Manual Attitude Control of the Lunar Module

ROBERT F. STENGEL*
Charles Stark Draper Laboratory, Massachusetts Institute of Technology, Cambridge, Mass.

During critical phases of the Apollo mission, spacecraft attitude is manually controlled by the astronaut-pilot. In such instances, the pilot's ability to perceive the state of his vehicle and to apply alternatives to the nominal control is an asset. The primary guidance, navigation, and control system uses digital control which is executed on a time-shared basis within the Lunar Guidance Computer. Control laws have scheduled gains, are nonlinear, and follow conditional paths within the computer. Manual attitude control using a rate command/attitude hold mode employs parallel logical paths for fast, precise rate response and for attitude hold about uncommanded axes. Improved handling qualities afford reduced control jet usage and miss distance during the landing.

Introduction

THE Lunar Module of the Apollo program presents a control problem unlike that of any other manned vehicle. The Lunar Module (LM) operates as a VTOL craft in an airless environment with one-sixth earth's gravitation. It is a deboost and launch vehicle capable of leaving, and returning to, lunar orbit. It is a spacecraft that must rendezvous and dock with another spacecraft. Attitude control is provided by a reaction jet system, and, during automatic powered flight with the descent stage propulsion system, by a main engine whose gimbal system was designed to provide only bias acceleration trim. Primary control logic is entirely digital. Fuel available for attitude control is strictly budgeted. At critical times during the lunar mission, the Lunar Module is controlled directly by its crew.

Descending from lunar orbit, the LM follows a braking trajectory under explicit automatic guidance. The landing point can be retargeted manually during this phase,[1] but the pilot does not enter the control loop. Upon reaching the landing point with nulled horizontal velocity, the LM is aligned with the local vertical in a hover mode and descends to the lunar surface with a 3-fps sink rate. The automatic system is incapable, however, of sensing and avoiding obstructions, such as craters or rocks, nor can it decide at the last moment that a nearby site is a more favorable landing point; it cannot cope with the unexpected. For the terminal phase of the lunar landing, the pilot must have the option of flying his craft to an alternate location.

The Lunar Module performs attitude maneuvers during unpowered orbital flight. To achieve a precisely defined attitude, such as that required for initial thrust-vector positioning, an automatic maneuver usually is more efficient. For less precisely determined tasks, such as coarse alignment of the inertial measurement unit (IMU), docking, or station-keeping, manual control is better. In addition, IMU gimbal-lock avoidance is not automatic; it must be performed manually.

Two manual attitude control modes are implemented in the primary guidance, navigation, and control system.[2] The minimum impulse mode provides a single 14-msec thruster pulse each time the controller is deflected. This mode is useful when a small, steady rate is required, as in IMU alignment. The major mode, and the subject of this paper, is the rate command/attitude hold mode, which incorporates a number of features that enhance the rapidity and precision of control response.

This paper presents a brief summary of the LM's control characteristics, followed by a discussion of the rate command/attitude hold mode used for Apollo 9. The LM manual mode designed for Apollo 10 and later flights is then described. Test results, including lunar landing simulations, are presented.

Received August 11, 1969; presented as Paper 69-892 at the AIAA Guidance, Control, and Flight Mechaics Conference, Princeton, N. J., August 18–20, 1969; revision received May 5, 1970; this work has been performed under NASA Contract NAS 9-4065. The contributions of the Charles Stark Draper Laboratory's Control & Flight Dynamics Division, and particularly those of D. W. Keene, are gratefully acknowledged.

* Member of the Research Staff. Member AIAA.

Background

Control Characteristics of the Lunar Module

The LM, shown in Fig. 1, assumes three control configurations during the lunar mission; the LM-alone descent configuration, Fig. 1, with a fully-loaded mass of 15,000 kg; the LM-alone ascent configuration with an empty mass of 2600 kg; and the CSM-docked configuration, which has a fully-loaded mass of 42,000 kg. Principal axes of the spacecraft are close to the X,Y,Z body axes in all configurations.

The Reaction Control System (RCS) is composed of two 8-thruster systems. Each rocket produces a nominal 445-N thrust, although the down-firing jets produce secondary thrusts and torques due to exhaust impingement on the descent stage. Yaw jets each provide torques of 695 Nm^2, while the roll and pitch jets produce primary torques of 746 Nm^2. Total RCS propellant mass is 267 kg. Propellant is used at the rate of 0.16 kg/sec/jet.

Figure 1 indicates the orientation of the control axes with respect to the spacecraft. The X,Y,Z translational body-axis system is coincident with the P,Q,R rotational rate coordinates. X is parallel to the thrust axis and is vertical with respect to the crew; the Y-axis is to the astronauts' right, and the Z-axis points forward. The terms "yaw, pitch, and roll" are pilot-oriented; thus, a P rotation is yaw, while Q and R represent pitch and roll. The RCS thrusters are oriented 45° from the Q,R axes. To facilitate jet selection and command, the P,U,V system is defined with the U and V axes passing through the thruster assemblies. Yaw control is executed about the P axis alone. Roll and pitch control is implemented in the U',V' frame, control axes which are skewed from the U,V axes to account for inertial coupling effects.

Manual attitude commands use the RCS thrusters only. The trim gimbal system of the descent stage main engine is driven in its bias acceleration nulling mode during manual control of powered flight, as in the lunar landing. To pre-

