The choice between an iterative or noniterative method of computation depends, generally, on the extent to which perturbations affect the solution. Since no analytic expression completely describes the forces acting upon a vehicle traveling between the earth and the moon, targeting of such a trajectory involves first an analytic approximation, then orbital integration to determine the error, a second approximation to compensate for the error, and so on, bracketing the solution until either an imposed iteration limit is reached or the approximation converges on the desired solution.

The accuracy of any rendezvous computation depends upon a good knowledge of the state vectors of the two vehicles with respect to each other. Since coelliptic-sequence initiation is performed after injection or abort, the initial estimate of the LM state vector could be quite poor. Normally, ample time is available for repeated rendezvous navigation to improve the probability of good state-vector estimates before the CSI maneuver. Even in the off-nominal case, there would be sufficient time to take a certain minimum number of marks to ensure a good rendezvous.

Average G (see Section C.1.1.1), which improves knowledge of the state vector during powered flight, tends also to slightly degrade the estimate of that vector due to accelerometer uncertainties; thus rendezvous navigation is needed repeatedly to ensure the high quality of the state vector.

B.2 Ground-Targeted Maneuvers

All ground-targeted maneuvers are transmitted to the AGC via voice or telemetry uplink. Sufficient data could be transmitted to permit immediate execution of a powered-flight program but, instead, an onboard pseudo-targeting buffer program (P30) is executed prior to the maneuver. This pseudo-targeting approach has several advantages over direct maneuver execution: it provides meaningful (perhaps critical) displays to the astronaut; it can itself generate many of the inputs required by the guidance program, permitting a significant reduction in the required number of uplink variables (especially important for voice uplinks which must be entered via the DSKY); and it is designed to accept conceptually simple inputs for a crew-originated maneuver in an emergency situation when ground communication is unavailable. Furthermore, this approach serves as a backup for the onboard rendezvous-targeting programs in the highly unlikely event that the onboard primary systems in both the CM and LM fail.

RTCC ground targeting considers and accounts for such predictable effects as nonimpulsive burns, mass and thrust variations, guidance/steering rotation and Lambert aimpoint. Such unpredictable effects as thrust dispersions from nominal, or cg or V_G misalignments are ignored.

The buffer program provides inputs only to External-ΔV guidance. A similar program was devised for Lambert guidance, but it was never used and has since been deleted.

B.3 Rendezvous Maneuvers

There are two basic profiles which achieve rendezvous: direct transfer to intercept, and rendezvous using intermediate parking orbits. The AGC has the means for targeting both of these maneuvers. Each method has its own significant advantages and disadvantages. The direct transfer is fast; however, maneuver magnitudes can be quite large, imposing a possible fuel penalty, high closing rates and nonstandard lighting and intercept conditions. On the other hand, parking-orbit rendezvous allows the final phase to be standardized (thus simplifying crew training), and permits smaller maneuver magnitudes (minimizing the effects of a poor maneuver); it is, however, long and drawn out—taking several hours.

Gemini used the parking-orbit rendezvous because its computer did not have a navigation filter, and the final approach was planned to allow for easy crew monitoring. Early NASA incredulity concerning the new Apollo system favored the smaller maneuver magnitudes of the parking-orbit rendezvous to minimize the effects of a single bad burn—should one occur. Later, after confidence in the Apollo system was established, the parking-orbit rendezvous retained its precedence principally because of the ease in training the crew for the standardized final approach.

Nominally, the LM is the active vehicle throughout the entire rendezvous sequence. During this time the CSM is computing the maneuvers it would perform if it were the active vehicle—these being mirror images of the LM maneuvers. Should the LM experience a major failure (e.g., in its propulsion system), the CSM is instantly prepared to begin retrieval.

Any parking-orbit rendezvous must eventually target a direct transfer to intercept; the difference is that a parking-orbit rendezvous establishes conditions for a standardized transfer. The onboard capability for parking-orbit rendezvous is known as the concentric flight plan (CFP) and usually involves two parking orbits. Figure B.3-1 illustrates the concentric flight plan.

The first two maneuvers of the CFP are designed to place the active vehicle in an orbit that is coelliptic or concentric with the orbit of the passive vehicle (there is a constant altitude differential) and has a certain phase-angle/altitude-differential relationship defined by the elevation angle (E) at a specified time. These maneuvers are called the coelliptic-sequence initiation (P32) and constant differential height (P33). Coelliptic orbits, together with the proper elevation angle, produce

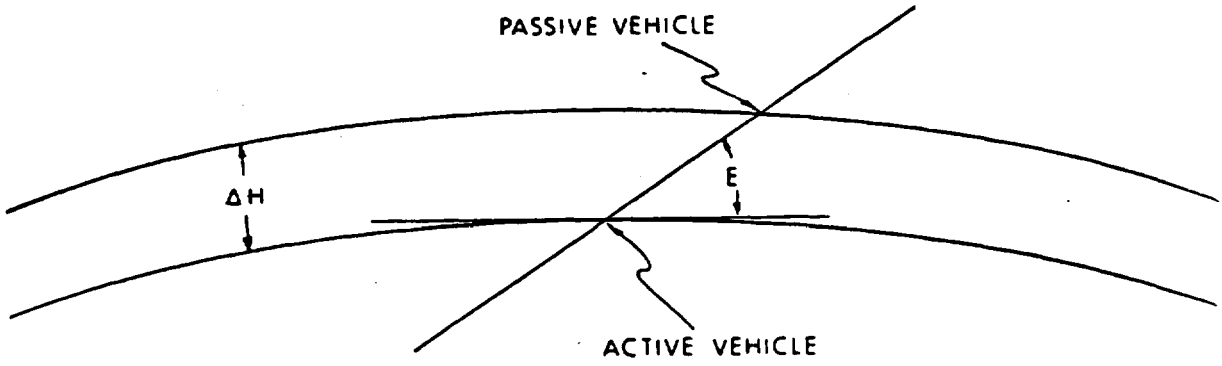

such desirable conditions as slow closing rates, easy astronaut takeover in the event of computer malfunction, and easy discovery of errors by monitoring elevation-angle changes.

Transfer to intercept, whether or not done in the context of parking-orbit rendezvous, involves simply planning the time and place of intercept and aiming to hit the spot at the proper time.

The third maneuver of the CFP is called transfer-phase initiation (P34), which targets intercept trajectories from the CDH parking orbit. Transfer-phase midcourse (P35) targets midcourse corrections (MCCs) to such trajectories.

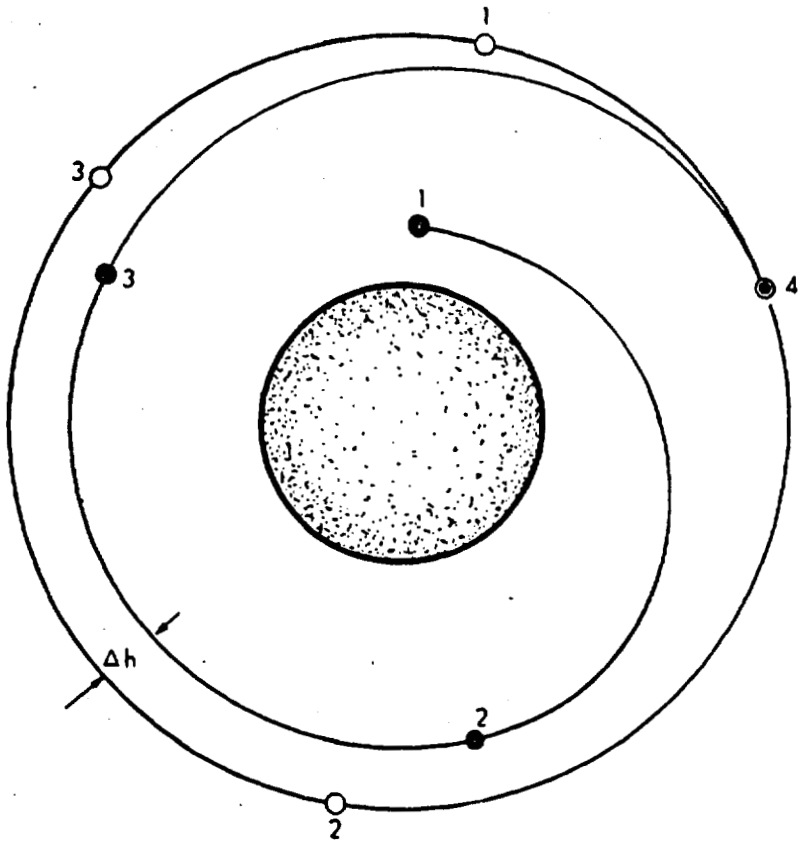

● — ACTIVE VEHICLE

○ — PASSIVE VEHICLE

Concentric Flight Plan Events

1. CSI Maneuver
2. CDH Maneuver
3. TPI Maneuver
4. Intercept

(Δh = differential altitude following CDH)

Figure B. 3-1 Concentric Flight Plan

Prior to actual intercept, a sequence of manual braking maneuvers is initiated to prevent collision and provide a proper attitude and rate of closure for docking.

All of these rendezvous programs can select either the LM or the CSM as the active vehicle and can be flown in either earth or lunar orbit.

B.3.1 Coelliptic Sequence Initiation (CSI) and Constant Differential Height (CDH)

The CSI computation iteratively solves a piecewise-continuous boundary-value problem which can be described as follows: given the current state vectors of both vehicles at CSI-ignition time $\left[t_{IG}(CSI)\right]$, what maneuver is required of the LM (assuming it is the active vehicle) to arrive at a point, TPI, at a specified time in the future, $t_{IG}(TPI)$*, and with an elevation angle of 27 deg** from the CSM?

Nominally, the ascent maneuver from the lunar surface targets the Lunar Module into a 9-by-45-mile parking orbit and the CSM is in a 60-mile circular orbit; consequently, the CSI maneuver done near apogee essentially coelliptisizes the parking orbit at about a 15-mile differential altitude and the CDH maneuver will be very small. Figure B.3-2 graphically illustrates a typical LM vertical displacement from the CSM as a function of its trailing distance. Since, during rendezvous, the LM is at a lower altitude than the CSM, it will be moving at a higher orbital velocity and will be constantly catching up with the CSM. On this graph the CSI maneuver occurs at about apogee; thus, as predicted, the CDH maneuver is very small. On this figure the acronym RR refers to rendezvous radar; AGS refers to the Abort Guidance System, which serves as a backup to the GN&CS.

The CSI computation yields two solutions—one each for CSI and CDH. In actual practice the CSI maneuver is burned from this solution, but usually the CDH maneuver ΔV is merely displayed.

*From $t_{IG}(TPI)$ to intercept, the vehicles will travel about 130 deg relative to the center of the moon (central angle of travel, also known as angle of transfer); consequently, $t_{IG}(TPI)$ is planned to occur at the midpoint of darkness to ensure good lighting conditions for docking.

**The 27-deg elevation angle is a backup requirement. In the event of a computer failure at TPI, this elevation angle permits a relatively simple manual maneuver to replace the automatic TPI. This manual maneuver is performed by aligning the LM'S thrust vector along the line-of-sight to the CSM.

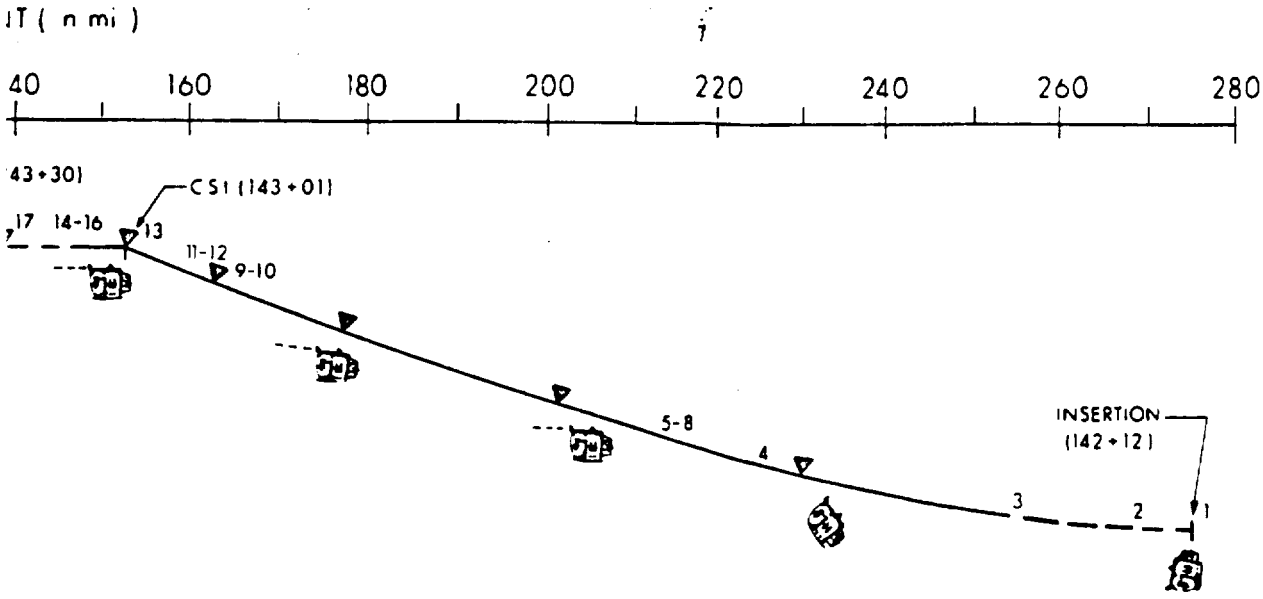

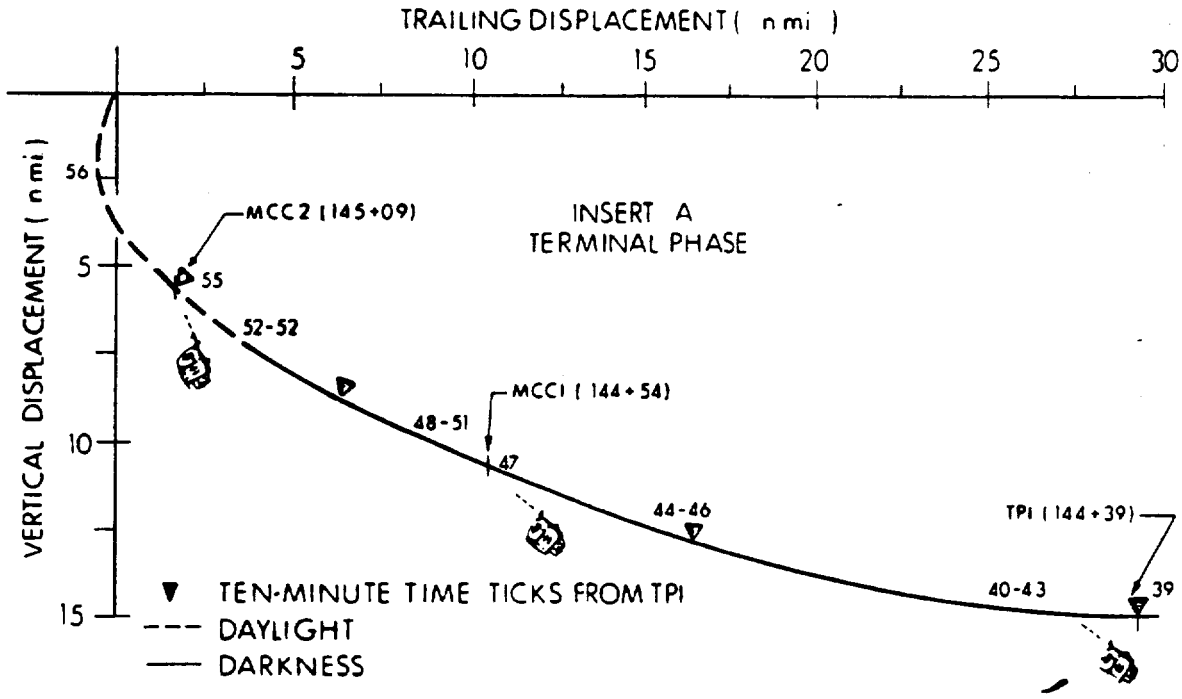

Figure B.3-2 Typical Mission Rendezvous: CSM-Centered Motion

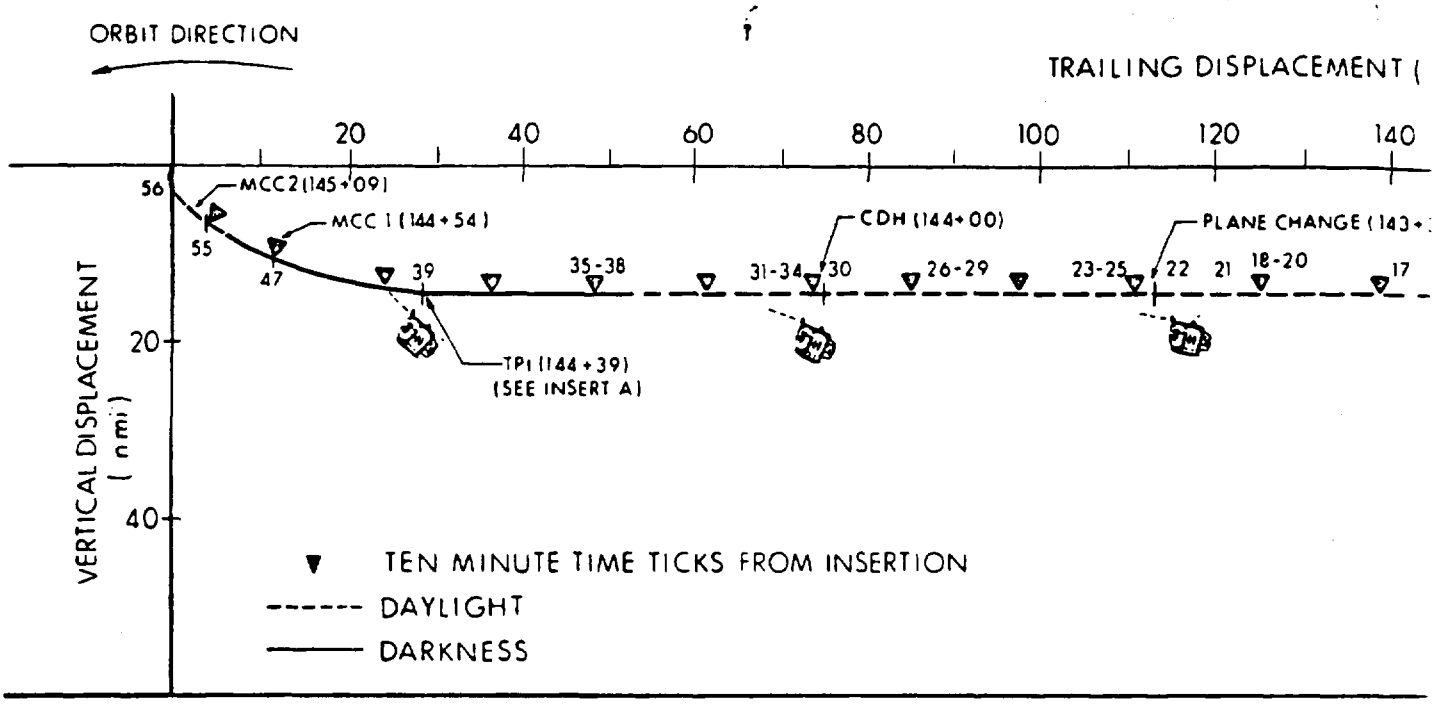

LM MAJOR EVENTS

	GROUND ELAPSED TIME	EVENT		GROUND ELAPSED TIME	EVENT
1.	142+12	INSERTION	29.	143+52	AGS FINAL CDH COMP
2.	142+13	AGS TARGET CSI	30.	144+00	CDH THRUSTING MANEUVER (P41)
3.	142+16	IMU REFSMMAT ALIGN (P52)	31.	144+02	REINITIALIZE W-MATRIX (V93)
4.	142+25	INITIALIZE W-MATRIX (10000, 10, 15)	32.	144+03	RESUME RR NAVIGATION
5.	142+26	INITIATE RR NAVIGATION (P20)	33.	144+03	TARGET TPI (GN&CS P34 AND AGS)
6.	142+26	GN&CS TARGET CSI (P32)	34.	144+05	RESUME AGS RADAR UPDATES
7.	142+27	AGS UPDATE AND ALIGN	35.	144+27	TERMINATE RR NAVIGATION (24 MARKS)
8.	142+28	INITIATE AGS RADAR UPDATES	36.	144+27	GN&CS FINAL TPI COMP
9.	142+49	TERMINATE RR NAVIGATION (23 MARKS)	37.	144+30	TERMINATE AGS RADAR UPDATES (9 UPDATES)
10.	142+49	GN&CS FINAL CSI COMP	38.	144+31	AGS FINAL TPI SEARCH COMP
11.	142+52	TERMINATE AGS RADAR UPDATES (9 UPDATES)	39.	144+39	TPI THRUSTING MANEUVER (P41)
12.	142+53	AGS FINAL CSI COMP	40.	144+41	REINITIALIZE W-MATRIX (V93)
13.	143+01	CSI THRUSTING MANEUVER (P41)	41.	144+41	RESUME RR NAVIGATION
14.	143+03	INITIALIZE W-MATRIX (2000, 2, 5)	42.	144+41	TARGET MCC1 (GN&CS P35 AND AGS)
15.	143+04	RESUME RR NAVIGATION	43.	144+41	RESUME AGS RADAR UPDATES
16.	143+04	TARGET CDH (GN&CS P33 AND AGS)	44.	144+51	TERMINATE RR NAVIGATION (9 MARKS)
17.	143+07	RESUME AGS RADAR UPDATES	45.	144+51	TERMINATE AGS RADAR UPDATES (6 UPDATES)
18.	143+21	TERMINATE RR NAVIGATION (17 MARKS)	46.	144+51	GN&CS AND AGS FINAL MCC1 COMP
19.	143+22	OBTAIN CSM ROOT FOR PLANE CHANGE	47.	144+54	MCC1 THRUSTING MANEUVER (P41)
20.	143+22	EXTERNAL ΔV FOR PLANE CHANGE (P30)	48.	144+56	REINITIALIZE W-MATRIX (V93)
21.	143+24	TERMINATE AGS RADAR UPDATES (7 UPDATES)	49.	144+56	RESUME RR NAVIGATION
22.	143+30	RCS PLANE CHANGE MANEUVER (P41)	50.	144+56	TARGET MCC2 (GN&CS P35 AND AGS)
23.	143+31	REINITIALIZE W-MATRIX (V93)	51.	144+56	RESUME AGS RADAR UPDATES
24.	143+32	RESUME RR NAVIGATION	52.	145+06	TERMINATE RR NAVIGATION (9 MARKS)
25.	143+32	TARGET CDH (GN&CS P33 AND AGS)	53.	145+06	TERMINATE AGS RADAR UPDATES (6 UPDATES)
26.	143+48	TERMINATE RR NAVIGATION (15 MARKS)	54.	145+06	GN&CS AND AGS FINAL MCC2 COMP
27.	143+48	GN&CS FINAL CDH COMP	55.	145+09	MCC2 THRUSTING MANEUVER (P41)
28.	143+51	TERMINATE AGS RADAR UPDATES (7 UPDATES)	56.	145+12	BRAKING (P47)

State-vector errors tend to propagate during the nontracking period of an orbital maneuver. To ensure the best possible state vectors for inputs to the CDH computation, further tracking is performed after the CSI burn; then the CDH maneuver is recomputed. However, if for some reason (such as loss of visibility) tracking cannot be performed between the two maneuvers, the astronaut will still call for the CDH maneuver—just as he would in the nominal case—but with degraded accuracy.

Other than not achieving the desired boundary conditions, there are a few other situations which could render unacceptable a CSI solution. One is when the various burns occur too closely together to permit tracking in between; another is when the LM is phased improperly with the CSM; still another is a too-large ΔV or a too-small ΔH. DSKY alarms notify the astronaut should one of these situations occur; he can then adjust his inputs to the CSI program until he achieves a satisfactory solution. Typically, the CSI and CDH burns are 180 deg apart, with a plane-change maneuver halfway in between to make the orbits coplanar.

The inputs the astronaut can adjust are t_{IG}(TPI), CDH apsidal crossing and the standard TPI line-of-sight elevation angle, E. He would not adjust the vehicles' state vectors, which are also inputs. The astronaut carries onboard charts which give him nominal as well as non-nominal inputs to the CSI computation. Although the onboard rendezvous-targeting computations are prime, the astronaut can always check his inputs with the CSM or the ground if he chooses.

When the CDH burn is completed, the elevation angle is small (e.g., 9 deg), the LM is catching up with the CSM and tracking is resumed. About 40 minutes later the elevation angle is 27 deg, indicating TPI time.

B.3.2 Transfer Phase Initiation (TPI) and Transfer Phase Midcourse (TPM)

If the CSI and CDH maneuvers successfully produce the desired boundary conditions at TPI time, the transfer phase will possess several desirable properties, such as a slow closing rate; predictable changes in the elevation angle, E; nearly-coincident TPI-thrust direction with the line-of-sight; and a $|\Delta V|$, measured in feet per second, which is nearly twice the $|\Delta H|$ measured in miles. These properties allow easy monitoring and easy astronaut takeover should the computer malfunction during or after the TPI maneuver.

The TPI computation solves the Lambert (intercept) problem directly. It solves for the maneuver required of the LM to intercept the Command Module at a given angle of transfer* from t_{IG}(TPI). In earth orbit, this Lambert solution uses offset aimpoints to compensate for the effects of earth oblateness.

Typically, the Lambert problem is computed two more times after the TPI burn to make transfer-phase midcourse corrections. Of course, time is allowed between all of these burns for further tracking.**

After the second midcourse correction (about 4 miles from the CSM), the astronaut begins a manual braking schedule during which he maintains a collision course and brakes his vehicle. Finally, the LM comes up in front of the CSM while pitching over, so that each docking hatch faces the other.

The GN&CS rendezvous function ends with the initiation of the manual braking schedule.

---

*Typically, 130 deg of CSM travel, or a transfer time of about 43 min.

**Historically, two other rendezvous routines were planned and subsequently discarded—TPI Search, and stable orbit rendezvous (SOR).

The TPI Search was intended to provide a fast return in an abort situation. It could be very costly in fuel consumption, but that compromise was deemed necessary, for example, in the case of a partial failure of the life-support systems. The program allowed the astronaut to manually iterate central angles of travel to find an early-rendezvous trajectory which would minimize fuel consumption. After the initial burn, the problem would be solved again to provide MCCs, as in the procedure used for the standard TPI.

TPI Search doesn't provide the easy monitoring and standard final transfer that the concentric flight plan does; it could be fuel costly but it didn't have to be, because there were fewer burns; and it could have large closing rates necessitating a busy braking schedule. As NASA confidence in the entire system grew, this type of contingency maneuver was discarded.

The SOR targeted an intercept with the CSM's orbit at a fixed distance behind or ahead of it. This aimpoint was used in a Lambert calculation to determine the necessary maneuver. After the SOR burn, navigation and midcourse corrections could further improve the accuracy of the maneuver. At the point where the two orbits cross, another maneuver would have been performed to match the two vehicles' orbital velocities. Needless to say, none of the burns mentioned here are impulsive. Final closure and docking would be accomplished manually.

B.4 Return to Earth

The principal task of the return-to-earth program (P37) is to provide an onboard targeting capability to return the spacecraft to a proper earth-reentry corridor in the event an abort occurs while there is a loss of communications with the ground. With this program, safe returns can be achieved from earth orbit, from trajectories resulting from translunar-injection SIVB-powered maneuver failure, from translunar coast (outside the lunar sphere of influence), and from transearth coast, including midcourse corrections (also outside the lunar sphere of influence). Figures B.4-1, B.4-2, and B.4-3 indicate typical RTE trajectories at TLI+6 hours, TLI+20 hours and TLI+50 hours, respectively.

B.4.1 Options

P37 originally provided the crew with three basic options—minimum-fuel return, minimum-time return and landing-site designation. Limited computer capacity has since dictated the elimination of the landing-site designation option; nevertheless, a limited amount of designation is still possible, since ΔV directly affects the return time which, due to earth rotation, influences the landing site. Three of the inputs for either remaining option are desired ignition time (t_{IG}), desired velocity change (ΔV_D), and desired reentry angle. The choice between minimum-fuel and minimum-time return is determined by the input ΔV_D (zero for the minimum-fuel option). The program provides a display of the landing-site latitude and longitude, required velocity change, the spacecraft's velocity magnitude at 400,000-ft entry altitude measured above the Fischer ellipsoid, the flight-path angle at 400,000-ft entry altitude, and the transit time to that point from the time of ignition. After the astronaut selects which propulsion system to use, SPS or RCS, the following quantities are displayed: middle-gimbal angle at ignition, time-of-ignition and time-from-ignition.

Two principal steps are used to compute the return trajectory. First, a two-body conic solution is generated which satisfies the constraints dictated by the inputs, thus yielding a fast two-body approximation which is displayed to the astronaut. Should he so elect, he can continue on to a precision solution recomputed to consider gravitational perturbations. The conic characteristics are then used as a basis for the development of a precision trajectory within the original constraints. The

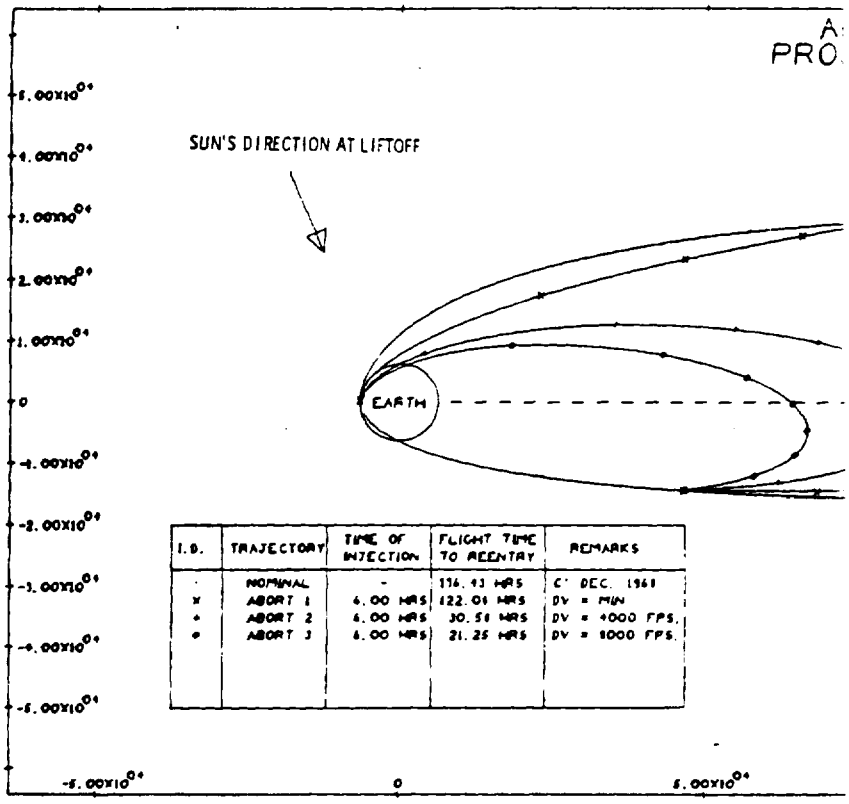

I.D.	TRAJECTORY	TIME OF INJECTION	FLIGHT TIME TO REENTRY	REMARKS
.	NOMINAL	-	136.43 HRS	C' DEC. 1961
x	ABORT 1	6.00 HRS	122.06 HRS	DV = MIN
+	ABORT 2	6.00 HRS	30.58 HRS	DV = 4000 FPS.
o	ABORT 3	6.00 HRS	21.25 HRS	DV = 1000 FPS.

POLLO TRAJECTORY IN EARTH CENTERED INERTIAL COORDINATES
JECTED INTO MOON ORBITAL PLANE. (EXPRESSED IN KILOMETERS)

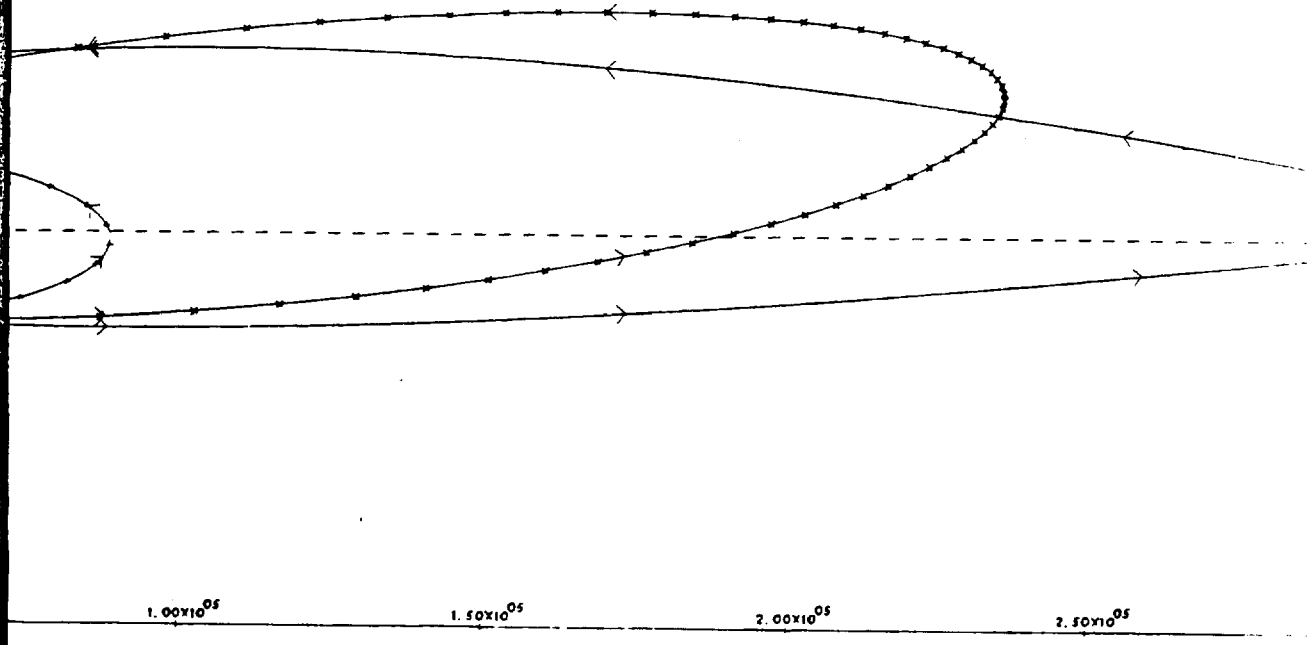

144-B

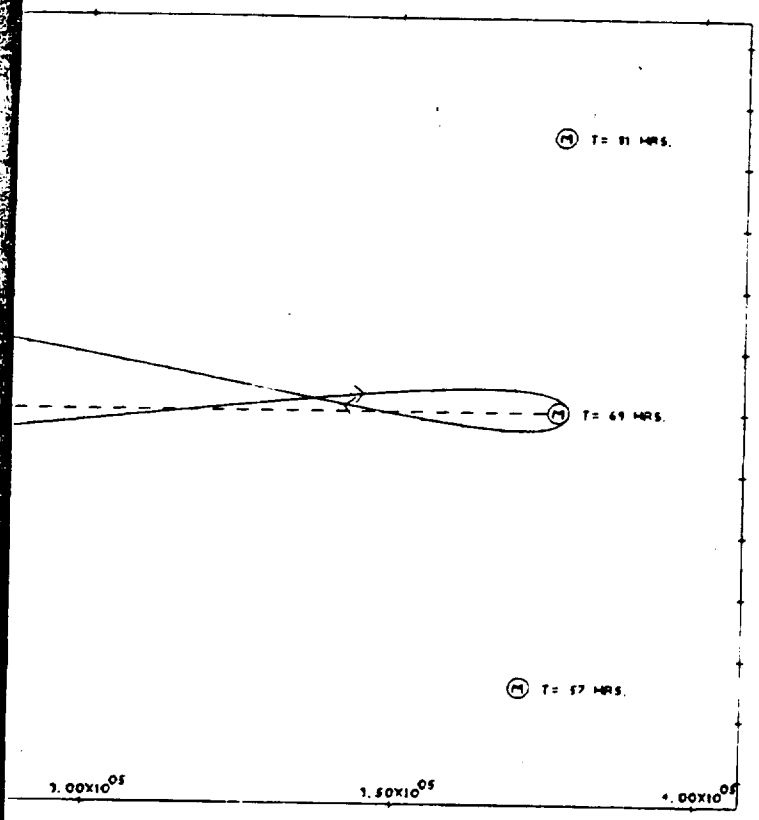

Figure B. 4-1 Typical Abort Trajectories for TLI+6 Hours

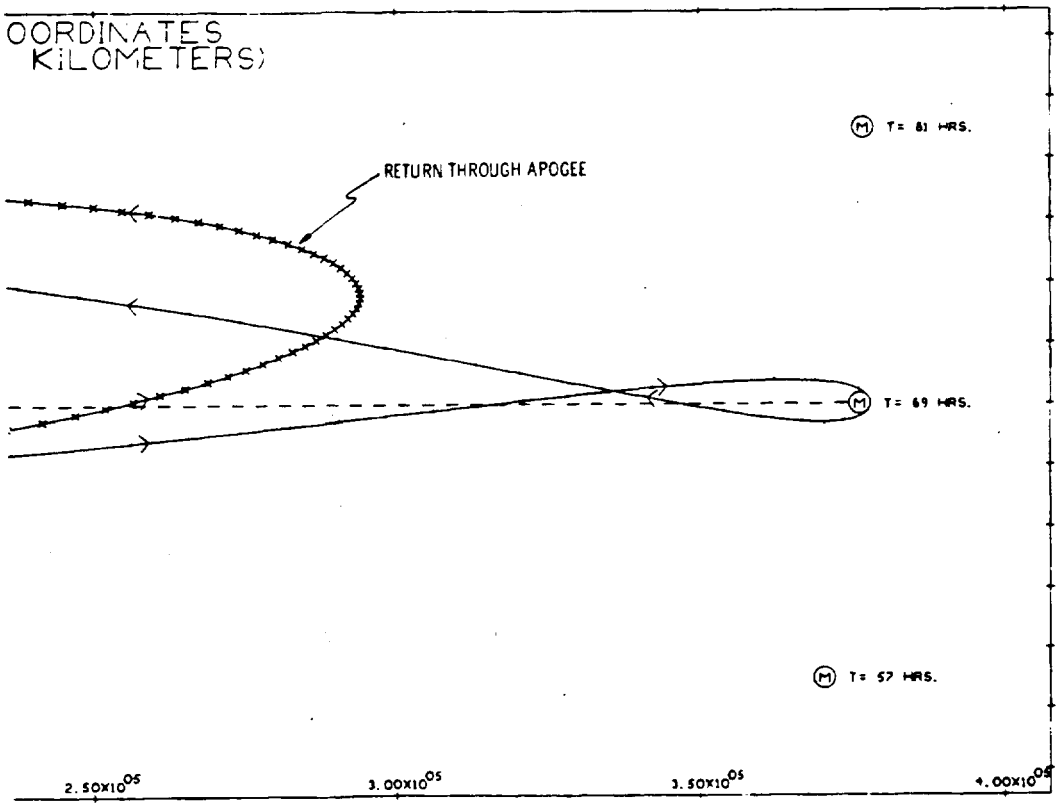

Figure B. 4-2 Typical Abort Trajectories for TLI+20 Hours

APOLLO TRAJECTORY IN EARTH CENTERED INERTIAL COORDINATES
PROJECTED INTO MOON ORBITAL PLANE. (EXPRESSED IN KILOMETE

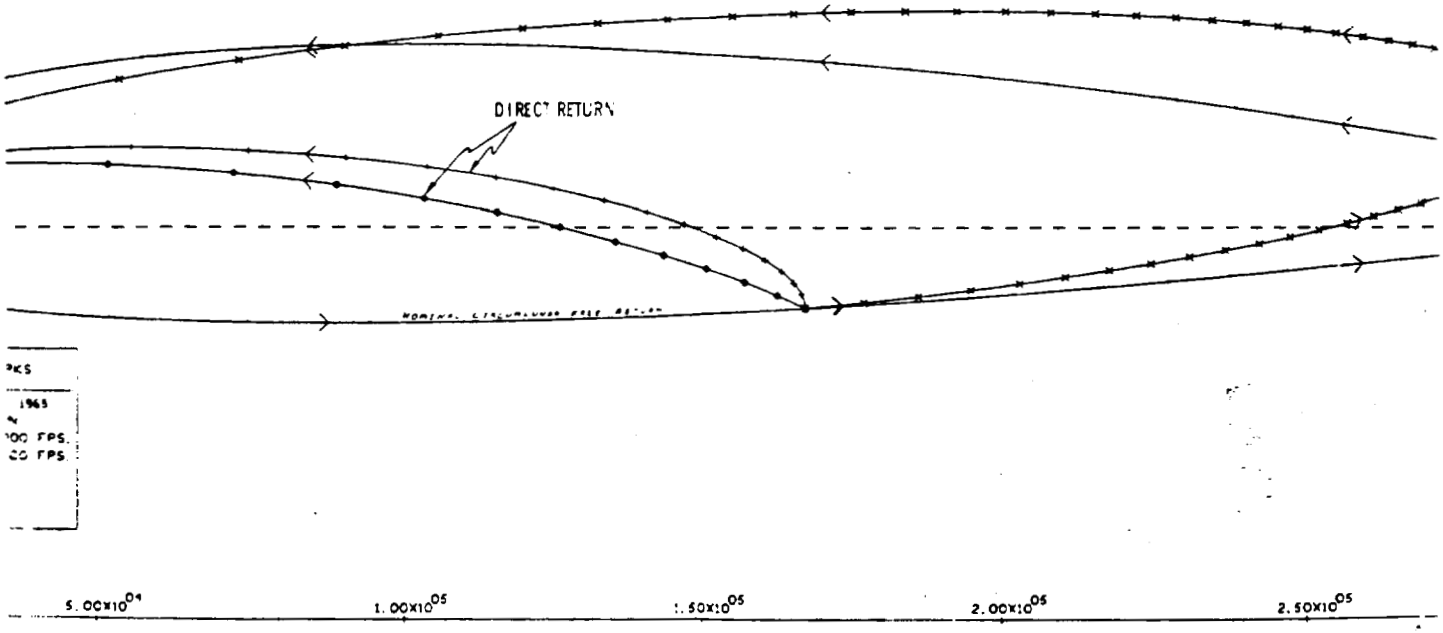

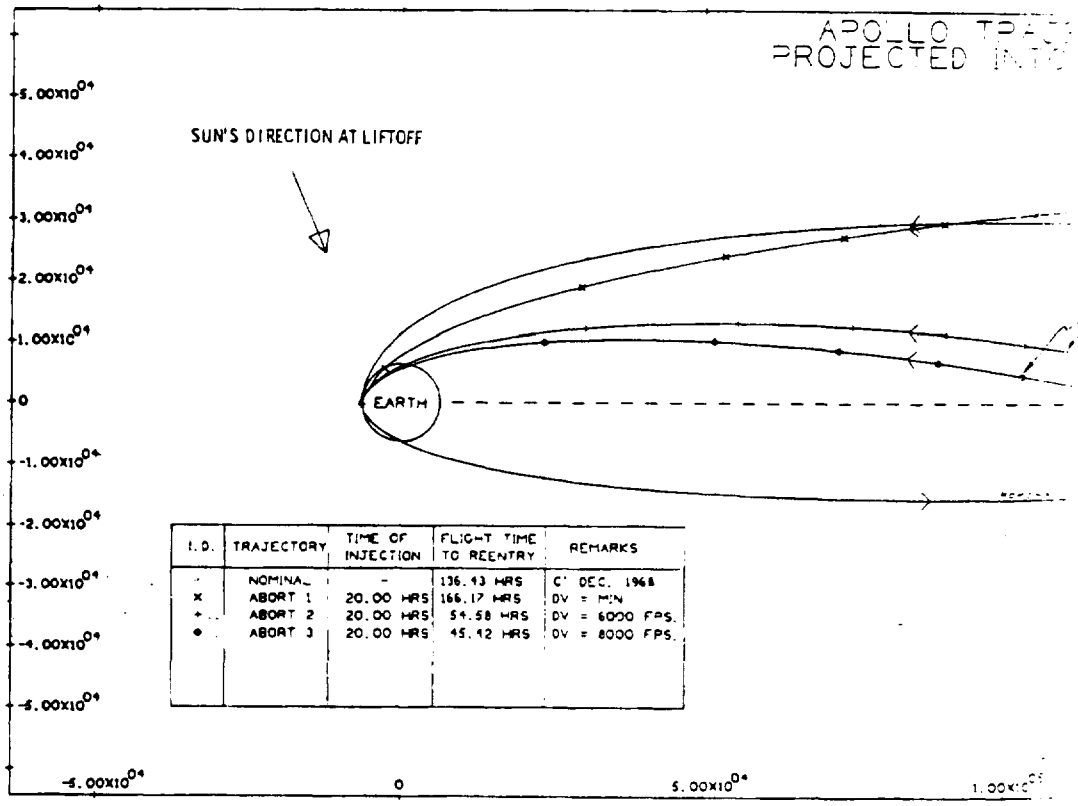

I.D.	TRAJECTORY	TIME OF INJECTION	FLIGHT TIME TO REENTRY	REMARKS
·	NOMINAL	-	136.43 HRS	0° DEC. 1968
×	ABORT 1	20.00 HRS	166.17 HRS	ΔV = MIN
+	ABORT 2	20.00 HRS	54.58 HRS	ΔV = 6000 FPS.
•	ABORT 3	20.00 HRS	45.42 HRS	ΔV = 8000 FPS.

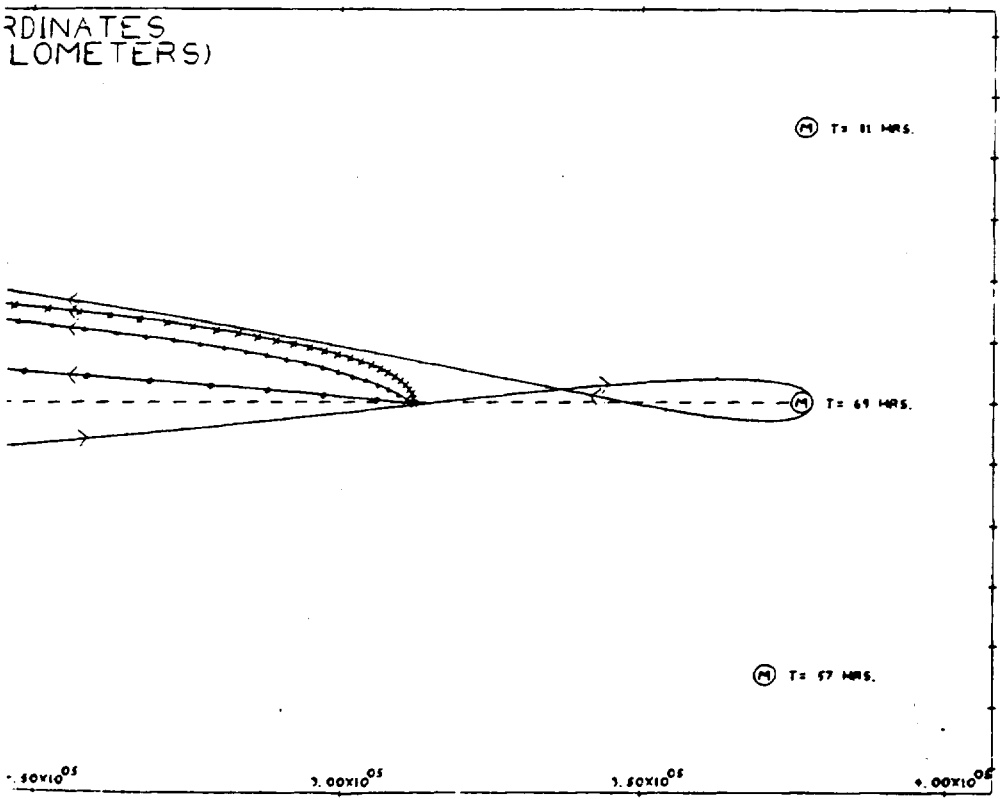

Figure B. 4-3 Typical Abort Trajectories for TLI+50 Hours

APOLLO TRAJECTORY IN EARTH CENTERED INERTIAL COORDINATES
OJECTED INTO MOON ORBITAL PLANE. (EXPRESSED IN KILOMETERS)

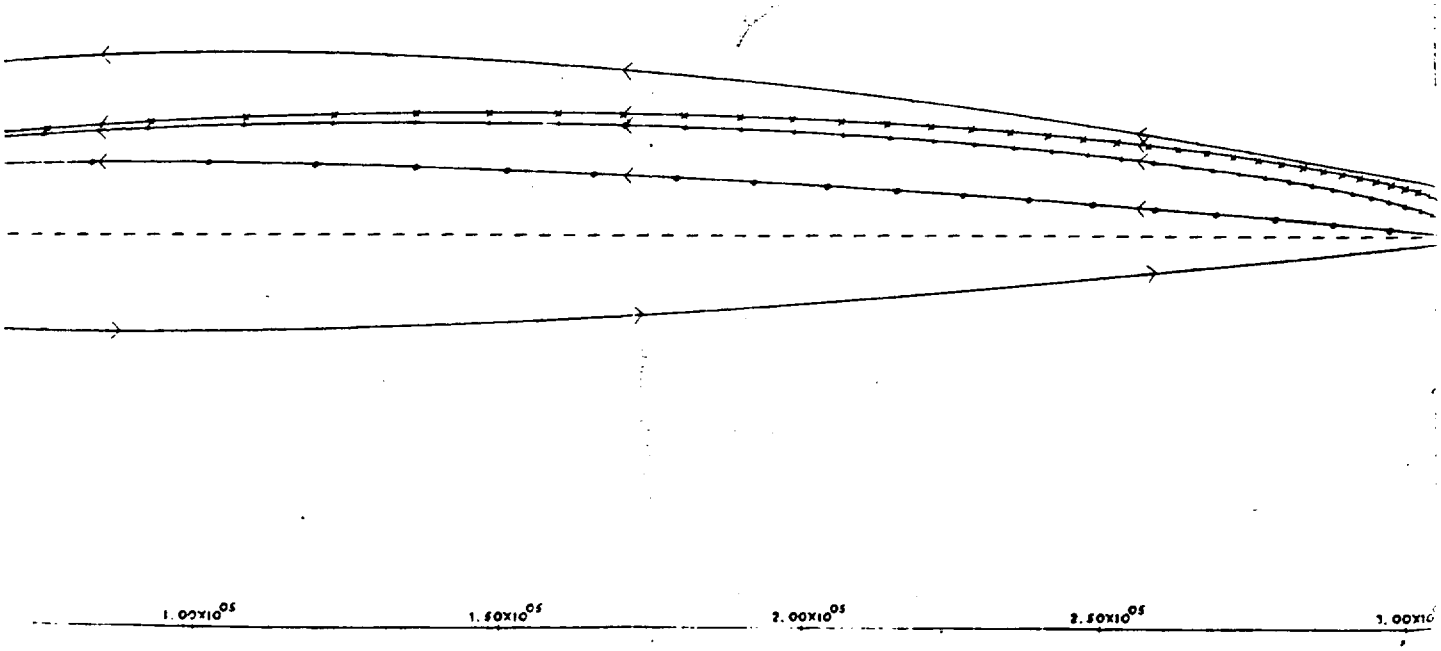

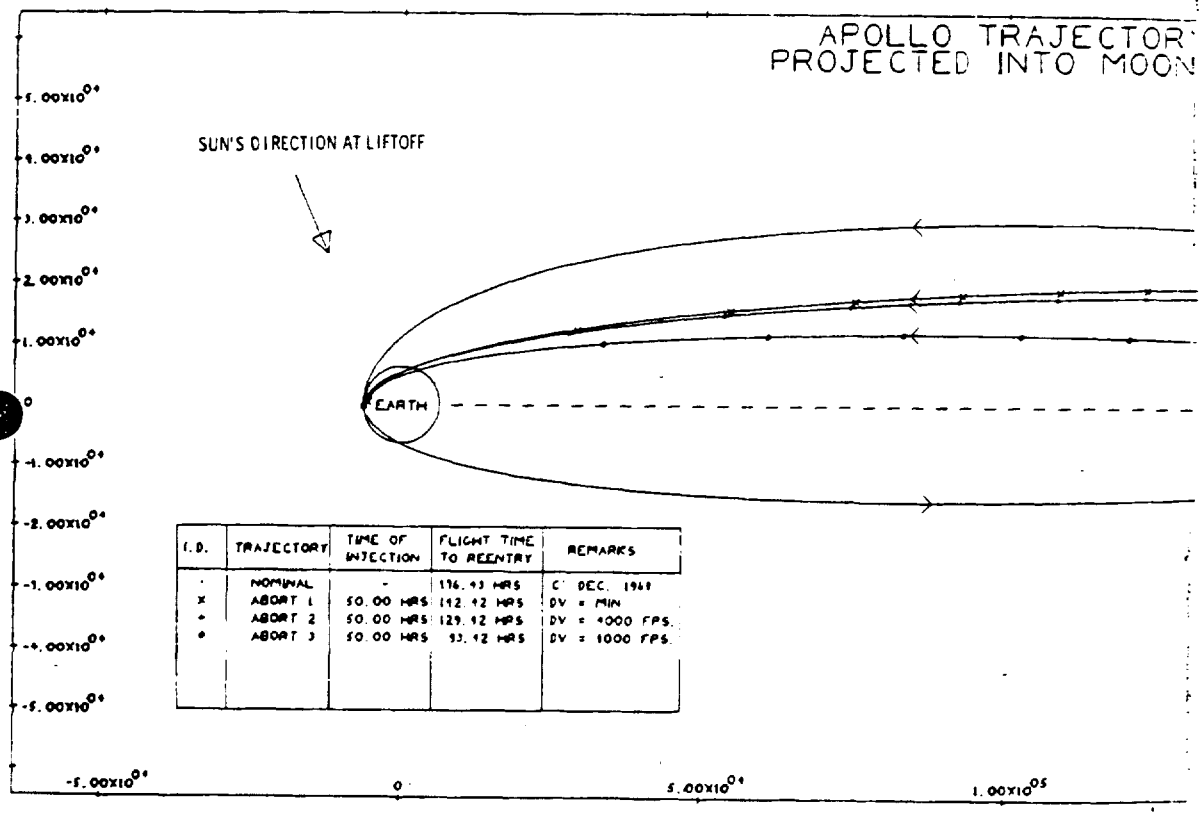

I.D.	TRAJECTORY	TIME OF INJECTION	FLIGHT TIME TO REENTRY	REMARKS
·	NOMINAL	-	176.93 HRS	C³ DEC. 1961
x	ABORT 1	50.00 HRS	142.42 HRS	DV = MIN
+	ABORT 2	50.00 HRS	129.42 HRS	DV = 4000 FPS.
*	ABORT 3	50.00 HRS	93.42 HRS	DV = 1000 FPS.

precision-trajectory displays can differ significantly from the conic displays, especially in the case of long transit-time returns. Thus the astronaut is obliged to check the new displays before accepting the precise solution.

