Risk Management for the Tiles of the Space Shuttle

M.-ELISABETH PATÉ-CORNELL
Department of Industrial Engineering and Engineering Management
Stanford University
Stanford, California 94305

PAUL S. FISCHBECK
Department of Engineering and Public Policy and Department of Decision and Social Sciences
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

The tiles of the space shuttle orbiter are critical to its safety at reentry, and their maintenance between flights is time-consuming. We performed a probabilistic risk analysis to identify the most risk-critical tiles and to set priorities in the management of the heat shield. The model is based on a multiple partition of the orbiter's surface. For the tiles in each zone, we used the following data: (1) the probability of debonding due either to debris hits or to a poor bond, (2) the probability of losing adjacent tiles once the first one is lost, (3) the probability of burn-through given the final size of the failure patch, and (4) the probability of failure of a critical subsystem under the skin of the orbiter if a burn-through occurs. A risk-criticality scale was designed based on the results of this model. It is currently used (along with temperature charts) to set priorities for the maintenance of the tiles. We found that 15 percent of the tiles account for about 85 percent of the risk and that some of the most critical tiles are not in the hottest areas of the orbiter's surface. We recommended that NASA inspect the bond of the most risk-critical tiles and reinforce the insulation of the external systems (external tank and solid rocket boosters) that could damage the

0091-2102/94/2401/0064$01.25

SPACE PROGRAM
RELIABILITY—FAILURE MODELS

INTERFACES **24:** 1 January–February 1994 (pp. 64–86)

high-risk tiles if it debonds at take-off. We computed that such improvements of the maintenance procedures could reduce the probability of shuttle accident attributable to tile failure by about 70 percent.

The space shuttle, at take-off, consists of the orbiter, which is attached to a large external tank (ET) containing liquid hydrogen and oxygen, and two solid-fuel rocket boosters (SRBs) (Figure 1). The solid rocket boosters provide additional thrust during the initial phase of flight and are jettisoned over the Atlantic where they are retrieved for re-use. The external tank supplies fuel for the orbiter's main engines, which provide the thrust necessary to reach orbit. After its fuel is depleted, the

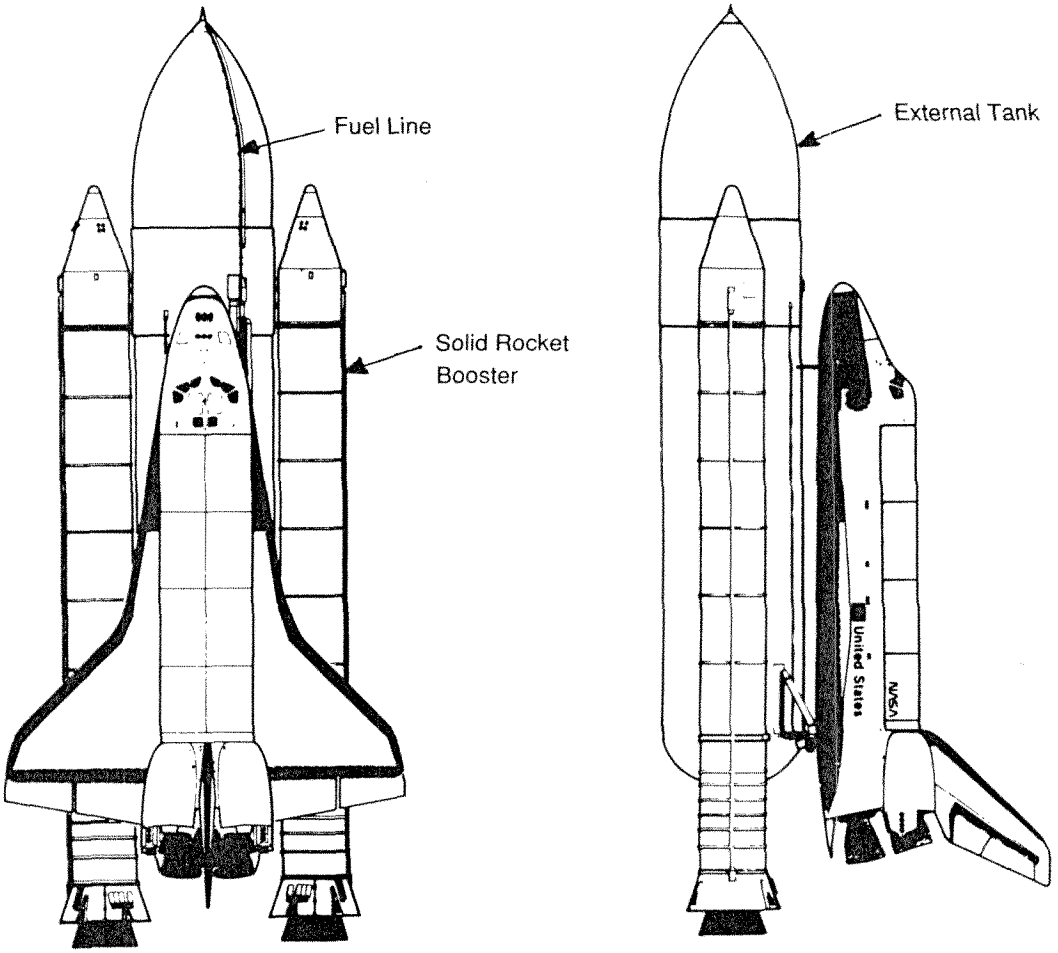

Figure 1: The space shuttle at take-off comprises the orbiter, the external tank, and the two solid-rocket boosters. The fuel line on the right side of the fuel tank brings liquid oxygen and liquid hydrogen to the orbiter's main engines.

external tank is jettisoned and burns up in contact with the atmosphere. Once the orbiter has completed its mission, it reenters the atmosphere and lands as an unpowered glider.

At reentry, the orbiter is subjected to high heat loads. Its heat shield is composed of reinforced carbon-carbon (an extremely strong material that can withstand up to 3,000°F) in the hottest areas (the nose and the edges of the wings), protective blankets in the coolest area (mostly the top side of the orbiter), and black tiles (about 25,000 of them) bonded to the bottom surface of the orbiter. The management of these black tiles between flights is both delicate and time-consuming and has often been on the critical path to the next launch. After every flight, some of the tiles are repaired; others have to be replaced.

Since the beginning of the shuttle program, the potential loss of any of the black tiles has been a major concern to the astronauts and to NASA. Once a tile is lost, adjacent tiles are more vulnerable to heat loads and aerodynamic forces. A gap in the heat shield (thermal protection system or TPS) could cause a burn-through in the aluminum skin of the orbiter during reentry, exposing and possibly crippling some of the critical subsystems and leading to the loss of the vehicle and crew (LOV/C).

The tiles are silicate blocks (roughly 8″ × 8″ × 2″) covered with a black glazing. They are bonded to a felt pad (strain isolation pad or SIP), which in turn is bonded to the orbiter's skin. Both bonds use a room-temperature vulcanized (RTV) material (Figure 2). Designed gaps between the tiles give the system flexibility and allow gases to vent during the ascent. Depending on the size of the gap, it is either left empty or filled by a gap filler. During reentry, the flow of gases around the orbiter shifts from laminar to turbulent: the later this shift, the less the total heat load. It is, therefore, important that the surface of the orbiter be relatively smooth to prevent local turbulence. To minimize the heat load, the gap fillers must fit perfectly in the interstices; the "step" and "gap" between tiles must be carefully controlled and measured so as to prevent surface irregularities.

To date, only two tiles have been lost. They failed because regular external loads exceeded the strength of a weakly bonded tile (on orbiter Columbia, the RTV weakened when it reacted with a waterproofing agent) and because a piece of debris struck with sufficient force to cause tile failure (also with Columbia in November 1987). Debris can come from several sources: ice that accumulates on the shuttle assembly, insulation from the external tank or the solid rocket boosters, debris on the ground or launch equipment, or space debris. Because of the risks involved, it is critical that tile installation and maintenance be done with the utmost care and that all sources of debris be controlled to the greatest extent possible.

