

INVESTIGATING THE ROLE OF RATE LIMITING IN PILOT-INDUCED OSCILLATIONS

David H. Klyde
Associate Fellow AIAA
Principal Research Engineer
Systems Technology, Inc.
Hawthorne, California

David G. Mitchell
Associate Fellow AIAA
Technical Director
Hoh Aeronautics, Inc.
Lomita, CA

ABSTRACT

From the Wright Flyer to fly-by-wire, the phenomenon of pilot-induced oscillations or PIO has been observed on prototype, experimental, and operational military and commercial aircraft. The introduction of irreversible control systems with surfaces driven by powered actuators brought many benefits along with increased system complexity and the introduction of additional nonlinearities. Chief among these nonlinearities are the hardware and software rate limits associated with the control surface actuators. Basic sizing tradeoffs conducted in the design process set the maximum rate of an actuator, while software rate limits are introduced to prevent overdriving the control surface when loads or structural limitations exist. As demonstrated in this paper, when operating as intended, there are usually no ill effects associated with rate limits, however, certain conditions can lead to a highly saturated condition. This results in the sudden introduction of significant added phase lags to the pilot-vehicle system. In many cases the end result is often PIO or other related loss of control events. One of the earliest well documented PIO events involving rate limiting occurred on the first flight of the X-15. This event is significant in that common linear systems analysis techniques do not reveal a susceptibility to PIO. The analysis of X-15 Flight 1-1-5 and the results of a more recent flight research program are presented in detail in this paper.

NOMENCLATURE

ATTAS – Advanced Technologies Testing Aircraft System
HQR – Cooper-Harper Handling Qualities Rating
NAVAIR – Naval Air Systems Command
NRC – National Research Council
PCRL – Phase Compensated Rate Limiter
PIO – Pilot-Induced Oscillation
PIOR – Pilot-Induced Oscillation tendency Rating
PIOS – Pilot-Induced Oscillation Suppression
RLE – Rate Limiting Element
USAF – United States Air Force
 A – actuator command input amplitude
 a_n – normal acceleration
 e – actuator model error signal ($e = \delta_c - \delta$)

e_L – saturation point
 F_{es} – cockpit stick force
 h – altitude
 K^* – RLE ratio of triangle output amplitude to sinusoidal maximum input amplitude (x_o/x_{imax})
 p – roll rate
 q – pitch rate
 T – linear system time constant ($T = 1/\omega_n$)
 T_{NL} – nonlinear system indicial response time constant
 t_D – RLE phase difference between triangle output and sinusoidal input ($t_D = t_o - t_i$)
 t_i – RLE time to maximum input amplitude
 t_o – RLE time to maximum output amplitude
 V_i – indicated airspeed
 V_L – actuator rate limit
 V_V – vertical speed
 x_i – RLE sinusoidal input amplitude
 x_{imax} – RLE maximum sinusoidal input amplitude
 x_o – RLE triangle output amplitude
 α – angle-of-attack
 β – angle-of-sideslip
 $\Delta\phi$ – additional phase lag due to rate limiting
 δ – actuator position
 δ_a – aileron position
 δ_c – actuator position command
 δ_h – horizontal stabilizer position
 $\dot{\delta}$ – actuator rate
 θ – pitch attitude
 $\tau_{p\theta}$ – pitch attitude phase delay
 ω – frequency
 ω_a – actuator bandwidth frequency
 $\omega_{BW\theta}$ – pitch attitude bandwidth frequency, defined as frequency for the lower of 6dB gain margin or 45 deg phase margin
 ω_i – input frequency
 ω_{onset} – saturation onset frequency
 ω_{PIO} – pilot-induced oscillation frequency
 $\omega_{u\theta}$ – pitch attitude neutral stability frequency

INTRODUCTION

PIO and Rate Limiting

Throughout the first 100 years of powered flight, aircraft handling qualities has often been addressed as an afterthought of the design process. Because of their often catastrophic nature, however, nothing brings attention to handling qualities like a high profile pilot-induced oscillation (PIO). Traditionally, the occurrence of such an event has led to significant research activities that are intended to alleviate the problem once and for all. Despite significant technical advances in this area, PIOs continue to occur with both flight test and operational aircraft.

The focus of this paper is on the role of rate limiting in PIO. The earliest well-documented PIO event that featured rate limiting was the first flight of the X-15 aircraft in 1959. This event, which is discussed in more detail later in this paper, is one of the most analyzed of all PIO events. Among the earliest is the analysis documented in Ref. 1 that features the development of describing function approximations for a rate limited actuator and an inverse describing function technique used to predict limit cycle oscillations and their frequencies. This work inspired the further development of the rate limiting describing function analyses of Ref. 2 that was conducted as part of the Air Force Unified Pilot-Induced Oscillation Theory program of the mid 1990's and is partially reported on herein.

Attention to the effects of rate limiting and other system nonlinearities surfaced again in conjunction with the severe PIOs that occurred with the YF-12 aircraft during aerial refueling. An excellent description of the nonlinear analysis of these events is contained in Ref. 3. More recent highly documented PIO events that involved rate limiting include the YF-22⁴ and JAS-39⁵ events. These events that resulted in significant damage to the YF-22 and the loss of two JAS-39 aircraft inspired a new thrust in handling qualities research that emphasized comprehension, prevention, and alleviation of the nonlinear effects associated with rate limiting. Specifically, the YF-22 event led to the above-mentioned Air Force program, while the JAS-39 events led SAAB to develop a patented control scheme designed to alleviate the effects of rate limiting. This scheme is described briefly later in this paper.

Categorizing PIO

In 1997, a summary report by a National Research Council (NRC) Committee on the Effects of Aircraft-Pilot Coupling on Flight Safety was published⁶. The NRC Committee separated PIOs by Category,

depending essentially on the degree of nonlinearity in the event:

- Category I: Linear pilot-vehicle system oscillations. These PIOs result from identifiable phenomena such as excessive time delay, excessive phase loss due to filters, improper control/response sensitivity, etc. They are the simplest to model, understand, and prevent. They are also the least common in operational flying.
- Category II: Quasi-linear events with some nonlinear contributions, such as rate or position limiting. For the most part, these PIOs can be modeled as linear events, with an identifiable nonlinear contribution that may be treated separately. The most common nonlinear contribution is the subject of this paper: rate limiting of a control effector actuator.
- Category III: Nonlinear PIOs with transients. Such events are difficult to recognize and rarely occur, but are always severe. Mode switching, that cannot be represented by a quasi-linear equivalent, is the common culprit.

Since the publication of the NRC's findings, several researchers have suggested that there may be other Categories that are distinct from the three defined above. "Category IV" PIO sometimes refers to events that are caused by, or have as a major contributor, structural modes and their interactions with the pilot.

