


NASA TECHNICAL NOTE



N73-25888
NASA TN D-7289

NASA TN D-7289

CASE FILE
COPY

APOLLO EXPERIENCE REPORT -
GUIDANCE AND CONTROL SYSTEMS

Digital Autopilot Design Development

by William H. Peters and Kenneth J. Cox

Manned Spacecraft Center

Houston, Texas 77058

1. Report No. NASA TN D-7289		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle APOLLO EXPERIENCE REPORT GUIDANCE AND CONTROL SYSTEMS — DIGITAL AUTOPILOT DESIGN DEVELOPMENT				5. Report Date June 1973	
				6. Performing Organization Code	
7. Author(s) William H. Peters and Kenneth J. Cox, MSC				8. Performing Organization Report No. MSC S-331	
9. Performing Organization Name and Address Manned Spacecraft Center Houston, Texas 77058				10. Work Unit No. 914-50-30-02-72	
				11. Contract or Grant No.	
				13. Type of Report and Period Covered Technical Note	
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, D. C. 20546				14. Sponsoring Agency Code	
15. Supplementary Notes The MSC Director has waived the use of the International System of Units (SI) for this Apollo Experience Report, because, in his judgment, the use of SI units would impair the usefulness of the report or result in excessive cost.					
16. Abstract <p>The development of the Apollo digital autopilots (the primary attitude control systems that were used for all phases of the lunar landing mission) is summarized. This report includes design requirements, design constraints, and design philosophy. The development-process functions and the essential information flow paths are identified. Specific problem areas that existed during the development are included. A discussion is also presented on the benefits inherent in mechanizing attitude-controller logic and dynamic compensation in a digital computer.</p>					
17. Key Words (Suggested by Author(s)) · Control Systems (Lunar Module) · Control Systems (Command Module) · Apollo Project · Spacecraft Guidance · Automatic Pilots · Attitude Control · Digital Computers · Spacecraft Control				18. Distribution Statement	
19. Security Classif. (of this report) None		20. Security Classif. (of this page) None		21. No. of Pages 47	22. Price 3.00

CONTENTS

Section	Page
SUMMARY	1
INTRODUCTION	2
GENERAL	2
Development Process	2
Functional Requirements	3
Preflight Testing and Verification	4
COMMAND-SERVICE MODULE DIGITAL AUTOPILOT DEVELOPMENT	6
Vehicle Design Constraints	6
Design Philosophy	7
Design Description	8
Design-Development Problems	10
LUNAR MODULE DIGITAL AUTOPILOT DEVELOPMENT	19
Vehicle Design Constraints	20
Design Philosophy	20
Design Description	21
Design Evolution	24
FLIGHT TEST RESULTS	28
CONCLUDING REMARKS	40
REFERENCES	42

TABLES

Table		Page
I	DIGITAL AUTOPILOT TESTING MATRIX	5
II	SUMMARY OF CSM DAP DETAILED TEST OBJECTIVES	29
III	SUMMARY OF LM DAP DETAILED TEST OBJECTIVES	30
IV	APOLLO 8 FLIGHT TEST RESULTS	31
V	APOLLO 9 FLIGHT TEST RESULTS	32
VI	APOLLO 10 FLIGHT TEST RESULTS	32
VII	APOLLO 9 LM FLIGHT BURN RESIDUALS	36
VIII	APOLLO 10 LM FLIGHT BURN RESIDUALS	37

FIGURES

Figure		Page
1	Digital autopilot developmental process	3
2	Original CSM/LM filter mechanization	11
3	Original guidance and control interface	15
4	Alternate guidance control interface (ramp steering)	15
5	Final guidance and control interface (rate steering)	16
6	Powered-flight automatic control	22
7	Input and output variables of the state estimator	22
8	The phase plane when the LM is in powered ascent or when either of the trim-gimbal-nulling drive times is greater than 2 seconds during powered descent	23
9	Apollo 9 yaw data for SPS burn 5	33
10	Ascent propulsion system burn to depletion, U-axis phase plane	36

ACRONYMS

APS	ascent propulsion system
CDU	coupling data unit
CM	command module
CMC	command module computer
CSM	command-service module
DAP	digital autopilot
DOI	descent orbit insertion
DPS	descent propulsion system
DTO	detailed test objective
EMS	entry monitor subsystem
G&N	guidance and navigation
GNCS	guidance, navigation, and control system
GSOP	guidance system operations plan
GTS	gimbal-trim system
ICS	interpretive computer simulator
IMU	inertial measurement unit
LGC	LM guidance computer
LM	lunar module
MSC	Manned Spacecraft Center
RCS	reaction control system
RHC	rotational hand controller
SCS	stabilization and control system
SPS	service propulsion system
TTCA	thrust/translation controller assembly
TVC	thrust vector control

APOLLO EXPERIENCE REPORT
GUIDANCE AND CONTROL SYSTEMS —
DIGITAL AUTOPILOT DESIGN DEVELOPMENT

By William H. Peters and Kenneth J. Cox
Manned Spacecraft Center

SUMMARY

The development of the Apollo digital autopilots, which are the primary attitude control systems, is summarized for the lunar module and the command-service module. The digital autopilots provide attitude control during all phases of the Apollo mission, including a backup mode for boost into earth orbit, coasting flight, velocity-change maneuvers, lunar landing, boost into lunar orbit, docking, and entry into earth atmosphere.

The development-process functions and the essential information flow paths are identified. This identification provides information on the particular NASA/contractor interface as well as the relationships between the many individual activities. A list of primary functional requirements is presented as a prelude to the discussions of command-service module and lunar module digital autopilot development. Specific problem areas that existed during the design-development phase are discussed in detail, including the solutions that evolved. The report concludes with a discussion of flight test data from the Apollo flights that are pertinent to the demonstration of performance of these attitude control systems.

The primary conclusion concerns the benefits inherent in mechanizing controller logic and signal-shaping dynamics in a digital computer. This mechanization provided the flexibility necessary to avoid expensive hardware changes and potential schedule delays and provided a large capacity for performing complicated control-system logic with essentially no hardware impact. The feasibility and reliability of this approach have been thoroughly demonstrated by the Apollo Program. Other significant conclusions include the following.

1. Requirements for monitoring system performance can place constraints on the control-system design.
2. The danger of preventing proper control computations (and subsequent actions) arises if logic branching is too sensitive to predicted plant parameters.
3. Refined analytical techniques are necessary to reduce the amount of reliance on large-scale simulations for design verification.

INTRODUCTION

The original Apollo Program concept for obtaining a highly reliable attitude control system for the command-service module (CSM) was to carry spare equipment and perform inflight maintenance if necessary. Because of several difficulties inherent in this approach and because of a realization that redundant attitude control systems could be provided by using baseline electronic equipment, on the condition that a complete attitude-controller logic could be implemented in the guidance and navigation (G&N) computer, a decision was made to use digital autopilots. At that time, only limited flight experience using digital control systems existed. Because it was not known at the time whether CSM/lunar module (LM) powered-flight attitude control could be accomplished by digital computer logic, a parallel development of a technological backup existed. This alternate approach would have relied on the digital computer only to provide backup attitude information. Subsequent studies demonstrated the feasibility of using digital autopilot (DAP) control for this phase of flight, and the technological backup was dropped.

A broad overview of the Apollo DAP development process and specific problem areas encountered during development of the Apollo digital autopilots are discussed. The first sections of the report are applicable to all the DAP designs that were used during different phases of the mission. In later sections, the material is separated into discussions relating to the CSM DAP and the LM DAP. The discussion places emphasis on the CSM and LM thrust vector control (TVC) DAP, because more problems were associated with the development of the digital controllers for powered flight than with the development of those for coasting flight and entry.

The control systems that are the subject of this report have been described in detail elsewhere (refs. 1 to 4), and the material presented herein is descriptive only to the point of illustrating gross concepts or particular design problems. The problem discussions are written for the control-system engineer, but they do not necessarily require an intimate knowledge of the Apollo control-system designs.

GENERAL

Before the specific DAP designs are discussed by mission phase, general comments that are applicable for all the CSM and LM control-system developments will be presented. Topics that will be treated are the NASA and contractor functions and aspects of the developmental process, the general DAP functional requirements, and the preflight testing and verification philosophy used in the design development.

Development Process

The significant elements of the developmental process are diagramed in figure 1. Much of the developmental activity involved the block labeled "Modifications," as indicated by the numerous interface loops. The flexibility for incorporating design modifications into a digital control system in a straightforward, easy, and inexpensive manner (compared to changes in conventional hardware systems), and also at a later

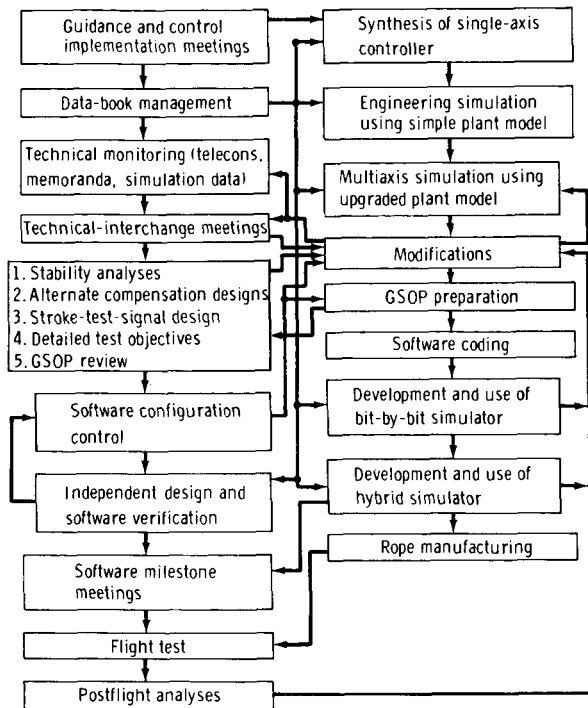


Figure 1. - Digital autopilot developmental process.

date in the developmental program than would normally be allowed, gives the designer freedom in the form of extended developmental time and the type of changes that will be tolerated at a particular point in the developmental schedule.

The activities diagramed in figure 1 depict the NASA functions on the left and the contractor functions on the right. Such a representation is not exact, however, for the flow lines represent only approximately the interactions that actually occurred. In the paragraphs that follow, some of the significant blocks in figure 1 are discussed in some detail.

The technical interchange meetings proved an invaluable extension of the technical monitoring function for tracking the developmental status and for planning future emphasis on the existing problems. The guidance system operations plan (GSOP) constituted the primary control documentation, and the GSOP review function provided the checkpoint for assuring that detailed requirements would be implemented.

The bit-by-bit simulations were digital programs that caused a general-purpose computer to function like the airborne computer hardware and provided detailed diagnostic information on flight software coding problems. The hybrid simulations coupled actual flight-type hardware into computer simulations of spacecraft dynamics. The independent design and software verification consisted primarily of simulations performed at the NASA Manned Spacecraft Center (MSC) and the spacecraft contractor laboratories, where both bit-by-bit and hybrid simulations were used.

After the software had been developed sufficiently that the schedule required configuration control, software milestone meetings gradually replaced technical-interchange meetings. Beyond this point, design-improvement changes were allowed only if justifiable to the software configuration control board.

Functional Requirements

A comprehensive and realistic set of quantitative requirements for any complex subsystem is difficult to list at the onset of a developmental program. This difficulty is caused primarily by an initial lack of knowledge concerning all the interactions between the particular subsystem and the total system environment with which the subsystem must ultimately be compatible. However, the quantitative requirements that gradually evolve during the developmental process are primarily design details that

can be considered subordinate to a set of more basic functional requirements. The hierarchy of functional requirements for the Apollo command module (CM), CSM, CSM/LM, LM/CSM, LM-descent, and LM-ascent control systems are as follows.

1. Maintenance of attitude orientation sufficiently near a commanded orientation during accumulation of velocity from a rocket-engine burn, during coasting flight, and during entry into the earth atmosphere
2. Provision of acceptable thrust-impulse economy. This requirement implies operation within the constraints of the total reaction-control propellant available
3. Provision of an attitude-control loop responsive enough to prevent dynamic interference with the guidance scheme
4. Provision of compensation for rigid-body motion, including time-varying gains to provide for inertia changes as a result of propellant consumption or stage-configuration changes
5. Provision of compensation for thrust-vector misalignment and other disturbance torques
6. Provision of compensation for propellant-sloshing and structural-bending resonances, a requirement implying operation within the constraints of structural load limits
7. Provision of stability margins of sufficient magnitude to guarantee the stability of guidance long-period modes, attitude short-period modes, propellant-sloshing resonances, and structural-bending resonances
8. Provision of acceptable performance in convergence from large transients and in the steady-state limit cycles
9. Maintenance of attitude errors and vehicle rates within bounds compatible with effective crew monitoring of successful performance
10. Provision of a capability for switching to a redundant controller in case of a malfunction

The terms "sufficiently near," "acceptable," and so on were not specified quantitatively in all cases. However, the requirements were closely monitored by cognizant NASA engineers during the developmental phase, and requirements for design improvements were negotiated on several occasions.