Fifteen percent of the tiles contributed to 85 percent of the risk.

Risk management at NASA has generally relied on a careful control of the process, the concept of safety factors, and the

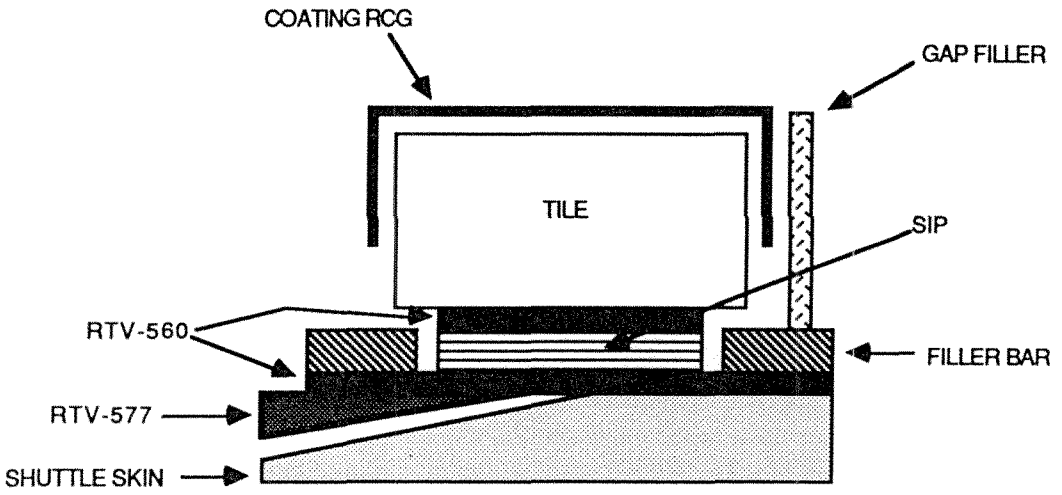

Note: Thickness exaggerated for clarity.

SIP: Strain-isolation pad

RTV-560: the tile bonding agent

RTV-577: screed used to smooth the orbiter's aluminum surface

RCG: Reaction cured glass

Figure 2: The tile system includes the tile coated by the black RCG bonded to a strain-isolation pad and installed on the orbiter's surface within a lattice formed by the filler bars. Gap fillers are used to smooth the surface by eliminating large spaces between the tiles.

setting of priorities based on criticality indexes (for example, the single failure of a component of "criticality one" implies loss of the vehicle and crew). The problem with these methods is that because they do not provide sufficient information about the relative contribution of each component to the overall risk, they do not allow for optimal resource allocation. Risk analysis, by quantifying these contributions, allows one to set optimal priorities. Despite these benefits, NASA has not used risk analysis. This is because, in the early 1960s, a consultant using risk analysis had computed a very small probability of success of NASA's mission to the moon. Fearing that such results would scare the public and discourage congressional funding, NASA forbid the use of any formal probabilistic risk analysis (PRA). Yet, after the Challenger accident, it became clear that the space shuttle system was much more vulnerable than NASA had been willing to admit—to itself and to the rest of the world—and that a realistic assessment of the risk was in order. A large number of modifications needed to be made, funding was becoming tighter, and priorities needed to be set to decide where to start. The Challenger accident had also revealed that many of the weaknesses that had finally doomed the mission were rooted in

the organization itself. Even though the eventual visible cause of the accident was a technical failure of a solid rocket booster O-ring [Presidential Commission Report 1986], this hardware failure was the direct result of management failures that included poor communication, misinterpretation of information, incentives to launch unless categorically proven unsafe, and excessive optimism under schedule pressures.

Classical risk analysis techniques are based on technical factors: an exhaustive identification of the different possible accident scenarios and the computation of the probability of failure of the whole system as a function of the probabilities of occurrence of the basic events of these scenarios. The basic events are generally technical in nature (for example, failure of various components). Organizational factors

Risk analysis is also a management tool.

are only implicitly part of the background. In recent years, we have extended these methods to include the effects of management on the inputs of the probabilistic risk analysis models, and therefore, on the risk of failure (see, for example, Paté-Cornell [1990]). This is the approach that we applied in this study of the safety of the tiles.

NASA (reluctantly) started a PRA effort in 1987. The Challenger accident had left the agency in a state of shock. Facing the risk of another accident was obviously painful, but addressing existing problems in proper order was necessary. However, instead of doing a complete top-down analysis, NASA opted for piecemeal risk assessments for the subsystems that seemed most worrisome (for example, the auxiliary power units). The goal of these studies was mostly to show that the risks from each of the major components were acceptable. At that time (1989), NASA also estimated that the overall risk of loss of the vehicle and crew was in the order of 10^{-2} per flight based on coarse statistics that included the Challenger accident and several near misses [NASA 1989]. The contribution of the tiles to this overall risk was unclear but believed to be substantial.

Within the framework of a cooperative research agreement between NASA and Stanford University, we did an extended probabilistic risk analysis for the black tiles based on the first 30 flights of the shuttle. Our goal was not only to compute the contribution of the tiles to the risk of loss of the vehicle and crew, but also to show how safety could be increased by improving the management and the procedures of tile processing [Paté-Cornell and Fischbeck 1993a, 1993b]. It was clear at the onset of the study that some management problems affected the quality of the tile work. For example, tile technicians were being paid less than machinists and electricians and had a high turnover rate. In general, there was no sense of priorities: procedures were meant to ensure that everything was done "perfectly." Several incidents, however, provided a simple reality check. During previous processing operations, it had been found that a few tiles had no primary bond with the orbiter (they were held in place only by the friction of the gap fillers). It was also discovered that under time pressure at least one technician had been spitting in the RTV to make it cure faster.

Addition of water to the bond agent is forbidden because, although it does accelerate the curing process, it also increases the probability that a catalytic reaction of the RTV will reverse causing the tile to debond under normal loads.

To set priorities among the tiles, we computed for each tile a risk-criticality index proportional to its contribution to the total probability of LOV/C. We showed how the same amount of maintenance resources (time, money, and attention) could be reallocated to provide more safety for the same global effort. We found that 15 percent of the tiles contributed to 85 percent of the risk. We also found a significant coupling between system failures. For example, debonding of insulation on the external tank (an event of minor concern) could damage critical tiles on the orbiter, eventually causing the loss of the vehicle and crew. In this perspective, probabilistic risk analysis is not only a simple assessment tool but also a management tool that, we believe, can lead to a fruitful improvement of risk-management practices in many other sectors as well.

The Probabilistic Risk Analysis Model

Our probabilistic risk assessment for the black tiles of the space shuttle was divided into two parts: the susceptibility of the tiles to damage and the effect of this tile damage on the performance of the shuttle. This two-phase approach allowed us to simplify the analysis considerably.

A tile fails when the loads on it exceed its capacity to withstand them. Originally, some tiles were partially burnt in some of the hottest zones. Design changes corrected these problems, and the main failure mode is now tile debonding. A tile debonds either because it receives an external load (a debris hit) that exceeds its design capacity or because it is unable to withstand a normal load, such as vibrations due to a weakening of one of its components (the tile material, the bonding agent, or the strain isolation pad). High load and low capacity can happen at the same time: in our analysis, we accounted for the case where a piece of debris hits a weakened tile and causes a failure that would not normally occur.

Historically, the source of this tile-damaging, larger-than-expected external load has been mostly from debris. The most common source of debris has been the insulation protecting the external tank and solid rocket boosters that has flaked off during launch hitting the underside of the orbiter. Other potential sources of debris include ice that forms on the external tank and solid-rocket boosters prior to launch and space debris. The severity of damage caused by debris impact varies considerably, but the debris problems are unavoidable. Every shuttle, upon return, has had some debris-damaged tiles. The number of tiles damaged in each of the first 33 flights ranged from a low of 53 to a high of 707 (with a mean of 179) out of about 25,000 tiles. The vast majority of this damage was minor, but on several flights, over 200 tiles had suffered major hits (the damaged area exceeded one square inch). Although it has not occurred yet, we included in our model the possibility that a single piece of debris

Every shuttle has had some debris-damaged tiles.

could significantly damage more than one tile. To reduce this problem, debris sources (and in particular the surface of the external systems) have to be properly maintained.

