GNC 2017 : 10th International ESA Conference on Guidance, Navigation & Control Systems
29 May - 2 June 2017, Salzburg, Austria
European Space Agency

EVALUATING THE STABILITY OF NASA'S SPACE LAUNCH SYSTEM WITH ADAPTIVE AUGMENTING CONTROL

Tannen S. VanZwieten(1), Michael R. Hannan(2), John H. Wall (3)

(1) NASA Engineering and Safety Center, Marshall Space Flight Center, Mail Code: EE04L, Huntsville, Alabama, 35812, 256.961.1509, tannen.vanzwieten@nasa.gov
(2) NASA Marshall Space Flight Center, Mail Code: EV41, Huntsville, Alabama, 35812, 256.544.1403, *mike.r.hannan@nasa.gov*
(3) Dynamic Concepts Inc., NASA Marshall Space Flight Center, Mail Code: EV41, Huntsville, Alabama, 35812, 256.544.6752, john.w.wall@nasa.gov

ABSTRACT

NASA's baseline Space Launch System (SLS) flight control system (FCS) design includes an adaptive augmenting control (AAC) component that modifies the attitude control system response to provide the classical gain-scheduled control architecture with additional performance and robustness. The NASA Engineering and Safety Center (NESC) teamed with the Space Launch System (SLS) Program to perform a comprehensive assessment of the stability and robustness of the FCS with AAC. This paper provides an overview of the approach, specific analysis techniques, and outcomes that were particularly relevant for the SLS Program. Multiple analysis techniques that specifically target the nonlinear AAC were commissioned as part of this assessment, which was completed outside of the Program's standard design analysis cycle. The following analyses were included, with each technique adding its own valuable insights: Lyapunov-based stability analysis, classical stability analysis with static AAC gain variations, circle criterion-based analysis of the FCS with a time-varying element, time-domain stability margin assessment, Monte Carlo simulations with expanded dispersions, and an extensive set of stressing cases. Several of the completed analyses focused on determining whether the inclusion of AAC introduced risk to the FCS, while others quantified the benefits of the adaptive augmentation.

1 INTRODUCTION

NASA's Space Launch System (SLS) Flight Control System (FCS) design is based on a classical gain-scheduled control design that has heritage with the Saturn Program [1] and Ares I-X. It consists of sensor blending, gain-scheduled proportional-integral-derivative control, bending filters, and a disturbance compensation algorithm, AAC, and optimal control allocation [2-4]. The adaptive component (AAC), central to this discussion, modifies the attitude control system response to provide the with additional performance and robustness. This algorithm was initially developed under the Constellation Program (CxP) [5], analyzed as a side-study for Space Launch System (SLS) Design Analysis Cycle (DAC)-1 (May 2012), and baselined as part of the SLS flight control system (FCS) architecture since DAC-2 (November 2012). The functionally intuitive design was shown to significantly enhance robustness in test cases without negatively impacting performance within the design envelope. The post-Preliminary Design Review (PDR) version of the SLS FCS flight software prototype, including the AAC, was flight tested on a piloted Fighter/Attack (F/A)-18 at NASA Armstrong Flight Research Center [6-10]. The aircraft acted as a surrogate launch vehicle by mimicking the pitch attitude error dynamics of the more massive, less responsive SLS for the completion of 100+ SLS-like trajectories.

Following the aforementioned algorithm development, maturation, and test activities, the NASA Engineering and Safety Center (NESC) and the Space Launch System (SLS) Program performed a comprehensive assessment of the stability and robustness of AAC [11]. This paper provides an overview of the approach, specific analysis techniques, and outcomes that were particularly relevant for the stability assessment. The standard launch vehicle flight control analyses performed prior to this joint SLS-NESC assessment, as part of SLS analysis cycles, were a combination of (1) frequency-domain stability analysis based on linear theory, and (2) high-fidelity Monte Carlo simulations. The former has shortcomings because it requires the linearization of the nonlinear AAC algorithm. The latter is of limited value because the core control algorithm (without AAC) is able to accommodate the dispersions and ACC is not substantially engaged. These analyses did not reveal any detrimental behavior, but neither did they fully exercise the adaptive algorithm. Thus, it was deemed prudent to commission a comprehensive, multifaceted analysis of the stability of the FCS *with* AAC.

Multiple techniques that specifically target the SLS AAC were commissioned, with each technique adding its own valuable insights. The following analyses were included: Lyapunov-based stability analysis, classical stability analysis with static AAC gain variations, circle criterion-based analysis of the FCS with a time-varying element, time-domain stability margin assessment, Monte Carlo simulations with expanded dispersions, and an extensive set of stressing cases. Several of the completed analyses focused on determining whether the inclusion of AAC introduced risk to the FCS, while others quantified the benefits of the adaptive augmentation. An overview of the analyses that were applied to the SLS AAC and major findings will be provided in the paper.

2 SLS FCS ARCHITECTURE

The primary elements of the SLS FCS, shown in Figure 1, are a blending of the rotational rates across three gyroscopes, classical proportional integral derivative (PID) control with bending filters, the disturbance compensation algorithm (DCA), and AAC.

These pieces are briefly described below [4], and readers are referred to the SLS CDR documentation [12] for a more in-depth discussion of each element and the associated FCS performance as a whole.

> Gyro blending: Optimizes attenuation of low-frequency structural modes by blending (weighting) the sensed rotational rate from multiple rate gyroscope assembly (RGA) sensors.
>
> PID + bending filters: Classical control architecture with bending filters specifically tailored to attenuate or phase-stabilize vehicle bending and slosh.
>
> Disturbance compensation and load relief: Generalization of Ares I-X Anti-Drift/Load-Relief algorithm designed to cancel external moments on the vehicle, minimize lateral drift, and alleviate wind loading. This is frequently referred to within the SLS Program as the DCA.

Control allocation: Allocates the desired control command to the engine TVC actuators and includes logic to accommodate engine-out, closed-loop throttling, and staging events. Often referred to as optimal control allocation (OCA).

AAC: The most novel addition to the SLS control architecture. It is designed to have minimal effect unless an extreme environmental or model dispersion is present, but modifies the total attitude control system response to provide the gain-scheduled control architecture with additional robustness when needed. The AAC increases performance when excessive tracking error is present and decreases the responsiveness when undesirable frequency content is observed in the control command as a result of internal dynamics (i.e., flexibility, fuel slosh, and actuators).

TVC actuators: SLS will use Shuttle heritage TVC actuators for the core-stage RS-25 engines and the SRBs. Understanding the dynamic performance of the actuation system is critical for any aerospace control analysis. As it relates to AAC, the actuator variation in phase lag at the crossover frequency is especially important. SLS models predict an acceptably low value of approximately 0.5 degrees (deg) of phase lag variability at the crossover frequency.

