HONEYWELL'S HISTORY AND PHILOSOPHY IN THE ADAPTIVE CONTROL FIELD

(Including a description of flight tests of the first Adaptive Flight Control System)

O. Hugo Schuck
Director of Research
Aeronautical Division
Minneapolis-Honeywell Regulator Company

HISTORY

It's rather hard to tell when we at Honeywell first became interested in adaptive controls. Perhaps we didn't use the right words, but certainly the need for adaptive controls has been recognized as being with us for a long time. This has been especially true in two of the areas in which we are vitally interested - flight control and process control.

By now in this symposium it isn't necessary to describe further the need for adaptive flight controls. In reprospect, it seems that we first consciously articulated the need in connection with our early work in automatic approach and landing. There, as you know, the ILS beam has a built-in convergence, leading to a system gain that varies inversely as the distance from the aircraft to the far end of the runway. This gain change plagued us considerably back in 1945 and 1946. If we had the gain set high enough for decent beam entry and following fifteen miles out, we went into a divergent oscillation as we neared the field. Even putting in a manual gain change at the outer marker wasn't enough, but we did manage to work out a useful compromise.

The obvious remedy was a continuous gain reduction as we neared the runway, a technique that we did use successfully in a relative humidity computer. So we were delighted with the advent of the Distance Measuring Equipment (the DME), and assumed that when SC31 recommended it for the Common System our troubles were over. But the DME was not adopted quickly, and anyway, the operators were somewhat less than enthusiastic about getting more electronic gear involved in the critical landing operation. So we sharpened our servomechanism theory pencils and learned to live with the convergence by pushing loci around on Nichols charts.

However, one aspect of this experience stayed with us and became, in fact, an integral part of our flight control system design philosophy. This was the idea of scheduling control system gains as pre-determined functions of measurable parameters. Now, of course, we're trying to get away from such scheduling.

Another trouble plagued us in the automatic landing problem. The beam converged, not with respect to a straight center line, but to a line that had bends, wiggles and other aberrations with respect to the straight line it was supposed to be. We learned to call these by the scientific name of " noise" - which is a good name because its definition is " anything you wish you didn't have" . We found that the ILS beams at different airports around the country had different degrees of noisiness. The CAA managed to keep the one at Indianapolis pretty clean, but the one at Minneapolis - where we did a lot of test flying - had some pretty bad bends. One, in fact, gave the uninitiated observer the idea that the aircraft was going to land in the third hangar.

Fortunately, the ILS beams all tended to straighten out as the touchdown point was approached. It was in this critical region that we wanted good control. On the other hand, we didn't want to shake up the passengers or stress the airplane unduly. So we saw the problem as one of following clean beams pretty closely, while smoothing out to a considerable extent the noise in those that had a lot of bends. In other words, we posed the objective of doing the best control job possible, considering the signal to noise ratio of the input. This implied a change in system filtering and gain as a function of the input noise.

One approach to this problem was through the use of non-linear techniques. The Air Force was developing an interest in non-linear mechanics, and gave us support for a research project to advance the non-linear mechanics art in its application to automatic controls. One of the specified areas of application was to beam following. After a considerable amount of analytical work, during which we explored many of the techniques being discussed here, we settled for a system that limited the second derivative of the localizer signal as a function of the beam noise. An alternate mode also lowered gain in the presence of high noise levels.

This system was built, it flew the airplane as it was supposed to, but it was unduly complicated. So it never left the Research Department.

In the meantime, we had been in contact with the Research Department in our Brown Instruments Division, and had found that they had their own brand of difficulty in the chemical process control field. Here the problem was not so much one of input noise, but one of changing conditions. They had harrowing stories of working all day in a chemical plant, adjusting gains and limits until the control system was working nicely and stably. Then the five o'clock whistle would blow, a few pressures and temperatures would change with the change in shift load, and the control would go divergent. So they were looking for an approach that would permit greater tolerance to the conditions of operation.

At this time we began to think more clearly in terms we now use, of self-adaptive control. We began to visualize a control system that would produce a desired result, even though there was input noise, and even though the conditions in which it was working changed greatly. After we had been working on this for a couple of years the Air Force found our approach of interest, and gave us support for a study in which adaptive control was one of three approaches to achieving better flight control of high-performance airplanes.

While we did try to keep our approach unbiased, our previous experience tended to point us away from statistical transfer function techniques. We had already become familiar with the discontinuous feedback concept pioneered by Flugge-Lotz of Stanford University. We looked further into its implications, and investigated its applicability to a typical flight control problem. Using an analog computer, we tried various ways of mode switching and evolved a concept combining a model and an intelligently switched bi-stable element. With the results looking quite promising we decided to concentrate on a more definitive check-out of this concept through actual flight test on an F-94C. Equipment was built, thoroughly tested in simulation, and installed. The flight test results were most gratifying. They confirmed the simulation results and led to several engineering programs.

FLIGHT TESTS

Let us look at this adaptive flight control system that was proven in the flight tests on the F-94C. Conceptually it is simple, deceptively so. As shown in Fig. 1, the input is applied to a model whose dynamic performance is what we wish the dynamic performance of the aircraft to be. The actual response of the aircraft is compared with the response of the model, and the difference is used as the input to the servo. If the gain of the servo is very high, the response of the aircraft will be identical to that of the model, no matter what the elevator effectiveness, so long as it is finite and has the right direction. Thus changes in altitude and airspeed have no effect on the response. By building the model to give ideal response, the aircraft is given ideal response characteristics.

Design of the model is fairly simple, provided one knows what kind of response is wanted. Regular network theory can be used. The big problem comes in connection with the need to make the gain of the servo very high. This is essential to making the inner loop sufficiently tight that the response of the aircraft is essentially that of the model. An ordinary linear servo system will not do. It simply cannot be given sufficiently high gain and still be stable. So we go in for non-linearity, the most extreme form of non-linearity, in fact - the bang-bang type. Full available power is applied one way, or the other, depending on the direction of a switching order.

Now it is well known that a simple bang-bang system is oscillatory. And we don't want an oscillatory aircraft. But the bang-bang principle does have the attractive advantage over any other system of providing full available power to correct even small discrepancies between aircraft response and model response. So we look for ways to tame it down, keeping its high-gain characteristics while reducing its oscillatory activity. This is accomplished through use of a number of techniques in combination, as can better be described in connection with a more detailed diagram.

In Fig. 2 we see the same model, whose output $\dot{\theta}_M$ is the desired pitch rate. Feedback is provided by a pitch rate gyro. It measures the actual pitch rate $\dot{\theta}$ of the aircraft, which is represented as a second-order system. Experience has shown that this representation is adequate to describe the short-period motion of the conventional rigid aircraft to elevator deflection. The damping ratio ζ_a, the natural frequency ω_a, the time constant T_a, and the elevator effectiveness M_{δ_e} are all known functions of the aircraft stability derivatives.

The input to the aircraft is elevator deflection δ_e, which is produced by a conventional servo and actuator shown in the block to the left of the aircraft block. It has the usual integration and second-order dynamics of such systems. In the case of the Lockheed F-94C aircraft, the natural frequency is 37 radians per second and the damping ratio is 0.7. The proportional plus integral term in the numerator results from the use of a high-pass network in the feedback loop internal to the servo-actuator system. Cancellation of this numerator term is the primary purpose of the lead-lag filter shown next to the left. Some small lead is introduced to compensate for the normal lost motion of slop in the control gearing.

