ELSEVIER

Reliability Engineering and System Safety 74 (2001) 345–352

RELIABILITY ENGINEERING & SYSTEM SAFETY

www.elsevier.com/locate/ress

Probabilistic risk analysis for the NASA space shuttle: a brief history and current work

Elisabeth Paté-Cornell[a,*], Robin Dillon[b]

[a]*Department of Management Science and Engineering, Stanford University, Stanford, CA 94305, USA*
[b]*Department of Management Science and Information Technology, Pamplin College of Business, Virginia Tech, Falls Church, VA22043, USA*

Received 2 June 2000; accepted 13 December 2000

Abstract

While NASA managers have always relied on risk analysis tools for the development and maintenance of space projects, quantitative and especially probabilistic techniques have been gaining acceptance in recent years. In some cases, the studies have been required, for example, to launch the Galileo spacecraft with plutonium fuel, but these successful applications have helped to demonstrate the benefits of these tools. This paper reviews the history of probabilistic risk analysis (PRA) by NASA for the space shuttle program and discusses the status of the ongoing development of the Quantitative Risk Assessment System (QRAS) software that performs PRA. The goal is to have within NASA a tool that can be used when needed to update previous risk estimates and to assess the benefits of possible upgrades to the system.

Keywords: Probabilistic risk analysis; space shuttle

1. Introduction

For a long time, NASA managers have viewed probabilistic risk analysis (PRA) and expected-utility decision analysis with some suspicion, and many still do. The agency often operates new or improved systems, and there is seldom an abundance of statistics to describe past performance. Probabilistic risk analysis, however, is most useful when little statistical data are available to assess the failure probability of a whole system. In these cases, it is often useful to decompose the system into subsystems and components to quantify the overall failure risk as a function of the system's architecture and of the probabilities of failure of the different elements for which more data are generally available.

At the onset of the Apollo program, NASA seemed to have accepted the notion that quantitative risk analysis could be useful for decision support [1]. But the failure probabilities computed for some missions of the Apollo program were largely overestimated because they were based on conservative estimates of subsystem failure risks. Because the results were so pessimistic and showed such a small probability of mission success, NASA at that time, turned away from quantitative risk assessment methods. A few years later, however, failure probabilities such as 10^{-5} or 10^{-6} per flight were often casually quoted without much justification, either for subsystems or for entire missions, to express and support NASA's confidence in their performances [2].

Rather than quantifying failure probabilities, the agency has generally preferred qualitative analyses such as Failure Mode and Effect Analysis (FMEA), Critical Item Lists (CILs), and Risk Matrices [3–5]. FMEA/CIL relies on the logical identification of a system's weak points and of failure/event combinations (cut-sets) leading to its catastrophic failure. Risk matrices usually include, for different components or subsystems, qualitative information and corresponding scale indices about the likelihood of failure events (e.g., high, medium or low) and the severity of their consequences (e.g., high, medium, or low). These matrices are often used as filters to decide which are the highest priority technical problems. A major difficulty when using risk matrices is to combine such information about the different components to characterize the robustness of the whole system.

Since the Challenger accident, however, the use of PRA at NASA has increased significantly, not only for the space shuttle but also for some unmanned space missions and for the space station [6–17]. Probabilistic models have been and are being developed to assess the risk contribution of

* Corresponding author.
E-mail addresses: mep@leland.stanford.edu (E. Paté-Cornell), dillon@vt.edu (R. Dillon).

0951-8320/01/$ - see front matter
PII: S0951-8320(01)00081-3

specific shuttle problems and the benefits of shuttle upgrades. These models are currently implemented through the software called QRAS (Quantitative Risk Assessment System), itself in its development phase. Therefore, the current shuttle PRA model is both a motivation and a product of this development. The goal is to have within NASA a tool that can be used when needed to update previous risk estimates and to assess the benefits of possible upgrades.

This paper describes briefly the history of shuttle PRA and the current efforts to develop a complete risk analysis model at the same time as the QRAS software. It is based in part on studies such as Fragola [1] and on the results of a recent panel review of current NASA PRA work for the shuttle [18].

2. PRA for NASA space shuttle: a brief history

2.1. Prior to the Challenger accident

As mentioned above, at the onset of the Apollo program, NASA generally accepted the notion of using risk analysis to choose some mission features and to compare the results to safety benchmarks [1]; but during the Apollo program, pessimistic risk estimates discouraged the agency to adopt quantitative risk analysis. As is often the case, the problem was that conservative values (as opposed to means) of future failure frequencies were used to account for uncertainties, instead of a full uncertainty analysis. In truth, the methods were in their infancy and the software needed did not exist. The conservative approach seemed prudent but the results were both alarming and discouraging. As usual when it is clear that they contradict the facts, wrong results became detrimental to the acceptance of the quantitative risk analysis method itself.

