Descoping

Martin S. Feather Steven L. Cornford Kenneth A. Hicks
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109, USA
Martin.S.Feather@Jpl.Nasa.Gov Steven.L.Cornford@Jpl.Nasa.Gov Kenneth.A.Hicks@Jpl.Nasa.Gov

Abstract

Descoping is the strategic abandonment and/or weakening of objectives. It is required whenever limited resources preclude satisfactory attainment of all those objectives. Potential causes of the need for descoping are numerous, and descoping is a recurring phenomenon during project planning and development.

We present an approach which facilitates descoping. It is founded upon a quantitative model of requirements attainment, resource consumption, and risk. Features of this model allow for the representation of interactions between objectives. Measures derived from this qualitative model support the identification and evaluations of descope options. Tool support for the model gives assistance to users in making their descope decisions.

1. Introduction

Determining the objectives (a.k.a. requirements) is widely recognized as one of the crucial early steps in project planning. In almost all cases the objectives must be balanced against the costs of attaining them - it is rare that objectives are committed to no matter what their cost. Cost limitations force selection of the subset of objectives to be pursued. Later in the lifecycle, deviations from the planned development process lead to the need to revisit this selection. Schedule slippages, cost overruns, and requirements changes can each contribute to this. Under fortuitous circumstances, an *increase* in objectives could be feasible, but the much more common situation is to need to descope further. For the purposes of this paper, the word "descope" is intended to cover both kinds of down-selection, both during initial planning, and during the course of development.

Section 2 describes the challenges of descoping, and past work in this area. Section 3 introduces the quantitative risk-based model that serves as the basis for our investigations of descoping. Section 4 presents the ways in which this model supports in-depth descope planning. Section 5 offers a conclusion, and some suggestions for future work.

2. Descoping Challenges

To descope effectively requires cost estimation (how much it will cost to attain a given set of objectives) and valuation (what is the end value of attaining a given set of objectives). These are both research areas in which there has been substantial progress. For example, COCOMO [Boehm et al, 2000] helps predict costs once the overall project characteristics (both product characteristics, and development process characteristics) have been estimated. Accord [Ullman, 2001] helps groups of people achieve consensus on the preferred set of objectives.

Descoping is complicated when objectives interact (i.e., when they are interdependent, so that an objective cannot be considered in isolation of all the other objectives). Such interaction appears to be commonplace. [Carlshamre et al, 2001] report a study in which they found interdependencies to be the norm in their setting (Ericsson Radio Systems). Robinson et al [Robinson et al, 1999] employ the term "requirements interaction management" in their survey of the broad range of studies in this area.

Other terms for what we are here calling "descoping" include "requirements prioritization", "requirements triage" [Davis, 2000], and (especially in the context of commercial software products) "release planning". Examples of tool-supported approaches that assist in this area include: the cost-and-value based approach [Karlsson & Ryan, 1997], the "negotiated win conditions" of [Boehm et al, 1994], the explicit treatment of non-functional requirements in evaluation alternatives as part of the i* approach [Mylopoulos et al, 2001].

Our setting, that of spacecraft design and operation, faces these same pressures. We are resource constrained – NASA's budget must be allocated to best achieve science return; launch vehicle capacities constrain mass and volume; solar panels can yield only so much electrical power. We too are often schedule constrained – albeit not because of economic pressures to be first to market, but because of cosmological factors that favor certain launch windows (e.g., proximity in orbit between Earth and Mars). Spacecraft introduce yet another complication - risk. Risk is unavoidable in our setting, due to the potential for irreparable hardware failures, unpredictable

aspects of the environment, lack of detailed and/or current knowledge of the spacecraft state (because of limited communications bandwidth and long round-trip light times), and the sheer complexity of their multi-disciplinary development. This forces the consideration of not only *which* objectives to select, but also *how diligently* to pursue them. [Greenfield 1998] recognized the need to trade risk itself as a resource, alongside other key factors in spacecraft development (e.g., cost, schedule, mass, power).