The second category of tile failure that we considered (a tile failing because its bond is weaker than the design capacity) can occur for several reasons: poor installation of the tile system, deterioration of the bond due to chemical reaction with external agents (for example, a waterproofing chemical or hydraulic fluid spills), and incorrect maintenance procedures. As with debris damage, this type of tile failure could affect a single tile or a group of tiles at the same time. To reduce this source of tile failure, NASA has to improve and monitor installation and maintenance procedures and needs to develop a reliable system for verifying the tile bonds. NASA has considered techniques for repairing damaged tiles in space in the past, but because of uncertainties about their effectiveness and the possibility of accidentally causing more damage to the tiles during repair, it has shelved all such plans.

In the second phase of our probabilistic risk analysis (assessment of the effects of tile damage on the shuttle's integrity), we started from the premise that a tile had been lost, and we systematically tracked the potential consequences of this initiating event during reentry and landing. The loss of a single tile would increase the load (heat and aerodynamic forces) on adjacent tiles and could cause these tiles to fail as well. Although it has not yet been experienced in flight, this zipper effect has the potential of opening a large patch of unprotected area on the orbiter's surface. It is important to note that it is not tile loss per se that causes the loss of the orbiter, but the increased heating or burn-through of the shuttle's aluminum skin, which in turn can lead to the failure of critical components, such as computers, hydraulic lines, flight controls, or fuel tanks. Whether or not the shuttle is lost depends on the location of the hot spot, the severity of damage, and the level of redundancy in the affected subsystems. Techniques for reducing the risk of losing a shuttle with a tile failure include hardening and relocating key components and increasing the redundancy of the overall system. If it were known before reentry that there was a gap in the tiles, it could be possible in some rare circumstances to reroute critical lines and drain some tanks prior to exposing the shuttle to intense heat loads.

Because of these variations of vulnerability and susceptibility, the black tiles cannot be considered as a uniform system. Certain tiles have higher probabilities of being damaged by debris, some receive greater heat loading during reentry, and some protect critical flight controls. Yet, it is not necessary to consider each tile as a separate entity with its own unique characteristics. To simplify the analysis, we grouped the tiles by levels of susceptibility to the initiating events and by degrees of vulnerability based on the shuttle components that they protect. We divided the orbiter's undersurface according to the density of debris hits, the vulnerability of adjacent tiles if one is already lost, the probability of burn-through, and the probability of loss of a critical subsystem following a burn-through. By superposing these four partitions, we obtained a finer

division of the orbiter's surface into minimal zones (called here "min-zones") of similar characteristics: all tiles in each min-zone have the same level of susceptibility to damage and the area that they cover has the same vulnerability if a tile is lost. A description of how we constructed these min-zones follows.

Figure 3 shows the structure of the probabilistic model used in the analysis in the form of an influence diagram. The model includes, for each min-zone (1) *initiating events* (probability distributions for the number of tiles initially lost either because of debonding under debris impacts or because of other factors that weakened the bond), (2) *final patch size* (probability distribution of the number of adjacent tiles lost conditional on the loss of the first tile), (3) *burn-through* (probability of burn-through conditional on a failure patch of a given size), (4) *system loss* (probability of failure of systems under the skin conditional on a burn-through), and (5) *loss of vehicle and crew* (probability of LOV/C conditional on failure of subsystems due to burn-through). We performed the analysis using the usual combination of probabilities estimated through past frequencies and subjective probabilities based on expert opinions (for example, the probabilities of failure of subsystems under the skin for which no formal PRA have been done). Very few statistical data were available. We knew, for example, how many tiles had been lost (two), the number and severity of debris hits, and the number of tiles that had been found during maintenance to have been poorly bonded. We needed expert opinions to assess, for instance, the probability of losing the orbiter given that a particular subsystem was exposed to high temperatures and the probability of tile debonding given a weak bond. The experts were (1) the tile (and other subsystem) specialists at Johnson Space Center, (2) the NASA tile maintenance experts at Kennedy Space Center, and (3) the tile maintenance contractor (Lockheed) at Kennedy Space Center. Whenever there were uncertainties about these probabilities, we made consistent use of means (estimated mean future frequencies) since the objective was to rank the tiles according to their contribution to the overall risk of accident. We then used Bayesian formulas to compute the probabilities of different scenarios.

The delineation of the min-zones is critical to the tractability of our analysis. The

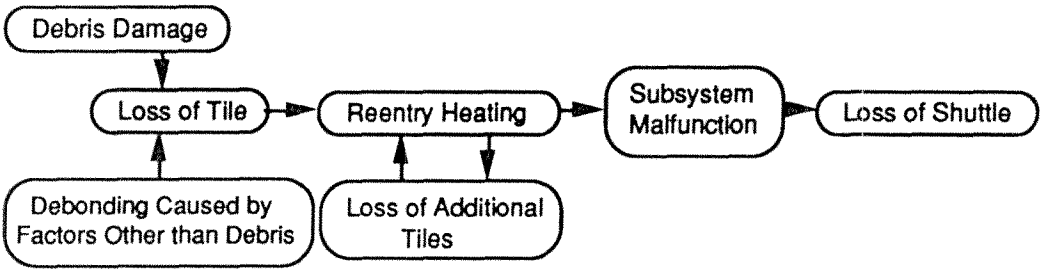

Figure 3: This influence diagram shows the structure of the probabilistic risk analysis for the loss of an orbiter due to failure of the black tiles. The blocks represent random variables and events, and the arrows show probabilistic dependencies between them.

more refined the definition of the factors determining a min-zone, the greater the complexity of the problem. Based on the reliability of the available data and the willingness of our experts to provide probabilistic estimates, we felt that an extremely detailed division was not defensible. Therefore, we divided the orbiter's surface according to the four zonal factors described above: (1) susceptibility to debris impact, (2) potential for loss of additional tiles following the loss of the first one (depending on heat and aerodynamic loads), (3) potential for burn-through given one or more missing tiles (heat loads), and (4) criticality of underlying systems.

We assumed that the probability of debonding caused by factors other than debris impact was uniform over the orbiter's surface and did not require a separate partition of this surface. In reality, the location of the various types and ages of the bond, the strain isolation pad, and the gap filler, as well as the temperature and pressure loads would affect the probability of debonding. We believe, however, that this simplification is adequate since one of the goals of our analysis is to determine the relative magnitude of the debonding problem to the debris problem. We estimated the probability of a weak bond by using maintenance records of replaced tiles (about 25 percent of installed tiles have been replaced so far) and a previous Lockheed study of bond verification [Welling 1989]. Based on this work, we concluded that chemical reversion of the RTV and weakening due to repeated exposure to load cycles are less likely to cause debonding than poor quality installation.

Of the four zonal factors, debris damage was perhaps the easiest to assess. By building a composite picture of all debris damage from the first 33 flights, we found that the undersurface of the orbiter could be divided into three areas: high, medium, and low density of debris hits. To a large extent, the high-debris density area corresponds to the area where insulation covering a fuel line on the side of the external tank hits the orbiter's surface if it fails during take-off. We estimated the degree of damage from historical data. By simply counting the number of hits per square meter, we estimated a probability of tile damage for each area. Paté-Cornell and Fischbeck [1990] give a detailed explanation of the procedures used in all the min-zone calculations for multiple tile debris impacts.

Whether or not tile loss leads to burn-through depends on three factors: the heat load during reentry, the number of tiles that are missing in a patch, and the ability of the shuttle's unprotected skin and underlying structure to dissipate the additional heat. Relying on expert judgment from NASA engineers at Johnson Space Center and on detailed data about recorded heat loads and aerodynamic forces at reentry, we partitioned the orbiter's undersurface into two secondary tile loss areas and assessed for each the probability of losing adjacent tiles given the loss of a first one.

