Development and Application of a Convolution Technique for Flying Qualities Research

H.T. Breul,* T.L. Keller,† and R.C. Weston‡
Grumman Corporate Research Center, Bethpage, New York

The mathematical concepts of convolution and superposition are used in a modeling scheme that creates a new capability for studying some modern flying qualities problems. The process computes solutions of very high-order six-degree-of-freedom linear vehicle dynamics on a minicomputer with input/output process times under 10 msec, substituting easily constructed, vehicle step-responses in the time domain for the usual differential equations. Thus, it offers a valuable way to study the flying qualities of nonclassical control responses produced by modern, highly augmented vehicles. Furthermore, standard non-real-time stability and control analysis software can be used to quickly produce real-time minicomputer models, making relatively inexpensive simulator flying qualities studies possible very early in the design/development cycle. A current application to the study of adverse lateral control responses at hover on a large-amplitude research simulator is reviewed.

Introduction

WITH the advent of fly-by-wire control schemes, control-configured design, direct force devices, and powerful onboard computers, modern aircraft can be made to respond to commands in dramatically new and different ways. Their controls can be blended in complicated ways, providing an unprecedented capability to design for the flying qualities of choice. The problem is that there is no complete understanding of what constitutes good flying qualities over this greatly expanded range of control performance, and, as control systems become more capable, they grow more complex, making it increasingly difficult to define flying qualities requirements in simple conventional terms. Thus, standard techniques for describing aircraft behavior are becoming less well suited to flying qualities research. The approach described here addresses this problem by adapting the mathematical concepts of convolution and superposition to a modeling scheme that provides a flight simulation technique relatively uninhibited by system complexity.

Approach

In the simulation technique to be described, linearity is implicit, permitting in most cases only small excursions of a vehicle away from some trim condition. This is not a severe restriction for studies examining flying qualities at single trim points or trim points that change slowly.

A linear, time-invariant dynamic system can be modeled in the time domain by specifying a set of "motion signatures" comprising the various responses to step disturbances in each of the degrees of freedom. There are two significant research advantages of this kind of time-domain modeling over the more usual differential equation approach.

First, a motion signature, which represents a time history of actual vehicle motion, is directly sensible by the pilot, and its shape can be easily and arbitrarily varied. Thus, the experimenter can directly manipulate the properties sensed by the pilot (cockpit kinematics) without any intervening transformation as in the differential-equation approach. In other words, the link between the experimenter and the arbiter (the pilot) in flying qualities research is more direct.

Second, in attempting to tease out the true cause of a pilot's reaction to a particular vehicle response, the experimenter can easily make a detailed change in a motion signature that would be next to impossible to duplicate in a differential equation because it would require the complicated adjustment of many parameters in a very high-order system. In other words, the experimenter has a much more flexible research tool at his/her disposal.

Convolution and superposition are the mathematical concepts that make it possible to create vehicle simulation models using only geometric shapes of step responses. Convolution provides a mechanism for modeling the dynamics of a system-variable in response to a stimulus using the step response. That is, if a system-variable's response to a step is known, the convolution algorithm describes how that variable responds to arbitrary inputs. Superposition permits the separate responses to various stimuli to be summed, yielding a total composite response. For example, the total roll rate response is obtained by first computing the various responses to each pilot command (rudder pedals, lateral stick, throttle, etc.) and then summing the individual outputs.

Computational Algorithm

The single-input/single-output convolution integral (Ref. 1) can be written as

$$G(t)=F(t)\ H(o)+\int_o^t F(\tau)\ \dot{H}(t-\tau)\mathrm{d}\tau \qquad (1)$$

where H is the step response—the "motion signature," $\dot{H}$ the impulse response, F the input or "command," and G the system's response to the command. The sampled-data form used with uniform sampling for solution on a digital computer becomes

$$G_N=F_N H_o+\sum_{K=0}^{K=N} F_K \dot{H}_{N-K}\Delta\tau \qquad (2)$$

where $F_K=F(K\Delta\tau)$, $\dot{H}_K=\dot{H}(K\Delta\tau)$, $\Delta\tau$ = the sampling interval, and N = the greatest integer $\leq t/\Delta\tau$. Only the current value of G is required during each program cycle $\Delta\tau$, and thus, the very large computational effort normally associated

Presented as Paper 85-3062 at the AIAA/AHS/ASEE Aircraft Design Systems and Operations Meeting, Colorado Springs, CO, Oct. 14–16, 1985; received Dec. 6, 1985; revision received Aug. 22, 1986.

