

A SHORT HISTORY OF THE SPACE SHUTTLE ORBIT FLIGHT CONTROL SYSTEM DEVELOPMENT AND APPLICATION EVOLUTION

Philip D. Hattis¹

The Space Shuttle Orbit Flight Control System (OFCS) provided all Guidance, Navigation, and Control capabilities from external tank separation during ascent until entry interface during return. Its development began in earnest in 1975 with first orbital flight application in 1981. The many operational requirements for the Shuttle and its unique vehicle design characteristics demanded specialized algorithms and numerous functional features despite very limited computer memory and speed by contemporary standards. This paper provides a short history of the Shuttle OFCS development and orbital application evolution.

THE OFCS GENESIS

In 1961 the Massachusetts Institute of Technology (MIT) Instrumentation Laboratory (IL) began work on the Guidance, Navigation, and Control (GN&C) system development for the Apollo program. This work included creation of applicable digital computers, GN&C software (SW), and navigation hardware (HW). All mission phases after separation from the launch vehicle were addressed utilizing a fly-by-wire approach where the digital computer flew the spacecraft, and the crew interacted through the computer. New digital approaches to in-space flight control were applied including for steerable orbit-change rocket engines as well as for Reaction Control System (RCS) thrusters used for attitude control, fine trajectory management, and rendezvous/docking between mission vehicles. Navigation was accomplished using digitally filtered data collected from high-precision Inertial Measurement Units (IMUs).

As the Apollo program wound down, the MIT IL supported NASA studies regarding how to design GN&C for a reusable Space Shuttle (henceforth referred to just as Shuttle). Reapplication of some GN&C design features from Apollo was proposed including fly-by-wire design, phase-plane control for RCS attitude control, closed-loop Thrust Vector Control (TVC) for orbital changes using Orbital Maneuvering System (OMS) engines, and Kalman filter-based navigation state estimation. A new design feature anticipated for the Shuttle was fault-tolerant digital computers operating the GN&C system.

In early 1972 the Nixon Administration decided to proceed with Shuttle development.² The baseline design had a winged, reusable orbiter, partly reusable solid rocket boosters, and a disposable External Tank (ET) to carry liquid oxygen and liquid hydrogen for main engine use during ascent to orbit. On July 25, 1972 the International Space Division of Rockwell received the

¹ Laboratory Technical Staff, Complexity Solutions Division, Draper, 555 Technology Square, Cambridge, MA, 02139.

² From <http://www.spaceline.org/rocketsum/shuttle-program.html>


Copyright © 2016 by the C.S. Draper Laboratory, Inc. This paper is released for publication to the American Astronautical Society in all forms.

contract to develop and manufacture the Orbiter, as well as to manage overall vehicle integration. The initial Rockwell contract included the orbiter GN&C development. However, in 1975 the Johnson Space Center (JSC) parsed off development of the Shuttle OFCS from Rockwell and awarded that work to the MIT IL successor, the Charles Stark Draper Laboratory, Inc. (subsequently referred to as Draper). Draper had separated from MIT in October 1973 because of concerns that MIT would restrict Draper's national security work due to Vietnam-era student protests.

Meanwhile, IBM was awarded a contract to develop the fault-tolerant computers for the Shuttle. This included the computer HW, the SW to manage fault-tolerant computer features, and the implementation of flight SW loads for the computers. Consequently, Draper's job was to develop an OFCS design that could fly the Shuttle from ET separation (just prior to orbit insertion) through entry interface (during final descent) using vehicle systems indicated in Figure 1³. Draper had to vet the OFCS design with high-fidelity simulations and then document the design with detailed requirements for how IBM was to implement the flight SW. Those OFCS design requirements were covered by a series of Functional Subsystem Software Requirements documents issued under Rockwell cover. When the IBM flight SW implementations were loaded into piloted Hardware-in-the-Loop (HWIL) simulations at Rockwell and JSC, the applied verification test procedures were based on Draper test guidelines. Cross-contractor teams, including on-going Draper participation, prepared the test procedures and reviewed test results.

FORWARD AVIONICS BAY

- GENERAL-PURPOSE COMPUTERS
- IMU
- MULTIPLEXERS-DEMULPLEXERS
- RCS DRIVERS (FWD)



*Separates early in insertion phase

Figure 1. Shuttle Features and Systems Impacting the OFCS Design.

