

IMECE2003-42749

RISK-BASED ANALYSIS AND DECISION MAKING IN MULTI-DISCIPLINARY ENVIRONMENTS

Martin S. Feather

Steven L. Cornford

Kelly Moran

Jet Propulsion Laboratory, California Institute of Technology

ABSTRACT

A risk-based decision-making process conceived of and developed at JPL and NASA, has been used to help plan and guide novel technology applications for use on spacecraft. These applications exemplify key challenges inherent in multi-disciplinary design of novel technologies deployed in mission-critical settings:

- 1) Cross-disciplinary concerns are numerous (e.g., spacecraft involve navigation, propulsion, telecommunications). These concerns are cross-coupled and interact in multiple ways (e.g., electromagnetic interference, heat transfer).
- 2) Time and budget pressures constrain development, operational resources constrain the resulting system (e.g., mass, volume, power).
- 3) Spacecraft are critical systems that must operate correctly the first time in only partially understood environments, with no chance for repair.
- 4) Past experience provides only a partial guide: New mission concepts are enhanced and enabled by new technologies, for which past experience is lacking.

The decision-making process rests on quantitative assessments of the relationships between three classes of information – objectives (the things the system is to accomplish and constraints on its operation and development), risks (whose occurrence detracts from objectives), and mitigations (options for reducing the likelihood and/or severity of risks). The process successfully guides experts to pool their knowledge, using custom-built software to support information gathering and decision-making.

Keywords: Risk, Decision-making, Design, Novel Technology

INTRODUCTION

NASA's Mission statement reads: *"To understand and protect our home planet. To explore the Universe and search for life. To inspire the next generation of explorers . . . as only*

NASA can." In April 2002 the NASA Administrator stated: *"... In broad terms, our mandate is to pioneer the future . . . to push the envelope . . . to do what has never been done before. ..."* [1]. These quotes emphasize the novel aspects of NASA's activities. Spacecraft design involves multiple disciplines (e.g., interplanetary navigation, spacecraft propulsion, telecommunications, properties of materials in deep space environments). Spacecraft development and operation are driven by budget and time concerns (e.g., planetary configurations dictate launch windows). Spacecraft operate remotely from earth; so third-party repair is almost always impossible. Their purpose is to yield information of distant, unknown environments, meaning they often must operate in those same unknown environments. The aim to improve the quantity and quality of the science information they gather and in return drives the use of novel technologies in novel ways.

At first glance, these would seem to imply that spacecraft development and operation has little in common with earthly activities. Certainly the specifics of the information involved (e.g., behavior of materials in a deep space environment; interplanetary navigation) are particular to space exploration. However, the fundamental challenges that underlie NASA's activities are familiar to a wide range of endeavors – the challenges stem from the need to do multi-disciplinary design of novel technologies intended for deployment in critical settings. They are:

- Multiple disciplines are involved. No single individual has depth of knowledge in every one of these disciplines. Furthermore, these disciplines are cross-coupled and interact in multiple ways. As a consequence, there is the need to pool the knowledge from multiple discipline experts, and to conduct decision-making taking the sum total of this knowledge into account.
- Time and budget pressures constrain development (an almost universal phenomenon, e.g., development of almost

any product these days is subject to time to market needs, and aspects of competitiveness). Limited operational resources constrain many systems (e.g., think of power, mass and volume constraints on hand-held devices).

- Depending upon the purposes for which systems are employed, the consequences of failure during their operation could range from a minor nuisance to catastrophic. Spacecraft represent an unusual extreme in which the entire system itself will be lost if things go badly wrong. On earth, the systems themselves may be more repairable, but the not necessarily the consequences of their operational failure.
- Spacecraft must operate in a harsh and often unpredictable physical environment. While more may be known about the physical operating environment for Earth-bound systems, considerable uncertainty arises from the unpredictability of the ways in which users will make use of systems. In the extreme, a deployed system will induce a change in the environment around it, leading to an unpredictable feedback loop (e.g., the phenomenon of “E-type Software” described by Lehman [2]).
- Past experience provides only a partial guide as to how to develop the new system. For NASA’s missions, new mission concepts are enhanced and enabled by use of new technologies, and/or of established technologies applied in novel ways. Past experience is lacking in either case. The same is true for almost any novel solution to a real world problem.

In response to these challenges, at JPL and NASA we have been developing and applying a risk-based approach to analysis and decision making for novel system applications. This paper summarizes the approach, and indicates how it addresses the challenges listed above.

NOMENCLATURE

Defect Detection and Prevention (DDP) – a risk-based decision-making process conceived of and developed at JPL and NASA.

Objectives – the things the system/technology is to accomplish and constraints on its operation and development.

Risks – the things whose occurrence would detract from attainment of Objectives.

Mitigations – options for reducing the likelihood and/or severity of Risks.

A RISK-BASED ANALYSIS AND DECISION-MAKING APPROACH

The focus of this paper is on a risk-based analysis and decision-making approach developed and applied at JPL and NASA. The approach is called “Defect Detection and Prevention (DDP)”. The name reflects its origins as a structured method for planning the quality assurance of hardware systems [3]. Since then its scope has expanded to also encompass decision-making earlier in the development lifecycle, and to be applicable to software, hardware and systems [4].

The approach has been generalized to aid in the decision-making during the early phases of advanced technology and system development. Decisions made in these early phases are important because they have the most leverage to influence the development to follow. However, decision-making during these phases is challenging because information on which to base those decisions is incomplete and uncertain, and in the case of

advanced technologies and systems, there is little past experience from which to extrapolate.

DDP’s Simple yet Quantitative and Detailed Risk Model

DDP offers a conceptually simple yet quantitative, detailed model of risk as foundation for reasoning:

- The simple conceptual core of the DDP risk model rests on three sets of information: *objectives* – the things we want the system/technology to achieve, *risks* – the things that, should they occur, detract from the attainment of objectives, and *mitigations* – the things that we could choose to do to reduce risks (by reducing their likelihood of occurrence and/or their severity should they occur).
- The quantitative treatment derives from information on how much each risk (should it occur) detracts from attainment of each objective, and by how much each mitigation (should it be applied) reduces each Risk.
- The detail arises from the ability to populate the DDP model with system/technology specific objectives, risks and mitigations. These are open-ended sets of information, and the level of detail to which these are populated is not pre-ordained, but rather is determined in the course of the DDP application process, to the extent necessary to support decision-making.

The combination of these aspects yields an approach that can encompass a wide range of concerns – many problems can be cast in terms of objectives (what do we want), Risks (what can get in the way) and mitigations (what can we do about it). Their quantitative treatment allows reasoning about their net combined effect. For example, the biggest risks are those with the greatest sum total adverse impact on attainment of objectives. The ability to incorporate problem-specific detail permits the representation and reasoning over system, technology and application-specific nuances.

The later sections will examine the ways in which the DDP information model helps address the key challenges of multi-disciplinary decision making. In preparation, we now expand upon the description of the DDP model.

Details of DDP’s Risk Model

In more detail, DDP deals with the following three sets of information:

- “**Objectives**” – these encompass all of the things the system/technology is to achieve, how and under what conditions it is to be operate, and how it is to be developed. In our area of spacecraft technologies, typical examples include science return objectives (e.g., quantity and resolution of data), aspects of its operational environment (e.g., ambient temperature, operating temperature, available electrical power), and nature of its development (e.g., must be ready for launch 2 years from now). Since not all objectives are equally important, DDP allows users to assigned “Weights” (numerical values of a unitless scale) to reflect their relative importance. For example, minimal success criteria (“must-have” capabilities) can be assigned high weights, while optional criteria (“nice to have” capabilities) can be assigned lesser weights.
- “**Risks**” – broadly speaking, all the problems that, should they occur, will adversely impact attainment of objectives. For hardware, these cover the gamut of failure modes of

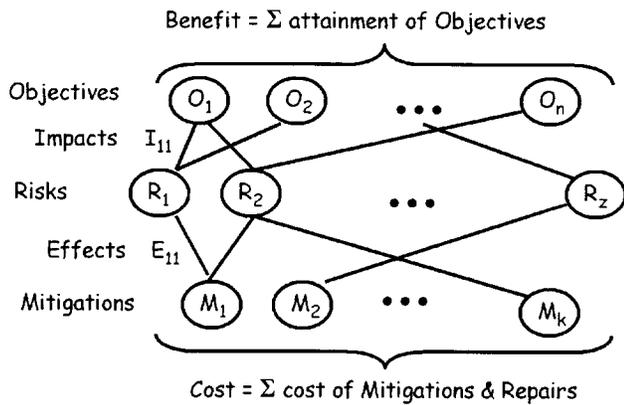


Figure 1. Topology of DDP information model

design, assembly (won't fit, too heavy, etc.) and operation (cracks, breaks, short-circuits etc.). For software, these include development risks (ambiguous requirements, flawed designs, poor contract arrangements, lack of configuration management) and risks in the developed software (erroneous results, deadlocks, timing errors, poor performance, mis-use of computing resources). Human and organizational concerns can also be incorporated here (e.g., key personnel unavailable; vendor goes out of business).

- “**Mitigations**” – the entire range of actions and options to reduce risks, including preventative measures (e.g., training of personnel, use of design standards), alleviations (e.g., hardware redundancy; software bounds checking that make a software component more resilient to erroneous inputs), and detections (e.g., design reviews, code inspections, analyses, tests of all kinds) that can potentially detect problems ahead of their manifestation in use, and so allow for their correction. There are almost always resource costs associated with mitigations (e.g., budget and schedule; for spacecraft hardware, power, volume and mass are perennial concerns; for spacecraft software, CPU and memory utilization), and one of the primary purposes of DDP is to help guide the selection of mitigations that reduce risk in a cost-effective manner.

In DDP, information in these three categories is *quantitatively* linked. Risks are quantitatively related to objectives to indicate *how much* a risk, should it occur, adversely impact an objective's attainment. Mitigations are quantitatively related to risks to indicate *how much* a mitigation, should it be applied, will reduce a risk (by preventing it from occurring in the first place, by detecting it and so allowing for its repair, or by reducing its impact on objectives).