*Head, Flying Qualities and Simulation Branch, System Sciences Directorate.
†Senior Staff Scientist, System Sciences Directorate. Member AIAA.
‡Research Engineer, System Sciences Directorate.

with a complete convolution (i.e., the calculation of G_n for all n, $0 \le n \le t$) is spread out over n program cycles. For the applications of interest here, the step response always starts at zero; that is, $H_o = 0$. Thus, Eq. (2) can be simplified to

$$G_N = \sum_{K=0}^{K=N} F_K \dot{H}_{N-K} \Delta\tau \tag{3}$$

This is the form used to compute the individual pieces of the vehicle-system output in response to each separate pilot command. Although the impulse response $\dot{H}$ is the quantity actually used to compute the system output, the step response H is deemed the "motion signature" and independent variable.§ This means that the independent variables in an experiment are described as elements of the step response, but the impulse response (derivative of the step response) is stored in the computer for the real-time computation of system output. Note that $\Delta\tau$, the period of the computation cycle, must be the same as the time between the samples of $\dot{H}$ stored in memory. Thus, the desired resolution of $\dot{H}$ determines $\Delta\tau$. During each program cycle, the current input must be sampled and the sum of all the products of F_k and $\dot{H}_{n-k}$ must be formed and output as an updated system response G_n. Equations (4) demonstrate how, despite the large number of computations required, a very small input/output response time is achieved. They are expansions of Eq. (3) at successive increments of time for $0 \le t \le m\Delta\tau$, where $\dot{H}$ is $n\Delta\tau$ long and $m > n$. Remember that the impulse response values $\dot{H}_o, \dot{H}_1, \ldots, \dot{H}_n$ are stored in memory.

$$G_o = F_o \dot{H}_o \Delta\tau \tag{4a}$$

$$G_1 = F_1 \dot{H}_o \Delta\tau + \{F_o \dot{H}_1\} \Delta\tau \tag{4b}$$

$$G_2 = F_2 \dot{H}_o \Delta\tau + \{F_1 \dot{H}_1 + F_o \dot{H}_2\} \Delta\tau \tag{4c}$$

$$\vdots$$

$$G_m = F_m \dot{H}_o \Delta\tau + \{F_{m-1} \dot{H}_1 + F_{m-2} \dot{H}_2 + \ldots + F_{m-n} \dot{H}_n\} \Delta\tau \tag{4d}$$

Starting at $t = 0$, $F(t)$ is sampled to get F_o, and the single product in Eq. (4a) is formed and output as G_o. F_o will be needed for products with the other elements of $\dot{H}_n$ to produce the outputs G at subsequent times. These are now formed and stored in a partial-answer array, and F_o is discarded. Next, at $t = \Delta\tau$, $F(t)$ is sampled to get F_1, and the first product in Eq. (4b) is formed, added to the second product (which had been formed at $t = 0$), and output as G_1. Now F_1 also will be needed to form products with elements of $\dot{H}_n$ for future outputs. These are formed and added to the appropriate products previously formed with F_o. The process continues in this manner so that, at any time $t = m\Delta\tau$ [see Eq. (4d)], the quantity in braces { } has already been computed. When F_m is obtained by sampling $F(t)$, only one multiplication/addition operation is needed to produce the current answer G_m. This new result is obtained quickly and output immediately. Then F_m is multiplied by the other $\dot{H}_n$, and these products are added to the appropriate terms in the partial-answer array that are needed for future outputs.

This technique allows almost all of the large number of convolution calculations required in each program cycle to be deferred until after the current answer has been output. The program thus achieves an input/output response time currently of the order of 10 msec for a full six-degrees-of-freedom (6-DOF) simulation with up to 72 separate convolutions. In fact, very little of this time is needed to compute the output. Almost all of it is used to perform routine overhead-type functions such as calibrating and scaling the sampled quantities, sorting and combining the individual convolution outputs, performing axis transformations, and scaling and limiting the outputs. This short response time is achieved on a relatively old (>5 yr) and slow minicomputer.¶ Newer minis are 6 to 10 times faster and would provide input/output times of 2 msec or less.