The PIO Trigger

For PIO to occur, there must be a trigger. As a result of the well-publicized crashes of the YF-22 and the JAS-39 aircraft in the early 1990's, it was speculated that rate limiting was a trigger for most PIOs on fly-by-wire aircraft. After the research efforts of the 1990's we can conclude that rate limiting *can* be a trigger for PIO. On the other hand, sometimes PIO can be the cause of the rate limiting. In addition, we also know that it is possible to encounter severe rate limiting without a PIO. All of these possible outcomes are illustrated later in this paper.

THE NATURE OF RATE LIMITING

Simplified Actuator Model with Rate Limiting

A simplified model of a rate-limited actuator that was previously analyzed^{2,7} is shown in Figure 1. As long as the error signal, e , remains below the saturation point, e_L , the system behaves as a linear first-order lag whose

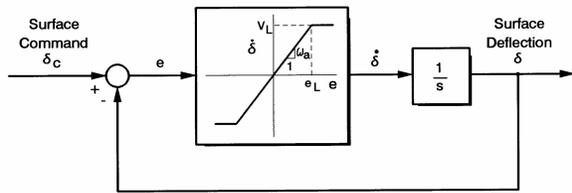


Figure 1. Simplified Actuator Model

response is entirely dictated by the time constant, $T = 1/\omega_a$, where ω_a is the actuator bandwidth. When saturation occurs the surface will move at its maximum rate, V_L , until the commanded magnitude, frequency, or both are reduced. Saturation occurs when the error signal exceeds the saturation point (i.e., $e > e_L = V_L / \omega_a$). It is also convenient to define a nonlinear system time constant, $T_{NL} = A/V_L$, where A is the amplitude of the input command. Although this time constant applies only to the indicial (step) response describing function, it has been shown to have an important role in characterizing the nonlinear system for other input forms.²

Figure 2 displays the closed loop actuator time responses for a saturation case with a time constant of $T = 0.05$ sec, a saturation point of $e_L = 2$ deg, and a sinusoidal input with an amplitude of 15 deg. The error signal still appears more or less sinusoidal even though the 2 deg saturation point is generally exceeded throughout the run length. The nonlinear nature is more evident in the actuator rate response that appears box car-like for this highly saturated case. In the lower panel, the actuator output position is compared to the input command. A triangle wave output response is displayed that reverses when approximately equal to the input (i.e., when the error signal passes through zero). For the linear case, the delay between the output and input is the linear time constant or 0.05 sec. In the nonlinear case shown in Figure 2, this delay has increased to 0.15 sec.

Exact Describing Function Representation

In general, rate saturation results in an amplitude reduction and a significant added phase lag. These characteristics are displayed in the magnitude and phase curve families of Figure 3 that represent exact sinusoidal describing function representations of the simplified model of Figure 1^{2,7}. Both sets of curves are plotted as a function of normalized frequency, ω/ω_a , and the linear to nonlinear time constant ratio, T/T_{NL} . In the sinusoidal input case T/T_{NL} is equal to $V_L/(\omega_a A)$. Thus, the two plots display the describing function magnitude and phase of the nonlinear system in terms of the actuator design parameters (V_L and ω_a) and the

input parameters (A and ω), all known quantities. There are several observations to note from the plots. First, the $T/T_{NL} = 1$ curve represents the linear case. Second, the more highly saturated cases represented by the smaller time constant ratio curves, depart from the linear curve at a normalized frequency that is equivalent to their time constant ratio. For example the $T/T_{NL} = 0.1$ curve departs from the linear curve at a normalized frequency of 0.1. Another more significant result is that known design and input parameters can be used to identify the added phase lag due to a rate limiting actuator.

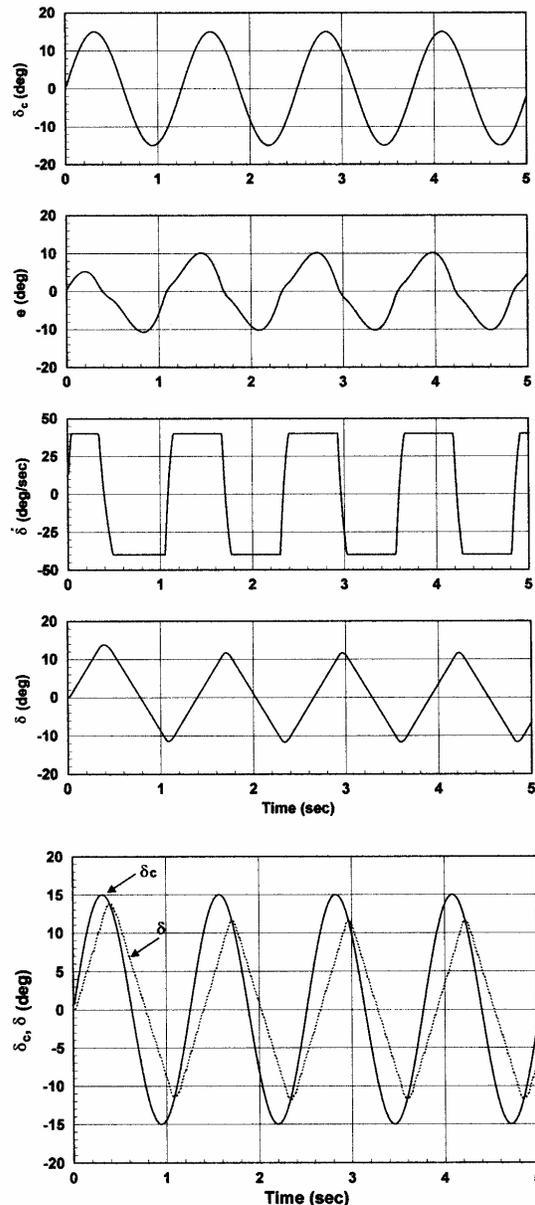


Figure 2. Highly Saturated Rate-limited Actuator Model Responses to a Sinusoidal Input