Furthermore, notions of failure risk and failure probability often clashed (and still do) with the engineering culture, primarily based on safety factors, which can be defined in many different ways, and include, for example, design for higher loads than anticipated. Useful as they may be in decreasing failure risks, the problem with safety factors is that they are not directly linked to the probability of system failure and the relationship may vary across systems. For example, a safety factor of two in one system may not imply the same safety level as in another. Therefore, safety factors alone are insufficient to support cost-effective allocation of upgrade resources and prioritization of retrofits across different systems.

Rather than quantitative risk estimates, shuttle program managers preferred the use of Failure Modes and Effects Analysis and Critical Items List (FMEA/CIL) to identify a system's weak points and to manage failure risks accordingly. FMEA is useful to the extent that it indicates which possible design changes might eliminate a failure mode, reduce its future frequency to a lower level, or mitigate its consequences. All single-point failures that cannot be eliminated are collected in a Critical Items List. The process by which this list is established is a static, qualitative, bottom-up approach, focused on the identification and reduction of the risk of single, independent component failures that can cause the loss of crew, vehicle, or mission [1,19–22]. In the critical item list, components are listed according to their level of criticality. Criticality 1 characterizes an element whose failure is sufficient to cause overall shuttle failure defined by loss of vehicle and crew (LOVC). Criticality 1R is when there is one redundancy to prevent LOVC. The problem is that the level of criticality alone is not an indicator of the contribution of a particular component to the overall failure risk. A Criticality 1 component with a low failure probability can be less threatening to a system's safety than several components in parallel, each with a high probability of failure, especially if these failures are highly dependent. Regardless of its shortcomings, FMEA is a first step towards a PRA and the complete description of all accident sequences. The next step is the assessment of the probability of system failure per time unit or per operation as a function of the probabilities of relevant events including component failures, accounting for dependencies and in particular, those caused by external events.

In addition to qualitative studies, some quantitative risk analyses had been performed for the space shuttle before the Challenger accident during flight 51L in January 1986, (see for example, Baker [23]). These were small-budget studies, highly constrained and limited by the definition of their scope and by restrictions on data sources. NASA mandated that in addition to the limited data available from past performances, the analysts should use as data the opinions of NASA's experts regarding the performance of specific systems. These experts, for example, estimated the probability of a Solid Rocket Booster (SRB) failure at 10^{-5} per flight without any formal systems analysis needed to support such an estimate. It is on that basis that early risk analyses for the shuttle system indicated a probability of LOVC in the order of one in several thousand per flight. Consequently, this time, they significantly *underestimated* the failure risk. Yet, the astronauts knew from their experience with minor failures in flight (e.g., of a particular switch in the crew cabin) that the risks were probably much higher. Feynman, in his Appendix to the Rogers Commission report [24] emphasized the problems with such estimates and the downside of overconfidence. Indeed, Kaplan [25] showed that a Bayesian analysis based on the shuttle 'near-misses' prior to the Challenger accident could have indicated a much higher failure probability.

2.2. Post-Challenger risk analyses

The Challenger accident drastically altered this optimism. NASA and its contractors performed a major review of all shuttle FMEAs and updated the Critical Item List. The immediate result was a large increase in the number of

critical items from 2,369 to 4,686 [26]. In addition, the Rogers Commission [24] that investigated the accident concluded that the perceptions of shuttle failure risks had been overly optimistic and that the risk assessment methods needed improvement. Furthermore, priorities needed to be set among risk mitigation measures given the limitations of NASA's resources. Qualitative methods that had provided general guidance for risk management were inadequate for prioritization because they did not allow quantification of the relative contributions of the different components to the probabilities of system failure.

At about the same time, a National Research Council (NRC) panel reviewed the risk assessment and management of the space shuttle program, and recommended quantitative approaches to set priorities among possible upgrades of critical items. The NRC panel found that previous quantification of the shuttle risks were based almost exclusively on subjective judgments and qualitative rationales, even though quantitative engineering analyses and test data relevant to risk assessment were available and could have been used [27].

In the late 1980s and early 1990s, probabilistic risk analysis (PRA) therefore seemed a better alternative to qualitative risk assessment. Yet, within NASA, there was still strong resistance. First, the cost of a complete PRA seemed high. The value of information as decision support was not well understood, and it was sometimes stated that instead, the same amount of money could be better invested in strengthening the system. The question of course is: where should that investment be made in priority, and how much will eventually be gained by replacing intuition by quantitative decision support? Second, the use of Bayesian probability, which is often the only option given the systems' novelty, was often considered at NASA too 'subjective' to be trusted for decision support. That was true until it became obvious that there was no better alternative because by definition, extensive data sets did not exist. At the very least, these methods allowed a systematic and consistent assessment and treatment of the risk components.

