February 13, 2004

MEMORANDUM TO: Stuart A. Richards, Chief
Division of Inspection Program Management
Office of Nuclear Reactor Regulation

FROM: Jeffrey Jacobson, Program Manager ***/RA/***
Division of Inspection Program Management
Office of Nuclear Reactor Regulation

SUBJECT: READ AND SIGN TRAINING FOR COLUMBIA SPACE SHUTTLE

Attachment: As stated

MEMORANDUM TO: Stuart A. Richards, Chief
Division of Inspection Program Management
Office of Nuclear Reactor Regulation

FROM: Jeffrey Jacobson, Program Manager
Division of Inspection Program Management
Office of Nuclear Reactor Regulation

SUBJECT: READ AND SIGN TRAINING FOR COLUMBIA SPACE SHUTTLE

Attachment: As stated

Accession Number:ML040440475

OFFICE	DIPM/IIPB
NAME	JJacobson
DATE	02/13/04

OFFICIAL RECORD COPY

Columbia Accident Investigation
Read and Sign Training

Introduction

In many respects, the causes of most modern technological failures are more similar than unique. The loss of the Columbia Space Shuttle and the subsequent Columbia Accident Investigation Board's report, contain valuable lessons learned that transcend the aerospace industry and can be used by others to improve the safety mission in their respective fields. Many parallels can be drawn between the organizational causes of the Columbia accident and significant nuclear power plant events. In particular, review of the Columbia report, highlights the importance of maintaining a questioning attitude toward safety, and the possible negative consequences that can occur when such a questioning attitude is lost or compromised.

The complete Columbia Accident Investigation Board report is several hundred pages and contains a wealth of valuable information and insights. This training module attempts to extract those portions of the report that are most pertinent to the NRC's regulation and oversight of the nuclear power industry. For those individuals who wish to pursue a more in-depth reading of the Columbia report it can be found at the following web address:
http://www.caib.us/news/report/default.html

Purpose of Training

The purpose of this training is to use the Columbia Accident Investigation Board report to:

1. illustrate the importance of maintaining a questioning attitude toward safety and the potential negative consequences that can occur when such a questioning attitude is lost or compromised

2. provide examples of how issues concerning an organization's safety culture can lead to technological failures

3. provide insights into investigation techniques that can be used to assess safety significant issues or events

4. illustrate the importance of a robust corrective action program and highlight the corrective action program weaknesses that contributed to the shuttle accident.

Learning Objectives

Completion of this training should provide the reader with an understanding of:

- the primary organizational causes of the Columbia accident

- how working safety issues outside the established corrective action program contributed to ineffective problem resolution with regard to the tank foam separation issue

- how uncertainties regarding the remaining life of the shuttle influenced decisions regarding safety upgrades and corrective actions to previous issues

- the importance of an independent safety oversight organization which pursues safety issues

- how longstanding technical issues can transform from significant safety concerns to less significant long-term study issues

- how an increasing backlog of corrective action items can create an unmonitored risk to safety

- how NASA's corrective action program's categorization scheme (in-family/out of family) may have hindered the evaluation and correction of the tank foam separation issue

- how NASA's lack of effectiveness reviews for corrective actions allowed significant issues to go uncorrected for extended time periods

- how NASA inappropriately used a flawed risk analysis to justify the continued violation of the shuttle's design basis

- how the individuals responsible for the space shuttle's safety were conditioned by past successes and were significantly influenced by schedule pressures

- how NASA's focus on process got in the way of safety during critical decisions that were made during the last flight of the Columbia

- the challenges involved with contracting out the responsibility for safety

- the importance of minority opinions with respect to safety conscious decision-making

- how fault tree analysis was effectively used by the Columbia Accident Investigation Board to eliminate other theories regarding the causes of the Columbia accident

<u>Practical Applications of This Training</u>

<u>For Inspectors</u>

- Be alert to issues being worked by the licensee outside of their corrective action program, particularly those issues that are repetitive or not well understood.

- Ensure proper focus is given to reviews of longstanding and/or repetitive issues that are not being effectively addressed by the licensee (recurring problems) or can not be logically explained due to inadequate root cause analysis

- During reviews of licensee root cause analyses, ensure both the technical and the organizational causes of the issues are clearly articulated.

- Be alert to the effects of schedule pressure on the licensee's willingness to identify and correct issues, or the effects of schedule pressure on the licensee's logic for addressing issues in the corrective action program.

- Be alert to how licensee decisions regarding plant life extension can influence corrective action determinations.

- Maintain cognizance of the limitations of risk analysis, including how the results of a risk analysis can be manipulated by unsubstantiated assumptions.

- When expressing safety concerns to management, strive to do so in a clear and direct manner. Do not shy away from presenting minority opinions on safety issues.

- During discussion of safety issues with the licensee, the focus should be on why a situation is safe, as opposed to why it is unsafe.

<u>For NRC Managers</u>

- Be sensitive to minority opinions during critical decision making and strive to create an environment where minority opinions are encouraged and valued.

- Be cautious when making important decisions based upon a licensee presentation that provides only a summary of the issues from the licensee's perspective ("engineering by viewgraphs"). Ensure that appropriate NRC staff have reviewed the full basis for the licensee's conclusions and that ample opportunities are created to engage the NRC staff regarding their opinions on the issues.

- Do not let process get in the way of safety during important decision making. While following established processes and protocols are important, our number one focus must always be on plant safety.

- Use this example as a way of reinforcing to your staff the importance of maintaining a questioning attitude and the negative consequences of what can happen when such a questioning attitude is lost or compromised.

- During discussion of safety issues with the staff, be sensitive to challenging the staff to prove why a situation is safe, as opposed to why it is unsafe.

<u>Time Estimate</u>

It is estimated that this training will take between two and three hours to complete. Time spent on this training should be charged to training.

contact: Jeff Jacobson, NRR (301 415-2977)

The following text is extracted from the Columbia Accident Investigation Report, Volume 1, dated August 2003. A brief description of the technical causes of the accident is contained in the Part 1 report synopsis which begins on the following page. This is followed by the Part 2 report synopsis which provides an introduction to the organizational causes of the accident.

Chapter 6.1 of the text explains in detail the history of the foam anomalies and the inadequate corrective actions that ultimately led to the accident. Chapter 6.3 contains an analysis of the decision making at NASA during the last flight of Columbia, after it was discovered that bipod foam had again likely impacted the shuttle. Chapter 6.2 of the text discusses schedule pressures that may have influenced decision making both prior to and during the Columbia mission.

Chapter 5 of the text discusses issues associated with cost reductions and staffing that were seen as contributors. Sections 7.4 and 7.5 of the report discuss issues concerning problem reporting and analysis. These sections of the report also talk about a flawed safety culture and its impact on decision-making.

Excerpts from Chapter 4 are included to provide insight into the methodology used by the accident board to conduct its investigation. Lastly, a compilation of the board's recommendations is included for information.

As an aid, the text that is most directly related to the learning objectives has been highlighted for the reader's convenience. Particularly insightful passages are highlighted in pink.

EXECUTIVE SUMMARY

The Columbia Accident Investigation Board's independent investigation into the February 1, 2003, loss of the Space Shuttle Columbia and its seven-member crew lasted nearly seven months. A staff of more than 120, along with some 400 NASA engineers, supported the Board's 13 members. Investigators examined more than 30,000 documents, conducted more than 200 formal interviews, heard testimony from dozens of expert witnesses, and reviewed more than 3,000 inputs from the general public. In addition, more than 25,000 searchers combed vast stretches of the Western United States to retrieve the spacecraft's debris. In the process, Columbia's tragedy was compounded when two debris searchers with the U.S. Forest Service perished in a helicopter accident.

The Board recognized early on that the accident was probably not an anomalous, random event, but rather likely rooted to some degree in NASA's history and the human space flight program's culture. Accordingly, the Board broadened its mandate at the outset to include an investigation of a wide range of historical and organizational issues, including political and budgetary considerations, compromises, and changing priorities over the life of the Space Shuttle Program. The Board's conviction regarding the importance of these factors strengthened as the investigation progressed, with the result that this report, in its findings, conclusions, and recommendations, places as much weight on these causal factors as on the more easily understood and corrected physical cause of the accident.

The physical cause of the loss of Columbia and its crew was a breach in the Thermal Protection System on the leading edge of the left wing, caused by a piece of insulating foam which separated from the left bipod ramp section of the External Tank at 81.7 seconds after launch, and struck the wing in the vicinity of the lower half of Reinforced Carbon-Carbon panel number 8. During re-entry this breach in the Thermal Protection System allowed superheated air to penetrate through the leading edge insulation and progressively melt the aluminum structure of the left wing, resulting in a weakening of the structure until increasing aerodynamic forces caused loss of control, failure of the wing, and break-up of the Orbiter. This breakup occurred in a flight regime in which, given the current design of the Orbiter, there was no possibility for the crew to survive.

The organizational causes of this accident are rooted in the Space Shuttle Program's history and culture, including the original compromises that were required to gain approval for the Shuttle, subsequent years of resource constraints, fluctuating priorities, schedule pressures, mis-characterization of the Shuttle as operational rather than developmental, and lack of an agreed national vision for human space flight. Cultural traits and organizational practices detrimental to safety were allowed to develop, including: reliance on past success as a substitute for sound engineering practices (such as testing to understand why systems were not performing in accordance with requirements); organizational barriers that prevented effective communication of critical safety information and stifled professional differences of opinion; lack of integrated management across program elements; and the evolution of an informal chain of command and decision-making processes that operated outside the organization's rules.

This report discusses the attributes of an organization that could more safely and reliably operate the inherently risky Space Shuttle, but does not provide a detailed organizational prescription. Among those attributes are: a robust and independent program technical authority that has complete control over specifications and requirements, and waivers to them; an independent safety assurance organization with line authority over all levels of safety oversight; and an organizational culture that reflects the best characteristics of a learning organization.

This report concludes with recommendations, some of which are specifically identified and prefaced as "before return to flight." These recommendations are largely related to the physical cause of the accident, and include preventing the loss of foam, improved imaging of the Space Shuttle stack from liftoff through separation of the External Tank, and on-orbit inspection and repair of the Thermal Protection System. The remaining recommendations, for the most part, stem from the Board's findings on organizational cause factors. While they are not "before return to flight" recommendations, they can be viewed as "continuing to fly" recommendations, as they capture the Board's thinking on what changes are necessary to operate the Shuttle and future spacecraft safely in the mid- to long-term.

These recommendations reflect both the Board's strong support for return to flight at the earliest date consistent with the overriding objective of safety, and the Board's conviction that operation of the Space Shuttle, and all human spaceflight, is a developmental activity with high inherent risks.

REPORT SYNOPSIS

PART ONE: THE ACCIDENT

Chapter 1 relates the history of the Space Shuttle Program before the Challenger accident. With the end looming for the Apollo moon exploration program, NASA unsuccessfully attempted to get approval for an equally ambitious (and expensive) space exploration program. Most of the proposed programs started with space stations in low-Earth orbit and included a reliable, economical, medium-lift vehicle to travel safely to and from low-Earth orbit. After many failed attempts, and finally agreeing to what would be untenable compromises, NASA gained approval from the Nixon Administration to develop, on a fixed budget, only the transport vehicle. Because the Administration did not approve a low-Earth-orbit station, NASA had to create a mission for the vehicle. To satisfy the Administration's requirement that the system be economically justifiable, the vehicle 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.

Chapter 2 documents the final flight of Columbia. As a straightforward record of the event, it contains no findings or recommendations. 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.

STS-107 was an intense science mission that required the seven-member crew to form two teams, enabling round-the-clock shifts. Because the extensive science cargo and its extra power sources required additional checkout time, the launch sequence and countdown were about 24 hours longer than normal. Nevertheless, the countdown proceeded as planned, and Columbia was launched from Launch Complex 39-A on January 16, 2003, at 10:39 a.m. Eastern Standard Time (EST).

At 81.7 seconds after launch, when the Shuttle was at about 65,600 feet and traveling at Mach 2.46 (1,650 mph), a large piece of hand-crafted insulating foam came off an area where the Orbiter attaches to the External Tank. At 81.9 seconds, it struck the leading edge of Columbia's left wing. This event was not detected by the crew on board or seen by ground support teams until the next day, during detailed reviews of all launch camera photography and videos. This foam strike had no apparent effect on the daily conduct of the 16-day mission, which met all its objectives.

The de-orbit burn to slow Columbia down for re-entry into Earth's atmosphere was normal, and the flight profile throughout re-entry was standard. Time during re-entry is measured in seconds from "Entry Interface," an arbitrarily determined altitude of 400,000 feet where the Orbiter begins to experience the effects of Earth's atmosphere. Entry Interface for STS-107 occurred at 8:44:09 a.m. on February 1. Unknown to the crew or ground personnel, because the data is recorded and stored in the Orbiter instead of being transmitted to Mission Control at Johnson Space Center, the first abnormal indication occurred 270 seconds after Entry Interface. Chapter 2 reconstructs in detail the events leading to the loss of Columbia and her crew, and refers to more details in the appendices.

In Chapter 3, the Board analyzes all the information available to conclude that the direct, physical action that initiated the chain of events leading to the loss of Columbia and her crew was the foam strike during ascent. This chapter reviews five analytical paths, aerodynamic, thermodynamic, sensor data timeline, debris reconstruction, and imaging evidence to show that all five independently arrive at the same conclusion. The subsequent impact testing conducted by the Board is also discussed.

That conclusion is that Columbia re-entered Earth's atmosphere with a pre-existing breach in the leading edge of its left wing in the vicinity of Reinforced Carbon-Carbon (RCC) panel 8. This breach, caused by the foam strike on ascent, was of sufficient size to allow superheated air (probably exceeding 5,000 degrees Fahrenheit) to penetrate the cavity behind the RCC panel. The breach widened, destroying the insulation protecting the wing's leading edge support structure, and the superheated air eventually melted the thin aluminum wing spar. Once in the interior, the superheated air began to destroy the left wing. This destructive process was carefully reconstructed from the recordings of hundreds of sensors inside the wing, and from analyses of the reactions of the flight control systems to the changes in aerodynamic forces.