Preflight Testing and Verification

The primary objectives of preflight testing were to validate the control-system performance during nominal, off-nominal, failure, and mission-related conditions. The types of simulation facilities used included the engineering digital simulator, the interpretive computer simulator (ICS), and the hybrid simulator.

Engineering digital simulators were used during the initial developmental and modification phases to provide dynamic validation and performance evaluation of the functional design under a broad spectrum of mission conditions. The bit-by-bit ICS modeled the detailed computer characteristics and was used to verify the proper translation of the functional designs into software programs that were compatible with the flight-computer characteristics and other interfacing software programs. Parameter studies associated with off-nominal performance generally cannot be run efficiently on the ICS. However, nominal performance verification was conducted on a mission-by-mission basis. The hybrid simulators were used to verify hardware/software interfaces and to provide overall system-performance validation. With respect to the digital autopilots, both the design-validation and mission-verification testing programs were conducted on hybrid simulators. All flight control modes were tested for both nominal and off-nominal conditions. The types of testing are listed in table I.

TABLE I. - DIGITAL AUTOPILOT TESTING MATRIX

Flight mode	Test condition
Coasting	Limit-cycle characteristics Transient behavior (acquisition and recovery) Manual/automatic attitude maneuver Translation/ullage maneuver Off-nominal performance (jet failures, mass mismatches, restarts, thrust misalignments)
Powered	Stability/controllability characteristics Initial start transient performance (misalignment) Transient behavior (acquisition and recovery) Limit-cycle characteristics Off-nominal performance (actuator degradations, failures, mass mismatches, restarts, degraded thrust)

In addition to the testing outlined in table I, special test conditions included (1) control-mode-switching transients, (2) structural loads and structural-stability evaluation with jet failures and off-nominal bending data, and (3) powered-flight transients for hardover actuator failures. The control systems were subjected to realistic flight-environment testing, which included evaluation of the effects of propellant shift and engine-mount compliance. Throughout the Apollo Program, verification-run matrices for both ICS and hybrid-simulator testing were extensive.

COMMAND-SERVICE MODULE DIGITAL AUTOPILOT DEVELOPMENT

The CSM DAP provides stabilization and control of the CSM during both coasting and powered flight. The CM DAP, on the other hand, is required to provide stabilization and control of the CM only during coasting flight. The development of the specific designs is reviewed by first outlining the vehicle design constraints and discussing the design philosophy. Later, brief design descriptions are given to provide general background; then significant design-development problems are discussed.

Vehicle Design Constraints

The basic Apollo vehicle configuration and control-system hardware configuration were established in advance of the bulk of the detailed digital-control-system design work. This early development of configurations forced the evolution of detailed digital-control-system designs that were subject to the constraints imposed by (1) vehicle structure, (2) propulsive devices, (3) total reaction-control propellant, (4) service module engine actuators, and (5) digital computer. Each constraint will be discussed in the sections that follow.

Vehicle structure. - The primary constraints imposed by the structure were the potential instabilities of propellant-sloshing and flexible-body resonances during main-engine operation and a potential for excessive structural loads because of reaction-jet torque pulses. The excessive structural loads were primarily the result of a failure-mode situation whereby excessive energy could be transferred to the structural resonances without a classical control-system instability. This condition was due to the fact that the disturbance torque present with a failed-on reaction-jet thruster was equal to the net restoring torque present when the control torque was applied. The resulting thruster duty cycle would be 0.5 and could conceivably contain high energy at the bending-resonance frequency.

Propulsive devices. - The reaction-jet thrusters had a duty-cycle constraint, a minimum tolerable "ontime" command of 14 milliseconds, and sizable pulse-on and pulse-off delays. The service propulsion system (SPS) had four long, cylindrical propellant tanks arranged so that considerable lateral movement of the vehicle center of mass occurred during a main-engine burn. Also, the engine throat was cooled by an ablative material that could erode in such a manner that thrust-vector alignment relative to the engine centerline could change with time during a burn. These SPS problems were both jerk disturbances. Several engine, engine-actuator, and structural characteristics contributed to a substantial uncertainty in the initial thrust-vector alignment

relative to the total vehicle center of mass. This alinement uncertainty produced an initial acceleration disturbance.

Total reaction-control propellant. - The total impulse capacity of the reaction-control propellants caused initial concern to the control-system designers because no method for accurately predicting analytically the effects of zero-g propellant-slosh dynamics existed. For example, the residual kinetic energy in the SPS propellants after main-engine shutdown might require substantial amounts of reaction-control-propellant expenditure to damp attitude motions resulting from the propellant forces on the tank walls. Fortunately, flight test results indicated that this expenditure was of no serious consequence, probably because attitude-error perturbations were filtered by the control-logic dead band. Considerable effort was expended on zero-g slosh models, but no tenable models evolved that were substantiated by Apollo flight results. The primary problem is caused by the long time constants involved and the many degrees of freedom that would be required for a realistic model, which would make the results almost entirely dependent on numerous initial conditions that cannot be determined.

Service module engine actuators. - The SPS engine actuators were electromechanical devices (motors, gears, and jackscrews) that had torque and rate limits. In addition, the engine was not mass balanced, and therefore a potential dynamic interaction existed between structural bending and TVC. This interaction is a type of dynamic coupling referred to in reference 5 as "dog-wags-tail."

Digital computer. - Fortunately, the computer capacity for control logic and computation rates did not result in significant constraints on the control-system design, nor did the data quantization of the digital-to-analog converters produce significant limit cycling. The limit cycling that did occur during TVC was primarily due to a quantization-induced limit cycle in the guidance loop.

Design Philosophy

As implied by previous statements, an original philosophy was that optimum reaction control system (RCS) propellant-consumption efficiency must be maintained at all times because of a limited rotational-impulse budget. As the total RCS propellant budget matured, it was realized that the percentage of the total budget being allocated to attitude hold was small, because most of the propellants were required for nominal translations and maneuvers and for emergency reserves. Therefore, a 20-percent-above-optimum-usage rate for attitude hold would represent a penalty that barely affected the noise level of budgeting uncertainty. Therefore, the optimum solutions could be detuned to make them more tolerant to plant-modeling uncertainty.

Another design philosophy used with success in the development of the reaction-jet attitude-control-system designs for the Apollo Program was plant modeling. This philosophy was neither the model-following adaptive technique nor cancellation compensation to change the plant-response characteristics, but simply the technique of basing control actions on a predicted plant response. This technique provided valuable lead time in the control loop, and this added time helped wash out the delays associated with data sampling, transport lags, and state estimation. This technique is valuable for reducing the control-sample frequency and attendant computer-load factor. Furthermore,

this technique is essentially trouble-free if the prediction is updated during each control cycle and if the control cycle is small compared to significant plant and disturbance-time constants. Each update tends to wash out any errors made in the previous prediction, so that the prediction error diminishes with the control-loop error, and the final prediction error is only that error associated with a prediction for a control action lasting less than one control sample period.

Early planning budgeted approximately 20 percent of computer time to the control task; and, like most things that are budgeted in advance, the control computations grew to consume the budget at some points in the mission. This total-computer-budget usage occurred only in portions of powered flight, and, even in these cases, economies could have been effected, if necessary. Some reserve capacity was used to build in additional flexibility as a hedge against changes in requirements. An example of this reserve capacity is the early use of a seventh-order compensation filter (later changed to sixth order) for control during CSM/LM-docked SPS burns, when studies had shown that a fourth-order filter would suffice for the plant dynamics known at that time. The filter coefficients were placed in the erasable memory, which permitted use of lower order filters, if desired.

Another area that could be called a design philosophy, but in reality concerned a specific design requirement, was the frequency-band-pass requirement for the pitch and yaw attitude control systems during a CSM/LM-docked SPS burn. This requirement should be as high as possible and consistent with an adequate high-frequency stability margin, but disagreement occurred on what margin was adequate. A conservative approach was taken initially, in which it was attempted to make the passband as low as possible and consistent with adequate low-frequency performance. Later, after an inflight test increased confidence in the structural characteristics and after several actual flight demonstrations, the passband was increased, and the higher performance gains were used for the lunar landing mission.

Design Description

Because the Apollo mission required a modular spacecraft that operated in many different configurations, several separate and unique control problems occurred. To compress developmental schedules, the contractor assigned different individuals the responsibilities of designing mission-control programs that would control (1) command module computer (CMC) backup to a Saturn rocket boost, (2) CSM or CSM/LM coast, (3) CSM (alone) powered, (4) CSM/LM powered, and (5) CM preentry and entry. Therefore, each of these phases can be viewed as independent control-system designs and implementations, because little program coding was shared by these control routines. The following material briefly describes the control system for each of these mission phases.

Command module computer backup to a Saturn rocket boost. - This mode performed the function of closing the Saturn control-system attitude-error loop in the CSM and permitted the astronaut to perform a guided function manually. This procedure provides a backup to the inertial platform in the Saturn instrument unit. The desired attitude is computed from a polynomial function of time during first-stage flight, and the resulting attitude errors are sent to the flight control computer in the instrument unit. The astronaut compares the CMC-displayed quantities of velocity magnitude,

altitude, and altitude rate against predetermined desired values and issues rate commands to the Saturn instrument unit by hand-controller motion.

Command-service module or CSM/LM coast. - The CSM coasting-flight digital autopilots use a digital filter algorithm to estimate attitude rate from noisy platform gimbal-angle measurements and estimates of applied control effectiveness. A phase-plane switching logic is used that contains hysteresis. Control torques are available from four quadrants (quads) of four reaction-jet thrusters. The thrusters in each quad provide positive and negative forces that lie approximately along lines parallel with two of the spacecraft body axes. Either pitch-plane quads or yaw-plane quads can be specified for roll control or X-translation commands (or both). Furthermore, any combination of quads can be specified as nonusable, and the DAP will attempt to use only the remaining usable quads. Attitude dead bands and maneuver rates can be changed through either astronaut or ground-station commands. Hand-controller inputs are given priority over automatic control and are "hot" at all times. A specified maneuver rate results from hand-controller action. Simultaneous three-axis maneuvers result when the automatic-maneuver routine is used. This routine rotates the spacecraft about the single axis normal to the plane including the current longitudinal axis and the longitudinal axis at the new desired orientation for all conditions except gimbal-lock avoidance. The phase-plane switching logic used during an automatic maneuver is the same as that used for attitude hold. The position constraint causes the actual rate to limit-cycle about the desired rate. For the CSM/coasting-flight autopilot, the control sample period is 100 milliseconds.

Command-service module (alone) powered. - The digital compensation filter to CSM powered flight is third order, with a zero-frequency gain that is initialized from computer knowledge of control effectiveness. During the SPS engine burn, the gain is decremented at 10-second intervals to account for the increase in control power resulting from reduced propellant load. The SPS engine gimbal-drive system was tested before a burn and positioned at the attitude of the expected trim position. The control sample period is 40 milliseconds.

Command-service module/lunar module powered. - The CSM/LM attitude controller for powered flight was originally seventh order (later changed to sixth order). The gain is also initialized and decremented every 10 seconds, as discussed for the CSM. In the original design concept, a burn would begin with the use of a control sample period of 40 milliseconds; but, after 7 seconds of burn time, the gain was changed, and the control sample period was increased to 80 milliseconds. The 40-millisecond DAP (termed MOD 40) was originally intended for use throughout the burn, but switching to the 80-millisecond DAP (termed MOD 80) increased the high-frequency stability margins.

Another feature of the CSM and CSM/LM powered-flight digital autopilots is a thrust-alignment tracking loop, which computes the average gimbal-position command and returns it to gimbal command as a bias. This feature prevents initial-thrust-alignment error from producing a bias in the attitude error, which would prohibit use of attitude error as an indicator of proper system performance (to give the astronauts real-time monitoring information).

Command module preentry and entry. - The CM digital autopilots exercised attitude control by selecting jets for firing from a set of six (or 12) reaction-control jets

on the CM. The CM was aerodynamically stable about the pitch and yaw axes during entry so that the entry DAP had to perform rate damping only for these two axes. The entry guidance supplied commands for rotating the lift vector about the velocity vector; therefore, transformations were performed to provide attitude control about the roll-stability axis. The exoatmospheric phase of entry DAP operation provided attitude control about all three axes, and mode switching occurred as a function of measured acceleration.

Design-Development Problems

Service propulsion system engine-actuator performance. - During the summer of 1965, the SPS engine actuator was not meeting performance specifications. The electromechanical (motor, gears, and jackscrews) actuator was supposed to be able to slew the SPS engine position at a steady-state rate of 0.36 rad/sec, but the first units delivered were capable of only 0.2 rad/sec. Severe heating problems resulted in a requirement to reduce the maximum actuator rate even further or to subject the actuator to major redesign. This added requirement caused NASA and contractor personnel to try to find the minimum engine-actuator rate that would be compatible with Apollo attitude-control requirements.

The factors necessary to establish this requirement were as follows.