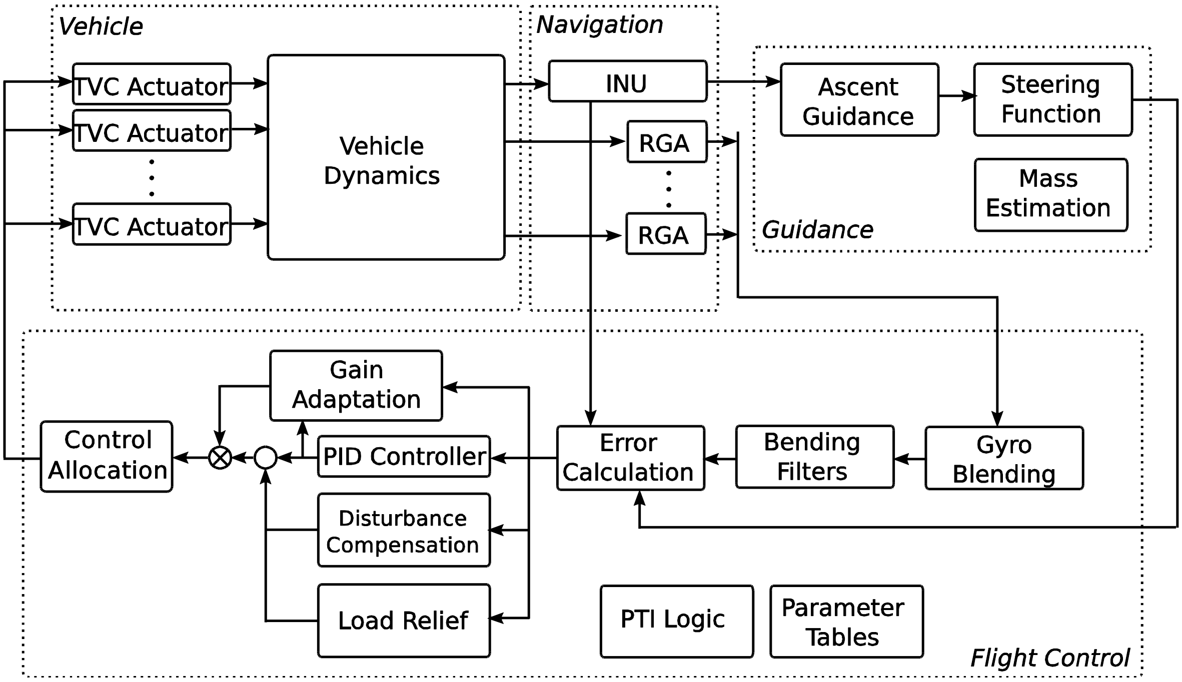

Figure 1. Simplified Block Diagram of SLS FCS (DAC-2)

3 AAC ARCHITECTURE

The AAC algorithm relies on the well-tuned gain-scheduled architecture for a vehicle and trajectory within the design envelope. When the output of the gain-scheduled architecture deviates from the expected (i.e., reference model) response, AAC modulates the performance in real time in a manner that balances attitude tracking with the mitigation of undesirable frequency content in the control path. The MSFC Flight Mechanics and Analysis Division developed the AAC algorithm to increase crew safety and vehicle survivability in the presence of mismodeled dynamics or in-flight anomalies by expanding the envelope under which the FCS is capable of reliably controlling the vehicle.

The AAC algorithm was designed with three primary objectives in mind:

1. Minimally adapt when the baseline control system is performing acceptably.
2. Increase performance and command tracking when extreme off-nominal conditions and disturbances produce large errors.
3. Decrease the system gain to prevent high-frequency content in the control loop from driving the system to instability.

Figure 2 shows the regions of adaptation and illustrates the idea behind the three objectives of AAC.

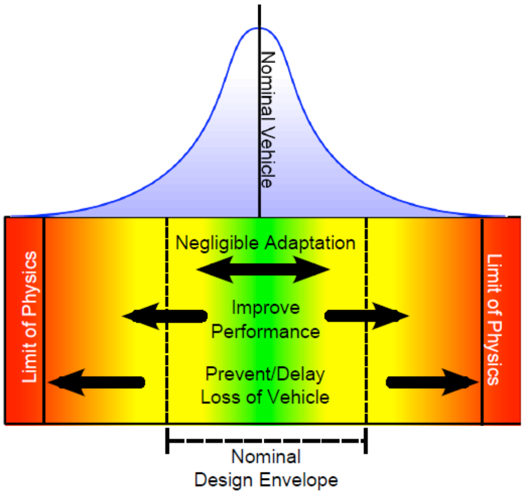

Figure 2. Conceptual Regions of Operation for Adaptive Augmentation

The SLS AAC uses sensed data to adjust the controller responsiveness on-line (see Figure 3). It increases responsiveness when the SLS response is sluggish (i.e., it does not match the reference model), which typically occurs at a lower frequency than the rigid-body gain crossover. AAC decreases responsiveness when high-frequency content is observed in the control command, typically attributed to flexible motion, fuel slosh, or actuator saturation, which occurs at a higher frequency than the rigid-body gain crossover.

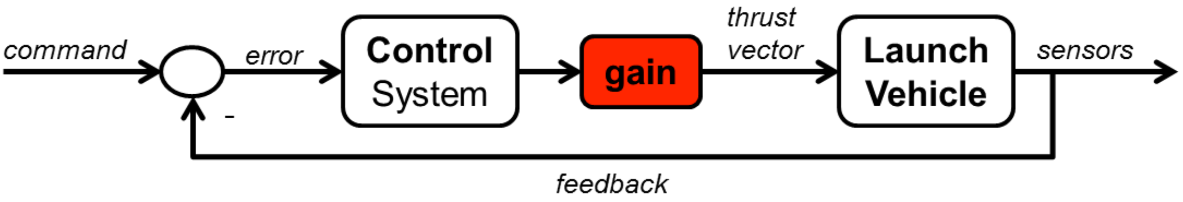

Figure 3. Simplified Vehicle-Control Interaction Diagram with Gain to be adjusted by AAC

The AAC architecture during Design Analysis Cycle (DAC)-1 and DAC-2 was as follows:

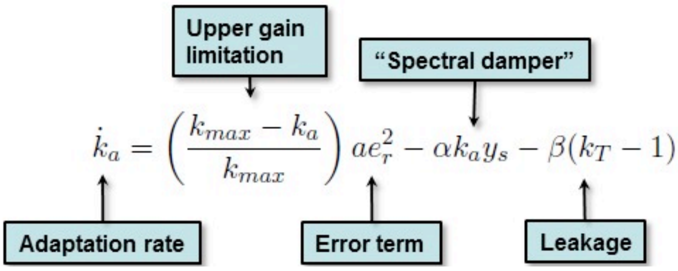

$$\dot{k}_a = \left(\frac{k_{max} - k_a}{k_{max}}\right) a e_r^2 - \alpha k_a y_s - \beta(k_T - 1)$$