By the time Columbia passed over the coast of California in the pre-dawn hours of February 1, at Entry Interface plus 555 seconds, amateur videos show that pieces of the Orbiter were shedding. The Orbiter was captured on videotape during most of its quick transit over the Western United States. The Board correlated the events seen in these videos to sensor readings recorded during re-entry. Analysis indicates that the Orbiter continued to fly its pre-planned flight profile, although, still unknown to anyone on the ground or aboard Columbia, her control systems were working furiously to maintain that flight profile. Finally, over Texas, just southwest of Dallas-Fort Worth, the increasing aerodynamic forces the Orbiter experienced in the denser levels of the atmosphere overcame the catastrophically damaged left wing, causing the

Orbiter to fall out of control at speeds in excess of 10,000 mph.

Chapter 4 describes the investigation into other possible physical factors that may have contributed to the accident. The chapter opens with the methodology of the fault tree analysis, which is an engineering tool for identifying every conceivable fault, then determining whether that fault could have caused the system in question to fail. In all, more than 3,000 individual elements in the Columbia accident fault tree were examined.

In addition, the Board analyzed the more plausible fault scenarios, including the impact of space weather, collisions with micrometeoroids or "space junk," willful damage, flight crew performance, and failure of some critical Shuttle hardware. The Board concludes in Chapter 4 that despite certain fault tree exceptions left "open" because they cannot be conclusively disproved, none of these factors caused or contributed to the accident. This chapter also contains findings and recommendations to make Space Shuttle operations safer.

PART TWO: WHY THE ACCIDENT OCCURRED

Many accident investigations do not go far enough. They identify the technical cause of the accident, and then connect it to a variant of "operator error" the line worker who forgot to insert the bolt, the engineer who miscalculated the stress, or the manager who made the wrong decision. But this is seldom the entire issue. When the determinations of the causal chain are limited to the technical flaw and individual failure, typically the actions taken to prevent a similar event in the future are also limited: fix the technical problem and replace or retrain the individual responsible. Putting these corrections in place leads to another mistake, the belief that the problem is solved. The Board did not want to make these errors.

Attempting to manage high-risk technologies while minimizing failures is an extraordinary challenge. By their nature, these complex technologies are intricate, with many interrelated parts. Standing alone, the components may be well understood and have failure modes that can be anticipated. Yet when these components are integrated into a larger system, unanticipated interactions can occur that lead to catastrophic outcomes. The risk of these complex systems is increased when they are produced and operated by complex organizations that also break down in unanticipated ways.

In our view, the NASA organizational culture had as much to do with this accident as the foam. Organizational culture refers to the basic values, norms, beliefs, and practices that characterize the functioning of an institution. At the most basic level, organizational culture defines the assumptions that employees make as they carry out their work. It is a powerful force that can persist through reorganizations and the change of key personnel. It can be a positive or a negative force.

In a report dealing with nuclear wastes, the National Research Council quoted Alvin Weinbergs classic statement about the "Faustian bargain" that nuclear scientists made with society. "The price that we demand of society for this magical energy source is both a vigilance and a longevity of our social institutions that we are quite unaccustomed to." This is also true of the space program. At NASAs urging, the nation committed to building an amazing, if compromised, vehicle called the Space Shuttle. When the agency did this, it accepted the bargain to operate and maintain the vehicle in the safest possible way. The Board is not convinced that NASA has completely lived up to the bargain, or that Congress and the Administration has provided the funding and support necessary for NASA to do so. This situation needs to be addressed if the nation intends to keep conducting human space flight, it needs to live up to its part of the bargain.

Part Two of this report examines NASAs organizational, historical, and cultural factors, as well as how these factors contributed to the accident. As in Part One, this part begins with history. Chapter 5 examines the post-Challenger history of NASA and its Human Space Flight Program. This includes reviewing the budget as well as organizational and management history, such as shifting management systems and locations. Chapter 6 documents management performance related to Columbia to establish events analyzed in later chapters. The chapter reviews the foam strikes, intense schedule pressure driven by an artificial requirement to de-liver Node 2 to the International Space Station by a certain date, and NASA managements handling of concerns regarding Columbia during the STS-107 mission.

In Chapter 7, the Board presents its views of how high-risk activities should be managed, and lists the characteristics of institutions that emphasize high-reliability results over economic efficiency or strict adherence to a schedule. This chapter measures the Space Shuttle Programs organizational and management practices against these principles and finds them wanting. Chapter 7 defines the organizational cause and offers recommendations. Chapter 8 draws from the previous chapters on history, budgets, culture, organization, and safety practices, and analyzes how all these factors contributed to this accident. This chapter captures the Boards views of the need to adjust management to enhance safety margins in Shuttle operations, and reaffirms the Boards position that without these changes, we have no confidence that other “corrective actions” will improve the safety of Shuttle operations. The changes we recommend will be difficult to accomplish and will be internally resisted.

Chapter 6 Decision Making At NASA

6.1 A HISTORY OF FOAM ANOMALIES

The shedding of External Tank foam, the physical cause of the Columbia accident had a long history. Damage caused by debris has occurred on every Space Shuttle flight, and most missions have had insulating foam shed during ascent. This raises an obvious question: Why did NASA continue flying the Shuttle with a known problem that violated design requirements? It would seem that the longer the Shuttle Program allowed debris to continue striking the Orbiters, the more opportunity existed to detect the serious threat it posed. But this is not what happened. Although engineers have made numerous changes in foam design and application in the 25 years that the External Tank has been in production, the problem of foam-shedding has not been solved, nor has the Orbiter's ability to tolerate impacts from foam or other debris been significantly improved.

The Need for Foam Insulation

The External Tank contains liquid oxygen and hydrogen propellants stored at minus 297 and minus 423 degrees Fahrenheit. Were the super-cold External Tank not sufficiently insulated from the warm air, its liquid propellants would boil, and atmospheric nitrogen and water vapor would condense and form thick layers of ice on its surface. Upon launch, the ice could break off and damage the Orbiter. (See Chapter 3.)

To prevent this from happening, large areas of the External Tank are machine-sprayed with one or two inches of foam, while specific fixtures, such as the bipod ramps, are hand-sculpted with thicker coats. Most of these insulating materials fall into a general category of "foam," and are outwardly similar to hardware store sprayable foam insulation. The problem is that foam does not always stay where the External Tank manufacturer Lockheed Martin installs it. During flight, popcorn- to briefcase-size chunks detach from the External Tank.

Original Design Requirements

Early in the Space Shuttle Program, foam loss was considered a dangerous problem. Design engineers were extremely concerned about potential damage to the Orbiter and its fragile Thermal Protection System, parts of which are so vulnerable to impacts that lightly pressing a thumbnail into them leaves a mark. Because of these concerns, the baseline design requirements in the Shuttle's "Flight and Ground System Specification-Book 1, Requirements," precluded foam-shedding by the External Tank. Specifically:

> **3.2.1.2.14 Debris Prevention**: The Space Shuttle System, including the ground systems, shall be designed to preclude the shedding of ice and/or other debris from the Shuttle elements during prelaunch and flight operations that would jeopardize the flight crew, vehicle, mission success, or would adversely impact turnaround operations.
>
> **3.2.1.1.17 External Tank Debris Limits**: No debris shall emanate from the critical zone of the External Tank on the launch pad or during ascent except for such material which may result from normal thermal protection system recession due to ascent heating.

The assumption that only tiny pieces of debris would strike the Orbiter was also built into original design requirements, which specified that the Thermal Protection System (the tiles and Reinforced Carbon-Carbon, or RCC, panels) would be built to withstand impacts with a kinetic energy less than 0.006 foot-pounds. Such a small tolerance leaves the Orbiter vulnerable to strikes from birds, ice, launch pad debris, and pieces of foam.

Despite the design requirement that the External Tank shed no debris, and that the Orbiter not be subjected to any significant debris hits, Columbia sustained damage from debris strikes on its inaugural 1981 flight. More than 300 tiles had to be replaced. Engineers stated that had they known in advance that the External Tank "was going to produce the debris shower that occurred" during launch, "they would have had a difficult time clearing Columbia for flight."

Discussion of Foam Strikes Prior to the Rogers Commission

Foam strikes were a topic of management concern at the time of the Challenger accident. In fact, during the Rogers Commission accident investigation, Shuttle Program Manager Arnold Aldrich cited a contractor's concerns about foam shedding to illustrate how well the Shuttle Program manages risk:

On a series of four or five external tanks, the thermal insulation around the inner tank … had large divots of insulation coming off and impacting the Orbiter. We found significant amount of damage to one Orbiter after a flight and … on the subsequent flight we had a camera in the equivalent of the wheel well, which took a picture of the tank after separation, and we determined that this was in fact the cause of the damage. At that time, we wanted to be able to proceed with the launch program if it was acceptable … so we undertook discussions of what would be acceptable in terms of potential field repairs, and during those discussions, Rockwell was very conservative because, rightly, damage to the Orbiter TPS [Thermal Protection System] is damage to the Orbiter system, and it has a very stringent environment to experience during the re-entry phase.

Aldrich described the pieces of foam as "… half a foot square or a foot by half a foot, and some of them much smaller and localized to a specific area, but fairly high up on the tank. So they had a good shot at the Orbiter underbelly, and this is where we had the damage."

Continuing Foam Loss

Despite the high level of concern after STS-1 and through the Challenger accident, foam continued to separate from the External Tank. Photographic evidence of foam shedding exists for 65 of the 79 missions for which imagery is available. Of the 34 missions for which there are no imagery, 8 missions where foam loss is not seen in the imagery, and 6 missions where imagery is inconclusive, foam loss can be inferred from the number of divots on the Orbiter's lower surfaces. Over the life of the Space Shuttle Program, Orbiters have returned with an average of 143 divots in the upper and lower surfaces of the Thermal Protection System tiles, with 31 divots averaging over an inch in one dimension. (The Orbiters' lower surfaces have an average of 101 hits, 23 of which are larger than an inch in diameter.) Though the Orbiter is also struck by ice and pieces of launch-pad hardware during launch, by micrometeoroids and orbital debris in space, and by runway debris during landing, the Board concludes that foam is likely responsible for most debris hits.

With each successful landing, it appears that NASA engineers and managers increasingly regarded the foam-shedding as inevitable, and as either unlikely to jeopardize safety or simply an acceptable risk. The distinction between foam loss and debris events also appears to have become blurred. NASA and contractor personnel came to view foam strikes not as a safety of flight issue, but rather a simple maintenance, or "turnaround" issue. In Flight Readiness Review documentation, Mission Management Team minutes, In-Flight Anomaly disposition reports, and elsewhere, what was originally considered a serious threat to the Orbiter came to be treated as "in-family," a reportable problem that was within the known experience base, was believed to be understood, and was not regarded as a safety-of-flight issue.

DEFINITIONS

In Family: A reportable problem that was previously experienced, analyzed, and understood. Out of limits performance or discrepancies that have been previously experienced may be considered as in family when specifically approved by the Space Shuttle Program or design project.

Out of Family: Operation or performance outside the expected performance range for a given parameter or which has not previously been experienced.

Accepted Risk: The threat associated with a specific circumstance is known and understood, cannot be completely eliminated, and the circumstance(s) producing that threat is considered unlikely to reoccur. Hence, the circumstance is fully known and is considered a tolerable threat to the conduct of a Shuttle mission.

No Safety-of-Flight-Issue: The threat associated with a specific circumstance is known and understood and does not pose a threat to the crew and/or vehicle.

Bipod Ramp Foam Loss Events

Chunks of foam from the External Tank's forward bipod attachment, which connects the Orbiter to the External Tank, are some of the largest pieces of debris that have struck the Orbiter. To place the foam loss from STS-107 in a broader context, the Board examined every known instance of foam-shedding from this area. Foam loss from the left bipod ramp (called the Y ramp in NASA parlance) has been confirmed by imagery on 7 of the 113 missions flown. However, only on 72 of these missions was available imagery of sufficient quality to determine left bipod ramp foam loss. Therefore, foam loss from the left bipod area occurred on approximately 10 percent of flights (seven events out of 72 imaged flights). On the 66 flights that imagery was available for the right bipod area, foam loss was never observed. NASA could not explain why only the left bipod experienced foam loss.

The first known bipod ramp foam loss occurred during STS-7,Challenger's second mission. Images taken after External Tank separation revealed that a 19- by 12-inch piece of the left bipod ramp was missing, and that the External Tank had some 25 shallow divots in the foam just forward of the bipod struts and another 40 divots in the foam covering the lower External Tank. After the mission was completed, the Program Requirements Control Board cited the foam loss as an In-Flight Anomaly. Citing an event as an In-Flight Anomaly means that before the next launch, a specific NASA organization must resolve the problem or prove that it does not threaten the safety of the vehicle or crew.

At the Flight Readiness Review for the next mission, Orbiter Project management reported that, based on the completion of repairs to the Orbiter Thermal Protection System, the bipod ramp foam loss In-Flight Anomaly was resolved, or "closed." However, although the closure documents detailed the repairs made to the Orbiter, neither the Certificate of Flight Readiness documentation nor the Flight Readiness Review documentation referenced correcting the cause of the damage, the shedding of foam.

