

Active Flutter Suppression Control Design for Wind-Tunnel Testing using μ -Synthesis

Logan S. Anderson* and Ryan J. Caverly†
University of Minnesota, Minneapolis, Minnesota, 55455

Jared A. Grauer‡
NASA Langley Research Center, Hampton, Virginia, 23681

An active flutter suppression (AFS) control law for the Integrated Adaptive Wing Technology Maturation (IAWTM) sub-project designed to enable wind-tunnel testing at dynamic pressures and Mach numbers beyond the flutter boundary in the Transonic Dynamics Tunnel (TDT) at the NASA Langley Research Center is presented. The test article is a flexible half-span model of a transport airplane based on the Common Research Model, with the aspect ratio of the wing increased to 13.5. The AFS control law is designed using μ -synthesis to be robust against a combination of aeroelastic excitation and structural model uncertainty. The proposed AFS controller is evaluated analytically and tested in a simulation of the system with the controller in the loop, with results demonstrating that the controller is able to extend the flutter boundary to higher Mach numbers and dynamic pressures than required for experimental testing as part of the IAWTM sub-project.

I. Introduction

BOTH economic and environmental factors for fixed-wing aircraft drive design objectives to improve the operational range, fuel efficiency, and performance of aircraft. NASA worked to advance these objectives through the Advanced Air Transport Technology (AATT) project [1–3]. The Integrated Adaptive Wing Technology Maturation (IAWTM) sub-project within the AATT is a collaboration between Boeing and multiple NASA centers, with wind-tunnel testing occurring in the Transonic Dynamics Tunnel (TDT) at the NASA Langley Research Center in Hampton, VA [2]. The IAWTM model is a half-span, side-wall-mounted model based on the NASA Common Research Model with the aspect ratio increased from 9 to 13.5 [2, 3] and is shown mounted in the TDT in Figure 1. The IAWTM sub-project has several objectives, with the focus of this report being the development and validation of control architecture for active flutter suppression (AFS). Additional objectives of the IAWTM sub-project include control design for gust load alleviation (GLA), maneuver load alleviation (MLA), and real-time drag minimization (RTDM) [2]. The GLA, MLA, and RTDM controllers are only considered in this work in the context of ensuring that the presented AFS control schemes do not interfere with those objectives and are able to suppress any modes excited by those controllers.

The development and experimental validation of closed-loop AFS control laws is one of the scientific goals of the IAWTM sub-project. While the TDT has numerous features to both minimize the risk of facility damage due to debris from a destructive failure of the test article and prevent destructive failures from occurring, the AFS control schemes help provide additional safety margin to prevent such failures from occurring in the first place [4, 5]. Preventing catastrophic aeroelastic instability is a significant concern for the wind-tunnel testing as damage to the test article or tunnel is likely to result in a significant loss of research potential and decrease facility availability for other research projects.

The IAWTM is a challenging system for AFS control design due to a combination of design elements. There are numerous structural modes at relatively low frequencies arising from the highly flexible wing. This is highlighted by three of the first four structural modes occurring between 7.5 and 8.3 Hz, with the tight spacing of the structural modes leading to strong modal coupling and inflating the size of the analytical models required for control design. In addition to the flexible dynamics, the aerodynamics introduce more complexity and uncertainty to the models. The wing must traverse a large range of conditions in the TDT to reach the test regime, with the allowable test envelope for the model extending to a maximum of Mach 0.85 and dynamic pressure of 230 lbf/ft². This creates a large range in aerodynamic forces, and therefore significant changes in the system dynamics, such that control laws must be robust to these changes.

*Ph.D. Candidate, Department of Aerospace Engineering and Mechanics. Student Member AIAA.

†Associate Professor, Department of Aerospace Engineering and Mechanics. Member AIAA.

‡Research Aerospace Engineer, Vehicle Systems and Control Branch, MS 410. Associate Fellow AIAA.




Fig. 1 The IAWTM model mounted in the TDT wind tunnel (credit: NASA / Mark Knopp).

To further compound these challenges, there are relatively large uncertainties associated with the aerodynamic models used in control design. The combination of these factors and other challenges in accurately representing the model and test conditions necessitates an AFS control design that is robust against large deviations from the nominal model and test condition. The IAWTM test article was installed in the TDT in July 2024 and a test series focused on validating the aerodynamic, structural, and actuator models was performed. The test article was reinstalled in the TDT for an additional test series in July 2025, where testing flutter suppression control laws was one of the test objectives.

Many AFS control laws have previously been developed in literature and applied to physical systems. An overview of the state-of-the-art methods for AFS, as well as a review of areas where the technology needs further maturation is provided in Ref. [6]. The list of methods for AFS control design that have been studied or implemented includes classical techniques, optimal and robust control methods, adaptive control design, neural networks, and deep and reinforcement learning based approaches. Several reports have noted that despite significant focus and effort on the design and tuning of modern and advanced control synthesis methods using large multiple-input multiple-output (MIMO) models that arise from aero-structural models of aeroelastic systems, resultant designs have not been found to be better than control schemes based on classical single-input single-output (SISO) models that typically produce much simpler designs [6–8]. This is highlighted by the NASA Active Flexible Wing program where a classical design was able to push the flutter boundary back for a test article in a wind tunnel through the use of accelerometer output feedback, whereas the modern control designs (modal rate feedback and linear quadratic methods) were insufficiently robust against modeling errors [8–11].

Active development for AFS control laws is ongoing based on multiple control frameworks including classical control and LQG. This work presents an AFS design method currently under development for the IAWTM sub-project. The method in this work achieves the design of a robust control law through μ -synthesis (also known as *DK*-iteration) [12, 13]. The past 30 years has seen significant maturation in modeling techniques to represent aeroservoelastic (ASE) aircraft as linear parameter-varying systems, as well as the tools to solve robust control problems using those models [14–16]. Historically, μ -synthesis has been applied to AFS with limited efficacy [6]; however, continued improvements in computing power and developments in modeling have led to highly successful use of μ -synthesis for GLA objectives [17], as well as AFS design using the X-56A as a model [18]. While both Ref. [18] and this work apply μ -synthesis to AFS design, the problem is approached in different ways. Reference [18] uses strain sensors to estimate the shape of the wing and drive the deformed shape to match virtual deformation reference signals, whereas this work focuses on suppressing wing-tip vibration using accelerometer measurements without directly considering the bending modes. This approach provides increased flexibility in the control design at the cost of the controller being less physically interpretable. Frameworks have been proposed to synthesize robust controllers using a limited number of physically

interpretable tuning parameters at the cost of restrictions on the uncertainty considered during synthesis [19]. The successful synthesis of an AFS controller is predicated on a combination of careful modeling of the uncertainty sources in the system and judicious model reduction to reduce the scale of the problem to fit within the available computing resources. As the development of both of these are reliant on advice of experts in the relevant fields and ad hoc intermediate heuristics, one of the goals of this work is to document and disseminate the reasoning and design process involved. In summary, the contributions of this work are the development of a robust control scheme for suppressing ASE flutter in the IAWTM model, as well as documentation of the motivation for the use of μ -synthesis and the design choices made that enable the use of μ -synthesis for this application.