B.4.2 Two-Body Problem

The fuel-critical two-body problem consists of generating a conic trajectory which returns the spacecraft to a specified reentry radius and flight-path angle under the constraint of minimizing the impulsive ΔV required to achieve the trajectory. The time-critical two-body problem is similar, except that the impulsive ΔV is prescribed rather than minimized. In either case, the premaneuver state vector is assumed to be known. Since reentry conditions do not constrain the trajectory plane, an in-plane maneuver is most efficient, thus reducing the problem to two dimensions.

During the program design, consideration was given to a closed-form solution to the minimum-fuel problem. This would have required solving for the roots of a general fourth-order polynomial. The logic necessary to eliminate imaginary and physically unrealizable roots made this approach unattractive. It was concluded that an iterative search, using the cotangent of the post-maneuver flight-path angle as an independent parameter, would be a better approach to the problem. The latter approach always converges on a physically realizable solution and greatly simplifies the equations. More important, the same equations can be used in the solution of the time-critical problem, and additional trajectory constraints can be readily included.

Computation of the two-body solutions to both the fuel-critical or time-critical problems is complicated by initial uncertainty in the reentry radius and flight-path angle. The reentry radius is defined to be 400,000 ft above the Fischer ellipsoid and therefore varies with reentry latitude. The flight-path angle is a function of the reentry speed and is represented by a polynomial curve fit. An initial rough estimate is made of both the radius and flight-path angle. Based on the resulting reentry latitude and speed, a better estimate of the proper radius and flight-path angle can then be made. This iteration is automatically continued until satisfactory terminal conditions are achieved. Then the landing site is computed, based on either a prestored entry range or the range resulting from a half-lift reentry profile.

The two-body solution does not reflect trajectory perturbations resulting from the earth's oblateness, the moon, or the finite duration of the maneuver. In most cases, however, it presents a good approximation to the maneuver size and comes within 10 to 15 deg of the true landing site. After the state vector is extrapolated to the ignition time, about 15 sec are required to compute the conic solution.

B.4.3 Precision Solution

A precision solution to the return-to-earth problem may be initiated after the user has an opportunity to evaluate the conic approximation. Although the precision solution still assumes an impulsive maneuver, it does compensate for trajectory perturbations due to the earth's oblateness and the moon. An attempt is made to satisfy the original time-critical or fuel-critical constraint, in addition to the reentry radius and flight-path angle requirements obtained from the two-body solution. The object of the precision phase is to compute a new post-impulse velocity which compensates for the perturbed gravitational field. The technique used to accomplish this also makes use of the two-body solution to determine the new post-impulse velocity, by rerunning the conic computations with an offset reentry radius.

The precision solution displayed to the crew contains no significant approximations. The accuracy of the solution is largely a function of the accuracy of the onboard state vector used by the program. The answers, therefore, are comparable to those obtained from a large ground computer.

The running time of the precision phase is significant, and is primarily a function of the time of flight from ignition to reentry. It typically ranges between one and two minutes per ten hours of flight time.

B.4.4 General Considerations

The return-to-earth program is included in the onboard computer primarily to provide return-targeting capability in the event that communication with the ground is lost. During translunar coast the availability of the LM for communications backup makes the probability of a communications loss very small. Therefore, the mission phase most likely to require use of the return-to-earth program is transearth coast. In this case, the only requirement is to provide targeting for small midcourse

corrections which ensure that the spacecraft reenters in the center of the entry corridor. Small minimum-fuel maneuvers, primarily in the horizontal direction, have negligible effect on the landing site; they are made several times during the return trip, and are preceded by periods of cislunar navigation.

Although the most probable use of the RTE program is in the fuel-critical mode for midcourse corrections, the crews have shown considerable interest during training in using the time-critical mode to achieve some degree of landing-site control. As explained earlier, landing-site longitude variability can be accomplished by varying the input ΔV. Optimum use of both the conic and precision phases of the program can minimize the computer time required for a solution which results in a particular longitude. Since the conic solution provides a fast approximate answer, it should be rerun until a ΔV is determined which results in a longitude close (within approximately 15 deg) to the desired longitude. Within this region the longitude is a fairly linear function of the ΔV. This solution should be continued through the precision phase to accurately establish the longitude error. Following this, the ΔV should be adjusted slightly and another conic solution computed. By comparing this solution to the previous conic solution, the sensitivity of longitude to changes in ΔV can be determined. A simple, manual computation can determine how much change in ΔV is required to bring the precision longitude to the desired longitude. This straightforward linear technique has proven to be quite efficient in producing a solution with a particular landing-site longitude.

APPENDIX C

MAJOR PROGRAM CAPABILITIES—
Powered-Flight Guidance and Navigation

C.1 Fundamentals of Powered-Flight Guidance and Navigation

The Apollo mission profile consists of many discrete segments, each with definable characteristics and objectives. (The segments are discussed in detail in Section 2.2.) The translunar-coast, transearth-coast, earth-orbit, and lunar-orbit phases represent examples of coasting flight—periods when the spacecraft trajectory is subject only to local gravitational forces. Connecting these coasting-flight segments are briefer, powered-flight segments, such as translunar injection (TLI), transearth injection (TEI), and midcourse corrections (MCCs)—periods when the spacecraft trajectory is subject not only to local gravitational forces, but also to purposeful application of thrusting and/or aerodynamic acceleration. In addition to those instances where it serves to accomplish a transfer between two coasting-flight phases, powered flight also occurs whenever either the initial or terminal state is on the earth or lunar surface. Thus, boost from and entry into the earth's atmosphere and descent to and ascent from the lunar surface are also powered-flight phases.

C.1.1 Powered-Flight Navigation

Powered-flight navigational requirements differ fundamentally from the navigational requirements of coasting flight.

During coasting flight, gravitational accelerations are the dominant influence upon the spacecraft's state vector. These accelerations are modeled quite closely, considering both the primary and secondary celestial bodies and the sun, and are integrated accurately to span the long coasting-flight intervals. Thus, precision integration greatly reduces the requirement for frequent state-vector updates either from the ground or computed onboard from optics and VHF ranging measurements.

During powered flight, spacecraft dynamics and scarcity of time preclude the use of optics and VHF ranging measurements; consequently, powered-flight navigation uses accelerometer measurements of aerodynamic and thrusting accelerations and a greatly simplified gravity model*.

C.1.1.1 Gravity Computation

The gravity calculations are performed in a straightforward manner. The equations of motion for a vehicle moving in a gravitational field are

$$\frac{d\underline{r}}{dt} = \underline{V} \quad ,$$

and

$$\frac{d\underline{V}}{dt} = \underline{g} + \underline{a}_T \quad ,$$

where $\underline{r}$ and $\underline{V}$ are the position and velocity vectors with respect to an inertial frame of reference. The measured acceleration vector of the vehicle, $\underline{a}_T$, is defined as the vehicle acceleration resulting from the sum of rocket thrust and aerodynamic forces, if any; $\underline{a}_T$ would be zero if the vehicle moved under the action of gravity alone. The vector sum of $\underline{a}_T$ and $\underline{g}$, the gravitational vector, represents the total vehicle acceleration.

Position and velocity are obtained as a first-order difference-equation calculation through a simple computational algorithm:

$$\Delta\underline{V}_a(t_n) = \underline{V}_a(t_n) - \underline{V}_a(t_{n-1}) \quad ,$$

$$\underline{r}(t_n) = \underline{r}(t_{n-1}) + \left[\underline{V}(t_{n-1}) + \frac{1}{2}\Delta\underline{V}_a\right]\Delta t + \frac{1}{2}\underline{g}_{n-1}(\Delta t)^2 \quad ,$$

$$\underline{V}(t_n) = \underline{V}(t_{n-1}) + \Delta\underline{V}_a(t_n) + \frac{1}{2}(\underline{g}_n + \underline{g}_{n-1})\Delta t \quad .$$

*For lunar-centered flight, a simple spherical force field is considered. For earth-centered flight, a first-order oblateness term is added. No "other body" effects are considered.

The vector $\underline{V}_a$ is the time integral of the nongravitational acceleration forces, the components of which are the outputs of three mutually orthogonal integrating accelerometers situated in the spacecraft's Inertial Measurement Unit. The gravitational vector $\underline{g}_n$ is a function of position at time t_n. This method is called "Average G" because velocity is updated by means of the average of the gravity vectors at the extremes of the measurement intervals.

A careful error analysis of a vehicle in earth orbit has demonstrated the above algorithm to yield errors of approximately 100 ft and 0.2 ft/sec after a period of 35 minutes, using a 2-sec time step and rounding all additions to eight decimal digits. (Errors increase for smaller time steps due to the effects of accumulated round-off errors; interestingly, errors also increase for larger time steps as truncation errors assume significance.) Compared to typical accelerometer scale-factor errors, the errors in the Average-G algorithm are smaller by several orders of magnitude.

C.1.2 Powered-Flight Guidance Using Cross-Product Steering

Powered-flight guidance and steering control the direction of applied thrusting acceleration to meet the targeted terminal conditions. The guidance law generates the necessary trajectory-control variables for the steering law, in addition to preparing the requisite displays to enable the astronauts to monitor the progress of the maneuver.

Steering controls the direction of vehicle thrusting accelerations according to the behavior of the guidance trajectory-control variables. Steering is implemented either manually or automatically. Manually, the astronaut implements steering by using his hand controller to apply translational accelerations via the Reaction Control System (RCS) Autopilot. Automatically, steering is implemented by a cross-product steering law which generates spacecraft-attitude commands for the Thrust Vector Control (TVC) Autopilot. (These autopilots are discussed in Appendix D.)

Powered-flight guidance depends upon the generation of an instantaneous required velocity ($\underline{V}_R$), corresponding to the present spacecraft position such that the targeted terminal conditions are achieved. Subtracting the current velocity ($\underline{V}$)

from the required velocity results in the instantaneous velocity to be gained ($\underline{V}_G$):

$$\underline{V}_G = \underline{V}_R - \underline{V} \quad .$$

Hence powered-flight guidance is commonly called $\underline{V}_G$ guidance. The steering law attempts to null this $\underline{V}_G$ (i.e., to cause the orbital velocity to approach the desired velocity) by controlling the direction of spacecraft thrusting acceleration.

For manual nulling of $\underline{V}_G$, the current $\underline{V}_G$ is an adequate control variable when expressed and displayed in a convenient coordinate system, e.g., along the orthogonal RCS jet axes. The astronaut selects X/Y/Z translation to null the corresponding displayed $\underline{V}_G$ components.

C.1.2.1 Cross-Product Steering

For automatic $\underline{V}_G$ nulling, the cross-product steering law generates vehicle-attitude commands for the TVC autopilot, which stabilizes vehicle attitude; in steady state the thrust vector will go through the spacecraft's center of gravity. Vehicle-attitude commands thus are essentially thrust-vector pointing commands. Cross-product steering therefore controls the direction of thrusting acceleration, $\underline{a}_T$, ensuring that the vector combination of properly oriented thrusting acceleration and inherent gravitational acceleration eventually nulls the guidance $\underline{V}_G$.

A sensible thrust-vector control system points the thrust vector in the general direction of $\underline{V}_G$. An autopilot rate-command signal, $\underline{\omega}_c$, proportional to $\underline{a}_T \times \underline{V}_G$,

$$\underline{\omega}_c = K(\underline{a}_T \times \underline{V}_G) \quad ,$$

causes the vehicle to rotate in a direction to align $\underline{a}_T$ and $\underline{V}_G$ as desired, and it vanishes (i.e., commands attitude hold) when the commanded alignment is achieved. This maneuver is depicted vectorially in Fig. C.1-1.

For burns that are brief relative to the orbital period, $\underline{a}_T \times \underline{V}_G$ steering proves adequate. For longer burns, where gravitational influences produce significant orbit "bending", however, a better control policy consists of aligning $-\underline{V}_G$ with $\underline{V}_G$ to permit limited—although not quite optimal—fuel economy. For burns of the Service

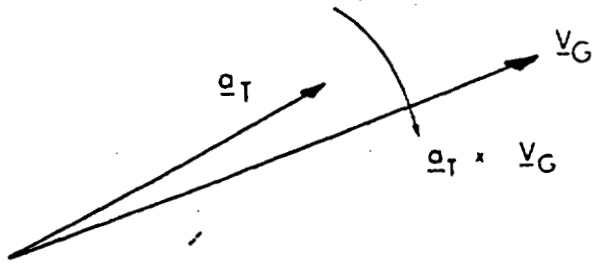

Figure C.1-1 $\underline{a}_T \times \underline{V}_G$ Steering Commands

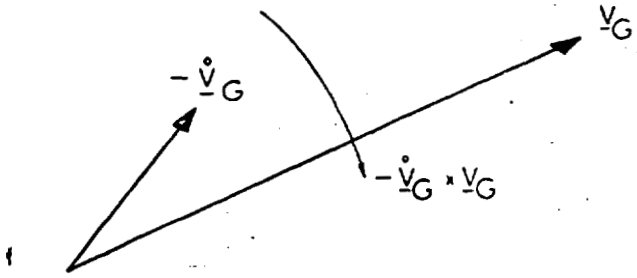

Figure C.1-2 $-\dot{\underline{V}}_G \times \underline{V}_G$ Steering Commands

Propulsion System, $\underline{a}_T$ is the dominant component of $-\dot{\underline{V}}_G$ and hence suggests the control law,

$$\underline{\omega}_c = K(-\dot{\underline{V}}_G \times \underline{V}_G) \quad .$$

The new term, $-\dot{\underline{V}}_G$, is generated as part of the guidance policy. This maneuver is depicted vectorially in Fig. C.1-2. When one considers the dominant role of $\underline{a}_T$ in the generation of $\dot{\underline{V}}_G$, the similarity of Fig. C.1-1 and Fig. C.1-2 becomes apparent.

Concern for control of attitude changes during the maneuver suggested that a linear combination of the two steering policies might take advantage both of the simplicity of $\underline{a}_T \times \underline{V}_G$ steering and of the fuel economy of $-\dot{\underline{V}}_G \times \underline{V}_G$ steering:

$$\underline{\omega}_c = K\left[(\underline{a}_T - \rho\dot{\underline{V}}_G) \times \underline{V}_G\right] \quad .$$

Implementation of this generalized cross-product steering policy is facilitated by the generation of a special $\underline{b}$ vector, closely related to $\dot{\underline{V}}_G$:

$$\dot{\underline{V}}_G = \underline{b} - \underline{a}_T \quad .$$

When the latter two equations are combined, the following generalized cross-product steering law results:

$$\underline{\omega}_c = K\left[(\underline{a}_T - c\underline{b}) \times \underline{V}_G\right] \quad .$$

This maneuver is depicted in Fig. C.1-3. Guidance supplies $\underline{V}_G$ and the $\underline{b}$ vector; Average-G accelerometer readings provide $\underline{a}_T$; c is specified for the particular maneuver; and K is again selected for a fast, stable response.

The result of the generalized cross-product steering policy is a vector rate of change of body attitude, with components lying along the three body axes. The roll component is ignored (since there is a separate roll autopilot which simply maintains zero roll rate), and the pitch and yaw components become commanded attitude rates for the pitch and yaw channels of the TVC autopilot.

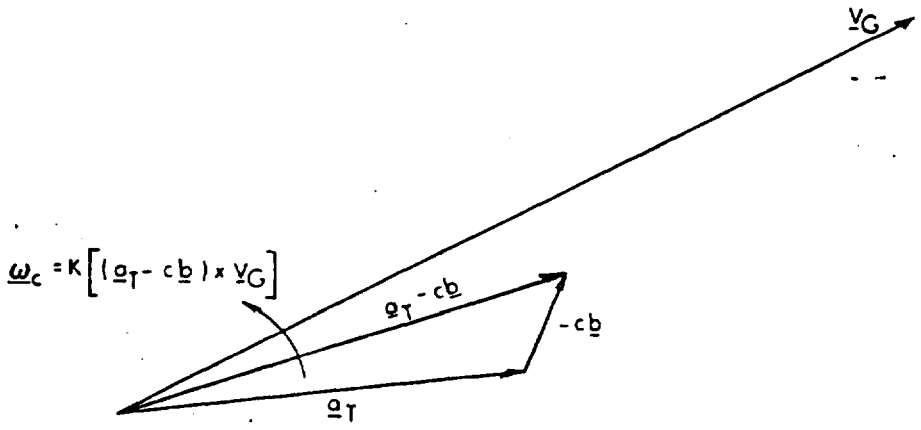

Figure C.1-3 Generalized Cross-Product Steering

C.1.2.2 Comparison of Explicit and Implicit Guidance Policies

Powered-flight guidance strategy employs two basic approaches, explicit and implicit.

Explicit guidance refers to guidance policies in which the output variables ($\underline{V}_R$, $\underline{V}_G$, $\underline{b}$, etc.) are computed from algorithms which explicitly involve present and terminal vehicle states. Implicit guidance refers to guidance policies in which the algorithms only implicitly involve vehicle states. For example, the direct computation of $\underline{V}_G = \underline{V}_R - \underline{V}$ involves $\underline{R}$, $\underline{V}$, and $\underline{R}_T$ explicitly, whereas in the integration of $\dot{\underline{V}}_G = -Q^* \underline{V}_G - \underline{a}_T$ (see below), nowhere do $\underline{R}$ and $\underline{V}$ appear explicitly. Direct $\underline{V}_G$ calculation is thus an explicit mechanization, while $\dot{\underline{V}}_G$ integration is an implicit mechanization of a $\underline{V}_G$ guidance policy.

Explicit calculation of $\underline{V}_R$ has several disadvantages. Computational times are long, the algorithms require large storage capacity, and a different set of algorithms is required for each type of maneuver. For some mission phases (e.g., TLI), no known closed-form solution exists for the steering vector $\underline{b}$, thus requiring simplified solutions involving "close-to-nominal" assumptions. With non-nominal flight (e.g., contingency aborts), such disadvantages become very severe.

The Lambert guidance policy implemented for Apollo consists of a mix of explicit and implicit forms. Between explicit $\underline{V}_G$ solutions, $\underline{V}_G$ is extrapolated by integrating the $\dot{\underline{V}}_G$ implicit-guidance equation. This permits use of a very crude $\underline{b}$ vector, and largely removes the close-to-nominal requirements that might otherwise exist.

Implicit guidance, as mentioned above, refers to guidance policies in which the output variables are computed from algorithms where the terminal state and the present state do not appear explicitly. Generally, the instantaneous required velocity is a function of present position and time,

$$\underline{V}_R = \underline{V}_R(\underline{r}, t)$$

The time derivative of V_R can be expressed by applying the total derivative law:

$$\dot{\underline{V}}_R = \frac{d\underline{V}_R}{dt} = \frac{\partial \underline{V}_R}{\partial t} + \left\| \frac{\partial \underline{V}_R}{\partial \underline{r}} \right\| \underline{V} \quad .$$

Substitution from the basic equation,

$$\underline{V}_G = \underline{V}_R - \underline{V} \ ,$$

and application of certain $\dot{\underline{V}}_R$ relationships along the solution trajectory (in particular, $\dot{\underline{V}}_R \Big|_{\underline{V}=\underline{V}_R} = \underline{G}$) yields, after some manipulation,

$$\dot{\underline{V}}_R = \underline{G} - Q^* \underline{V}_G \ ,$$

where Q^* is the 3x3 matrix of partial derivatives,

$$\left\| \frac{\partial \underline{V}_R}{\partial \underline{r}} \right\| \quad .$$

This equation states that the time rate of change of $\underline{V}_R$ along any path is equal to $\underline{G}$ (as it would be along the correct free-fall path) plus a correction if the actual path is not the free-fall path (as $\underline{V}_G \to 0$, $\dot{\underline{V}}_R \to \underline{G}$).

Substitution of this result into the time derivative of $\underline{V}_G = \underline{V}_R - \underline{V}$ yields an expression for $\dot{\underline{V}}_G$:

$$\dot{\underline{V}}_G = -Q^* \underline{V}_G - \underline{a}_T \quad .$$

This simple first-order equation can serve as a guidance policy, involving only the thrusting acceleration inputs (from the Average-G accelerometer readings) and the matrix Q^*. Nowhere do present vehicle position or velocity, present velocity required, or terminal (targeted) state appear explicitly. Terminal-state, present velocity, and present required velocity are all implicit in the current velocity to be gained. Present position is implicit in the Q^* matrix. A single initial explicit calculation of $\underline{V}_R$ to obtain the initial $\underline{V}_G$ is required. From then on the equation for $\dot{\underline{V}}_G$ is self-sufficient, requiring only the external generation of the elements of Q^*. Implicit

guidance is often referred to as $\dot{\underline{V}}_G$ guidance, whereas explicit guidance, as mentioned above, is called $\underline{V}_G$ guidance. All necessary inputs to the cross-product routine are readily available: $\underline{V}_G$ from the integration of the $\dot{\underline{V}}_G$ guidance equation, $\underline{a}_T$ from the Average-G accelerometer readings, and $\underline{b} = \dot{\underline{V}}_G + \underline{a}_T$.

In general, the Q^* elements are functions of position and time. Where flight is close to some nominal trajectory, the position dependence can be removed, and the elements can be represented simply as functions of time. Further simplifications can be made, such as using two or three straight-line segments with perhaps one or two discontinuities; using simple straight lines; using simple constants; and even using zero for selected elements.

The advantage of implicit guidance is its computational simplicity. One disadvantage is the requirement for much preflight simplification and approximation to determine adequate representations for the Q^* elements. Another disadvantage is the severe restriction to near-nominal flight—greater here than it was for explicit guidance.

The interrelationships of explicit/implicit guidance and cross-product steering should be noted. Initial $\underline{V}_R$ and $\underline{V}_G = \underline{V}_R - \underline{V}$ calculations are made in both the explicit and implicit approaches. Thereafter, only in the explicit case are periodic resolutions for instantaneous $\underline{V}_R$ and $\underline{V}_G$ required. Implicit guidance simply integrates the $\dot{\underline{V}}_G$ equation. To do so, however, periodic updates of the Q^* matrix (partial derivatives,

$$\left\| \frac{\partial \underline{V}_R}{\partial \underline{r}} \right\| ,$$

functions of time and position) are required. Note in the $\dot{\underline{V}}_G$ expression that definition of a $\underline{b}$ vector,

$$\underline{b} = -Q^* \underline{V}_G ,$$

yields an equation

$$\dot{\underline{V}}_G = \underline{b} - \underline{a}_T ,$$

which was introduced above, during cross-product steering discussions. Both explicit and implicit guidance laws use the cross-product steering law. For implicit guidance, the $\underline{b}$ vector used by cross-product steering is already available, appearing as the $-Q^*\underline{V}_G \triangleq \underline{b}$ in the $\dot{\underline{V}}_G$ equation; no additional computations are required. For explicit guidance, however, the cross-product steering requirement for $\underline{b}$ presents an added computational burden; analytic expressions for $\underline{b}$ must be implemented, solely for inputs to cross-product steering.

There are two ways to circumvent the computational burden, however. First, one can revert to simple $\underline{a}_T \times \underline{V}_G$ steering (i.e., by setting $c = 0$ in the generalized cross-product steering law). Then, the value of $\underline{b}$ is immaterial and need not be computed. The implementation of Lambert-targeted ASTEER guidance employs such a simplification (see Section C.1.2.4). Fuel penalties for the non-optimal $c = 0$ are minimal in the situations where ASTEER is used.

Alternatively, the $\underline{b}$ vector can be obtained with adequate accuracy by a simple numerical differentiation using a first-order back difference. If the equation for the $\underline{b}$ vector ($\underline{b} = -Q^*\underline{V}_G$) is substituted into the equation for $\dot{\underline{V}}_R$ ($\dot{\underline{V}}_R = \underline{G} - Q^*\underline{V}_G$), the following expression for $\underline{b}$ results:

$$\underline{b} = \dot{\underline{V}}_R - \underline{G}$$

A simple first-order back difference of $\underline{V}_R$ is suggested:

$$\dot{\underline{V}}_R(\tau) = \frac{\underline{V}_R(\tau-\Delta\tau) - \underline{V}_R(\tau)}{\Delta\tau}$$

The $\underline{b}$-vector equation then becomes simply:

$$\underline{b}(\tau) = \frac{\underline{V}_R(\tau) - \underline{V}_R(\tau-\Delta\tau)}{\Delta\tau} - \underline{G}(\tau)$$

Note that $\underline{G}(\tau) = G[\underline{R}(\tau)]$ now involves present vehicle state explicitly (obtained from Average G). Time- and memory-consuming explicit $\underline{b}$-vector calculations are no longer required. Preprogramming and close-to-nominal trajectory restrictions inherent in the explicit $\underline{b} = -Q^*\underline{V}_G$ implementation are avoided. There are some

geometries for which the above equation gives degraded $\underline{b}$ vectors, but generally the results are acceptable.

The final form of powered-flight guidance is a combination of explicit and implicit guidance laws. Periodically, the explicit calculation of $\underline{V}_G = \underline{V}_R - \underline{V}$ is made; this involves direct computation of $\underline{V}_R$ (e.g., via the iterative Lambert subroutine, discussed below). The necessary velocity computations are slow, requiring several seconds (2 to 4 normally, but 10-sec cycles have been observed). During the explicit computational interval, implicit guidance is used to maintain a current estimate of $\underline{V}_G$ for steering. The required $\underline{b}$ vector is obtained from simple back differences of $\underline{V}_R$, with local gravitation taken into consideration.

C.1.2.3 Lambert Powered-Flight Guidance

Lambert powered-flight guidance solves the problem of reaching a specified target position at a specified intercept time. The period of free-fall coasting flight which follows the termination of the active guidance phase may last a matter of hours (one or more orbits in rendezvous) or days (during return to earth following a translunar orbit or a transearth midcourse correction).

Computations for Lambert powered-flight guidance can be done either onboard the spacecraft or telemetered from the ground. Lambert targeting (TPI or TPM) generates the target vector, the intercept time, the time of ignition, and various other parameters and displays. Return-to-earth targeting generates these same Lambert-type inputs for the Lambert guidance. Guidance prepares the necessary trajectory-control variables during the active transfer phase between the two periods of coasting flight. The $\underline{V}_G$ is nulled either manually (via the astronaut's hand-controller activity) or automatically (via the cross-product steering law).

Lambert powered-flight guidance involves both explicit and implicit guidance policies. Briefly, the explicit task involves the direct calculation of required velocity, $\underline{V}_R$, and velocity-to-be-gained:

$$\underline{V}_G = \underline{V}_R - \underline{V} ,$$

where $\underline{V}$ is the current velocity. The implicit task involves the extrapolation of $\underline{V}_G$ during the explicit $\underline{V}_G$ computation interval.

Lambert's theorem states that the time of flight (t_f) along the solution trajectory depends only upon the length of the semimajor axis, a, of the solution conic (valid for general conics); the sum of the distances to the initial (r_1) and final (r_2) points of the arc from the center of force; and the length of the arc (c) joining these points:

$$t_f = t_f(a, r_1 + r_2, c) \quad .$$

Time of flight is the known specified quantity:

$$t_f = t_2 - t_1 \quad .$$

The task of the Lambert subroutine is to determine iteratively the solution conic, in particular the parameter a. Required velocity is then simply the velocity associated with the solution conic at the present position:

$$\underline{V}_R = \underline{V}_R(\underline{R}_1, a) \quad .$$

The present velocity to be gained is simply the difference between present velocity required and present orbital velocity. The $\underline{b}$ vector, required only for the steering task, is computed by back-differencing $\underline{V}_R$ and accounting for local gravitation, as described above.

C.1.2.4 Lambert ASTEER Guidance

Lambert ASTEER guidance is an explicit guidance policy which solves the problem of reaching a specified target vector at approximately a specified time. ASTEER guidance utilizes a single initial solution of Lambert's theorem by the Lambert iterator, e.g., for the semimajor axis, a, of the desired coasting-flight conic:

$$t_f = f(a, \underline{R}_1, \underline{R}'_T) \quad .$$

Since the variation of a during the brief rendezvous powered maneuvers directed by the AGC on the LM is negligible, a can be assumed constant (leading only to minor error in the time of actual intercept). As in the nominal Lambert case discussed

in the previous section, Lambert targeting provides the offset target vector, $\underline{R}_T$, and the time of flight, $t_f = t_T - t_1$. The initial $\underline{V}_{R_1}$ is computed explicitly for the initial a:

$$\underline{V}_{R_1} = f(a_1, \underline{R}_1, \underline{R}'_T) \quad .$$

Thereafter, explicit $\underline{V}_R(\tau)$ updates are computed as functions only of $\underline{R}(\tau)$, using the initial a:

$$\underline{V}_R(\tau) = f(a_1, \underline{R}_\tau, \underline{R}'_T) \quad .$$

Initial and updated explicit $\underline{V}_G$'s are computed via the equation,

$$\underline{V}_G(\tau) = \underline{V}_R(\tau) - \underline{V}(\tau) \quad .$$

By assuming constant $a = a_1$, the Lambert-iterator resolutions can be avoided and time saved, thus permitting $\underline{V}_G$ updates every 2-sec Average-G steering cycle.

ASTEER avoids the steering $\underline{b}$-vector problem simply by setting $\underline{b} = \underline{0}$. This is convenient from the vantage of avoiding unnecessary $\underline{b}$-vector computation time, and empirically justified for those situations where ASTEER is used (e.g., short burns such as LM-active rendezvous maneuvers, where $\underline{a}_T \times \underline{V}_G$ steering might be expected to work, anyway).

C.1.2.5 External ΔV Powered-Flight Guidance

External ΔV powered-flight ($\dot{\underline{V}}_G$) guidance is an implicit guidance policy designed for ground targeting of lunar-orbit insertion (LOI), descent-orbit initiation (DOI), transearth injection (TEI), midcourse correction and abort maneuvers and for onboard targeting of certain rendezvous maneuvers. External ΔV was implemented for several reasons:

a. The explicit $\underline{V}_R$ and $\underline{b}$ computations planned originally for the LOI, DOI, and TEI maneuvers required extensive fixed- and erasable-memory storage, and consumed too much computer time.

b. External ΔV is a simple, fast, reliable guidance/steering mode, compared

to the slower—and more complex—Lambert and ASTEER modes. Considerable experience with external ΔV maneuvers had been gained during the Gemini program, whereas Lambert maneuvers were a new development.

c. External ΔV is characterized by an inertially-fixed burn attitude, which simplifies the problems of burn monitoring and manual takeover in the event of partial or complete primary-system failure.

d. Reduced computational burdens permit more time for nominal activity (displays, autopilots, compensation, etc.) and extended verb activity.

e. External ΔV is conceptually simpler, e.g., for crew-originated maneuvers. Fewer inputs are required, with fewer demands made upon the crew.

Disadvantages of External ΔV are a fuel penalty on long burns and the requirement for ground targeting and uplink communication of the targeting parameters—time of ignition and initial $\underline{V}_G$ (in local vertical coordinates); but these disadvantages are offset by the tremendous savings in computer-memory capacity which an onboard targeting program would have required. Voice-linked loads of planned and contingency maneuvers are also transmitted in case of a communications failure. Fuel penalties, even during the long LOI, TEI, and translunar aborts, are minimal.

As previously mentioned, External ΔV is an implicit $\dot{\underline{V}}_G$ policy. There is no reference to a required velocity; targeting supplies the initial $\underline{V}_G$ direction, and $\underline{V}_G$ is extrapolated from

$$\dot{\underline{V}}_G = \underline{b} - \underline{a}_T ,$$

discussed above. The $\underline{b}$ vector for External ΔV maneuvers is identically zero, meaning that $\underline{V}_G$ is affected only by thrusting accelerations.

A unique feature of External ΔV prethrust computations is an approximate compensation for finite burn times, in which the in-plane component of the initial $\underline{V}_G$ is rotated in the direction of the angular-momentum vector, by half the predicted central angle of the burn. Current vehicle mass and a nominal thrust level (2- or 4-jet RCS, or the SPS) are also used in the calculation. No attempt is made to account for mass variation during the burn.

C.1.2.6 Thrust-Cutoff Sequencing

Automatic termination of thrust is provided for powered-flight maneuvers, since the time constants of computation, display and astronaut-reaction would result in an unacceptable ΔV for a manual termination. The guidance/steering calculations generate an estimate of time-to-cutoff, which is then displayed for burn-monitoring purposes. When time-to-cutoff drops below four sec, the steering commands are set to zero, the steering and time-to-cutoff calculations are disabled, and a task is scheduled to issue the engine-cutoff discrete at the proper time. The guidance $\underline{V}_G$ calculations continue (in case they should be needed for manual $\underline{V}_G$ trimming) following engine shutdown.

The time-to-cutoff calculations involve the present $\underline{V}_G$, the just-measured $\Delta\underline{V}_a$ (Average G, Section C.1.1), and the most recent $\underline{b}$ vector (identically zero for ASTEER and External ΔV steering). As depicted in Fig. C.1-4, the desired effect is the minimization of cutoff $\underline{V}_G$ magnitude, assuming that the direction and magnitude of $-\dot{\underline{V}}_G$ are constant.

The equation for time-to-cutoff (TTC), measured from the time of the accelerometer reading associated with $\underline{V}_G$ and $\dot{\underline{V}}_G$,

$$TTC = \frac{\underline{V}_G \cdot \mathrm{UNIT}(-\dot{\underline{V}}_G)}{\dot{V}_G}\left[1 - \frac{\underline{V}_G \cdot \mathrm{UNIT}(-\dot{\underline{V}}_G)}{2V_E}\right] - \Delta T_{TO} \quad ,$$

contains a first-order variable-mass approximation (V_E is the exhaust velocity) and a biasing impulsive ΔV tailoff time, ΔT_{TO}. The vector dot product represents the projection of $\underline{V}_G$ onto the current $-\dot{\underline{V}}_G$ direction.

Cutoff residuals comprise the cross-axis component shown in Fig. C.1-4, plus axial cutoff errors due to the approximations associated with the TTC algorithm. As cutoff approaches, the approximations become quite good $\left(\underline{b} \rightarrow \underline{0} \text{ as } \underline{V} \rightarrow \underline{V}_R, \text{ and } V_E >> \underline{V}_G \quad \mathrm{UNIT}(-\dot{\underline{V}}_G)\right)$; axial cutoff errors are therefore on the order of the theoretical maximum of two accelerometer pulses, provided that engine parameters (V_E, constant thrust, ΔT_{TO}, etc.) are modeled accurately.

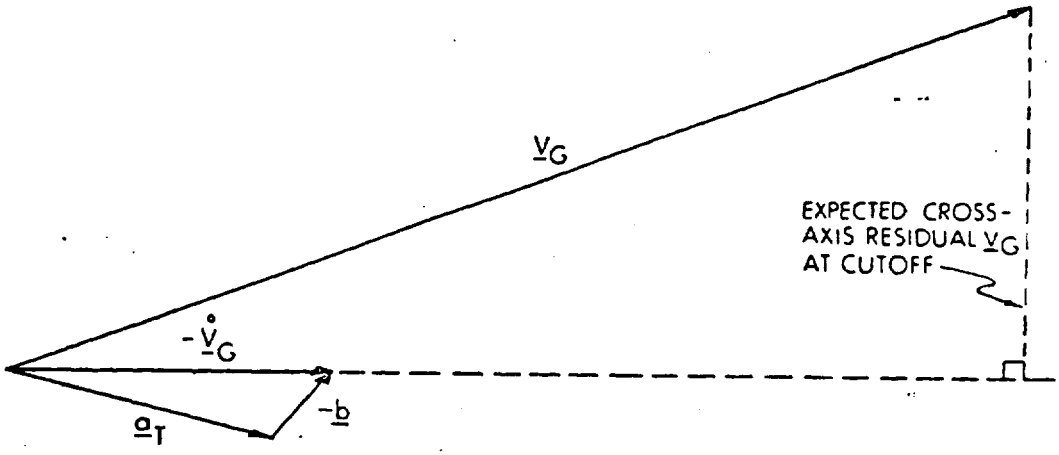

Figure C.1-4 Time-to-Cutoff Geometry

The cross-axis component of residual $\underline{V}_G$ at cutoff, however, can be quite large, especially for short burns during which the steering/autopilot loop dynamics prevent recovery from initial vehicle and thrust-pointing errors, the dynamics associated with a moving center of gravity, or $\underline{b}$-vector dynamics in a severe-burn geometry. However, the post-burn trimming maneuver can accommodate these situations.

C.2 Thrust Monitor Program

The thrust monitor program monitors and displays velocity changes applied to the spacecraft during manual maneuvers and during maneuvers not controlled by the GN&C System. The program initially suspends state-vector updating by the VHF range-link (i.e., resets the update and track flags), and advances the vehicle-state vector to the current time. This operation continues until the state vector is extrapolated ahead to the ignition time. The Average-G routine discussed in the previous section is then initiated, thus allowing the earliest possible ignition. Average G remains in operation until the program is terminated upon completion of the maneuver. The primary output of the thrust-monitor program is the measured maneuver $\Delta\underline{V}$, expressed in vehicle coordinates.

The thrust-monitor program is normally employed during two major maneuvers—the translunar-injection maneuver controlled by the Saturn guidance system and the manually-controlled terminal rendezvous maneuvers required for a CSM retrieval of the LM. During the TLI maneuver, the astronaut can call up the display of inertial velocity altitude and altitude rate. During the active CSM terminal rendezvous maneuvers, relative range, range rate, and the angle of the vehicle X axis relative to the horizontal plane are available for display.

C.3 Earth-Orbit Insertion Monitor Program

The AGC program for earth-orbit insertion monitors the performance of the Saturn launch vehicle from the detection of the liftoff discrete to the accomplishment of earth orbit. In the event of a Saturn inertial-platform failure during boost, however, the Command Module AGC provides a backup capability to guide the Saturn along a prescribed trajectory which is discussed in some detail in Section D.5.

Since a close relationship must exist between the guidance and navigation equipment on both the Saturn launch vehicle and the Command Module, it is important to detail that relationship, beginning prior to launch.

As explained in Section 2.2.2.1, two sets of guidance equipment are prepared for launch: The Saturn guidance equipment in the Saturn Instrument Unit controls the launch vehicle, while the Apollo guidance equipment in the Command Module provides a monitor of Saturn guidance during launch. The Saturn and CM are both inertially aligned to a common vertical and launch azimuth reference. During countdown, both systems are gyro-compassed to an earth-frame reference. Near liftoff, they both respond to discrete signals to switchover from the earth reference to the nonrotating inertial reference used during boost.

Prior to liftoff, the AGC interrogates the liftoff discrete every half-second, and when liftoff is detected, the computer initiates the earth-orbit insertion program. Should the discrete fail or malfunction, the astronaut can initiate the program. At liftoff, the current clock reading updates the universal time previously stored in the AGC, and the AGC clock itself is then zeroed.

During first-stage flight, the Saturn Instrument Unit controls the flight by swiveling the engine's outer four rockets. For inertial vertical flight, the vehicle is rolled from its launch azimuth to the flight-path azimuth. Saturn guidance then controls the vehicle in an open-loop, preprogrammed pitch maneuver designed to pass safely through the critical period of high aerodynamic loading.

Throughout this period, the AGC's powered-flight navigation routine (Average G, Section C.1.1) calculates the position and velocity of the Saturn vehicle. This routine is initialized with the position and velocity at liftoff, computed using knowledge of the earth-rotation vector, observed time of launch, and preloaded values of launch-site longitude and geodetic latitude. (To relate the stable-member (SM) orientation at liftoff to the basic reference-coordinate system, the direction of Z_{SM} is calculated first. The Z_{SM} axis, which is aligned along the true direction of gravity at the launch pad, defines the astronomical latitude. The astronomical latitude is the derived angle between the zenith and the equator, measured by the observed angle between the horizon and the polar axis. The geodetic latitude, which is also defined as the angle between the zenith and the equator, is calculated on the basis

of the earth's reference ellipsoid. Since the astronomical latitude is approximately equal to the geodetic latitude at the launch pad, the geodetic latitude is substituted in the numerical calculation of the direction of gravity.)

The earth's polar axis is described in reference coordinates for each space mission by storing in the AGC the quantities necessary to determine precession and nutation. On this basis, the Z_{SM} may also be computed in the reference system.

The CMC assumes that X_{SM} and Y_{SM} are in the plane of the horizon at the launch site. During prelaunch, the X_{SM} is kept aligned with the launch azimuth, an angle in the horizon plane measured positive from North in an easterly direction to X_{SM}. Since North and South are determined by the line of intersection between the horizon plane and the plane determined by the earth's polar axis and the zenith, X_{SM} can be computed in reference coordinates. X_{SM} and Z_{SM} in reference coordinates determine Y_{SM} in reference coordinates, thus allowing the transformation (REFSMMAT) from earth-reference to stable-member coordinates.

Both the Saturn and CM guidance systems continuously measure vehicle motion and compute position and velocity. In addition, the GN&C System compares the actual Saturn trajectory with that to be expected from the AGC, using a sixth-order polynomial approximation—thus generating an attitude-error display for the crew. During boost (first-stage only), these attitude errors are available to the Saturn Instrument Unit for nulling, should a takeover be required. Should a takeover be required during second or third-stage boost, the only commands the Instrument Unit can receive are from the Rotational Hand Controller which the astronaut manipulates as he compares DSKY displays of velocity, altitude and altitude rate with a nominal trajectory profile available on printed card. In this fashion the spacecraft can be flown safely into earth orbit with relatively minor errors in apogee and perigee.

C.4 Entry Guidance and Mission Control Programs

The phase of the mission beginning with the jettison of the Service Module by the Command Module and ending with safe arrival at the designated landing site comprises three major control functions—entry guidance, mission control programs, and Entry Digital Autopilot. (For a summary of the entry phase, see Section 2.2.2.14.)

The entry guidance directs the CM to a safe return at the designated landing site. The mission control programs (P61 through P67) inform the crew of their location along the entry trajectory, as determined by the Entry Guidance. The Entry DAP uses the output of the entry guidance (i.e., roll command) to perform automatically all maneuvers necessary for reentry.

This section of Appendix C discusses entry guidance and mission control programs. The Entry DAP is discussed in Appendix D, which describes all of the Digital Autopilots.

C.4.1 Entry Guidance

A spacecraft returning from a lunar mission reenters the earth's atmosphere at velocities exceeding escape velocity. To cope with the sensitive dynamics involved, automatic entry guidance and control are employed. Although escape velocity is only 40 percent greater than orbital velocity, the sensitivity of the reentry range capability to variations in the critical entry flight-path angle* is several orders of magnitude greater for entry from lunar missions than from earth-orbital missions. The automatic guidance system is expected to provide quick response and to minimize the effects of variation of the actual environmental characteristics (such as CM aerodynamics and the atmosphere) from the design standards.

The automatic entry-guidance system was designed with two main objectives, having the following priority: (1) a safe return to the earth's surface, and (2) landing-point control. For the CM to splashdown successfully at the designated landing site, these objectives must be met. Safe return means that the deceleration during reentry should never exceed a prescribed limit, nor should the spacecraft skip back out of the atmosphere at greater than orbital velocity. The reentry trajectory may include a free-fall, ballistic lob portion out of the atmosphere in order to reach a distant landing point, but this must be done at subcircular velocity. The midcourse guidance phase has the initial responsibility for a safe return in that the spacecraft must be steered into an acceptable entry corridor from which a safe return is possible.

*The entry flight-path angle is defined as the angle between the inertial velocity vector and local horizontal at the entry interface altitude of 400,000 ft above the reference geoid.

The entry guidance must achieve the objective of range-control without interfering with the objective of a safe return. Acceptable reentry-angle values (about the nominal) define a region called the reentry corridor. A nominal entry-angle value for a lunar mission is approximately -6.5 deg. (A nominal entry-angle value for an earth-orbital mission is approximately -1.5 deg.)

To fulfill its main objectives, the entry guidance must cope with errors. The distinction is made between navigation errors and steering errors.

Navigation errors are inaccuracies in the determination of the spacecraft's own position and velocity. These errors cannot be removed by guidance design. It is sometimes convenient to think of a navigation error as the error in where the spacecraft thinks the target is. Causes of navigation errors include:

(1) Errors in the indicated initial position and velocity at the start of reentry.
(2) Initial misalignment of the IMU.
(3) IMU gyro and accelerometer errors.

Steering errors, on the other hand, depend on guidance design and represent the spacecraft's inability to reach the position where it thinks the target is. In a well-designed system, the expected miss-distance should be approximately the same as the expected navigation error near the end of the reentry trajectory. A subtle type of steering error involves control actions which are taken early in entry, while the velocity is still supercircular. The danger is that during this phase, while the sensitivities are high, control actions based on navigation errors may be improper and result in a large enough trajectory deviation such that later control actions are incapable of compensating sufficiently for the early mistakes. Thus, steering errors are a function of navigation errors. Specifically, the chief troublemaker is the error in indicated rate-of-climb during the supercircular phase. The effect of this error is minimized by the use of a reference trajectory during the pull-up portion, Upcontrol. On the other hand, navigation errors are not a function of steering errors or the steering scheme, except indirectly.

A controlled-entry flight to the designated landing point is achieved by taking advantage of the aerodynamic lift capabilities of the CM. The CM is a wingless, axially symmetric, reentry vehicle constructed so that its center of gravity is

displaced from its axis of symmetry. When flying hypersonically in the atmosphere, it trims with a low, constant ratio of lift to drag. The only means for perturbing the trajectory in a controlled manner is to roll about the wind axis (velocity vector) with reaction-control jets, permitting the lift vector to be pointed anywhere in the plane perpendicular to the wind axis. The roll angle defines the orientation of the lift vector relative to the trajectory plane—i.e., the plane containing the wind axis and the position vectors. The component of lift in the trajectory plane is the means of trajectory control for the down-range flight, while the component of lift out of the trajectory plane is for control of cross-range flight.

The entry-guidance equations regularly compute the desired lift direction currently necessary if the trajectory is to reach the designated landing site. In actual flight, the Entry DAP causes the CM to roll about the wind axis so that the actual lift direction is forced into agreement with the desired lift direction. This results in achieving the desired in-plane component for down-range control. The rolling maneuver yields an out-of-plane component of lift that is used for lateral steering. After the cross-range error is removed, the lateral lift component is an unwanted by-product of the steering, and its effect is constrained to an acceptably small value by the guidance, which causes the CM to roll periodically so as to reverse the sign of the lateral lift and null out lateral drift. Since the in-plane component is the fundamental controlled quantity, in that it controls down-range flight, its sign normally remains unchanged during the nulling process (lateral switching). In effect, the restriction on the sign of the in-plane component during lateral switching demands that the CM roll through the smaller of the two possible angles—through the so-called shortest path—when lateral switching takes place; the Entry DAP normally causes this type of maneuver. But in certain instances where such a maneuver would cause the vehicle to fall short of the target, the entry guidance insists on rolling over the top without regard for the angle size, and a program switch is set to inform the roll DAP that the CM is being commanded to roll through the larger angle. In summary, the entry guidance continuously provides a single steering quantity—the commanded roll angle—that is to be achieved by the CM.