³ Unless otherwise noted, this and other figures in the paper came from "A Review of the Space Shuttle Orbital Flight Control System," by Philip Hattis, Draper Report CSDL-P-1786, October 31, 1983. The original source of this image is NASA document SD72-SH-105-1 Book2-1B.

EARLY OFCS DESIGN CHALLENGES

The original fault-tolerant computer architecture for the Shuttle was four redundant General Purpose Computers (GPCs) to fly the entire mission with a single SW load. IBM developed that AP-101 computer with 64K words (256K bytes equivalent) of magnetic core memory with a speed of about 400K instructions/second. This version of the GPC was used for the Approach and Landing Test (ALT) series that carried the Enterprise orbiter in 1977 atop a modified 747, initially for captive tests, and then for release at altitude to glide to a runway at Edwards Air Force Base. However, well before the ALT flights it was apparent that the GPCs had insufficient capacity to handle the SW needed for the entire Shuttle mission. Consequently, in 1976 a series of meetings known as the Haise Scrub (led by astronaut Fred Haise of Apollo 13 fame) was held to reduce the flight SW memory and throughput demands to levels that could be accommodated by the GPCs. This was accomplished in part by limiting SW functionality to only that needed by mission phase. The outcome was a number of major changes to the GPCs, the SW design, and the in-flight handling of SW loads. The GPC was expanded to 104K words (416K bytes equivalent) of magnetic core memory. Three separate SW sets were created: 1) An ascent/entry load that took the Shuttle through orbit insertion, and subsequently from deorbit through entry and landing; 2) an on-orbit load that flew the Shuttle through its on-orbit mission; 3) a System Management (SM) load that performed payload-related functions while the vehicle was in orbit. To enable this SW functionality split by mission phase required that at least some of the GPCs have their SW overwritten in flight. Two redundant tape drives were added on board the Shuttle to enable in-flight GPC SW load changes. Because in-flight SW load changes posed new risks, the following strategy was applied that included risk mitigation measures:

- The nominal ascent/entry flight control load would operate with four redundant, primary GPCs.
- The on-orbit flight control load would operate with 2 redundant GPCs, with one separated GPC used in parallel for SM. These SW loads would be installed in-flight on three of the primary GPCs used during ascent, and would be reloaded with the ascent/entry SW loads before deorbit.
- One primary GPC with the ascent/entry load was shut down while on-orbit with the ascent/entry SW retained. This way, if the tape drives failed to reload SW before deorbit, then at least one GPC had the right SW already loaded for entry/landing.

The reduced OFCS GPC redundancy on-orbit was accepted since the environment was more forgiving – there would be time in orbit to reconfigure and recover from the effects of GPC failures. However, because failures were very time-critical during ascent and entry, the quad-redundant GPCs could assure essential capability maintenance and vehicle survival despite failures.

Both this GPC implementation and SW management strategy were used on all Shuttle flights for almost one decade. In 1990 NASA began a retrofit of the orbiter avionics which included new AP-101S GPCs with 1M bytes of semiconductor memory and 1.2 M instructions/second processing speed which served as the Shuttle GPC for the remainder of flight operations (through 2011). Along with the AP-101S upgrade, solid state memory transfers for in-flight SW load changes replaced tape drive use.

Fault-tolerant computing was a new discipline when the Shuttle was being developed, and there were concerns that common faults in the GPC redundant set management SW across the GPCs could cause a collapse of the redundant, “primary” GPC operations in flight. Consequently, a fifth, backup GPC was added. This GPC only supported ascent/entry/landing operations, and never underwent SW load changes in flight. Its SW functionality emulated the GN&C features of the

primary ascent/entry GPCs, but it was independently coded. If needed, the backup GPC was able to be engaged upon crew command at any time during ascent and entry/landing, but that never occurred during any Shuttle mission.

Because the AP-101 throughput was so limited, much attention of the Haise Scrub was applied to limiting the processing rates of computationally demanding functionality. Fundamental differences in the ascent/entry and on-orbit dynamics and objectives enabled trades of what GN&C instrumentation was needed, and its processing rates. This resulted in two separate versions of GN&C for in-space operations.

