the MIT Interferometer test bed,[9] the damping is set to 1% for structures with articulation and significant wiring (MACE); otherwise, it is 0.1% for tight trusses. The result is a state–space model of the I/O behavior of the structure as viewed from the compensator.

Next, experimental data are taken in the form of I/O transfer functions and compared to those predicted by the model to illustrate errors and suggest improvements. This process is called updating.

Updating can be of two distinct varieties. First, the experimental frequency responses can be used directly to update the model.[3] Second, estimates of modal frequencies and damping ratios, identified from experimental data, also can be used to update the FEM.[4] Although not used concurrently, these two methods can be used successively to improve the FEM. The net result of the update process is a model that is more closely correlated to the experimental 1-*g* data.

Note that updating must be performed using physical parameters of the FEM (e.g., stiffnesses, masses, lengths). This is because the gravity and suspension effects are buried in a nontrivial way in the natural frequencies and mode shapes of the measurement model, making them impossible to remove.

The FEM also can be used to improve the measurement model. The FEM indicates the minimum order of the model, because measurement models are prone to include multiple modes to represent a single mode, especially in a multi-input multi-output (MIMO) system. The FEM also reveals whether lightly damped zeros are minimum phase. This is a problem because these zeros have low magnitude and are often below the noise floor of the sensor, making

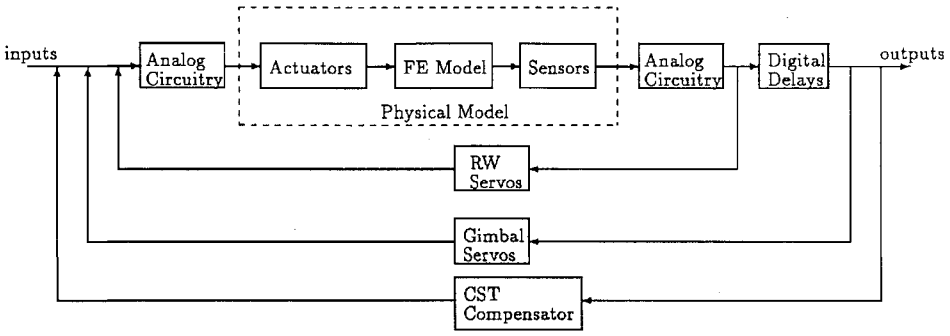

Fig. 1 I/O model.

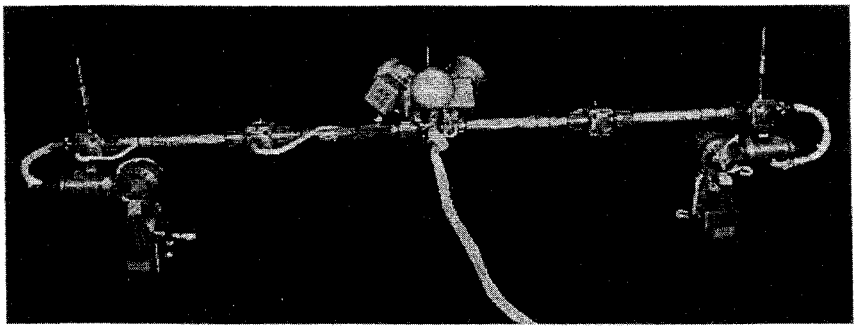

Fig. 2 MACE test article suspended in 1 *g*.

the minimum phase determination from the experimental frequency responses difficult.

Once the model has been updated, it must be verified. This is important because if the wrong parameters in the 1-*g* FEM have been adjusted to improve the match with the data, there is no reason to believe that removing the suspension and gravity effects from the 1-*g* model will result in a good prediction of the 0-*g* behavior. Traditionally, this is performed by testing a system with altered mass, stiffness, or boundary conditions, altering the model in a similar fashion, and comparing the results. A good model will capture the perturbed behavior with similar accuracy.

A controlled structure allows perturbations of a different sort, namely changing the structure's behavior through closed-loop control. Because the goal of MACE is to design controllers for implementation in 0 *g*, this closed-loop verification process forms an essential part of the approach.

Because a measurement model is a nearly perfect representation of the data, controllers designed using this model, in the absence of nonlinearity, should give the best closed-loop performance. Thus, the results of measurement-based control provide a yardstick against which the finite element-based control results can be judged. If the finite element performance is very close to that of the measurement model, this indicates that the FEM is good enough and no further updating is necessary.

To form the 0-*g* model, the suspension system and gravity effects are removed from the FEM. The resulting 0-*g* frequencies and mode shapes are combined with models of the sensors and actuators, signal conditioning, and servo controllers to produce a nominal I/O model. Because the model is to be used for robust control design, some knowledge of what errors are likely to be present in 0 *g* is also desired. This is a so-called uncertainty model[10] and also must be determined from a combination of ground testing and analytical modeling. The rest of this paper details how this modeling approach was applied to MACE.

Hardware Description and Control Problem

The MACE program used three generations of hardware: the development model (DM), the engineering model (EM), and the flight model (FM). The purpose of the DM was to evaluate preliminary designs of the major components and mature the control design process. It consisted of a structural bus, a reaction-wheel assembly, and a single two-axis gimbal. The purpose of the EM, which added a second two-axis gimbal, was to evaluate the final designs of the major components and was nearly identical to the FM, which flew on STS-67 in March 1995.

