

Maintainability Issues in the Disaster of Space Shuttle “Columbia”

Background

Space Shuttle Columbia was destroyed on February 1, 2003 while returning to the earth. The disaster claimed the lives of all seven of its crew. Designated STS-107, this was Columbia’s 28th flight. The STS-107 mission was a multi-disciplinary science mission modeled after STS-90 and also included placing Triana, a deployable Earth-observing satellite, into the orbit. This disaster was the second tragedy in the 113 flights of the Space Shuttles. The first was the explosion of Space Shuttle Challenger soon after the launch in January 1996. NASA immediately established the Columbia Accident Investigation Board within two hours of the loss of signal from Columbia based on procedures established by NASA responding to the Challenger accident 17 years earlier. The Board finished the investigation and arrived at finding and recommendations for reducing the chances of further accidents in nearly seven months. The first volume of the final report could be accessed at http://caib.nasa.gov/news/report/pdf/vol1/full/caib_report_volume1.pdf. In addition to identifying the direct physical failures that led directly to the destruction, the final report from the Board also identified underlying weaknesses, revealed in NASA’s organization, operations, and history, that can pave the way to catastrophic failure. One weakness resides in the maintenance design and operations of Space Shuttles.


Physical Cause of the Disaster

The direct cause of the disaster was that a piece of insulting foam that separated from the left bipod ramp of the External Tank and struck the leading edge of the left wing during ascent. Columbia re-entered Earth’s atmosphere with the pre-existing breach in the leading edge of its left wing in the vicinity of Reinforced Carbon-Carbon (RCC) panel 8. This



Figure 1: Columbia sitting at Launch Complex 39-A. The upper circle shows the left bipod (–Y) ramp on the forward attach point, while the lower circle is around RCC panel 8-left.

breach was large enough to allow superheated air to penetrate the cavity behind the RCC panel and eventually enter the interior to destroy the left wing. The locations of the left bipod ram and the RCC panel 8-left are shown in Figure 1. The detailed location of RCC panel-8 is shown in Figure 2. There are totally 22 panels of RCC on each wing.



Findings of the Board for the Disaster Causes

Board investigators reviewed Columbia's maintenance records, including the recovery from STS-109 and preparation for STS-107, and relevant areas in NASA's Problem Reporting and Corrective Action database, which contained 16,500 Work Authorization Documents consisting of 600,000 pages and 3.9 million steps. This database maintains critical information on all maintenance and modification work done on the Orbiters (as required by the Orbiter Maintenance Requirements and Specifications Document). It also maintains Corrective Action Reports that document problems discovered and resolved, the Lost/Found item database, and the Launch Readiness Review and Flight Readiness Review documentation.

Post-flight RCC component inspections for cracks, chips, scratches, pinholes, and abnormal discoloration are primarily visual, with tactile evaluations (pushing with a finger) of some regions. Boeing personnel at the Kennedy Space Center make minor repairs to the silicon carbide coating and surface defects. With the goal of a long service life, panels 6 through 17 are refurbished every 18 missions, and panels 18 and 19 every 36 missions. The remaining panels have no specific refurbishment requirement. At the time of STS-107, most of the RCC panels on Columbia's left wing were original equipment, but panel 10-left, T-seal 10-left, panel 11-left, and T-seal 11-left had been replaced (along with panel 12 on the right wing). Panel 10-left was tested to destruction after 19 flights. Minor surface repairs had been made to panels 5, 7, 10, 11, 12, 13, and 19 and T-seals 3, 11, 12, 13, 14, and 19. Panels and T-seals 6 through 9 and 11 through 17 of the left wing had been refurbished.

Some of the Board's findings are as follows.

- The original design specifications required the RCC components to have essentially no impact resistance.
- Current inspection techniques are not adequate to assess structural integrity of the RCC components.
- After manufacturer's acceptance non-destructive evaluation, only periodic visual and touch tests are conducted.
- RCC components are weakened by mass loss caused by oxidation within the substrate, which accumulates with age. The extent of oxidation is not directly measurable, and the resulting mission life reduction is developed analytically.
- date, only two flown RCC panels, having achieved 15 and 19 missions, have been destructively tested to determine actual loss of strength due to oxidation.
- Contamination from zinc leaching from a primer under the paint topcoat on the launch pad structure increases the opportunities for localized oxidation.
- NASA has an inadequate number of spare Reinforced Carbon-Carbon panel assemblies.