In the following years (1990s), a number of pilot studies and a first attempt at a comprehensive shuttle PRA study were undertaken. In an initial attempt to incorporate probabilistic risk analysis methods in its decision support, NASA commissioned two 'proof-of-concept' studies. Their objective was to determine if a PRA could identify high-risk areas that traditional FMEA/CIL and hazard analysis techniques could not. One of these studies focused on the auxiliary power units (APUs), and the other on the main propulsion pressurization system [28,29]. These two studies showed that the probabilities of failure of a small number of CIL items represented most of the shuttle failure risk, and that in addition, several important failure scenarios had not been identified by NASA's previous analyses. Furthermore, the APU study demonstrated that the number of redundancies had to be weighed against the increase of risk of fire and explosion caused by the possibility of a hydrazine leakage (reference to original study, see Fragola [1]). These studies also showed that components of criticality higher than 1 (i.e. less critical to LOVC) could contribute as much as 30% of the overall failure risk. This implies that limiting a PRA to Criticality 1 items is likely to lead to an underestimation of the risk.

Shortly after the completion of these proof-of-concept studies, NASA was required to fund an independent shuttle PRA to support the approval of the launch of the Galileo mission because plutonium fuel was present on the spacecraft [30]. This study was limited to the ascent portion of the mission and focused primarily on scenarios that presented a risk to the nuclear payload. It included a propagation of uncertainties about the future frequencies of component failures, thus providing a probability distribution for the future frequency of a shuttle accident or flight abort. Despite NASA's conclusion that the probability of shuttle failure could be high, (median estimate 1/78), this study concluded that the risk to the public caused by plutonium contamination was low.

Prior to approval of the launch of the Ulysses spacecraft from the shuttle, the shuttle risk analysis that was required for the Galileo mission needed updating to consider any variations associated with this spacecraft since the previous study. In 1993, NASA thus commissioned an update of the Galileo study, using Bayesian techniques to integrate the former risk estimates with the new evidence that had been gathered since the original report [31].

Around the same time (the early 1990s), damage to several of the tiles of the shuttle heat shield during previous missions prompted NASA's management to commission a study of the thermal protection system (TPS). This study showed that the contribution of the black tiles that protect the underside of the shuttle orbiter at re-entry was about 10% of the overall shuttle failure risk [32]. Only two tiles had failed in flight thus far, without causing damage to the orbiter's skin. One failed because of a weak bond, and the other because it was hit by a piece of debris probably coming from the insulation of the external tank. That study showed that 15% of the tiles were the source of 85% of the probability of a shuttle accident induced by TPS failure. More importantly, on the basis of a simple first-order risk analysis, it allowed ranking the tiles by order of risk contribution, and therefore, setting priorities in the tile inspection before each flight [32,33].

Then, in 1995, NASA funded the first attempt at a comprehensive quantitative risk assessment including all phases of a shuttle mission [34]. The method used was similar to the PRA framework developed by the US Nuclear Regulatory Commission [35–40] (Master Logic Diagram, fault tree and event tree analyses, etc.) to obtain the probability of a major accident as a function of the probabilities of component and subsystem failures [41]. Because of resource limitations, however, a number of components were assumed to contribute negligible additional risks and were not included in the analysis. In addition, some external

accident initiators were left out, for example, the penetration by micrometeroids of systems other than the tiles. The results showed a probability of a shuttle accident (LOVC) between 1/76 and 1/230, which seemed consistent with the limited experience available at that time. Following that study, NASA managers (as the Nuclear Regulatory Commission had done a few years before) decided to use PRA as one of the bases for the support of decisions regarding improvements in shuttle safety. They needed a tool to routinely perform shuttle PRAs, which had to be updated regularly to monitor risk variations and to evaluate the effects of changes in design and operation procedures. This is the effort that is now underway and which we describe further.

Many lessons were already learned in these early PRAs; for example:

1. Conservative estimates should not be mixed with probabilities that represent (for instance) mean future frequencies of failures. Otherwise the results are meaningless and possibly counterproductive.
2. Guessing the probability of failure of a complex system such as the SRBs is unlikely to lead to an accurate figure when the system can be analyzed to provide a better result.
3. Near-miss events and partial failures can provide valuable information for the assessment of system failure risk, especially when a catastrophic failure has not yet happened.
4. Restricting a PRA to Criticality 1 items is likely to lead to an underestimation of the failure risks.
5. Adding redundancies does not always improve the safety of the system (APUs, for example, introduce an added risk of hydrazine leakage that has to be weighed against the value of an extra redundancy).
6. A top-down analysis is needed to capture the dependencies among system failures, for example between the debonding of debris from the insulation of the external tank and their effects on the tiles of the heat shield.