Next, to the left is the limiter, operating on the output of the relay. It sets the magnitude of the relay's output voltage. Some adjustment of this magnitude seemed desirable, and this is the purpose of the gain changer in the upper block. Its operation is described by the equations in the lower right corner. If the system error is large-larger than B in absolute magnitude, full output voltage is obtained from the relay. After the system error has been reduced below B the output of the relay is decreased exponentially with time in accordance with the second equation. Fig. 3 shows graphically this decrease in available input to the filter.

Returning to Fig. 2, we see that instead of feeding the error signal directly to the relay, a somewhat modified input is provided by the lead-lag network in the switching logic block. Ideally the denominator time constant is zero; the numerator constant is about 0.2 seconds. A further modification of the input to the relay is made by the introduction of a high-frequency sinusoidal dither signal. Its frequency of 2000 cycles per second is so high that it does not appear in the output motion. An averaging process takes place, so that the output of the relay is linearized for very small signals. Use of a sinusoidal dither signal gives the arc-sin characteristic shown in Fig. 4. Obviously a mechanical relay cannot follow a 2000 cycle per second dither input. An electronic relay can, and such was used in the F-94C flight test equipment.

Returning again to Fig. 2, we see that the pitch-rate inner loop consists of the switching logic's lead network, the relay with dither, a limiter controllable from the error magnitude by the gain changer, a lead-lag filter, the servo and actuator, and the aircraft, with feedback provided by a pitch-rate gyro. The dynamics of the gyro are included for completeness, but they can be neglected. The loop gain is very high, and consequently the actual pitch rate can be maintained acceptably close to the output of the model $\dot{\theta}_M$.

For the F-94C flight test program a model was used having a natural frequency of 3 radians per second and a damping ratio of 0.7. These values have been established by the NACA and Cornell Aeronautical Laboratories as acceptable for manned aircraft of the type used.

To obtain pitch attitude control the switch shown at the left in Fig. 2 is closed. Since the inner loop transfer function is unity, the attitude control system can be described by a third-order transfer function: second order from the model and an integration from pitch rate to pitch attitude.

This, then, is the nature of the system we're talking about. Before going on to describe its mechanization for the flight test program, I want to admit - in fact, emphasize - that it did not reach this form by a process of pure, abstract cerebration. Each portion is there because it was found necessary in the course of a long simulation program, and the various constants were worked out before the mechanization was undertaken. It was the apparently successful performance obtained in simulation that led to the decision to undertake flight test verification. Furthermore, since this was a research project and the objective of the flight tests was primarily to provide a feasibility check of the adaptive flight control technique that had been evolved, minimum modifications were made to the aircraft and its equipment. Most of the equipment used was from the old Honeywell E-10 Autopilot development program of several years before.

Fig. 5 shows the amplifier that was built for the flight tests. It incorporated the various circuits that were unique to the adaptive system. All of this equipment was given careful hanger testing, with analog simulation of aircraft flight dynamics, before initiating the airborne tests. It is interesting to note that these tests showed the previous simulations of the hardware to be quite valid.

Fig. 6 shows the flight test engineer's test panel. In designing the flight test program we decided to go beyond just verifying the simulation, and try to get some comparisons with the operation of the standard E-10 Autopilot - both quantitative and pilot reaction. Available to the flight test engineer, accordingly, there are on the test panel various switches and controls to allow a considerable variety of configurations to be set up. These included changes in some of the significant parameters of the adaptive system, such as the dither amplitude, filter characteristics, and gain changer operation. Provision was also made for the introduction of standardized pitch rate and pitch attitude commands.

Let us look at a few of the results of the flight test program, as carried out at Minneapolis in April of 1958. The first series covered operation without the gain changer, and with limiter output set to provide a maximum elevator deflection rate of 4.6 degrees per second. In Fig. 7 we see a typical recording. At the top is the command input, third down are the model response and the aircraft response. We find it convenient to record in opposite sense those quantities that are to be directly compared.

The performance shown is typical of the results obtained throughout most of the flight envelope. In the steady-state condition the traces are acceptably smooth, that is, the residual motion is quite small - on the order of what is obtained with the usual linear control system. The model-following capability is excellent, even at relatively high control frequencies. In other words, the aircraft's response does agree with the model's response. This is true for small-scale maneuvers; it is also true for violent maneuvers.

Fig. 8 shows flight test results of one of a series of violent maneuvers. This is a complete loop. During the short time taken to carry out the maneuver the altitude changed between 10,000 and 20,000 feet and the Mach number between 0.4 and 0.8. Note that the pitch rate is quite constant throughout the loop, as specified by the output of the model. Even cutting out the afterburner at the top of the loop produced only a small bump in the pitch rate record. We see here a very significant feature of an adaptive system of this type, that it can adapt quickly to rapid changes in operating conditions. There is no waiting for a scheduling device or computer to " catch up".

In Fig. 2 a switch was shown that provided pitch attitude control. This mode of operation was also checked in the flight testing program. The pitch attitude feedback gain was set by use of third-order charts to give an overshoot of 30%. As it turned out, the pilots' thought this overshoot a bit excessive, and we wished we had designed for a lower value. However, that was what we specified, and that was what we got, as the recordings in Fig. 9 show. At least, we got it at 0.6 Mach. At 0.4 Mach the initial overshoot was still about 30%, but the transient took a somewhat longer time to damp out.

This is not surprising, since the performance of the pitch-rate inner loop was expected to deteriorate somewhat at the lower dynamic pressures. The object of all this work was, of course, to have the aircraft response always agreeing with the model response, no matter what the flight conditions were. A large part of the flight test program was, accordingly, devoted to determining just how the agreement did depend on flight condition. For small variations the measured discrepancies were found to be quite small. However, at the extreme ranges of the flight envelope, certain deteriorations were noted.

Fig. 10 shows results obtained at landing speeds. Under such low dynamic pressure conditions the response is that of a system with lower damping than that of the model. When the input is cycled rapidly there are both amplitude changes and phase shifts.

Fig. 11 shows results obtained at high speeds. With high dynamic pressures the model-following capability is quite good, but the amount of residual motion in the steady state is only marginally acceptable.

Fig. 12 shows a summary of the first tests covering the flight envelope. Throughout most of the operating range excellent control and stability were obtained. However, at low dynamic pressure there is a small following error, and at high dynamic pressure there is an objectionable amount of residual motion.

The results described so far all show operation without the gain changer, that is, the output of the relay was always limited at 4.6 degrees per second elevator rate. This was the compromise value that had been worked out in the simulation studies. Better results at low speed would be expected if the limiting rate were higher. Flight test results confirming this are shown in Fig. 13. At low dynamic pressures a limiting rate of 9.2 degrees per second gives noticeably more precise following of the model. Likewise, at high dynamic pressures a lower limiting rate of 2.3 degrees per second reduces the steady-state residual motion to an entirely acceptable degree.

These flight test results confirm the simulation results, and show that a four-to-one change in limiting elevator rate will provide fully acceptable performance of the adaptive flight control system over the entire envelope of the F-94C. It was for this reason that the gain changer previously mentioned was added in the simulation studies. With properly chosen constants it produced the desired results in simulation. Unfortunately, the aircraft's availability schedule did not permit adequate investigation of gain-changer performance in the flight-test program reported here. Subsequent tests at Wright Field by the Flight Control Laboratory showed that the gain changer did, in fact, do what it was supposed to do in cleaning up performance at the edges of the flight envelope. A slightly different way of accomplishing this function is used in the later equipment described in a companion paper.