The spectral damper output signal y_s, is formed from the controller gimbal command output u_G as

$$\begin{aligned} y_{HP} &= H_{HP}(s) u_G \\ y_s &= H_{LP}(s) y_{HP}^2 \end{aligned}$$

where H_{HP} is a linear high-pass filter and H_{LP} is a linear low-pass filter. The total loop gain is formed by the sum of a fixed minimum gain and the adaptive gain

$$k_T = k_0 + k_a$$

In this version of AAC, the adaptive update law included the error-driven "up gain," the spectral damper "down gain" that is driven by presence of high-frequency content in the control loop, and a leakage term that attracted the loop gain multiplier back toward unity. The adaptive gain can be viewed as a knob that tunes the controller on-line by increasing or decreasing the responsiveness when needed and gradually returning to the response of the gain-scheduled controller response with a multiplier of $k_T = 1$ when augmentation is no longer merited. The lower limit was defined by k_0, and the upper limit was defined as $k_{\max}$. These were set to be 0.5 and 2.0 respectively, corresponding to ±6 decibel (dB) gain margin guideline. The implementation has evolved prior to and during the life of the SLS Program. The aforementioned update law was adjusted prior to the program's Critical Design Review (CDR), and changes included an alternate implementation of leakage and addition of filters that decrease the rate of adaptation and provide frequency separation between the "error term" and the "spectral damper".

4 CONTRIBUTIONS AND LIMITATIONS OF EACH ANALYSIS TECHNIQUE

Typically, launch vehicle control systems schedule the control gains *a priori* (i.e., non-adaptively), and the stability analysis techniques are based on the assumption that the entire system is linear and time invariant (LTI). This stability analysis provides the SLS Program with confidence that there is adequate system margin. While the SLS AAC algorithm was intentionally designed to augment the existing classical architecture, the introduction of an adaptive component violates the theoretical requirement (i.e., that the system be LTI) for classical stability analysis. This creates a need for the development and application of alternative analysis techniques to supplement the classical methods that are the industry standard.

Each analysis technique considered as part of the NESC-SLS assessment had something specific to contribute and offered a unique viewpoint to be considered toward the development of a holistic understanding of the algorithm stability [11]. A brief description of the analyses completed, their intended contribution, and limitations of the analyses are summarized in Table 1. They are described in more detail in subsequent sections, as indicated in the "section" column.

Table 1. Summary of Analyses Completed, Contributions, and Limitations

Section	Description	Contribution	Limitations
5	Nonlinear (Lyapunov) Stability	Provides *nonlinear* stability proof	Applies to the error-driven "up-gain" and leakage terms only Does not analyze spectral damper Requires modifications to the error-driven component of the adaptive architecture for a stability proof to exist
6	Generalized Gain Margins (GGMs) based on the Circle Criterion	Available gain margin for nominal system provides a guideline for the establishment of appropriate saturation constraints on the adaptive gain	AAC does not meet the "memoryless" assumption since it is a function of the state. AAC dynamics not considered (i.e., worst case scenario is analyzed)
7	Classical Analysis with AAC Gain Modulation	Provides insight regarding SLS robustness and gain-scheduled FCS to AAC gain modulations	AAC gains assumed to statically vary the PD gain with a nominal, reduced-order LTI model
8	Time-domain Stability Margin Assessment	Provides gain and phase margins based on high-fidelity linear models	Calculations in the time-domain completed across ascent, with various starting points so the vehicle could "fly through" brief periods of instability
9.1	Monte Carlo Results	Performance impacts of AAC across the design envelope	Time domain only Does not fully exercise the algorithm
9.2	AAC Gain Variations across DACs	AAC variation within the design envelop; impacts instantaneous gain margins	Time domain only Does not fully exercise the algorithm
9.3	Expanded Dispersions	Assesses AAC's impact on performance metrics across a wider envelope	Does not include expanded flex and TVC dispersions
10	Stressing Cases	Expansive suite of stressing cases designed to fully exercise the algorithms to assess performance and limitations	Failure scenarios could exist that have not been identified

5 NONLINEAR (LYAPUNOV) STABILITY ANALYSIS

In this aspect of the assessment, the stability of the SLS with AAC is investigated using Lyapunov theory. When adaptive control laws are developed in academia, they are typically derived from Lyapunov stability analysis. With the SLS algorithm, however, the adaptive update law was developed based on an understanding of typical adaptive control techniques, coupled with practical engineering judgment rather than directly implementing a Lyapunov-based design. Lyapunov stability analysis was considered as part of the NESC-SLS stability assessment to prove the convergence of the launch vehicle error dynamics and provide suggestions for minor algorithm modifications.

The general nonlinear Lyapunov-based Model Reference Adaptive Control theory that was developed maintains a globally asymptotically stable plant under adaptive control with bounded adaptive feedback gains. However, the existing Lyapunov analysis does not lend itself to analyzing the spectral damper, or down-gain, component of AAC; therefore, the Lyapunov stability analysis focused on the error-driven up-gain portion of the algorithm. The FCS block diagram at the onset of the assessment is shown in Figure 4 with the up-gain AAC dynamics only and the filters excluded for simplicity. This simplified block diagram reveals two deviations from an architecture that can be analyzed using Lyapunov techniques. First, the adaptive gain is multiplied by the PD control components, but a different error signal is used to drive the adaptive update law. The same signal needs to be used in both locations for a stability proof to exist. After discussions with the SLS flight controls team, it was determined that the Lyapunov stability analysis should proceed after making the change so that the multiplicative adaptive gain is first applied to the same signal that drives the adaptive update law. The second modification made prior to achieving significant analysis results was to square the error in the adaptive update law.