The other time-critical convolution calculation is the simple number-crunching computation of the term in braces in Eq. (4d) that occurs after the current answer is output. Because the "current" brace contains only "old" inputs, i.e., values sampled during previous program cycles, it is most easily produced by a process of accumulation. A partial-answer vector resides in an array in which n future values of the brace are accumulating ($\dot{H}$ is $n\Delta\tau$ long). The elements are adjusted and shifted each program cycle according to Eq. (5).

$$PA_i = F_m \times \dot{H}_i + PA_{i+1} \quad i = 1, \ldots, n \tag{5}$$

where PA is the partial-answer array, F_m the current sample of the input, $\dot{H}$ the array of impulse response values, and i the array index. In this way the brace value [Eq. (4d)] needed to form the new output during the next program cycle accumulates over the previous cycles and appears when needed in the first element (PA_1) of the array. The calculation of Eq. (5) is performed for each convolution in progress and typically uses up more than half of the program cycle time. Thus, there is a tradeoff between cycle time and the total number of values of all impulse responses.

The problem of fitting a large number of long convolution calculations into a program cycle has been significantly alleviated by taking advantage of the dual nature of many typical aircraft responses: rapid variation initially, and then a long, slow decay. Breaking these responses into two separate responses, one (containing only the high-frequency part) being operated on at every program cycle and the other (containing only the long, slow part) at every nth cycle, with n typically 4 or 5, has reduced total convolution calculation times as much as 70% in some cases.

Creating Motion Signatures

In a flying qualities study program where realistic 6-DOF simulation is used, there is a myriad of possible combinations of step responses that could be implemented. To invest the simulation with realism and reduce the number of possible configurations to a manageable size, a set of motion signatures characterizing an existing vehicle design at a specific trim point is selected. These are generated off line by a large mainframe computer solving the standard equations of motion. From the resulting set (two dozen or so) of basic motion signatures, only one or two are chosen to be systematically modified in a flying qualities experiment. The rest remain unchanged, providing a real-aircraft context in which to conduct the flying qualities trials. Values of the signatures are generated at time intervals (e.g., one msec) much shorter than required so that the fine structure of each signature is clearly revealed. These "raw" responses are not immediately usable and must be transformed in several ways. For example, a real-time convolution is practical only if the motion signature decays to zero in a reasonable length of time. Some typical aircraft responses meet this criterion, and some others come close enough that they can be forced into compliance. Those that do not usually have derivatives that

§This is somewhat arbitrary, but intuitively it seems easier to think of the control time history as a result of the pilot's making a series of incremental steps in control rather than a series of discrete pulses.

¶The minicomputer currently in use is a Hewlett Packard 1000F with 256 kilobytes of memory and a 12-megabyte hard disk.

do; these can be used instead, requiring that the corresponding outputs be integrated. (Differentiation noise is not a problem here because the responses are generated as solutions of differential equations or are manually constructed and are inherently quite smooth. Furthermore, subsequent integration of their outputs effectively removes any small residual noise.) A large package of offline software has been developed to facilitate the creation and management of motion signatures.

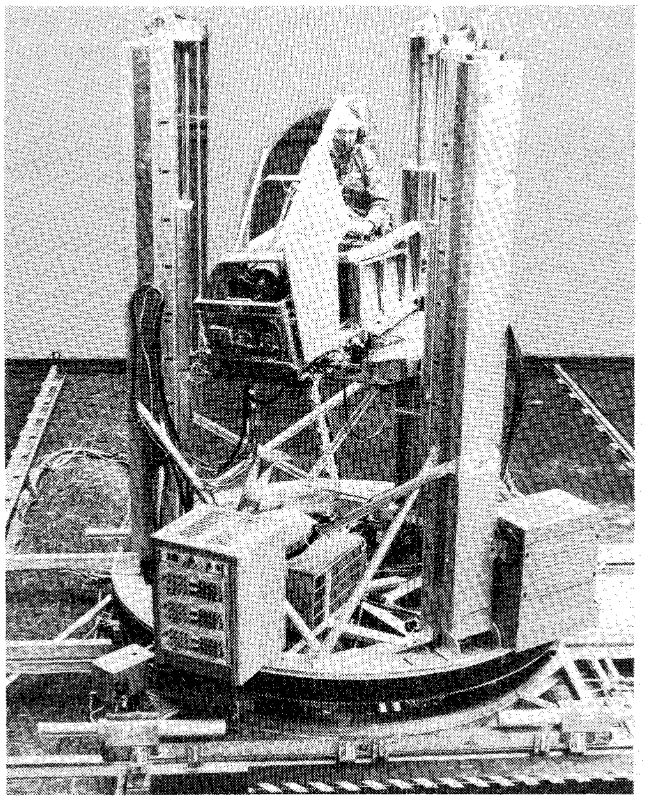

Fig. 1 The Large Amplitude Research Simulator (LARS).