The FM test article (Fig. 2), consists of a segmented straight tubular polycarbonate structural bus with a two-axis gimbal-supported pointing payload at the left end, called the primary gimbal. A second two-axis gimbal, called the secondary gimbal, which disturbs the system by undergoing various scanning profiles, is mounted on the right end. Each two-axis gimbal actuates rotation about two perpendicular axes via two dc torque motors. A reaction-wheel assembly is attached to the center of the bus to provide attitude-control torques and vibration suppression. The reaction-wheel assembly comprises three orthogonally mounted dc servo motors with an inertia wheel mounted on each rotor. One segment of the structural bus consists of an active member, which allows the actuation of bending strain through the use of piezoelectric materials. The structure is suspended in 1 *g* using three pneumatic–electric low-frequency suspension devices.

The test article is outfitted with a variety of sensors to measure the structural motion. Each gimbal axis has a laser rotary encoder to measure the angle of the gimbal relative to the structural bus. A two-axis rate-gyro platform is mounted in the payload attached to the primary gimbal, measuring inertial angular rates. Three rate gyros are mounted on the reaction-wheel assembly with their sensitive axes aligned with the reaction-wheel axes. Eight strain gauge pairs (two per strut) are used to provide a measurement of the vertical and horizontal bending strains.

All signal conditioning for the sensors and power conditioning for the actuators is performed in the experiment support module (ESM). The signal conditioning consists of antialiasing Bessel filters and notch filters to attenuate internal rate-gyro dynamics. The ESM also contains a real-time control computer, operating at a 500-Hz sampling rate, on which all of the structural-control algorithms are implemented.

The goal of closed-loop control for MACE is to achieve a 40-dB improvement in the pointing accuracy of the primary gimbal in the presence of disturbances generated by the scanning motion of the secondary gimbal. The control scheme relies on a layered approach. The first layer consists of local single input/single output servo loops closed around all axes of the gimbals using the laser rotary encoders for feedback. These servo loops have bandwidths of approximately 3 Hz and achieve performance levels of nearly 12 dB over purely

open-loop behavior. The second layer, which is the primary focus of the MACE control designs, consists of MIMO compensators that use the primary gimbal and reaction-wheel assembly inputs as actuators and the primary-payload and reaction-wheel rate gyros as sensors. Because of phase rolloff considerations from time delays and sensor dynamics, the bandwidth of these MIMO controllers is limited to approximately 60 Hz. Note that the actuators and the sensors can be separated into nearly collocated pairs. The control designs, however, do not rely solely on this collocation. Instead, the control designs make significant use of the noncollocated feedback channels. Additionally, linear quadratic Gaussian control designs make use of the disturbance-to-performance and disturbance-to-sensor transfer functions. For these reasons, it is seen that the control designs place great emphasis on accurately modeling the distributed nature of the entire structure over a fairly large frequency range.

High-Fidelity 1-*g*, I/O Model

Most of the initial parameters of the FM had been updated using DM and EM data and were already fairly accurate.[6,11] Because each generation involved a change in hardware, these different generations represent a step in the process of model verification.

The first piece of the I/O model is the FEM, which includes models of the bare structure, suspension system, and gravity effects. The suspension system, consisting of the pneumatic–electric devices and the load rods, provides a bounce frequency of 0.2 Hz. These devices are modeled as springs and dashpots constrained to move only in the vertical direction. The 5-m-long graphite–epoxy load rods result in pendular frequencies of 0.23 Hz. These rods are modeled as Bernoulli–Euler beams with the proper cross-sectional and material properties.

Gravity effects, such as sag and predeformation, preloading of the members, pendular effects on articulating appendages, stiffening of rigid-body modes, and suspension violin behavior are captured through a nonlinear, static, stress-stiffening procedure in the finite element code.[2] Gravity is mathematically applied to the structure in an iterative procedure, and the resulting stresses are used to modify the stiffness matrix.

The next piece is the damping model, a combination of modal and concentrated damping. This model was selected because certain, potentially significant, sources of damping are only present in the suspended structure, namely the suspension devices. Because the model of each device contains a dashpot, the resulting damping can be included in the 1-*g* model but excluded from the 0-*g* model.

At this point, a modal state–space model of the system can be constructed, as in

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= Cx \end{aligned} \tag{1}$$

where $x = [\eta \ \dot{\eta}]^T$, η are the modal displacements, u are the actuator and disturbance inputs, and y are the sensor outputs and performance variables. The system matrix A contains only the frequency and damping information, whereas the input matrix B and the output matrix C contain the mode-shape eigenstructure (actuator–sensor/disturbance–performance) relationships. For the MACE model, A, B, and C take on the following forms:

$$A = \begin{bmatrix} 0 & I \\ -\Omega^2 & -2\zeta\Omega - \Phi^T \tilde{C}_s \Phi \end{bmatrix} \tag{2}$$

$$B = \begin{bmatrix} 0 \\ \Phi^T b \end{bmatrix} \tag{3}$$

$$C = \begin{bmatrix} c_d \Phi & 0 \\ 0 & c_r \Phi \end{bmatrix} \tag{4}$$

where Ω^2 is a diagonal matrix containing the square of the modal frequencies ω_i^2 on the diagonal. The matrix $\zeta\Omega$ contains the product of the damping ratio and modal frequency, $\zeta_i\omega_i$, on the diagonal. $\tilde{C}_s$ is the concentrated damping matrix attributable to the suspension devices, and Φ is the matrix of mode shapes. The matrix b and matrices c_d and c_r are pointing matrices that describe the behavior of the actuators, displacement, and rate sensors, respectively. This is called a physical model because the inputs and outputs are in their natural physical units (e.g., Nm and rad/s).