2.3. The current work on shuttle PRA and the QRAS software

The studies mentioned above were all completed by independent consultants outside of NASA. In July 1996, the NASA administrator requested that an independent quantitative analysis of the risk of a shuttle accident be conducted by internal NASA experts, and that supporting software be developed. The long-term objective is to use the results as decision support for shuttle upgrades. The chosen approach is to develop the Quantitative Risk Analysis System (QRAS) software to perform PRA, permit its updating, and allow real-time support of decisions ranging from retrofit to launch under specified circumstances. This ongoing study involves, in parallel, the improvement of the first version of the QRAS software and the performance of a PRA. It will result in a model that will provide an overall shuttle failure probability and will allow estimation of the risk changes associated with proposed shuttle upgrades (e.g., an upgrade of the main engine turbopumps).

Two separate teams are currently developing these risk analysis models. The first one, at Johnson Space Center (JSC) analyzes the orbiter and its main propulsion systems including the auxiliary power systems, hydraulic system, thrust vector control, and main propulsion system. The second team, at Marshall Space Flight Center (MSFC), is in charge of the other shuttle subsystems of the main engines, the external tank, the solid rocket boosters, and the reusable solid rocket motors. These studies are designed to be limited to Criticality 1 and 1R items, and generally assume that failures of such items inevitably lead to an accident or mission failure. For analytical purposes, the system, as well as the PRA, have been divided into 'modules'. The analysis is being done 'bottom-up' on the basis of these modules. Some links have been included to ensure that an accident sequence that cuts across modules, and across analytical teams, are accounted. Yet, the current exercise is facing some of the classical difficulties of coordinating a PRA study when the system has been divided for analytical purposes. Both teams are supposed to rely on QRAS while it is still in its development phase. Therefore, at this stage, some elements of the PRA (e.g., fault tree analysis results, especially for the analysis of the orbiter) are computed 'off-line' independently from the existing software. A panel of experts recently reviewed the current shuttle PRA efforts [18]. Some of the comments of this panel are described below.

The computer software QRAS originated at NASA Headquarters in conjunction with the University of Maryland in 1998 [42–44]. It is currently being developed by NASA Headquarters and its contractors and subcontractors, including Allied Signal and L&M Technology. QRAS aggregates subsystem failure mode probabilities from the bottom-up to produce intermediate and top-level catastrophic failure probabilities and bounds on the uncertainties. It is based on the identification of a set of scenarios represented by event sequence diagrams (ESDs), starting with an initiating event and ending with an accident, a flight abort, or a benign outcome, either directly or through a sequence of intermediate ('pivotal') events. Among the results is a prioritization of the subsystem failure modes that contribute most to the overall risk, and an evaluation (and ranking) of space shuttle potential upgrades, both from a safety and a cost point of view [20]. The first version of this software is currently used at JSC to assess the failure risks of the shuttle orbiter, and at MSFC to assess the failure risks of the other shuttle subsystems.

Once the QRAS software is completed and available, NASA will be able to develop and upgrade PRA models at very detailed levels, integrating physical models of failure processes into the logic model and the probabilistic analysis.

For the moment, however, the two space centers that are charged of the development of PRA models using the first version of the QRAS software have had mixed experiences with it because it still misses important features. For example, in its current version, QRAS treats accident sequences independently from others even though some may have common events. It does not include proper treatment of external events and common causes of failures, and it does not have the capability of building and analyzing fault trees. Therefore, some of the current PRA work has to be done off-line, for instance using fault tree analyses that are not part of the current software, before integrating the results in the QRAS models. A consortium of industry contractors, the United Space Alliance (USA), has been charged with the space shuttle operations and is monitoring the shuttle PRA work done both at JSC and MSFC. Its objective is to use the results to support recommendations for upgrades of the shuttle design as well as improvements of maintenance and processing operations.

3. Some characteristics of the current PRA modeling efforts for the space shuttle

In a recent review of the space shuttle PRA, a panel of experts [18] concluded first and foremost that the PRA models currently developed by NASA were an important step towards improving the risk management process. It is essential at this stage that NASA adopt current risk analysis methods to be able to improve its systems in a cost-effective way. Yet, it was also found that the current models exhibited a number of characteristics that left space for improvement.

3.1. The effect of organizational dispersion

The coordination and communication among the teams that perform the shuttle PRAs at JSC and MSFC may not be sufficient. For example, the two groups use different 'ground rules' and assumptions, possibly because they interpreted differently NASA's initial directions. The studies were to be limited to the most significant of Criticality 1 items. In addition, a common assumption was that failure of these items inevitably leads to a system failure. The first question is to choose the items to be included in the analysis, and the two teams adopted different procedures to choose the events included in their models, the level of detail of their studies, and the treatment of quantitative data. Therefore, the results obtained in the two centers are not directly comparable at this stage. In addition, when a system is divided at the onset of a PRA without an overarching model to ensure completeness and consistency, issues can surface in the treatment of dependencies across subsystems, common causes of failures and performance of the interfaces.