The Board made the following recommendations.

- Develop and implement a comprehensive inspection plan to determine the structural integrity of all Reinforced Carbon-Carbon system components. This inspection plan should take advantage of advanced non-destructive inspection technology.
- Initiate a program designed to increase the Orbiter's ability to sustain minor debris damage by measures such as improved impact-resistant Reinforced Carbon-Carbon and acreage tiles. This program should determine the actual impact resistance of current materials and the effect of likely debris strikes.
- To the extent possible, increase the Orbiter's ability to successfully re-enter the Earth's atmosphere with minor leading edge structural sub-system damage.
- In order to understand the true material characteristics of Reinforced Carbon-Carbon components, develop a comprehensive database of flown Reinforced Carbon-Carbon material characteristics by destructive testing and evaluation.
- Improve the maintenance of launch pad structures to minimize the leaching of zinc primer onto Reinforced Carbon-Carbon components.

- Obtain sufficient spare Reinforced Carbon-Carbon panel assemblies and associated support components to ensure that decisions related to Reinforced Carbon-Carbon maintenance are made on the basis of component specifications, free of external pressures relating to schedules, costs, or other considerations.
- Develop, validate, and maintain physics-based computer models to evaluate Thermal Protection System damage from debris impacts. These tools should provide realistic and timely estimates of any impact damage from possible debris from any source that may ultimately impact the Orbiter. Establish impact damage thresholds that trigger responsive corrective action, such as on-orbit inspection and repair, when indicated.

Despite original design requirements that the External Tank not shed debris, and the corresponding design requirement that the Orbiter not receive debris hits exceeding a trivial amount of force, debris has impacted the Shuttle on each flight. Over the course of 113 missions, foam-shedding and other debris impacts came to be regarded more as a turnaround or maintenance issue, and less as a hazard to the vehicle and crew.

Assessments of foam-shedding and strikes were not thoroughly substantiated by engineering analysis, and the process for closing In-Flight Anomalies is not well-documented and appears to vary. Shuttle Program managers appear to have confused the notion of foam posing an “accepted risk” with foam not being a “safety-of-flight issue.” At times, the pressure to meet the flight schedule appeared to cut short engineering efforts to resolve the foam-shedding problem.

The Board is convinced that the factors that led to the Columbia accident go well beyond the physical mechanisms. The causal roots of the accident can also be traced, in part, to the turbulent post-Cold War policy environment faced by NASA. The budget squeeze also came at a time when the Space Shuttle Program exhibited a trait common to most aging systems: increased costs due to greater maintenance requirements, a declining second- and third-tier contractor support base, and deteriorating infrastructure. Maintaining the Shuttle was becoming more expensive at a time when Shuttle budgets were decreasing or being held constant. Budget constraints have impacted personnel and resources required for maintenance and upgrades. Only in the last few years have those budgets begun a gradual increase. In 1999, NASA chartered the Shuttle Independent Assessment Team to examine Shuttle sub-systems and maintenance practices. The Shuttle Independent Assessment Team report stated that the Shuttle “clearly cannot be thought of

as ‘operational’ in the usual sense. Extensive maintenance, major amounts of ‘touch labor’ and a high degree of skill and expertise will always be required.” However, “the workforce has received a conflicting message due to the emphasis on achieving cost and staff reductions, and the pressures placed on increasing scheduled flights as a result of the Space Station. The same ambiguity about investing in Shuttle upgrades has also affected the maintenance of Shuttle Program ground infrastructure, much of which dates to Project Apollo and 1970s Shuttle Program construction. Most ground infrastructure was not built for such a protracted lifespan. Maintaining infrastructure has been particularly difficult at Kennedy Space Center, where it is constantly exposed to a salt water environment.

By July 2002, the Shuttle and Space Station Programs were facing a schedule with very little margin. Two setbacks occurred when technical problems were found during routine maintenance on Discovery. STS-107 was four weeks away from launch at the time, but the problems grounded the entire Shuttle fleet. The longer the fleet was grounded, the more schedule margin was lost, which further compounded the complexity of the intertwined Shuttle and Station schedules.

