Design and Testing of Drag Minimization and Maneuver Load Alleviation Control Laws for the IAWTM Wind Tunnel Test

Christopher J. Forte*
KBR Wyle, Inc., Moffett Field, CA 94035

Nhan T. Nguyen†
NASA Ames Research Center, Moffett Field, CA 94035

John F. Quindlen‡
Boeing Technology Innovation, Huntington Beach, CA, 92647

Heather S. Hussain§
Boeing Technology Innovation, Charlotte, NC, 28202

This paper presents the design and experimental testing of drag minimization and maneuver load alleviation (MLA) controllers for high aspect ratio wing wind tunnel model. The controllers are developed for the Integrated Adaptive Wing Technology Maturation (IAWTM) wind tunnel tests at Langley Research Center Transonic Dynamics Tunnel. The IAWTM wind tunnel tests are a joint effort between NASA and Boeing to mature technologies associated with flexible high aspect ratio transport aircraft wings. Two drag minimization controllers are developed, including a peak-seeking control law that performs model-free optimization to minimize drag. A second drag minimization controller is developed that performs a real-time model identification and optimizes the control surface profile to minimize drag. Similarly, two MLA controllers are developed, including a feedforward MLA design that is designed to utilize distributed control surfaces to shift the load inboard to minimize wing root bending moment. The second MLA controller utilizes an onboard surrogate model to provide optimal control surface and angle of attack commands to minimize wing root bending moment. The design and experimental results for the wind tunnel testing of the drag minimization and MLA controllers are summarized.

I. Introduction

The aircraft industry has long pursued improved energy efficiency by designing airframes that are both aerodynamically and structurally optimized. Manufacturers achieve this goal through energy-efficient engines and lightweight advanced composites, as reducing operational weight is a primary driver of efficiency. While these materials maintain load-carrying capacity, they often provide less structural rigidity, increasing structural flexibility and making aircraft more susceptible to aeroelastic effects such as gusts and maneuver-induced loads. Higher wing aspect ratios can exacerbate these issues, and significant wing twist under load may cause the optimal lift distribution for minimal drag to vary across the flight envelope.

To compensate for these aeroelastic effects, cruise drag minimization control is required. Currently, this is typically implemented using table-lookup methods based on validated analytical models for specific aircraft geometries, which must be verified through wind tunnel and flight tests. However, variations in aircraft production and wide-ranging operating conditions, including gross weight, airspeed, and altitude, can lead to significant differences in aircraft performance. Advanced control strategies can provide greater flexibility than these traditional approaches. Similarly, maneuver load alleviation (MLA) control can reduce wing loading during maneuvers, further enhancing structural performance and safety. Drag minimization and MLA control laws are an area of active research with several proposed control solutions[1–5]. These developments motivate the control strategies formulations presented in the following sections.

*Aerospace Engineer, Intelligent System Division, Member AIAA.
†Senior Research Scientist, Technical Group Lead, Intelligent Systems Division, Associate Fellow AIAA.
‡GNC Engineer, Guidance, Navigation, Control & Autonomy, Member AIAA.
§Associate Technical Fellow, Guidance, Navigation, Control & Autonomy, Member AIAA.

The Integrated Adaptive Wing Technology Maturation effort is a collaborative program between NASA and Boeing aimed at advancing technologies for high-aspect-ratio wing design. Key technologies include control strategies that mitigate the structural and aeroelastic challenges associated with lightweight, high-aspect-ratio wings, enabling both efficient and safe operation. Control objectives encompass drag minimization, maneuver load alleviation, gust load alleviation, and active flutter suppression. To support these objectives, a wind tunnel test article was developed based on the 13.5-aspect-ratio Common Research Model (CRM) [6–8], and testing was conducted in the Transonic Dynamics Tunnel (TDT) at NASA Langley Research Center to validate the performance of the proposed control strategies. The wind tunnel test article installed in the TDT test section is shown in Fig. 1.

Under IAWTM, Boeing and NASA developed drag minimization and MLA control laws that were tested during the second wind-tunnel entry in 2025. One drag minimization control law is a peak-seeking controller that performs model-free optimization to reduce drag. Another drag minimization controller performs real-time model identification and optimizes the control surface profile accordingly to minimize drag.

Two MLA controllers were also developed. The first is a feedforward MLA design that uses distributed control surfaces to shift aerodynamic loads inboard, thereby reducing wing-root bending moment. The second uses an online estimation strategy to generate onboard surrogate model to generate optimal control solution. This paper provides details on these control system designs and summarizes the wind tunnel test results.

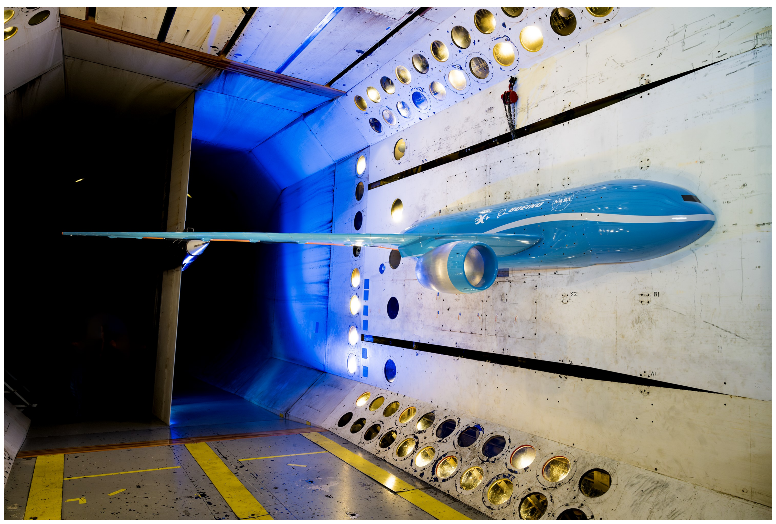

Fig. 1 IAWTM Test Article in TDT Test Section

II. Wind Tunnel Test Description

This section is an overview of the NASA Transonic Dynamics Tunnel within the scope of the IAWTM drag minimization and Maneuver load alleviation tests. Additional details about the full range of tests and IAWTM program can be found in [7, 9].

A. Test Article

The IAWTM test article, shown in the detailed diagram in Fig. 2, is a 10.8%-scale, aeroelastically scaled, half-span, wall-mounted representation of the 13.5-aspect-ratio CRM. The wind-tunnel model incorporates a dynamically scaled wing, whereas the fuselage length is shortened to satisfy the geometric constraints of the test section. A flow-through nacelle is included to capture the effects of nacelle mass and inertia on the wing's structural response. Overall, the model emulates the structural and aerodynamic characteristics of a high-aspect-ratio transport wing within the constraints of wind-tunnel scaling.

The IAWTM model in Fig. 2 has 10 control surfaces. Three control surfaces are high speed hydraulically actuated ailerons designated AILA-AILC and are intended for gust load alleviation and active flutter suppression controls. The 7 remaining control surfaces are slower acting electric actuators, called mini plain flaps (MPFs), and are designated MPFA-MPFG. The MPFs and hydraulic ailerons are used in conjunction for the drag minimization and maneuver load alleviation controls. This high degree of surface segmentation enables finer control of the spanwise lift distribution. Control surface deflections are limited to $\pm 10°$ under standard conditions and reduced to $\pm 5°$ at elevated tunnel dynamic pressures. The MPFC was not utilized during the closed-loop tests due to hardware malfunctions.

The model is instrumented with ten full-bridge strain gauges and sixteen accelerometers to measure structural responses. The accelerometer and strain-gauge locations are shown in Fig. 3. The model is mounted via a mounting-box adapter to the TDT electric turntable (ETT) and a five-component balance (NASA balance 1637S), providing force and moment measurements and enabling the determination of lift and drag. Additional details of the test article are provided by Heaney[7].

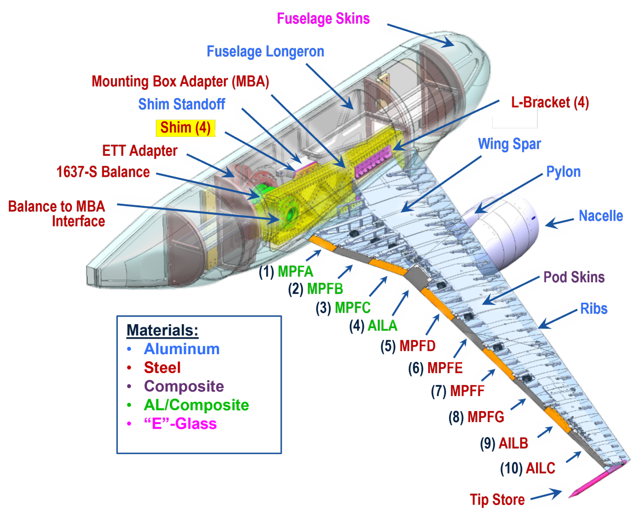

Fig. 2 IAWTM Test Article

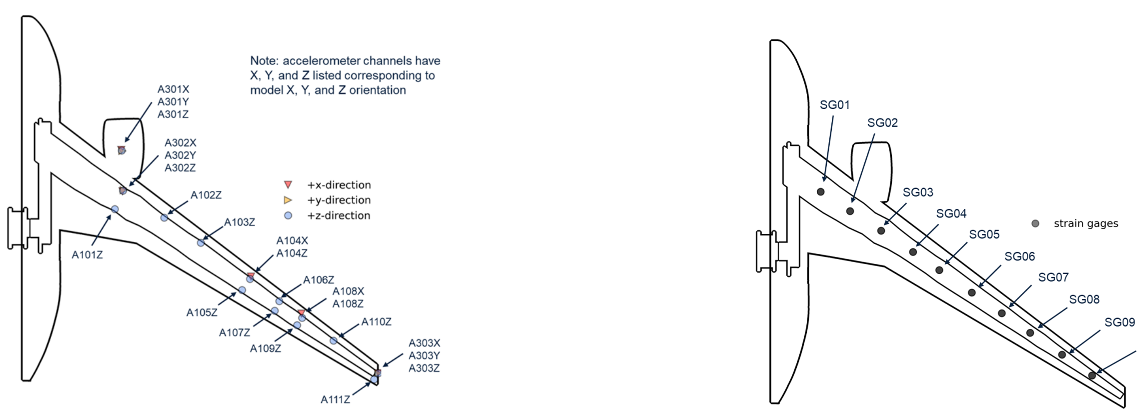

(a) IAWTM Test Article Accelerometer Locations **(b) IAWTM Test Article Strain Gage Locations**

Fig. 3 Sensor locations on the IAWTM wind-tunnel model.

B. Aerodynamic Modeling

The aerodynamic models used to support controller development employ high-fidelity CFD simulations performed with the FUN3D solver[10–13]. FUN3D solves the steady-state Reynolds-averaged Navier-Stokes (RANS) equations and is coupled to a structural model represented by 24 modes extracted from a GVT-validated NASTRAN^{TM} finite-element model. This high-fidelity aeroelastic framework enables accurate evaluation of control surface sensitivities and aerodynamic derivatives required for drag minimization and MLA controller design. Databases are further refined with wind-tunnel measurements from the 2024 and 2025 wind tunnel entries, providing accurate aerodynamic and control surface data for controller development.

C. Transonic Dynamics Tunnel

The TDT, shown in Fig. 4, is used to conduct the wind-tunnel experiments. The facility operates in both air and heavy gas (R134a), which allows accurate matching of the full-scale vehicle's aeroelastic response. Hydraulic power is supplied to the high-speed, hydraulically actuated ailerons, and the tunnel's airstream oscillation system (AOS) provides controlled gust excitations for gust load alleviation testing. Bypass valves are installed to rapidly unload the model by reducing the dynamic pressure in the test section in the event of instability.

For these control experiments, lift and drag are obtained from the five-component balance measurements, while the ETT maintains the model pitch angle. Both of these components are shown in Fig. 2. In addition, the strain gage located nearest the wing root serves as a surrogate measurement of the wing-root bending moment. These facility capabilities provide the experimental environment in which the drag minimization and MLA controllers are evaluated.

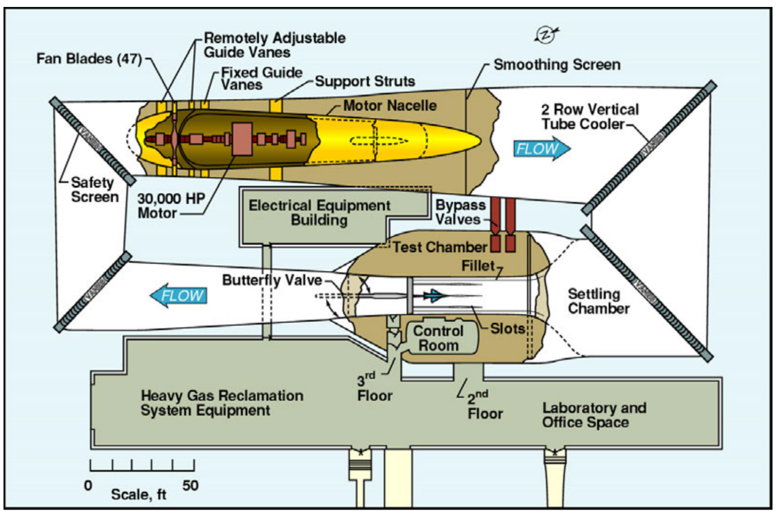

Fig. 4 TDT Overview

D. Controller Implementation Considerations

Lift and drag measurements are subject to considerable noise because of the model's high flexibility and its sensitivity to tunnel turbulence. Figure 5 shows the input to aileron A during the open-loop aerodynamic characterization conducted prior to control law testing, and Fig. 6 presents the corresponding lift and drag coefficients computed from the wind-tunnel balance data. While the lift measurements exhibit a clear trend, the drag data contain a substantial amount of noise. Signal quality is improved by applying low-pass filtering with a low cutoff frequency to remove high-frequency structural content, and by using long data-acquisition periods with time averaging to further reduce noise and increase measurement accuracy. These mitigation efforts are summarized in more detail in the following control design sections.

Another significant design consideration is that the model pitch is controlled manually by a wind-tunnel operator rather than by the flight-control system, so any controller that relies on angle-of-attack (AOA) measurements or commands must account for this constraint. To accommodate this, the controllers are designed with sufficiently low rate commands for the operator to track reliably, or included logic to hold until the commanded set point is reached. Additional considerations included flap freeplay and potential inaccuracies in flap-position tracking, both of which introduce uncertainty into the measured surface positions.

