Towards more accurate life cycle risk management through integration of DDP and PRA

Steven L.Cornford, Todd Paulos, Leila Meshkat, Martin Feather

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109

Abstract

Risk assessment is one of the key inputs to an informed decision making process. Risk management is the continuous application of this throughout the lifecycle. Risk management based on quantitative assessment of risk has been a mainstay of decision making for several decades.

The term "Probabilistic Risk Assessment" (PRA) encompasses a mature and widely used set of techniques for conducting quantitative risk assessments. Originating in the nuclear power industry, they are now regularly applied in nuclear, process, chemical, petroleum, aerospace, and other industries. PRA techniques are especially suited to assessing the risk of a complex system operating in a rich and changing environment. They are able to take into account knowledge of the system's components' reliabilities, knowledge of the environment in which it operates, and knowledge of the design of the system itself. The net result is an assessment and understanding of the reliability of the whole system in its operational context. This gives insight into the strengths and weaknesses of the system – for example, revealing where vulnerabilities exist (thus suggesting key areas for improvement). Information such as this can be key to informed decision making.

One of the challenges for PRA techniques is application early in the life cycle, most especially in the formative stages of development, when the system design is incomplete, immature and/or still in flux. However, this is a time of key decision-making. Those early design decisions will have major consequences that pervade all the subsequent phases: development, test, deployment, operation, maintenance and, ultimately, decommissioning.

NASA's Defect Detection and Prevention (DDP) process has been created to address this need for early lifecycle risk assessment and management. DDP employs a simple form of quantitative analysis, one designed for application at very early stages.

The focus of this paper is on the integration of PRA and DDP. The intent is twofold: to extend risk-based decision though more of the lifecycle, and to lead to improved risk modeling (hence better informed decision making) wherever it is applied, most especially in the early phases as designs begin to mature.

The sections that follow:

1. Summarize the salient points of PRA and DDP.
2. Discuss the relative strengths of each approach, and advantages of their integration.
3. Describe a case study in which separate applications of PRA and DDP on the same system were then compared.
4. Outline the route we are following to integration.

1. INTRODUCTION

In this section, we provide a brief introduction to the Probabilistic Risk Assessment (PRA) and Defect Detection and Prevention (DDP) techniques.

1.1 Probabilistic Risk Assessment

Probabilistic Risk Assessment is a term that is not simply defined, but instead is a term that has to be

explained. Probabilistic Risk Assessment is a comprehensive, structured and logical analysis that can be used to identify and assess risks in complex systems. A PRA examines the reactions of a system to variations in its normal environment; this includes not only perturbations and failures to the nominal mission, but the events needed for mission success as well.

It is well known that the first comprehensive PRA of a technical study is the WASH 1400 [5] study, whose stated purpose was to quantify the risks to the public from commercial nuclear power plant operation. Today, the methodology has evolved and is used to analyze technical systems in nuclear, process, chemical, petroleum, aerospace, and other industries. A good, recent reference on the use of PRA in the aerospace industry has been developed by NASA HQ [3], [4].

Probabilistic Risk Assessment is a scenario based methodology. Scenarios are strings of events that begin with an initiator and lead to some sort of a conclusion, or end state. In between the initiator and end state are pivotal events in the scenario. Pivotal events may either be protective, mitigative, aggravative, or benign. Scenarios can be modeled in many different fashions, but are most commonly modeled through the use of fault trees and event trees. The best way to describe the difference between event trees and fault trees is that event trees show the logical progression of events, while fault trees are snapshots in time, and are used to model events in the event tree.

Event trees are said to be based on inductive, or forward, logic; i.e., the forward thinking represents the possible conditional events in the scenario based on the preceding event, or the possible events that can occur given an initiator. Fault trees are said to be deductive in nature, i.e., they are used to identify all of the possible failure causes of an event from a top down approach. There is no one single way to develop a PRA model and the trade off is that the larger the event tree, the smaller the fault trees, and vice versa. The use of event trees and fault trees and their sizes is up to the analyst, but their sizes are typically decided based upon the PRA methodology used (large event tree versus small event tree), and to facilitate defining a complex world with competing risks into a model with binary decision points.

After the PRA model is complete, it is integrated, quantified, and sanity checked for completeness and accuracy. It should be noted that PRA is an iterative process. No one can work through a PRA model from start to finish without making changes to event trees, fault trees, or end states, etc. It is part of the PRA process.

