

Towards Cost-Effective Reliability through Visualization of the Reliability Option Space

Martin S. Feather, Jet Propulsion Laboratory, California Institute of Technology, Pasadena

Key Words: Cost-benefit Tradeoffs, Heuristic Search, Visualization, Human-Computer Interaction, Optimization.

SUMMARY & CONCLUSIONS

In planning a complex system's development there can be many options to improve its reliability. Typically their sum total cost exceeds the budget available, so it is necessary to select judiciously from among them. Reliability models can be employed to calculate the cost and reliability implications of a candidate selection. However, there will typically be many such candidate selections, so we employ heuristic search techniques to explore reliability cost-benefit spaces, and visualization to present the results of these machine-conducted searches back to the expert designers to assist them in their decision making. This helps them understand (1) the overall cost-benefit trade space (how much reliability can be attained for a given level of expenditure, and how much additional reliability can be had for a modest increase in expenditure), and (2) the contribution of individual options (when an individual option (e.g., a test) contributes cost-effectively to the overall reliability). The net result is that expert human designers can make cost and benefit informed decisions on attainment of reliability. This approach is demonstrated on data drawn from planning the developments of advanced technologies for spacecraft.

1. INTRODUCTION

The context in which we work is the design and development of NASA's deep space probes. High reliability is required of these complex systems. Furthermore, high reliability is required of the development of those systems – that it be accomplished within budget, that it be completed within schedule (the spacecraft must be ready for launch within a certain time period when celestial mechanics allow for a time and energy efficient route to destination), and that the resulting system adhere to strict resource limitations (mass, volume, power, etc). Finally, NASA is continually pursuing new scientific objectives, for which new mission concepts are devised (e.g., sample-return missions), new technologies are developed (e.g., autonomous spacecraft control), and existing technologies are employed in novel settings (e.g.,

electronics subject to the harsh temperature cycles of planetary environments).

In response to these needs, at JPL and NASA a process has been developed to perform life cycle risk management [1]. One of the key aims of this process is to utilize the experience and insight of spacecraft experts drawn from the full range of discipline areas (e.g., science, navigation, propulsion, software, materials, communications). The complexity of the systems involved and the pressures on their development render purely manual planning and decision making problematic. Our risk management process therefore utilizes automation to assist those experts in pooling their knowledge, in deriving information from their pooled knowledge (via automated calculation and search), and in decision-making based on the sum total of that knowledge. This process is not, however, automatic. There is continued need for interchange between the experts and the automation, so that they can understand the results of that automation, and use those results to guide them (not replace them) in their decision-making. Visualization plays a key role in effective communication of information to those expert users.

The focus of this paper is on the visualization we use to convey automated search results, and their implications, to expert users. We employ heuristic search to help locate cost-effective solutions within a (large) space of reliability options. The large search space arises from the many options available to attain reliability – design options, materials options, test and analysis options, etc. At the heart of our process is a reliability model we use to calculate, for a given selection of options, the cost and benefit of that selection. We show how appropriate visualization of the search results is able to convey both the overall nature of the cost-benefit trade-space, and the contributions of individual reliability options. The net result is help to the expert users in making cost-benefit informed decisions concerning the reliability of the spacecraft systems whose development they are planning.

As we have stated, our work takes place in the context of spacecraft development. Nevertheless, we feel there are numerous terrestrial applications to which the same approach would apply. Many applications have

high reliability needs and have budget, time and resource constraints on their development. The reliability model that underpins our work is quite general in nature (it emphasizes spanning the breadth of concerns rather than the intricate details of the systems themselves). Moreover, whatever reliability model is used our method requires only that it be capable of calculating the cost and benefit of a candidate solution. Hence we are optimistic that our use of heuristic search for cost-effective reliability solutions, coupled with cogent visualization to report the results of that search, has broad applicability.

The remainder of this paper is structured as follows:

In section 2 we briefly summarize our reliability model. In section 3 we discuss the use of heuristic search to expose the cost-benefit reliability option space for a problem expressed in our reliability model, and show the use of visualization to convey the overall nature of the space. In section 4 we show how further visualization is able to reveal the contribution of individual reliability options. Finally, in section 5 we discuss status and mention some related work.