1. The attitude-control passband requirement (set by the guidance requirement)
2. Flexible-body and propellant-slosh stability (To guarantee a stability margin by an analytical process, significant nonlinearities had to be avoided.)
3. Rigid-body recovery from large transients
4. Sound system-design principles
5. Cost and schedule factors

Item 1 established only mild requirements for fast actuator response. Spacecraft guidance is primarily a low-frequency phenomenon and is fairly insensitive to temporary saturation of inner-loop signals, unless these lead to inner-loop instabilities.

The requirement for linear response to structural oscillations is more difficult to ascertain but can be estimated by assuming that only one or two of the low-order modes are excited to near the point of structural failure. This requirement is then interpreted in terms of actuator command as a function of attitude-sensor location and the dynamic gain in the controller. Because of the attenuation being proposed for the CSM/LM TVC DAP at frequencies near and above the first-mode resonance, the actuator command was essentially independent of the bending motions. In cases where bending or propellant-slosh resonances are phase stabilized, it is important that the maximum expected excitation not saturate TVC, because a large signal instability can result.

Item 3 proved the major constraint on the SPS-actuator rate requirement. Analog simulations of CSM attitude control, using the backup control system and an actuator

rate limit of 0.1 rad/sec, demonstrated a large signal instability; that is, vehicle attitude would not converge properly from large initial tumbling rates. Additional studies indicated that an actuator rate capability of 0.15 rad/sec was desirable to avoid changes in the control-system hardware.

Item 4 stipulates that a designer must not overconstrain the design problem. This requirement implies conservatism when constraints are being adjusted, because all possible ramifications are difficult to predict. For the case under discussion, the cost and schedule implications (payoff function) were heavily weighted toward accepting some risk of later difficulty in other areas (e.g., incurring a requirement for future changes to hardware in the backup control-system hardware). As a result of these conflicting factors, the minimum specification finally agreed upon by NASA and contractor personnel (0.1 rad/sec) was not conservative; however, actual delivered hardware performed considerably better than this minimum requirement. The final hardware was capable of slewing the SPS engine at approximately 0.15 rad/sec, and backup control-system gains had been lowered so that an adequately large signal-stability margin was obtained.

Thrust-vector-control filter mechanization. - The single-axis attitude-control loop compensation for the LM attached to the CSM was synthesized by the guidance and control contractor to be a seventh-order filter. The first attempt at a digital mechanization of these dynamics was a straightforward, recursive computation performed by shifting all data by one step each computation cycle before multiplying by the appropriate coefficients and summing. This process is shown schematically in figure 2, where z^{-1} represents a delay of one computational interval. The gains are simply the coefficients of the numerator and denominator polynomials of the z-transform version of the filter dynamics (ref. 6). The reference designates this method as "direct programing."

The straightforward approach did not work satisfactorily for the reasons discussed in detail in reference 7. Briefly, the cause was a result of repetitive truncation error in the computational process, because each intermediate result had to be stored in a finite word-length register. The effect was to shift the effective location of some filter poles into the right half plane, causing the filter to be unstable. After some study of this problem by the guidance and control contractor, the second approach discussed in reference 4, called "iterative programing," was adopted. This technique separates the factored form of the filter transfer function into a product of two or more (three in the case

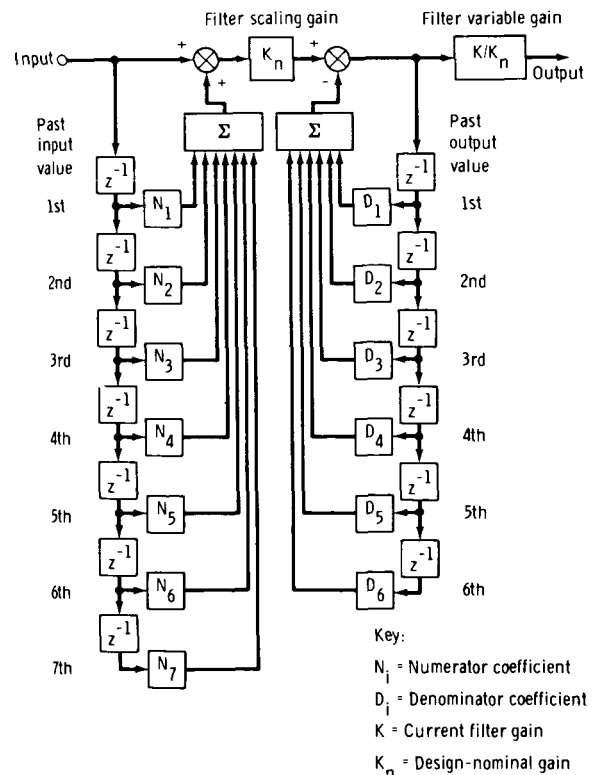


Figure 2. - Original CSM/LM filter mechanization.

under discussion) smaller filters, implements these, and allows the output of the first to be the input of the second, and so on. The final Apollo TVC DAP filter mechanization is discussed in references 1 and 2; the earlier version, in reference 3.

Plant-model deficiencies. - A straightforward synthesis of filter dynamics necessary for stable operation during an SPS engine burn would have been possible only if complete confidence in the mathematical model of the plant dynamics had existed. This model includes not only the mathematical formulation of the model, but also the programming of that model into the digital computer simulations being used to verify the filter design. Generally, the required mathematical formulations have existed for some time, but a problem exists because of the multitude of variables and a desire to eliminate unnecessary coupling terms to reduce simulator complexity. This reduction results in simplified equation sets being published in various reports, and whether or not a specified reduced set is applicable to a specified situation is impossible to determine unless the missing terms are available for inspection and evaluation. Two examples of this type of problem that arose during the development of the Apollo CSM TVC DAP concerned the stability of propellant-slosh resonances and a potential instability in the airframe and engine dynamics.

Early evaluation of the single-plane dynamics of the Apollo CSM/LM, performed independently by MSC personnel and by a support contractor, consistently disagreed with the guidance and control contractor results regarding stability of propellant-slosh resonances. The guidance and control contractor results showed the CSM propellant sloshing to be the more severe but with an acceptable stability margin, whereas the other analyses indicated the LM propellant sloshing to be more severe and to have unacceptable stability margins. A thorough investigation of the models showed that coupling with the translational degree of freedom produced this disparity of results. The MSC and support contractor models had both been derived under the assumption that this coupling was small enough to be ignored for evaluation of high-frequency rotational dynamics.

Another early plant-model deficiency was the omission of coupling terms that can give rise to a type of flutter instability. The MSC studies had shown that the combined airframe and actuator dynamics could be unstable without any attitude-controller inputs to the actuator. The energy source was the thrust force, and the feedback path was inertial forces on the engine because of structural oscillations, causing engine-gimbaling motion. This effect has been termed "dog-wags-tail" as opposed to the other airframe and engine dynamic coupling known as "tail-wags-dog." Early contractor models had omitted these terms. This coupling does reduce damping of the first two structural resonances of the Apollo CSM/LM vehicle, but not to the point of instability. If the coupling had reduced the damping to the point of instability, then the adopted DAP filter design approach of providing large attenuation to bending signals (gain stabilization) would not have produced an acceptable design. Active compensation at the first bending-mode frequency would have been required.

Structural-bending data. - The availability of data to be used in the models was the overriding problem of the entire TVC DAP developmental process. The difficulties began in the spring of 1965 when a decision was reached that a combined CSM/LM structural-dynamics test, which had been planned, could be eliminated from the program. This deficiency was later compounded when another decision was made to designate the inline responsibility for definition of combined-vehicle structural dynamics to MSC

personnel instead of to one of the Apollo spacecraft prime contractors. The result was a period of confusion regarding data to define airframe dynamics.

The idea of an inflight dynamics test was proposed in October 1965. The structural-dynamics specialists had hopes of getting some data from this test that could be used to calibrate their analytical models in lieu of the ground testing that had been deleted, but the required additional telemetry measurements were not sanctioned by program management. A limited version of the originally planned test was performed.

An MSC contractor computed the first set of three-dimensional mode shapes for the combined Apollo vehicle in February 1966, 2 years after the initial TVC controller synthesis. This effort predicted a first-mode frequency approximately 1 octave lower than former predictions (roughly 1.0 hertz as opposed to 2.0 hertz). Obviously, this lower prediction seriously degraded the confidence in the quality of any bending data in existence at that time and probably was the most significant factor in getting dynamics ground testing of the combined vehicle back into the program.

Ground testing was reinstated; but, because of scheduling problems, test preparation, test performance, data reduction, and other considerations, the results were not available until December 1968. In the interim, the detailed structural-dynamics modeling being performed by NASA personnel produced the first set of realistic three-dimensional mode shapes, part of which were made available in June 1967. This essential lack of credible bending data for the 18-month period encompassing 1966 and the first half of 1967 precipitated the next two problems.

Change from MOD 40 to MOD 80. - During the summer of 1966, just before publication of the GSOP that would specify the software to be used in the first mission containing a CSM/LM-docked SPS engine burn (ref. 3), the guidance and control contractor decided that additional conservatism regarding the high-frequency stability margin was required. (The estimates of the first bending-mode frequency existing at that time was nearly 1 octave below the nominal first-mode frequency used in the initial DAP filter synthesis.) The design change that was implemented at that point is a good example of the flexibility available to the digital-control-system designer; this flexibility would not be available if the control-system design were being translated into hardware.

The change represented a radical departure from the former concept of a single-stage compensation filter with fixed dynamics and time-varying gain. The original design had the filter coefficients selected for a sample interval of once each 40 milliseconds (MOD 40). The new concept was, after starting the burn with the originally designed filter, to switch gains and sampling interval after a partial nulling of the start transients. A thrust-misalignment-correcting inner loop was also added.

The postswitch filter dynamics, because of a sampling interval of 80 milliseconds (MOD 80), were essentially the same as those of the MOD 40, except that the real frequency where a given shaping effect would occur was now 1 octave lower. This characteristic of a digital filter can be explained by reference to the relationship between real frequency ω and the fictitious w-plane frequency v . In reference 6, this relationship

is shown to be $v = \tan(\omega T/2)$ where T is the sample period in seconds. A specific value of v , say v_0 , is uniquely related to the real frequency only through the sampling interval. Hence

$$v_0 = \tan \frac{\omega_0 T_0}{2} = \tan \left[\frac{1}{2} \left(\frac{\omega_0}{C} \right) (C T_0) \right] \quad (1)$$

implies that the specific real frequency where a given amount of attenuation and phase shift would occur is shifted by the inverse of any scale factor C applied to the sampling interval. This shift is an important flexibility that is available in a digital filter but is not available in filters implemented by hardware components. Hardware filters have an inherent sense of real time based upon the physical and geometrical aspects of the hardware, whereas a digital filter is a mathematic entity and, therefore, is insensitive to real time. For this reason, extremely low frequency filtering can be done digitally in situations where the same job would be impossible to perform with hardware.

Mathematical proof that real-frequency scaling and sample-period scaling are related is beyond the scope of this report. However, a convincing argument is available. Because the digital filter is synthesized in the w -plane (poles and zeros arranged to supply a specific filtering action for a specific value of w -plane frequency v_0) and because the bilinear transformation into the z -plane is not a function of sampling interval, then a given set of digital filter coefficients will provide the same filtering for a given v_0 regardless of the sampling interval involved.

This particular change is listed as one of the TVC DAP developmental problems because the passband of the MOD 40 design (approximately 0.6 rad/sec) was already low, and this change brought the passband down to approximately 0.3 rad/sec. The resulting performance was significantly degraded, and a considerable amount of effort was required in tuning up the guidance and control interface and in trying to contain the attitude errors within reasonable bounds. The result was a NASA directive to the contractor to produce an uprated TVC design based on the assumption that the current estimates of bending frequencies were accurate within ± 30 percent. The low-gain DAP was flown on the Apollo 9 mission, but the uprated design was used on all subsequent missions.

Guidance and control interaction. - The low-gain attitude-control loop first implemented for the CSM TVC DAP would permit large-peak-attitude-error transients and steady-state offset attitudes, resulting from initial-thrust-torque bias. The design change implemented to attack this problem, still with the assumption that the low gain was necessary for conservatism on the high-frequency stability margin, was to add an inner loop that would extract the average commanded engine position and return this information to the engine position command as a bias. The digital filter used for this purpose had large time constants to avoid interaction with the attitude short-period mode, resulting in additional phase lags at frequencies that now coupled with guidance (outer-loop) performance. Also, the long time period required to initialize properly the thrust-position filter (i. e., to find the trim position) reduced the effectiveness of

this approach for limiting the initial attitude excursions. Therefore, large guidance errors that could not be completely removed in short burns could still accumulate during the first few seconds of a burn.

The thrust-misalignment correction loop was not completely ineffective. The main purpose of the correction loop was to remove large steady-state attitude errors required to provide acceleration trim so that an attitude error could be used by a crewman to determine if the maneuver were progressing normally.

Mechanization of the guidance and control interface can produce, or aggravate, guidance and control interaction problems. The evolution of the interface used in the Apollo guidance and control system is shown in figures 3 to 5. The guidance law produces a desired acceleration-vector rate command that is proportional to the sensed acceleration-direction error. The steady-state acceleration direction is very nearly the attitude of the vehicle longitudinal axis, but this case is not true when control torques are being commanded. Therefore, a requirement exists to provide stable-attitude inner-loop control by using attitude-position feedback.