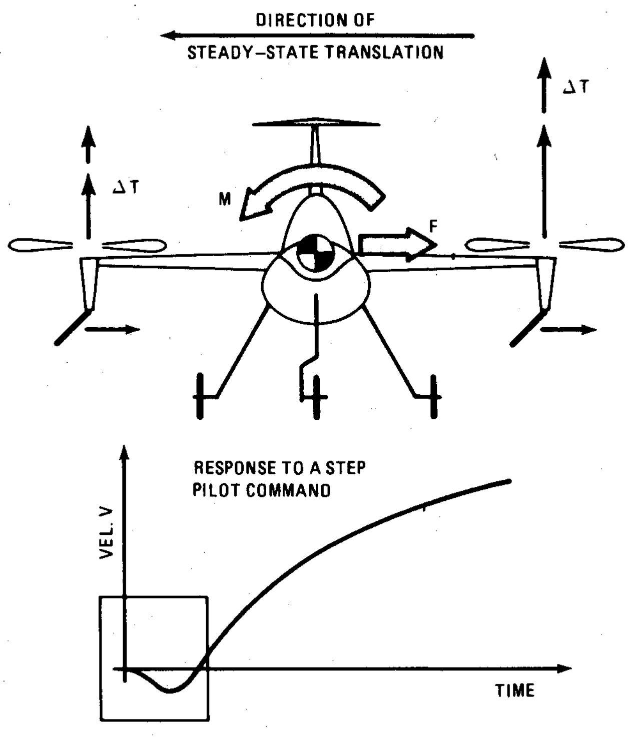

Fig. 2 Adverse lateral control response at hover.

Accrued Advantages

Several advantages accrue to the convolution technique ("Convo") as it is currently applied to flight simulation:

1) Sets of motion signatures describing an extensive range of aircraft dynamics can be stored on disk, accessible "on the fly" at the touch of a button. Because the exchange of signatures (implementation of a new set of dynamics) takes only a fraction of a second and is barely perceptible by the pilot, back-to-back comparison of configurations becomes exceptionally facile and carries the highest degree of validity. The immediacy of the comparison allows the pilot to make fine distinctions that would be otherwise impossible or highly suspect.

2) Sequences of motion signatures can be easily constructed (literally "drawn" with the aid of a graphics tablet) with any strange shape variations that the experimenter may require in attempting to elucidate a kinematic feature of concern to the pilot. This would generally be completely impractical in simulations using the equations-of-motion approach.

3) The time between the sampling of state variables and the issuance of commands to the cockpit drives is very short compared to that for the normal technique involving differential equations (which typically must complete the solution before any responses are produced). Thus, there is effectively no lag introduced by the convolution computation.

4) A serendipitous benefit to aircraft designers of using offline-generated motion signatures in conjunction with the convolution technique has come to light. In the development of a new aircraft, the equations of motion, implemented on a mainframe computer, are typically one of the first analytical tools used. The convolution approach to flight simulation described here is generic in the sense that the software and the hardware are the same regardless of what vehicle is being simulated. Thus, the motion signatures produced early in the design cycle can be used to mount a meaningful, inexpensive, man-in-the-loop simulation well before a normal, expensive, specialized, "all-up" simulation can be implemented.

5) Air turbulence impacts aircraft piloting tasks so profoundly that no simulation for flying qualities studies would be complete without it. Convo provides a capability to study separately the usually comingled effects of control system design and aircraft response to turbulence on flying qualities. In a convolution simulation, air turbulence is treated like any other input. Included in the large array of motion signatures that define vehicle kinematics are some that define responses to step changes in the three components of wind velocity. A suitably constructed turbulence model produces the required quasirandom components of wind velocity, which are then convolved with the appropriate motion signatures to produce the turbulence-induced components of vehicle motion. The advantage that accrues to the flying qualities researcher is that the vehicle response to turbulence, which is normally strongly tied to the control system configuration, can be varied independently of the controls-induced responses. From the pilot's point of view, this is a natural separation. The vehicle's response to turbulence is really a component of the flight task, an outside stimulus to the man-vehicle loop, as opposed to the controls response, which is an intimate part of the loop.