The simplified steps of a PRA include:

1. Defining the end states or the objective of the PRA.
2. Developing a list of initiating events or mission events.
3. Developing scenarios using event trees and fault trees.
4. Performing a data analysis for the basic events in the PRA model.
5. Integrating the model and quantifying the scenarios.
6. Perform a sanity check on the model considering both logic and quantified results.
7. Perform an uncertainty analysis.
8. Communicate the risk by providing the risk results and insights to program management.

Sadly, the seemingly most over-emphasized aspect of doing a PRA is the "P," or determining what the "number" is. This is the single largest misconception about what a PRA provides. There is an incredible amount of uncertainty in the PRA results, or any reliability analysis, and that needs to be considered and understood. Every PRA has uncertainty, and no PRA is complete without an uncertainty analysis.

The reason for doing PRA is not because it gives the "number," but instead because of the identification and prioritization of the risk contributors, and because of the systems analysis process that forces the engineers and operators to consider the scenarios, how to improve them, how to mitigate them, and how to incorporate or develop recovery actions where possible. The identification and prioritization of risks is critical in determining the most cost-effective solution for mitigating risks. This is the real power of PRA, and it is these advantages that are emphasized that are emphasized in both the NASA PRA guidebook and course [3],[4].

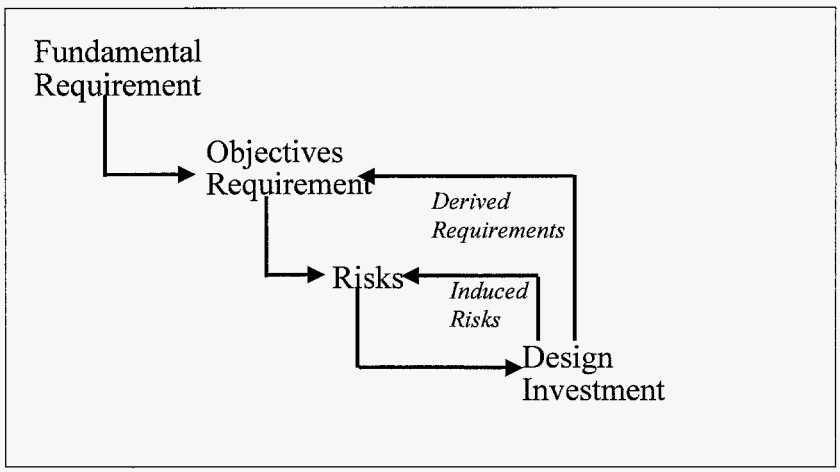

Figure 1 : Requirements Flow Down and Ripple Effects of Option Selection.

1.2 Defect Detection and Prevention

"Defect Detection and Prevention" (DDP), is a simple quantitative risk model designed for application early in the lifecycle, when information is sparse yet the capability to influence the course of the development to follow is large. Dr. Steve Cornford at JPL originally conceived of DDP specifically to facilitate assurance planning [9]. The core idea of DDP is to relate three sets of information:

1. "Objectives" (what you want to achieve).
2. "Risk Elements" (what can get in the way of attaining those objectives).
3. "Investments" (what you can choose to do to overcome the problems).[12]

In DDP, relationships between these items are *quantitative* (e.g., *how much* a Risk Element, should it occur, detracts from an Objective's attainment). Such a quantitative treatment is key to DDP's realization of the vision of "risk as a resource", as espoused in [9]. This is one of key the ways that DDP differs from many of the purely qualitative approaches (e.g., QFD [12]) usually employed early in the life cycle.

Cornford's initial experiments used Microsoft Excel® spreadsheets to manually explore the utility of the process. Positive results then led to development of custom software for the DDP process [1]. Supported by this software, DDP has been applied to assess the viability of, and planning for, the development of novel technologies and systems for use on space missions [Cornford et al, 2001], [8].

Defect Detection and Prevention is a life cycle risk management technique that aims to provide a detailed description of the success requirements of a mission (OBJ's), the failure modes (RE's), the options available for decreasing the probability of occurrence of the failure mode or solution options (INV's), and the relationships between them. Further, DDP provides an optimization scheme that determines the optimal combination of INV's to employ for attaining an optimal risk and cost level based on the preferences and constraints established by the decision maker. The preferences, in turn, are captured in the weights allocated to the requirements and the consequences of various failure modes by the decision maker. The constraints are placed on the optimizer.