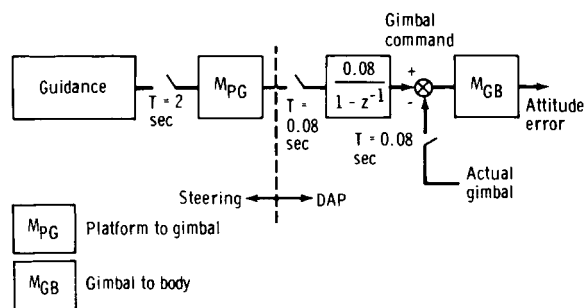
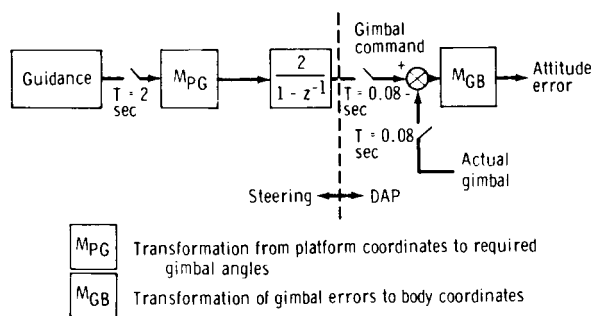


Figure 3. - Original guidance and control interface.

Figure 4. - Alternate guidance and control interface (ramp steering).

The first interface mechanization performed a simple digital integration of the guidance-rate command (fig. 3) to form a vehicle-attitude command. This integration caused step changes in the attitude command during each guidance-computation period (2.0 seconds), and these changes caused significant interaction with propellant-slosh dynamics.

The next approach was to smooth the attitude commands by performing the integration at the DAP sample rate (fig. 4). This change was effective in reducing the slosh excitation, but some concern arose about the fact that gimbal-angle errors were being subjected to a coordinate transformation. This transformation of incremental attitude errors cast some doubt on large-signal stability. Subsequently, the coordinate-transformation problem was solved, as shown in figure 5, by transforming guidance-rate command to body coordinates where the command is compared with a derived body rate before the integration to produce body-attitude error.

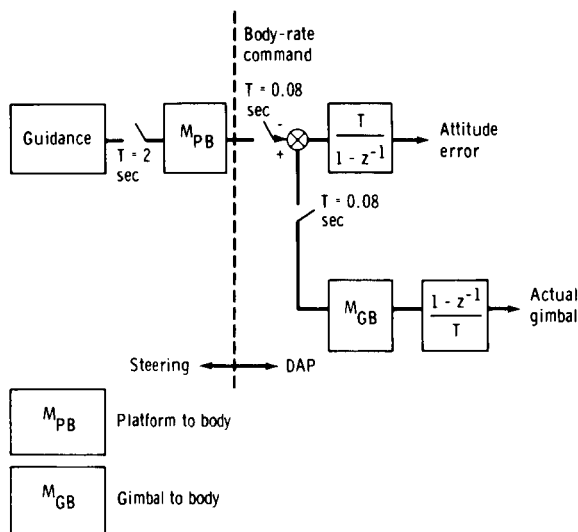


Figure 5. - Final guidance and control interface (rate steering).

Initialization. - No specific aspect of the CSM DAP initialization caused any hardships that would be classed as a problem, but enough minor changes were implemented in this area that the subject is worthy of mention. In the software system, just as in the case of hardware systems, the engineering attention required of the interface between two components may sometimes exceed the attention required for the selection or design of the components. Therefore, much of the software associated with the digital autopilots performed the following functions.

1. Permitted crew selection of operating mode
2. Permitted crew selection of reaction-jet quads (e. g., jet quads to use for roll control)
3. Permitted crew selection of maneuvering rate
4. Permitted crew selection of attitude dead bands
5. Permitted crew monitoring of mass and engine-trim initialization
6. Permitted crew change of mass and engine-trim initialization if not satisfied with quantities in the computer
7. Computed control effectiveness based upon stored functions of vehicle mass properties, mass-properties initialization data, and the available information on which reaction-jet motors can be used
8. Permitted integration of the DAP into a complete maneuver program
9. Provided computer restart protection

Items 1 to 6 correspond closely with the proper configuration of several switches on a control panel in the cockpit of a conventional aircraft, with the following exception. The flight control systems of aircraft (and previous spacecraft) seldom underwent the complete metamorphosis required of the Apollo CSM DAP for (1) coasting-flight control of the Saturn IVB using Saturn IVB thrusters, (2) coasting-flight control of the CSM, (3) powered-flight control of the CSM, (4) coasting-flight control of the CSM/LM, (5) powered-flight control of the CSM/LM, and (6) entry of the CM. This problem was solved through the use of a special software routine designated RO3. Verb 48 would call this routine, which would successively cycle through nouns 46, 47, and 48, each displaying either coded octal numbers that represented switch settings or decimal numbers that represented CSM mass, LM mass, or SPS-engine-gimbal pitch- and yaw-trim

settings. As these numbers were displayed, they could be changed simply by punching new numbers into the keyboard before proceeding to the next noun.

Knowledge of the vehicle mass (item 7) was necessary because the vehicle mass was used to compute vehicle inertia and control effectiveness. The Apollo digital autopilots successfully used knowledge of the expected vehicle response to a control action to provide valuable lead information.

Item 8 on the list of peripheral software entails an external program automatically configuring the DAP to select the proper operating mode and to provide the necessary initialization data and the correct sequence of commands. For example, the SPS thrusting program (P40) tasks that affected DAP operations included (1) performance of SPS engine gimbaling check, (2) selection of coasting-flight autopilot with narrow-dead-band attitude hold as engine ignition time approaches, (3) initiation of ullage, (4) command of SPS thrust-on, (5) termination of ullage, (6) switch from reaction-jet control to TVC, (7) termination of thrust, and (8) reconfiguration of DAP from TVC to wide-dead-band reaction-jet control.

One design change in the initialization process that was made late in the program should be discussed. Originally, the TVC DAP would initialize with zero-attitude errors, regardless of the true attitude at the time. In simulations, the disturbance torque during ullage was observed always to cause the attitude error at thrust-on to be on the order of 1.0° . The DAP initialization would wash out this information, and the initial thrusting directions would be in error.

The guidance loop would eventually compensate if given sufficient time, but velocity errors would result for short-duration burns. The faster acting attitude-control loop would prevent these velocity errors if the TVC DAP were to be initialized with the more correct attitude reference. A change was made to initialize the TVC DAP with the non-zero-attitude errors, if the configuration were CSM alone, at the expense of accepting an initial engine-gimbaling and attitude transient. The engine-gimbaling transient was too severe for the larger inertia configuration of the CSM/LM without a technique for ramping the error in slowly. Therefore, rather than risk an unnecessary software impact, the modification was made applicable to the one configuration only.

Roll divergence. - During independent design-verification testing of off-nominal conditions, the roll attitude was discovered to diverge in the presence of a failed-on roll thruster during an SPS velocity-change maneuver. This problem was caused by a combination of two separate design features. First, the time at which a new roll-control pulse could follow a roll of the same polarity was intentionally constrained to be greater than one control sample period (0.5 second). This constraint avoided the use of a high-frequency pulse train of small pulses to hold attitude against a disturbance torque. The constraint worked quite well for smaller disturbances; but, in the presence of the large angular-acceleration disturbance that existed with a constant 100-pound thrust acting through a 7-foot lever arm and acting on the small roll inertia (as small as $15\,000\text{ slug-ft}^2$), the attitude-error limit in the control law was too small to prevent divergence.

The mechanics of this phenomenon were as follows. Restoring torque thrusters would fire until the divergence rate was nulled, in addition to an incremental amount of

thrust ontime proportional to attitude error. During the forced delay, at which time the control torque could not be reapplied, the disturbance torque would add angular momentum in the direction of divergence. If the restoring-control momentum impulse could have possibly continued to grow with the increasing attitude error, a steady-state limit-cycle condition would have been reached. Instead, the original design had a limit point beyond which the increasing attitude error no longer affected the terminal restoring momentum. Therefore, with a disturbance torque of sufficient magnitude to null the terminal restoring momentum (i. e., to revert to zero-attitude rate) in less than half the jets-off delay, the attitude would diverge. This problem was solved by adding hysteresis to the phase-plane logic, so that restoring momentum would continue to grow with increasing attitude error, but the desirable feature of partially nulling large convergence rates was retained.

Command-service module RCS DAP rate-estimator gains. - Just before the launch of the Apollo 7 mission, the CSM deorbit burn was discovered to be incurring a velocity error as a result of an initial thrusting error caused by the TVC DAP attitude error initialization process. (The DAP ignored attitude error existing at SPS thrust initiation.) The four-jet ullage burn that preceded the SPS engine ignition in these simulations produced approximately 200 ft-lb of disturbance torque about the negative pitch axis of the CSM vehicle at a time when the pitch inertia was lowest (SPS propellants near depletion). The resulting disturbance acceleration would produce sufficient bias in the DAP estimated rate (with rate information being derived in a digital filter from noisy angular measurements) to degrade significantly the attitude-hold capability. The attitude dead band was 0.5° , but the attitude errors at ignition were on the order of 1.2° . An investigation showed that an insufficient margin at this flight condition existed because slightly higher disturbance torques would have caused attitude divergence. This problem was solved by using a two-jet ullage for the Apollo 7 deorbit ullage maneuver and by raising the gains in the rate-estimation filter to reduce the rate bias resulting from a given acceleration disturbance.

Command-service module DAP/hand-controller interface. - Another significant design change that was implemented in the CSM coasting-flight autopilot involved the switching logic used in response to rotational hand controller (RHC) motions. The only information available to the CMC regarding hand-controller status is 6 bits of an input channel such that proportional rate control (vehicle-rate command proportional to amount of hand-controller deflection) was not possible.

The original design simply fired jets to null a rate error, defined by the difference between the estimated vehicle rate and a preselected maneuver rate, until the rate error was less than a small value. This small value was computed as a function of vehicle inertia, but the value corresponded closely with the vehicle rate-change that would be expected for a minimum impulse of control torque.

The disadvantages of this design were twofold. First, for control about a light inertia axis, where the minimum impulse rate could be nearly as large as the lowest desired maneuver rate (0.05 deg/sec), performance was unacceptable; second, unnecessary time delays and design complexity were involved in using a different switching logic when the hand controller was out of detent. The solution was to form the rate error as before, integrate the rate errors to form an attitude error, and then use the attitude and rate errors in the same phase-plane switching logic that was used for attitude-hold and automatic maneuvers.

Entry. - The CM DAP used during the final phase of the Apollo missions was relatively uncomplicated and trouble-free. The CM was aerodynamically stable in pitch and yaw deviations from the stability axes, and because the guidance scheme did not require alteration of the trim angle of attack, only the rate damping needed to be provided about these axes. The guidance scheme commanded rotations about the velocity vector, and because of the low guidance gains, guidance-commanded maneuvers were infrequent. This lack of maneuvers implies that the roll attitude-control requirements were not critical.

This statement does not imply that the entry DAP design was trivial. The design did have several modes and interfaced with several external programs. This particular DAP, like most of the others, was developed independently by separate designers. As such, the entry DAP provides another independent data point on a special-purpose digital control system that was used on the Apollo Program and that did the job well.

One aspect of the entry DAP development emphasizes the need for close scrutiny of simulator models. The amount of reaction-control propellant consumed during entry of the Apollo 7 CM was observed to be approximately 50 percent higher than predicted. This overconsumption was later shown to be the result of an insufficient atmospheric model in the simulator environment. Inclusion of a realistic wind profile in the simulator produced results that more closely matched the flight data.

LUNAR MODULE DIGITAL AUTOPILOT DEVELOPMENT

The LM DAP provides stabilization and control of the LM during both coasting and powered flight in three configurations — descent, ascent, and docked with the CSM. During the preliminary spacecraft-design phase, many fundamental decisions were made that impacted the control-system design. For the LM, three basic propulsion-force and torque systems were established — RCS, descent propulsion system (DPS), and ascent propulsion system (APS). Characteristics that influenced the control task included the type of actuation system, the geometrical location and number of thrusters or jets, and the type of thrust-variation system.

During DPS-powered flight, the LM DAP design provides yaw control with the RCS jets and pitch/roll attitude control with a combination of the RCS and the gimbal-trim system (GTS). The geometrical location of the RCS jets is significant in establishing the fundamental design approach.

During APS-powered flight, the primary purpose of the RCS is to provide attitude stabilization and control. However, whenever feasible, a design requirement exists to fire only the upward-thrusting RCS jets so that velocity changes from the main engine are augmented. Because the APS is a nongimbaled, fixed-throttle system, the RCS control laws associated with this mode must accommodate large time-variant disturbance torques.

The RCS provides automatic or manual rotation and small translation control for all LM configurations during coasting flight. For coasting flight, the design problem is characterized by the presence of low-disturbance torques (except for an RCS jet-on failure). The LM DAP was also required to control the entire Apollo spacecraft with the LM docked to the CSM, while performing a velocity-change maneuver using the LM DPS.