- 1) A Transition Digital Autopilot (TDAP): TDAP was in the ascent/entry SW load to fly the orbiter from Space Shuttle Main Engine (SSME) shutdown through completion of orbital insertion. This included separation from the ET, use of a pair of OMS engines with 26,690 N thrust each for orbit insertion completion, and use of 38 Primary RCS (PRCS) jets with 3,870 N thrust each for attitude control and fine orbital adjustments. The TDAP was re-invoked at the end of orbital flight when the ascent/entry SW load was restored. It used the OMS to deorbit and the RCS to control the vehicle attitude into the atmosphere when the entry autopilot took over which initially using the PRCS jets and eventually transitioned to aerodynamic surfaces for control. TDAP angular rate data was derived at 12.5 Hz using a quad-redundant set of Rate Gyro Assemblies (RGAs) that had the high dynamic range needed during ascent and entry. Attitude was determined by mid-value selection from a triplex set of IMUs with data derived at 1.04 Hz and with RGA data used to extrapolate attitude changes between IMU updates. IMU processing rates were kept low since IMU data reduction was a significant burden on the GPCs.
- 2) An On-Orbit Digital Autopilot (OODAP): OODAP was in the on-orbit SW load to fly the orbiter through its on-orbit mission. Over the Shuttle program duration, on-orbit missions included complex attached payload operations, payload deployments, and rendezvous/docking missions with tight orbit and attitude management requirements. The OMS was used for sizable orbit changes, and the PRCS was used for small orbit changes. Much of the time while on-orbit, the Vernier RCS (VRCS) system was applied for precision attitude control using 6 jets each with 107 N thrust. The RGAs were turned off for the entire on-orbit flight phase to conserve power as were two of the IMUs. One IMU with data read and processed at 6.25 Hz was used for attitude determination with a Kalman filter-based state estimator to process IMU data into estimated angular rates. To enhance the performance of the state estimator, and to reduce transients in the estimates, expected rate changes from thruster firings were calculated and fed forward into the estimator.

The GPC SW load functionality distribution across flight phases following the Haise Scrub was as shown in Figure 2. The TDAP performed ET separation, and subsequently two OMS burns to circularize the orbit. During on-orbit flight the OODAP flew the vehicle with payload bay doors open both to enable mission/payload operations and to utilize radiators inside the doors to dissipate vehicle-derived heat into space. The TDAP then performed deorbit and set up entry interface.




Figure 2. GN&C SW Functionality Separation by Flight Phase Was Driven by the 104K Word GPC Memory Limit on Initial Shuttle Orbital Flights.

Differences in angular state processing between the TDAP and OODAP are shown in Figure 3. The IMU processing burden on the GPCs was reduced during Transition by using classically filtered data from RGAs to derive attitude and angular rate readings. Kalman-filter-based state estimator functionality was only available during the on-orbit phase, without access to any RGA data.

SOME OFCS ALGORITHM FEATURES

The OFCS role was to manage the orbit of the Shuttle during exo-atmospheric flight. This was done with the RCS and OMS engines. The location of the Shuttle orbital propulsion components could not be symmetric in all axes as is approximated in most spacecraft. This was because the orbiter's configuration was much like an airplane. Figure 4⁴ shows the Shuttle RCS jet thrust directions. Figure 1 shows the OMS positions. The following factors and thruster locations impacted the OFCS algorithm design.

- ET separation: The orbiter had to separate from the ET after SSME shutdown during ascent without risk of re-contact (that could damage protective thermal tiles on the orbiter). This was done firing all ten down-firing PRCS jets (with IDs ending with the letter D), but with closed-loop attitude control cycling jet firings to maintain safe orbiter/ET relative orientation during separation.

⁴ Also from "A Review of the Space Shuttle Orbital Flight Control System," by Philip Hattis, Draper Report CSDL-P-1786, October 31, 1983, with the original source of this image being NASA document STS 81-009.




Figure 3. Transition and On-Orbit IMU Data Processing Differences Were Driven by Limitations in GPC Throughput Capacity.




Figure 4. Shuttle Orbiter RCS Jet Locations and Plume Directions with VRCS Jets Identified Using Dashed Circles.