DDP uses a variety of visual tools to illustrate the various aspects of the problem and aid the user to understand the often complex interrelationships between the OBJ's, RE's, INV's and risk and cost measures. These visual aids include weighted trees for each of the OBJ's, RE's and INV's, charts

---

[2] In previous papers on DDP these three sets of information were referred to as "Requirements", "Failure Modes" and "PACTs" respectively. The switch of terminology reflects application of DDP to areas more broad than implementation phase assurance planning. Investments refer to all of the possible activities which can detect, prevent (reduce probability of occurrence) and alleviate (reduce impact of occurrence).

showing the effects of each on the other and a variety of other matrices, charts and dashboards.
The risk associated with a risk element is defined as the product of the likelihood of occurrence of the failure mode and its consequence. Risk elements are classified into programmatic, technical, infrastructure, management and constraints and each of these classes are shown using a different color. The main classes of solution options are considered to be technology investments, design/architectural options, tests, analyses, process controls, and operational solutions. Failure is used in its broadest sense. A failure implies any event that could result in not meeting a goal or requirement. Requirements/objectives, in turn, reflect the goals of the mission. **Figure 1** shows the process of designing a mission (or other object) using the DDP process. We start by identifying the fundamental requirements. The objectives of the project and lower level requirements are derived from these fundame ntal requirements. The events that can lead to the non-fulfillment of the requirements or the risk elements are then identified. Design choices are made to reduce the identified risks. These design choices, in turn, may introduce new risks and/or derived requirements. Therefore the mission design process is more cyclic than hierarchical and it takes a few cycles to refine the initial design and produce an acceptable design. The mission design process is dynamic in nature and DDP is capable of easily capturing the refinements and modifying its initial model.

In particular, one of the most powerful aspects of the DDP process is the explicit inclusion of the investments that can be used to reduce the likelihood and/or impact of the various risk elements. The users can now explicitly examine the planned activities to ensure they are focused on the right elements of the design and explore various combinations of activities to mitigate the risks. Each of these investments has resource costs (e.g., mass, cost, power) associated with them and the tool provides a running total of the resources allocated to various investments. The DDP tool has been used as a front-end to provide quick, near real-time identification of a prioritized risk element list as well as the most promising investments. The tools allows the users to identify areas for additional work and it is this feature which will be exploited to identify areas most benefiting from more detailed, PRA analysis.

Once an acceptable design is created, it may be desirable to apply the sophisticated analytical techniques used in PRA analyses to further study this design. This process is illustrated in **Figure 2**.

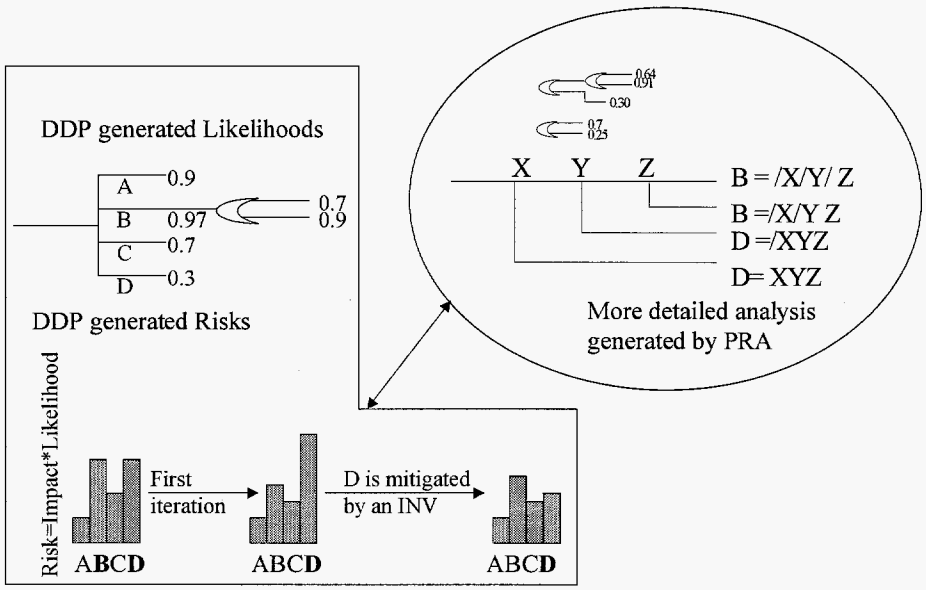

Figure 2 Illustration of the initial risk prioritization process using DDP and then refined Fault Trees and likelihood estimates generated by PRA and integrated back into the DDP evaluation.

