Qualitative Differences Between Shuttle On-Orbit and Transition Control

Philip D. Hattis*
The Charles Stark Draper Laboratory, Inc., Cambridge, Massachusetts

Two separate reaction control system digital autopilots evolved from one original Space Shuttle orbital autopilot concept. A computer overload forced this evolution. Part of the overload problem was due to unique performance requirements imposed on the reaction control system during each of several different flight regimes. The two resultant reaction control system digital autopilots yield different effector responses because they rely on different sources of sensory input and they process data differently. This paper describes the evolution of the two and illustrates their behavioral differences. The transition autopilot, used in orbital insertion and deorbit, is sensitive to Orbiter flexure due to its feedthrough character. The on-orbit autopilot is sensitive to transient rate control degradation from large disturbances due to feed-forward rate estimation. Simulation results and flight data are used to illustrate performance differences between the two autopilots under various conditions. These include computer failures where electronic stringing and procedural reconfiguration differences affect autopilot behavior.

Introduction

THE Space Shuttle reaction control system (RCS) consists of 38 primary jets, each producing about 870 lbf, and six vernier jets, each producing about 24 lbf. All thrusters have 80-ms cycle-time granularity. The thrusters are located in three pods—forward, left aft, and right aft—with Fig. 1 depicting the general layout.[1] The primary jets are divided into 14 clusters around the vehicle with translation and rotation control possible and multiple failure tolerance available for the most essential capabilities. The vernier jets each fire in a different direction. They only control rotation and have no failure tolerance if any one of four of the thrusters malfunction. Fig. 2 illustrates the jet cluster configuration with an explanation of the identification nomenclature.[2]

The primary jets are essential to control large disturbances such as residual torques after shutdown of Space Shuttle main engines, auxiliary power unit venting, and orbital maneuvering system failures. Time-critical maneuvers, such as separation from the external tank and rotation to orbit change burn attitude, also require the large thrusters.

The vernier jets are suitable for long-term attitude hold with small disturbances such as gravity gradients, and they perform more efficiently for this purpose. The vernier jets make finer pointing and tracking control possible. The small vernier control torques are also more compatible with the operation of some highly flexible payload deployment devices.

Large venting forces and flight-critical maneuvers are most prevalent in orbital flight shortly after ascent and shortly before re-entry. Long-term attitude control, tracking, pointing, and payload operations are required primarily while the Shuttle is on-station in orbit.

Five identical digital general purpose computers operate the Shuttle in a fly-by-wire configuration. One computer is dedicated to backup software in case generic problems with the programs in the other computers arise. The remaining primary computers can be processed synchronously with any number of them making identical computations for redundancy. However, the sensors and actuators for all electrically operated systems have dedicated control through only one computer at a time in a manner determined in part by system stringing procedures. An active primary software autopilot always resides in at least one of the primary computers. Only those computers with an autopilot resident can communicate with the avionics in a manner determined by the crew-assigned stringing.

The Shuttle reaction control system autopilots have inherited some features from the Apollo and Skylab programs.[3-5] The Apollo design included a second-order Kalman filter for rate estimation from inertial measurement unit data with feed-forward of expected accelerations of commanded jets, a per rotation axis phase plane control law (which used straight switch lines for simplicity), and a table lookup jet selection. In the Skylab program, accommodation was made to consider specific interaxis rotation coupling due to a significant center-of-gravity offset from the thrusters in the stacked spacecraft/laboratory configuration. In Apollo and Skylab, all nominal actuators were controlled through one computer.

While the Shuttle autopilot design maintains some of the Apollo/Skylab heritage, the differences in the Shuttle vehicle and operations ultimately led to many changes. The Shuttle thruster layout has only one plane of symmetry increasing rotation coupling problems. Two different types of thrusters require separate control with the small number of vernier jets unamenable to table lookup selection due to complex coupling effects. Payload operations unique to the Shuttle era introduce requirements for new control modes. The large read/write capacity of the Shuttle computers permits complete program changes in flight. Also, simultaneous multiple computer usage introduces new failure modes unlike any considered in Apollo or Skylab.