3. A risk-based cost-benefit model

The basis for our investigations is a quantitative risk-based model that we have been developing at JPL. This model, called "Defect Detection and Prevention (DDP)", has been applied to help assess and plan developments of novel spacecraft technologies and systems [Cornford et al, 2001], [Cornford et al, 2002]. We have reported on this model in other publications, with an emphasis on how it is used to assess risk and plan how to best reduce risk. Here we focus on descoping., where the aim somewhat different, namely to identify the objectives to abandon.

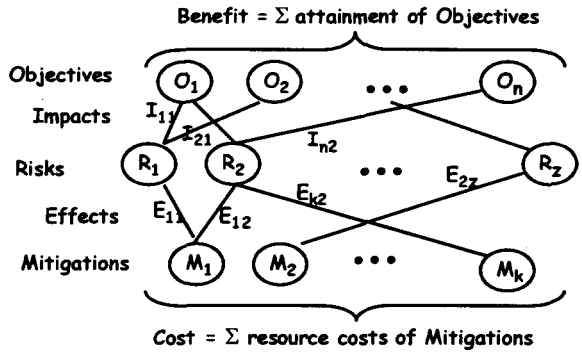

Figure 1. Topology of DDP's risk-centric cost-benefit model

The topology of DDP's risk-based model is sketched in Figure 1. *Objectives* (O1, O2, ...) are given weights, reflecting their relative importance. *Risks* (R1, R2, ...) are all the things that, should they occur, have adverse *Impacts* on Objectives. These Impacts are assigned numerical strengths, indicating how much of the Objective would be lost should the Risk occur. *Mitigations* (M1, M2, ...) are all the things that should they be applied, have a reducing *Effect* on the likelihood and/or impact of Risks. These Effects are assigned numerical strengths, indicating by how much the likelihood and/or impact of the Risk will be reduced should the Mitigation be applied. On occasion a Mitigation may make certain Risks *worse*, either by "aggravating" the Risk (making its impacts on Requirements more severe) or by "inducing" the Risk (increasing its likelihood of occurrence). The DDP model's Effect links accommodates these phenomena.

Overall, the *cost* of a DDP model is the sum of the costs of the Mitigations selected for application. The *benefit* of a DDP model is the sum of attainment of its Objectives, calculation of which takes into account the Risks' impacts on those Objectives, moderated by the reducing effects of the selected Mitigations on those Risks.

In practice, the DDP data for a given application is voluminous and coupled. This can be seen in Figure 2, which shows data from an actual application drawn in this "topological" presentation style. Each of the squares in the top row represents an Objective; there are 32 in all. Each of the squares in the middle row represents a Risk; there are 69 of them in all. 352 Impact links connect Objectives and Risks. Each of the squares in the bottom row represents a Mitigation; there are 99 of them in all. 440 Impact links connect Mitigations and Risks. The DDP tool uses a variety of alternate presentations that make it possible to elicit and scrutinize this kind of information [Feather et al, 2000].