In this figure, an initial DDP analysis on the system produces rough estimates of risk in four areas: A, B, C & D. These estimates indicate that areas B and D are somewhat high risk. A detailed PRA analysis of these two areas determines that area B is somewhat lower in risk than the initial rough estimate produced by DDP and area D is somewhat riskier. This information is fed back into DDP and the decision maker selects an investment option to mitigate the risk imposed by D. Once this mitigation scheme is taken into consideration, the relevant information is fed back into PRA and the exact likelihood measures are computed and fed into DDP. DDP in turn updates it's risk profile accordingly. This process continues until the decision maker is satisfied with the risk profile.

2. Risk Management for Large Scale Systems

Large scale systems are composed of many interdependent components and modules. Risk management for such systems entails the identification of the risk elements, the options available for mitigating these risks, the interdependencies between the risks and the solution options and the preferred solution options to be employed throughout the life cycle of the system, taking into consideration the many dimensions of these systems. Because of the inherent complexity of large scale systems and the fact that these systems are dynamic, risk management should also be an evolutionary process. Therefore, there is the need for a process that can provide clear and continuous insight into the evolving risk landscape of such systems.

The DDP process is a life cycle risk management technique that captures the many dimensions of large scale systems. The user identifies the requirements for system success, the factors that could lead to these requirements not being met, and the techniques that could be employed for increasing the probability of requirement satisfaction. Each of the requirements, risk elements, and solution options have many different classes within themselves. The various dimensions of each of these elements and the interaction within and between these elements are captured in trees, charts, matrices and dashboards within the context of DDP.

Probabilistic Risk Assessment techniques, on the other hand, provide a more formal approach for risk analysis that analyzes the interactions within large systems, and how the system operates. This approach is based on considering the various events and failures that could occur and the scenarios that would result from these occurrences. This methodology does provide the likelihood of each scenario, but it is really the systems analysis and risk prioritization that is the power of doing PRA. The earlier a PRA is developed for the project, the more it can help. However, there will be more uncertainty than towards the end of the project when the design is more complete, the system interactions are better understood, and more effort can be placed in the data analysis. It makes sense to incorporate a qualitative perspective from DDP when dealing with these prioritized risks due to the uncertainty, and due to the fact that the model results will be highly sensitive to the single point failures that always exist in aerospace flight projects. It is very difficult to work a PRA, or any other quantitative analysis, real time like a qualitative analysis, and working risk management issues using DDP with the project real time can be very helpful in identifying, prioritizing, and managing risks.

Using the multiple perspectives of DDP and PRA deepens our insight about the system and allows for a broad and formal analysis of the inherent risks involved in large scale systems. In the following section, we study the interplay between DDP and PRA and construct a unified approach for risk management of complex systems which builds on the strengths of both these techniques.