RCS Digital Autopilot Evolution and Flight Regimes

Originally, a single orbital autopilot was built to control the Space Shuttle from the end of nominal main engine powered flight to the beginning of atmospheric re-entry. Control

Presented as Paper 82-1577 at the AIAA Guidance and Control, Atmospheric Flight Mechanics, and Astrodynamics Conference, San Diego, Calif., Aug. 9-11, 1982; submitted Aug. 27, 1982; revision received March 18, 1983.

*Member, Technical Staff.

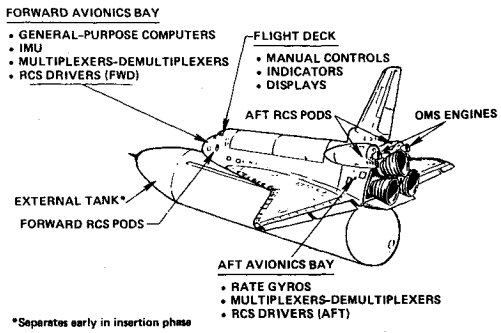

Fig. 1 Location of Orbiter systems.

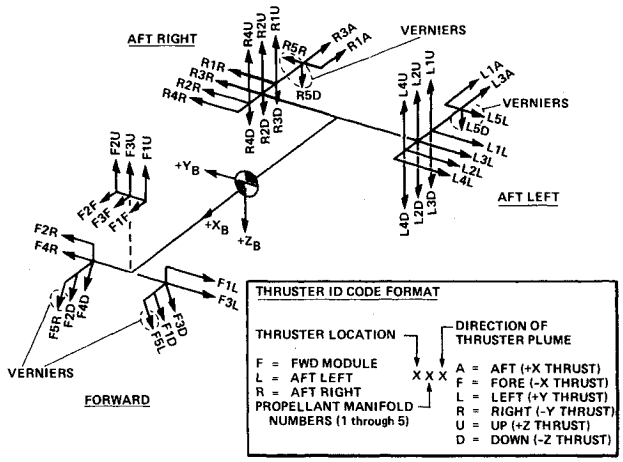

Fig. 2 RCS jet locations and plume directions.

requirements were included for the mated external tank (ET)/Orbiter flight after main engine cutoff (MECO), ET/Orbiter separation before orbital insertion, propellant burn-off for center-of-mass control during deorbit, vernier jet control, options that affect primary jet control (including capabilities for rendezvous, proximity operations, and propellant balancing between RCS tanks), and various automatic and manual control modes. The resulting autopilot size and execution time, when combined with software to control ascent with abort modes, entry, and maintenance of Orbiter system status while in orbit, led to the creation of two autopilots. The transition digital autopilot (DAP) was developed for the orbital insertion and deorbit flight phases when four synchronized general purpose computers (GPC) simultaneously process the DAP and its required input/output (I/O). The on-orbit DAP was developed for operation during orbit when no more than two GPCs simultaneously process the DAP with the system management assigned to a separate computer.

The transition DAP operates for only a few hours at each end of a flight, and it provides only a primary jet capability. It operates at the start of the orbital flight, spanning ET/Orbiter mated coast, ET separation, and guided orbital maneuvering system (OMS) burns to achieve stable orbit. It also operates at the end of the orbital flight, spanning the deorbit OMS burn.

The on-orbit DAP has all the in-orbit flight control capabilities, including vernier jet control and a state estimator, but it does not contain the logic required just after ascent and just before entry.

Sensory Data Sources

The transition DAP uses a fully powered sensor configuration that includes four rate gyro assemblies (RGAs) and three inertial measurement units (IMUs) to handle critical flight phases, such as OMS burns to attain and leave stable orbits. Each of the four synchronous GPCs has no more than one RGA or IMU latched to its I/O. Since no payload operations are to be done during transition phases, the attitude and rate control requirements are reduced, which allows a state estimator to be eliminated. Thus the RGAs give rates directly, and IMU updates give attitude at low rates to save GPC time.

To conserve power, the on-orbit DAP does not use RGAs and often may use only one IMU with the I/O latched to one GPC. To save GPC time, the IMUs are processed at an intermediate rate.

Sensory Data Processing Differences

The differences in the angular rate and attitude computations, and in the sensors generating the data, cause substantial qualitative differences in control system responses.