Entry guidance is further described by using the illustrative deceleration profile of Fig. C.4-1. At point 1, the CM lift is initially directed upward for nominal-to-steep entry angles; for shallow angles, lift is directed downward. A function of the guidance is to establish the initial lift by roll-angle command, which

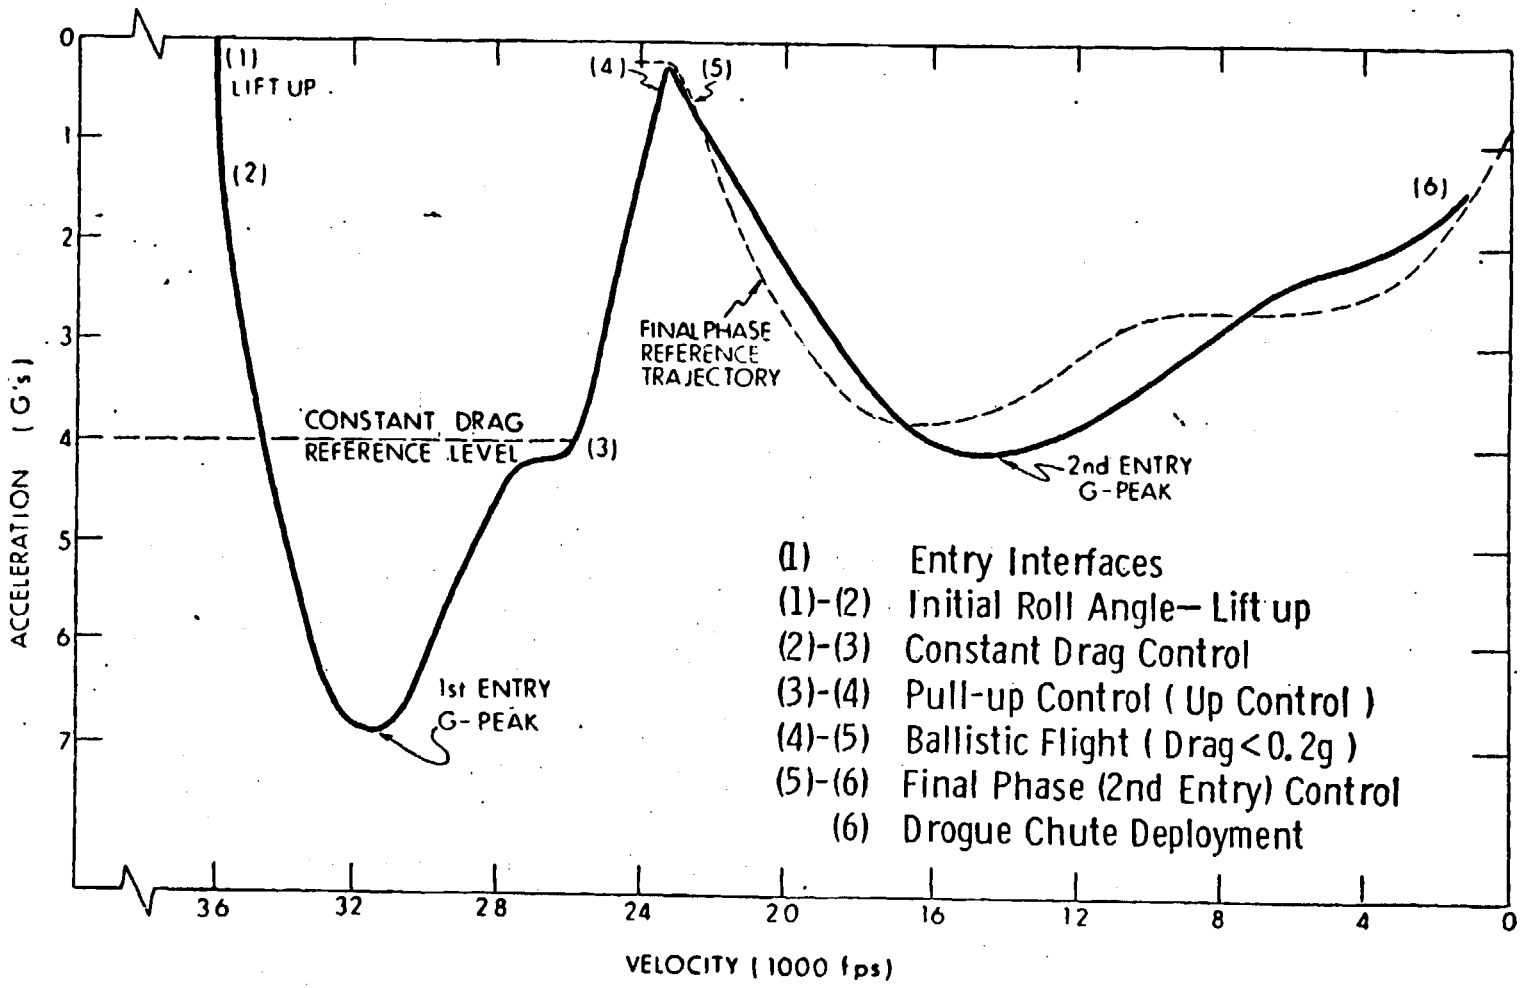

Figure C. 4-1 Typical Deceleration Profile

is held constant until the CM is in the atmosphere and the deceleration level reaches 0.05g. Prior to 0.05g, the Entry DAP holds the CM in the attitude of hypersonic trim at the specified roll angle. After 0.05g, the atmosphere holds the CM in trim, with rate-damping assistance in pitch and yaw by the DAP. At this point, the guidance begins the lateral-range calculations which will permit a small (15 deg) roll-angle deviation. The g-level continues to increase until about 1.3g, when constant drag control begins—attempting to maintain a constant deceleration (point 2). Velocity and range-to-go decrease. When the rate of descent reaches 700 ft/sec, a trajectory-search, in addition to constant drag, begins to determine if the predicted range for a constant lift-to-drag ratio (L/D) flight from the spacecraft's current location would yield the desired range. The predicted range is made up of analytic expressions based on a candidate reference trajectory consisting of segments for pull-up, ballistic lob and final (second-entry) phase. If the predicted range is not within 25 nmi of the desired range, the constant drag control continues to be flown. When the predicted range is within 25 nmi of the desired range, the guidance begins to steer the CM along the pull-up reference (point 3). When the terminal conditions of the Upcontrol reference are met, a ballistic lob is flown if the drag becomes less than 0.2g (point 4). The guidance maintains a roll command constant at the most recent value. Should the drag become less than 0.05g, the roll command is set to zero and, in addition, the inadequate aerodynamic stability in this low-dynamic pressure region requires that pitch and yaw hypersonic trim again be maintained by the three-axis DAP. On reaching the peak altitude, the CM falls back toward the earth and reenters the atmosphere at point 4. At 0.05g, the DAP relinquishes. The final phase of guidance is entered when the acceleration builds up and exceeds a 0.2g threshold. The final-phase guidance (point 5) steers the CM to the landing site, using a prestored reference trajectory. Short-range trajectories omit the ballistic and Upcontrol phases; very long-range trajectories omit the constant-drag phase.

Two remaining functions of entry guidance are the g-limiter logic, which modifies the roll commands during the final phase to avoid exceeding the maximum 8g limit, and the lateral logic, which periodically switches the desired lift direction from one side of the trajectory plane to the other to allow the roll angle to control both down-range and lateral range.

Three entry-guidance functions are used throughout the trajectory—Navigation, Targeting and Mode Selector. The first and second of these functions begin in Program P61; the third begins in P63. Navigation updates the position and velocity vectors using acceleration data. Targeting computes the current landing site vector based on earth rotation during an estimated flight time and calculates the present range and lateral range for the landing site. Mode Selector selects each entry-guidance phase on the basis of the current position along the entry trajectory.

C.4.2. Entry Mission Control Programs

The entry trajectory is broken into major segments, each identified with a mission program number, such as P61, shown in Fig. C.4-2 for both orbital and lunar entry. After the crew selects P61, the programs run in numerical sequence. Selection of the next major mode is automatically determined by the mission control program prior to beginning P63 and by entry guidance thereafter. While each phase of the entry guidance is operating, the DSKY displays the corresponding mission control program, P61 through P67, to inform the crew of the spacecraft's location along the entry trajectory. The mission control programs form the framework within which the crew is able to monitor the phases of flight beginning with the maneuver to CM/SM separation attitude and continuing until droge-chute deployment.

The mission control programs for entry are summarized in the following sections:

C.4.2.1 Entry Preparation, P61

The Entry Preparation Program starts the entry sequence for the CM. P61 initiates onboard navigation and checks the IMU alignment. Next, it calculates data so that the crew can initialize the Entry Monitor System (EMS)* and also monitor certain future trajectory check points. The data from the P61 displays are also compared with the ground-supplied EMS data and serve as a check on the operation of the AGC. On completion of the displays, P62 is begun.

---

*The EMS is a backup system used to monitor entry performance; it provides a roll-attitude indication, a range-to-go counter and a velocity and acceleration plot to compare with acceptable velocity/acceleration profiles.

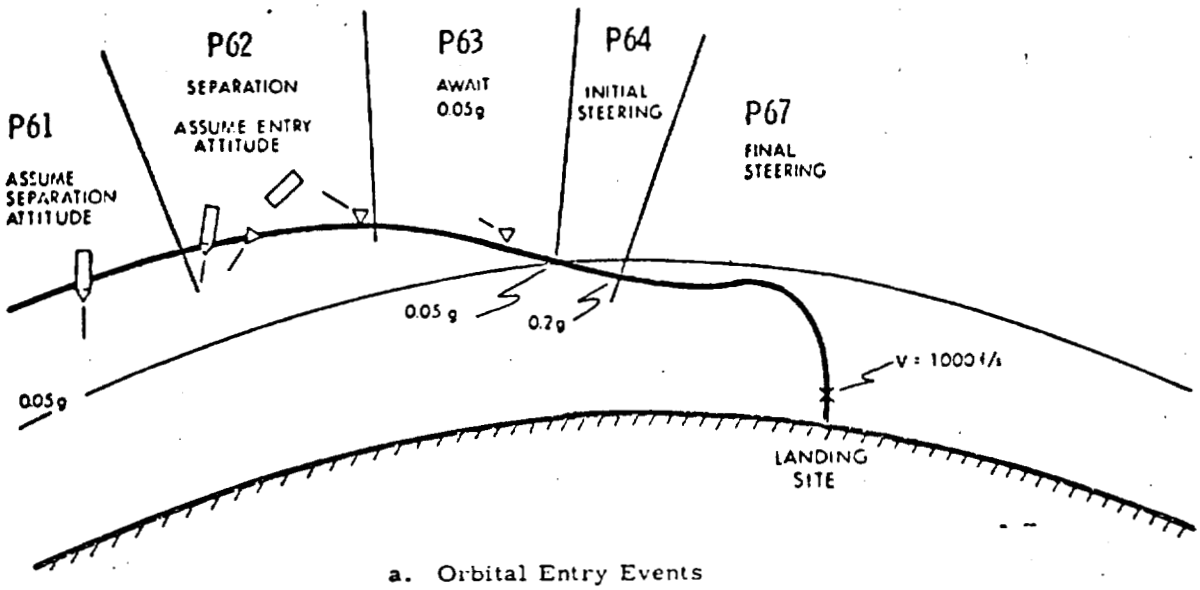

a. Orbital Entry Events

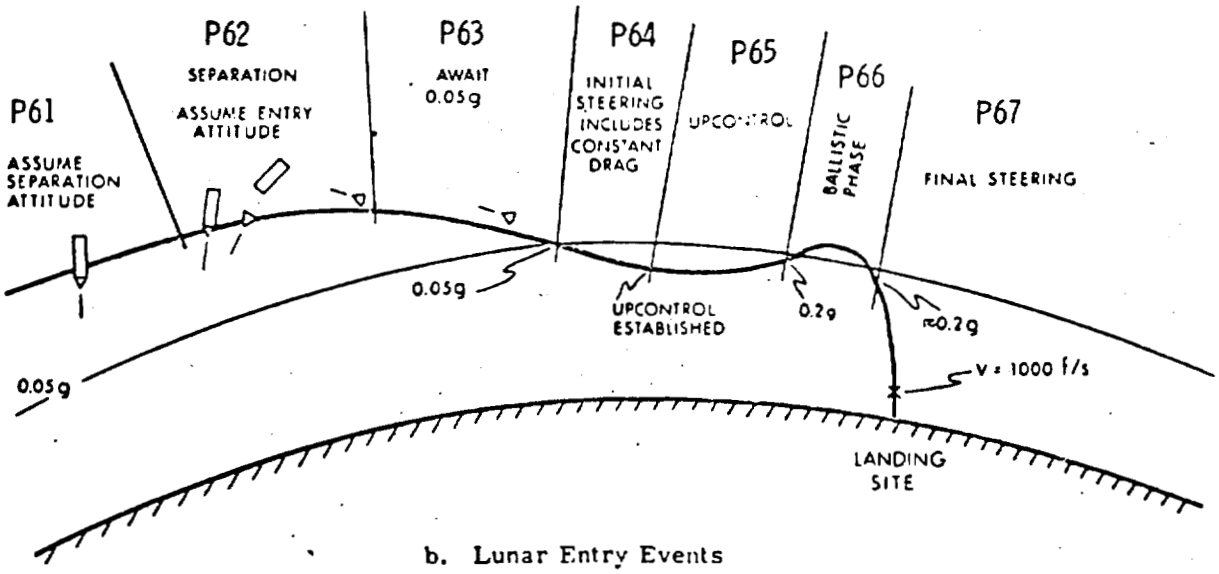

b. Lunar Entry Events

Figure C. 4-2 Typical Entry Events along Trajectory

C.4.2.2 CM/SM Separation and Pre-entry Maneuver, P62

The CM/SM Separation and Pre-entry Maneuver Program establishes the trajectory-monitoring displays for the crew, accepts crew notification that CM/SM separation has occurred, activates the Entry DAP, and orients the CM to the correct attitude of hypersonic trim. When the DAP first becomes active, it establishes attitude hold, following separation, until the crew verifies the correctness of the display of initial roll-attitude specification and target location. After verification, the DAP initiates and performs the maneuver to bring the CM into entry hypersonic-trim attitude at the specified roll attitude.

When the ENTRY DAP determines that the maneuver to entry attitude is essentially completed, P63 is begun.

C.4.2.3 Entry Initialization, P63

The Entry Initialization Program activates the entry guidance and establishes the pre-0.05g guidance-monitoring display. During P63, the DAP continues to hold the CM to the hypersonic trim with respect to the computed relative wind axis, as the CM approaches the atmosphere. When the guidance senses that the atmospheric drag level has exceeded 0.05g, P64 is begun.

P63 is the last P60 mission control program (until P67) to actually perform a program-sequence function. By comparison, P64, P65, P66, and the initial portion of P67 are merely milestones established by entry guidance along the trajectory.

C.4.2.4 Post-0.05g, P64

The Entry Post-0.05g Program establishes the guidance-monitoring display for the crew, initiates the lateral guidance and performs the trajectory planning function by finding a candidate reference trajectory that is compatible with the range to be flown. The guidance decides that the current entry is from an earth orbit if the velocity is less than 27,000 ft/sec at 0.05g (the start of P64); lateral steering is activated and, when the drag exceeds 0.2g, P67 is begun. The guidance decides that the entry is from lunar orbit if the velocity is greater than 27,000 ft/sec at 0.05g; lateral steering is activated, the trajectory search is begin, and the guidance

directs the CM to fly to a constant drag level, about 4g. The guidance looks for a reference trajectory that satisfies the range requirement, using some or all of the trajectory segments: Upcontrol phase, ballistic phase, and final phase. Throughout P64, the DAP follows roll commands issued by entry guidance, and performs coordinated rolling with rate-damping pitch and yaw. Except as mentioned in P66, the DAP continues this behavior throughout the flight. When constant-drag flight has reduced the predicted range to within 25 nmi of the desired range, P65 is begun. If an Upcontrol reference trajectory is not found, for which the terminal velocity exceeds 18,000 ft/sec, then P67 is begun.

C.4.2.5 Upcontrol, P65

The Entry Upcontrol Program establishes the guidance-monitoring displays for the crew. P65 is flown for all lunar entries except for those where the landing-site range is short. During P65, the guidance steers the CM along the pull-up reference trajectory toward a possible controlled ballistic lob. When the termination conditions of the Upcontrol reference are met, the phase is over:

a. When the drag acceleration becomes less than the terminal drag condition, P66 is begun.

b. If the rate of descent becomes negative while the velocity has nearly reached the terminal velocity conditions, then pull-up flight is ended and P67 is begun.

(For intermediate-range trajectories, the terminal drag condition may be greater than 0.2g, and the net effect of satisfying the terminal drag condition is a direct transfer to P67 by way of P66.)

C.4.2.6 Ballistic, P66

The Entry Ballistic Program establishes the guidance-monitoring displays for the crew. P66 signifies the ballistic lob portion of the trajectory and lasts as long as the atmospheric drag level is less than 0.2g. Because of the rarified atmosphere, the guidance ceases to produce new values for the desired lift-vector direction. While the drag level remains greater than 0.05g, the CM desired roll attitude is the most recent roll command (from P65); pitch and yaw hypersonic

trim are still maintained by the atmosphere. If the drag level becomes less than 0.05g, pitch and yaw attitude control by the DAP is activated to maintain the CM in hypersonic trim (as during P62 and P63); in addition, a wings-level, lift-up roll attitude is established for the CM by providing a zero roll-angle command. When the drag again exceeds 0.05g, the DAP reverts to pitch and yaw rate dampers. Whenever the atmospheric drag exceeds 0.2g, P67 is begun.

C.4.2.7 Final Phase, P67

The Entry Final Phase Program establishes the guidance-monitoring displays for the crew. P67 is the terminal portion of all entry trajectories. The guidance steers the CM to the landing site by generating the desired lift direction based on perturbations away from an onboard stored reference trajectory. The guidance also prevents the load factor from exceeding 8g. When the relative velocity decreases to 1000 ft/sec, the guidance ceases to generate new steering commands, but maintains the most recent value for the CM terminal roll attitude. P67 is terminated at crew option following drogue-chute deployment, although for telemetering coverage, termination is delayed as long as possible.

C.5 Lunar-Landing Guidance and Navigation

One of the most important phases of the Apollo mission is the guidance and navigation of the Lunar Module during the deceleration maneuvers prior to touchdown on the lunar surface. This section discusses the guidance and navigation capabilities onboard the LM for this powered-landing phase, which is presented in the context of the complete mission in Section 2.2.2.7.

The basic function of the LM guidance and navigation systems during powered landing is to take the spacecraft from a nominal initial altitude of about 50,000 ft and a velocity of approximately 5600 ft/sec and safely land it at an assigned site on the moon with virtually zero touchdown velocity. Several conditions and constraints govern the means by which the powered landing is accomplished:

a. The Descent Propulsion System (DPS) propellant must be utilized in an efficient manner, i.e., the required velocity increment should be as small as possible.

b. The selected landing site must be visible to the astronaut through the window of the LM for a time interval of at least 75 sec immediately prior to touchdown.

c. New state-vector estimates and steering commands for the LM cannot be obtained more frequently than once every 2 sec.

d. The DPS must be operated either at a fixed high-throttle setting (close to maximum thrust) or as a continuously-throttleable engine over a limited range of lower throttle settings—with the direction of applied thrust essentially parallel to the longitudinal axis of the LM.

e. The astronaut must have the capability of manually redesignating the landing site during the interval when the site is visible. (Conditions prevailing at 500-ft altitude must be "comfortable" such that manual takeover can be accomplished with ease.)

The navigation and guidance systems each perform different functions during the powered landing maneuver, as indicated in Fig. C.5-1. The navigation system basically determines (estimates) the state of the vehicle, i.e., its position and velocity. The guidance system uses the navigation information to compute specific-force commands for use in steering the vehicle. The following two parts of this section describe the guidance and navigation systems which enable the LM to accomplish the stated objectives during the powered-landing maneuver.

C.5.1 Guidance System Description

The mission requirements of efficient fuel utilization and a 75-sec landing-site visibility interval during the powered-landing maneuver are in direct conflict. For best fuel management during the powered maneuver, the Descent Propulsion System should be operated at the largest permissible thrust level, with the thrust direction slightly above the local horizontal. For the astronaut to have visibility of the landing site through the LM window, however, the longitudinal axis of the vehicle (and hence, the thrust direction) must be in a nearly vertical orientation. Continuous throttle control is required during the final part of the landing maneuver, moreover, to properly shape the trajectory (to meet visibility requirements) and to achieve the desired terminal conditions (position and velocity) at touchdown.

To accomplish these objectives in an efficient manner, the landing maneuver is divided into three major phases, as indicated in Fig. C.5-2. For convenience

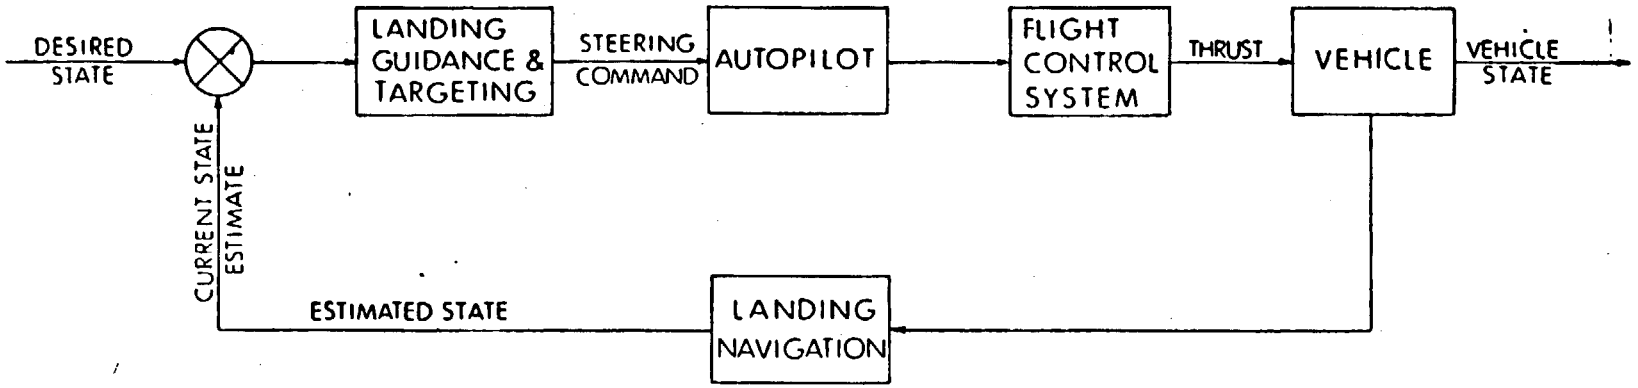

Figure C. 5-1 Functional Diagram Showing Relationships between LM Guidance and Navigation Systems

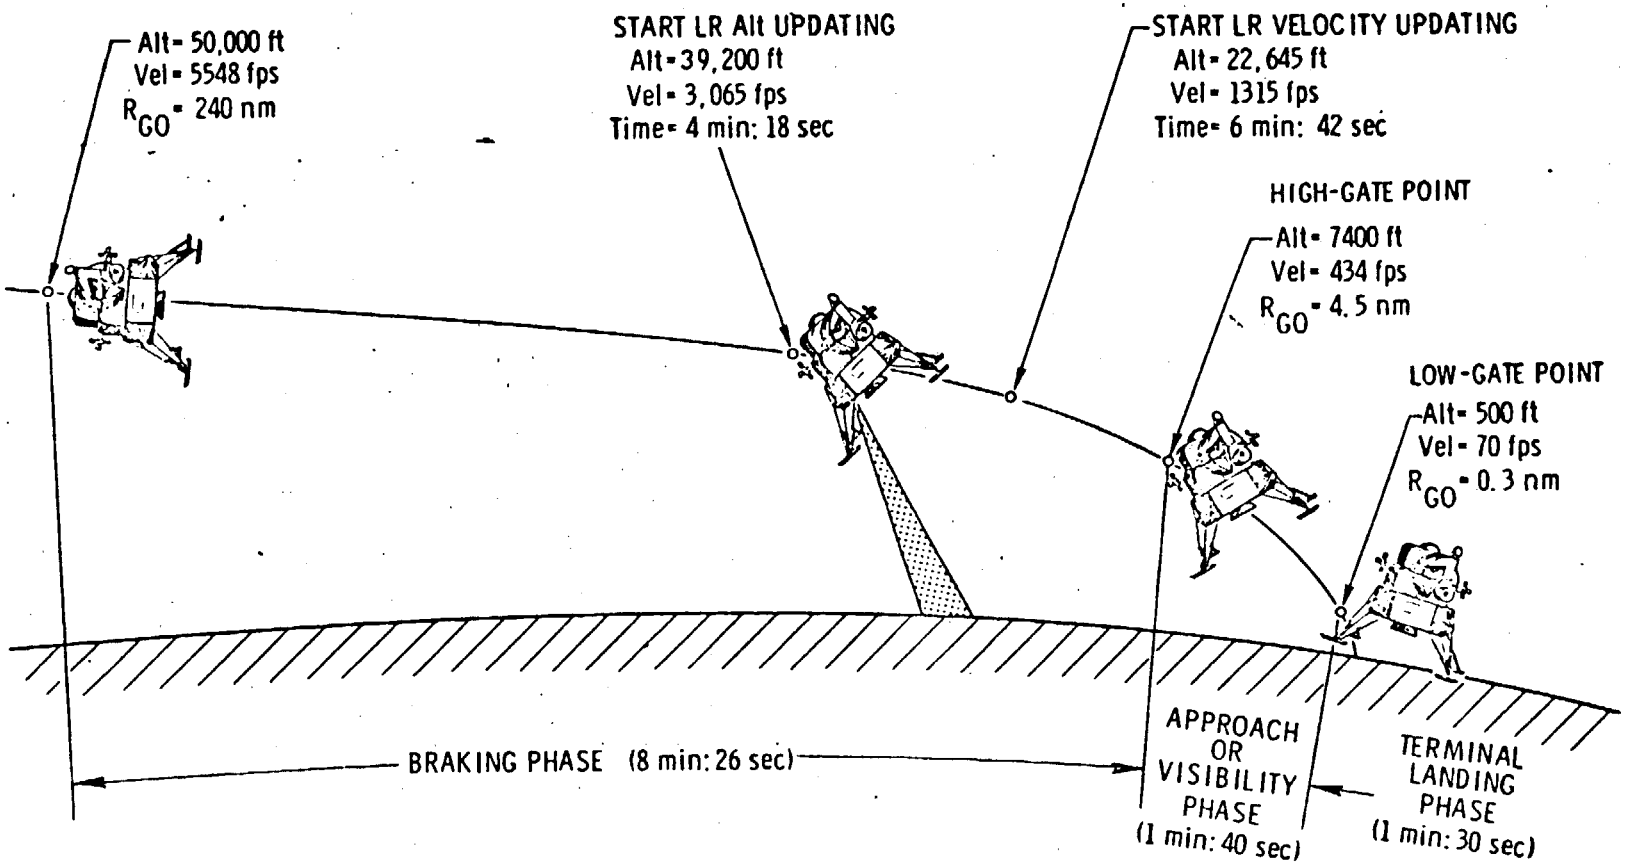

Figure C. 5-2 Powered Lunar-Landing Trajectory Phases

these three phases are referred to as the "braking" and "visibility" phases, and the terminal-descent (hover) maneuver. The major deceleration of the vehicle is accomplished during the braking phase, which is typically about 470 sec in duration. The braking phase is the longest with respect to time and range covered. The major objective of this phase is to establish the desired initial altitude and velocity conditions for the following visibility phase within efficient propellant usage limits; the bulk of the spacecraft's orbital velocity is removed during the braking phase. The visibility (approach) phase is designed such that the commanded vehicle attitude at reduced engine-throttle setting allows the astronaut to view the landing area for the first time. During this second phase the astronaut can—if he so desires—redesignate the landing site toward which the GN&C System is controlling the Lunar Module. A furthur objective of the visibility phase is to establish and maintain a trajectory from which the astronaut can easily take over control when he desires. During the third and final terminal-descent phase, the astronaut manually controls the LM through the final vertical descent to lunar touchdown. For convenience, the braking-phase terminal point is referred to as the "High-Gate"* point, and the visibility-phase terminal point is called the "Low-Gate" point. The terminal point for the hover is, of course, the landing site.

The guidance system solves the powered-landing guidance problem essentially as successive two-point boundary-value guidance problems. First of all, the guidance system executes an algorithm which determines the exact time and spacecraft attitude required at ignition. Ignition nominally occurs at the periapsis of the prior orbit. The spacecraft is maneuvered to that desired attitude and 35 sec prior to lighting the engine, the DSKY is blanked for five sec to notify the astronaut that everything is proceeding smoothly. Ullage is commanded for 75 sec prior to ignition—the beginning of the braking phase. For 26 sec the engine operates at minimum thrust, about 12 percent of the engine's rated thrust of 10,500 pounds.

*The "target" for the braking phase lies very near the lunar surface, projected forward about 62 sec past High Gate, but the phase actually ends at High Gate. The reason for avoiding the designated target (other than the obvious intersection with the lunar terrain) is that as time approaches zero with the guidance equations, the gain of the guidance equations approaches infinity. To avoid that gain variation, the targets of the braking phase are projected about a minute downstream from the desired braking-phase terminus. For the same reason, the target for the visibility phase lies some 10 sec beyond Low Gate.

With sufficient propellant, the guidance can permit substantial landing-site redesignation; as much as 7000 ft forward or to either side, or about 4500 ft backward. The guidance system requires no fixed landing site and indeed provides a relatively gentle standard approach to Low Gate whether the landing site has been redesignated or not.

Nominally, the descent trajectory is planar; however, redesignation results in a nonplanar trajectory. The descent trajectory during both the braking and visibility phases is provided by a vector polynominal which determines a three-dimensional line in space.

During almost the entire visibility phase, the guidance system maintains spacecraft attitude so that the landing site nearly coincides with the reticle (Landing Point Designator) etched on the Lunar Module window; the computer displays on the DSKY a number which informs the astronaut where along this reticle to look to see where the computer is taking him.

When the Low-Gate aim conditions have been met, the Lunar Module begins the so-called terminal-descent maneuver, which is nearly vertical. The point where this maneuver is started is dependent upon mission ground rules, crew option, and the erasable load provided for the specific mission.

The terminal-descent program in Apollo 11 and 12 automatically nulled the horizontal components of velocity and provided a 3 ft/sec rate of descent. In Apollo 13 and subsequent flights, the astronaut must specify the rate of descent by means of his rate-of-descent (ROD) controller; there is also an attitude-hold mode in which spacecraft attitude (and hence horizontal velocity) can be manually controlled—not necessarily nulled—to provide the desired translation across the lunar terrain.

The fuel allotment of the DPS provides for hovering immediately prior to touchdown. If the spacecraft does not enter this terminal descent phase, an abort is initiated on the ascent stage (see the next section of this Appendix).

Thrust magnitude and orientation time histories are shown in Fig. C.5-3 for a typical simulated landing trajectory. As can be seen, the throttle is operated at maximum thrust for most of the braking phase. During this period the DPS is

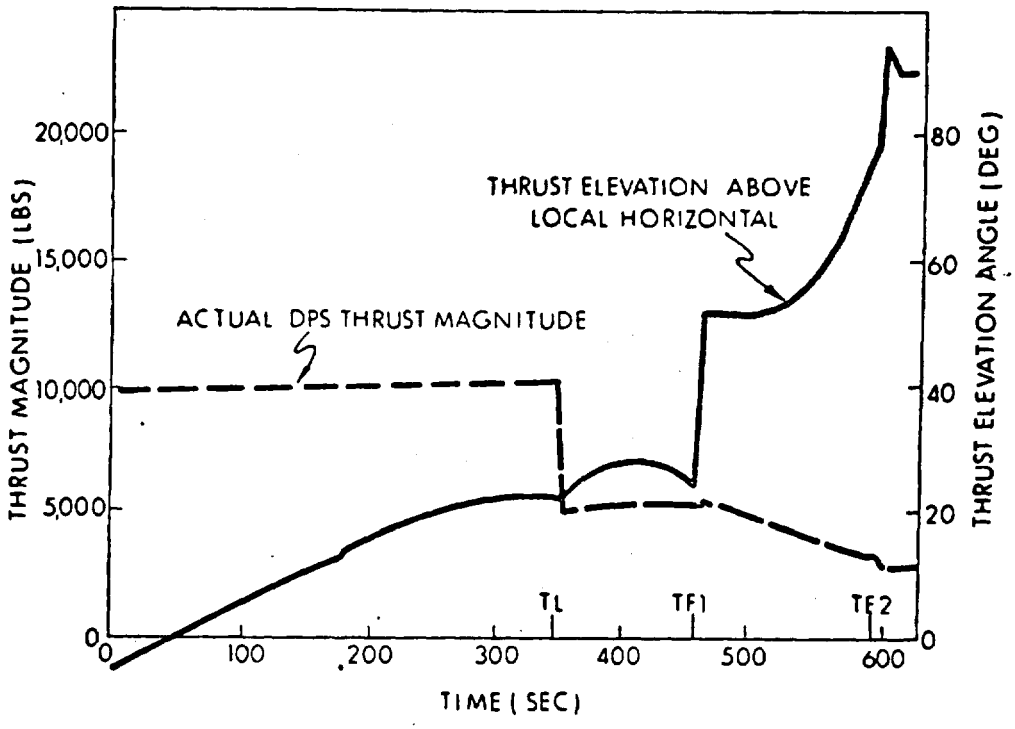

Figure C. 5-3 Thrust Profile for Reference Trajectory

oriented along the direction of the command specific force even though, due to limitations*, the throttle cannot follow the thrust command. The DPS is operated thereafter in the continuously-throttleable range between 12 and 58 percent of maximum thrust.

C.5.2 Navigation System Description

The navigation concept used by the LM GNC&S controls the braking and visibility phases of the lunar-landing maneuver and then continues to provide altitude and velocity data to the astronaut in the final hover phase. This approach employs an inertial navigation system updated by a doppler landing radar. The first three to four minutes of the braking phase is completely controlled by the inertial navigation system. When the altitude has decreased to between 40,000 and 30,000 ft, as illustrated in Fig. C.5-2 for the Apollo 11 mission, the landing radar is activated and initial altitude updates are accepted by the navigation system. When the velocity has decreased to less than 2000 ft/sec, landing-radar velocity corrections are next incorporated. It is important that the landing radar altitude and velocity updates be achieved during the braking phase for two reasons: first, to guarantee a safe altitude condition at the start of the approach phase, and second, to make the major trajectory changes before the visibility phase, so that the commanded vehicle altitude will be relatively free of control transients during this second phase, thus allowing the astronaut to visually determine and evaluate the landing area and have time to redesignate the landing site, if necessary. The visibility phase lasts only 100 sec or less, so major navigation-update changes resulting in large commanded vehicle attitude changes are undesirable.

The landing-radar sensor used to update the inertial navigation system during the landing maneuver is a four-beam doppler radar with the beam configuration shown in Fig. C.5-4. Three beams are used for velocity determination and one for range. The landing radar antenna is mounted at the base of the LM descent stage and can be oriented in one of two fixed positions. The first antenna position is used during the braking phase when the vehicle attitude is essentially horizontal,

*Due to a hard mechanical stop, the engine is incapable of delivering more than 94 percent of the rated thrust.

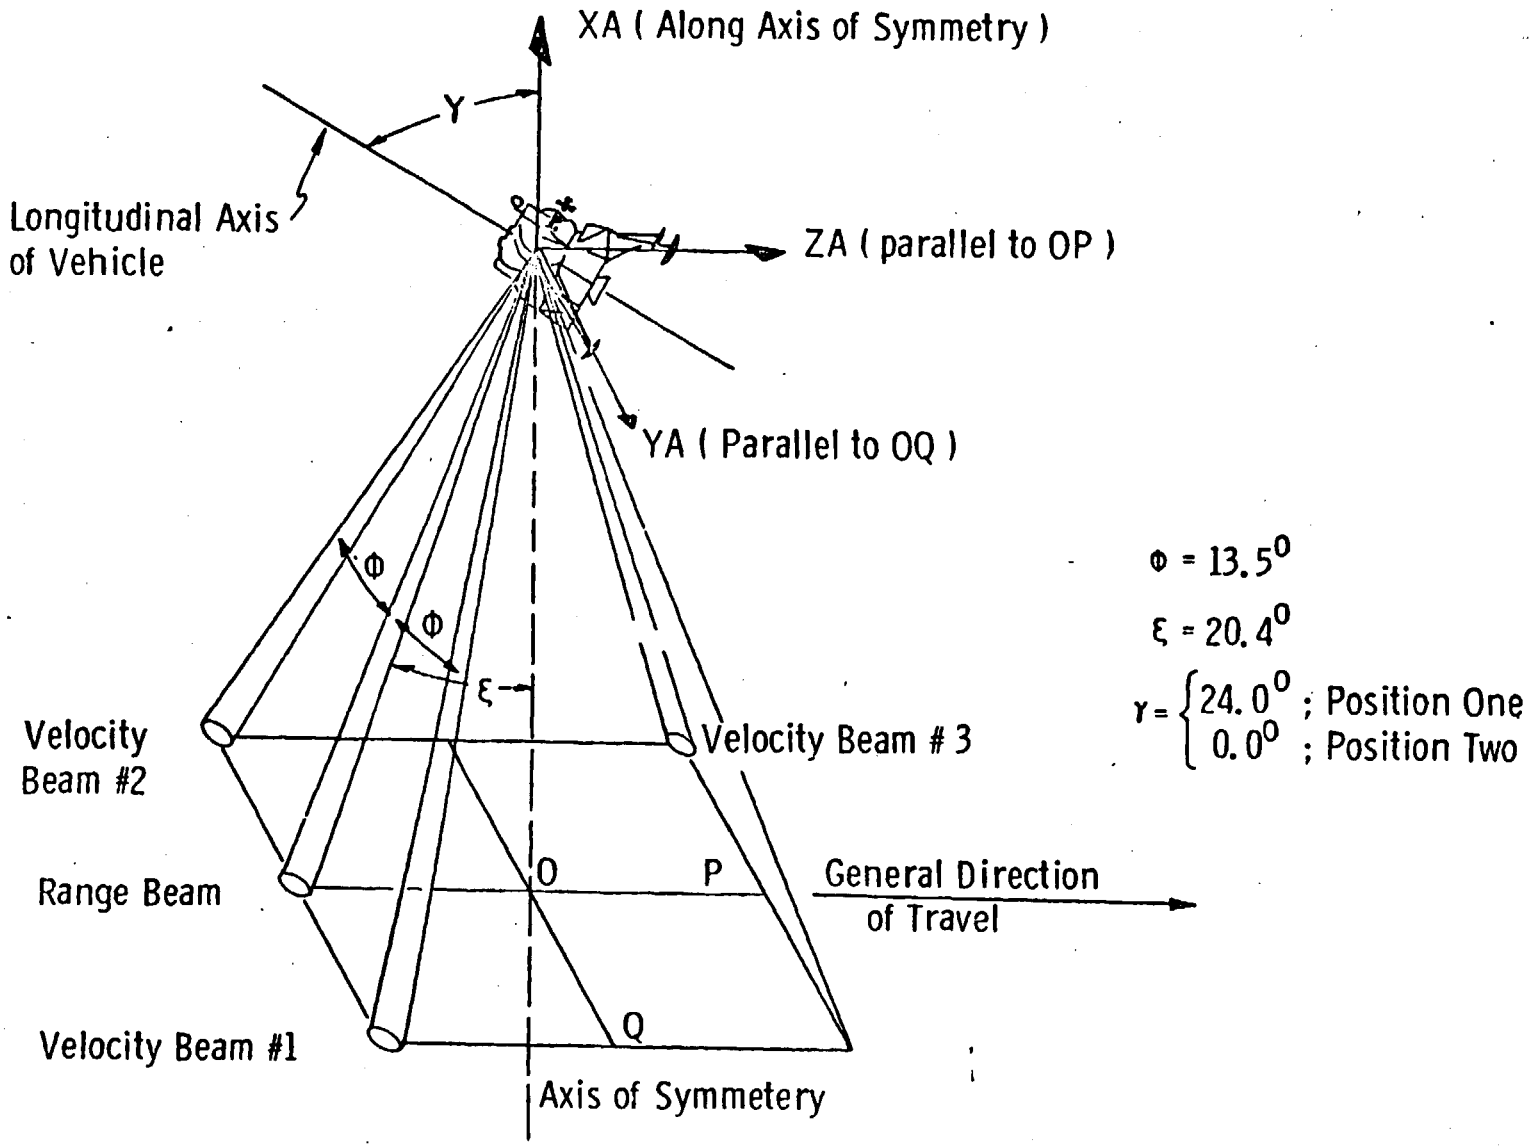

Figure C. 5-4 Landing-Radar Beam Geometry

and the second is used in the approach and landing phases as the LM orientation nears a vertical attitude.

The navigation problem during the powered landing is significantly different from that during the coasting cislunar and rendezvous navigation phases discussed in Appendix A. In addition to gravitational forces, a large thrust acceleration acts on the vehicle at all times during the landing maneuver. Another significant difference is that the guidance system controlling the vehicle thrust and attitude continuously uses navigation data throughout the maneuver, as opposed to the intermittent updates maneuvers typical for cislunar and rendezvous phases. Despite these major differences, the general navigation concept used during landing is very similar to that used during coasting phases.

A simplified landing navigation and control functional diagram is illustrated in Fig. C.5-5. During the landing maneuver, the inertial measurement unit is active at all times and provides the specific force data necessary for extrapolation of the state-vector estimates. During the initial part of the braking phase, the IMU is the only navigation sensor employed, and the estimated vehicle-state vector is used to command required vehicle attitudes and and engine throttle levels so that the desired terminal conditions for the mission phase are achieved. When the landing radar is activated later in the braking phase, the vehicle altitude above the lunar surface is estimated in the computer and then compared with the landing-radar measured range data converted to altitude. The difference between these two parameters is automatically checked in a landing radar (LR) data test to verify that the LR is operating normally. If the data check satisfactorily, a correction term is computed by applying a precomputed weighting factor to the altitude difference*, as shown in Fig. C.5-5. This correction is then used to update the state-vector estimate and to correct the landing trajectory in the next guidance-equation computation cycle. After velocity lock-on is achieved by the LR at a lower altitude and velocity, a similar operation is conducted for velocity updates to the state vector. When both altitude and velocity updating are being done, the navigation-measurement data from the LR are controlled by the onboard computer. Altitude is updated every two sec, and each component of velocity every six sec.

*If the LR data test is failed a given consecutive number of times, the DSKY alerts the astronaut to this fact; he can force acceptance of the rejected data should he desire, but this option would not normally be used.

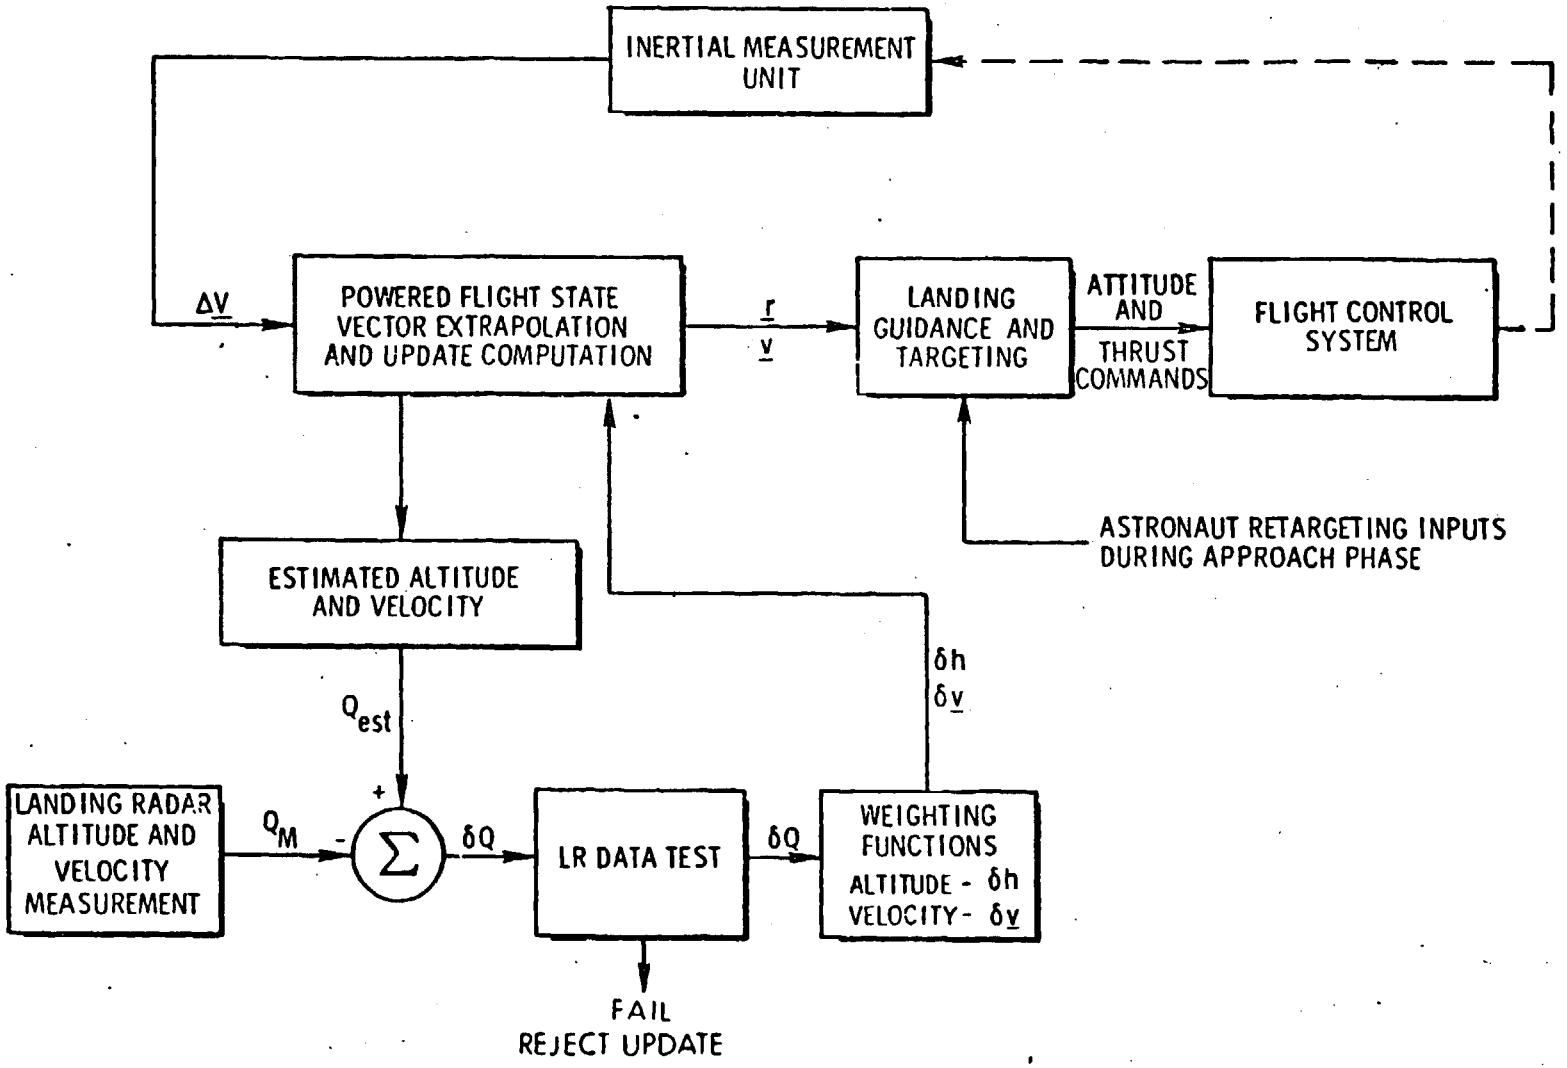

Figure C. 5-5 Simplified Functional Diagram of Lunar-Landing Navigation and Control

With reference to Fig. C.5-5, it should be noted that LR measurement data are used to update only four of the six components of the estimated state vector, i.e., altitude and the three velocity components. The down-range and cross-track position estimation errors are not updated by LR measurements. The astronaut can correct these two horizontal-plane position errors during the visibility phase by incrementally retargeting the guidance equations. In this operation, the landing site to which the GN&CS is controlling the vehicle is displayed to the astronaut by a DSKY number referenced to a grid pattern on the LM window. As previously mentioned, it is important that the major altitude and velocity corrections to the state vector be completed before the visibility phase so that the astronaut can effectively assess the landing area and correct the cross-track and down-range deviations, if required.

The weighting functions used to compute the state-vector updates in Fig. C.5-5 are significantly different from the time-varying statistical weighting functions used in the cislunar and rendezvous navigation phases. These powered-landing weighting functions are linear approximations to the statistical optimum navigation-filter weighting functions based upon inertial and LR sensor accuracies, lunar terrain uncertainties, and measurement bias errors. Since computation time is a critical parameter during the landing maneuver, the LR weighting functions are precomputed and stored. The altitude weighting function is a linear function relative to estimated altitude. Since altitude updates are typically started some 90 nmi from the landing site, lunar terrain altitude variations relative to the landing site were important factor in shaping this function. The velocity weighting function is linearized with respect to estimated velocity; it is truncated to fixed small values at the low-velocity conditions to minimize commanded attitude changes during the more critical terminal phases of the landing maneuver, and to avoid incorporating large LR velocity updates in the velocity region where LR dropouts can occur due to near-zero velocity conditions along various radar beams. The LR measurement weighting functions are uncoupled or noncorrelated with any measurement direction other than that along which the LR measurement is being made. Whereas correlation in the navigation-update weighting functions is very important in the cislunar and rendezvous navigation concepts, it is intentionally avoided in the landing maneuver navigation since implementation of such a correlation function was not considered practical due to modeling uncertainties and G&N system computation time limitations in the landing maneuver.

Several important differences exist between the operations of power-landing navigation and those of cislunar and rendezvous navigation. First, the state-vector updating procedure and monitoring in powered landing is completely automatic, since the navigation is time-critical. Next, the navigation-measurement weighting function used for the landing navigation is a precomputed linearized approximation to a statistically optimum weighting function, and is uncorrelated between state-vector components. Finally, the landing-navigation function updates only four of the six state-vector components, and the astronaut is required to manually correct the remaining horizontal position deviations during the final two phases of the landing maneuver.