- Orbit circularization, major orbit adjustment, and deorbit: This was done with the OMS engines. Each OMS engine had two-axis gimbals and actuators to enable steering the thrust direction. Before an orbit change, the PRCS would rotate the vehicle to the desired OMS burn direction. During a burn, a TVC algorithm would keep the thrust vector pointing near the orbiter center of mass while managing the vehicle attitude control in at least two axes. With a one-engine OMS burn, the PRCS would handle vehicle roll control (around the orbiter longitudinal axis). With two-engine OMS burns the OMS TVC would use some of the gimbal pitch authority for differential motion to manage the vehicle roll orientation.
- Attitude and rate determination: This was accomplished as shown in Figure 3.
- Attitude control: A variety of manual, automatic, and hybrid attitude control algorithms could be selected by the crew. Manual control through a crew stick could pulse jets to rotate in a given control axis, or with a stick rotation to its stop, would sustain a jet firing in a given rotation axis until the stick was released from its stop. A hybrid mode fired the jets in the direction of crew stick motion until a designated discrete rate was realized, and the rate would be maintained closed-loop until the crew released the stick. Automatic modes allowed crew designation of a desired change in vehicle attitude or rotation rate in either inertial or local vertical coordinates, starting at a specified time, using thrusters to perform the commanded maneuver. In the TDAP all these modes required use of the PRCS jets. In the OODAP these modes could be executed using either the PRCS or VRCS jets. Use of the VRCS would be designated when desired maneuver rates were low (which could reduce vehicle propellant usage) and/or when attitude control precision was needed. All closed-loop attitude control used a per-body-axis phase plane algorithm which had special features unique to each flight phase:
 - Special TDAP phase plane feature. The RGAs used for angular rate determination during transition flight were analog devices connected to Multiplexor-Demultiplexor (MDM) cards to read data. The MDM technology used on the Shuttle had very limited bits available for data reads. Based on the dynamic range requirements for atmospheric flight, the granularity of the pitch/yaw rate read by the RGAs was .04 deg/s, and in the roll axis it was .08 deg/s. Furthermore noise on the RGA MDM card had a $1-\sigma$ value equal to an MDM card 1-bit quanta. This posed a risk of fuel-wasting jet chatter when the vehicle attitude was near a phase plane attitude-error switch line, since the TDAP could make a determination that the vehicle was outside the switch line even after a jet firing had actually brought it back in. This problem was mitigated as shown in Figure 5, by moving the switch line out by the expected attitude change over 80 ms due to an MDM 1-quanta change whenever a jet fired on that line. The line extension was retracted once the estimated attitude error moved well away from that line.
 - OODAP phase plane feature: During on-orbit flight closed-loop attitude control could be maintained for hours at a time while exposed to low-level accelerations from effects such as gravity gradients, aerodynamic torques due to the thin upper atmosphere, and gradual outgassing from the vehicle and payload(s). Without account for these small acceleration effects, the vehicle orientation could stay near a phase plane switch line with a short limit cycle between firings. To enable longer limit cycles (and lower propellant consumption), a disturbance acceleration estimation feature was included in the state estimator, with a 56-second step-response time constant. A current estimate of the disturbance acceleration was used to bias a phase plane rate error switch line as shown in Figure 6.




Figure 5. The TDAP Phase Plane Had a Moving Attitude Error Switch Line to Reduce Jet Chatter Due to RGA-Related Rate Signal Noise.




Figure 6. The OODAP Phase Plane Had a Biased Switch Line Based on Estimated Disturbance Accelerations to Prevent Short Limit Cycles.