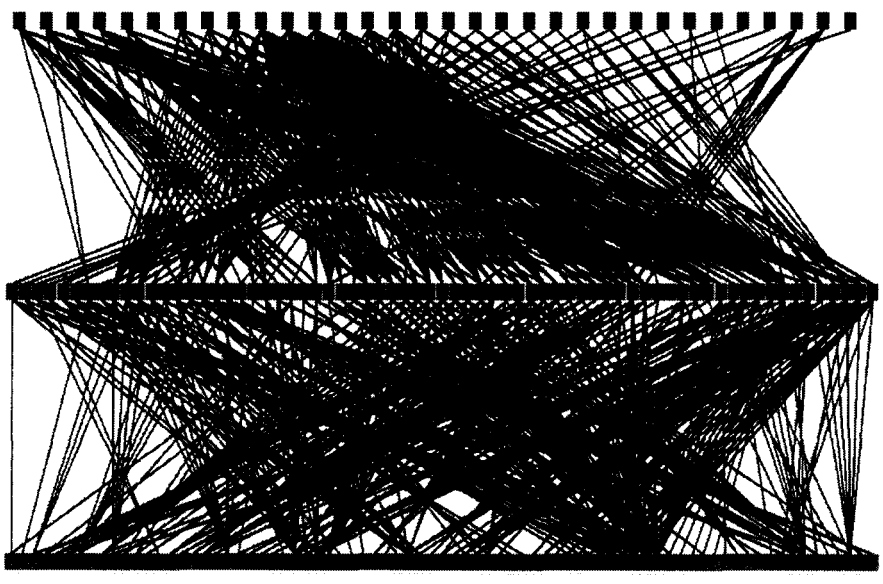

Figure 2. Topology of an actual DDP model

4. DDP support for descoping

DDP allows experts to pool their knowledge, and gain insight into the ways in which Objectives can (or cannot) be attained by suitable selection of Mitigations.

DDP was originally conceived of to guide the judicious selection of quality assurance activities [Cornford, 1998]. In most such applications, the sum total cost of all the possible activities (which become Mitigations in the DDP model) far exceeds the resources available. DDP is used to guide experts to selecting of a subset of those Mitigations that maximize

attainment of Objectives while remaining within resource limits. Objectives are used primarily to provide a measure of the benefit of a given solution.

In our applications of DDP to a variety of problems, we have found that it also gives insight into the Objectives themselves. This can be useful when, as is often the case, resources are so scarce as to preclude the satisfactory attainment of them all, thus necessitating descoping. The subsections that follow detail the key ways in which the DDP model and its software facilitate such descope decision making:

- The model's explicit and detailed treatment of interactions.
- The various quantitative measures that reveal different aspects of descope needs.
- Visualization to permit users to see the ramifications of those measures.
- Optimization to direct users towards descope options worth of particular attention.

4.1. An explicit and extensible model of interactions

Interactions arise in our DDP model through the Impact and Effect connections. These cross-couple the Objectives, Risks and Mitigations. A Risk may impact multiple Objectives (to different extents); an Objective may be impacted by multiple Risks. Similarly, a Mitigation may effect multiple Risks (to different extents); a Risk may be effected by multiple Mitigations.

This explicit treatment of cross-coupling, using risks as the intermediary, is central to the DDP model and sets it apart from other approaches. Typical requirements prioritization approaches have tended to follow the route of eliciting pairwise couplings directly among the objectives themselves (e.g., the "interdependencies" of [Carlshamre et al, 2001]). Typical risk management approaches have tended to ask users to estimate the risk (likelihood and severity) that remains after taking into account all the planned risk-reducing activities. In contrast, DDP *derives* a risk's severity from the sum total of its Impacts on Objectives moderated by the mitigations whose effects reduce its impacts, and *derives* a risk's likelihood by starting with its a-priori likelihood, and taking into account the mitigations whose effects reduce its likelihood.

The disadvantage of the DDP approach is that it requires more information, which takes longer to gather from the experts involved in the study. However, the following advantages stem from DDP's more thorough approach:

- Rederivation of risk when the situation changes. During descoping, risk can be rederived from the modified information. A risk impacting an objective that is removed or downgraded during descoping will be commensurately less severe.
- Guidance on selection of mitigations. Because mitigations are explicitly represented, DDP is able to guide users to cost-effectively select which mitigations to apply. During descoping, users will likely need to revisit this selection. Furthermore, the ability to trace from objectives to risks to mitigations allows users to understand which of their objectives are proving the most costly to attain.

The DDP model is extensible. Users can and do add instances of Objectives, Risks and Mitigations pertinent to the application they are studying. Of course, when they add them, they must also input the Impact and Effect values that relate them.