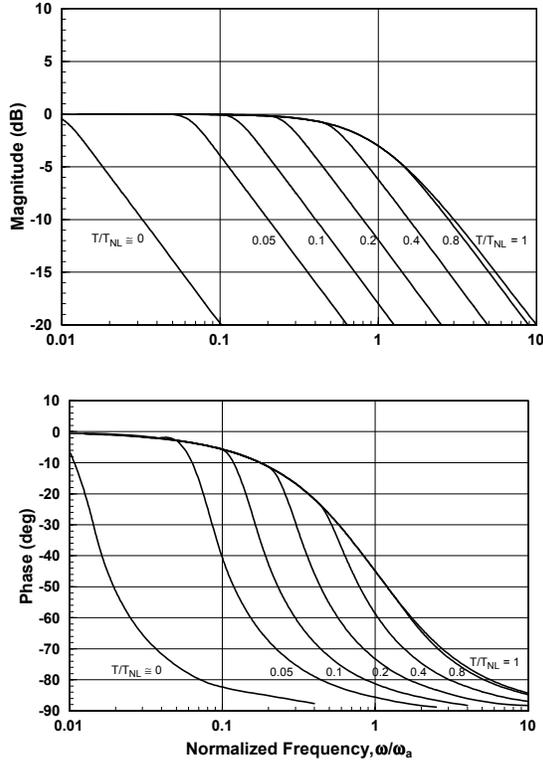


Figure 3. Frequency Response of Simplified Nonlinear Actuator Model

The Software Rate Limit

Often times software rate limits are placed in series with an actuator to insure that the actual physical limit is not encountered, thus preventing potential damage to the flight hardware. It is not unusual for these nonlinearities to be found in the feedback path as well. As described in Ref. 2 the primary difference between the software limit and the hardware limit is the significant loss in actuator bandwidth that accompanies the hardware limit when saturated. A software rate limit can be represented by the actuator model described above where the bandwidth is effectively infinite (i.e., $T/T_{NL} \rightarrow 0$). Although Figure 3 displays both an amplitude reduction and added phase lag due to rate limiting, it is the added phase lag that is the dominant effect. Thus when encountered in an automatic control system or a pilot-vehicle system, the primary effect of actuator or software rate saturation is to consume available phase margin that can then lead to loss of control.

A describing function approximation of a software rate limiting element can be obtained by assuming a sinusoidal input/triangle output rate limiting element as was done by Hanke in Ref. 8 and shown in Figure 4. This method does not consider the servo loop explicitly,

since only the wave forms of the input and output of the rate limiting element are considered. It does, however, implicitly take into account the servo loop by requiring a reversal of the output whenever the servo error, $e = x_i - x_o$, becomes zero. This nonlinear element model is exact when applied to a control system software limiter that contains no dynamics. It is also used to approximate an infinite bandwidth actuator (i.e., $\omega_d \gg \omega$).

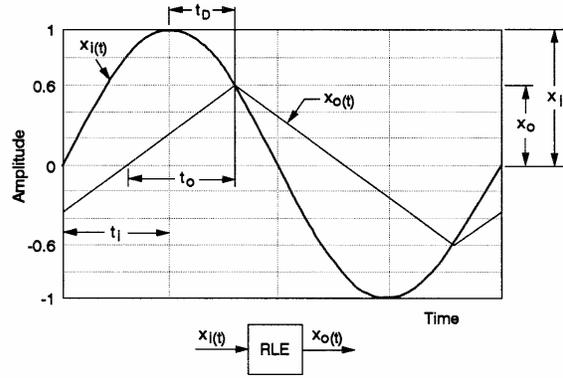


Figure 4. Rate Limiting Element Time Response for a Sinusoidal Input (from Ref. 8)

The actual output/input magnitude is defined by taking the ratio of the constant output to maximum input rates and then solving for x_o/x_{imax} . Defining K^* as x_o/x_{imax} , this yields the following result,

$$K^* = \frac{x_o}{x_{imax}} = \frac{\pi}{2} \frac{\dot{x}_o}{\dot{x}_{imax}}$$

This result can be written in terms of the Figure 1 variables, as shown below by recognizing that the output rate when saturated is V_L and the maximum input rate is $A\omega$,

$$K^* = \frac{\pi}{2} \frac{V_L}{A\omega}$$

The K^* parameter is next used to define the describing function magnitude and phase. The describing function magnitude is obtained by multiplying K^* , which represents the actual peak magnitude of the triangle wave, by the Fourier fundamental of the triangle wave (i.e., $8/\pi^2$) as shown by,

$$\left| \frac{\delta(j\omega)}{\delta_c(j\omega)} \right| = \frac{8}{\pi^2} K^* = \frac{4}{\pi} \frac{V_L}{A\omega}$$

The phase difference between the output and input is represented by t_D in Figure 4. It is also noted in the figure that the input and output amplitudes are equal at

$t = t_i + t_D$. Thus to obtain t_D , set the input relation equal to the output and use the above substitution for t to produce the following,

$$x_i \sin[\omega(t_i + t_D)] = x_o$$

This equation is simplified by substituting K^* for x_o/x_{imax} , expanding $\sin[\omega(t_i + t_D)]$ and noting that $\omega t_i = \pi/2$. This results in,

$$\cos(\Delta\phi) = K^*$$

where $\Delta\phi = \omega t_D$ is the phase angle between the input and output. Finally, the phase difference ($\Delta\phi$) is obtained by solving for the argument,

$$\Delta\phi = \cos^{-1}(K^*)$$

The K^* parameter also provides a measure of the severity of the rate limiting, where $0 \leq K^* \leq 1$. As $K^* \rightarrow 1$, rate limiting diminishes and the nonlinear system becomes increasingly linear. As $K^* \rightarrow 0$, rate limiting increases and the corresponding amplitude reduction and added phase lag also increase as shown in Figure 5.

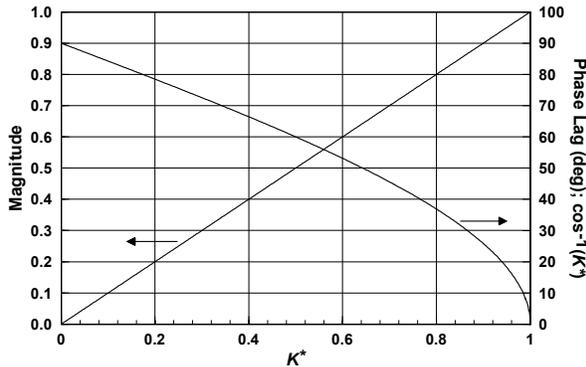


Figure 5. Sinusoidal Input/Triangle Output Describing Function

THE X-15 REVISITED

Program Description

The X-15 program began in 1952 as a joint venture between the military and NACA to investigate the basic problems associated with human space flight. As described in Ref. 9, the objectives of the aircraft program were to investigate aerodynamic forces, heating, stability and control, reentry characteristics, and human physiology at extremely high speeds and altitudes. The North American Aviation built X-15, shown in Figure 6, was released from a B-52 carrier vehicle at 45,000 ft. The aircraft would accelerate to speeds from Mach 2 to 6 and achieve altitudes as high

as 350,000 ft. The program concluded in 1968 after almost 200 flights that provided over 18 hours of high speed research.



Figure 6. X-15 Aircraft (NASA Photo)

FLIGHT 1-1-5 LANDING FLARE PIO

The X-15 landing flare PIO occurred on 8 June 1959 with pilot Scott Crossfield at the controls. This first flight (designated as Flight 1-1-5) was an unpowered glide flown using the side-located controller and with the pitch damper off. Additional details of the flight and subsequent changes to the aircraft are provided in Ref. 10. As shown in the flare time history traces of Figure 7, severe longitudinal oscillations developed near the end of the flare cycle and rate limiting is clearly evident in the horizontal stabilizer angle (δ_h) trace. The “triangle-wave” response of the δ_h time trace in the PIO region indicates that the actuator was operating in the highly saturated region. From the pitch rate (q) trace, a PIO frequency of approximately 3.3 rad/sec is estimated. For this flight the maximum control surface rate was limited to 15 deg/sec.