The remaining portions of this paper are structured as follows. Section II presents details about the TDT test environment and physical IAWTM model. Section III contains the numeric modeling used to create plant models for control synthesis and simulation. Section IV presents the μ -synthesis framework, objective weighting, and uncertainty modeling. Section V presents the resulting AFS control law and stability evaluation. Section VI includes simulation results for the AFS control law in feedback with the uncertain system. Section VII concludes this paper.

II. Test Article and Experimental Facilities

The design and construction of the IAWTM test article, as well as the capabilities of the TDT and how they are related to aeroelastic testing, have been previously reported [2–5]. A short overview is included for completeness.

The test article is a side-wall-mounted fuselage and wing model with 10 trailing edge control surfaces, of which 3 are hydraulically-driven ailerons and 7 are electronically actuated miniature plain flaps (MPF). The fuselage is 13.3 ft long and the half-span distance of the wing is 12.9 ft, giving the wing an aspect ratio of 13.5. The model features a flow-through nacelle mounted below the wing, and omits any form of empennage. There are three model configurations used for different parts of the wind-tunnel testing. These configurations differ through the addition or removal of a small ballast mass at the wing-tip and the inclusion or removal of the strain gauge balance from the load path. This work is only concerned with the flutter configuration, which includes the ballast mass and removes the balance from the load path. Figure 1 shows the test article mounted in the TDT in a configuration without the wing-tip ballast. In addition to the trailing-edge control surfaces, there is a pitching mechanism that controls the angle of attack of the wing, and four upstream gust vanes in the TDT. The combination of trailing edge control surfaces, pitch actuation, and gust vanes are the 15 inputs used in this work.

The wing has numerous embedded sensors that provide measurement data accessible by the controllers. Three tri-axial accelerometers and 13 single axis accelerometers provide acceleration data from points along the wing. Figure 2 shows the locations of the accelerometers on a planform of the wing. Strain data along the main wing spar is provided by 10 full-bridge strain gauges. Additional strain information comes from three fiber-optic strain sensors (FOSS) that are embedded on the wing, which provide strain measurements at a total of 117 locations along the wing.

The tests were performed in the TDT 16 ft by 16 ft test section. The TDT can operate up to Mach 1.2 and dynamic pressures up to 320 lbf/ft², and either normal air or heavy gas R-134a can be used to enable aerodynamic scaling. The bypass valves in the tunnel allow for rapidly reducing the dynamic pressure in the test section in the event that unsafe dynamic behavior is observed during testing. All wind-tunnel tests for the IAWTM sub-project were performed in heavy gas R-134a, with the highest energy test condition initially planned for the July 2025 tests being Mach 0.85 with a dynamic pressure of 230 lbf/ft², while the highest condition achieved was Mach 0.72 with a dynamic pressure of 163 lbf/ft².

The sensor data from the IAWTM model is streamed to computers in the TDT control room in real time using an interface built through a combination of Simulink[®], MATLAB[®], and dSPACE[®] products. This interface provides an environment that enables control laws designed in MATLAB[®] to provide commands to the control surfaces in real time. The setup of this interface for the IAWTM wind-tunnel tests requires that the control laws be discretized to run at 200 Hz. Historically, AFS control laws tested in this setup are designed in continuous time and then discretized after satisfactory performance is achieved in continuous time simulations.

III. Numeric Modeling

Performing control synthesis requires a mathematical representation of the physical model. The development of a numeric model for this system is non-trivial and is an ongoing effort using the data from the testing in July 2024. The

*The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

models reported in this work are the most current models as of January 2025. The rest of this section is divided to present each aspect of the numeric model that makes up the overall plant model. The structural model of the wing is presented in Section III.A, followed by the aerodynamic models in Section III.B, the actuator models in Section III.C, and the sensor models in Section III.D. The nominal model of the overall system is then presented in Section III.E and the generalized plant formulation needed for control synthesis is presented in Section III.F. In equations throughout this work, signals and systems labeled in bold type are multi-channel or MIMO, while standard type indicates single-channel or SISO respectively. Additionally, $\mathbf{0}$ is used to indicate a matrix of all zeros and $\mathbf{1}$ indicates an identity matrix.

A. Structural Model

The structural dynamics were modeled in MSC Nastran[®] using finite element methods (FEM) for each of the three configurations; however, only the flutter configuration is included in this work. The FEM models were subsequently tuned and validated through ground vibration testing with the model mounted in the TDT and data from the 2024 wind-tunnel test series. More detailed discussion on the development of these models is available in Refs. [2] and [20]. This produced a linear model including 20 vibration modes starting at 2.7 Hz and ranging up to approximately 100 Hz. The first 4 modes are presented in Table 1 and shown in Fig. 3. Analysis early in the IAWTM sub-project indicated that the listed modes were the most concerning for destructive flutter, and their frequency range factored into subsequent control design.

The numeric structural model is tightly integrated with the aerodynamic models discussed in Section III.B. The complete model features two states per structural mode, accounting for 48 total states in the model.

Table 1 Descriptions of the first 4 vibration modes in flutter configuration

Mode	rad/s	Hz	Description
1	16.9	2.7	Wing 1st vertical bending
2	50.5	8.0	Wing 2nd vertical bending, ballast pitch
3	51.6	8.2	Wing 1st fore-aft bending, vehicle yaw
4	64.1	10.2	Wing 2nd vertical bending, ballast pitch, vehicle pitch




Fig. 2 Linear accelerometer layout for the IAWTM test article [20]. The hydraulically-powered ailerons are labeled as δ_A , δ_B , and δ_C , while the unlabeled control surfaces are the MPFs.




Fig. 3 The first 4 vibration modes in flutter configuration. Red indicates upward deflection, while blue indicates downward. Flutter ballast visible protruding from the wing-tip. Figure from Ref. [20], reproduced with permission.