This physical model is not very useful for control design because the real-time control computer measures and commands volts through the A/D and D/A, respectively. To convert the model to volts, the physical input and output matrices are scaled by the electronic gains for each actuator and sensor. Then, models of the rate-gyro dynamics, Bessel and notch filters for the sensors, and smoothing filters for the actuators are included. Analog servo controllers, such as the reaction-wheel speed control, also must be included. These servos (Fig. 1) feed the tachometer signals back through a gain and a low-pass filter to keep the wheels spinning at a constant reference speed. To generate attitude-control torques, the reaction-wheel speeds are varied.

The last piece of the I/O model is the real-time computer. This model consists of the software gains, time delays introduced by sampling, and any servo controllers implemented using the digital computer. The time delays introduced by sampling, which have been measured at 3.6 ms, are included in the model using a Pade approximation on the outputs. The gimbal servo loops (Fig. 1) feed back the encoder signals, which are digitally integrated, yielding the angle and angular rates of the gimbals relative to the bus, which are then fed back to point the gimbals at the desired relative angles. Time delays are included in these loops because they are closed digitally.

With the model of the real-time computer and servos, the 1-*g* I/O model is complete. The dynamic characteristics of this model can be compared with experimental data to improve the correlation of the model through a process called updating.

Model Updating

Model updating, because it is an optimization process, requires two things: a cost for minimizing and update parameters. Several candidate costs have been considered: minimizing the errors between the FEM and identified frequencies and damping ratios; minimizing the difference between the magnitudes of modeled and measured transfer functions; and minimizing the difference between the complex logarithms of these transfer functions. Each cost has its advantages and disadvantages. Minimizing the frequency and damping error is computationally simple, but requires them to be identified using a measurement model. Minimizing the difference between transfer-function magnitudes does not require a measurement model, but it is computationally intensive, requiring computation of the model's transfer functions each time a parameter is varied. Another disadvantage, which is addressed in the third cost, is the fact that comparing the magnitudes places emphasis on regions of large magnitude—the poles—and neglects the zeros, which are often very important in control design. In the logarithmic cost, the poles and zeros are weighted equally. An additional advantage of the complex logarithm is that it weights the phase of the frequency response, which is also important for control design. The computational expense of the logarithmic cost is only slightly more than the magnitude cost. Weighing the advantages and disadvantages of each cost, the logarithmic cost, given in Eq. (5), was chosen for MACE updating:

$$J = \sum_{\omega} \sum_{\text{actuators}} \sum_{\text{sensors}} \left\| \log\left(\frac{G_{\text{data}}(j\omega)}{G_{\text{model}}(j\omega)} \right) \right\|^2 \tag{5}$$

This cost can be written in vector form, such that $J = j^T j$, with

$$j = \begin{bmatrix} J_1 \\ J_2 \\ \vdots \\ J_n \end{bmatrix} \tag{6}$$

where J_i is the cost evaluated at each frequency point for a given I/O pair. This vector cost can be expanded to include multiple actuators and sensors by stacking the vectors. The frequency grid for this cost is determined by the frequency grid for the experimental data, typically linearly spaced between a lower and an upper frequency

limit. For MACE, the lower limit was 0.1 Hz, which is the frequency resolution of the identification experiments. The upper frequency limit for MACE was chosen to be 60 Hz, because the majority of the structural modes lie below this frequency and this is above the bandwidth for both 1- and 0-g controllers.

With the cost defined, update parameters are required. Every parameter that makes up the model is uncertain to some degree. Using every parameter in the update is computationally prohibitive. To reduce (or eliminate) the uncertainty in as many parameters as possible, extensive component testing has been performed. The predominant uncertainties that remain are in the component stiffnesses. These result from statistical variation in material properties, complex cross-sectional geometries, the details of the gravity loading in the joints, and amplitude nonlinearities. Because of these uncertainties, the Young's moduli (11 in all) used in the FEM were chosen as the update parameters for errors in frequency and mode shape. Because structural damping is not physically modeled, the unservoed structural damping ratios also were used as update parameters.

Finally, the sensitivity of the cost to these parameters must be addressed. Because these are parameters of the 1-g FEM, their effect on the cost is buried beneath the iterative procedure used to add the suspension and gravity effects. This makes the analytical computation of the sensitivities very difficult, if not impossible. For this reason, a three-step finite difference technique is used to compute the sensitivities. In the first step, an update parameter is perturbed in the FEM, producing a set of perturbed frequencies and mode shapes. Second, a new perturbed I/O model is formed to compute a perturbed-frequency response in the third step, and consequently a perturbed cost. The sensitivity is then found from

$$\frac{\partial j}{\partial \alpha_i} \approx \frac{\Delta j}{\Delta \alpha_i} = \frac{(j_i - j_0)\alpha_i}{\Delta \alpha_i} \tag{7}$$

with j_i the perturbed cost, j_0 the nominal cost, and $\Delta\alpha_i$ the change in the ith parameter of the FEM.