2. RELIABILITY MODEL

In this section we briefly summarize our reliability model, and identify the characteristics of the model upon which the rest of our approach relies. We have published widely on the model – for a relatively complete account, the interested reader is referred to [2].

2.1 Summary of Reliability Model and its Applications

Stated in the most general terms possible, the heart of our reliability model is very straightforward: for whatever system is being modeled, we capture within the model (1) the characteristics required of the system and of its development, (2) the key obstacles that, were they to arise, would impede attainment of those characteristics, and (3) the options available for preventing, removing or working around those obstacles. We also capture within the model quantitative values for how much each obstacle (were it to arise) would impede attainment of the characteristics, and for how much each option (were it chosen) would effectively reduce each obstacle. Finally, we capture the cost of each option, and the relative value of each characteristic.

When applied to spacecraft development, the characteristics required are often termed “objectives” or “requirements”, the obstacles are termed “risks” or “failure modes”, and the options “mitigations”. Most of our applications have been to individual technologies intended for use on space missions, for which: (1) the objectives encompass the science objectives driving the use of the technology (e.g., data accuracy and volume),

environmental constraints on resources available to the technology (e.g., RAM, power), and environmental constraints on the extent to which the technology can impact its environment (e.g., electromagnetic fields), (2) the risks encompass potential development problems (inability to construct, test, repair, operate and maintain the system) as well as the multitude of ways the operating system can fail to meet requirements, and (3) the mitigations encompass preventative measures that can be employed to reduce the likelihood of risks occurring (e.g., coding standards, training, use of qualified parts), to detect the presence of risks prior to fielding and use of the system (e.g., inspections, reviews, analyses, tests), and to alleviate the severity of risks (e.g., array bounds checking coupled with appropriate responses).

We have also applied this model to study designs of spacecraft apparatus. Currently the model is in use as part of the risk management process for an entire mission. One of our colleagues has explored the use of the model to assist activity selection across an entire program of NASA Earth Science Missions [Tralli, 2003]. We have even explored its use for scrutinizing a portfolio of research activities aimed at fulfilling the needs of practitioners, where the practitioners were encoded as the model’s requirements, practitioners’ needs as obstacles, and researchers activities as the mitigations [4].

The diagrams in this paper are computed from actual data gathered in the application of our model to spacecraft technologies. These applications are perhaps the closest to the reliability concerns of this symposium.

2.2 Key Characteristics of Reliability Model

Our use of search and visualization to explore the consequences of our reliability model relies on the model being “evaluatable”, in the sense that for a given selection of mitigations, the model can be automatically evaluated to yield a measure of cost and of benefit.

Costs arise in our model through association with mitigations (e.g., the cost of performing a test), and with repairs (e.g., the cost of repairing a problem that a test reveals). For a given selection of mitigations, our software automatically calculates the total of these costs.

Benefits arise in our model from the extent to which the objectives are attained. Individual objectives are given numerical weights to reflect their relative importance. For a given selection of mitigations, our software automatically calculates the total expected attainment of those objectives, taking into account the likelihood and severity of the risks, calculation of which in turn takes into account the effectiveness of the selected mitigations at reducing those risks.

The methods described in the following sections rely only on knowledge of the total cost and benefit for a selection of mitigations. Thus any reliability model capable of being evaluated to yield these measures would fit within our scheme.

3. USE OF SEARCH AND VISUALIZATION TO REVEAL THE OVERALL COST-BENEFIT TRADE SPACE

In our technology applications there can be many mitigations (dozens, possibly hundreds). Most of these are independent choices, so the combinatorics of selecting from among these imply a space containing huge numbers of possible candidate solutions. In a typical one of our applications, 58 mitigations represent design and development choices whose costs range from the low thousands of dollars to, in a few cases, hundreds of thousands of dollars. Since there are 58 mitigations, there are in principle 258 (approximately 1017) different selections from among them.