The transition DAP, without a state estimator, uses a complicated quad midvalue selection of RGAs at 12.5 Hz to measure angular rates, and integrates the rates for attitude change measurements. The IMUs are read at 1.04 Hz for RGA bias correction. The RGAs have a first-order roll-off analog filter at 50 rad/s. They are quantized at 0.04 deg/s in pitch and yaw and 0.08 deg/s in roll. They are also prone to noise; the 1-σ noise value of any one gyro is close to one quantum. The noise and quantization adversely affect the knowledge of vehicle state. Oscillating disturbances near or below the roll-off frequency can pass through to the DAP. Large low-frequency or constant disturbances are perceived quickly.

The on-orbit DAP uses a two-part state estimator. One part is an extrapolator run at 12.5 Hz, which uses feed-forward estimates of next cycle RCS-induced rate changes and estimates of undesired accelerations. The other part incorporates IMU measurements at 6.25 Hz. The estimator can be crudely viewed as a second-order roll-off filter with different gains for primary and vernier jets. Figure 3 shows the filter design and frequency response. There are some long time constants associated with the state estimator, as seen in the acceleration step response in Fig. 4. The low-pass nature of the filter protects the flight control system from the effects of high-frequency oscillating disturbances, but the slow response time corrupts the knowledge of vehicle states during the transient effects of large unidirectional disturbances.

Controller (Phase Plane and Jet Selection) Differences

The two DAPs vary in the details of the phase plane and jet selection construction. Design alterations due to RGA effects are made in the transition DAP and alterations are made due to undesired acceleration and vernier-jet cross coupling in the on-orbit DAP.

The transition DAP is exposed to substantial noise from the RGAs. To reduce jet chatter attributable to perceived attitude and rate changes as a result of measurement noise, the deadband switch curve is moved a little further from the origin when the attitude error magnitude has caused a corrective firing. The result protects against noise spikes.

The on-orbit DAP operates the vernier jets which cannot control rotation without significant coupling between axes. The phase-plane design (Fig. 5) includes a hysteretic region in rate and attitude to prevent large amounts of jet activity. A desired direction of rate and/or attitude correction is achievable in an uncommanded axis when in the hysteresis region while another axis is commanded. Since the vernier algorithm attempts to find the jets most closely aligned with the commanded direction (and more than one jet may contribute to that direction), the ability to specify a preference in uncommanded axes permits more fuel-efficient jet selection and fewer duty cycles. The phase-plane position is used to compute the preferences, which are never weighted more than 80% of an explicit command.

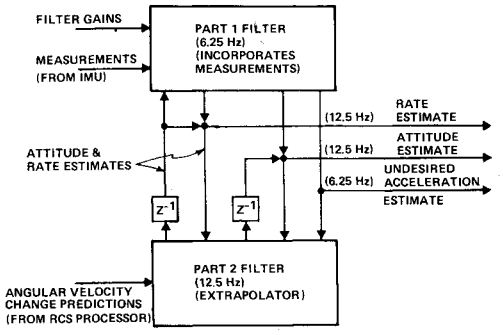

Fig. 3a State estimator design.

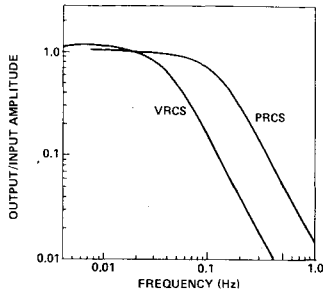

Fig. 3b State estimator frequency response for angular rate with vernier jets (VRCS) and primary jets (PRCS).

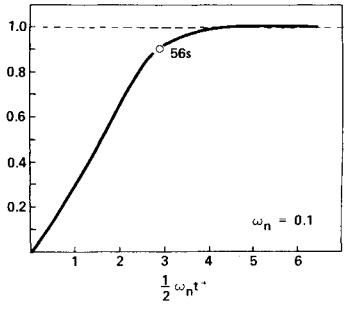

Fig. 4 State estimator acceleration step response.

The on-orbit phase plane also obtains an undesired acceleration estimate from the state estimator, which is used to bias the OFF-command switch curve in the hysteretic region to improve operational efficiency. The undesired acceleration is also considered when computing vernier off-axis preferences, since it can slightly influence the net phase-plane error correction obtained during the time the jets are firing.