Once these sensitivities have been computed, they can be scrutinized to determine whether some of the selected parameters can be excluded from the update or whether, perhaps, several parameters should be combined into a superparameter. Two tests can be performed to examine the alignment of these sensitivities with 1) the initial cost vector j_0 and 2) each other. The first test shows whether the initial cost is sufficiently sensitive to the parameter. The second test shows the orthogonality between the parameters. If several parameters have sensitivities in the same direction, indicating a reduction in rank of the sensitivity matrix, they can be combined. Otherwise, the optimization will fail, because there is not a unique set of parameters that minimizes the cost.

These tests begin by forming the matrix of sensitivities ΔJ as

$$\Delta J = \left[\frac{\Delta j}{\Delta \alpha_1} \cdots \frac{\Delta j}{\Delta \alpha_N} \right] \tag{8}$$

and the matrix of normalized sensitivities $\Delta \bar{J}$ as

$$\Delta \bar{J} = \left[\frac{\Delta j}{\Delta \alpha_1} \Big/ \left\| \frac{\Delta j}{\Delta \alpha_1} \right\| \cdots \frac{\Delta j}{\Delta \alpha_N} \Big/ \left\| \frac{\Delta j}{\Delta \alpha_N} \right\| \right] \tag{9}$$

The alignment of the sensitivities with the initial cost (test 1) can be found by taking the inner product of the initial cost with the matrix of sensitivities $|j_0^T \Delta J|$. The absolute value is used because only the magnitude is important (a negative sign only indicates the sensitivity direction). If a parameter is orthogonal to the initial cost, the inner product will be small and the parameter should be discarded. Alternatively, if the parameter is aligned with the initial cost but only has a very weak influence, the inner product again will be small and the parameter should be discarded. Thus, large values of the inner product indicate parameters that should be retained.

In test 2, the inner product of the normalized sensitivity of one parameter is taken with that of another. This can be written for all parameters (i.e., a series of inner products) as the matrix product $|\Delta \bar{J}^T \Delta \bar{J}|$. Again, only the magnitude is important. The diagonal of this matrix will be unity, because the sensitivity vectors have been normalized, and the magnitude of the off-diagonals will lie between zero and unity. Off-diagonal values near zero indicate that

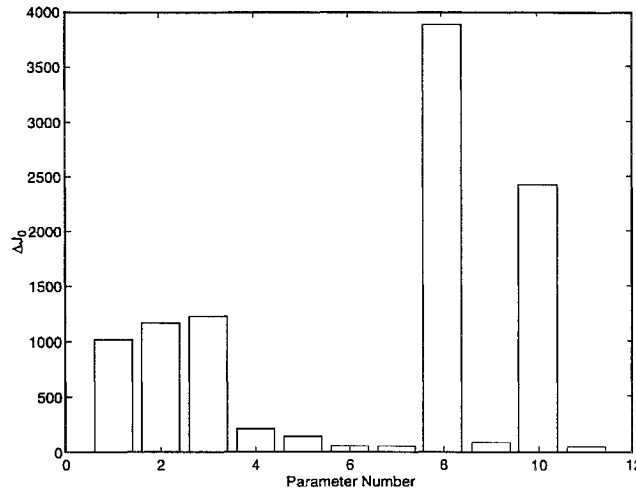

Fig. 3 Alignment with initial cost.

	1	2	3	4	5	6	7	8	9	10	11
1	1.0	.8	.3	.4	.3	.3	.2	.0	.3	.1	.0
2		1.0	.3	.3	.3	.2	.1	.0	.2	.0	.1
3			1.0	.4	.8	.4	.4	.2	.6	.1	.5
4				1.0	.4	.4	.1	.1	.1	.0	.0
5					1.0	.7	.5	.0	.8	.0	.5
6						1.0	.3	.1	.5	.1	.5
7							1.0	.5	.3	.0	.8
8			Symmetric					1.0	.1	.6	.0
9									1.0	.1	.3
10										1.0	.1
11											1.0

Fig. 4 Orthogonality test.

the sensitivities of the two parameters are nearly orthogonal. The ideal situation would be that this matrix is the identity matrix, indicating that the sensitivity matrix is full rank and efficiently spans the parameter space. If any of the off-diagonal values are near unity, indicating that the corresponding parameters have nearly the same sensitivity vector, the possibility for an improper update is greatly increased because the optimization is unsure which parameter to change to improve model correlation. In this case, two courses of action are possible: choose the parameter that has the most influence on the initial cost; or combine the two parameters into a single superparameter.

Figures 3 and 4 graphically depict these two tests for the FM. From Fig. 3, only 5 of the 11 Young's moduli are significantly aligned with the initial cost. These parameters correspond to the stiffnesses that contribute the most to the flexibility of the suspended structure. From Fig. 4, the orthogonality matrix is predominantly diagonal, with a few exceptions. Several of these exceptions correspond to the same material, but with differing cross-sectional properties. These pairs might be suitable for the creation of superparameters. The rest of the exceptions do not correspond to like materials. Fortunately, in these cases, at least one of the two correlated parameters does not affect the initial cost very much and would be discarded from test 1.

Once the update parameters have been selected, a least-squares optimization is performed to minimize the cost. Several iterations are usually required to recalculate the sensitivities because their range of validity is often quite small. Updating is halted when the change in cost is below a desired threshold.