The overall structure of a DDP information model is relatively simple – see Fig. 1. However in practice there can be numerous items and links among them. For example, the topology of the data in a recently completed DDP study (of a technology intended for spacecraft application) is shown in Fig. 2. This comprises 29 objectives, 58 risks, 36 mitigations, and some 900 links among them.

MULTI-DISCIPLINARY CONCERNS

The single most important aspect of the DDP approach is that it supports multiple experts pool their knowledge and allows them to take the sum total of their pooled knowledge

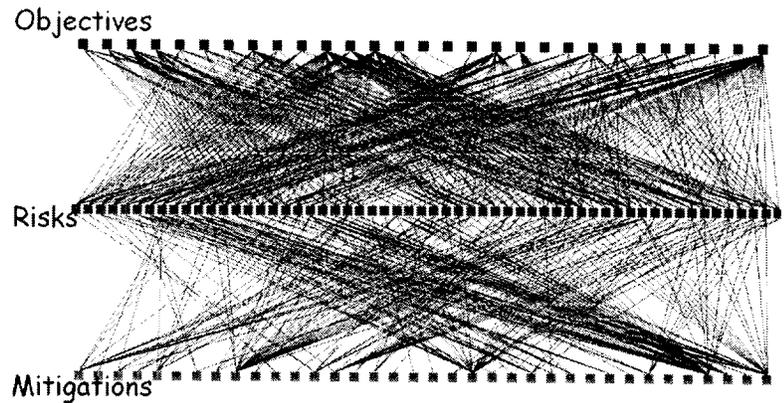


Figure 2. Topology of data in a completed DDP model

into account as they make decisions. This helps address two key concerns of multi-disciplinary studies, namely:

- No single individual has depth of knowledge in every one of the disciplines, and
- Disciplines are cross-coupled and interact in multiple ways.

DDP enables pooling of multiple experts' knowledge by utilizing:

- a structured group process to elicit in-depth information from experts,
- software support to perform calculations and search over the entirety of the gathered information (e.g., calculate aggregate risk information, search for near-optimal solutions), and
- cogent visualizations to present the information back to the experts and so retain their engagement in a risk-informed decision-making process.

The DDP process is conducted in a group setting in which all the experts relevant to the technology/system under consideration are simultaneously involved. It is crucial the experts' combined areas of expertise span the concerns of the study, and that those experts are able and willing to contribute their expertise. In this respect DDP is akin to traditional risk assessment methods, which gather their risk information from representatives of all relevant disciplines. However, DDP differs by seeking more fundamental kinds of information than do typical risk assessment methods. This provides opportunities for experts to contribute specific and detailed knowledge. The DDP software tool performs the calculations to pool their separately contributed items, and present that information back to the experts.

DDP's use of fundamental risk information

Most risk assessment methods take as starting point a technology/system *incorporating* all the risk reducing measures to be applied to make that technology/system reliable. They focus on the risks *remaining* in that technology/system. In particular, these methods ask experts to provide, for each risk, estimates of the remaining likelihood and impact of that risk. The product of these is the usual definition of risk (sometimes referred to as “risk exposure”), formula (1):

$$\text{risk} = \text{likelihood} \times \text{severity (a.k.a. consequence)} \quad (1)$$

In contrast, DDP takes as starting point a technology/system *separate* from the risk reducing measures expected to be applied. Those are represented separately,

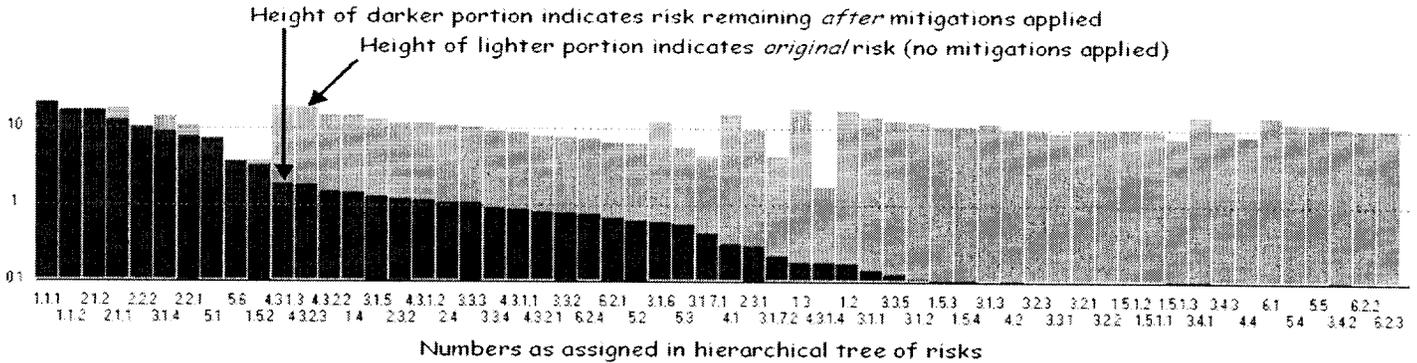


Figure 3. Graphical presentation of risks sorted into descending order of remaining risk levels

explicitly accounting for their risk reducing effects. The risks remaining in the design are *calculated* from the a-priori risks, and the risk reductions effects of the mitigations. In detail, DDP asks experts to provide, for each risk, estimates of:

- what would be its likelihood of occurrence were nothing done to prevent it, and
- what adverse impacts it would have on objectives should it occur, and
- what likelihood- and severity- reducing effects each of the mitigations would have on that risk were they to be applied.

The first two are combined to calculate the *unmitigated* contribution of a risk, using the formula (2):

$$\text{Unmitigated-Risk}(R) = \text{A-Priori-Likelihood}(R) \times \sum (O \in \text{Objectives}) : \text{Impact}(R, O) \times \text{Weight}(O) \quad (2)$$

where A-Priori-Likelihood(R) is risk R's likelihood were nothing done to prevent it, Impact(R, O) is the proportion of the objective O that would be lost were R to occur, and Weight(O) is the weighting (importance) ascribed to objective O. The summation of Impact \times Weight over all objectives adds up the risk's total severity were it to occur.

Note that the information on a risk's impacts on objectives is explicit and separate from information on the objectives' weights. This gives an opportunity to different discipline experts to contribute their own knowledge. Typically the mission scientists will be the ones who ascribe relative weights to the various science objectives (e.g., the relative importance of measuring the water content in soil vs. the importance of measuring atmospheric composition). The inventors of a novel piece of technology may be the most knowledgeable in estimating the adverse effect of a particular type of failure. The discipline engineers may be the most knowledgeable in estimating the a-priori likelihood of a given fault (e.g., single event upsets caused by radiation).

DDP then goes on to *calculate* the risk remaining after taking into account the risk-reducing effects of mitigations, using formula (3):

$$\text{Mitigated-Risk}(R) = \text{Unmitigated-Risk}(R) \times \prod (M \in \text{Mitigations}) : (1 - \text{Effect}(M, R)) \quad (3)$$

where Effect (M, R) is the proportional risk reduction that Mitigation M has on Risk R. The product of (1 - Effect) over all mitigations is DDP's formula for calculating the combined effect of multiple mitigations against the same risk. The

intuition is that a mitigation acts like a "filter", reducing a risk by some proportion; multiple mitigations act like filters in series. For example, a mitigation with an effect of 0.8 against a risk reduces that risk by that proportion, leaving $(1 - 0.8) = 0.2$ of the risk remaining; a subsequent mitigation with an effect of 0.7 against that same risk reduces the 0.2 that's left of that risk further, leaving $0.2 \times (1 - 0.7) = 0.2 \times 0.3 = 0.06$ of the risk remaining. For a more detailed discussion of DDP's risk model the reader is referred to [5].

Note that the information on a mitigation's effect at reducing a risk is explicit and separate from information on the risk's impact. This provides another opportunity to different discipline experts to contribute their own knowledge. For example, the testing and quality assurance experts may have the best knowledge of how effective a given practice is at detecting problems which can then be repaired (e.g., formal inspections of software requirements may uncover defects which can quickly and easily be corrected while early in the software development process).

DDP's calculation and presentation of aggregate risk

Custom software has been developed to support all the steps of the DDP process. In particular, DDP supports pooling the fundamental risk data that the experts provided, and presenting the aggregate risk information back to those experts via cogent visualizations.

As described in the previous section, the sum total impact of a risk on objectives is *calculated* from the objectives' relative weights, how much the risk adversely impacts those objectives (taking into account the mitigations that reduce its impact), and how likely it is to occur (taking into account the mitigations that reduce its likelihood). This is done for each of the risks.

As example is shown in Fig. 3, where both the Unmitigated-Risk and the Mitigated-Risk values have been calculated for each of the 58 risks, and the results plotted in a bar chart. This is a screenshot taken from the DDP software operating on actual data taken from one of the DDP conducted risk studies. The data shows risks at an intermediate stage, when only some of the mitigations have yet been identified; by the conclusion of the study risks had been reduced much further than this chart shows.

DDP also calculates the sum total risk against each objective, which when presented in a similar bar chart allows experts to see the status on an objective-by-objective basis.