Since the PIO occurred with the pitch SAS off, the X-15 bare airframe data of Ref. 11 could be used to generate relevant longitudinal transfer functions. Corrections to the Ref. 11 data were made to accommodate the PIO flight condition and aircraft weight. A first order model for the horizontal stabilizer actuator ($\omega_a = 25$ rad/sec) was obtained from Ref. 12. Using these data the Bode and Nichols frequency response survey of Figure 8 for the θ/δ_h transfer function was generated. The transfer function gain was arbitrarily set so that the frequency response would pass through 0 dB at -110 deg of phase. Several key Category I PIO indicators (e.g., bandwidth frequency, phase delay, and average phase rate) are identified on the plots. Not only do all of the applied Category I criteria indicate that the X-15 would not be susceptible to PIO but also the aircraft was found to be Level 1 for

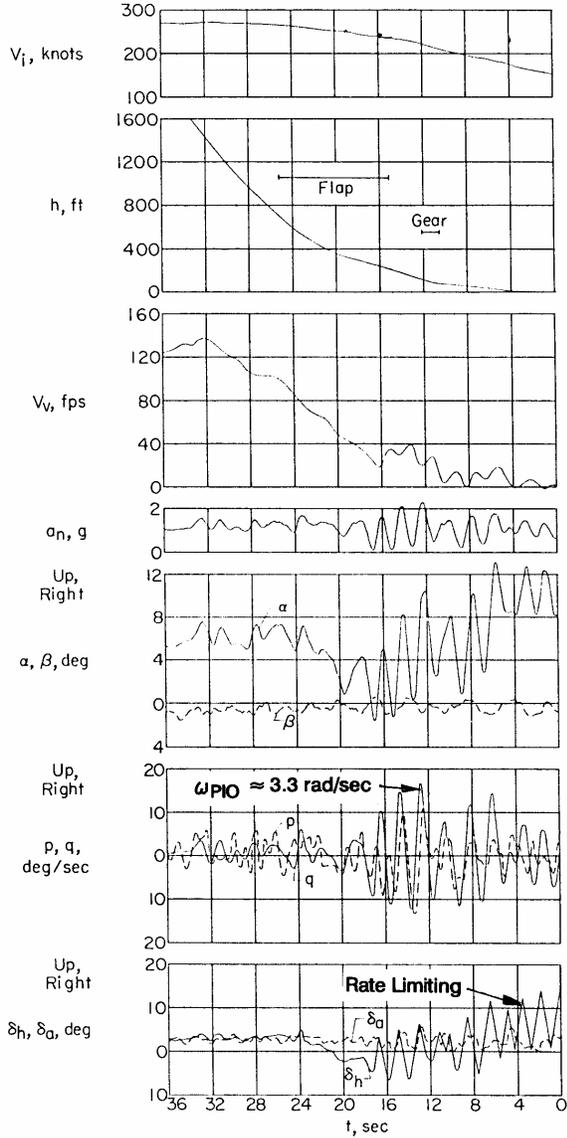


Figure 7. Time History of the X-15 Flight 1-1-5 Landing Flare (from Ref. 10)

most of the applied handling qualities measures. Figure 8 also indicates that the instability frequency for the linear system with a synchronous pilot loop closure is 5.31 rad/sec. This is almost twice that of the observed PIO frequency of 3.3 rad/sec.

Impact of Rate Limiting

A complete describing function analysis of the X-15 PIO is provided in Ref. 2 and Ref. 7 and is therefore not included here. Instead the impact of rate limiting on the airplane bandwidth handling qualities/PIO criteria¹³ is illustrated. Rate limiting will occur when the input frequency exceeds what Ref. 2 refers to as the “saturation frequency” and what Ref. 8 refers to as the

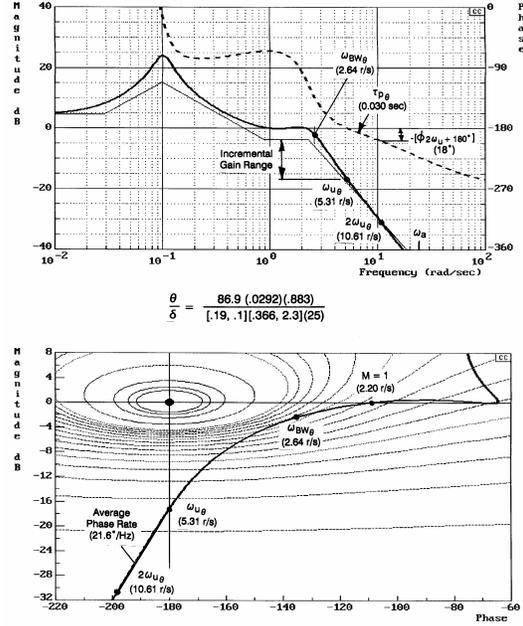


Figure 8. X-15 Pitch Attitude Frequency Response at the Flight 1-1-5 Landing Flare Condition

“onset frequency.” This frequency can be defined in terms of K^* as follows,

$$\omega_{onset} = \frac{2K^*}{\pi} \omega_i$$

Noting that K^* is a function of input amplitude, A , and maximum rate, V_L , onset frequencies were computed for the X-15 example where $V_L = 15$ deg/sec, $\omega_i = 3.3$ rad/sec (i.e., the PIO frequency), and $A = 3, 6, 9, 12$, and 15 degrees. The results are plotted in Figure 9 with a simplified version of the Ref. 13 PIO boundaries included. When $A = 3$ deg, the onset frequency is above the PIO frequency and no rate limiting occurs. Thus, the bandwidth and phase delay parameters represent the linear system values for this case. As the input amplitude increases to 6 deg, the onset frequency moves below the PIO frequency. Here, we are in what is referred to in Ref. 2 as the near saturation region and the effect of the rate limiting is minimal as evident by the small shift in the bandwidth/phase delay parameter plane. For the remaining three input amplitude cases, the aircraft has entered the highly saturated region and is clearly susceptible to PIO. Note the significant jump in phase delay for these points that results from the significant added phase lag from the rate limiting. Beyond $A = 9$ deg the bandwidth frequency is determined by the gain margin frequency, and thus also shows a dramatic shift. The impact of the rate limiting is to greatly reduce bandwidth frequency and increase phase delay, both contributing factors to a poor handling aircraft that is also highly susceptible to PIO.