C.6 Lunar Ascent and Abort Guidance

Once the Lunar Module has separated from the Command Module for the descent to the lunar surface, a means must exist for the LM to ascend back into a parking orbit preparatory to the rendezvous and docking of the two spacecraft. In the context of the Apollo mission, the ascent can occur in the planned circumstance, after a safe touchdown on the lunar surface; or in the unplanned circumstance, when an emergency situation exists. Whether the ascent be nominal or abort, the guidance equations which determine the ascent maneuvers are identical; only the initial conditions fed into the guidance equations are different.

By definition, nominal ascent can occur only after the crew of the LM has signalled to the computer an "acceptable" landing on the surface. By extension, nominal ascent guidance is the program used for any ascent from the lunar surface after an acceptable landing—even if an emergency ascent is deemed necessary prior to the planned time of liftoff.

Also by definition, an abort guidance program is used for ascent in two general situations: (1) when an emergency develops during descent, prior to touchdown on the lunar surface; and (2) when the crew does not acknowledge the touchdown as acceptable, choosing instead to ascend as quickly as possible.

The ascent and abort guidance programs are essentially open loop, in that they receive no positional updates from navigation. The programs perform their guidance function exclusively from signals received from the Lunar Module's Inertial

Measurement Unit. Initialization for ascent to the required parking orbit is dependent upon prevailing conditions.

On the basis of the initial conditions, the computer determines which of the LM's propulsion systems will be employed during ascent. Two systems are available, and either or both may be employed for a particular ascent—depending, of course, on the fuel available in each and the estimated fuel which will be required to reach the planned parking orbit. The Ascent Propulsion System (APS) is used alone for an ascent from the lunar surface or for an abort occurring very close to touchdown. The Descent Propulsion System (DPS) is used alone only when its remaining fuel volume is sufficient to propel the LM into the parking orbit—up to about 300 sec into the braking phase of the lunar descent. When an abort is called at a later time, the DPS will be flown to depletion; parking-orbit injection will then be accomplished using either the APS or the RCS. The foregoing procedures must, of course, be modified should the abort be caused by a propulsion-system failure.

A flag (LETABORT) internal to the lunar descent programs permits aborts to be called once the lunar programs have been activated. Throughout the descent phase, the crew can call up abort guidance; this remains true until after touchdown. Once the crew decides that a safe landing has occurred and that surface operations can commence, they confirm it to the AGC by calling the lunar-surface program, which cancels the LETABORT—and leaves the nominal ascent program as the sole means of achieving ascent.

In a nominal ascent, the target of the ascent is planned well ahead of time, based on the state vector of the orbiting Command Module. The LM must lift off the lunar surface at such a time and with such velocity that it will insert into a parking orbit with a favorable position and velocity relative to the CM. Knowing the state vector of the Command Module, the ascent program selects the proper liftoff time.

Initialization for an ascent from the lunar surface is straightforward because it is known in advance. But initialization for an abort during descent is quite involved because neither the necessity nor the time of an abort can be known in advance. In addition, the phase angle—the angular relationship between the LM and the CM—at the time of the abort is equally unknown and must be estimated at initialization.

Consequently a variable targeting system must be employed for an abort, based upon the past information of the CM state vector and an estimate of the LM's current position.

Initialization establishes the initial conditions, sets up initial displays, decide on which targets to use, and brings up the guidance equations into the servicer loop. In all three ascent programs, the initialization routine does this by putting an address into an erasable location toward which the end of the servicer routine goes and which it recognizes as an address; the routine then continues on with the guidance.

Abort initialization also establishes the parking-orbit injection in the radial and cross-range directions. (With a fixed-thrust engine, if these two positional components are controlled, the third component—down-range position—is indirectly controlled.) Depending on the propulsion system which is used to begin an abort, two abort guidance programs are available. The descent abort program begins with the DPS, and the abort-stage program uses only the APS.

No matter which ascent program is called, once the initialization of the program is completed, the initialization routine simply slips out of the loop.

The ascent trajectory for either a nominal launch or an abort is preplanned: the LM thrusts to a vertical ascent for about 0.8 sec, after which it begins a pitch-over maneuver and accepts new data from the IMU. The importance of the 0.8-sec vertival ascent becomes apparent in the case where, upon touchdown on a boulder-strewn surface, the LM begins to topple or if, during descent, the LM begins to tumble.

Once the guidance equations take over control of the program, they produce a commanded thrust vector based on the inertial platform's accelerometer readings and the newly-inputted CM state vector. The thrust-vector components are converted into changes in the CDU angles by a subroutine called FINDCDUW; these angles are then transmitted to the Digital Autopilot. Thus, the LM DAP is the control system during either nominal or abort ascent.

Ascent into the parking orbit required for rendezvous ends when the target velocity commanded (in radial and cross-range positions) is achieved parallel to

the plane of the Command Module orbit. At that point, the rendezvous porgrams take over.

C.7 FINDCDUW – A LM Powered-Flight Guidance/Autopilot Interface Routine

As explained in the previous sections of this Appendix, the LM navigation and guidance systems perform complementary functions during the powered-flight maneuvers. Basically, navigation estimates the position and velocity of the Lunar Module, and guidance uses this information to compute the commands needed to steer the vehicle. The guidance-system commands are computed in terms of thrust vectors and must be translated into angular commands to be accepted by the LM DAP. The routine devised to accomplish this conversion is called FINDCDUW. (The "W" in FINDCDUW represents "window control".

Specifically, FINDCDUW aligns the LM thrust vector with the commanded direction during powered maneuvers and aligns the reticle in the window in the direction commanded by the guidance elevations. During descent, the reticle is aligned with the landing site; during ascent, it is aligned in the forward direction of motion.

FINDCDUW receives from the guidance program thrust and window-pointing commands which determine a required spacecraft attitude. Gimbal-angle changes corresponding to this attitude are computed and limited in magnitude to the maximum which may be commanded in one guidance-program computation period (20 deg in 2 sec). The limited gimbal-angle commands are then divided into 20 increments, each one of which is equal to the change which can be expected in one DAP cycle (0.1 sec). Finally, the gimbal-angle increments are sent to the autopilot, along with the corresponding attitude-rate commands and a bias angle whose magnitude is computed from the commanded rates and the control authority. The bias angle is utilized to smooth transient behavior in the autopilot. All of these constraints upon the autopilot commands result in a continuous, rate-limited attitude profile extending over whatever duration is required to complete the maneuver.

There is further constraint imposed upon the angle of the middle gimbal to avoid gimbal lock*. Provided the middle gimbal angle is not initially in gimbal lock, and provided the combination of input commands does not yield a terminal attitude in gimbal lock, FINDCDUW renders it impossible to pass through gimbal lock by constraining the middle gimbal angle to the range between its initial and terminal values. Furthermore, by avoiding gimbal lock, FINDCDUW makes the GN&CS abort guidance not dependent upon manual intervention. Other schemes for gimbal-lock avoidance produce similar outcomes, but are computationally less direct.

*A three-degree-of freedom gimbal system such as the one used on the Apollo IMU can cause problems due to a phenomenon called "gimbal lock". Gimbal lock occurs when the IMU's outer axis becomes parallel to its inner axis. In this position, all three axes of gimbal freedom lie in a plane, and no axis is in a direction to absorb instantaneously rotation about an axis perpendicular to this plane. Thus, at gimbal lock the inner stable member can be pulled away from its inertial alignment. Even though a three-degree-of-freedom gimbal system geometrically allows any relative orientation, the required outer-gimbal angular acceleration needed at gimbal lock to maintain stabilization exceeds servo capability. One direct solution to gimbal-lock problems is to add a fourth gimbal and axis of freedom which can be driven so as to keep the other three axes from getting near a common plane. However, the cost in complexity and weight for a fourth gimbal is considerable. Fortunately, Apollo IMU operations are such that gimbal lock can generally be avoided by constraining spacecraft motion relative to the gimbals. If just one of these gimbal-angle changes is constrained, the problem is entirely averted.

APPENDIX D

MAJOR PROGRAM CAPABILITIES—
Digital Autopilots

D.1 Developmental History of the Digital Autopilots

Early in 1964, MIT was asked by NASA to determine the feasibility of implementing digital autopilots (DAPs) in the guidance computers. This request was occasioned by rather widespread and growing dissatisfaction with the degree of flexibility offered by the analog autopilot then being employed. (By chance, a switchover was impending from the Block I G&N system to the Block II system, with its enlarged computer memory and improved electronics CDUs.) Furthermore, it was felt that a digital autopilot as the primary control system, backed up by the preexisting analog autopilot, could provide far greater mission reliability than the analog autopilot alone, and that a loss of the nonredundant analog autopilot might have meant mission failure—or worse, loss of the crew and spacecraft.

But of the advantages to digital control, perhaps the greatest was the above-mentioned flexibility: the coding effort required to modify or replace autopilot functions, such as manual operational modes, could be bought more cheaply than corresponding changes to an analog system. In addition, the ease of implementing nonlinear functions, such as deadzones, parabolic curves and counters, reinforced the arguments for a digital autopilot. Although a digital autopilot had not yet been flown in any manned vehicle, digital controllers had been discussed in the control literature and had been flight-tested in at least one unmanned guided missile.

During the summer of 1964, MIT was given the go-ahead on DAP design and implementation. This decision followed MIT studies showing that digital controllers were not only feasible, but offered improvements in control performance as well. This conclusion led to the choice of the digital autopilot as the primary control system.

D.1.1 CSM DAPs

Control of the CSM and CSM/LM vehicles, via the CM Apollo Guidance Computer, involves three separate DAPs, one each for coasting flight, powered flight, and atmospheric entry. In addition, the AGC onboard the Command Module provides for takeover of Saturn steering during boost.

The coasting-flight DAP, which fires the Reaction Control System (RCS) jets for attitude control, provided significant improvements in both performance and flexibility over the Block I system. This was achieved by a number of new features, among them an improvement in the automatic-maneuver routine and the addition of several manual-control modes. By early 1965, the basic RCS autopilot functions were laid out, including phase-plane and jet-select logic, a new maneuver routine, and interfaces for the various manual modes. Along with the development of these functions, some additional features were implemented, such as the sharing of rotational and translational jets during some maneuvers (e.g., ullage), and a rate estimator employing Kalman filtering. Though the RCS DAP functions were not optimized in a rigorous sense, a primary concern was to use as little RCS fuel as possible for attitude maneuvers, since the fuel supply was a limited fixed quantity. Throughout the design effort and even into the flights, the RCS DAP design remained flexible to accommodate many modifications incorporated to improve the capability and performance of the system. RCS DAP functions are discussed in detail in Section D.2 below.

For powered flight, the problem was to maneuver the vehicle according to commands received periodically from the guidance loop. Changes are effected in pitch and yaw by thrust-vector control (TVC), i.e., by deflecting the thrust vector relative to the spacecraft by gimballing the Service Propulsion System engine mounted on the Service Module. Roll control, which could not be achieved with the SPS engine, requires the use of RCS jets with attendant control logic. Because the SPS engine-positioning servos are linear devices*, the TVC loops could be designed around linear constant-coefficient compensation filters. This design allowed analytical determination of the stability of the bending and sloshing modes, which

*The position and rate limits (4.5 deg and 9.0 deg/sec, respectively) on engine motion were large enough to ensure they would not normally be encountered.

proved to be a great asset during the course of the development. The two vehicle configurations, CSM alone and CSM/LM, were sufficiently different to warrant separate TVC autopilot designs. Naturally, each of these DAPs went through several iterations as better bending and slosh information became available. In fact, it was not until the Apollo 10 mission that the final TVC DAP design was flown. Section D.3 discusses the TVC DAP used for pitch and yaw changes and RCS roll control during powered flight.

The Entry DAP controls vehicle attitude from CM/SM separation, occurring about 20 minutes before atmospheric entry, through to drogue-chute deployment. This DAP entails three automatic control modes: (1) three-axis attitude maneuver and stabilization prior to encountering the atmosphere; (2) coordinated roll maneuvers in the atmosphere, which control lateral and longitudinal range by rotating the lift vector; and (3) pitch and yaw rate damping about the aerodynamically-stable trim attitude. The first Entry DAP, which was designed off-line during the Block I flights, was developed in an attempt to reduce the RCS fuel used in the atmospheric roll maneuvers. This consideration was important because roll maneuvers normally used more RCS fuel than the extra-atmospheric and rate-damping modes combined. The design was successful, and by the time the decision had been made to use digital autopilots, this first Entry DAP had proved it could significantly reduce RCS fuel usage significantly. Further improvements, especially in the phase-plane logic, allowed even greater savings, until by the first Block II flight (Apollo 7), the atmospheric-roll maneuvers used only about one-eighth of the RCS fuel that the Block I system would have used. The combined performance of the guidance/autopilot system during entry has been excellent, as measured by the small target miss-distances (averaging roughly one nautical mile). The Entry autopilot is discussed in Section D.4.

The Apollo Guidance Computer onboard the CM also includes provision for takeover of Saturn steering during boost, should inertial reference fail in the Saturn Instrument Unit. Until Apollo 13, takeover was provided as a manual function; more recently an automatic capability has been added. Section D.5 discusses AGC takeover of Saturn steering.

D.1.2 LM DAP

The design of the LM DAP differed considerably from that of the CSM DAPs, due mainly to differences in the configuration of the two vehicles. Aside from the fact that the LM would not experience atmospheric flight, the major difference was that all modes of LM coasting and powered flight, both for the descent and ascent stages, used essentially the same basic control logic; i.e., the LM DAP had to be far more integrated than the CSM autopilots. One of the main factors necessitating this integrated design was that neither the descent engine nor the ascent engine was intended to control vehicle attitude. The ascent engine was rigidly mounted, while the descent engine could be gimballed only at the very slow fixed rate of 0.2 deg/sec. This meant that RCS jet firings would be required for attitude control in LM powered flight as well as coasting flight.

Given this level of integration, the individual control modes had their own developmental histories. The state estimator, used to provide the DAP with information on angular rates and accelerations, was first implemented as a Kalman filter with time-varying gains. These gains were stored in tabular fashion in fixed memory. However, early testing revealed estimation errors resulting from perturbations such as propellant slosh and CDU quantization. To reduce these errors, the estimator was changed by replacing the table of time-varying gains with several constant gains, and introducing an additional nonlinear filter to reduce the perturbation on the attitude measurements.

The descent-engine control system was originally designed merely to keep the engine pointing through the vehicle center of gravity. This acceleration-nulling mode required only a knowledge of the vehicle's angular acceleration and the control effectiveness of the engine. The RCS jets would then be used for attitude control. It was soon noticed, however, that given a suitable control law, the engine's Gimbal Trim System was capable of providing full attitude control when vehicle rates and accelerations were low, thus saving RCS fuel. The control law chosen was a modified time-optional law.

As in the CSM, the LM RCS control laws for powered or coasting flight were based on phase-plane logic. This logic, which varies with vehicle configurations and flight conditions, includes such features as variable switch curves, biasing of

the deadbands, and separate firing logic for large and small phase-plane errors. Section D.6 provides descriptions of the various functions of the LM autopilot.

D.2 CSM Reaction Control System (RCS) Autopilot

The Reaction Control System Digital Autopilot (RCS DAP), an integral part of the CSM GN&C System, provides automatic and manual attitude control and stabilization and manual-translation control of the Apollo spacecraft. The autopilot is designed to control four spacecraft configurations during the so-called coasting phases of flight—CSM/SIVB, CSM/LM, CSM/LM ascent stage, and CSM alone. For the latter three configurations, control forces and moments are provided by the Service Module Reaction Control System, which employs 16 rocket thrusters mounted in groups of four, known as quads; for the CSM/SIVB configuration, control forces and moments are provided by the SIVB RCS system. CSM control is achieved by jet-on and jet-off command signals supplied to the solenoids of the thrusters.

The RCS DAP receives and processes data from various internal and external sources. Measurements of spacecraft attitude are provided by the Inertial Measurement Unit (IMU) through Coupling Data Units (CDUs), which serve as gimbal-angle encoders. Attitude rate information is derived from these measurements. Manual attitude commands are generated by the Rotational Hand Controller (RHC) and interface with the computer through an input channel. The inputs are discrete* in nature, specifying the direction of rotation required about each of the spacecraft axes. The Translational Hand Controller (THC) interfaces in a similar fashion and provides translational-acceleration commands along each of the spacecraft axes. The minimum-impulse controller (MIC) interfaces with another input channel; each deflection of the MIC produces a single short rotational impulse (14 msec) of the SM RCS jets about the appropriate axes.

Selection of the various autopilot control modes is governed by a panel switch. Using the DSKY, the crew may also specify several autopilot control parameters, such as angular deadbands, maneuver rates, thruster-quad status and spacecraft-

*The AGC has 15 input and output channels whose bits are individually distinct (i.e., discrete). Each bit either causes or indicates a change of state, e.g., liftoff, zero optics, SPS-engine on, RCS-jet on, etc.

mass properties. Automatic attitude maneuvers are performed using internal steering commands provided by the guidance and navigation programs and initiated by keyboard request. In response to these steering commands, the RCS DAP issues jet-on and jet-off commands to the Reaction Control System and generates attitude error signals for display on the flight-director attitude-indicator (FDAI) error meters.

D.2.1 Modes of Operations

The RCS DAP may operate in one of three modes selected by the crew via the Command Module AGC MODE switch*. These three modes can be summarized as follows:

D.2.1.1 Free Mode

RHC commands are treated as minimum-impulse commands. Each time the RHC is moved out of detent, a single 14-msec firing of the control jets results on each of the axes commanded. If there are no RHC commands, MIC commands are processed. If neither RHC nor MIC commands are present, the spacecraft drifts freely.

D.2.1.2 Hold Mode

If there are no RHC commands, the vehicle is held about the attitude reached upon switching to hold or upon termination of a manual rotation. RHC commands override attitude hold and result in rotations at a predetermined rate on each of the appropriate control axes for as long as the RHC remains out of detent. When the RHC returns to detent in all three axes, all angular rates are driven to within a deadband; attitude hold is then established about the new spacecraft attitude. In the Hold mode, all MIC commands are ignored.

---

*THC commands are accepted in any mode and are combined with the rotation commands whenever possible. In the event of a quad failure, however, rotation control will assume priority over translation; i.e., translations are ignored if they would induce rotations that could not be compensated by RHC (or automatic) commands or if they were to cancel a desired rotation. The crew is responsible for performing ullage with the THC and by the selection and management of ±X-translation quads, as described below.

D.2.1.3 Auto Mode

If there are no RHC commands, the DAP accepts rate and attitude commands from the guidance programs to bring the vehicle to the desired attitude at the specified rate. RHC commands override automatic maneuvers and are interpreted as rate commands, as in the Hold mode. When RHC commands cease, automatic maneuvers are not immediately resumed, but rather attitude hold is established about the new orientation. The automatic maneuver can be resumed only by astronaut action via the DSKY. In the absence of automatic-maneuver commands, the DAP functions exactly as in the attitude-hold mode. In the Auto mode, as well, MIC commands are ignored.

D.2.2 Crew Control of the RCS DAP Configuration

D.2.2.1 DSKY Operation

D.2.2.1.1 Data Loading

Most of the autopilot variables over which the crew has control are loaded by means of Verb 48, which is normally executed before the DAP is turned on for the first time and anytime thereafter when the crew wishes to change or update the data. Verb 48 displays, successively, three nouns (only two are needed for the RCS autopilot) for loading and verification.

Noun 46 permits the loading of data relating to current spacecraft configuration; the choice of quads to be fired for X-axis translations; the size of the angular deadband for maneuvers in the Hold and Auto modes; and the specified rate for RHC activity in the Hold and Auto modes and for automatic maneuvers supervised by a special routine for coasting-flight attitude maneuvers, KALCMANU. In addition, Noun 46 includes information on jet selection for roll maneuvers and on the operational status of the quads.

Noun 47 allows the loading of the current weight (in pounds) of the CSM and of the LM. The spacecraft moments of inertia and other pertinent parameters, such as propellant loading and cg offset, are stored in the CM AGC as a function of these keyed-in weights.

Noun 48 allows the loading of data pertinent only to the TVC DAP.

D.2.2.1.2 Other DSKY Operations

After Verb 48 has been completed, Verb 46 may be used to establish autopilot control of the spacecraft. If the specified configuration is CSM alone, CSM/LM, or CSM/LM (ascent-stage only), and if the Thrust Vector Control DAP is not running, the RCS DAP begins initialization. If the specified configuration is SIVB/CSM, RHC commands are sent to the SIVB control system for manual rate control.

Verb 49 calls up R62, a routine that permits the crew to specify a final vehicle attitude. This attitude can be achieved by means of an attitude maneuver supervised by the special routine, KALCMANU.

Routine 61, the tracking-attitude routine, provides the RCS DAP with the information required to automatically track the LM during rendezvous navigation. Whenever R61 requires an attitude maneuver resulting in any gimbal-angle change of 10 deg or more, it will, after an appropriate interval, check a flagbit and, if appropriate, call the attitude-maneuver routine, R60.

Verb 79 calls R64, the barbecue-mode routine which is closely related to utilization of the RCS autopilot. R64 enables the crew to perform (1) passive thermal control, a roll about the X-axis of the stable member; (2) an orbital rate-drive procedure about the Y stable-member axis; and (3) deadband changing without the requirement for direct erasable loading. With Noun 16, the astronaut informs the computer of the time (ground-elapsed time) at which he wishes the X- or Y-axis maneuvers to begin. For that maneuver Noun 79 displays the commanded rate, the RCS DAP deadband, and the stable-member axis about which the maneuver will occur.

D.2.2.2 Attitude-Error Displays

The RCS autopilot generates three types of attitude errors for display on the FDAI error meters, all of which are updated every 200 msec.

Mode 1 displays autopilot phase-plane errors as a monitor of the RCS DAP and its ability to track automatic-steering commands. In this mode, the display is zeroed when the MODE switch is placed in the Free position.

Mode 2 displays total attitude errors to assist the crew in manually maneuvering the spacecraft to the spacecraft attitude (gimbal angles) specified in Noun 22. The attitude errors with respect to these angles and the current CDU angles are resolved into RCS control axes.

Mode 3 displays total astronaut attitude errors with respect to the spacecraft attitude (gimbal angles) in Noun 17 to assist the crew in manually maneuvering the spacecraft. The attitude errors with respect to these angles and the current CDU angles are resolved into RCS control axes.

D.2.3 RCS DAP Implementation

The RCS autopilot was designed to perform a number of functions during a mission:

a. Attitude hold and stabilization
b. Automatic maneuvering, including
   (1) large-angle spacecraft reorientations
   (2) automatic tracking of the LM
   (3) passive thermal control
   (4) orbital-rate drive
c. Manual attitude-rate control
d. Manual rotational minimum-impulse control

D.2.3.1 Attitude Hold and Stabilization

The RCS autopilot is formulated as a sample data system which, in the attitude-hold mode, attempts to null the computed set of body-attitude errors and the spacecraft angular velocity. Figure D.2-1 depicts the logic of this RCS mode. The input to the attitude-hold logic is a set of reference angles corresponding to the desired outer-, inner-, and middle-IMU gimbal angles. In addition to computing estimated body rates, during each 0.1-sec sampling interval, the DAP compares these desired angles with current CDU angles. These rates are derived from IMU-gimbal-angle differences, which are transformed into corresponding body-angle differences and smoothed by a second-order filter. For greater accuracy, the commanded angular acceleration of the RCS jets is included in the computation.

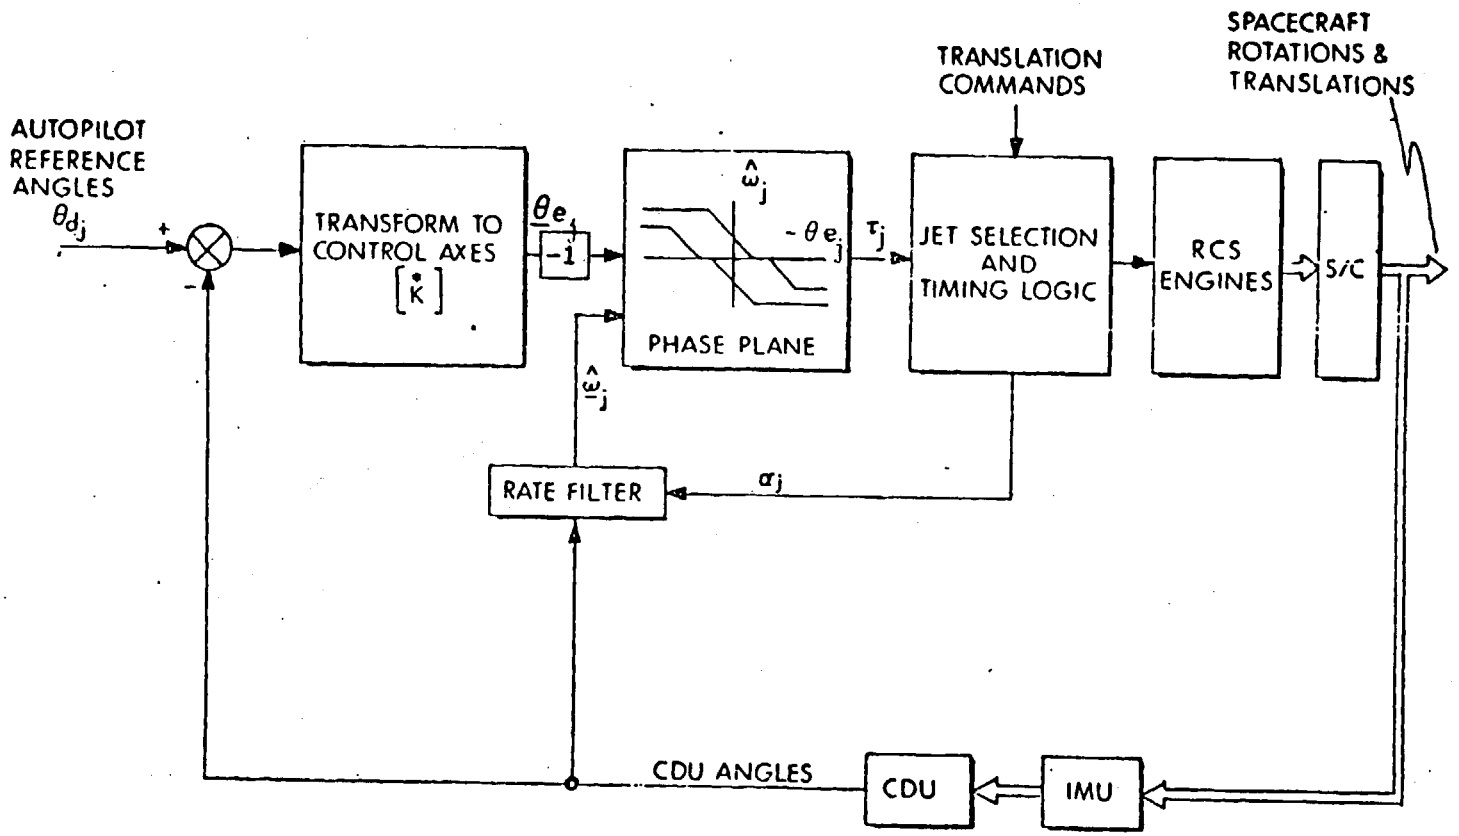

Figure D. 2-1 Functional Diagram of CSM RCS Attitude-Hold Logic

As a function of the attitude error and the attitude rate, nonlinear switching logic (phase-plane logic) is used to generate RCS jet-on times, indicating the required polarity and duration of the thruster torque for each control axis. A jet-selection logic combines the rotation commands with the translation commands from the THC and selects the individual jets to be fired.

D.2.3.2 Automatic Maneuvering

The RCS autopilot can perform several different types of automatic rotations. Figure D.2-2 is a functional diagram of the automatic-control logic. The first of the automatic rotations is the ability to perform large-angle reorientations of the spacecraft, as required for main-engine (SPS) thrust-axis alignments prior to powered-flight thrusting maneuvers. A special steering routine performs these maneuvers in a fuel-efficient manner. The RCS DAP can also orient the CSM such that the LM is continuously within the optics field-of-view during rendezvous navigation. In addition, the autopilot can be used to establish the thermal-balancing roll rotation required to maintain uniform solar heating of the spacecraft during extended periods of drifting flight. The autopilot can be used to produce a rotation at an astronaut-specified rate about the Y stable-member axis. In earth and lunar orbit, if the Y stable-member axis is aligned normal to the orbital plane and if the specified rate corresponds to orbital rate, this capability provides a pseudo-local-vertical tracking mode.

D.2.3.3 Manual Attitude-Rate Control

The Rotational Hand Controller interfaces with the computer by means of discrete inputs which indicate a positive, negative, or zero rotation command for each control axis, in accordance with the duration of RHC deflection and some phase-plane switching logic. With rate-command control, the RCS autopilot causes a vehicle attitude rate to be generated in response to these commands. The magnitude of the command rates is specified by the astronaut in the autopilot data-load procedures (see Section D.2.2.1.1). Four rates are available for selection: 2.0, 0.5, 0.2 and 0.05 deg/sec. Figure D.2-3 provides the logic for RCS manual-rate control.

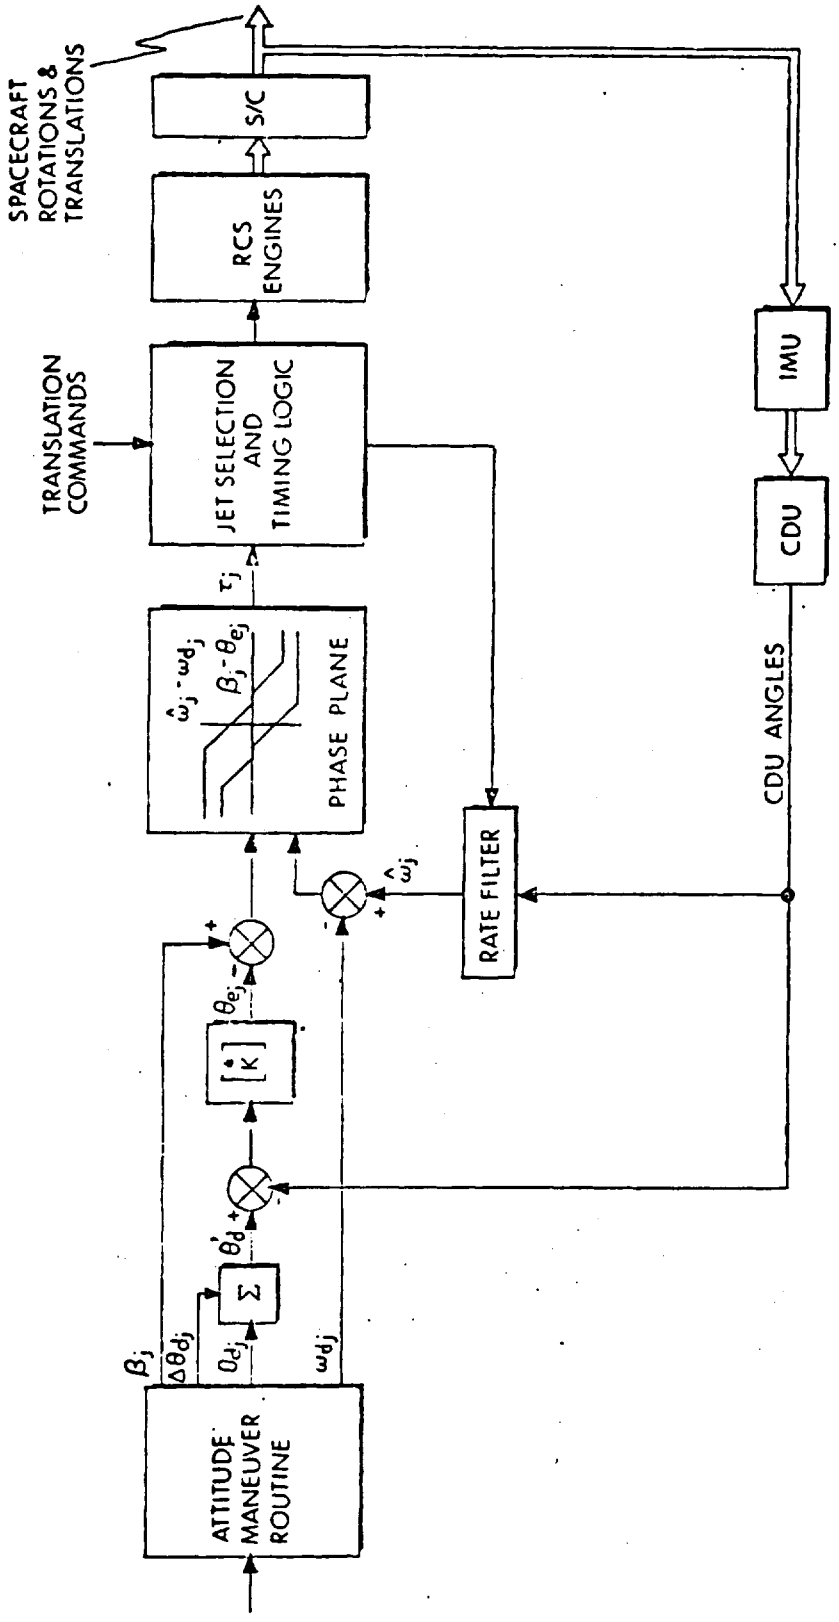

Figure D.2-2 Functional Diagram of CSM RCS Automatic Control Logic

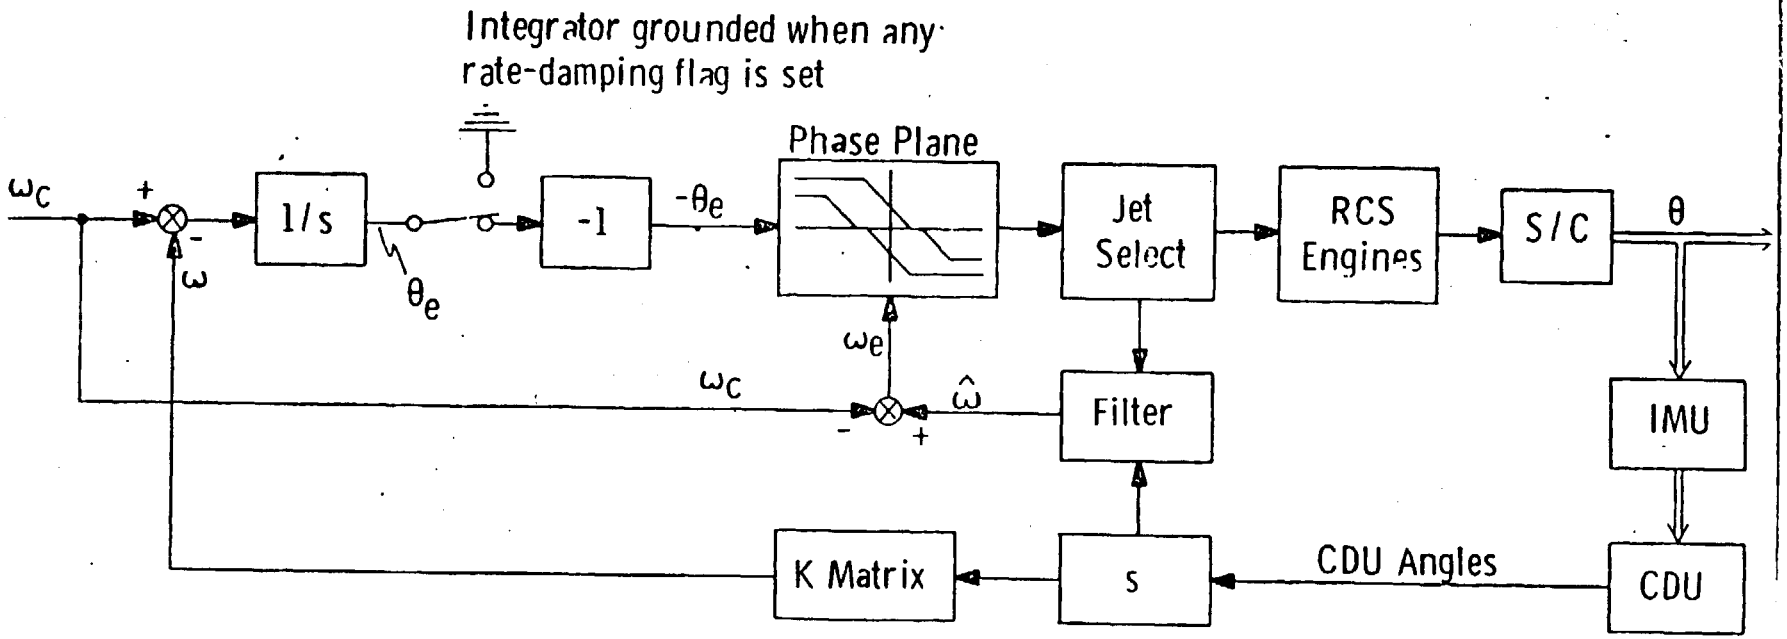

Figure D.2-3 Functional Diagram of CSM RCS Manual-Rate Control

D.2.3.4 Manual Rotational Minimum Impulse Control

With manual rotational minimum-impulse control, a deflection of the MIC or the RHC at the navigation station produces small rotational commands to the RCS jets about each of the commanded axes. Each time the controller is deflected, the autopilot generates a jet-on command which lasts for 14 msec, thereby producing a small change in the spacecraft angular velocity about the appropriate axis. No other control action is taken until the controller is again deflected to produce another minimum impulse. This control is particularly useful for navigation sightings using the onboard sextant.

D.2.4 Restart Behavior of the RCS DAP

Should a restart occur during RCS DAP operation, any jets that happen to be on will be turned off, and reinitialization of the RCS DAP will be scheduled. This reinitialization is the same as the initialization caused by RCS-DAP turn-on using Verb 46, with the exception that the attitude reference angles are not changed. Automatic maneuvers governed by R60 that were in progress at the time of a restart will not automatically be resumed; rather, attitude hold will be established following reinitialization; automatic maneuver can be resumed by appropriate DSKY action.

D.3 CSM Thrust Vector Control (TVC) Autopilot

D.3.1 Summary Description

During powered-flight, the pitch and yaw control of the spacecraft is achieved through the deflection of a gimballed engine of the CSM Service Propulsion System. Attitude control about the roll axis is provided by the Reaction Control System jets. The computation of gimbal-servo commands in response to computed errors between commanded and measured attitudes is the function of the Thrust Vector Control Digital Autopilot, implemented in the Apollo Guidance Computer onboard the CM.

D.3.1.1 TVC Pitch and Yaw Control

The following is a summary outline of TVC pitch and yaw control: (Figure D. 3-1 is a functional block diagram of the TVC pitch and yaw control.)

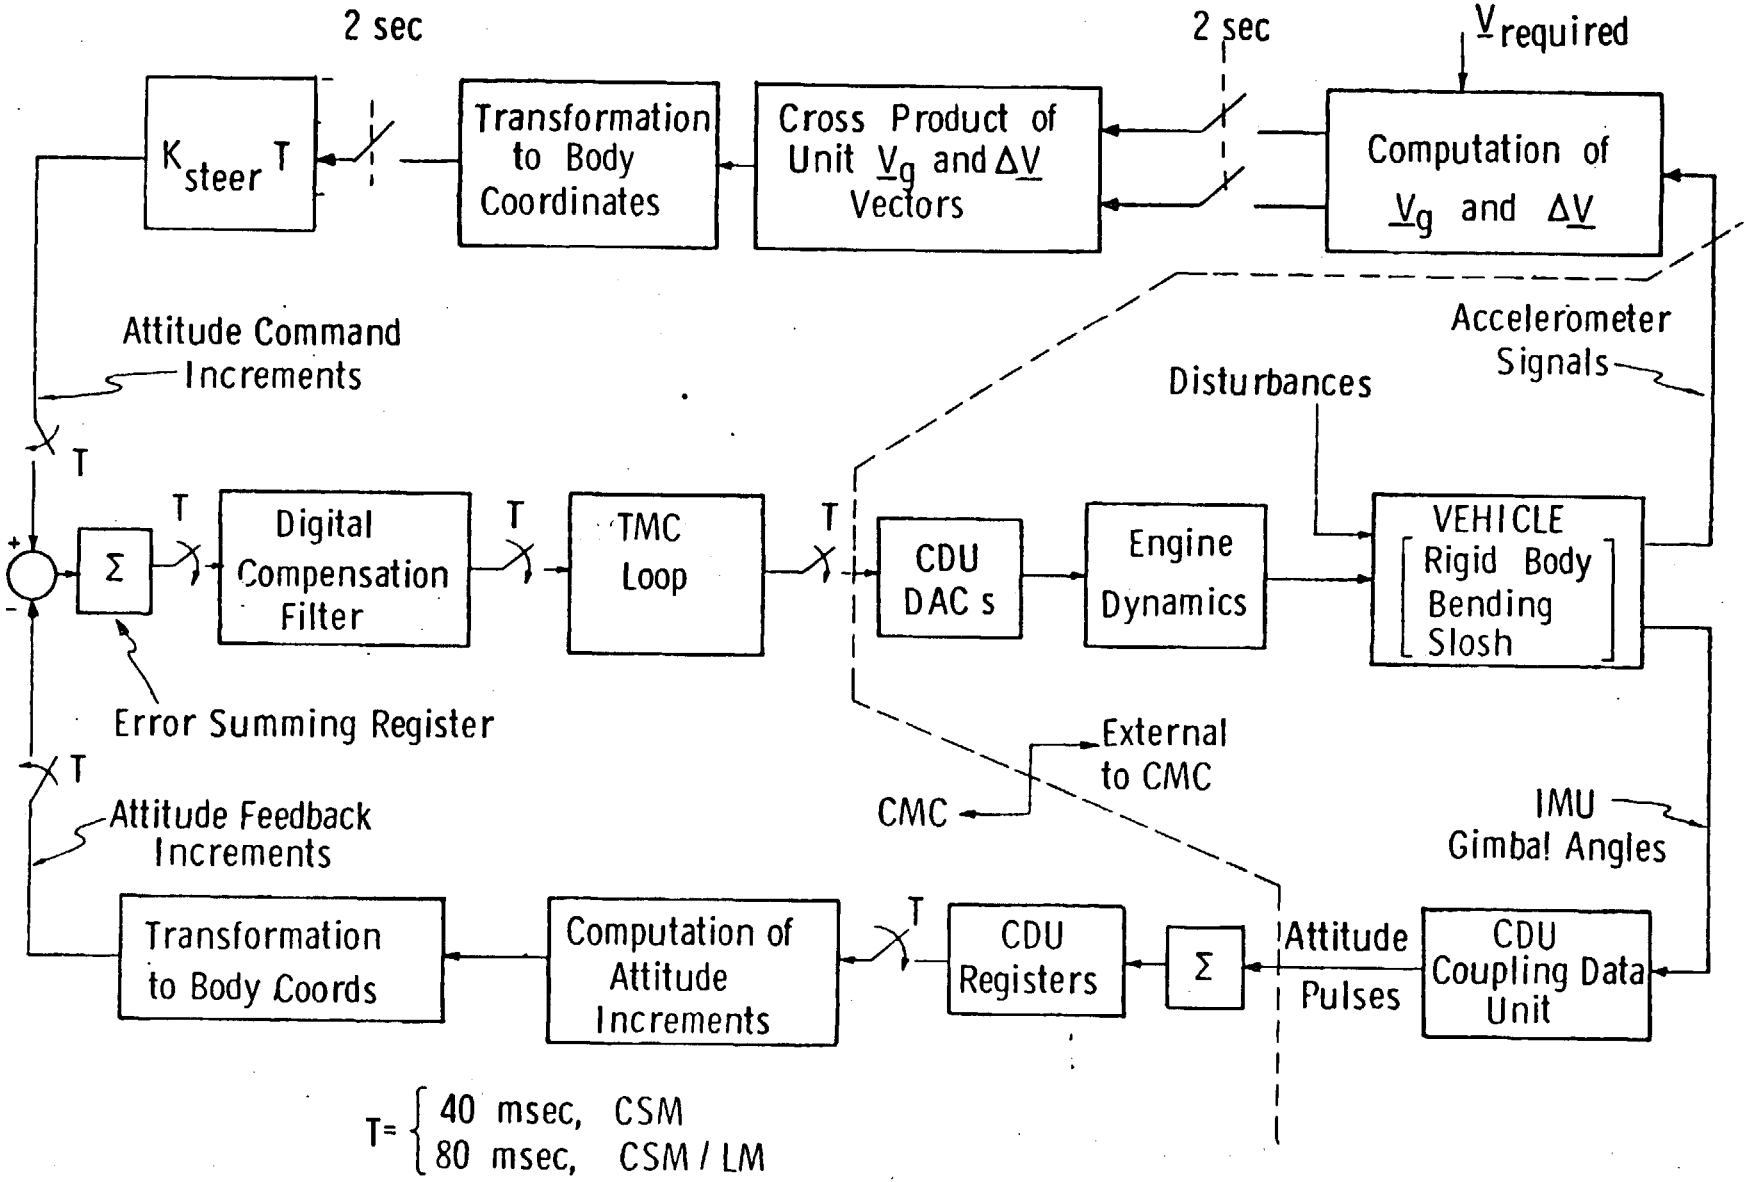

Figure D. 3-1 Functional Diagram of Thrust Vector Control System (Pitch or Yaw)

The AGC steering-loop computations generate incremental attitude commands in inertial coordinates and transform them into body coordinates.

A CDU measures and digitizes the gimbal angles of the IMU and transmits this information to the AGC, where they are stored in a CDU register. The CDU register is sampled regularly by the DAP program which back-differences the CDU angles to obtain the incremental changes over each sampling interval. The CDU increments over each sampling interval are then transformed into body coordinates and subtracted from the commanded increments generated by the steering program. The resultant differences are summed and represent attitude errors, expressed in body coordinates. The respective attitude errors are fed to the pitch and yaw compensation filters, whose outputs, together with estimated trim signals, are the commands to the engine-gimbal servos in pitch and yaw. The compensated signals control rapid transients and determine system bandwidth and the stability of the vehicle slosh and bending.

The CSM/LM DAP has two compensation filters. First is a high-bandwidth filter which stabilizes transients at engine ignition and slosh for the nominal bending effects and propellant loading expected during a lunar mission; this high bandwidth is achieved, however, at a very slight expense to the slosh-stability margin. Second is a low-bandwidth filter which provides poorer transient response, but stabilizes slosh for all propellant loadings. The autopilot begins a TVC burn in the high-bandwidth mode, and remains in that mode unless the astronaut initiates a switchover to the low-bandwidth mode. To retain the maximum advantage of the high-bandwidth filter, this switchover is performed only in the highly improbable case when an observed slosh instability leads to excessive engine oscillations. In the undocked CSM, where bending and slosh are less problematic, the DAP can utilize an even higher-bandwidth filter with considerable success.

As shown in Fig. D.3-2, the command signal to each engine-gimbal servo is comprised of the compensation-filter output and a bias (or trim) from a Thrust Misalignment Correction (TMC) loop. This loop trims the compensated command such that a zero output from the compensation filter will still cause the thrust vector to pass exactly through the center of gravity (cg), when there is no cg movement and no motion of the thrust vector relative to the commanded angle.

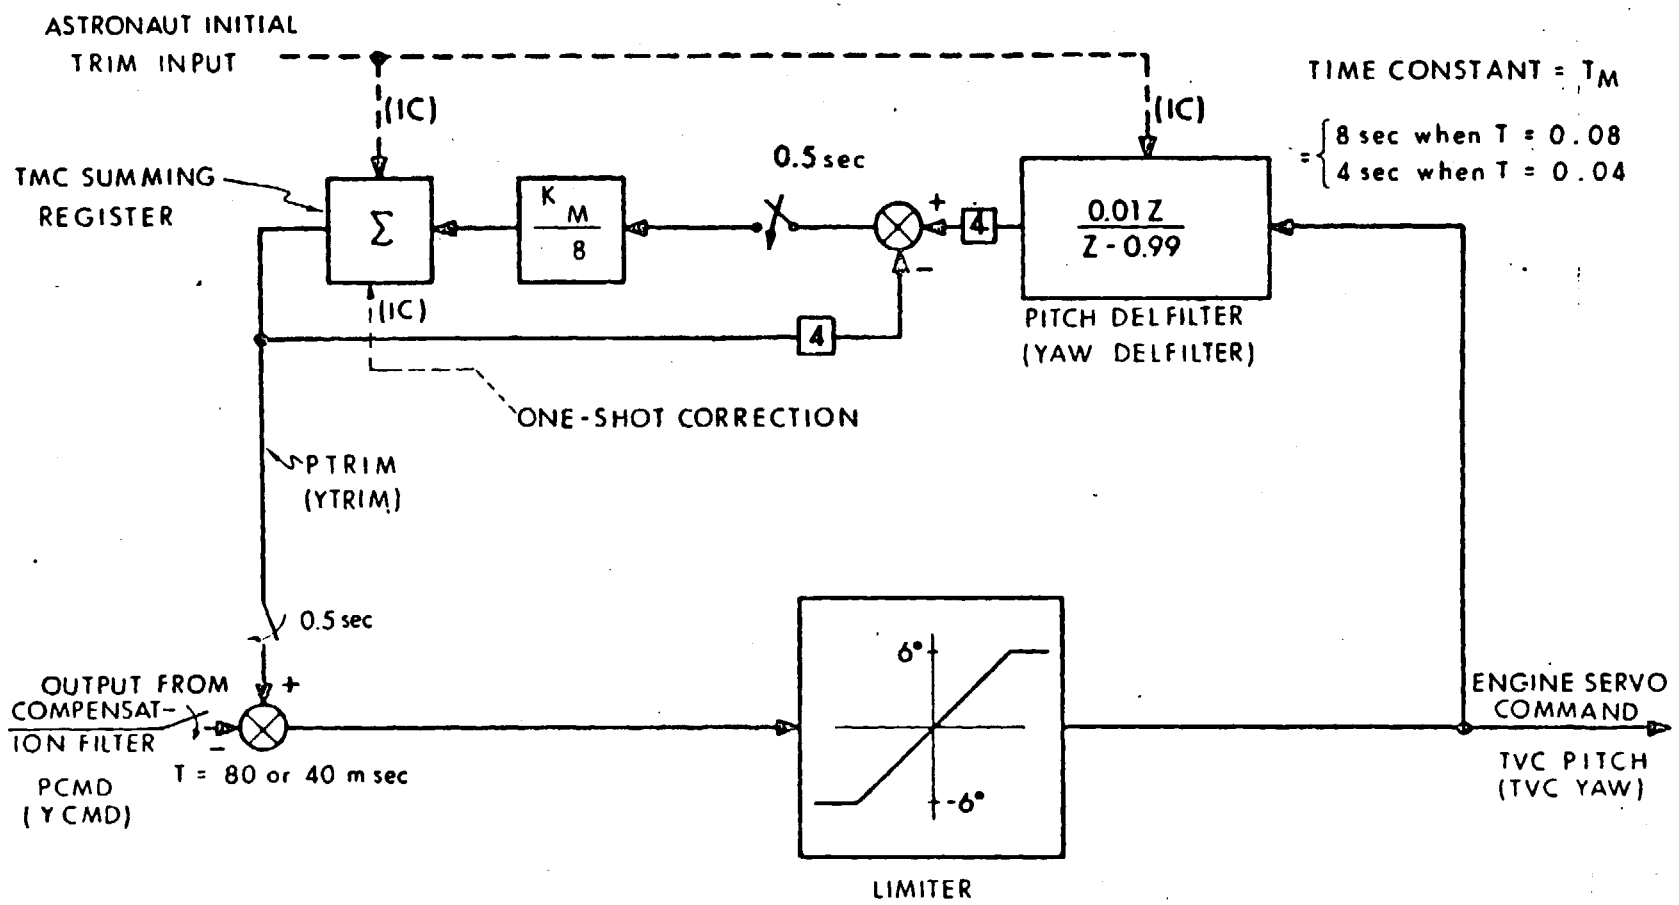

Figure D. 3-2 Functional Diagram of Thrust Misalignment Correction (TMC) Loop

The two major elements of the TMC loop are its summing register, which supplies the bias, and its low-pass filter, DELFILTER, which tracks the total command signal at autopilot sampling frequencies. The difference between the bias and the DELFILTER output is slowly integrated to correct for thrust-to-cg misalignment. This action is roughly equivalent to inserting a proportional-plus-integral transfer function between the compensation filter and the total command signal. The TMC loop is designed so that its dynamics have no effect on the vehicle bending and slosh modes and very little effect on the rigid-body stability.