Straightforward incremental approaches to selection of mitigations are unlikely to lead to optimal solutions. For example, what might appear to be a promising-looking cost-effective mitigation (e.g., one that costs relatively little and significantly reduces a major risk) need not necessarily be part of an optimal solution. There could be another more expensive mitigation that reduces both that risk, and some other risk for which there are no alternative mitigations. Thus that more

expensive mitigation may well need to be selected anyway, rendering the promising-looking mitigation irrelevant. This phenomenon arises even within our relatively simple reliability modeling framework, because within our models it is typical for a given mitigation to reduce several risks (and in turn for a risk to threaten several objectives). Our models also encompass the phenomenon of mitigations that reduce some risks, but make others worse (e.g., a vibration test may serve to detect certain kinds of flaws, but also has some likelihood of causing flaws; software patches to fix one set of bugs can themselves introduce their own bugs). These further complicate judicious selection of mitigations. We speculate that reliability models that include design details such as redundancy, spares, etc., can readily introduce even more complexity.

In response, we use automated heuristic search for optimization. This is a widely-accepted approach to locating near-optimal solutions to complex design problems. For example, see [5] for an overview of this kind of work. We have implemented simulated annealing (a form of heuristic search), included as part of our software, and use it to locate near-optimal solutions. We have also explored genetic algorithms, and machine learning [6] for this same purpose. Using heuristic search we can search for a specific optimum. For example, for a given budget, find the set of mitigations that maximize the benefit (as calculated by our reliability model) while costing (as calculated by our reliability model) no more than that budget.

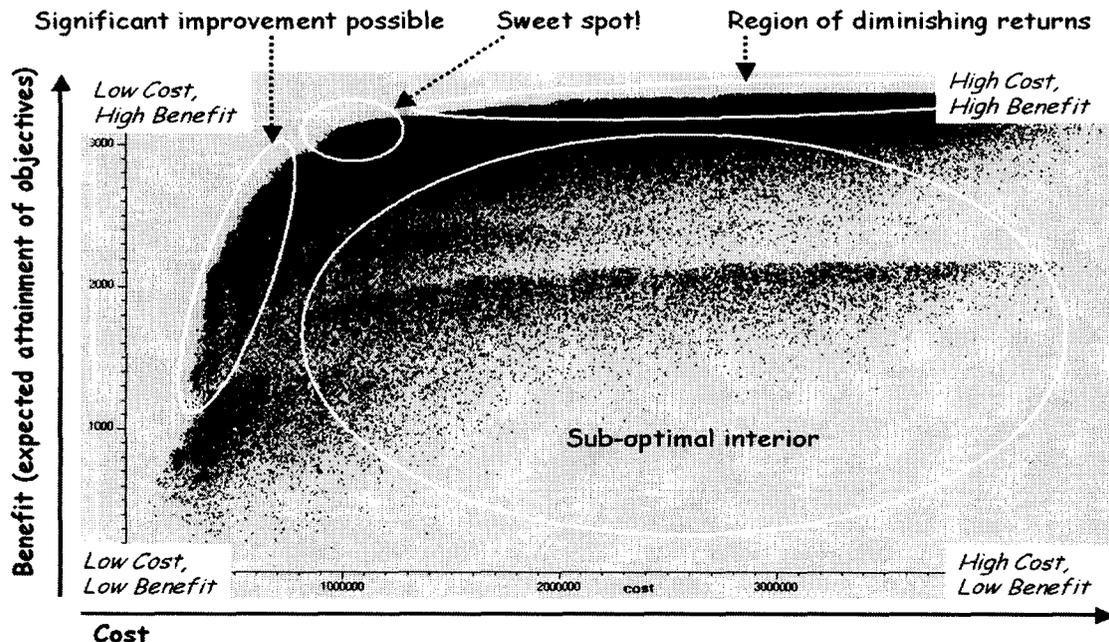


Figure 1. Visualization of the overall cost-benefit reliability option space

In order to reveal the overall cost-benefit trade space we use a series of individual cost-bounded optimal searches at successive cost levels. The result of such a series of searches, as performed on one of our technology studies, is shown in Figure 1. The sum total cost of all mitigations (approximately \$4,750,000) determines the rightmost value of the x-axis, and the sum total value of all objectives (approximately 3,600) determines the topmost value of the y-axis. Each of the approximately 300,000 individual points in the black “cloud” corresponds to a solution (i.e., selection of mitigations). For a given solution, our quantitative reliability model was used to calculate cost and benefit. A small black point corresponding to the solution was then drawn on the plot – solution cost determines horizontal position, solution benefit vertical position. The upper-left frontier of the cloud is thus the “optimal” boundary, also referred to as the “Pareto front” [5]. Note that we plotted a point for every solution investigated by the search, not just the “near-optimal” solution points on the boundary. The simulated annealing search is designed to concentrate towards this optimal boundary, so in fact there are many sub-optimal solutions not explored by this search, and so not plotted (were they to be plotted they would fall somewhere within the interior).