Differences in Response to Jet Failures

The transition and on-orbit DAPs react in fundamentally different ways to failed jets. For OFF failures, the transition DAP has a benign response, since it will close the loop quickly to counteract the effect of reduced acceleration. For ON failures, the transition DAP will also close the loop quickly to counteract the effect of the undesired acceleration, but it will cause jet cycling to control the disturbance torque. The jet chattering can induce Orbiter flex at the RGA location. This oscillating motion is perceived by the DAP, and under conditions tested early in the program, including very large flex-model tolerances, it can cause extra jet cycling activity with an associated propellant usage penalty. The simulated effect is seen in an example phase-plane plot in Fig. 6. The high initial rate causes the phase-plane trajectory to land in the drift channel. When the jet turns off, the Orbiter structure unloading causes enough ringing at the RGA location to

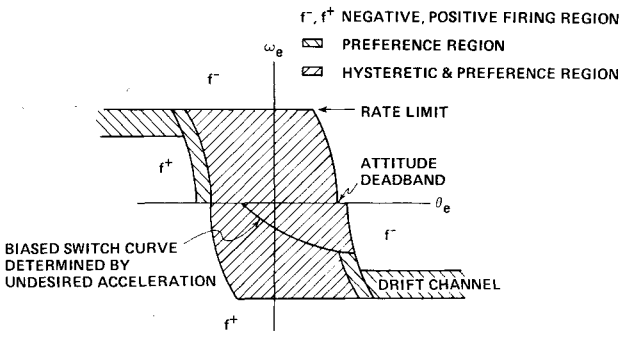

Fig. 5 On-orbit phase plane regions.

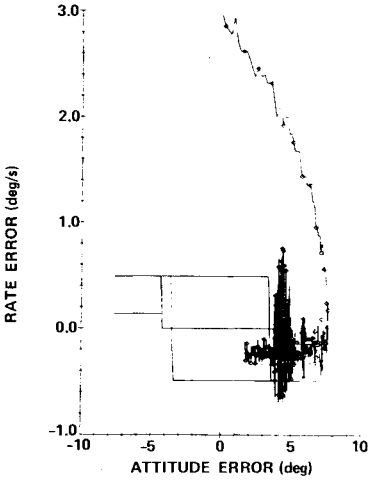

Fig. 6 Transition phase plane, roll axis with jet F3L failed on. Bending mode slopes are increased 25% of the second natural frequency value. Modal amplitudes are increased 30% each. Damping is 0.01. Frequencies are reduced 10%. RGA analog roll-off filter break frequency is increased to 85 rad/s. Rate limits are 0.5 deg/s. Initial rates are 3 deg/s/axis.

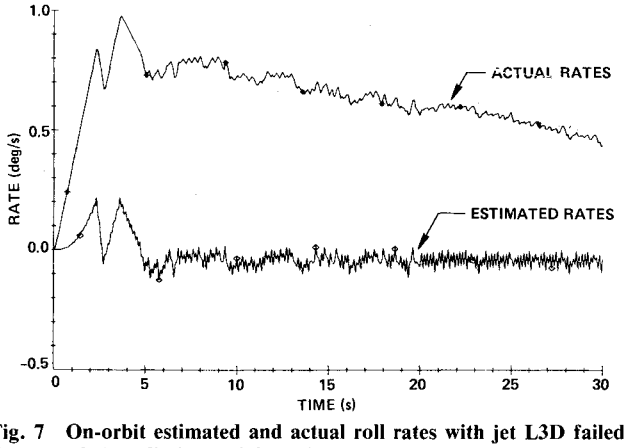

Fig. 7 On-orbit estimated and actual roll rates with jet L3D failed ON at $t = 0$. Rate limits are 0.2 deg/s.

induce the measured rate to go out of the channel upper boundary periodically and cause short one-sided thruster firings. These pump up the flexure enough to cause the measured rates to go outside of the channel on both sides, and two-sided RCS firings result. Since the rigid-body rates cause a drift toward a smaller attitude error during the entire jet chattering time, the error eventually becomes small enough to land inside the deadband, which requires that the peak-to-peak amplitude of the flexure exceed double the rate limit to sustain jet cycling. This is not the case; the bending damps out. Flight data indicate that damping on the Orbiter is higher