Vehicle Design Constraints

Numerous constraints influenced the LM DAP design, the most predominant class of which related to weight restrictions associated with the lunar landing program. Weight considerations constrained the system design through (1) structural characteristics of the LM/CSM — that is, structural-bending modes are significant; (2) propellant-sloshing dynamics — that is, slosh baffles were removed early in the program; and (3) unbalanced couple control — that is, requirements were established for APS-powered flight.

Another class of constraints, generally identified late in the design-development phase, involved restrictions on RCS-jet firing. These restrictions included (1) duty-cycle constraints — propulsion instabilities; (2) exhaust-contamination constraints — particles on windows or optics; and (3) thermal constraints — RCS-exhaust-plume impingement heating.

A third class of constraints that influenced the design problem was associated with propulsion-system characteristics. The slow-speed characteristic of the DPS trim-gimbal actuator was established for crew safety to avoid hardover actuator failures during powered descent of the LM. A special gear drive was developed to restrict the trim-gimbal-drive rate of ± 0.2 deg/sec. Unlike the classical actuator used for the CSM TVC system, the DPS actuator cannot fail at a higher drive rate. A second propulsion-system constraint was associated with the decision to have a nongimbaled APS engine. This decision imposed significant limits on the allowable center-of-mass characteristics during powered ascent flight. Unfortunately, effective control of the mass center of gravity is difficult in a program such as Apollo. Another propulsion-system constraint was associated with the decision to locate the RCS jets 45° from the body axes. This geometry significantly influenced the interaction between the RCS mode (control axes) and the GTS mode (body axes) during DPS-powered flight.

The fourth class of constraints that impacted the design problem included computer-oriented restrictions. The LM guidance computer (LGC) is limited in both fixed and erasable memories; in addition, definite timing restrictions are placed upon the programs required to provide the control functions.

Design Philosophy

An important question with respect to design philosophy was how to use the inherent flexibility associated with a spacecraft digital computer. This question was significant because the LM DAP represented a first-generation DAP design development. Emphasis was placed upon using digital capabilities such as logic (switching and branching), nonlinear computations, and function generation.

The concept of performance margin was an area of design philosophy that influenced DAP development. This concept emphasized the principle that the acceptability of the design should be based upon performance of the system during extreme, but required, degraded conditions. Acceptable performance during off-nominal conditions, such as single undetected jet failures and large control-effectiveness uncertainties (thrust magnitude, inertia properties, thrust misalignment, actuator drive rates, etc.), was difficult to achieve. The performance-margin concept identified a general trade-off

between lowering the requirements for nominal system performance to achieve acceptable performance during degraded conditions, while maintaining the highest possible measure of nominal performance.

The philosophy of providing system-design flexibility to accommodate developmental problems or future contingencies was related to the concept of performance margin. An example was the guideline to stabilize the GTS control loop independently of the RCS control loop. Three years after this design was initiated, additional thermal constraints (which essentially restricted all X-axis RCS-jet firings) were identified for the LM/CSM configuration during DPS-powered flight. The design was flexible enough to accommodate this restriction without significantly affecting the Apollo Program.

The final design philosophy for LM DAP development was associated with the RCS propellant performance requirements. Design emphasis for the achievement of efficient propellant usage should be placed upon those control functions that require the largest percentage of the entire propellant budget over a complete mission profile. For the LM DAP, these control functions included manual translations, manual and automatic attitude maneuvers, and maneuvers associated with powered-flight guidance. In accordance with this philosophy, the importance of efficient RCS propellant performance for coasting-flight attitude control could be downgraded. Therefore, it could be asked why the design complexity and associated verification cost should be increased to save 20 percent of the performance on an item that uses 5 percent of the total mission propellant. A definite trade-off exists between design complexity and performance-improvement payoff.

Design Description

The baseline LM DAP design that was initially flight-tested on the manned Apollo 9 mission is partially described to provide a background for a discussion on developmental problems. Primary emphasis is given to reviewing the design of the DPS-powered-flight mode. This mode is more complicated than the coasting-flight modes and provides significant insight into the design principles.

A block diagram of the DPS-powered-flight automatic control is given in figure 6. The major design elements include a G&N outer loop, a mass-properties and control-law parameter routine, an attitude state estimator, RCS control laws, jet-selection logic, and trim-gimbal control laws. A timing-and-control-logic interaction between the RCS control and the GTS control is also required.

The integrated G&N outer loop interfaces with the LM DAP through a steering routine. A description of the attitude variables, which the steering routine inputs to the control system, is given in reference 4.

The basic input/output diagram for the attitude state estimator is presented as figure 7. The direct measurements available to the recursive state estimator are the three gimbal angles. The estimator predicts angular attitude, velocity, and bias acceleration. Nonlinear threshold logic is used to reject low-level measurement noise. Angular-acceleration information caused by RCS-jet firings and trim-gimbal activity is an additional set of inputs to the state estimator.

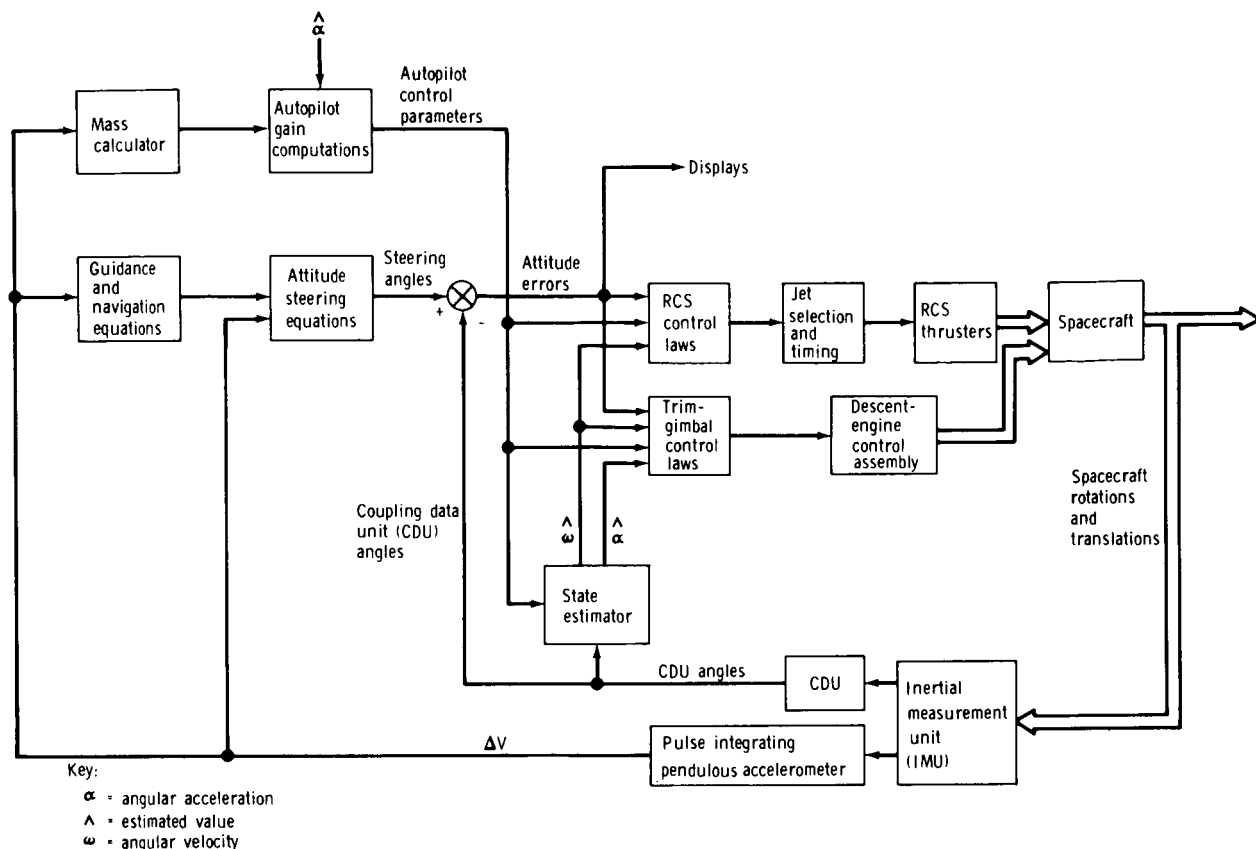


Figure 6. - Powered-flight automatic control.

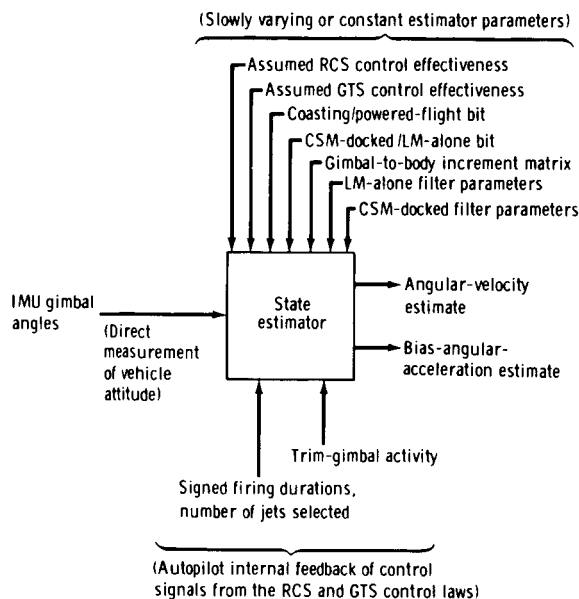


Figure 7. - Input and output variables of the state estimator.

The RCS control laws compute the requirements for rotational impulses by using information based upon attitude phase-plane errors, control effectiveness, and phase-plane targeting logic. These control laws are predictive in nature and are related to the classical two-point boundary value problem. To some extent, this predictive design is inherently sensitive to the uncertainties in control effectiveness and unmodeled disturbances. Angular-error/error-rate phase planes are used to establish jet-firing durations in each control axis. The attitude and rate errors are used to establish the estimated state location in the phase plane. The acceleration inputs required by the RCS control laws include net angular acceleration (jet acceleration plus offset acceleration) and coasting acceleration (acceleration caused by offset alone). The basic shape of the target parabolas and switching-line

parabolas are set by this angular-acceleration information. Additional logic establishes the intersection of these parabolas with $\dot{E} = 0$, based upon crew-selected dead band and the magnitude of the disturbance acceleration.

The phase-plane configuration for the DPS-powered-flight mode is given in figure 8. The control logic is developed by (1) dividing the phase plane into coasting and firing zones, and (2) computing jet-firing durations for each zone to accomplish a desired control action (i. e., null rates, drive-to-target parabolas, and command minimum impulse).

The jet-selection logic combines the required rotational impulses with the commanded translation inputs and selects appropriate jets for control action. Additional information used by the jet-selection logic includes the desired number of jets to be fired and the identification of disabled jets. A normal jet-selection policy in the three translational and three rotational axes applies for nominal conditions. When a required jet has been disabled, an alternate jet-selection policy is used. However, if multiple jet failures for a given control axis have occurred, a computer-program alarm is lighted and an alarm code informs the astronauts that a rotational translation failure exists.




Figure 8. - The phase plane when the LM is in powered ascent or when either of the trim-gimbal-nulling drive times is greater than 2 seconds during powered descent.

Two slow-speed actuators are used to gimbal the descent engine. The control modes developed for commanding these trim actuators are an attitude-control mode and an acceleration-nulling mode.

The GTS control law associated with the attitude-control mode has been developed to be a function of errors in attitude, rate, and acceleration. The control-law equations are basically a modification of a time-optimal solution. The control output commands the sign of the change in angular acceleration.

The GTS control law associated with the acceleration-nulling mode is designed to regulate the offset (disturbance) acceleration from the descent thrust. The primary dynamic environment that causes the disturbance acceleration to change is a shifting center of mass as a result of propellant depletion, DPS-engine-mount compliance, and DPS-engine ablation effects. This control law is structured in the form of a trim-gimbal drive-time equation.

Design Evolution

The history of the LM DAP development will be presented by discussing the design-formulation phase (September 1964 to December 1966), and the SUNBURST flight program phase (December 1966 to August 1967). Where applicable, comparisons will be made to the SUNDANCE baseline design previously discussed. The significant problems encountered will also be discussed.

Apollo 5, the first (unmanned) LM mission, was launched into earth orbit January 22, 1968, and used the SUNBURST flight program. After this mission, a decision was made to simplify the DAP logic, and a significant redesign of the control system was begun, resulting in the SUNDANCE digital program. This design version was flight-tested on the first manned LM mission (Apollo 9), launched March 3, 1969. Subsequent lunar orbital and lunar landing missions (Apollo 10 and 11) were flown with slightly modified SUNDANCE digital autopilots in the LUMINARY flight program series.

SUNBURST DAP. - Many modifications in design philosophy and in control-system implementation occurred during the design-formulation phase of the LM DAP development. This early development culminated in the design that was incorporated into the first LM digital flight program (SUNBURST). Much of the material presented is given a more detailed treatment in reference 4, which provides descriptions of the designs that evolved to solve the developmental problems.

