



Piloted Evaluation of a Control Allocation Technique to Recover from Pilot-Induced Oscillations

Diana M. Acosta*

NASA Ames Research Center, Moffett Field, California 94035

Yildiray Yildiz†

University Affiliated Research Center, Moffett Field, California 94035

Robert W. Craun‡

Mission Critical Technologies, Inc., Moffett Field, California 94035

Steven D. Beard§

NASA Ames Research Center, Moffett Field, California 94035

and

Michael W. Leonard,[¶] Gordon H. Hardy,^{**} and Michael Weinstein^{††}

Science Applications International Corporation, Moffett Field, California 94035

DOI: 10.2514/1.C032576

This paper describes the maturation of a control allocation technique designed to assist pilots in recovery from pilot-induced oscillations. The control allocation technique to recover from pilot-induced oscillations is designed to enable next-generation high-efficiency aircraft designs. Energy-efficient next-generation aircraft require feedback control strategies that will enable lowering the actuator rate limit requirements for optimal airframe design. A common issue on aircraft with actuator rate limitations is they are susceptible to pilot-induced oscillations caused by the phase lag between the pilot inputs and control surface response. The control allocation technique to recover from pilot-induced oscillations uses real-time optimization for control allocation to eliminate phase lag in the system caused by control surface rate limiting. System impacts of the control allocator were assessed through a piloted simulation evaluation of a nonlinear aircraft model in the NASA Ames Research Center's Vertical Motion Simulator. Results indicate that the control allocation technique to recover from pilot-induced oscillations helps reduce oscillatory behavior introduced by control surface rate limiting, including the pilot-induced oscillation tendencies reported by pilots.

I. Introduction

TRENDS in next-generation aircraft design relax stability requirements to gain improvements in energy efficiency and environmental compatibility corresponding to reduced fuel burn, emission production, and noise generation through reduced drag. Stability refers to an aircraft's ability to passively return to an equilibrium state (typically, original speed and orientation) after encountering a disturbance. Traditional transport aircraft are designed to be stable, but these design choices are the result of compromises in the bare airframe that degrade the aircraft's energy efficiency and environmental compatibility [1,2]. Although studies [3–6] have shown the fuel burn and emissions benefits of reduced stability or instability, work remains to ensure that the aircraft stability can be adequately

augmented through feedback control systems. As the trend in aircraft design leads to marginally stable or unstable but controllable airframes, high-control power and feedback control augmentation are required to improve flying qualities and maintain closed-loop stability. Of these two requirements, the high-control power requirement poses a challenge for next-generation transport and mobility aircraft.

Experience has shown that it is possible to stabilize unstable fighter aircraft with sufficient control power by using larger control surfaces and fast actuators. As the sizes of aircraft increase, the moment forces needed from the control surfaces also increase, resulting in the need for larger control surfaces. Moving these large control surfaces at a high rate required to meet convective stability and control design standards has been termed technologically challenging [5].

Control solutions are needed that will stabilize marginally stable and unstable airframe designs of next-generation transport aircraft. These control solutions will involve the efficient use of many multiaxis control surfaces that are optimized for minimum size such that, together, the surfaces provide sufficient control authority in each axis. Efficient use of the control surfaces will require balancing the available control authority in all axes while minimizing control surface deflections and rates. Research in recent years has begun to address these needs through the use of advanced control allocation techniques. Some of these include optimization methods for performance in the presence of control effector rate limits [7], optimization methods for desired computational speed and implementation requirements [8], investigations of the effects of control effector interactions for systems with many surfaces [9], and minimization of control surface deflections to make more control power available as surfaces approach position saturation [10]. The control allocation technique to recover from pilot-induced oscillations (CAPIO) is one of the first known techniques to address the problems that arise from stringent actuator rate limits for multi-input/multi-output applications without the use of ganging. There exist various other methods for pilot-induced oscillation (PIO) recovery and mitigation. Similar-

Presented as Paper 2013-4926 at the AIAA Guidance, Navigation, and Control Conference, Boston, MA, 19–22 August 2013; received 17 September 2013; revision received 7 April 2014; accepted for publication 23 April 2014; published online 7 October 2014. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 1542-3868/14 and \$10.00 in correspondence with the CCC.

*Aerospace Engineer, Intelligent Systems Division, Mail Stop 269-1.

†Associate Scientist, Mail Stop 269-1; currently Assistant Professor, Department of Mechanical Engineering, Bilkent University, 06800 Ankara, Turkey. Member AIAA.

‡Aerospace Engineer, Mail Stop 269-1; currently First Lieutenant, U.S. Marine Corps, Corpus Christi, TX 78418.

§Simulation Engineer, Aerospace Simulation Research and Development Branch, Mail Stop 243-1. Senior Member AIAA.

¶Senior Simulation Engineer, Mail Stop 243-6; currently Senior Member of Consulting Staff, Mentor Graphics, Inc., 46871 Bayside Parkway, Fremont, CA 94538.

**Senior System Analyst, Mail Stop 243-6.

††Senior Simulation Engineer, Mail Stop 243-6.

ities and differences of CAPIO from these methods and the advantages of CAPIO over them are discussed in previous publications [11–13]. In this paper, CAPIO is presented along with piloted simulation results of its effectiveness.

The objective of CAPIO is to enable energy-efficient and environmentally compatible next-generation aircraft with technologically achievable control surface rate limits. To do this, CAPIO seeks to allow aircraft to fly within a nominal flight envelope that includes cases when control surfaces are functioning at their rate limit. Traditionally, a nominal flight envelope will not include cases when control surfaces are functioning at their rate limits since this leads to phase lag associated with pilot-induced oscillations [14]. CAPIO actively detects and eliminates phase lag introduced by control surface rate limiting. By doing so, CAPIO allows the pilot to maintain or regain closed-loop control of the aircraft.

This paper describes the first human-in-the-loop implementation and evaluation of CAPIO accomplished through a motion-based simulation. Specifically, the pilot–aircraft–control system was evaluated using a demanding task to observe the systems characteristics, including performance with respect to task performance criteria, handling qualities ratings, and pilot-induced oscillation ratings. Section II describes the PIO phenomenon. Section III follows with an overview of CAPIO as engineered to execute in real time on the nonlinear aircraft simulation system. The simulation evaluation is described in Sec. IV, and Sec. V presents and discusses the results from the evaluation.

II. Pilot-Induced Oscillation

A PIO is a sustained or uncontrollable, inadvertent oscillation resulting from the pilot's efforts to control the aircraft [15]. The pilot reacts to the motion of the aircraft, creating a closed-loop feedback control system. The oscillations can therefore be identified as closed-loop instabilities of a feedback control system [16].

During a PIO, there are phase lags between the pilot's commands and the aircraft's responses. A typical PIO is characterized as "an oscillation at a frequency where the attitude response lags the stick inputs by approximately 180 degrees" [17]. The onset of a typical PIO caused by rate-saturated actuators can be recognized in the conceptual diagram in Fig. 1, which depicts the phase lag between a pilot's control signal and a control surface response. Although PIOs can be easily identified during postflight data analysis, pilots often do not know they are in a PIO; from their perspective, the aircraft appears to have broken [17].

Three contributors interact to induce a PIO; these contributors are the aircraft, the pilot, and the trigger. The aircraft can contribute to a PIO by having unstable or marginally stable modes or by having quasi-linear and nonlinear characteristics. An aircraft's quasi-linear and nonlinear characteristics can be the result of lags in the control system, stick command shaping, mode switching in the software, actuator rate and position limiting, aerodynamic properties, or a combination of these factors [16]. Pilots help sustain the oscillations through changes in their dynamic behavior: for example, by inputting higher than normal gain to the control system [16]. Finally, the piloted task or trigger event is the impetus for changes in the pilot or vehicle dynamics and starting the PIO [16].

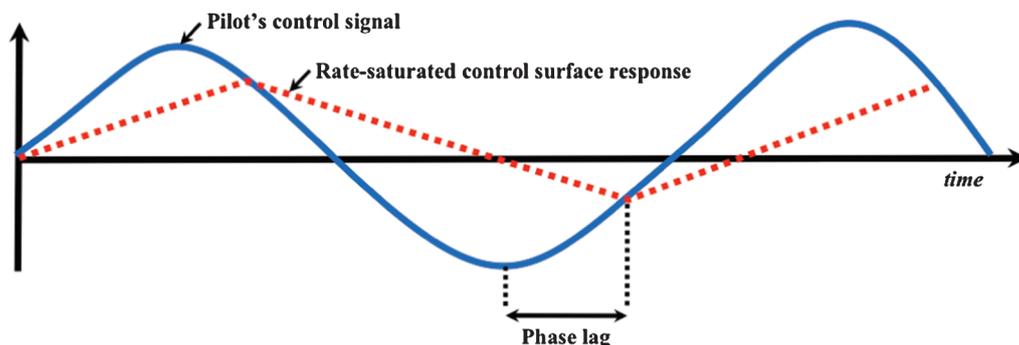


Fig. 1 Diagram depicting phase lag between a pilot's control signal and a rate-saturated control surface response.