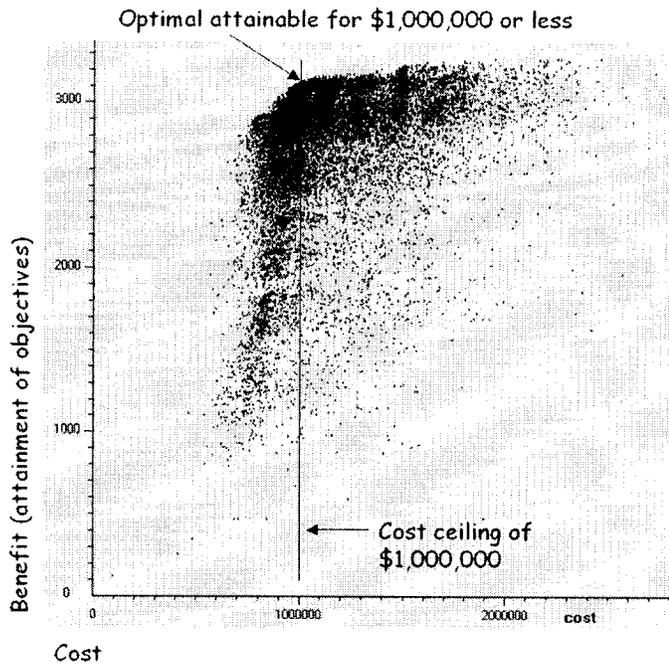


Figure 4. Search for an optimal cost-bounded solution

CONSTRAINED DEVELOPMENT

Time and budget pressures constrain the development of systems and technologies, and limited operational resources constrain systems themselves. The fundamental observation is that it takes an expenditure of resources to apply the mitigations that reduce risks, and hence the selection of those mitigations must be driven by consideration of both the benefits (risks reduced) and costs (resources consumed). Decision-making involving cost/benefit tradeoffs in this constrained setting is facilitated in DDP by:

- Capturing the cost information in the DDP data model
- Automatic calculation of total costs and benefits for given selections of mitigations
- Heuristic search to locate near-optimal cost/benefit solutions
- Calculation and presentation of the cost/benefit tradeoff space

The DDP data model includes a place to record resource cost information associated with mitigations (of course, this requires that the experts provide this information). For a given selection of mitigations, the DDP software computes the sum of their costs. When experts make decisions as to which mitigations to apply, they take into account both their costs, and their benefits (namely, the increase in attainment of objectives that results from the effect of mitigations at reducing risks). The DDP model also accommodates another important source of cost, namely the cost of *repair* of problems detected during development. For example, a fatal problem discovered at test time will necessitate some kind of repair, part replacement, rebuild etc. In the DDP model information on such costs is associated with the risks themselves – the cost of repairing the problem (risk) depends on what it is, not how it was discovered (i.e., which mitigation

detected it). Furthermore, DDP accommodates the escalation of repair costs that occurs when repairs are made later rather than sooner.

For non-trivial DDP applications, the search space of possible mitigation selections is huge. For example, in the application whose data is pictured in Figure 1, there are 36 individual mitigations, so the number of possible selections of such is 2^{36} (more than 10^{10}). In other DDP applications we have seen even greater numbers of mitigations, with correspondingly larger search spaces. This makes cost-effectively selecting mitigations a considerable challenge. In response, DDP uses the heuristic search technique of simulated annealing to locate near-optimal selections of mitigations. Figure 4 shows an example of DDP's search for a near-optimal solution in a study with 58 mitigations: the search has been directed to look for maximal attainment of objectives while costing no more than \$1 million. The black cloud on the figure is composed of thousands of points, each one corresponding to a distinct selection of mitigations. The location of each point on the horizontal axis is determined by the cost of that solution (as calculated by DDP), and location on the vertical axis by the benefit, i.e., attainment of objectives (again, as calculated by DDP). The cost ceiling of \$1 million is indicated by the vertical line, about a third of the way from the left. All the points on or to the left of the line cost less than or equal to that cost ceiling, so the optimal is the point highest up in that region, the vicinity of which is pointed to.

In addition to simulated annealing, we have also explored the use of other heuristic search techniques for finding near-optimal solutions. Further information on this work is reported in [6].

From an amalgamation of a series of such searches DDP can reveal the entire cost/benefit trade space, as shown in Figure 5. The upper left boundary of the widespread black cloud of points represents the optimal boundary (a.k.a. the

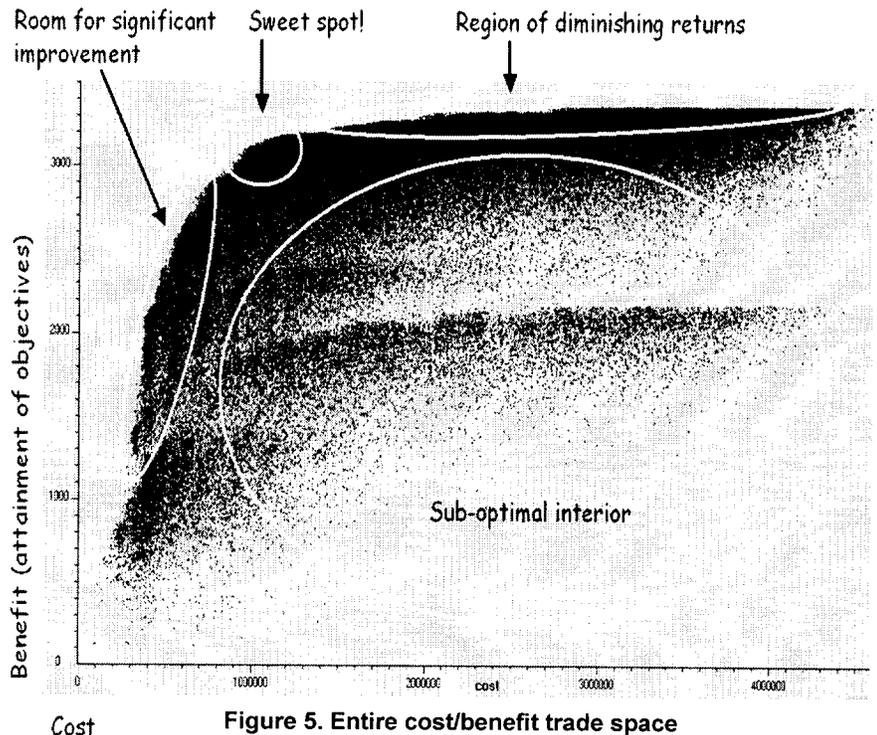


Figure 5. Entire cost/benefit trade space

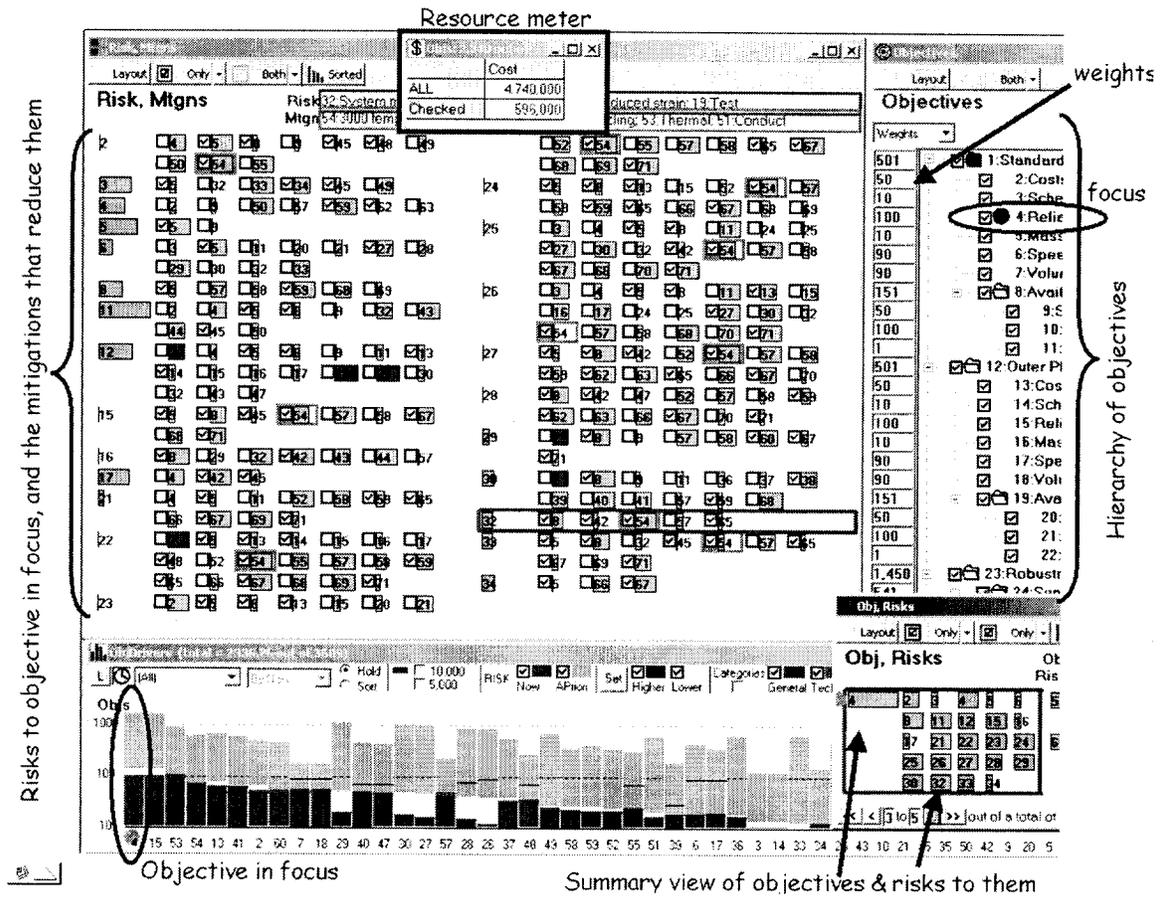


Figure 6. Screenshot of DDP in use to examine detailed risk information

“Pareto front”). Experts can look at this overall space to help decide the region of spending they would prefer. Depending on how risk averse they are, they may wish to extend partway to the right of what is here labeled as the “sweet spot” – the purpose of this kind of search and plot is to make the human decision makers aware of the nature of the overall cost/benefit trade space.

CRITICAL NATURE, UNKNOWN/UNPREDICTABLE ENVIRONMENT, AND LACK OF EXPERIENCE BASE

Spacecraft are critical systems that must operate correctly in only partially understood environments with no chance for repair. Furthermore, spacecraft often employ new technologies to yield better science return, etc. Past experience provides only a partial guide in these cases.

When knowledge and experience is lacking, no process can claim to yield perfect answers. The advantage that DDP conveys is that it encourages a more detailed decomposition of the problem, so that well understood aspects can be separated from the less understood aspects. Because of the critical systems nature of our studies, where aspects are poorly understood, we err on the side of caution. We use a pessimistic estimate of the impact of a risk on an objective, and the pessimistic estimate of the effectiveness of a mitigation at reducing a risk. We are often find it advantageous to decompose an item (e.g., a risk) into subcases, those that we do understand well, and can accurately score the impacts of, etc.,

and those that we understand less well, for which we adopt the pessimistic posture.