In an analogous fashion, we partitioned the tiles into three burn-through areas and assessed the probability of burn-through in each area. It is interesting to note, that in the two cases in which tiles have been lost in the past, burn-through did not occur (in one case, the tile was lost over a service

hatch and the extra structure in the shuttle's frame was able to distribute the increased heating).

Finally, we determined criticality zones by studying the layout of key components under the orbiter's skin. If burn-through (or excessive heating) occurs, certain of these systems could be damaged. Whether the loss (or damage) of a particular sub component leads to the loss of the orbiter depends on its criticality. The probabilities of failure that we used were based on expert opinions. Ideally, a probabilistic risk analysis should be done for each system to determine its relationship to the shuttle's overall reliability.

NASA seems to have grown from a can-do organization to a large bureaucracy.

By overlaying these four partitions, we defined 33 min-zones and assigned to each of them an identification number. Of these 33 min-zones, 21 are unique with different sets of susceptibility and vulnerability indices. Several zones have the same combinations of indices and appear at different locations on the orbiter. Figure 4 shows the final layout of the min-zones and the results of the risk analysis. We determined the probability of failure of the orbiter attributable to each zone by calculating this probability for both types of initiating events (debris hits and debonding due to other causes) and then summing to obtain the results. We simplified the boundaries of the min-zones and approximated the number of tiles in each area (we did not make an actual count). For each tile in each min-zone, we defined and computed a risk-criticality factor proportional to the relative contribution of this tile to the overall probability of LOV/C, accounting not only for the loads applied to this tile, but also for the consequences should it fail. This risk-criticality factor (and its allocation between the two failure modes) is the point of reference that we used to set priorities among different management measures designed to improve tile and shuttle reliability. The min-zones in Figure 4 are shaded according to this risk-criticality index (the darker areas being more critical). According to our numerical analysis, the total probability of losing the orbiter on any given mission due to failure of the thermal protection system is on the order of 10^{-3} with approximately 40 percent of this probability attributable to debris-related problems and 60 percent to problems of debonding caused by other factors. For the total risk, including both initiating events (debris hits and debonding for other causes), 85 percent of the risk can be attributed to about 15 percent of the tiles. Among the 33 min-zones shown in Figure 4, eight min-zones contain these most risk-critical tiles (#1, 2, 3, 4, 5, 9, 10, 11) and are outlined in Figure 4. Complete numerical results can be found in Paté-Cornell and Fischbeck [1993a]. Derivation of the theoretical underpinnings can be found in the appendix of this paper.

Management Factors: PRA as a Management Tool

One of the objectives of this study was to assess the effect of management practices on the reliability of the thermal protection system. In the course of this work, we spent considerable time at Kennedy

Space Center and Johnson Space Center interviewing managers, tile technicians, and engineers in order to identify potential management problem areas. We iteratively passed on our research results back to NASA through a series of visits to the Kennedy Space Center to ensure that the information was properly communicated to

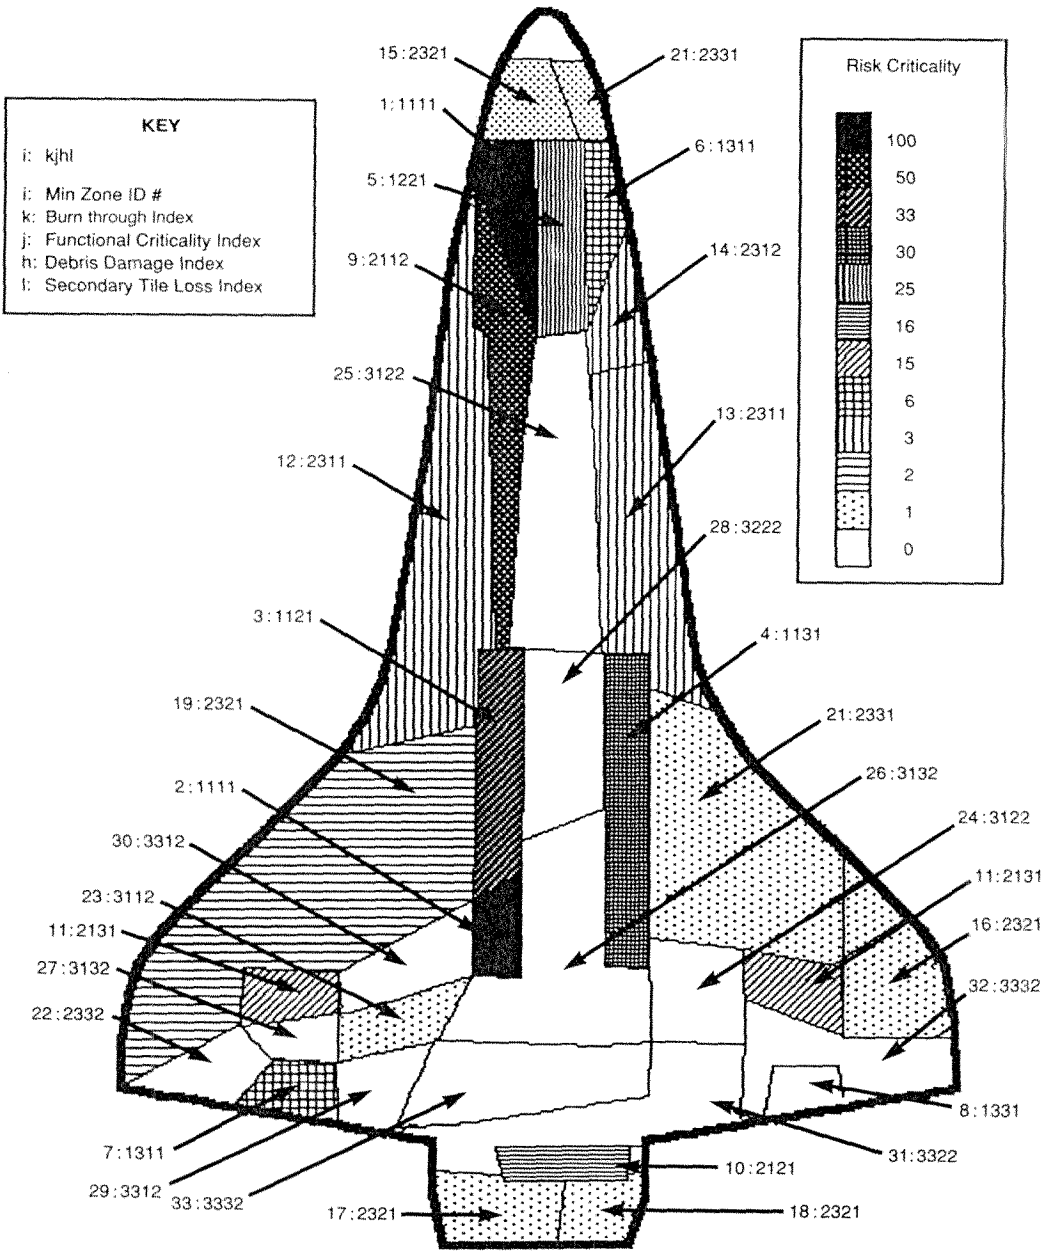

Figure 4: This map of the orbiter, showing the min-zones and the risk criticality of each tile, represents the main results of the analysis.

tile managers and inspectors. We describe the changes that have occurred since the publication of our original report [1990] below as an epilogue to this study.

The management factors that affect the tiles are often part of more general organizational characteristics of NASA. Space shuttle operations involve NASA, its headquarters, its space centers, its contractors (which are generally more closely associated with one of the space centers), the United States Congress, which votes the funds for the space programs, the media, and the public that eventually decides through the electoral process which programs are more desirable. Under political pressures, NASA became a fragmented organization, divided among space centers as well as space programs (space station, shuttle, unmanned planetary probes, and so forth). Rivalries among centers are unavoidable because they must share a budget that is perceived within NASA as increasingly tight. At the same time, NASA seems to have grown from a can-do organization to a large bureaucracy in which the influence of the scientists has markedly decreased.