PHILOSOPHY

What I've been giving is a highly summarized review of our activities up to and through the F-94C flight tests. I haven't tried to mention all the disappointments and blind alleys. Some of these are described in our various reports, others we would just as soon quietly forget. I do think, however, that it's worth pointing out the research role, since it is over five years since we consciously set our hands to learning how to make a self-adaptive automatic control system. Since that time there has been maintained, on company funds, a consistent effort in this direction. Furthermore, it's worth pointing out the value of trying to help the military solve its problems, in that the support thus made available facilitated getting an early physical proof of the research theory.

How, then, do we view the future? We see adaptive controls as typifying the nature of the future of automatic control. From a practical point of view we see adaptive control advantageous primarily as a means of achieving a desired level of equipment and performance with superior reliability. After reliability and reliability and reliability, there come weight, size, and cost. These are all considerations that apply strongly to airborne equipment, but they're also of significance in ship-borne and ground-borne equipment, including process control and environment control.

From a theoretical point of view, adaptive controls form a particular class in the broad field of non-linear controls. We feel we must continue to do research work in this field, and particularly in the adaptive class. Our current research program is so aligned, and we are hoping to expand it. There are two prongs to our research effort. One is to extend our theoretical understanding of the present approach in regard to higher order systems. The other is to continue to investigate approaches other than the one we are now using.

As a sort of side line, we have been intrigued for quite a while with the relationship of our mechanistic controls to those highly refined controls we find in the biological organism. A recent paper by our Charles Johnson, delivered at the Eleventh Annual Conference on Electrical Techniques in Medicine and Biology, discussed this relationship. Johnson showed that, when one lines up the characteristics of the human being as a controller, and those of the adaptive servomechanism, that has been described here, the parallelism is inescapable.

It is easy to say that the automatic control art has much to learn from biology, whose control art has been refined over millions of years. But it is somewhat dangerous, we feel, to take such a statement too seriously. The mechanisms involved are different, radically different. True, the biological controls are non-linear, they are digital, they are adaptive. But they use physical elements that are not appropriate to our control devices. One cannot hope to design a flight control system by dissecting a bumble bee, and there is no known biological organism that is capable of space flight. Where we can learn from biology, and perhaps support the biologists themselves, is in terms of concepts. Conceptually, our adaptive control is parallel to that of the organism, but the mechanisation is quite different. Our aim, then, is to learn conceptually, and mechanize practically.

That statement tends to summarize Honeywell's philosophy in regard to automatic control and particularly adaptive control. We want to learn from every possible source, and we want to apply the knowledge we obtain to produce improved performance in practical situations. And the primary measure of practicality is reliability.

REFERENCES

1. Markusen, David L., Norton, John S., Pomeroy, Orville, P., " The Application of Some Non-Linear Techniques for the Improvement of Aircraft Beam Following" , WADC Technical Report 53-74, Minneapolis-Honeywell Regulator Company, Minneapolis, Minnesota, February 1953.

2. Young, Turrittin, Loud, Hess, Culmer and Putzer, " Non-Linear Control System Applications to Beam Following and Other Flight Control Systems" , WADC Technical Report 53-520, Minneapolis-Honeywell Regulator Company, Minneapolis, Minnesota, December 1953.

3. Stone, C. R. (ed), " A Study to Determine an Automatic Flight Control Configuration to Provide a Stability Augmentation Capability for a High-Performance Supersonic Aircraft" , WADC Technical Report 57-349 or Minneapolis-Honeywell Aero Report 48312, Final, Minneapolis, Minnesota, 30 May 1958.

4. Johnson, Charles W., Adaptive Servomechanisms, Paper No. 68 presented at the 11th Annual Conference on Electrical Techniques in Medicine and Biology, Minneapolis-Honeywell Regulator Company, Minneapolis, Minnesota, 25 November 1956.

ACKNOWLEDGMENTS

Many individuals have contributed to the developments discussed in this paper. Among those whose contributions were most direct must be mentioned: R. C. Alderson, D. L. Markusen, R. W. Bretoi, C. Ling, A. L. Ljungwe, J. T. Van Meter, R. C. McLane, L. T. Prince, R. C. Lee, R. C. Stone, C. W. Johnson, D. L. Mellen.

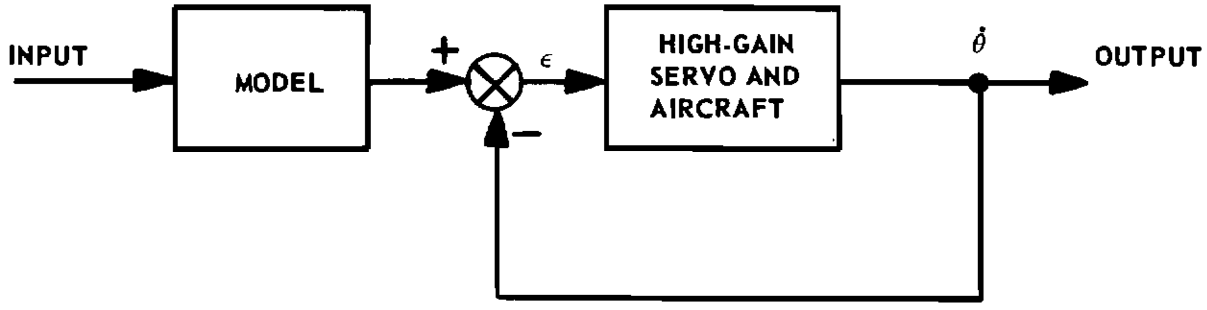

Figure 1. Adaptive Flight Control System (Elevator Channel Only)