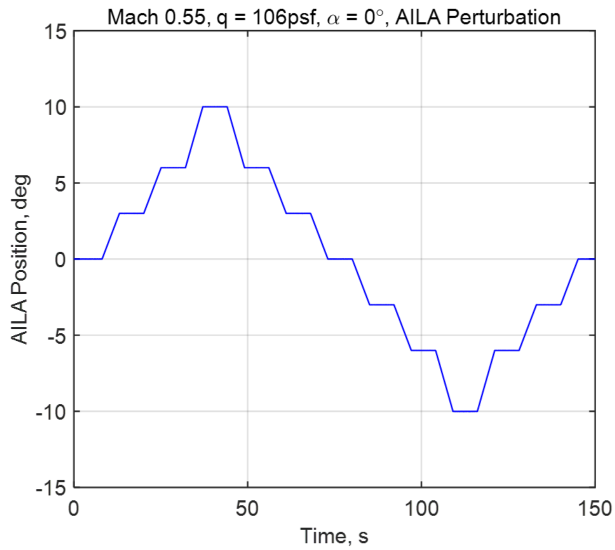

Fig. 5 Aileron A Pitch-Pause Control History

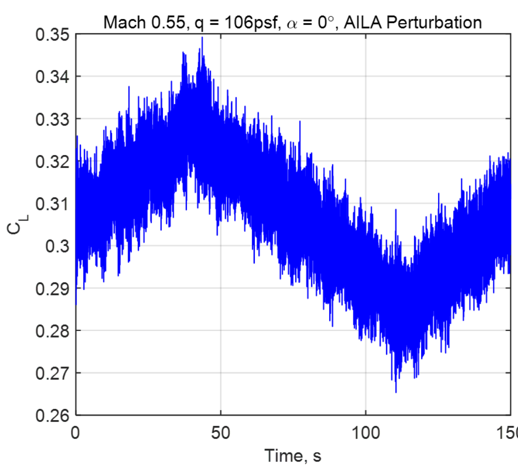

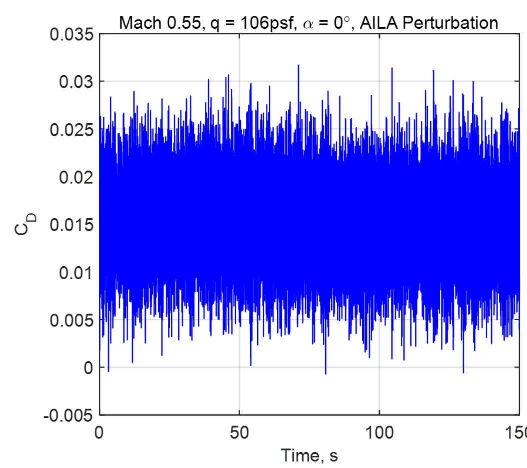

Fig. 6 Lift and drag coefficients.

III. Control Law Design

Boeing and NASA Ames each developed drag minimization and maneuver load alleviation (MLA) control strategies for the second IAWTM wind-tunnel entry conducted in June-July 2025. These control laws were integrated with NASA Langley's TDT dSPACE hardware-in-the-loop control system computer for real-time implementation. The drag minimization tests focused on reducing drag while maintaining a prescribed lift target, whereas the MLA tests sought to alleviate structural loads during a maneuver. This section summarizes the control strategies employed in these experiments.

A. NASA Ames Real-Time Drag Optimization and Maneuver Load Alleviation

A drag minimization and maneuver load alleviation controller are designed using a real-time system identifcation approach to generate an onboard surrogate model. This surrogate model is then employed to compute the optimal control surface schedule and angle of attack for two control objectives: lift-constrained drag minimization and maneuver load alleviation.

The real-time drag-optimization control framework, illustrated in Fig. 7, consists of four stages. First, the model is excited with an angle-of-attack perturbation, followed by control surface perturbations. The resulting responses are processed using a recursive least-squares algorithm to estimate the sensitivities of lift, drag, and wing-root bending moment, thereby forming the onboard surrogate model. This model is expressed in Eqs. 1-3.

$$C_L = C_{L_0} + C_{L_\alpha}\alpha + C_{L^2_\alpha}\alpha^2 + C_{L_\delta}\delta + \delta^T C_{L^2_\delta}\delta \tag{1}$$

$$C_D = C_{D_0} + C_{D_\alpha}\alpha + C_{D^2_\alpha}\alpha^2 + C_{D_\delta}\delta + \delta^T C_{D^2_\delta}\delta \tag{2}$$

$$M_y = M_{y_0} + M_{y_\alpha}\alpha + M_{y^2_\alpha}\alpha^2 + M_{y_\delta}\delta + \delta^T M_{y^2_\delta}\delta \tag{3}$$

Then an lift-constrained optimization is performed, computing optimal control surface profile and angle of attack for a target lift. Lastly, this optimal profile is sent to the flight controller and angle of attack operator for actuation.

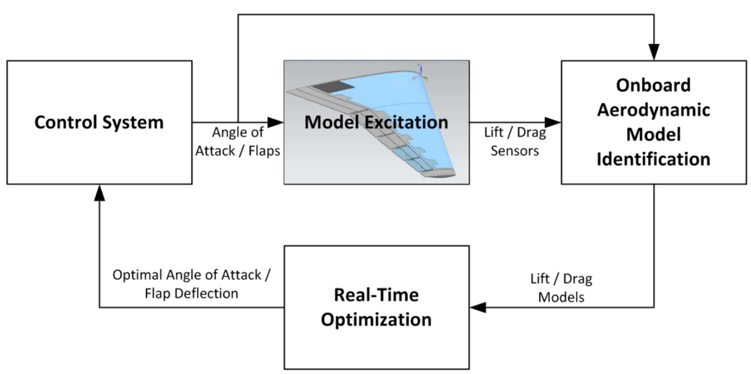

Fig. 7 Real-Time Drag Optimization Framework

1. *Online Parameter Estimation*

The model angle of attack is first perturbed by $\pm\Delta\alpha$, where $\Delta\alpha = 0.5°$. Notably during this stage, the control surface positions are held at constant values $\bar{\delta}$. A linear regression is performed by minimizing the error squared:

$$\min\ J = \sum_{i=1}^{3n}\left(\begin{bmatrix}\Delta C_L\\ \Delta C_D\\ \Delta M_y\end{bmatrix}_i - \begin{bmatrix}C_{L_0}\\ C_{D_0}\\ M_{y_0}\end{bmatrix} - \begin{bmatrix}C_{L_\alpha}\\ C_{D_\alpha}\\ M_{y_\alpha}\end{bmatrix}\alpha_i - \begin{bmatrix}C_{L_{\alpha^2}}\\ C_{D_{\alpha^2}}\\ M_{y_{\alpha^2}}\end{bmatrix}\alpha_i^2\right)^2 \tag{4}$$

where n is the number of iterations and

$$\begin{bmatrix}\Delta C_L\\ \Delta C_D\\ \Delta M_y\end{bmatrix} = \begin{bmatrix}C_L\\ C_D\\ M_y\end{bmatrix} - \begin{bmatrix}C_{L_\delta}\\ C_{D_\delta}\\ M_{y_\delta}\end{bmatrix}\bar{\delta} - \begin{bmatrix}\bar{\delta}^\top C_{L_{\delta^2}}\bar{\delta}\\ \bar{\delta}^\top C_{D_{\delta^2}}\bar{\delta}\\ \bar{\delta}^\top M_{y_{\delta^2}}\bar{\delta}\end{bmatrix} \tag{5}$$

ΔC_L, ΔC_D and ΔM_y are evaluated using current estimates of the control surface sensitivities.

A second set of perturbations are performed where angle of attack is held constant and the control surfaces δ are individually perturbed by $\pm\Delta\delta$, where $\Delta\delta = 4°$. However if one of these perturbations were to violate the saturation limits, two different deflections in the same direction are used. The difference between the perturbed lift, drag, and bending moment, C_L, C_D, and M_y, and the respective trim values $\bar{C}_L$, $\bar{C}_D$, and $\bar{M}_y$ during a single control surface perturbation is examined.

$$\begin{bmatrix} \delta C_L \\ \delta C_D \\ \delta M_y \end{bmatrix} = \begin{bmatrix} C_L \\ C_D \\ M_y \end{bmatrix} - \begin{bmatrix} \bar{C}_L \\ \bar{C}_D \\ \bar{M}_y \end{bmatrix} \tag{6}$$

A linear regression can then be performed by minimizing the error squared using the cost function given by[14]:

$$\min\ J = \sum_{i=1}^{2n} \left[z_i - \Theta_\delta^\top \Phi\left(\Delta\delta_k\right) \right]^2 \tag{7}$$

where

$$z_i = \begin{bmatrix} \delta C_L \\ \delta C_D \\ \delta M_y \end{bmatrix}_i \tag{8}$$

$$\Theta_{\delta_k} = \begin{bmatrix} C_{L_{\delta_k^2}} & C_{D_{\delta_k^2}} & M_{y_{\delta_k^2}} \\ C_{L_{\delta_k}} & C_{D_{\delta_k}} & M_{y_{\delta_k}} \end{bmatrix} \tag{9}$$

$$\Phi(\Delta\delta_i) = \begin{bmatrix} 2\bar{\delta}_{k,i}\Delta\delta_{k,i} + \Delta\delta_{k,i}^2 \\ \Delta\delta_{k,i} \end{bmatrix} \tag{10}$$

2. Drag and Maneuver Load Optimization

Following system excitation and RLS-based parameter identification, the optimal control surface deflections and angle of attack are computed using the onboard surrogate model. For the drag minimization objective, the optimal control commands are obtained by solving the following cost function:

$$\min\ J = C_D + \lambda_L\left(C_L - \bar{C}_L\right) \tag{11}$$

where λ_L is the adjoint variable for lift.

Evaluating the partial derivatives of the cost function gives:

$$\frac{\partial J}{\partial \alpha} = \frac{\partial C_D}{\partial \alpha} + \lambda_L \frac{\partial C_L}{\partial \alpha} \tag{12}$$

$$\frac{\partial J}{\partial \delta^\top} = \frac{\partial C_D}{\partial \delta^\top} + \lambda_L \frac{\partial C_L}{\partial \delta^\top} \tag{13}$$

$$\frac{\partial J}{\partial \lambda_L} = C_L - \bar{C}_L \tag{14}$$

Setting $\frac{\partial J}{\partial \alpha} = 0$ yields:

$$\lambda_L = -\left(\frac{\partial C_L}{\partial \alpha}\right)^{-1}\left(\frac{\partial C_D}{\partial \alpha}\right) \tag{15}$$

At each iteration, λ_L is evaluated using previous values for α and δ. By setting $\frac{\partial J}{\partial \delta^\top} = 0$ and using the updated valued for λ_L, the optimal control surface deflections are obtained. First, Eq. 13 is expanded:

$$\frac{\partial J}{\partial \delta^\top} = \frac{\partial C_{D_\delta}}{\partial \delta^\top} + \lambda_L \frac{C_{L_\delta}}{\partial \delta^\top} + 2\frac{\partial C_{D_{\delta^2}}}{\partial \delta^\top}\delta + 2\lambda_L \frac{C_{L_{\delta^2}}}{\partial \delta^\top}\delta \tag{16}$$

Then, the optimal control surface deflections are obtain by solving for δ. This is written as:

$$\delta = -\left(2\frac{\partial C_{D_{\delta^2}}}{\partial \delta^\top} + 2\lambda_L \frac{\partial C_{L_{\delta^2}}}{\partial \delta^\top}\right)^{-1}\left(\frac{\partial C_{D_\delta}}{\partial \delta^\top} + \lambda_L \frac{\partial C_{L_\delta}}{\partial \delta^\top}\right) \tag{17}$$

By setting $\frac{\partial J}{\partial \lambda_L} = 0$ the lift constraint equation is obtained:

$$\bar{C}_L = C_{L_0} + C_{L_\alpha}\alpha + C_{L_{\alpha^2}}\alpha^2 + C_{L_\delta}\delta + \delta^\top C_{L_{\delta^2}}\delta \tag{18}$$

from which the angle of attack α is computed using the values for δ obtained by Eq. 17. The controller then issues the optimal commands to the flight-control system for actuation, and the algorithm is iterated in real time until convergence.

For MLA, the optimization objective is to minimize wing-root bending moment while maintaining the prescribed lift. Accordingly, the following cost function is minimized to compute the optimal control surface deflections and angle of attack:

$$\min\ J = M_y^2 + \lambda_L\left(\bar{C}_L - C_L\right) \tag{19}$$

Evaluating partial derivatives of this cost function yields:

$$\frac{\partial J}{\partial \alpha} = 2\frac{\partial M_y}{\partial \alpha}M_y + \lambda_L\frac{\partial C_L}{\partial \alpha} \tag{20}$$

$$\frac{\partial J}{\partial \delta^\top} = 2\frac{\partial M_y}{\partial \delta^\top}M_y + \lambda_L\frac{\partial C_L}{\partial \delta^\top} \tag{21}$$

$$\frac{\partial J}{\partial \lambda_L} = C_L - \bar{C}_L \tag{22}$$

The MLA algorithm leverages the identified surrogate model to schedule optimal flap profiles and corresponding angles of attack as a function of the target lift, $\bar{C}_L$. The optimal flap deflections, angle of attack, and adjoint variable λ_L are obtained by setting $\frac{\partial J}{\partial \alpha} = 0$, $\frac{\partial J}{\partial \delta^\top} = 0$, and $\frac{\partial J}{\partial \lambda_L} = 0$. The resulting system in δ, α, and λ_L is then solved using the Newton-Raphson method. The controller subsequently applies the computed control inputs to minimize the wing-root bending moment for a given commanded lift coefficient, $C_{L_{cmd}}$.

B. Boeing Drag Controller Design

Boeing developed a fuel-burn minimization controller for the nominal full-scale aircraft simulation model. The fuel-burn optimization control law commanded the wing trailing edge surfaces to reduce fuel burn, which implicitly reduced drag. The baseline longitudinal and throttle control laws ran in parallel with the fuel-burn optimization controller and these baseline controllers would command elevator to re-trim the aircraft to maintain level flight. The end result was that the aircraft would trim at the same lift coefficient, but at a different Angle of Attack with a lower drag coefficient and fuel burn rate.

Fuel flow measurements are not available in this wind tunnel model, so Boeing designed a drag optimization control law for wind tunnel testing. Lift and drag measurements are directly available through NASA's balance and thus the controller seeks to explicitly minimize drag measurements. Figure 8 pictures the high-level architecture of the drag optimization control law. Note that the control law includes two separate controllers: one for gradient-based optimization of drag coefficient and a second for lift coefficient tracking. The former contains the actual drag optimization component while the latter is intended to replace the baseline longitudinal and throttle controllers that are missing from a wind tunnel model.