NASA is a high-visibility organization, uncertain about its future funding and therefore, strongly influenced by its need for public relations. This high visibility makes it difficult for the organization to learn. To keep its funding, NASA has often led the public and Congress to expect a higher flight frequency and a higher level of safety than it could actually deliver. Soon after the shuttle's introduction, the agency shifted from a conservative attitude of "launch if proven safe" to an attitude of "launch unless proven unsafe." This optimism was more common among managers than among engineers and scientists who were more in touch with the system's characteristics [Feynman 1988]. Engineers, however, often have only a partial view of the system since different parts are managed in different locations, sometimes by different organizations. The 1986 Challenger accident obviously shook NASA and its culture of invulnerability and created new needs for appropriate methods to sort potential problems.

To some extent, these same organizational factors affected the processing of the tiles and, in particular, their maintenance between flights, which often took place under tight schedule constraints. Some of the problems that surface at the maintenance stage (for example poorly bonded tiles) are rooted in the history of the tiles, which starts with their design, manufacture, and installation on the orbiter. Because the system was new and subjected to a severe and poorly understood space environment, the first orbiter (Columbia) experienced problems that required prompt correction. For example, the initial design had a weak link between the bond and the felt pad (SIP) and the tiles had to be "densified" to strengthen that area. Also, moisture that had been trapped in the tiles froze in orbit damaging the material. To solve this problem, the tiles were waterproofed, but the first waterproofing agent reacted with the bond holding the tile in

The agency shifted to an attitude of "launch unless proven unsafe."

place causing it to partially liquefy. Many tiles were replaced before an effective waterproofing agent was developed. It should also be mentioned that the tiles that were originally installed on Columbia were put in place under severe schedule constraints, which may have affected the quality of the work. Finally, tiles in some particularly hot areas had to be redesigned. For example, the tiles in the elevon cove (the gap between the rear flaps) and on the edges of the main landing-gear doors had to be reinforced. After these adjustments, normal heat loads were not considered a problem for well-bonded tiles. This is why our analysis focuses mostly on the possibility of tile debonding because of a weak bond or the impact of controllable debris.

After each flight, the tiles are inspected several times; some are fixed and some are replaced. The process involves (1) a sequence of tile damage inspections and decisions as to what must be done, (2) tile replacement, (3) bond verification using pull tests, (4) step and gap measurement, and (5) decision whether or not to use a gap filler. The replacement of the tile is a delicate operation that involves several fits, densification, cleaning and priming of the cavity, inspections, bonding of the tile to the felt pad, bonding of the tile/felt pad system to the cavity, and verification of the bond. Several problems can occur in the process. First, the primary bond holding the tile may be partial or in some cases nonexistent. These tiles go undetected during their post-installation inspection because the gap fillers hold them in place with sufficient strength to pass the pull test. If the tiles are not precisely located in the cavity, they may rest on filler bars (pads that form a lattice on the aluminum surface and allow for venting of gases during ascent; see Figure 2) reducing the contact surface. Also, the cavity itself may not be sufficiently clean to provide proper bonding, the bonding agent (RTV) may be allowed to dry before pressure is applied, or the bond may deteriorate because of the addition of water to accelerate the curing process. All these factors can weaken the bond and lead to loss of tiles under normal loads that include vibrations and aerodynamic forces or under unintended external loads like debris impact. To minimize these risks, the main contractor at Kennedy Space Center thoroughly inspected all tiles following the maintenance work. In addition, NASA inspects in great detail about 10 percent of the tiles every three flights to try to detect signs of bond weakening. They use the wiggle test and look for signs of slumping or burning. Our objective was to improve this process to reduce the risk of accidental loss of the orbiter.

In general, studies of the effects of management on system safety start with the organization. In this study our approach was the reverse. We started from the elements of the technical risk analysis model represented by the influence diagram of Figure 3. For each of these technical variables, we identified decisions and actions that could adversely affect it. These processing errors (some of which are described above) may be obvious mistakes or simply questionable judgments that may eventually cause a failure (often in conjunction with other events). Management characteristics, in turn, affect information availability, the incentive structure, and the resource constraints (time, money, personnel

attention, and so forth). We examined whether and to what degree the organization caused or encouraged detrimental decisions and actions through its structure, procedures, and culture. For example, through extensive interviews with technicians and managers, we determined that time pressures due to the launch schedule were at the root of occasional (forbidden) short cuts.

We focused our process study on two key elements of tile safety: (1) ensuring a strong bond (for example, finding out if there are other tiles on the surface of all orbiters that still have a very weak bond even though they passed the initial pull tests) and (2) preventing controllable debris from hitting critical tiles on the orbiter's surface. New problems will unavoidably surface later: How will the bond age? When will it be time to replace the whole system? Does the pull test weaken the bond in the long run? These uncertainties about the long-term performance of the tiles will require future decisions. The probabilistic risk analysis model provides a way for NASA to prioritize these efforts, for example, instead of using random pull tests, to focus on areas where it is essential that the tiles be well bonded (that is, in the most risk-critical zones shown in Figure 4).

Figure 5 shows the influence of organizational factors on the technicians' decisions and actions and their relationship to the elements of the risk analysis model. Most of the links are self-explanatory; for example, the debonding caused by factors other than debris ("spontaneous debonding") is a function of the quality of the tile's bond, which itself depends on the quality of the maintenance work and the inspection procedure. The quality of the tile's work, in turn, can be traced back to hiring and training procedures and to the constraints under which the technicians operate. The quality of inspection may be the most critical variable since the number of inspectors is sometimes a limiting factor. Testing existing, already installed tiles also remains a priority problem. Developing means of nondestructive testing has been a concern of NASA for many years. So far, no technology has been successfully developed. The best technique available remains the wiggle test. Some skilled technicians are able to feel by a slight movement of the hand whether or not there is some slack in the bond. This test, however, is

NASA is strongly influenced by its need for public relations.

time-consuming and costly because there are few tile technicians who can do the job. A large proportion of the risk is attributable to weak bonds in general and in particular to a few tiles that may still be in place but have little (if any) direct bond to the orbiter's skin. Even if these tiles have held so far, it is essential that they be identified and replaced in the most risk-critical areas. Therefore, we proposed to use the nondestructive "wiggle test" to inspect the tiles in these areas in order of risk criticality.

We found a high level of turnover among tile technicians. Although dedicated people have worked on the thermal protection system for a long time (and therefore have had the opportunity to learn