B. Aerodynamic Models

The aerodynamic models are developed using panel methods and assuming potential flow in ZAERO™ by ZONA Technology Inc. [21]. Rational function approximations are fit to 24 generalized aerodynamic forces with 12 lag states each. Dynamic pressure is used to scale the resultant models over a grid of test conditions, with a total of 468 grid nodes distinguished by the combination of Mach number and dynamic pressure. In total, the grid involves 18 unique Mach numbers between 0.3 and 0.95 and 26 dynamic pressures between 50 and 300 lbf/ft². These aerodynamic models use the angular positions, velocities, and accelerations of each of the 15 control surfaces as inputs. The details about the control surfaces are discussed in Section III.C.

The aerodynamic model for each condition is then combined with the previously described structural model. This results in a state-space model for each condition with 45 inputs and 346 states, with 48 of the states accounting for structural dynamics and the remaining 298 states accounting for the aerodynamic effects in the system. The output dimension of this system is entirely dependent on which sensors are being considered, with details of the sensor modeling available in Section III.D. Throughout the remainder of this work the combined structural and aerodynamic model is referred to as the aeroelastic model, while the combined actuator, structural, aerodynamic, and data acquisition system model is referred to as the plant.

The aeroelastic model conditions are shown in Fig. 4. The open-loop stability of the aeroelastic models is shown by the color and shape of the markers. Blue dots indicate open-loop stable models, and the red triangles and black square indicate unstable conditions where flutter is expected to occur. The black square indicates the model with Mach 0.85 free stream velocity and 230 lbf/ft², which is the primary focus for the ongoing flutter suppression efforts as it is analogous to the cruising condition of commercial transport aircraft. Figure 4 also shows the boundaries where structural loading constraints limit the maximum control surface deflection to 10 and 20 deg as gold dashed lines. The lines at 300 lbf/ft² and Mach 0.95 additionally indicate the total structural limit for the wing.

C. Actuator Models

The full aeroservoelastic model developed for the IAWTM sub-project features a total of 15 actuators. There are four upstream gust vanes in the TDT, a pitching mechanism that controls the angle of attack of the wing, and 10 trailing edge control surfaces. Of these actuators, only three of the trailing edge surfaces are available for use in AFS. Ailerons A, B, and C are labeled as δ_A , δ_B , and δ_C in Fig. 2, with the control surface ordering beginning with the surface closest to the fuselage and increasing towards the wing-tip.

The actuator models used in this work are based on a combination of bench-top tests and prior rounds of testing conducted in the TDT. The actuators are assumed to be irreversible for modeling purposes. The three hydraulic ailerons have sufficient bandwidth across all dynamic pressures used for AFS testing. However, the linkage connecting the hydraulic actuator to aileron B reduces the responsiveness for small deflections, with commanded deflections below 0.2 deg magnitude resulting in no deflection of the control surface. Third-order transfer functions are used to model the actuator dynamics, with input delays to account for the lag between a control signal being sent and the control surface moving [20].




Fig. 4 Aerodynamic modeling conditions of the IAWTM test article are defined by the combination of dynamic pressure and Mach number. The operating conditions where the control surface deflections are limited to 10 and 20 degree deflections are indicated by the yellow dashed lines. The black square is the test condition that is the primary focus for flutter suppression, with the blue dots and red triangles indicating open-loop stable and unstable models, respectively.

The actuator model treats the 15 actuators as decoupled, with the command signal for each as inputs and the angular acceleration, velocity, and position of each actuator as outputs. This leads to a state-space realization with 15 inputs, 10 of which have input delays, 45 outputs, and 45 states. A 3rd-order Padé approximation is used to replace the input delays with additional states for compatibility with standard linear time-invariant (LTI) synthesis tools. The use of a 3rd-order approximation was found to be an acceptable tradeoff between accurately representing the original model with delays and inflating the size of the resulting state-space model. This results in the delay-free state-space version of the actuator model with 15 inputs, 45 outputs, and 75 states. A Bode plot of the input to angular position response of each of the ailerons is included in Fig. 5. While the model of the ailerons is retained as the actuators used in AFS control, the 12 other actuators are also retained as they provide a convenient framework for informing the control synthesis process how disturbances can excite the system.

D. Sensor Models

Both the actual sensors and the TDT data acquisition system (DAS) play an important role in modeling the sensor data available for use by the control laws. The AFS control development in this work only uses the embedded sensors within the wing. The wing has 16 embedded accelerometers, of which 13 are single-axis accelerometers aligned with either the body-axis x or z directions and 3 are three-axis accelerometers aligned with the body-axis. This results in a total of 22 acceleration measurements from locations across the full span of the wing. There are 3 embedded fiber-optic strain sensors, which collectively provide strain data at 117 locations along the wing. Additionally, there are 10 strain gauges providing strain data for both the upper and lower surface of the wing resulting in 20 output channels. Between the three types of sensors there are 159 sensor channels available to the AFS controller. Due to high bandwidths, sensors were modeled as effective time delays.

The simplest model of sensor noise that enables controller design is used to avoid inflating the order of the resulting plant model. The accelerometers are modeled as having zero-mean Gaussian white noise with a standard deviation of 0.001 g. The FOSS strain measurements are modeled as being corrupted by zero-mean Gaussian white noise with a standard deviation of 10^{-5} in/in, while the strain gauges are modeled with zero-mean Gaussian white noise corruption with standard deviation of 5×10^{-6} in/in. These noise values are based on a combination of bench-top testing and the manufacturer specifications for the sensors, with the specific values based on the maximum expected noise magnitude between 2 and 200 Hz, the frequency range of interest for flutter. The actual sensor dynamics are not included in the modeling, as the bandwidth of these dynamics is indicated to be a minimum of two decades above the flutter band.




Fig. 5 Control signal to angular position response for each hydraulic aileron. The blue, orange, and yellow lines are ailerons A, B, and C, respectively. The red dashed line at 6.4 Hz indicates the expected frequency of the first flutter mode.

The data acquisition system acts as the interface to provide the sensor data to the control policy. For this test, the DAS in the TDT was configured to use four-pole anti-aliasing filters with break frequencies of 400 Hz on the sensor measurements by default. This was not included in the model used for control synthesis, as including it inflates the number of states by 4 per sensor used in both the system model and resulting controllers. Furthermore, the performance and control cost weight functions in the control design (see Section III.F) penalize control response to high frequency signals that would be hidden by the anti-aliasing filters enough that the resulting controllers have minimal response at those high frequencies. While they are omitted from the control design, the anti-aliasing filters are included in the simulations presented in Section VI.

E. Nominal Plant

The nominal plant model is assembled by combining the actuator models, aeroelastic model, and DAS model in series as shown in Fig. 6. The control signals for the actuators, δ_U , act as the input, and become the control surface deflections, rates, and accelerations, δ , after passing through the actuator model. These are fed into the aeroelastic model, which feeds state information into the DAS to ultimately produce sensor measurements, y .