The overall purpose of gathering the information within DDP is to help guide decision-making. It may well turn out that even with pessimistic assumptions, a number of well-understood risks still emerge as more serious than certain less understood risks. For example, compare risks that relate to the challenging time of planetary landing phase with risks that relate to operation of a rover on the planetary surface. We might have a much greater understanding of the relevant aspects of former environment (e.g., atmospheric density) than of the latter (e.g., composition of the soil over which a rover will move). However, a failure during landing is almost certainly going to be fatal, while a failure to move up a slope on the surface may simply bring the rover to a halt. This may direct our decision-making as to where to appropriately expend our mitigation resources. When risks do show up as significant, the DDP tool gives the experts the ability to scrutinize those risks in detail, by looking at the objectives they impact, and the mitigations that could be applied to reduce those risks. If it turns out that risks cannot be satisfactorily reduced with the resources available, DDP can help experts understand how to “descope” their mission (i.e., select which objectives to abandon so that resources can be better directed to quell the risks impacting the remaining objectives, and not squandered on what turns out to be an especially challenging objective, say).

The key to all these is the simultaneous ability to view the “big picture” of risks, objectives’ attainment, cost of selected

mitigations, along with the ability to drill down in detail into the data that underlies these calculations. DDP offers a number of cogent visualizations that support such investigations. To convey a feel for DDP's capabilities in this regard, Figure 6 shows (most of) a screenshot of DDP in use to examine detailed risk information. Within this screen are the following elements:

- The bar chart in the lower left pane is showing the status of risk against each of the objectives (akin to the bar chart of risks seen earlier in Fig. 3, but here each bar corresponds to an objective, and the heights correspond to various measures of risk to that objective).
- The hierarchy of objectives is shown in the upper right pane (in this image we have truncated the right of the screenshot to hide the actual names of objectives, since this is data of a proprietary nature). One of the objectives has been brought "in focus", meaning the tool highlights it, and displays information particular to that objective. Alongside each objective is its user-assigned weight.
- The fragment of the pane in the lower right shows a summary view of the risks on that objective, where each risk is portrayed as a small rectangle whose width is proportional to the risk's impact on that objective.
- The large pane in the upper left, filled with a multitude of rectangles and checkboxes, shows in greater detail the risks that impact the in focus objective. Alongside each risk are tiny checkboxes coupled with rectangles. Each represents a mitigation which (if applied) reduces the risk it is listed alongside. The checkbox is checked if and only if the mitigation is currently selected for application. The user can click to toggle this checked status, causing DDP to recalculate risks, costs, etc., in response.
- The resource meter is shown floating over the top middle of the screen. This keeps a running total of the resource costs of the checked (i.e., selected for application) mitigations.

At first sight this appears a dauntingly complex display, but in practice it serves to convey the risk-related information to the experts in such a way as to aid them in their decision making. The responsiveness of the DDP software helps in this regard – when the user clicks a checkbox to toggle the status of a mitigation, complete recalculation and display update is fast. For a typical DDP dataset with a hundred or so in total of objectives, risks and mitigations, and on the order of a thousand links among them, this takes on the order of one second, operating on a modern-day PC laptop. Thus it is possible for experts to quickly try "what if" experiments – e.g., "what would be the change to cost and risk if we turned off that mitigation and turned on that other one?" The bar chart displays can help in these "what if" explorations, by visual presentation of the *changes* to risks from some baseline. In fact this is just barely discernable on the bar chart in Fig. 6: some of the bars have three shades of gray, the lightest indicating how much the risk against that objective has dropped compared to a previously set baseline.

SUMMARY AND CONCLUSIONS

DDP has now been applied to make risk-informed decisions in over 20 different studies of advanced technologies intended for spacecraft use. The nature of the decisions has varied. In some cases the primary outcome has been selection of a suite of mitigations that in concert will cost-effectively

reduce risk to sufficiently low levels. The ability to search and reveal the cost/benefit tradeoff space has been of assistance here. In other cases the primary outcome has been selection from among major design alternatives, using DDP to study the risks of each and indicate the cost of mitigating the risks in each option. In almost all cases the process leads to clarification of requirements. In one of the studies, a particularly problematic requirement was identified. It was found to be problematic because the cumulative risk information contributed by the engineering experts showed the requirement to be significantly at risk, and furthermore showed that sufficient mitigation of those risks would be very expensive in terms of schedule and budget. The net result was that the mission scientists were motivated to rethink their objectives, and make the problematic one a much lower priority.

We have not performed any formal experiments to measure the overall effectiveness of DDP applications. We do know the cost – predominantly the time of the experts involved in populating the study with data and making decisions. For example, a typical DDP application, with 15 participants in all four of the 4-hour sessions, consumes 240 hours (6 work weeks) of time. We require the participation of experts. Such people are always in demand, so this is by no means a trivial investment of time. Estimating the benefit of DDP applications is much more subjective. We are encouraged as to the validity and utility of the findings that DDP yields for the following reasons:

1. Most of the results calculated by DDP, based on the information gathered from the multiple discipline experts, match those experts' overall intuitions. For example, DDP's list of most significant risks that remain despite application of mitigations is usually in agreement with what the experts tell us they would have expected. Overall this suggests that the detailed information we are gathering from those experts, and the ways we combine that information in DDP's risk calculations, is reasonably valid.
2. In almost every case study there is some result calculated by DDP that is a "surprise". That is, it does not match the experts' intuitions (e.g., a risk shows up as more significant than they would have anticipated). Furthermore, when the experts look at the detail underpinning that result (and the capability of DDP to let them explore the details as well as see the big picture is crucial in this regard), they concede that the "surprise" is a genuine finding. Overall this suggests that the approach we follow is capable of findings that would be overlooked by even a well-qualified set of experts.
3. The decisions arrived at with the aid of DDP have led to much clarified in problem statements (objectives), better substantiated cost estimates, well-organized rationales for development plans, and, in some cases, very significant cost savings by stimulating design revisions motivated by the realization that resources were being mis-applied to (over)address a relatively tiny risk area, while other risk areas were in more dire need of attention. Perhaps the most important aspect shared by these decisions is that they are being made relatively early in the technology or system lifecycle, when the costs of revising choices are still relatively small. In contrast, fundamental design changes later in development incur much higher costs.

4. DDP is not a mandated activity, yet is seeing increasing use. The earlier DDP applications were supported by research funding (as was the development of the DDP process and support software itself). Specifically, the discipline experts' time was covered by this funding. More recently, there has been a trend to the use of DDP paid for predominantly by the technology/system under scrutiny. While we continue to use research funds to expand DDP's capabilities, the bulk of the cost of performing the risk studies themselves, namely the cost of the discipline experts' time, is no longer being borne by the research funding. This, plus the increasing frequency of DDP based studies, is indicative of an expanding acceptance of the net value of the DDP process.

Related work

Early decision-making is often assisted by *qualitative* decision support methods. For example Quality Function Deployment (QFD) has been used in a wide variety of settings [7]. DDP's effect and impact matrices are reminiscent of the Relationship Matrix used in many forms in QFD. DDP is distinguished by its foundation upon a *quantitative* risk model, which gives meaning to DDP's cost and benefit calculations.

As a design matures there are other decision support techniques that better capitalize upon knowledge of design details. For example, probabilistic risk assessment techniques (e.g., fault tree analysis, Bayesian methods) compute overall system reliability from design knowledge of how the system is composed of those components, and estimates of individual component reliabilities. The origins of these approaches lie in applications to assess risk in the nuclear power industry [8], with its need to estimate the probability of catastrophic failure (e.g., meltdown) from knowledge of the power system's design, and reliability measures for the components used in that design. Fault Tree Analysis [9] is now applied to a wide variety of systems, both hardware and software (e.g., software fault tree analysis [10]) including some NASA missions and their hardware and software components [11]. In contrast, DDP aims to fill the niche of early decision making for advanced technology and system development. We are currently exploring means to connect DDP and PRA techniques [12].

ACKNOWLEDGMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its

endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

- [1] O'Keefe, S. "Pioneering the Future", address by the NASA Administrator, Maxwell School of Citizenship & Public Affairs, Syracuse University, April 12, 2002 ftp://ftp.hq.nasa.gov/pub/pao/okeefe/2002/pioneering_the_future.pdf
- [2] Lehman, M.M. "Some Characteristics of S-type and E-type Software" in *Proceedings of the Second FEAST Workshop*, Dept. of Computing, Imperial College of Science Technology and Medicine, London, UK.
- [3] Cornford, S.L. "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.
- [4] S.L. Cornford, M.S. Feather & K.A. Hicks. "DDP – A tool for life-cycle risk management", *IEEE Aerospace Conference*, Big Sky, Montana, Mar 2001, pp. 441-451. <http://www.doc.ic.ac.uk/~mml/feast2/papers/pdf/2feastwks.pdf>
- [5] Feather, M.S. & Cornford, S.L.. "Quantitative risk-based requirements reasoning", to appear in *Requirements Engineering* (Springer), published online 25 February 2003, DOI 10.1007/s00766-002-0160-y. Available from: <http://eis.jpl.nasa.gov/~mfeather/AvailablePublications/>
- [6] Cornford, S.L., Feather, M.S., Dunphy, J.R., Salcedo, J. & Menzies, M. "Optimizing Spacecraft Design – Optimization Engine Development: Progress and Plans", *Proceedings, 2003 IEEE Aerospace Conference*, Big Sky, Montana, March 2003.
- [7]. Akao, Y. 1990 "*Quality Function Deployment*", Productivity Press, Cambridge, Massachusetts.
- [8] *Reactor Safety Study*, Report WASH-1400, Nuclear Regulatory Commission, 1975.
- [9] Vesely, W.E., Goldberg, F.F., Roberts, N.H. & Haasl, D.F., "*Fault Tree Handbook*", U.S. Nuclear Regulatory Commission NUREG-0492, 1981.
- [10] N.G. Leveson. "*Safeware: system safety and computers*". Addison Wesley, Reading, MA, 1995.
- [11] *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.
- [12] S.L. Cornford, T. Paulos, L. Meshkat & M.S. Feather. "Towards More Accurate Life Cycle Risk Management Through Integration of DDP and PRA". *IEEE Aerospace Conference*, Big Sky MT, Mar 2003.