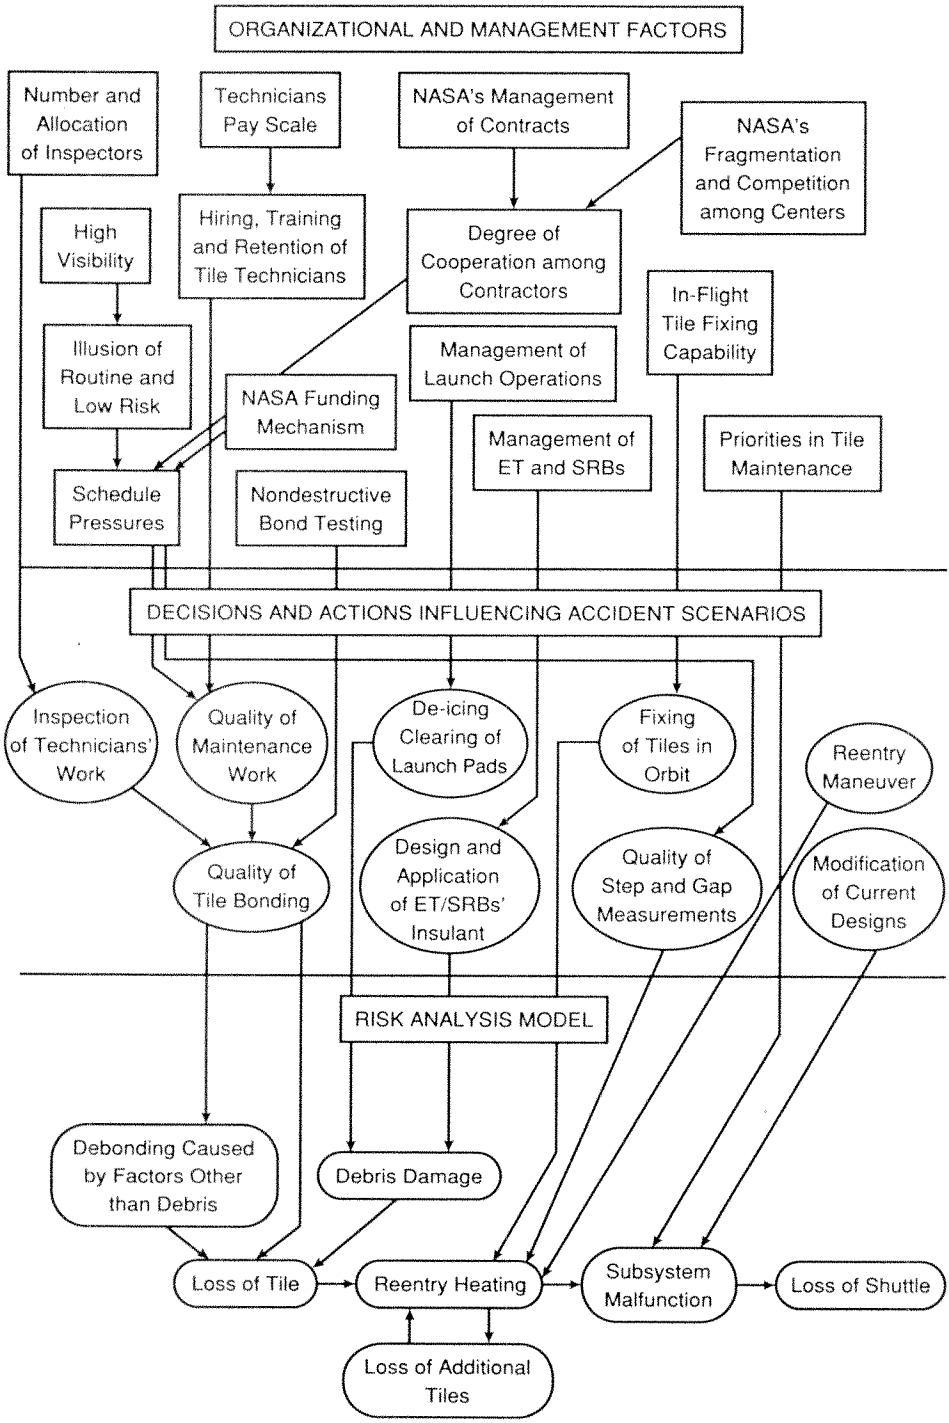

Figure 5: Organizational factors influence decisions and actions which in turn, affect the basic events of accident scenarios.

about it), a significant percentage of them have moved as soon as possible to better paying positions. This was caused by a discrepancy in the pay scale that NASA had inherited early on from the Department of Defense between material technicians and machinists or electricians. Not only was it costly to train technicians who then chose to leave, but also the quality of tile processing could only be negatively affected by the loss of experience. In addition, the tile technicians had the impression that others had no respect for the tiles and tended not to pay sufficient attention to their integrity when performing other tasks, thus adding to the work load of tile repair and to the time pressures on the tile crews. Therefore, we recommended in our initial report [Paté-Cornell and Fischbeck 1990] that this pay-scale gap be closed.

Another source of management-controlled reliability problems is the way that different systems of the orbiter are processed independently, ignoring the risk caused by their interaction. For example, we showed that one key source of debris damage is the insulation covering the external tank and the solid-rocket boosters. Some areas on the surface of these external systems are more critical than others because they may be sources of debris that could hit the orbiter in its most risk-critical min-zones. NASA had already worked on the simulation of the trajectories of such debris. We recommended that it use these simulations and the results of our study to identify areas of the external systems where the insulation should be treated with special care.

Liability concerns and rivalries among the main contractors may have also at times added to the schedule pressures, which in turn may have affected the quality of tile work. Relatively harmonious relations have been instituted among the people who work on the tiles; but the two main contractors are in a competitive situation that does not provide incentives for them to make each other's work easier. For legal and contractual reasons, it has occasionally been in their interest to withhold (or delay communication of) technical information useful, if not essential, to the other. Clearly, this competition provides a strong motivation to detect and correct errors. Yet, contracts that affect the same subsystems should be written and managed so as to foster cooperation when necessary.

Our report to NASA suggested a wide range of improvements that can be summarized as follows:

—To perform a follow up PRA for the tiles (our study was only a first-order analysis and used expert opinions whenever statistics were unavailable. Better inputs can be provided by lab tests (for example, of weakened tiles and bonds) and by proper PRA for the subsystems under the orbiter's skin);

—To relieve the time pressures on tile workers and in particular to avoid rigidly setting the number of tiles to be processed daily;

—To improve the information flow among the contractors and to give the maintenance crew direct access to the data that they need;

—To adjust the technicians' pay scale in order to reduce the turnover and keep the benefits of their experience; and

More important,

—To set priorities in NASA's final tile inspection and to focus initially on the most risk-critical tiles with the understanding that this should not imply that the others could be neglected; and
—To set priorities in the management of the insulation of the external tank and the solid rocket boosters, for example, by identifying the zones that could affect the most risk critical tiles (the technical problems of securing the insulation have not yet been not completely solved).

We then proceeded to assess (coarsely) the benefits of the last two recommendations based on the corresponding reduction of the probabilities of the two main failure modes.

Risk Reduction Benefits of Process Modifications

NASA [1989] has computed the probability of loss of the vehicle and crew to be about 10^{-2} per flight. The cost of losing an orbiter is on the order of $5 billion and the lives of the astronauts. There are about eight to 10 launches per year at this time. The contribution of the thermal protection system (TPS) to the risk of LOV/C (TPS risk) was found to be about 10 percent (on the order of 10^{-3} per flight) with 40 percent attributable to debris and 60 percent to debonding under normal loads. Decreasing the TPS risk can be achieved either by decreasing the controllable loads or by increasing the tiles' capacities. The benefits of in-depth safety measures, such as increasing the technicians experience base, are potentially important but difficult to quantify. We focused here on two direct risk-reduction measures: (1) improvement and inspection of the insulation of the external propulsion systems (external tank and solid rocket boosters) in the most risk-critical areas and (2) external inspection ("wiggle test") of the most risk-critical tiles.

The benefits of backmapping the most risk-critical zones of the orbiter onto the surface of the external tank and the solid rocket boosters and securing the insulation in these areas obviously depend on the final probability of insulation debonding. We found that 80 percent of the debris-initiated part of the TPS risk is attributable to eight percent of the tiles; therefore, reducing to zero the risks of debonding by debris of these eight percent risk-critical tiles would reduce the TPS risk by 32 percent *(80 percent of 40 percent) and the overall* probability of loss of vehicle and crew by about 3.2 percent (or 3.2×10^{-4}), a modest gain but one that could probably be obtained at a minor cost.

Detecting the unbonded tiles in the most risk-critical areas is the most efficient way to increase the system's capacity. About 130,000 tiles have been installed on the first four orbiters. At the time of our study, about 25 percent of these tiles had been replaced and about half had been inspected. Among those inspected, 12 had been found on the four orbiters then in service to have no bond other than through the gap fillers. We therefore assumed that about half of the unbonded tiles had been detected at the time of the study and that about 12 unbonded tiles (an average three per orbiter) remained. Detecting these unbonded tiles in the most risk-critical areas would reduce both the probability of spontaneous debonding and the probability of tile loss under debris impact.