The overall nominal plant at each tunnel condition has 3 control input channels and 12 disturbance input channels corresponding to the remaining actuators not used in flutter suppression. There are 493 states after the Padé approximation is applied to the component elements. There are a maximum of 159 sensor output channels from this model; however, attempts to synthesize robust control laws using all available sensors were found to be numerically intractable and fail to generate a stabilizing controller. The selection process to find a small but useful subset of sensors is discussed as part of the robust control synthesis in Section IV.C.

F. Generalized Plant

While the nominal model shown in Fig. 6 is a useful representation of the dynamics of the aeroelastic wing, multiple components needed for designing a controller are still missing. Specifically, this model does not provide any information about how disturbances and measurement noise interact with the system or what performance objectives the controller




Fig. 6 Block diagram of the plant model for IAWTM simulation and control synthesis.

is to optimize. These useful elements for control design are included through expansion from the nominal plant to the generalized plant. The control synthesis algorithms consider the exogenous inputs to be white-noise, with weighting filters used to color the white-noise to match the known magnitude and frequency properties of the true signals. The desired closed-loop frequency response can be designed and used to build the performance weights, which inform the control synthesis algorithms of what the performance requirements are.

A block diagram of the generalized plant is shown in Fig. 7a. The $G(s)$ block represents all of the nominal plant model from Fig. 6, while the $W_d(s)$ and $W_n(s)$ blocks are filtering the exogenous white noise, and the $W_u(s)$ and $W_e(s)$ are the performance weighting filters. A consolidated feedback interconnection between the generalized plant and the controller is shown in Fig. 7b. The exogenous input signal $w(s)$ is composed of the disturbance and noise signals, while the performance signal $z(s)$ is composed of the $z_1(s)$ and $z_2(s)$ outputs. This layout highlights the specific control objective, optimizing the performance ($z(s)$) of the generalized plant-controller feedback interconnection in response to the exogenous white noise input ($w(s)$).




Fig. 7 Generalized plant framework used for control synthesis.

The design of the assorted weighting filters has a significant influence on the synthesized controller, with $W_n(s)$ and $W_d(s)$ informing the synthesis algorithm of the expected frequency and magnitude content of the noise and disturbances, while $W_e(s)$ and $W_u(s)$ are used to specify what signals need to be minimized and at what frequencies. It is generally advisable to use the lowest order filters that can sufficiently represent the expected inputs or desired performance characteristics. The noise weighting filters $W_n(s)$ are based on the noise characteristics associated with each type of sensor as discussed in Section III.D. As the noise is being modeled uniformly in the frequency domain, this is the simplest of the filters with a static gain defined for each of the three classes of sensors.

The disturbances on the system are modeled in this work as stochastic deflections of the 15 actuators with a maximum of 1 deg standard deviation. The expected flutter frequency band is between 2 and 20 Hz, with no correlation modeled between the disturbances acting through each of the actuators. This leads to a 15 input, 15 output disturbance filter $W_d(s)$, composed of 15 identical SISO filters. Noting that the actuator model uses radians as inputs and the filter frequencies are written in radians per second, each of these filters is then given by

$$W_d(s) = 0.0175 \frac{s}{s + 2\pi(2)} \frac{2\pi(20)}{s + 2\pi(20)}, \quad (1)$$

as shown in Fig. 8a. Higher-order filters with faster attenuation outside of frequency range of interest were examined; however, were not found to improve the performance of the resulting controllers in simulation.

There are two additional filter blocks in the generalized plant that are weighting the performance signals. The performance filters inform the control synthesis algorithms what frequency bands of the performance signals are of

concern, and the relative importance of each aspect of the performance signal. In this framework, the two filters are acting on the amount of vibration in the wing $W_e(s)$ and the total control effort $W_u(s)$. This combination of weights serves to guide the control synthesis to suppress flutter (i.e., vibrations), while also respecting the physical limits of the control surfaces. Three frequency bands are of interest for AFS performance, consisting of a low-frequency band where flutter is not a concern and control action should be left to the MLA and GLA controllers, the mid-range where flutter is expected, and the high frequency band where no modes are expected to result in damaging flutter.




Fig. 8 Bode plots of the weighting filters for the generalized plant. The vertical black bars indicate the primary frequency range of interest for active flutter suppression of 2 to 20 Hz, whereas the first flutter mode is expected to occur at approximately 6.4 Hz and is indicated by the red vertical dashed line.

The vibration performance output filter uses vibration in the wing as a measure of flutter to be minimized. This filter is also referred to as the error filter as the ideal vibration signal is uniformly zero for all time, thus making any vibration an error. The error filter $W_e(s)$ is designed to have large magnitude around frequencies where flutter suppression is a high priority, and small magnitude elsewhere. There are three cost regions in the error filter, with the guiding design principle being that the allowable magnitude of flutter at any specific frequency is proportional to the inverse of the error filter. These three frequency bands are divided by the vertical black bars in Fig. 8b. The filter is described as a SISO

system in this section; however, in practice the $W_e(s)$ filter block repeats the SISO filter for each point on the wing used as a vibration measurement point. This formulation is tuned for the vertical accelerations at locations corresponding with vertical accelerometers. The same design principles apply for designing a filter to use other vibration metrics; however, the magnitude scalings need adjusting to account for differing units (e.g., strain instead of acceleration), and normalization between dissimilar measurements may be required.

The low-frequency band consists of anything below 2 Hz and suppression of oscillations in this range should primarily be left to the MLA and GLA control systems, but the AFS control law is not forbidden from operating in this range. The center-frequency band for this filter ranges between 2 and 30 Hz and represents the primary band where aeroelastic flutter is expected. As the desired performance of an AFS controller is to suppress flutter in this band, the magnitude of the error filter is large here and the inverse (i.e., the allowable flutter magnitude) is small. Ultimately, wing-tip accelerations exceeding 10g are considered to be indicative of flutter and likely to lead to damage to the test article. The highest frequency band is above 30 Hz in the error filter, with this threshold chosen to ensure that both synthesis and analysis encompass at least the first 4 predicted modes and extend approximately half a decade past the 4th modal frequency. Above the 30 Hz boundary where the high frequency band begins, none of the expected modes are predicted to exhibit dangerous unbounded growth. As the limit cycle oscillations in this band are not predicted to be dangerous and the control surfaces have decreased responsiveness at higher frequencies, the nominal cost is set to be 1/10th the low-frequency cost. The error filter is composed of a first-order biproper high-pass filter and a second-order biproper low-pass filter, with a one decade span between the break frequencies for each of the component filters, and the entire filter scaled by a constant to the center band permitted a 10g limit. The spacing of the break frequencies of the high- and low-pass filters are chosen to maximize the filter gain in the center band without raising the gain in the low or high bands more than necessary through an ad-hoc gradient descent, while the filter orders are empirically found to be a useful mix of minimizing the order of the generalized plant and yield sharp enough boundaries between the frequency bands. Ultimately, this results in the low-pass filter $L_P(s)$ and hp filter $H_P(s)$, which are defined as