Most of the significant problems associated with the preliminary design were identified as a result of extensive simulation testing. The problem of estimating rate and acceleration when undetected jet failures existed during powered descent proved difficult. Consideration was given to the use of Kalman filters to estimate (from spacecraft dynamics) which of the 16 jets had failed and to adjust the control functions accordingly. A second approach (subsequently implemented) was to use the Kalman filter equations only when the GTS control law was operative or when the RCS jets were inhibited. However, disabling control during powered flight for the time needed to obtain a good filter estimation was considered unacceptable, and this technique was discarded.

A second problem was that convergence to minimum-impulse attitude limit-cycle operation was not achieved by using the initial design. Design-verification studies indicated that this problem was caused by rate-estimation inaccuracies and quantization effects. The phase-plane logic was then modified to improve the performance characteristics.

Significant design problems were identified with respect to the Kalman filter performance. In simulation testing, this estimator was shown to be sensitive to slosh disturbances and large initial conditions. Furthermore, during the DPS start transient, the filter performance exhibited poor convergence because of engine-mount compliance, propellant shift, and initial engine mistrim conditions.

Rate-overshoot performance was indicated to be a problem during coasting-flight maneuvers. The command-maneuver logic did not account for the finite time required to accelerate or decelerate to the desired maneuver rate, and additional jet firings resulted. To solve this problem, lag angles were provided to prevent overshoot when initiating or terminating an automatic maneuver.

The original LM DAP design attempted to incorporate information from the guidance-commanded DPS thrust into the logic for switching control from the GTS to reaction-jet control. This incorporation was thought to be necessary because the disturbance acceleration would be changed, with a change in throttle setting, because of the change in thrust-force magnitude as well as a thrust-vector rotation resulting from engine-mount compliance. Therefore, a logic was configured whereby a sensed throttle change of more than 5 percent of full thrust would force mandatory use of reaction-jet control until throttle activity subsided. Simulation testing showed that errors in computer knowledge of actual vehicle mass could prevent return to trim-gimbal control. This design discrepancy was resolved for the later flight programs by relinquishing the objective of providing the estimation filter and control-mode logic with lead information from the throttle commands.

The performance of the LM control system during the DPS-start-transient period was of sufficient concern to require design modifications of the SUNBURST flight program before the Apollo 5 mission. The major problem was caused by errors introduced in the open-loop drive-time equation and by the subsequent poor convergence from RCS control to GTS control. The effect of a drive-time error is to maintain a residual offset disturbance torque while the system is in the RCS mode. If this offset is large, the RCS jets converge the attitude and rate errors slowly to the region in which return to the GTS control is made. During this period, the jets must fire to combat the sustained offset disturbance, and unnecessary reaction-jet propellant is consumed.

Factors that significantly contribute to the error in open-loop drive time are as follows.

1. Propellant shift during ullage and the low-throttling time period
2. Actuator and engine-mount compliance
3. Uncertainties in the assumed values of mass-property variables

4. Errors in estimate of offset acceleration caused by insufficient measurements, propellant-slosh dynamics, attitude-rate initial conditions, and measurements noise

Simulation testing indicated that these factors could seriously degrade the performance of the control system during the DPS-start-transient period. Design modifications were made to improve the RCS and GTS logic interface and the quality of the estimate of offset acceleration. The interface logic provided significantly improved convergence characteristics at the expense of permitting large attitude errors during the transfer of control modes. An additional logic of ensuring a minimum number of measurements for the Kalman filter was inserted because of the transient characteristic of the estimator.

The following four additional design modifications were made to improve the DPS-start-transient performance: (1) the Kalman filter weighting values were modified; (2) the maximum open-loop drive time was limited to 15 seconds; (3) the GTS control (and Kalman filter estimates) were scheduled at specific times when operating in the low-throttle region; and (4) the GTS attitude control law during the low-thrust-start-transient period was modified.

In August 1967, a decision was made to redesign the LM DAP. The objectives of the redesign were to reduce memory storage requirements, improve off-nominal performance, and reduce computer execution time. The four major changes that resulted included simplification of the jet-selection logic, simplification of the RCS control law logic, improvement in the GTS and RCS interface design, and development of an integrated state estimator design. This improved design was incorporated in the SUNDANCE flight program.

SUNDANCE and LUMINARY DAP. - Apollo 5, the first unmanned LM mission, was launched into earth orbit in January 1968 and used flight program SUNBURST (which included design corrections for most of the preliminary design problems previously discussed). Later a significant redesign was performed, which resulted in the SUNDANCE digital program. This design version was flight-tested on the first manned LM mission (Apollo 9). Subsequent lunar landing missions were flown with slightly modified SUNDANCE digital autopilots in the LUMINARY flight program series. The significant design problems that occurred during the SUNDANCE and LUMINARY developmental phase will be reviewed in this section.

Several structural-bending problems associated with the LM/CSM RCS were identified. The original design was unstable in bending because of (1) state-estimator time lags and (2) jet-firing logic that sometimes inhibited jet commands for whole sampling cycles. This dynamics problem was identified late in the design because simulation techniques were used exclusively to verify the stability characteristics, and an unfortunate simulator error was made in the representation of the bending dynamics. These problems resulted in erroneous testing results. This design and verification problem illustrates the necessity of performing stability analysis in addition to detailed simulation testing. Acceptable bending stability was obtained by decreasing the state-estimator gains and by increasing the attitude dead band.

A structural-loading problem was identified when the LM/CSM vehicle was subjected to an RCS-jet failure condition. Significant unmodeled disturbance torques (such as a jet failed "on" that was not detected) caused the vehicle to limit-cycle on one side

of the control law phase plane. Simulations were obtained in which the control-jet pulsing frequency was close to the natural frequency of the bending mode. Test results indicated conditions in which acceptable structural load limits were exceeded. In many cases, large oscillations were sustained even after the disturbance torque excitation was removed. A design modification was made to reduce the control torque energy in the frequency spectrum of the first structural-bending mode. This reduction was accomplished by inhibiting jet firings about a given axis for a time interval after the previous firing for that axis.

An interesting design problem occurred in the area of inertial cross-coupling effects during ascent powered flight. With the jet-firing calculations established in the control (U/V) axis system, an RCS torque applied around the U-axis produces not only an acceleration around the desired U-axis, but also, in general, a coupled acceleration about the V-axis. The same situation applies to an RCS torque applied around the V-axis. The simplified equations of motion that demonstrate the effect of inertial cross coupling are as follows.

$$\dot{\omega}_U = \frac{I_{yy} + I_{zz}}{2(I_{yy})(I_{zz})} M_U + \frac{I_{yy} - I_{zz}}{2(I_{yy})(I_{zz})} M_V \quad (2)$$

$$\dot{\omega}_V = \frac{I_{yy} - I_{zz}}{2(I_{yy})(I_{zz})} M_U + \frac{I_{yy} + I_{zz}}{2(I_{yy})(I_{zz})} M_V \quad (3)$$

where $\dot{\omega}$ = angular velocity, rad/sec

M_U, M_V = control torques available, ft-lb

I_{yy}, I_{zz} = principal moments of inertia, lb-sec²-ft

The off-diagonal terms represent coupling between control axes that was sufficient to cause undesirable limit-cycle performance during powered ascent.

A design modification was made to eliminate the inertial cross-coupling effects. A nonorthogonal set of control axes (U', V') was defined that had the property that a pure U-torque produces no observable acceleration in the V'-direction, and a pure V-torque produces no observable acceleration in the U'-direction. This solution involved the computation of incremental angles through which the control axes (U, V) should be offset to become the true direction of angular acceleration during application of U-axis or V-axis torques. The control errors were then transformed to the new control axes (U', V'). These transformation angles were functions of the vehicle inertia and were periodically updated through mass-property computations.

A guidance and control interaction problem was identified for the terminal lunar landing mission phase. Large guidance and steering system time delays, together with a time-variant guidance loop gain, resulted in low-frequency attitude oscillations. The

time lags were caused by a 2-second delay inherent in the ramp-steering design concept and a computational time delay. Design modifications that were incorporated into the lunar landing mission programs included the implementation of lead compensation into the guidance loop and a decrease of the attitude dead band at small values of time to go.

A computational scaling design problem occurred for the LM/CSM GTS control system. This gimbal control system would not trim at low thrust levels because of an underflow caused by a coding implementation. Testing indicated that the powered-flight start-transient performance was extremely sensitive to initial engine mistrims. A re-scale of the coding for the GTS control law to eliminate the underflow characteristics resulted in an acceptable design.

The final design problem to be discussed was in the area of sloshing stability during DPS-powered flight. Both simulation and flight test data results have indicated that the LM/CSM and the LM configurations were sensitive to sloshing dynamics during powered flight. Detailed analyses have demonstrated that the GTS control law alone is unstable in slosh for certain conditions. However, in this situation, the total control system maintains stability because of the slosh-damping provided by the RCS control system. The performance penalty is primarily that of unnecessary jet firings. Factors that influence the GTS control system slosh stability are (1) large time lags in the rate estimator because of the nonlinear threshold logic, (2) the acceleration estimate term in the GTS control law that is filtered and does not provide effective control, and (3) large sampling lags that occur at the slosh frequencies with the GTS control law computed only five times per second.

FLIGHT TEST RESULTS

The CSM and LM flight results discussed will include performance data from the manned and unmanned missions, and results to satisfy detailed test objectives (DTO's) established preflight.

Typical flight data results are presented to indicate performance trends. The ability to match the preflight-simulation-test results closely with the actual flight data is dependent upon the quality of the telemetered data and upon the accuracy of the spacecraft environment that was modeled in the simulator. In general, powered-flight maneuvers and coasting-flight attitude maneuvers can be closely duplicated, but attitude-hold limit-cycle behavior is more difficult to match in the postflight analytic process.

Detailed test objectives were established for the LM and CSM digital autopilots; in general, these DTO's were scheduled during the early missions in an attempt to verify basic performance capabilities formally. Tables II and III summarize the DTO's, descriptions, and results. Rather than being segregated on a mission basis, the DTO's are grouped into CSM and LM tests.

TABLE II. - SUMMARY OF CSM DAP DETAILED TEST OBJECTIVES

Title	Description	Remarks
Guidance, navigation, and control system (GNCS) attitude control	Set of 15 RCS DAP functional tests: attitude hold (2 dead bands), automatic maneuvers (4 rates), manual maneuvers (4 rates), minimum-impulse controller inputs, RCS translations (3 axes), and RHC inputs in free mode	<p>One 3.6° automatic trim maneuver was recorded</p> <p>One maximum-dead-band attitude hold was performed for less than 5 minutes</p> <p>The stabilization and control system (SCS) control was used in all other RCS activity except four ullages</p>
GNCS ΔV control	Evaluation of SPS burns for long- and short-burn velocity errors, engine mistrims, start transients, center of gravity, tracking, and slosh damping	One long SPS burn and five short burns were performed and all objectives were met
GNCS entry	Evaluation of GNCS guidance for entry from earth orbit	Accurate guidance was achieved, and control was as expected except in the transonic region
Propellant slosh damping	Obtaining data on propellant slosh damping after SPS and RCS burns for RCS fuel budgeting	No unusual slosh amplitudes were noted
GNCS entry lunar return	Performance of GNCS entry from a lunar return and evaluation of the entry monitor subsystem (EMS) monitoring capability	<p>Excellent agreement between G&N and EMS data was indicated</p> <p>Nominal G&N performance was achieved</p>
TVC DAP stability margin	Stroking test	Successfully accomplished

TABLE III. - SUMMARY OF LM DAP DETAILED TEST OBJECTIVES

Title	Description	Objective verified
GNCS attitude stabilization and control during coast periods	<p>Demonstration of RCS attitude stabilization using the GNCS for attitude reference</p> <p>Verification of simple attitude and translation commands and attitude hold at minimum and maximum dead bands</p>	<p>Coasting-attitude hold (maximum dead band)</p> <p>Attitude maneuvers</p> <p>Translation commands</p>
LM GNCS/DAP performance and thrust performance	Performance of a medium-duration DPS firing to include manual throttling with CSM and LM docked and short-duration DPS burn with undocked LM	<p>Docked DPS burn</p> <p>Short DPS insertion burn</p>
GNCS attitude/translation control	Demonstration of RCS translation and attitude control of the staged LM using automatic and manual GNCS control	<p>Automatic-attitude maneuver</p> <p>Attitude hold of staged LM (maximum and minimum dead bands)</p> <p>Manual rotational and translational commands (staged LM)</p> <p>Automatic translation</p>
GNCS-controlled APS burn	Performance of a GNCS/DAP-controlled long-duration APS burn	APS burn to depletion
GNCS undocked DPS performance	<p>Flight test objective 1 —</p> <p>Evaluation of the capability of GNCS to execute DPS high-thrust undocked maneuver</p> <p>Flight test objective 2 —</p> <p>Evaluation of the capability of GNCS to execute the descent orbit insertion (DOI) maneuver</p>	<p>DPS phasing burn</p> <p>DOI burn</p>

Command-service module TVC DAP performance data . - Apollo 7 was the first manned Apollo flight to use a DAP. This flight consisted of the Apollo CSM only. (A different digital filter is required when the LM is attached.) The flight contained eight main engine burns. Five burns occurred under TVC DAP control for the entire burn, and a sixth occurred that was initiated by the primary control system, with a switch to a backup manual control mode after 35.8 seconds. Two burns were minimum-impulse burns, and three burns ranged between 7.8 and 10.9 seconds in duration. Thrust mistrim at the beginning of these burns was less than 0.2° in all cases, and peak attitude errors (except roll) were all less than 0.5° during DAP control.