The second bipod ramp foam loss occurred during STS-32R, Columbia's ninth flight, on January 9, 1990. A post-mission review of STS-32R photography revealed five divots in the intertank foam ranging from 6 to 28 inches in diameter, the largest of which extended into the left bipod ramp foam. A post-mission inspection of the lower surface of the Orbiter revealed 111 hits, 13 of which were one inch or greater in one dimension. An In-Flight Anomaly assigned to the External Tank Project was closed out at the Flight Readiness Review for the next mission, STS-36, on the basis that there may have been local voids in the foam bipod ramp where it attached to the metal skin of the External Tank. To address the foam loss, NASA engineers poked small "vent holes" through the intertank foam to allow trapped gases to escape voids in the foam where they otherwise might build up pressure and cause the foam to pop off. However, NASA is still studying this hypothesized mechanism of foam loss. Experiments conducted under the Board's purview indicate that other mechanisms may be at work. As discussed in Chapter 3, the Board notes that the persistent uncertainty about the causes of foam loss and potential Orbiter damage results from a lack of thorough hazard analysis and engineering attention.

The third bipod foam loss occurred on June 25, 1992, during the launch of Columbia on STS-50, when an approximately 26- by 10-inch piece separated from the left bipod ramp area. Post-mission inspection revealed a 9-inch by 4.5-inch by 0.5-inch divot in the tile, the largest area of tile damage in Shuttle history. The External Tank Project at Marshall Space Flight Center and the Integration Office at Johnson Space Center cited separate In-Flight Anomalies. The Integration Office closed out its In-Flight Anomaly two days before the next flight, STS-46, by deeming damage to the Thermal Protection System an "accepted flight risk." In Integration Hazard Report 37, the Integration Office noted that the impact damage was shallow, the tile loss was not a result of excessive aerodynamic loads, and the External Tank Thermal Protection System failure was the result of "inadequate venting." The External Tank Project closed out its

In-Flight Anomaly with the rationale that foam loss during ascent was “not considered a flight or safety issue.” Note the difference in how the each program addressed the foam-shedding problem: While the Integration Office deemed it an “accepted risk,” the External Tank Project considered it “not a safety-of-flight issue.” Hazard Report 37 would figure in the STS-113 Flight Readiness Review, where the crucial decision was made to continue flying with the foam-loss problem. This inconsistency would reappear 10 years later, after bipod foam-shedding during STS-112.

The fourth and fifth bipod ramp foam loss events went undetected until the Board directed NASA to review all available imagery for other instances of bipod foam-shedding. This review of imagery from tracking cameras, the umbilical well camera, and video and still images from flight crew hand held cameras revealed bipod foam loss on STS-52 and STS-62, both of which were flown by Columbia. Because these incidents of missing bipod foam were not detected until after the STS-107 accident, no In-Flight Anomalies had been written. The Board concludes that NASA’s failure to identify these bipod foam losses at the time they occurred means the agency must examine the adequacy of its film review, post-flight inspection, and Program Requirements Control Board processes.

The sixth and final bipod ramp event before STS-107 occurred during STS-112 on October 7, 2002. At 33 seconds after launch, when Atlantis was at 12,500 feet and traveling at Mach 0.75, ground cameras observed an object traveling from the External Tank that subsequently impacted the Solid Rocket Booster/External Tank Attachment ring. After impact, the debris broke into multiple pieces that fell along the Solid Rocket Booster exhaust plume.15 Post-mission inspection of the Solid Rocket Booster confirmed damage to foam on the forward face of the External Tank Attachment ring. The impact was approximately 4 inches wide and 3 inches deep. Post-External Tank separation photography by the crew showed that a 4- by 5- by 12-inch (240 cubic-inch) corner section of the left bipod ramp was missing, which exposed the super lightweight ablator coating on the bipod housing. This missing chunk of foam was believed to be the debris that impacted the External Tank Attachment ring during ascent. The post-launch review of photos and video identified these debris events, but the Mission Evaluation Room logs and Mission Management Team minutes do not reflect any discussions of them.

STS-113 Flight Readiness Review: A Pivotal Decision

Because the bipod ramp shedding on STS-112 was significant, both in size and in the damage it caused, and because it occurred only two flights before STS-107, the Board investigated NASA’s rationale to continue flying. This decision made by the Program Requirements Control Board at the STS-113 Flight Readiness Review is among those most directly linked to the STS-107 accident. Had the foam loss during STS-112 been classified as a more serious threat, managers might have responded differently when they heard about the foam strike on STS-107. Alternately, in the face of the increased risk, STS-107 might not have flown at all. However, at STS-113’s Flight Readiness Review, managers formally accepted a flight rationale that stated it was safe to fly with foam losses. This decision enabled, and perhaps even encouraged, Mission Management Team members to use similar reasoning when evaluating whether the foam strike on STS-107 posed a safety-of-flight issue.

At the Program Requirements Control Board meeting following the return of STS-112, the Intercenter Photo Working Group recommended that the loss of bipod foam be classified as an In-Flight Anomaly. In a meeting chaired by Shuttle Program Manager Ron Dittemore and attended by many of the managers who would be actively involved with STS-107, including Linda Ham, the Program Requirements Control Board ultimately decided against such classification. Instead, after discussions with the Integration Office and the External Tank Project, the Program Requirements Control Board Chairman assigned an “action” to the External Tank Project to determine the root cause of the foam loss and to propose corrective action. This was inconsistent with previous practice, in which all other known bipod foam-shedding was designated as In-Flight Anomalies. The Program Requirements Control Board initially set December 5, 2002, as the date to report back on this action, even though STS-113 was scheduled to launch on November 10. The due date subsequently slipped until after the planned launch and return of STS-107. The Space Shuttle Program decided to fly not one but two missions before resolving the STS-112 foam loss.

The Board wondered why NASA would treat the STS-112 foam loss differently than all others. What drove managers to reject the recommendation that the foam loss be deemed an In-Flight Anomaly? Why did they take the unprecedented step of scheduling not one but eventually two missions to fly before the External Tank Project was to report back on foam losses? It seems that Shuttle managers had become conditioned over time to not regard foam loss or debris as a safety-of-flight concern. As will be discussed in Section 6.2, the need to adhere to the Node 2 launch schedule also appears to have influenced their decision. Had the STS-113 mission been delayed beyond early December 2002, the Expedition 5 crew on board the Space Station would have exceeded its 180-day on-orbit limit, and the Node 2 launch date, a major management goal, would not be met.

Even though the results of the External Tank Project engineering analysis were not due until after STS-113, the foam-shedding was reported, or "briefed," at STS-113's Flight Readiness Review on October 31, 2002, a meeting that Dittemore and Ham attended.

The conclusion of the analysis that (added for clarity) the fact that ramp closeout work was "performed by experienced practitioners" or that "application involves craftsmanship in the use of validated application processes" in no way decreases the chances of recurrent foam loss. The statement that the "probability of loss of ramp Thermal Protection System is no higher/no lower than previous flights" could be just as accurately stated "the probability of bipod foam loss on the next flight is just as high as it was on previous flights." With no engineering analysis, Shuttle managers used past success as a justification for future flights, and made no change to the External Tank configurations planned for STS-113,and, subsequently, for STS-107.

Along with this chart, the NASA Headquarters Safety Office presented a report that estimated a 99 percent probability of foam not being shed from the same area, even though no corrective action had been taken following the STS-112 foam-shedding. The ostensible justification for the 99 percent figure was a calculation of the actual rate of bipod loss over 61 flights. This calculation was a sleight-of-hand effort to make the probability of bipod foam loss appear low rather than a serious grappling with the probability of bipod ramp foam separating. For one thing, the calculation equates the probability of left and right bipod loss, when right bipod loss has never been observed, and the amount of imagery available for left and right bipod events differs. The calculation also miscounts the actual number of bipod ramp losses in two ways. First, by restricting the sample size to flights between STS-112 and the last known bipod ramp loss, it excludes known bipod ramp losses from STS-7, STS-32R, and STS-50. Second, by failing to project the statistical rate of bipod loss across the many missions for which no bipod imagery is available, the calculation assumes a "what you don't see won't hurt you" mentality when in fact the reverse is true. When the statistical rate of bipod foam loss is projected across missions for which imagery is not available, and the sample size is extended to include every mission from STS-1 on, the probability of bipod loss increases dramatically. The Board's review after STS-107, which included the discovery of two additional bipod ramp losses that NASA had not previously noted, concluded that bipod foam loss occurred on approximately 10 percent of all missions.

During the brief at STS-113's Flight Readiness Review, the Associate Administrator for Safety and Mission Assurance scrutinized the Integration Hazard Report 37 conclusion that debris-shedding was an accepted risk, as well as the External Tank Project's rationale for flight. After conferring, STS-113 Flight Readiness Review participants ultimately agreed that foam shedding should be characterized as an "accepted risk" rather than a "not a safety-of-flight" issue. Space Shuttle Program management accepted this rationale, and STS-113's Certificate of Flight Readiness was signed.

The decision made at the STS-113 Flight Readiness Review seemingly acknowledged that the foam posed a threat to the Orbiter, although the continuing disagreement over whether foam was "not a safety of flight issue" versus an "accepted risk" demonstrates how the two terms became blurred over time, clouding the precise conditions under which an increase in risk would be permitted by Shuttle Program management. In retrospect, the bipod foam that caused a 4- by 3-inch gouge in the foam on one of Atlantis' Solid Rocket Boosters just months before STS-107 was a "strong signal" of potential future damage that Shuttle engineers ignored. Despite the significant bipod foam loss on STS-112, Shuttle Program engineers made no External Tank configuration changes, no moves to reduce the risk of bipod ramp shedding or potential damage to the Orbiter on either of the next two flights, STS-113 and STS-107, and did not update Integrated Hazard Report 37. The Board notes that although there is a process for conducting hazard analyses when the system is designed and a process for re-evaluating them when a design is changed or the component is replaced, no process addresses the need to update a hazard analysis when anomalies occur. A stronger Integration Office would likely have insisted that Integrated Hazard Analysis 37 be updated. In the course of that update, engineers would be forced to consider the cause of foam-shedding and the effects of shedding on other Shuttle elements, including the Orbiter Thermal Protection System.

STS-113 launched at night, and although it is occasionally possible to image the Orbiter from light given off by the Solid Rocket Motor plume, in this instance no imagery was obtained and it is possible that foam could have been shed.

The acceptance of the rationale to fly cleared the way for Columbia's launch and provided a method for Mission managers to classify the STS-107 foam strike as a maintenance and turnaround concern rather than a safety-of-flight issue. It is significant that in retrospect, several NASA managers identified their acceptance of this flight rationale as a serious error.

The foam-loss issue was considered so insignificant by some Shuttle Program engineers and managers that the STS-107 Flight Readiness Review documents include no discussion of the still-unresolved STS-112 foam loss. According to Program rules, this discussion was not a requirement because the STS-112 incident was only identified as an "action," not an In-Flight Anomaly. However, because the action was still open, and the date of its resolution had slipped, the Board believes that Shuttle Program managers should have addressed it. Had the foam issue been discussed in STS-107 pre-launch meetings, Mission managers may have been more sensitive to the foam-shedding, and may have taken more aggressive steps to determine the extent of the damage.

The seventh and final known bipod ramp foam loss occurred on January 16, 2003, during the launch of Columbia on STS-107. After the Columbia bipod loss, the Program Requirements Control Board deemed the foam loss an In-Flight Anomaly to be dealt with by the External Tank Project.

Other Foam/Debris Events

To better understand how NASA's treatment of debris strikes evolved over time, the Board investigated missions where debris was shed from locations other than the External Tank bipod ramp.

The number of debris strikes may be small, but a single strike could damage several tiles. One debris strike in particular foreshadows the STS-107 event. When Atlantis was launched on STS-27R on December 2, 1988, the largest debris event up to that time significantly damaged the Orbiter. Post-launch analysis of tracking camera imagery by the Intercenter Photo Working Group identified a large piece of debris that struck the Thermal Protection System tile at approximately 85 seconds into the flight. On Flight Day Two, Mission Control asked the flight crew to inspect Atlantis with a camera mounted on the remote manipulator arm, a robotic device that was not installed on Columbia for STS-107. Mission Commander R.L. "Hoot" Gibson later stated that Atlantis "looked like it had been blasted by a shotgun." Concerned that the Orbiter's Thermal Protection System had been breached, Gibson ordered that the video be transferred to Mission Control so that NASA engineers could evaluate the damage.

When Atlantis landed, engineers were surprised by the extent of the damage. Post-mission inspections deemed it “the most severe of any mission yet flown.” The Orbiter had 707 dings, 298 of which were greater than an inch in one dimension. Damage was concentrated outboard of a line right of the bipod attachment to the liquid oxygen umbilical line. Even more worrisome, the debris had knocked off a tile, exposing the Orbiter’s skin to the heat of re-entry. Post-flight analysis concluded that structural damage was confined to the exposed cavity left by the missing tile, which happened to be at the location of a thick aluminum plate covering an L-band navigation antenna. Were it not for the thick aluminum plate, Gibson stated during a presentation to the Board that a burn-through may have occurred.

The Board notes the distinctly different ways in which the STS-27R and STS-107 debris strike events were treated. After the discovery of the debris strike on Flight Day Two of STS-27R, the crew was immediately directed to inspect the vehicle. More severe thermal damage, perhaps even a burn-through may have occurred were it not for the aluminum plate at the site of the tile loss. Fourteen years later, when a debris strike was discovered on Flight Day Two of STS-107, Shuttle Program management declined to have the crew inspect the Orbiter for damage, declined to request on-orbit imaging, and ultimately discounted the possibility of a burn-through. In retrospect, the debris strike on STS-27R is a “strong signal” of the threat debris posed that should have been considered by Shuttle management when STS-107 suffered a similar debris strike. The Board views the failure to do so as an illustration of the lack of institutional memory in the Space Shuttle Program that supports the Board’s claim, discussed in Chapter 7, that NASA is not functioning as a learning organization.