Thus performing a study using DDP is a non-trivial effort. In our applications of DDP (e.g., to assess and plan for the maturation of technologies to make them ready for use on spacecraft) it has been typical to require a team of 10 – 20 experts to provide information in several half-day long sessions. The following advantages stem from DDP's extensibility:

- Ability to capture problem-specific information. Objectives are particularly problem-specific of course, but so are risks, and mitigations. For example, we used DDP's extensibility to complement a more closed-form software quality assurance planning tool [Kurtz & Feather, 2000].
- Ability to adjust the information as the situation changes. Descoping that is motivated by a change in the availability of resources (time, schedule, facilities, personnel) can be explored by adjusting the corresponding entries in the DDP model.

4.2. Quantitative measures of attainment, and their role in descope planning

The DDP model defines several quantitative measures, which the DDP tool automatically computes from the user-provided data. The key such measures relevant to descoping are outlined next.

- **Objective's degree of attainment**: defined for each objective as the proportion of that objective attained. Its definition takes into account the adverse impact of all extant Risks and the reducing effects on those of all selected Mitigations.
- **Sum total attainment of all objectives**: defined for the entire model as the sum, over all objectives, of each objective's weight times its degree of attainment. This is the overall "value" measure of a DDP model.

Objective's degree of risk: defined for each objective as the proportion of that objective impacted by all extant Risks, taking into account the reducing effects of all selected Mitigations. In the DDP model, it is possible (indeed common) for an objective's degree of risk to be *greater* than 1.0! This indicates an objective adversely impacted by several Risks, so much so that they more than completely eliminate attainment of that objective.

The **objective's degree of attainment** measure gives an indication of how well an individual objective is being attained. This measure is used to understand which objectives are being attained, and by how much, given a DDP model's configuration (set of Risks and selection of Mitigations). Objectives that are being completely, or nearly completely, attained are low-risk items that we can have confidence will likely be met.

The **sum total attainment** measure gives an indication of the overall value of a DDP model as currently configured. This can be used to compare major alternative descope options.

The **objective's degree of risk** measure gives an indication of how much work needs to be done to attain an objective. For an objective that is less than totally eliminated by Risks, this is the complement of its degree of attainment measure. That is, under those circumstances, degree of attainment = (1 – degree of risk). However, when an objective is more than totally eliminated by risks, its degree of attainment will be zero, while its degree of risk will be greater than 1. This degree of risk measure is important to understanding how implausible an objective really is. One that is just slightly over 1.0 at risk is a candidate for improvement (by selection of additional Mitigations to reduce the Risks impacting that objective), but one that is well over 1.0 at risk is a strong contender for descoping.

4.3. Visualization support for exploring descope options

The DDP software supports users in their exploration of descope options. Users can change the selection of mitigations and see the ramifications on all the automatically calculated measures. Following a change recalculation of these measures is automatic and rapid. For example, on the DDP model shown in Figure 2, a typical change (e.g., turning off an objective) requires less than a second on a 2MHz PC for recomputation. Speed is therefore not a problem. More challenging is the presentation of this information to users. The DDP tool provides information visualization via cogent displays designed to convey the status of these measures and the information that underpins them.