III. Control Allocation Technique to Recover from Pilot-Induced Oscillations

CAPIO is designed to assist in the recovery from category II PIOs, which occur when control surfaces are operating at their maximum rate limit. Traditional control allocation techniques are susceptible to these category II PIOs due to the phase lag between the desired and achieved rotational accelerations during rate limiting. CAPIO augments traditional control allocation techniques and seeks to reduce this phase lag and the associated risk of category II PIOs.

Traditional control allocation systems seek to command control surface deflections to satisfy the desired rotational accelerations, denoted herein as v_d , from the stability and control augmentation system. The allocator commands a total control effort v , also referred to as the assumed achieved rotational acceleration in roll, pitch, and yaw. For the purposes of the control allocation derivation, actuator dynamics are neglected and control surfaces are viewed as pure moment generators. These assumptions lead to the following approximate aircraft model:

$$\dot{x} = Ax + B_u u = Ax + B_v v \quad (1)$$

$$v = B_u u \quad (2)$$

where x is the aircraft state vector of length $n \geq 3$, u is the control surface deflection vector of length $m > 3$, $B_u = B_v B$, and

$$B_v = \begin{bmatrix} I_{3 \times 3} \\ 0_{(n-3) \times 3} \end{bmatrix}$$

Given the desired rotational accelerations requested by the control system v_d , the control surface deflection is a solution to the system

$$B_u u = v_d \quad (3)$$

which is underdetermined. Standard noniterative techniques can be used to solve Eq. (3) for a unique minimum-norm solution (see, e.g., [18]). The imposition of bound constraints to ensure that the solution is within magnitude and rate limits, however, requires that the determination of the commanded control surface deflections be recast as an optimization problem.

The objective function of the optimization problem for the traditional control allocation system is represented by

$$J = \|W_p^{1/2}(Bu - v_d)\|_2^2 + \epsilon \|u\|_2^2 \quad (4)$$

where the operation

$$\|\bullet\|_2^2$$

is the square of the Euclidean norm, $\epsilon > 0$, and W_p is a weighting term. Simple bounds on u are imposed to ensure the solution is within magnitude and rate limits. The solution to this optimization problem u^* becomes the minimum-norm solution to the system $Bu = v_d$ as ϵ approaches zero. Whereas the solution u^* represents the desired

actual control surface deflection, nonlinear actuator dynamics represented by the actuator model $a(\bullet)$, which would require the determination of u_c such that $a(u_c) = u^*$, are ignored and the commanded control surface deflection u_c is set equal to u^* .

Under this allocation scheme, control surfaces functioning at their rate limit can lead to the assumed achieved rotational acceleration falling out of phase with the desired rotational acceleration. For example, this can happen when the desired rotational acceleration peaks and changes from increasing to decreasing, whereas the assumed achieved rotational acceleration may still be increasing to match the desired rotational acceleration (similar to Fig. 1). In that case, the derivatives of the two signals have opposite signs and the result perceived by the pilot is a sluggish response.

CAPIO, in addition to minimizing the error between v (or Bu) and v_d , seeks to reduce phase lag by minimizing the error in the derivative of these signals. The objective function for CAPIO is represented by

$$J = \|W_p^{1/2}(Bu - v_d)\|_2^2 + \|W_D^{1/2}(B\dot{u} - \dot{v}_d)\| + \varepsilon\|u\|_2^2 \quad (5)$$

where the operation

$$\|\bullet\|_2^2$$

is the square of the Euclidean norm, $\varepsilon > 0$, and W_p and W_D are weighting terms. Similar to the optimization problem for the traditional control allocation system, simple bounds on u are imposed to ensure the solution is within magnitude and rate limits, and the solution to the control allocation problem, the commanded control surface deflection u_c , is set equal to u^* .

The two weights, W_p and W_D , provide a manner in which CAPIO balances the objectives of having the control system follow the acceleration commands and their derivatives. In the application presented in this paper, W_p is always a positive constant and W_D is set either to zero or a positive constant value to disengage or engage the derivative following behavior of CAPIO. These two weights can also be used for axis prioritization.

Within CAPIO, two subsystems observe the signal dynamics of v and v_d to detect phase lag in real time and determine when to engage or disengage derivative following. These subsystems are described in the next two subsections. The commanded control surface deflections, which are solutions to the objective function in Eq. (5), are calculated online by a real-time optimizer. This real-time optimization is also described next.

A. Real-Time Detection of Phase Lag

Real-time detection of oscillatory behavior was designed into CAPIO to indicate when phase lag was present and needed to be addressed. The detection scheme used within the CAPIO system was a real-time reimplementation of the concept Mitchell and Hoh [19] developed for the PIO detector called ROVER (real-time oscillation verifier). ROVER detects PIOs by analyzing the smoothed signals of angular rate and pilot control input to identify three key characteristics of a PIO. Those characteristics are an oscillation frequency of the angular rate signal within the range of a PIO, large peak-to-peak amplitudes for both the signals, and large phase differences between the two signals [20]. The reimplementation of the ROVER concept described here considers the latter two characteristics. Since the

purpose of real-time detection of phase lag is intended to detect the early onset of phase lag and possible PIO, modest peak-to-peak amplitudes for the signals of interest are acceptable.

Phase lag is detected based on the assumed achieved rotational accelerations v and the desired rotational accelerations v_d . The calculated phase lag represents the phase lag apparent to the control allocation system, as seen in Fig. 2, where v is calculated as $v = Ba(u_c)$ from the commanded control surface deflection u_c .

Each of these signals are analyzed to identify peaks and compared to calculate the approximate phase difference. Peaks are identified by looking for zero crossings in the derivatives of the acceleration signals. The signals v and v_d are not filtered before calculating \dot{v} and \dot{v}_d to avoid introducing additional lag. To avoid the detection of false peaks due to noise, the magnitude of acceleration at a new peak is required to be outside of a deadband around the previous detected peak magnitude to be counted as a real peak. Additionally, peaks in the signal v must follow a peak in the signal v_d to ensure all calculated phase differences are positive.

The phase difference between v and v_d is calculated following the identification of a peak in v . To calculate the phase difference, the system assumes the signals exhibit sinusoidal behavior and have the same frequency. The frequency is calculated by treating the difference in time between the last two peaks of v_d as half of a period:

$$f = 2\left(\frac{1}{t_{v_d2} - t_{v_d1}}\right) \quad (6)$$

where t_{v_d1} and t_{v_d2} are the times of the last two peaks in v_d , as shown in Fig. 3. Using this frequency f and the time Δt between the latest peaks of v and v_d , the phase difference $\Delta\varphi$ is calculated as

$$\Delta\varphi = 360f\Delta t = 360\left(\frac{1}{2(t_{v_d2} - t_{v_d1})}\right)(t_v - t_{v_d2}) \quad (7)$$

where t_v is the time of the latest peak in v .

This is done for the rotational accelerations in all three body-fixed axes. If the value of the phase difference in any axis is above a preset threshold, corresponding flags are sent to the subsystem of CAPIO that engages and disengages derivative following to indicate a significant phase lag is present in the system.

B. Derivative-Following Engagement and Disengagement

CAPIO balances two objectives: having the control system follow 1) the acceleration commands and 2) their derivatives. The derivative following behavior is engaged and disengaged through updates to W_D .

In theory, the derivative following mode is engaged any time a significant phase lag is detected and then disengaged when the phase lag is no longer significant, based on the flag input from the real-time phase lag detector. The engagement and disengagement of the derivative following mode are done for each of the rotational axes individually, so significant phase lag in roll acceleration engages derivative following for the roll axis, leaving the pitch and yaw axes to follow accelerations. The derivative following mode causes W_D to increase from zero to a preset value that is determined by tuning. When W_D is sufficiently large, the second term in the objective

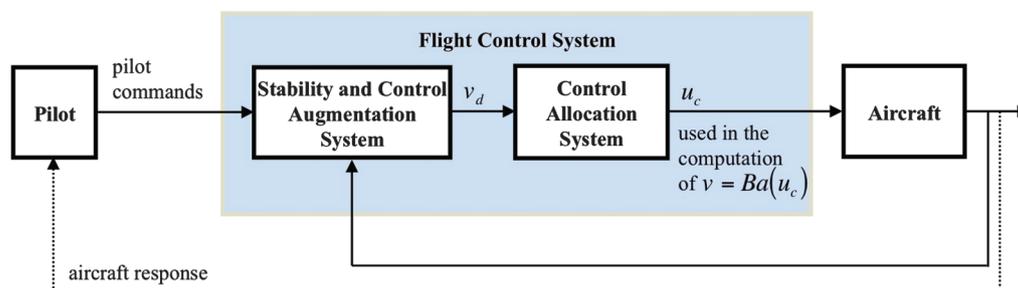


Fig. 2 Conceptual system diagram.