The TMC summing register and DELFILTER are initialized at the beginning of a TVC burn by the astronaut. Later the summing register is reinitialized to implement a so-called one-shot correction.

In the case of the CSM DAP, the TMC summing register is not incremented until a one-shot correction is made 3 sec after engine ignition; this one-shot correction adds to the contents of the register twice the change which has occurred in the DELFILTER output (the factor of two being required to compensate for the transient lag of DELFILTER). Following this correction, the TMC summing register is incremented every 0.5 sec, as shown in Fig. D.3-2.

The CSM/LM DAP increments the TMC summing register from the beginning of the burn. However, the summing register is reinitialized in the event of a low-bandwidth switchover to the current value of DELFILTER (based on the assumption that the switchover occurs beyond the initial transient of DELFILTER). At switchover, the low-bandwidth filter is zeroed, so that the entire burden of supplying the servo command is shifted to the TMC summing register. Following switchover, the TMC loop continues to operate with the same gain, sampling frequencies and DELFILTER time constant.

The operation of the TVC DAP cannot be considered separately from that of the steering loop which interacts with the autopilot. The steering-loop operations are shown in Fig. D.3-1 for the case of External ΔV guidance (see Section C.1.2.5). This form of guidance—which has been used for the SPS burns in all the lunar missions to date—is based on achieving a commanded velocity change, ΔV_c, which is specified prior to each burn. This desired velocity change is inserted into the steering program as an initial value of the computed velocity-to-be-gained vector,

V_G, which is computed in inertially-fixed basic reference coordinates. The V_G vector is reduced during the course of the burn by subtracting the accelerometer-measured increments, ΔV. The ΔV increments are computed at 2-sec intervals by summing the accelerometer pulses accumulated over these intervals, and by transforming the resulting vector components from IMU coordinates to basic reference coordinates. The purpose of the steering loop is to align the vehicle thrust vector with the current velocity-to-be-gained vector. This is achieved by commanding a vehicle-turning rate which is proportional to the normalized cross-product of these two vectors. This cross product is computed in basic reference coordinates and transformed into body coordinates every 2 sec. (Section C.1.2.1 discusses cross-product steering for powered-flight guidance.)

The pitch- and yaw-axis components of the vector cross-product are multiplied by a proportionality constant, K_{steer}, to obtain the commanded rates about these axes—and are multiplied by the autopilot sampling frequency, T, to obtain attitude-command increments which are supplied to the autopilot every T seconds. (K_{steer} has three different values, employed respectively with the CSM DAP and the two modes of the CSM/LM DAP.)

D.3.1.2 TVC Roll Control

The TVC Roll DAP is designed to provide attitude and rate control about the roll axis using the RCS jets. Its function is strictly attitude hold. The orientation of the CSM about the roll axis is held within a specified deadband throughout the burn. The outer-gimbal angle of the IMU, which is parallel to the vehicle roll axis, is read and processed to yield approximate roll-attitude and roll-rate measurements. A switching logic in the phase plane is then used to generate commands to the RCS jets. (The Roll DAP will not be described further in this report so that adequate space can be given to the pitch and yaw autopilots, which have the major role in thrust vector and velocity control.)

D.3.2 Design Requirements of the TVC DAP

D.3.2.1 General Design Considerations

The primary requirement which the TVC pitch and yaw DAP programs must fulfill is to provide, in conjunction with the external-guidance loop, satisfactorily

small velocity-pointing errors at thrust cutoff. The DAPs must also limit excursions in vehicle attitude and in thrust-vector orientation to minimize propellant usage and gimbal-servo clutch wear and to facilitate pilot monitoring. The DAP programs must function satisfactorily with uncertain initial conditions and vehicle characteristics that vary during the burn. These factors are discussed below:

D.3.2.2 Initial Conditions and Time-Varying Thrust Misalignment

The TVC DAPs will experience several initial perturbations at SPS engine-ignition time:

a. Non-zero initial attitude errors in pitch and yaw—To avoid bending excitation in the CSM/LM configuration, the TVC DAP neglects initial attitude errors which could result from an ullage maneuver. In the undocked CSM, bending is less problematic and these errors are correctly initialized.

b. Non-zero initial attitude rates in pitch and yaw—Off-nominal RCS-jet performance during ullage may lead to attitude rates at SPS ignition time of up to 1 deg/sec.

c. Initial lateral slosh-mass displacement—Flight results have shown initial slosh-produced oscillations of up to 0.1 deg in the vehicle attitude.

d. Initial longitudinal propellant displacements (in the event of a no-ullage ignition.)

e. Thrust-vector misalignment—Before thrust initiation, the AGC supplies trim signals to the engine-gimbal servos to orient the anticipated thrust vector through the estimated cg position. Very likely, however, some error will occur in the alignment. There are two sources of thrust-vector misalignment—uncertainties in the thrust-vector orientation and uncertainties in the estimation of the cg position. These sources can yield a 3σ misalignment angle ranging from 1.4 deg (full) to 0.98 deg (empty) for the CSM, and 1.25 deg (full) to 0.71 deg (empty) for the CSM/LM.

In addition, due to propellant consumption, the cg position will vary with time. There is also a possibility of angular deflections of the thrust vector within the nozzle, as a result of uneven erosion.

The maximum predicted rates of change in the thrust-to-cg angle for the CSM/LM are about 0.003 deg/sec, in the pitch and yaw planes. For the CSM, the figures are 0.0083 deg/sec in the pitch plane, and 0.014 deg/sec in the yaw plane. Thrust-vector deflection due to nozzle erosion is estimated to be less than ±0.2 deg in any 20-sec interval. The total erosion deflection over a long burn will be within ±0.3 deg.

D.3.2.3 Vehicle Characteristics

The dynamic characteristics of the CSM/LM are sufficiently different from those of the CSM-alone to require separate autopilot programs tailored to the characteristics of each configuration. The three principal differences resulting from the two spacecraft configurations are as follows:

a. Bending-mode frequencies of the CSM/LM are as low as approximately 2 Hz; bending-mode frequencies of the CSM-alone begin at approximately 5 Hz.

b. Both in the moment arm from the gimballed engine to the center of gravity and in the vehicle moment of inertia, the two configurations differ substantially. Thus, a given deflection of the gimballed engine of the CSM-alone can produce as much as four times the angular acceleration as the same deflection in the CSM/LM vehicle.

c. The fuel and oxidizer slosh behavior in the CSM/LM vehicle differs from that in the CSM-alone because of the additional slosh masses in the LM tanks, the effects of the increased mass and moment of inertia of the overall vehicle, and the differences in cg location.

The stabilization of slosh and the avoidance of its excitation are of equal importance for both vehicles. The stabilization of the bending modes and the avoidance of their excitation is primarily a problem in the docked CSM/LM configuration.

D.3.2.4 Design Approach

The discrete, quantized nature of the digital computations have not presented any major impediments to the design of the TVC autopilots. The effects of quantization at the D/A and A/D interfaces have been found to be negligible in the cases of the CSM and CSM/LM TVC DAPs. The effects of finite word length and fixed decimal-point arithmetic in the Command Module AGC have been reduced to negligible levels in these DAPs by (1) careful attention to the manner in which the control equations are solved; (2) proper scaling of the computer variables; (3) use of double-precision variables where required; and (4) proper attention to the manner in which (1), (2), and (3) are combined with the selection of the sampling rates to maximize both the linear operating range and the precision of each digital-filtering operation. Conventional design techniques employing Z and W transforms, root loci and frequency response characteristics were found to be adequate for the design of the TVC autopilots.

D.3.3 TVC DAP Implementation

D.3.3.1 Compensation Filters

A basic TVC autopilot program provides a generalized sixth-order filter which consists of three cascaded second-order sections. The CSM/LM configuration uses all three sections; the CSM only two. The second-order factors for the cascade sections have been selected in such a way as to minimize the transient excursions of the signals between sections, thus allowing these signals to be scaled to take advantage of the available digital word lengths—and thereby minimizing the effects of round-off errors. Generally, it has been found best to group zeros and poles having similar frequencies together in the same cascade section. Where this has not been possible (as in the CSM/LM low-bandwidth third cascade), an attempt has been made to keep the steady-state gains of the numerator and denominator from becoming too dissimilar. (The steady-state gain of either the numerator or the denominator is found simply by setting $z = 1$ and adding the coefficient values.)

D.3.3.1.1 Switchover from High Bandwidth to Low Bandwidth

In the CSM/LM mode, provision is made for manual switchover from the normal high-bandwidth filter to a slower low-bandwidth filter. Once this switchover is

commanded, the computer calls a section of coding which is designed to (1) zero the filter storage locations; (2) update the TMC summing registers with new values of PTRIM and YTRIM based on the DELFILTER outputs; (3) load the low-bandwidth coefficients from fixed memory into erasable memory; and (4) load new values for the DAP gain, the TMC loop gain, and the steering gain.

No provision is made during a burn for returning to the high-bandwidth mode once the switchover has taken place. On the next burn, however, the TVC initialization logic reloads the high-bandwidth coefficients from erasable memory.

D.3.3.2 TVC DAP Variable Gains

The DAP gains are established initially and updated periodically using a small routine which is called every 10 sec and computes a piecewise-linear approximation to the curve for $I_{AVG}/T\ell_x$ versus SPS propellant weight. This value is then multiplied by the gain constant KTLX/I to obtain the TVC DAP gain, K_z. Consequently, the gain relationship for the TVC DAP is given by

$$K_z = (KTLX/I)\,(I_{AVG}/T\ell_x)$$

D.3.3.3 Trim Estimation

Three sources of trim information are provided to the TVC DAPs: (1) the initial values of PTRIM and YTRIM from a DSKY entry prior to the first burn; (2) a single-shot correction shortly after ignition (for CSM only), plus repetitive corrections during the burn (both for CSM and CSM/LM); and (3) an end-of-burn update at the engine-off command. In addition, a one-shot trim update is made in the CSM/LM mode if the high-bandwidth to low-bandwidth switchover is executed.

At the time of the engine-off command, a current update of the trim estimates is made. (This consists of picking off the pitch and yaw DELFILTER values and loading them into the trim registers.) The basis for this final update is that, for CSM/LM burns of less than about 25 sec, DELFILTER tracks the actual engine position faster than the full TMC loop—and therefore provides a better trim estimate for the next burn. Before each burn the astronaut reviews the computer-stored trim values for acceptability. Ordinarily he will not alter these trims unless the vehicle configuration has changed since the last burn.

D.3.3.4 Restart Protection

Much of the computer logic, including that of the TVC DAP, must be protected against restarts. (As discussed in Section 2.1.4.2, restarts are caused by such events as power transients or a parity fail on a memory-read instruction.)

Briefly, restart protection involves (1) storing the results of certain computations in temporary locations; (2) setting a flagword to indicate that the computation is completed; and then (3) performing a copy cycle to copy the computed results from the temporary registers into their normal registers. In this way, restarts occurring during a computation operation cause that operation to be repeated, while restarts occurring during a copy cycle require only that the copy cycle be repeated. In the TVC DAP there are copy cycles for the pitch and yaw channels, as well as for the DAP-related routines.

D.3.3.5 Computer Storage and Time Requirements

The total AGC memory used by the TVC DAP, including the Roll DAP, is about 1500 words. This breaks down to 1320 words of fixed memory, 26 words of nonsharable erasable memory (which must be preserved throughout the mission), and 154 sharable erasable words used for scratch-pad computation and temporary storage during TVC only. The pitch and yaw channels together require only about 500 words of fixed memory; the remaining 820 words are used by the Roll DAP and by DAP-related logic, such as the TVC initialization and monitoring routines, the mass-properties routine, and the restart routine.

The computer time used by the TVC DAPs is as follows: for the CSM, about 7 msec per channel per 40-msec sample, or about 35 percent of the available computer time; and for the CSM/LM, about 8 msec per channel per 80-msec sample, or about 20 percent of the available computer time. In addition to this, the time required for the combined Roll DAP and monitoring operations is about 10 msec every 0.5 sec.

D.3.3.6 Selection of Sampling Frequencies

The sampling frequencies of the TVC DAPs and their associated steering loops were selected as follows:

a. The sampling frequencies employed in the compensation filters and feedback loops of the autopilots were made high enough to ensure that none of the major bending modes of these vehicles would be subject to the so-called "folding effect". The CSM sampling frequency of 25 Hz (T = 40 msec) is roughly five times the minimum bending frequency of that vehicle. The CSM/LM sampling frequency of 12.5 Hz (T = 80 msec) is about five times the minimum bending frequency of the CSM/LM.

b. The TMC loop's low-pass filter, DELFILTER, requires a fairly high sampling frequency to attenuate the high-frequency components of the engine-servo command. It is convenient to operate this filter at the sampling frequencies employed by the autopilot compensation filters. The same DELFILTER coefficients are used for both autopilots. These coefficients produce a 4-sec time constant at the CSM sampling frequency and an 8-sec time constant at the sampling frequency of the CSM/LM. These time constants represent a compromise between the conflicting requirements of (a) attenuation of slosh oscillations, and (b) accurate tracking of slow variations in the servo command.

c. The TMC summing register is incremented at a low sampling frequency of 2 Hz, or approximately 12 rad/sec. This sampling frequency is adequate for the TMC loop, whose active frequency range is below 2 rad/sec.

d. The steering-loop computations generate the attitude rate command ($\dot{\theta}_c$) once every 2 sec—i.e., at a sampling frequency of 0.5 Hz or about 3 rad/sec. This sampling frequency is well beyond the requirements of the CSM and CSM/LM steering loops, whose open-loop crossover frequencies are 0.15 rad/sec and less.

e. The staircase waveform of $\dot{\theta}_c$ is smoothed by the process of generating command increments, $\dot{\theta}_c T$, to be summed at each sampling period. This smoothing process has been found essential for preventing an adverse interaction between steering-loop sampling and the slosh-mode oscillations—both of which occur in the same frequency range.

D.3.3.7 Effects of Computational Time Delays

All computational time delays have been neglected in deriving the transfer functions representing the autopilot and steering loops. The effects of these delays are examined below.

In the case of the autopilot loop, there is a computational delay of 3.3 msec between the time the IMU gimbal angles are read and the time the engine commands are released. This delay has a negligible effect on autopilot stability. For example, the maximum frequency at which the delay could have been of any importance is the 7.5 rad/sec maximum lead frequency of the CSM autopilot. At this frequency, the computational lag produces a phase shift of only 1.4 deg.

The TMC loop has a computational delay equal to one autopilot sampling period, T. This delay produces a negligible effect in the low-frequency range where the TMC loop is effective. It can be shown that the effects of this delay are so small as to be imperceptible on the plots of the open-loop characteristics of the autopilots.

A larger computational delay of about 0.4 sec occurs in the steering loop. However, the effects of this delay are negligible in the low-frequency range, where the steering loop interacts with the autopilot. For example, the delay has the largest effect in the case of the CSM/LM high-bandwidth mode, where it results in the open-loop characteristics being altered by less than 0.8 dB and 3 deg at any frequency between 0.1 rad/sec and 0.5 rad/sec.

D.3.4 TVC DAP Operation

D.3.4.1 Pre-burn Initialization

During an SPS burn, the functioning of the TVC DAPs is automatic, but there are several interfaces that must be properly established prior to ignition.

First, a small routine called the DAP Data Load may be called by the astronaut several minutes or more before the burn. This routine displays such information as (1) the masses of the CSM and LM which are used by the AGC to compute the autopilot gains; (2) the current engine trim angles to place the thrust vector through

the center of gravity; and (3) two flagwords which tell the computer whether the LM is docked or not, and which of the RCS jets to use for translational and roll thrusting. These displayed quantities may be accepted as is, or changed by keyboard entry if so desired. (For example, if the LM had just undocked from the CSM, the Command Module pilot would call up this routine and key in a new value for the flagword that indicates the docked/undocked configuration.)

The second important pre-burn function is the initialization of the digital-to-analog converters that transmit the pitch and yaw commands to test the SPS engine servos. This initialization procedure involves zeroing the D/A converters and energizing a relay to complete the electrical paths to the engine servos. The test entails commanding a sequence of ±2 deg deflections, both in the pitch and the yaw gimbal servos, which the astronaut can monitor on the SPS gimbal-angle indicator dial. Upon completion of the test, the trim values are commanded in preparation for ignition. Should the astronaut bypass the test, the trim angles will be commanded directly. Thus, at the end of this gimbal-trim routine, the SPS engine will have been aligned for ignition, and the servos will be energized and ready for the TVC commands during the burn.

A third pre-burn activity is the ullage maneuver in which the astronaut fires the +X translational RCS jets for about 20 sec preceding ignition to settle the liquid propellant in the tanks. (This maneuver is not required if the tanks are nearly full.)

D.3.4.2 Start-up Sequence

Following the pre-burn initialization, the TVC DAPs are started by a call from the thrusting program to the TVC initialization sequence about 0.4 sec after the ignition command. This delay is provided to accommodate for the delay between the ignition command from the AGC and the achievement of full engine thrust.

The DAP initialization provides for all remaining TVC preparation: (1) it zeros the erasables used for storing past values of the filter variables; (2) it loads the DAP coefficients and gain, the TMC loop gain, and the steering gain; (3) it initializes the TMC loop, the attitude-error integrators, and the DAP commands; and (4) it prepares the attitude-error needle display on the FDAI with a special initialization call.

With the initialization completed, the DAPs are ready to operate. The TVC program provides for time-separating the pitch and yaw computations, so that the Pitch DAP calls the Yaw DAP one-half sample period later, and vice-versa. This separation of the two channels ensures that other computer functions (e.g., steering computations and telemetry) can function at their assigned rates throughout the burn.

In conjunction with the regular cyclical operations of the Pitch and Yaw DAPs, the DAP-monitoring routine TVCEXECUTIVE comes up at 0.5-sec intervals to: (1) call the Roll DAP; (2) call the attitude-error needle display and provide it with an update; (3) update the DAP variable gains at 10-sec intervals; and (4) perform the one-shot and repetitive corrections for the TMC loop.

D.3.4.3 Shutdown Sequence

The engine-shutdown sequence originates in the thrusting program after the steering computations predict that the time-to-go to reach the desired cut-off velocity is less than 4 sec. At that time the engine cut-off time is computed, allowing for the expected thrust tail-off characteristics. Following the cut-off command from the AGC, the TVC DAP continues to function for about 2.5 sec, after which the RCS DAP is called in its attitude-hold mode.

D.4 CM Entry DAP

As explained in Section C.4, the aerodynamic-lift capabilities of the Command Module permit a controlled-entry flight to a designated landing point. The CM is a wingless, axially symmetric, reentry vehicle constructed with its center of gravity displaced from the axis of symmetry. When flying hypersonically in the atmosphere, the CM trims with a constant low ratio of lift to drag. The sole means for perturbing the trajectory in a controlled fashion is to roll about the wind axis (velocity vector) with the reaction-control jets, permitting the lift vector to be pointed anywhere in the plane perpendicular to the wind axis. The roll angle defines the orientation of the lift vector relative to the trajectory plane- i.e., the plane containing the wind axis and the position vectors. Down-range control is achieved via the component of lift in the trajectory plane; and cross-range control via the component of lift out of the trajectory plane.

In actual flight, an Entry DAP causes the CM to roll about the wind axis so that the actual lift direction is forced into agreement with the desired lift direction commanded by the entry-guidance equations. This results in achieving the desired in-plane component for down-range control. The rolling maneuver also yields an out-of-plane component of lift used for lateral-range control. When the cross-range error is minimized, the lateral-lift component becomes an unwanted by-product of the steering, and its effect is constrained to an acceptably small value by the guidance, which causes the CM to roll periodically so as to reverse the sign of and null the lateral drift. Since the in-plane component is the fundamental controlled quantity—in that it controls down-range flight—its sign normally remains unchanged during this nulling (lateral switching) process. The restriction on the sign of the in-plane component during lateral switching in effect requires that the CM roll through the smaller of the two possible angles (i.e., through the so-called shortest path) during lateral switching. The Entry DAP normally commands this type of maneuver. (In certain instances, where such a maneuver would cause the spacecraft to fall short of its target, the entry guidance demands a roll through the larger of the two angles, and accordingly informs the roll DAP.)

Entry DAP design is simplified by the fact that, within the atmosphere, the CM is aerodynamically stable. Since stability is no problem, only rate dampers are used in pitch and yaw. Furthermore, aerodynamic forces are utilized to do most of the work during a coordinated roll.

D.4.1 Exoatmospheric and Atmospheric Entry DAPs

The Entry autopilots are designed to perform automatically all maneuvers for all phases of entry flight starting with positioning the CM in the entry attitude prior to entry interface and continuing until drogue-chute deployment. Such capability requires several modes of operation, as illustrated functionally in Fig. D.4-1. The two basic modes are exoatmospheric and atmospheric. The atmosphere is defined to begin when the G forces acting upon the spacecraft exceed 0.05g.

In the exoatmospheric mode, the CM has three-axis attitude control based on the Euler set* R, β, α. The Euler set defines the angular attitude of the CM body axes (identical to standard aircraft axes, from the pilot's viewpoint) with respect

* The Euler set is trajectory-related and independent of inertial reference.

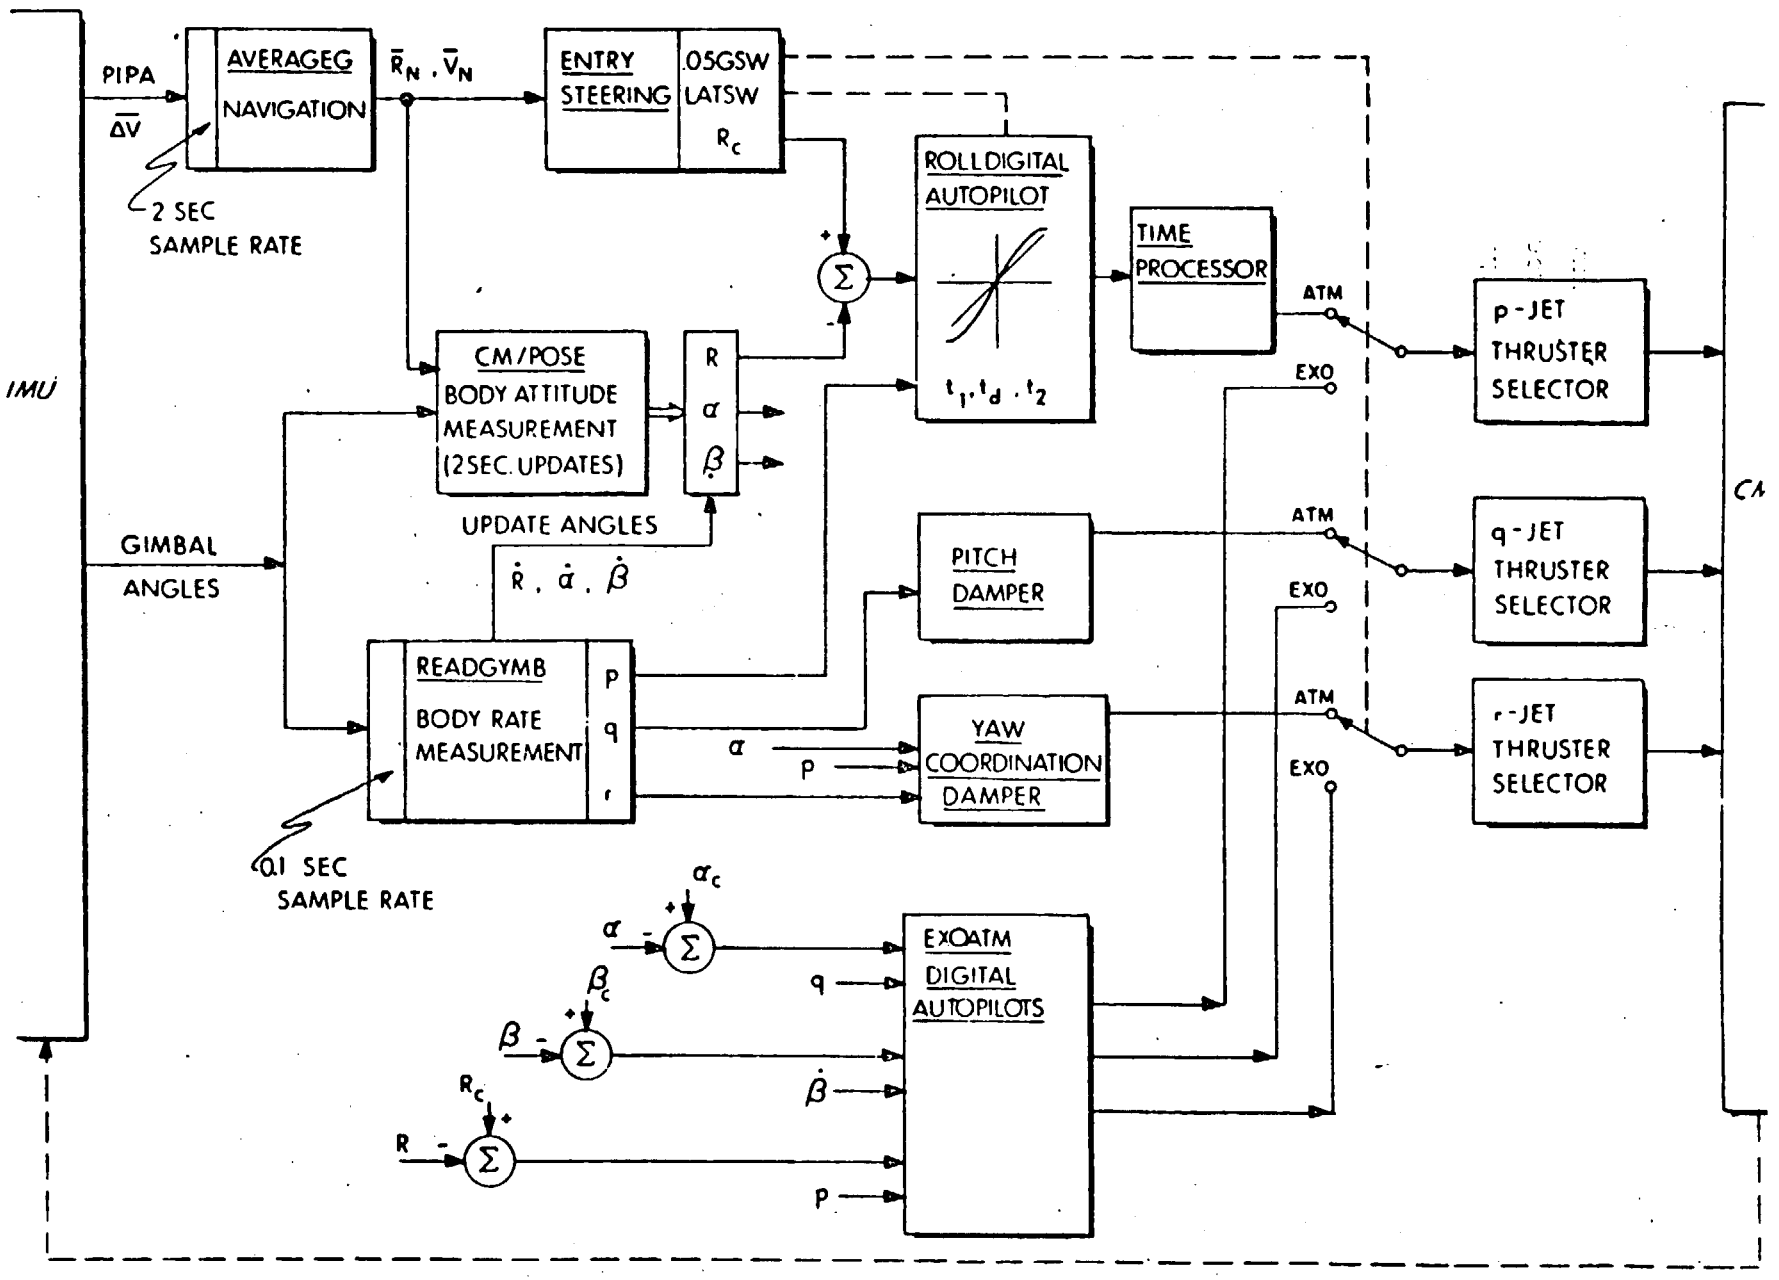

Figure D. 4-1 Functional Diagram of Entry Digital Autopilot

to a right-handed vector triad defined as having its X axis along the negative of the computed wind axis, and its Y axis normal to the trajectory plane, i.e., the cross-product of wind axis and position vector. The Euler sequence is roll R about the X axis, yaw β about the Z axis, and pitch α about the Y axis.

The exoatmospheric DAP maneuvers the CM from an arbitrary attitude into a local trim attitude relative to the local wind axis (computed from the current state vector). The maneuver is basically a pitch-over until α is about -20 deg—the hypersonic trim angle of attack.

At the same time, the CM yaws until β goes to 0 deg. The roll angle is held constant until $|\alpha|$ becomes less than 45 deg, at which time the CM does a coordinated roll maneuver about the wind axis until R goes to either 0 deg or 180 deg, as specified by the crew.

A policy based on the Euler attitude-rate equations is used to drive the Euler errors to zero. In the pitch DAP, $\dot{\alpha}$ is considered to be equal to pitch rate, q; $\alpha_c - \alpha$ and $\dot{\alpha}$ are used to command the pitch (q-axis) jets. In the roll and yaw DAPs, the attitude rates, $\dot{R}$ and $\dot{\beta}$, are considered to be orthogonal axes rotated through the angle α with respect to the orthogonal axes, p and r. To decouple the roll and yaw axes, the attitude rates, $\dot{R}$ and $\dot{\beta}$, are assigned to the nearest jet axis, p or r. Thus, for $|\alpha| > 45$ deg, the roll DAP, using $R_c - R$ and $\dot{R}$, fires its yaw (r-axis) jets, and the yaw DAP, using $\beta_c - \beta$ and $\dot{\beta}$, fires the roll (p-axis) jets. When $|\alpha| < 45$ deg, the roll DAP, using $R_c - R$ and $\dot{R}$, fires the roll (p-axis) jets, and the yaw DAP, using $\beta_c - \beta$ and $\dot{\beta}$, fires the yaw (r-axis) jets.

Each navigation cycle, the Euler angles are computed from state-vector data, and are compensated for computation time delay. Between the navigation cycles, the angles are updated at each 0.1-sec DAP cycle by integrating the Euler attitude rates obtained by resolution from the CM body rates. The atmospheric mode is selected whenever the atmospheric drag exceeds 0.05g. The CM measurement continues as before.

The exoatmospheric DAP drives α and β to the commanded hypersonic-trim values. In the atmosphere, these angles are essentially angle of attack and angle of side slip, and the CM is controlled by orienting the lift vector. The amount of

lift available for control purposes (i.e., the ratio of lift to drag), is determined by the cg location. Preflight ballast procedures ensure that the Z_{cg} offset will give the desired lift-to-drag ratio, L/D.

Within the atmosphere, aerodynamic forces tend to keep β essentially zero, and α essentially at trim α_t, so that attitude control is required only in roll. As the CM rolls, its X axis is constrained to roll about the wind axis at the angle α_t. This coordinated maneuver requires that yaw rate r be equal to p tan α_t. For both the pitch and yaw axes, rate damping maintains pitch and yaw angular rates within prescribed limits.

D.4.2 Phase-Plane Logic

The Entry DAP is, in reality, six separate digital attitude controllers—three for the exoatmospheric mode and three for the atmospheric mode. There is some consolidation, in that several controllers use common phase-plane logic.

Three axes of the exoatmospheric DAP operate each 0.1 sec and use the phase-plane logic of Fig. D.4-2. The X axis is attitude error, such as $\beta_c - \beta$; the Y axis is attitude rate, such as $\dot{\beta}$. The biased deadzone is utilized to obtain a minimum limit-cycle frequency; however, since the logic is used on a sampled basis, the deadzone is effectively enlarged by one sample time and proportionally as shown in the lower portion of Fig. D.4-2. The logic is constructed so that errors are reduced at a rate between 2 deg/sec and 4 deg/sec.

The behavior of the logic is illustrated by the upper portion of Fig. D.4-2. If, at the DAP sample time, the error and error rate correspond to a point in the shaded area, the indicated jet is turned on; should the point lie in the clear area, the jet is turned off. A typical trajectory is illustrated, and the dots represent updates. Between DAP updates, the jets remain on, if already on, and off, if already off.

The foregoing DAP logic is valid for both single-ring* and dual-ring thruster operation. However, dual-ring operation uses about twice the propellant as single-ring.

* The CM has two independent sets (rings) of thrusters to provide the redundancy necessary to assure safe operation of the jets.

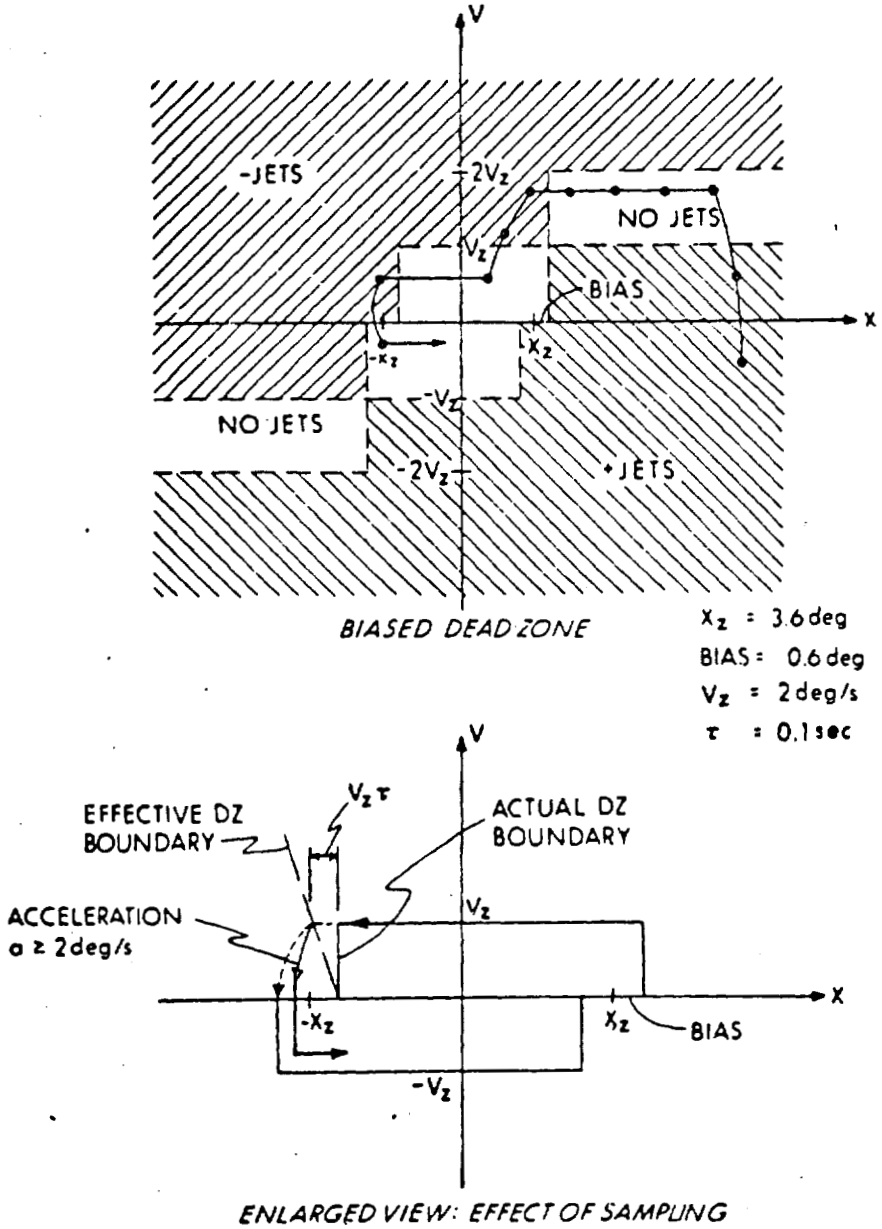

Figure D. 4-2 Exoatmospheric Phase-Plane Logic

When $|\alpha|$ becomes less than 45 deg, the roll-axis logic changes over to the 2-sec predictive phase-plane logic described below for the atmospheric DAP.

The atmospheric logic is such that when the pitch rate, q, exceeds 2 deg/sec at a DAP update, the proper q jet is fired; for yaw coordination and damping, when the combination $r - p \tan \alpha_t$ exceeds 2 deg/sec, the proper yaw (r-axis) jet is fired. For both pitch and yaw, if the rate is less than 2 deg/sec, the jet is turned off. As with the exoatmospheric DAP, the jets are changed only at DAP updates.

Unlike the pitch and yaw axes, the roll axis is controlled by a 2-sec predictive DAP, which becomes active during the exoatmospheric mode, when $|\alpha|$ became less than 45 deg; thus, it is already operative when the atmosphere is encountered. Every 2 sec during the entry phase, entry guidance provides a roll command R_c. To do this, it examines current vehicle position and velocity and also landing-point position, decides on the proper orientation of the lift vector, and generates the commanded roll attitude necessary to achieve that orientation. The roll autopilot uses the command R_c, the present roll attitude R, and roll rate p to generate firing times for the jet thrusters. In general, three time intervals are generated each 2 sec—two are thrust durations, and one is quiescent duration.

The vehicle roll attitude, in response to applied roll RCS torque, is modeled adequately by considering only the torque due to moment of inertia, thus permitting use of a phase plane wherein motion can be described using only straight lines and parabolas. The (X, V) phase plane of Fig. D.4-3 is used where the roll attitude error is $X = R_c - R$, and the attitude rate is $V = \dot{R}$. When the CM is at the commanded attitude, both X and V are zero. Note that roll rate p is related to $\dot{R}$ by $\cos \alpha$ in a coordinated roll. The inclusion of the $\cos \alpha$ term ensures a coordinated roll such that $p \rightarrow 0$ as $X \rightarrow 0$. The X axis is selected for the $\cos \alpha$ term; consequently, $V = p$ and $X = (R_c - R)\cos \alpha$ in the following discussion. This choice, though originally made for simplicity so that only the error X depends on allowing the velocity limits and deadzone limits to be expressed in units of body rate, is also desirable when platform misalignment is considered. In this case, as the state-vector error begins to accumulate, errors are introduced into the indicated body-attitude angles such that the indicated α can go through 90 deg. In this region, the $\cos \alpha$ in the DAP approaches zero and turns off the roll-axis attitude control. Since, in this region, the effect of p on indicated roll attitude R is insignificant, the DAP is not used to

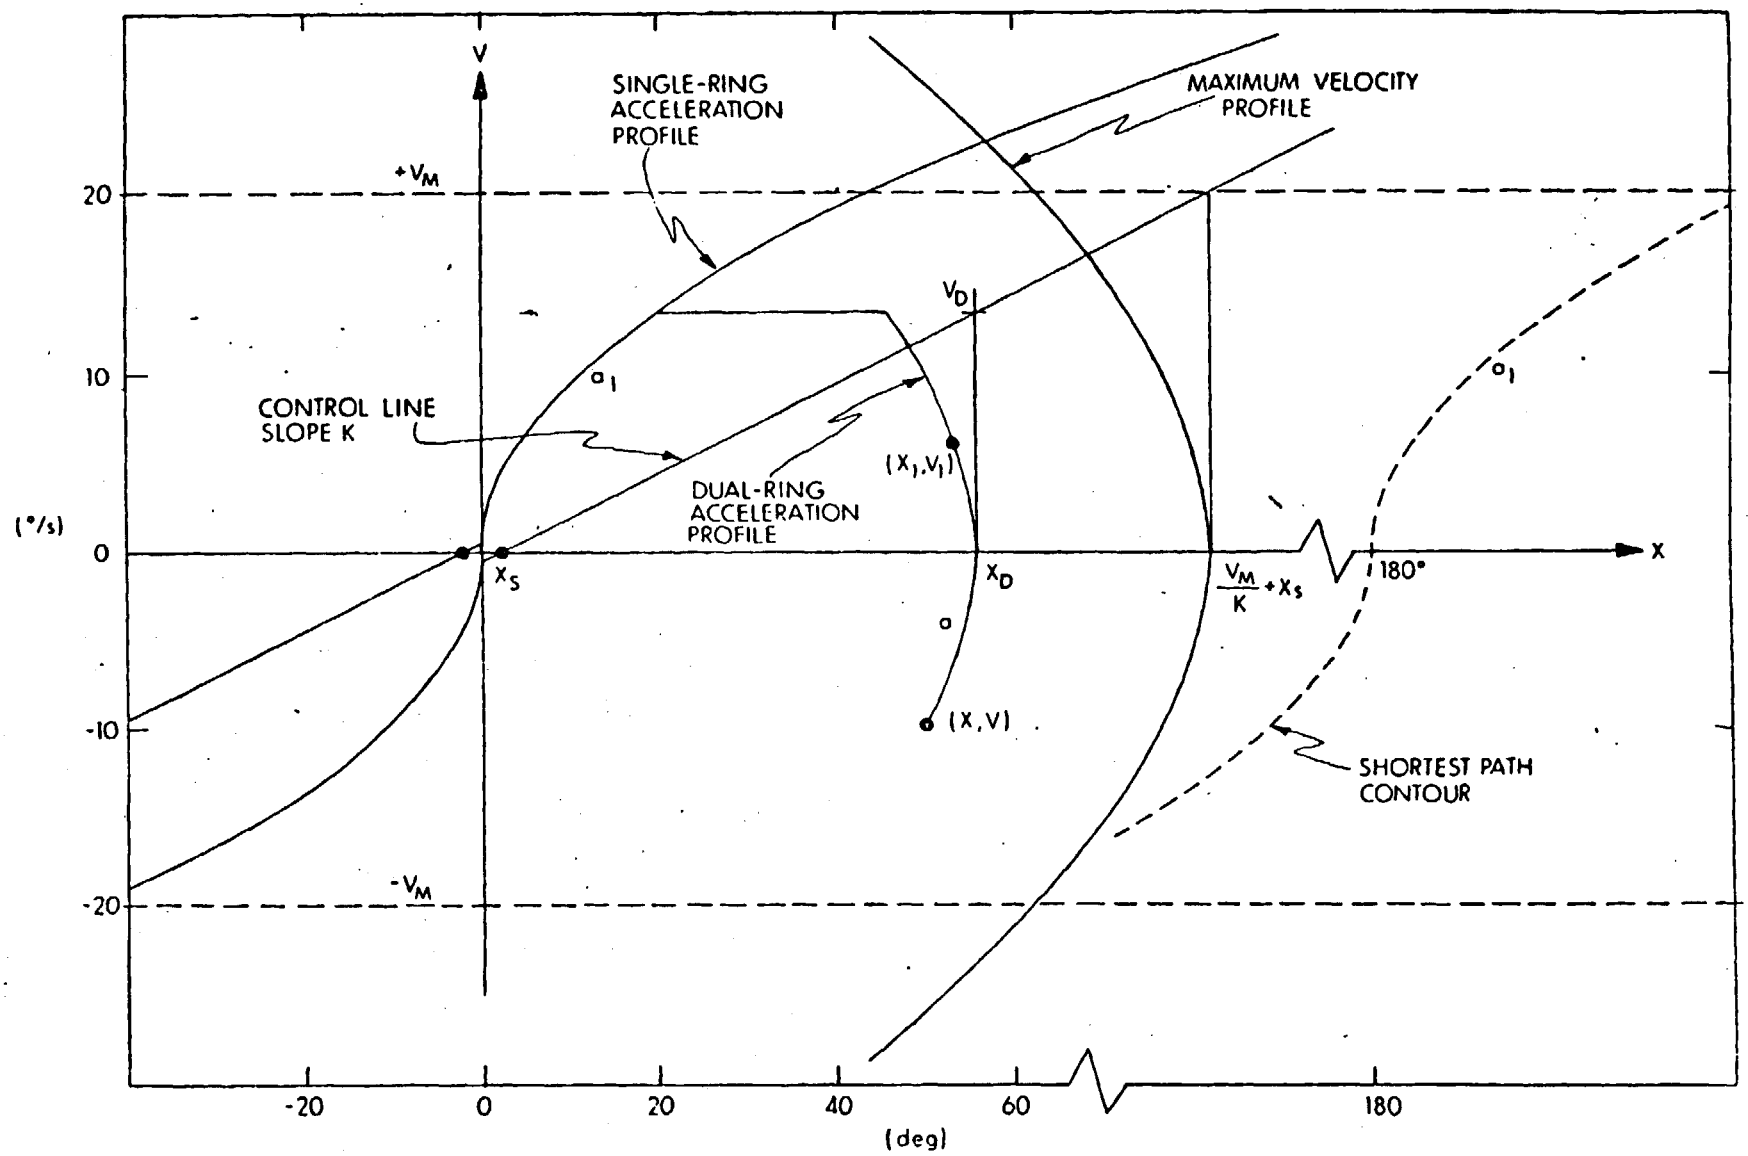

Figure D. 4-3 Hybrid-Gain Roll-Attitude Phase Plane

generate roll rates, and the CM will drift at deadzone values of p. When α increases further, cos α becomes nonzero, and the DAP resumes control. The sign reversal of cos α maintains stability. When the DAP is used in the exoatmospheric mode, inclusion of α allows a coordinated roll maneuver to be generated, even though α is changing.

The roll phase-plane logic illustrates two important design considerations. The first feature of the control technique is the use of a line of slope K to determine the drift rate at which an error is reduced. A fundamental and motivating advantage of this method is that fuel consumption (ΔV) becomes somewhat proportional to error, in that the CM responds rapidly to large error and more slowly to small error. Slope K is chosen to be as small as is consistent with the rate of response needed by the guidance during entry. Also, the straight line simplifies the equations for the jet-firing times and also allows the DAP to follow ramp inputs efficiently by establishing the necessary rate. The second feature of the control technique is the construction of the predicted trajectory from the DAP update point to the origin, using parabolas of different acceleration, to minimize the sensitivity of the control system to actual jet-thruster acceleration. This latter design consideration ensures not only stable operation, but comparable transient response behavior—even in the presence of a possible error of 100 percent in the control authority. Such a contingency could occur as follows: To allow rapid crew detection of jet failure during entry, only one ring of thrusters is active at a time. However, should circumstances dictate the use of both thruster rings, the DAP is required to perform in a stable and comparable fashion—even though it has no knowledge of whether single or dual rings are operative. This requirement is met by imposing a hybrid phase-plane profile: the first thrusting interval assumes dual-ring acceleration; a nonthrusting drift interval follows; and then a second thrusting interval assumes single-ring acceleration.

Figure D.4-3 is a simplified illustration of the roll logic. The prediction is based on the typical trajectory drawn from the point X, V to the origin. Assuming an initial acceleration a, the X-axis intercept, X_D, is obtained and is projected on to the control line to yield the drift velocity V_D. Using dual-ring acceleration, a, and the velocity difference of V_D and V, the first firing time, t, is computed for the +p-jet thruster. The firing time t_2 for the -p-jet thruster uses acceleration a_1

(single-ring acceleration) and velocity V_D. The drift time t_D is determined from the velocity V_D and the error to be covered at this velocity. The procedure is the same for any point below the control line. For example, the point (X_1, V_1) in the figure yields the same X_D and hence the same V_D as above; however, because the velocity correction needed, $(V_D - V)$, is smaller, the first firing time t_1 is smaller.

If the point (X, V) is such that $(V_D - V)$ is negative, the first burn is omitted. If the point (X, V) lies to the right of the maximum velocity trajectory, V_D is defined to be V_M (20 deg/sec) to provide a roll-rate limit. Once the time intervals t_1, t_2 and t_D are computed, only the first 2 sec of the trajectory are implemented. Each subsequent 2-sec DAP update will compute a new trajectory.

Since the actual acceleration of the CM is always either a or a_1, the roll response differs from the predicted trajectory as follows:

First, consider the dual-ring response to this phase-plane logic illustrated in Fig. D.4-4. The actual initial acceleration equals the assumed. V continues to rise to the V_D computed from the X intercept, X_D. Because of the dual-ring acceleration, both t_1 and t_D were computed correctly from the beginning. However, at the point where the trajectory intersects the single-ring profile, t_2 is computed based upon half the prevailing acceleration. The trajectory goes barreling down, but is stopped—before it goes too far—by a regular 2-sec update. Since the trajectory is again below the control line, a new t_1 is computed, and the trajectory rises to a new V_D determined by the new X intercept; the trajectory drifts toward the ordinate until it again intersects the single-ring profile; and then this time it goes down twice the required distance because of the dual acceleration and the absence of an update. This process is repeated until the trajectory is securely within a deadzone (to be described later).

Second, consider the single-ring response to the phase-plane logic, illustrated in Fig. D.4-5. In this case, which is the nominal one, the actual acceleration is half the assumed. Each vertical rise is not only half what is predicted, but the drift period is entered at a lower velocity. By the third update, the trajectory is above the control line, and the trajectory drifts toward the single-ring acceleration profile. At intersection, the trajectory follows this profile directly to the origin, without overshoot.