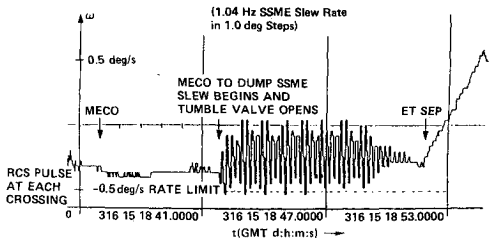

Fig. 8 STS-2 mated coast pitch rates from 12.5 Hz sampled data downlisted by the backup computer. The timescale is in Greenwich mean time (GMT).

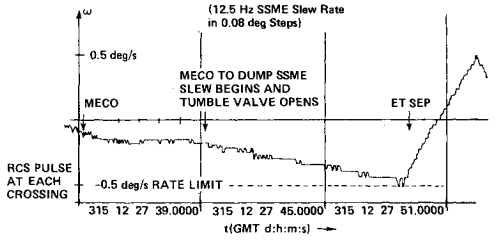

Fig. 9 STS-5 mated coast pitch rates from 12.5 Hz sampled data downlisted by the backup computer.

than that modeled before space transportation system flight 1 (STS-1). High damping, combined with large off-nominal tolerances required to induce flex effects, implies that chances are extremely low that the DAP-interaction-with-flexure problem will occur in flight. As extra protection, digital second-order roll-off filters were incorporated on STS-5.

The on-orbit DAP, with a state estimator bandwidth lower than the Orbiter natural frequencies, will not see Orbiter flex effects from cycling jets. However, both ON-and OFF-failed thrusters cause large errors in feed-forward estimates of expected RCS rate changes. The sluggishness of the undesired acceleration estimate causes an erroneous rate estimate transient at the time of reconfiguration of an ON failure. Poor rate control can result during the failure due to corrupted knowledge of the state, and the rate estimate transient causes fuel usage penalities. Figure 7 shows divergence of real rate and estimated rate for a simulated ON failure (without reconfiguration). The actual rate is a result of commanded jet activity and the failed-ON jet. The estimated rate is a composite of feed-forward rate-change data and the integrated effect of an undesired acceleration. (Refer to Fig. 4 for the convergence time scale.) The DAP controls the estimated rates to within the rate limits but the attitude drifts out of the deadband. Many deadband-induced short jet firings result, based on IMU attitude data, even though actual rates are well outside the rate limit and fewer, longer, rate-damping jet pulses would be more efficient and provide better control. While the performance is not close to optimum for the large disturbance example, it is the result of a conscious tradeoff to assure superior state estimator accuracy and DAP performance when small, slowly varying disturbances prevail throughout on-orbit operation.

Response to In-Flight Disturbances

Significant disturbances not attributable to failures, which cause unusual DAP responses as a result of specific DAP design characteristics, have been observed in flight. These include Space Shuttle main engine (SSME) slewing, to which the transition DAP is sensitive (due to the high RGA bandwidth), and vernier jet plume impingement to which the on-orbit DAP is sensitive due to the long state estimator transient response time.

During orbital insertion while still attached to the ET, the SSMEs are moved from the MECO position to a position where residual SSME propellant is dumped. On STS 1-4, the SSME actuators stepped discretely at 1.04 Hz with large force transients, which excited the structure and caused substantial flexure rates at the RGA locations. During the motion, the ET tumble-valve flow accelerates the mated configuration. Cyclic jet firng, where the rigid-body effects are mostly negated by the tumble-valve torque, was the response seen on STS-2. The amplitude of the rate excitation at the RGAs is seen in Fig. 8. (The data shown are actually from the backup computer, because of higher resolution from downlist data. The backup computer selects midvalue data from three RGAs. The controlling primary computer software selects midvalue data from four RGAs. The controlling software saw rate-limit and null-rate crossings with resulting ON or OFF jet commands on a few occasions when the plot seems to show that they were not quite achieved.) The peak rate magnitudes are large partly because the flex frequencies are so low that the RGA analog-roll-off filter provides very little attenuation. Before STS-1, a possible SSME slew/DAP interaction justified changing the mated coast rate limits from 0.3 deg/s, which the ET separation dynamics community preferred, to 0.5 deg/s. Flight data show that the decision was appropriate. With STS-5, SSME slewing procedures changed, resolving the mated coast flexure/autopilot interaction as seen in Fig. 9.