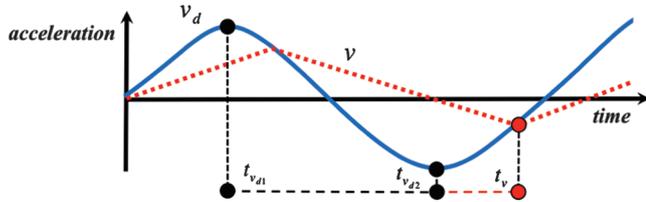


Fig. 3 Phase lag detection diagram.

function [Eq. (5)] becomes dominant, which forces the derivative of the assumed achieved accelerations to follow the derivative of the desired accelerations and eliminate the phase lag. When the phase lag drops below the preset threshold, W_D is set to zero.

In practice, there are certain situations when derivative following is not desired despite the detection of significant phase lag. The next scenarios outline exceptions used by CAPIO in determining when to engage and disengage derivative following.

1. Scenario One

One situation where CAPIO will disengage derivative following is when the desired acceleration levels off at a value in an asymptotic manner. Since no new peak is detected, a new phase difference cannot be calculated and derivative following will remain on, possibly with a large steady-state error. To alleviate this, CAPIO looks at \dot{v}_d and \dot{v} , and if the difference is below a threshold, derivative following is disengaged in that axis. This situation is illustrated in Fig. 4.

2. Scenario Two

Another situation where CAPIO will disengage the derivative following exists when the stability and control augmentation system imposes maximum (or minimum) limits on the desired accelerations. If the desired acceleration is at a maximum (or minimum), the signal will flatline and due to derivative following, the assumed achieved acceleration would also flatline. This is a problem because the pilot will perceive the flat response as a loss in control power when requesting full control authority. CAPIO, therefore, will disengage derivative following if the desired acceleration is at a maximum (or minimum), as seen in Fig. 5.

3. Scenario Three

The final situation when CAPIO will not allow derivative following despite a flag from the phase lag detection subsystem is when the desired and assumed achieved accelerations have different signs. As illustrated in Fig. 6, if the desired acceleration is positive and its derivative is near zero while the assumed achieved acceleration is still negative, the aircraft will accelerate in a direction opposite of the pilot's intentions. Pilots may interpret this response as

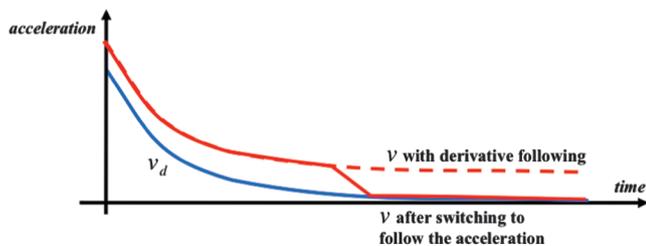


Fig. 4 Situation when desired acceleration v_d levels off in an asymmetric manner.

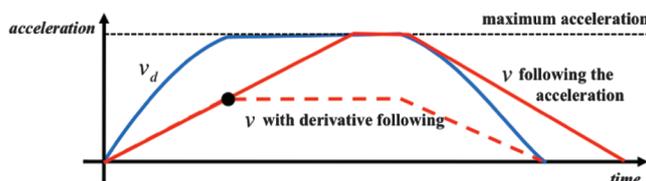


Fig. 5 Situation when desired acceleration v_d is at a maximum limit.

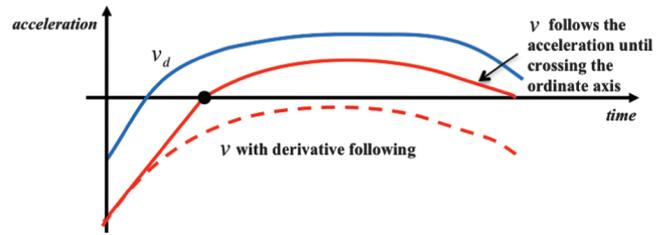


Fig. 6 Situation when desired acceleration v_d and assumed achieved acceleration v have different signs.

a loss of control. To alleviate this problem, CAPIO will only allow derivative following when the desired and assumed achieved accelerations have the same sign.

C. Real-Time Optimization

CAPIO calculates the commanded control surface deflections using a real-time iterative optimization algorithm. Real-time optimization helps to ensure the efficient use of many multi-axis control surfaces such that, together, the surfaces provide sufficient control authority in each axis while balancing between v following v_d , and \dot{v} following \dot{v}_d , with respect to the weighting matrix W_D .

For use in the real-time simulation environment, the optimal solution is computed at discrete times t_k . In this formulation, the variable u_k denotes the unknown column vector of desired actual control surface deflections at time t_k , and $\Delta t = t_k - t_{k-1}$. After eliminating the constant terms, the discrete objective function, based on Eq. (5), is given by

$$q(u_k) = \frac{1}{2} u_k^T H u_k + c^T u_k \quad (8)$$

where the matrix H and the vector c are determined to be

$$H = 2(\Delta t^2 B^T W_P B + B^T W_D B + \epsilon I) \quad (9)$$

$$c = -2(\Delta t^2 v_d^T W_P B + u_{k-1}^T B^T W_D B + \Delta t \dot{v}_d^T W_D B) \quad (10)$$

when \dot{u} is approximated by $\dot{u} \approx (u_k - u_{k-1})/\Delta t$ and \dot{v}_d is calculated as $\dot{v}_d \approx (v_{d,k} - v_{d,k-1})/\Delta t$. The solution u_k is the minimizer of $q(u_k)$ subject to bounds on u_k that ensure the solution is within magnitude and rate limits. These bounds take into consideration the actual achieved actuator positions at time t_{k-1} . If known, the actual control surface deflection can be sent as feedback to the control allocator. Otherwise, the actual control surface deflection can be approximated by a nonlinear function a of the actuator model as $a(u_{k-1})$. Whereas the solution u_k represents the desired actual control surface deflection, nonlinear actuator dynamics represented by the actuator model are ignored and the commanded control surface deflection u_c is set equal to u_k .

The primary challenge with real-time optimization is to ensure satisfactory termination of the iterative process within the allowable timeframe. Several steps are taken to meet this challenge. First, a first-derivative method is chosen in order to limit the complexity of each optimization iteration. Second, within the optimization method, advantage is taken of the fact that H is positive definite by using an exact line search (see, for example, [21]) instead of an iterative search. Third, explicit formation of H and c is avoided in the computation of $q(u_k)$ and $\nabla q(u_k)$. Instead, these computations are economized to take advantage of the outer product formulation of H and c .

The minimization of a quadratic function with positive definite Hessian can be accomplished by a variety of first-derivative optimization methods, though the stipulation of bounds narrows the field of choices. The optimization method selected was the limited-memory Broydon-Fletcher-Goldfarb-Shanno bound-constrained (L-BFGS-B) algorithm (see [22]). An implementation of this method, available as open-source Fortran code [23], was integrated into the CAPIO

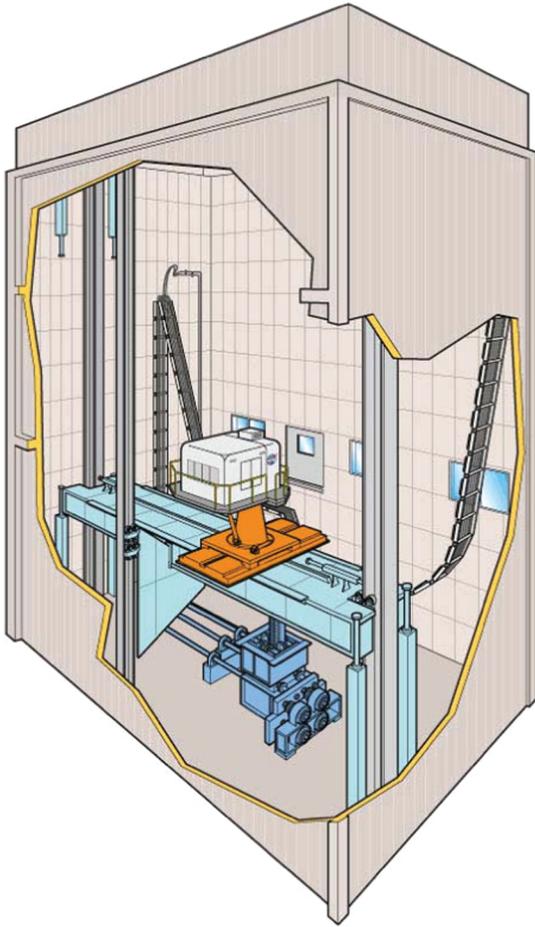


Fig. 7 Vertical Motion Simulator.

algorithm. Only minor modifications were required to adapt L-BFGS-B to this context: L-BFGS-B was modified to perform an exact line search whenever possible, with a reduced number of function evaluations allowed before termination of the search, and a preset limit on the number of function evaluations was observed. At each frame, the optimizer used the solution of the last problem as its initial point. For more details on the implementation, see [24].

With this implementation, the simulation of the research aircraft math model with CAPIO was executed with an average increase in computational time per frame of only 17% relative to the math model without CAPIO. This outcome was sufficient to allow real-time execution of the CAPIO system at the required frame rate.