Zones 1, 2, 3, 4, 9, 10, and 11 from Figure 4 contribute 85 percent of the risk of LOV/C through debonding due to factors other than debris, and 51 percent (85 percent of 60 percent) of the total TPS risk through this failure mode alone, even though they represent only 14 percent of the tiles. Zones 1, 2, 5, and 9 are particularly susceptible to debris hits and represent almost 80 percent of the TPS risk due to debris hits even though they represent only seven percent of the tiles (securing these tiles simply increases our confidence in the benefits of preventing the corresponding loss of insulation debris and is not assumed here to add to that figure). We assumed in our model that the two failure modes (debris hits and spontaneous debonding) were mutually exclusive and that the benefits of reducing their contribution to the TPS risks are additive. Therefore, the combination of detecting all loose tiles in zones 1, 2, 3, 4, 5, 9, 10, and 11 and securing the insulation on the corresponding zones of the external propulsion systems can reduce the TPS risk by about 80 percent (50 percent for spontaneous debonding and 30 percent for debris hits) and therefore, the overall probability of loss of vehicle and crew by about eight percent at a relatively low cost.

NASA cannot afford financially or politically to lose another orbiter.

This implies that since the initial probability of LOV/C due to failure of the tiles is 10^{-3}, applying our recommendations should reduce this risk by about 8×10^{-4}. Assuming that the loss of an orbiter costs about $5 billion (and the lives of five to seven astronauts), the expected value of the monetary component of the risk reduction benefits is on the order of $4 million per flight, plus a reduction of the individual risk by eight percent, starting from a high 10^{-2} per flight. This is probably an upper bound of the benefit figure. Obviously, if the effectiveness of the proposed measures is not 100 percent, the benefits are reduced by the same factor. Also, some of the weaker bonds would be discovered otherwise through the current inspection process, but much more slowly than if efforts were concentrated now on the 15 percent most critical tiles. Even if this combination of factors reduced the benefits by 50 percent, their expected value would remain high given the annual number of launches.

Epilogue and Conclusions

Since NASA received our report, it has instituted a number of improvements along the lines of our original recommendations. Our results were well received (in particular by the launch director), and the tile group recognized that prioritization can improve the maintenance process. The map that we have developed and presented here has been used along with a temperature chart by the tile safety officer at Kennedy Space Center to select the first 10 percent of the tiles that he inspects in detail after every third flight (as shown by the inspection map that was used on the shop floor in April 1993). The salaries of the tile workers have been brought in line with those of the machinists and electricians, the tile technicians training now puts more emphasis on the location of the dif-

ferent components under the aluminum skin of the orbiter, and the group working on the tiles has been given a greater role in the scheduling work loads. Communication difficulties among the contractors and problems of technicians' access to data bases have been alleviated. Our study also contributed to an effort by the agency to reinforce the insulation of the external systems, and more work has been done on the simulation of the debris trajectories. Since our study, the average number of major debris hits on the orbiter has decreased from an average 40 to 14 per flight.

Finally, NASA recently decided to conduct a probabilistic risk analysis for the whole orbiter. Our study was one of the factors that showed NASA that PRA could effectively be used as a management tool. The use of risk analysis to focus attention on the hot spots is clearly in line with the current efforts of the agency, which has seen its budget shrink. NASA must find new ways of being cost-effective because it simply cannot afford financially or politically to lose another orbiter.

Acknowledgment

This study was funded in part by a cooperative research agreement between NASA and Stanford University. We thank Ben Buchbinder at the NASA headquarters for his support of this work and our points of contact in the different space centers: David Weber (Lockheed), Frank Jones, Susan Black, Carol Demes, and Joy Huff (NASA) at Kennedy Space Center; James A. Smith, Robert Maraia, Carlos Ortiz, and Raymond Gomez (NASA) at Johnson Space Center; B. J. Schell, Frank Daniels, and Jack McClymonds (Rockwell) in Downey, California. We also thank Peter Banks and Michael Wiskerchen for their support and feedback in the course of this study.

APPENDIX: Equations of the PRA Model

(This appendix is adapted from Paté-Cornell and Fischbeck [1993a])

Our analysis of this problem is a variant of a conventional first-order PRA. Throughout this appendix, the factors that determine the min-zones are indexed as follows:

i = index of min-zones,
h = index of debris areas,
j = index of functional criticality areas,
k = index of burn-through areas, and
l = index of secondary tile loss areas.

Note that a double subscript (for example, ji) represents parameter j (criticality in this case) of min-zone i and that the term debonding refers to debonding due to factors other than debris impact.

n = total number of black tiles on the orbiter,
n_i = number of tiles in min-zone i,
N = total number of min-zones,
N_i = number of failure patches in min-zone i,
q = index for the failure patches in any min-zone,
M = final number of tiles in any failure patch,
m = index for the number of tiles in a failure patch,
Ft = initiating failure of a tile,
$Fa|Ft$ = failure of any adjacent tile given initiating failure,
D = number of adjacent tiles in initial debris area,
S = number of adjacent tiles in initial debonding area,
L = loss of vehicle and crew (LOV/C),
$P(X)$ = probaility of event X,
$P(X|Y)$ = probability of event X conditional on event Y,
$P(X, Y)$ = joint probability of event X and event Y, and

$EV(Z)$ = expected value of random variable Z.

This analysis follows closely the structure of variables described in Figure 3. Two types of initiating events are considered: debris impact and spontaneous debonding due to other causes, mainly a weak bond. It is assumed that the two types of initiating events are probabilistically independent. Since each min-zone has its own set of characteristics, they are treated as separate entities. Tiles in each specific min-zone have the same probability of being initially damaged and causing a larger failure patch, burn-through, damage to a critical system, and the loss of the vehicle. Because of these assumptions, the analysis determines first the probability of losing the vehicle for each type of initiating event and each min-zone. The overall failure probability is then computed as the sum of the failure probabilities for all zones and initiating events (debris impacts and debonding).

For each initiating event and for each min-zone, the structure of the equations is the following:

(1) Final patch size: probability distribution of the number of adjacent tiles lost as a function of the number of tiles initially lost and of the potential loss of additional tiles.

(2) Number of failure patches: number of combinations of groups of lost tiles in a specified min-zone.

(3) Probability of failure of the orbiter due to a specified failure patch as a function of the patch size, the probability of burn-through, and the functional criticality of the min-zone.

(4) Probability of failure of the orbiter for a specified min-zone (sum of the p(LOV/C) for all possible failure patches in the min-zone).

Then for the whole orbiter and both failure modes:

(5) Probability of failure of the orbiter for all tiles and both failure modes: (sum of the p(LOV/C) for all min-zones and both failure modes).

Initiating Event: Debris Impact

To determine the probability that a specific tile in min-zone i starts a patch due to debris impact, one must also consider the size of the initial damage. We will demonstrate the procedure for the case where a single tile is initially damaged (multiple tile damage is similarly calculated). It should be remembered that the probability of initial tile failure in min-zone i, $P_i(Ft)$, should be read as $P_i(Ft \mid D = 1)$.

Once the first tile in min-zone i is lost due to debris, there is the potential for adjacent tiles to also fail. The probability that the final patch size reaches M depends on the secondary loss index of the min-zone (l_i) and is given by the following geometric distribution (which means that $M - 1$ additional tiles fail and no adjacent tile afterwards):

$$P_i(M \mid Ft) = P_{li}(Fa \mid Ft)^{M-1} \times [1 - P_{li}(Fa \mid Ft)]. \quad (1)$$

This equation assumes that the probability that adjacent tiles debond does not change as the patch grows.