After the STS-27R damage was evaluated during a post-flight inspection, the Program Requirements Control Board assigned In-Flight Anomalies to the Orbiter and Solid Rocket Booster Projects. Marshall Sprayable Ablator (MSA-1) material found embedded in an insulation blanket on the right Orbital Maneuvering System pod confirmed that the ablator on the right Solid Rocket Booster nose cap was the most likely source of debris. Because an improved ablator material (MSA-2) would now be used on the Solid Rocket Booster nose cap, the issue was considered “closed” by the time of the next mission’s Flight Readiness Review. The Orbiter Thermal Protection System review team concurred with the use of the improved ablator without reservation.

An STS-27R investigation team notation mirrors a Columbia Accident Investigation Board finding. The STS-27R investigation noted: “it is observed that program emphasis and attention to tile damage assessments varies with severity and that detailed records could be augmented to ease trend maintenance” (emphasis added).22 In other words, Space Shuttle Program personnel knew that the monitoring of tile damage was inadequate and that clear trends could be more readily identified if monitoring was improved, but no such improvements were made. The Board also noted that an STS-27R investigation team recommendation correlated to the Columbia accident 14 years later: “It is recommended that the program actively solicit design improvements directed toward eliminating debris sources or minimizing damage potential.”

Making the Orbiter More Resistant To Debris Strikes

If foam shedding could not be prevented entirely, what did NASA do to make the Thermal Protection System more resistant to debris strikes? A 1990 study by Dr. Elisabeth Paté-Cornell and Paul Fishback attempted to quantify the risk of a Thermal Protection System failure using probabilistic analysis. The data they used included (1) the probability that a tile would become debonded by either debris strikes or a poor bond, (2) the probability of then losing adjacent tiles, (3) depending on the final size of the failed area, the probability of burn-through, and (4) the probability of failure of a critical sub-system if burn-through occurs. The study concluded that the probability of losing an Orbiter on any given mission due to a failure of Thermal Protection System tiles was approximately one in 1,000. Debris-related problems accounted for approximately 40 percent of the probability, while 60 percent was attributable to tile debonding caused by other factors. An estimated 85 percent of the risk could be attributed to 15 percent of the “acreage,” or larger areas of tile, meaning that the loss of any one of a relatively small number of tiles pose a relatively large amount of risk to the Orbiter. In other words, not all tiles are equal , losing certain tiles is more dangerous. While the actual risk may be different than that computed in the 1990 study due to the limited amount of data and the underlying simplified assumptions, this type of analysis offers insight that enables management to concentrate their resources on protecting the Orbiters' critical areas.

Two years after the conclusion of that study, NASA wrote to Paté-Cornell and Fishback describing the importance of their work, and stated that it was developing a long-term effort to use probabilistic risk assessment and related disciplines to improve programmatic decisions. Though NASA has taken some measures to invest in probabilistic risk assessment as a tool, it is the Board's view that NASA has not fully exploited the insights that Paté-Cornell's and Fishback's work offered.

Impact Resistant Tile

NASA also evaluated the possibility of increasing Thermal Protection System tile resistance to debris hits, lowering the possibility of tile debonding, and reducing tile production and maintenance costs. Indeed, tiles with a "tough" coating are currently used on the Orbiters. This coating, known as Toughened Uni-piece Fibrous Insulation (TUFI), was patented in 1992 and developed for use on high-temperature rigid insulation. TUFI is used on a tile material known as Alumina Enhanced Thermal Barrier (AETB), and has a debris impact resistance that is greater than the current acreage tile's resistance by a factor of approximately 6-20. At least 772 of these advanced tiles have been installed on the Orbiters' base heat shields and upper body flaps. However, due to its higher thermal conductivity, TUFI-coated AETB cannot be used as a replacement for the larger areas of tile coverage. (Boeing, Lockheed Martin and NASA are developing a lightweight, impact-resistant, low-conductivity tile.) Because the impact requirements for these next-generation tiles do not appear to be based on resistance to specific (and probable) damage sources, it is the Board's view that certification of the new tile will not adequately address the threat posed by debris.

Conclusion

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.

NASA's lack of understanding of foam properties and behavior must also be questioned. Although tests were conducted to develop and qualify foam for use on the External Tank, it appears there were large gaps in NASA's knowledge about this complex and variable material. Recent testing conducted at Marshall Space Flight Center and under the auspices of the Board indicate that mechanisms previously considered a prime source of foam loss, cryopumping and cryoingestion, are not feasible in the conditions experienced during tanking, launch, and ascent. Also, dissections of foam bipod ramps on External Tanks yet to be launched reveal subsurface flaws and defects that only now are being discovered and identified as contributing to the loss of foam from the bipod ramps.

While NASA properly designated key debris events as In-Flight Anomalies in the past, more recent events indicate that NASA engineers and management did not appreciate the scope, or lack of scope, of the Hazard Reports involving foam shedding. Ultimately, NASA's hazard analyses, which were based on reducing or eliminating foam-shedding, were not succeeding. Shuttle Program management made no adjustments to the analyses to recognize this fact. The acceptance of events that are not supposed to happen has been described by sociologist Diane Vaughan as the “normalization of deviance.” The history of foam-problem decisions shows how NASA first began and then continued flying with foam losses, so that flying with these deviations from design specifications was viewed as normal and acceptable. Dr. Richard Feynman, a member of the Presidential Commission on the Space Shuttle Challenger Accident, discusses this phenomena in the context of the Challenger accident. The parallels are striking:

> The phenomenon of accepting … flight seals that had shown erosion and blow-by in previous flights is very clear. The Challenger flight is an excellent example. There are several references to flights that had gone before. The acceptance and success of these flights is taken as evidence of safety. But erosions and blow-by are not what the design expected. They are warnings that something is wrong … The O-rings of the Solid Rocket Boosters were not designed to erode. Erosion was a clue that something was wrong. Erosion was not something from which safety can be inferred … If a reasonable launch schedule is to be maintained, engineering often cannot be done fast enough to keep up with the expectations of originally conservative certification criteria designed to guarantee a very safe vehicle. In these situations, subtly, and often with apparently logical arguments, the criteria are altered so that flights may still be certified in time. They therefore fly in a relatively unsafe condition, with a chance of failure of the order of a percent (it is difficult to be more accurate).

Chapter 6.3 Decision Making During The Flight of STS 107

Summary: Mission Management Decision Making

Discovery and Initial Analysis of Debris Strike

In the course of examining film and video images of *Columbia's* ascent, the Intercenter Photo Working Group identified, on the day after launch, a large debris strike to the leading edge of *Columbia's* left wing. Alarmed at seeing so severe a hit so late in ascent, and at not having a clear view of damage the strike might have caused, Intercenter Photo Working Group members alerted senior Program managers by phone and sent a digitized clip of the strike to hundreds of NASA personnel via e-mail. These actions initiated a contingency plan that brought together an interdisciplinary group of experts from NASA, Boeing, and the United Space Alliance to analyze the strike. So concerned were Intercenter Photo Working Group personnel that on the day they discovered the debris strike, they tapped their Chair, Bob Page, to see through a request to image the left wing with Department of Defense assets in anticipation of analysts needing these images to better determine potential damage. By the Board's count, this would be the first of three requests to secure imagery of *Columbia* on-orbit during the 16-day mission.

Upon learning of the debris strike on Flight Day Two, the responsible system area manager from United Space Alliance and her NASA counterpart formed a team to analyze the debris strike in accordance with mission rules requiring the careful examination of any “out-of-family” event. Using film from the Intercenter Photo Working Group, Boeing systems integration analysts prepared a preliminary analysis that afternoon. (Initial estimates of debris size and speed, origin of debris, and point of impact would later prove remarkably accurate.)

As Flight Day Three and Four unfolded over the Martin Luther King Jr. holiday weekend, engineers began their analysis. One Boeing analyst used Crater, a mathematical prediction tool, to assess possible damage to the Thermal Protection System. Analysis predicted tile damage deeper than the actual tile depth, and penetration of the RCC coating at impact angles above 15 degrees. This suggested the potential for a burn-through during re-entry. Debris Assessment Team members judged that the actual damage would not be as severe as predicted because of the inherent conservatism in the Crater model and because, in the case of tile, Crater does not take into account the tile's stronger and more impact-resistant “densified” layer, and in the case of RCC, the lower density of foam would preclude penetration at impact angles under 21 degrees.

On Flight Day Five, impact assessment results for tile and RCC were presented at an informal meeting of the Debris Assessment Team, which was operating without direct Shuttle Program or Mission Management leadership. Mission Control's engineering support, the Mission Evaluation Room, provided no direction for team activities other than to request the team's results by January 24. As the problem was being worked, Shuttle managers did not formally direct the actions of or consult with Debris Assessment Team leaders about the team's assumptions, uncertainties, progress, or interim results, an unusual circumstance given that NASA managers are normally engaged in analyzing what they view as problems. At this meeting, participants agreed that an image of the area of the wing in question was essential to refine their analysis and reduce the uncertainties in their damage assessment.

Each member supported the idea to seek imagery from an outside source. Due in part to a lack of guidance from the Mission Management Team or Mission Evaluation Room managers, the Debris Assessment Team chose an unconventional route for its request. Rather than working the request up the normal chain of command – through the Mission Evaluation Room to the Mission Management Team for action to Mission Control – team members nominated Rodney Rocha, the team's Co-Chair, to pursue the request through the Engineering Directorate at Johnson Space Center. As a result, even after the accident the Debris Assessment Team's request was viewed by Shuttle Program managers as a non-critical engineering desire rather than a critical operational need.

When the team learned that the Mission Management Team was not pursuing on-orbit imaging, members were concerned. What Debris Assessment Team members did not realize was the negative response from the Program was not necessarily a direct and final response to their official request. Rather, the "no" was in part a response to requests for imagery initiated by the Intercenter Photo Working Group at Kennedy on Flight Day 2 in anticipation of analysts' needs that had become by Flight Day 6 an actual engineering request by the Debris Assessment Team, made informally through Bob White to Lambert Austin, and formally through Rodney Rocha's e-mail to Paul Shack. Even after learning that the Shuttle Program was not going to provide the team with imagery, some members sought information on how to obtain it anyway.

Debris Assessment Team members believed that imaging of potentially damaged areas was necessary even after the January 24, Mission Management Team meeting, where they had reported their results. Why they did not directly approach Shuttle Program managers and share their concern and uncertainty, and why Shuttle Program managers claimed to be isolated from engineers, are points that the Board labored to understand. Several reasons for this communications failure relate to NASA's internal culture and the climate established by Shuttle Program management, which are discussed in more detail in Chapters 7 and 8.

A Flawed Analysis

An inexperienced team, using a mathematical tool that was not designed to assess an impact of this estimated size, performed the analysis of the potential effect of the debris impact. Crater was designed for "in-family" impact events and was intended for day-of-launch analysis of debris impacts. It was not intended for large projectiles like those observed on STS-107. Crater initially predicted possible damage, but the Debris Assessment Team assumed, without theoretical or experimental validation, that because Crater is a conservative tool – that is, it predicts more damage than will actually occur – the debris would stop at the tile's densified layer, even though their experience did not involve debris strikes as large as STS-107's. Crater-like equations were also used as part of the analysis to assess potential impact damage to the wing leading edge RCC. Again, the tool was used for something other than that for which it was designed; again, it predicted possible penetration; and again, the Debris Assessment Team used engineering arguments and their experience to discount the results.

As a result of a transition of responsibility for Crater analysis from the Boeing Huntington

Beach facility to the Houston-based Boeing office, the team that conducted the Crater analyses had been formed fairly recently, and therefore could be considered less experienced when compared with the more senior Huntington Beach analysts. In fact, STS-107 was the first mission for which they were solely responsible for providing analysis with the Crater tool. Though post-accident interviews suggested that the training for the Houston Boeing analysts was of high quality and adequate in substance and duration, communications and theoretical understandings of the Crater model among the Houston-based team members had not yet developed to the standard of a more senior team. Due in part to contractual arrangements related to the transition, the Houston-based team did not take full advantage of the Huntington Beach engineers' experience.

At the January 24, Mission Management Team meeting at which the "no safety-of-flight" conclusion was presented, there was little engineering discussion about the assumptions made, and how the results would differ if other assumptions were used.

Engineering solutions presented to management should have included a quantifiable range of uncertainty and risk analysis. Those types of tools were readily available, routinely used, and would have helped management understand the risk involved in the decision. Management, in turn, should have demanded such information. The very absence of a clear and open discussion of uncertainties and assumptions in the analysis presented should have caused management to probe further.

Shuttle Program Management's Low Level of Concern

While the debris strike was well outside the activities covered by normal mission flight rules, Mission Management Team members and Shuttle Program managers did not treat the debris strike as an issue that required operational action by Mission Control. Program managers, from Ron Dittemore to individual Mission Management Team members, had, over the course of the Space Shuttle Program, gradually become inured to External Tank foam losses and on a fundamental level did not believe foam striking the vehicle posed a critical threat to the Orbiter. In particular, Shuttle managers exhibited a belief that RCC panels are impervious to foam impacts. Even after seeing the video of Columbia's debris impact, learning estimates of the size and location of the strike, and noting that a foam strike with sufficient kinetic energy could cause Thermal Protection System damage, management's level of concern did not change.

The opinions of Shuttle Program managers and debris and photo analysts on the potential severity of the debris strike diverged early in the mission and continued to diverge as the mission progressed, making it increasingly difficult for the Debris Assessment Team to have their concerns heard by those in a decision-making capacity. In the face of Mission managers' low level of concern and desire to get on with the mission, Debris Assessment Team members had to prove unequivocally that a safety-of-flight issue existed before Shuttle Program management would move to obtain images of the left wing. The engineers found themselves in the unusual position of having to prove that the situation was unsafe – a reversal of the usual requirement to prove that a situation is safe.

Other factors contributed to Mission management's ability to resist the Debris Assessment Team's concerns. A tile expert told managers during frequent consultations that strike damage was only a maintenance-level concern and that on-orbit imaging of potential wing damage was not necessary. Mission management welcomed this opinion and sought no others. This constant reinforcement of managers' pre-existing beliefs added another block to the wall between decision makers and concerned engineers.