Tables IV to VI are presented to summarize the significant flight test results for the TVC DAP from the Apollo 8 to 10 missions, respectively. Data from the Apollo 7 and 11 missions are omitted because these flights were similar to the Apollo 8 and 10 missions, respectively.

The Apollo 8 lunar flight used essentially the same control systems that were used by the Apollo 7 mission with equally good results. The Apollo 8 mission provided additional flight experience, especially in the area of long-duration burns. Control-system dynamics interaction with guidance performance was insignificant for the CSM-only configuration. However, this interaction was not insignificant for the CSM/LM configuration, as used on the Apollo 9 mission.

The Apollo 9 mission was the main developmental flight for the TVC DAP design. The control-loop performance had been degraded to the point of potentially significant guidance and control interaction, primarily as a safety measure for this first CSM/LM-docked flight to guard against high-frequency instability. Also, an inflight dynamics test had been designed to determine experimentally the frequency response of the airframe dynamically coupled to the SPS engine actuation system (ref. 5). This test consisted of a special routine in the CMC acting as a function generator to supply excitation to the SPS engine pitch command during powered flight. Telemetered response data

TABLE IV. - APOLLO 8 FLIGHT TEST RESULTS

SPS burn	Vehicle configuration	Total ΔV , ft/sec	Peak pitch or yaw error, deg	ΔV residuals before nulling, ft/sec		
				\dot{X}	\dot{Y}	\dot{Z}
1	CSM	24.8	0.4944	4.4	-0.1	0.1
2	CSM	3000.0	.4504	-1.4	0	.2
3	CSM	134.8	.1634	(a)	(a)	(a)
4	CSM	3520.0	.3955	-.5	.4	-.1

^aNo data.

TABLE V. - APOLLO 9 FLIGHT TEST RESULTS

SPS burn	Vehicle configuration	Total ΔV , ft/sec	Peak pitch or yaw error, deg	ΔV residuals before nulling, ft/sec		
				\dot{X}	\dot{Y}	\dot{Z}
1	CSM/LM	36.8	0.48	1.6	0.5	-0.2
2	CSM/LM	850.6	2.31	0	1.0	.2
3	CSM/LM	2570.7	4.36	2.7	-2.5	-2.3
4	CSM/LM	299.1	2.77	.2	3.9	2.7
5	CSM/LM	571.8	7.21	1.9	11.4	1.7
6	CSM	38.8	.41	1.1	-.6	-.3
7	CSM	653.3	1.85	-1.3	1.0	-.2
8	CSM	321.4	1.35	7.5	.63	-2.0

TABLE VI. - APOLLO 10 FLIGHT TEST RESULTS

SPS burn	Vehicle configuration	Total ΔV , ft/sec	Peak pitch or yaw error, deg	ΔV residuals before nulling, ft/sec		
				\dot{X}	\dot{Y}	\dot{Z}
1	CSM/LM	19.7	0.5	1.0	0.3	0.7
2	CSM/LM	48.7	.5	-.9	-.1	.3
3	CSM/LM	2982.4	.3	0	-.2	0
4	CSM/LM	138.9	.39	.5	-.4	.4
5	CSM	3630.3	.6	.3	1.6	-.1

were then analyzed to determine the total plant transfer function response to verify the dynamic-stability margin and to provide confidence in the analytical mathematical models. The essential result of this test was an experimental verification of the coupled-plant-transfer-function response in the flight environment in a frequency band encompassing the first two predominant structural resonances. The predicted amplitude ratios at the first two frequencies were -1.5 and -6.6 decibels, while those determined from the flight data were -6.0 and -12.0 decibels, respectively.

The fifth SPS burn on Apollo 9 exhibited a serious deviation of the spacecraft attitude. Data for this burn are shown in figure 9. The attitude error was quite large throughout the burn, peaking at just over 7°. Analysis of these data showed that this type of response was proper for the TVC controller being used on this mission. The

main disturbance inputs were center-of-mass motion (a jerk disturbance) as well as a 0.2° initial thrust misalignment (acceleration disturbance). In view of the fact that even larger initial thrust misalignments can occur, this performance is considered unacceptable.

Table V provides performance data for the eight SPS burns that were conducted on the Apollo 9 mission, and these data are in the form of peak attitude errors, peak body rates, and residual velocity errors at the end of the burns. The velocity error at the end of SPS burn 5 was 11.4 ft/sec, demonstrating the severity of the control dynamics interacting with guidance performance.

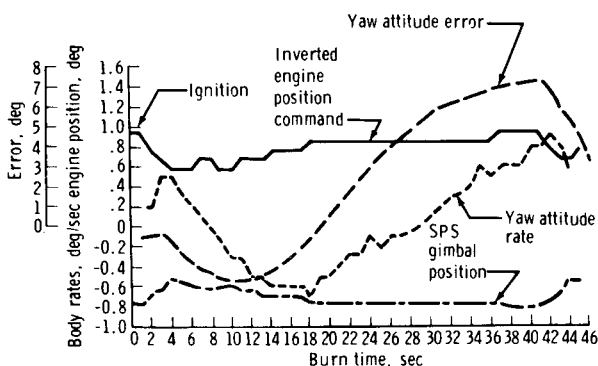


Figure 9. - Apollo 9 yaw data for SPS burn 5.

The updated CSM/LM TVC DAP design had been accepted for incorporation into the flight software that was used on the first lunar mission of the complete Apollo vehicle. As a result, guidance performance was greatly improved. The Apollo 10 performance is summarized in table VI.

Detailed test objective results for CSM. - As mentioned previously, a summary of the CSM DTO's is given in table II. In this section, the individual test objectives will be discussed in more detail.

The purpose of the GNCS attitude control test was to verify the ability of the CSM RCS DAP to perform automatic- and manual-attitude maneuvers, to maintain attitude hold, and to control RCS translations. Unfortunately, nearly all RCS maneuvers were performed in the SCS control mode, and the primary guidance, navigation, and control system maneuvers were not thoroughly tested on this mission.

No manual-attitude rotational maneuvers were performed under computer control. One automatic-attitude maneuver was performed, but the initial conditions were too nearly ideal for a valid test; four maneuver rates were called for in the DTO. A maximum-dead-band attitude hold partially satisfied one test objective, but the minimum-dead-band test did not satisfy another objective. Four ullages were performed, which demonstrated a failure of the RCS DAP to maintain minimum dead

band during ullage. This problem had been observed in preflight simulations, and a correction was made that partially verified the ability of the RCS DAP to perform attitude control during translations. No indication was given in the postflight analysis of any RHC inputs while the spacecraft was in the CMC free mode.

The objective of the GNCS delta-V control test was to verify the TVC DAP performance and guidance system accuracy during both long and short SPS burns. Six GNCS-controlled burns were performed during the Apollo 7 mission, and all test objectives were achieved. The cross-axis velocity errors for short burns were consistent with preflight simulation results and demonstrated that the primary cause of these velocity errors was the combination of poor RCS attitude hold and the failure to initialize the TVC DAP with the RCS errors at the end of ullage. A TVC DAP change was implemented in the Apollo 11 spacecraft to correct this initialization problem for the undocked CSM configuration (no correction was made for the CSM/LM-docked configuration). Two minimum-impulse burns were performed to demonstrate the ability of the TVC DAP to perform 0.5-second burns. Significant and variable slosh effects were noted in both GNCS- and SCS-controlled burns. No bending was detected. Large actuator transients caused by engine gimbal compliance were noted, but no stability problem was evident. The center-of-gravity tracking function of the DAP was demonstrated during the long burn. Steering during both long and short burns was shown to be consistent with simulation results. The ability to remove residual velocity errors was demonstrated using the RCS jets.

The purpose of the GNCS entry test was to evaluate the ability of the GNCS to guide the CSM entry from earth orbit and to provide adequate display-and-keyboard and flight-director-attitude-indicator displays. The entry maneuver was started in the SCS mode and was switched to computer control using the entry DAP. Steering accuracy was demonstrated by a 3-mile splashdown error. The DAP properly performed an over-the-top roll maneuver and limited the roll rate to the 20-deg/sec maximum. Aerodynamic effects were greater than expected in the transonic region and caused increased jet firings and high RCS propellant consumption during the last 2 minutes of the entry phase. The offset error between computer DAP rates and the rate gyroscope outputs was noted, but DAP performance was consistent with the computed values.

The purpose of the propellant-slosh-damping test was to obtain data on propellant slosh damping after SPS cutoff and after RCS burns. The main objective was to obtain data on the amount of RCS propellants that may be expended as a result of residual kinetic energy in the SPS propellants at the end of an SPS engine burn. Describing the zero-g dynamics analytically proved to be an insurmountable task; therefore, the availability of flight test data became desirable for use in budgeting RCS propellants and in determining procedures for initializing RCS attitude hold following main engine burns. All the specific procedures called for in the DTO were not performed; but reaction-jet activity was minimal after all eight SPS burns on the Apollo 7 mission, and no positive evidence of troublesome effects from zero-g propellant dynamics was found.

The purpose of the GNCS-entry lunar-return test was to verify GNCS performance in controlling an entry at lunar-return conditions and to evaluate the EMS monitor capability. The EMS load factor/velocity (G/V) trace agreed closely with the CMC-computed trace. The commander reported that the EMS range-to-go counter agreed with the G/V trace when the 50-nautical-mile line was crossed. Steering and DAP control during the automatic entry were nominal.

The attitude control of CSM/LM configuration in powered-flight test consisted primarily of the stroking test that was discussed in the previous section of this report. A complete description of the stroking test and the details of the postflight analysis are documented in reference 5.

Lunar module DAP performance data. - The LM flight performance results discussed include test data from the unmanned Apollo 5 mission and the manned Apollo 9, 10, and 11 missions.

Only the DAP coasting-flight modes were exercised in the Apollo 5 flight. Flight data for an automatic 5-deg/sec attitude maneuver showed close agreement with simulation data. The Apollo 5 mission provided some unplanned limit-cycle data during coasting ascent because of a mass-mismatch condition. This situation arose because, although the spacecraft was actually in an ascent configuration, the DAP computed the vehicle inertia to be that of the unstaged LM. As a consequence of the 300-percent-inertia mismatch condition, a virtually continuous firing limit cycle resulted, but the narrow-dead-band attitude-hold logic did maintain the desired attitude. After this operation, one RCS propellant system was allowed to fire to depletion, and data were taken at various lower thrust levels as the propellant was being depleted. Almost immediately, the limit-cycle characteristics began to improve, and eventually the attitude-hold function settled into a minimum-impulse condition. This condition implied that the DAP performed properly when its estimate of control effectiveness was normal.

Limit-cycle data were also analyzed during the descent coast phase of the Apollo 5 mission. Unexplained limit-cycle trajectories in both pitch and roll phase planes, which were asymmetrical in computed error rate and symmetrical in attitude error, were observed. During a 2-hour period, 125 jet firings occurred, approximately 30 of which had durations of from 50 to 110 milliseconds. Preflight verification testing indicated that 16-millisecond (minimum-impulse) firings should occur at the dead-band extremities. An extended effort was made to match the flight test through simulation testing. Inertia coupling, aerodynamic torques, and diagonal firing logic were all examined, but the observed limit-cycle phenomenon was only partially explained. The effect of these firings on total RCS propellant consumed for attitude control was negligible.

The Apollo 9 mission, during which the LM was manned for the first time, was flown in earth orbit. All powered- and coasting-flight DAP modes were exercised during the mission, and the control-system performance was generally excellent. No anomalous or unexpected control-system conditions occurred. Data examined in the postflight analysis included peak-to-peak rates, attitude-dead-band excursions, general limit-cycle characteristics (including existence of disturbance torques), and trim-gimbal performance. A phase-plane plot of the limit-cycle performance during a powered-ascent firing is presented in figure 10. The intent of the plot is to trace the shape of the limit-cycle trajectory. Because of the data-sampling limitations, only discrete data points in the phase plane were available. The plot does indicate on a qualitative basis that the results were within a range consistent with preflight simulation results.

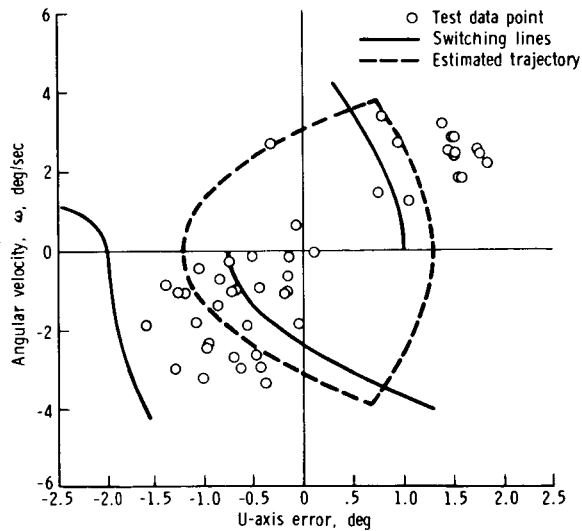


Figure 10. - Ascent propulsion system burn to depletion, U-axis phase plane.