Our software automatically performs the heuristic search, calculating the cost and benefit for thousands of solutions as it progresses, and generates the plot of their points on the background of the cost and benefit axes seen in Figure 1. This is comprised of some 300,000 points. Its generation took approximately of 10 hours running on a 1.8 GHz PC. The primary determinant of the time is the time it takes the implementation of the reliability model to evaluate a given solution’s cost and benefit values.

For purposes of exposition, we have manually annotated the figure, using several white ellipses to highlight regions of interest:

- Points within the “Sub-optimal interior” indicate solutions that are far from optimal; for a given such solution, there are less expensive solutions that achieve the same level of benefit, and for that solution’s cost, there are more effective solutions that achieve more benefit.
- Points within the “Region of diminishing return” indicate solutions that, while close to or on the near-optimal boundary, achieve very little more benefit than less expensive solutions to their left. Their selection would be warranted only for the most risk-averse, wealthy, applications!
- Points within the “Sweet spot!” region indicate where we would like to be, budget permitting.

- Points within the “Significant improvement possible” region indicate solutions that, while close to or on the near-optimal boundary, could be significantly improved upon by a small increase in cost. If we find that for the budget available, solutions fall within this region, then there is a strong case to be made for additional budget. Asking for more money is common; the distinction here is that the justification for doing so is clearly evident.

It is possible that the available budget does not permit attainment of the requisite level of benefit. This could be because for that budget, solutions fall within the “Significant improvement possible” region. It could be because all the solutions, even the most expensive ones, fail to reach the requisite level of benefit (reliability). In response it may be appropriate to abandon some of the objectives, leaving a smaller set whose attainment can be more effectively achieved. While this will not necessarily increase the expected level of attainment of benefit, it may nevertheless be much preferred for reasons not included with the reliability model (e.g., highly negative public reaction to mission failure). This exemplifies the kind of strategic decision making (e.g., selection of a less ambitious mission with more certainty of success) that is informed by kind of cost-benefit trade space information we yield. That is, we use the power of automated search and visualization to convey information that is of value to decision makers, rather than striving to make the decisions for them.

4. USE OF VISUALIZATION TO REVEAL THE CONTRIBUTION OF INDIVIDUAL OPTIONS

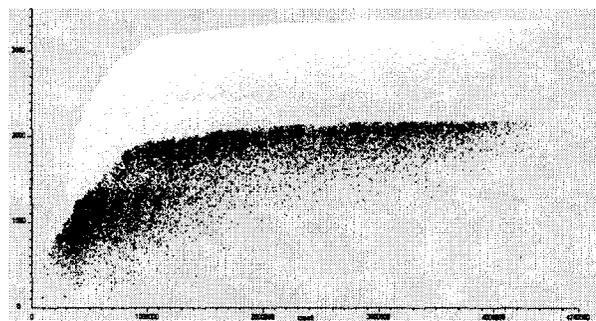


Figure 2. Visualizing the contribution of an individual mitigation

This section shows how the contribution of individual options within the overall cost-benefit trade space can be revealed through further visualization. An extreme example is shown in Figure 2 for a chosen one of the mitigations involved in the same reliability study as was plotted in Figure 1. This plots the same points as were in Figure 1, but colors each point:

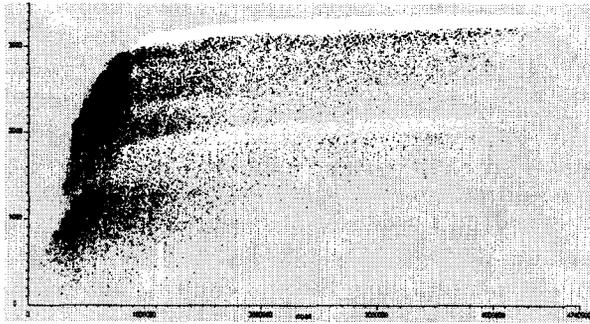


Figure 3. Another mitigation's visualization

- white if the chosen mitigation is included in the solution represented by that point, and
- black otherwise.