Another factor that enabled Mission management's detachment from the concerns of their own engineers is rooted in the culture of NASA itself. The Board observed an unofficial hierarchy among NASA programs and directorates that hindered the flow of communications. The effects of this unofficial hierarchy are seen in the attitude that members of the Debris Assessment Team held. Part of the reason they chose the institutional route for their imagery request was that without direction from the Mission Evaluation Room and Mission

Management Team, they felt more comfortable with their own chain of command, which was outside the Shuttle Program. Further, when asked by investigators why they were not more vocal about their concerns, Debris Assessment Team members opined that by raising contrary points of view about Shuttle mission safety, they would be singled out for possible ridicule by their peers and managers.

A Lack of Clear Communication

Communication did not flow effectively up to or down from Program managers. As it became clear during the mission that managers were not as concerned as others about the danger of the foam strike, the ability of engineers to challenge those beliefs greatly diminished. Managers' tendency to accept opinions that agree with their own dams the flow of effective communications.

After the accident, Program managers stated privately and publicly that if engineers had a safety concern, they were obligated to communicate their concerns to management. Managers did not seem to understand that as leaders they had a corresponding and perhaps greater obligation to create viable routes for the engineering community to express their views and receive information. This barrier to communications not only blocked the flow of information to managers, but it also prevented the downstream flow of information from managers to engineers, leaving Debris Assessment Team members no basis for understanding the reasoning behind Mission Management Team decisions.

The January 27 to January 31, phone and e-mail exchanges, primarily between NASA engineers at Langley and Johnson, illustrate another symptom of the "cultural fence" that impairs open communications between mission managers and working engineers. These exchanges and the reaction to them indicated that during the evaluation of a mission contingency, the Mission Management Team failed to disseminate information to all system and technology experts who could be consulted. Issues raised by two Langley and Johnson engineers led to the development of "what-if" landing scenarios of the potential outcome if the main landing gear door sustained damaged. This led to behind-the-scenes networking by these engineers to use NASA facilities to make simulation runs of a compromised landing configuration. These engineers – who understood their systems and related technology – saw the potential for a problem on landing and ran it down in case the unthinkable occurred. But their concerns never reached the managers on the Mission Management Team that had operational control over Columbia.

A Lack of Effective Leadership

The Shuttle Program, the Mission Management Team, and through it the Mission Evaluation Room, were not actively directing the efforts of the Debris Assessment Team. These management teams were not engaged in scenario selection or discussions of assumptions and did not actively seek status, inputs, or even preliminary results from the individuals charged with analyzing the debris strike. They did not investigate the value of imagery, did not intervene to consult the more experienced Crater analysts at Boeing's Huntington Beach facility, did not probe the assumptions of the Debris Assessment Team's analysis, and did not consider actions to mitigate the effects of the damage on re-entry. Managers' claims that they didn't hear the engineers' concerns were due in part to their not asking or listening.

The Failure of Safety's Role

As will be discussed in Chapter 7, safety personnel were present but passive and did not serve as a channel for the voicing of concerns or dissenting views. Safety representatives attended meetings of the Debris Assessment Team, Mission Evaluation Room, and Mission Management Team, but were merely party to the analysis process and conclusions instead of an independent source of questions and challenges. Safety contractors in the Mission Evaluation Room were only marginally aware of the debris strike analysis. One contractor did question the Debris Assessment Team safety representative about the analysis and was told that it was adequate. No additional inquiries were made. The highest-ranking safety representative at NASA headquarters deferred to Program managers when asked for an opinion on imaging of Columbia. The safety manager he spoke to also failed to follow up.

Summary

Management decisions made during Columbia's final flight reflect missed opportunities, blocked or ineffective communications channels, flawed analysis, and ineffective leadership. Perhaps most striking is the fact that management – including Shuttle Program, Mission Management Team, Mission Evaluation Room, and Flight Director and Mission Control – displayed no interest in understanding a problem and its implications. Because managers failed to avail themselves of the wide range of expertise and opinion necessary to achieve the best answer to the debris strike question – "Was this a safety-of-flight concern?" – some Space Shuttle Program managers failed to fulfill the implicit contract to do whatever is possible to ensure the safety of the crew. In fact, their management techniques unknowingly imposed barriers that kept at bay both engineering concerns and dissenting views, and ultimately helped create "blind spots" that prevented them from seeing the danger the foam strike posed.

Because this chapter has focused on key personnel who participated in STS-107 bipod foam debris strike decisions, it is tempting to conclude that replacing them will solve all NASA's problems. However, solving NASA's problems is not quite so easily achieved. Peoples' actions are influenced by the organizations in which they work, shaping their choices in directions that even they may not realize.

6.2 SCHEDULE PRESSURE

Countdown to Space Station "Core Complete:" A Workforce Under Pressure

During the course of this investigation, the Board received several unsolicited comments from NASA personnel regarding pressure to meet a schedule. These comments all concerned a date, more than a year after the launch of Columbia, that seemed etched in stone: February 19, 2004, the scheduled launch date of STS-120. This flight was a milestone in the minds of NASA management since it would carry a section of the International Space Station called "Node 2." This would configure the International Space Station to its "U.S. Core Complete" status.

At first glance, the Core Complete configuration date seemed noteworthy but unrelated to the Columbia accident. However, as the investigation continued, it became apparent that the complexity and political mandates surrounding the International Space Station Program, as well as Shuttle Program management's responses to them, resulted in pressure to meet an increasingly ambitious launch schedule.

In mid-2001, NASA adopted plans to make the over-budget and behind-schedule International Space Station credible to the White House and Congress. The Space Station Program and NASA were on probation, and had to prove they could meet schedules and budgets. The plan to regain credibility focused on the February 19, 2004, date for the launch of Node 2 and the resultant Core Complete status. If this goal was not met, NASA would risk losing support from the White House and Congress for subsequent Space Station growth.

By the late summer of 2002, a variety of problems caused Space Station assembly work and Shuttle flights to slip beyond their target dates. With the Node 2 launch endpoint fixed, these delays caused the schedule to become ever more compressed.

Meeting U.S. Core Complete by February 19, 2004, would require preparing and launching 10 flights in less than 16 months. With the focus on retaining support for the Space Station program, little attention was paid to the effects the aggressive Node 2 launch date would have on the Shuttle Program. After years of downsizing and budget cuts (Chapter 5), this mandate and events in the months leading up to STS-107 introduced elements of risk to the Program. Columbia and the STS-107 crew, who had seen numerous launch slips due to missions that were deemed higher priorities, were further affected by the mandatory Core Complete date. The high-pressure environments created by NASA Headquarters unquestionably affected Columbia, even though it was not flying to the International Space Station.

February 19, 2004 "A Line in the Sand"

Schedules are essential tools that help large organizations effectively manage their resources. Aggressive schedules by themselves are often a sign of a healthy institution. However, other institutional goals, such as safety, sometimes compete with schedules, so the effects of schedule pressure in an organization must be carefully monitored. The Board posed the question: Was there undue pressure to nail the Node 2 launch date to the February 19, 2004, signpost? The management and workforce of the Shuttle and Space Station programs each answered the question differently. Various members of NASA upper management gave a definite "no." In contrast, the workforce within both programs thought there was considerable management focus on Node 2 and resulting pressure to hold firm to that launch date, and individuals were becoming concerned that safety might be compromised. The weight of evidence supports the workforce view.

Employees attributed the Node 2 launch date to the new Administrator, Sean O'Keefe, who was appointed to execute a Space Station management plan he had proposed as Deputy Director of the White House Office of Management and Budget. They understood the scrutiny that NASA, the new Administrator, and the Space Station Program were under, but now it seemed to some that budget and schedule were of paramount concern. As one employee reflected:

> I guess my frustration was … I know the importance of showing that you … manage your budget and that's an important impression to make to Congress so you can continue the future of the agency, but to a lot of people, February 19th just seemed like an arbitrary date … It doesn't make sense to me why at all costs we were marching to this date.

The importance of this date was stressed from the very top. The Space Shuttle and Space Station Program Managers briefed the new NASA Administrator monthly on the status of their programs, and a significant part of those briefings was the days of margin remaining in the schedule to the launch of Node 2, still well over a year away. The Node 2 schedule margin typically accounted for more than half of the briefing slides.

NASA Headquarters stressed the importance of this date in other ways. A screen saver) was mailed to managers in NASA's human spaceflight program that depicted a clock counting down to February 19, 2004, U.S. Core Complete.

While employees found this amusing because they saw it as a date that could not be met, it also reinforced the message that NASA Headquarters was focused on and promoting the achievement of that date. This schedule was on the minds of the Shuttle managers in the months leading up to STS-107.

CHAPTER 7.4 ORGANIZATIONAL CAUSES: A BROKEN SAFETY CULTURE

Safety Information Systems

Numerous reviews and independent assessments have noted that NASA's safety system does not effectively manage risk. In particular, these reviews have observed that the processes in which NASA tracks and attempts to mitigate the risks posed by components on its Critical Items List is flawed. The Post Challenger Evaluation of
Space Shuttle Risk Assessment and Management Report (1988) concluded that:

> The committee views NASA critical items list (CIL) waiver decision-making process as being subjective, with little in the way of formal and consistent criteria for approval or rejection of waivers. Waiver decisions appear to be driven almost exclusively by the design based Failure Mode Effects Analysis (FMEA)/CIL retention rationale, rather than being based on an integrated assessment of all inputs to risk management. The retention rationales appear biased toward proving that the design is "safe," sometimes ignoring significant evidence to the contrary.

The report continues, "… the Committee has not found an independent, detailed analysis or assessment of the CIL retention rationale which considers all inputs to the risk assessment process." Ten years later, the Shuttle Independent Assessment Team reported "Risk Management process erosion created by the desire to reduce costs …"The Shuttle Independent Assessment Team argued strongly that NASA Safety and Mission Assurance should be restored to its previous role of an independent oversight body, and Safety and Mission Assurance not be simply a "safety auditor."

The Board found similar problems with integrated hazard analyses of debris strikes on the Orbiter. In addition, the information systems supporting the Shuttle intended to be tools for decision-making are extremely cumbersome and difficult to use at any level.

The following addresses the hazard tracking tools and major databases in the Shuttle Program that promote risk management.

- **Hazard Analysis**: A fundamental element of system safety is managing and controlling hazards. NASA's only guidance on hazard analysis is outlined in the Methodology for Conduct of Space Shuttle Program Hazard Analysis, which merely lists tools available. Therefore, it is not surprising that hazard analysis processes are applied inconsistently across systems, sub-systems, assemblies, and components.

  United Space Alliance, which is responsible for both Orbiter integration and Shuttle Safety Reliability and Quality Assurance, delegates hazard analysis to Boeing. However, as of 2001, the Shuttle Program no longer requires Boeing to conduct integrated hazard analyses. Instead, Boeing now performs hazard analysis only at the sub-system level. In other words, Boeing analyzes hazards to components and elements, but is not required to consider the Shuttle as a whole. Since the current Failure Mode Effects Analysis/Critical Item List process is designed for bottom-up analysis at the component level, it cannot effectively support the kind of "top-down" hazard analysis that is needed to inform managers on risk trends and identify potentially harmful interactions between systems.

  The Critical Item List (CIL) tracks 5,396 individual Shuttle hazards, of which 4,222 are termed "Criticality 1/1R." Of those, 3,233 have waivers. CRIT 1/1R component failures are defined as those that will result in loss of the Orbiter and crew. Waivers are granted whenever a Critical Item List component cannot be redesigned or replaced. More than 36 percent of these waivers have not been reviewed in 10 years, a sign that NASA is not aggressively monitoring changes in system risk.

  It is worth noting that the Shuttle's Thermal Protection System is on the Critical Item List, and an existing hazard analysis and hazard report deals with debris strikes. As discussed in Chapter 6, Hazard Report #37 is ineffectual as a decision aid, yet the Shuttle Program never challenged its validity at the pivotal STS-113 Flight Readiness Review.

  Although the Shuttle Program has undoubtedly learned a great deal about the technological limitations inherent in Shuttle operations, it is equally clear that risk, as represented by the number of critical items list and waivers has grown substantially without a vigorous effort to assess and reduce technical problems that increase risk. An information system bulging with over 5,000 critical items and 3,200 waivers is exceedingly difficult to manage.

- **Hazard Reports**: Hazard reports, written either by the Space Shuttle Program or a contractor, document conditions that threaten the safe operation of the Shuttle. Managers use these reports to evaluate risk and justify flight. During mission preparations, contractors and Centers review all baseline hazard reports to ensure they are current and technically correct.

Board investigators found that a large number of hazard reports contained subjective and qualitative judgments, such as "believed" and "based on experience from previous flights this hazard is an 'Accepted Risk.' " A critical ingredient of a healthy safety program is the rigorous implementation of technical standards. These standards must include more than hazard analysis or low-level technical activities. Standards must integrate project engineering and management activities. Finally, a mechanism for feedback on the effectiveness of system safety engineering and management needs to be built into procedures to learn if safety engineering and management methods are weakening over time.

Dysfunctional Databases

In its investigation, the Board found that the information systems that support the Shuttle program are extremely cumbersome and difficult to use in decision-making at any level. For obvious reasons, these shortcomings imperil the Shuttle Program's ability to disseminate and share critical information among its many layers. This section explores the report databases that are crucial to effective risk management.

- Problem Reporting and Corrective Action: The Problem Reporting and Corrective Action database records any non-conformances (instances in which a requirement is not met). Formerly, different Centers and contractors used the Problem Reporting and Corrective Action database differently, which prevented comparisons across the database. NASA recently initiated an effort to integrate these databases to permit anyone in the agency to access information from different Centers. This system, Web Program Compliance Assurance and Status System (WEBPCASS), is supposed to provide easier access to consolidated information and facilitates higher-level searches.