In considering the effectiveness of updating in the three generations of MACE hardware, several observations can be made. In the DM generation of the model, updating based on engineering insight was the most effective method, because the errors were very large and the experience with the model was limited. Engineering-insight updating is where the modeler alters boundary conditions, element types, material properties, and so forth in an ad hoc fashion based on discrepancies between the model and the data. In the EM, however, the improved correlation of the model with the data allowed updating based on optimization of the correlation with experimental data to be used very effectively, achieving a 28% reduction in

the cost.[6] For the FM, optimization could once again be used effectively. However, the large improvements made in the EM could not be repeated because of the high accuracy of the initial FM.

Model Verification

In model verification, the structure is changed in some way and the model is examined to see whether it can predict the perturbed behavior. The nominal model, although it might correspond well with a certain set of experimental data, might do so with incorrect parameter values. Therefore, this model verification step is very important because the final step in the MACE approach represents a substantial structural perturbation, namely the removal of gravity and the suspension.

As noted previously, the evolution of MACE hardware lends itself to the verification process extremely well. Each change of hardware allows the model to be verified. Also, within each generation of hardware, several opportunities to verify the model exist. The perturbations that can be made to the hardware can be grouped according to whether the perturbations are in the mass, stiffness, or suspension of the structure. This spans the types of perturbations caused by removing gravity and were therefore deemed appropriate and sufficient. Other perturbations include the number of actuators and sensors, and what type of servos are implemented on the structure.

Mass perturbations consisted of the addition of a gimbal or the modification of the reaction-wheel assembly. In another example, extra mass was added to the EM disturbance gimbal to evaluate the robustness of a controller. Experimental data were collected and compared favorably with the closed-loop predictions of the perturbed model.

In the case of stiffness perturbations, experimental data were taken both with and without the active strut in the EM and the FM structure. Both sets of data compared well with the corresponding model, despite a 15% change in frequency caused by the greater stiffness of the active strut.

For the suspension perturbation, masses were placed on the outer two suspension rods of the FM in an effort to separate the cable violin frequencies. Because of the symmetry of the MACE test article, these modes are not distinct in frequency, causing difficulties in model updating and control design. When these masses were added to the rods, the data showed that the modes did indeed separate. This effect also was seen in the perturbed FEM, confirming the accuracy of the suspension model.

The number and type of sensors and actuators affect how many different I/O combinations are available to provide information about the behavior of the structure. In the case of the FM, there are 180 different sensor–actuator combinations. In addition, their type and location are important (i.e., strain vs inertial).

A final type of perturbation is the application of control. Because the servos couple much of the structural dynamics (including damping) together, closing a servo loop places further scrutiny on the structural model. Moreover, because some of the servos involve the real-time computer, they help verify the closed-loop impact of time delay and sensor dynamics.

Multivariable control couples structural elements in ways that are different from the passive structure. Thus, the predictions of the closed-loop behavior, when compared to the closed-loop data, identify those model errors, though small in open-loop performance, that have substantial impact on closed-loop performance.

An additional use of the closed-loop results is to suggest improvements in the model, called closed-loop updating. Generally, these improvements are not made in a methodical way, but rather are made through engineering insight. In closed-loop updating, controllers are designed using the 1-*g* I/O FEM and implemented on the suspended structure.[11] Errors in the model will cause the controllers to go unstable at some level of control authority. Information on which modes are driven unstable is used either to examine the model assumptions for likely causes of the errors or to place added emphasis on the problem modes in the automated update.

Once the model has been updated on the basis of one set of closed-loop results, controllers are designed using the new model, and closed-loop updating occurs again. The question of when to stop updating is answered using the measurement model. Controllers are designed using the measurement model and implemented on the structure, which, because the measurement model is the most accurate representation of the experimental data, should achieve the highest performance. The performance achieved using the measurement model is compared with that achieved using the FEM. When these performances are comparable, the FEM updating has converged. This is the case with MACE FEM,[12] which achieved improvements of 40 dB in the pointing performance of both axes of the primary payload.

Table 1 summarizes the final 1-*g* model including the modal frequencies, damping ratios, and their errors, where a negative sign on

Table 1 Modal frequencies and damping ratios for 1- and 0-*g* FEMs and frequency and damping errors for 1 and 0 *g*