- Position control: For the TDAP, PRCS translation control was mostly used to trim-out OMS burn errors or to make small orbit adjustments. This was done manually. However, in addition to the automatic ET separation discussed above, the TDAP had an automatic PRCS burn option that would utilize the aft-firing jets (with a designation ending in A) located along the longitudinal access. This x-axis closed-loop burn capability was a backup for deorbit in case both OMS engines were deemed unusable. For the OODAP, all PRCS position control was manual, with two unique position control modes along the z- axis available during rendezvous and docking (with the top of the orbiter being the rendezvous approach direction since the crew could look out overhead cockpit windows while flying the vehicle using interfaces in the aft cockpit). One unique OODAP position control mode allowed all nine up-firing PRCS thrusters (with the designation ending in U) to be used together for emergency braking in the rendezvous direction (which was never applied on actual missions). The other unique OODAP position control mode used the two forward-firing and two aft-firing jets together (the jets with designations ending in F and A). While these jets were located along the longitudinal x-axis, they were canted a little upward along the z-axis. By firing them together, a low-level z-axis braking acceleration was applied in the direction of rendezvous. This “low-z” position control mode was used during all Shuttle rendezvous missions.
- GPC processing time challenges: With the limited GPC processing rates, some mathematical operations were burdensome. A single dot product operation on two 3-component vectors took about 0.1 ms on the AP-101 used on initial flights. That was relevant to the OFCS since the VRCS jet selection algorithm was based on the value of dot products between a desired rate change command and the expected acceleration vector from a jet firing. The algorithm could select one, two, or three jets at a time to generate a propellant-efficient approximation of the desired rotation vector. Much thought went into designing the algorithm to achieve its propellant efficiency goal with a minimum number of dot product operations per thruster selection cycle. Analysis of the ideal thruster selection solution results combined with heuristic studies resulted in an algorithm that required one dot product per thruster per OFCS cycle (6 dot products total) without direct consideration of thruster combination dot products. This was accomplished by requiring that specific, relative dot product contribution thresholds be met for second and third jet firings based on comparison of the per-jet-derived dot product values. With assignment of the right, heuristically determined threshold values, performance came close to the ideal (and more processing-intensive) solution.
- Fault-tolerance features: A feature of the primary GPC fault-tolerant architecture was that only one GPC could actually read data from and/or send commands to a given component. Consequently, while every OFCS GPC shared and processed data from all active navigation sensors, and derived commands for all systems managed by the OFCS, each GPC only operated a subset of those components. This was called an electronic string. If a GPC failed, all components on that GPC’s string became inoperative. An inoperative string could be reassigned manually by the crew to a healthy GPC to be restored to operation, but only during quiescent mission phases. Shuttle requirements demanded that the OFCS remain fully functional after any one failure, and remain safe (at least able to complete ascent/entry) after any two failures. This made the fault-tolerance design of the OFCS a complicated challenge in the face of GPC failures. An example design concern is illustrated by Figure 7. To turn on an OMS engine, at least one valve set in each 1A/1B and 2A/2B pair must open. Likewise, to shut down an engine, both valve sets on one of the 1A/1B or 2A/2B pairs must close. Also, to verify that an engine was actually firing, its chamber pressure sensor had to be monitored. The figure 7 diagram tells us the following:

- If strings 1 and 3 became inoperative, then only the left engine was usable and its pressure sensor could not be read. Likewise if strings 1 and 4 became inoperative then only the right engine could be used and its pressure sensor could not be read.
- If strings 3 and 4 became inoperative then the pressure cannot be read from either engine.




Figure 7. Electronic Stringing of OMS Propellant Valves and Pressure Sensors.

Based on these considerations, there were some dual-string losses that prevented safe use of OMS engines, which motivated the previously mentioned mode for using the aft-firing PRCS jets as a backup for deorbit (to be fail-safe after the second failure). However, under these circumstances two of the four aft firing jets were also inoperative, so the deorbit had to be accomplished with a much lower thrust, longer burn, with additional PRCS propellant used to maintain vehicle orientation during the burn. Properly strung OMS/RCS propellant cross feed valves were provided to enable the aft PRCS jets to access OMS propellant tanks, providing propellant needed for this contingency.

- **Simultaneous PRCS jet usage constraints:** Some of the PRCS feed lines in the aft vehicle were quite long. During PRCS ground-test firings it was found that opening thruster valves caused expansion waves, and closing valves cause compression waves that could take 10s of milliseconds to dissipate (like water hammer effects in radiator lines). It was also found that simultaneous PRCS firings involving common feed lines could intensify the pressure fluctuations that would be further amplified by intersection with reflecting waves. Sufficiently amplified, the waves could cause auto-detonation in the fuel in the feed lines. Two types of constraints were imposed on the OFCS design to mitigate this issue: 1) Minimum thruster cycle times were increased from an original goal of 40 ms to 80 ms to allow more time for

pressure waves to dissipate; and 2) a maximum number of simultaneous PRCS thruster firings on a common feedline was set, resulting in algorithm changes that adjusted thruster selection criteria when combined rotation/translation axis firing commands were simultaneously received.