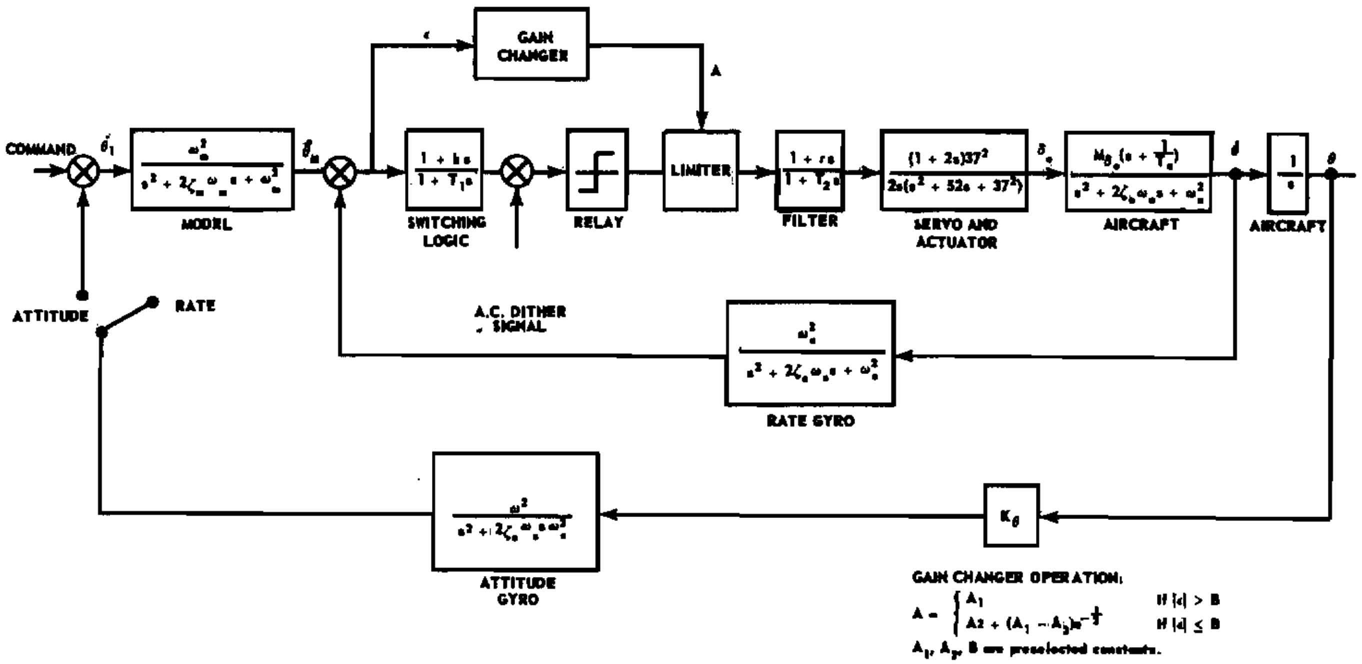

Figure 2. Block Diagram of Adaptive Control System for F-94C

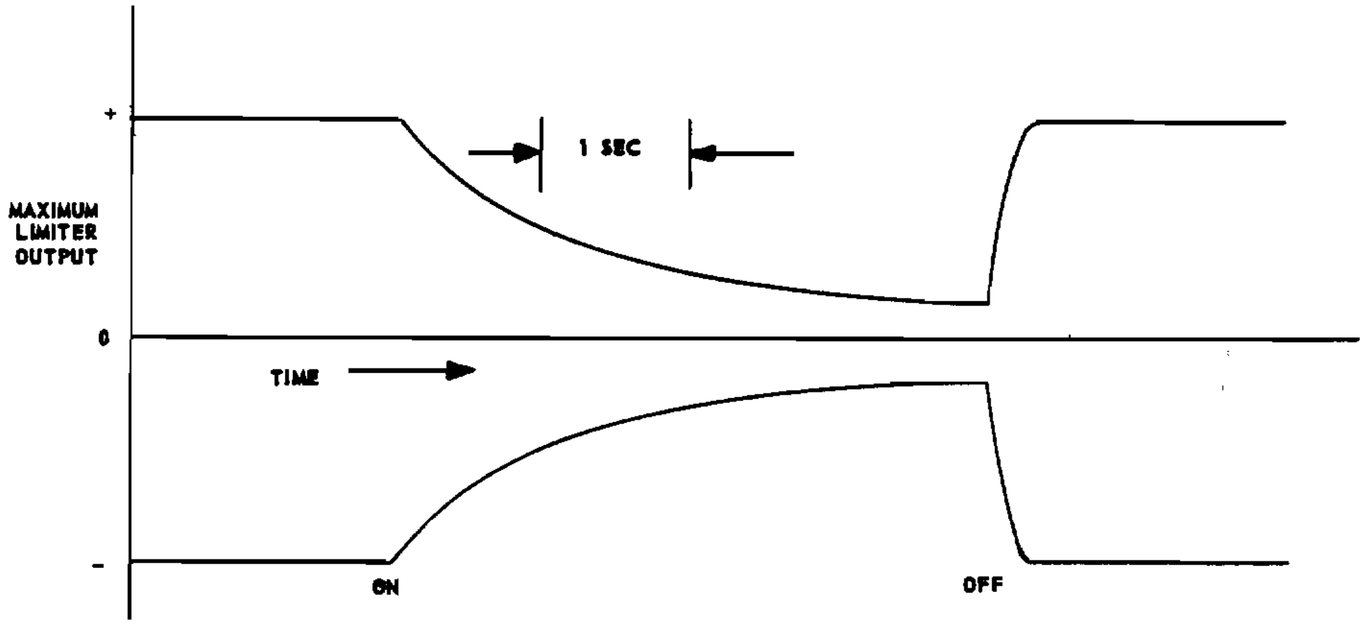

Figure 3. Effect of Gain Changer on Limiter

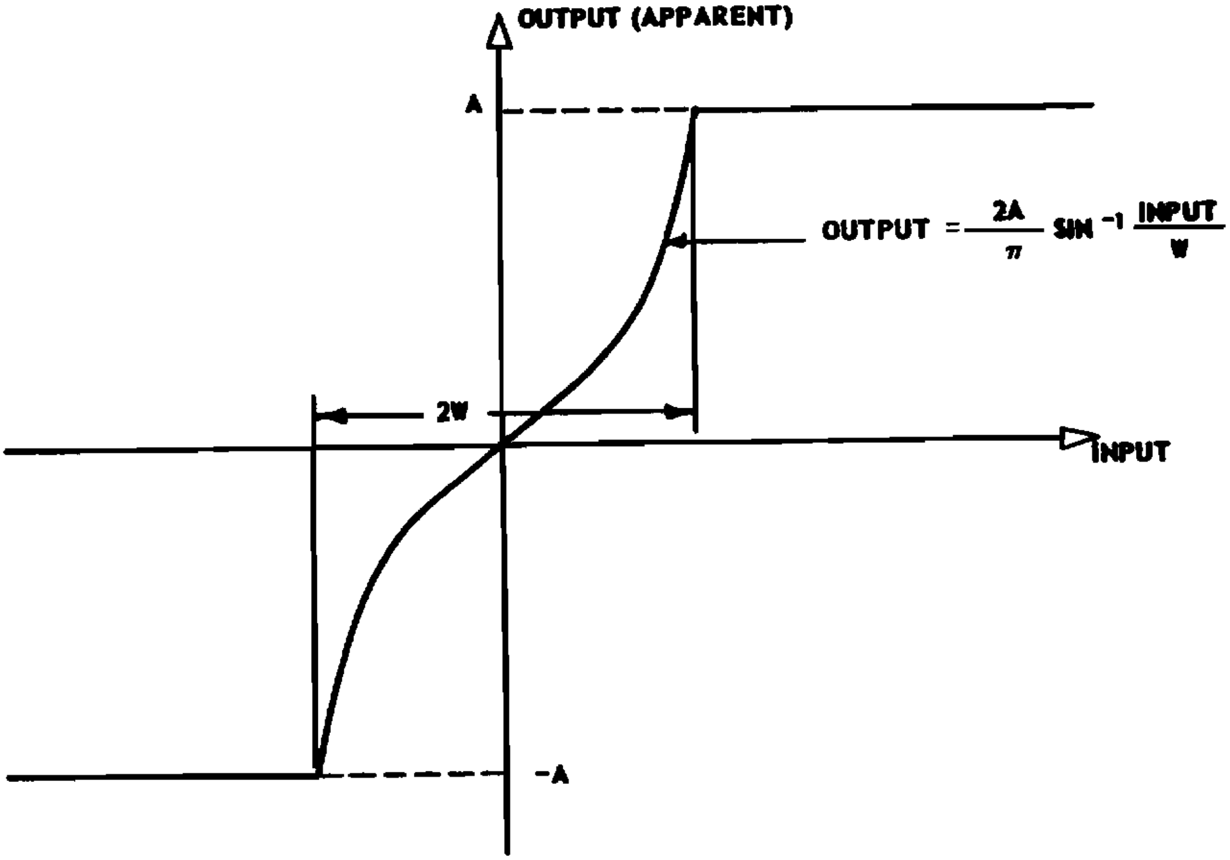

Figure 4. Relay Characteristic - Sine-wave Dither