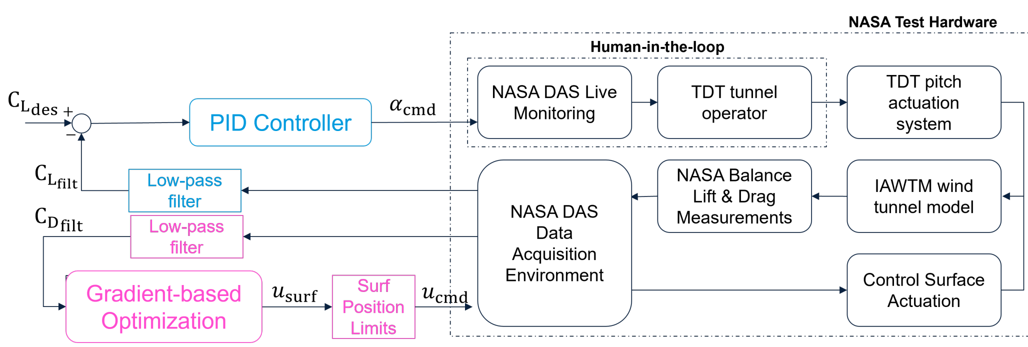

Fig. 8 High-level drag optimization control architecture. There are two separate controllers: one for drag coefficient optimization and a second to maintain constant lift coefficient.

Since the wind tunnel model is not a free-flying aircraft and lacks any horizontal stabilizer, the baseline longitudinal control laws no longer apply. In their place, a separate lift controller is designed to track a desired lift coefficient through explicit pitch angle commands. As shown in Figure 8, this lift controller has unique design considerations since there is no direct connection between the pitch angle commands and the actual wind tunnel model pitch angle. There is a human-in-the-loop element where the pitch angle commands are sent to the TDT tunnel operators, who then manually adjust the TDT pitch actuation system to produce the desired pitch angle.

Figure 9 shows a block diagram of the lift controller and its elements. The controller itself is a straightforward discrete-time PID controller, but the implementation is complicated by various components to account for the human TDT operators and avoid exciting aeroelastic modes. First, it is easy to task-saturate the human TDT operators with rapidly changing pitch angle commands. In order to slow down the response, the lift controller operates at 0.1 Hz so the operators would only receive a pitch angle command update once every 10 seconds. The intent is that would provide sufficient time to change the model's pitch angle and minimize any potential overshoot. Additionally, the pitch angle commands are rate limited to 0.27°/sec, which is the rate limit for the TDT pitch actuation system when operating in "slow mode." This also has the benefit of slowing down the pitch angle response, which will help avoid the excitement of aeroelastic modes.

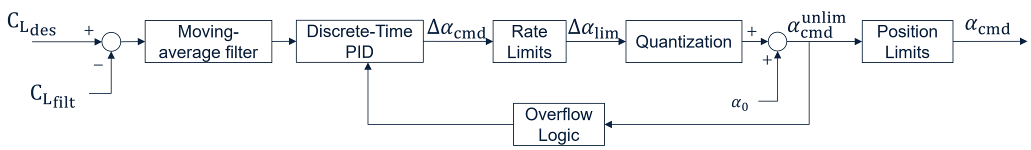

Fig. 9 Lift controller component of the drag optimization control law.

The pitch angle commands are limited between −2° and +3°. These limits were approved in pre-test analysis to keep the model within approved NASA balance loads. As shown in Figure 9, the lift controller also incorporates integrator management protections for when α_{cmd}^{unlim} exceeds the aforementioned pitch angle deflection limits. The overflow detections are passed to the integrator's anti-windup elements.

Lastly, the drag optimization control law includes switches to selectively disable the lift controller. This was used extensively during wind tunnel testing to iteratively test controller performance. The drag controller was tested by itself to confirm control surface movement before enabling the lift controller.

The drag controller component utilizes a peak-seeking control architecture[**?**] to perform model-free optimization on the drag coefficient measurements. This peak-seeking controller estimates the control surface derivatives and moves the control surfaces in the direction projected to minimize drag coefficient. The controller implementation also includes additional saturation elements to ensure the control surface commands stay within desired limits. Figure 10 pictures the drag controller and its major elements.

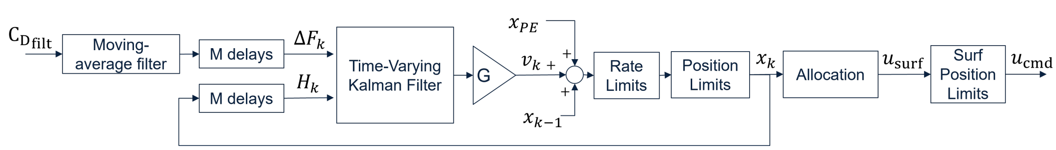

Fig. 10 Drag controller component of the drag optimization control law.

The peak-seeking controller models drag as a Taylor series expansion around the current control surface commands (x_k), where $H.O.T$ represents the higher-order terms, $C_D(x_{k-1})$ is the previous timestep's drag measurement, and x_{k-1} is the control surface position command at the previous timestep.

$$C_D(x_{k-1}) \approx C_D(x_k) + b_k^T(x_{k-1} - x_k) + H.O.T. \tag{23}$$

The term b_k corresponds to the control surface gradient at the current timestep, where N is the number of active control surfaces.

$$b_k^T = \begin{bmatrix} \frac{\partial C_D}{\partial x_{1k}} & \frac{\partial C_D}{\partial x_{2k}} & \dots & \frac{\partial C_D}{\partial x_{Nk}} \end{bmatrix} \tag{24}$$

The time-varying Kalman filter estimates these control surface derivatives.

The peak-seeking controller multiplies the Kalman filter gradient estimate by a gradient descent gain matrix, G. The gradient descent output v_k is then added to the previous command x_{k-1} as well as a persistently exciting signal, x_{PE}. The persistently exciting signal adds extra perturbations to the system to improve observability of the drag measurements and potentially prevent the Kalman filter from prematurely converging to local optimums. The net control surface commands are fed through position and rate saturation limits. The drag controller's surface position commands are limited to ±5°. These limits prevent large control surface deflection limits in order to avoid large changes in the balance loads. The rate saturation limits are set to 0.5°/sec to prevent the control surfaces from moving too fast. This is done to minimize the excitation of aeroelastic modes and limit transient effects on the aerodynamic data. The rate limits also help prevent overloading on the NASA balance since the low rate limits should give the human TDT tunnel operators ample time to terminate the controller if the balance loads continue to approach their approved limits. The position- and rate-limited commands are fed back to the tapped delays.

In order to reduce the dimensionality of the problem and speed up convergence, the peak-seeking algorithm uses pseudo-surface groupings rather than individual control surfaces. Due to the large amount of noise on the drag measurements, it would be difficult to observe the change in drag caused by a single mini-plain flap. Grouping the MPFs into pairs is expected to increase the observability of the drag gradient. The MPF groupings also help improve the convergence of the model-free optimization since it reduces the number of dimensions the system must perturb. Figure 11 shows the allocation used by the drag controller. MPFA and MPFB are grouped into a single pseudo surface. MPFD and MPFE form a second grouping while MPFF and MPFG form a third. Inboard aileron AILA is used by itself since it is a larger surface and thus produces a more noticeable change in drag when deflected. Outboard ailerons AILB and AILC are disabled for the drag controller. The initial plan was to test the drag controller at higher Mach numbers. Data from the first wind tunnel entry and CFD showed that the outboard ailerons would experience control reversal at higher conditions, so the surfaces were locked-out. Inboard MPFC was disabled due to the challenges encountered with MPFC during the first wind tunnel entry. Seen in Figure 10, the pseudo-surface commands x_k are sent through the allocation matrix to produce commands for the individual surfaces. Note that these commands are sent through another set of position limits. These position limits correspond to the individual control surface's position limits. Some of the surfaces, such as MPFD, have upper or lower deflection limits that are slightly more restrictive than the ±5° saturation limit of the pseudo-surface groupings. This last step will prevent the individual control surface commands u_{cmd} from exceeding those limits.

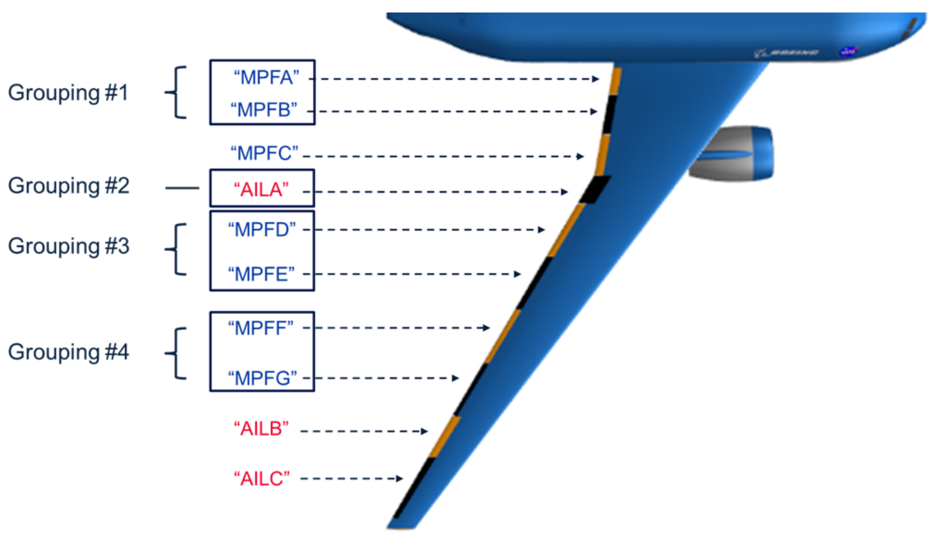

Fig. 11 Control surface groupings used by the peak-seeking drag controller.

Figure 12 shows the lift and drag coefficient increments associated with each of the control surface groupings plus ailerons AILB and AILC for comparison. Groupings MPFA/B and MPFD/E produce the largest change in both plots. Aileron AILA also has noticeable effectiveness, but the MPFF/G grouping is the least effective out of the 4 drag controller groupings. The plots also illustrate why ailerons AILB and AILC are not useful for drag optimization.

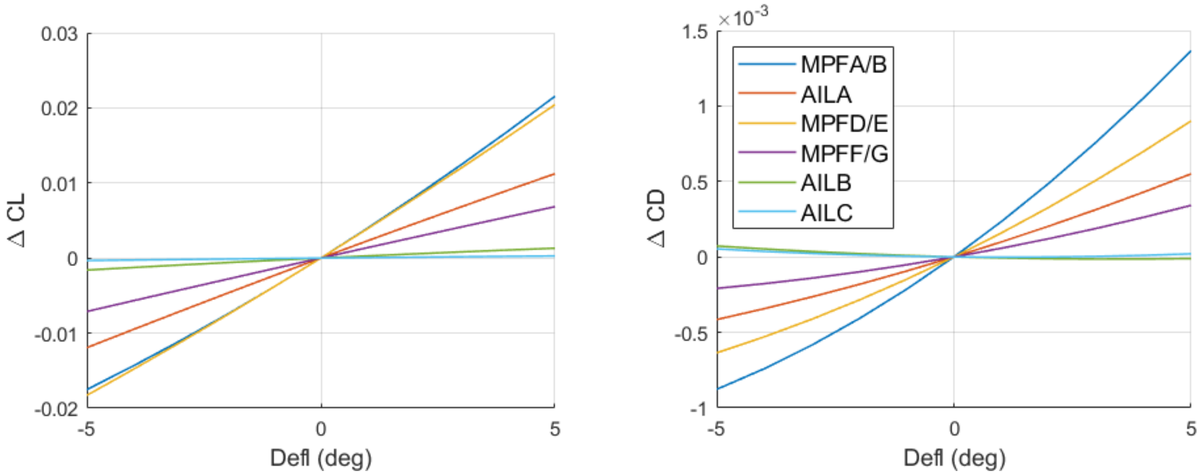

(a) Lift coefficient **(b) Drag coefficient**

Fig. 12 Lift and drag increments associated with each control surface grouping at $\bar{q}_\infty$ = 110 psf, α = 2°.

The drag optimization control law was tested in the simulation environment prior to the second wind tunnel entry. The simulation environment included Gaussian noise applied to each measurement with the noise magnitudes derived from data in the first wind tunnel entry. The tests indicated that the drag controller would reduce the drag coefficient even in the presence of noisy drag measurements. The lift controller would also maintain constant lift coefficient. At the conclusion of a 250 second simulation at Mach 0.55, $\bar{q}_\infty$ = 110 psf, α = 0°, the drag coefficient dropped from C_D = 0.01515 at t=0 sec to C_D = 0.01448, a 4.4% reduction in drag.

C. Boeing MLA Controller Design

Boeing designed a feedforward maneuver load alleviation control law. In a nominal full-scale aircraft, the MLA control law is based upon load factor. As load factor increases, the inboard control surfaces deflect trailing-edge down while the outboard surfaces deflect trailing-edge up in order to shift the load inboard and reduce the wing root bending moment. The reverse is true for a nose-down pitching maneuver. Figure 13 shows the control surface allocation schedule for the nominal full-scale aircraft.

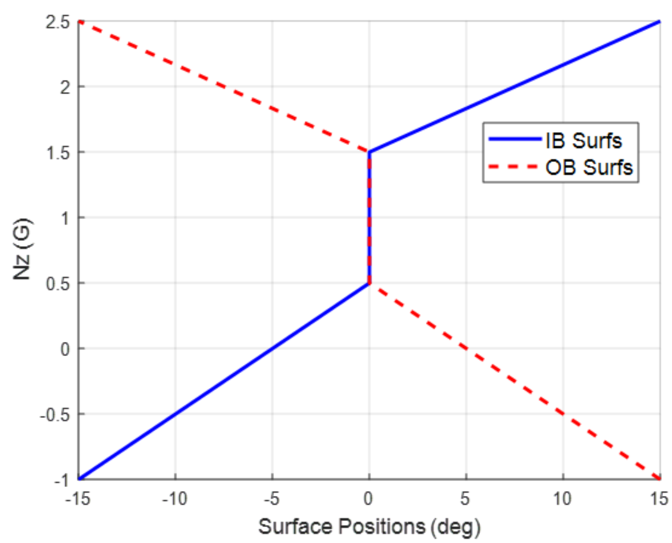

Fig. 13 Maneuver load alleviation control surface schedule for the nominal, full-scale CRM13 aircraft configuration.