One of the key displays is the bar chart of Objectives,

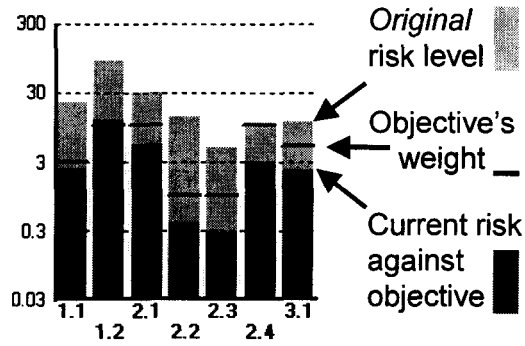

Figure 3. Fragment of DDP display of objective attainment measures

which shows several of the measures associated with each objective. Figure 3 presents an annotated fragment of such a bar chart. Each column represents a separate objective, the number below corresponding to its place in the objectives tree. Within a column, three measures are indicated: a horizontal bar denotes the objective's user-assigned weight, a lightly shaded bar indicates the objective's original degree of risk (i.e., how much it would be at risk if *no* mitigations were to be applied) and a dark bar indicates the objective's current degree of risk (i.e., how much it is at risk taking the risk-reducing effects of the currently selected Mitigations into account). In the tool itself, color is used to make these distinctions more vivid. For the purposes of this paper, they have been rendered in grayscales. Thus for Objective number 3.1, it was originally more than totally at risk, but that risk has been reduced somewhat by the currently selected Mitigations. Note that the vertical axis is logarithmic, hence the reduction in risk is fairly significant. We use a logarithmic scale because in our setting reducing risks to relatively small levels is our usual goal, for which a logarithmic scale is better suited than a linear scale.

Figure 4 shows the entire bar chart display of Objectives taken from actual DDP data, at a stage of partial risk mitigation. From this display it is evident that significant risk remains, and in some cases there are Objectives which are still more than totally at risk, e.g., numbers 4.1 and 4.2.

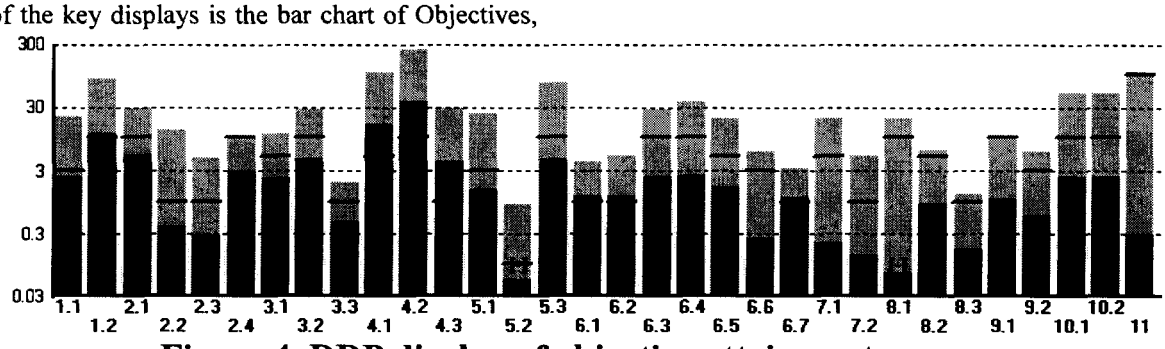

Figure 4. DDP display of objective attainment measures

4.4. Optimizing the attainment of objectives

DDP's calculation of costs and benefits permit treatment of the cost/benefit tradeoff as an *optimization* problem. The goal could be to select mitigations that maximize the sum total of objectives' attainment while staying within a cost ceiling. Alternately, the goal could be to select mitigations that minimize cost while achieving attainment of objectives at or above a benefit floor. However, the voluminous and intertwined nature of typical DDP models makes this challenging. For example, in a DDP model with 50 Mitigations, there are 2^{50} possible ways of selecting which of those Mitigations to perform.

We use heuristic search techniques to automatically explore the large search space of Mitigation selections. We have had success using both genetic algorithms, and simulated annealing. We have also collaborated with Tim Menzies to apply his machine-learning based technique, which has proven capable of both finding near-optimal solutions, and identifying the Mitigations whose selection (or non-selection) are most critical to get those solutions [Feather & Menzies, 2002]. More details on our use of heuristic search can be found in [Cornford et al, 2003].

Shown in Figure 5 is an illustration of an amalgam of simulated annealing searches applied to optimizing an actual DDP spacecraft technology model across its range of cost levels. The horizontal axis of the plot corresponds to cost, with greater costs to the right. The vertical axis corresponds to benefit (sum total attainment of objectives), with greater benefits to the top. Thus the optimum is to be found towards the upper left corner, where costs are low and benefits high. Each tiny point corresponds to a selection of Mitigations, which the DDP tool has evaluated to determine cost and benefit, and plotted at the corresponding position on the chart. A series of simulated annealing searches, at successive cost levels, have been combined to reveal the overall cost-benefit profile of the model's data. The searches tend to focus their attention at or near the optimal boundary. Of course, their searches explore numerous points in the interior. Hence the scattering of points across the chart, with a concentration towards the optimal boundary.