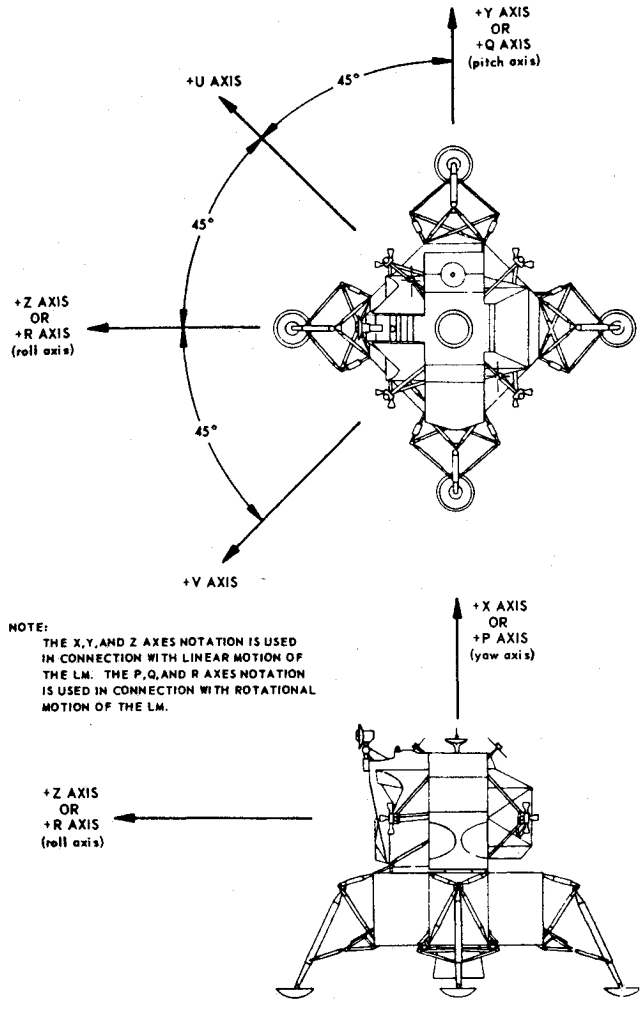

Fig. 1 Control axes of the LM.

vent short RCS firing, a minimum impulse of 14-msec is imposed within the Lunar Guidance Computer. This produces a finite resolution in rate change which varies with moment-of-inertia.

Characteristics of the attitude controller assembly (ACA) are illustrated in Fig. 2. This 3-axis hand controller has an 800-Hz output proportional to deflection, beginning 2° from the center position. The voltage output is digitized to 42 counts at the 10° soft stop deflection; the soft stop is considered "full-scale," although voltage and counts increase until the hard stop is reached at 13°. The detent switch closes at 1.25°, ±0.75°. It is possible to close this switch, informing the guidance computer of a deflection, without issuing a command voltage; control can switch from automatic to manual logic with zero commanded rate in this case.

Attitude control, jet select logic for translations using the RCS thrusters, and attitude estimates are computed in that portion of the Lunar Guidance Computer (LGC) program called the digital autopilot.[3] Manual attitude commands are also processed in this coding. The digital autopilot (DAP) occupies 11% of the 38,912-word LGC memory. The DAP is executed 10 times/sec. Computation time varies between 6 and 25 msec, depending upon the mode of operation and the number of interrupts which occur during the execution.

Estimates of bias acceleration, rate, and attitude are calculated in a recursive system of equations similar to a Kalman filter; however, filter gains are not the Kalman optimal values. The state is extrapolated using past estimates plus computed contributions due to descent engine rotation and RCS thruster firings. The extrapolation is corrected by a nonlinear function of the difference between the extrapolated attitude and the measurement of IMU gimbal angles.[3]

Attitude control using RCS thrusters is implemented with phase plane logic about three independent axes (P,U',V'). The switch curves, which are parabolic for the LM-alone case, determine command times and firing duration for the thrusters. Attitude is controlled within a deadband; thus there is a coast zone between switch curves. The LM-alone switch curves are shaped as a function of the mass estimate. In powered flight, they are positioned by the bias acceleration estimate to allow limit cycling with one-sided RCS firings.

Manual Rate Command Design Considerations

The LGC program for Apollo 9 contains manual rate command logic which is a digital realization of early reaction control rate command systems, such as those used in design studies on the Lunar Landing Research Vehicle (LLRV), and with the Lunar Landing Research Facility (LLRF). This logic[4] provides rate command with a rate error deadband when the hand controller is out-of-detent. Automatic attitude hold is maintained with the controller in-detent.

This manual rate command programing shares the jet-select and skip logic with the autopilot logic, but computes its own thruster on-times, taking into account rate error, vehicle inertia, and bias acceleration. The on-time is targeted for zero rate error; however, once the rate error is less than the rate deadband, the RCS thrusters are commanded off. Unlike corresponding analog systems, rate error is frequently nulled well within the deadband. Firing duration is not limited to multiples of the sampling period but can be timed within milliseconds. The jet selection policy choses thruster pairs to provide rotation of the correct sense and to account for detected jet failures. Pure couples are commanded when the jets are available, except in powered ascent, when only $+X$-firing jets are used for small rate errors. The skip logic prevents on-switching of the same jet more often than 5 times/sec.

The manual rate command mode used for Apollo 10 and all later missions, has been designed to meet the following objectives: 1) to reduce drift about uncommanded axes, 2) to provide precise rate control, 3) to improve handling qualities, 4) to assure positive return to automatic modes from manual control, and 5) to reduce on-time of the $+X$-firing thrusters during the lunar landing.

Although rate command with deadband is an acceptable mode from a handling qualities viewpoint, it is open-loop for small, secular errors. Because the switch curves are independent of attitude error, it is possible for the spacecraft to have an uncorrected drift rate that is incrementally smaller than the deadband.

Two factors complicate the drift problem. If the hand controller is out-of-detent about any axis, all three axes use the manual logic. Consequently, the spacecraft can drift about uncommanded axes. Also, an unmodelled (bias) acceleration can cause the phase point to chatter along the

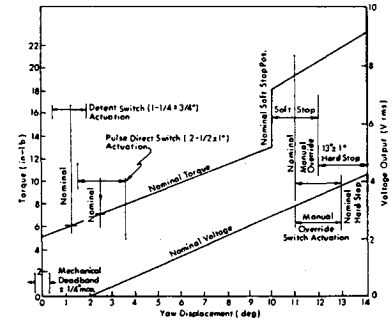

Fig. 2 Attitude controller assembly characteristics.