  However, NASA safety managers have complained that the system is too time-consuming and cumbersome. Only employees trained on the database seem capable of using WEBPCASS effectively. One particularly frustrating aspect of which the Board is acutely aware is the database's waiver section. It is a critical information source, but only the most expert users can employ it effectively. The database is also incomplete. For instance, in the case of foam strikes on the Thermal Protection System, only strikes that were declared "In-Fight Anomalies" are added to the Problem Reporting and Corrective Action database, which masks the full extent of the foam debris trends.

- **Lessons Learned Information System**: The Lessons Learned Information System database is a much simpler system to use, and it can assist with hazard identification and risk assessment. However, personnel familiar with the Lessons Learned Information System indicate that design engineers and mission assurance personnel use it only on an ad hoc basis, thereby limiting its utility. The Board is not the first to note such deficiencies. Numerous reports, including most recently a General Accounting Office 2001 report, highlighted fundamental weaknesses in the collection and sharing of lessons learned by program and project managers.

Conclusions

Throughout the course of this investigation, the Board found that the Shuttle Program's complexity demands highly effective communication. Yet integrated hazard reports and risk analyses are rarely communicated effectively, nor are the many databases used by Shuttle Program engineers and managers capable of translating operational experiences into effective risk management practices. Although the Space Shuttle system has conducted a relatively small number of missions, there is more than enough data to generate performance trends. As it is currently structured, the Shuttle Program does not use data-driven safety methodologies to their fullest advantage.

7.5 ORGANIZATIONAL CAUSES: IMPACT OF A FLAWED SAFETY CULTURE ON STS-107

In this section, the Board examines how and why an array of processes, groups, and individuals in the Shuttle Program failed to appreciate the severity and implications of the foam strike on STS-107. The Board believes that the Shuttle Program should have been able to detect the foam trend and more fully appreciate the danger it represented. Recall that "safety culture" refers to the collection of characteristics and attitudes in an organization promoted by its leaders and internalized by its members that makes safety an overriding priority. In the following analysis, the Board outlines shortcomings in the Space Shuttle Program, Debris Assessment Team, and Mission Management Team that resulted from a flawed safety culture.

Shuttle Program Shortcomings

The flight readiness process, which involves every organization affiliated with a Shuttle mission, missed the danger signals in the history of foam loss.

Generally, the higher information is transmitted in a hierarchy, the more it gets "rolled-up," abbreviated, and simplified. Sometimes information gets lost altogether, as weak signals drop from memos, problem identification systems, and formal presentations. The same conclusions, repeated over time, can result in problems eventually being deemed non-problems. An extraordinary example of this phenomenon is how Shuttle Program managers assumed the foam strike on STS-112 was not a warning sign.

During the STS-113 Flight Readiness Review, the bipod foam strike to STS-112 was rationalized by simply restating earlier assessments of foam loss. The question of why bipod foam would detach and strike a Solid Rocket Booster spawned no further analysis or heightened curiosity; nor did anyone challenge the weakness of External Tank Project Manager's argument that backed launching the next mission. After STS-113's successful flight, once again the STS-112 foam event was not discussed at the STS-107 Flight Readiness Review. The failure to mention an outstanding technical anomaly, even if not technically a violation of NASA's own procedures, desensitized the Shuttle Program to the dangers of foam striking the Thermal Protection System, and demonstrated just how easily the flight preparation process can be compromised. In short, the dangers of bipod foam got "rolled-up," which resulted in a missed opportunity to make Shuttle managers aware that the Shuttle required, and did not yet have a fix for the problem.

Once the Columbia foam strike was discovered, the Mission Management Team Chairperson asked for the rationale the STS-113 Flight Readiness Review used to launch in spite of the STS-112 foam strike. In her e-mail, she admitted that the analysis used to continue flying was, in a word, "lousy" (Chapter 6). This admission that the rationale to fly was rubber-stamped is, to say the least, unsettling.

The Flight Readiness process is supposed to be shielded from outside influence, and is viewed as both rigorous and systematic. Yet the Shuttle Program is inevitably influenced by external factors, including, in the case of the STS-107, schedule demands. Collectively, such factors shape how the Program establishes mission schedules and sets budget priorities, which affects safety oversight, workforce levels, facility maintenance, and contractor workloads. Ultimately, external expectations and pressures impact even data collection, trend analysis, information development, and the reporting and disposition of anomalies. These realities contradict NASA's optimistic belief that pre-flight reviews provide true safeguards against unacceptable hazards. The schedule pressure to launch International Space Station Node 2 is a powerful example of this point.

The premium placed on maintaining an operational schedule, combined with ever-decreasing resources, gradually led Shuttle managers and engineers to miss signals of potential danger. Foam strikes on the Orbiter's Thermal Protection System, no matter what the size of the debris, were "normalized" and accepted as not being a "safety-of-flight risk." Clearly, the risk of Thermal Protection damage due to such a strike needed to be better understood in quantifiable terms. External Tank foam loss should have been eliminated or mitigated with redundant layers of protection. If there was in fact a strong safety culture at NASA, safety experts would have had the authority to test the actual resilience of the leading edge Reinforced Carbon-Carbon panels, as the Board has done.

Debris Assessment Team Shortcomings

Chapter Six details the Debris Assessment Team's efforts to obtain additional imagery of Columbia. When managers in the Shuttle Program denied the team's request for imagery, the Debris Assessment Team was put in the untenable position of having to prove that a safety-of-flight issue existed without the very images that would permit such a determination. This is precisely the opposite of how an effective safety culture would act. Organizations that deal with high-risk operations must always have a healthy fear of failure operations must be proved safe, rather than the other way around. NASA inverted this burden of proof.

Another crucial failure involves the Boeing engineers who conducted the Crater analysis. The Debris Assessment Team relied on the inputs of these engineers along with many others to assess the potential damage caused by the foam strike. Prior to STS-107, Crater analysis was the responsibility of a team at Boeing's Huntington Beach facility in California, but this responsibility had recently been transferred to Boeing's Houston office. In October 2002, the Shuttle Program completed a risk assessment that predicted the move of Boeing functions from Huntington Beach to Houston would increase risk to Shuttle missions through the end of 2003, because of the small number of experienced engineers who werewilling to relocate. To mitigate this risk, NASA and United Space Alliance developed a transition plan to run through January 2003.

The Board has discovered that the implementation of the transition plan was incomplete and that training of replacement personnel was not uniform. STS-107 was the first mission during which Johnson-based Boeing engineers conducted analysis without guidance and oversight from engineers at Huntington Beach.

Even though STS-107's debris strike was 400 times larger than the objects Crater is designed to model, neither Johnson engineers nor Program managers appealed for assistance from the more experienced Huntington Beach engineers, who might have cautioned against using Crater so far outside its validated limits. Nor did safety personnel provide any additional oversight. NASA failed to connect the dots: the engineers who misinterpreted Crater, a tool already unsuited to the task at hand were the very ones the Shuttle Program identified as engendering the most risk in their transition from Huntington Beach. The Board views this example as characteristic of the greater turbulence the Shuttle Program experienced in the decade before Columbia as a result of workforce reductions and management reforms.

Mission Management Team Shortcomings

In the Board's view, the decision to fly STS-113 without a compelling explanation for why bipod foam had separated on ascent during the preceding mission, combined with the low number of Mission Management Team meetings during STS-107, indicates that the Shuttle Program had become overconfident. Over time, the organization determined it did not need daily meetings during a mission, despite regulations that state otherwise.

Status update meetings should provide an opportunity to raise concerns and hold discussions across structural and technical boundaries. The leader of such meetings must encourage participation and guarantee that problems are assessed and resolved fully. All voices must be heard, which can be difficult when facing a hierarchy. An employee's location in the hierarchy can encourage silence. Organizations interested in safety must take steps to guarantee that all relevant information is presented to decision-makers. This did not happen in the meetings during the Columbia mission (see Chapter 6). For instance, e-mails from engineers at Johnson and Langley conveyed the depth of their concern about the foam strike, the questions they had about its implications, and the actions they wanted to take as a follow-up. However, these e-mails did not reach the Mission Management Team.

The failure to convey the urgency of engineering concerns was caused, at least in part, by organizational structure and spheres of authority. The Langley e-mails were circulated among co-workers at Johnson who explored the possible effects of the foam strike and its consequences for landing. Yet, like Debris Assessment Team Co-Chair Rodney Rocha, they kept their concerns within local channels and did not forward them to the Mission Management Team. They were separated from the decision-making process by distance and rank.

Similarly, Mission Management Team participants felt pressured to remain quiet unless discussion turned to their particular area of technological or system expertise, and, even then, to be brief. The initial damage assessment briefing prepared for the Mission Evaluation Room was cut down considerably in order to make it “fit” the schedule. Even so, it took 40 minutes. It was cut down further to a three-minute discussion topic at the Mission Management Team. Tapes of STS-107 Mission Management Team sessions reveal a noticeable “rush” by the meeting’s leader to the preconceived bottom line that there was “no safety-of-flight” issue (see Chapter 6). Program managers created huge barriers against dissenting opinions by stating preconceived conclusions based on subjective knowledge and experience, rather than on solid data. Managers demonstrated little concern for mission safety.

Organizations with strong safety cultures generally acknowledge that a leader’s best response to unanimous consent is to play devil’s advocate and encourage an exhaustive debate. Mission Management Team leaders failed to seek out such minority opinions. Imagine the difference if any Shuttle manager had simply asked, “Prove to me that Columbia has not been harmed.”

Similarly, organizations committed to effective communication seek avenues through which unidentified concerns and dissenting insights can be raised, so that weak signals are not lost in background noise. Common methods of bringing minority opinions to the fore include hazard reports, suggestion programs, and empowering employees to call “time out” (Chapter 10). For these methods to be effective, they must mitigate the fear of retribution, and management and technical staff must pay attention. Shuttle Program hazard reporting is seldom used, safety time outs are at times disregarded, and informal efforts to gain support are squelched. The very fact that engineers felt inclined to conduct simulated blown tire landings at Ames “after hours,” indicates their reluctance to bring the concern up in established channels.

Safety Shortcomings

The Board believes that the safety organization, due to a lack of capability and resources independent of the Shuttle Program, was not an effective voice in discussing technical issues or mission operations pertaining to STS-107. The safety personnel present in the Debris Assessment Team, Mission Evaluation Room, and on the Mission Management Team were largely silent during the events leading up to the loss of Columbia. That silence was not merely a failure of safety, but a failure of the entire organization.

CHAPTER 5

FROM CHALLENGER TO COLUMBIA

The Board is convinced that the factors that led to the Columbia accident go well beyond the physical mechanisms discussed in Chapter 3. The causal roots of the accident can also be traced, in part, to the turbulent post-Cold War policy environment in which NASA functioned during most of the years between the destruction of *Challenger* and the loss of *Columbia*. The end of the Cold War in the late 1980s meant that the most important political underpinning of NASAs Human Space Flight Program, U.S.-Soviet space competition was lost, with no equally strong political objective to replace it. No longer able to justify its projects with the kind of urgency that the superpower struggle had provided, the agency could not obtain budget increases through the 1990s. Rather than adjust its ambitions to this new state of affairs, NASA continued to push an ambitious agenda of space science and exploration, including a costly Space Station Program.

If NASA wanted to carry out that agenda, its only recourse, given its budget allocation, was to become more efficient, accomplishing more at less cost. The search for cost reductions led top NASA leaders over the past decade to downsize the Shuttle workforce, outsource various Shuttle Program responsibilities including safety oversight and consider eventual privatization of the Space Shuttle Program. The programs budget was reduced by 40 percent in purchasing power over the past decade and repeatedly raided to make up for Space Station cost overruns, even as the Program maintained a launch schedule in which the Shuttle, a developmental vehicle, was used in an operational mode. In addition, the uncertainty of top policymakers in the White House, Congress, and NASA as to how long the Shuttle would fly before being replaced resulted in the delay of upgrades needed to make the Shuttle safer and to extend its service life.

The Space Shuttle Program has been transformed since the late 1980s implementation of post-*Challenger* management changes in ways that raise questions, addressed here and in later chapters of Part Two, about NASAs ability to safely operate the Space Shuttle. While it would be inaccurate to say that NASA managed the Space Shuttle Program at the time of the *Columbia* accident in the same manner it did prior to *Challenger*, there are unfortunate similarities between the agencys performance and safety practices in both periods.

5.1 THE *CHALLENGER* ACCIDENT AND ITS AFTERMATH

The inherently vulnerable design of the Space Shuttle, described in Chapter 1, was a product of policy and technological compromises made at the time of its approval in 1972. That approval process also produced unreasonable expectations, even myths, about the Shuttles future performance that NASA tried futilely to fulfill as the Shuttle became "operational" in 1982. At first, NASA was able to maintain the image of the Shuttle as an operational vehicle. During its early years of operation, the Shuttle launched satellites, performed on-orbit research, and even took members of Congress into orbit. At the beginning of 1986, the goal of "routine access to space" established by President Ronald Reagan in 1982 was ostensibly being achieved. That appearance soon proved illusory. On the cold morning of January 28, 1986, the Shuttle Challenger broke apart 73 seconds into its climb towards orbit. On board were Francis R. Scobee, Michael J. Smith, Ellison S. Onizuka, Judith A. Resnick, Ronald E. McNair, Sharon Christa McAuliffe, and Gregory B. Jarvis. All perished.

Space Flight Operations Contract

By the middle of the decade (1990's), spurred on by Vice President Al Gore's "reinventing government" initiative, the goal of balancing the federal budget, and the views of a Republican led House of Representatives, managers throughout the government sought new ways of making public sector programs more efficient and less costly. One method considered was transferring significant government operations and responsibilities to the private sector, or "privatization." NASA led the way toward privatization, serving as an example to other government agencies.