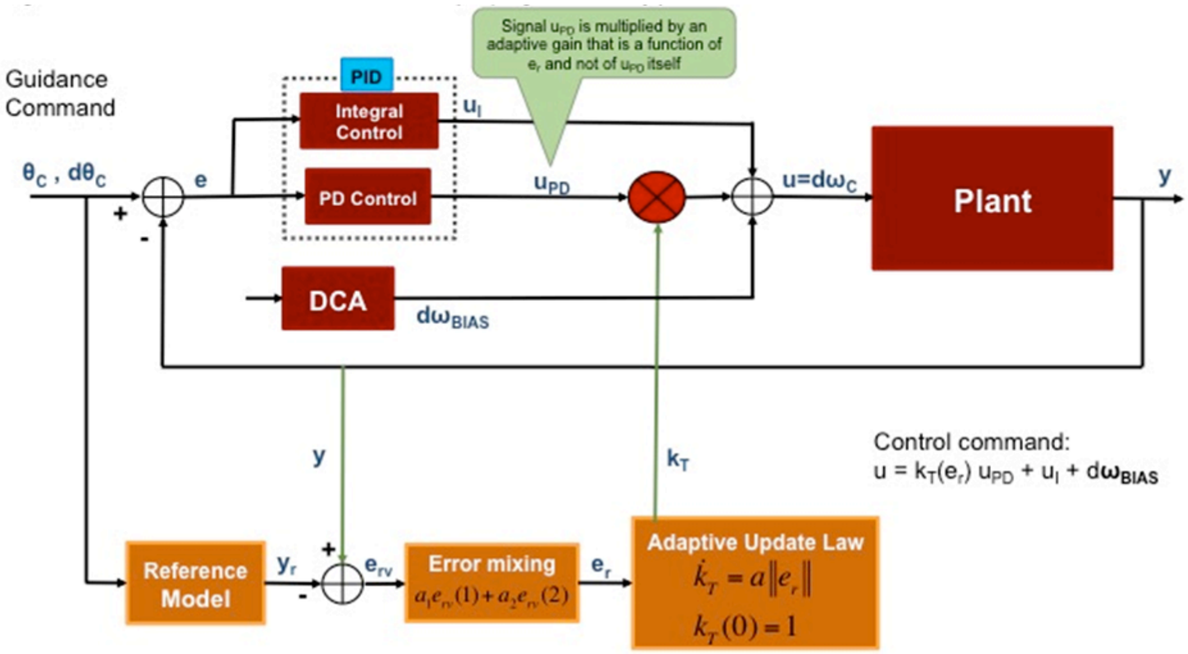

Figure 4. AAC Up-gain Architecture Prior to Assessment, Excluding Filters and Spectral Damper

Minor structural adjustments were also suggested by the assessment team to enhance the stability properties. Integral feedback control, introduced to counteract set-point changes or step disturbances, was included in the original architecture (Figure 4) in a manner that introduces a zero at the origin and violates the hypothesis of the nonlinear stability theorem. However, if this were considered through the viewpoint of adaptive disturbance mitigation, then the nonlinear stability proof would be preserved if the integral control was implemented as shown in Figure 5. This is only a slight modification to the architecture, but the deeper benefit is the applicability of a MRAC theorem, whereas the original structure does not satisfy the theorem. The updated architecture may be understood as a fixed-gain integral control or as direct adaptive disturbance rejection using the aforementioned theorem with a unity basis function. The adaptive theory can be applied to filtered signals, as long as the same filtered signal is used in both the update law and the control law. An alternate option is to include a first order filter on the output, which is equivalent to the inclusion of a "leakage" term in the update law and produces error dynamics that are ultimately bounded (i.e., a weaker stability result).

The aforementioned modifications result in the alternate block diagram that is provided in Figure 5. Note that this specific formulation provides a representation most closely matched with the existing architecture, but other forms exist for which there is a stability proof. A more common approach would be to make use of the error signal before the PD controller, but the proposed architecture in Figure 5 is most closely aligned to the SLS AAC architecture prior to the assessment.

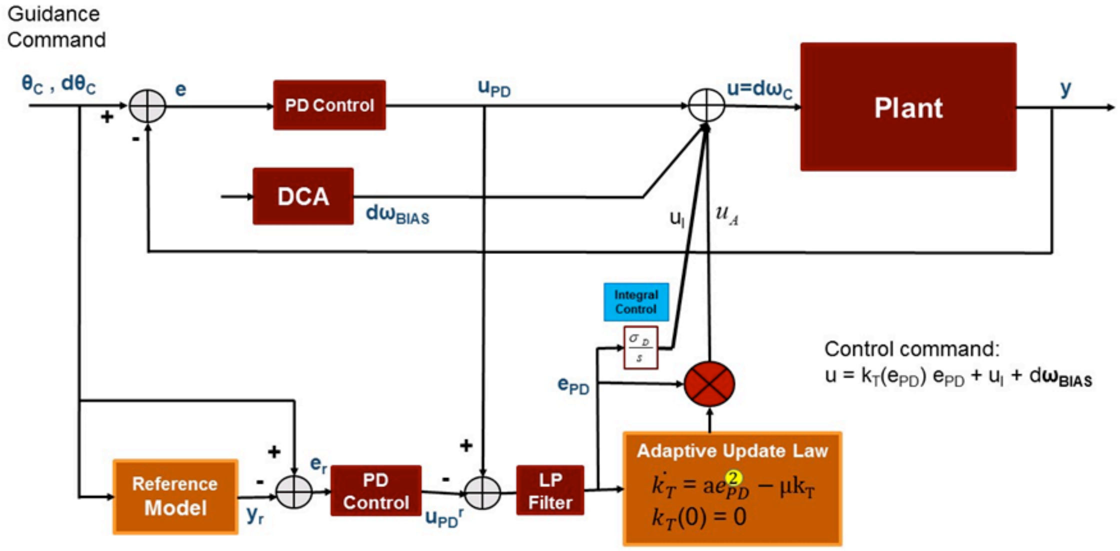

Figure 5. Alternate AAC "up-gain" Architecture with Stable Integral Control and Filters

MRAC update laws typically include reference state and reference model input terms, but their exclusion in this case is justified since the adaptive controller is meant to regulate excursions from the operating region where the fixed gain controller provides adequate control with little or no AAC contribution. Applicability of the assumptions on the launch vehicle dynamics and FCS for the stability proof to hold were evaluated and determined to be reasonable.

From a theoretical point of view, the adaptive gains are guaranteed to be bounded, but saturation constraints have not been explicitly considered in the Lyapunov analysis. The SLS FCS has saturation constraints: (1) on the position and rate control signals that relate to the physical geometric constraints and actuator limitations, respectively; and (2) on the adaptive gain that constrains the adaptive law from altering the loop gain by more than 6 dB. It is possible for the adaptive gains to remain bounded and also exceed these saturation constraints. These types of limitations should be assessed through simulations and can be mitigated, if desired, by adjusting the gain weighting in the adaptive law.

In conclusion, Lyapunov analysis is the foundational theoretical tool for proving the stability of many adaptive control systems, which are inherently nonlinear and cannot be reasonably approximated through linearization (i.e., $e^2 \rightarrow 0$ when linearized). Algorithm corrections were identified and a corresponding proof documented that can serve as a theoretical foundation for AAC's ability to drive the error exponentially to zero (i.e., stable) or to some small value (with the inclusion of leakage) even if the gain-scheduled FCS does not provide stability. This is the heart of the adaptive update law that proves the error-driven and leakage components are designed in a manner that enhances stability.

6 CLASSICAL ANALYSIS WITH AAC GAIN MODULATION

A select set of transfer functions were calculated for a simplified model of the SLS with FCS provide insight into the system's ability to follow a reference command, the effect of a load disturbance, and the sensitivity to plant uncertainty, external disturbances, and measurement noise. The investigation was completed over the full range of permissible AAC gain values (i.e., 0.5 to 2.0). This approximates the action of the AAC from a quasi-static analysis standpoint. Since a nominal launch vehicle model is considered, AAC is not acting in response to mismodeled dynamics (i.e., gain action is unnecessary, and minimal adaptation would be anticipated). Some stability margin degradation occurred in this scenario, as expected, since the gain-scheduled FCS should be optimized for this situation *a priori*. Completion of similar classical analysis with the full-scale SLS model would provide insight into the worst-case gain margin degradation, enhance understanding of the relationship between proportional and derivative gain scaling and other controller elements, and potentially provide insights into how to balance the design to minimize gain-margin losses.