IV. Simulation Evaluation Description

The simulation evaluation was intended to assess the viability of CAPIO to enable reduced control surface rate requirements for energy-efficient next-generation aircraft and identify areas of future research for the maturation of the technology. As such, the evaluation was designed to meet two objectives as follows:

- 1) Demonstrate that CAPIO does not degrade system characteristics (measured by task performance criteria, handling qualities ratings, and pilot-induced-oscillations ratings) of a nominal aircraft.
- 2) Demonstrate that CAPIO improves system characteristics (measured by task performance criteria, handling qualities ratings, and/or pilot-induced-oscillations ratings) of an aircraft with reduced actuator rate limits when compared to a baseline control system.

Since the derivative following characteristics of CAPIO are activated after phase lag is detected, the point of interest is where flight control system commands to the aircraft exceed the control surface rate limits. To support this interest, the simulation evaluation required realistic motion cues for the pilots, a representative next-generation research aircraft model, extreme research aircraft configurations, and a demanding task. Details on these attributes, as well as the experimental procedure and data collected, are described next.

A. Vertical Motion Simulator Facility

The Vertical Motion Simulator (VMS) is the ideal facility to test the CAPIO system because of its large motion envelope. Schroeder and Chung [25] concluded that larger simulator motion envelopes provide more accurate handling qualities ratings (HQRs) and PIO ratings than smaller simulator motion envelopes when compared to the same ratings taken in the actual aircraft. Schroeder and Chung also found that a large motion simulator was the only platform that induced markedly divergent PIOs. Additionally, pilots gave large motion higher confidence factor ratings and achieved lower touchdown velocities compared to small motion simulators.

The VMS motion system is an uncoupled six-degree-of-freedom motion simulator, shown in Fig. 7. It is located in, and partially supported by, a specially constructed 120 ft tower. The VMS system motion capabilities are provided in Table 1 [26]. Included in the table are two sets of limits: system limits that represent the absolute maximum levels of attainable under controlled conditions, and operational limits that represent attainable levels for normal piloted operations.

The cab serves as the aircraft cockpit. The evaluation pilot occupied the right seat, with the test engineer in the left. A computer image generation system creates the out-the-window visual scene for the six-window collimated display, with the head-up display superimposed on the center window. Additional aircraft information was provided on three head-down displays at both pilot stations.

B. Research Aircraft Model

The research aircraft math model flown in this evaluation was the Speed Agile Concept Demonstrator (SACD), which is a short takeoff and landing mobility concept being developed by industry under the U.S. Air Force Research Laboratory's (AFRL) Advanced Joint Air Combat System studies. The SACD program seeks to mature technology in the areas of high lift, efficient transonic flight, and flight control for future integrated mobility configurations that are intended to carry larger, heavier payloads than the C-130; fly efficiently across a wide range of speeds; cruise above Mach 0.8; and routinely operate from short, unprepared runways [27]. As part of the SACD program sponsored by the AFRL, an aircraft math model of the SACD and corresponding flight control system were developed and delivered to NASA.

The flight control system for the SACD model is a full-authority control system that accepts pilot inputs and calculates desired control

Table 1 VMS motion system performance limits [26]

Degree of freedom	Displacement		Velocity		Acceleration	
	System limits	Operational limits	System limits	Operational limits	System limits	Operational limits
Longitudinal	±4 ft	±4 ft	±5 ft/s	±4 ft/s	±16 ft/s ²	±10 ft/s ²
Lateral	±20 ft	±15 ft	±8 ft/s	±8 ft/s	±13 ft/s ²	±13 ft/s ²
Vertical	±30 ft	±22 ft	±16 ft/s	±15 ft/s	±22 ft/s ²	±22 ft/s ²
Roll	±0.31 rad	±0.24 rad	±0.9 rad/s	±0.7 rad/s	±4 rad/s ²	±2 rad/s ²
Pitch	±0.31 rad	±0.24 rad	±0.9 rad/s	±0.7 rad/s	±4 rad/s ²	±2 rad/s ²
Yaw	±0.42 rad	±0.24 rad	±0.9 rad/s	±0.8 rad/s	±4 rad/s ²	±2 rad/s ²

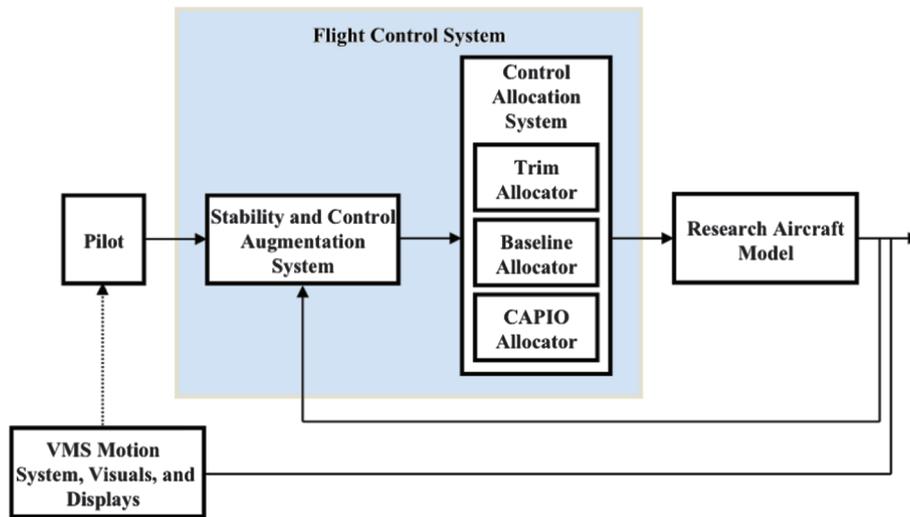


Fig. 8 Flight control system diagram.

surface deflections commands in order to stabilize, trim, and maneuver the aircraft. The flight control system architecture comprises a stability and control augmentation system (SCAS) in series with a control allocation system, as shown in Fig. 8. The SCAS calculates rotational acceleration commands in the aircraft's body-fixed axis. Next, these commands are passed to the control allocation system, which calculates the necessary control surface deflections. The separation of the SCAS from the control allocation effectively decouples the two systems, allowing the two designs to mature independently such that changes can occur in one system without necessitating the redesign or retuning of the other system. For more information on the aircraft concept and flight control system, see [28,29].

For this simulation evaluation, the SCAS remained identical to the original system delivered and only the control allocation system was modified. The control allocation system is composed of three allocators with different functions operating in parallel. The three allocators are the trim allocator, baseline allocator, and CAPIO allocator. The purpose of the trim allocator is to determine the control surface deflections needed to achieve the trim commands. The purpose of both the baseline allocator and the CAPIO allocator is to determine the control surface deflections needed for stabilization and maneuvering. These latter two allocators are used interchangeably or together, depending on a mode specified by the user. For example, the baseline allocator can be responsible for calculating the deflections needed for commands in all three axes, whereas the CAPIO allocator is responsible for none, or vice versa; or, the baseline allocator can be responsible for calculating the deflections needed for commands in one body-fixed axis, whereas the CAPIO allocator is responsible for the remaining two body-fixed axes, or vice versa. Output signals from the trim allocator, baseline allocator, and CAPIO allocator are merged and sent as one set of control surface deflection commands to the research aircraft math model.

C. Research Aircraft and Flight Control System Configurations

To support the research objectives of the evaluation, four configurations of the research aircraft and flight control system were flown. The factors that differentiated the four configurations were the control allocator used and the rate limiting imposed on the control surfaces.

To recognize the impact CAPIO had on system characteristics, two versions of the control allocator were flown. The first version employed the baseline allocator in all axes and served as the reference for measuring performance. The second version, referred to as the CAPIO allocator, employed the CAPIO allocation algorithm in the pitch and roll axes, with the baseline allocator operating in the yaw axis. Roll was the axis of interest in the piloted simulation task described in the following subsection.

The control surface rate limits were set at a nominal value and a reduced value for actuators in the axis of interest for the task. The reduced actuator rates were chosen such that actuators would encounter their limits to accomplish the task. This would introduce phase lag that would induce PIO tendencies, thus providing sufficient differences in system characteristics between the nominal and reduced rate-limited aircraft configurations.

Table 2 summarizes each of the resulting sets of configurations for each task. Other factors, such as the atmospheric conditions, initial location, and trim airspeed, remained consistent across each of the task's four configurations.

D. Piloted Simulation Task

The pilot-in-the-loop task was designed and deployed in this simulation to engage the pilot-aircraft system in high-gain precision maneuvers to expose any PIO tendencies in the system in a controlled and repeatable manner. The task was an offset approach and landing task.