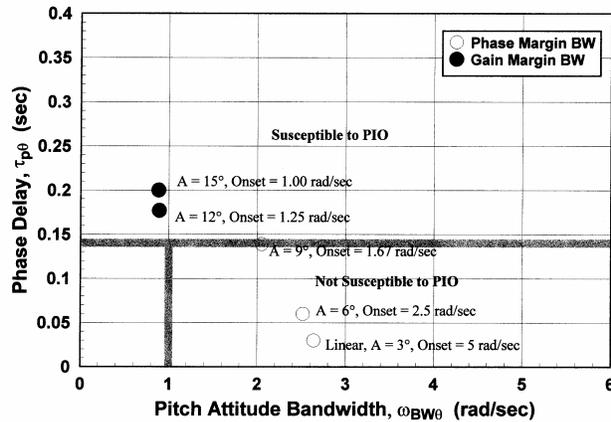


Figure 9. X-15 Example Bandwidth/Phase Delay Variations with Increasing Rate Saturation

THE HAVE LIMITS FLIGHT TEST PROGRAM

Description

In 1997, students of the USAF Test Pilot School at Edwards AFB, California, conducted a project called HAVE LIMITS.¹⁴ This student project provided the first systematic, quantitative information on the interchange between rate limiting and PIO. Three pilots flew a total of nine sorties on a variable-stability NT-33A, operated for the Air Force by Calspan (now Veridian). The experiment focused entirely on the pitch axis, with three basic aircraft models and seven levels of actuator rate limiting from 10 deg/sec to 157 deg/sec. The task was a pitch and roll attitude tracking task using commands displayed on a head-up display (HUD).

HAVE LIMITS Configurations

Three configurations from the HAVE LIMITS program illustrate the interaction of rate limiting and aircraft dynamics on the occurrence of PIO.

Configuration 2D — This configuration was designed with a very high pitch attitude Bandwidth and low Phase Delay without augmentation. In the flight tests, it was PIO-proof until actuator rate was extremely low. Selected signals for a 20-second run segment from a HUD attitude tracking task are shown in Figure 10 for configuration 2D with an actuator rate limit of 20 deg/sec. The pilot for this run considered the combination of dynamics and rate limiting to be very good, with an assigned HQR of 2 and PIOR of 2. The baseline (no limit) case was rated a 4 and a 2, respectively.

The top responses in Figure 10 are cockpit stick force, F_{es} , in pounds, and pitch rate, q , in deg/sec. The middle set shows commanded and “actual” elevator positions (the “actual” surface was a simulated elevator for the configuration, modeled in the variable-stability system), both in deg of surface. The bottom set is derived rate for the commanded and “actual” actuators, in deg/sec. Only momentary differences between commanded and actual actuator position and rate can be observed in Figure 10. This configuration was sufficiently high-bandwidth that the pilot could control it with small, short control inputs that seldom reached the limit. Rate limiting actually served to lessen the initial abruptness of the configuration.

The pitch rate and stick force traces show no real evidence of a sustained oscillation, or of out-of-phase character. There is no hint of PIO. This is significant since configuration 2D represents an airplane that is made to be Level 1 by inherent aerodynamics, not with feedbacks.

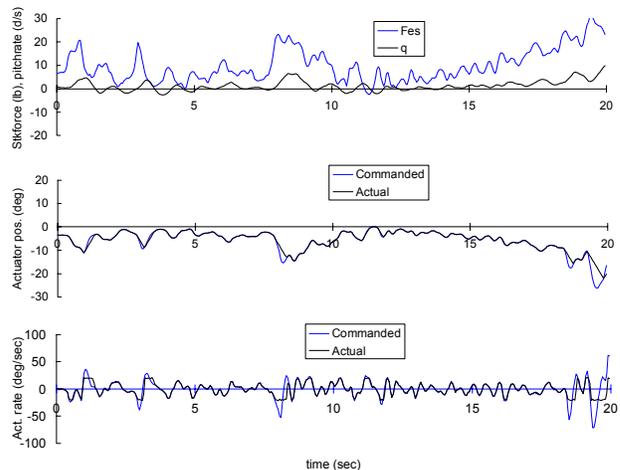


Figure 10. Example Time History Segment for Configuration 2D with 20 deg/sec Rate Limit (no PIO Tendencies)

Configuration 2P — Configuration 2P, created by inserting a prefilter on configuration 2D, had deficiencies even with no limiting and received PIO ratings of 4 and 5 for the lowest rate limits. The time history in Figure 11 is for a run with configuration 2P that received an HQR of 7 and a PIOR of 4. This pilot rated the no-limit case a 4 and a 2, respectively, so rate limiting has obviously increased the tendency for PIO.

Evidence of an incipient PIO can be seen in the stick force and pitch rate traces in Figure 11 between about nine and 13 seconds. The large aft stick input at eight seconds drives the actuator onto its limit, and the result is a cycle of an oscillation where stick force and pitch

rate are clearly out of phase with each other. This does not develop into a full-blown PIO, however, either because the task demands were lower or because the pilot backed out of the control loop. Nevertheless, there is flight evidence of a potential for PIO.

Configuration 2DU — With augmentation, configuration 2DU was identical to the PIO-resistant 2D. But the good dynamics were obtained by wrapping the augmentation around a bare airframe that was very unstable with an aperiodic divergence in the short period. Because the bare airframe was highly unstable, any amount of rate limiting was detrimental. The pilot for this example rated the no-limit baseline case HQRs of 5 and 4 (two evaluations) and PIORs of 3 and 2 because of pitch bobbles and abruptness.

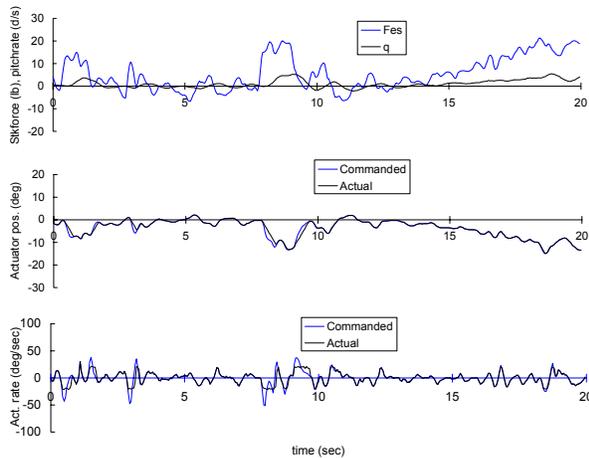


Figure 11. Example Time History Segment for Configuration 2P with 20 deg/sec Rate Limit (Tendency to PIO Evidenced by Phasing Between Stick Force and Pitch Rate at $t = 9-13$ sec)

The time history in Figure 12 is typical of all of this pilot’s runs with configuration 2DU with rate limiting. It is also a classic rate-limited PIO. A small PIO begins at about three seconds, with intermittent rate saturation of the actuator. Demand (commanded) is not much greater than achieved (actual) until around 10 seconds, where the PIO becomes divergent. At approximately 15 seconds in this segment the safety pilot took over control of the airplane as the divergence became too large for flight safety.

Results for configuration 2DU are significant: this was a basically good-flying airplane, with no inherent characteristics indicative of a PIO-prone design, that degenerated into a highly PIO-prone airplane solely in the presence of rate limiting. The cause, of course, is the use of active augmentation to provide the good flying qualities on an airplane that otherwise was almost unflyable.