IMECE2003-42749

RISK-BASED ANALYSIS AND DECISION MAKING IN MULTI-DISCIPLINARY ENVIRONMENTS

Martin S. Feather

Steven L. Cornford

Kelly Moran

Jet Propulsion Laboratory, California Institute of Technology

ABSTRACT

A risk-based decision-making process conceived of and developed at JPL and NASA, has been used to help plan and guide novel technology applications for use on spacecraft. These applications exemplify key challenges inherent in multi-disciplinary design of novel technologies deployed in mission-critical settings:

- 1) Cross-disciplinary concerns are numerous (e.g., spacecraft involve navigation, propulsion, telecommunications). These concerns are cross-coupled and interact in multiple ways (e.g., electromagnetic interference, heat transfer).
- 2) Time and budget pressures constrain development, operational resources constrain the resulting system (e.g., mass, volume, power).
- 3) Spacecraft are critical systems that must operate correctly the first time in only partially understood environments, with no chance for repair.
- 4) Past experience provides only a partial guide: New mission concepts are enhanced and enabled by new technologies, for which past experience is lacking.

The decision-making process rests on quantitative assessments of the relationships between three classes of information – objectives (the things the system is to accomplish and constraints on its operation and development), risks (whose occurrence detracts from objectives), and mitigations (options for reducing the likelihood and/or severity of risks). The process successfully guides experts to pool their knowledge, using custom-built software to support information gathering and decision-making.

Keywords: Risk, Decision-making, Design, Novel Technology

INTRODUCTION

NASA's Mission statement reads: *"To understand and protect our home planet. To explore the Universe and search for life. To inspire the next generation of explorers . . . as only*

NASA can." In April 2002 the NASA Administrator stated: *"... In broad terms, our mandate is to pioneer the future . . . to push the envelope . . . to do what has never been done before. ..."* [1]. These quotes emphasize the novel aspects of NASA's activities. Spacecraft design involves multiple disciplines (e.g., interplanetary navigation, spacecraft propulsion, telecommunications, properties of materials in deep space environments). Spacecraft development and operation are driven by budget and time concerns (e.g., planetary configurations dictate launch windows). Spacecraft operate remotely from earth; so third-party repair is almost always impossible. Their purpose is to yield information of distant, unknown environments, meaning they often must operate in those same unknown environments. The aim to improve the quantity and quality of the science information they gather and in return drives the use of novel technologies in novel ways.

At first glance, these would seem to imply that spacecraft development and operation has little in common with earthly activities. Certainly the specifics of the information involved (e.g., behavior of materials in a deep space environment; interplanetary navigation) are particular to space exploration. However, the fundamental challenges that underlie NASA's activities are familiar to a wide range of endeavors – the challenges stem from the need to do multi-disciplinary design of novel technologies intended for deployment in critical settings. They are:

- Multiple disciplines are involved. No single individual has depth of knowledge in every one of these disciplines. Furthermore, these disciplines are cross-coupled and interact in multiple ways. As a consequence, there is the need to pool the knowledge from multiple discipline experts, and to conduct decision-making taking the sum total of this knowledge into account.
- Time and budget pressures constrain development (an almost universal phenomenon, e.g., development of almost

any product these days is subject to time to market needs, and aspects of competitiveness). Limited operational resources constrain many systems (e.g., think of power, mass and volume constraints on hand-held devices).

- Depending upon the purposes for which systems are employed, the consequences of failure during their operation could range from a minor nuisance to catastrophic. Spacecraft represent an unusual extreme in which the entire system itself will be lost if things go badly wrong. On earth, the systems themselves may be more repairable, but the not necessarily the consequences of their operational failure.
- Spacecraft must operate in a harsh and often unpredictable physical environment. While more may be known about the physical operating environment for Earth-bound systems, considerable uncertainty arises from the unpredictability of the ways in which users will make use of systems. In the extreme, a deployed system will induce a change in the environment around it, leading to an unpredictable feedback loop (e.g., the phenomenon of “E-type Software” described by Lehman [2]).
- Past experience provides only a partial guide as to how to develop the new system. For NASA’s missions, new mission concepts are enhanced and enabled by use of new technologies, and/or of established technologies applied in novel ways. Past experience is lacking in either case. The same is true for almost any novel solution to a real world problem.

In response to these challenges, at JPL and NASA we have been developing and applying a risk-based approach to analysis and decision making for novel system applications. This paper summarizes the approach, and indicates how it addresses the challenges listed above.

NOMENCLATURE

Defect Detection and Prevention (DDP) – a risk-based decision-making process conceived of and developed at JPL and NASA.

Objectives – the things the system/technology is to accomplish and constraints on its operation and development.

Risks – the things whose occurrence would detract from attainment of Objectives.

Mitigations – options for reducing the likelihood and/or severity of Risks.

A RISK-BASED ANALYSIS AND DECISION-MAKING APPROACH

The focus of this paper is on a risk-based analysis and decision-making approach developed and applied at JPL and NASA. The approach is called “Defect Detection and Prevention (DDP)”. The name reflects its origins as a structured method for planning the quality assurance of hardware systems [3]. Since then its scope has expanded to also encompass decision-making earlier in the development lifecycle, and to be applicable to software, hardware and systems [4].

The approach has been generalized to aid in the decision-making during the early phases of advanced technology and system development. Decisions made in these early phases are important because they have the most leverage to influence the development to follow. However, decision-making during these phases is challenging because information on which to base those decisions is incomplete and uncertain, and in the case of

advanced technologies and systems, there is little past experience from which to extrapolate.

DDP’s Simple yet Quantitative and Detailed Risk Model

DDP offers a conceptually simple yet quantitative, detailed model of risk as foundation for reasoning:

- The simple conceptual core of the DDP risk model rests on three sets of information: *objectives* – the things we want the system/technology to achieve, *risks* – the things that, should they occur, detract from the attainment of objectives, and *mitigations* – the things that we could choose to do to reduce risks (by reducing their likelihood of occurrence and/or their severity should they occur).
- The quantitative treatment derives from information on how much each risk (should it occur) detracts from attainment of each objective, and by how much each mitigation (should it be applied) reduces each Risk.
- The detail arises from the ability to populate the DDP model with system/technology specific objectives, risks and mitigations. These are open-ended sets of information, and the level of detail to which these are populated is not pre-ordained, but rather is determined in the course of the DDP application process, to the extent necessary to support decision-making.

The combination of these aspects yields an approach that can encompass a wide range of concerns – many problems can be cast in terms of objectives (what do we want), Risks (what can get in the way) and mitigations (what can we do about it). Their quantitative treatment allows reasoning about their net combined effect. For example, the biggest risks are those with the greatest sum total adverse impact on attainment of objectives. The ability to incorporate problem-specific detail permits the representation and reasoning over system, technology and application-specific nuances.

The later sections will examine the ways in which the DDP information model helps address the key challenges of multi-disciplinary decision making. In preparation, we now expand upon the description of the DDP model.

Details of DDP’s Risk Model

In more detail, DDP deals with the following three sets of information:

- “**Objectives**” – these encompass all of the things the system/technology is to achieve, how and under what conditions it is to be operate, and how it is to be developed. In our area of spacecraft technologies, typical examples include science return objectives (e.g., quantity and resolution of data), aspects of its operational environment (e.g., ambient temperature, operating temperature, available electrical power), and nature of its development (e.g., must be ready for launch 2 years from now). Since not all objectives are equally important, DDP allows users to assigned “Weights” (numerical values of a unitless scale) to reflect their relative importance. For example, minimal success criteria (“must-have” capabilities) can be assigned high weights, while optional criteria (“nice to have” capabilities) can be assigned lesser weights.
- “**Risks**” – broadly speaking, all the problems that, should they occur, will adversely impact attainment of objectives. For hardware, these cover the gamut of failure modes of

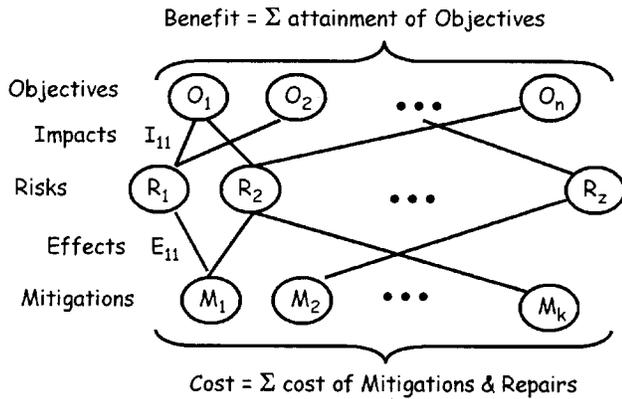


Figure 1. Topology of DDP information model

design, assembly (won't fit, too heavy, etc.) and operation (cracks, breaks, short-circuits etc.). For software, these include development risks (ambiguous requirements, flawed designs, poor contract arrangements, lack of configuration management) and risks in the developed software (erroneous results, deadlocks, timing errors, poor performance, mis-use of computing resources). Human and organizational concerns can also be incorporated here (e.g., key personnel unavailable; vendor goes out of business).

- **“Mitigations”** – the entire range of actions and options to reduce risks, including preventative measures (e.g., training of personnel, use of design standards), alleviations (e.g., hardware redundancy; software bounds checking that make a software component more resilient to erroneous inputs), and detections (e.g., design reviews, code inspections, analyses, tests of all kinds) that can potentially detect problems ahead of their manifestation in use, and so allow for their correction. There are almost always resource costs associated with mitigations (e.g., budget and schedule; for spacecraft hardware, power, volume and mass are perennial concerns; for spacecraft software, CPU and memory utilization), and one of the primary purposes of DDP is to help guide the selection of mitigations that reduce risk in a cost-effective manner.

In DDP, information in these three categories is *quantitatively* linked. Risks are quantitatively related to objectives to indicate *how much* a risk, should it occur, adversely impact an objective's attainment. Mitigations are quantitatively related to risks to indicate *how much* a mitigation, should it be applied, will reduce a risk (by preventing it from occurring in the first place, by detecting it and so allowing for its repair, or by reducing its impact on objectives).