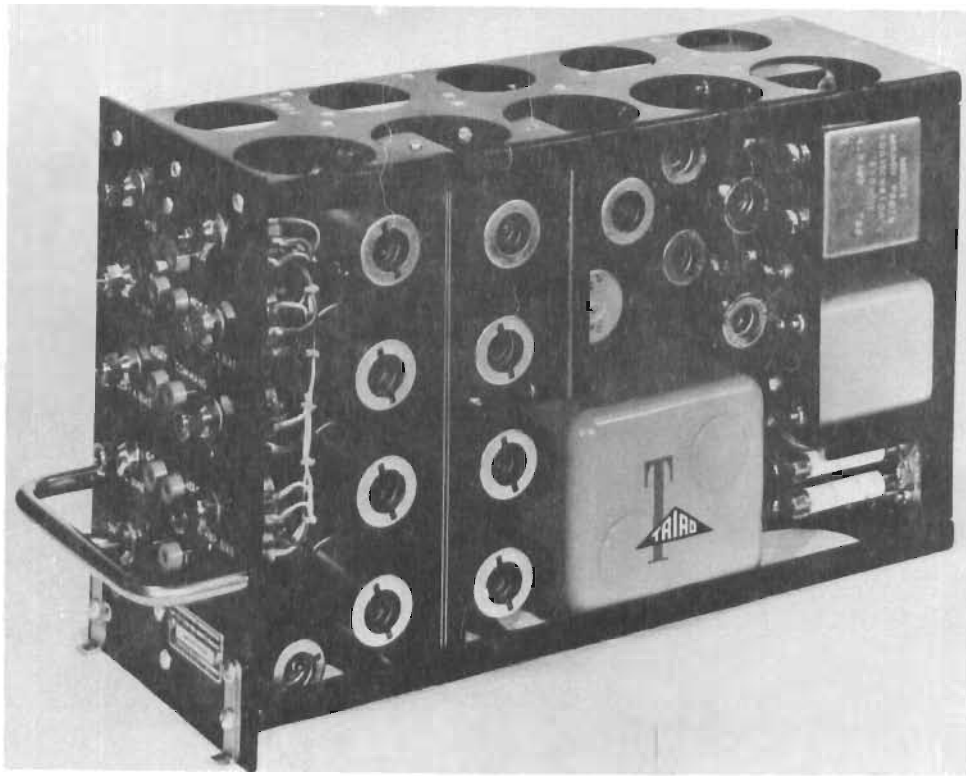

Figure 5. Adaptive Control System Amplifier

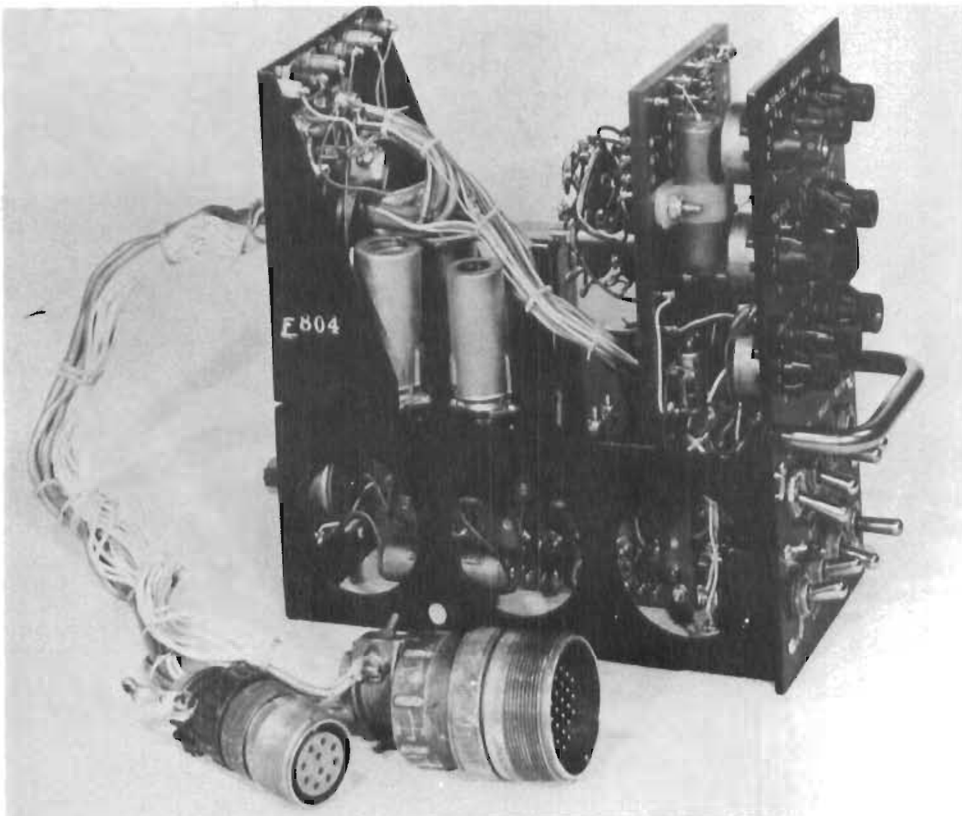

Figure 6. Engineer's Test Panel

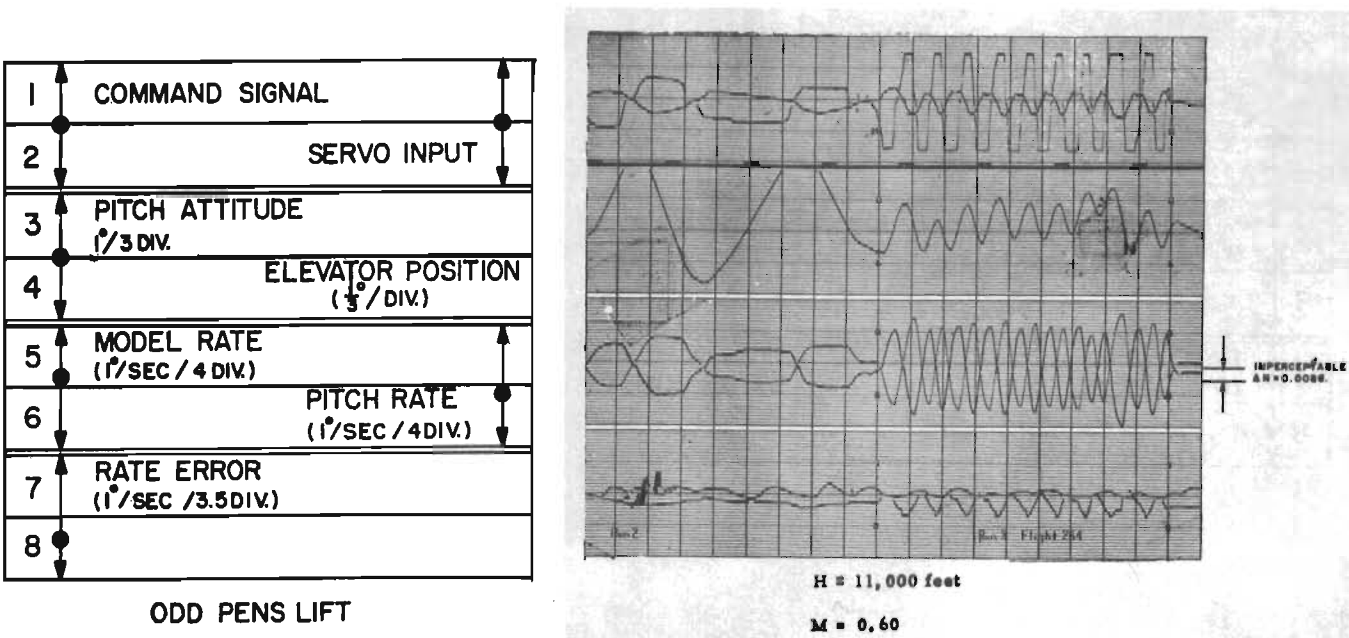

Figure 7 Typical Flight Test Response of F-94C with Adaptive Control System; Pitch Rate Commands