The broad swathe of white points that dominate the upper portion of the figure indicate that the chosen mitigation is key to nearly all optimal solutions. Only for very low cost levels do black points appear on the optimal frontier, indicating that the chosen mitigation would not be appropriate. We have chosen here a mitigation that plays a cost-effective role across almost the entire cost range of solutions. In fact, its cost is \$160,000 (to put this in context, the rightmost extreme of the horizontal axis corresponds to approximately \$4,750,000).

What appears to be a black "shadow" of the white region is, in fact, exactly that! Consider a solution found by our heuristic search that involves the chosen mitigation; its point will be located appropriately on the chart, based on that solution's calculated cost and benefit, and colored white. Our heuristic search method will try mutations of solutions, and so is likely to try a mutation of this solution in which the chosen mitigation is turned off. The resulting solution will be evaluated for cost and benefit, and its point located on the chart, colored black (since it does not involve the chosen mitigation). Compared to the first solution, it will cost \$160,000 less, so its point will be shifted a small distance to the left, and, since the solution contributes greatly to the benefit attainment, in its absence the

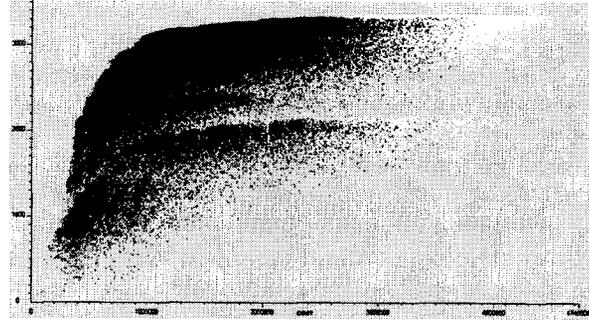


Figure 4. An expensive mitigation

benefit attainment will be a lot less, so its point will be shifted down. This phenomenon occurs for most of the swathe of white points, which we can see are shifted slightly to the left and significantly downwards to give rise to the black region.

Figures 3 and 4 show application of this same visualization technique to two different mitigations (to save space, these figures are half the height and width of the earlier figures). The mitigation shown in Figure 3 is a little more expensive than the mitigation of Figure 2 (\$200,000 rather than \$160,000) but from this visualization can be seen to be cost-effective (i.e., appears along the optimal boundary) only for solutions starting at approximately \$1,000,000. The mitigation shown in Figure 4 is significantly more expensive (\$700,000), and not surprisingly plays a role in optimal solutions at only the more expensive end of the spectrum (starting at approximately the \$3,000,000 level).

Figures 5 and 6 show two more mitigations. From their visualizations, it is clear that they are not cost-effective. The mitigation of Figure 5, even though relatively low cost, appears in very few places (there's a hard to discern at this scale small patch of white on the optimal frontier around the \$1,500,000 level). In fact, this is a mitigation that (according to the experts who provided the data in the reliability model) reduces some risks, but makes some other worse. This visualization suggests avoiding its use in this application, in keeping with the experts' intuition, it turns out. Figure 6 shows