Optimization runs such as this reveal to users the cost/benefit space. The points along the upper boundary give an indication of the feasible near-optimal profileGuided by this information, users are assisted to:

- Predict the level of benefit that can be had for a given cost (or equivalently, predict the cost it takes to attain a given level of benefit).
- Justify the need for additional funding when the profile indicates that there is still a significant gain in objectives to be had for modest increases in spending.
- Avoid overspending. For example, it is evident from the figure that close to maximum possible attainment of objectives can be had for approximately one-quarter of the maximum cost. Above that level of spending, the additional benefit gain is very gradual. Conversely, below that level of spending, benefits start to drop of markedly.

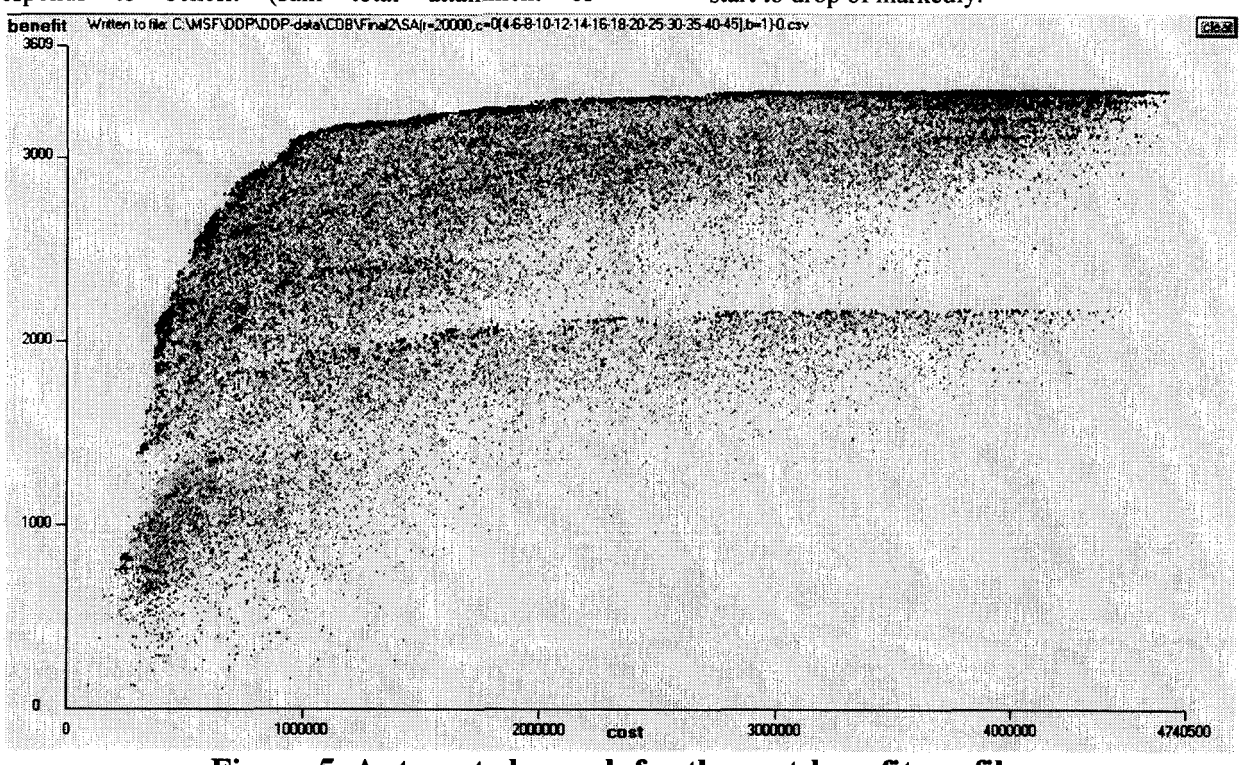

Figure 5. Automated search for the cost-benefit profile

Users can view specific points in this solution space in through the DDP tools' usual displays. For descope purposes, they focus on objectives attainment of a given solution. To do so, they use the objectives attainment, they use the bar chart display shown earlier to see which objectives are being attained and to what extent, and which objectives are being left unattained (i.e., still significantly at risk).