IN-FLIGHT INTEGRATED SYSTEM DYNAMICS SURPRISES ON EARLY MISSIONS

Some dynamics surprises were experienced in early Shuttle missions that challenged the proper operation of the OFCS and resulted in subsequent design revisions to mitigate resulting issues. The following are two significant examples.

- **SSME stowing motion/ET vent combined risk to ET separation:** After the SSME powered flight was completed during ascent, at Main Engine Cutoff (MECO), the engines needed to be stowed to their reentry position as soon as possible. The slewing of the engines was done hydraulically, with the pressure provided by pumps driven by Auxiliary Power Units (APUs). The APUs had limited hydrazine fuel supplies, so needed to be turned off as soon as possible since they (and their fuel) would be needed again during entry to power aerosurface hydraulics. Consequently the SSME slew to an entry stow position occurred in the short interval between their shutdown and ET separation. At the same time a valve on the ET was opened to vent residual propellant. This was done to tumble the tank during entry to facilitate its complete breakup over the Indian Ocean, with the valve opening having to occur while the orbiter was still attached to provide power. On the initial Shuttle missions, the SSME stow maneuvers occurred in a series of 1 deg gimbal steps done at 1.04 Hz. While on the ET, the Shuttle had a pitch rocking vibration mode due to ET flexure, with a frequency near 4 Hz. Ground vibration tests had been performed with full-scale Shuttle and ET mockups to determine that rocking frequency with a derived uncertainty bound of about 15%. Analysis indicated that the SSME slew and tank venting effects could not significantly excite the rocking vibration anywhere within the defined frequency bounds for that vibration mode. However, flight experience on the STS-2 flight indicated that the pitch rocking vibration frequency was actually more than 20% off from the predicted value. This resulted in the SSME stow slew step frequency being almost an exact 4:1 subharmonic of the vibration mode. As seen in Figure 8, orbiter thruster firings before ET separation in response to ET venting were influenced by the SSME slewing, exciting the pitch rocking vibration. The reason this unexpected dynamic response mattered was that the ET separation could not safely occur if the orbiter pitch rate at the moment of separation exceeded a magnitude of 0.5 deg/s. The response in Figure 8 shows that bound was reached until just short of separation. While the problem was the result of a complex interaction, the solution was simple. The SSME slew to stow position was subsequently changed to occur in smaller 80 ms steps, and no similar excitation occurred on any subsequent flights.




Figure 8. STS-2 Orbiter Pitch Rates Measured at 12.5 Hz Between MECO and ET SEP Shown in Greenwich Mean Time (GMT)

- Excessive VRCS pulsing due to plume impingement: Two of the aft VRCS thrusters fired in the down direction (thrusters R5D and L5D in Figure 4). Each VRCS jet was designed to generate 107 N of thrust. The state estimator anticipated the effects of these jet firings with feed forward predictions of their resulting torques on the vehicle. However, upon their use during the STS-1 flight the state estimator over-predicted their torque (shown in Figure 9), with about a 15 second estimate recovery time. This resulted in a large increase in the duty cycles of these thrusters. The cause was found to be thruster plume impingement on the orbiter body flap. The solution to mitigate the problem was to change the estimated thrust to about 59 N, with a change in the resulting torque vector direction. Before the STS-2 mission, detailed analysis of the plume impingement provided specific correction factors for the thruster torque values used in the feed forward estimates, fully resolving this problem for the remainder of the program.




Figure 9. State Estimator Overshoot Due to VRCS Plume Impingement.

EVOLUTION OF OFCS CAPABILITIES OVER THE SHUTTLE PROGRAM LIFE

A primary objective from the inception of the Shuttle program was to operate, deploy, and retrieve payloads in orbit. These operations were all performed using the OODAP for GN&C. Because the knowledge of the status of payload operations was in the GPC running the SM SW, not the GN&C SW, much analysis and some new, special GN&C procedures and SW features were needed to assure that the vehicle maintained stable dynamics during payload operations without harming fragile payload components.