$$L_P(s) = 0.1 \frac{s + 2\pi(133)}{s + 2\pi(13.3)} \quad (2a)$$

$$H_P(s) = 10 \frac{s + 2\pi(0.4)}{s + 2\pi(4)} \quad (2b)$$

and the overall filter is defined as

$$W_e(s) = 1.97 H_P(s) L_P^2(s). \quad (3)$$

The $W_e(s)$ filter is shown in Fig. 8b; however, the frequency band indicated by the vertical black bars is the 2 to 20 Hz band, not the 2 to 30 Hz used in the design of this filter to avoid premature roll-off in the filter.

The control cost filter $W_u(s)$ is used to inform the control synthesis of how the control surfaces can be used. As this feeds into the cost function the control synthesis is optimizing, this filter is designed to have large magnitude in the frequency bands where control effort should be minimized, and small magnitude where larger control effort is permissible. Additionally, the weights must have non-zero magnitude at all frequencies, which creates a requirement that biproper transfer functions are used, while the inverse of the control cost filter is directly related to the magnitude of the allowable control effort. Low-frequency control effort below the center frequency band is undesirable in the AFS controller, as it leads to interference with the MLA and GLA control objectives, while both actuator performance limits and signal rate limits for this test make control effort above 50 Hz inadvisable and above 100 Hz impossible. The lower threshold for the pass-band of this filter was shifted down from the 2 Hz boundary between the lower and center frequency bands to 1 Hz to reduce roll-off within the flutter band. Maximizing control effort in the flutter band while minimizing the effort everywhere else was the primary driver of the design of the control cost filters. Additionally, bench-top testing found that the mechanical linkage between the hydraulic actuator and the surface of aileron B significantly reduced the available bandwidth compared to ailerons A and C, which is accounted for by increasing the cost of use by a factor of 5 after the filter shaping is otherwise finished. The maximum allowable control surface deflection magnitude at the AFS test condition is 10 deg, and the control costs are designed so that the resulting controller respects this limit within the flutter band and has smaller response elsewhere. Evaluating the magnitude of the costs in the low and flutter frequency bands indicates that a third-order biproper low-pass filter is needed to achieve the desired magnitude difference within a half-decade. A biproper high-pass filter of the same order is then designed with the first corner frequency chosen based on the response frequency limits of each of the ailerons, and the second corner frequency is chosen to ensure the high-frequency band reaches an arbitrarily large, but finite, cost. The Bode plots of each of the control cost filters are shown in Fig. 8c, with the cost for aileron A in blue, aileron B in red, and aileron C in yellow.

IV. Robust Control Synthesis

As discussed throughout this work, robustness against measurement noise and uncertainty in the dynamics of the system is a key requirement for an AFS control law to be functional. Robust control synthesis is used in this work to design a control law that satisfies this requirement. The control synthesis was performed using μ -synthesis, also referred to as DK -iteration [12, 13, 22]. While tools exist to solve the μ -synthesis problem [23], the set-up process is nontrivial; this section is focused on explaining the design choices that were made to facilitate the synthesis or driven by requirements of the μ -synthesis process. Finally, μ -synthesis is a non-convex problem and is not guaranteed to find a global optimum. Depending on the starting condition, it is possible to recover multiple unique controllers when running the algorithm on the same uncertain generalized plant.

The remainder of this section is focused on specific aspects of this synthesis framework. Section IV.A presents a brief overview of the μ -synthesis algorithm and framework. Section IV.B discusses the sources of uncertainty in the dynamics that are being considered in this work and how the dynamic uncertainty is included in the synthesis process. Section IV.C presents an overview of what sensor combinations were examined and what factors drove the sensor selection process.

A. Overview of μ -Synthesis

With the large number of inputs and outputs in the generalized plant and the dynamic uncertainty in the system, the specific formulation used for modeling the uncertainty has a large influence on the controller found by μ -synthesis. Testing early in the development of this work found that input multiplicative uncertainty (IMU) framework yielded a tighter bound on the uncertainty with lower order filters than additive or output multiplicative frameworks. Dropping the (s) arguments for cleanliness of notation, the IMU framework is given by $\mathbf{G}_\Delta = \mathbf{G}(\mathbf{1} + \mathbf{W}_1\Delta\mathbf{W}_2)$, where $\mathbf{G}_\Delta$ is the uncertain plant, Δ is a bounded-input bounded-output (BIBO) stable uncertainty block with gain less than one, $\mathbf{W}_1$ and $\mathbf{W}_2$ are shaping filters for the uncertainty, and $\mathbf{G}$ is the nominal plant. As is standard, the Δ -block is assumed to be stable with an $\mathcal{H}_\infty$ -norm of 1 or less. The IMU framework is shown in Fig. 9a. By including the $\mathbf{W}_1$ and $\mathbf{W}_2$ weights and the corresponding inputs and outputs in the generalized plant, the uncertain system can be rearranged in the form of Fig. 9b. This is the standard linear fractional transformation (LFT) framework used to evaluate the robust stability and performance properties of the system.

The red dashed line in Fig. 9b indicates the combined plant-controller loop, which is labeled as $\mathbf{M}$. A controller is said to robustly stabilize the system if the $\mathbf{M} - \Delta$ interconnection is BIBO stable for all uncertainty in the set Δ . By using the assumption that the Δ -block is BIBO stable and has gain less than one, the small-gain theorem can be leveraged to prove that the $\mathbf{M} - \Delta$ loop is stable for any uncertainty included in the uncertainty set, provided that the $\mathcal{H}_\infty$ -norm of $\mathbf{M}$ from the input $\mathbf{u}_\Delta$ to the output $\mathbf{y}_\Delta$ is less than or equal to one. A controller is said to meet robust performance specifications if for any uncertainty in the set Δ the closed-loop $\mathcal{H}_\infty$ from $\mathbf{w}$ to $\mathbf{z}$ remains below the target value of one. This condition is met if and only if the structured singular value, μ , of $\mathbf{M}(j\omega)$ is less than one for all $\omega \in \mathbb{R}$. Further technical details on these robust stability and robust performance conditions, as well as the structured singular value, are found in [22, 24].




Fig. 9 Uncertainty modeling framework.