$$\dot{\delta}_{e_{max}} = 4.6 \text{ deg/sec}$$

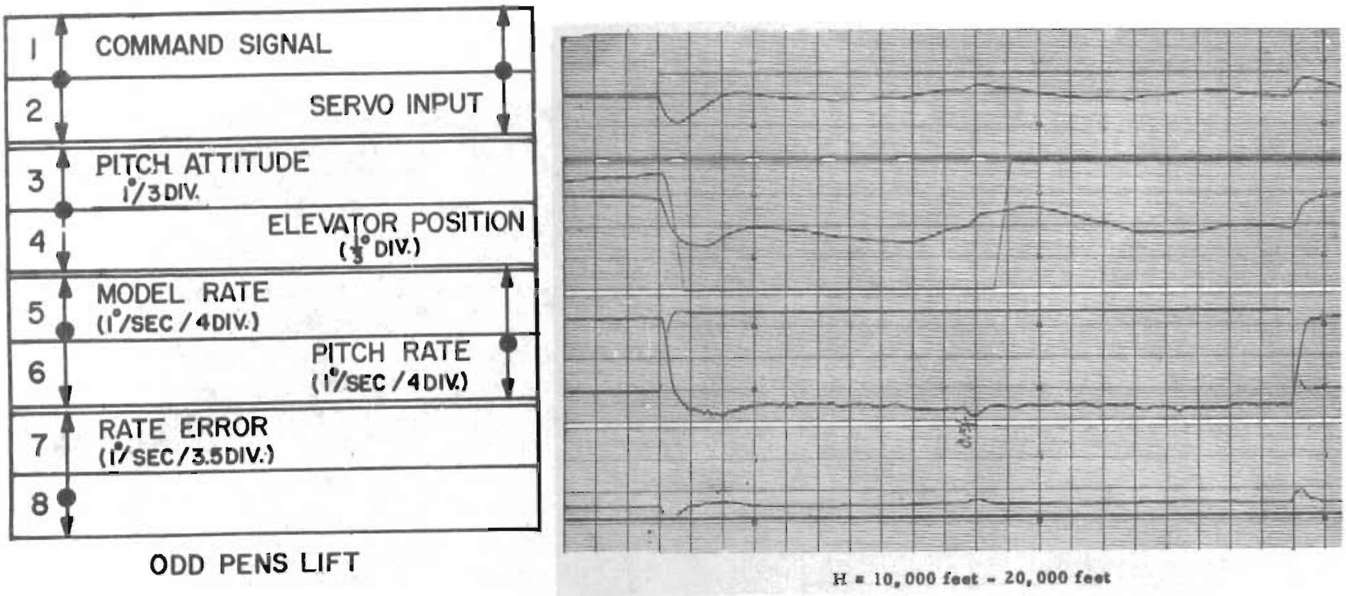

Figure 8 Typical Flight Test Loop

$$\dot{\delta}_{e_{max}} = 4.6 \text{ deg/sec}$$

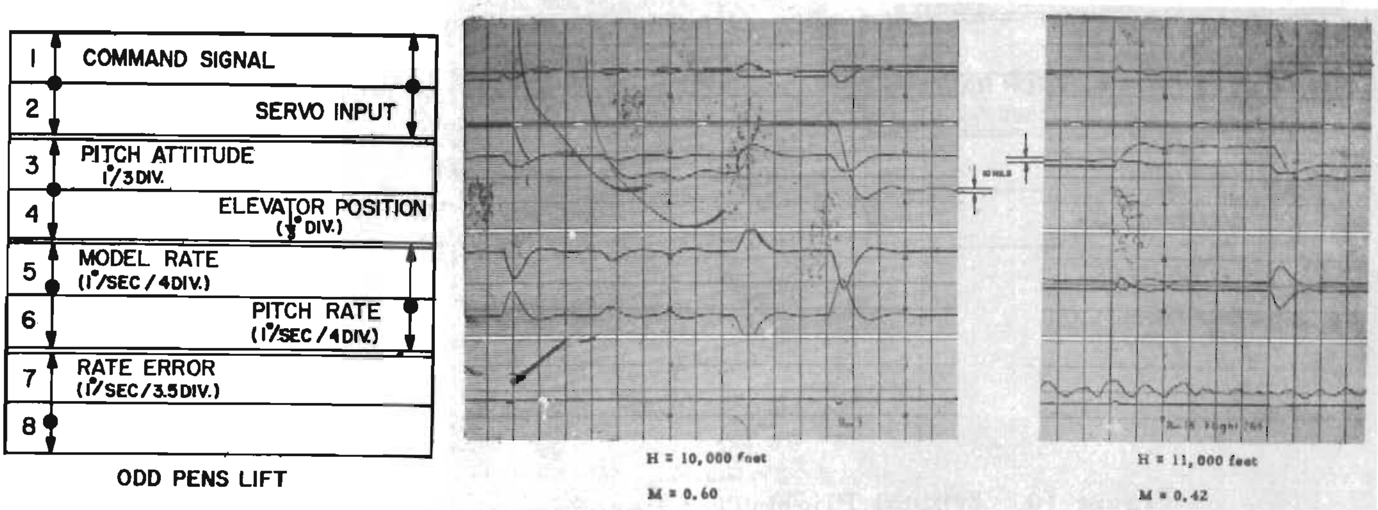

Figure 9 Typical Flight Test; Pitch Attitude Response

$\dot{\delta}_{e_{max}}$ = 4.6 deg/sec