	Frequency					Damping ratio			
	FEM, Hz		Change,	Error, %		FEM, %		Error, %	
Description of mode	1 *g*	0 *g*	%	1 *g*	0 *g*	1 *g*	0 *g*	1 *g*	0 *g*
Suspension rotary pendulum	0.24	—	—	4.3	—	9.1	—	9.3	—
Suspension tilt	0.44	—	—	−2.3	—	7.8	—	29.1	—
Suspension compound pendulum	1.55	—		0.0	—	4.5	—	37.2	—
First vertical bending	2.13	2.10	−1.4	−0.5	−2.1	2.7	2.4	7.0	−1.8
First horizontal bending	3.74	3.31	−12.7	−0.5	−0.1	17.8	10.2	52.3	−3.0
Gimbal pendulum	3.38	3.07	−10.1	−0.5	−1.4	37.8	39.9	−27.7	9.4
Gimbal pendulum	5.26	5.11	−2.9	−0.9	−13.7	46.3	63.6	−14.6	5.5
Gimbal pendulum	5.32	5.37	0.9	−1.1	−4.3	72.2	73.9	2.6	3.8
Gimbal pendulum	5.50	5.47	−0.5	1.9	−2.7	76.1	75.9	7.6	6.8
First violin	7.26	—	—	−1.4	—	1.3	—	−53.3	—
First violin	7.34	—	—	0.1	—	2.4	—	−35.8	—
Second vertical bending	8.98	8.76	−2.5	−1.9	1.5	29.0	41.0	37.6	−12.5
Second horizontal bending	10.12	10.09	−0.3	−2.3	−1.5	3.7	2.4	54.1	26.3
Third vertical bending	10.45	10.11	−3.4	0.6	3.7	40.1	44.5	21.1	−6.4
First violin	11.65	—	—	0.3	—	1.4	—	−20.6	—
Third horizontal bending	12.16	11.86	−2.6	1.8	n/a[a]	43.2	55.1	31.8	n/a
Third horizontal bending	14.09	14.07	−0.2	−3.8	−3.1	31.0	40.4	4.8	−22.6
Second violin	14.68	—	—	1.9	—	2.1	—	−15.8	—
Second violin	14.71	—	—	−0.8	—	1.3	—	−38.1	—
Fourth vertical bending	15.82	15.68	−0.9	2.4	1.6	9.0	8.3	3.2	17.0
Fourth horizontal bending	17.81	17.84	0.1	−0.4	1.3	1.0	0.9	−18.4	−22.6
Fifth horizontal bending	24.92	25.05	0.5	−0.4	1.0	2.0	2.3	−27.0	−22.5
Fifth vertical bending	39.74	39.89	0.4	−1.7	3.9	2.1	2.1	5.0	7.6
Sixth vertical bending	49.29	49.71	0.8	2.6	2.1	2.2	2.2	5.0	92.0
Sixth horizontal bending	49.22	50.03	1.6	1.4	n/a	1.8	2.0	−31.4	n/a
Seventh horizontal bending	56.39	57.38	1.7	−3.9	−13.3	1.6	1.7	−11.0	5.5

[a]n/a indicates that the mode was undetectable in the data.

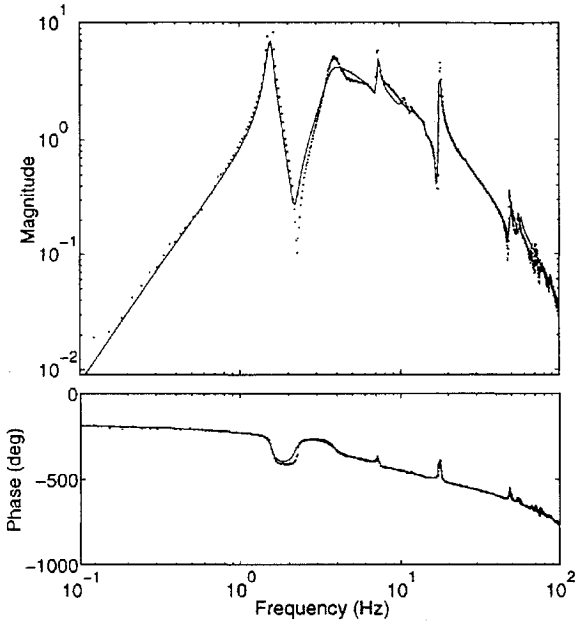

1-g reaction wheel X to bus rate gyro X

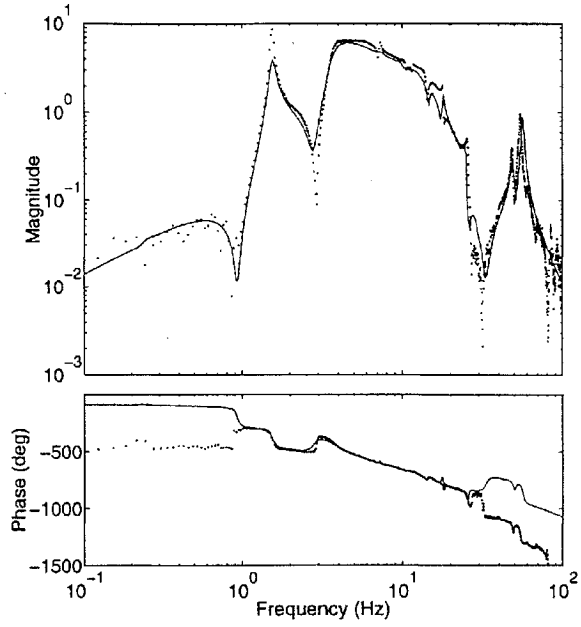

1-g primary gimbal X to bus rate gyro X

Fig. 5a Comparison of the final 1-g model (——) with measured 1-g data (· · · · ·).