Disturbances also can result from nominal thruster performance. STS-1 data showed that a substantial unmodeled plume impingement effect occurs on the aft down-firing vernier jets due to wing, fuselage, and body-flap interference. The net thrust vector is reduced and rotated when compared to preflight models. Since the state estimator relies on feed-forward estimates of expected rate changes and those estimates for STS-1 were based on preflight models, a large error existed in the predictions. The system response was an overshoot in the rate-estimate and, consequently, occasional commanded jet reversals occurred to correct the fictitious excess rate change. The response caused extra jet duty cycles and propellant consumption. This effect is another manifestation of the slow state estimator response to disturbances.

GPC Failure Effects

All analog sensors and actuators on the Shuttle communicate to the GPCs through multiplexor/demultiplexor (MDM) devices, each of which can be I/O linked to one GPC at a time. Eight flight-critical MDMs operate in orbital flight. Four MDMs interface with forward Orbiter systems, and four interface with aft systems. No sensor or actuator used in orbital flight has I/O through more than one MDM. The MDMs are grouped in four strings. Each string has one forward and one aft system MDM. All systems on one string must be tied to one GPC, but one GPC can communicate with an arbitrary number of strings. A GPC failure can, therefore, freeze or cause spurious operations of many systems at once. Table 1 shows an example of some of the systems affected on a single string.

The number of GPCs nominally processing the active autopilot, and therefore the number of systems tied to any one GPC, determines the influence of a single GPC failure. Also, the permissible crew actions to clear failure effects and reconfigure affected systems vary with flight phase; thus they alter the effects of otherwise identical failures.

In the transition phases, four GPCs nominally process the DAP with one string per GPC. In this configuration, any single GPC failure affects no more than one of each type of sensor (all of which are nominally powered up). Also, one primary jet RCS manifold per pod is disabled by the failure. (No more than one jet per cluster is affected.) This keeps the system within the fail-operational margin. Also, one pair of OMS actuators on one engine can freeze, but RCS wraparound control can correct the resulting disturbances. During orbit insertion, MDM recycling to clear frozen residual I/O values can be performed, and a string transfer can be done to recover the lost capability when the flight-

Table 1 String 3 lashup of systems used in orbital flight

MDM	System
FF3	Jets on RCS manifolds F4, F5
	IMU No. 3
	Selection of right secondary OMS gimbals
FA3	Jets on manifolds R2, L2
	RGA No. 3
	Right OMS secondary gimbals—status and control
	Left OMS arm-A
	Left OMS chamber pressure indication
	Left OMS propellant valve 2-A
	Right OMS propellant valve 1-A
	Left and right RCS tank isolation valves 3/4/5-A
	Left and right RCS cross-feed valves 1/2-A

critical maneuvers are completed. In deorbit, if flight-critical maneuvers are imminent, MDM recycling and GPC restringing may be forbidden, which would cause all systems on the affected string to be inactive for the duration of the flight.

In the on-orbit phase, no more than two GPCs process the autopilot simultaneously; this requires two strings per GPC. If only one IMU is running, one GPC failure can cause loss of all attitude data. In any case, one GPC removes about half the thrusters from operation and in the case of the vernier jets eliminates attitude control authority. While the resulting system status from a GPC failure is unacceptable, critical maneuvers while on-orbit can be postponed without crew risk, if necessary, to allow string reconfiguration and sensory-data and control capability recovery. All strings could be transferred to one GPC, or a second good GPC, not previously applied to autopilot processing, could be transferred to that function.

Dual GPC failures affect DAP performance more adversely. In the on-orbit phase, all system control can be lost until strings and GPCs are reconfigured. In transition, at least one OMS engine can freeze and/or shut down, making RCS attitude control necessary during critical OMS burns when the RCS control authority is also severely degraded from loss of half the thrusters' capabilities. Prompt crew reconfiguration is desirable but not possible in the most critical burn phases (i.e., if the deorbit burn has progressed past the commit to entry point). However, in the most adverse circumstances, a transfer to the backup computer system recovers all strings and automatically assures control, and operational capability is not reduced substantially.