7 GENERALIZED GAIN MARGINS (GGMs)

The main intent of the GGM analysis is to determine whether the existing bounds are reasonable or, if not, aid in the definition of logically derived bounds that have physical significance but do not excessively restrict the nonlinear adaptive algorithm. This is assessed from a "do no harm" perspective, where the bounds represent a worst-case gain response for the nominal system since the modeled gain variation occurs in the absence of plant model errors or extreme environmental disturbances that would typically result in AAC providing a gain change. Rather than performing a static analysis, this analysis calculated the maximum time-varying gain modulation that would be allowable for a nominal controller and launch vehicle configuration without causing instability.

Standard gain-scheduled FCSs require a minimum of ±6 dB of gain margin. This accounts for uncertainty in the launch vehicle dynamics, as conceptually depicted in Figure 6. The AAC is a complementary approach for managing uncertainties, with the capability to increase robustness outside the nominal parameter uncertainties. Limits are applied to the adaptive gains and are defined based on the amount of gain margin that is to be "allocated" to AAC. The AAC gain in each axis is nominally set to unity, but the in-flight adaptation varies this gain—and, therefore, the FCS responsiveness—between 0.5 and 2. These gain limits are approximately equivalent to ±6 dB in total loop gain variation. It would be precisely ±6 dB if the gain modulation were applied to the entire control signal. This is good for a quick and direct tieback to classical gain margin, but it does not account for the time-varying nature of the adaptive gain.

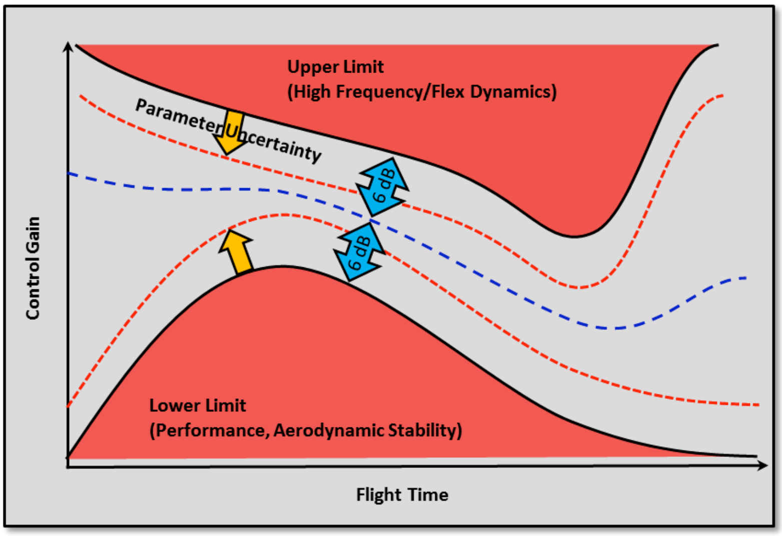

Figure 6. Fixed Gain Scheduled throughout Flight as it Relates to Parameter Uncertainty

The GGM analysis built on research that was supported through MSFC Technical Excellence [13]. The AAC was assumed to provide a time-varying, memoryless multiplicative adjustment to the total loop gain. This is in accordance with the assumptions needed to apply the Circle Criterion which provides conservative conditions that are sufficient but not necessary. Compared with the standard launch vehicle stability margins that assume fixed (non-adaptive) control gains, this comes a step closer to including the adaptive dynamics. This work was completed using a simplified launch vehicle model that included the effects of slosh and elasticity. The existing results were compared with those from an additional methodology that was based on Linear Matrix Inequalities (LMIs). The results using the simple model compared well with one another and remained within the linear gain margins, with an exception late in flight where the LMI-based margins were underconservative (see Figure 7). This exception was attributed to numerical issues associated with calculating a near-optimal value that was just outside the set of feasible margin estimates.

The previous research on GGMs [13] was applied to SLS using the systems produced by the primary frequency-domain analysis tool (i.e., FRACTAL) with the intention of providing additional SLS-specific insight as to the appropriate amount of margin to “allocate” to the AAC. The results with the full-scale model were used to inform the reasonableness of existing bounds on the adaptive gain. Continued use of this tool during design and verification cycles would provide a check that the bounds constraining AAC, which act as a fail-safe measure, remain appropriate. If future GGM analysis indicates the saturation constraints on the adaptive gain exceed the available margin, then reduced constraints should be considered. This tool has been integrated into the mainline SLS frequency-domain analysis program.

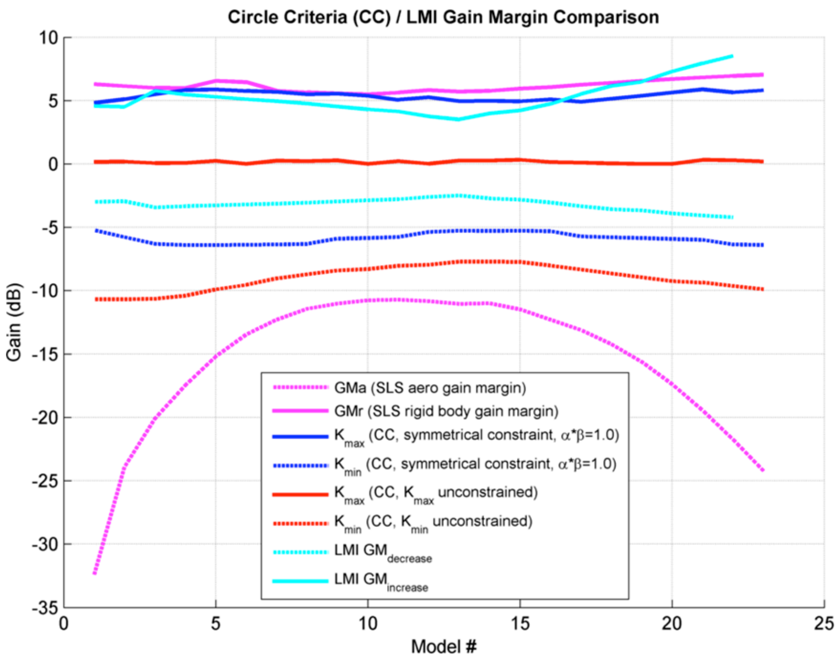

Figure 7. Linear, GGM, and LMI-based Gain Margin Results for Simplified Launch Vehicle Model

8 TIME DOMAIN STABILITY MARGINS

While the classical stability analysis with AAC gain modulation and the GGMs evaluated SLS' stability with a worst-case AAC, the intention of including AAC is to *increase* the high- and low-frequency gain margins. Theoretically, the high- or low-frequency gain margin could be increased by up to approximately 6 dB since it is permitted to adaptively adjust the gain between 0.5 and 2 times the fixed gain. The primary objective of the TDSM study was to use high-fidelity time-domain simulations to evaluate the margins by shifting the fixed-gain or phase (time delay) until instability is reached. As a check, TDSMs with AAC de-activated were compared with the gain and phase margins calculated using the program's primary frequency-domain analysis tool. Good agreement was found between the FRACTAL ("baseline") gain margins and the MAVERIC time-domain gain margins, with a difference of 0.2 dB or less. The phase margins were more challenging to match in the time domain, particularly during boost phase of flight when they were changing rapidly.