The offset approach and landing task was designed to assess handling qualities and reveal PIO tendencies in the lateral axis under tight, aggressive control. To accomplish this objective, pilots were asked to land the airplane given drastic situational circumstances and demanding performance standards. The task began with the aircraft configured in a trimmed approach descending through cloud cover, which obscured the pilot's view of the runway. While the glideslope and localizer indicated a proper approach, the aircraft was offset 300 ft laterally and 300 ft longitudinally; this scenario represented a simulated navigational equipment failure. At 230 ft above ground level, the aircraft broke out of the cloud ceiling and the pilot had to locate the runway. The pilot was then required to maneuver the aircraft through moderate turbulence and attempt to land within a touchdown box painted on the runway. The pilot was expected to fly the aircraft in a manner that met specific performance standards. The performance standards for the precision landing at touchdown were longitudinal distance from the threshold, lateral offset from runway

Table 2 Research aircraft and flight control system configurations

Name	Configuration			Actuator rate limit in axis of interest, deg /s
	Control allocation system			
	Roll axis	Pitch axis	Yaw axis	
Baseline 100	Baseline	Baseline	Baseline	100
CAPIO 100	CAPIO	CAPIO	Baseline	100
Baseline 25	Baseline	Baseline	Baseline	25
CAPIO 25	CAPIO	CAPIO	Baseline	25

centerline, deviation from runway heading and track, deviation from target airspeed, and sink rate.

E. Experimental Procedure

The simulation was conducted over the course of two weeks. During those two weeks, seven test pilots participated in the study. The pilot schedule allotted two days for each pilot to complete the test matrix, with a maximum of two pilots per day, alternating sessions to reduce pilot fatigue. Time at the conclusion of the matrix was reserved for repeat sessions to be used at the discretion of the research engineers.

Orientations were held to brief each pilot on the experiment's background, objective, tasks, procedures, and aircraft system. Pilots were allowed one 1 h warm-up session, and before collecting data for each configuration, a series of practice runs was conducted until the pilots felt they were achieving consistent results. Following the practice runs, a minimum of three data runs were flown for each configuration. At the conclusion of each practice and data run, performance feedback was provided to the pilot via an end-of-run display in the cockpit. Subjective and objective data were collected throughout and upon the completion of the data runs, as described in the following subsection.

Each pilot flew and evaluated all four configurations at least once, resulting in a comprehensive dataset representing nearly 200 data runs. The pilots were not told what configuration they were flying.

F. Collection of Objective and Subjective Data

Both objective and subjective data were collected during the simulation evaluation. The objective data recorded digitally during the simulation evaluation were in three formats. These formats include the simulation time history data with performance standard data, end-of-run pilot displays, and video with audio recording. The subjective data collected for each configuration, when appropriate, were in the form of handling qualities ratings, PIO ratings, and pilot comments.

The handling qualities ratings were collected with the Cooper-Harper handling qualities rating scale [16]. The HQR scale provides numerical data on how the pilot perceived the required workload and aircraft performance for a given task.

The PIO ratings were captured with the PIO rating scale [16], shown in Fig. 9. The PIO scale provides a numerical rating meant to reflect the pilot's perception of the aircraft's performance and flight characteristics for a given task, with particular attention on oscillation and undesirable motion.

Verbal comments expressed by pilots were recorded, and written comments were captured on pilot comment cards unique to the task. These comments provide insight into how each pilot viewed the task and perceived the system's overall performance.

V. Evaluation Results and Discussion

The results of the motion-based piloted simulation conducted in the VMS are presented and discussed with respect to the research objectives.

A. Simulation Results in the Presence of PIOs

Analysis of the simulation data revealed that CAPIO contributed to a reduction in PIO characteristics, and it contributed to an improvement in the pilots' perception of PIO tendencies and in the pilots' PIO ratings, as compared to the baseline allocator. This benefit is particularly pronounced for the aircraft configuration with reduced actuator rate limits.

The lateral stick force and roll rate signals were analyzed to detect PIOs experienced during the simulation. The analysis identified key characteristics of a PIO, as described by Mitchell and Klyde in [20]. The system was determined to be in a PIO when the phase difference between the lateral stick force and roll rate signals exceeded 60 deg, when the oscillation frequency of the roll rate signal was between 1 and 5 rad/s, and when the peak-to-peak amplitudes for both the

PIO RATING SCALE

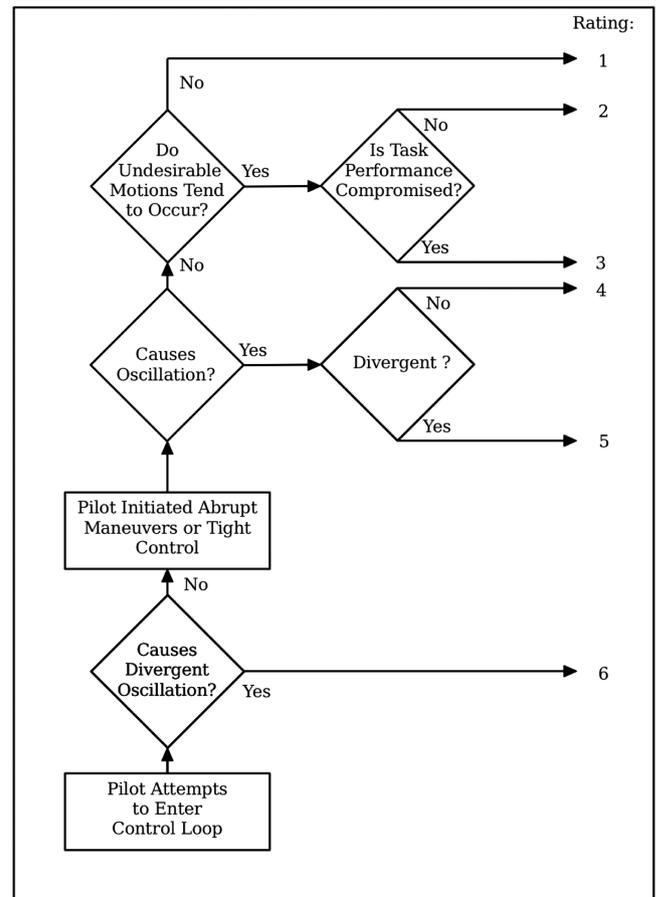


Fig. 9 PIO rating scale.

lateral stick force and roll rate signals were above a predetermined threshold.

Whereas the system with nominal rate limiting experiences few PIOs, the system with reduced actuator rates was susceptible to oscillatory behavior at two points during the task: upon breakout below the ceiling when the pilot first banks after visually identifying the runway, and during final alignment with the runway when the pilot inputs a roll reversal. Since CAPIO does not prevent the occurrence of PIOs (but rather helps the recovery), PIOs are seen at both susceptible points for the baseline and CAPIO allocators with reduced actuator rates. Example time histories for each of the four configurations are shown in Figs. 10–13. The time histories are all data runs from the same pilot and are plotted at the same scale, with y axis values removed to conceal details of the research aircraft model.

The lateral stick force and roll rate signals shown in Figs. 10 and 11 are for the system with nominal actuator rate limiting, and with the baseline allocator and CAPIO allocator, respectively. In both of these cases, large peak-to-peak amplitudes for the signals are present following the simulation time of approximately 25 s, when the pilot performs the aggressive maneuver to align the aircraft with the runway. Whereas phase lag is apparent between the lateral stick force and roll rate signals, the magnitude of the phase difference does not characterize these oscillations as PIOs.

The lateral stick force and roll rate signals shown in Figs. 12 and 13 are for the system with reduced actuator rate limiting, and with the baseline allocator and CAPIO allocator, respectively. As seen with nominal actuator rate limiting, these cases also exhibit large peak-to-peak amplitudes for the lateral stick force and roll rate signals as the pilot performs the aggressive maneuver to align the aircraft with the runway. With the reduced actuator rate limits, however, both the frequency of the oscillations and the phase difference between the signals have increased. These oscillations are considered to be PIOs and are indicative of the data collected throughout the simulation. As

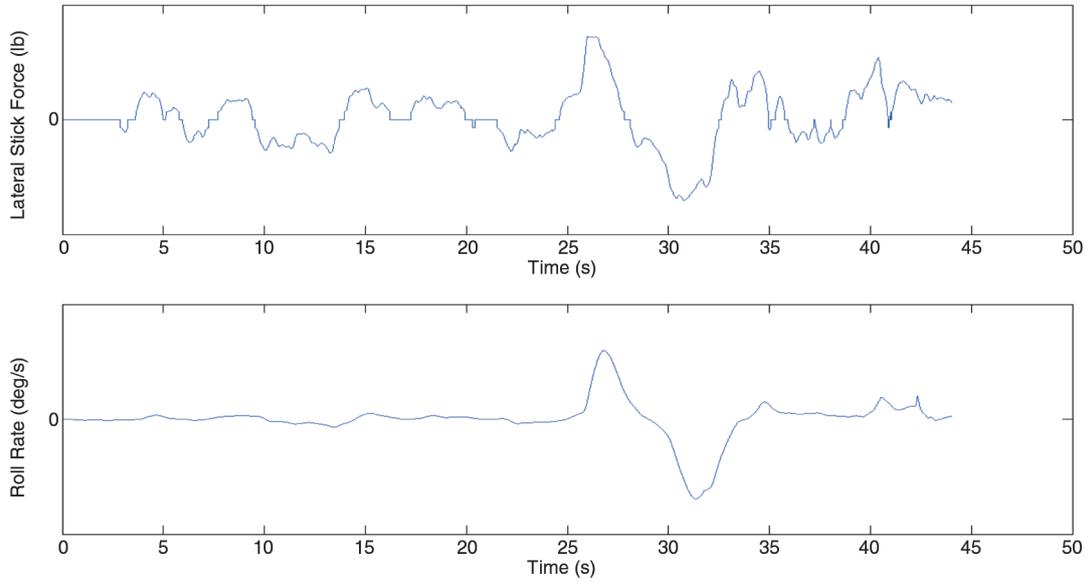


Fig. 10 Time histories for the baseline 100 configuration.