In each min-zone, there is the possibility of several patches starting. The probability that the number of patches reaches N_i in min-zone i is

$$P_i(N_i) = \frac{n_i!}{N_i!(n_i - N_i)!} P_i(Ft)^{N_i} \times [1 - P_i(Ft)]^{n_i - N_i}. \quad (2)$$

This formulation assumes that the initial tile failures are independent and that there will be no overlapping of patches because the probability of an initiating event (Ft) is small compared to the number of tiles in each min-zone (n_i). The product $EV(N_i) \times EV(M)$, which equals the total number of tiles lost in each min-zone, is considered negligible compared to n_i. Also, N_i (num-

ber of patches) and M (size of patches) are considered independent random variables. Based on these assumptions, the expected number of patches is approximately

$$EV(N_i) \approx n_i \times P_i(Ft), \tag{3}$$

and the size of each patch is given by the mean of the distribution of M,

$$EV(M) = 1/[1 - P_i(Fa|Ft)]. \tag{4}$$

Given this result, it is now possible to calculate the probability that the orbiter will fail due to debris that impact one tile only; using j as the index of the criticality areas and k as the index of the burn-through areas, we define the probabilities of orbiter failure due to a patch of size M, in min-zone i, initiated by debris impact ($D = 1$) as follows:

$$P_i(L|M = m) = p_{jki,m}. \tag{5}$$

It must be remembered that any given min-zone could have several patches in it, and each patch could be of a different size. To calculate the probability of orbiter loss due to the specific number of patches (N_i) in min-zone i, the following definition is necessary. Let p_i' be the probability that an arbitrary patch in min-zone i causes a failure.

$$p_i' = \sum_{m=1}^{\infty} p_{jki,m} \times P_{li}(Fa|Ft)^{m-1} \times [1 - P_{li}(Fa|Ft)]. \tag{6}$$

Therefore, using q as the number of patches in a given min-zone, the failure probability for a specific number of patches in a min-zone is

$$P_i(L|N_i = q) = p_i' \times q. \tag{7}$$

Equation (7) assumes again that the probabilities are small and that the patches will not interfere with each other (they are assumed to be separate and independent). These assumptions are valid providing that each min-zone has a sufficiently large number of tiles and that the size of the patches is relatively small.

Based on Equation (7), the probability of orbiter failure given all patches that occur in min-zone i becomes

$$\begin{aligned} P(L, \text{min-zone } i) &= \sum_{q=0}^{\infty} P_i(L|N_i = q) \\ &\quad \times P_i(N_i = q) \\ &= p_i' \times n_i \times P_i(Ft). \end{aligned} \tag{8}$$

This result represents only the cases of debris impact causing the initial failure of a single tile. A more complete rewriting of Equation (8) highlights this fact:

$$\begin{aligned} &P(L, \text{min-zone } i, D = 1) \\ &\quad = p_i'(D = 1) \times n_i \times P_i(Ft|D = 1). \end{aligned} \tag{9}$$

In order to expand this model to include the possibility that the initial debris impact damages more than one tile, it is necessary to modify some of the above equations. It is assumed that if a large enough piece of debris hits the orbiter, several adjacent tiles may be knocked loose at once. Each of these missing tiles may in turn cause their adjacent tiles to fail and a specific number of additional tiles can fail in multiple ways. Therefore, additional summations are required in order to account for the increased number of exposed tiles. This compounded problem requires that equation (1) be rewritten to account for this potentially larger patch growth rate. If the initial damage involves two tiles, the probability that the final patch reaches size M is

$$\begin{aligned} &P_i(M|Ft, D = 2) = (M - 2 + 1) \\ &\quad \times P_{li}(Fa|Ft)^{M-2} \times [1 - P_{li}(Fa|Ft)]^2. \end{aligned} \tag{10}$$

If three tiles are damaged initially,

$$\begin{aligned} &P_i(M|Ft, D = 3) = [\sum_{i=1}^{M-3+1} i] \\ &\quad \times P_{li}(Fa|Ft)^{M-3} \times [1 - P_{li}(Fa|Ft)]^3. \end{aligned} \tag{11}$$

If four tiles are damaged initially,

$$P_i(M|Ft, D = 4) = [\sum_{k=1}^{M-4+1} \sum_{i=1}^{k} i] \times P_{ti}(Fa|Ft)^{M-4} \times [1 - P_{ti}(Fa|Ft)]^4. \quad (12)$$

This set of equations can be extended to include greater initial damage; historical evidence, however, supports limiting the analysis to this level. It must be remembered that the value M of the final patch size must always be at least equal to the size of the initial damage area, D. Equation (2) in its most general form is written

$$P_i(N_i|D = d) = \frac{N_i!}{n_i!(N_i - n_i)!} \times P_i(Ft|D = d)^{N_i} \times [1 - P_i(Ft|D = d)^{n_i - N_i}, \quad (13)$$

and equation (3) becomes

$$EV(N_i) \approx n_i \times P_i(Ft|D = d). \quad (14)$$

Because all the initial damage probabilities are very small, it is possible to approximate the probability of debris causing loss of an orbiter for all damage areas in a particular min-zone by

$$P(L, \text{min-zone } i, \text{debris hit}) = \sum_{d=1}^{\text{Max } d} P(L, \text{min-zone } i, D = d). \quad (15)$$

Once this probability is determined, the probability of orbiter failure for all min-zones due to debris impact is simply the sum of the probabilities of failure for all min-zones since all min-zones and initiating events are assumed to be independent

$$p(L, \text{debris}) = \sum_{i=1}^{N} P(L, \text{min-zone } i, \text{debris hit}). \quad (16)$$

Initiating Event: Debonding (Due to Causes Other than Debris Impact)

A similar set of equations is used to compute the probability of orbiter failure due to all potential failure patches in min-zone i that started from a debonding damage area of initial size S. The size of the initial patch given debris hit (D) is replaced by the size of the initial debonding patch (S). These new equations are simpler because we assume a uniform distribution of the debonding probability over all tiles. The total probability of shuttle failure for this type of damage is

$$P(L, \text{debonding}) = \sum_{i=1}^{N} P(L, \text{min-zone } i, \text{debonding}). \quad (17)$$

Total Probability of Failure

Finally, assuming independence of initiating events (debris and debonding due to other causes), the overall probability of shuttle failure per flight due to tile damage is

$$P(L, \text{tile problem}) = P(L, \text{debonding}) + P(L, \text{debris hit}). \quad (18)$$

References

Feynman, R. P. 1988, "An outsider's inside view of the Challenger Inquiry," *Physics Today*, Vol. 41, No. 2 (February), pp. 26–37.

NASA 1989, *Independent Assessment of Shuttle Accident Scenario Probabilities for the Galileo Mission*, Vol. 1, Washington DC.

Paté-Cornell, M. E. 1990, "Organizational aspects of engineering system safety: The case of offshore platforms," *Science*, Vol. 250 (30 November), pp. 1210–1217.

Paté-Cornell, M. E. and Fischbeck, P. S. 1990, "Safety of the thermal protection system of the space shuttle orbiter: Quantitative analysis and organizational factors," *Report to the National Aeronautics and Space Administration*, Stanford University (December), Stanford, California.

Paté-Cornell, M. E. and Fischbeck, P. S. 1993a, "Probabilistic risk analysis and risk-based priority scale for the tiles of the space shuttle," *Reliability Engineering and System Safety*, Vol. 40, No. 3, pp. 221–238.

Paté-Cornell, M. E. and Fischbeck, P. S. 1993b, "PRA as a management tool: Organizational

factors and risk-based priorities for the maintenance of the tiles of the space shuttle orbiter," *Reliability Engineering and System Safety*, Vol. 40, No. 3, pp. 239–257.
Presidential Commission on the Space Shuttle Challenger Accident 1986, *Report to the President*, Washington DC, June.
Welling, R. 1989, "The tile bond verification shuttle inspection system," Lockheed Research and Development Division, Palo Alto, California, March.

Benjamin Buchbinder, Risk Management Program Manager, Office of Safety and Mission Assurance, National Aeronautics and Space Administration, Washington, DC 20546, writes "The paper by Professor Elisabeth Paté-Cornell describing the application of probabilistic risk assessment (PRA) to Space Shuttle thermal protection system processing reflects the results of a highly successful project sponsored by NASA. This project, led by Professor Paté-Cornell, has made an outstanding contribution to NASA management decision-making and the use of risk assessment in general. It has helped the Space Shuttle Program to identify the critical thermal protection tiles, and thus concentrate maintenance effort on the areas of significant risk. Others have applied PRA to the space program, but no one to my knowledge has combined the estimation of safety risk with consideration of the management process, to produce such useful results.

"We are in the midst of a long term effort to change the NASA culture and use PRA and related disciplines to improve programmatic decision-making. We have experienced some success in achieving an orderly, evolutionary change in approach throughout the agency, and this excellent work is one of our real success stories thus far."