4.5. Combination of features

These features of the DDP model and its software work well in combination. The detailed and extensible quantitative risk model is the basis. The DDP software supports populating this model with information elicited from experts in group sessions in real time. DDP automatically computes several measurements from the accumulated information, including several that have particular relevance to guiding descope decisions. Cogent visualizations present the information in ways palatable to human scrutiny. Automated heuristic search techniques are provided to help explore the large space of options, the results of which can be used to guide strategic decisions of how much overall to spend, and specific solutions points can be scrutinized in detail.

5. Conclusions and future work

We have outlined the descope problem. In our context of spacecraft development, we face many of the same pressures on project development as are common elsewhere (severely limited schedules, budgets and other resources). In addition, we must explicitly deal with risk. All these factors combine to make descope a recurring need.

The quantitative risk-centric model we have used for risk management is well suited to in-depth consideration of descope options and their implications. It has a detailed model of interactions among objectives, risks and mitigations. Quantitative measures defined in terms of that model give insight into opportunities for, and consequences of, descopes. Cogent visualizations inform project managers of this information, facilitating their strategic decision-making. Automated optimization (using heuristic search techniques) finds descope opportunities along the near-optimal boundary of the cost-benefit trade space.

Taken together, these capabilities provide significant support for strategically planning descopes. We see the need for further work in the areas of:

- Enhanced interplay between the optimization / search techniques, and human-guided decision-making. For example, allow users to impose additional constraints on the solution sets they are willing to accept, and re-optimize taking those into account. Out collaborative work with Tim Menzies [Feather & Menzies, 2002] has an aspect of this. The work of [Menzies & Hu, 2001] suggests opportunities for more such benefits.
- Exploration of descope options that, rather than discard objectives, change their relative weights. For example, suppose a mission with primary and secondary science return objectives needs to be descoped; rather than discarding one of those objectives, perhaps the descope needs can be met be reversing their prioritization? The New Millenium missions [Minning at al, 2000], each designed to flight validate advanced technologies, would be promising application areas for this.

The status of our work is that the DDP model has been used on numerous studies of spacecraft technologies. Although we did not approach those studies with descoping of objectives in mind, it is interesting to note that some of them led to descope decisions. In retrospect, we see descoping as a recurring phenomenon, and are encouraged by DDP's ability to assist in this. Future work will, we hope, further extend its capabilities in this direction.

We are currently performing some experiments to *automatically* explore how descope options might lead to alternate selections of mitigations. Given a DDP model for which the optimal cost-benefit profile has already been computed (as in Figure 5), we systematically turn off one of the objectives, and re-perform the optimal search at a given cost level or levels. The aim is to find a solution (selection of Mitigations) that attains the *remaining* objectives more effectively. The existence of such a solution indicates that the turned-off objective is a candidate for descoping. Conversely, the lack of such a solution indicates that the turned-off objective was being attained anyway, so there is little or no point to its descoping. Initial results in running this on actual DDP application data shows that there are surprisingly few instances where turning off a single objective makes a radical difference. We are extending these experiments to turn off multiple objectives. Unfortunately, the naive way we are approaching this problem is computationally expensive – each time we turn off an objective, we have to re-run the heuristic search to see whether new solutions emerge. Compounding this, if we wish to explore turning off pairs, triples, etc., of objectives, the number of possible combinations grows combinatorially. This is clearly an area where a more sophisticated approach is sorely needed.