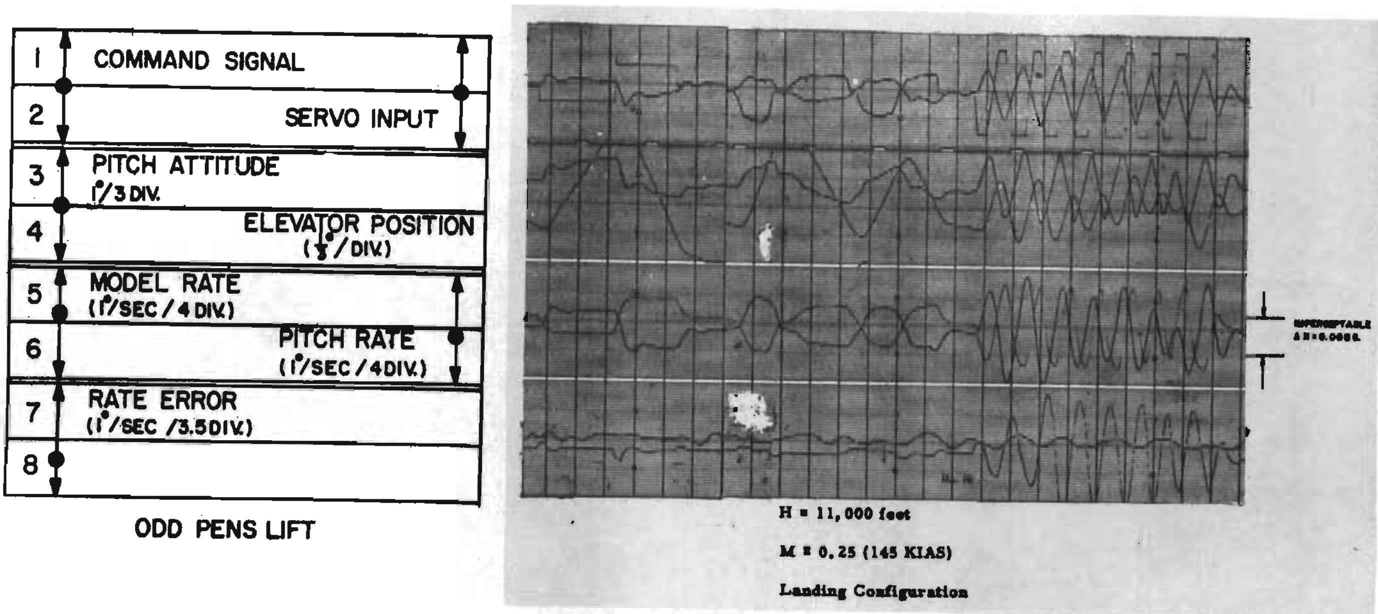

Figure 10 Typical Flight Test Response of F-94C with Adaptive Control System; Pitch Rate Commands

$$\dot{\delta}_{e_{max}} = 4.6 \text{ deg/sec}$$

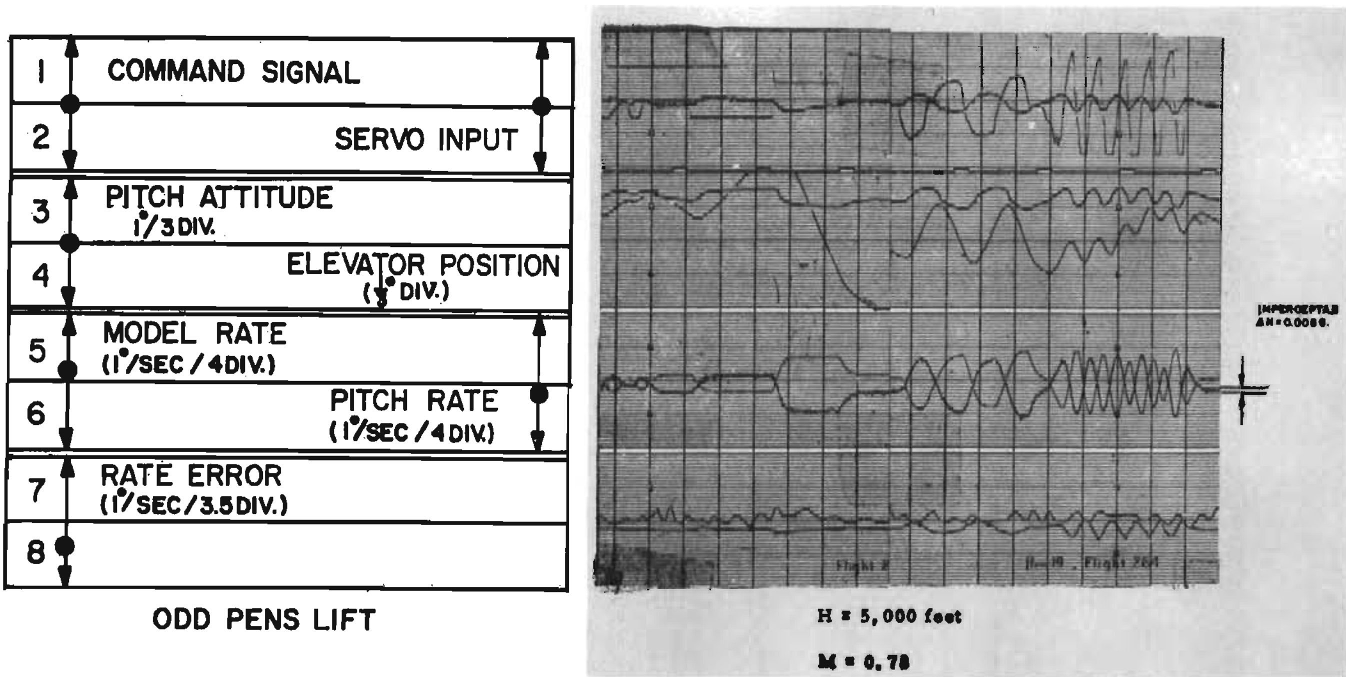

Figure 11 Typical Flight Test Response of F-94C with Adaptive Control System; Pitch Rate Commands