An Application

An examination of adverse lateral control responses at hover currently being performed on the Large Amplitude Research Simulator (LARS, Fig. 1) illustrates the kind of flying qualities problem that benefits from Convo. Despite the success of modern control system technology in manipulating control forces to produce desired steady-state responses, there sometimes exist in VTOL (Vertical Takeoff and Landing) vehicles unwanted transient characteristics that are difficult to remove. For instance, when lateral velocity is to be

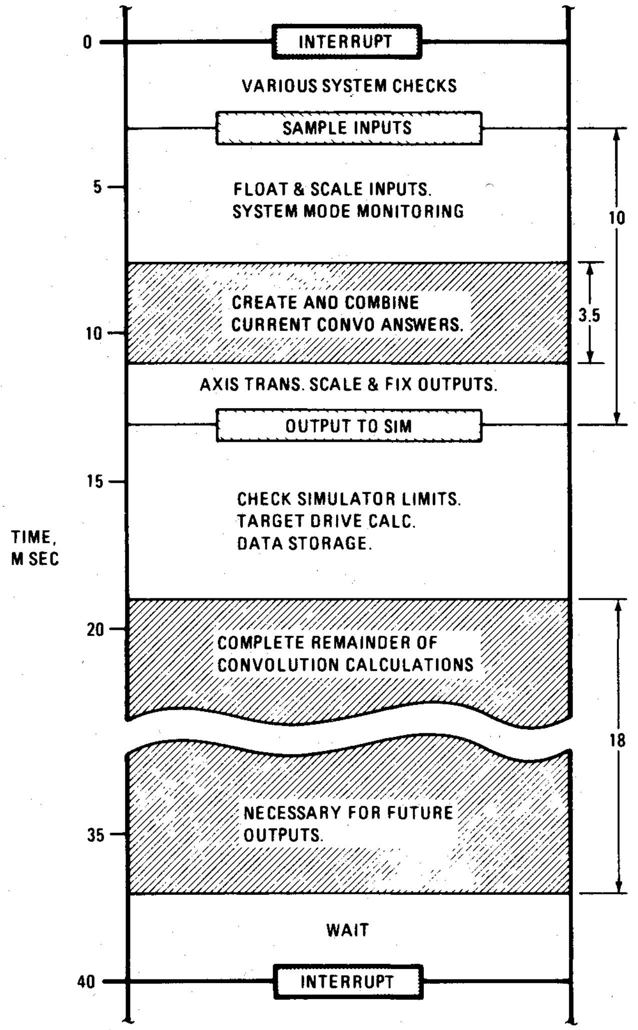

Fig. 3 The program cycle timing.

produced by bank angle but rolling moments are produced by an unbalanced force below the C.G. rather than by a pure couple (see Fig. 2), an initial, wrong-way, lateral velocity precedes the commanded steady-state value. Control system technology can reduce or eliminate this adverse response only to the extent that other control forces are available and maybe then at the expense of delaying or otherwise altering the commanded steady-state value. The designers may therefore be forced to accommodate an anomalous behavior in the interests of practicality, and the anomaly becomes a prime candidate for detailed study. The present research effort examines this problem in the context of lateral tracking of the landing pad of a small ship in a heavy sea. The authors know of no previous research that has addressed this topic in detail.