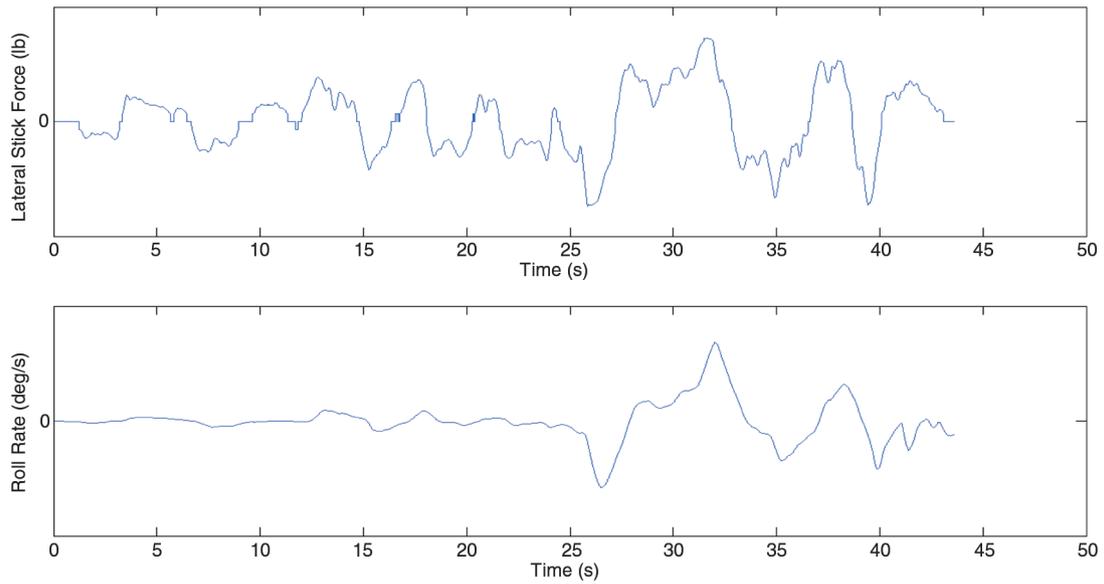


Fig. 11 Time histories for the CAPIO 100 configuration.

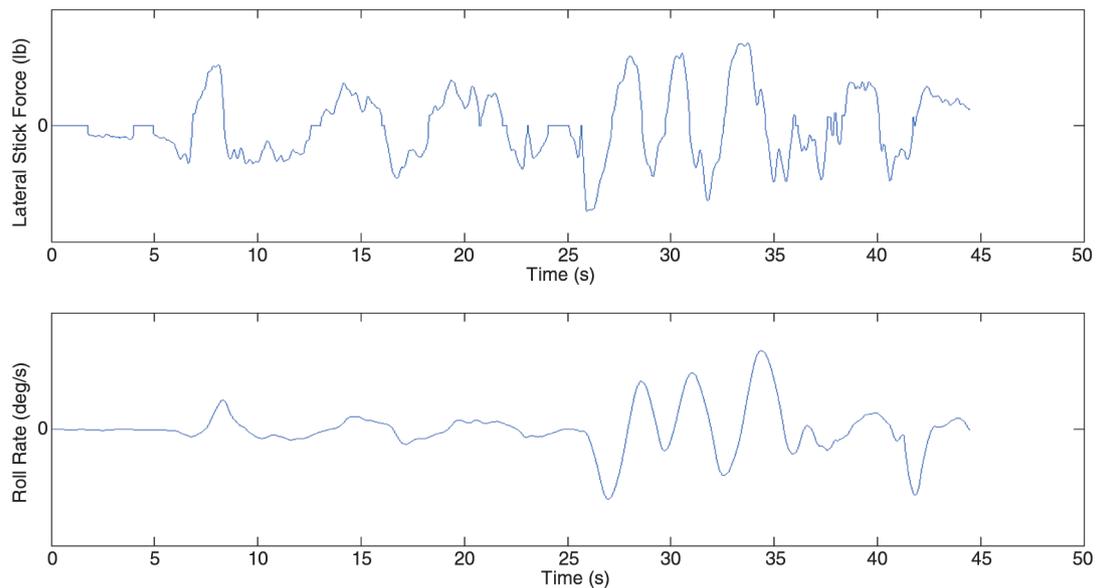


Fig. 12 Time histories for the baseline 25 configuration.

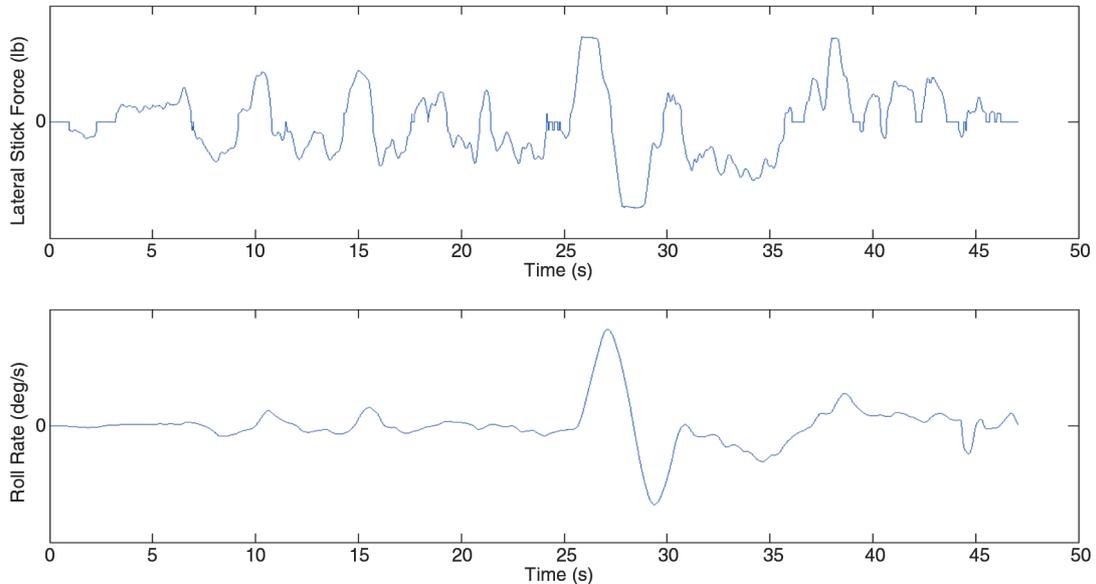


Fig. 13 Time histories for the CAPIO 25 configuration.

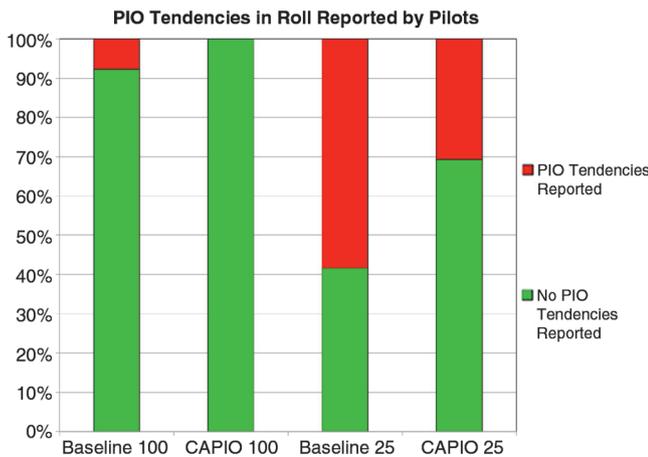


Fig. 14 PIO tendencies in roll axis reported by pilots for configurations.

seen in these time histories, the duration of the PIO experienced by the system with the CAPIO allocator is noticeably shorter than with the baseline allocator.

Pilots' perception of the PIO tendencies and their PIO ratings are consistent with the PIO characteristics identified in the simulation data. With reduced actuator rate limits, pilots indicated fewer PIO tendencies for the CAPIO allocator: over half of the pilot reports indicated PIO tendencies in roll with the baseline allocator, but only about 30% of the pilot reports indicated PIO tendencies in roll with the CAPIO allocator. The percentage of pilot reports that did and did not indicate PIO tendencies in roll is given in Fig. 14, represented by the colors of each bar corresponding to a configuration. Table 3 provides the reports for PIO tendencies in roll for each configuration by pilot.