6. Acknowledgements

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. Discussions with Julia Dunphy (JPL), Tim Menzies (WVU), Burton Sigal (JPL) and Stephen Prusha (JPL) have been most useful in helping us formulate our ideas.

7. References

[Boehm et al, 2000] B. Boehm, C. Abts, A.W. Borwn, S. Chulani, B.K. Clarl, E. Horowitz, R. Madachy, D. Reifer & B. Steece. "Software Cost Estimation with COCOMO II", Prentice-Hall, July 2000.

[Boehm et al, 1994] B. Boehm, P. Bose, E. Horowitz, and M. Lee, 1994, "Software Requirements as Negotiated Win Conditions", Proceedings 1st International Conference on Requirements Engineering, Colorado Springs, Colorado, pp. 74-83.

[Carlshamre et al, 2001] P. Carlshamre, K. Sandahl, M. Lindvall, B. Regnell & J. Natt och Dag. "An Industrial Survey of Requirements Interdependencies in Software Product Release Planning", *5th IEEE Int. Symposium on Requirements Engineering*, 2001, pp. 84-91.

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

[Cornford et al, 2001] 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.

[Cornford et al, 2002] S.L. Cornford, J. Dunphy and M.S. Feather, 2002, "Optimizing the Design of end-to-end Spacecraft Systems using Risk as a Currency", IEEE Aerospace Conference, Big Sky, Montana, Mar 2002.

[Cornford et al, 2003] S.L. Cornford, J. Dunphy, M.S. Feather and J. Salcedo, 2002, "Heuristic Optimization for Risk Reduction in Spacecraft Designs", submitted to the 2003 IEEE Aerospace Conference, Big Sky, Montana.

[Davis, 2000] Davis, Al. "The Most Ignored Piece of Requirements Engineering: Requirements Triage" *ICRE 2000 Tutorial.* Omni-vista inc., Colorado Springs, CO.

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

[Feather & Menzies, 2002] Feather, M.S. & Menzies, T. "Converging on the Optimal Attainment of Requirements", to appear in the Proceedings of the IEEE Joint International Conference on Requirements Engineering, Essen, Germany, Sept. 2002.

[Greenfield, 1998] M.A. Greenfield "Risk Management 'Risk As A Resource' " http://www.hq.nasa.gov/office/codeq/risk/

[Karlsson & Ryan, 1997] J. Karlsson & K. Ryan. A Cost-Value Approach for Prioritizing Requirements. *IEEE Software*, Sept./Oct. 1997, 67-74.

[Kurtz & Feather, 2000] T. Kurtz & M.S. Feather: "Putting it All Together: Software Planning, Estimating and Assessment for a Successful Project", in *Proceedings of 4th International Software & Internet Quality Week Conference*, Brussels, Belgium, 2000.

[Menzies & Hu, 2001] T. Menzies & Y. Hu. "Constraining Discussions in Requirements Engineering via Models", *1st International Workshop on Model Based Requirements Engineering*, San Diego, California, Dec 2001.

[Minning at al, 2000] C. P. Minning, M. G. Buehler, T. Fujita, F. Lansing, G. Man, A. Aljabri, and C. Stevens, "Technology Selection and Validation on New Millennium Flight Projects," IEEE Aerospace Conference, Big Sky, Montana, Mar 2001

[Mylopoulos et al, 2001] J. Mylopoulos, L. Chung, S. Liao, H. Wang, and E. Yu, 2001, "Exploring Alternatives during Requirements Analysis", IEEE Software 18(1), pp. 92-96.

[Robinson et al, 1999] W.N. Robinson., S. Pawlowski & S. Volkov. "Requirements Interaction Management", GSU CIS Working Paper 99-7, Georgia State University, Atlanta, GA, August, 30, 1999.

[Ullman, 2001] D. Ullman. "12 Steps to Robust Decisions: Building Consensus in Product Development and Business", Trafford Publishing, 2001.