The simulated vehicle is a current tilt-fan VTOL design (not yet flying) that uses an attitude command system. For the experiments reported here, the 6-DOF simulation was performed with altitude trimmed and three controls (lateral and longitudinal stick and rudder pedals) available to the pilot. Eighteen motion signatures ranging from 2 to 7 s long and comprising a total of 1688 sampled-data values were required to represent vehicle dynamics. Single-speed convolutions were used and the program cycle time was 40 ms, short enough to properly represent the highest frequency in any of the impulse responses and long enough to complete all necessary calculations which, in addition to the required convolutions, included creation of drive signals for the tracking target and various data storage operations. Figure 3 shows the time needed for the major computational events during the 40 msec program cycle. Variations in the initial, wrong-way part of the lateral velocity response to the lateral stick command were the experimental variables. These variations (Fig. 4) are characterized by three features that, it was thought, pilots would perceive independently: onset, magnitude, and duration of adverse velocity. The somewhat stylized shapes shown result from an effort to make the experimental variables physically independent and their variations free of irregular perturbations of jerk (time rate-of-change of acceleration). Thus, the step responses in Figs. 4b and 4c were constructed using a single value of jerk. It should be noted that the simple, independent variation of response parameters shown here would be nearly impossible in conventional simulations, requiring a very high-order set of differential equations and a complicated variation of many coefficients to change from one configuration to another. The convolution technique permits the manipulation of these response characteristics without regard for the physical systems needed to produce them. It must be emphasized that this research approach is purposefully not inhibited by any consideration of physical realizability.** It attempts to define the pure fundamentals of pilot motion preference, unmindful of how the motions might be generated.

Two pilots flew the 13 configurations in over 150 trials. Each trial comprised lateral tracking of a target for 82 s, then the rating of the configuration on the Cooper-Harper pilot-opinion scale (Ref. 2). The motion of the target was made similar to the sway of the landing pad of a destroyer in a heavy sea (sea state 5, Ref. 3). The motion sequence was repeated every 82 s in order to provide the same kinematic milieu for all trials. To minimize memorization by the pilot, a sequence exhibiting no outstanding characteristic (such as a particularly large peak or long period of inactivity) that could be readily remembered was selected, and trials were begun at random points in the sequence. Finally, to minimize the effects of learning, the experiments began only after a lengthy period of familiarization.

The essence of the experimental results can be seen in the plots of pilot rating averages shown in Fig. 5. In this study, pilot GH was by far the more experienced rater, having participated in some exploratory work and a lengthy prior experiment of different design but involving the same dynamic configurations. His responses are therefore considered to be the more reliable and are given greater weight in the following discussion.

In Fig. 5, the configuration numbers along the abscissa refer to the lateral velocity responses of Fig. 4. It is clear that the "duration" parameter (Fig. 5b) is strongly negative, as would be expected. The "onset" parameter (Fig. 5a) appears to be weak, with one pilot reporting it slightly negative and the other slightly positive. The effects of the "magnitude" parameter (Fig. 5a) seem paradoxical. Small amounts actually improve flying qualities, up to a point, but larger amounts ultimately make them worse. This curious result makes sense if we characterize the experimental configurations by their acceleration response time histories (Fig. 6) instead of by their velocity response time histories (Fig. 4). What appear as variations in the magnitude of the adverse velocity response of Fig. 4c can be seen in Fig. 6c to be variations in two parameters: a kind of pulse at the beginning of the response, and a period of time for the acceleration to become positive. The latter is identical to the redefined duration parameter shown in Fig. 6b. Note that, as the magnitude of the adverse velocity response in

**In fact, the dynamic responses represented by the motion signatures used are all "hardware-realizable" in the sense that they dictate real simulator motion (hence may contain no anomalies such as velocity discontinuities). There is no fundamental reason why they could not be similarly realizable in flight hardware if appropriate sensors, actuators, and force generators could be adopted. The practicality of this would depend upon the context.