The overall structure of a DDP information model is relatively simple – see Fig. 1. However in practice there can be numerous items and links among them. For example, the topology of the data in a recently completed DDP study (of a technology intended for spacecraft application) is shown in Fig. 2. This comprises 29 objectives, 58 risks, 36 mitigations, and some 900 links among them.

MULTI-DISCIPLINARY CONCERNS

The single most important aspect of the DDP approach is that it supports multiple experts pool their knowledge and allows them to take the sum total of their pooled knowledge

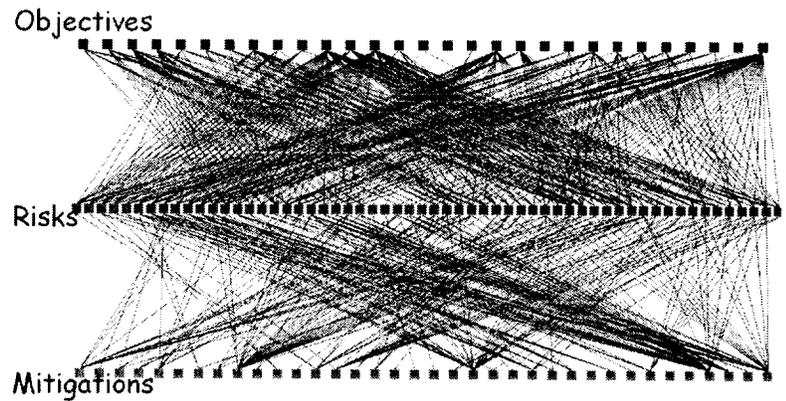


Figure 2. Topology of data in a completed DDP model

into account as they make decisions. This helps address two key concerns of multi-disciplinary studies, namely:

- No single individual has depth of knowledge in every one of the disciplines, and
- Disciplines are cross-coupled and interact in multiple ways.

DDP enables pooling of multiple experts' knowledge by utilizing:

- a structured group process to elicit in-depth information from experts,
- software support to perform calculations and search over the entirety of the gathered information (e.g., calculate aggregate risk information, search for near-optimal solutions), and
- cogent visualizations to present the information back to the experts and so retain their engagement in a risk-informed decision-making process.

The DDP process is conducted in a group setting in which all the experts relevant to the technology/system under consideration are simultaneously involved. It is crucial the experts' combined areas of expertise span the concerns of the study, and that those experts are able and willing to contribute their expertise. In this respect DDP is akin to traditional risk assessment methods, which gather their risk information from representatives of all relevant disciplines. However, DDP differs by seeking more fundamental kinds of information than do typical risk assessment methods. This provides opportunities for experts to contribute specific and detailed knowledge. The DDP software tool performs the calculations to pool their separately contributed items, and present that information back to the experts.

DDP's use of fundamental risk information

Most risk assessment methods take as starting point a technology/system *incorporating* all the risk reducing measures to be applied to make that technology/system reliable. They focus on the risks *remaining* in that technology/system. In particular, these methods ask experts to provide, for each risk, estimates of the remaining likelihood and impact of that risk. The product of these is the usual definition of risk (sometimes referred to as “risk exposure”), formula (1):

$$\text{risk} = \text{likelihood} \times \text{severity (a.k.a. consequence)} \quad (1)$$

In contrast, DDP takes as starting point a technology/system *separate* from the risk reducing measures expected to be applied. Those are represented separately,

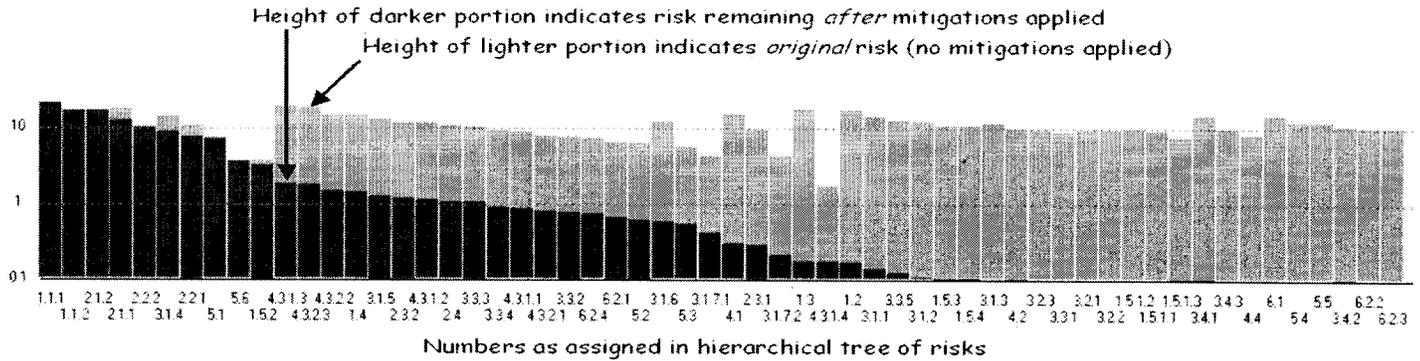


Figure 3. Graphical presentation of risks sorted into descending order of remaining risk levels

explicitly accounting for their risk reducing effects. The risks remaining in the design are *calculated* from the a-priori risks, and the risk reductions effects of the mitigations. In detail, DDP asks experts to provide, for each risk, estimates of:

- what would be its likelihood of occurrence were nothing done to prevent it, and
- what adverse impacts it would have on objectives should it occur, and
- what likelihood- and severity- reducing effects each of the mitigations would have on that risk were they to be applied.

The first two are combined to calculate the *unmitigated* contribution of a risk, using the formula (2):

$$\begin{aligned} \text{Unmitigated-Risk}(R) = \\ \text{A-Priori-Likelihood}(R) \times \\ \sum (O \in \text{Objectives}) : \text{Impact}(R, O) \times \text{Weight}(O) \end{aligned} \quad (2)$$

where A-Priori-Likelihood(R) is risk R's likelihood were nothing done to prevent it, Impact(R, O) is the proportion of the objective O that would be lost were R to occur, and Weight(O) is the weighting (importance) ascribed to objective O. The summation of Impact \times Weight over all objectives adds up the risk's total severity were it to occur.

Note that the information on a risk's impacts on objectives is explicit and separate from information on the objectives' weights. This gives an opportunity to different discipline experts to contribute their own knowledge. Typically the mission scientists will be the ones who ascribe relative weights to the various science objectives (e.g., the relative importance of measuring the water content in soil vs. the importance of measuring atmospheric composition). The inventors of a novel piece of technology may be the most knowledgeable in estimating the adverse effect of a particular type of failure. The discipline engineers may be the most knowledgeable in estimating the a-priori likelihood of a given fault (e.g., single event upsets caused by radiation).

DDP then goes on to *calculate* the risk remaining after taking into account the risk-reducing effects of mitigations, using formula (3):

$$\begin{aligned} \text{Mitigated-Risk}(R) = \\ \text{Unmitigated-Risk}(R) \times \\ \prod (M \in \text{Mitigations}) : (1 - \text{Effect}(M, R)) \end{aligned} \quad (3)$$

where Effect (M, R) is the proportional risk reduction that Mitigation M has on Risk R. The product of (1 - Effect) over all mitigations is DDP's formula for calculating the combined effect of multiple mitigations against the same risk. The

intuition is that a mitigation acts like a "filter", reducing a risk by some proportion; multiple mitigations act like filters in series. For example, a mitigation with an effect of 0.8 against a risk reduces that risk by that proportion, leaving $(1 - 0.8) = 0.2$ of the risk remaining; a subsequent mitigation with an effect of 0.7 against that same risk reduces the 0.2 that's left of that risk further, leaving $0.2 \times (1 - 0.7) = 0.2 \times 0.3 = 0.06$ of the risk remaining. For a more detailed discussion of DDP's risk model the reader is referred to [5].

Note that the information on a mitigation's effect at reducing a risk is explicit and separate from information on the risk's impact. This provides another opportunity to different discipline experts to contribute their own knowledge. For example, the testing and quality assurance experts may have the best knowledge of how effective a given practice is at detecting problems which can then be repaired (e.g., formal inspections of software requirements may uncover defects which can quickly and easily be corrected while early in the software development process).

DDP's calculation and presentation of aggregate risk

Custom software has been developed to support all the steps of the DDP process. In particular, DDP supports pooling the fundamental risk data that the experts provided, and presenting the aggregate risk information back to those experts via cogent visualizations.

As described in the previous section, the sum total impact of a risk on objectives is *calculated* from the objectives' relative weights, how much the risk adversely impacts those objectives (taking into account the mitigations that reduce its impact), and how likely it is to occur (taking into account the mitigations that reduce its likelihood). This is done for each of the risks.

As example is shown in Fig. 3, where both the Unmitigated-Risk and the Mitigated-Risk values have been calculated for each of the 58 risks, and the results plotted in a bar chart. This is a screenshot taken from the DDP software operating on actual data taken from one of the DDP conducted risk studies. The data shows risks at an intermediate stage, when only some of the mitigations have yet been identified; by the conclusion of the study risks had been reduced much further than this chart shows.

DDP also calculates the sum total risk against each objective, which when presented in a similar bar chart allows experts to see the status on an objective-by-objective basis.