The fundamental goal of μ -synthesis is to design a controller that ensures robust stability and robust performance for all of the uncertainties contained in the Δ -block. This is achieved by solving for a controller that drives the structured singular value of $\mathbf{M}(j\omega)$ below one at all frequencies. Due to the challenges in computing the structured singular value, μ -synthesis is a non-convex control synthesis problem that is typically solved by DK -iteration, which involves iterating between solving for an $\mathcal{H}_\infty$ -optimal controller that drives down the structured singular value and solving for D -scalings that improve the computed upper bound on the structured singular value. This process proceeds until the structured singular value of $\mathbf{M}(j\omega)$ is below one at all frequencies or no further progress is made in reducing the peak value of the structured singular value. Further details on μ -synthesis and DK -iteration are found in Refs. [12, 13, 22].

B. Uncertainty Modeling

Up to this point, there has not been significant discussion on what actually contributes to the uncertainty that the Δ -block represents. The three primary uncertainty sources considered in this work are the aerodynamic modeling, the model reduction required to run μ -synthesis, and operating between conditions with a numeric model. The approach used in this work is to use unstructured input multiplicative uncertainty to overbound an uncertainty set constructed from the full actuator-wing models of the condition of interest and the 8 surrounding test conditions shown in Fig. 4. The aerodynamics introduce uncertainty to the system as they are based on models developed in ZAERO™ that have only had limited validation from the first round of wind-tunnel testing during 2024. Several avenues were explored to attempt to bound the uncertainty introduced through the aerodynamic modeling. An early method tested in the development of this work was to use the differences in the models from the previous aerodynamic model iterations to quantify how much uncertainty there was with different aspects of the numeric model. This approach initially showed promise for uncertainty modeling; however, it was abandoned as the aerodynamic model iterations were strongly correlated with each other. This made it impossible to determine if the differences were actually indicative of meaningful uncertainty in the aerodynamics or simply changes in the panel set-up and modeling. While this method was not ultimately used in this work, it is included as an object lesson for anyone attempting μ -synthesis control design for aerodynamic systems. Applying μ -synthesis to large systems is not trivial on typical desktop or laptop computers, as the process involves repeatedly solving a $\mathcal{H}_\infty$ -control synthesis problem and performing a robustness analysis to estimate the robust performance and find dynamic scalings. This process is repeated until a stopping condition is met, typically either the robust performance stops improving or a set number of iterations is reached. As the workstation laptop available for this work cannot reliably run the $\mathcal{H}_\infty$ synthesis in MATLAB on the original generalized plant with approximately 600 states, model reduction is mandatory to enable the more demanding μ -synthesis computations. Several approaches were examined, with the most effective solution ultimately found to be to excise the aerodynamic elements from the aero-structural model, thus directly reducing the actuator-wing model from 493 states after the Padé approximation to remove time delays down to 161 states. This reduction brought the generalized plant down to a number of states that both $\mathcal{H}_\infty$ synthesis and μ -synthesis were able to run reliably in MATLAB. As the aerodynamic effects were one of the primary uncertainty sources to begin with, this approach was found to be an efficient way to sufficiently reduce the system dimension without adding to the uncertainty associated with the aerodynamic effects.

The aerodynamic and model reduction uncertainties, along with the fact that the TDT conditions will not always be aligned with one of the modeled conditions shown in Fig. 4, led to the design of the uncertainty set. Specifically, since the testing is occurring in subsonic conditions and the models form a relatively dense grid, the aerodynamics are expected to evolve smoothly between model conditions. Additionally, the changes in aerodynamic response between neighboring grid points form a crude sensitivity approximation, with the larger changes indicating increased sensitivity and associated uncertainty in the response.

Ultimately the uncertainty set is constructed through the use of the `ucover` function in MATLAB, using the reduced actuator-wing model of the target condition as the nominal model, and the set of 9 full actuator-wing models from the surrounding model conditions as the sampled uncertainty set to be covered. As the problem formulation is using unstructured dynamic uncertainty, `ucover` generates filters that scale the Δ -block to over-bound the uncertainty set as tightly as possible with a stable, minimum-phase transfer function. IMU leads to acceptably tight overbounding of the uncertainty set using second order weights for both the $\mathbf{W}_1$ and $\mathbf{W}_2$ blocks as shown in Fig. 9a. This leads to both weight blocks having 15 inputs and outputs and 30 states. The coverage obtained from this process is shown in Fig. 10, where singular value plots of the uncertainty set are shown in blue, and the overbound achieved by the worst case Δ with unity gain in $\mathbf{G}(\mathbf{I} + \mathbf{W}_1\Delta\mathbf{W}_2)$ is shown in red. The uncertain generalized plant arising from this formulation has 302 states, which is small enough to use with the `musyn` function.




Fig. 10 Singular value plot of the uncertainty set (blue) and over bound weight used for μ -synthesis (red) at the target condition for AFS.

C. Sensor Selection

The sensor selection has a significant impact on the AFS controller arising from the μ -synthesis algorithm. Several combinations of sensors were examined as part of the control design process. The number of sensors determine the dimension of the input to the controller, while which specific sensors are selected determines what information is available. Control design through μ -synthesis is known to produce controllers with large state dimension, while the input dimension directly scales the size of two of the matrices of the resulting controller.

Initial control designs using all available sensor information leads to controllers with approximately 150 input channels. This approach is found to be numerically intractable in synthesizing controllers, with small changes in the performance weights resulting in controllers that are hyper-responsive or inert, with no apparent middle ground. Controllers are instead designed using both all the accelerometers and all the z-axis accelerometers. This reduces the input dimension to 22 and 14 respectively; however, the resulting controllers between the two sets of sensors have nearly identical performance. This leads to the question of how many of the z-axis accelerometers are actually needed? Testing indicates that the 8 accelerometers used in the ILAF control scheme from Ref. [20], and labeled in Fig. 2, work reliably for use with μ -synthesis. Analysis and results presented in Sections V and VI are all based on the 8 accelerometer control design.

V. Analytic Control Evaluation

The controller arising from the μ -synthesis process must be evaluated for stability properties before it can be considered for experimental validation on the test hardware in the TDT. The primary considerations at this stage are determining if the size of the controller is amenable to running in real time on the TDT systems, and whether the system is robustly stable with the controller in the loop. A singular value plot of the controller is provided in Fig. 11.

The control law resulting from the μ -synthesis described in this work has 8 input channels corresponding with the 8 z-axis accelerometers (see Section IV.C) and 3 output channels as deflection commands for the hydraulically-driven ailerons (see Section III.C). The controller also has 332 states, which is within the capabilities of the TDT hardware to run in real time.