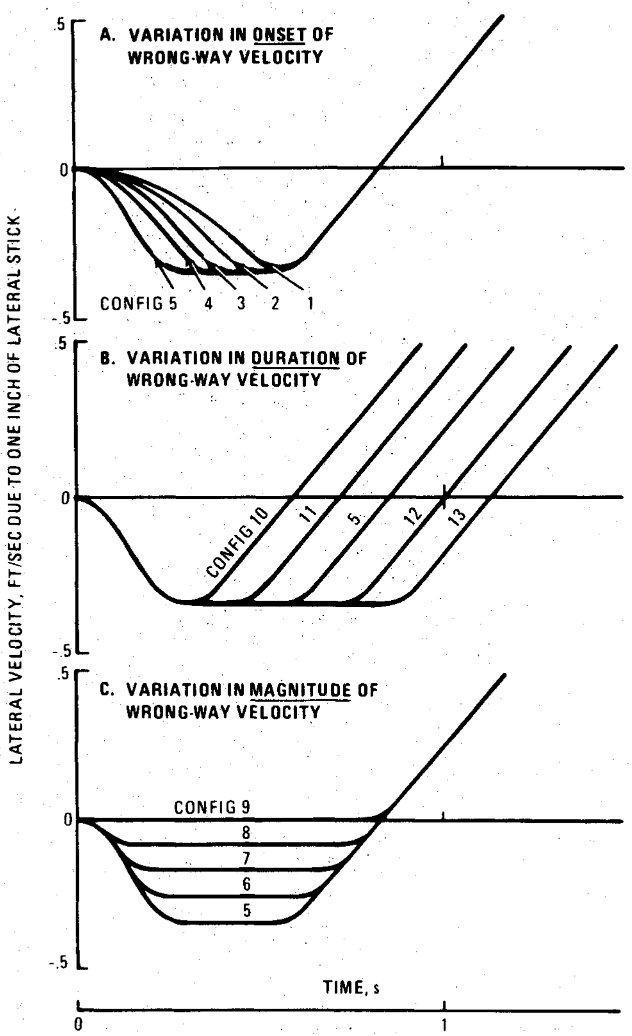

Fig. 4 Variations of the velocity response used in the lateral tracking experiment.

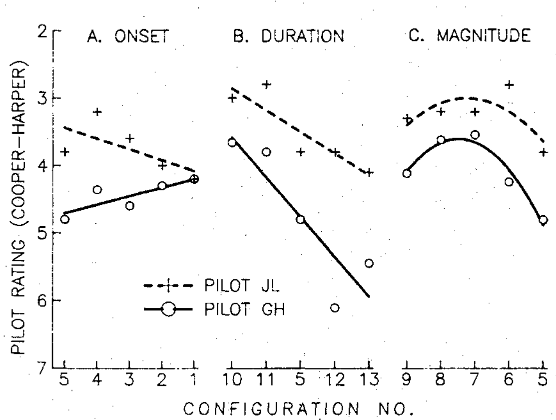

Fig. 5 Pilot opinion results.

Fig. 4c increases, the two parameters of Fig. 6c vary inversely, i.e., the adverse acceleration pulse gets larger and the duration gets smaller. This inverse relationship explains the pilot opinion variations shown in Fig. 5c. While the adverse acceleration pulse remains small, the improved flying qualities for a small increase stem from the reduced time-to-positive-acceleration. Larger magnitudes, however, cause the adverse acceleration pulses to overwhelm the accompanying beneficial reductions in duration, finally causing the flying qualities to deteriorate. The different shapes of the initial pulse variations in Fig. 6a that produce the velocity onset variations of Fig. 4a