The problem of dispersion of work across centers with insufficient communications is a common one that had already been identified in the past as one of the safety problems of the shuttle system. For example, after the Challenger accident, the Rogers commission [24] showed that poor communications were at the source of the fatal decision to launch on that day. In the same way, during the earlier analysis of the tiles contribution to failure risks, it was shown that part of the risk of tile failure could be attributed to debris hits caused by the debonding of parts of the insulation of the external tank [32]. Yet, the two systems are managed independently, the external tank at Marshal Space Flight Center and the tiles at Kennedy Space Center, and it took the chemical analysis of a missing tiles cavity before the link was established.

3.2. Analytical modules as opposed to an overarching model

The role of an overarching model is to ensure the completeness and the accuracy of a PRA, the inclusion of dependencies across systems and of common events across accident sequences, and the proper treatment of external events that can affect simultaneously several subsystems. This requires a top-down approach starting from a systematic analysis of accident sequences, or conjunctions of events leading to failure. An overarching model can be based on different tools such as the Master Logic Diagram developed and used in the nuclear power industry, a complete event tree, or an influence diagram. Influence diagrams are particularly helpful because they can process both probabilistic dependencies and also deterministic functions such as the Boolean analysis involved in fault trees. They also provide a graphical display of interdependencies among events. What is important, in any case, is less the nature of the tool itself than the completeness of the set of scenarios and the analysis of dependencies that are included in the PRA.

The PRA models for the different parts of the shuttle are currently constructed mostly bottom-up. The system has been divided into modules that are then analyzed. This structure has no doubt facilitated the division of work, and some accident sequences that cut across modules have been included. The decomposition of the system, however, is generally one of the steps of the analysis, based on logic and if resources are constrained, on the value of information of further decomposition. The definition of modules as a starting point in the analysis can lead to missing failure dependencies and commonality of elements among accident sequences. Therefore, it can hide the true risk contribution of an element that affects several modules if no integration mechanism permits assessing the role of this component across the system.

3.3. A simplified approach to consistency in the level of analytical depth and detail

It is generally impossible to include all components and all event scenarios in a PRA, and an adapted screening procedure is necessary. This screening procedure is meant to filter out the scenarios that are low contributors to the overall risk while retaining the important ones. For simplicity, the current PRAs for the space shuttle are limited to

failure scenarios involving Criticality 1 and 1R items, and among these scenarios, a decision is made a priori as to which ones are sufficient risk contributors to be included in the analysis.

As mentioned above, the criticality level is only loosely coupled to an item's contribution to the probability of failure and it was shown for the shuttle that items of Criticality 2 and above were significant risk contributors. Therefore, the simplicity of this choice probably leads to excluding some risky items that are possibly more dangerous than some that were included. More importantly, perhaps, it might eliminate from the ranking of upgrades some improvements that could be more cost-effective than those considered. This choice, by itself, would lead to an *underestimation* of the overall probability of failure.

The definition of Criticality 1 items does not imply that their failure inevitably leads to systems failure, only that it *can* cause an accident. Several intermediate ('pivotal') events can occur following such a failure. Some sequences (or conjunctions) can lead to an accident, others to a safe flight abort or to a correction by human intervention that permits completion of the mission. Yet, in order perhaps to produce conservative results or to balance the exclusion of other components, it is generally assumed in the current studies that the occurrence of an initiating event of Criticality 1 inevitably leads to an accident. Therefore, this time, the simplifying assumption may lead to an *overestimation* of the consequences of a Criticality 1 event. It may be that the choice of Criticality 1 items (and only of some of them) compensates for the assumption that they lead to system failure; but it is impossible to tell without further information whether the overall results reflect an overestimation or underestimation of the failure risk.

The definition of initiating events is only a starting point. The choice of analytical depth and of adequate level of detail in the different parts of a system is critical to ensure first the best use of the resources spent for the analysis, and second, the consistency of results across subsystems. The analytical depth in the current shuttle PRAs is simply determined by the hierarchy of components and subsystems. A certain form of consistency has thus been obtained.

Alternatively, consistency in analytical depth could be based on the value of the additional information that one might expect from pursuing the analysis further down in some parts of the system. Therefore, another rule could be to stop the analysis when it does not bring additional information that is likely to make a difference in the results and in the decisions that they support [45,46]. When there are sufficient failure data at a subsystem level, it does not need to be analyzed further, unless one seeks to evaluate the contribution of one of its components to the overall failure probability. By contrast, when there is little statistical evidence about a subsystem and when data are available at the component level only, the analysis has to be done at the component level. Sometimes, it may even be necessary to go further and decompose the failure of a component into its failure modes (for example, a valve can be stuck open or closed). One can then take advantage of the powers of Bayesian treatment of the evidence and of systems analysis to aggregate the risk contribution of the different elements.