A summary of the Apollo 9 performance flight test results is as follows.

1. Descent configuration

a. The automatic attitude-hold function before the DPS insertion burn was verified as nominal.

b. Residual velocities for the translational burns were acceptable.

2. Ascent configuration

a. The automatic attitude-hold function was nominal.

b. Automatic maneuvers appeared nominal.

c. The APS burn-to-depletion performance was acceptable.

The Apollo 9 powered-flight burn residuals are presented in table VII.

TABLE VII. - APOLLO 9 LM FLIGHT BURN RESIDUALS

Burn	Desired ΔV , ft/sec	Components of residual velocity, ft/sec		
		\dot{X}	\dot{Y}	\dot{Z}
Docked DPS	1745.0	4.2	0.1	0.1
Insertion	90.7	0	-.1	-.2
Coelliptic insertion	39.9	.1	0	0

The Apollo 10 and 11 manned missions were lunar orbital and lunar landing missions, respectively. The LM control-system performance during those operational missions was generally excellent. No anomalous or unexpected control-system conditions were observed. A summary of the Apollo 10 performance flight test results is as follows.

1. Descent configuration

- a. The automatic attitude-hold function before the DOI was verified as nominal.
- b. The automatic maneuver at 2 deg/sec to the DOI burn attitude appeared nominal.
- c. Residual velocities following the DOI and DPS phasing burns were acceptable.
- d. The maximum spacecraft rate during the phasing burn was 1.2 deg/sec and attitude-error dead band was maintained satisfactorily.

2. Ascent configuration

- a. The automatic attitude-hold function was nominal.
- b. Automatic maneuvers were performed nominally.
- c. The APS insertion burn performed nominally.
- d. The maximum rate during the APS insertion burn was 2.5 deg/sec.
- e. The transfer-phase initiation burn appeared nominal.

The Apollo 10 powered-flight burn residuals are presented in table VIII.

TABLE VIII. - APOLLO 10 LM FLIGHT BURN RESIDUALS

Burn	Desired ΔV , ft/sec	Components of residual velocity, ft/sec		
		\dot{X}	\dot{Y}	\dot{Z}
DOI	71.2	-0.1	-0.3	-0.5
Phasing	176.9	.2	-.5	-.9
Insertion	220.9	-1.5	0.3	-1.2

A summary of the Apollo 11 performance flight test results is as follows.

1. Descent configuration

- a. The DOI burn delta-V residuals appeared nominal.
- b. Attitude transients, because of the powered descent initiation, were nominal and quickly damped.
- c. Slosh oscillations were apparent 233 seconds into the burn and caused peak-to-peak amplitudes of 3.0 deg/sec.
- d. Powered descent required 88 pounds of RCS propellant compared with a budget of 40 pounds. The majority (66 pounds) was expended during periods of manual control because of terrain avoidance near touchdown.

2. Ascent configuration

- a. All ascent control functions appeared nominal.
- b. The RCS translation burns were nominal.

Lunar module DTO results. - A summary of the LM DTO's is given in table III. In this section, the individual test objectives will be discussed in more detail.

The objective of the GNCS attitude stabilization and control test was to demonstrate attitude stabilization at various inertia configurations using the RCS. This demonstration would include commanding simple attitude and translation commands in all axes in addition to maintaining attitude hold at minimum and maximum dead bands.

The jet-selection logic and the T_{jet} calculations were verified for coasting-flight attitude-hold periods by use of the available down-link data. The LGC-estimated rates, the IMU gimbal angles, and desired gimbal angles were used to construct phase planes that were used in the investigations. The attitude-hold capability was verified only for the maximum dead band because no data were available for periods of minimum-dead-band attitude hold. On subsequent missions, the attitude-hold capability for periods of minimum dead band was evaluated.

Two automatic-attitude maneuvers were investigated and found to be satisfactory. Both of the maneuvers were performed at approximately the prescribed angular rates of 5 deg/sec, and the attitude errors during these maneuvers were within $\pm 1^\circ$ of the prescribed dead band. The angular rates were determined from three sources: the LGC-estimated rates, the rate gyroscope outputs, and a matrix operation that permits a finite difference technique of the gimbal transformed to body coordinates. The agreement in rates obtained from all three sources appeared reasonable.

The RCS capability of maintaining attitude hold during translation commands was verified only for an X-axis translation. The RCS did maintain the appropriate dead bands during the periods investigated. The translation capability of the RCS was verified for the other axes on subsequent missions.

The objectives of the LM GNCS/DAP performance and thrust-performance test were to perform a medium-duration DPS firing to include manual throttling with CSM and LM docked and a short-duration DPS firing with an undocked LM and approximately half-full DPS propellant tanks. An LM/CSM-docked DPS burn was performed during the Apollo 9 mission, where the throttle profile as specified in the DTO was used. The manual-throttle profile used during the latter portion of the burn verified the variable thrust ability of the DPS engine. Satisfactory DPS engine response to thrust commands was obtained. The RCS for this burn was turned off, and attitude control was maintained solely by the GTS. Peak angular rates following the throttleup to the fixed throttle position were quite low (0.14 deg/sec about the pilot pitch (Q) axis and 0.2 deg/sec about the pilot roll (R) axis). Angular rates and attitude excursions obtained during the steady-state portion of the burn were low. The peak angular rates that occurred during the throttling profile near the end of the burn were -0.35 deg/sec for the Q-axis and 0.53 deg/sec for the R-axis. The steady-state position of the DPS engine agreed with the predicted position when compliance was taken into consideration. Slosh was evident during the burn and the magnitude and frequency of the oscillations agreed with predicted values. Velocity errors at cutoff were reasonable.

A short-duration DPS firing with an undocked LM and approximately half-full DPS propellant tanks was also performed. The peak attitude errors about the U and V control axes were less than 1.7° . The peak angular rates about these two axes were 1.18 and -0.67 deg/sec, respectively. The dead band for the insertion burn was 1° , and the RCS in combination with the GTS performed well in maintaining this dead band. The GTS did trim the descent engine to the expected value. The slosh that was evident during the burn was anticipated, and the velocity residuals were nominal.

The objective of the GNCS attitude/translation control test was to demonstrate RCS translation and attitude control of the staged LM using automatic and manual GNCS control. An automatic maneuver was performed at 2 deg/sec with command rates about the Q- and R-axes. A review of the down-link rates, CDU angles, desired rates, and final desired CDU angles verified that the overall performance was satisfactory. The desired rates were obtained, and the final CDU angles indicated that the automatic maneuver provided the appropriate vehicle rotation. The vehicle configuration was the staged LM (heavy ascent).

The attitude-hold capability for the ascent configuration was verified for a period of time immediately after the APS burn to depletion. During this period, the dead band started out at 5° and then was automatically switched to 0.3° . The DAP functioned properly, and the transition from maximum to minimum dead band was smooth. The configuration was the light-staged LM.

Eight minutes of flight data were analyzed during a period just before docking. The DAP was in a 0.3° dead-band attitude-hold mode (staged LM, heavy ascent), but a profusion of manual RHC and thrust/translation controller assembly (TTCA) activity occurred at this time. Rates induced by RHC activity were as high as 2 deg/sec, and the DAP was nulling the rates and maintaining a 0.3° attitude dead band after the release of the RHC. Attitude errors induced by TTCA activity were also nulled at the completion of the activity. The translational capability of the DAP was also verified during the coelliptic-sequence-initiation burn and immediately after the burn. The burn itself was a four-jet positive-X translation. The DAP was verified to hold the appropriate dead band during the burn. The velocity errors at cutoff (1 ft/sec in the

X-axis and -1.36 ft/sec in the Z-axis) were nulled by manual translation commands. Some cross coupling was obtained between the spacecraft axes during the manual commands, but this effect was minimized by a subsequent design modification.

The objective of the GNCS-controlled APS burn test was to perform a GNCS/DAP-controlled long-duration APS burn. During the APS burn to depletion, the peak angular rates about the Q- and R-axes were of the order of ± 5 deg/sec which agreed with the predicted values. The larger angular error excursions were of the order of $\pm 2.5^\circ$. The limit-cycle frequency was calculated during a steady-state portion of the burn. The observed limit-cycle frequency was 0.37 hertz for the pitch axis and 0.375 hertz for the roll axis. These values compared favorably with predicted values.

The DAP performed well and maintained the appropriate dead bands within anticipated limits. Cross-axis velocity errors at cutoff were nominal. Some cross coupling between the control axes was evident during this burn. This cross coupling caused some extra firings of the RCS, but these extra firings were against the offset acceleration and did not represent an inefficient use of fuel. Overall RCS-propellant consumption was as anticipated.

The objectives of the GNCS undocked DPS performance test were to evaluate the capability of the GNCS to execute a DPS high-thrust-level, undocked maneuver and to execute an undocked DOI maneuver. For the DPS phasing burn, the gimbal drive actuators appeared to have worked nominally in steering and tracking the center of gravity, even though a gimbal-fail indication occurred during the burn. The peak angular rates about the U and V control axes during the burn were -1.22 and 0.78 deg/sec, respectively. The peak attitude errors for both axes were less than 1.75° . The burn dead band was 1° . In general, the maximum attitude errors and attitude errors at maximum thrust for this burn were comparable to predicted results. The burn used less RCS propellant than was anticipated.

The DOI burn was performed behind the moon, and no telemetry data were available. The real-time data recorded for DOI are presented in table VIII. The burn was nominal with no computer alarms associated with control-system performance.

CONCLUDING REMARKS

The history of the design and development of first-generation Apollo digital control systems has been presented in this report. Because of the design flexibility inherent in digital systems, increased emphasis will probably be placed upon digital-control-systems techniques for future applications.

Each digital autopilot flown on the Apollo vehicles is unique, either by reason of completely different logic design or of a completely different approach to the programming. Therefore, each design represents an independent example of a digital autopilot that performed the design goals with minimal developmental problems. The number of versions flown (a total of 16) is impressive, and this degree of success provides important information concerning the soundness of the basic approach of implementing all flight control-system logic and dynamic compensation functions in a digital computer.

Experience gained during the development of the Apollo digital autopilots may be used to avoid future design problems. Logical decision techniques should be applied with care in design development because conditions may exist in which these techniques may unexpectedly lock out entire system functions. The use of logic in avoiding degraded performance has to be traded off with unintended restrictions.

Further research effort should be expended to develop additional analytical techniques for digital-control-system design. Adaptive design techniques making use of the inherent flexibility available in digital systems should also be established.

Design requirements should include the requirement to preserve the capability for monitoring system effectiveness. For example, efficient use of the delta-V capability of the Apollo service propulsion system placed only mild constraints on maintaining small vehicle-attitude errors and rates during the start transient; however, a design goal was to produce a system that minimized these start transients for nominal operation so that the transients would be useful indicators of off-nominal conditions that could be of serious consequence.

Caution must be exercised in the use of state variable estimation techniques because unmodeled process disturbances may allow divergence of the controlled quantity.

Caution must also be exercised in placing too much reliance on simulation results for design verification without a full appreciation of the approximations that have been made in developing the process models and in implementing these models in the simulations. The latter becomes particularly important when simulating high-frequency dynamics in a digital computer.

Manned Spacecraft Center
National Aeronautics and Space Administration
Houston, Texas, November 20, 1972
914-50-30-02-72

REFERENCES

1. Anon.: Guidance System Operations Plan for Manned CM Earth Orbital and Lunar Missions Using Program COLOSSUS 2C (Comanche Rev. 67) and COLOSSUS 2D (Manche 72, Rev. 3). Section 3, Digital Autopilots (revision 9). MIT/IL R-577, Dec. 1969.
2. Stubbs, Gilbert S.; Penchuk, Alexander; and Schlundt, Robert W.: Digital Autopilots for Thrust Vector Control of the Apollo CSM and CSM/LM Vehicles. Preprint 69-847, AIAA, Aug. 1969.
3. Anon.: Guidance System Operations Plan AS-278. Vol. 1, CM GNCS Operations. NASA CR-93682, 1966.
4. Anon.: Guidance System Operations Plan for Manned LM Earth Orbital Missions Using Program SUNDANCE (Rev. 306). Section 3, Digital Autopilot (revision 1). NASA CR-92222, 1968.
5. Peters, William H.; and Marcantel, Bernard: Inflight Dynamics Testing of the Apollo Spacecraft. NASA TM X-58089, 1972.
6. Tou, Julius T.: Digital and Sampled-Data Control Systems. McGraw-Hill Book Co., Inc., 1959.
7. Whitman, Charles Lewis: The Implementation of Digital Filters in Computers of Small Word Length. M.S. Thesis, MIT, T-443, Feb. 1966.