For the wind tunnel MLA control law, the general architecture is the same, but numerous modifications were necessary for implementation in the wind tunnel environment. Figure 14 shows the high-level control architecture of the wind tunnel implementation. The main difference from the full-scale control law is that there is no direct measurement of load factor available. Instead, the wind tunnel MLA control allocation schedule is fed estimated load factor ($\hat{n}_z$). Figure 15 summarizes the computation of the estimated load factor. Filtered lift coefficient measurements are used to compute the ratio with a chosen "1g" lift coefficient as the estimated $\hat{n}_z$. In order to simplify the testing process, it is assumed that the model is at a nominal 1g condition when the MLA controller is engaged. This value of $C_{L_{filt}}$ when MLA is triggered is considered the 1g lift coefficient. Then, whenever the model pitches up or down, the resulting change is lift coefficient is treated as a change in load factor. Additional protections and saturation limits are applied to the outputs to prevent singularities.

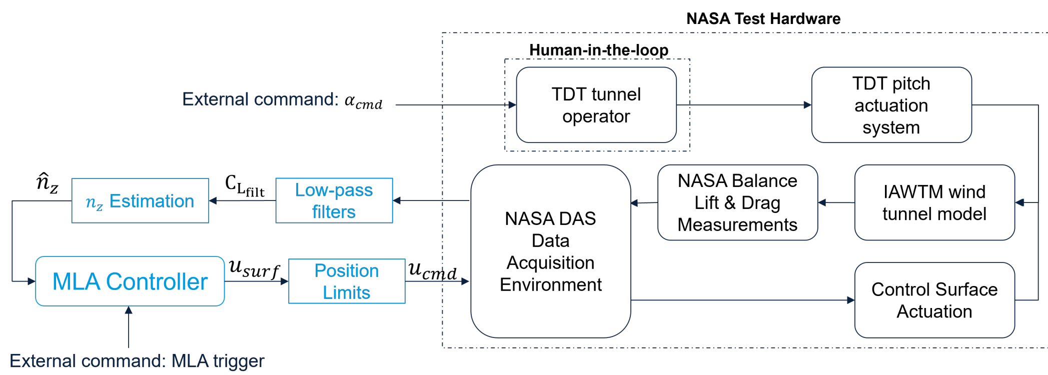

Fig. 14 High-level maneuver load alleviation control architecture of the Boeing feedforward controller.

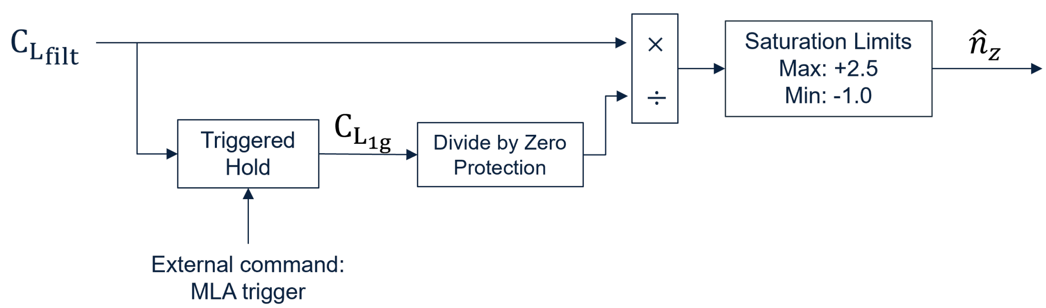

Fig. 15 Estimation of load factor ($\hat{n}_z$) using filtered lift coefficient ($C_{L_{filt}}$). The controller assumes the model is started at 1g when MLA is enabled.

The estimated load factor is passed to the same control allocation schedule as shown in Figure 13. For the wind tunnel implementation, the inboard and outboard surface allocation is shown in Figure 16. Inboard mini-plain flaps MPFA and MPFB are grouped with inboard aileron AILA as the "inboard" surfaces. Outboard mini-plain flaps MPFD, MPFE, MPFF, and MPFG and grouped with outboard aileron AILB as the "outboard" surfaces. Note that outboard aileron AILC is locked-out from the outboard surface commands as it would produce control reversal at some of the higher test conditions. This was also expected with the full-scale nominal aircraft. Inboard mini-plain flap MPFC was also disengaged, but for another reason. During testing in the first wind tunnel entry, MPFC experienced a number of issues and would sometimes freeze at nonzero deflections, requiring the tunnel to come down. It was easiest to remove MPFC from the MLA control allocation in order to prevent this from happening during closed-loop testing.

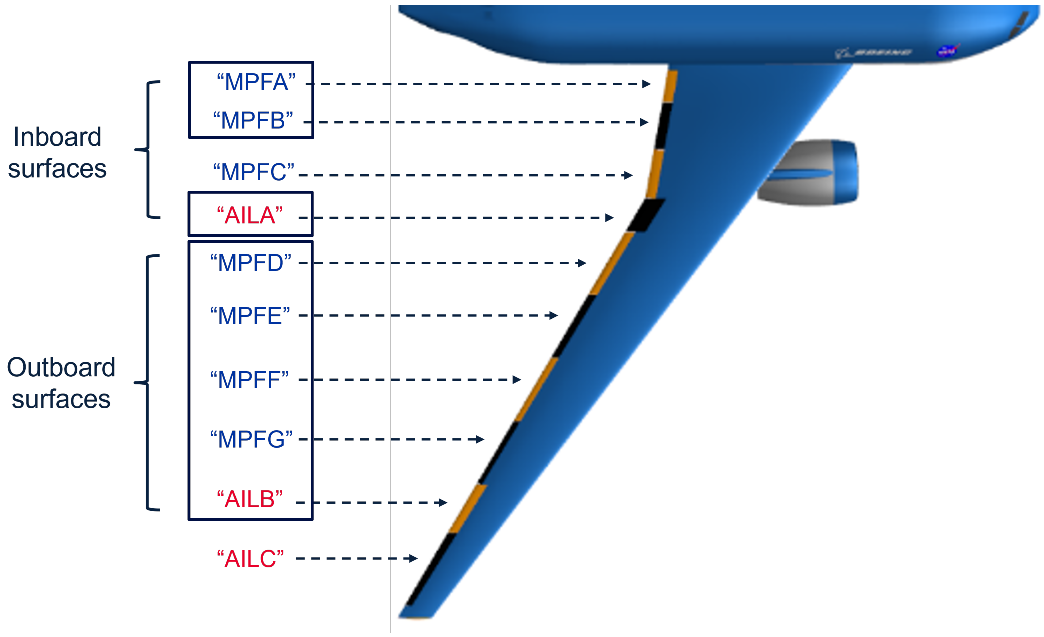

Fig. 16 Maneuver load alleviation control surface allocation for the wind tunnel model. Outboard aileron AILC is locked-out due to avoid control reversal while inboard MPF MPFC is disabled due to issues with the actuator.

Lastly, the MLA control surface commands are routed through the previously-described control surface saturation limits. Since the MPFs have different upper and lower deflection limits, the non-uniform saturation limits will produce a non-uniform spanwise distribution of control surface positions as the magnitude of the estimated load factor increases. Regardless, the wind tunnel testing results in Section IV.E show that the resulting control surface distribution still manages to decrease inboard strain gage SG01 measurements by a significant amount.

IV. Test Results

A. NASA Real-Time Drag Optimization Test Results

The NASA Real-Time Drag Optimization control law is tested at the primary test condition of Mach 0.55 and $\bar{q}_\infty$ = 110 psf. During the perturbation phase, the angle of attack is perturbed by ±0.5°, and each control surface is perturbed by ±2°. Saturation limits are set at ±5° for the MPFs and ±10° for the hydraulic ailerons. The target lift coefficient is set to $C_L = 0.388$.

The test is also initialized from a high-drag control surface configuration within the defined saturation limits. This approach improves the signal-to-noise ratio of the drag measurements, thereby increasing the likelihood of a successful test. Additionally, a low-pass filter with low cutoff frequency is used to mitigate the effect of tunnel turbulence and structural dynamic response of the model on the lift and drag measurements. The low-pass filtering is combined with long command time durations enhance the accuracy of aerodynamic measurements via a time-averaging process of the filtered data. Sufficient settling time is allowed after each new set point to ensure the tunnel reaches steady-state.

The control surface position command generated by the real-time drag optimization controller is sent to the servoactuators controlled by the dSPACE system. The angle of attack is controlled manually by a human operator, therefore additional delays are incorporated to allow sufficient time for the operator to achieve the commanded set point.

Fig. 17 depicts the drag time history for the full test duration. The raw drag signal is filtered using a moving average for legibility and also depicts the low-pass-filtered drag measurement. The dashed black line indicates the drag at the initial condition, while the dashed blue line represents the final drag at the optimized angle of attack and flap profile. This is a 12.2% reduction in drag which is slightly less than the predicted drag reduction of 14.4%.

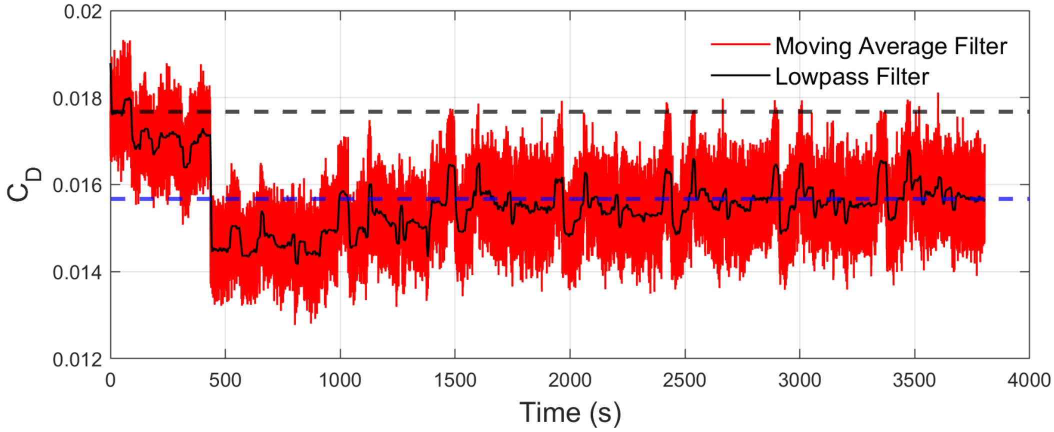

Fig. 17 Drag Time History

Fig. 18 depicts the lift measurement for the full test duration. The optimized lift reaches $C_L = 0.390$, which is 0.52% above the target. Early iterations show larger deviations from the target lift, which gradually converge as parameter identification improves. From approximately 0 to 400 seconds, the first set of perturbations occurs, with the larger response being angle of attack perturbations followed by the individual control surface perturbations. The lift and drag sensitivities are initialized at zero for the system identification algorithm. As such, the onboard model starts with relatively low accuracy which is why the drop in lift from 500 to 1500 seconds occurs. The estimation improves with each controller iteration as additional data is added to the recursive least-squares parameter identification, which results in the lift closely converging to the target lift.

Figure 19 shows the angle of attack measurement with an initial value of 0.16° and a final optimized value of 1.47°. The perturbations applied for system identification are also clearly visible. Figure 20 presents the control surface history, showing perturbations and reasonable convergence. Given the significant noise in lift and drag measurements, perfect convergence is not expected.

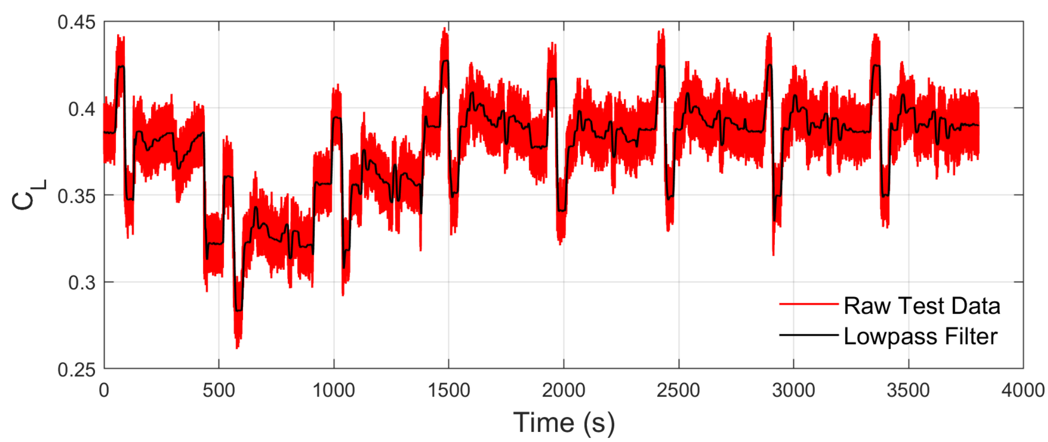

Fig. 18 Lift Time History

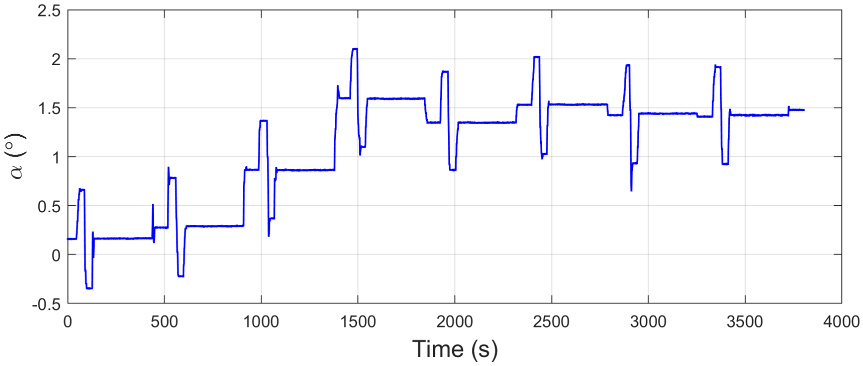

Fig. 19 Angle of Attack Time History

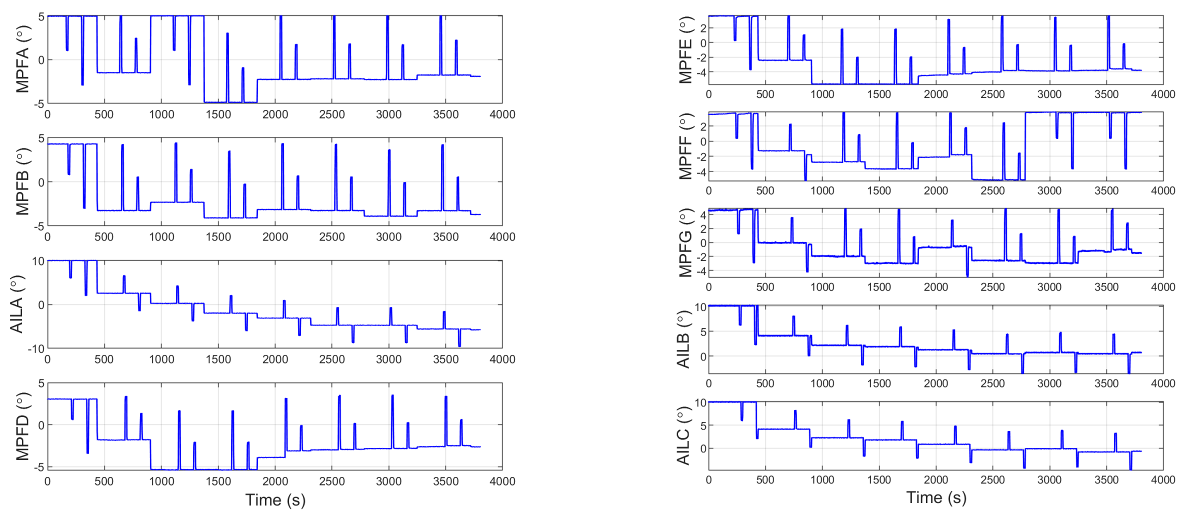

Fig. 20 Control Surface Position Time History

B. NASA Maneuver Load Alleviation Results

The NASA maneuver-load-alleviation (MLA) control law is also tested at the primary condition of Mach 0.55 and $\bar{q}_\infty = 110$ psf. The controller receives a commanded lift increment of $\Delta C_L = 0.4$, starting from $C_L = 0.15$, which corresponds to a maneuver from approximately $-2°$ to $3°$ when increasing angle of attack only. An offline model is used to compute the optimal control surface position and angle of attack for a target C_L. These optimal profiles are then scheduled with C_L to provide control commands to the dSPACE system and angle of attack operator for a C_L command. Two saturation configurations are evaluated: $\pm5°$ for the first run and $\pm10°$ for the second.