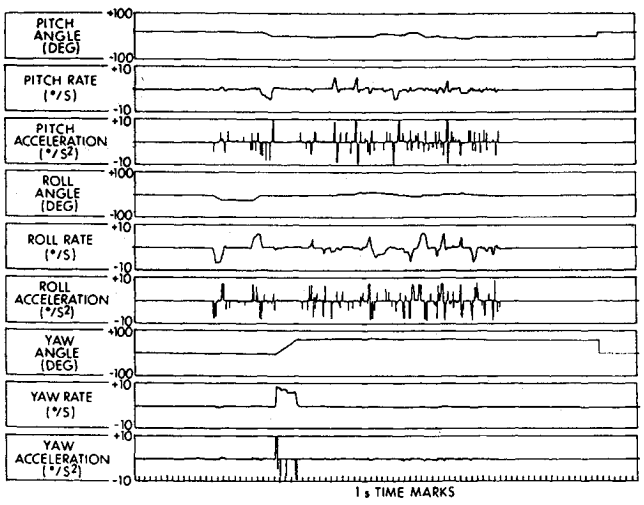

Fig. 3 Time history of a manual lunar landing simulation.

switch curve. Sampling and state estimation delays compound this drift. To limit drift inexorably, attitude errors must be incorporated in the control computations.

The rate deadband determines the resolution of rate control. Targeting jet on-time for zero rate error and using inertia and bias acceleration estimates often result in rate step response with errors initially smaller than the deadband, but such precision cannot be guaranteed. Once the rate error is within the deadband at a sampling instant, firing ceases. Furthermore, uncertainty in command response is large: if the rate error is incrementally smaller than the negative limit and a positive change is requested, the rate error must traverse the entire dead zone before a firing occurs. To obtain precise rate control, integral compensation is required.

Rate uncertainty is reflected in the manually commanded rates obtained during 15 lunar landing simulations[5,6] using NASA's Lunar Module Mission Simulator (LMS) and the Apollo 9 manual control mode. As shown in Fig. 3, vehicle rates during the landing are best characterized as impulses. This indicates that integrals of angular rate, e.g., angular attitude and horizontal translation, are the quantities of highest concern to the pilot.

Histograms of the maximum pitch and roll rates obtained in the 15 landings are presented in Fig. 4. It is evident that rates within the deadband are excluded. The pilot does not need high rates often; no commands greater than 10° per sec are issued in any of the tests, although the Maximum Commanded Rate available is 20° per sec. Three to 4° per sec rates occur most frequently in all quadrants, and the smallest available rates, 2 to 3° per sec, are not used often. In summary: a) very small rates are excluded by the rate deadband; b) there is a hesitance to use rates just outside the deadband; c) with the exception of the "excluded middle," there appears to be a desire to use the smallest rates that can be obtained with certainty.

Tightening the rate loop is insufficient to improve the pilot's assessment of handling qualities in the landing task. One astronaut's reaction to the change was that small attitude changes still could not be commanded reliably. The difficulty lay in the deflection required to obtain hand controller output and in the controller sensitivity. In "shirtsleeves" simulator tests, the pilot can feel and hear the detent switch click, yet, depending upon the particular unit, the controller must be deflected up to 1.5° more before the output voltage begins to build up. Having reached the voltage ramp at 2° deflection, the commanded rate increases 2.5° per sec for each degree of controller deflection, in steps of nearly 0.5° per sec. Using small, smooth hand motions, small attitude changes are indeed difficult to command with the Apollo 9 manual mode control, because the pilot cannot predict when the voltage build-up will begin.

Once the hand controller has been returned to detent, control should be passed from manual rate command to automatic attitude hold positively and with a minimum transient. When the rate error about an axis has been reduced to a small value, that axis should be considered to have passed the damping test; however, return to automatic control should be assured whenever the hand controller is returned to detent, even if any or all components of angular rate fail to damp within a short time.

Tests of an early version of the Apollo 11 LGC program uncovered an excessive total on-time of the RCS thrusters during manual landing simulations. The additional RCS propellant usage was discomforting, but the primary concern was the cumulative heating of the descent stage which would be caused by exhaust impingement of the down-firing ($+X$) RCS thrusters. Inhibiting the $+X$ jets for small rate errors was proposed as a solution; however, the deterioration in handling qualities was unacceptable. In simulations, pilots were forced to use larger rates more often, bringing the $+X$ jets back into use. As a result, on-time savings were unpredictable. This was one indication that handling qualities were at the base of the problem. Prior research indirectly indicated that improved handling qualities, through reduced controller sensitivity, might alleviate the problem.[7] This proved to be the correct solution, as will be shown in a later section.

In addition to the preceding objectives, the manual rate command mode must operate satisfactorily in the presence of detected and undetected jet failures (on and off), with mass estimate errors in the guidance computer, and in the presence of unmodelled accelerations.

Manual Rate Command Logic

Two modes of control are employed in the new manual rate command: the direct rate and the pseudo-auto modes. Neither mode alone meets the requirements of the previous section; however, by alternating between the 2 policies, fast, precise response is obtained.

The block diagram of Fig. 5 traces the major paths of this digital control system for all three axes. The hand controller deflection is scaled to a rate command, which is differenced with the estimated vehicle rate to form a rate error. The command is integrated and is subtracted from an integrated sum of IMU gimbal angle increments transformed to body axes, providing an attitude error. Attitude and rate errors are referenced to body axes; therefore, the attitude errors are meaningful only when they are relatively small. One of the features of the dual mode operation is that these errors rarely become large.

The commanded rate, ω_c, is supplied to the control logic to perform a switching function. If the change in commanded rate between successive autopilot cycles exceeds a breakout level, the direct rate mode is used to null rate error, comput-

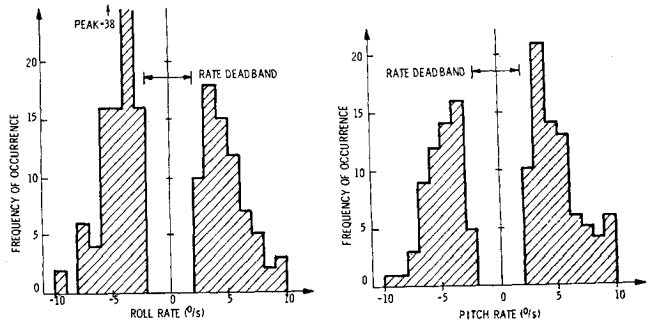

Fig. 4 Frequency of manually commanded rates.