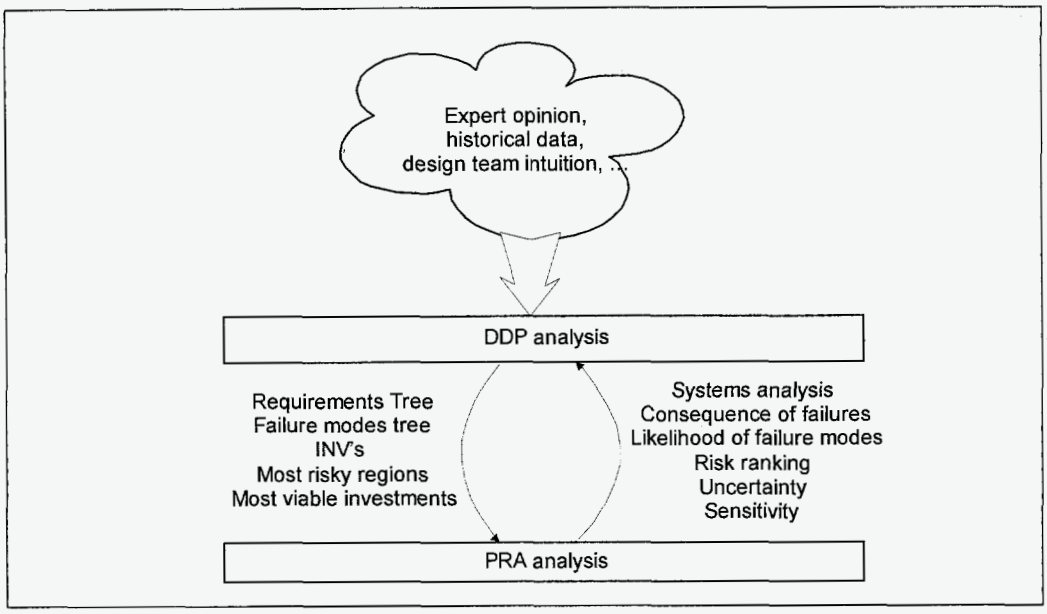

Figure 3: High level interplay of PRA and DDP for Risk Management

3. Best of both worlds: Combining DDP and PRA

In this section, we propose combining the techniques used in the Defect Detection and Prevention process and the Probabilistic Risk Assessment process for conducting a life cycle risk management for large scale systems. First, we study the interplay between DDP and PRA from a high level and then we consider the steps involved in each and identify the areas where each could lend to the other. This leads us to a unified approach for conducting risk management for large scale systems by building on the strengths of both PRA and DDP.

Figure 2 shows the high level process of risk management for large scale systems using DDP and PRA. Initially, the available information obtained from various sources including expert opinion, historical data, and design team intuition is entered in the DDP software. The DDP analysis then yields various trees and charts and helps identify the most risky regions of the project. These regions are then further analyzed for a better estimate of the probability of occurrence and consequences using a PRA study. The likelihood measures are then fed back into DDP. DDP then refines the project portfolio based on this new information. This process continues until the decision maker is satisfied with the investment portfolio created in DDP and the levels of risk that are determined by the combination of DDP and PRA analysis.

Figure 3 shows our unified approach for risk analysis within the context of DDP and PRA from a lower level of detail. Initially, the requirements for the success of the project and their respective weights are identified and input to DDP. The factors that could lead to the non-fulfillment of these requirements or the failure modes of the project are also identified in DDP. For the most important failure modes, the PRA can analyze these in more detail and provide a better understanding of what the risk is. These likelihood and consequences are then fed back into DDP. Solution options, or techniques that are used to reduce or eliminate the probability of failures and the effects of the RE's, OBJ's and SO's on each other are determined in DDP. Again, for the most risky areas identified in DDP, the associated events and failure modes are identified and a more formal analysis for finding their respective probability of occurrence is conducted using PRA techniques.