In keeping with his philosophy that NASA should focus on its research and development role, Goldin wanted to remove NASA employees from the repetitive operations of various systems, including the Space Shuttle. Giving primary responsibility for Space Shuttle operations to the private sector was therefore consistent with White House and congressional priorities and attractive to Goldin on its own terms. Beginning in 1994, NASA considered the feasibility of consolidating many of the numerous Shuttle operations contracts under a single prime contractor. At that time, the Space Shuttle Program was managing 86 separate contracts held by 56 different firms. Top NASA managers thought that consolidating these contracts could reduce the amount of redundant overhead, both for NASA and for the contractors themselves. They also wanted to explore whether there were functions being carried out by NASA that could be more effectively and inexpensively carried out by the private sector.

An advisory committee headed by early space flight veteran Christopher Kraft recommended such a step in its March 1995 report, which became known as the "Kraft Report." (The report characterized the Space Shuttle in a way that the

Board judges to be at odds with the realities of the Shuttle Program).

Although the Kraft Report stressed that the dramatic changes it recommended could be made without compromising safety, there was considerable dissent about this claim. NASA's Aerospace Safety Advisory Panel independent, but often not very influential was particularly critical. In May 1995, the Panel noted that "the assumption [in the Kraft Report] that the Space Shuttle systems are now 'mature' smacks of a complacency which may lead to serious mishaps. The fact is that the Space Shuttle may never be mature enough to totally freeze the design." The Panel also noted that "the report dismisses the concerns of many credible sources by labeling honest reservations and the people who have made them as being partners in an unneeded 'safety shield' conspiracy. Since only one more accident would kill the program and destroy far more than the spacecraft, it is extremely callous" to make such an accusation.

NASA leaders accepted the advice of the Kraft Report and in August 1995 solicited industry bids for the assignment of Shuttle prime contractor. In response, Lockheed Martin and Rockwell, the two major Space Shuttle operations contractors, formed a limited liability corporation, with each firm a 50 percent owner, to compete for what was called the Space Flight Operations Contract. The new corporation would be known as United Space Alliance.

The company was responsible for 61 percent of the Shuttle operations contracts. Some in Congress were skeptical that safety could be maintained under the new arrangement, which transferred significant NASA responsibilities to the private sector. Despite these concerns, Congress ultimately accepted the reasoning behind the contract.

The Space Flight Operations Contract was designed to reward United Space Alliance for performance successes and penalize its performance failures. Before being eligible for any performance fees, United Space Alliance would have to meet a series of safety "gates," which were intended to ensure that safety remained the top priority in Shuttle operations. The contract also rewarded any cost reductions that United Space Alliance was able to achieve, with NASA taking 65 percent of any savings and United Space Alliance 35 percent.

NASA and United Space Alliance formally signed the Space Flight Operations Contract on October 1, 1996. Initially, only the major Lockheed Martin and Rockwell Shuttle contracts and a smaller Allied Signal Unisys contract were transferred to United Space Alliance. The initial contractual period was six years, from October 1996 to September 2002. NASA exercised an option for a two year extension in 2002, and another two year option exists. The total value of the contract through the current extension is estimated at $12.8 billion. United Space Alliance currently has approximately 10,000 employees.

Workforce Transformation and the End of Downsizing

Workforce reductions instituted by Administrator Goldin as he attempted to redefine the agency's mission and its overall organization also added to the turbulence of his reign. In the 1990s, the overall NASA workforce was reduced by 25 percent through normal attrition, early retirements, and buyouts- cash bonuses for leaving NASA employment. NASA operated under a hiring freeze for most of the decade, making it difficult to bring in new or younger people.

NASA Headquarters was particularly affected by workforce reductions. More than half its employees left or were transferred in parallel with the 1996 transfer of program management responsibilities back to the NASA centers. The Space Shuttle Program bore more than its share of Headquarters personnel cuts. Headquarters civil service staff working on the Space Shuttle Program went from 120 in 1993 to 12 in 2003.

While the overall workforce at the NASA Centers involved in human space flight was not as radically reduced, the combination of the general workforce reduction and the introduction of the Space Flight Operations Contract significantly impacted the Centers' Space Shuttle Program civil service staff. Johnson Space Center went from 1,330 in 1993 to 738 in 2002; Marshall Space Flight Center, from 874 to 337; and Kennedy Space Center from 1,373 to 615.

Kennedy Director Roy Bridges argued that personnel cuts were too deep, and threatened to resign unless the downsizing of his civil service workforce, particularly those involved with safety issues, was reversed.

By the end of the decade, NASA realized that staff reductions had gone too far. By early 2000, internal and external studies convinced NASA leaders that the workforce needed to be revitalized. These studies noted that "five years of buyouts and downsizing have led to serious skill imbalances and an overtaxed core workforce. As more employees have departed, the workload and stress [on those] remaining have increased, with a corresponding increase in the potential for impacts to operational capacity and safety." NASA announced that NASA workforce downsizing would stop short of the 17,500 target, and that its human space flight centers would immediately hire several hundred workers.

5.5 WHEN TO REPLACE THE SPACE SHUTTLE?

In addition to budget pressures, workforce reductions, management changes, and the transfer of government functions to the private sector, the Space Shuttle Program was beset during the past decade by uncertainty about when the Shuttle might be replaced. National policy has vacillated between treating the Shuttle as a "going out of business" program and anticipating two or more decades of Shuttle use. As a result, limited and inconsistent investments have been made in Shuttle upgrades and in revitalizing the infrastructure to support the continued use of the Shuttle.

Even before the 1986 Challenger accident, when and how to replace the Space Shuttle with a second generation reusable launch vehicle was a topic of discussion among space policy leaders. In January 1986, the congressionally chartered National Commission on Space expressed the need for a Shuttle replacement, suggesting that "the Shuttle fleet will become obsolescent by the turn of the century." Shortly after the Challenger accident (but not as a reaction to it), President Reagan announced his approval of "the new Orient Express". This reusable launch vehicle, later known as the National Aerospace Plane, "could, by the end of the decade, take off from Dulles Airport, accelerate up to 25 times the speed of sound attaining low-Earth orbit, or fly to Tokyo within two hours." This goal proved too ambitious, particularly without substantial funding. In 1992, after a $1.7 billion government investment, the National Aerospace Plane project was canceled.

This pattern, optimistic pronouncements about a revolutionary Shuttle replacement followed by insufficient government investment, and then program cancellation due to technical difficulties was repeated again in the 1990s.

In 1994, NASA listed alternatives for access to space through 2030.

- Upgrade the Space Shuttle to enable flights through 2030
- Develop a new expendable launcher
- Replace the Space Shuttle with a "leapfrog" next-generation advanced technology system that would achieve order-of-magnitude improvements in the cost effectiveness of space transportation.

Reflecting its leadership's preference for bold initiatives, NASA chose the third alternative. With White House support, NASA began the X-33 project in 1996 as a joint effort with Lockheed Martin. NASA also initiated the less ambitious X-34 project with Orbital Sciences Corporation. At the time, the future of commercial space launches was bright, and political sentiment in the White House and Congress encouraged an increasing reliance on private-sector solutions for limiting government expenditures. In this context, these unprecedented joint projects appeared less risky than they actually were. The hope was that NASA could replace the Shuttle through private investments, without significant government spending.

Both the X-33 and X-34 incorporated new technologies. The X-33 was to demonstrate the feasibility of an aerospike engine, new Thermal Protection Systems, and composite rather than metal propellant tanks. These radically new technologies were in turn to become the basis for a new orbital vehicle called VentureStar™ that could replace the Space Shuttle by 2006. The X-33 and X-34 ran into technical problems and never flew. In 2001, after spending $1.3 billion, NASA abandoned both projects.

In all three projects National Aerospace Plane, X-33, and X-34 national leaders had set ambitious goals in response to NASA's ambitious proposals. These programs relied on the invention of revolutionary technology, had run into major technical problems, and had been denied the funds needed to overcome these problems assuming they could be solved. NASA had spent nearly 15 years and several billion dollars, and yet had made no meaningful progress toward a Space Shuttle replacement.

In 2000, as the agency ran into increasing problems with the X-33, NASA initiated the Space Launch Initiative, a $4.5 billion multi-year effort to develop new space launch technologies. By 2002, after spending nearly $800 million, NASA again changed course. The Space Launch Initiative failed to find technologies that could revolutionize space launch, forcing NASA to shift its focus to an Orbital Space Plane, developed with existing technology, that would complement the Shuttle by carrying crew, but not cargo, to and from orbit. Under a new Integrated Space Transportation Plan, the Shuttle might continue to fly until 2020 or beyond.

As a result of the haphazard policy process that created these still-born developmental programs, the uncertainty over Shuttle replacement persisted. Between 1986 and 2002, the planned replacement date for the Space Shuttle was consistent only in its inconsistency: it changed from 2002 to 2006 to 2012, and before the Columbia accident, to 2020 or later.

Safety Concerns and Upgrading the Space Shuttle

This shifting date for Shuttle replacement has severely complicated decisions on how to invest in Shuttle Program upgrades. More often than not, investments in upgrades were delayed or deferred on the assumption they would be a waste of money if the Shuttle were to be retired in the near future.

In 1995, for instance, the Kraft Report embraced the principle that NASA should "freeze the design" of the Shuttle and defer upgrades due to the vehicle's "mature" status and the need for NASA to "concentrate scarce resources on developing potential replacements for the Shuttle." NASA subsequently halted a number of planned upgrades, only to reverse course a year later to "take advantage of technologies to improve Shuttle safety and the need for a robust Space Shuttle to assemble the ISS."

In a June 1999 letter to the White House, NASA Administrator Daniel Goldin declared that the nation faced a "Space Launch Crisis." He reported on a NASA review of Shuttle safety that indicated the budget for Shuttle upgrades in Fiscal year 2000 was "inadequate to accommodate upgrades necessary to yield significant safety improvements." After two "close calls" during STS-93 in July 1999 Goldin also chartered a Shuttle Independent Assessment Team (SIAT) chaired by Harry McDonald, Director of NASA Ames Research Center. Among the team's findings, reported in March 2000:

- "Over the course of the Shuttle Program … processes, procedures and training have continuously been improved and implemented to make the system safer. The SIAT has a major concern … that this critical feature of the Shuttle Program is being eroded." The major factor leading to this concern "is the reduction in allocated resources and appropriate staff … There are important technical areas that are 'one-deep.' " Also, "the SIAT feels strongly that workforce augmentation must be realized principally with NASA personnel rather than with contractor personnel."
- The SIAT was concerned with "success-engendered safety optimism … The SSP must rigorously guard against the tendency to accept risk solely because of prior success."
- "The SIAT was very concerned with what it perceived as Risk Management process erosion created by the desire to reduce costs … The SIAT feels strongly that NASA Safety and Mission Assurance should be restored to its previous role of an independent oversight body and not be simply a 'safety auditor.' "

- “The size and complexity of the Shuttle system and of NASA/contractor relationships place extreme importance on understanding, communication, and information handling … Communication of problems and concerns upward to the SSP from the ‘floor’ also appeared to leave room for improvement.”

The Shuttle Independent Assessment Team report also 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.”

Responding to NASA’s concern that the Shuttle required safety-related upgrades, the President’s proposed NASA budget for Fiscal Year 2001 proposed a “safety upgrades initiative.” That initiative had a short life span. In its Fiscal Year 2002 budget request, NASA proposed to spend $1.836 billion on Shuttle upgrades over five years. A year later, the Fiscal Year 2003 request contained a plan to spend $1.220 billion, a 34 percent reduction. The reductions were primarily a response to rising Shuttle operating costs and the need to stay within a fixed Shuttle budget. Cost growth in Shuttle operations forced NASA to “use funds intended for Space Shuttle safety upgrades to address operational, supportability, obsolescence, and infrastructure needs.”

At its March 2001 meeting, NASA’s Space Flight Advisory Committee advised that “the Space Shuttle Program must make larger, more substantial safety upgrades than currently planned … a budget on the order of three times the budget currently allotted for improving the Shuttle systems” was needed. Later that year, five Senators complained that “the Shuttle program is being penalized, despite its outstanding performance, in order to conform to a budget strategy that is dangerously inadequate to ensure safety in America’s human space flight program.”

5.8 CONCLUSION

Over the last decade, the Space Shuttle Program has operated in a challenging and often turbulent environment. As discussed in this chapter, there were at least three major contributing factors to that environment:

- Throughout the decade, the Shuttle Program has had to function within an increasingly constrained budget. Both the Shuttle budget and workforce have been reduced by over 40 percent during the past decade. The White House, Congress, and NASA leadership exerted constant pressure to reduce or at least freeze operating costs. As a result, there was little margin in the budget to deal with unexpected technical problems or make Shuttle improvements.

- The Shuttle was mischaracterized by the 1995 Kraft Report as “a mature and reliable system … about as safe as today’s technology will provide.” Based on this mischaracterization, NASA believed that it could turn increased responsibilities for Shuttle operations over to a single prime contractor and reduce its direct involvement in ensuring safe Shuttle operations, instead monitoring contractor performance from a more detached position. NASA also believed that it could use the “mature” Shuttle to carry out operational missions without continually focusing engineering attention on understanding the mission-by-mission anomalies inherent in a developmental vehicle.

- In the 1990s, the planned date for replacing the Shuttle shifted from 2006 to 2012 and then to 2015 or later. Given the uncertainty regarding the Shuttle’s service life, there has been policy and budgetary ambivalence on investing in the vehicle. Only in the past year has NASA begun to provide the resources needed to sustain extended Shuttle operations. Previously, safety and support upgrades were delayed or deferred, and Shuttle infrastructure was allowed to deteriorate.

The Board observes that this is hardly an environment in which those responsible for safe operation of the Shuttle can function without being influenced by external pressures. It is to the credit of Space Shuttle managers and the Shuttle workforce that the vehicle was able to achieve its program objectives for as long as it did.