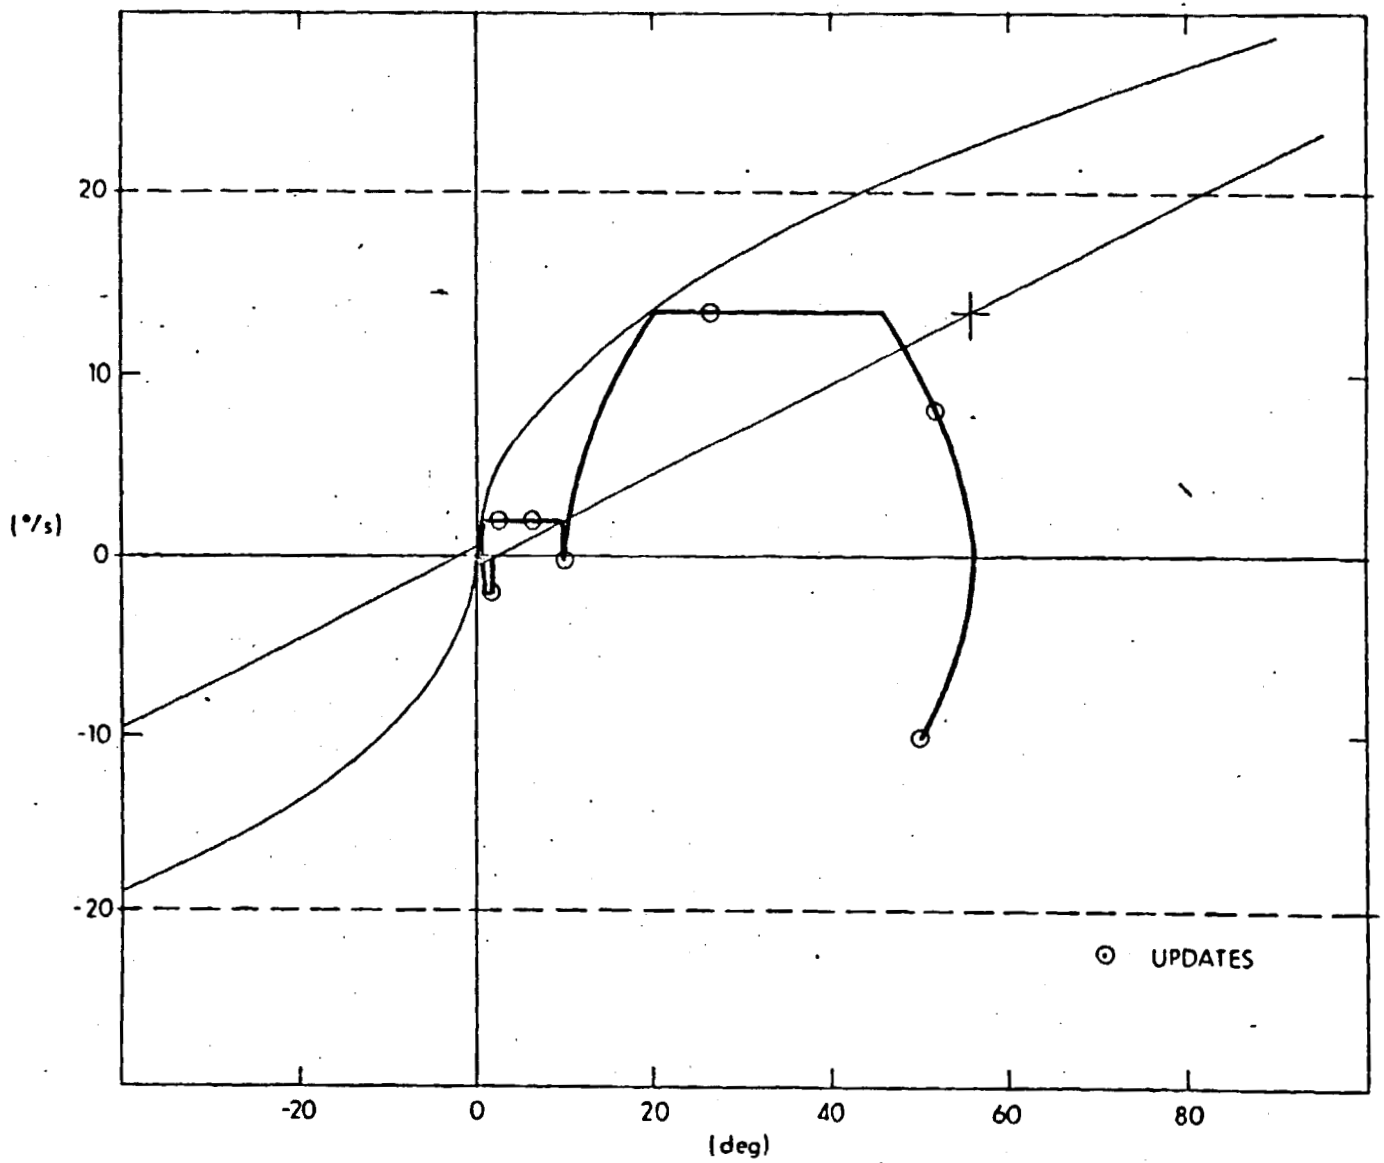

Figure D. 4-4 Dual-Ring Response to Roll-Attitude Phase Plane

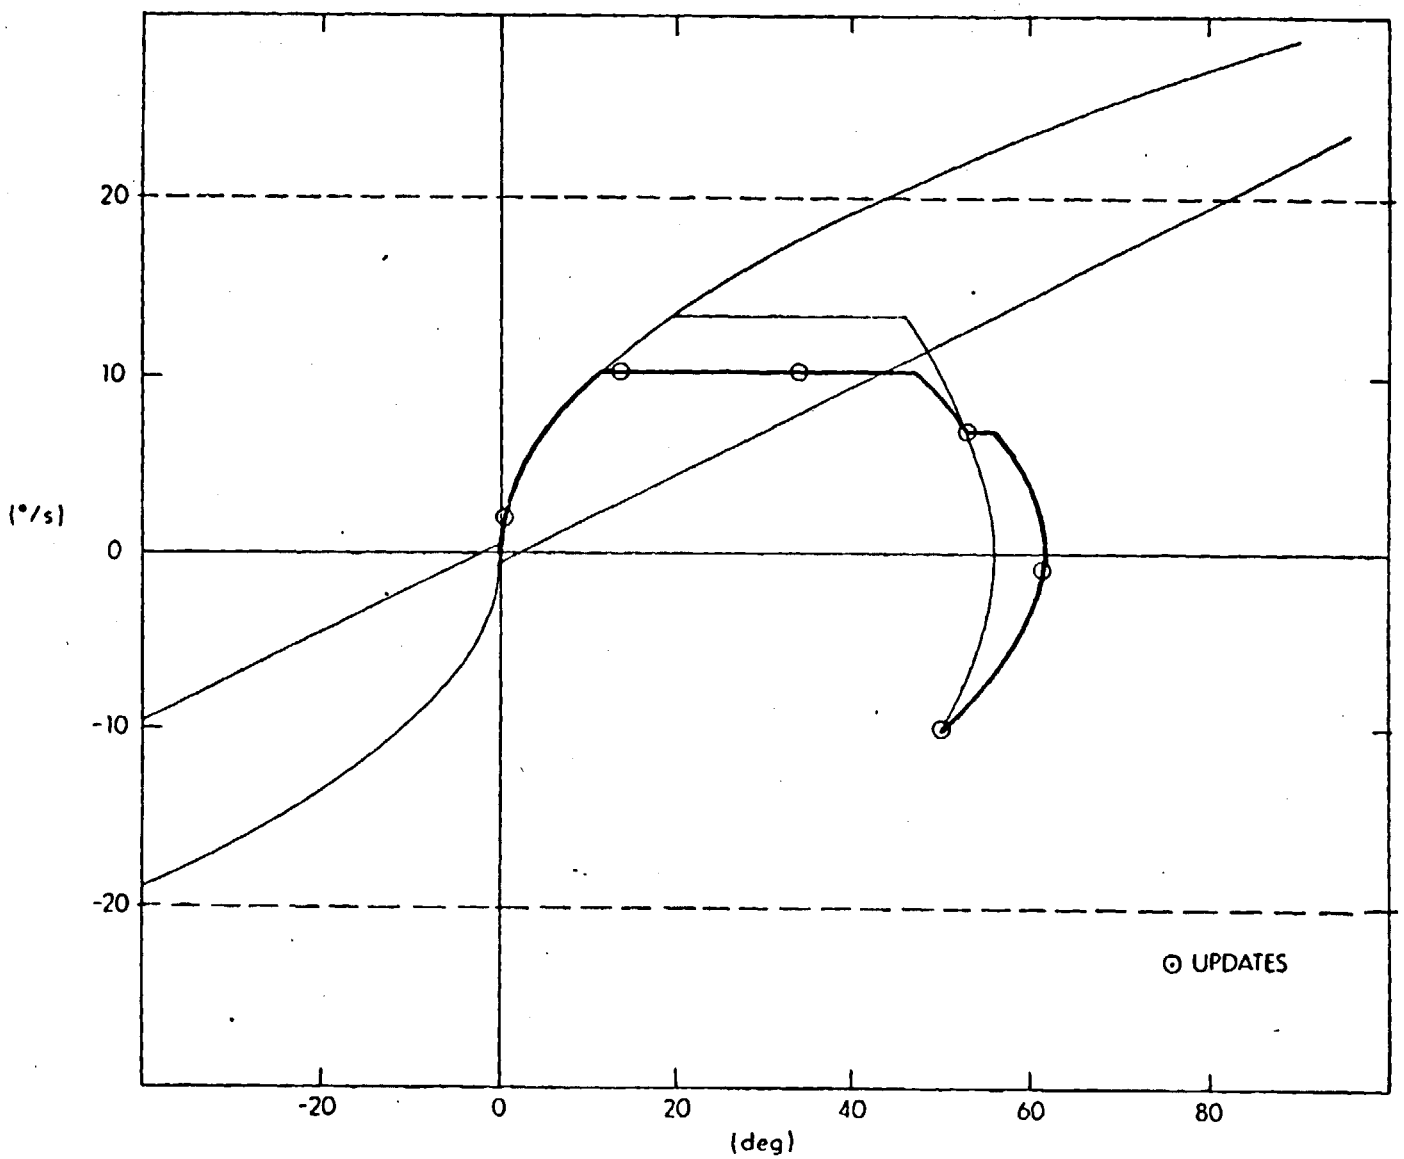

Figure D. 4-5 Single-Ring Response to Roll-Attitude Phase Plane

Figure D.4-6 illustrates the transient response of the DAP with predicted, dual-ring and single-ring conditions. Despite the wide variance in conditions, the DAP produces comparable solution times. Furthermore, simulations and actual flight* experience show that the fuel penalty for using dual rings is only about 20 percent—rather than the 100 percent one might anticipate. Figure D.4-6 also clearly illustrates that both dual-ring and single-ring response are slower than the predicted, but this is the price paid for indifference to actual acceleration. However, what is bought is stable operation over a two-to-one acceleration variation—a property that does not exist if both predicted thrust intervals are based on the same acceleration.

D.4.2.1 Shortest-Path Logic

As mentioned earlier, the roll DAP normally chooses the smaller of the two possible angles through which to roll to null the roll error. This is performed by using the shortest-path contour shown in Fig. D.4-3 at 180 deg and the corresponding contour at -180 deg.

The shortest angular path test consists of determining whether the point (X, V) lies within or without the contours at ± 180 deg. Any point (X, V) inside the contours considers the origin as its terminal point. Points outside these contours consider ± 360 deg as their origin. Such points X are shifted by -360 sgn(X) and thereby appear inside the contours as far as the phase-plane logic is concerned. Such a shift is necessary since, physically, -360, 0 and 360 are the same attitude. Furthermore, it is necessary that the contour dividing the regions of stable nodes be dynamic, rather than geometric—otherwise, there exists a region between the dynamic and the geometric contours where points will initially head for the origin as node; but the trajectory will carry across the geometric contour and then head for 360 as the proper node. This results not only in taking the longer path, but also in taking a longer time to do so.

D.4.2.2 Buffer Zone and Deadzone of the Roll-Attitude Phase Plane

As shown in Fig. D.4-7, the DAP has a deadzone at the origin, for the purpose of eliminating a high-frequency limit cycle. Its shape was chosen to provide a

*Apollo 7 flew a dual-ring reentry.

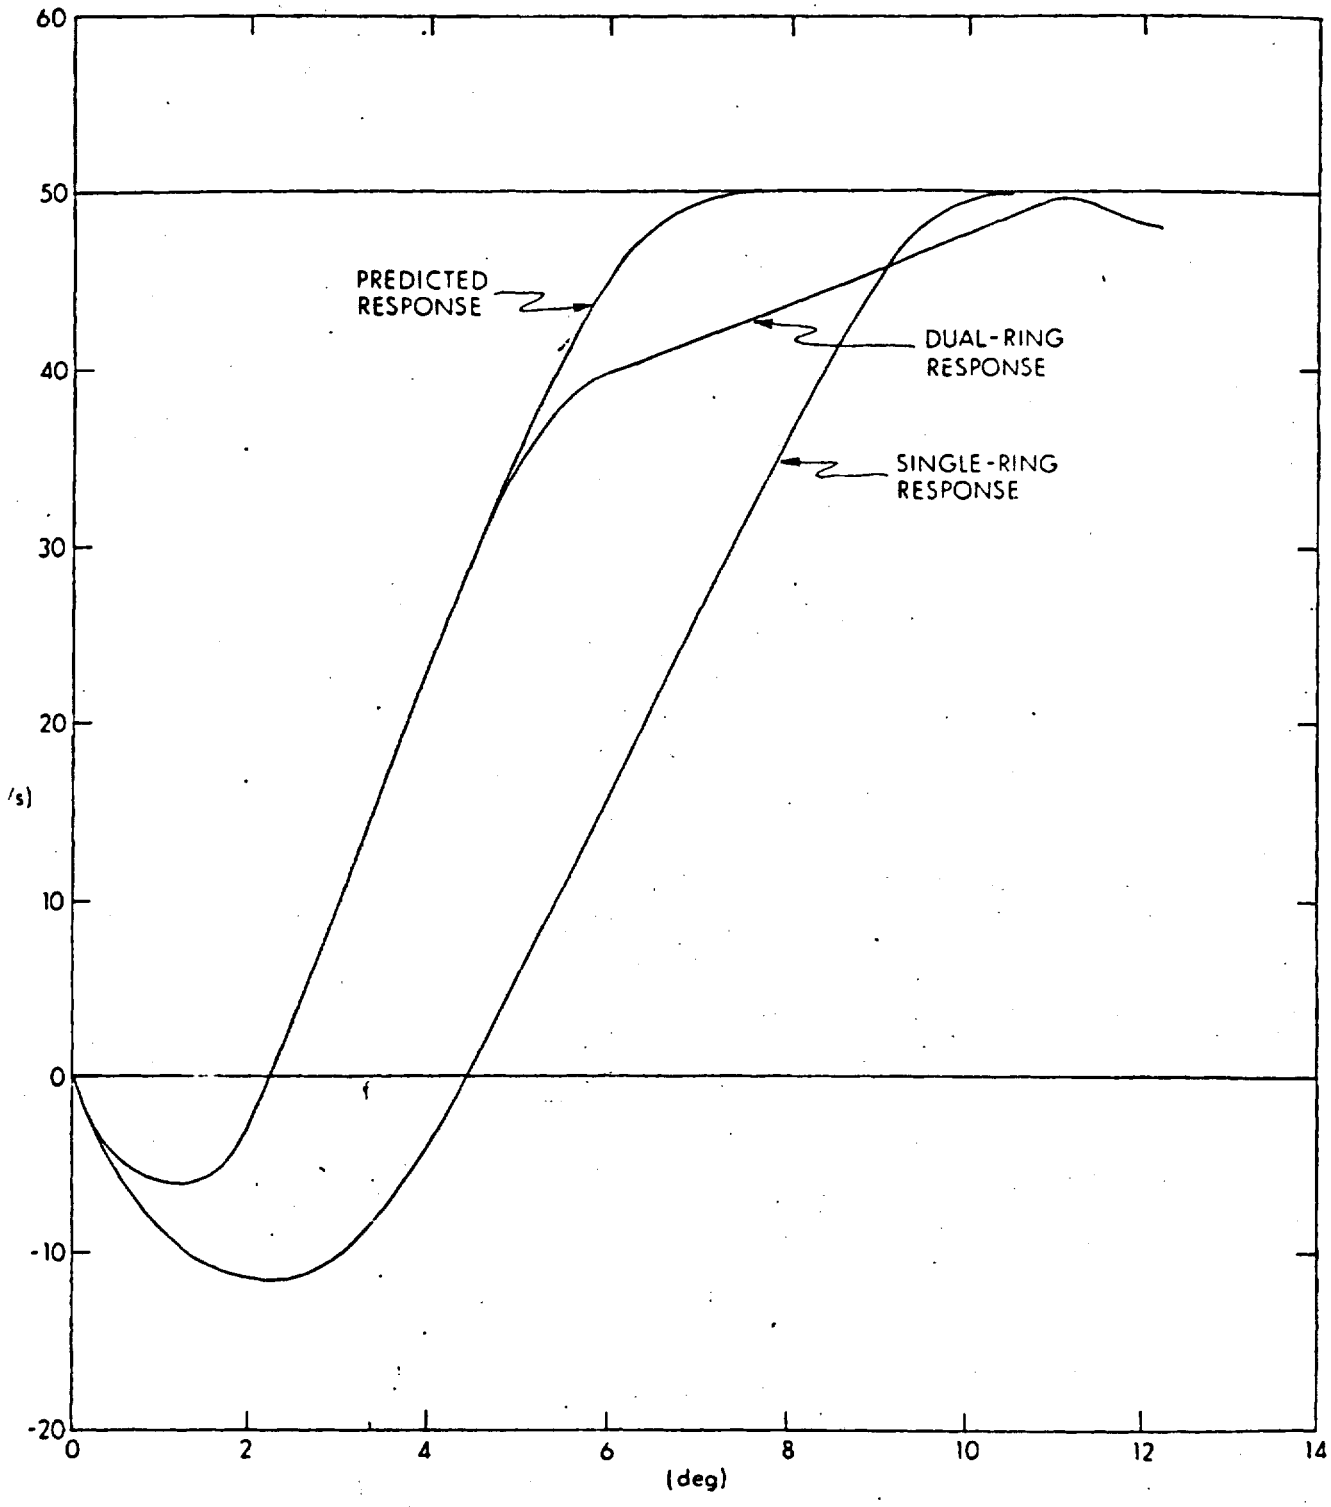

Figure D. 4-6 Transient Response of Roll-Attitude Phase Plane for Dual-Ring and Single-Ring Conditions

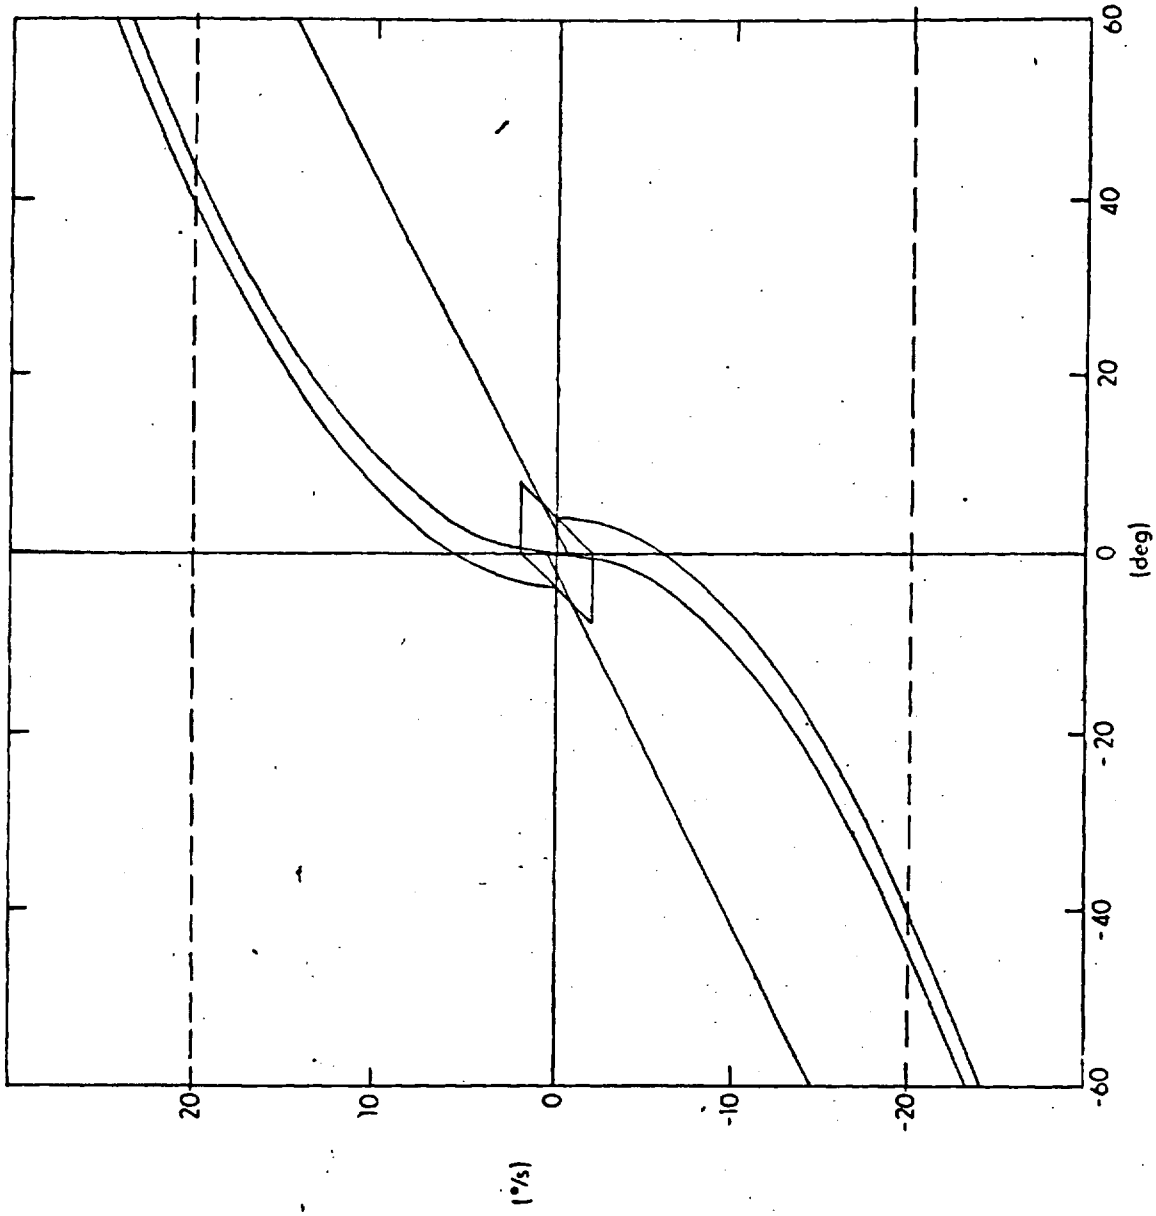

Figure D. 4-7 Buffer Zone and Deadzone of the Roll-Attitude Phase Plane

smaller limit-cycle amplitude for a given zone width. Also, the terminal trajectory has a band or buffer zone along which it channels trajectories to the deadzone. The buffer zone overcomes the effect of noise in body-rate measurements and deviations in the actual acceleration from the nominal DAP design value, a_1.

D.4.3 Entry DAP Displays

The DSKY displays associated with the entry mission-control programs were mentioned briefly in Section C.4.2. Additionally, the autopilot provides the following displays to the FDAI attitude-error needles:

Exoatmospheric DAP	The three attitude errors, $R_c - R$, $\beta_c - \beta$, and $\alpha_c - \alpha$, are presented each 0.2 sec. The error used by the autopilot to fire the roll jets is displayed on the roll-error needle, the error used to fire the pitch jets on the pitch-error needle, and the error used for the yaw jets on the yaw-error needle. If $\|\alpha\| > 135$ deg, the roll DAP does rate damping only, and the roll-error needle is zeroed every 2 sec.
Atmospheric DAP	The only presentation is roll error, corrected for shortest angular path, on a 0.2-sec basis. The pitch- and yaw-error needles are not driven and are at null.

To avoid hitting the needle limits, the maximum deflection allowed is 67.5 deg for roll and 16.875 deg for both pitch and yaw.

D.4.4 Manual Override

No provision exists for the use of manual controls, i.e., hand controllers, in the primary GN&C system during entry; consequently, if the astronaut chooses to perform a manual maneuver, e.g., to avoid gimbal lock, he must switch to the backup control system to override the primary system. In this event, the Entry DAP merely ignores the override and continues to provide FDAI attitude-error displays and jet commands based upon prevailing CM attitude and rates. Hardware in the backup system prevents the GN&CS jet commands from reaching the solenoid drivers until the override is terminated. Consequently, during a manual override, the Entry DAP remains continually prepared (whenever GN&CS operation is restored) to resume control with an uninterrupted knowledge of the trajectory and prevailing conditions.

D.5 AGC Takeover of Saturn Steering

Prior to the flight of Apollo 10, MIT was requesied by NASA to provide a boost-takeover capability in the event of a Saturn-Launch-Vehicle (LV) stable-platform failure. The task consisted of supplying an attitude reference and certain guidance information to ensure that, for a failure anytime after liftoff, the spacecraft could achieve earth orbit with help both from the Command Module AGC and from the crew.

For a takeover, the Saturn Instrument Unit (IU) must first sense that its platform has failed, and must call for the AGC backup mode by lighting the LV guidance-failure light in the cockpit. Once notified of a failure, the crew switches the LV guidance to GN&CS control.

With the LV guidance failure light illuminated and the LV guidance switch thrown, the computer can send attitude-error signals to the Saturn autopilot. It is significant that with this procedure, takeover cannot be effected unless the IU initiates the action. Simulation testing has shown that this simple backup scheme can place the spacecraft safely into earth orbit with errors in apogee and perigee of no greater than 10 to 15 nmi.

D.5.1 Generation of Guidance Commands

For first-stage flight, the IU commands an open-loop attitude profile that pitches the vehicle about 60 deg from the vertical in less than 3 minutes. To provide backup here, the AGC calculates a sixth-order polynomial fit to the desired pitch profile. The attitude-error signals are then the difference between these pitch commands and the attitude feedback from the IMU. During takeover, the delay between the failure detection and the first AGC command may be several seconds; however, even in the region of maximum dynamic pressure, transients in attitude, angle-of-attack, and engine angle do not produce excessive structural loads. In addition to the pitch maneuver, the IU normally commands a yaw and a roll maneuver. The yaw maneuver provides additional clearance from the launch tower, and the roll maneuver rotates the vehicle from the launch-pad azimuth to the down-range azimuth, an angular change of about 18 deg. In the backup mode the AGC neglects the yaw maneuver, but it does command a constant roll rate to achieve the desired down-range azimuth.

For second- and third-stage flight, the IU provides a closed-loop guidance scheme to achieve the desired orbit parameters. However, computer-storage limitations preclude such a scheme for AGC backup; thus the task of guiding the vehicle into orbit was given to the crew. This requires that the pilot compare DSKY displays of altitude, altitude rate, and velocity against nominal values tabulated on a card. Using this information, as well as the attitude display, the pilot can then use the Rotational Hand Controller in a rate-command mode to fly the desired trajectory. The manual mode is enabled by a keyboard entry.

D.6 LM Autopilot

D.6.1 Integrated Design

The Lunar Module DAP provides attitude control of the LM spacecraft during both coasting and powered flight. The autopilot is designed to control three spacecraft configurations: LM descent, LM ascent, and CSM-docked. The modules comprising these configurations are shown in Fig. D.6-1.

Torques for attitude control may be generated by the Reaction Control System and by the Descent Propulsion System. The LM Reaction Control System employs 16 jets mounted in clusters of four on outriggers equally spaced around the LM ascent stage. Each jet has a thrust of 100 lb. The Descent Propulsion System (DPS) has a single engine throttleable from a maximum thrust of 10,000 lb down to 12 percent of the maximum thrust. This engine is mounted in a gimbal system with actuators; thus the angle of the thrust vector relative to the spacecraft center of mass can be controlled. The actuators can change the engine angle at the constant rate of 0.2 deg/sec. The Ascent Propulsion System has a single 3500-lb engine, mounted rigidly to the ascent stage. Since the thrust vector of this engine cannot be rotated to pass through the spacecraft center of mass, attitude control during powered flight in the ascent configuration must be maintained by use of the RCS jets.

The LM DAP is an integral part of the LM Primary Guidance, Navigation and Control System. Inputs to and outputs from the LM guidance computer which are associated with the control function (the autopilot) are shown in Fig. D.6-2. The autopilot directly commands the firing of each thruster of the RCS, and the

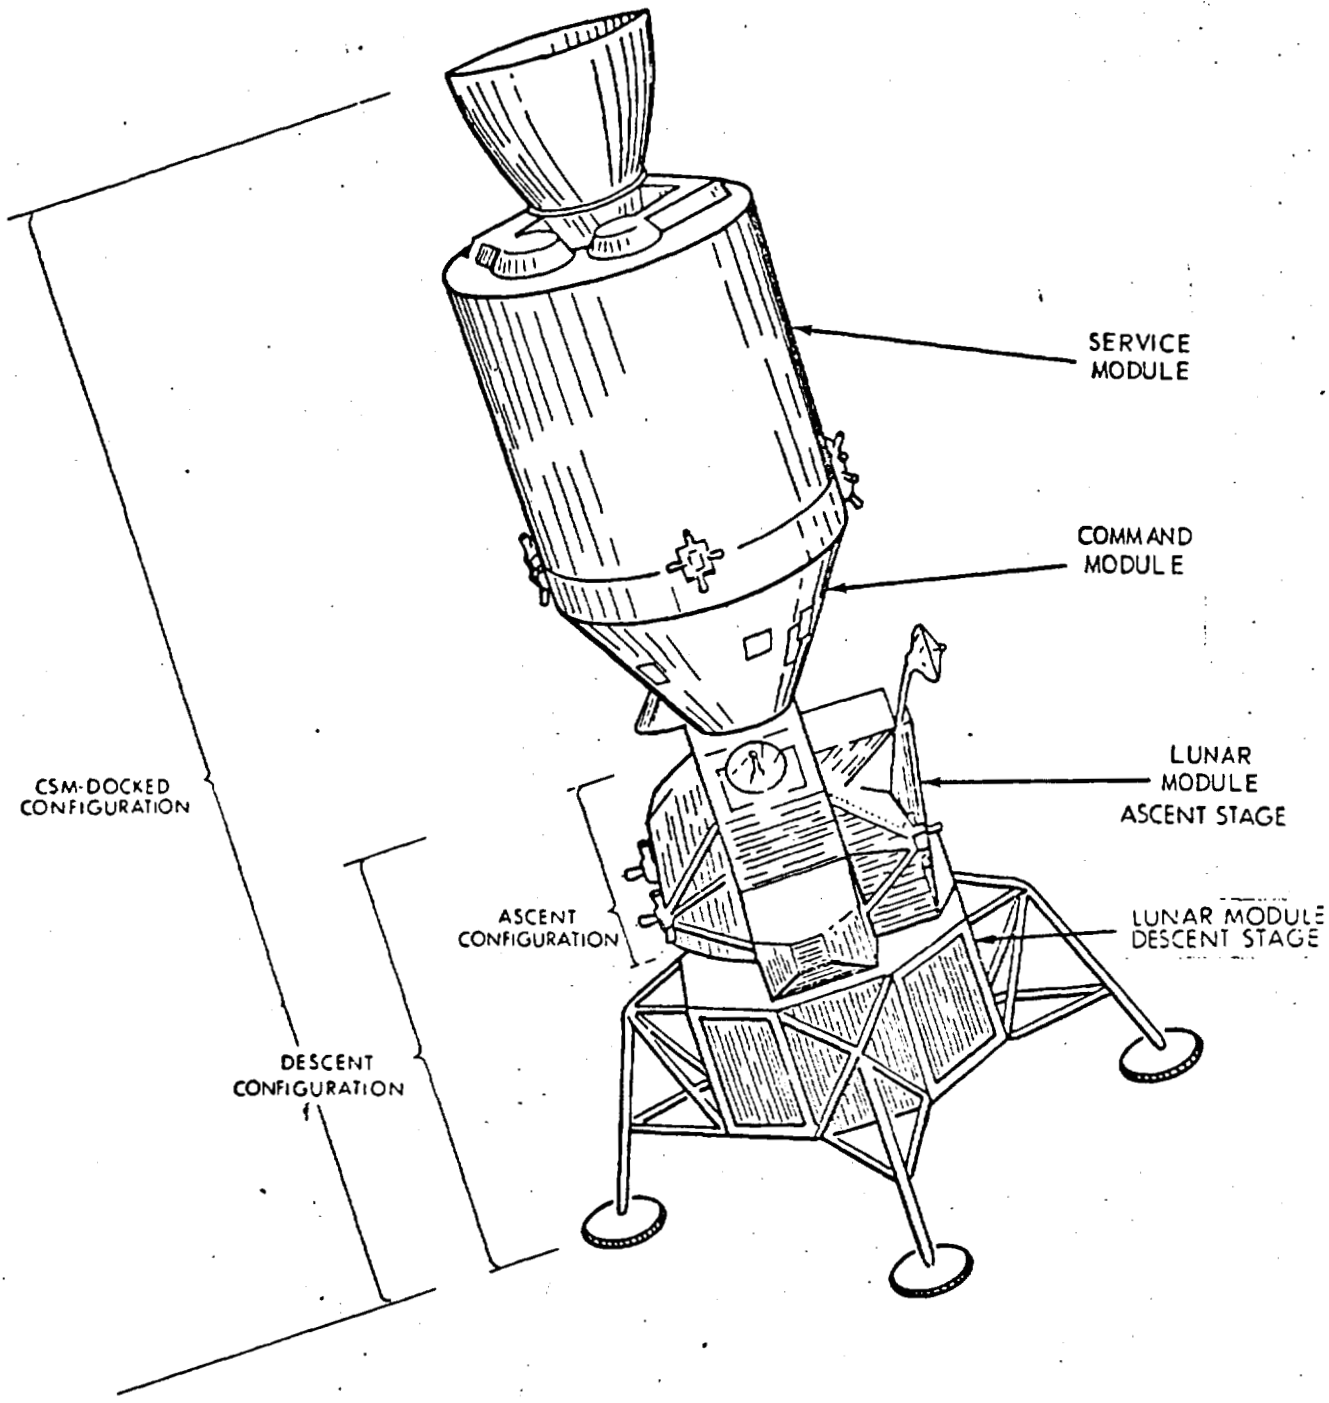

Figure D. 6-1 Spacecraft Configurations Controlled by the Lunar Module Autopilot

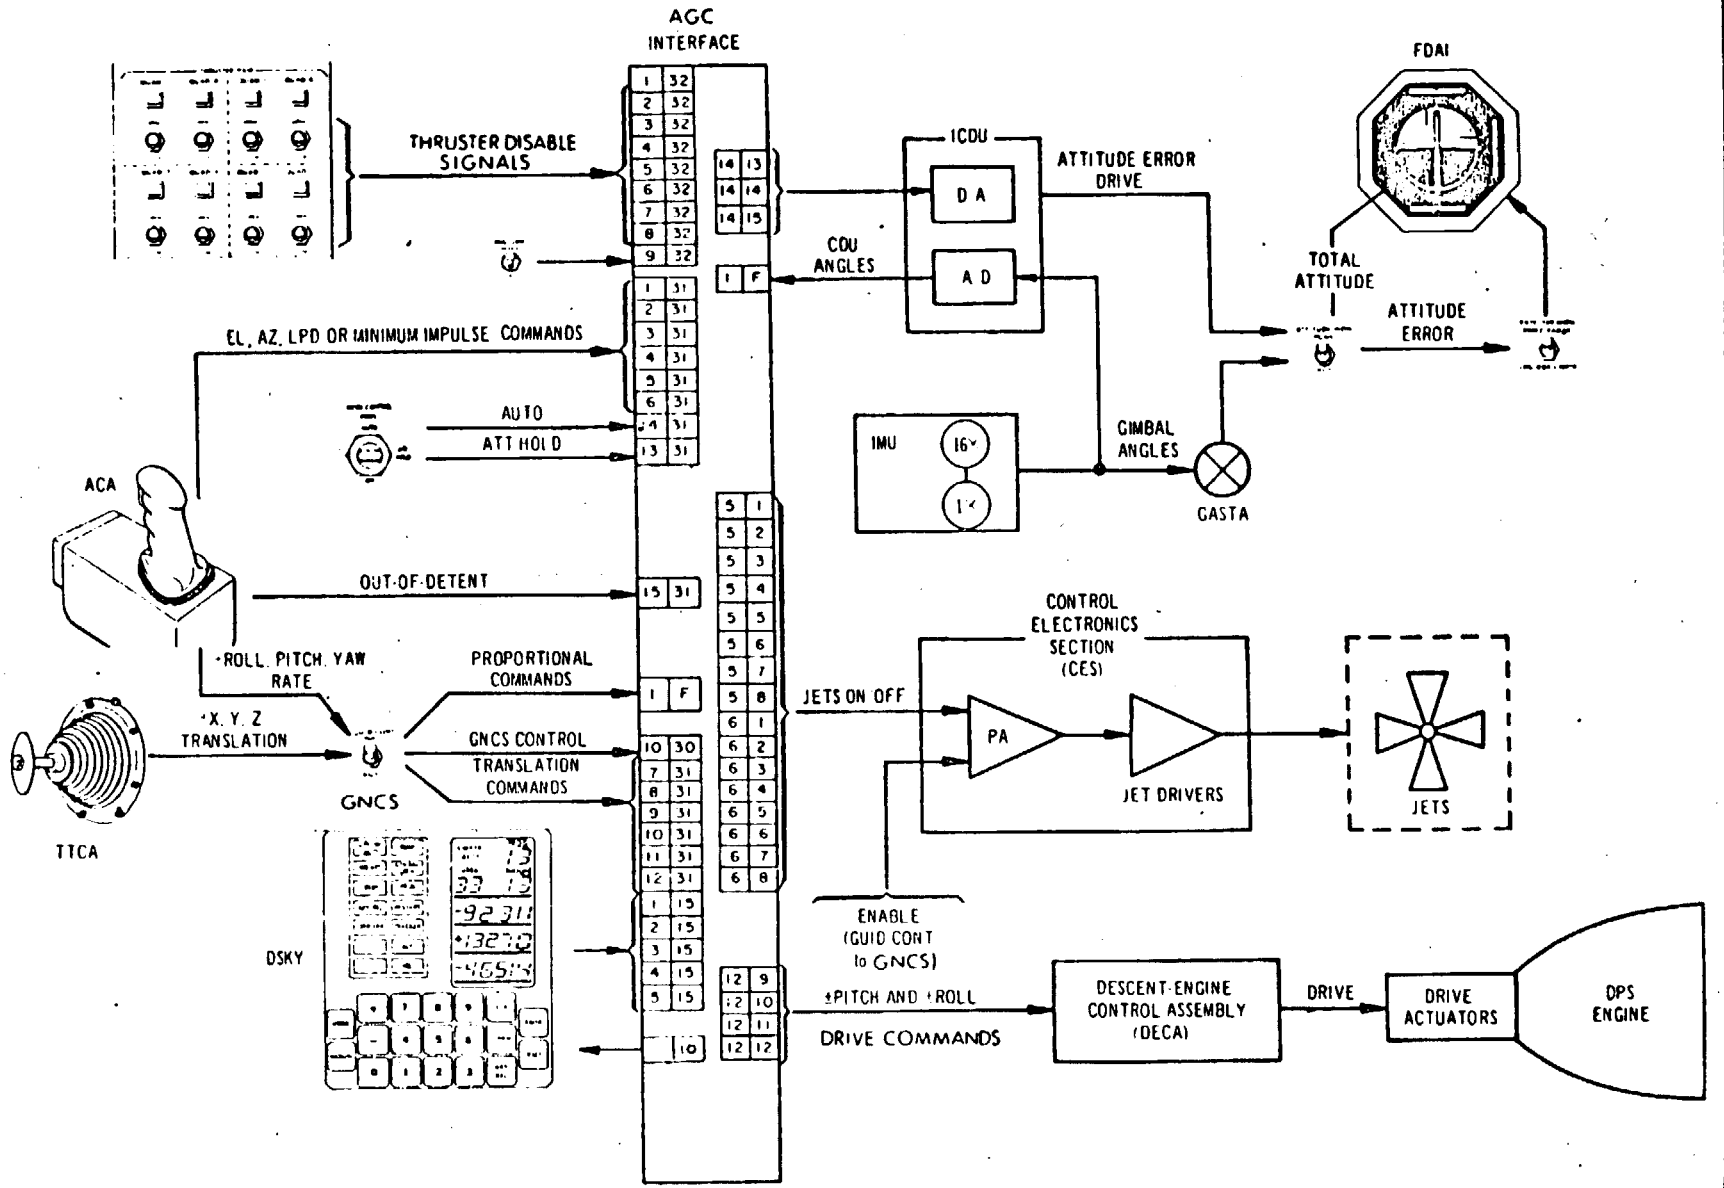

Figure D.6-2 LM DAP/AGC Interfaces

movement of the descent-engine Gimbal Trim System (GTS). In addition, computed attitude errors are displayed on the FDAI. Among the inputs to the LM DAP are the following:

Measurements of the spacecraft attitude with respect to inertial space, as indicated by the three gimbal angles of the gyro-stabilized platform of the IMU.

Hand-controller signals for providing both manual rotational-control commands and manual translational-control commands.

Mode-switch discretes for selecting the autopilot control modes.

Eight thruster-pair disable-switch discretes for modifying the RCS selection logic and optimizing the autopilot performance in the presence of failed jets.

Keyboard inputs for specifying autopilot control parameters, such as angular deadbands, maneuver rates, hand-controller functions, and spacecraft-mass properties.

Internal steering commands for providing automatic attitude control in both coasting and powered flight.

An internal discrete for providing automatic ullage (+X translation) prior to main-engine ignition.

Internal discretes for switching the autopilot control modes and configurations.

The LM autopilot provides various control modes for coasting and powered flight. For coasting flight, the following modes are available (as discussed in Sections D.6.2 and D.6.3):

Attitude-hold mode, in which the inertial attitude is held constant.

Automatic attitude-maneuver mode, in which the vehicle rotates at a constant angular velocity from some initial attitude to some desired final attitude.

Rate-command/attitude-hold mode, in which the vehicle rotation rate is brought to the desired rotation rate indicated by the Rotational Hand Controller.

Minimum-impulse mode, in which single small firings of the RCS jets are commanded in response to each deflection of the Rotational Hand Controller.

These modes are available during powered flight of the Lunar Module (as discussed in Sections D.6.4, D.6.5 and D.6.6):

> Automatic steering, in which the vehicle follows attitude commands provided by the guidance equations. This is the mode which is used during most of the powered descent to the lunar surface.
>
> Rate command/attitude hold, which has the same characteristics as in coasting flight. This is the mode which is selected by the LM pilot during the final phase of the lunar landing to guide the vehicle manually over what could be a rocky terrain.
>
> X-axis override mode, a combined automatic and manual mode in which the LM pilot controls the vehicle attitude about the X-axis (the thrust direction) by manual rate commands, while the direction of the thrust vector is maintained automatically in response to the guidance commands.

D.6.1.1 Design Approach and Structure of the Autopilot

Certain performance requirements guided the design of the LM autopilot. Attitude control must be maintained with a minimum expenditure of RCS propellant. Concomitantly, the number of RCS firings must be minimized to achieve greater thruster reliability. Attitude control must be maintained even in the presence of a disabled single jet, quad or RCS system (8 jets). The DAP must be stable in the presence of bending modes in the CSM/LM docked configuration, slosh in all configurations, and transients due to ignition, abort stage, transfer from the backup Abort Guidance System and switching of DAP modes. Finally, attitude control must be essentially unaffected by off-nominal vehicle, thruster and sensor characteristics.

In its most general form, the requirements for the LM DAP posed a formidable multi-input, multi-output problem: control logic had to be synthesized to relate the measurement of the three-axis spacecraft attitude with the firing of 16 RCS thrusters and the gimballing of the descent engine about two axes. The design approach employed for the LM autopilot was to separate the total synthesis problem into a set of smaller design problems. Accordingly, four major subsections of the LM DAP were defined, each of which could be approached somewhat independently:

1. Certain DAP parameters are functions of vehicle mass and can be expected to change slowly; others need to be computed infrequently. These

calculations are done in a DAP subsection called 1/ACCS (so named because the reciprocals of accelerations are computed), which is executed every two seconds during powered flight and after major transients (such as LM-configuration changes). Its major outputs are the estimates of RCS jet-control authority (angular-rate change induced by a jet firing), Gimbal Trim System control authority (angular-acceleration changes induced by a trim gimbal drive), and related quantities. GTS control can be executed under direction of 1/ACCS when the control mode requires only infrequent reevaluation of GTS drive status.

2. The state* estimator estimates the spacecraft's angular-velocity and angular-acceleration vectors from sampled measurements of the vehicle attitude. The state estimator was designed with the recursive structure of a Kalman filter. However, nonlinear logic was incorporated to reject small-amplitude high-frequency disturbances from vibration and the CDUs; these are recognized and incorporated into the estimates only when they exceed certain threshold magnitudes. Filter gains were chosen to make the rate estimate respond rapidly to changes, and cause the offset angular-acceleration estimate to respond much more slowly. In this way, rapid maneuvers can be performed accurately, satisfying the astronaut's need for quick, flexible, and reliable response during manual control—particularly for emergency maneuvers. On the other hand, the slow acceleration estimate is largely insensitive to slosh oscillations, responding with a greatly damped oscillation in the presence of slosh. In the case of CSM/LM docked powered flight with the LM active, the rate- and acceleration-estimator filter gains are reduced even further to additionally buffer the estimates from the disturbance induced by bending oscillations. This state estimator derives the angular velocity and angular acceleration of the LM, based only on measurement of spacecraft attitude and assumed control response, thus demonstrating that rate gyros are not required sensors for this autopilot.

3. The RCS control laws fire the RCS jets in response to the vehicle-attitude state, the attitude commands, and the translation commands. The RCS control

*The word "state" used in this section refers to the spacecraft's attitude, attitude rate (and sometimes angular acceleration).

laws in the LM-alone configuration employ parabolic switch curves in their phase-plane logic. The critical parameters in the RCS control laws are adapted in response to the varying moment of inertia of the spacecraft and the bias angular acceleration due to the thrusting main engine. The control-law design permits rapid response to commands with a minimum of jet firings. In coasting flight, steady-state attitude control is maintained with a minimum-impulse limit cycle. In ascent powered flight, a larger-pulse low-frequency limit cycle is employed to hold attitude against the bias angular acceleration. The RCS control laws maintain satisfactory attitude control of even the lightest ascent configuration* even though the control action is reevaluated at most 10 times per sec. Satisfactory control is possible because the RCS control laws compute the exact firing time required to achieve a desired rate change.

4. The trim-gimbal control laws drive the orientation of the descent engine about its two axes in response to the vehicle-attitude state and attitude commands. A third-order minimum-time control law is used to control the vehicle attitude by means of the thrusting descent engine. This permits attitude control often without the assistance of the RCS jets.

Within the RCS control laws and the trim-gimbal control laws, further design simplifications were made by separating these control laws into logically distinct control channels. The choice of these control channels followed from the natural control axes of the LM, shown in Fig. D.6-3.

The descent engine may be gimballed under computer control about the pitch (Q) axis and the roll (R) axis. Therefore, the descent-engine trim-gimbal control laws have been separated into two channels (Q and R). The computation of the proper trim-gimbal drive for each channel is based on independent single-plane control laws.

The RCS jets mounted on the LM ascent stage are skewed 45 deg away from the spacecraft's coordinate frame (in the QR plane) to avoid jet impingement on the pads. This rotational jet-coordinate frame is designated the P, U and V axes. The

*In the lightest ascent configuration, single-jet firing for 0.1 sec (one DAP pass) could produce an attitude-rate change of 1.5 deg/sec.