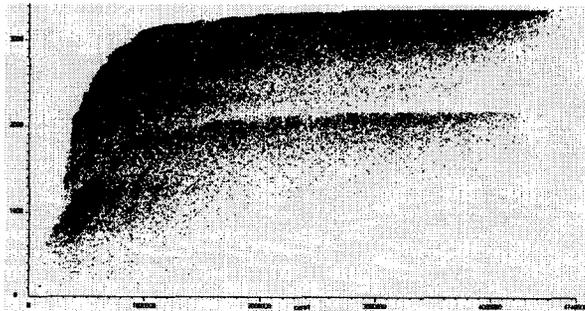


Figure 5. An unworthy mitigation

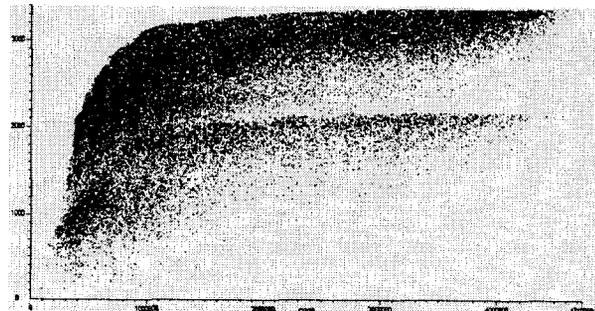


Figure 6. A dubious mitigation.

been gathered in the course of an actual technology study performed at JPL – the figures in this paper have all been generated from this study's actual data. It has subsequently been used in another technology study. In the latter, we performed the experiment of inspecting the visualizations of each of the mitigations, and on that basis estimating whether (for a chosen cost level) the mitigation appeared to be cost-effective or not. We then manually selected those and only those mitigations we deemed to be cost-effective, and using our tool to evaluate the reliability model on that solution, determined where in the cost-benefit trade space that solution would fall. The result was a point that was indeed located on the optimal cost-benefit boundary in the region of our chosen cost level, thus confirming that the insights gained from the individual visualizations “add up” when used to compose an entire solution.

An interesting approach to simultaneous visualization of multiple design variables within complex system designs is presented in [7]. Their approach allows the user to pick several design variables at once, and see the values those variables take on in design solutions spread through a multi-dimensional design space. Their approach applies to designs whose variables can take on values over continuous ranges (e.g., thickness of material). In contrast, our approach has been applied to only binary choices, where those are choices from among options of steps that increase reliability.

An alternative to visualization is to use some kind of analytic method to discern the utility (or otherwise) of mitigations. An interesting and apparently very general approach to this is the work of Tim Menzies (<http://tim.menzies.us>). His method is able to identify for a model which, out of its set of choices, are the ones most influential in determining the outcome. It is based on random sampling of model inputs, using the model to evaluate each set of inputs, and using his machine-learning based approach to identify significant contributors, including both those that is important to include in the solution, and those it is important to exclude from the solution. Menzies has applied this

The benefit of these forms of visualization stem from the insights they provide the experts. Cost-effective design of high-reliability systems requires a blend of input, guidance and insight that the human experts can provide, and the storage, calculation and search capabilities that the software can provide. Visualization is key to conveying information from the software to those experts.

The status of our work is that the visualizations shown in this paper are generated from our implemented software. The capability of visualizing the effect of individual mitigations provides additional information to assist those experts in their decision-making. In particular, it reveals which of the design choices are crucial, both choices of options crucial to perform, and choices of options crucial to not perform. This indicates to the designers the restrictions on their freedom of choice if they want to obtain solutions close to the near-optimal design frontier.

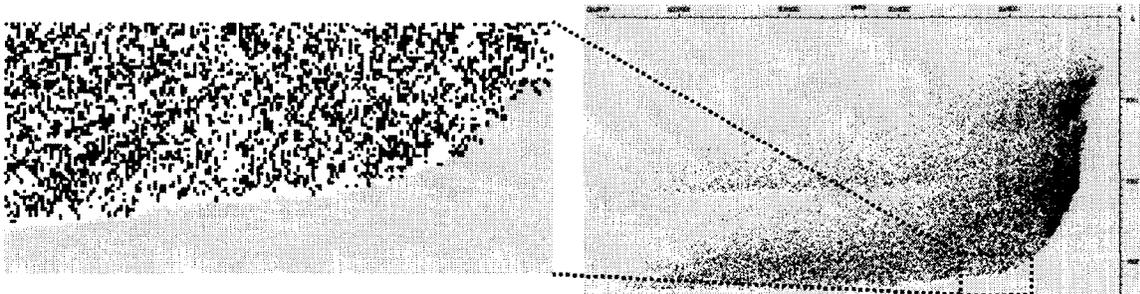
The overall cost-benefit trade space visualization has been utilized to help experts make costing decisions for technology developments.