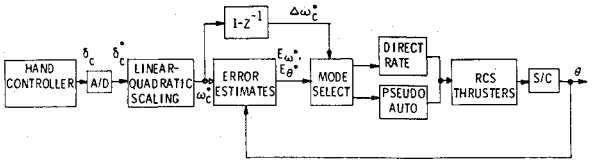

Fig. 5 Manual rate control logic.

ing a firing time for the RCS thrusters as,

$$T_{\text{jet}} = -E_\omega/a \tag{1}$$

where, E_ω = actual rate minus commanded rate, a = net available rotational acceleration, including estimated bias acceleration.

The sign of T_{jet} determines the sense of the commanded rotation. If the firing time exceeds 150 msec, the jets are commanded on, and a new firing time is computed on the next DAP cycle. If the firing time is less than 150 msec (but not zero), the jets are commanded on for the correct time, but the computation is skipped on the next cycle. This procedure takes place: a) every time the breakout level is exceeded, b) until E falls within the target deadband, or c) until a time limit is exceeded.

The direct rate mode has two important attributes. It provides immediate response to urgent commands, while ignoring slow controller inputs. The commanded rate need not be large, but if it is requested quickly, the direct rate mode is used. It also allows large rates to be commanded without significant rate overshoot. With constant acceleration, phase plane trajectories are parabolas; however, the phase plane logic of the automatic system, shown in Fig. 6, is always entered (from direct rate) below the target deadband, with attitude error initialized to zero, i.e., in the coast zone.

The pseudo-auto mode, which uses the phase plane logic shown in Fig. 6, limits drift about uncommanded axes and makes precision rate control possible. Attitude cannot drift beyond the attitude deadband, and small rate errors are integrated to the coast zone boundary, where a corrective RCS firing is commanded. If the rate command is applied slowly, the pseudo-auto mode is entered immediately. If the direct rate mode is used initially, pseudo-auto control begins when the rate error is less than the target deadband or when the time limit is exceeded. If the mode is entered with jets on, the firing is continued to zero rate error, even though the phase point is in the coast zone.

Direct rate/pseudo-auto switching is carried out independently in the P axis and Q,R-axes; thus it is possible to hold attitude in yaw while urgent commands are issued in pitch or roll and vice versa. The P-axis coding is separated from the Q,R-axes coding to allow a manual yaw maneuver during powered flight. The logic is quite similar in the 2 frames, except that additional branches are required for the multi-axis, Q,R control. Executive control of the rate command mode, including initialization, reading, zeroing, and enabling of attitude controller assembly counters, and the decision to continue rate damping or to return to automatic control, occurs in the P-axis coding.

The breakout level and target deadband are currently set to the same value, 0.6° per sec for the LM-alone, and the time limit is 4 sec. The pseudo-auto attitude deadband is 0.3° during the lunar landing.

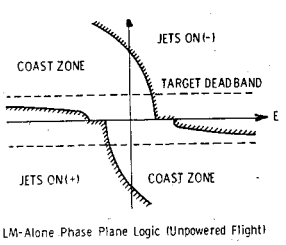

Fig. 6 Phase plane control laws.

Hand Controller Scaling

The sensitivity of commanded rate to hand controller deflection is the most important manual control parameter once rotational control acceleration ("control power") is fixed. A range of controller sensitivities that provides stable human pilot-loop closures often can be defined, but, as a consequence of the pilot's adaptive ability, optimization of the sensitivity within that range is a subjective process. The choice of scaling can be affected strongly by control power, vehicle and control system dynamics, external disturbances, the control task, and the individual pilot's ability to perceive and react.

Early research on manual lunar landing explored the relationship between controller scaling and control power in fixed- and moving-base simulators.[7-12] Cooper Ratings, numbers assigned to subjective evaluations of handling qualities, are the primary figures of merit in these tests, although one fixed-base test[7] estimated RCS propellant consumption of the LLRV as a function of controller sensitivity and rate deadband. Evaluations based on landing point accuracy are notably absent.

Hewes summarizes control sensitivity/control power results obtained for pitch control with the LLRV and the LLRF.[12] He finds that when control power and the Maximum Commanded Rate (MCR) are equal, i.e., when the MCR can be reached in 1 second, the ratings are good. Hewes suggests that the combination of 12.5° per sec MCR and 12.5° per sec control power is optimum.

Jarvis[7] presents flight test and fixed-base RCS propellant consumption as a function of controller sensitivity and rate deadband. Reduced controller sensitivity has a striking effect on the RCS propellant used; there is a monotonic reduction with decreasing sensitivity.

With descent engine propellant fully used, the 4-jet control powers of the Lunar Module are 11, 9, and 10° per sec² in pitch, roll, and yaw; during the terminal phase of landing, control powers are 10 to 20% less. The Hewes data suggest that the 20° per sec MCR of Apollo 9 is a factor of 2 higher than the optima for these control powers. Jarvis's fuel-usage data also point in the direction of reduced sensitivity.

These trends are confirmed in LMS simulations using the new manual rate command logic. Although Cooper Ratings were not recorded in all the tests, a reduction of Maximum Commanded Rate from 20 to 14° per sec improved one experienced pilot's rating by 1; the consensus of several pilots was that handling qualities were improved. The emphasis in these tests was placed on accuracy in flying to a designated site and on reducing $+X$-jet on-time; however, handling qualities ratings continued to improve as the MCR was reduced to a minimum value of 8° per sec. The effect on $+X$-jet on-time is shown in Fig. 7. Flight to three landing sites with six MCR's was evaluated by two pilots. The average for the two hardest tasks shows a marked decrease in jet on-time with decreasing MCR. The scatter in these points emphasizes the subjective nature of choosing controller sensitivity, although the favorable effect of reduced sensitivity is apparent. Reduced MCR also improved landing point accuracy in this series.

In spite of the improvements resulting from reduced controller sensitivity, there remained one conflict: reduced sensitivity made small rates and small angle changes easier to obtain, but there was concern that the MCR was insufficient for emergency conditions. The solution that has been adopted is linear-quadratic scaling of the hand controller output.