Table 3 PIO tendencies in roll axis reported by pilots for configurations

Pilot	PIO tendency in roll			
	Baseline 100	CAPIO 100	Baseline 25	CAPIO 25
Pilot 1	No, No	No	Yes, Yes	Yes, No
Pilot 2	No	No, No	Yes	No
Pilot 3	No, No	No, No	Yes	No, Yes
Pilot 4	No, No	No	No, No	No, No
Pilot 5	No, Yes	No, No	No, No	Yes, No
Pilot 6	No, No	No, No	Yes, No	No, No
Pilot 7	No, No	No, No	Yes, Yes	No, Yes

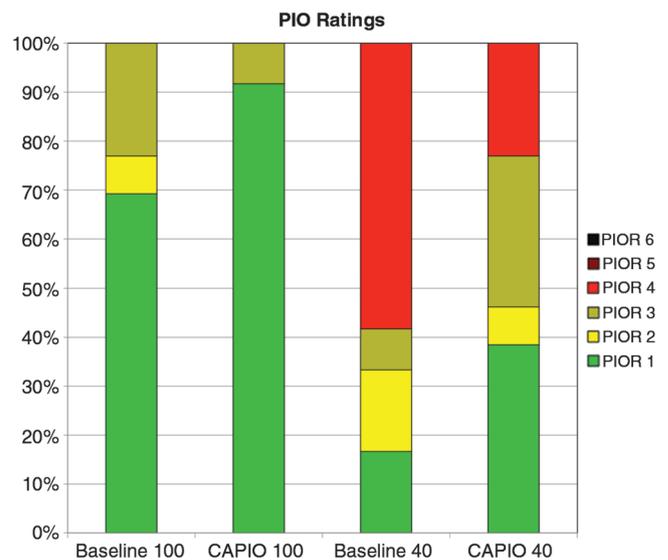


Fig. 15 PIO ratings (PIOR) for configurations.

The pilots' PIO ratings are given in Fig. 15 and Table 4. Each column of Fig. 15 represents a configuration, and the colors of the bars represent the percentage of rating values received from pilots. The majority of pilots gave the systems with nominal actuator rate limits a PIO rating of 1, corresponding to the observation that no undesirable motions tended to occur. When actuator rate limits were reduced, only approximately 16 and 36% of ratings were a value of 1 for the baseline and CAPIO allocators, respectively. Almost 60% of the PIO ratings for the baseline allocator configuration were 4, corresponding to the observation of nondivergent oscillations as a

Table 4 PIO ratings for configurations

Pilot	PIO ratings			
	Baseline 100	CAPIO 100	Baseline 25	CAPIO 25
Pilot 1	1, 1	1	4, 4	3, 1
Pilot 2	2	1, 1	4	3
Pilot 3	1, 1	1, 1	4	1, 4
Pilot 4	1, 1	1	1, 1	1, 1
Pilot 5	3, 3	3, 1	2, 2	4, 2
Pilot 6	3, 1	1, 1	4, 3	3, 3
Pilot 7	1, 1	1, 1	4, 4	1, 4

result of abrupt maneuvers or tight control. The PIO ratings for the CAPIO allocator configuration were better, with only 23% of the ratings being 4 and over one-third of the ratings being 3, corresponding to the observation of undesirable motions that compromised task performance.

B. Simulation Results of Overall System Characteristics

Analysis of the simulation data revealed that the overall system characteristics were strongly influenced by the control surface rate limits. The baseline allocator and CAPIO allocator received comparable overall pilot assessments, handling qualities ratings, and task performance for the offset approach and landing task.

Pilot comments captured on the pilot comment cards convey that flight dynamics of the system degrade as control surface rate limits are reduced. For the nominal aircraft configuration, the pilots on average declared the ability to control and track bank angle to be fair and the ability to correct lateral offset to be between fair and difficult. With the rate limits reduced, the average pilot rating for both of these characteristics dropped to difficult. This degradation is present, regardless of the control allocator. The average predictability of roll response ranges between satisfactory and unsatisfactory, with a slight decrease toward unsatisfactory for configurations with reduced actuator rate limits. The difference in predictability offered by the CAPIO and baseline allocator, however, is not significant. Pilot comments indicate that the sluggish response or lag in response from the systems contributes to the unpredictability, especially at reduced actuator rate limits.

The distinction between the nominal aircraft and the aircraft with reduced actuator rate limits is also prominent in the Cooper–Harper ratings. The nominal aircraft had mostly level 2 handling qualities for either allocator given a majority of Cooper–Harper ratings of 4, 5, and 6. At these ratings, pilots found the system deficiencies to range from minor but annoying to very objectionable but tolerable due to the extensive compensation required. Based on pilot comments collected during the simulation, it was found that the difficulty of the task contributed to the need for extensive compensation and poor system rating. The handling qualities for the aircraft with reduced actuator rate limits were worse at mostly level 3, given a majority of Cooper–Harper ratings of 7, 8, and 9. (See Fig. 16 and Table 5). The reduced actuator rate limits imposed another challenge on the pilots in addition to the difficult task, resulting in a system that pilots deemed to have major deficiencies.

Despite the poor handling qualities ratings and pilot comments, on average, pilots achieved performance within the desired and satisfactory performance standards threshold for the final approach and touchdown. The desirable and satisfactory performance is achievable because, for all configurations, pilots were able to retain tight control over the system by operating at very high workloads and recovering from undesired oscillatory behavior before touchdown.

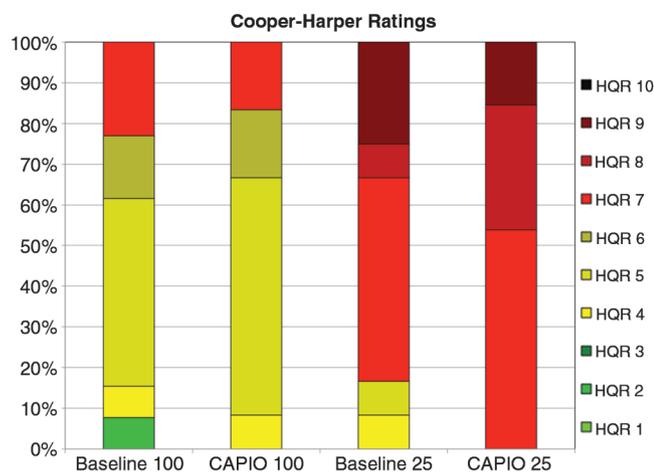


Fig. 16 Cooper–Harper ratings for configurations.

Table 5 Cooper–Harper ratings for configurations

Pilot	Cooper–Harper ratings			
	Baseline 100	CAPIO 100	Baseline 25	CAPIO 25
Pilot 1	6, 5	5	7, 7	7, 7
Pilot 2	5	4, 5	7	7
Pilot 3	4, 5	5, 5	7	7, 7
Pilot 4	5, 5	7	4, 7	8, 8
Pilot 5	6, 7	5, 6	5, 7	7, 7
Pilot 6	7, 2	5, 6	9, 9	9, 9
Pilot 7	7, 5	7, 5	9, 8	8, 8

C. Summary Observation

The simulation evaluation was intended to assess the viability of CAPIO to mitigate the effects of factors that make next-generation aircraft more susceptible to PIOs as airframe designs become more energy efficient. In particular, the simulation evaluation assessed the effects of reducing actuator rate limits. As the actuators of the aircraft functioned at their reduced rate limits, phase lag was introduced. In [16], McRuer linked phase lag to poor pilot ratings. The experimental results are consistent with expectation, revealing degradation in pilot ratings from the configuration with nominal actuator rate limits to the configuration with reduced actuator rate limits. Although CAPIO was meant to reduce the phase lag, it was designed to only do this after significant phase lag was detected. Due to this design feature, CAPIO did not impact the overall pilot perception.

This experiment did validate that CAPIO helped to reduce PIO behavior introduced by control surface rate limiting, and that pilots recognized the improvement in their ability to maintain or regain closed-loop control of the aircraft with respect to PIO behavior. CAPIO accomplished this by adjusting the control allocation mode in flight, causing changes in the effective vehicle dynamics that can be expected to degrade pilot perceptions [16, 17]. Since pilot perception of overall system characteristics remained consistent between the baseline and CAPIO allocators, it appears that the changes in the effective vehicle dynamics caused by mode switching in CAPIO did not adversely impact pilot perception. This result also suggests that the real-time optimization algorithm used within CAPIO offered similar control performance to the baseline allocator, which did not use real-time optimization.

In summary, the simulation evaluation effectively meets the research objectives by demonstrating that CAPIO does not degrade system characteristics of a nominal aircraft, and that CAPIO improves PIO recovery characteristics of an aircraft with reduced actuator rate limits when compared to the baseline control allocator.

VI. Conclusions

A motion-based piloted simulation evaluation was completed to assess the impact of the CAPIO allocator on the pilot–aircraft–control system characteristics. CAPIO is intended to address problems that arise from stringent actuator rate limits for multi-input/multi-output applications, such as phase lag and PIOs. The system does this by actively detecting and eliminating phase lag introduced by control surface rate limiting. Results from the simulation evaluation confirm that CAPIO successfully contributed to an improvement in the pilots' perception of PIO tendencies and in the pilots' PIO ratings, as compared to the baseline allocator. Since CAPIO is currently only designed to improve system characteristics after phase lag is detected, it received a comparable overall pilot assessment, as well as comparable handling qualities ratings and task performance, to the baseline allocator.