Analytical judgment and flexibility are thus required to ensure that the main risk contributors are identified and included in the analysis, down to the same level of risk contribution. This choice may or may not correspond to the classical hierarchy of subsystems and components. A simplified uniform approach to analytical depth can be ineffective because a detailed analysis at the component level can be useful in some places but unnecessary in others.

3.4. Human decisions and action in a risk analysis

Human decisions and actions are key factors in system failure risks. Yet, they are sometimes ignored or poorly treated. It is important to note that they can include not only catastrophic errors, but also operators actions to correct a dangerous condition.

Errors can occur in manufacturing, system assembly, inspection and maintenance, and operation (mission). When these errors are already included in the database used for risk estimation, they are de facto included in the analysis and do not need to be addressed further. Existing statistical data, however, may not include rare errors that can have catastrophic consequences. It seems that in the PRAs that are currently performed by NASA, there is no analysis of process errors that can affect the different subsystems. This type of errors, for instance, were analyzed and included in the 1993 study (mentioned above) of the LOVC risks due to failure of the tiles [32,33]. These errors included, for example, failure to center a tile in its cavity during maintenance operations, or letting the bond dry before applying pressure. Both can significantly reduce the strength of the bond causing a tile to debond in operation and leaving the aluminum skin exposed to heat loads at re-entry.

Errors can also affect the operations phase. Yet, there seems to be an implicit assumption in the NASA studies that astronauts make no mistakes (with the possible exception of an error at landing). Clearly, this apparent omission of human errors tends to underestimate the risk of catastrophic failures.

But there is a positive side that human intervention can reduce the risks of an accident or stop the propagation of an accident sequence. Again, it may be that skilled interventions compensate for the possibility of human error, but it is impossible to determine without further information the net effect of these two omissions on the overall failure risk.

3.5. Mixed methods in data analysis

A risk analysis is most useful when there are few statistical data of different nature and from different sources such that the situation requires Bayesian treatment of the evidence. The frequentist approach to classical statistics

has the advantage of being commonly used but requires a large amount of information. Furthermore, the confidence levels that qualify the results are difficult to interpret as characteristics of uncertainties. The Bayesian analysis is more powerful in this respect and does not require a specified amount of data (i.e., the quantity of information is reflected in the results of the uncertainty analysis). But it requires the use of prior probabilities (e.g., uniform distributions may be used to reflect complete ignorance when that is truly the case), which injects an element of subjectivity in the analysis. In the case of space systems, one generally does not have the choice because flight data are rare by definition. Yet many different types of data can be used as input to a risk study.

The current PRAs for the space shuttle often use both frequentist and Bayesian analyses (hence possible inconsistencies), but not always all available information. Possible data include test data, flight data, surrogate data and expert opinions when appropriate. Therefore, the results could probably be improved by adopting consistently a Bayesian approach, using all existing data. For example, surrogate data can be used as priors to be updated based on additional experience and new flight data. In any case, failure probabilities must be assessed differently if they represent marginal or conditional probabilities, in which case the events on which they depend must be considered.

Finally, in the current PRAs, the simplifying choice was made to compute first-order probabilities only. The results are thus represented by the probability (or mean future frequence) of different potential system states, based on the probabilities of different hypotheses or models, and of parameter values given these different models [47,48]. In contrast, in a second-order uncertainty analysis, the uncertainties about the possible underlying hypotheses or models are propagated throughout the analysis. The results are probability distributions of the probabilities (or future frequencies) of different system and subsystem states. A first-order probability analysis is sufficient to set priorities when the ranking criterion is the mean future frequency. Yet, an assessment of uncertainties in the input (i.e., failure frequency distributions for the basic components) and consistent propagation of these uncertainties in the analysis could permit, in addition, an assessment of the effects of uncertainties on the results and on priorities.

Many of these simplifying assumptions will be unnecessary after the completion of the QRAS software. QRAS is currently being updated to eventually involve features such as fault tree analysis, an overarching model, external events, human errors and adequate Bayesian treatment of all available information. The current experience is probably a necessary step towards the realization that such features (among others) are needed to provide results that are credible in absolute terms, and in relative terms, permit ranking of upgrades by order of cost-effectiveness.

NASA's experience in this respect is not unique. The US Nuclear Regulatory Commission has taken a long time to develop its PRA models to the point where they can be used, along with other types of information, to make safety decisions. Given the unique nature of its systems, NASA will probably need to go through a similar exercise and the use of quantitative risk analysis will be of great value in assisting decisions in all phases of space systems life, from design, processing, operations and upgrading. It is important however, that fundamental issues be recognized and resolved quickly.