Fig. 21 illustrates the lift versus the root-strain magnitude. The blue data denote the baseline maneuver using angle of attack only, with no MLA control active. The red data denote the MLA-active case with $\pm5°$ saturation limits, and the green data denote the MLA-active case with $\pm10°$ limits. In both MLA-active cases, the wing-root strain decreases at a given C_L relative to the nominal case, with further reduction achieved using the larger saturation limits.

Fig. 22 shows the optimal control surface profiles at $C_L = 0.5$ for both saturation configurations. The inboard surfaces deflect positively, while the outboard surfaces deflect negatively, effectively shifting the aerodynamic load inboard to reduce the wing-root bending moment. At $C_L = 0.5$, the wing-root strain decreases by approximately 7% for the $\pm5°$ saturation case and by 13.2% for the $\pm10°$ case relative to the nominal maneuver.

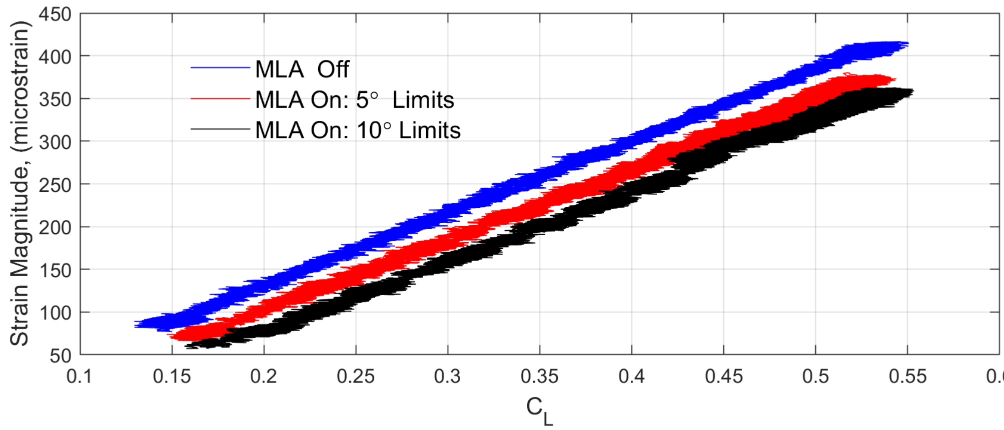

Fig. 21 MLA Controller Performance

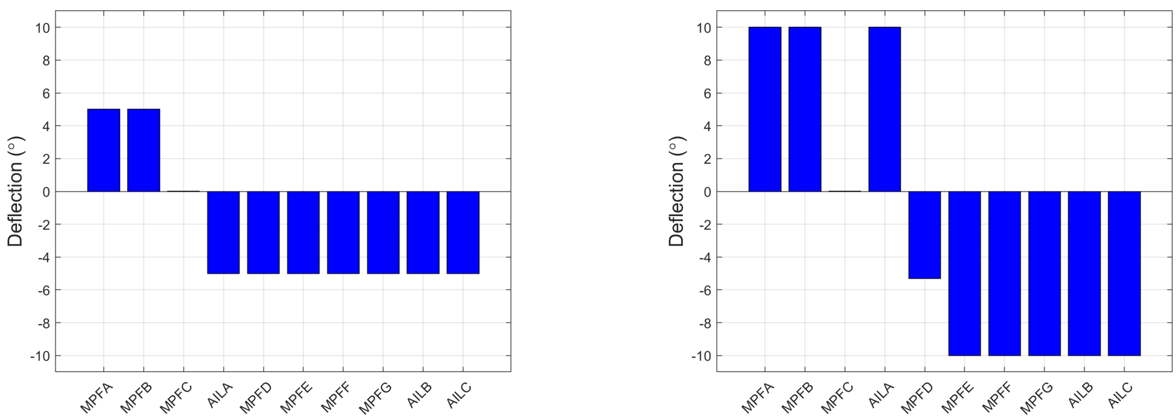

Fig. 22 MLA Optimal Profile Corresponding to C_L = 0.5

C. Boeing Drag Test Results

The Boeing drag optimization control law was tested at the primary test condition Mach 0.55, $\bar{q}_\infty = 110$ psf with $\alpha = +2°$. A +2° pitch angle was chosen for the tests to produce the largest drag measurements in an attempt to mitigate some of the effects of the large amount of noise on the drag signal channel. Before the test, the intent had been to test the controller at different Mach numbers and Angles of Attack; however, delays truncated the test window and it was only possible to test at the primary test condition. Since the drag optimization control law converges slowly by design, the tests are performed in 10 minute data records, unless where otherwise noted.

After testing was complete, a couple of issues were discovered which complicated analysis of the results. First, it was not known that the dSPACE diagnostic signals would be affected by quantization limits. NASA opened up 14 dSPACE-out channels for use as diagnostic signals. There had been a lot of discussion about upper saturation limits on the channels, but there was no discussion on the analog-to-digital resolution. Unfortunately, signals with small magnitudes, such as lift and drag coefficients were subject to quantization resolution roundoff, which limited their use. For example, the filtered drag coeffcicient signal used by the controller was recorded as one of those diagnostic channels; however, the data stored in the data records is not useful for analysis. This is seen in Figure 23 where the onboard filtered signal for drag is heavily affected by the analog-to-digital conversion resolution. Luckily, the raw drag measurements are recorded on analog channels and are not affected by this. Therefore, the lift and drag measurements seen in upcoming plots are the raw analog recordings subjected to post-processing filtering.

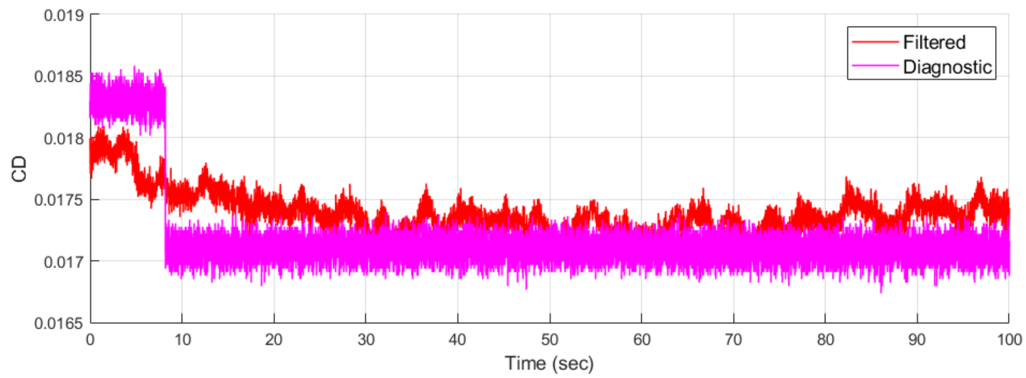

Fig. 23 Post-processed drag coefficient ("filtered") versus the onboard filtered signal recorded as a diagnostic channel. Note that the analog-to-digital conversion's resolution has adversely affected the diagnostic signal.

Likewise, the other controller diagnostic signals are also subject to the analog-to-digital conversion resolution issue. There were a couple of diagnostic channels to record the time-varying Kalman filter predictions, but these are rendering virtually unusable by the quantization. Most critically, the recordings of control surface gradients are mostly lost, as seen in Figure 24. It is impossible to determine the true gradient and whether the controller is actively converging towards a solution. It is possible to infer rough direction of how the controller intended to move a surface based upon the magnitude and direction of the diagnostic signals, but that only applies to situations where the controller attempts to command large deflections.

A compounding issue is that the small number of diagnostic channels (14) prevented all relevant controller signals from being recorded. As a result, only a subset of the drag controller's internal state information could be stored, providing an incomplete picture. It was determined to store only MPFA/B and AILA gradient information since those would be the first two surfaces used. At the time, the decision to store AILA information over MPFD/E was made based upon aileron vs. MPF actuator behavior. The electric servos used by the MPFs were subject to larger uncertainty about position as well as potential burnout/lock-up after extended duration. It therefore seemed useful to save AILA information since there was more confidence that the commanded position equaled the actual position. In the event some or all of the electric servos locked-up, it would still be possible to run drag optimization with just aileron AILA even if such operation of the drag controller would be suboptimal.

A second compounding issue is that the reset flags in the drag controller were not properly hooked in with dSPACE. The scaling parameters on the PE signals for each grouping worked as intended, where the PE perturbations could be disabled for each surface, but the resets to disable individual surface groupings did not. This allowed surfaces to slowly move as the previous gradients and surface positions were still present in the Kalman filter. It would have been possible to reset all the Kalman filter properties by resetting the dSPACE environment, but that could not be performed at the test conditions for safety reasons. The tunnel would have to be brought down, which would have required too much time from an already truncated test window.

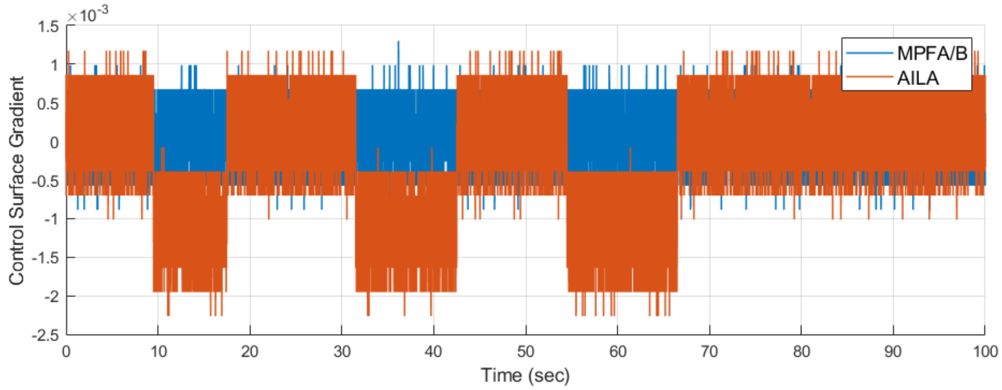

Fig. 24 Drag controller predictions of control surface gradients recorded on the diagnostic channels during Data Record 423. Unfortunately, the affect of the analog-to-digital conversion resolution has removed much, if not all, of the informativeness of these signals.

1. Session 1

The drag control law testing was broken into two sessions. The first testing session was affected by an error that prevented the pitch control law from being activated. The channel number for the lift force measurement was incorrectly labeled and the lift controller was not using the correct signal. As a result, pitch control was disabled in the first test session until the channel numbering issue could be fixed during tunnel downtime. In the meantime, the drag optimization control law was still capable of normal operation, but there was no parallel controller that attempts to maintain constant lift. Table 1 lists the important switches that were triggered during the closed-loop testing.

Table 1 Summary of the first session of closed-loop drag optimization testing at Mach 0.55, $\bar{q}_\infty$ = 110 psf, α = 2°

Data Record #	Pitch Control (y/n)	Aero Data Source: Surrogate Mdl. (y/n)	Aero Data Source: Exp. Data (y/n)	Persistency of Excitation Active: MPFA/B (y/n)	Persistency of Excitation Active: AILA (y/n)	Persistency of Excitation Active: MPFD/E (y/n)	Persistency of Excitation Active: MPFF/G (y/n)
422	N	Y	N	Y	N	N	N
423	N	Y	N	Y	Y	Y	Y
424	N	N	Y	Y	N	N	N
425	N	N	Y	Y	Y	Y	Y

The first two runs, Data Records 422 and 423, used the surrogate aero model instead of actual lift and drag measurements. The drag optimization control law implementation had a switch that could bypass the real measurements and replace them with lift and drag predictions from the aerodynamic model derived from Entry 1 and 2 experimental data. For Data Record 422, all 4 pseudo-surface groupings were active; however, only the MPFA/B grouping was perturbed by a persistency of excitation (PE) signal. Data Record 423 expands the PE signal perturbations to all 4 control surface groupings. After the completion of Data Record 423, the next two test points switched from the surrogate aerodynamic model to the raw experimental data. The lift and drag signals used by the controller are now the true measurements. The same PE signal triggers are used: Data Record 424 only applies PE perturbations to the MPFA/B grouping while Data Record 425 sequentially applies PE perturbations to the other 3 groupings.

Another item to note while looking at the results from the first session is that the Kalman filter reset switch had not been correct at this point. The issue was only discovered during this session (Test Point 423) and therefore there was no way of resetting the Kalman filter estimates and control surface initializations. As a result, the control surface commands would immediately migrate back to their final positions from the preceding test point whenever the drag controller was reengaged. This did not affect Test Point 422 since it was the first, but it's readily apparent in all subsequent data records.

2. Test Point 422

Figures 25 - 29 show the drag optimization controller results for Data Record 422. This case was using the surrogate aero model and only the MPFA/B grouping was sent PE perturbation signals. As a result, the numerical gradients for the AILA, MPFD/E, and MPFF/G groupings are set to zero and those control surfaces remain fixed to zero even though the controller has not explicitly deactivated those surfaces. Only the MPFA/B grouping commands nonzero deflections. One note is that the controller is using the surrogate aerodynamic model so the lift and drag signals used by the control law are "cleaner" and lack potential real-world nonlinearities and stochasticity.