In August 2002, the Shuttle Program realized it would be unable to meet the Space Station schedule with the available Shuttles. Columbia had never been outfitted to make a Space Station flight, so the other three Orbiters flew the Station missions. But Discovery was in its Orbiter Maintenance Down Period, and would not be available for another 17 months. All Space Station flights until then would have to be made by Atlantis and Endeavour. As managers looked ahead to 2003, they saw that after STS-107, these two Orbiters would have to alternate flying five consecutive missions, STS-114 through STS-118. To alleviate this pressure, and regain schedule margin, Shuttle Program managers elected to modify Columbia to enable it to fly Space Station missions. Those modifications were to take place immediately after STS-107 so that Columbia would be ready to fly its first Space Station mission eight months later. This decision put Columbia directly in the path of Core Complete.

As the autumn of 2002 began, both the Space Shuttle and Space Station Programs began to use what some employees termed “tricks” to regain schedule margin. Employees expressed concern that their ability to gain schedule margin using existing measures was waning.

The industrial safety programs in place at NASA and its contractors are robust and in good health. However, the scope and depth of NASA's maintenance and quality assurance programs are troublesome. Though unrelated to the Columbia accident, the major deficiencies in these programs uncovered by the Board could potentially contribute to a future accident.

United Space Alliance technicians must document an estimated 730,000 tasks to complete a single Shuttle maintenance flow at Kennedy Space Center. Nearly every task assessed as Criticality Code 1, 1R (redundant), or 2 is always inspected, as are any systems not verifiable by operational checks or tests prior to final preparations for flight.

Nearly everyone interviewed at Kennedy indicated that the current inspection process is both inadequate and difficult to expand, even incrementally. One example was a longstanding request to add a main engine final review before transporting the engine to the Orbiter Processing Facility for installation. This request was first voiced two years before the launch of STS-107, and has been repeatedly denied due to inadequate staffing. In its place, NASA Mission Assurance conducts a final "informal" review. Adjusting government inspection tasks is constrained by institutional dogma that the status quo is based on strong engineering logic, and should need no adjustment. This mindset inhibits the ability of Quality Assurance to respond to an aging system, changing workforce dynamics, and improvement initiatives. The Board believes that NASA should adopt and maintain a Shuttle flight schedule that is consistent with available resources. Although schedule deadlines are an important management tool, those deadlines must be regularly evaluated to ensure that any additional risk incurred to meet the schedule is recognized, understood, and acceptable.

To satisfy the Administration's requirement that the system be economically justifiable, *Columbia* had to capture essentially all space launch business, and to do that, it had to meet wide-ranging requirements. These sometimes-competing requirements resulted in a compromise vehicle that was less than optimal for manned flights. NASA designed and developed a remarkably capable and resilient vehicle, consisting of an Orbiter with three Main Engines, two Solid Rocket Boosters, and an External Tank, but one that has never met any of its original requirements for reliability, cost, ease of turnaround, maintainability, or, regrettably, safety. Designated STS-107, this was the Space Shuttle Program's 113th flight and *Columbia*'s 28th. The flight was close to trouble-free. Unfortunately, there were no indications to either the crew

onboard *Columbia* or to engineers in Mission Control that the mission was in trouble as a result of a foam strike during ascent. Mission management failed to detect weak signals that the Orbiter was in trouble and take corrective action. *Columbia* was the first space-rated Orbiter. It made the Space Shuttle Program's first four orbital test flights. Because it was the first of its kind, *Columbia* differed slightly from Orbiters *Challenger*, *Discovery*, *Atlantis*, and *Endeavour*. Built to an earlier engineering standard, *Columbia* was slightly heavier, and, although it could reach the high-inclination orbit of the International Space Station, its payload was insufficient to make *Columbia* cost-effective for Space Station missions. Therefore, *Columbia* was not equipped with a Space Station docking system, which freed up space in the payload bay for longer cargos, such as the science modules Spacelab and SPACEHAB. Consequently, *Columbia* generally flew science missions and serviced the Hubble Space Telescope.