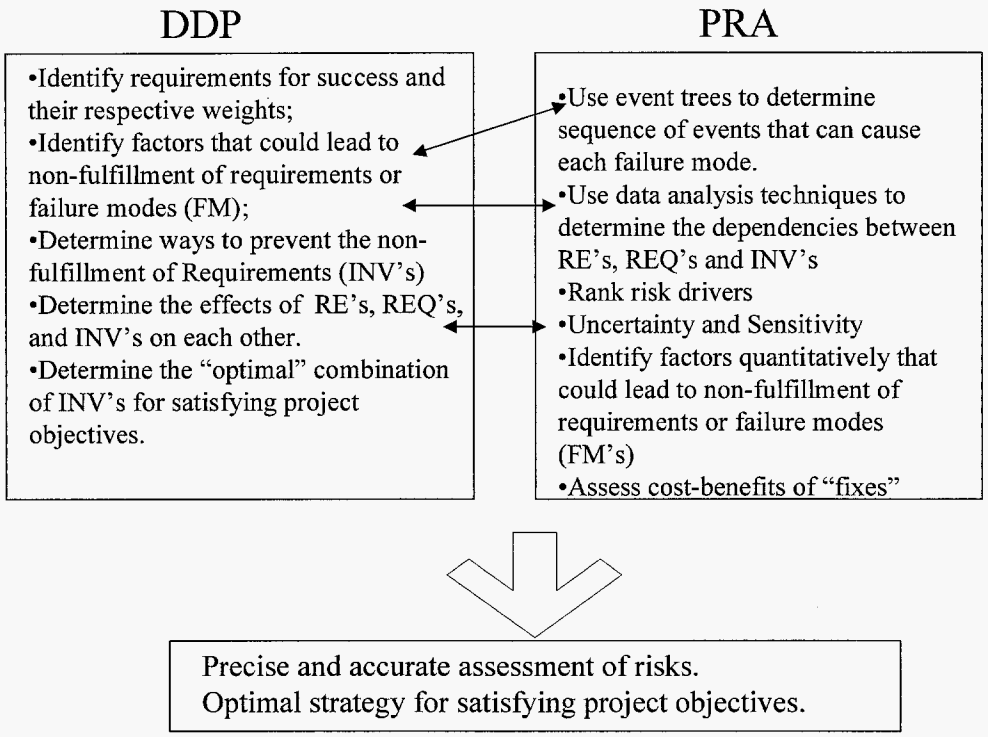

Figure 4: Unified Approach for Risk Analysis using DDP and PRA

One difference to consider between PRA and DDP is that DDP is basically the 'breadth" of the risk management analysis, and PRA is the 'depth.' The PRA analysis can be partially directed by DDP, forcing the analysis to an appropriate level of detail for the current project phase. In turn, PRA can analyze the entire list of risks, better prioritize the list by a more complete examination of the failures, and to ensure that the risk list is as complete as possible. This is the depth of the risk analysis, but being careful not to go into unnecessary detail. This is a difficult topic to teach to PRA practitioners: when is the detail sufficient [3,4]. There are many examples of PRAs that have gone into too much detail, and this is one reason why PRA is attacked by many critics for being too costly or not providing the information that is needed. This isn't the fault of the methodology, but the fault of the inexperienced or under-worked practitioner. It is the desire of this effort, that by merging DDP and PRA we can optimize the efforts of each.

3.1. Example applications

For a pilot project, the Mars Smart Lander (Mars '09, or MSL) project was used as a test case of merging PRA and DDP. The first step in the process was for the respective PRA and DDP camps to perform their analysis as their usual mode of doing business. As a first cut, the MSL PRA used the Mars Exploration Rover (Mars '03, or MER) PRA model as a starting point, and incorporated changes where changes where known to exist. For example, MSL is a nuclear powered rover that is three-axis stable, and thus has a slightly different Entry, Descent and Landing Sequence and a different design. Where design details did not exist, such as deployment mechanisms, or specific propulsion system characteristics, the MER model was used. This allowed for a very rapid PRA model to be developed. This also allowed for all of the MER lessons learned to be carried through immediately.

The DDP camp, on the other hand, met with the program management and the COG engineers and listed the perceived risks and elicited the levels of risks as perceived by them. They further discussed the risk management plans for the next year and discussed different ways of balancing the risk within the constraints of the budget.

Once the two analyses were completed, the two analyses could be compared and reiterated. For example, from the DDP analysis, it was learned that the LIDAR mirror could be destabilized during powered descent; this was not considered in the original MER PRA analysis since MER relies mostly on a radar altimeter setup. Other issues that came forth in the DDP analysis were long life issues with certain devices, such as brushless motors, lubricants, bearings, dust sealant systems, actuators, and electronic packaging. MSL will have a two year life on Mars, while the MER model only considered a 90 day mission. These are issues that will have to be addressed in MSL PRA updates, but for the initial PRA analysis, there was insufficient knowledge and time to adequately study these issues; however, the DDP program allows them to be tracked even though they didn't show up as a drive in the PRA analysis.

On the other hand, the DDP analysis did not determine that engine out scenarios during powered descent could pose control problems per se; instead the DDP analysis focused on generic issues such as "engine issues," "throttle valve issues," and "hi-flow regulator issues." These piece part issues did not consider the vehicle affects like the PRA did, which showed that they all contribute to the Entry, Descent and Landing risk. The PRA also had a much more detailed model of the Deployment Phase, which due to the number of pyro events that must be performed in sequence, and the number of mechanisms involved, is on the MER list of risk drivers. Instead the DDP analysis looked at piece part or top level issues that were based on limited design knowledge since that part of the design had not been developed yet.