The best achieved robust $\mathcal{H}_\infty$ performance achieved by this controller is 29.7. That indicates that the closed-loop

gain of the M block shown in Fig. 9b remains below 29.7 for up to $1/29.7$ of the uncertainty included in the synthesis. As this value is larger than 1, this result indicates that the controller cannot be certified for use directly from the robust performance metric, and other stability measures must be considered. The multi-loop disk margin with independent variations in all input and output channels is one such metric and has previously been used in AFS stability assessment [25]. Assessing the disk margin of the feedback loop of the controller and system at the design condition with the DAS included yielded a disk margin of 0.0014. While this is a smaller than desired margin, it is still sufficient to indicate the controller's ability to stabilize the system if the models of the actuator response and signal delays are sufficiently accurate. These analytic results are validated through simulation in Section VI.

VI. Simulation Validation

Simulations are performed to validate the stability and assess the performance of controllers using a MATLAB test environment. All presented simulations feature a doublet input sequence on the ailerons. Additionally, the presented simulations all include a 4th-order Butterworth low-pass filter with a 400 Hz break frequency acting on the measurement data provided to the controller.

The doublet sequence is used as a standard check for all AFS controllers on the IAWTM system and provides a way to compare the performance of multiple controllers in a meaningful way on a level playing field. The doublets are all generated as a high pulse for 0.08 s followed by a low pulse for the same duration, with the pulses having ± 2 degree magnitude. This signal is subsequently filtered through a low-pass filter described by $F(s) = 50/(s + 50)$, resulting in the control signals shown in Fig. 12. The doublets are staggered with the beginning of each pulse occurring at times of 1, 3, and 5 seconds in the simulations and applied to ailerons A, B, and C sequentially.

Every synthesized controller is initially tested in simulation with the aileron doublet commands, with no disturbance or noise acting on the system. This provides a simple test to validate that the controller meets the minimum performance requirements and limits the wing-tip accelerations to below 10g magnitude without saturating the ailerons. This simulation is shown for the reduced actuator-wing model in Fig. 13a and the full actuator-wing model in Fig. 13b. In both simulations, the excitation from each of the doublet inputs is suppressed such that the wing-tip acceleration is below 0.1g magnitude before the next doublet occurs. All simulation results are presented in the same format, with the top subplot showing the vertical wing-tip acceleration at the a303z and a111z sensor locations (sensor locations shown on the planform in Fig. 2). The a303z location is recorded in blue, while the a111z location is in red. The middle subplot shows the position of each of the ailerons measured in degrees throughout the duration of the simulation, while




Fig. 11 Singular value plot of the developed robust control law.

the bottom subplot shows the control signals sent to the ailerons. The position and control signal plots use a consistent color scheme, with ailerons A, B, and C, represented by the blue, red, and yellow lines respectively.

Subsequent simulation introduces disturbances acting through random deflections of the MPFs and measurement noise. The disturbances are generated as zero-mean unit-variance white noise sampled at 10 kHz, and passed through the same low-pass filter as was used in the doublet signal generation to produce frequency-weighted noise. The colored noise is then scaled to have a standard deviation of 2 degrees over the duration of the 8 second simulation. The results of simulating the controller with the full actuator-wing model in this disturbed environment are shown in Fig. 14a. The energy introduced by the MPF excitement is larger than the aileron doublet sequence, and leads to the controller using nearly the full range of allowable aileron deflection to keep the system stable.

The noise is generated by sampling a zero-mean, 10^{-3} g standard deviation normal distribution at 10 kHz. A zero-order hold is then applied to the generated noise, which is added to the true acceleration values at the accelerometer locations in the simulation. The corrupted signals are then passed through the Butterworth filter acting as an approximation of the DAS system before being passed onto the controller. A simulation with noise and no disturbance is shown in Fig. 14b. The control signals and aileron deflections occurring in the first second of the simulation are a result of the controller responding to the measurement noise. The system remaining stable with the added noise is a positive result, as the generated noise is the maximum amount the μ -synthesis design expects to encounter based on the generalized plant weight W_n (see Section III.F).

The final case for simulation is both disturbance and noise simultaneously. This is presented in Fig. 15. Despite the adverse condition of both large energy disturbances and sensor noise, the controller maintains safe wing-tip accelerations without exceeding the 10 degree magnitude limits on the ailerons.

VII. Conclusion

An AFS control scheme was developed through μ -synthesis. The synthesis process was tuned so that the resulting controller would maintain safe wing-tip accelerations without saturating the ailerons. An uncertainty set was designed for control synthesis using the aero-structural models from proximal test conditions to overbound the uncertainties from the aerodynamic models, model reduction, and operation at conditions between model design points. The resulting controller has smaller analytic margins than desired; however, it suppresses flutter effectively at the Mach 0.85 and dynamic pressure of 230 lbf/ft² design condition in simulation.

Developing methods and problem formulations to increase the robustness margin of the controllers arising from the μ -synthesis process is an important next step. While these margins are currently smaller than desirable, this is not unique to this AFS design scheme, rather every AFS scheme under development has encountered similar challenges to varying degrees. The IAWTM model has been described on at least one occasion as “the most flexible wing ever put in the TDT,” which highlights both the difficulty of the problem and the need for further research to further mature this technology.




Fig. 12 Doublet commands for the hydraulic ailerons used in all presented simulations.




Fig. 13 Closed-loop simulation responses to the aileron doublet commands. No disturbance or sensor noise was included in these simulations.




Fig. 14 Closed-loop simulation responses to the aileron doublet commands. Figure 14a is the response to normally distributed random inputs on the 7 MPFs with standard deviation of 2 degrees, but no sensor noise. Figure 14b is the response with no disturbance, but normally distributed white noise with a standard deviation of 0.001g as described in Section III.D.

Acknowledgments

This work was supported by a NASA Space Technology Graduate Research Opportunity. The NASA IAWTM sub-project was supported by the NASA Advanced Air Transport Technology (AATT) project. The contributions by the IAWTM team members at NASA Ames Research Center (ARC), NASA Armstrong Flight Research Center (AFRC), NASA Langley Research Center (LaRC), and Boeing Research and Technology are gratefully acknowledged.




Fig. 15 Closed-loop simulation responses to the aileron doublet commands with both disturbance from the MPFs and sensor noise.