Measuring controller deflection, $\delta°$, from the centered position, the linear scaling law of the Apollo 9 program is

$$\omega_c = \begin{cases} \dfrac{\text{MCR}}{8}\,\text{sgn}(\delta)\,[|\delta| - 2], & |\delta| > 2 \\ 0, & |\delta| \leqq 2 \end{cases} \tag{2}$$

The maximum commanded rate is realized when $|\delta|$ is 10°. The revised scaling law is,

$$\omega_c = \frac{\text{MCR}}{40} \operatorname{sgn}(\delta) \left[|\delta| - 2 + \frac{(|\delta| - 2)^2}{2} \right] \quad (3)$$

with MCR equal to 20° per sec in normal scaling, and 4° per sec in fine scaling.

Note that the linear portion of Eq. (3) would yield a 4° per sec MCR at 10° maximum deflection, one-fifth of the former value. This characteristic is compared with several linear scalings in Fig. 8. It is interesting that the 10° per sec linear law, which is near-optimum according to the Hewes data, intersects the new law at 4° per sec, which is in the most-frequently-used range of Fig. 4.

Equation (3) is not an exact relation, for the hand controller output is quantized and ω_c is "staircased." Calling δ_c the number of counts (an integer), the commanded rate is

$$\omega_c = \text{MCR} \ [0.00045335 \ \delta_c(|\delta_c| + 10.5)] \quad (4)$$

and the increase in ω_c per count is approximately,

$$|\Delta\omega_c| = \text{MCR}[0.00045335(2|\delta_c| + 10.5)] \quad (5)$$

Little has been said about scaling for the docking task. Inertias of the nearly-empty ascent configuration are quite low; control power and minimum impulse are 4 to 7 times larger than those of the landing configuration. In the only published account of LM-active docking simulation, successful docking was achieved with 20° per sec MCR, although the possibilities for over-control were great.[13] This report indicates that attitude control is used to set up a proper attitude relative to the passive target. Once this is achieved, translational control is the primary concern, and the attitude control receives little exercise. It is important, however, that precision attitude control be provided, for several of the simulated docking attempts aborted as a result of overcorrections just prior to contact. The new manual rate command provides a 4° per sec MCR in fine scaling, and Eq. (4) is used to give low sensitivity at small deflection.

The relationship of hand controller scaling to direct rate breakout level is designed to emphasize the most useful traits of the rate command system. In the normal scaling used for the landing, the ratio of controller sensitivity-to-breakout level is high; small deflection changes are considered urgent, and, as the mean deflection increases, the change becomes more urgent. Equation (5) indicates that every count triggers the direct rate mode when δ_c is 28 or more, i.e., when the commanded rate is greater than 9.77° per sec. The breakout level, 0.6° per sec, is unchanged for fine scaling, in order to encourage control to remain with the pseudo-auto mode. Each time the direct rate mode is used, the pseudo-auto attitude reference is reinitialized. So long as direct rate is not used, absolute attitude reference is maintained; therefore, a small pitch command can be issued without disturbing roll reference. This is a useful attribute in the docking task, for the pilot usually sets up each axis individually, attaining the proper orientation in a sequential process.[13]

Multiaxis Control

Because the RCS thrusters are located 45° from the Y,Z axes, the same eight jets are used for both pitch and roll control. Referring to Fig. 1, it can be seen that pure pitch or pure roll rotations use four jets, while combined rotations occur about either the U or V axis, using only two thrusters. When a combined pitch-roll rotation occurs, control power about both axes is reduced by a factor of 2. Figure 9 indicates that there are two approaches to nulling multi-axis errors. It is possible to respond about both axes immediately with two jets, nulling the smaller error entirely, then

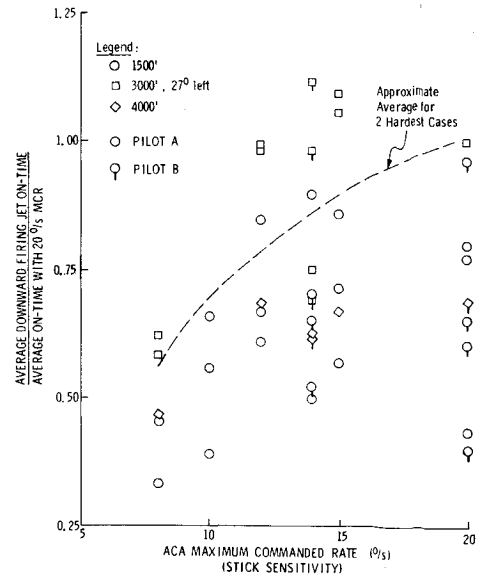

Fig. 7 Effect of controller sensitivity on RCS firing time.

nulling the remaining error with four jets (Path no. 1) The secondalternative is to reduce the larger error with four jets (Path no. 2). When both errors are equal, both axes are nulled using two jets.

The latter method is intrinsic to U,V control axes and is the minimum-integral-error path. It is more natural to the pilot as well. Should the pilot issue equal commands in pitch and roll and then sense that roll control is more urgent, he would increase roll deflection of the hand controller. With the second logic, roll control would increase immediately. In the first case, pitch error would be nulled entirely before full control power was applied to the critical axis.

Test Results

All LGC programing receives extensive digital and hybrid computer simulation testing prior to release. The digital simulation provides trajectory computation and step-by-step execution of LGC coding in a batch processing mode.[14] The real-time hybrid simulation provides a fixed-base cockpit environment, utilizing flight display hardware as well as IMU, radar, and LGC functional simulators.[15] Digital trajectory computations, combined with analog vehicle simulation and "window optics," are used for mission design and verification studies, including star tracking and lunar landing.

Figure 10 is representative of standard response of the LM to a sequence of manually commanded step inputs. Measurement noise and jet switching delays are included in this digital simulation of the descent configuration, with fully-loaded ascent propellant tanks and 90%-loaded descent propellant tanks. The figure is a time history of LGC-

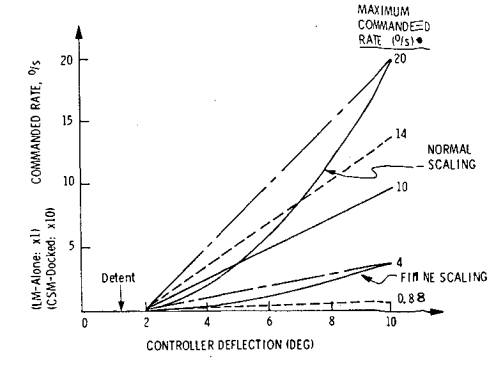

Fig. 8 Luminary rotational hand controller scaling.