Conclusions

The two DAPs that have evolved from the original orbital autopilot concept have very different responses to failures and disturbances as a result of features uniquely required during the flight phases in which each operates.

The transition reaction control system digital autopilot attitude control relies on rate gyro data inputs with first-order analog filtering. Without additional filtering, this resulted in a system with rapid response to angular rate changes from any source, but also had the potential of forced RCS oscillations due to resonant flexure at the rate gyro locations. Simulated worst-case tolerance jet failures in some instances resulted in temporary two-sided jet firings.Also, main engine slew-induced structural resonance caused jet pulsing during mated coast prior to STS-5. With the introduction of second-order digital roll-off filters and a modified main engine slewing procedure on STS-5, most of these potential Orbiter flexure/autopilot interactions have been precluded.

The on-orbit reaction control system digital autopilot has substantial Orbiter flex attenuation due to the inherent filter properties of the state estimator designed to use inertial measurement unit data. Feed-forward rate change estimates from the jet selection algorithm to the state estimator supplement the attitude data from the sensor platform to reduce the effects of the long estimator time constants. With large errors in the feed-forward due to unmodeled effects, rate control can be degraded and propellant consumption increased since the unmodeled errors are subject to the estimator lags. Failed-ON jets introduce transients at failure initiation and reconfiguration. Improper preflight jet plume impingement models introduce transients at each commanded thruster firing. Both are corrected by undesired acceleration estimate feedback into rate estimation. The potential performance penalty is accepted as a design tradeoff, yielding good estimator and reaction control system performance with disturbances having long time constants compared to estimator lags (e.g., gravity gradient and aerodynamic torques).

Computer failure effects vary in different flight phases because of differing electronic stringing rules and failure reconfiguration procedures. The transition autopilot operates with four redundant computers, making it impossible for one failure to remove an entire jet cluster or gimbal control of both orbital maneuvering engines. Loss of the gimbal actuators on one engine can result, however, and require thruster attitude control wraparound. Dual computer failures can badly degrade transition performance, with sluggish wraparound thruster attitude control due to loss of couple capability, and the potential of a maneuvering engine shutdown. In severe situations, the backup computer can reinstate systems lost due to primary computer failures.The on-orbit autopilot usually operates with two redundant computers, with one failure being like two in transition. However, few on-orbit operations are time-critical enough to preclude an opportunity to effect recovery of the failed or unused computer.

Acknowledgments

The material presented in this paper is the result of a multiyear collaborative effort including flight control/autopilot expert staff at The Charles Stark Draper Laboratory, Inc., NASA, Rockwell International, Lockheed, and Honeywell. At the Johnson Space Center, E. T. Kubiak manages the transition autopilot design and K. L. Lindsay manages the on-orbit autopilot design.

This paper was prepared by the Charles Stark Draper Laboratory, Inc., under Contract NAS 9-16023 with the National Aeronautics and Space Administration. The work was supported by NASA Johnson Space Center.

Publication of this report does not constitute approval by NASA of findings or conclusions contained herein. It is published for the exchange and stimulation of ideas.

References

[1]NASA, Requirements/Defintion Document, Flight Control, Part 1, Configuration, Performance and Functional Requirements, SD72-SH-0105-1, Book 2-1B, Rev. B, Dec. 5, 1980.

[2]NASA, Space Shuttle Orbiter Operational Level C Functional Subsystem Software Requirements, Guidance, Navigation, and Control, Part C, Flight Control, Orbit DAP, STS-81-0009, May 1, 1981.

[3]Johnson, M. S. and Giller, D. R., *The Software Effort, Apollo Guidance, Navigation, and Control—MIT's Role in Project Apollo,* The Charles Stark Draper Lab. Rept. R-700, Vol. V, July 1971.

[4]Battin, R. H., "Attitude Control of the Apollo Spacecraft," AGARD Lecture Series No. 45 on Attitude Stabilization of Satellites in Orbit, Sept. 1971.

[5]Turnbull, J. F., "Digital Autopilot for the Skylab Orbital Assembly," *Navigation: Journal of the Institute of Navigation,* Vol. 19, No. 3, Fall 1972, pp. 281-289.