4. Conclusions

The Probabilistic Risk Analysis method has gone through ups and downs at NASA. From the hopes of the early times of the Apollo program, to the disappointments of pessimistic then optimistic results (and were wrong in both cases), it is slowly being improved and incorporated in the NASA thinking about risk ranking and prioritization of upgrades. Where it is resisted, it is often because it runs against the engineering tradition of safety factors and suspicion about the use of Bayesian probability. Yet, if one wants to assess the risk, the Bayesian approach is unavoidable because there are seldom enough data for a classical statistical analysis. PRA has now been adopted as one of the decision supports for the management of the space shuttle, of the space station and of some unmanned space missions. In the long term, this decision will improve the consistency and the efficiency of the management of NASA's space systems. As usual, in this early phase of the PRA modeling, several problems still need to be addressed. Some of them are essentially organizational (i.e., the work is divided among several space centers). But the most important issues are the need for an overarching model and fundamental consistency in the choice of method of problem structure, analytical depth and treatment of data. As always, the value of the analysis will be determined by the use of the information that it provides, and NASA should realize an improvement in decision-making based on quantitative assessment rather than intuition and guesses at the system level.

References

[1] Fragola JR. Risk Management in US Manned Spacecraft: From Apollo to Alpha and Beyond. Proceedings of ESA Product Assurance Symposium and Software Product Assurance Workshop, Noordwijk, Netherlands, March 19–22, 1996

[2] Feynman R. Personal Observations on the Reliability of the Shuttle, Appendix IIF. In: Rogers, et al., 1986.

[3] Bowles JB. The New SAE FMECA Standard. Proceedings of the Annual Reliability and Maintainability Symposium 1998:48–53.

[4] Littlefield ML. FMEA/CIL Implementation for the Space Shuttle New Turbopumps. Proceedings of the Annual Reliability and Maintainability Symposium 1996:48–52.

[5] Onodera K. Effective Techniques of FMEA at Each Life-Cycle Stage. Proceedings of the Annual Reliability and Maintainability Symposium 1997:50–6.