$$\dot{\delta}_{e_{max}} = 4.6 \text{ deg/sec}$$

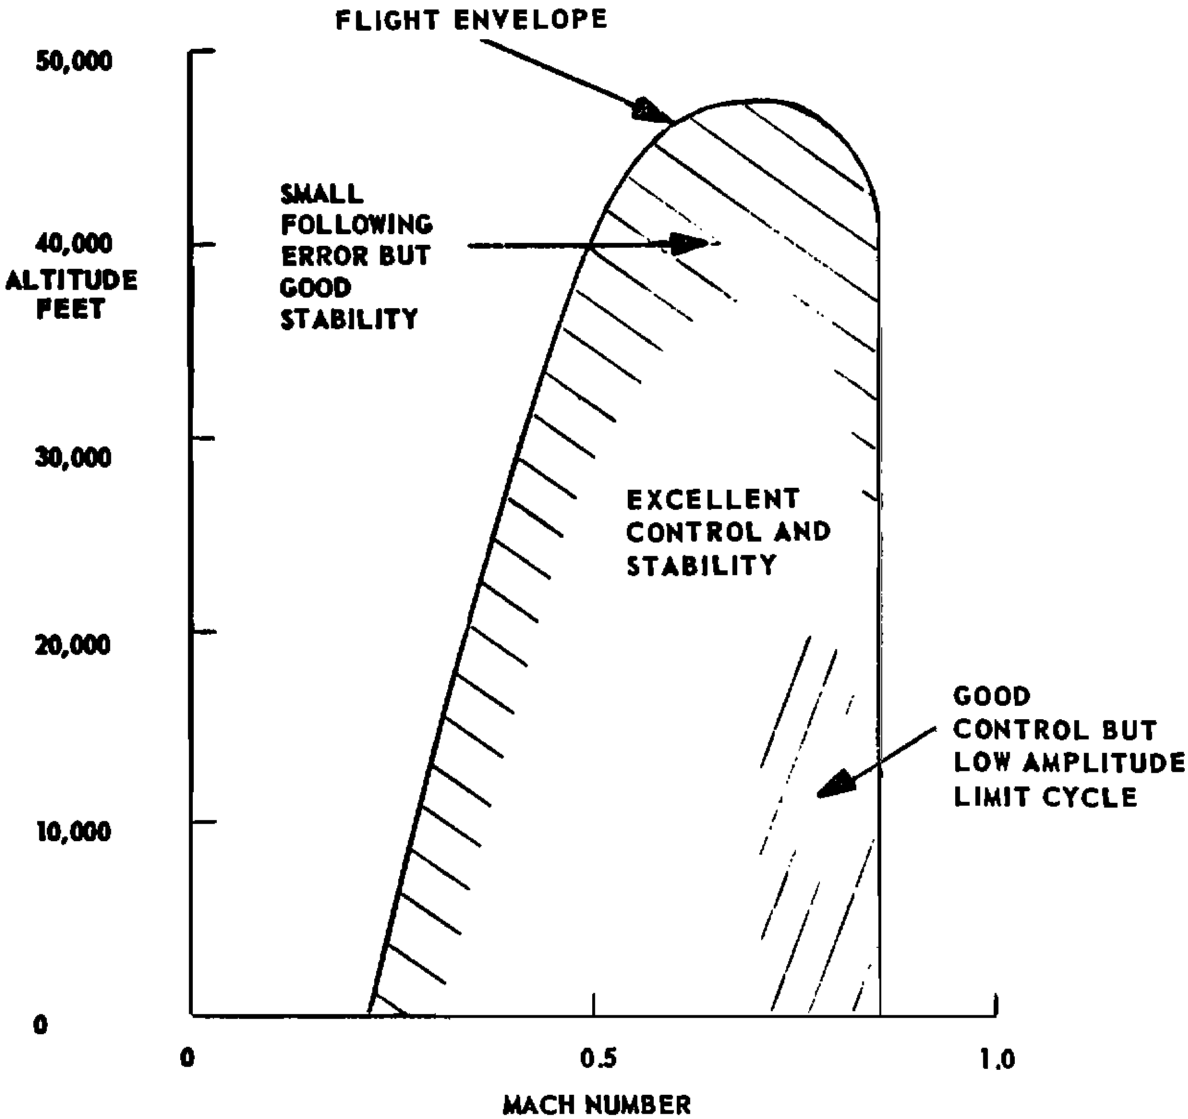

Figure 12 Summary of Adaptive Pitch-rate Control F-94C Flight Test Results (Minimum System - No Gain Changer)

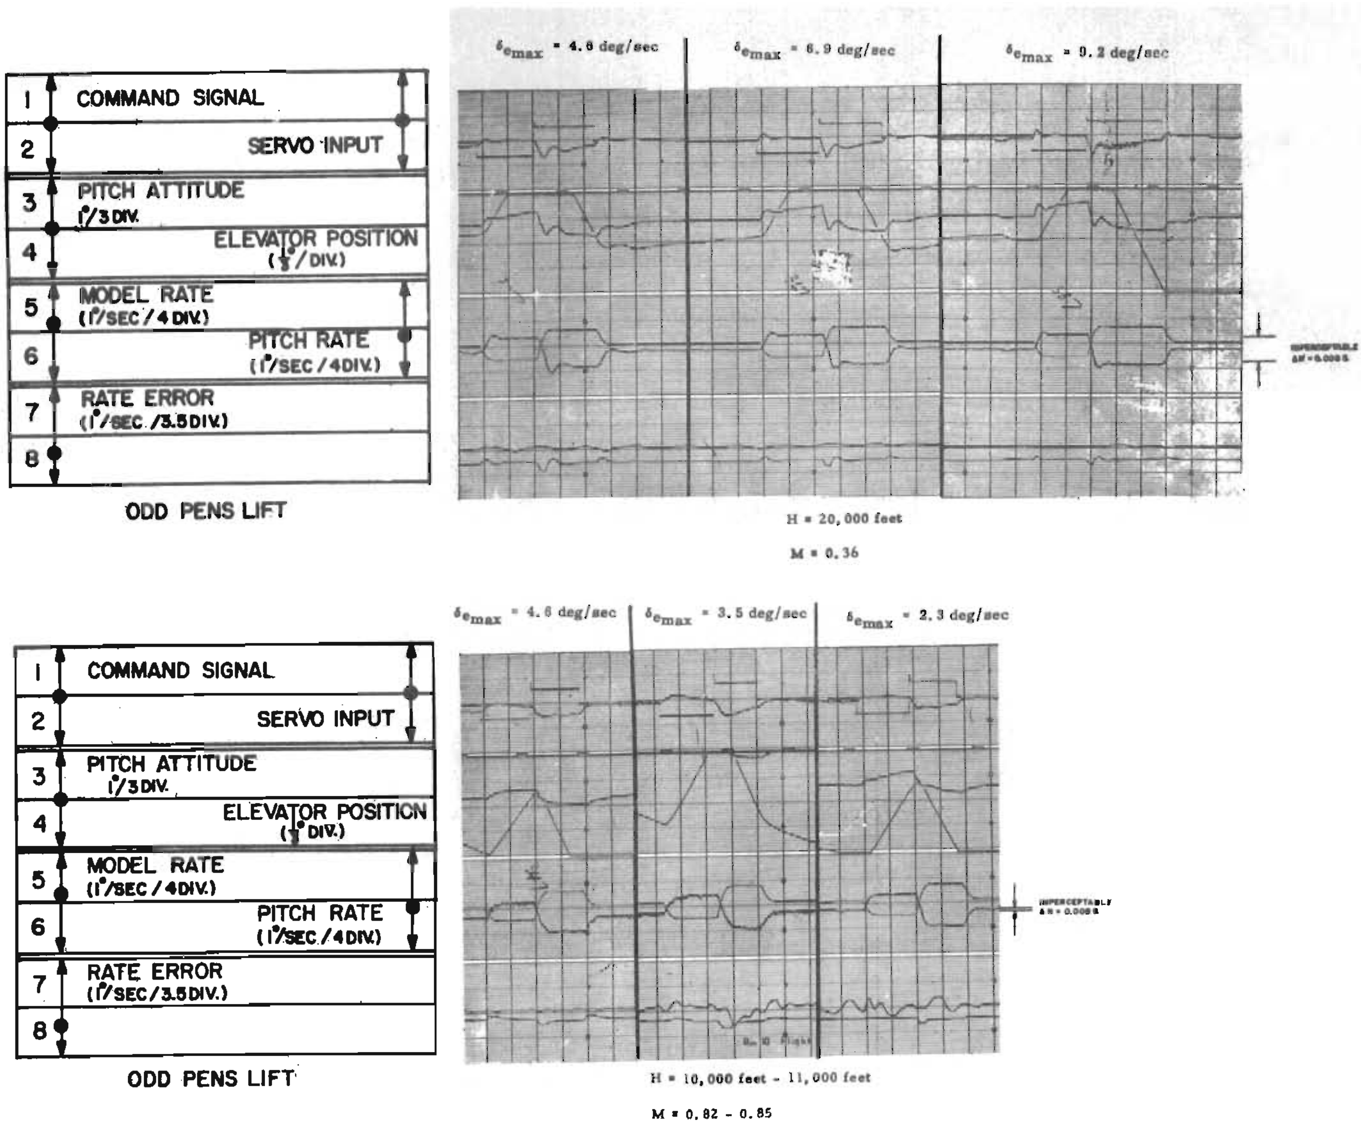

Figure 13. Flight Test Results of Effect of Limiter Amplitude