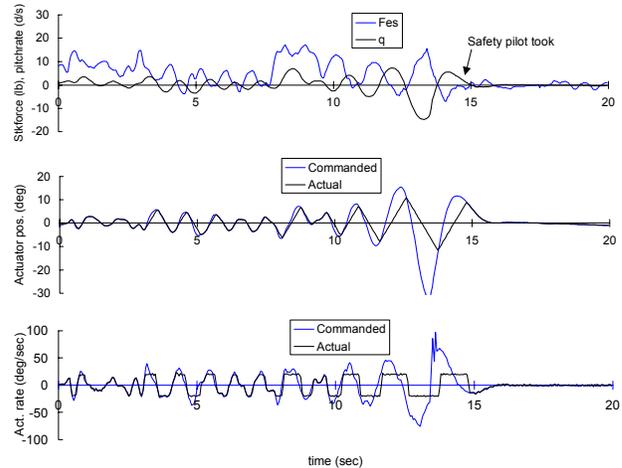


Figure 12. Example Time History Segment for Configuration 2DU with 20 deg/sec Rate Limit (Evidence of PIO Starting at $t = 4$ sec, Divergent at $t = 10$ sec)

Measures of Actuator Usage

The HAVE LIMITS data provide a base for assessment of actuator demands, and for possible methods for avoiding excessive demands. Measures of actuator usage were made by computing cumulative actuator rates for the flight segments represented by Figure 10 through Figure 12. Figure 13 shows a comparison of the cumulative actuator rates for two configurations (2D and 2DU) and two values of rate limiting. Actuator rate usage plots are not shown for configuration 2P. The plots for this case are very similar to those for 2D and do not show any significant new information.

The curves in Figure 13 show actuator rates as a percentage of time for the 20-second run segments. The lowest curve on this figure is for configuration 2D with no added rate limit (labeled “2DU/157” – 157 deg/sec was the reported rate limit of the NT-33A’s physical elevator surface). For the 20-second time slice, actuator rates are below 20 deg/sec for almost 90% of the run.

Just above the no-limit curve is the cumulative curve for the configuration 2D run of Figure 10. Rate demands are almost unchanged from the no-limit case. The differences are probably due either to run-to-run variations in pilot gain, or to the slightly greater demand resulting from the small amount of time the actuator was on its rate limit. In either case, the pilot did not report any PIO tendencies.

The curves for configuration 2DU in Figure 13 reflect the high demands on elevator for stabilizing the unstable airframe. For the unlimited case (“2DU/157”),

actuator rate is below 20 deg/sec for over 90% of the segment. For the final 10%, the demand increases sharply. When rate limiting at 20 deg/sec is introduced, the demand cannot be met, the actuator is rate-limited for most of the run, and a divergent PIO occurs.

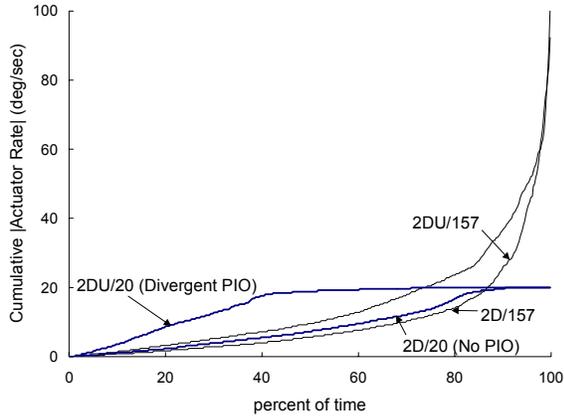


Figure 13. Actuator Rates for Configurations 2D & 2DU (Numbers on Labels Indicate Rate Limits in deg/sec)

Pilot Rating Results

The effects of rate limiting on pilot opinion for the HAVE LIMITS configurations are illustrated in Figure 14. Mean Cooper-Harper Handling Qualities Ratings and PIO Tendency Ratings are plotted against elevator rate limit. For the baseline configuration 2D, no pilot reported PIO, even when rate limits were as low as 10 deg/sec, though handling qualities were degraded (mean HQR of 5). Because this configuration was slightly abrupt, a small amount of rate limiting actually improved the ratings, as evidenced by the dips in HQR and PIOR with a rate limit of 30 deg/sec. Configuration 2P had degraded handling qualities and a slight tendency to PIO as rate limits were decreased. For Configuration 2DU, any decrease in rate limits resulted in at least one divergent PIO.

Pilot Rating Variability

In the HAVE LIMITS flight experiment, the limited number of configurations (three, with seven rate limits) and pilots (three) did not result in significant variations in pilot opinion. A previous USAF Test Pilot School project (HAVE GRIP,¹⁵) found significant variations in pilot ratings and concluded that the evaluation task in that experiment (precision offset landing) was “insufficient to consistently uncover handling qualities deficiencies” resulting from rate limiting. What the authors of Ref. 15 actually discovered, however, was the variability in pilot opinion resulting from the introduction of a highly nonlinear phenomenon.

A more recent ground-based simulation^{13,16} confirmed that the nonlinear nature of rate limiting makes its impact very dependent upon piloting technique. As an example, Figure 15, reproduced from Ref. 16, shows the pilot ratings from nine different pilots for three configurations evaluated in the simulation.

The configuration labeled 2DR30 in Figure 15 is the same as HAVE LIMITS configuration 2D, with a 30 deg/sec rate limit (the case with the best pilot ratings from flight, Figure 14). All nine pilots gave this configuration good HQRs and PIORs (Pilot B objected to a pitch bobble tendency).

Configuration 1DL3 in Figure 15 was a low-short-period model with added lags to make it susceptible to Category I PIO; all pilots except Pilot B reported at least one PIO with this configuration. Their HQRs and PIORs are reasonably consistent.

Configuration 2DUR20 is 2DU from HAVE LIMITS with a 20 deg/sec rate limit. In the flight program (Figure 14), this configuration was unflyable with divergent PIO. All three pilots in HAVE LIMITS assigned an HQR of 10 and a PIOR of 5 or 6. In the simulator, of the seven pilots who flew this configuration, six also experienced divergent PIOs and assigned HQRs of 10 and PIORs of 5 or 6. The seventh, Pilot D, did not experience rate limiting or divergent PIO, and he flew the configuration twice. In Ref. 13 and Ref. 16 an analysis of closed-loop control behavior shows that Pilot D adopted the lowest bandwidth of all the pilots.

If Pilot D were the chief test pilot for an airplane manufacturer, it is conceivable that his simulator experiences with an airplane like configuration 2DUR20 could be convincing enough to persuade the company to build a prototype. Based on the experiences of the other six pilots, and of the three pilots in the HAVE LIMITS flight program, that would be a serious mistake.