An examination of the Shuttle Program's history from Challenger to Columbia raises the question: Did the Space Shuttle Program budgets constrained by the White House and Congress threaten safe Shuttle operations? There is no straightforward answer. In 1994, an analysis of the Shuttle budget concluded that reductions made in the early 1990s represented a "healthy tightening up" of the program. Certainly those in the Office of Management and Budget and in NASA's congressional authorization and appropriations subcommittees thought they were providing enough resources to operate the Shuttle safely, while also taking into account the expected Shuttle lifetime and the many other demands on the Federal budget. NASA Headquarters agreed, at least until Administrator Goldin declared a "space launch crisis" in June 1999 and asked that additional resources for safety upgrades be added to the NASA budget. By 2001, however, one experienced observer of the space program described the Shuttle workforce as "The Few, the Tired," and suggested that "a decade of downsizing and budget tightening has left NASA exploring the universe with a less experienced staff and older equipment."

It is the Board's view that this latter statement is an accurate depiction of the Space Shuttle Program at the time of STS-107. The Program was operating too close to too many margins. The Board also finds that recent modest increases in the Shuttle Program's budget are necessary and overdue steps toward providing the resources to sustain the program for its now extended lifetime. Similarly, NASA has recently recognized that providing an adequately sized and appropriately trained workforce is critical to the agency's future success.

An examination of the Program's management changes also leads to the question: Did turmoil in the management structure contribute to the accident? The Board found no evidence that the transition from many Space Shuttle contractors to a partial consolidation of contracts under a single firm has by itself introduced additional technical risk into the Space Shuttle Program. The transfer of responsibilities that has accompanied the Space Flight Operations Contract has, however, complicated an already complex Program structure and created barriers to effective communication. Designating the Johnson Space Center as the "lead center" for the Space Shuttle Program did resurrect some of the Center rivalries and communication difficulties that existed before the Challenger accident. The specific ways in which this complexity and lack of an integrated approach to Shuttle management impinged on NASA's performance during and before the flight of STS-107 are discussed in Chapters 6 and 7.

As the 21st century began, NASA's deeply ingrained human space flight culture, one that has evolved over 30 years as the basis for a more conservative, less technically and organizationally capable organization than the Apollo-era NASA remained strong enough to resist external pressures for adaptation and change. At the time of the launch of STS-107, NASA retained too many negative (and also many positive) aspects of its traditional culture: "flawed decision making, self deception, introversion and a diminished curiosity about the world outside the perfect place." These characteristics were reflected in NASA's less than stellar performance before and during the STS-107 mission, which is described in the following chapters.

CHAPTER 4
Other Factors Considered

During its investigation, the Board evaluated every known factor that could have caused or contributed to the Columbia accident, such as the effects of space weather on the Orbiter during re-entry and the specters of sabotage and terrorism. In addition to the analysis/scenario investigations, the Board oversaw a NASA "fault tree" investigation, which accounts for every chain of events that could possibly cause a system to fail. Most of these factors were conclusively eliminated as having nothing to do with the accident; however, several factors have yet to be ruled out. Although deemed by the Board as unlikely to have contributed to the accident, these are still open and are being investigated further by NASA. In a few other cases, there is insufficient evidence to completely eliminate a factor, though most evidence indicates that it did not play a role in the accident. In the course of investigating these factors, the Board identified several serious problems that were not part of the accidents causal chain but nonetheless have major implications for future missions.

4.1 FAULT TREE

The NASA Accident Investigation Team investigated the accident using “fault trees,” a common organizational tool in systems engineering. Fault trees are graphical representations of every conceivable sequence of events that could cause a system to fail. The fault trees uppermost level illustrates the events that could have directly caused the loss of Columbia by aerodynamic breakup during re-entry. Subsequent levels comprise all individual elements or factors that could cause the failure described immediately above it. In this way, all potential chains of causation that lead ultimately to the loss of Columbia can be diagrammed, and the behavior of every subsystem that was not a precipitating cause can be eliminated from consideration.

NASA chartered six teams to develop fault trees, one for each of the Shuttles major components: the Orbiter, Space Shuttle Main Engine, Reusable Solid Rocket Motor, Solid Rocket Booster, External Tank, and Payload. A seventh “systems integration” fault tree team analyzed failure scenarios involving two or more Shuttle components. These interdisciplinary teams included NASA and contractor personnel, as well as outside experts.

Some of the fault trees are very large and intricate. For in-stance, the Orbiter fault tree, which only considers events on the Orbiter that could have led to the accident, includes 234 elements. In contrast, the Systems Integration fault tree, which deals with interactions among parts of the Shuttle, includes 295 unique multi-element integration faults, 128 Orbiter multi-element faults, and 221 connections to the other Shuttle components. These faults fall into three categories: induced and natural environments (such as structural inter-face loads and electro-mechanical effects); integrated vehicle mass properties; and external impacts (such as debris from the External Tank). Because the Systems Integration team considered multi-element faults that is, scenarios involving several Shuttle components, it frequently worked in tandem with the Component teams.

In the case of the Columbia accident, there could be two plausible explanations for the aerodynamic breakup of the Orbiter: (1) the Orbiter sustained structural damage that undermined attitude control during re-entry; or (2) the Orbiter maneuvered to an attitude in which it was not designed to fly. The former explanation deals with structural damage initiated before launch, during ascent, on orbit, or during re-entry. The latter considers aerodynamic breakup caused by improper attitude or trajectory control by the Orbiter’s Flight Control System. Telemetry and other data strongly suggest that improper maneuvering was not a factor. Therefore, most of the fault tree analysis concentrated on structural damage that could have impeded the Orbiter’s attitude control, in spite of properly operating guidance, navigation, and flight control systems.

When investigators ruled out a potential cascade of events, as represented by a branch on the fault tree, it was deemed “closed.” When evidence proved inconclusive, the item remained “open.” Some elements could be dismissed at a high level in the tree, but most required delving into lower levels. An intact Shuttle component or system (for example, a piece of Orbiter debris) often provided the basis for closing an element. Telemetry data can be equally persuasive: it frequently demonstrated that a system operated correctly until the loss of signal, providing strong evidence that the system in question did not contribute to the accident. The same holds true for data obtained from the Modular Auxiliary Data System recorder, which was recovered intact after the accident.

The closeout of particular chains of causation was examined at various stages, culminating in reviews by the NASA Orbiter Vehicle Engineering Working Group and the NASA Accident Investigation Team. After these groups agreed to close an element, their findings were forwarded to the Board for review. At the time of this report's publication, the Board had closed more than one thousand items.

The open elements are grouped by their potential for contributing either directly or indirectly to the accident. The first group contains elements that may have in any way contributed to the accident. Here, "contributed" means that the element may have been an initiating event or a likely cause of the accident. The second group contains elements that could not be closed and may or may not have contributed to the accident. These elements are possible causes or factors in this accident. The third group contains elements that could not be closed, but are unlikely to have contributed to the accident. Appendix D.3 lists all the elements that were closed and thus eliminated from consideration as a cause or factor of this accident.

Some of the element closure efforts will continue after this report is published. Some elements will never be closed, because there is insufficient data and analysis to unconditionally conclude that they did not contribute to the accident. For instance, heavy rain fell on Kennedy Space Center prior to the launch of STS-107. Could this abnormally heavy rainfall have compromised the External Tank bipod foam? Experiments showed that the foam did not tend to absorb rain, but the rain could not be ruled out entirely as having contributed to the accident. Fault tree elements that were not closed as of publication are listed in Appendix D.4.

Recommendations

It is the Board's opinion that good leadership can direct a culture to adapt to new realities. NASA's culture must change, and the Board intends the following recommendations to be steps toward effecting this change. Recommendations have been put forth in many of the chapters. In this chapter, the recommendations are grouped by subject area with the Return-to-Flight [RTF] tasks listed first within the subject area. Each Recommendation retains its number so the reader can refer to the related section for additional details. These recommendations are not listed in priority order.

PART ONE THE ACCIDENT
Thermal Protection System

R3.2-1 Initiate an aggressive program to eliminate all External Tank Thermal Protection System debris-shedding at the source with particular emphasis on the region where the bipod struts attach to the External Tank. [RTF]

R3.3-2 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. [RTF]

R3.3-1 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. [RTF]

R6.4-1 For missions to the International Space Station, develop a practicable capability to inspect and effect emergency repairs to the widest possible range of damage to the Thermal Protection System, including both tile and Reinforced Carbon-Carbon, taking advantage of the additional capabilities available when near to or docked at the International Space Station. For non-Station missions, develop a comprehensive autonomous (independent of Station) inspection and repair capability to cover the widest possible range of damage scenarios. Accomplish an on-orbit Thermal Protection System inspection, using appropriate assets and capabilities, early in all missions. The ultimate objective should be a fully autonomous capability for all missions to address the possibility that an International Space Station mission fails to achieve the correct orbit, fails to dock successfully, or is damaged during or after undocking. [RTF]

R3.3-3 To the extent possible, increase the Orbiter's ability to successfully re-enter Earth's atmosphere with minor leading edge structural sub-system damage.

R3.3-4 In order to understand the true material characteristics of Reinforced Carbon-Carbon components, develop a comprehensive database of flown Rein-forced Carbon-Carbon material characteristics by destructive testing and evaluation.

R3.3-5 Improve the maintenance of launch pad structures to minimize the leaching of zinc primer onto Reinforced Carbon-Carbon components.

R3.8-1 Obtain sufficient spare Reinforced Carbon-Car-bon panel assemblies and associated support components to ensure that decisions on Rein-forced Carbon-Carbon maintenance are made on the basis of component specifications, free of external pressures relating to schedules, costs, or other considerations.

R3.8-2 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.

Imaging

R3.4-1 Upgrade the imaging system to be capable of providing a minimum of three useful views of the Space Shuttle from liftoff to at least Solid Rocket Booster separation, along any expected ascent azimuth. The operational status of these assets should be included in the Launch Commit Criteria for future launches. Consider using ships or aircraft to provide additional views of the Shuttle during ascent. [RTF]

R3.4-2 Provide a capability to obtain and downlink high-resolution images of the External Tank after it separates. [RTF]

R3.4-3 Provide a capability to obtain and downlink high-resolution images of the underside of the Orbiter wing leading edge and forward section of both wings' Thermal Protection System. [RTF]

R6.3-2 Modify the Memorandum of Agreement with the National Imagery and Mapping Agency to make the imaging of each Shuttle flight while on orbit a standard requirement. [RTF]

Orbiter Sensor Data

R3.6-1 The Modular Auxiliary Data System instrumentation and sensor suite on each Orbiter should be maintained and updated to include current sensor and data acquisition technologies

R3.6-2 The Modular Auxiliary Data System should be redesigned to include engineering performance and vehicle health information, and have the ability to be reconfigured during flight in order to allow certain data to be recorded, telemetered, or both as needs change.

Wiring

R4.2-2 As part of the Shuttle Service Life Extension Program and potential 40-year service life, develop a state-of-the-art means to inspect all Orbiter wiring, including that which is inaccessible.

Bolt Catchers

R4.2-1 Test and qualify the flight hardware bolt catchers. [RTF]

Closeouts

R4.2-3 Require that at least two employees attend all final closeouts and intertank area hand-spraying procedures. [RTF]

Micrometeoroid and Orbital Debris

R4.2-4 Require the Space Shuttle to be operated with the same degree of safety for micrometeoroid and orbital debris as the degree of safety calculated for the International Space Station. Change the micrometeoroid and orbital debris safety criteria from guidelines to requirements.

Foreign Object Debris

R4.2-5 Kennedy Space Center Quality Assurance and United Space Alliance must return to the straightforward, industry-standard definition of “Foreign Object Debris” and eliminate any al-ternate or statistically deceptive definitions like “processing debris.” [RTF]

PART TWO – WHY THE ACCIDENT OCCURRED

Scheduling

R6.2-1 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. [RTF]

Training

R6.3-1 Implement an expanded training program in which the Mission Management Team faces potential crew and vehicle safety contingencies beyond launch and ascent. These contingencies should involve potential loss of Shuttle or crew, contain numerous uncertainties and unknowns, and require the Mission Management Team to assemble and interact with support organizations across NASA/Contractor lines and in various locations. [RTF]

Organization

R7.5-1 Establish an independent Technical Engineering Authority that is responsible for technical requirements and all waivers to them, and will build a disciplined, systematic approach to identifying, analyzing, and controlling hazards throughout the life cycle of the Shuttle System. The independent technical authority does the following as a minimum:

- Develop and maintain technical standards for all Space Shuttle Program projects and elements
- Be the sole waiver-granting authority for all technical standards
- Conduct trend and risk analysis at the sub-system, system, and enterprise levels

• Own the failure mode, effects analysis and hazard reporting systems
• Conduct integrated hazard analysis
• Decide what is and is not an anomalous event
• Independently verify launch readiness
• Approve the provisions of the recertification program called for in Recommendation R9.1-1.

The Technical Engineering Authority should be funded directly from NASA Headquarters, and should have no connection to or responsibility for schedule or program cost.

R7.5-2 NASA Headquarters Office of Safety and Mission Assurance should have direct line authority over the entire Space Shuttle Program safety organization and should be independently resourced.

R7.5-3 Reorganize the Space Shuttle Integration Office to make it capable of integrating all elements of the Space Shuttle Program, including the Orbiter.

PART THREE – A LOOK AHEAD

Organization

R9.1-1 Prepare a detailed plan for defining, establishing, transitioning, and implementing an independent Technical Engineering Authority, independent safety program, and a reorganized Space Shuttle Integration Office as described in R7.5-1, R7.5-2, and R7.5-3. In addition, NASA should submit annual reports to Congress, as part of the budget review process, on its implementation activities. [RTF]