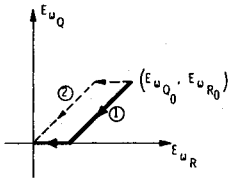

Fig. 9 Multiaxis control paths.

estimated and actual rotational rates; the latter trace is marked by X's.

Each control response is initiated in the direct rate mode. The P-axis hand controller deflection at 15 sec is sensed by the digital autopilot at 15.05 sec. At 15.15 sec, the direct rate mode is entered with $-4°$ per sec rate error, since the breakout level has been exceeded by the step command. Rate error falls within the target deadband at 15.95 sec, when the actual and estimated rates are 3.45 and 3.41° per sec. The firing is continued by the pseudo-auto mode until 16.18 sec; actual and estimated rates are 4.08 and 4.00° per sec on the next DAP cycle. As the state estimate improves and the attitude error is driven to the coast zone boundary, the actual rate is corrected to 3.99° per sec. Throughout the P command, the Q,R-axes are maintained in pseudo-auto attitude hold, employing minimum impulses.

The remaining commands are executed in a similar way. Although each axis has rotated nearly 40°, the final yaw, pitch, and roll angles are 0.25°, 1.02°, and 0.26° away from their original positions. It should be remembered that absolute attitude reference is not maintained in this sequence. As long as there are no bias accelerations, there is equal probability that the attitude error at reinitialization is positive or negative, and the sample mean approaches zero as the number of reinitializations becomes large; however, there is no guarantee against short-term trends in attitude drift. In the presence of a bias, a net drift in a given direction can be expected with each reinitialization.

A total of 6.72 kg of RCS propellant is used for this test sequence. This is 8.5% above the theoretical minimum; the excess can be attributed to attitude hold limit cycling and to corrections necessary to obtain the commanded rate.

Raising the commanded rate to 20° per sec, similar performance is noted, although response time increases to 4 sec. The minimum angle traversed in achieving a rate, starting

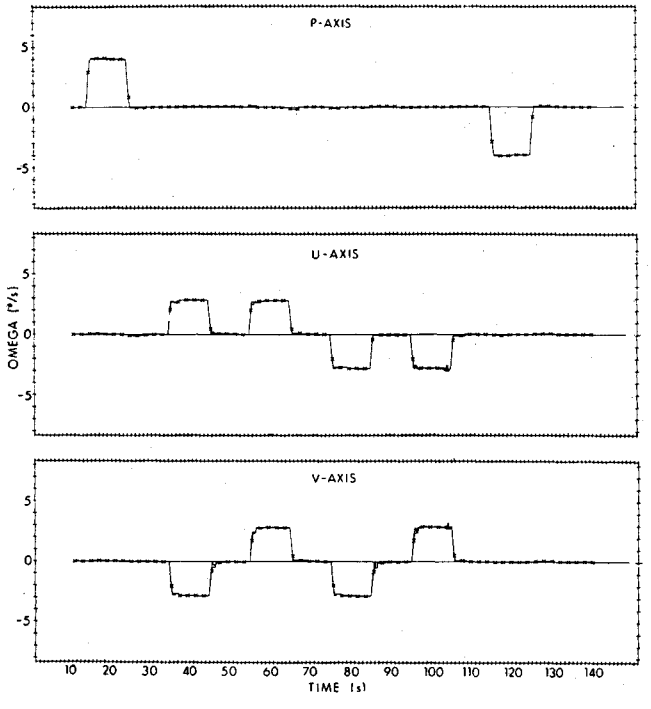

Fig. 10 Rate step response of the descent configuration.

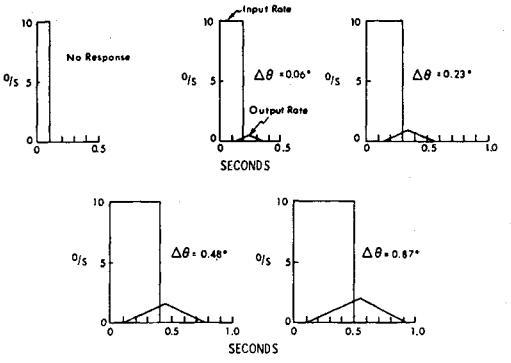

Fig. 11 Impulse response of LM (heavy descent).

from and returning to zero rate, is

$$\Delta\theta_{\min} = \omega^2_{\max}/a \tag{6}$$

where a is the rotational acceleration. For this case, $\Delta\theta_{\min}$ is 80°; for the landing configuration, it is half that. It is clear that the normal MCR is available only for large angle changes in either case.

Very small rates can be commanded: a single hand controller count provides a commanded rate which is 0.5% of the MCR. Over long-term averages, the system will supply this low rate within the resolution of the minimum impulses used to maintain a stable limit cycle.

The impulse response of the manual rate command system plotted in Fig. 11 is nonlinear and contains transport lag. The 10° per sec impulse, lasting from 0.1 to 0.5 second, is a saturating input. Vehicle response would be identical for any commanded rate greater than the maximum achieved rate. Response is obtained by "bang-bang" firing of the RCS thrusters, with no coast time before thrust reversal. There is no response to the 0.1-sec impulse; initialization lag obscures the input. For longer saturating impulses, lasting T sec (in 0.1-sec increments), the maximum angular rate is approximately,

$$\omega_{\max} \cong a(T - 0.1) \tag{7}$$

and the net angle change is

$$\Delta\theta \cong a(T - 0.1)^2 \tag{8}$$

The smallest angle change obtained with the heavy descent configuration in digital simulations is 0.06°; with landing inertias, this increment is doubled. Thus, the availability of small attitude changes is not dependent on low sensitivity. Using saturating inputs, the angle change is nearly time-optimal and its magnitude is independent of the amount of controller deflection, depending only on the duration of the command.