In the experimental drag recordings shown in Figure 26, the noise on the drag measurements is so high it is very difficult to perceive any appreciable change. Figure 27 shows the same drag data, but only shows the filtered results after a smoother has been applied in post-processing. This new figure shows the experimental drag measurements do indeed converge towards a lower drag value while the MPFA/B control surface commands hover around $-4°$. The average of the first 3 seconds of raw drag measurements indicated a drag coefficient of roughly $C_D = 0.0179$ while the average of the final 3 seconds of raw drag measurements indicated $C_D = 0.0173$, about a 3.3% reduction. The lift measurements in Figure 25 confirms the controller follows engineering intuition: in the absence of a requirement to maintain constant lift, the easiest method to reduce drag is to dump lift.

Figure 28 shows the control surface gradients for MPFA/B and AILA computed by the drag controller. Since the controller is starting from zero and only MPFA/B has persistency of excitation applied, AILA does not move and therefore the gradient remains fixed at zero. The MPFA/B gradient does match expectations for how the controller would reduce drag. The gradient is multiplied by the gradient descent gain to produce incremental control surface commands. The positive gradient will command movement in the trailing-edge up direction. The spike in MPFA/B gradient between t=50-60 seconds corresponds exactly to where the MPFA/B commands rapidly descend in the trailing-edge up direction. When the gradient recedes back towards zero, this means the incremental commands are trending towards zero and the surface commands are relatively static. This is seen in Figure 29 where the MPFA/B commands fluctuate around $-4°$. The fluctuations around $-4°$ are due to the PE perturbations inserted into the commanded positions.

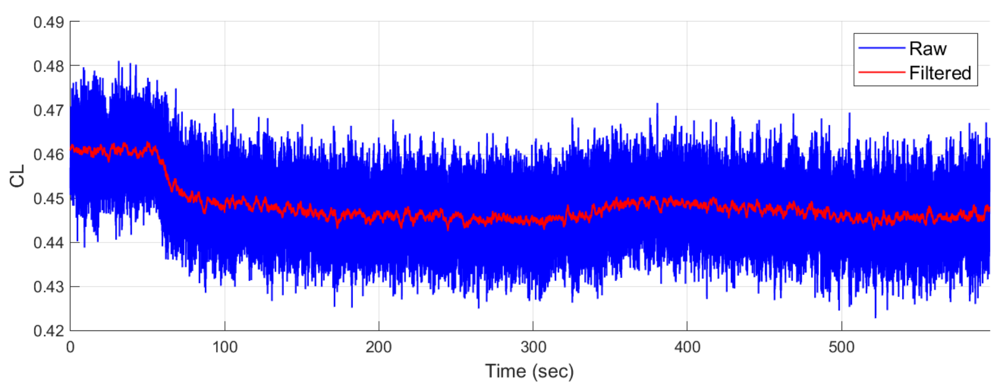

Fig. 25 Raw vs filtered lift coefficient (C_L) for Data Record 422.

3. Test Points 423 - 425

The remaining test points in the first section through various switches to test drag controller performance with different combinations of data sources and control surface perturbations.

In order to limit the paper to a reasonable size, the following paragraphs will summarize the remaining data records from the first test session without extensive analysis and plots. Data Record 423 uses the same process as 422, but now routes PE signals through the remaining 3 control surface groupings. Unfortunately, the controller did not converge to a lower drag state during this test point. It is not clear whether this occurred because the simultaneous excitation of 4 groupings obfuscated identification of the control derivatives or if it would have converged towards a lower drag state after a larger time duration. The computed gradients for MPFA/B and AILA were shown earlier in Figure 24. These gradients are not very useful for debugging the response since the digital-to-analog resolution issue muddled the data. As suggested, the simultaneous movement of MPFA/B, AILA, and MPFD/E could be obfuscating the actual gradients. The extensive pre-test simulation analysis had not encountered this behavior.

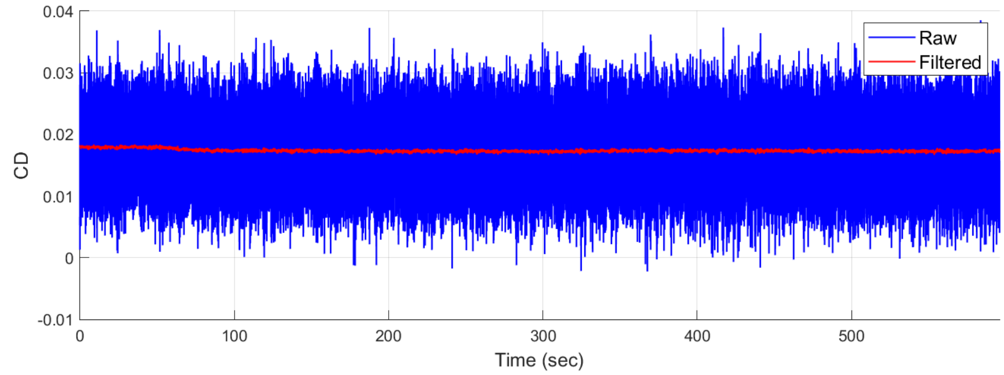

Fig. 26 Raw vs filtered drag coefficient (C_D) for Data Record 422. The magnitude of the noise on the drag measurement makes it difficult to perceive any appreciable change.

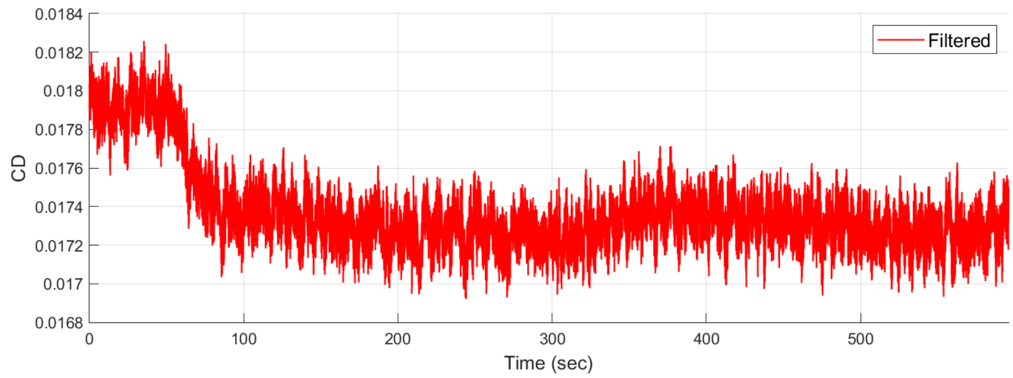

Fig. 27 Filtered drag coefficient (C_D) for Data Record 422 after a post-processing smoother has been applied to improve observability of trends in the noisy drag measurements.

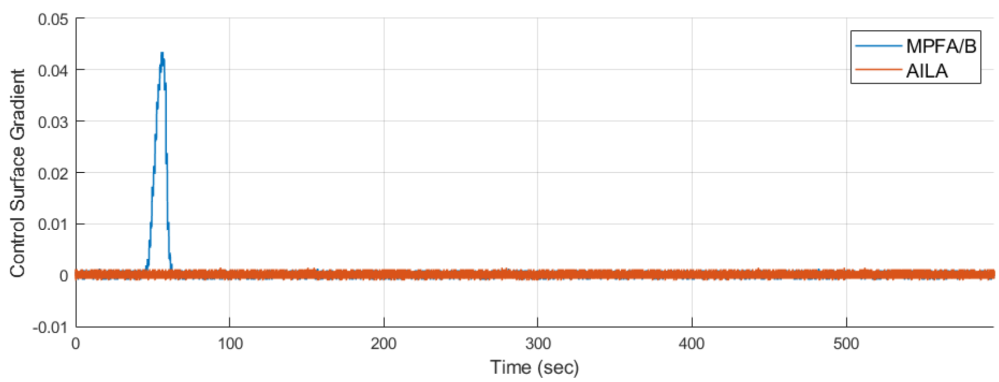

Fig. 28 Predicted control surface gradient computed by the drag controller during Data Record 422. The gradient is used to produce incremental commands. The positive gradient translates into a trailing-edge up surface command when multiplied by the gradient descent gain.

The source of the aerodynamic data was switched from the surrogate model to real experimental measurements for Data Records 424 and 425. Test Point 424 reverted the controller PE signals back to Test Point 422, where only the MPFA/B grouping is sent random perturbations. In this test, the drag measurements did converge towards the same $C_D = 0.0173$ value as seen in Data Record 422. Data Record 425 reactivates the persistency of excitation signals on

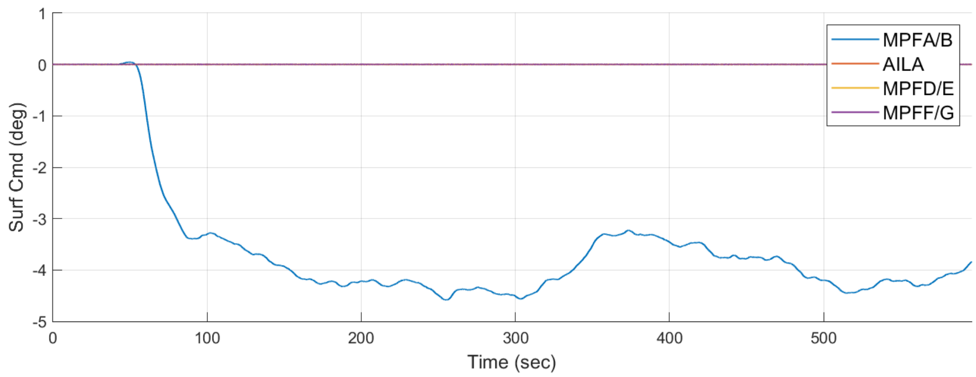

Fig. 29 Control surface commands for Data Record 422. Note that the persistency of excitation signals have been disabled on everything but the MPFA/B control surface grouping.

all 4 surface groupings. The drag coefficient results oscillate near the minimum drag values encountered during the previous tests, but hasn't converged by the end of the 10 minute data record. It was unclear whether the controller would have converged towards a local optimum given additional time, but the MPFA/B and MPFD/E groupings moved in the trailing-edge up direction that will reduce drag. As seen in Figure 12, those two surfaces have the steepest control surface gradient and their motion in the trailing-edge up direction would decrease net drag. While the MPFA/B and MPFD/E groupings have moved in the expected direction, AILA and MPFF/G did not converge. This isn't necessarily surprising as these surfaces produce the least amount of drag increment and the controller could have difficulty computing a gradient for those if the other surfaces are moving at the same time.

D. Session 2

The lift measurement signal was corrected for the second test session of the day. This allowed the lift controller to be activated. Table 2 lists the important switches that were triggered during this second round of closed-loop testing. In addition to the correction of the lift measurement signal channel, a Kalman filter switch was also added to the dSPACE page used to control the drag testing. This switch drained the Kalman filter estimates from the controller in an attempt to prevent the drag controller from commanding nonzero control surface position at the start of each test. This did work as intended, but there was insufficient time to redesign and recompile the controller to fully disable control surface groupings.

Table 2 Summary of the second drag testing session at Mach 0.55, $\bar{q}_\infty$ = 110 psf, α = 2°

		Aero Data Source		Persistency of Excitation Active			
Data Record #	Pitch Control (y/n)	Surrogate Mdl. (y/n)	Exp. Data (y/n)	MPFA/B (y/n)	AILA (y/n)	MPFD/E (y/n)	MPFF/G (y/n)
450	Y	N	Y	Y	Y	Y	Y
451	N	N	Y	Y	Y	Y	Y
452	N	N	Y	Y	Y	N	N
453	Y	N	Y	Y	Y	N	N
454	N	N	Y	Y	Y	Y	N

1. Test Point 450

Figures 30 - 34 show the drag optimization control law results for Data Record 450. In this test, both the drag and lift controllers are active. In the plots, it is clear the lift controller did not operate as intended and the system was oscillating between pitch angles. This is visible in Figure 33, where the Angle of Attack commands oscillate and eventually saturate at the +3° upper bound. The effects of the Angle of Attack variations are clearly visible in the lift and drag plots.

The drag controller's control surface commands did not converge to a particular set of values. The large jumps in Angle of Attack would make it difficult for the controller to identify whether a change in drag was due to surface movement or a pitch angle change. This is confirmed by the control surface gradients seen in Figure 32. At the start of the data record, the surfaces move slowly in random directions caused by the PE perturbations. However, the large drop in drag caused by the rapid initial Angle of Attack decrease complicates estimation of the gradients. The controller would attribute the change in drag to the movement of the surface commands and adjust the commands them accordingly. Indeed, the MPFA/B drag gradient estimates hover around -0.005 for a few seconds immediately following the drop in Angle of Attack. Those gradient values would produce trailing-edge down control surface commands, which is exactly what is seen in Figure 34.

There are a couple spots where the controller incorrectly attributes a large change in AoA-induced drag to control surface commands, but the resulting control surface commands are inadvertently in the correct direction that will reduce drag. The spikes in the MPFA/B gradient estimates around t=190 seconds, t=320 seconds, and t=400 seconds show this accidentally-correct behavior. Those spikes all correspond to points where the preceding MPFA/B commands had been slowly moving trailing-edge down, but a simultaneous large change in Angle of Attack produced a large increase in drag. The Kalman filter would have computed a large gradient and sent the surfaces in the opposite direction. It appears to be a coincidence that the surfaces move in the direction that actually reduces drag. The large downward spike in AILA gradient around t=320 seconds is a case where the controller computes an incorrect gradient and commands deflections in the direction that would inadvertently increase drag. The remaining MPFA/B and AILA gradient estimates are difficult to analyze since information is lost in the digital-to-analog quantitization. The random fluctuations seen in Figure 34 are the effects of the PE perturbations.

One observation stood out when observing the Angle of Attack commands in Figure 33 - the large jump in commanded AoA at the start of test. This jump occurred right as the control law was activated and before the drag controller commanded nonzero surface positions. This immediate jump in Angle of Attack indicated that the AoA initialization was set incorrectly. In Figure 9, the lift controller commands are rate-limited to ±0.27°/sec commands. The results in Figure 33 clearly show a faster change then that limit. Although there is no way of retroactively checking, the initial α_0 seen in Figure 9, must have been set incorrectly. These values were set by the human dSPACE operator prior to executing a test, and thus it is very likely the α_0 was set to a value much lower than +2°, which produced the large jump in Angle of Attack. The lift controller was not robust to this jump and did not converge once the surfaces started moving.