Even with the accomplishments from the simulation evaluation, further maturation of CAPIO is still needed to enable lowering the actuator rate limit requirements. Analytical studies are needed to enhance and guarantee the closed-loop stability properties of an aircraft, and to extend CAPIO to influence and improve overall system characteristics. The system integration and impact of CAPIO will also need to be demonstrated through piloted flight studies. If successful, CAPIO will allow aerodynamically efficient airframe

designs that are currently unattainable due to technologically prohibitive control power requirements.

Acknowledgments

The authors extend their appreciation to two projects within the NASA Aeronautics Research Mission Directorate, the Fundamental Aeronautics Program's Subsonic Fixed Wing Project and the Integrated Systems Research Program's Environmentally Responsible Aviation Project, for their support of this research. This research was accomplished with the contributions of many individuals and organizations. The authors thank our project pilot and coauthor, Gordon H. Hardy, for his guidance throughout the research and evaluation of the control allocation technique to recover from pilot-induced oscillations. In addition to Gordon H. Hardy, the authors thank the pilots who participated in and provided insight during the preliminary and final simulation evaluations, including Dana Purifoy, Richard Ewers and Jim Smolka from the NASA Dryden Flight Research Center; Carl Ott from the U.S. Army Aeroflightdynamics Directorate; Dan Dugan from NASA Ames Research Center; Jim Lindsey from Monterey Technologies, Inc., at NASA Ames Research Center; Jim Martin; and George Tucker. The authors are also thankful for the exceptional service and assistance provided by the NASA Ames Research Center Vertical Motion Simulator team.

References

- [1] Greitzer, E. M., Bonnefoy, P. A., De la Rosa Blanco, E., Dorbian, C. S., Drela, M., Hall, D. K., Hansman, R. J., Hileman, J. I., Liebeck, R. H., Levegren, J., Mody, P., Pertuze, J. A., Sato, S., Spakovszky, Z. S., Tan, C. S., Hollman, J. S., Duda, J. E., Fitzgerald, N., Houghton, J., Kerrebrock, J. L., Kiwada, G. F., Kordonow, D., Parrish, J. C., Tylko, J., Wen, E. A., and Lord, W. K., "N3 Aircraft Concept Designs and Trade Studies, Final Report," NASA CR-2010-216794/VOL1, Dec. 2010.
- [2] Bruner, S., Baber, S., Harris, C., Caldwell, N., Keding, P., Rahrig, K., Pho, L., and Wlezian, R., "NASA N + 3 Subsonic Fixed Wing Silent Efficient Low-Emissions Commercial Transport (SELECT) Vehicle Study," NASA CR-2010-216798, Nov. 2010.
- [3] Bradley, M. K., and Droney, C. K., "Subsonic Ultra Green Aircraft Research: Phase I Final Report," NASA CR-2011-216847, April 2011.
- [4] D'Angelo, M. M., Gallman, J., Johnson, V., Garcia, E., Tai, J., and Young, R., "N + 3 Small Commercial Efficient and Quiet Transportation for Year 2030-2035," NASA CR-2010-216691, May 2010.
- [5] Cameron, D., and Princen, N., "Control Allocation Challenges and Requirements for the Blended Wing Body," *AIAA Guidance, Navigation and Control Conference and Exhibit*, AIAA Paper 2000-4539, Aug. 2000.
- [6] Roman, D., Allen, J. B., and Liebeck, R. H., "Aerodynamic Design Challenges of the Blended-Wing-Body Subsonic Transport," *AIAA Applied Aerodynamics Conference*, AIAA Paper 2000-4335, Aug. 2000.
- [7] Durham, W. C., and Bordignon, K. A., "Multiple Control Effector Rate Limiting," *Journal of Guidance, Control, and Dynamics*, Vol. 19, No. 1, 1996, pp. 30-37. doi:10.2514/3.21576
- [8] Bodson, M., "Evaluation of Optimization Methods for Control Allocation," *Journal of Guidance, Control, and Dynamics*, Vol. 25, No. 4, 2002, pp. 703-711. doi:10.2514/2.4937
- [9] Oppenheimer, M. W., and Doman, D. B., "A Method for Including Control Effector Interactions in the Control Allocation Problem," *AIAA Guidance, Navigation and Control Conference and Exhibit*, AIAA Paper 2007-6418, Aug. 2007.
- [10] Frost, S. A., Bodson, M., and Acosta, D. M., "Sensitivity Analysis of Linear Programming and Quadratic Programming Algorithms for Control Allocation," *AIAA Infotech@Aerospace Conference*, AIAA Paper 2009-1936, April 2009.
- [11] Yildiz, Y., and Kolmanovsky, I. V., "A Control Allocation Technique to Recover from Pilot-Induced-Oscillations (CAPIO) Due to Actuator Rate Limiting," *Proceedings of the 2010 American Control Conference*, IEEE, Baltimore, MD, June-July 2010.
- [12] Yildiz, Y., Kolmanovsky, I. V., and Acosta, D., "A Control Allocation System for Automatic Detection and Compensation of Phase Shift Due to Actuator Rate Limiting," *2011 American Control Conference*, IEEE, San Francisco, CA, June-July 2011.
- [13] Yildiz, Y., and Kolmanovsky, I., "Stability Properties and Cross-Coupling Performance of the Control Allocation Scheme CAPIO," *Journal of Guidance, Control, and Dynamics*, Vol. 34, No. 4, 2011, pp. 1190-1196. doi:10.2514/1.50310
- [14] "Handling Qualities of Unstable Highly Augmented Aircraft," AGARD Advisory Rept. 279, May 1991.
- [15] "Flying Qualities of Piloted Aircraft," U.S. Dept. of Defense MIL-HDBK-1797, Dec. 1997.
- [16] McRuer, D., "Pilot-Induced Oscillations and Human Dynamic Behavior," NASA CR-4683, July 1995.
- [17] "Flight Control Design Best Practices," Research and Technology Organization, NATO Rept. RTO-RT-29, Neuilly-sur-Seine Cedex, France, Dec. 2000.
- [18] Golub, G. H., and Van Loan, C. F., *Matrix Computations*, 2nd ed., John Hopkins Univ. Press, Baltimore, MD, 1989, pp. 256-258.
- [19] Mitchell, D. G., and Hoh, R. H., "Development of Methods and Devices to Predict and Prevent Pilot Induced Oscillations," SBIR Phase II AFRL Rept. AFRL-VA-WP-TR-2000-3046, Dec. 2000.
- [20] Mitchell, D. G., and Klyde, D. H., "Identifying a Pilot-Induced Oscillation Signature: New Techniques Applied to Old Problems," *Journal of Guidance, Control, and Dynamics*, Vol. 31, No. 1, 2008, pp. 215-224. doi:10.2514/1.31470
- [21] Fletcher, R., *Practical Methods of Optimization*, 2nd ed., Wiley, New York, 1987, p. 21.
- [22] Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C., "A Limited Memory Algorithm for Bound Constrained Optimization," *SIAM Journal on Scientific Computing*, Vol. 16, No. 5, 1995, pp. 1190-1208. doi:10.1137/0916069
- [23] Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J., "Algorithm 778: L-BFGS-B: FORTRAN Subroutines for Large Scale Bound Constrained Optimization," *ACM Transactions on Mathematical Software*, Vol. 23, No. 4, 1997, pp. 550-560. doi:10.1145/279232.279236
- [24] Leonard, M., "Real-Time Optimization for Use in a Control Allocation System to Recover from Pilot Induced Oscillations," *AIAA Guidance, Navigation, and Control Conference*, AIAA Paper 2013-5239, Aug. 2013.
- [25] Schroeder, J., and Chung, W., "Simulator Platform Motion Effects on Pilot-Induced Oscillation Prediction," *Journal of Guidance, Control, and Dynamics*, Vol. 23, No. 3, 2000, pp. 438-444. doi:10.2514/2.4578
- [26] Aponso, B., Beard, S., and Schroeder, J., "The NASA Ames Vertical Motion Simulator—A Facility Engineered for Realism," *Royal Aeronautical Society Spring 2009 Flight Simulation Conference*, Paper AF2009223, June 2009.
- [27] Zeune, C. H., "Enabling Speed Agility for the Air Force," *AIAA Aerospace Sciences Meeting*, AIAA Paper 2010-0349, Jan. 2010.
- [28] Hyde, D. C., Shweyk, K. M., Brown, F., and Zeune, C., "Flight Control Development for a Speed-Agile Powered-Lift Transport Aircraft," *AIAA Atmospheric Flight Mechanics Conference*, AIAA Paper 2012-4578, Aug. 2012.
- [29] Shweyk, K. M., Hyde, D. C., Levensgood, K., and Zeune, C., "Validation of the Flight Control System of a Conceptual, Powered-Lift, Speed-Agile, Transport Aircraft," *AIAA Atmospheric Flight Mechanics Conference*, AIAA Paper 2012-4579, Aug. 2012.