TDSMs calculated using SLS' high-fidelity time-domain simulation with and without the AAC active showed an average of 5-dB added gain margin with the inclusion of AAC. The exception to this is during the use of Programmable Test Inputs (PTIs) when the sensitivity of AAC is reduced substantially, yielding as little as 2.8 dB of additional gain margin over the gain-scheduled design. There is an accompanying slight decrease in phase margin throughout flight by approximately 5 deg against a 30-deg design guideline, but the degradation does not materialize until a severe time delay is introduced. The TDSM analysis revealed insightful sensitivities, provided an important verification of the margins reported by frequency-domain tools, and supported that the AAC behaves as intended.

9 MONTE CARLO SIMULATIONS

In the standard sets of SLS Monte Carlo simulations, AAC has minimal impact on performance for both boost and core stage flight. The minimal influence of AAC across the Monte Carlo simulation is a direct consequence of its design intent to adapt only when needed. This further indicates that the gain-scheduled SLS control law is adequate for all expected dispersion combinations. This robustness to dispersions reflects the extensive work done in areas like slosh damping requirements, bending mode requirements, and rate gyro blending. The early requirements work has rendered the SLS amenable to achieving adequate margins for expected model variations and flight conditions. For a typical SLS Monte Carlo simulation, approximately 450 system and environmental parameters are dispersed from their nominal values. This includes the dispersion of hundreds of parameters for models of propellant slosh, SRB thrust, core engine thrust, dry mass of structural and cargo elements, loaded propellant, flexible dynamics, actuator, aerodynamics, sensor error, winds, atmospheric environment, and others [16].

To more fully exercise the AAC algorithm, Monte Carlo simulations with expanded dispersions were performed. The nominal parameter dispersion magnitudes were doubled, except cases where scaling of the dispersed parameters would yield physically unrealistic subsystem model dynamics where care was taken to limit the expansions accordingly. Also, only the nominal data sets were used for the more complex flexible dynamics and the SRB thrust dispersions which were implemented based on sets of pre-dispersed data provided by the model developers.

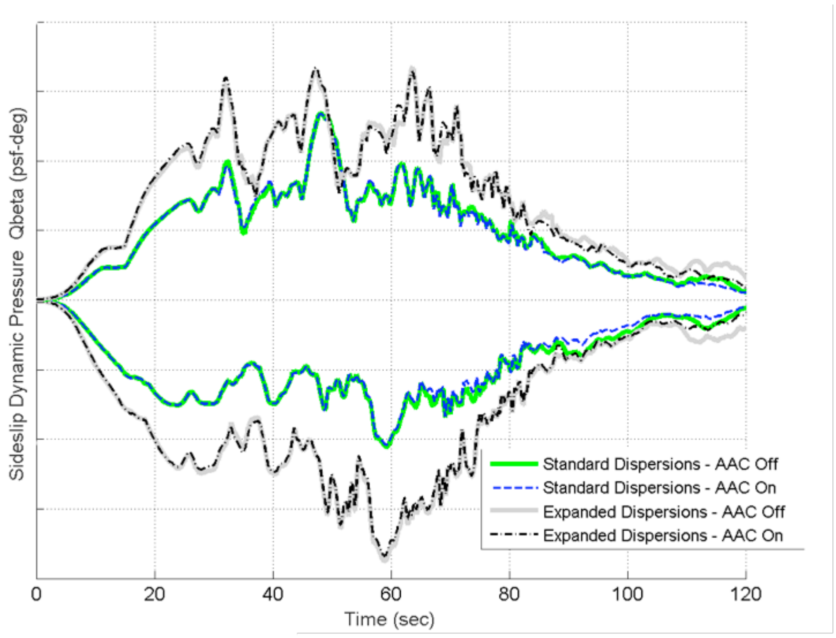

Figure 8. Bounds on Dispersed Sideslip Dynamic Pressure for Monte Carlo Simulations (Boost Phase)

The performance impact for both standard and expanded dispersions was relatively minor (see Figure 8). The notable exception to this is during solid rocket boosters (SRB) tail-off, where AAC improves the performance envelope in the yaw axis that is otherwise degraded due to potential mismatches between the port and starboard thrust tail-off prior to SRB separation.

The AAC gain activity itself was considered as an indicator of changes to the instantaneous gain margins, and exhibited variations in magnitude across DACs. The activity within the design envelope generally decreased as the algorithm matured, the FCS filters transitions improved, and the flex modeling artifacts were eliminated. However, algorithm updates made prior to Critical Design Review (CDR) resulted in larger adaptive gain variations. Therefore, additional balancing within the AAC parameter and filter design was competed as a result of an assessment recommendation to reduce gain variations within the design envelope. From a broader perspective, this meant optimizing the design in light of conflicting objectives to adapt only when needed, eliminate the risk of LCOs, and enhance the robustness to severe off-nominal conditions.

10 STRESSING CASES

More than 40 single-run scenarios were simulated to specifically test, evaluate, and demonstrate the ability of the AAC algorithm to meet its three chief objectives. Stressing cases were specially formulated and ascent trajectories developed in the mainline time-domain analysis simulation (i.e., MAVERIC). A subset of these cases may be run in the Software Integration Laboratory and potentially two other time-domain tools (i.e., SAVANT and CLVTOPS). The purpose of the suite of adaptive control system stressing cases is to:

- Demonstrate the ability of AAC to meet its three primary objectives.
- Establish AAC's overall robustness characteristics as implemented in the flight software prototype.
- Uncover scenarios in which the adaptive control system cannot recover the vehicle performance/stability to understand its weaknesses and limitations.

The stressing cases were built upon the existing nominal or dispersed trajectories and represent modest to severe off-nominal scenarios where various parameters of the plant, controller, and/or disturbances are adjusted to stress the FCS. The initial set of stressing cases was based on the F/A-18 flight test scenarios that were developed using a lower fidelity model. The set of test cases was expanded, resulting in 48 stressing cases at the time the assessment report was produced [11]. The task to build and run the suite of stressing cases was achieved through MATLAB® scripts integrated into the MAVERIC simulation. These scripts perform source and data file modifications to the baseline simulation using data-driven input instructions for each stressing case. Comparison of the vehicle response with and without the adaptive control system in each of the stressing cases demonstrated the effect of the adaptive algorithm to meet its intended objectives. Most of the stressing cases that were simulated, and certainly the cases for which AAC provides the most benefit, involve modeling deficiencies such as extreme aerodynamic environments or inaccurate elastic model predictions that represent plausible risk scenarios.