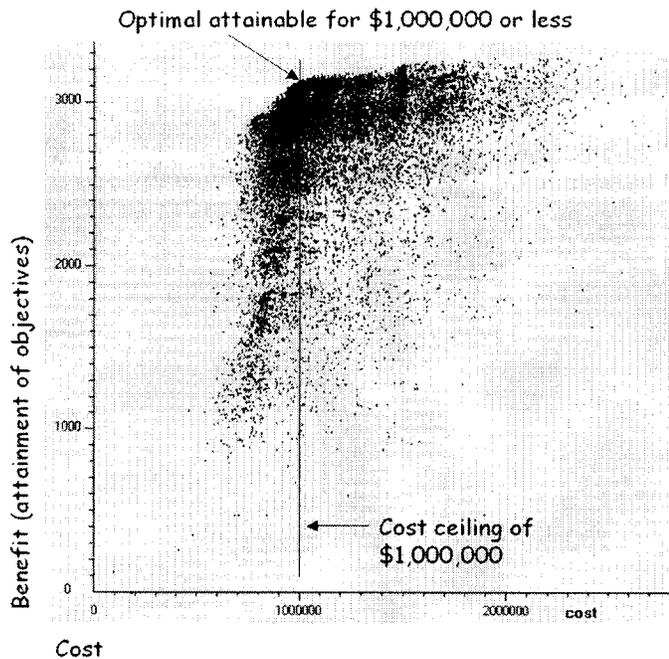


Figure 4. Search for an optimal cost-bounded solution

CONSTRAINED DEVELOPMENT

Time and budget pressures constrain the development of systems and technologies, and limited operational resources constrain systems themselves. The fundamental observation is that it takes an expenditure of resources to apply the mitigations that reduce risks, and hence the selection of those mitigations must be driven by consideration of both the benefits (risks reduced) and costs (resources consumed). Decision-making involving cost/benefit tradeoffs in this constrained setting is facilitated in DDP by:

- Capturing the cost information in the DDP data model
- Automatic calculation of total costs and benefits for given selections of mitigations
- Heuristic search to locate near-optimal cost/benefit solutions
- Calculation and presentation of the cost/benefit tradeoff space

The DDP data model includes a place to record resource cost information associated with mitigations (of course, this requires that the experts provide this information). For a given selection of mitigations, the DDP software computes the sum of their costs. When experts make decisions as to which mitigations to apply, they take into account both their costs, and their benefits (namely, the increase in attainment of objectives that results from the effect of mitigations at reducing risks). The DDP model also accommodates another important source of cost, namely the cost of *repair* of problems detected during development. For example, a fatal problem discovered at test time will necessitate some kind of repair, part replacement, rebuild etc. In the DDP model information on such costs is associated with the risks themselves – the cost of repairing the problem (risk) depends on what it is, not how it was discovered (i.e., which mitigation

detected it). Furthermore, DDP accommodates the escalation of repair costs that occurs when repairs are made later rather than sooner.

For non-trivial DDP applications, the search space of possible mitigation selections is huge. For example, in the application whose data is pictured in Figure 1, there are 36 individual mitigations, so the number of possible selections of such is 2^{36} (more than 10^{10}). In other DDP applications we have seen even greater numbers of mitigations, with correspondingly larger search spaces. This makes cost-effectively selecting mitigations a considerable challenge. In response, DDP uses the heuristic search technique of simulated annealing to locate near-optimal selections of mitigations. Figure 4 shows an example of DDP's search for a near-optimal solution in a study with 58 mitigations: the search has been directed to look for maximal attainment of objectives while costing no more than \$1 million. The black cloud on the figure is composed of thousands of points, each one corresponding to a distinct selection of mitigations. The location of each point on the horizontal axis is determined by the cost of that solution (as calculated by DDP), and location on the vertical axis by the benefit, i.e., attainment of objectives (again, as calculated by DDP). The cost ceiling of \$1 million is indicated by the vertical line, about a third of the way from the left. All the points on or to the left of the line cost less than or equal to that cost ceiling, so the optimal is the point highest up in that region, the vicinity of which is pointed to.

In addition to simulated annealing, we have also explored the use of other heuristic search techniques for finding near-optimal solutions. Further information on this work is reported in [6].

From an amalgamation of a series of such searches DDP can reveal the entire cost/benefit trade space, as shown in Figure 5. The upper left boundary of the widespread black cloud of points represents the optimal boundary (a.k.a. the

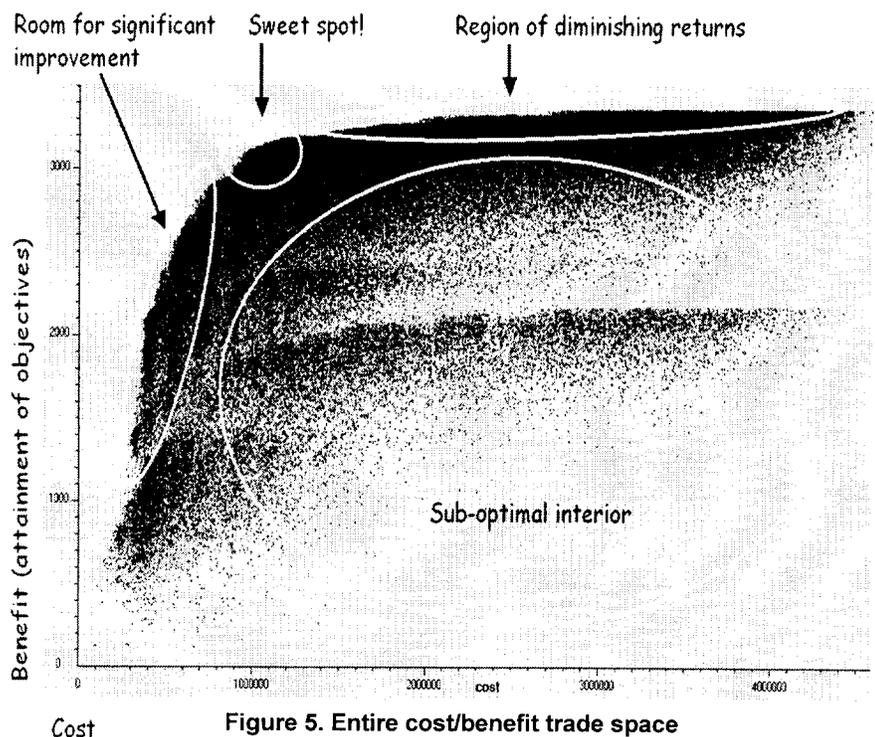


Figure 5. Entire cost/benefit trade space

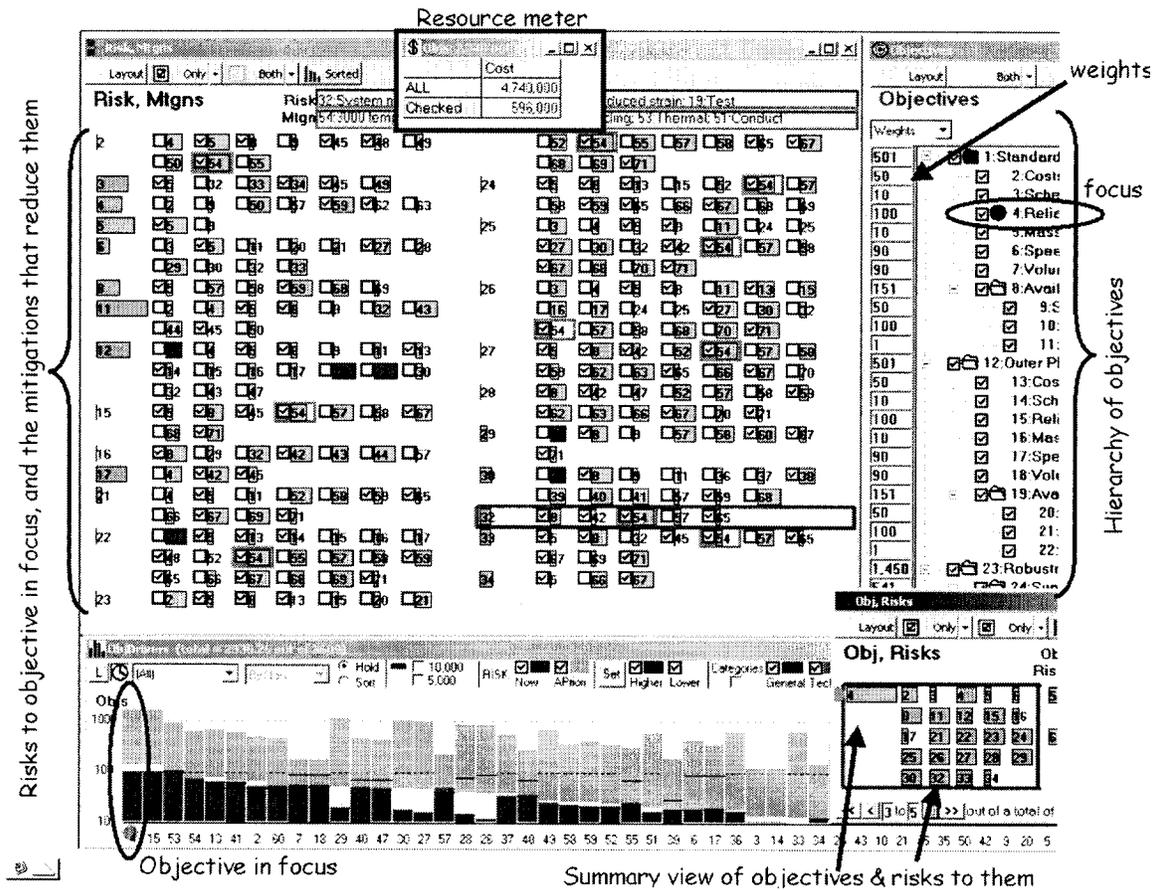


Figure 6. Screenshot of DDP in use to examine detailed risk information

“Pareto front”). Experts can look at this overall space to help decide the region of spending they would prefer. Depending on how risk averse they are, they may wish to extend partway to the right of what is here labeled as the “sweet spot” – the purpose of this kind of search and plot is to make the human decision makers aware of the nature of the overall cost/benefit trade space.

CRITICAL NATURE, UNKNOWN/UNPREDICTABLE ENVIRONMENT, AND LACK OF EXPERIENCE BASE

Spacecraft are critical systems that must operate correctly in only partially understood environments with no chance for repair. Furthermore, spacecraft often employ new technologies to yield better science return, etc. Past experience provides only a partial guide in these cases.

When knowledge and experience is lacking, no process can claim to yield perfect answers. The advantage that DDP conveys is that it encourages a more detailed decomposition of the problem, so that well understood aspects can be separated from the less understood aspects. Because of the critical systems nature of our studies, where aspects are poorly understood, we err on the side of caution. We use a pessimistic estimate of the impact of a risk on an objective, and the pessimistic estimate of the effectiveness of a mitigation at reducing a risk. We are often find it advantageous to decompose an item (e.g., a risk) into subcases, those that we do understand well, and can accurately score the impacts of, etc.,

and those that we understand less well, for which we adopt the pessimistic posture.