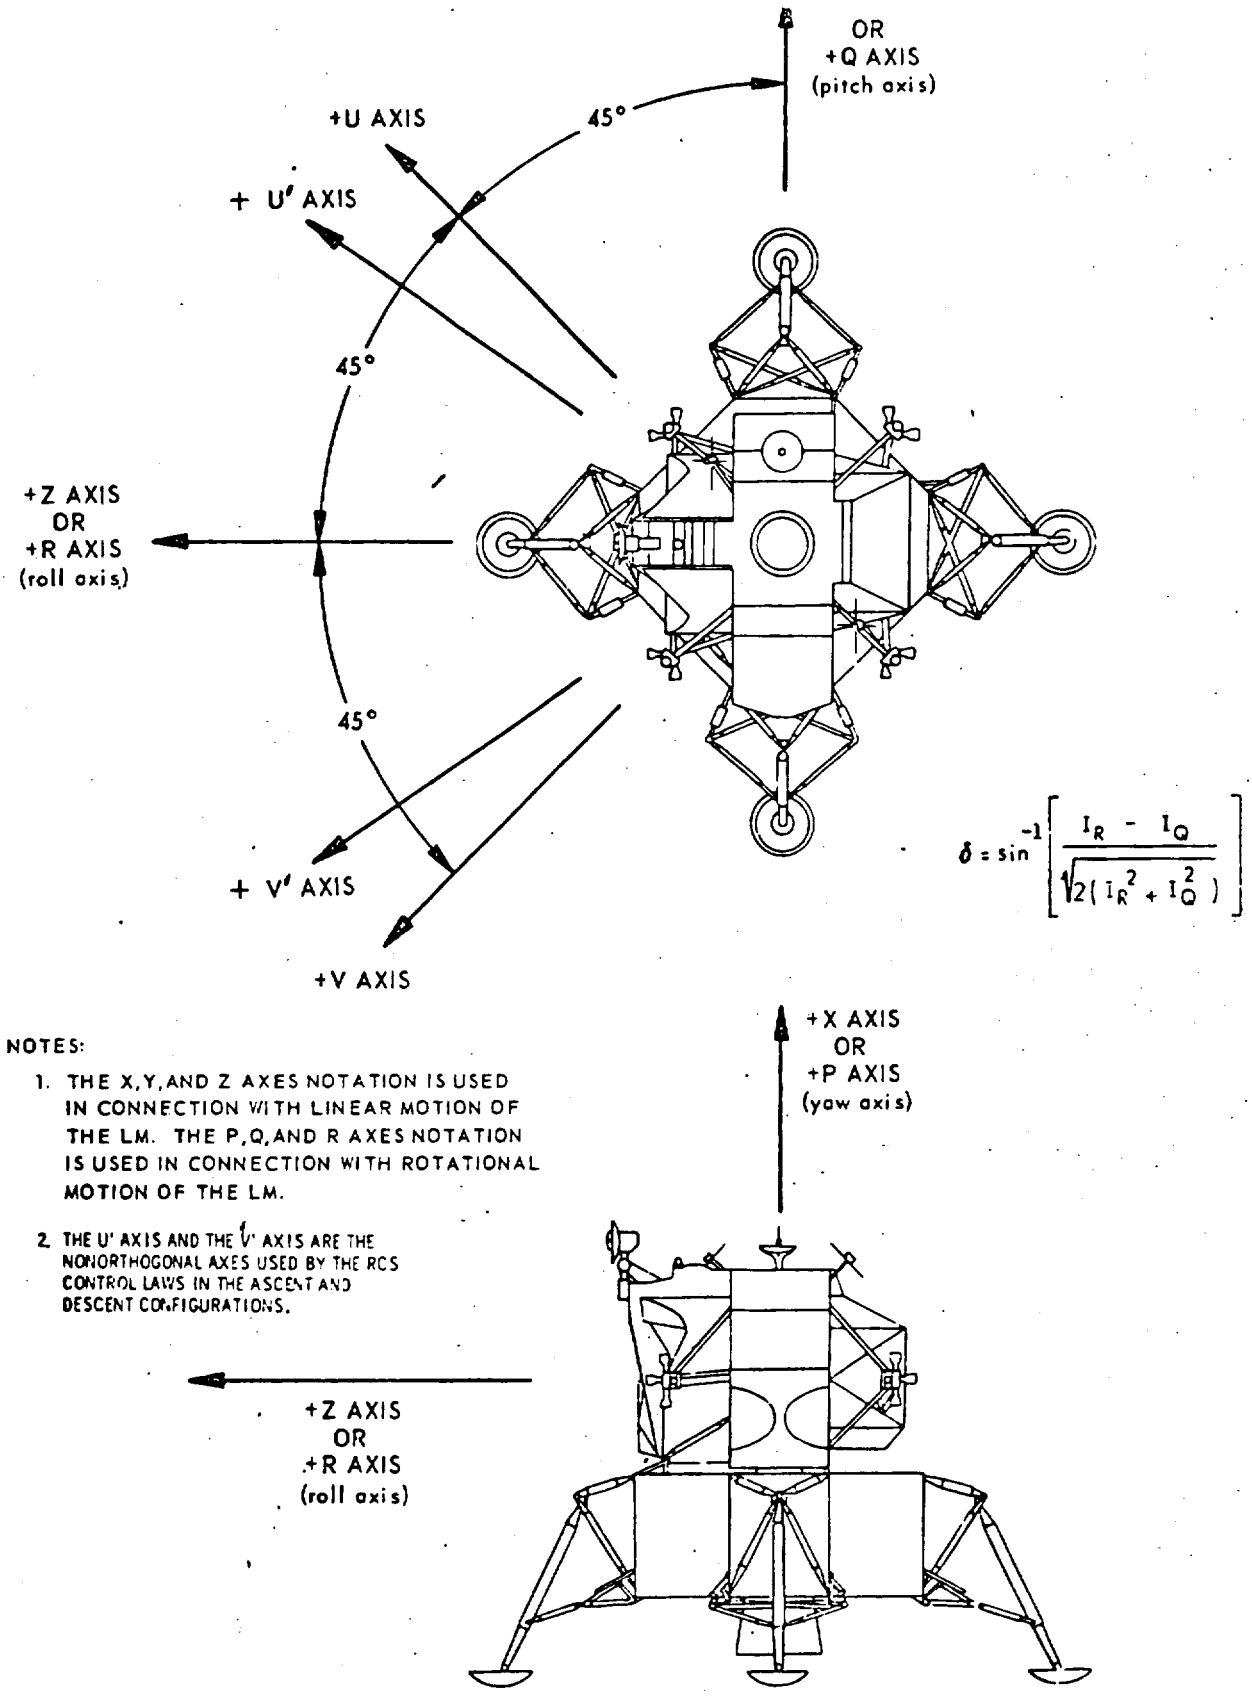

Figure D. 6-3 The Control Axes of the LM

locations and orientations of the RCS jets are such that if the spacecraft center of gravity lies near the geometric center of the 16 RCS jets, then:

a. The eight jets that thrust only in the Y or Z directions produce torques about the P axis only. Accordingly, these jets are termed the P jets.

b. Four of the jets that thrust in the ±X direction produce torques about the U axis only. Accordingly, these jets are termed the U jets.

c. The other four jets that thrust in the ±X direction produce torques about the V axis only. Accordingly, these jets are termed the V jets.

Due to the existence of significant cross-inertia between the U and V axes, cross-coupled acceleration between the U and V axes is introduced whenever a U jet or a V jet is fired. To decouple the RCS control channels and thereby reduce the number of RCS firings and hence propellant consumption, a nonorthogonal set of control axes called U' and V' was introduced. The U'- and V'-axis directions are determined as follows: first, cross-coupling between the P axis and any axis in the Q,R plane is assumed negligible and hence is ignored—therefore, the U' and V' axes are constructed to lie in the Q,R plane and correct for cross-coupled acceleration in that plane only; second, acceleration vectors are produced by applying a torque around U and V, as illustrated in Fig. D.6-4; third, perpendiculars are drawn to each of these acceleration vectors to produce the new U' and V' axes. It can be seen that if an RCS torque is applied about U, there will be a component of acceleration along U' and no component of acceleration along V'. Similarly, an RCS torque applied about V will produce no component of acceleration along U'. Therefore, if a U-axis torque is commanded to achieve U'-axis control and a V-axis torque is commanded to achieve V'-axis control, no cross-coupled acceleration will result. The angle δ by which the U' and V' axes are skewed away from the U and V axes is computed from the Q, R moments of inertia.

D.6.2 Manual Modes of the LM DAP

During certain critical phases of the Apollo mission, LM attitude is manually controlled by the LM crew. To achieve a precisely defined attitude, such as that required for initial thrust-vector positioning, an automatic maneuver is usually more

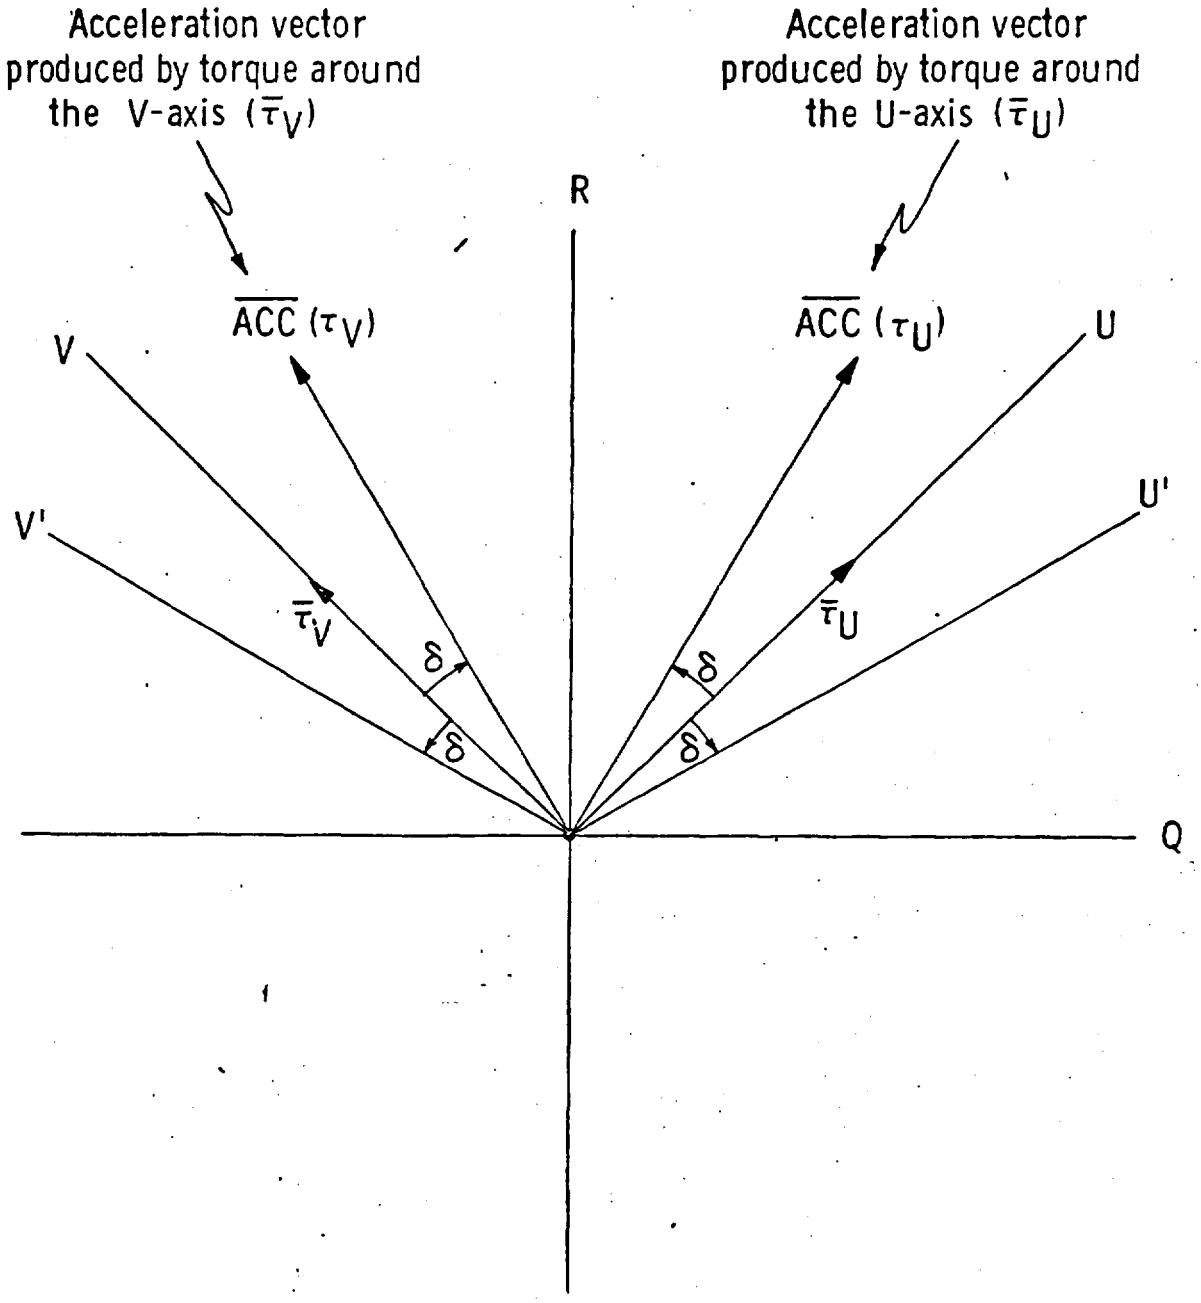

Figure D. 6-4 Nonorthogonal LM U', V' Axis System

efficient. For less precisely determined tasks, however, manual control is better. For instance, station keeping and gimbal-lock avoidance can best be performed in a manual mode. In such instances, the pilot's ability to perceive the state of his vehicle and to select alternatives to the nominal control is clearly an asset.

In the GN&CS, two manual modes are implemented through the use of a Rotational Hand Controller. Two such controllers are available (although they cannot be used simultaneously), one for each of the astronauts onboard the LM. The Minimum-Impulse mode provides a single, 14-msec thruster pulse* each time the controller is moved out of detent; each pulse results in an angular-velocity increment whose magnitude is a function of vehicle inertia.

The second—and by far, the more significant—manual mode is Rate Command/Attitude Hold, which incorporates a number of features that enhance the rapidity and precision of control response. This closed-loop mode provides angular rates which are a function of the degree of RHC deflection. The remainder of this section will deal exclusively with the evolution of the Rate Command/Attitude Hold mode.

D.6.2.1 Rate-Command/Attitude-Hold Mode

Development of the Rate-Command/Attitude-Hold mode has been both lengthy and complex. The LM AGC program for Apollo 9 (SUNDANCE) contained manual rate-command logic which was a digital realization of earlier reaction-control rate-command systems. This logic provided rate command within the resolution of a rate-error deadband when the RHC was held out of detent. Automatic attitude hold was maintained when the controller was in detent.

In the manual rate-command logic of SUNDANCE, RHC output was scaled to a maximum commanded rate of 20 deg/sec in "Normal" and 4 deg/sec in "Fine". Normal would be used for a lunar landing, in which high maneuver rates might be necessary; Fine would be used for all other operations, where a premium would be placed on precision and fuel-saving low rates.

---

*The RCS jets require a minimum pulse width of about 14 msec to ensure a proper mixture ratio in the combustion chamber.

The SUNDANCE program provided acceptable manual control for the LM-alone configuration, as evidenced both by simulated and actual earth-orbital flight. But improvements in the manual rate command were deemed necessary to meet the following objectives for an actual lunar landing (program LUMINARY):

a. to reduce drift about uncommanded axes
b. to provide more precise rate control
c. to assure positive return to attitude-hold mode after rate commands
d. to make manual rate command available for coasting flight in the CSM-docked configuration
e. to reduce the on-time of the +X firing thrusters during the lunar landing.

The following sections discuss each of these objectives in terms of the LM DAP manual rate-command mode.

D.6.2.1.1 Reduction of Drift

Although SUNDANCE's rate-command-with-deadband was an acceptable mode from a handling-qualities standpoint, it was open-loop for small, secular errors. Because the thruster switch curves were independent of attitude error, it was possible for the spacecraft to have an uncorrected drift rate just barely within the deadband.

Two factors complicate the drift problem. If the controller were out-of-detent about any axis, all three axes used the manual logic. Consequently, the spacecraft could drift about uncommanded axes (up to nearly 2 deg/sec with normal scaling). Also, a bias acceleration could cause the phase point to chatter along the switch curve. Sampling and state-estimation delays compounded this drift. To limit drift, attitude errors were incorporated in LUMINARY's control computations.

D.6.2.1.2 Precise Rate Control

The rate deadband determines the resolution of rate control. Targeting jet-on time for zero rate error and using inertial and bias acceleration estimates often result in rate-step response with errors initially smaller than the deadband. Such precision cannot be guaranteed, however, because of error in knowledge of vehicle rate or of control authority. Once the rate error is within the deadband at a sampling

instant, firing ceases. A heavy configuration requiring more than a 0.1-sec firing for a rate change equal to the deadband will never have rate error nulled entirely, and firing may stop just within the deadband for lighter vehicles (depending on initial rate error). Furthermore, uncertainty in command response extends to twice the error-angle deadband width, measured from zero. If the rate error is just barely within the negative limit and a positive change is requested, the rate error must traverse the entire deadzone before a firing occurs. To obtain precise rate control, LUMINARY's manual rate command applies integral compensation.

Tightening the rate loop alone proved insufficient to improve the pilot's estimation of handling qualities in the lunar-landing task. (It was still virtually impossible to achieve small attitude changes during simulations.) The difficulty lay in the small amount of deflection required to obtain RHC output and in the sensitivity of the controller. (During simulations, pilots "felt" or "heard" the detent switch click, yet, depending upon the particular RHC, up to 1.5 deg additional deflection was necessary to obtain any output signal.) With small, smooth hand motions, small attitude changes clearly are difficult to command, since the pilot cannot predict when the voltage buildup begins. For LUMINARY, therefore, RHC sensitivity was modified to resolve these difficulties.

D.6.2.1.3 Return to Attitude-Hold Mode

Once the RHC is returned to detent, control should be passed from manual rate command to attitude hold positively and with a minimum transient. The latter requirement is met if rates are damped before the switch to automatic attitude hold; if the rate error were large, the attitude-hold phase-plane logic could command oscillatory response in seeking to null the attitude error in minimum time, consuming RCS propellant unnecessarily. Small rate error is not imperative for mode change, however. Return to automatic control should be assured whenever the RHC is returned to detent, even if any or all components of angular rate fail to damp within a short time. Once damping about an axis has reduced the rate error to a small value, that axis should be considered to have passed the damping test. Chattering about more than one axis can delay—and possibly prevent—return to attitude hold. In the SUNDANCE logic, the requirement for the return, after the RHC is returned to detent, was that all rate errors be less than the rate deadband simultaneously. Phase-point chattering out-of-phase will fail this test. If this occurs as a result of

an undetected jet failure or a mass uncertainty, the only ways to return to attitude hold are to inform the computer of the failure or to momentarily switch the DAP mode. In LUMINARY, rates about all axes are normally damped before the switch to automatic mode; however, the return to automatic mode is forced if damping has not been completed by the end of a brief interval.

D.6.2.1.4 Availability for CSM-Docked Configuration

CSM-docked rate command was not a requirement of SUNDANCE. To accommodate the use of this mode for coasting flight in Apollo 10, the minor required changes were made in LUMINARY.

D.6.2.1.5 Reduction of +X-Thruster On-Time

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

D.6.2.1.6 RHC Scaling

The sensitivity of commanded rate to RHC deflection is the most important manual-control parameter, once rotational-control acceleration ("control power") is fixed. A range of controller sensitivities that provides stable human-pilot loop

*The later addition of jet-plane deflectors (see Section D.6.5) somewhat alleviated this problem, but a heating constraint still exists for the deflectors.

closures can often be defined; as a consequence of the pilot's adaptive ability, however, optimization of the sensitivity within that range is a subjective process. Choice of scaling can be affected by control power, vehicle and control-system dynamics, external disturbances, the control task, and the individual pilot's ability to perceive and react.

It has been found that reduced controller sensitivity has a striking effect on the consumption of RCS propellant: there is a monotonic reduction with decreasing sensitivity. A reduction of maximum commanded rate (MCR) from SUNDANCE's 20 deg/sec to 14 deg/sec produced improved handling qualities, according to several test pilots. (The emphasis in these tests were placed on accuracy in flying to a designated site and on reducing the +X-jet on-time.) Moreover, handling qualities continued to improve as the MCR was reduced to a final value of 8 deg/sec. Reduced MCR improved jet-on time and RCS fuel consumption—and also landing-point accuracy.

In spite of the improvements resulting from reduced controller sensitivity, one conflict remained: reduced sensitivity made small rates and small angle changes easier to obtain, but there was concern that the MCR was insufficient for emergency conditions. A 20-deg/sec MCR was deemed mandatory by the astronauts. The solution adopted is nonlinear scaling of the RHC output.

D.6.3 Coasting Flight

The LM DAP has four coasting-flight modes which may be utilized in the LM-ascent, LM-descent or CSM-docked configurations. These modes are Rate Command/Attitude Hold, Minimum-Impulse Command, Attitude Hold and Automatic Maneuvering. Each mode controls vehicle attitude with the 16 RCS jets located on the LM ascent stage. Since Rate Command/Attitude Hold and Minimum-Impulse Command are discussed in Section D.6.2, this section discusses only the coasting-flight performance of the LM DAP in the Attitude-Hold and Automatic-Maneuvering modes.

D.6.3.1 Attitude-Hold Mode

The Attitude-Hold mode stabilizes the spacecraft about each of the Inertial Measurement Unit's reference axes to hold the spacecraft to within a specified

deadband. The principal design objectives included minimization of RCS propellant consumption, minimization of the number of RCS jet firings, acceptable operation in the presence of detected and undetected RCS jet failures, and rapid recovery from large attitude-error and attitude-rate excursions.

D.6.3.1.1 Ascent and Descent Configurations

In the ascent and descent configurations, the number of jet firings and the RCS propellant consumption are both minimized when a minimum-impulse limit cycle is attained about each of the three control axes. (A minimum-impulse limit cycle about a given axis in the absence of disturbing torques is defined as a limit cycle in which a single torque impulse of the smallest available duration, i.e., 14 msec RCS firing, reverses the attitude rate whenever the attitude error drifts out of the deadband.) Although only one jet is fired for each U- or V-axis minimum-impulse torque correction, the displacement of the vehicle center of gravity from the RCS jet plane necessitates the use of two jets fired as a force couple for each P-axis minimum-impulse firing to prevent P-axis firing from disturbing U- and V-axis limit cycles. The RCS control-law phase plane for ascent and descent coasting flight, as illustrated in Fig. D.6-5, was designed for minimum-impulse limit-cycle operation in the steady state. For example, when the state is in the coast zone, Zone 4, with a small positive error rate, $\dot{E}$, the error E increases until Zone 3 is entered. When the state is in Zone 3, the LM DAP commands a 14-msec minimum-impulse RCS firing, which induces a small negative rate and causes the state to reenter Zone 4. The error then decreases until Zone 3 is reentered on the other side of the phase plane. At this point, another minimum-impulse firing induces a small positive error rate and completes the cycle. In coasting-flight minimum-impulse limit cycles, the maximum attitude errors are determined by the deadband size; the maximum attitude-error rates are determined by the control authority of the RCS jets; and the limit-cycle frequency is a function of both the control authority and the deadband size.

The LM DAP was designed to provide acceptable control of the ascent and descent configurations in the presence of detected and undetected jet failures. When jets are disabled due to detected failures, the LM DAP jet-selection logic modifies the selection of jets to exclude those which have been disabled. Single undetected jet-off failures cause rate and attitude undershoot when selected for attitude control,

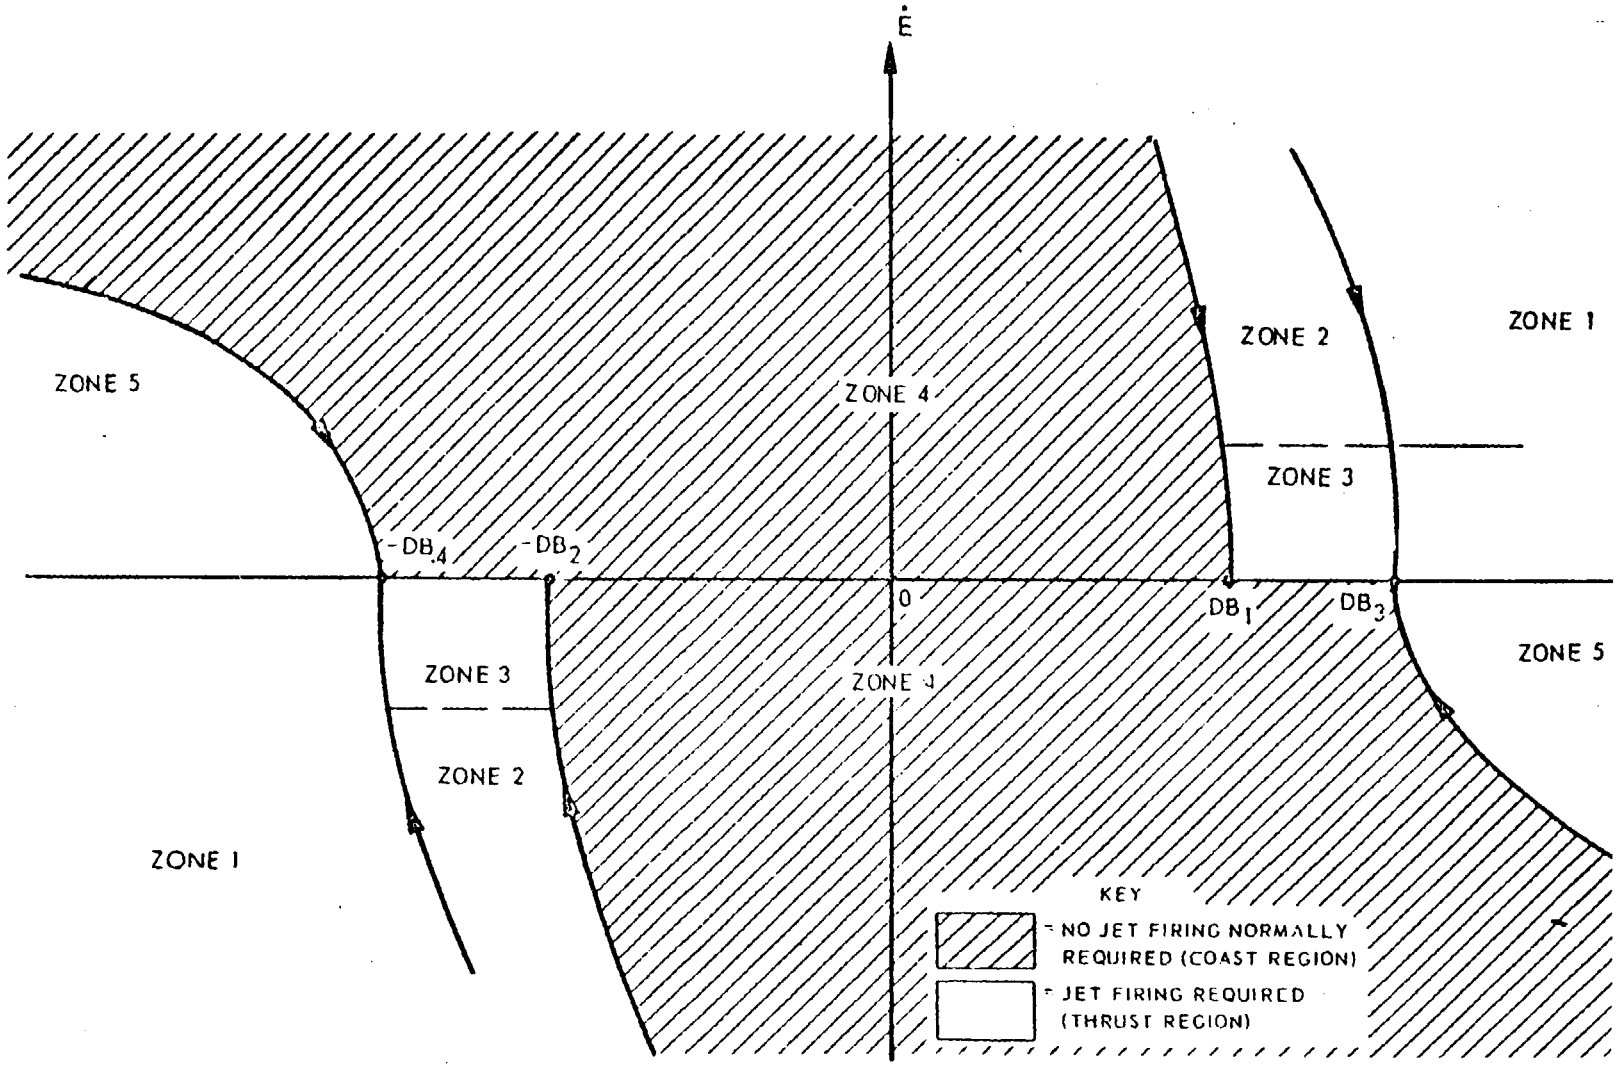
Ė
ZONE 5
ZONE 4
ZONE 2
ZONE 1
ZONE 3
-DB4
-DB2
0
DB1
DB3
ZONE 3
ZONE 4
ZONE 5
ZONE 2
ZONE 1
KEY
= NO JET FIRING NORMALLY REQUIRED (COAST REGION)
= JET FIRING REQUIRED (THRUST REGION)

but result in no fuel penalties. Simulation results indicate that minimum-impulse limit cycles are attained by the LM DAP with any single jet pair disabled and with any single failed-off jet undetected. Undetected jet-on failures degrade the LM DAP performance considerably; however, the degradation lies within the limits NASA deems acceptable.

Rapid recovery from large attitude errors and error rates was provided by the Zone 1 and Zone 5 logic of the RCS control law (FINELAW), illustrated in Fig. D.6-5; and by the special control law (ROUGHLAW) for very large attitude errors and error rates illustrated in Fig. D.6-6. If the state is in Zone 1 or Zone 5 of the FINELAW phase plane, the DAP commands the jets to fire until the state crosses one of the Zone 4 boundary-target parabolas and enters the coast zone. For example, for large negative attitude errors and error rates, the Zone 1 logic would cause positive-torquing jets to fire until the state passed through Zone 5, and crossed the target parabola to enter Zone 4 with a positive error rate. The attitude error would then increase, causing the state to drift across Zone 4 until Zone 2 was entered. The Zone 2 logic would cause the jets to torque negatively until the state had crossed Zone 3 and entered Zone 4, again with a very small negative error rate. The attitude error would then slowly decrease, causing the state to drift across the coast zone into Zone 3. In Zone 3, a single minimum-impulse jet firing would occur, producing a small positive error rate and initiating a minimum-impulse limit cycle.

If the attitude-error magnitude exceeds 11.25 deg or the error rate exceeds 5.625 deg/sec, the ROUGHLAW control-law phase plane in Fig. D.6-6 applies. If the state is in Zone A or Zone D, jets are fired until the rate magnitude is 6.5 deg/sec. No jets are fired by the DAP as the state drifts across Zone C. When Zone B is entered, the jets are fired continuously until the FINELAW region is entered.

D.6.3.1.2 CSM-Docked Configuration

Since the CSM-docked LM DAP was intended for use only as a backup system, the simple control law illustrated in Fig. D.6-7, which differs considerably from the ascent- or descent-configuration phase planes, was employed merely to minimize

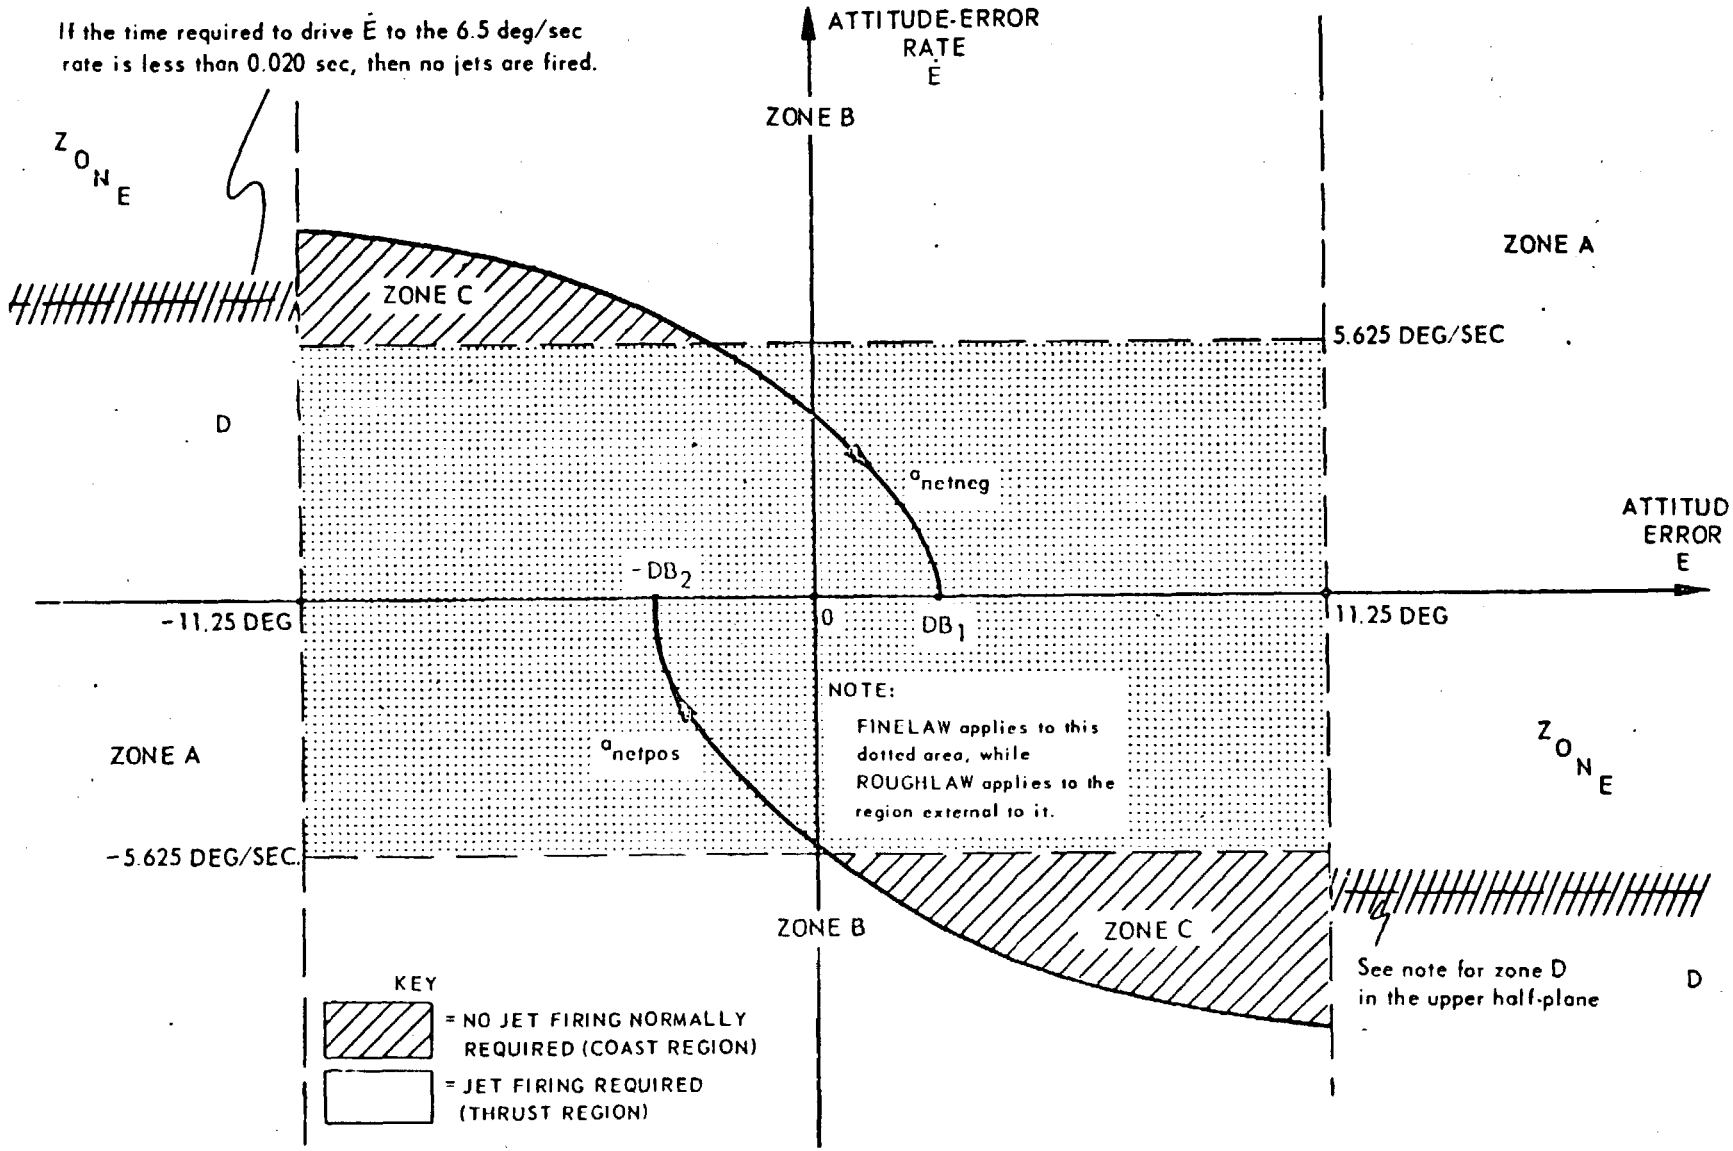

Figure D. 6-6 RCS Control-Law Phase Plane for Large Attitude Error and Error Rates (Ascent and Descent Configurations)

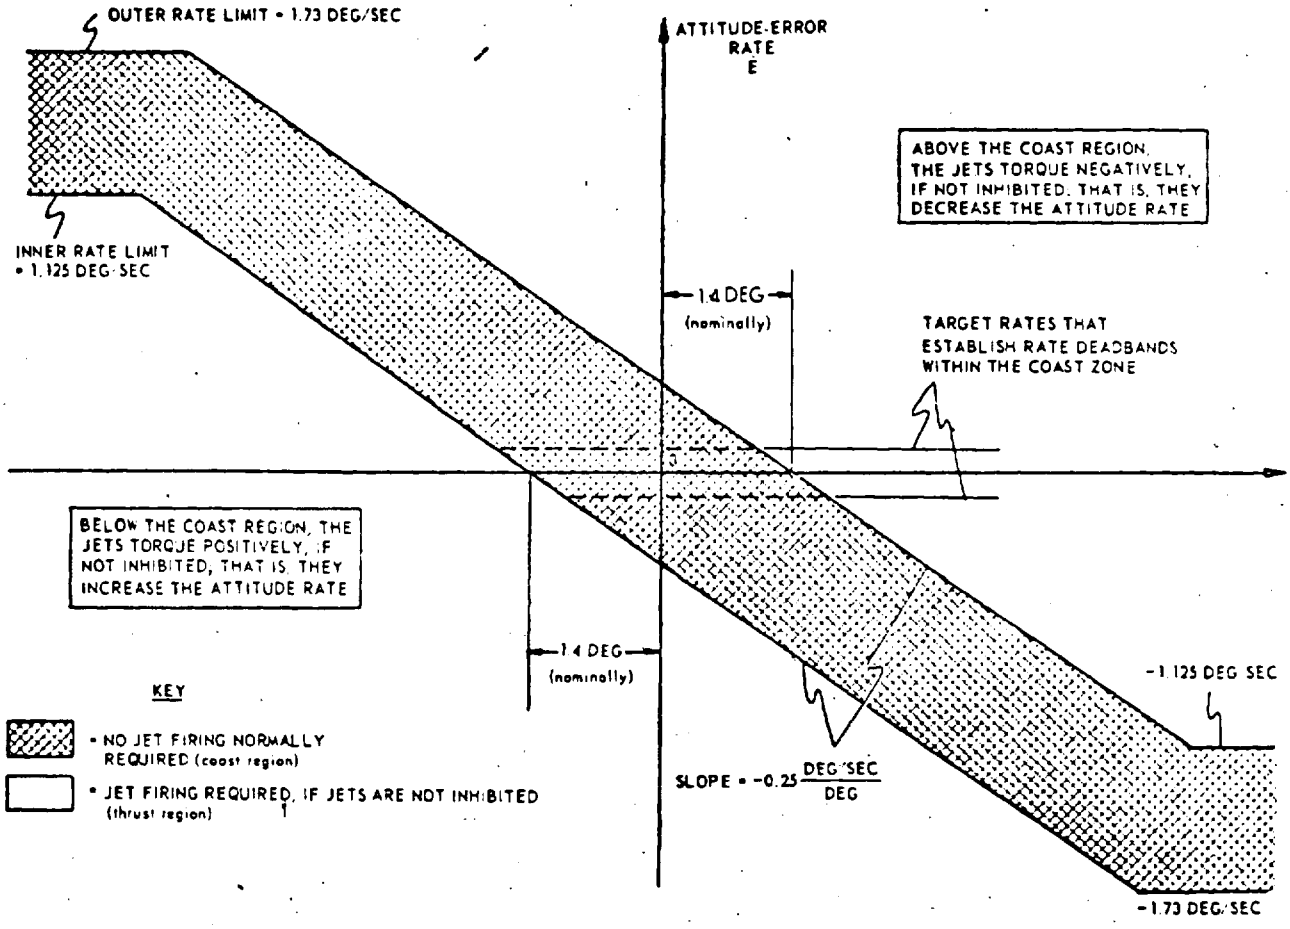

Figure D.6-7 Control-Law Phase Plane (CSM-Docked Configuration)

A NOTE ON SOURCES

By the time this record of MIT's Apollo software efforts entered the germination stage, many of the souls who had participated in those efforts had begun to disperse to other projects—both within and without the Draper Laboratory. Nonetheless, virtually all wanted to ascertain that that part of the history in which they played such important roles was finally, indeed, recorded. Some personally documented their accomplishments; others supplied bits and pieces that eventually interlocked to permit the construction of a unified whole. A considerable amount of the information recorded within these pages could be gleaned from documents that already existed*, but a surprising—indeed, exasperating—amount had never before been documented. For the latter, memories had to be tapped and taped; 57 transcribed interviews, countless conversations and mountains of notes bear testimony to the cooperation and enthusiasm which I encountered along this historical path.

Considerably more information was gathered than could be presented within any single cohesive text. But what I hope has remained is an insight into the team which carried the concept of Apollo software from a hopeful infancy, through an oft turbulent adolescence, to its magnificently successful goal. The Apollo software team was a heterogeneous, sometimes colorful lot, one which demonstrated two basic characteristics: competence and perseverence. The pressures imposed by the schedule sometimes revealed frailities, but more often demonstrated elemental strengths. The epoch which this history records was a significant time in all the participants' professional and personal lives, and, in recalling this period, no one felt dispassionate. Despite the crushing schedules, the fantastic amount of mental and physical exertion which the project came to demand—the goal which was to be reached seemed to energize us all.

Not every member of the Apollo software team contributed to this report, but a great many did, some extensively, some less so. In listing these persons below,

* A compendium of abstracts of all Laboratory reports pertaining to Project Apollo appears as an Appendix to Volume I of this report.

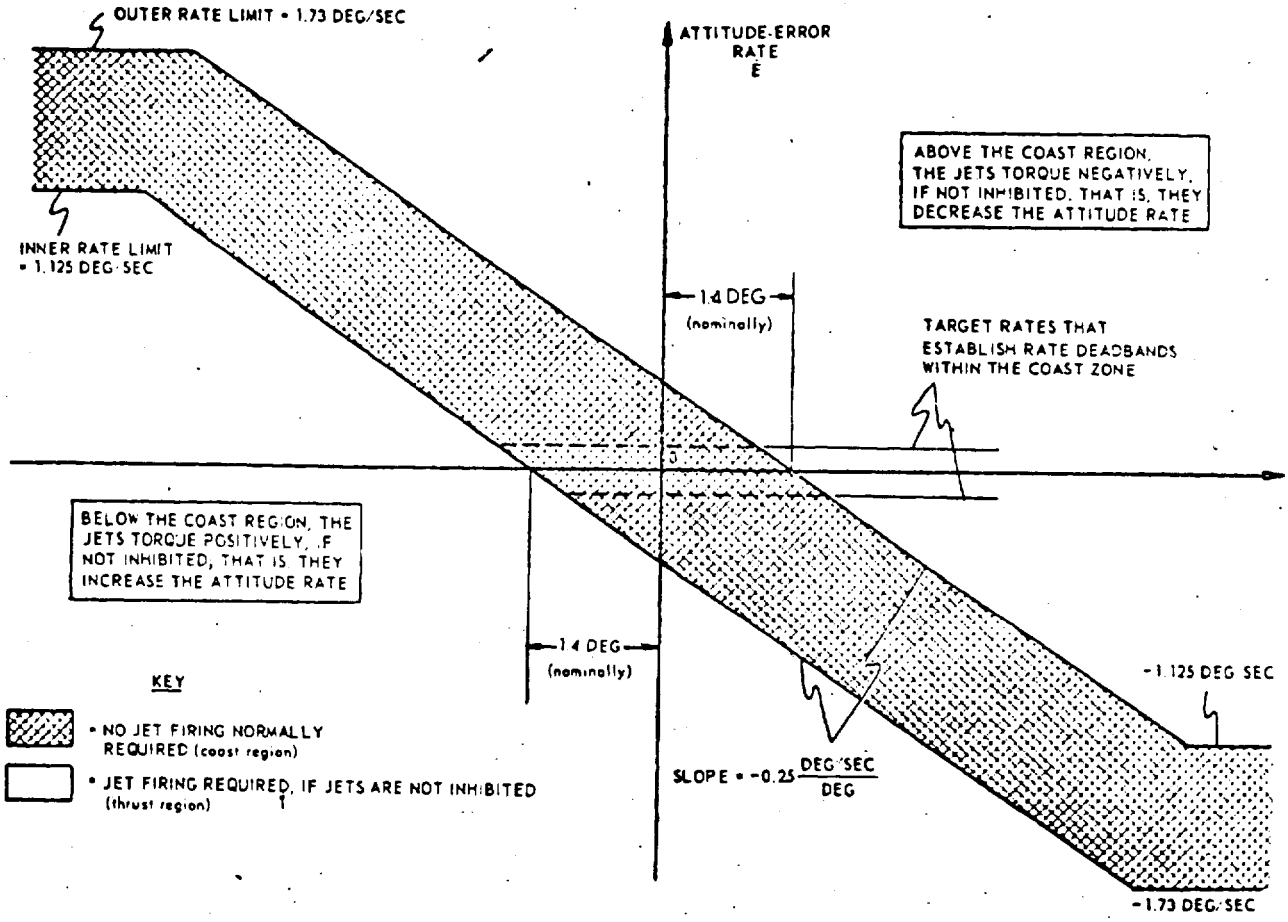

Figure D.6-7 Control-Law Phase Plane (CSM-Docked Configuration)

coding requirements for this configuration*. Each 0.1 sec, the DAP determines the location of the rotational state in the control law and turns jets on or off for a full DAP cycle. If the state is outside the cross-hatched coast zone, the jets are fired until the coast zone is entered. When the state is in the coast zone with the jets still firing, the jets remain on if the trajectory has not crossed the error-rate target (E=0), and are turned off if the trajectory has crossed the error-rate target. If the state is in the coast zone with the jets off, the jets remain off until the state drifts out of the coast zone. Since, in the CSM-docked configuration, the shortest jet-firing time commanded by the LM DAP is 0.1 sec and, in the absence of disabled jets or X-axis translational commands, jets are always fired in pairs about each control axis, minimum-impulse limit cycles cannot be attained. The smallest RCS propellant consumption and the least number of RCS jet firings are obtained when each limit cycle consists of one positive- and one negative-torquing two-jet, 0.1-sec firing as illustrated by the fine-line trajectory in Fig. D.6-8. Due to the large inertia in the CSM-docked configuration, the peak limit-cycle error rates attained in this type of limit-cycle are comparable to the peak error rates obtained in a minimum-impulse limit cycle in the lightest descent configuration.

Acceptable performance of the LM DAP in the CSM-docked configuration with disabled jets is achieved (for most situations) by the same means as in the ascent and descent configurations. That is, the jet-selection logic is automatically modified to select only jets which have not been disabled. Undetected jet failures and, for some mass loadings, disabled -X thrusting jets, however, created challenging problems which were unique to the CSM-docked configuration. In tests of early LM DAP designs, it was discovered that an undetected jet-on failure fixed the vehicle state at one of the coast-zone boundaries of the RCS control-law phase plane, requiring rapid on-off pulsing of the jets to maintain attitude control. Since the jet-pulsing frequency under these circumstances was often close to the natural frequencies of the vehicle bending modes, large bending oscillations could develop. In many of the cases tested, the magnitude of these oscillations became large enough to cause the state estimate to move from one side of the RCS control-law coast zone to the other, resulting in alternate positive and negative torquing of the jets at the bending

*The LM-alone control law could not be employed because the large inertias in the CSM-docked configuration caused scaling problems.

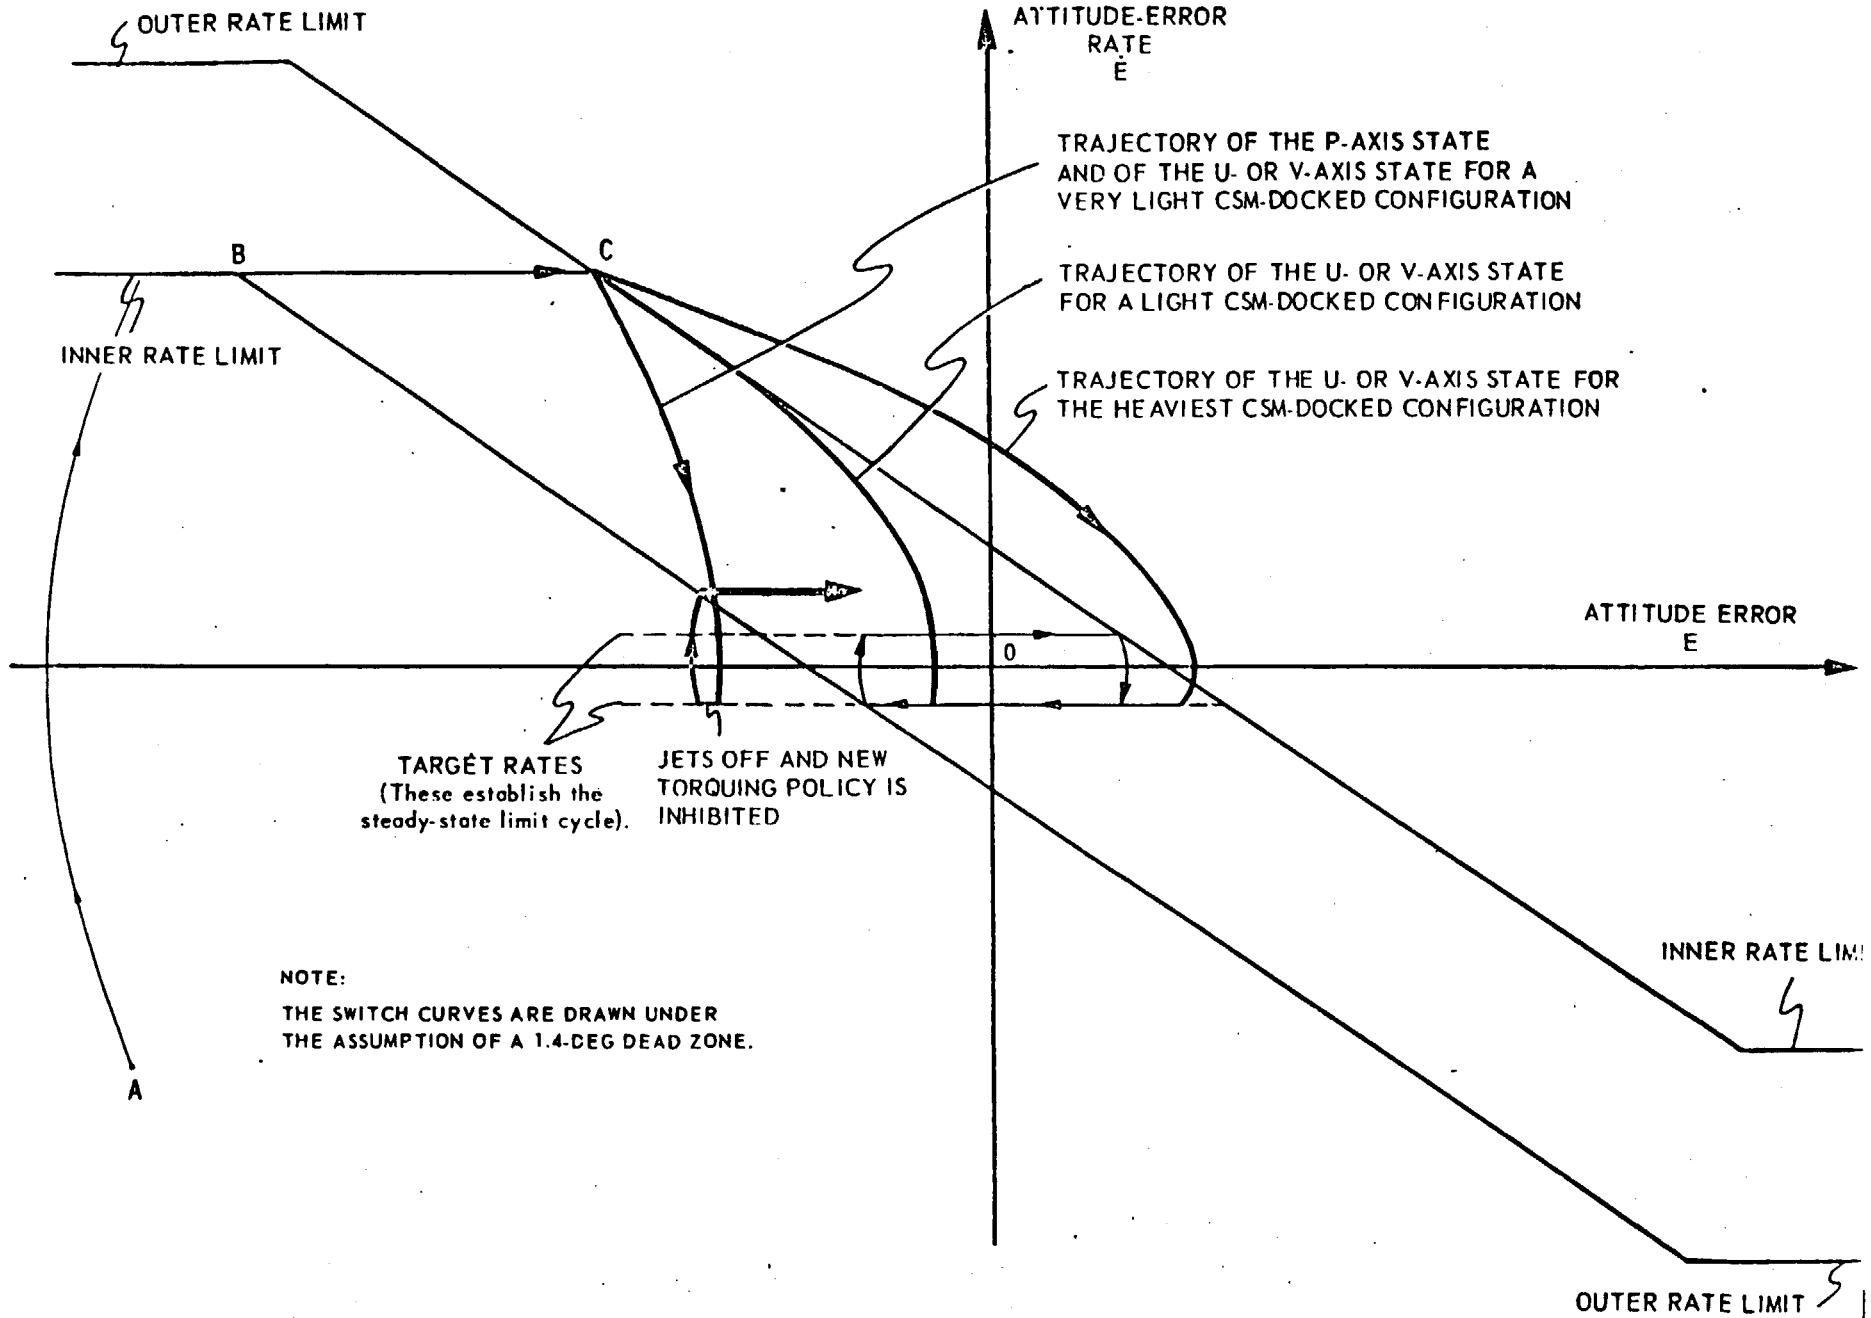

Figure D. 6-8 Typical Error-State Trajectories of the CSM-Docked Configuration

frequency. Due to sampling lags and lags in the state estimator, these jet firings were in phase with the bending oscillations and would sustain them even after the failed-on jet had been detected and turned off. Such bending instabilities were highly undesirable, since they produced forces significantly larger than the maximum load capability of the LM/CSM docking tunnel. Consequently, the LM DAP was modified to improve the bending stability in the CSM-docked configuration. Two approaches were utilized. The first approach reduced the bending excitation due to RCS jet firings by inhibiting all jet firings about an axis for a predetermined time interval each time the jets firing about that axis were turned off. This jet-inhibition scheme significantly reduced the bending excitation by ensuring that the jet-firing frequency was always lower than the resonant frequencies of the most significant bending modes. The second means of improving the bending stability in the CSM-docked configuration prevented the oscillations which did develop from becoming self-sustaining. With this approach, the jets were turned off and left off for a predetermined time interval each time a reversal of torquing direction was commanded about an axis in which jets were still firing. Simulation results and theoretical worst-case analyses indicated that these modifications reduced to a safe level the maximum bending-oscillation magnitudes and bending moments on the docking tunnel in the presence of a jet-on failure.