Sample test case results are provided in Figures 9 and 10, with and without AAC. In this case shown in Figure 9, a significant loss of low-frequency "aero" gain margin is simulated during core-stage flight, a condition modeled through a reduction in the controller's fixed gains. At the onset of the low-frequency instability, the AAC gain increases to a near-constant amplitude and maintains the resulting vehicle response at a small limit cycle. This is expected response, where unstable low-frequency modes involving aerodynamics or other phenomena that reduce margins below the rigid-body bandwidth are mitigated by the gain increasing action of AAC.

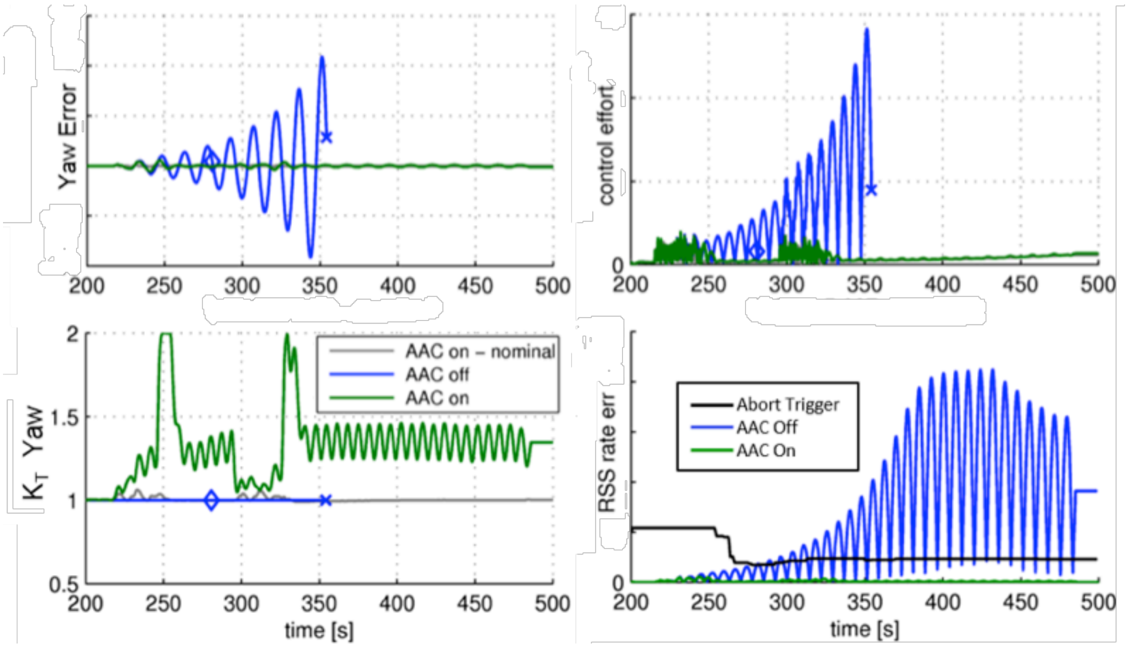

Figure 9. AAC Mitigates Low-frequency Aero Instability

Figure 10 shows an example of sustained unstable control-structure interaction instability in core-stage flight. With AAC off, the core TVC rates (upper right) reach a sustained limit cycle, which increases duty cycle (upper right) until its demise limit (denoted with ×). The rate error abort (denoted with a diamond) trigger captures this (bottom left), but future abort signal filtering will depress its effectiveness. With AAC on, the gain is decreased (bottom right), and the resulting duty cycle is drastically reduced, avoiding demise and exceeding only the duty cycle specification limit (denoted with circles) near the end of core-stage flight. If the core-stage duty cycle limit was reached twice, then this could result in demise. The attitude error and control effort (middle row) further illustrate the stabilizing effect of AAC in this stressing scenario.

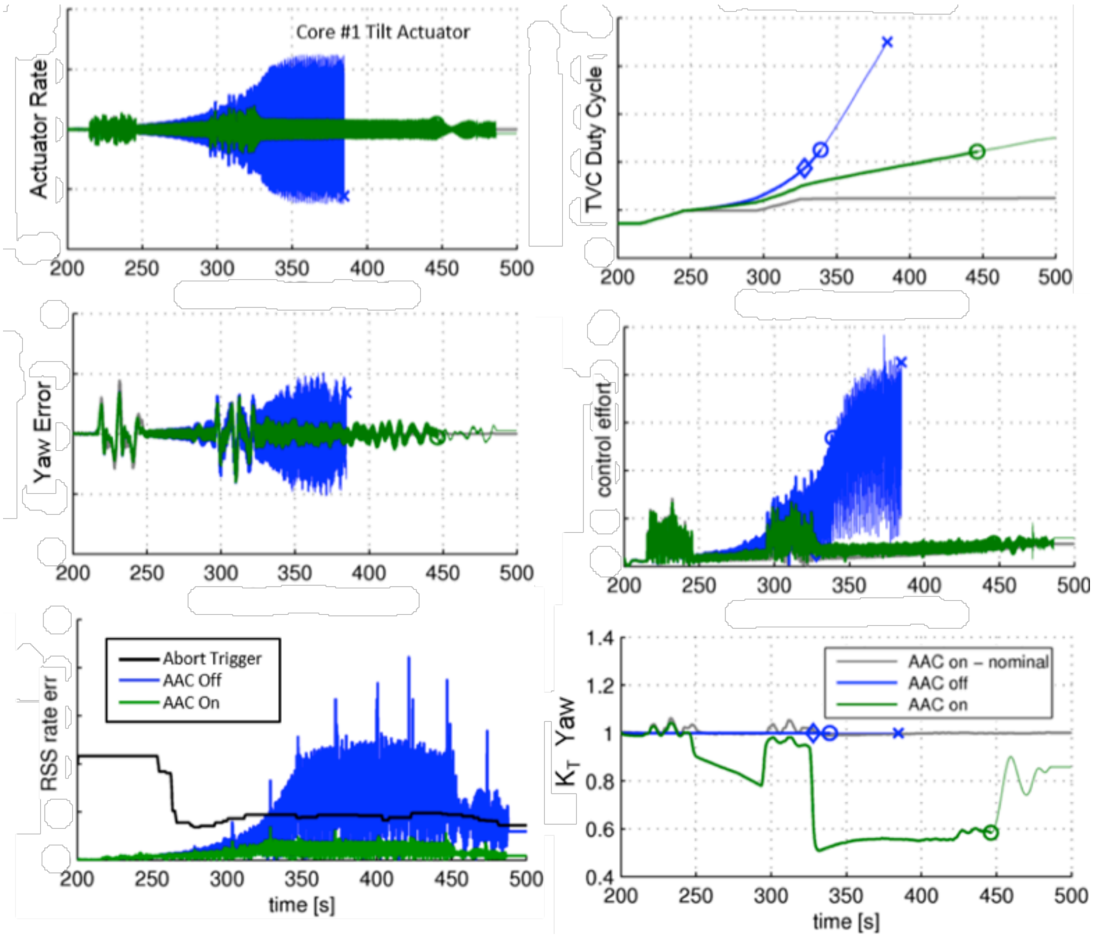

Figure 10. AAC Mitigates Sustained Unstable Control/Structure Interaction

In summary, the AAC algorithm provides additional robustness and performance for the already robust and highly performing SLS FCS in the presence of extreme off-nominal conditions. In scenarios near or beyond the ability of the vehicle to recover, AAC responds without any adverse effects. The expansive list and detailed analysis of stressing cases demonstrates the significant benefit that AAC provides to SLS control and serves as supporting evidence for its flightworthiness in EM-1, the first SLS test flight. This suite of test cases is already integrated into the mainline simulation tool, so the algorithm can be fully exercised and its performance reevaluated when there are changes to the algorithm or the vehicle configuration. It is anticipated that this analysis will be performed as a standard evaluation during each design and verification cycle.