The initial large payload deployment from the Shuttle was on the STS-6 mission, involving the release of an Inertial Upper Stage (IUS) and an attached Tracking and Data Retrieval Satellite (TDRS) with a mass of about 17 tonnes upon deployment. As shown in Figure 10⁵, the payload was rotated 58 deg from a horizontal launch position to a deployment position. Prior to release, the payload cradle had a rocking mode in pitch that could interact with the dynamics of the PRCS algorithm. The PRCS was used closed-loop until just before release to maintain proper release orientation, though free drift was used at release. An innovative application of describing function analysis was applied to determine what OODAP attitude and rate error bounds could be used for closed-loop control without risk of sustained PRCS-firing limit cycles which would cost fuel and could damage to the payload.




Figure 10: IUS TDRS Prior to Deployment.

Canada provided a robotic Remote Manipulator System (RMS) arm for use from the Shuttle payload bay. The RMS was used to extract payloads from the payload bay, position them for release, and in some instances to recover them. The RMS was first used with a small payload (about 155 kg) on the STS-3 mission, with a payload over 1.4 tonnes on the STS-7 mission, and later with payloads well over 10 tonnes, most famously for a series of mission to deploy and upgrade the Hubble Space Telescope (HST) as shown in Figure 11. Even with the small STS-3 payload it became clear that the combined payload/RMS/orbiter dynamics made it unacceptable to utilize thrusters while the RMS moved a payload relative to the orbiter. However, it was necessary to control the Shuttle/payload orientation when the RMS was stationary. Even with the RMS stationary, its primary flexure mode could be near or within the OODAP bandwidth as shown in Figure 12. Extensive simulation analysis was done, including detailed models of the very complex RMS mechanical dynamics, to determine how to use the OODAP safely with large payloads on the RMS. An OODAP feature was added to provide insight into the expected, combined mass properties at pre-planned RMS/payload relative positions, and procedures were implemented to assure VRCS thruster pulse patterns for attitude control did not excite the mission-specific payload/RMS configurations.

⁵ From "A Frequency Domain Stability Analysis of a Phase Plane Control System," by Alexander N. Penchuk, Phillip D. Hattis, and Edward T. Kubiak, AIAA Journal of Guidance and Control. Vol. 8, No. 1, pp. 50-55, Jan.-Feb, 1985.



Figure 11. The HST Being Deployed from the Shuttle Using the RMS. (Courtesy NASA)




Figure 12. RMS Primary Mode Frequency as a Function of Payload Mass.

An advanced Orbital Experiment (OEX) autopilot flew in addition to the baseline OODAP on two Shuttle missions in 1985, and was successfully used during a portion of the orbital operations to demonstrate the benefits of a linear optimal RCS selection algorithm that enabled combined PRCS/VRCS operations and closed-loop orbital translation control that could be used during rendezvous.^{6,7} As an in-flight experiment, a unique feature was a crew-controlled “guillotine”

⁶ “More tests for Draper’s new autopilot,” D-Notes, (a Draper public outreach publication), November 22, 1985.

⁷ “Interview: Bergmann discusses autopilot’s performance,” D-Notes, December 31, 1985

switch to revert immediately to the OODAP in case of unexpected control behavior. However, no problems arose to necessitate use of that feature.

The Shuttle served as the construction vehicle for most of the International Space Station (ISS). Given large changes of ISS mass and inertias as modules were added during assembly, and very large, flexible, and delicate structural features such as the ISS solar panels and radiators, much OODAP design update attention was applied to how to safely operate the Shuttle with the ISS during each phase of ISS construction. The initial demonstration of some of the concepts was accomplished during Shuttle flights to the Russian Mir space station in the 1990s (Figure 13). Specific, new OODAP capabilities added for Shuttle flights to Mir and ISS are shown in Table 1.⁸ For both Mir and ISS operations, the new OODAP features were used to enable the Shuttle to control the entire stack while the Shuttle was attached, and in the case of the ISS to raise its orbit. The success of the ISS assembly process and associated, added OODAP capabilities is highlighted by an ISS completed assembly photo with the Shuttle attached (Figure 14).




Figure 13. Shuttle Atlantis Docked to the Russian Mir Space Station. (Courtesy NASA)



Figure 14. Shuttle Endeavour Docked to the Fully Assembled ISS. (Courtesy NASA)

Table 1. OODAP Upgrades to Support ISS Assembly/Docked Operations.