Maintenance Operations for Space Shuttles

The Johnson Space Center in Houston, Texas, is home to both the Space Shuttle and International Space Station Program Offices. The center includes the training, simulation, and mission control centers for the Space Shuttle Program. The Kennedy Space Center located on Merritt Island, Florida, provides maintenance and overhaul services for the Space Shuttle Program. Personnel at Kennedy support for the Orbiters, assemble and check-out the integrated vehicle prior to launch, and operate the Space Station Processing Facility where components of the orbiting laboratory are packaged for launch aboard the Space Shuttle.

During the heyday between 1982 and early 1986, the Shuttle demonstrated its capabilities for space operations with nine missions in 1985. However, the Space Shuttle Program was proving difficult to operate even at that period, with more maintenance required between flights than had been expected. 10 working days was projected during the Shuttle's design phase in 1975 to process a returned Orbiter for its next flight. However, an average of 67 days elapsed before the Shuttle was ready for launch by the end of 1985. The problem became even worse later. For example, *Columbia*'s depot-level maintenance for the STS-107 mission took six months longer than originally planned, primarily to correct problems encountered with Kapton wiring. This resulted in the STS-109 Hubble Space Telescope service mission with *Columbia* being launched before STS-107 because it was considered more urgent. The STS-107 was scheduled for launch

on January 11, 2001, but the actual launch happened on January 16, 2003 after 13 delays over two years.

The Orbiter Major Modification process in which orbiters are removed from service for inspections, maintenance, and modification occurs every eight flights or three years. Orbiter Major Modifications combine with Orbiter flows (preparation of the vehicle for its next mission) and include Orbiter Maintenance Down Periods (not every Orbiter Maintenance Down Period includes an Orbiter Major Modification). The primary differences between an Orbiter Major Modification and an Orbiter flow are the larger number of requirements and the greater degree of intrusiveness of a modification (a recent comparison showed 8,702 Orbiter Major Modification requirements versus 3,826 flow requirements). Ten Orbiter Major Modifications had been performed until August 2003, with an eleventh in progress. They have varied from 6 to 20 months. Because missions do not occur at the rate the Shuttle Program anticipated at its inception, it is endlessly challenged to meet numerous calendar-based requirements. These must be performed regardless of the lower flight rate, which contributes to extensive downtime. The Shuttle Program has explored the possibility of extending Orbiter Major Modification cycles to once every 12 flights or six years. This initiative runs counter to the industry norm of increasing the frequency of inspections as systems age, and should be carefully scrutinized, particularly in light of the high-performance Orbiters' demands.

Throughout the history of Orbiter Major Modifications, a major area of concern has been their wide variability in content and duration. Columbia's last Orbiter Major Modification is just the most recent example of overruns due to technical surprises and management difficulties. It exceeded the schedule by 186 days. While many factors contributed to this delay, the two most prominent were the introduction of a major wiring inspection one month after Orbiter Major Modification roll-in, and what an internal NASA assessment cited as "poor performance on the parts of NASA, USA [United Space Alliance], and Boeing." Estimating the "right" amount of work required on each Orbiter continues to be a challenge. For example, 20 modifications were planned for Discovery's modification; the number has since grown to 84. Such changes introduce turmoil and increase the potential for mistakes.

An Air Force “benchmarking” visit in June 2003 highlighted the need for better planning and more scheduling stability. It further recommended improvements to the requirements feedback process and incorporating service life extension actions into Orbiter Major Modifications.

The Board believed that the Space Shuttle Program Office must make every effort to achieve greater stability, consistency, and predictability in Orbiter Major Modification planning, scheduling, and work standards (particularly in the number of modifications). Endless changes create unnecessary turmoil and can adversely impact quality and safety. NASA and United Space Alliance managers must understand workforce and infrastructure requirements, match them against capabilities, and take actions to avoid exceeding thresholds. NASA should continue to work with the U.S. Air Force, particularly in areas of program management that deal with aging systems, service life extension, planning and scheduling, workforce management, training, and quality assurance. The Space Shuttle Program Office must determine how it will effectively meet the challenges of inspecting and maintaining an aging Orbiter fleet before lengthening Orbiter Major Maintenance intervals.