After comparing notes, both analyses could be updated to consider all of the risks, and the focus could begin on risk prioritization through better systems and data analysis. There are many cases where project management felt that certain risks are critical, so the PRA model has been tagged to analyze them in more detail. As the project design detail emerges, the risk level can be continually updated, monitored and compared to determine the most cost-effective solutions to all of the risks.

Similarly, the DDP efforts can begin to incorporate all of the risks, primarily the technical risks which are so often forgotten this early in the program. The DDP program also has the flexibility to consider all risks, technical, cost and programmatic, so that its resources can address risks across the board. The PRA does not explicitly handle cost, programmatic or technology development risks. Those risks are much better suited to be handled in DDP.

Since the project is in the early phase of development, its risk becomes one of technical and cost considerations. As the program progresses, the risk focus will change as well to be more focused on technical issues as the vehicle goes through more detailed design. Then towards the later phases, the projects risk focus will again change back to schedule and cost as it gets closer to its launch date. This process should continue through the life of the project.

4. Future directions

Although we have presented an approach for combining DDP and PRA, there are still many unanswered questions and this is an ongoing research activity. In the future, we plan to elaborate on our approach and implement this approach in DDP.

5. Summary & Conclusions

Reducing risk is an important concern in development of complex systems. Risks encompass both those in the developed system itself (e.g., that it fails in some catastrophic manner), and those in the development activity (e.g., that it overruns budget and schedule).

There are three kinds of tactics to reduce risk: *prevention* (reducing the likelihood of the risk arising in the first place), *alleviation* (reducing the severity of the risk should it occur) and *detection & repair* (detecting the presence of a risk and repairing it, thus decreasing the likelihood of that risk persisting).

Our vision is to combine aspects of both DDP and PRA so as to apply probabilistic reasoning in estimation of the efficacy and cost of risk reduction strategies that employ a mix of these risk reduction tactics. The reasoning would take into account the logical fault tree structure of risks, and estimates of the efficacy and cost of individual risk reduction tactics. The advantages of this approach is that by

combining the two methods, both the qualitative and quantitative analyses will be more thorough and unite the two risk camps, provide a means for more effective risk communication, better understand the cost-benefits of various risk reduction strategies, and give the project the ability to optimize such strategies (e.g., minimize risk for a given cost) and treat risk as a tradable parameter (e.g., accept more risk in order to reduce cost).

6. References:

[1] M.S. Feather, S.L. Cornford, M. Gibbel. "Scalable Mechanisms for Goals Interaction Management", *Proceedings 4th IEEE International Conference on Requirements Engineering,* Schaumburg, Illinois, 19-23 Jun 2000, IEEE Computer Society, pp 119-129

[2] Martin S. Feather, Steven L. Cornford, Julia Dunphy, "Cost-Benefit Based Assurance Planning", *International Workshop on Model Based Requirements Engineering Proceedings*, 2001.

[3] NASA's Third Workshop for Probabilistic Risk Assessment Methods (PRAM-3) for Managers and Practitioners, given at the Embassy Suites Hotel, Arcadia, California, June 3-6, 2002.

[4] Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners, version 1.1, prepared for the Office of Safety and Mission Assurance, NASA HQ, Washington, DC, August 2002.

[5] *Reactor Safety Study*, Report WASH-1400, Nuclear Regulatory Commission, 1975.

[6] S.L. Cornford, J. Dunphy, and M.S. Feather: "Optimizing the Design of end-to-end Spacecraft Systems using risk as a currency", *IEEE Aerospace Conference*, Big Sky, Montana, 2002

[7] S.L. Cornford, M.S. Feather & K.A. Hicks. "DDP – A tool for life-cycle risk management", *Proceedings, IEEE Aerospace Conference*, Big Sky, Montana, Mar 2001, pp. 441-451.

[8] S.L. Cornford, M.S. Feather, J.C. Kelly, T.W. Larson, B. Sigal & J.D. Kiper: "Design and Development Assessment", *Proceedings, 10th IEEE International Workshop on Software Specification and Design*, San Diego, California, 5-7 Nov 2000, pp 105-204.

[9] S.L. Cornford: "Managing Risk as a Resource using the Defect Detection and Prevention process", *Proceedings, 4th International Conference on Probabilistic Safety Assessment and Management*, 13-18 September 1998, New York City, NY, International Association for Probabilistic Safety Assessment and Management.