The capability for visualization of individual options was developed and tested on data that had previously

5. STATUS AND RELATED WORK

an instance of a mitigation of dubious utility. The plot has a “speckely” look, because there is a close mixture of white and black points, with black points slightly dominating. It turns out to be a low cost (\$5,000), low benefit mitigation. For cases such as these, “zooming in” on the optimal frontier can show whether or not the mitigation is of net benefit. This is seen in Figure 7, where another mitigation of uncertain utility (again, with a “speckely” look, albeit slightly more light colored overall) is shown alongside a zoom-in to a portion of the optimal frontier. From the zoom-in, it can be seen that this mitigation does appear almost uniformly along the upper portion of the frontier, indicating that it has an overall positive (albeit small) role to play in near-optimal solutions.

Figure 7. Zoom-in on a mitigation



approach to, among other things, Boehm's well known COCOMO cost/risk model for software cost and risk estimation [8], and the Software Engineering Institute's recommended set of practices included in level 2 of their Capability Maturity Model (CMM) for software development [9]. We have had some opportunity to try his approach on our reliability models, with positive results [10]. We are motivated to seek a more direct comparison of his analytic approach with our visualization method.

ACKNOWLEDGEMENT

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. I am pleased to thank Steve Cornford for his leadership of the development of the reliability model used in this paper, and Kenneth Hicks for his work to guide the application of that model to technology applications for spacecraft use.

REFERENCES

1. S.L. Cornford, M.S. Feather and K.A. Hicks: "DDP - A tool for life-cycle risk management", *Proceedings, IEEE Aerospace Conference*, Big Sky, Montana, Mar 2001, pp 441-451.
2. M.S. Feather, & S.L. Cornford. "Quantitative risk-based requirements reasoning", to appear in *Requirements Engineering* (Springer), published online 25 February 2003, DOI 10.1007/s00766-002-0160-y. (See <http://eis.jpl.nasa.gov/~mfeather/Publications/> for an advance copy)
3. D.M.Tralli. Programmatic Risk Balancing. *Proceedings, IEEE Aerospace Conference*, Big Sky Montana, Mar 2003, pp. 2.775-2.784.
4. M.S. Feather, T. Menzies & J. Connelly. "Matching Software Practitioner Needs to Researcher Activities", to appear in the *2003 Asia-Pacific Software Engineering Conference* (APSEC 2003), Chiangmai, Thailand, Dec 2003.
5. P. Sen & J-B. Yang. *Multiple Criteria Decision Support in Engineering Design*, Springer-Verlag, 1998.
6. S.L. Cornford, M.S. Feather, J.R. Dunphy, J. Salcedo & T. Menzies. "Optimizing Spacecraft Design -

Optimization Engine Development: Progress and Plans", *Proceedings, IEEE Aerospace Conference*, Big Sky, Montana, Mar 2003, pp. 7.3361-7.3368.

7. G.M. Stump, M. Yukish, T.W. Simpson & E.W. Harris. "Design Space Visualization and its Application to a Design by Shopping Paradigm", to appear in the *ASME Design Engineering Technical Conferences - Design Automation Conference*, Chicago IL Sept 2-6, ASME Paper No. DETC2003/DAC-48785.

8. T.J. Menzies & E. Sinsel. "Practical Large Scale What-if Queries: Case Studies with Software Risk Assessment", *Proceedings 15th IEEE International Conference on Automated Software Engineering*, Grenoble, France, Sept 2000, pp 165-173.

9. T. Menzies & J.D. Kiper. "Better reasoning about software engineering activities", *Proceedings 16th International Conference on Automated Software Engineering*, San Diego, California, Nov 2001, pp 391-394.

10. M. Feather & T. Menzies. "Converging on the Optimal Attainment of Requirements", *Proceedings of the IEEE Joint International Conference on Requirements Engineering*, Essen, Germany, Sept. 2002, pp 263-270.

BIOGRAPHY

Martin S. Feather, PhD.
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109-8099, USA

e-mail: Martin.S.Feather@Jpl.Nasa.Gov

Martin S. Feather is a Principal in the Software Quality Assurance group at JPL. He works on developing research ideas and maturing them into practice, with particular interests in the areas of software validation (analysis, test 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>