The overall purpose of gathering the information within DDP is to help guide decision-making. It may well turn out that even with pessimistic assumptions, a number of well-understood risks still emerge as more serious than certain less understood risks. For example, compare risks that relate to the challenging time of planetary landing phase with risks that relate to operation of a rover on the planetary surface. We might have a much greater understanding of the relevant aspects of former environment (e.g., atmospheric density) than of the latter (e.g., composition of the soil over which a rover will move). However, a failure during landing is almost certainly going to be fatal, while a failure to move up a slope on the surface may simply bring the rover to a halt. This may direct our decision-making as to where to appropriately expend our mitigation resources. When risks do show up as significant, the DDP tool gives the experts the ability to scrutinize those risks in detail, by looking at the objectives they impact, and the mitigations that could be applied to reduce those risks. If it turns out that risks cannot be satisfactorily reduced with the resources available, DDP can help experts understand how to “descope” their mission (i.e., select which objectives to abandon so that resources can be better directed to quell the risks impacting the remaining objectives, and not squandered n what turns out to be an especially challenging objective, say).

The key to all these is the simultaneous ability to view the “big picture” of risks, objectives’ attainment, cost of selected

mitigations, along with the ability to drill down in detail into the data that underlies these calculations. DDP offers a number of cogent visualizations that support such investigations. To convey a feel for DDP's capabilities in this regard, Figure 6 shows (most of) a screenshot of DDP in use to examine detailed risk information. Within this screen are the following elements:

- The bar chart in the lower left pane is showing the status of risk against each of the objectives (akin to the bar chart of risks seen earlier in Fig. 3, but here each bar corresponds to an objective, and the heights correspond to various measures of risk to that objective).
- The hierarchy of objectives is shown in the upper right pane (in this image we have truncated the right of the screenshot to hide the actual names of objectives, since this is data of a proprietary nature). One of the objectives has been brought "in focus", meaning the tool highlights it, and displays information particular to that objective. Alongside each objective is its user-assigned weight.
- The fragment of the pane in the lower right shows a summary view of the risks on that objective, where each risk is portrayed as a small rectangle whose width is proportional to the risk's impact on that objective.
- The large pane in the upper left, filled with a multitude of rectangles and checkboxes, shows in greater detail the risks that impact the in focus objective. Alongside each risk are tiny checkboxes coupled with rectangles. Each represents a mitigation which (if applied) reduces the risk it is listed alongside. The checkbox is checked if and only if the mitigation is currently selected for application. The user can click to toggle this checked status, causing DDP to recalculate risks, costs, etc., in response.
- The resource meter is shown floating over the top middle of the screen. This keeps a running total of the resource costs of the checked (i.e., selected for application) mitigations.

At first sight this appears a dauntingly complex display, but in practice it serves to convey the risk-related information to the experts in such a way as to aid them in their decision making. The responsiveness of the DDP software helps in this regard – when the user clicks a checkbox to toggle the status of a mitigation, complete recalculation and display update is fast. For a typical DDP dataset with a hundred or so in total of objectives, risks and mitigations, and on the order of a thousand links among them, this takes on the order of one second, operating on a modern-day PC laptop. Thus it is possible for experts to quickly try "what if" experiments – e.g., "what would be the change to cost and risk if we turned off that mitigation and turned on that other one?" The bar chart displays can help in these "what if" explorations, by visual presentation of the *changes* to risks from some baseline. In fact this is just barely discernable on the bar chart in Fig. 6: some of the bars have three shades of gray, the lightest indicating how much the risk against that objective has dropped compared to a previously set baseline.

SUMMARY AND CONCLUSIONS

DDP has now been applied to make risk-informed decisions in over 20 different studies of advanced technologies intended for spacecraft use. The nature of the decisions has varied. In some cases the primary outcome has been selection of a suite of mitigations that in concert will cost-effectively

reduce risk to sufficiently low levels. The ability to search and reveal the cost/benefit tradeoff space has been of assistance here. In other cases the primary outcome has been selection from among major design alternatives, using DDP to study the risks of each and indicate the cost of mitigating the risks in each option. In almost all cases the process leads to clarification of requirements. In one of the studies, a particularly problematic requirement was identified. It was found to be problematic because the cumulative risk information contributed by the engineering experts showed the requirement to be significantly at risk, and furthermore showed that sufficient mitigation of those risks would be very expensive in terms of schedule and budget. The net result was that the mission scientists were motivated to rethink their objectives, and make the problematic one a much lower priority.

We have not performed any formal experiments to measure the overall effectiveness of DDP applications. We do know the cost – predominantly the time of the experts involved in populating the study with data and making decisions. For example, a typical DDP application, with 15 participants in all four of the 4-hour sessions, consumes 240 hours (6 work weeks) of time. We require the participation of experts. Such people are always in demand, so this is by no means a trivial investment of time. Estimating the benefit of DDP applications is much more subjective. We are encouraged as to the validity and utility of the findings that DDP yields for the following reasons:

1. Most of the results calculated by DDP, based on the information gathered from the multiple discipline experts, match those experts' overall intuitions. For example, DDP's list of most significant risks that remain despite application of mitigations is usually in agreement with what the experts tell us they would have expected. Overall this suggests that the detailed information we are gathering from those experts, and the ways we combine that information in DDP's risk calculations, is reasonably valid.
2. In almost every case study there is some result calculated by DDP that is a "surprise". That is, it does not match the experts' intuitions (e.g., a risk shows up as more significant than they would have anticipated). Furthermore, when the experts look at the detail underpinning that result (and the capability of DDP to let them explore the details as well as see the big picture is crucial in this regard), they concede that the "surprise" is a genuine finding. Overall this suggests that the approach we follow is capable of findings that would be overlooked by even a well-qualified set of experts.
3. The decisions arrived at with the aid of DDP have led to much clarified in problem statements (objectives), better substantiated cost estimates, well-organized rationales for development plans, and, in some cases, very significant cost savings by stimulating design revisions motivated by the realization that resources were being mis-applied to (over)address a relatively tiny risk area, while other risk areas were in more dire need of attention. Perhaps the most important aspect shared by these decisions is that they are being made relatively early in the technology or system lifecycle, when the costs of revising choices are still relatively small. In contrast, fundamental design changes later in development incur much higher costs.

4. DDP is not a mandated activity, yet is seeing increasing use. The earlier DDP applications were supported by research funding (as was the development of the DDP process and support software itself). Specifically, the discipline experts' time was covered by this funding. More recently, there has been a trend to the use of DDP paid for predominantly by the technology/system under scrutiny. While we continue to use research funds to expand DDP's capabilities, the bulk of the cost of performing the risk studies themselves, namely the cost of the discipline experts' time, is no longer being borne by the research funding. This, plus the increasing frequency of DDP based studies, is indicative of an expanding acceptance of the net value of the DDP process.

Related work

Early decision-making is often assisted by *qualitative* decision support methods. For example Quality Function Deployment (QFD) has been used in a wide variety of settings [7]. DDP's effect and impact matrices are reminiscent of the Relationship Matrix used in many forms in QFD. DDP is distinguished by its foundation upon a *quantitative* risk model, which gives meaning to DDP's cost and benefit calculations.

As a design matures there are other decision support techniques that better capitalize upon knowledge of design details. For example, probabilistic risk assessment techniques (e.g., fault tree analysis, Bayesian methods) compute overall system reliability from design knowledge of how the system is composed of those components, and estimates of individual component reliabilities. The origins of these approaches lie in applications to assess risk in the nuclear power industry [8], with its need to estimate the probability of catastrophic failure (e.g., meltdown) from knowledge of the power system's design, and reliability measures for the components used in that design. Fault Tree Analysis [9] is now applied to a wide variety of systems, both hardware and software (e.g., software fault tree analysis [10]) including some NASA missions and their hardware and software components [11]. In contrast, DDP aims to fill the niche of early decision making for advanced technology and system development. We are currently exploring means to connect DDP and PRA techniques [12].

ACKNOWLEDGMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its

endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

- [1] O'Keefe, S. "*Pioneering the Future*", address by the NASA Administrator, Maxwell School of Citizenship & Public Affairs, Syracuse University, April 12, 2002 ftp://ftp.hq.nasa.gov/pub/pao/okeefe/2002/pioneering_the_future.pdf
- [2] Lehman, M.M. "Some Characteristics of S-type and E-type Software" in *Proceedings of the Second FEAST Workshop*, Dept. of Computing, Imperial College of Science Technology and Medicine, London, UK.
- [3] Cornford, S.L. "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.
- [4] S.L. Cornford, M.S. Feather & K.A. Hicks. "DDP – A tool for life-cycle risk management", *IEEE Aerospace Conference*, Big Sky, Montana, Mar 2001, pp. 441-451. <http://www.doc.ic.ac.uk/~mml/feast2/papers/pdf/2feastwks.pdf>
- [5] Feather, M.S. & Cornford, S.L.. "Quantitative risk-based requirements reasoning", to appear in *Requirements Engineering* (Springer), published online 25 February 2003, DOI 10.1007/s00766-002-0160-y. Available from: <http://eis.jpl.nasa.gov/~mfeather/AvailablePublications/>
- [6] Cornford, S.L., Feather, M.S., Dunphy, J.R., Salcedo, J. & Menzies, M. "Optimizing Spacecraft Design – Optimization Engine Development: Progress and Plans", *Proceedings, 2003 IEEE Aerospace Conference*, Big Sky, Montana, March 2003.
- [7]. Akao, Y. 1990 "*Quality Function Deployment*", Productivity Press, Cambridge, Massachusetts.
- [8] *Reactor Safety Study*, Report WASH-1400, Nuclear Regulatory Commission, 1975.
- [9] Vesely, W.E., Goldberg, F.F., Roberts, N.H. & Haasl, D.F., "*Fault Tree Handbook*", U.S. Nuclear Regulatory Commission NUREG-0492, 1981.
- [10] N.G. Leveson. "*Safeware: system safety and computers*". Addison Wesley, Reading, MA, 1995.
- [11] *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.
- [12] S.L. Cornford, T. Paulos, L. Meshkat & M.S. Feather. "Towards More Accurate Life Cycle Risk Management Through Integration of DDP and PRA". *IEEE Aerospace Conference*, Big Sky MT, Mar 2003.