[10] Steven L. Cornford, Martin S. Feather, Kenneth A. Hicks, "DDP – A Tool for Life-Cycle Risk Management", *2001 IEEE Aerospace Conference Proceedings*, Big Sky, Montana, January 2001

[11] Steven. L.Cornford, Julia Dunphy, and Martin S. Feather, "Optimizing the Design of Spacecraft Systems using risk as currency", January 2002, IEEE.

[12] Y. Akao. "*Quality Function Deployment*", Productivity Press, Cambridge, Massachusetts, 1990.

7. Biographies

Steven L.Cornford
MS 179-224
(818)354-1701
Steven.Cornford@jpl.nasa.gov

Steven Cornford *is a Senior Engineer in the Strategic Systems Technology Program Office at NASA's Jet Propulsion Laboratory. He graduated from UC Berkeley with undergraduate degrees in Mathematics and Physics and received his doctorate in Physics from Texas A&M University in 1992. Since coming to JPL he focused his early efforts at JPL on establishing a quantitative basis for environmental test program selection and implementation. As Payload Reliability Assurance Program Element Manager, this evolved into establishing a quantitative basis for evaluating the effectiveness of overall reliability and test programs as well as performing residual risk assessments of new technologies. This has resulted in the Defect Detection and Prevention (DDP) process is the motivation for this paper. He received the NASA Exceptional Service Medal in 1997 for his efforts to date. He has been an instrument system engineer, a test-bed Cognizant Engineer and is currently involved with*

improving JPL's technology infusion processes as well as the Principal Investigator for the development and implementation of the DDP software tool.

Martin S. Feather
MS 125-233
(818)354-1194
Martin.S.Feather@jpl.nasa.gov

Martin Feather *is a Principal in the Software Quality Assurance group of NASA's Jet propulsion Laboratory. He works on developing research ideas and maturing them into practice, with particular interests in the areas of software validation (analysis, text automation, V&V techniques) and of early phase requirements engineering and risk management. He obtained his BA and MA degrees in mathematics and computer science from Cambridge University, England, and his PhD degree in artificial intelligence from the University of Edinburgh, Scotland. Prior to joining JPL, Dr. Feather worked on NSF and DARPA funded research while at the University of Southern California's Information Sciences Institute. For further details, see* http://eis.jpl.nasa.gov/~mfeather

Leila Meshkat
MS 301-180
(818)393-7378
Leila.Meshkat@jpl.nasa.gov

Leila Meshkat is a Mission Systems Engineer at the Mission & Systems Architecture Section of the Jet Propulsion Laboratory. Prior to joining JPL, she was a Research Associate at the University of Southern California's Information Sciences Institute and a Lecturer at the USC School of Engineering. She holds a Ph.D in Systems Engineering from the University of Virginia, an MS in Operations Researh from the George Washington Univesity and a B.S. in Applied Mathematics from the Sharif University of Technology. Her current research interests include Reliability and Risk Analysis and she is a member of IEEE and Omega Rho.

Dr. Todd Paulos
282 Cliffwood Drive
Simi Valley,
CA 93065
tpi@ix.netcom.com
805-522-9300

Todd Paulos *is an independent consultant with over 10 years of experience in the field of Reliability Engineering and PRA. Dr. Paulos received his B.S. degree in Engineering from Harvey Mudd College, and M.S. and Ph.D. degrees in Mechanical Engineering from the University of California at Los Angeles. Dr. Paulos helped developed a PRA guidebook and PRA course for NASA Headquarters, and managed PRAs for several JPL programs such as the Mars Exploration Rover (Mars '03), the joint Mars '07 mission with the French space agency CNES, Mars Smart Lander (Mars '09), Mars Sample Return (Mars '13, maybe), CloudSat, GRACE, and Herschel/Planck missions.*

Prior to becoming an independent consultant, Dr. Paulos worked at the Lockheed Martin Skunk Works where he led the PRA efforts on the X-33 and Reusable Launch Vehicle Programs, and on other aerospace programs.

Dr. Paulos is a member of American Institute for Astronautics and Aeronautics (AIAA), the American Society for Mechanical Engineers (ASME), the Aircraft Owners and Pilots Association (AOPA), and the Southern California Society for Risk Analysis (SCSRA).