11 CONCLUSION

In addition to the program's standard time and frequency domain analyses, the stability and robustness of the SLS FCS with AAC was assessed with Lyapunov analysis, GGMs, classical stability metrics under varying gains, TDSMs, Monte Carlos with expanded dispersions, and stressing cases. Several of the analyses focused on providing confidence that the AAC would *not harm* the system:

- **Classical stability metrics** were evaluated using a simplified model to assess the impact of scaling the FCS proportional and derivative gains between their minimum and maximum.

- **GGMs** made use of a more complex Circle Criterion analysis to assess if nonlinear gain variations in the allowable range could drive the launch vehicle to instability. Saturation constraints define the allowable range in order to mitigate potential risk if the performance of AAC during flight is not as intended.
- **Monte Carlo simulations** were used to consider AAC's impact and the gain variation within nominal and expanded dispersion envelopes. Performance characteristics were not significantly impacted even under expanded dispersions. The only visible outcome from the adaptation was the AAC gain variations themselves, which are noted due to their impact on the standard frequency-domain stability assessments.

In addition to evaluating the stability risk introduced by the AAC, its ability to *enhance* stability was considered through the following analyses:

- **Nonlinear (Lyapunov) stability analysis,** which is the foundational theoretical tool for proving the stability of many adaptive control systems. Nonlinear stability analysis was completed only for error-driven and leakage terms and omits the spectral damper component. A stability proof is documented for a modified version of the algorithm, and specific NESC recommendations were made to achieve this stability guarantee.
- **TDSMs** calculated using SLS' high-fidelity time-domain simulation with and without the AAC active, which showed an average of 5-dB added gain margin with the inclusion of AAC.
- **Stressing cases** completed in a high-fidelity SLS simulation environment demonstrated that the AAC provides enhanced robustness and performance for the SLS flight control system in the presence of extreme off-nominal conditions.

In summary, the stability and robustness of AAC was assessed from both a "do no harm" perspective as well as a "do some good" perspective. The analyses completed resulted in several recommendations regarding the design, parameter tuning, and completion of future analysis that were provided to the SLS program. Substantial benefits of the adaptive augmentation remain clear. Subject to the recommendations in the NESC Final Report [11], a nonlinear stability proof is available that ensures the error-driven and leakage components of the algorithm (two of three main elements) will adjust the FCS response in a manner that drives the system toward stability. Time-domain stability margins nearly doubled the minimum gain margins with AAC in comparison with the gain-scheduled response. The simulated response to a diverse array of stressing cases consistently demonstrated that the ACC robustness-enhancement and performance-improvement objectives were met when operating well outside the design envelope. An additional outcome of this assessment was the development of analysis software, which has been integrated into production SLS tools and provides a suite of stressing cases, circle criterion analyses, expanded Monte Carlo simulations, and time-domain stability margins (TDSMs).

12 REFERENCES

[1] Frosch, J. and Vallely, D.: "Saturn AS-501/S-IC Flight Control System Design," *J. Spacecraft*, Vol. 4, No. 8, 1967, pp. 1003–1009.

[2] Wall, J., Orr, J., and VanZwieten, T.: "Space Launch System Implementation of Adaptive Augmenting Control," in *AAS Guidance, Navigation, and Control Conference Proceedings*, Breckenridge, CO, AAS 14-051, 2014.

[3] Orr, J., Wall, J., and Hall, C.: "Space Launch System Ascent Flight Control Design," Technical Memo, Revision 2, June 2015.

[4] VanZwieten, T.: "Control of NASA's Space Launch System," *The Impact of Control Technology*, 2nd edition, T. Samad and A.M. Annaswamy (eds.), IEEE Control Systems Society, 2014, available at www.ieeecss.org.

[5] Orr, J. and VanZwieten, T.: "Robust, Practical Adaptive Control for Launch Vehicles," AIAA-2012-4549, August 2012.

[6] Hanson, C.: "Capability Description for NASA's F/A-18 TN 853 as a Testbed for the Integrated Resilient Aircraft Control Project," Tech. Rep. DFRC-IRAC-CAP-002, NASA Dryden Flight Research Center, January 2009.

[7] Miller, C.: "Nonlinear Dynamic Inversion Baseline Control Law: Architecture and Performance Predictions," in *AIAA Guidance, Navigation, and Control Conference Proceedings*, Portland, OR, 2011.

[8] Miller, C.: "Nonlinear Dynamic Inversion Baseline Control Law: Flight-Test Results for the Full-scale Advanced Systems Testbed F/A-18 Airplane," in *AIAA Guidance, Navigation, and Control Conference Proceedings*, Portland, OR, 2011.

[9] "Flight Testing of the SLS Launch Vehicle AAC Algorithm on an F/A-18," NESC Final Report, NESC-RP-13-00847, July 31, 2014.

[10] VanZwieten, T., Gilligan, E., Wall, J., Orr, J., Miller, C., and Hanson, C.: "Adaptive Augmenting Control Flight Characterization Experiment on an F/A-18," in *AAS Guidance, Navigation, and Control Conference Proceedings*, Breckenridge, CO, AAS 14-052, 2014.

[11] "Stability of the Space Launch System (SLS) Flight Control System (FCS) with Adaptive Augmentation," NESC Final Report, NESC-RP-14-00964, July 14, 2016.

[12] "Integrated Guidance Navigation and Control (GN&C) Performance Assessment, Volume 1: Ascent Performance and Stability," SLS-RPT-080-01, 2016.

[13] Zhu, J.: "A Graphical Method for Determination of the Absolute Stability Gain Margin by Circle Criterion," Final Report, Product ID: ESSSA-FY13-968, Ohio University, March 19, 2013.

[14] Mellen, D.: "The Development and Flight Test of an Adaptive Flight Control System for the X-15 Vehicle," Technical Report, Minneapolis-Honeywell, February 1963.

[15] "A Comprehensive Analysis of the X-15 Flight 3-65 Accident," NESC Final Report, NESC-RP-14-00957, September 4, 2014.

[16] "Space Launch System Program (SLSP) Six Degrees of Freedom (6-DOF) Dispersed Trajectory Analysis (TD)," SLS DRD 1406VM-005, 2015.