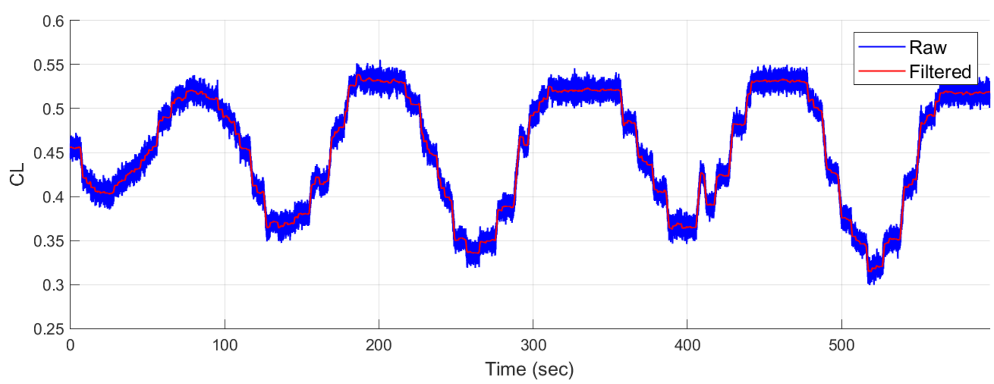

Fig. 30 Raw vs filtered lift coefficient (C_L) for Data Record 450.

2. Test Point 454

After the failure of the lift controller in Data Record 450, Data Record 451 disabled the lift controller and reran the test. The results were inconclusive and the control surfaces and drag measurements did not converge to a minimum drag steady state. The conclusion was the same as Test Points 423, 425, and 450: it appeared as if the drag controller encountered difficulty computing the gradient for the control surfaces when all 4 groupings moved simultaneously. Due to time constraints, the remaining tests were truncated to a 5 minute test window instead of the original 10 minute window. Data Record 452 and 453 reverted the PE signal excitations back to only MPFA/B and AILA. The combination of these two surfaces did reduce drag in Data Record 452. Data Record 453 reactivated the lift controller and the lift

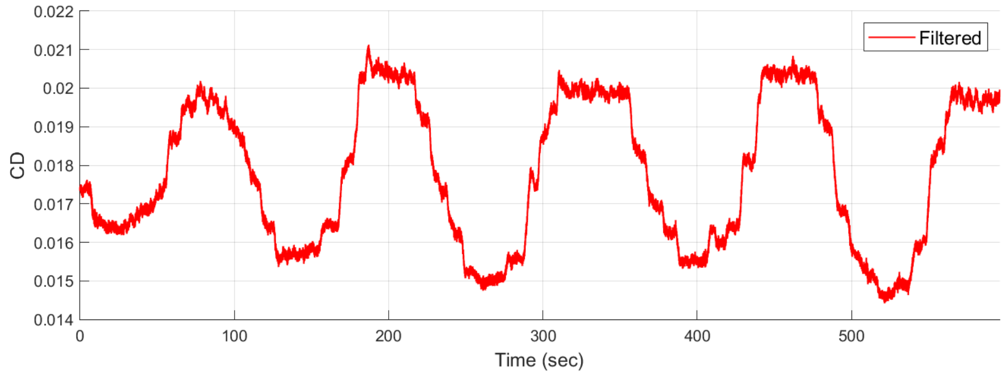

Fig. 31 Filtered drag coefficient (C_D) for Data Record 450 after a post-processing smoother has been applied to improve observabilty of trends in the noisy data measurements.

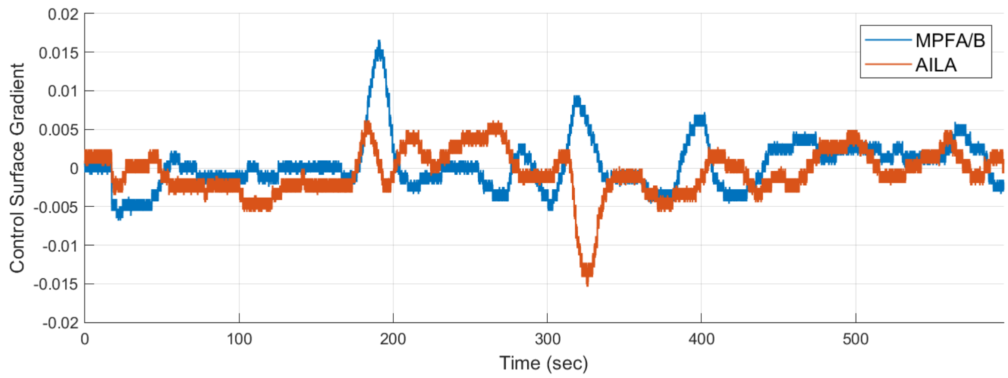

Fig. 32 Predicted control surface gradients computed by the drag controller during Data Record 450. The signals are affected by the analog-to-digital conversion resolution, but the gradients do spike enough at certain points enough to provide some basic observations.

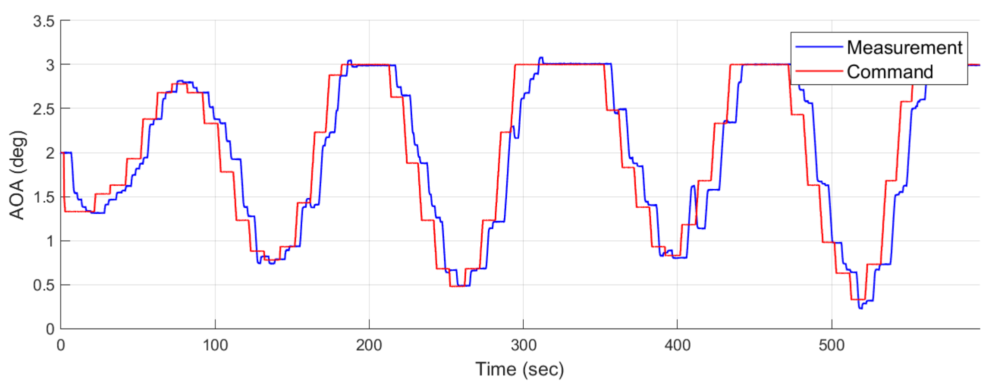

Fig. 33 Measured vs commanded Angle of Attack for Data Record 450.

measurements demonstrate the controller was indeed attempting to maintain constant C_L as intended. The drag results were inconclusive: the MPFA/B and AILA surfaces moved in the direction that would minimize drag, but a longer data record was required to confirm the drag controller converges towards a lower drag steady state while the lift controller varies Angle of Attack.

Data Record 454 is the last test point with the Boeing drag optimization control law. The lift controller was

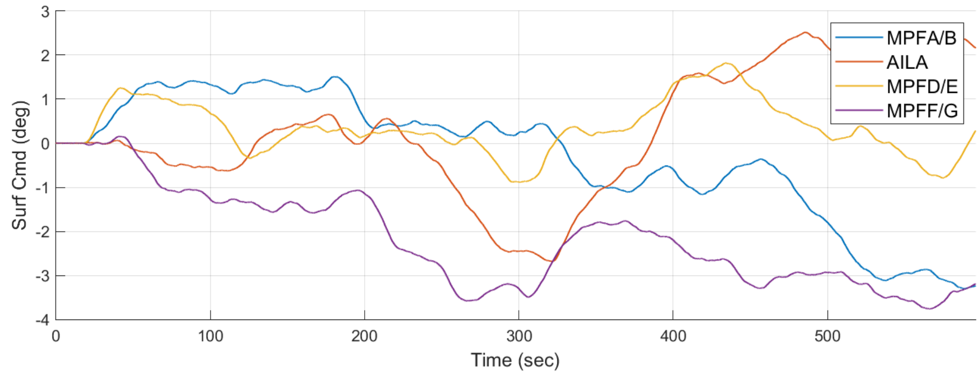

Fig. 34 Control surface commands for Data Record 450. Note that persistency of excitation signals are applied to all 4 groupings.

disengaged again, so there is no active control of the Angle of Attack and lift force will fluctuate. For this data record, PE perturbations are sent to three groupings: MPFA/B, AILA, and MPFD/E. Figures 35 - 38 plot the corresponding controller performance results. The drag results in Figure 36 are showing a net reduction in drag.

In Figure 38, the MPFA/B, AILA, and MPFD/E groupings all move in the direction that would reduce drag. They do not converge to the $-5°$ lower bound, which is where the drag would be minimized the most. It's not immediately obvious why they do not converge towards $-5°$, but it could be caused by the aforementioned difficulty computing gradients when multiple surfaces move simultaneously. The three surface commands fluctuated with the random increments inserted by the PE signal perturbations, but the simultaneous trend in the minimum drag direction is likely more than just coincidence. The MPFF/G surface commands did increase like MPFD/E did in Figure **??** and is also likely caused by the difficulty in computing gradients with all surfaces moving at once. This would be more acute for MPFF/G since that surface grouping has the smallest drag gradient and would produce the smallest change in drag. The controller would have difficulty identifying that the slight increase in MPFF/G command is increasing drag because the other three surfaces are producing a net reduction in drag.

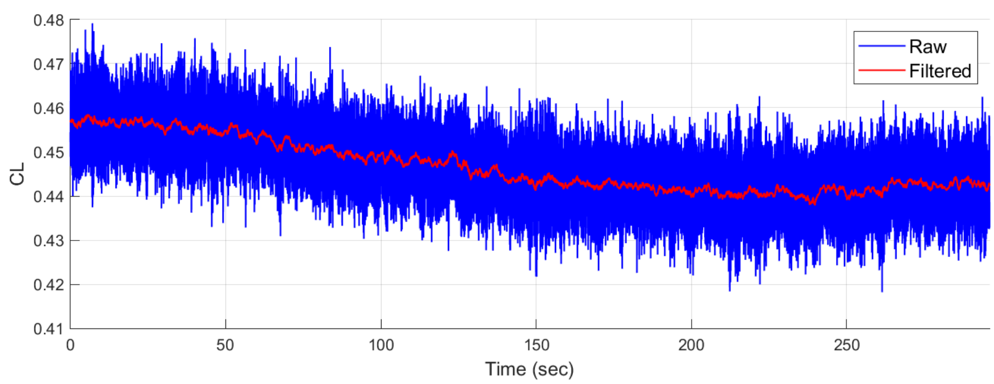

Fig. 35 Raw vs filtered lift coefficient (C_L) for Data Record 454.

E. Boeing MLA Test Results

The Boeing feedforward maneuver load alleviation control laws was tested at the primary test condition of Mach 0.55, $\bar{q}_\infty = 110$ psf. In order to execute a simulated pull-up maneuver, the wind tunnel model was set to $\alpha = -2°$ before the MLA controller was engaged. This resulted in the simulated 1g lift coefficient of roughly $C_L = 0.15$. This low Angle of Attack setting allowed NASA test engineers to pitch the wind tunnel model up through it's full range of motion and demonstrate large changes in lift coefficient. The wind tunnel model was then pitched up through a series of 0.5° step until $\alpha = 2°$, as shown in Figure 39.

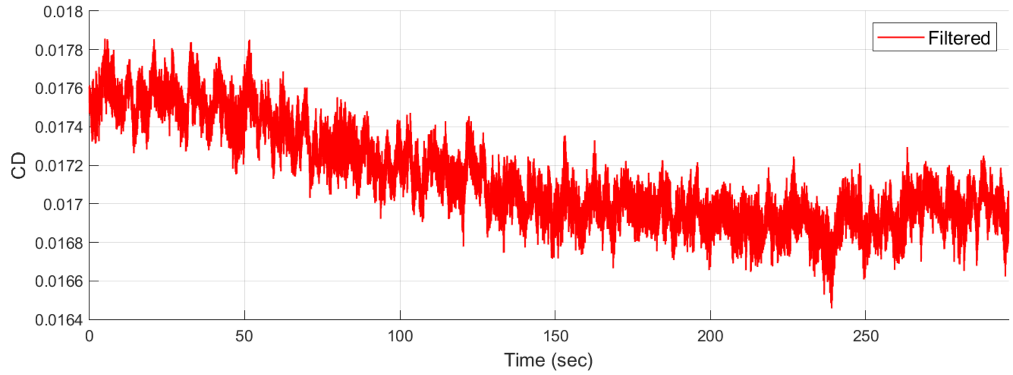

Fig. 36 Filtered drag coefficient (C_D) for Data Record 454 after a post-processing smoother has been applied to improve observabilty of trends in the noisy data measurements.

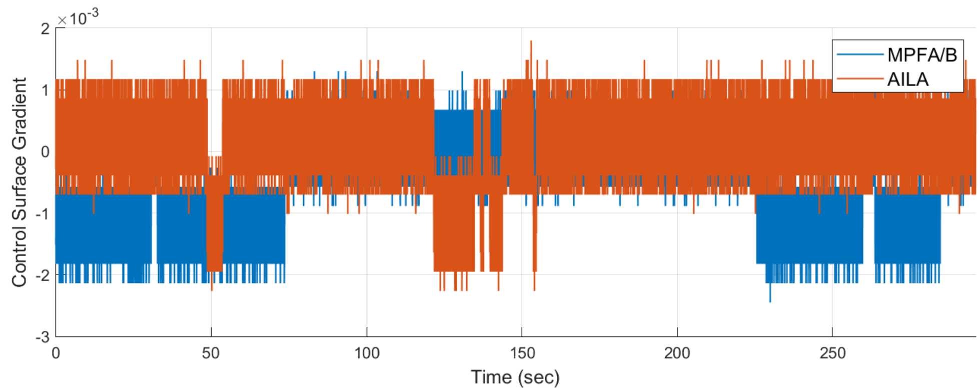

Fig. 37 Predicted control surface gradients computed by the drag controller during Data Record 454. The signals are affected by the analog-to-digital conversion resolution and lost any informativeness.

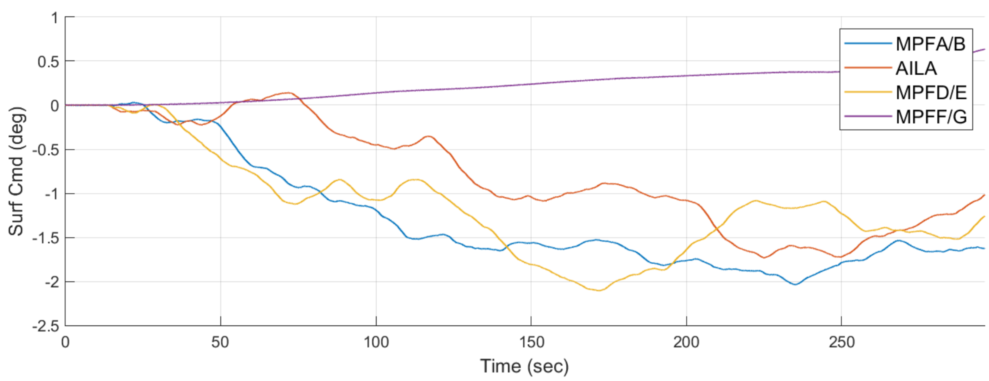

Fig. 38 Control surface commands for Data Record 454. Note that persistency of excitation signals are applied to the MPFA/B, AILA, and MPFD/E groupings.