For some CSM-docked mass loadings, the presence of a disabled or failed-off -X thrusting jet can cause a serious control instability when the LM DAP attempts to control pitch or roll attitude. The problem can be explained using Fig. D.6-9. Typically, when a jet pair is selected to induce a commanded clockwise rotation about the cg, the 100-lb downward-thrusting jet impinges upon the jet-plume deflector, producing a force of 89 lb in the +X direction acting on the moment arm D_1 and a force of 59 lb perpendicular to the X axis acting on the moment arm D_2. The net moment due to the firing of the downward-thrusting jet is

$$M_{+X} = (89\ \text{lb})D_1 - (59\ \text{lb})D_2$$

which, for many mass loadings, can be negative, thus commanding the vehicle to rotate counterclockwise. Normally the -X thrusting jet can counteract this moment, but should it fail-off or be disabled, a grave instability results. Even in a normal situation, propellant consumption is excessive for the amount of net torque gained.

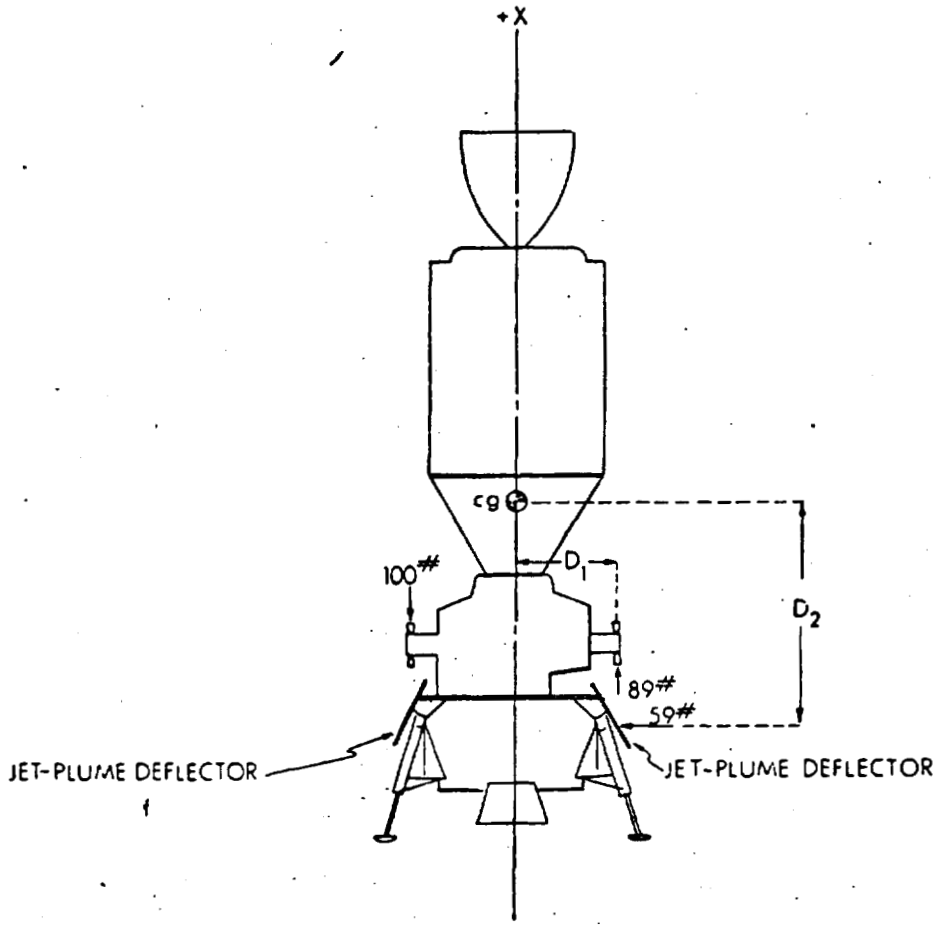

Figure D. 6-9 CSM/LM Docked Configuration with Jet-Plume Deflectors

To eliminate this problem, MIT proposed modification of the DAP jet-select logic such that downward-thrusting jets producing a net torque in the wrong direction are avoided. NASA rejected this software fix and instead used a procedural work-around in which the deflected jets are manually disabled when they produce a net torque in the wrong direction.

The CSM-docked control law illustrated in Fig. D.6-7 included provisions for rapid recovery from large attitude errors and rates. In Fig. D.6-8, for example, jets would be fired continuously to move the state of the spacecraft from its initial position at A to the inner rate limit at B. At B the jets would be turned off to allow the state to drift across the coast zone to C. At C, the jets would fire until the error rate reached zero. Usually, for U- or V-axis control, the state would then begin limit cycling. However, for U- or V-axis control of a very light configuration, or for P-axis control, the higher control authority would cause the state to undershoot to a point outside the coast zone. In such a case, the inhibition logic would prevent jet firings for a short time as the state drifted slightly. The jets would then torque in the opposite direction until the state entered the coast zone, after which a short firing would occur to produce a slightly negative error rate and initiate the limit cycle. (The extra two jet firings needed to recover from large errors and rates, in this case, are the penalty for using a single fixed phase plane to control all axes for all mass loadings in the CSM-docked configuration.)

D.6.3.2 Automatic-Maneuvering Mode

Since the main engine is not thrusting during coasting flight, automatic maneuvers required for certain mission functions (e.g., preburn alignments, rendezvous tracking, etc.) are performed using RCS jets. These maneuvers are controlled by the LM DAP according to a set of steering variables computed by the attitude-maneuver routine, KALCMANU. These steering variables affect only the attitude and rate errors used by the RCS control-law phase planes. A principal design objective of the Automatic-Maneuvering mode is accurate tracking of desired gimbal angles and desired spacecraft rates; in addition, the four design objectives of the Attitude-Hold mode cited in Section D.6.3.1 also apply to Automatic Maneuvering. In every other respect, the Automatic-Maneuvering mode is identical to the Attitude-Hold mode.

D.6.4 Descent Powered Flight

In the powered descent to the lunar surface, the LM DAP must provide considerable precision in attitude control, with rapid response to attitude deviations. Control must be precise because seemingly insignificant deviations from the requested (preprogrammed) trajectory can result in landing-point errors of several miles.

The terminal descent and touchdown of a lunar landing can be done completely automatically; however, some manual intervention is permitted to allow for late landing-site redesignation. This manual-override capability has priority over the automatic guidance requirements. In either the automatic or manual-override control modes, the descent powered-flight DAP can alter its control procedures and/or jet-selection logic in the unlikely event of an RCS jet failure—on or off, recognized or not—or a failure of the descent-engine Gimbal Trim System (GTS).

As discussed in Section 6.1.1, the design of the powered-flight portion of the LM DAP was approached through four overlapping pathways—estimation of slowly-varying parameters, state estimation, RCS jet selection and timing, and GTS-drive selection and timing. However, the GTS-drive selection is particular to descent powered flight and is presented here in greater detail.

The LM descent engine can be rotated about the Q and R axes at a constant rate of 0.2 deg/sec by the Gimbal Trim System, thus directing the thrust vector through the spacecraft center of mass. Consequently, the RCS jet-control system is relieved of the burden of a bias angular acceleration due to an offset thrust vector. This mode of operation is called the GTS acceleration-nulling mode. Further analysis showed that the GTS could also control the LM angular rates and attitude, but only when the flight conditions placed relatively mild requirements on the DAP. This option was implemented in the LM DAP and is called the GTS attitude-control mode. The time-optimal control law was modified in the LM application so that delays attributable to the 0.1-sec cycle period of the DAP and to the mechanical lags in the trim gimbals would not generate an unacceptable large, steady-state, control-limit cycle.

The existence of two independent Q,R axis attitude-control laws (RCS and GTS) created the possibility of conflicting control torques being applied to the LM, so a

set of criteria were specified defining the precedence of the two systems, in an effort to obtain maximum benefit from the GTS (saving RCS firings and fuel) while preserving the rapid response of the RCS system.

The RCS/GTS interface was organized to implement these specifications. RCS alone exerts control over LM attitude, using the RCS phase plane and coast zone, until GTS is declared usable by a ΔV monitor and the gimbal monitor. Every two seconds during powered descent, the ΔV monitor and the gimbal monitor verify that ignition has been achieved, that adequate thrust is present, that the astronaut is not signaling a GTS failure, and that the astronaut has not indicated the onset of ascent. When these specified conditions are met, GTS control is admissible.

During a nominal automatic lunar descent, the GTS alternates control with the RCS on successive DAP passes. Considerable care was exercised to ensure that conflicting commands and chatter were avoided. However, should the astronaut select a manual override of the automatic system, the GTS becomes limited to acceleration-nulling control.

Within the GTS portion of the DAP, a flag is checked to determine whether the thrust vector has yet been brought to within one degree of the center of mass. Until this condition is met, acceleration nulling will be executed every two seconds under 1/ACCS. If that flag is set, GTS interrogates the jets before exercising control on every other DAP pass to see whether any RCS jets are firing. If RCS jets are firing, the jets retain primary attitude control, and GTS is temporarily limited to acceleration nulling, with one nulling drive being executed immediately from GTS if attitude control was being exercised by GTS on the preceding GTS DAP pass. If the RCS jet interrogation shows all jets off, GTS executes attitude control every 0.2 sec, controlling the angular acceleration, the rate errors, and the attitude errors as long as the spacecraft attitude remains within the RCS coast zone.

D.6.5 Ascent Powered Flight

The LM ascent configuration in powered flight requires a much higher duty cycle for RCS jet activity than does the powered-descent LM, since, unlike the DPS engine, the APS engine is rigidly mounted, unable to control the angle of the thrust vector relative to the spacecraft cg. Attitude control of the LM ascent configuration

is achieved solely by the RCS jets commanded by an appropriate RCS control law.

Since the ascent engine produces only a small torque about the vehicle P axis (the engine is nominally canted only 1.5 deg away from the P axis), the control problem about this axis is relatively straightforward. Indeed, the RCS control law assumes that no disturbing torque exists about the P axis; the major control problem for the ascent LM, therefore, is one of controlling attitude about axes perpendicular to the vehicle P axis—in particular the autopilot U', V' axes.

D.6.5.1 Autopilot Single-Jet Control Boundary

To avoid diminishing the effective thrust of the ascent engine, an important requirement for powered-ascent control is that, whenever possible, only upward-forcing RCS jets be used. In the Q,R plane (which contains the U,V and U',V' axis systems) a locus of vehicle cg positions exists, within which, theoretically, no more than a single U-axis or V-axis RCS-jet firing would be necessary to maintain attitude control of the vehicle. This locus is called the theoretical single-jet control boundary. For cg positions lying on the single-jet control boundary, at least one upward-forcing jet must fire continuously to maintain control; and for the particular axis (U or V) about which the jet is firing continuously, the net torque produced about that axis by the ascent engine and by this continuously-firing jet is zero. For cg positions outside this boundary, control cannot be maintained unless at least one downward-forcing jet is used.

In describing the single-jet control boundary, it is useful to employ a quantity called the "effective" cg displacement from the thrust axis of the ascent engine. The components of the effective cg displacement from the thrust axis along the Q and R axes can be defined as follows:

$$cg_{Q(eff)} = \frac{\text{torque about the R axis produced by ascent engine}}{\text{ascent-engine thrust}}$$

$$cg_{R(eff)} = \frac{\text{torque about the Q axis produced by ascent engine}}{\text{ascent-engine thrust}}$$

In Fig. D.6-10, the locus of effective cg displacements which defines the theoretical single-jet control boundary is drawn in the Q,R plane. In addition, locations of the upward-forcing RCS jets is shown together with the U,V and U',V' axis systems.

In calculating the theoretical control boundary shown in Fig. D.6-10, it is assumed that the component of the ascent-engine thrust along the vehicle X axis has a magnitude of 3500 lb, and that a U jet produces a torque of $550\sqrt{2}$ ft-lb about the U axis, and a V jet produces a torque of $550\sqrt{2}$ ft-lb about the V axis. Under actual operating conditions, these assumptions do not hold exactly true, and some deviation from the theoretical control boundary can be expected. For the purposes of this discussion, however, it will be assumed that the single-jet control boundary can be represented as in Fig. D.6-10.)

During powered ascent, the autopilot determines whether the cg lies within the single-jet control boundary by estimating the net control acceleration produced by a single upward-forcing jet about the appropriate U' or V' axis. (Net control acceleration is the angular acceleration about a given control axis, U' or V', which results from the combined torques produced by the commanded RCS jet and by the ascent engine.) Due to the way in which the U', V' system was constructed, a net control acceleration of zero about the U' axis is equivalent to a net torque of zero about the U axis; and a net control acceleration of zero about the V' axis is equivalent to a net torque of zero about the V axis. As noted above, the determination of a zero net torque about the U or V axis is required to establish the single-jet control boundary; and thus the computation of net acceleration in the nonorthogonal U', V' system allows the autopilot to make this determination.

It is not desirable, however, for the autopilot to allow net control acceleration to become as small as zero. (This would mean continuous jet firing to merely maintain attitude.) When the autopilot determines that the net single-jet acceleration about a control axis is less than $\pi/128$ rad/sec^2, a decision is made to use two-jet control about that axis. (This number is determined from scaling considerations in the AGC to avoid overflow.) The autopilot, therefore, assumes a single-jet control boundary which lies within the theoretical boundary.

Nominal cg positions for the LM ascent configuration during powered ascent lie well within the single-jet control boundary and would be expected not to require two-jet firings about a control axis.

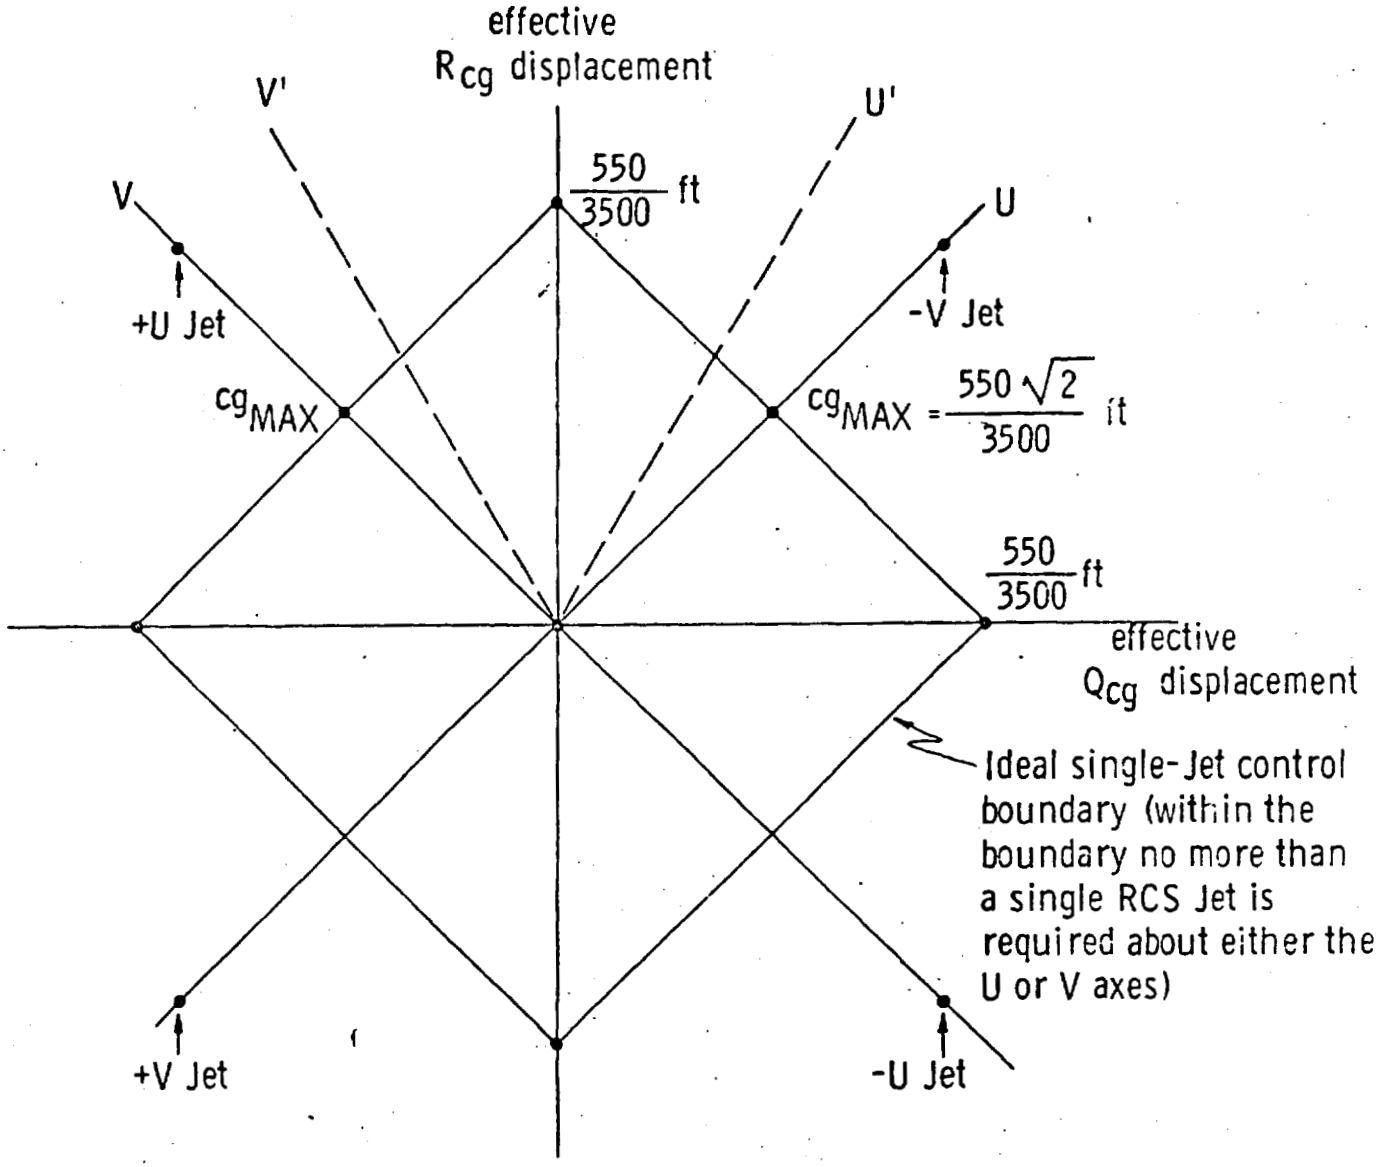

Figure D. 6-10 Ideal Single-Jet Control Boundary

During powered ascent, certain conditions (which may adversely affect guidance), such as an excessive attitude-error angle or error-angle rate, may require the use of two-jet control for the U and V axes, even for cg positions within the single-jet control boundary. The use of two-jet control is mandatory for sufficiently large errors of this type, even though the use of downward-forcing jets is ordinarily deemed undesirable. However, simulations show that in steady-state operation, error-angle excursions are sufficiently small to avoid two-jet control for all cg positions within the DAP single-jet control boundary.

D.6.5.2 Effect of Incorrect Knowledge of Inertia

The vehicle's moment of inertia is not used explicitly by the DAP. However, a priori knowledge of this quantity is assumed in the AGC subroutine which computes single-jet acceleration as a function of mass (for the P, Q and R axes). If the function is not accurate or if the computer's knowledge of mass is inaccurate, then this is equivalent to an error in knowledge of vehicle moment of inertia. Since the autopilot does not have a filter for estimating single-jet acceleration, no correction can be made to the value computed in the AGC; however, the DAP's recursive state estimator does estimate offset acceleration (offset acceleration is defined here as that part of the angular acceleration which cannot be explained by commanded jet firings). An important characteristic of the state estimator is that any error in the computed single-jet control acceleration tends to produce a compensating error in the state-estimator value for offset acceleration. (This compensation does not occur during coasting flight, since offset acceleration is not estimated in this case.) The result is that a compensating error is made in the DAP's computation of the net control acceleration about a control axis.

To test the effect of incorrect knowledge of moment of inertia on autopilot performance, all-digital simulations of the powered ascent LM were run in which extreme "mass-mismatches" existed (i.e., a large error existed in the AGC's knowledge of mass). For the purpose of this test, a special LM subprogram was used in which an effective cg displacement was selected of 2 in. from the thrust axis along the vehicle Z axis, as illustrated in Fig. D.6-11. In one simulation run, a mass-mismatch case was studied in which the DAP determined the mass of the LM ascent configuration to be 2775 kg and the actual mass was 4933 kg. As a

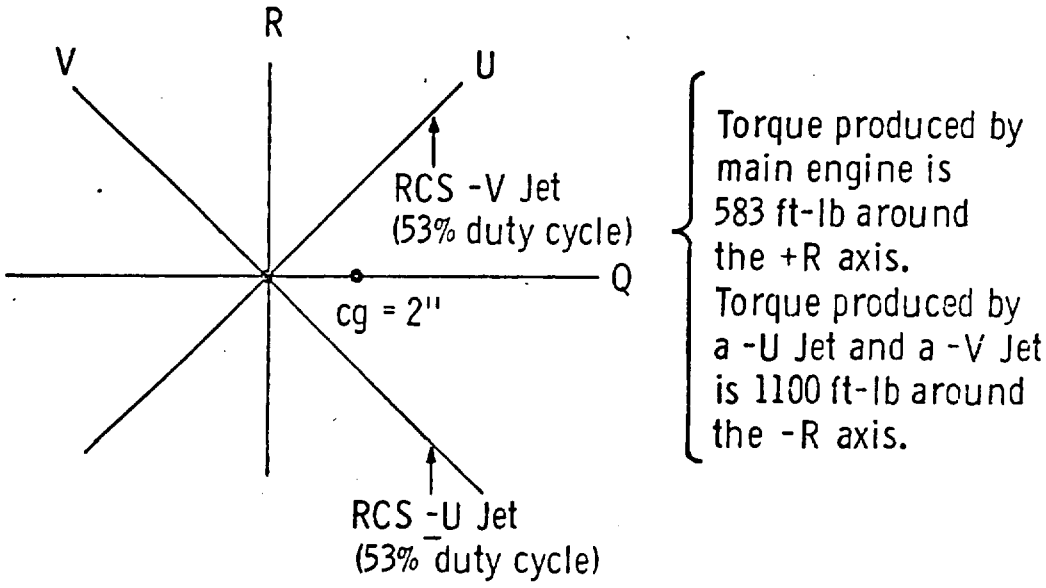

Figure D. 6-11 Effective cg Displacement for Simulations of Mass-Mismatch

result, the DAP overestimated the single-jet accelerations about the P,Q and R axes. For example, at approximately 20 sec after ignition, the DAP overestimated the single-jet acceleration about the R axis by 9.1 deg/sec^2 for each of the two jets required to maintain control, for a total overestimate of 18.2 deg/sec^2 for the two jets. This error, however, was partially compensated for by an overestimate in offset acceleration about the R axis of 9.9 deg/sec^2. To illustrate the effectiveness of the autopilot in maintaining control for this case, a phase-plane plot (attitude-error rate vs attitude error) of a typical steady-state limit cycle is given in Fig. D.6-12. The phase plane is shown for the autopilot U' axis. Because of the errors in the DAP's knowledge of the RCS single-jet acceleration, the U' and V' axis directions determined by the DAP differed from the theoretical U' and V' directions; thus, acceleration cross-coupling existed between the DAP U' and V' control axes. The effect of this cross-coupling, however, was not a serious one. On the average, the extra firings required per limit cycle were 1.0 for the U axis and 0.7 for the V axis. Phase-plane attitude-error angle excursions were deemed reasonable. (A second simulation was run in which the DAP determined the mass of the LM ascent configuration to be 4933 kg, while the actual mass was 2775 kg. Again, as in the previous case, no serious complications resulted from cross-coupling. On the average, the extra firings required per limit cycle were 1.3 for the U axis and 0.8 for the V axis.)

D.6.5.3 Effect of an Undetected Jet Failure

The problem of undetected jet failure is similar to that of mass-mismatch, although there are important differences. The case of an undetected failed-on jet presents no special control problem for the autopilot. This is because the angular acceleration produced by the failed-on jet can be considered equivalent to an additional offset acceleration produced by the ascent engine. In either case, the DAP state-estimator filter estimates the acceleration as an offset acceleration, and no error is introduced into the DAP's knowledge of net control authority.

The more difficult case for the autopilot is that of an undetected failed-off jet. When this jet is commanded to fire, the state estimator "sees" the lack of response as resulting from an opposing step in offset acceleration. (The effect is similar to that of the mass-mismatch case where the DAP overestimates the single-jet accelerations about the P,Q and R axes.) If the DAP designates the failed-off jet for single-jet control, then control will temporarily be lost about that axis. This

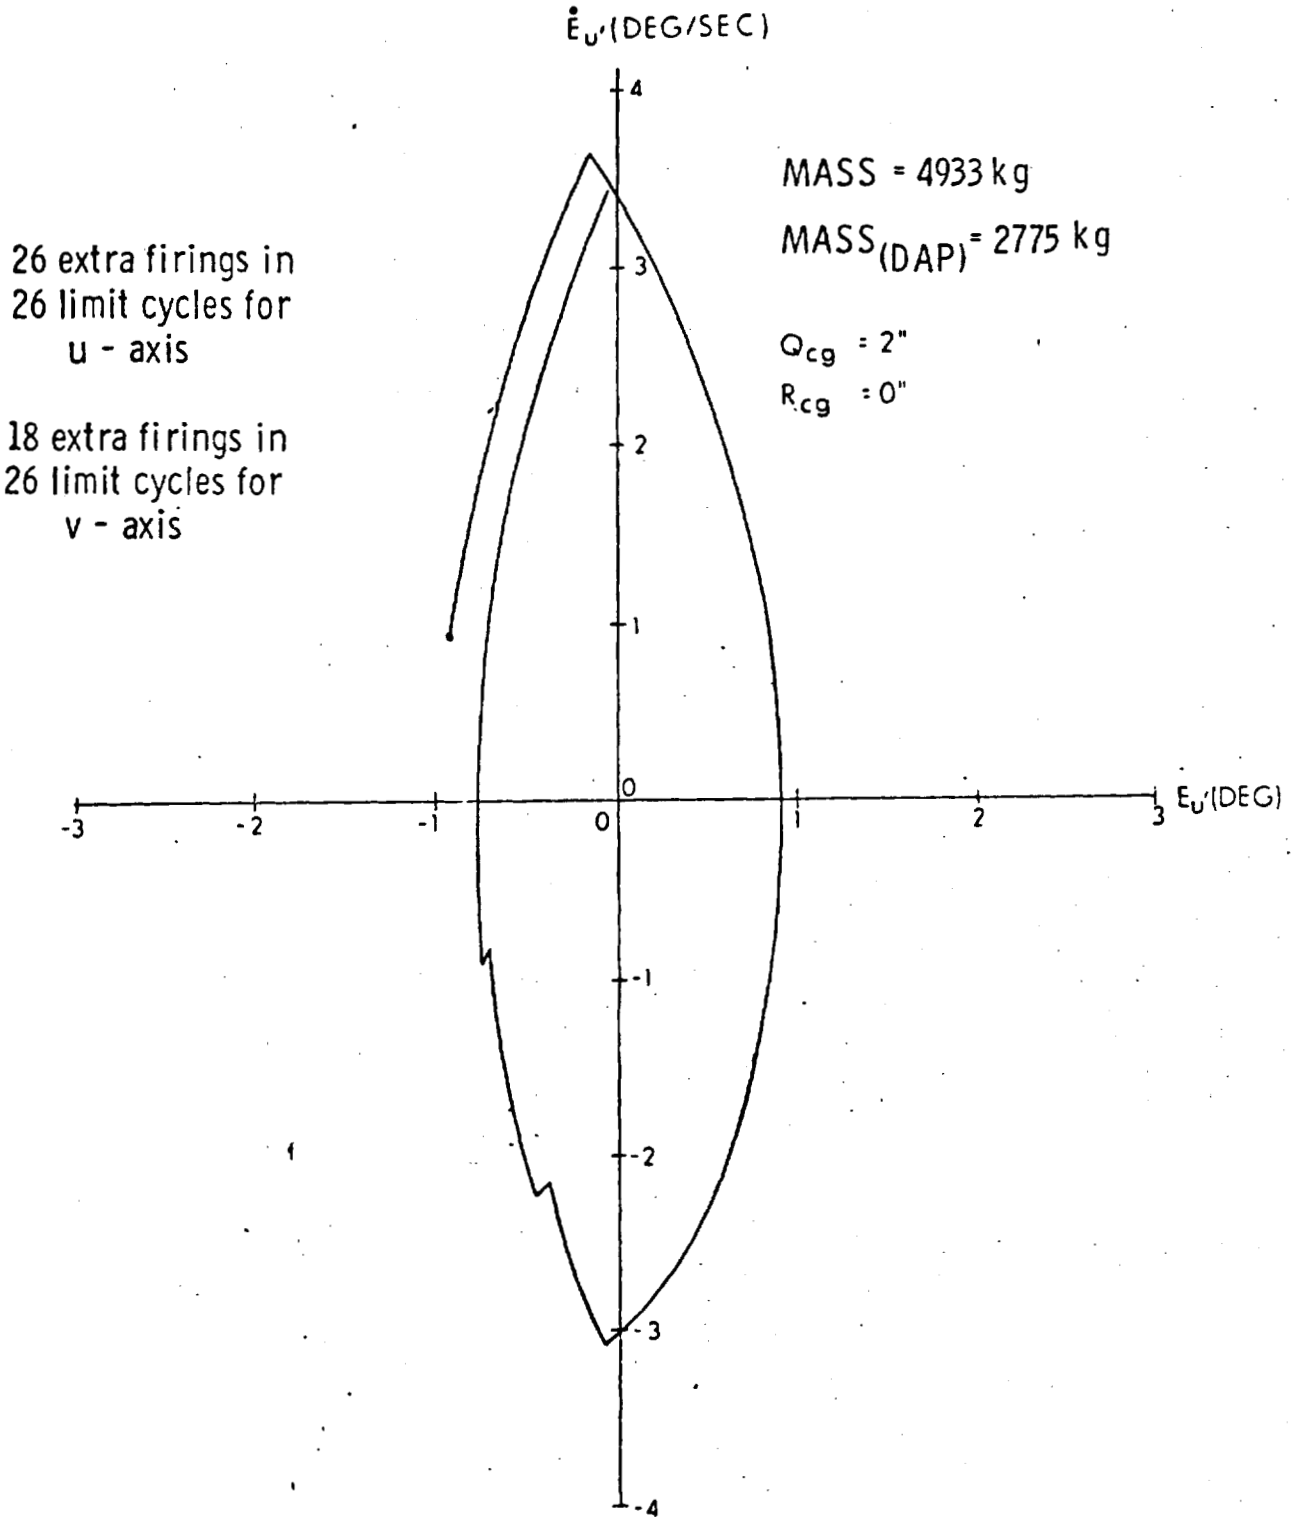

Figure D. 6-12 Limit-Cycle Behavior for Mass-Mismatch Heavy Vehicle–Powered Ascent

loss of control will exist until the attitude error or attitude-error rate is sufficiently large to require mandatory two-jet control or until the estimate of net angular acceleration becomes sufficiently small that the autopilot concludes that the cg lies outside the single-jet control boundary. This situation is illustrated in Fig. D.6-13. The V' axis phase-plane limit cycles in this figure are taken from a digital-simulator run in which the effective cg displacement from the thrust vector is constrained to be along the vehicle U axis, with Q and R components of 0.5 in. On paths A-B, F-G and H-I, single-jet control (using the failed jet) is commanded. Thus, when the phase-plane state lies on these paths, the only angular acceleration acting about the V' axis is the V' component of offset acceleration, and control is temporarily lost. Paths B-C and G-H lie in a region of the phase plane which calls for mandatory two-jet control. On these paths, since one operable jet is firing, control of the vehicle is regained. On paths C-D and I-J, the estimated value of net angular acceleration is sufficiently small to cause the autopilot to conclude that the single-jet control boundary has been exceeded; two-jet control is then commanded. During periods of the limit cycle in which the failed jet is not commanded to fire (c.g paths D-F and H-K), the state estimator tends to reduce its offset-acceleration estimate for the V' axis toward the correct value. This causes the estimate of net single-jet angular acceleration to increase. At points E and K, the estimate of net angular acceleration has become sufficiently large to cause the autopilot to conclude that the cg again lies within the single-jet control boundary. It can be seen from the above that, for the case being tested, the autopilot alternates between periods in which it concludes that the cg lies within the single-jet control boundary and periods in which it concludes that the cg lies outside the single-jet control boundary. This results in "loose" limit-cycle activity, as can be seen in Fig. D.6-13. It will be noted that the existence of mandatory two-jet control regions in the phase plane does prevent unreasonably large attitude errors from building up. For a sufficiently large offset torque about an axis requiring a failed-off jet for single-jet control, the DAP will always determine that the offset acceleration lies outside the single-jet control boundary. The offset torque about the axis in question must exceed 275 2 ft-lb (half the torque produced by one RCS jet) for this to be true. Since single-jet control is never called for, in this case, reasonably tight limit cycles can be expected.

D.6.5.4 Velocity Errors

In general, the average attitude error for an autopilot phase-plane limit cycle is non-zero and will depend upon the values of offset acceleration and RCS control

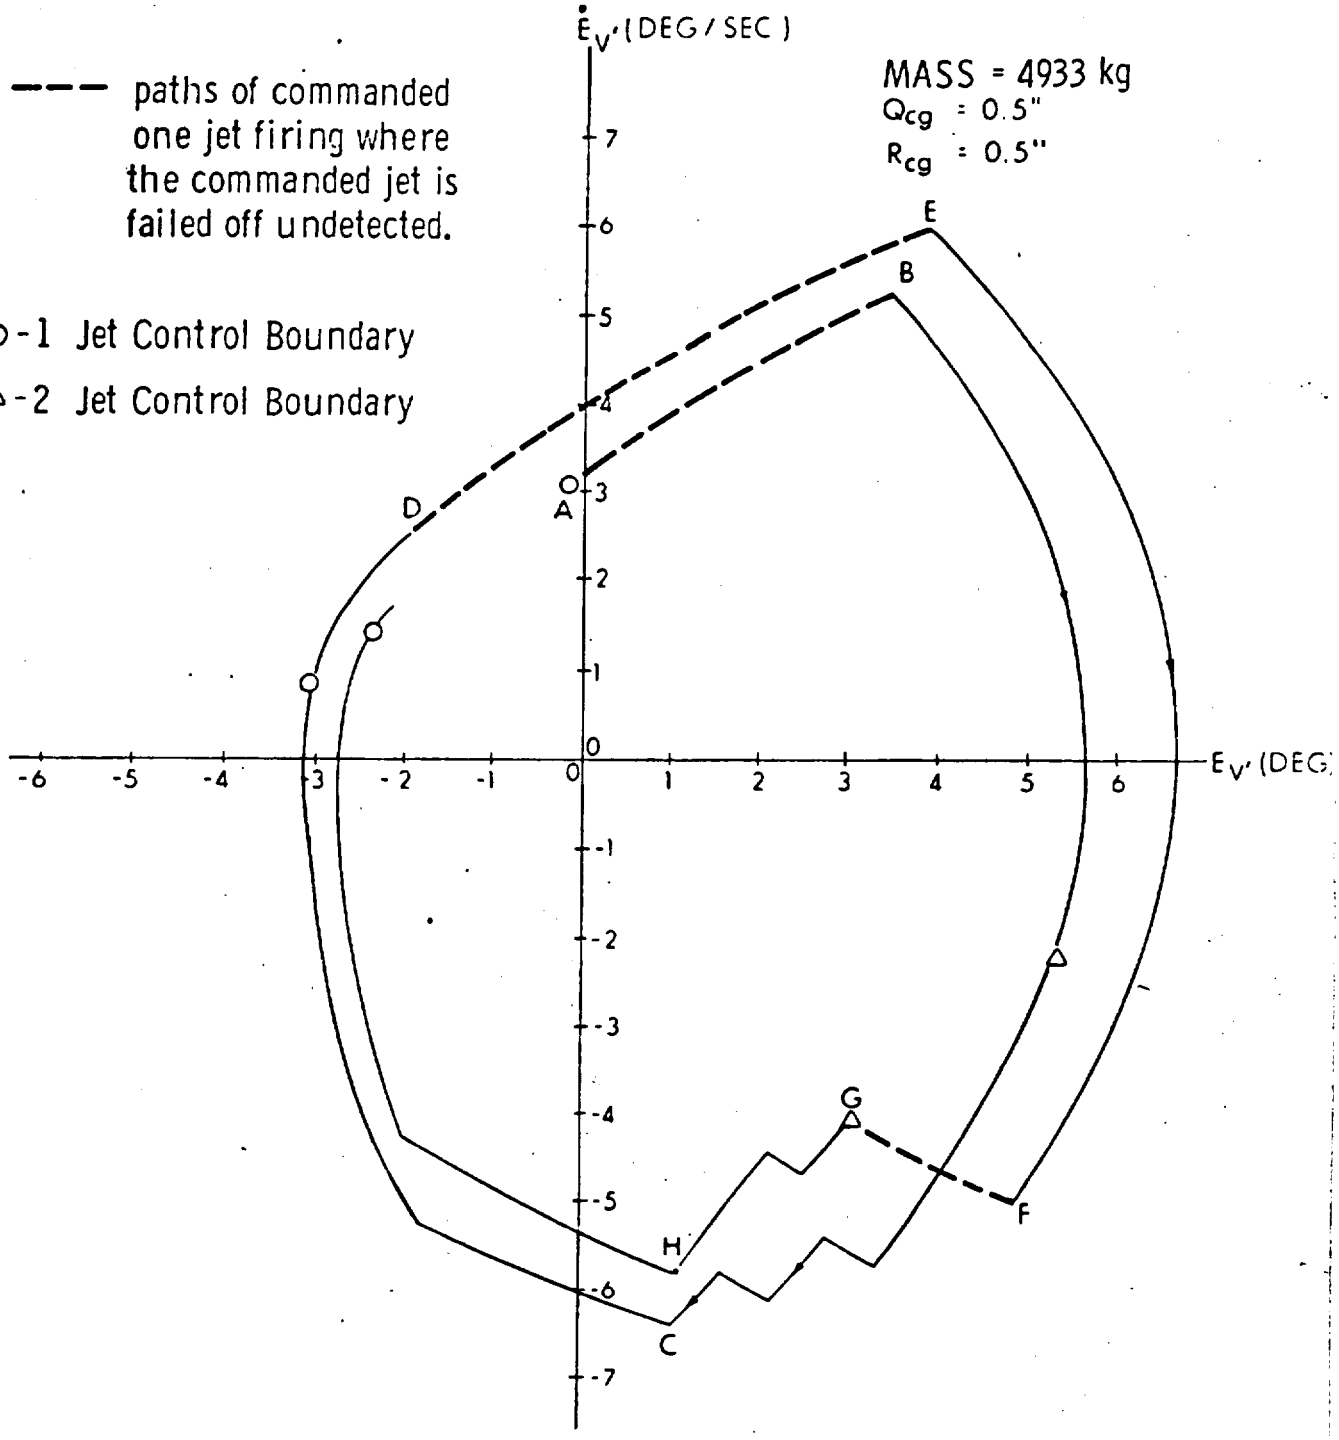

Figure D. 6-13 Limit-Cycle Behavior for -V Jet Failed Off (Undetected)-- Powered Ascent (Nonorthogonal Control Axes)

acceleration. During each 2-sec guidance period, this average attitude error will tend to produce "cross-axis" velocity, i,e., a component of velocity normal to the commanded direction of the ascent-engine thrust vector. An additional source of cross-axis velocity is the inaccuracy in the knowledge (by the guidance law) of the direction of the ascent-engine thrust vector with respect to the vehicle axes. One of the functions of the FINDCDUW Guidance/Autopilot Interface Routine (see Section C.7) is to respond to a bias in thrust direction and to modify the commanded direction of thrust provided by the guidance law to correct the bias. The part of the FINDCDUW routine which performs this function is the Thrust-Direction Filter. In the presence of a fixed bias in the direction of the LM ascent-engine thrust vector with respect to the desired X-axis orientation, the Thrust-Direction Filter acts as a first-order filter with a time constant of 8 sec. In general, therefore, it can be expected that for a fixed bias in the direction of the thrust vector with respect to the DAP's desired vehicle X-axis direction, 24 sec will be required for the filter to correct for 95 percent of the bias.

The effects of autopilot bias and thrust-vector misalignment during powered ascent on velocity errors at engine cutoff have been tested on the All-Digital Simulator. The tests used the APS powered-flight guidance program, P42, with external ΔV guidance. A light ascent vehicle was assumed and ΔV was supplied as an input to the program. It was concluded from this study that the Thrust-Direction Filter was effective in reducing the cross-axis velocity resulting from either of the causes discussed above. Attitude-pointing errors produced by the cant of the ascent engine from the vehicle X axis are easily corrected, since the engine cant is a fixed bias throughout the run. Cross-axis velocity at cutoff for all of the runs was generally within 1 ft/sec. (In two test cases, however, cross-axis velocity exceeded 1 ft/sec. An extreme off-nominal cg close to the maximum cg displacement which can be controlled by the DAP produced a 2 ft/sec error at cutoff. In a case where the cg was constrained to lie along the thrust vector, a 1.25 ft/sec error resulted. In this latter case, a slow oscillation of the FINDCDUW thrust-axis direction results. This is a consequence of the flow limit-cycle period in the phase plane for the case in which the offset acceleration is zero. Since the limit cycle period is much larger than the 2-sec guidance period, the Thrust-Direction Filter follows the slow oscillation.)

D.6.6 CSM-Docked Powered Flight

The contingency of a CSM Service Propulsion System failure prior to the initiation of lunar landing would require that the LM descent configuration become the active spacecraft, pushing the CSM/LM configuration into a trajectory for the flight back to earth. Design considerations for LM DAP control of the CSM-docked configuration are complicated by radically altered control authorities and mass distributions and four physical constraints—bending and torsion at the docking terminal; slosh interactions between the two vehicles; RCS jet-plume impingement; and engine-on and throttling transients. These constraints are discussed below.

D.6.6.1 Bending and Torsion Constraints

Bending oscillations—and torsion, to a lesser degree—could jeopardize the structural integrity of the docking tunnel which joins the two comparatively rigid vehicles. Accordingly, appropriate care must be taken in the LM DAP design to ensure that the RCS jets (the principal instigators of bending excitation at critical frequencies) reinforce bending amplitudes.

D.6.6.2 Slosh Constraint

When the vehicles are in the CSM-docked configuration, consideration must be taken for the slosh modes imposed by the presence of both CSM and LM propellant. It must be verified that the LM DAP design ensures that the Gimbal Trim System (the principal agitator of slosh) does not cause slosh amplitudes to exceed acceptable levels.

D.6.6.3 RCS Jet-Plume Impingement Constraints

Thermal constraints imposed by RCS jet-plume impingement are even more severe in the CSM-docked configuration than in the LM-alone. (Section D.6.3 discusses these constraints in the context of LM coasting flight.) In the docked configuration, the LM RCS jets can impinge on both vehicles. The problem is further aggravated, because the larger pitch and roll inertias associated with the CSM-docked configuration necessitate longer jet-on times to achieve desired performance.

To combat the serious thermal constraint imposed by the jet-plume impingement, the crew can initiate, via the DSKY, LM DAP routines to inhibit the firing of U- and V-torquing jets during CSM-docked powered flight. In addition to this software approach, jet-plume deflectors were added to the LM; their effect upon LM DAP control of the CSM-docked configuration is discussed in Section D.6.3.

D.6.6.4 Constraints Related to Engine-On and Throttling Transients

Because all but the yaw-axis jets are inhibited during CSM-docked powered flight (thus preventing impingement on either the CSM or the LM jet-plume deflectors), the GTS must assume control of the spacecraft about the pitch and roll axes. Under these circumstances, engine-on and throttling transients become a serious GTS design consideration. These transients stem from three factors—computational underflow, initial mistrims of the descent-engine bell, and compliance of the descent-engine mount.

Until the computational-underflow problem was solved in the LUMINARY IA program (for Apollo 11), the scaling of the GTS control law was such that attitude errors could not be detected for the CSM-docked configuration at low thrust. Attitude errors could thus not be steered out during the ten-percent thrust period (26 sec) prior to throttle-up. A 40-percent manual throttle-up a few seconds after ignition avoided these errors*; at this thrust level, the more significant attitude errors could be detected by the GTS and steered out. Because it yielded superior performance, this intermediate manual throttle-up procedure was retained even after the computational-underflow issue was resolved.

Transients are also produced by mistrims of the descent engine prior to ignition. (A "mistrim" is here defined as a deviation of the engine-bell orientation.) Because the engine-bell gimbal drive rate is only about 0.2 deg/sec for the descent engine, substantial time is required to correct large initial mistrims, during which significant excursions in attitude error, rate and acceleration might occur.

*All four of these transient factors are kept to a minimum by this 40-percent manual throttle-up procedure.

The descent-engine mount has a thrust-proportional misalignment (compliance) associated with abrupt throttle changes. This association effectively introduces a mistrim at throttle-up—or any other throttle manipulation—with an effect comparable to that noted in the previous paragraph.

A NOTE ON SOURCES

By the time this record of MIT's Apollo software efforts entered the germination stage, many of the souls who had participated in those efforts had begun to disperse to other projects—both within and without the Draper Laboratory. Nonetheless, virtually all wanted to ascertain that that part of the history in which they played such important roles was finally, indeed, recorded. Some personally documented their accomplishments; others supplied bits and pieces that eventually interlocked to permit the construction of a unified whole. A considerable amount of the information recorded within these pages could be gleaned from documents that already existed*, but a surprising—indeed, exasperating—amount had never before been documented. For the latter, memories had to be tapped and taped; 57 transcribed interviews, countless conversations and mountains of notes bear testimony to the cooperation and enthusiasm which I encountered along this historical path.

Considerably more information was gathered than could be presented within any single cohesive text. But what I hope has remained is an insight into the team which carried the concept of Apollo software from a hopeful infancy, through an oft turbulent adolescence, to its magnificently successful goal. The Apollo software team was a heterogeneous, sometimes colorful lot, one which demonstrated two basic characteristics: competence and perseverence. The pressures imposed by the schedule sometimes revealed frailities, but more often demonstrated elemental strengths. The epoch which this history records was a significant time in all the participants' professional and personal lives, and, in recalling this period, no one felt dispassionate. Despite the crushing schedules, the fantastic amount of mental and physical exertion which the project came to demand—the goal which was to be reached seemed to energize us all.

Not every member of the Apollo software team contributed to this report, but a great many did, some extensively, some less so. In listing these persons below,

* A compendium of abstracts of all Laboratory reports pertaining to Project Apollo appears as an Appendix to Volume I of this report.