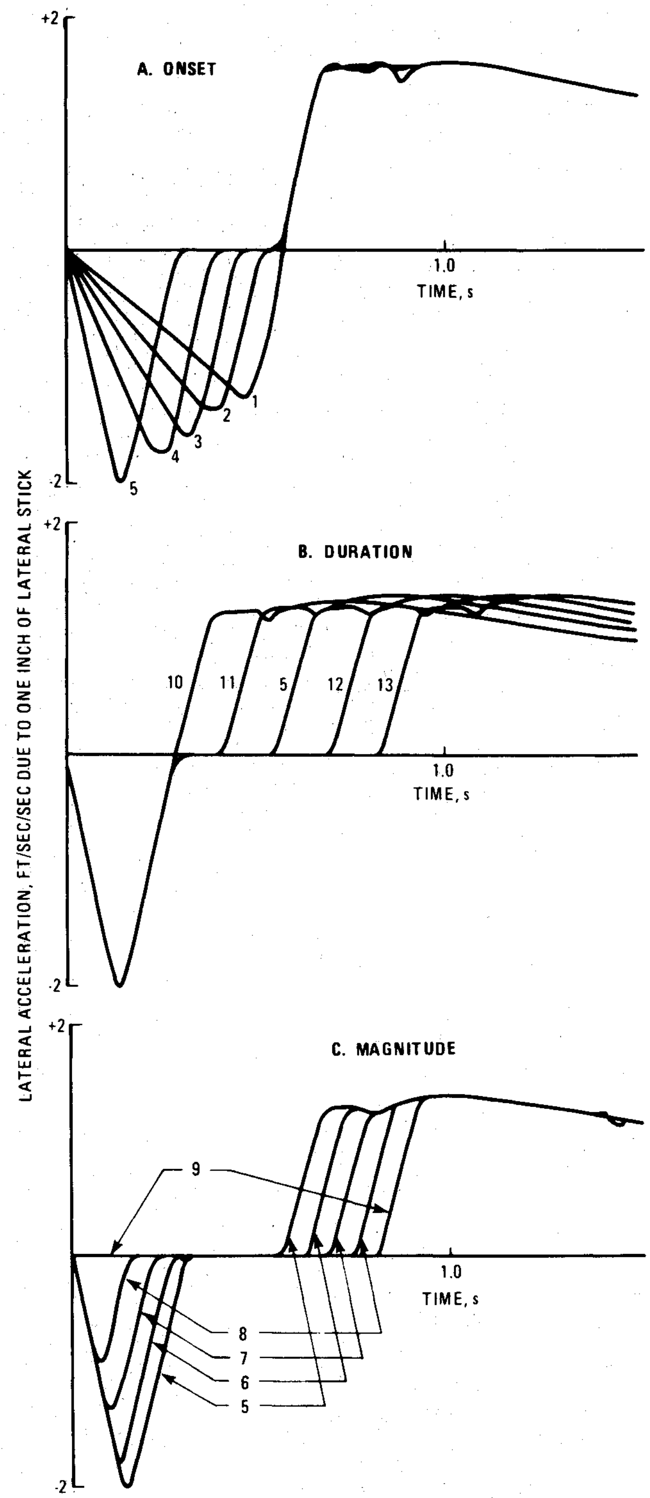

Fig. 6 Acceleration responses corresponding to Fig. 4.

all integrate to the same maximum adverse velocity, and, as noted, cause little variation in flying qualities (Fig. 5a). We infer, therefore, that the deterioration of flying qualities caused by the larger initial acceleration pulse of the magnitude variation results solely from the increased wrong-way velocity it produces. We conclude that within the context of these experiments there are two independent features of the wrong-way lateral response that affect flying qualities: maximum wrong-way velocity and the time for the acceleration to start in the commanded direction. The effect of the former was confounded in the present experiment, and it is therefore difficult to assess the relative strengths of the two parameters. This work is continuing, and further study of these and other effects is planned.

Conclusions

A linear-system modeling scheme based on the mathematical concepts of convolution and superposition has been developed for simulator flying qualities studies. For the important class of fixed trim-point simulations it has a number of advantages over conventional modeling schemes:

1) It is very easy to "draw" arbitrary motion signatures, and therefore to independently vary the time domain characteristics in search of response details that are meaningful to the pilot.

2) Accurate, real-time representation of high-order, 6-DOF vehicle dynamics at rep rates suitable for piloted simulator drives can be done inexpensively on a minicomputer.

3) The algorithm used has an inherently short input-to-output delay time (less than 2.0 msec on a modern minicomputer).

4) Changing configurations is essentially instantaneous, allowing pilots to make accurate back-to-back comparisons of different dynamic characteristics.

5) The dynamics of new vehicles can be easily created from step responses produced by standard non-real-time stability analysis programs. Thus, relatively cheap, high-quality, man-in-the-loop simulations are possible early in the design cycle.

6) Two important features of control system performance, gust suppression and response shape, can be separated and studied independently. An extensive set of operational software has been developed to mechanize and support the convolution simulation modeling scheme, and it is now being applied to the study of hover control problems.

A preliminary investigation of the effect on pilot opinion of initial wrong-way velocity responses to lateral commands in a tilt-fan-type VTOL aircraft at hover has been conducted in the LARS, using the convolution technique for dynamics computation. Three specific shape features of the wrong-way response have been examined in detail. One had little effect on pilot opinion, one had a strong monotonic effect as expected, and one had a significant but ambiguous effect. A hypothesis explaining these results, to be tested in future experiments, has been propounded.

Acknowledgments

The authors wish to express their appreciation to Grumman test pilots Gerald Haser and Julius Longshore, who flew the LARS and provided the expert opinion on which the research results are based.

References

[1]Gardner, M.F. and Barnes, J.L., "Transients in Linear Systems," Vol. 1, John Wiley & Sons, Inc., New York, 1942, pp. 233-235.

[2]Cooper, G.E. and Harper, R.P. Jr., "The Use of Pilot Rating in the Evaluation of Aircraft Handling Qualities," NASA TN D-5153, April 1969.

[3]Brown, R.G. and Camaratta, F.A., "NAVAIRENGCEN Ship Motion Computer Program—Theory, Documentation and Users Guide," NAEC MISC-903-8.