Ramp inputs are not especially useful in controlling the Lunar Module. If, however, a slow ramp is commanded, the direct rate mode is not used, and there is initial delay as the pseudo-auto phase point drives toward a switching curve. Command buildup and vehicle response are quadratic

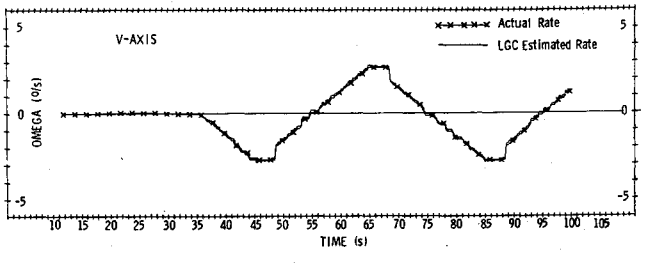

Fig. 12 Response to a ramp input, time history.

and quantized. If the controller deflection is then reversed, there is a lag as the attitude error shifts from one switching boundary to the opposite These effects are apparent in the time history of Fig. 12, and where there is a flat spot at the peak rate (the quadratic shape is absent because linear scaling was used in this test). The phase trajectory of Fig. 13 provides the reason for this behavior. During the rate buildup, the phase point bounces along the switch curve "flat" (see Fig. 6). When the input is reversed, the point traverses the coast zone, and the rate error is nulled.

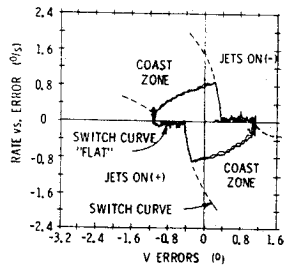

Fig. 13 Response to a ramp input, phase trajectory.

Manual Control of the Lunar Landing

The most challenging manual task of the lunar mission is landing the LM on the moon. The maneuver is comparable to a helicopter landing on earth, but the differences are at least as significant as the similarities. The LM's rotational control power is only a fraction of helicopter control power. Translational control is obtained by tilting the main engine thrust vector, which entails tilting the whole spacecraft. In hover, thrust magnitude is one-sixth what it would be on earth, so the vehicle must be rotated six times as far to obtain an equal translational acceleration. Visibility is restricted in the Lunar Module, and external motion cues are more difficult to obtain than in a bubble-canopied helicopter. The main engine depletes its propellant within less than 4 min. of the beginning of manual control, forcing the landing to be made quickly.

The LM is well-instrumented, providing all the flight status information necessary for a successful landing. In addition to an all-axis attitude indicator, there are meters indicating horizontal velocity components, altitude, altitude rate, and thrust. The second crew member can call additional information on the computer display registers, including total horizontal velocity, mass estimate, and the slant range to the landing site. Although the main engine can be throttled manually, rate-of-descent is usually maintained by the guidance computer, with the astronaut specifying sink rate through a "click" switch. The latter allows the pilot to concentrate on control of attitude and horizontal translation.

A nominal manual landing sequence begins at an altitude of 500 ft. Forward velocity and sink rate are 58 fps and 15 fps, respectively. The LM is pitched up 17°. In principle, it is not difficult for the astronaut to guide the spacecraft from this point to touchdown, provided there are no obstacles to be avoided. Lateral velocity, roll angle, and yaw angle are small and may not need adjustment. Rate-of-descent is gradually decreased to 3 fps, positive pitch angle is maintained to null the forward velocity, and the vehicle is then aligned with the local vertical for a low-rate descent to the surface. In such an instance, the landing maneuver is a programed sequence requiring negligible control.

The real landing is somewhat different. Even small angle and velocity biases propagate into velocity and position errors which must be controlled. If an alternate landing site is desired, angles and velocity components must be perturbed to shape the trajectory, then nulled before touchdown. Horizontal position control is a 2-axis, third-order task requiring a great deal of piloting skill.

Although landing the LM is a multi-variable control task, it appears that the pilot controls these variables sequentially. Croston's data[5] show few multi-axis rate commands, and MIT's hybrid simulation experience agrees. In order to fly to a crossrange target, two techniques are used. If the lateral distance is small, a hover technique is used. In this case, the yaw axis is ignored, and the pilot sideslips to the landing site. For greater crossrange, the technique is closer to flying an aircraft. As shown in Fig. 13, the pilot rolls the spacecraft to establish a side velocity and then nulls the roll angle. Next, the vehicle is yawed to the direction of motion, and the craft flies to the target with negligible side velocity.

Conclusions

The LM rate command system provides rapid response to urgent inputs, as well as precise response to all commands. This is made possible by using direct rate and pseudo-auto modes of control and by computing RCS firing times as functions of vehicle inertia, bias acceleration, and rate error. Urgency is gauged not by the size of the rate error but by the speed of hand controller deflection. Attitude errors are included in the control computation to limit drift about uncommanded axes and to provide integral compensation for rate control. Quadratic scaling of the hand controller output minimizes the difficulty in commanding small rates, retaining high command rate capability. Improvements in the LM's handling qualities result in more precise landings and less fuel usage, in addition to better pilot ratings.

In a sense, the rate command mode is a "second generation" system, for it draws upon analysis of extensive fixed-base and flight simulations of more basic control concepts. The earlier version programed in the LGC approximated these analog control systems. The improved manual mode is quintessentially digital, making freer use of the logical branches, counters, and nonlinearities which are so readily (and reliably) programed in the digital computer.

References

[1] Klumpp, A. R., "A Manually Retargeted Automatic Landing System for the Lunar Module (LM)," *Journal of Spacecraft and Rockets*, Vol. 5, No. 2, Feb. 1968, pp. 129–138.

[2] Miller, J. E. et al., ed., *Space Navigation, Guidance and Control*, AGARDograph 105, Technivision Ltd, Maidenhead, England, 1966.

[3] *Guidance System Operations Plan for Manned LM Earth Orbital and Lunar Missions Using Program LUMINARY*, Sec. 3, Digital Autopilot, IL R-567, Dec. 1968, MIT, Cambridge, Mass.

[4] *Guidance System Operations Plan for Manned LM Earth Orbital Missions Using Program SUNDANCE (Rev. 306)*, Sec. 3, Digital Autopilot (Rev. 1), IL R-557, Sept. 1968, MIT, Cambridge, Mass.

[5] Croston, R. C. and Thrasher, J. R., "LMS/LUMINARY 9 Lunar Landing Autopilot Tests," ASD TIR 580-S-8183, June 1968, General Electric Co., Houston, Texas.