[6] Agarwala AS. Reliability Engineering in Defense and Aerospace – A

Transition to the Commercial World. Communications in Reliability, Maintainability, and Supportability 1994;1(1):14–9.
[7] Davison M, Vantine WL. Understanding Risk Management: A Review of the Literature and Industry Practice. European Space Agency Risk Management Workshop, ESTEC, March 30–April 2, 1998:253–6.
[8] Frank M. A Survey of Risk Assessment Methods from the Nuclear, Chemical, and Aerospace Industries for Applicability to the Privatized Vitrification of Hanford Tank Wastes. Report to the Nuclear Regulatory Commission, August, 1998.
[9] Frank M. Assessment of the Cassini Mission Nuclear Risk with Aleatory and Epistemic Uncertainties. Proceedings of the 4th International Conference on Probabilistic Safety Assessment and Management. September 13–18, 1998.
[10] Frank M. Personal correspondence describing NASA project work. October, 1998.
[11] Guarro S, Bream B, Rudolph LK, Mulvihill RJ. The Cassini mission risk assessment framework and application techniques. Reliability Engineering and System Safety 1995;49:293–302.
[12] Jet Propulsion Laboratory (JPL). Cassini Recertification Review, JPL Internal Document D-11715, 2. Pasadena, California: Jet Propulsion Laboratory, 1994.
[13] Miles R. Personal correspondence describing NASA project work, July, 1998.
[14] Mulvihill RJ. Personal correspondence describing NASA project work, July, 1999.
[15] Railsback J. Personal correspondence describing NASA project work, July, 1998.
[16] Shemanski T, Silke K. Reliability Growth Model Overview. Reliability Bulletin 92-02, General Dynamics Space Systems Division, 1992.
[17] Silke K, Bennett J. Launch Vehicle Reliability Assessment. Reliability Bulletin 92-01, General Dynamics Space Systems Division, 1992.
[18] Paté-Cornell ME, Frank MV, Mulvihill RJ, Fragola JR. On the current status of Probabilistic Risk Analysis for the US Space Shuttle, Report to the National Aeronautic and Space Administration, Code Q, Washington D.C., February, 2000.
[19] Maggio G. Space Shuttle Probabilistic Risk Assessment: Methodology and Application. Proceedings of the Annual Reliability and Maintainability Symposium 1996:121–32.
[20] Rutledge P, Weinstock R. Quantitative Risk Assessment System (QRAS). Proceedings of the 4th International Conference on Probabilistic Safety Assessment and Management, September 13–18, 1998.
[21] Safie FM. An Overview of Quantitative Risk Assessment of Space Shuttle Propulsion Elements. Proceedings of the 4th International Conference on Probabilistic Safety Assessment and Management, September 13–18, 1998.
[22] Frank M. Applications of Technical Risk Assessment in Aerospace. European Space Agency Risk Management Workshop, ESTEC, March 30–April 2, 1998:43–66.
[23] Baker J. Space Shuttle Range Safety Hazards Analysis. Technical Report 81-1329, prepared for NASA, KSC, J. Baker (author), John Wiggins Inc., 1981.
[24] Rogers W. et al. Report of the Presidential Commission on the Space Shuttle Challenger Accident, Washington D.C., 1986.
[25] Kaplan S. On the Inclusion of Precursors and Near-Miss Events in Quantitative Risk Assessments: A Bayesian Point of View and a Space Shuttle Example. Reliability Engineering and System Safety 1990;27:103–15.
[26] Pinkus RL, Shuman LJ, Hummon NP, Wolfe H. Engineering Ethics: Balancing Cost, Schedule, and Risk- Lessons Learned from the Space Shuttle. Cambridge: Cambridge University Press, 1997.
[27] National Research Council (NRC). Post-Challenger Evaluation of Space Shuttle Risk Assessment and Management. Committee on Shuttle Criticality Review and Hazard Analysis Audit of the Aeronautics and Space Engineering Board, National Academy of Sciences, National Research Council, National Academy Press, Washington, DC, January, 1988.
[28] Slay, et al. Space Shuttle Risk Assessment Proof-of-Concept Study, Auxiliary Power Unit and Hydraulic Power Unit Analysis Report. McDonnell Douglas Corp., December 18, 1987.
[29] Plistiras J., et al. Space Shuttle Main Propulsion Pressurization System Probabilistic Risk Assessment, Final Report. Lockheed Corporation, Palo Alto, CA, 1988.
[30] Buchbinder B. Independent Assessment of Shuttle Accident Scenario Probabilities for the Galileo Mission, Volume 1. NASA/HQ Code QS, Washington DC, 20546, April, 1989.
[31] SAIC. Probabilistic Risk Assessment of the Space Shuttle Phase 1: Space Shuttle Catastrophic Failure Frequency Final Report, 1993.
[32] Paté-Cornell ME, Fischbeck PS. Probabilistic risk analysis and risk-based priority scale for the tiles of the space shuttle. Reliability Engineering and System Safety 1993;41:221–38.
[33] Paté-Cornell ME, Fischbeck PS. 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 1993;41:239–57.
[34] SAIC. Probabilistic Risk Assessment of the Space Shuttle, 1995.
[35] U.S. Nuclear Regulatory Commission (USNRC). Reactor Safety Study: Assessment of Accident Risk in U.S. Commercial Nuclear Plants, WASH-1400 (NUREG-75/014). Washington, DC: U.S. Nuclear Regulatory Commission, 1975.
[36] U.S. Nuclear Regulatory Commission (USNRC). PRA Procedures Guide, NUREG/CR-2300. Washington DC: U.S. Nuclear Regulatory Commission, 1983.
[37] U.S. Nuclear Regulatory Commission (USNRC). Procedural and Submittal Guidance for the Individual Plant Examination of External Events (IPEEE) for Severe Accident Vulnerabilities, Final Report. Washington, DC: U.S. Nuclear Regulatory Commission, 1991.
[38] U.S. Nuclear Regulatory Commission (USNRC). A Technique For Human Error Analysis (Atheana). Washington, DC: Division of Systems Technology, Office of Nuclear Regulatory Research, 1996.
[39] Vesely WE. Fault Tree Handbook. Washington, DC: Office of Nuclear Regulatory Research, 1981.
[40] Mosleh A. Procedure For Analysis Of Common-Cause Failures In Probabilistic Safety Analysis. Washington DC: Division of Safety Issue Resolution, Office of Nuclear Regulatory Research, Nuclear Regulatory Commission, 1993.
[41] Fragola J.R. Space Shuttle Probabilistic Risk Assessment. Proceedings of PSAMIII, Crete, Greece, 1996.
[42] Mosleh A. Personal correspondence describing NASA project work, September, 1998.
[43] Mosleh A. Quantitative Risk Assessment System: Software Requirement, University of Maryland, CTRS A5-5.1, May, 1998.
[44] Mosleh A. Quantitative Risk Assessment System: Software Design, University of Maryland, CTRS A5-5.2, May, 1998.
[45] Howard RA. Information Value Theory in The Principles and Applications of Decision Analysis. Howard RA, Matheson JE (eds.) Palo Alto, CA: Strategic Decisions Group, 1989.
[46] Matheson JE. The Economic Value of Analysis and Computation. In: Howard RA, Matheson JE, editors. The Principles and Applications of Decision Analysis, Palo Alto, CA: Strategic Decisions Group, 1989.
[47] Helton JC. Treatment of uncertainty in performance assessment for complex systems. Risk Analysis 1994;14:483–511.
[48] Paté-Cornell ME. Uncertainties in risk analysis: Six levels of treatment. Reliability Engineering and System Safety 1996;54:95–111.