ALLEVIATING THE EFFECTS OF RATE LIMITING

Despite the best intentions of flight control systems designers, actuator rate limiting may be inevitable for any aircraft. Even if the actuation system of a particular aircraft is designed to provide a margin from surface rate limiting, it is likely that at some point, changes to the software will result in the implementation of a software rate limiter, or changes in the aircraft’s mission statement will result in a configuration that will overburden the surface rate limits. As this paper has demonstrated, for many aircraft the consequences of rate limiting will be negligible.

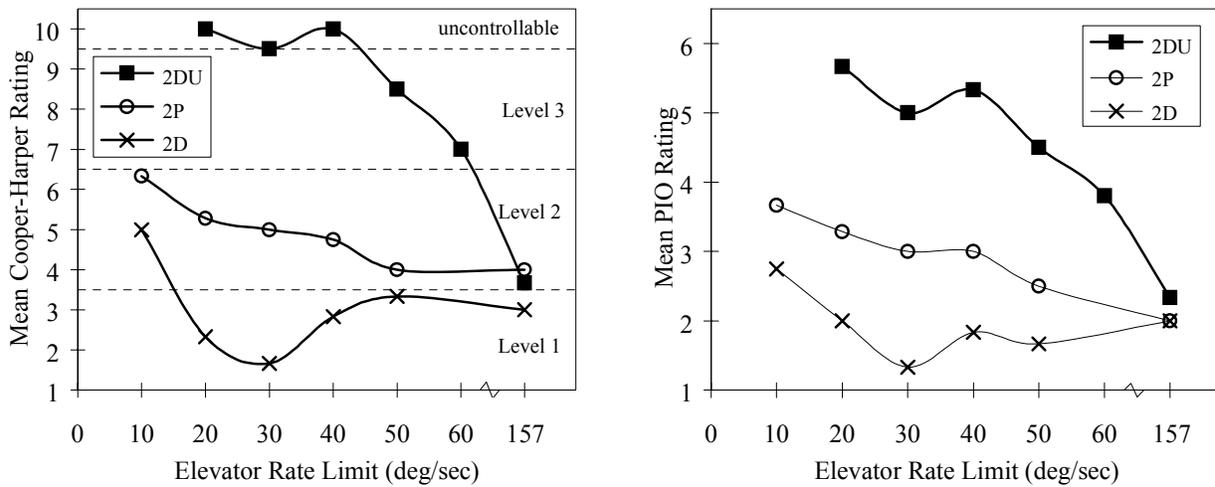


Figure 14. Mean Cooper-Harper Handling Qualities Ratings and PIO Tendency Ratings for HAVE LIMITS Configurations

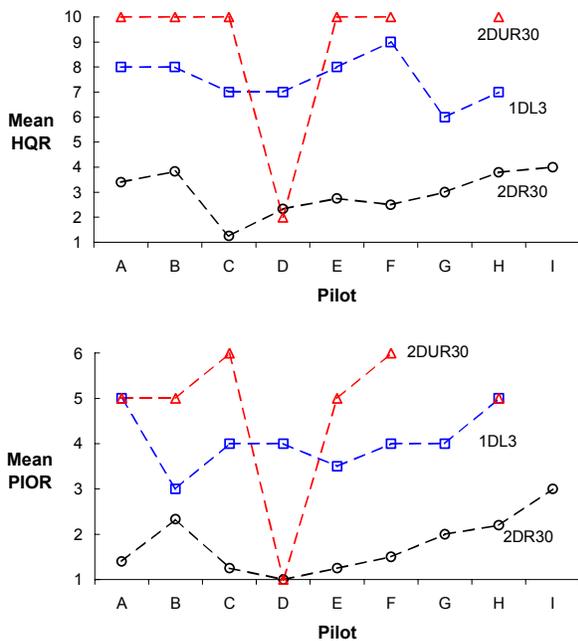


Figure 15. Pilot Rating Variability (Reproduced from Ref. 16)

In the past, when the consequences have not been negligible, and when it has been impractical to significantly change the control laws or the surface actuation system, solutions have been pursued through software. These solutions take the form of filters that are meant to reduce the likelihood of rate limiting, or to mitigate the negative effects if rate limiting is encountered. As such, the filters represent only a “patch” to minimize the existing deficiencies, and not a real solution to the problem.

PIO Suppression Filters

Repeating an observation from the introduction of this paper, nothing seems to inspire significant new handling qualities research like a high-profile PIO event. A PIO during landing on the Space Shuttle Enterprise ALT-5 in 1977 is a prime example. This PIO resulted from a combination of basic shuttle handling qualities, time delay through the digital flight control computers, and rate limiting of the elevator actuators.¹⁷ Subsequent analysis and simulation¹⁸ identified the unusual pilot location relative to the aircraft’s center of rotation as a significant factor as well.

Methods for mitigating PIO on the Space Shuttle were investigated by NASA. Reference 17 describes a PIO suppression (PIOS) filter that was designed to reduce pilot gain when potential for PIO is high, while minimizing any additional phase lag. To achieve the desired gain reduction, the filter modifies the stick shaping function as a function of the amplitude and frequency of the pilot’s input, thereby reducing the amount of rate limiting.

The PIOS filter was implemented in the Shuttle control laws, and no pitch PIOs have been reported in the open literature since the 1977 ALT-5 event. Subsequent control law changes resulted in the removal of the PIOS filter.

Based on the apparent success of the application of the Shuttle PIO suppression filter, two types of filters were evaluated by NASA Dryden Flight Research Center in three flight test programs.¹⁹ These evaluations found that PIO suppression filters can improve handling qualities of fighter aircraft that are susceptible to PIO

because of excessive time delay. Of course, such delays are the primary degrading effect of rate limiting.

The U.S. Navy's Phase Compensating Rate Limiter

In the late 1980's, engineers with the Naval Air Systems Command (NAVAIR) of the U.S. Navy developed a phase compensating filter to reduce the effects of rate limiting.²⁰ The filter is time-domain-based and reverses the direction of the actuator when 1) the actuator is on its rate limit, and 2) the actuator is traveling in opposite direction from its command. This filter, patented by the U.S. Government, is similar to that tested by NASA on variable-stability aircraft in the 1990's.²¹ Some of the problems encountered with that version of the filter include sensitivity to input noise and a tendency to develop an offset bias over time. Most of the problems with the filter can be overcome with additional complexity.

The SAAB System

The phase compensation technique developed by SAAB to overcome the deficiencies associated with rate limiting on the JAS-39 Gripen⁵ uses a feedback signal through a lowpass filter that will almost immediately reverse the direction of the rate limit output when the input reverses direction. The technique also features a bypass circuit that insures that only the low-frequency components of the input are limited with phase compensation. Unlike the Navy's PCRL, the Gripen design introduces some additional effective delay through the lag filter. After successful simulation and flight test evaluations, the rate limiters with feedback and bypass were qualified for use in the Gripen production flight control system software.