When the MLA controller is engaged, the Angle of Attack increase will produce a corresponding change in estimated load factor. Figure 40 depicts this change in estimated load factor. Note that the estimated load factor hits the upper saturation limit of $\hat{n}_z = 2.5$g. Figures 41 and 42 show the resulting control surface commands for the inboard and outboard control surfaces. The control surfaces move in unison until the sequentially hit their individual deflection limits.

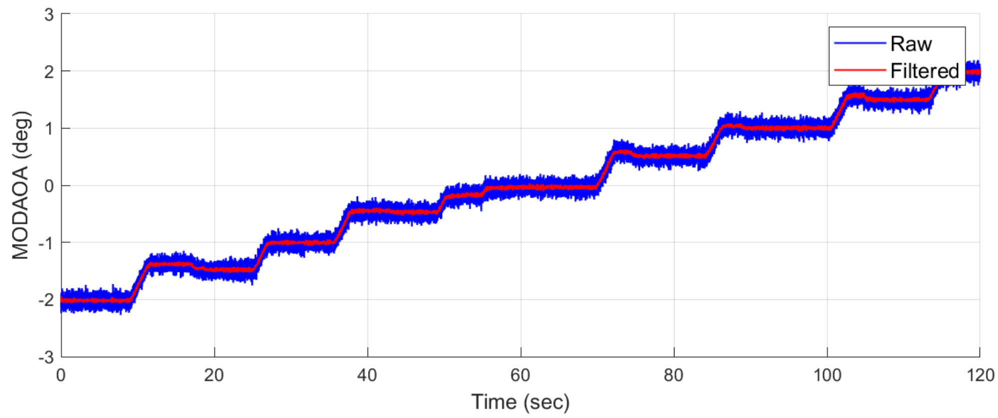

Fig. 39 Angle of Attack increase over time given $0.5°$ step commands.

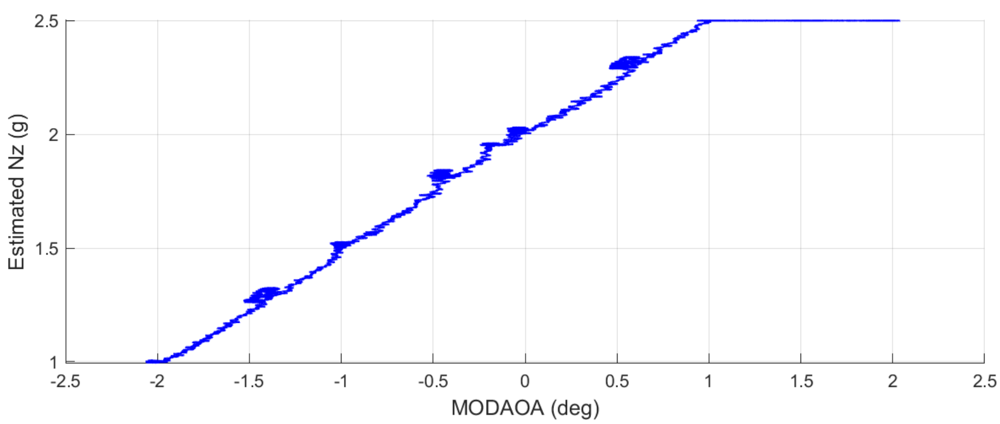

Fig. 40 Estimated Nz versus Angle of Attack during the simulated pull-up maneuver.

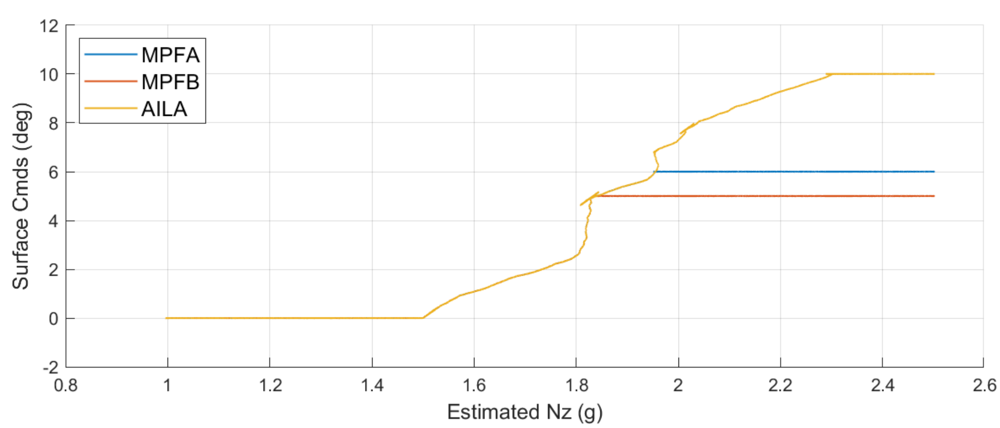

Fig. 41 Inboard control surface commands against estimated Nz for the simulated pull-up maneuver.

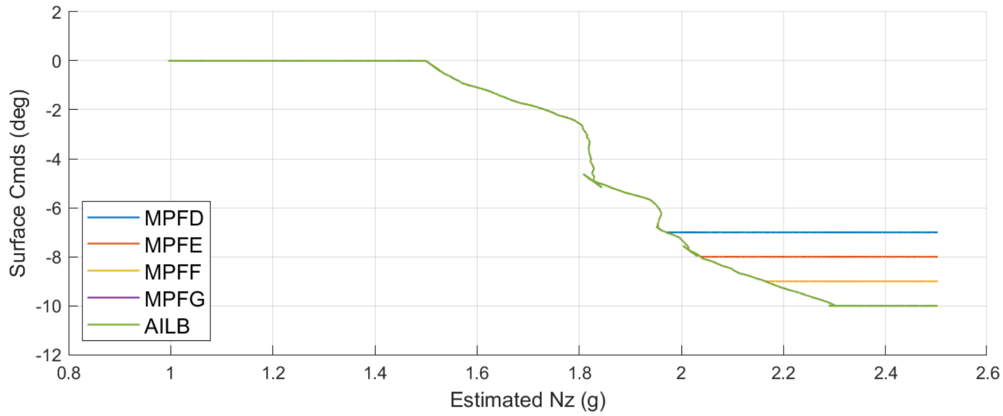

Fig. 42 Outboard control surface commands against estimated Nz for the simulated pull-up maneuver.

Figure 43 shows the raw strain gage SG01 measurements plotted against lift coefficient for the open- and closed-loop system. It is most appropriate to compare the strain gage measurements at the same lift coefficient since the MLA control surface commands will affect the lift coefficient. A performance improvement will only occur if the MLA controller can produce a lower strain at the same lift coefficient than the open-loop system. Indeed, Figure 43 clearly illustrates that when the MLA controller engages, the resulting strain is lower for the same lift coefficient as the open-loop system. The initial region from $C_L = 0.15$ to $C_L = 0.25$ corresponds to the piece in the allocation table from Figure 13 where n_z falls between 1 and 1.5g. As soon as the estimated load factor exceeds 1.5g, the effect of the MLA control surface commands are clearly visible. Table 3 summarizes the reduction in filtered SG01 produced by the MLA controller. As intended the feedforward controller consistently reduces the strain gage readings when compared against the open-loop system.

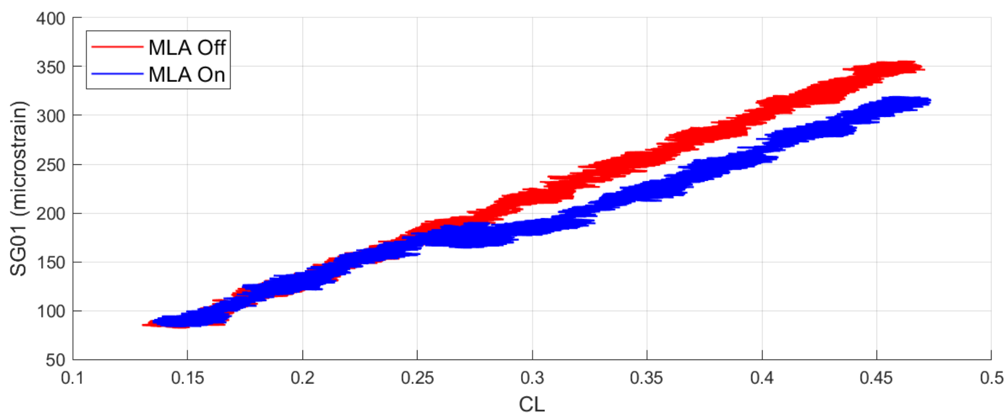

Fig. 43 Comparison of SG01 measurements against lift coefficient when the Boeing feedforward MLA controller is engaged vs disengaged.

Table 3 Reduction in filtered strain gage SG01 signals produced by the MLA controller at different lift coefficients.

Lift Coefficient C_L	Estimated $\hat{n}_z$	SG01: MLA Off	SG01: MLA On	Percent reduction
0.30	2.0g	218	190	12.84%
0.35	2.33g	266	231	13.16%
0.40	2.67g	311	273	12.22%
0.45	3.0g	347	310	10.66%

V. Conclusion

This paper presents the development and wind-tunnel validation of drag minimization and maneuver-load alleviation control strategies for the Integrated Adaptive Wing Technology Maturation (IAWTM) program. The controllers are implemented and tested during the June-July 2025 entry in the NASA Langley Research Center Transonic Dynamics Tunnel, using the 13.5 aspect ratio variant of the NASA Common Research Model equipped with multiple distributed control surfaces.

Both Boeing and NASA develop lift-constrained drag minimization control laws that leverage the distributed control surfaces to reshape the spanwise lift distribution and reduce overall drag. Complementary control strategies also address maneuver-load alleviation, with each approach effectively shifting aerodynamic loading inboard to reduce wing-root bending moments.

The results from this wind-tunnel campaign demonstrate measurable improvements in drag performance and load alleviation, marking a meaningful step toward enabling the practical use of highly flexible, high–aspect-ratio wings in future transport aircraft.

VI. Acknowledgments

This work was supported by the NASA Advanced Air Transport Technology (AATT) Project and the Advanced Air Vehicles Program (AAVP) as part of a Cooperative Agreement with Boeing. The authors would like to acknowledge significant contributions to this project and the work presented in this paper over the course of many years by team members from Boeing, NextGen Aeronautics, and NASA.

References

[1] Schaefer, J. R., and Brown, N. A., "Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Architecture and Performance Predictions," *AIAA Guidance, Navigation, and Control Conference*, Boston, MA, 2013. https://doi.org/https://doi.org/10.2514/6.2013-4856.

[2] Brown, N. A., and Schaefer, J. R., "Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Flight-Test Results," *AIAA Guidance, Navigation, and Control Conference*, Boston, MA, 2013. https://doi.org/https://doi.org/10.2514/6.2013-5171.

[3] Nguyen, N. T., and Hashemi, K. E., "Multi-Objective Flight Control for Ride Quality Improvement for Flexible Aircraft," *AIAA SciTech Forum*, 2020. https://doi.org/https://doi.org/10.2514/6.2020-1623.

[4] Forte, C., Nguyen, N., Xiong, J., and Sager, J., "Real-Time Drag Optimization and Maneuver Load Alleviation Control for a High Aspect Ratio Wing Wind Tunnel Model," *AIAA SciTech Forum*, San Diego, CA, 2022. https://doi.org/https://doi.org/10.2514/6.2022-0715.

[5] Nguyen, N., and Xiong, J., "Real-Time Drag Optimization of Aspect Ratio 13.5 Common Research Model with Distributed Flap System," *AIAA SciTech Forum*, 2021. https://doi.org/https://doi.org/10.2514/6.2021-0069.

[6] 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/https://doi.org/10.2514/1.J056603.

[7] Heaney, P. S., and Quindlen, J. F., "Overview of the Integrated Adaptive Wing Technology Maturation Wind-Tunnel Test Objectives," *AIAA SciTech Forum*, Orlando, FL, 2024. https://doi.org/https://doi.org/10.2514/6.2024-2509.

[8] Vassberg, J., Dehaan, M., Rivers, M., and Wahls, R., "Development of a Common Research Model for Applied CFD Validation Studies," *26th AIAA Applied Aerodynamics Conference*, Honolulu, HI, 2008. https://doi.org/https://doi.org/10.2514/6.2008-6919.

[9] Quindlen, J. F., Ortega, D. M., Clark, G. J., Hussain, H. S., Forte, C. J., Nguyen, N. T., Xiong, J., Heaney, P. S., Bentley, P. J., Coulson, D. A., Giles, I. M., McHugh, G. R., Pinkerton, J. L., Ramey, J. M., and Stevens, W. L., "Summary of the Integrated Adaptive Wing Technology Maturation Wind Tunnel Tests of a High Aspect Ratio Commercial Transport Aircraft," *AIAA SciTech Forum*, Orlando, FL, 2026.

[10] Xiong, J., Nguyen, N. T., Stanford, B., and Bartels, R. E., "Aeroelastic Modeling and CFD Simulation of Wind-Tunnel Scale Aspect Ratio 13.5 Common Research Model," *AIAA SciTech Forum*, 2021. https://doi.org/https://doi.org/10.2514/6.2021-2527.

[11] Xiong, J., Nguyen, N. T., and Bartels, R. E., "Wind-tunnel Wall Interference Investigation of an Aspect Ratio 13.5 Common Research Model Using FUN3D," *AIAA SciTech Forum*, National Harbor, MD, 2023. https://doi.org/https://doi.org/10.2514/6.2023-1571.

[12] Xiong, J., and Nguyen, N. T., "Aeroelastic Simulation of Wind Tunnel Scale Aspect Ratio 13.5 Common Research Model with Trailing Edge Mini-Flaps," *AIAA SciTech Forum*, Orlando, FL, 2024. https://doi.org/https://doi.org/10.2514/6.2024-2678.

[13] Xiong, J., Nguyen, N. T., and Bartels, R. E., "Study of Trailing Edge Mini Flaps Interactions on a Wind Tunnel Scale Aspect Ratio 13.5 Common Research Model," *AIAA SciTech Forum*, Orlando, FL, 2025. https://doi.org/https://doi.org/10.2514/6.2025-0873.

[14] Forte, C. J., and Nguyen, N. T., "Design and Experimental Validation of Real-Time Drag Optimization and Maneuver Load Alleviation Controls for a High Aspect Ratio Flexible Transport Wing," *AIAA SciTech Forum*, 2026.