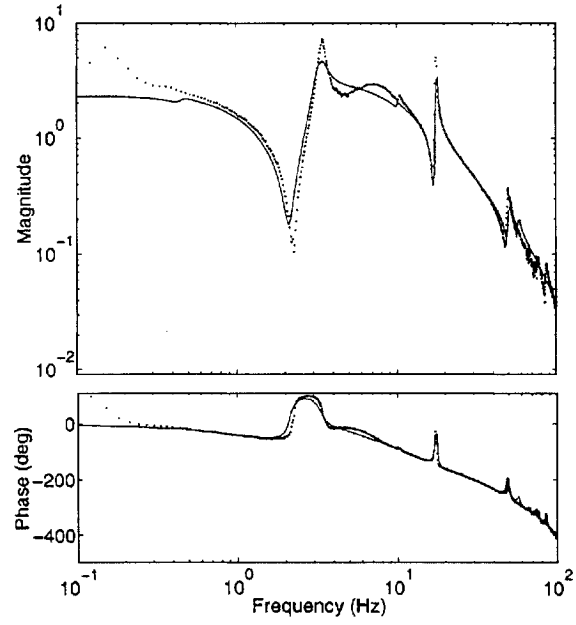

0-g reaction wheel X to bus rate gyro X

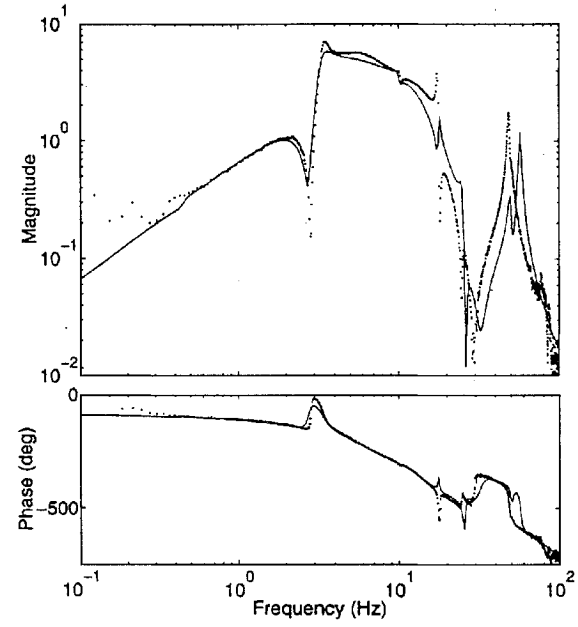

0-g primary gimbal X to bus rate gyro X

Fig. 5b Comparison of the final 0-g model (——) with measured 0-g data (· · · · ·).

the errors indicates that the actual value was smaller than the FEM value. This table shows that the largest frequency error is only 4%, with the majority of the errors being less than 2%. This correlation of predicted and experimental frequencies is excellent for a model of this size over such a large frequency range. The table also shows many damping ratios larger than 3%, with several modes having damping ratios larger than 10%. This is primarily because of the influence of the gimbal servo loops. Figure 5a shows the agreement between the 1-g model and ground data with the local servo loops closed. The agreement between the two is very good. The modes overlay very well, both in frequency and in damping, as do the zeros. This is important because improperly modeled zeros can cause serious problems in high-authority control.

High-Fidelity 0-g I/O Model

To form the 0-g model, the nodal points, masses, and elements that make up the model of the suspension system are removed. Because the test article does not have gravity preloading while in orbit, the iterative procedure used for the 1-g model is no longer needed. Thus, the 0-g frequencies and mode shapes are simply obtained from a linear eigensolution. Because there is no method for determining the damping in 0 g, the structural damping ratios used in the 1-g model are again used for the 0-g model. Finally, the electronics and time delays are added to arrive at a 0-g I/O model.

The resulting 0-g modal frequencies and damping ratios are given in Table 1, along with their errors. Note that there are fewer flexible modes in 0 g, because of the absence of suspension modes. It is also seen from column 4 (frequency change) that most of the frequency shifts are rather small, with one notable exception, the first horizontal bending mode. This mode undergoes a large shift because it is strongly coupled with the suspension system, which is absent in 0 g.

Table 1 shows that the 0-g frequency errors are similar in magnitude to those in 1 g. The small errors in 0 g indicate the success of the MACE approach in predicting the 0-g behavior of the test article. It is also seen in Table 1 that two of the modes present in

the 0-g model were not detected in the data. One of these modes is lightly damped and could present problems in control design. The model predicted two modes in the 50-Hz range, but the flight data indicated only one. An explanation is that the two modes present in the model have merged into a single mode.

Table 1 also shows that the 0-g damping was also predicted fairly accurately. In comparing the 1- and 0-g damping errors, it appears that many of the 0-g errors are actually smaller in magnitude than those in 1 g. This is a reflection of the difficulty in modeling damping, where there is no technique for extrapolating 0-g damping ratios from 1-g testing, which is further complicated by the fact that the servos dominate the damping of many of the modes. From the MODE results,[8] it was seen that the damping increased from 1 to 0 g. Thus, the large 1-g errors can be explained as an attempt to try to lead the 0-g damping by making the 1-g model more heavily damped than the ground data.

To examine the changes in the I/O behavior of the structure, Fig. 5 can be used to make the comparison between 1- and 0-g models. Above 15 Hz, the general shapes of the frequency responses do not change from 1 to 0 g. Below this frequency, however, the 0-g model exhibits some significant differences. The 1.5-, 7-, and 14-Hz suspension modes have disappeared from the 0-g model. The frequency of the first horizontal bending mode has undergone a significant drop from 3.7 Hz in 1 g to 3.3 Hz because of the large interaction with the suspension. The zero below 2 Hz in the gimbal response has disappeared because it was attributable to the suspension system. Finally, the dc value for the reaction-wheel response changes from zero in 1 g to a nonzero value in 0 g, because the suspension is no longer present to restrain the bus from rotating in response to a speed command into the reaction wheels.