Upgrade Feature	Function	Description
Alt PRCS	Control while limiting loads	Implemented frequency-constrained pulsing of the PRCS jets to enable control with reduced structural loads.
Minimum Angle Jet Selection	Control	Improved control while docked by better aligning resultant thruster acceleration with desired stack acceleration direction
Notch Filters	Stability during control	Assured closed-loop control stability by minimizing structurally-induced feedback from the stack in the OODAP rate estimate
Reboost	ISS operability	Provided an automated reboost capability within structural limits
Enhanced Feed Forward	Stability during control	Improved OODAP closed-loop performance when using the Notch Filters while docked to the stack
Onboard Mass Props Estimation	Stack operability	Onboard thruster acceleration calculation to limit need for ground uplink of the changing stack configuration data

⁸ Derived from a table in “Space Shuttle GN&C Development History and Evolution,” by Douglas Zimpfer, Phil Hattis, John Rupper, and Don Gavert, presented at the AIAA Space 2011 Conference and Exposition; 26-29 Sep. 2011.

RETROSPECTIVE OBSERVATIONS

About 6 years elapsed from the initiation of OFCS development to first Shuttle flight. Some benefit was derived from prior Apollo design experience, but many new and complex design features were successfully implemented for flight during that comparatively short period, including accommodation of separated SW loads during different flight phases and robust support for fully fault-tolerant avionics. Over the subsequent 30 years of Shuttle flights, the OFCS functionality and procedures evolved to support many classes of orbital operations, including deployment/retrieval of payloads up to and over 20 tonnes, as well as assembly and docked control of the ISS for all configurations used during its build-up. This was done with GPC throughput that was extremely limited and memory capacity that was very small by current standards. Despite a small number of surprises experienced during flights, the OFCS enabled the Shuttle to achieve its originally intended orbital capabilities and more, meeting all mission demands on every flight during which it was invoked.

SOME APPLICABLE PAPERS IN THE PUBLIC RECORD

1. "Digital Autopilot for Thrust Vector Control of the Shuttle Orbital Maneuvering System," by A. Penchuk and S. Croopnick, *AIAA Journal of Guidance and Control*, Vol. 6, No. 6, pp. 436-441, Nov.-Dec., 1983.
2. "Minimum Impulse Limit Cycle Design to Compensate for Measurement Uncertainties," by E.T. Kubiak and M. W. Martin, *AIAA Journal of Guidance and Control*. Vol. 6, No. 6. Pp. 432-435, Nov.-Dec. 1983.
3. "Dynamic Interaction of the Shuttle On-Orbit Flight Control System with Deployed Flexible Payloads," L. Sackett and C.B. Kirchwey, presented at the AIAA Guidance and Control Conference, June 1982.
4. "Stability of the Shuttle On-Orbit Flight Control System for a Class of Flexible Payloads," Christopher B. Kirchwey and Lester L. Sackett, presented at the AIAA Guidance and Control Conference, August, 1983.
5. "A Survey of the Payload Deployment and Retrieval System Environment," by P. Hattis, L. Lepanto, D. Sargent, and J. Turnbull, presented at the Shuttle Payload Dynamic Environment and Loads Prediction Workshop, January 24-26, 1984.
6. "Qualitative Differences Between Shuttle On-Orbit and Transition Control," by Philip D. Hattis, *AIAA Journal of Guidance and Control*. Vol. 7, no. 1, pp. 4-8, Jan-Feb. 1984.
7. "A Frequency Domain Stability Analysis of a Phase Plane Control System," by Alexander N. Penchuk, Phillip D. Hattis, and Edward T. Kubiak, *AIAA Journal of Guidance and Control*. Vol. 8, No. 1, pp. 50-55, Jan.-Feb, 1985.
8. "Combined Solution of Spacecraft Rotational and Translational Maneuvers," by David C. Redding, Bruce A. Persson, and Edward V. Bergmann, presented at the 1986 AIAA Astrodynamics Conference, August 18-20, 1986.
9. "Some Shuttle Orbit Flight Control Design Lessons," by Kenneth J. Cox and Philip D. Hattis, *Proceedings of the IEEE*, Vol. 75, Issue 3, pp. 336-355, March 1987.
10. "Space Shuttle GN&C Development History and Evolution," by Douglas Zimpfer, Phil Hattis, John Rupper, and Don Gavert, presented at the AIAA Space 2011 Conference and Exposition; 26-29 Sep. 2011.