References

- [1] Gipson, L., “Advanced Air Transport Technology (AATT) Project Technical Challenges,” <https://www.nasa.gov/aeronautics/advanced-air-transport-technology-aatt-project-technical-challenges/>, Apr. 2020. Accessed Jan. 2024.
- [2] Heaney, P. S., and Quindlen, J. F., “Overview of the Integrated Adaptive Wing Technology Maturation Wind-Tunnel Test Objectives,” *AIAA SciTech Forum*, 2024, p. 2509. <https://doi.org/10.2514/6.2024-2509>.
- [3] Brooks, T. R., Kenway, G. K. W., and Martins, J. R. R. A., “Benchmark Aerostructural Models for the Study of Transonic Aircraft Wings,” *AIAA Journal*, Vol. 56, No. 7, 2018, pp. 2840–2855. <https://doi.org/10.2514/1.J056603>.
- [4] Cole, S. R., Noll, T. E., and Perry, B., “Transonic Dynamics Tunnel Aeroelastic Testing in Support of Aircraft Development,” *Journal of Aircraft*, Vol. 40, No. 5, 2003, pp. 820–831. <https://doi.org/10.2514/2.6873>.
- [5] Ivanco, T. G., “Unique Testing Capabilities of the NASA Langley Transonic Dynamics Tunnel, an Exercise in Aeroelastic Scaling,” *AIAA Ground Testing Conference*, 2013, p. 2625. <https://doi.org/10.2514/6.2013-2625>.
- [6] Livne, E., “Aircraft Active Flutter Suppression: State of the Art and Technology Maturation Needs,” *Journal of Aircraft*, Vol. 55, No. 1, 2018, pp. 410–452. <https://doi.org/10.2514/1.C034442>.
- [7] Schmidt, D., and Chen, T., “Frequency Domain Synthesis of a Robust Flutter Suppression Control Law,” *Journal of Guidance, Control, and Dynamics*, Vol. 9, No. 3, 1986, pp. 346–351. <https://doi.org/10.2514/3.20112>.
- [8] Adams, W. M., Jr., Christhilf, D. M., Waszak, M. R., Mukhopadhyay, V., and Srinathkumar, S., “Design, Test, and Evaluation of Three Active Flutter Suppression Controllers,” Tech. rep., NASA Langley Research Center, October 1992.
- [9] Adams, W. M., Jr., and Christhilf, D. M., “Design and Multifunction Tests of a Frequency Domain-Based Active Flutter Suppression System,” *Journal of Aircraft*, Vol. 32, No. 1, 1995, pp. 52–60. <https://doi.org/10.2514/3.46683>.
- [10] Mukhopadhyay, V., “Flutter Suppression Control Law Design and Testing for the Active Flexible Wing,” *Journal of Aircraft*, Vol. 32, No. 1, 1995, pp. 45–51. <https://doi.org/10.2514/3.46682>.
- [11] Waszak, M. R., and Srinathkumar, S., “Flutter Suppression for the Active Flexible Wing-A Classical Design,” *Journal of Aircraft*, Vol. 32, No. 1, 1995, pp. 61–67. <https://doi.org/10.2514/3.46684>.
- [12] Doyle, J. C., “Structured Uncertainty in Control System Design,” *IEEE Conference on Decision and Control*, Fort Lauderdale, FL, USA, 1985, pp. 260–265. <https://doi.org/10.1109/CDC.1985.268842>.
- [13] Dullerud, G. E., and Paganini, F., *A Course in Robust Control Theory: A Convex Approach*, Texts in Applied Mathematics, Vol. 36, Springer Science & Business Media, 2013. <https://doi.org/10.1007/978-1-4757-3290-0>.

- [14] Pfifer, H., Moreno, C. P., Theis, J., Kotikalpudi, A., Gupta, A., Takarics, B., and Seiler, P., “Linear Parameter Varying Techniques Applied to Aeroservoelastic Aircraft: In Memory of Gary Balas,” *IFAC-PapersOnLine*, Vol. 48, No. 26, 2015, pp. 103–108. <https://doi.org/10.1016/j.ifacol.2015.11.121>.
- [15] Kotikalpudi, A., Pfifer, H., and Seiler, P., “Sensitivity of Robust Flutter Boundary to Model Uncertainties in Aeroservoelastic Systems-Invited,” *AIAA Atmospheric Flight Mechanics Conference*, 2016, p. 1752. <https://doi.org/10.2514/6.2016-1752>.
- [16] Hjartarson, A., Seiler, P., and Balas, G., “LPV Aeroservoelastic Control Using the LPVTools Toolbox,” *Proceedings of the AIAA Atmospheric Flight Mechanics Conference*, 2013. <https://doi.org/10.2514/6.2013-4742>.
- [17] Stalla, F., Looye, G., Kier, T. M., Michel, K., Schmidt, T., Hanke, C., Kirmse, T., and Pusch, M., “Wind Tunnel Assessment of a Robust Gust Load Alleviation Controller Designed Using μ -Synthesis,” *AIAA SciTech Forum*, 2025. <https://doi.org/10.2514/6.2025-0904>.
- [18] Suh, P. M., Chin, A. W., and Mavris, D. N., “Virtual Deformation Control of the X-56A Model with Simulated Fiber Optic Sensors,” *AIAA Atmospheric Flight Mechanics Conference*, American Institute of Aeronautics and Astronautics, Boston, MA, USA, 2013. <https://doi.org/10.2514/6.2013-4844>.
- [19] Theis, J., Pfifer, H., and Seiler, P., “Robust Modal Damping Control for Active Flutter Suppression,” *Journal of Guidance, Control, and Dynamics*, Vol. 43, No. 6, 2020, pp. 1056–1068. <https://doi.org/10.2514/1.G004846>.
- [20] Grauer, J. A., and Waite, J., “Design of a Collocation-Based Active Flutter Suppression Control Law for the IAWTM Wind Tunnel Model,” *AIAA SciTech Forum*, 2024, p. 2201. <https://doi.org/10.2514/6.2024-2201>.
- [21] Zona Technology Inc., *ZAERO v. 9.3 Theoretical Manual*, January 2019. URL <http://www.zonatech.com>.
- [22] Skogestad, S., and Postlethwaite, I., *Multivariable Feedback Control: Analysis and Design*, John Wiley & Sons, 2005.
- [23] Balas, G., Doyle, J. C., Glover, K., Packard, A., and Smith, R., *μ -Analysis and Synthesis Toolbox*, MUSYN Inc. and The MathWorks, Natick, MA, USA, 1993.
- [24] Zhou, K., Doyle, J. C., and Glover, K., *Robust and Optimal Control*, Prentice-Hall, Upper Saddle River, NJ, 1996.
- [25] Theis, J., Pfifer, H., and Seiler, P., “Robust Control Design for Active Flutter Suppression,” *AIAA Atmospheric Flight Mechanics Conference*, American Institute of Aeronautics and Astronautics, 2016. <https://doi.org/10.2514/6.2016-1751>.