Figure 5b shows the comparison of the 0-g model with the flight data with the local servo loops closed. These 0-g comparisons with data show results similar to those on the ground, providing stronger evidence of the success of the MACE approach, with only a few notable exceptions. First, at low frequency near 0.2 Hz, there is an indication of a mode in the data but not in the model, indicating an unmodeled 0-g suspension mode. The other exception is in the 50-Hz range, where the data indicate that the two modes in the model have merged into a single mode.

Once the nominal 0-g model has been obtained, the last step in the approach is the estimation of the likely errors needed for control design. This so-called uncertainty model must rely on ground testing and the analytical insights gained from the FEM. It must not only contain information about the errors in the nominal model, but must also include information on the variation in the experimental data caused by disassembly/reassembly, sensor noise, nonlinearities, and so forth. The determination of this uncertainty model consists of projecting the 1-g errors into 0 g, using the perturbation method to derive the projection operator.[10] This uncertainty model completes the approach laid out in the beginning of this paper, resulting in a 0-g model suitable for robust control design, which has been refined using ground test data.

The final and perhaps strongest evidence of the success of the MACE approach and the accuracy of the resulting models is in the closed-loop results. Both on the ground and on orbit, the I/O models allowed high-authority controllers to be designed and successfully implemented. These models enabled 38- and 35-dB pointing improvement in 1 and 0 g, respectively.

Concluding Remarks

Prediction, with confidence, of the on-orbit behavior of a structure requires an approach that combines both analytical modeling techniques and ground-based experimentation. When the eventual goal is a model to be used in control design, this approach must also include closed-loop testing. An example of such an approach has been implemented on the MACE. The result of this approach is a 0-g model that has been updated using 1-g experimental data and verified through 1-g closed-loop experimentation. A 0-g model with less than 4% error in modal frequencies enabled controllers to be designed prior to the launch of MACE, achieving a 35-dB improvement in 0-g pointing accuracy over standard industry practice.

Acknowledgments

This work was supported by the NASA IN-STEP Program and NASA Langley Research Center Controls-Structures Interaction Office, with Gregory Stover and Jerry Newsom as contract monitors, under Contract NAS1-18690.

References

[1]Ljung, L., *System Identification: Theory for the User*, Prentice–Hall, Englewood Cliffs, NJ, 1987.

[2]Rey, D., Crawley, E., Alexander, H., Glaese, R., and Gaudenzi, P., "Gravity and Suspension Effects on the Dynamics of Controlled Structures," *Proceedings of the AIAA Structures, Structural Dynamics, and Materials Conference* (La Jolla, CA), AIAA, Washington, DC, 1993, pp. 3156–3171.

[3]Visser, W., and Imregun, M., "Technique to Update Finite Element Models Using Frequency Response Data," *Proceedings of the International Modal Analysis Conference* (Florence, Italy), Society for Experimental Mechanics, Bethel, CT, 1991, pp. 462–468.

[4]Flanigan, C. C., "Test/Analysis Correlation Using Design Sensitivity and Optimization," *Proceedings of the Aerospace Technology Conference and Exposition* (Anaheim, CA), Society of Automotive Engineers, Warrendale, PA, 1988 (Paper 871743).

[5]Campbell, M. E., Grocott, S. C. O., and How, J. P., "Overview of Closed Loop Results for MACE," *Proceedings of the International Federation of Automatic Control World Congress* (San Francisco, CA), Elsevier Science Ltd., Oxford, England, UK (to be published).

[6]Glaese, R. M., "Development of Zero-Gravity Structural Control Models from Ground Analysis and Experimentation," M.S. Thesis, Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology, Cambridge, MA, 1994.

[7]Liu, K., Jacques, R. N., and Miller, D. W., "Frequency Domain Structural System Identification by Observability Range Space Extraction," *Proceedings of the American Control Conference* (Baltimore, MD), Inst. of Electrical and Electronics Engineers, Piscataway, NJ, 1994, pp. 107–111.

[8]Crawley, E. F., Barlow, M. S., van Schoor, M. C., and Bicos, A. S., "Variation in the Modal Parameters of Space Structures," *Proceedings of the AIAA Structures, Structural Dynamics, and Materials Conference* (Dallas, TX), AIAA, Washington, DC, 1992, pp. 1212–1228.

[9]Anderson, E. H., Blackwood, G. H., and How, J. P., "Passive Damping in the MIT SERC Controlled Structures Testbed," *Proceedings of the International Symposium on Active Materials and Adaptive Structures* (Alexandria, VA), 1991, pp. 13–18.

[10]Campbell, M. E., Grocott, S. C., How, J. P., Miller, D. W., and Crawley, E. F., "Verification Procedure for On-Orbit Controllers for the MIT Middeck Active Control Experiment," *Proceedings of the American Control Conference* (Seattle, WA), Inst. of Electrical and Electronics Engineers, Piscataway, NJ, 1995, pp. 3600–3605.

[11]How, J. P., Glaese, R. M., Grocott, S. C., and Miller, D. W., "Finite Element Model Based Robust Controllers for the Middeck Active Control Experiment," *IEEE Transactions on Control Systems Technology* (to be published).

[12]How, J. P., and Miller, D. W., "Assessment of Modelling and Robust Control Techniques for Future Spacecraft: Middeck Active Control Experiment," *Proceedings of the AAS Guidance, Navigation, and Control Conference* (Keystone, CO), American Astronomical Society, San Diego, CA, 1994 (AAS Paper 94-053).