[6] Croston, R. C., "LM RCS Duty Cycle Analysis for Lunar Landing," ASD TIR 727-S-8222, Aug. 1968, General Electric CO., Houston, Texas.

[7] Jarvis, C. R., "Flight Test Evaluation of an On-Off Rate Command Attitude Control System of a Manned Lunar-Landing Research Vehicle," TN D-3903, April 1967, NASA.

[8] Matranga, G. J. et al., "Handling Qualities and Trajectory Requirements for Terminal Lunar Landing," TN D-181, Aug. 1963, NASA.

[9] Cheatham, D. C. and Hackler, C. T., "Handling Qualities for Pilot Control of Apollo Lunar-Landing Spacecraft," *Journal of Spacecraft and Rockets*, Vol. 3, No. 5, May 1966, pp. 632–638.

[10] Mallick, D. L., Kluever, E. E., and Matranga, G. J., "Flight Results Obtained with a Non-Aerodynamic Variable Stability, Flying Platform," Society of Engineering Test Pilots, 10th Symposium and Banquet, Los Angeles, Sept. 1966.

[11] Matranga, G. J., Mallick, D. L., and Kluever, E. E., "An Assessment of Ground and Flight Simulators for the Examination of Manned Lunar Landing," AIAA Paper 67-238, Cocoa Beach, Fla., 1967.
[12] Hewes, D. E., "Interim Report on Flight Evaluations of Lunar Landing Vehicle Attitude Control Systems," AIAA Paper 67-239, Cocoa Beach, Fla., 1967.
[13] Hatch, H. G., Jr., Pennington, J. E., and Cobb, J. B., "Dynamic Simulation of Lunar Module Docking with Apollo Command Module in Lunar Orbit," TN D-3972, June 1967, NASA.
[14] Mimno, P., "Digital Simulation Manual," IL R-599, Jan. 1968, MIT, Cambridge, Mass.
[15] Goetzinger, J. D., "Pre-Simulation Report: LM Dynamics Simulation," IL E-2079, Sept. 1966, MIT, Cambridge, Mass.

AUGUST 1970 J. SPACECRAFT VOL. 7, NO. 8

Network Support Simulation and Sensitivities Analysis for the Saturn-Apollo Missions

WALKER H. LAND JR.*
IBM Corporation, Huntsville, Ala.

The approach is to develop, by simulation and analytical techniques, network sensitivity timeline error graphs. These graphs depict the nominal and perturbed coverage times, nominal and perturbed acquisition of signal (AOS), and loss of signal (LOS) as a function of a continuous range of trajectory errors at parking orbit insertion. The superposition principle, as used here, states that the total motion, $X(t)$, when subjected to a total error result, $F(t)$, is the sum of separate motions $X_1(t), X_2(t), \ldots X_n(t)$, arising from the individual error sources, $F_1(t), F_2(t), \ldots F_n(t)$, so long as the separate motions are contained within the linear region of the solution space. The resulting sensitivities graphs are not restricted to the parking orbit for a particular mission for the Saturn V. Another advantage is that a real-time computer solution is not required to reconfigure the network, as previously was the case. This approach has been experimentally verified by simulation and used in both the Apollo 7 and 8 missions. Also, included are the results of a post-flight data analysis of theoretical versus actual network coverage for the AS-204 and AS-501 missions. This analysis considers the effects of refractive bending of RF signals, multipath, signal dropout, and polarization diversity.

Method of Approach and Superposition Principle

THE error timeline shows the coverage time, AOS and LOS times for the nominal, and the errors from the nominal, for each of the separate trajectory variables. The abscissa of each error graph depicts the time (in seconds) from parking orbit insertion while the ordinate shows the error at parking orbit insertion in the stated trajectory variable. For no error, the resulting AOS and LOS times would be those of the nominal. Within the linear region of the AOS and LOS error curves, the total error may be obtained by the algebraic addition of the separate AOS and LOS errors.

Without some advanced knowledge of the magnitude of the errors at parking orbit insertion, a large number of timelines could possibly exist, all of which could depict the actual network. That is, there are five possible errors (excluding ascent-to-orbit time error which is not a trajectory error) in the trajectory variables at parking orbit insertion and a possibly infinite number (not infinite range) of values that these error variables may have. This means that $5X^5$ different combinations of errors (where X is a very large number) may exist a priori. For example, if X were 4, a small number indeed, 5120 timelines are required to describe all possible combinations of these errors. Using the superposition principle, this large number of required timelines is reduced to five, regardless of the size of X.

Examples of sensitivity graphs constructed using simulation techniques (in the vicinity of 5800–6500 sec for the fixed azimuth of 72° of the AS-205 mission) are depicted in Fig. 1.

These graphs are to be used as follows: a) determine the errors at parking orbit insertion from the real-time trajectory data by the methods depicted in the mathematical appendix; b) for these errors determine the change in the timeline for each error variable using Fig. 1; c) then algebraically add the separate deviations for each trajectory error in AOS and LOS to obtain the total AOS and LOS error for each station. Use these adjusted AOS and LOS times, for each station, to obtain the updated network. The variable launch azimuth case, to be described in the AS-503 sensitivity timeline error graph preparation, is handled in essentially the same manner.

AS-205 Sensitivity and Postflight Data Analysis

Simulation Validation of Superposition Principle

The superposition principle was experimentally verified using the perturbation cases 1–6 in Table 1.

Table 2, which shows the results of a portion of this analysis, was developed as follows. The AOS and LOS times were generated for each of the trajectory errors that were taken separately and also for these trajectory errors taken collectively. (Errors here mean deviations from the nominal.)

Presented as Paper 69-936 at the AIAA Aerospace Computer, Systems Conference Los Angeles, Calif., September 8–10, 1969; submitted September 22, 1970; revision received April 27, 1970. The author acknowledges A. McNair, NASA I-MO Contract Monitor; R. E. Gordon, who suggested the original problem; and N. K. Denmark, who worked with the author in the AS-204 and AS-501 data analysis. This work was done under NAS 8-14000 contract.

* Staff Engineer, Federal Systems Division.