The DLR System

One final rate limiting alleviation scheme to discuss is a phase compensation technique developed by DLR. As described in Ref. 22 a simple phase compensation algorithm for a rate limiting element is derived directly from the describing function relationships of Hanke⁸ shown earlier in this paper. The compensator does not use feedback or logic, but instead reduces the input amplitude as a function of frequency by the same amount that the amplitude is reduced by the rate limiter. Hence, the rate limiter can follow the input signal without added phase lag. The concept has been successfully demonstrated in-flight for a wide range of system inputs using the DLR Advanced Technologies Testing Aircraft System or ATTAS.

SUMMARY AND CONCLUSIONS

- *Rate limiting of control effector actuators can have dire effects on handling qualities and PIO.* The magnitude of the effects depends on three basic factors:
 1. How long the actuator is rate-limited.
 2. How much more the pilot/flight control system demands of the airplane.
 3. The consequences on aircraft dynamics of encountering the limit.
- *Rate limiting can cause PIO.* Fundamentally, rate limiting introduces phase lag into the aircraft's response. This alone can be sufficient to lead to pilot-induced oscillations, depending upon the characteristics of the aircraft. Flight 1-1-5 of the X-15 provides a perfect example of this effect. If the aircraft is augmented, reaching a rate limit also means a sudden change in response dynamics as the aircraft transitions from the closed-loop to the open-loop response. If the open-loop response is unstable, divergent PIO is likely, especially when the additional phase lag is introduced.
- *Rate limiting does not necessarily cause PIO.* If the response characteristics of the basic aircraft are of sufficiently high bandwidth the additional phase loss due to rate limiting will not lead to PIO.
- *PIO can lead to rate limiting.* While not specifically addressed in this paper, it is apparent that if an aircraft is susceptible to PIO even when there is no rate limiting, the occurrence of PIO can push control surface effectors onto their rate limits.
- *Rate limiting is a nonlinear phenomenon and hence is highly dependent upon pilot technique.* It is possible for two pilots to have vastly different opinions of the same aircraft simply depending upon their control strategies. Despite our best attempts at devising consistent, repeatable evaluation tasks, some pilots will naturally adopt strategies that will minimize the chance of ever encountering the rate limiting. This is why it is critically important that aircraft, and highly-augmented aircraft especially, be evaluated by as many pilots as possible.

REFERENCES

1. Ashkenas I. L., H. R. Jex, and D. T. McRuer, *Pilot-Induced Oscillations: Their Causes and Analysis*, Northrop-Norair Rept., NOR 64-143, June 1964.
2. Klyde, D. H., D. T. McRuer, and T. T. Myers, *Unified Pilot-Induced Oscillation Theory Volume I: PIO Analysis with Linear and Nonlinear Vehicle Characteristics, Including Rate Limiting*, WL-TR-96-3028, Dec. 1995.
3. Smith, J. W., and D. T. Berry, *Analysis of Longitudinal Pilot-Induced Oscillation Tendencies of YF-12 Aircraft*, NASA TN D-7900, Feb. 1975.
4. Dornheim, M. A., "Report Pinpoints Factors Leading to YF-22 Crash," *Aviation Week and Space Technology*, 9 Nov. 1992, pp. 53-54.
5. Rundqwist, L., and R. Hillgren, "Phase Compensation of Rate Limiters in JAS 39 Gripen," AIAA-96-3368 presented at the *Atmospheric Flight Mechanics Conference*, San Diego, CA, 29-31 July 1996.
6. Aviation Safety and Pilot Control: Understanding and Preventing Unfavorable Pilot-Vehicle Interactions, Committee on the Effects of Aircraft-Pilot Coupling on Flight Safety, National Academy Press, Washington D.C., 1997.
7. Klyde, D. H., D. T. McRuer, T. T. Myers, "Pilot-Induced Oscillation Analysis and Prediction with Actuator Rate Limiting," *Journal of Guidance, Control, and Dynamics*, Volume 20, Number 1, Jan. - Feb. 1997, pp. 81 to 89.
8. Hanke, D., "Handling Qualities Analysis on Rate Limiting Elements in Flight Control Systems," *Flight Vehicle Integration Panel Workshop on Pilot Induced Oscillations*, AGARD-AR-335, Feb. 1995.
9. Wallace, L. E., Flights of Discovery: 50 Years at the NASA Dryden Flight Research Center, NASA SP-4309, NASA History Office, Washington, D.C., 1996.
10. Matranga, G. J., *Analysis of X-15 Landing Approach and Flare Characteristics Determined from the First 30 Flights*, NASA TN-D-1057, July 1961.
11. Heffley, R. K., and W. F. Jewell, *Aircraft Handling Qualities Data*, NASA CR-2144, Dec. 1972.
12. Taylor, L. W., Jr., and J. W. Smith, *An Analysis of the Limit-Cycle and Structural-Resonance Characteristics of the X-15 Stability Augmentation System*, NASA TN D-4287, Dec. 1967.
13. Mitchell, D. G., and R. H. Hoh, *Development of Methods and Devices to Predict and Prevent Pilot-Induced Oscillations*, AFRL-VA-WP-TR-2000-3046, Dec. 2000.
14. Kish, B. A., Capt., USAF, et al., *A Limited Flight Test Investigation of Pilot-Induced Oscillation due to Elevator Rate Limiting (HAVE LIMITS)*, AFFTC-TR-97-12, June 1997.
15. Peters, P. J., Capt., USAF, et al., *Limited Investigation of the Effects of Elevator Rate Limiting and Stick Dynamics on Longitudinal Pilot Induced Oscillations (HAVE GRIP)*, AFFTC-TR-96-37, Dec. 1996.
16. Mitchell, D. G., and B. K. Stadler, "Simulation Investigation of Category I and II PIO," AIAA-99-4003 presented at the *Atmospheric Flight Mechanics Conference Proceedings*, Portland, OR, August 1999, pp. 1-11.
17. Smith, J. W., and J. W. Edwards, *Design of a Nonlinear Adaptive Filter for Suppression of Shuttle Pilot-Induced Oscillation Tendencies*, NASA TM-81349, 1980.
18. Teper, G. L., R. J. DiMarco, I. L. Ashkenas, and R. H. Hoh, *Analyses of Shuttle Orbiter Approach and Landing Conditions*, NASA CR-163108, July 1981.
19. Shafer, M. F., R. E. Smith, J. F. Stewart, and R. E. Bailey, *Flight Test Experience with Pilot-Induced Oscillation Suppression Filters*, NASA TM-86028, Jan. 1984.
20. Koper, J., *An Approach for Compensating Actuator Rate Saturation*, NADC-87120-60, Aug. 1987.
21. Deppe, P. R., C. R. Chalk, M. F. Shafer, *Flight Evaluation of an Aircraft with Side and Center Stick Controllers and Rate-Limited Ailerons*, NASA CR-198055, Nov. 1996.
22. Hanke, D., *Phase Compensation: A means of Preventing Aircraft-Pilot Coupling Caused by Rate Limitation*, DLR-Forschungsbericht 98-15, 1998.