Paper #160

Postoptimality Analysis in the Selection of Technology Portfolios

Virgil Adumitroaie, Kacie Shelton, Alberto Elfes and Charles R. Weisbin
Jet Propulsion Laboratory, California Institute of Technology
MS 171-264
4800 Oak Grove Drive
Pasadena, CA 91109
Virgil.Adumitroaie@jpl.nasa.gov

Abstract

This paper describes an approach for qualifying optimal technology portfolios obtained with a multi-attribute decision support system. The goal is twofold: to gauge the degree of confidence in the optimal solution and to provide the decision-maker with an array of viable selection alternatives, which take into account input uncertainties and possibly satisfy non-technical constraints. The analysis is presented in the context of the assessment of capability development portfolio for the NASA Aeronautics program. The results underscore the importance and the usefulness of the postoptimality study in augmenting the level of confidence in the technology portfolio recommendations.

Introduction

Owing to the increased need for consistent, transparent and auditable decision-making processes and tools (Silberglitt and Sherry 2002), our team is developing and utilizing START (STrategic Assessment of Risk and Technology), a quantitative multi-attribute decision support system (Weisbin et al. 2004, Weisbin et al. 2005, Elfes et al. 2006), to perform prioritization of advanced technology portfolios. Project investments are selected through optimization of net mission value as a function of capability level achieved, subject to cost and time constraints. The underlying data set, which quantitatively characterizes requirements (performance, cost, schedule, risk) and proposed technological solutions (achievable capabilities, resource requirements, degree of maturity, schedule), is replete with uncertainty. This inherent uncertainty of the input data must be combined into a global confidence range, which provides the decision maker with an overall sense of quality and likelihood of success of the investment strategy.

We use two complementary methods to take a first step in evaluating the degree of confidence about the standard optimal investment portfolio and determining how the choice of capabilities is affected by variations in the information provided by the capability developers: parametric sensitivity analysis and k-best sets analysis.

The parametric sensitivity analysis reveals whether a given uncertainty in a cost or expected utility might lead to a portfolio recommendation differing from the initial portfolio, and ultimately allows to categorize capabilities as “robustly chosen”, “robustly rejected”, or “trade candidates” (i.e., capabilities that were chosen or rejected with significant uncertainty). In addition to the parametric screening, a k-best analysis is performed to identify competitive portfolios and their common set of capabilities. This common set is in turn compared to the set of robustly chosen capabilities, while the k-best portfolios are presented as options to the optimal recommendation.

The application of the postoptimal analysis presented here originates from a study conducted for NASA's Aeronautics Research Mission Directorate (ARMD). The United States has set a goal of enabling a Next Generation Air Transportation System (NGATS 2004) by the year 2025 to provide for substantially increased capacity while improving or keeping constant any harmful effects on the environment (emissions, noise), safety, and security. The Joint Program Development Office facilitates the multi-agency support of this effort. NASA contributes primarily as an R&D provider of enhanced capabilities, and its Aeronautics Research Mission Directorate (ARMD) has initiated an activity to formulate and assess the return on investment (ROI) for candidate capability-development tasks deemed necessary for the realization of the new system.

Three scenarios were identified as potential elements of an overall architecture to address the country's air transportation needs during the next several decades: (1) Linear extrapolation of today's capabilities; (2) More large regional airports with more large airplanes using them; (3) A highly decentralized system in which considerable traffic is handled by small planes travelling directly point-to-point.

Comparing the relative merits of these three approaches was not among the study's objectives; our analysis included recommendations for capability investments (consistent with the data made available) for each of the three scenarios. Candidate capability areas were derived from programs in vehicle systems development, airspace control, safety, and security. A total of 38 capabilities were specified and quantified in terms of state-of-the-art vs. required performance and maturity; system-level importance; estimated cost; time required for development; and uncertainties in meeting the technical performance objective (assuming full funding) and associated acceptance. The portfolio analysis targeted identifying the best set of capabilities that would support the implementation of desirable future scenarios that contribute to the high-level Joint Planning and Development Office (JPDO) goals, subject to performance requirements, and budget and development time constraints. Further details of this study are found in (Weisbin et al. 2005, Manvi et al. 2005).

Methodology

The starting point in this analysis is the optimal portfolio for a given investment budget level. For each capability, the capability utility, probability of development success, and the probability of acceptance are combined to compute an overall expected utility of the capability (Elfes et al. 2006). The expected utility, together with the development cost of the capability, are the key quantities used in computing an optimal portfolio. The optimal portfolio selection problem is to determine the set of capabilities that provide the maximum composite value while fitting within the available budget. In the START decision support system the solution is obtained by employing a knapsack algorithm (Martello and Toth 1990).

Given the preponderance of input uncertainties and political constraints the optimal solution in itself is not very useful without other qualifying information. Generally, the decision maker needs to know about the robustness of the optimal solution and if there are alternative selections close to the optimal point (perhaps satisfying a non-technical preference).

Parametric Screening Method. We employed two approaches to examine the robustness of our results. First, we changed incrementally the cost and utility, one at a time, for each capability until a change in the resulting portfolio was observed with respect to the nominal solution. This approach yielded the range within which the portfolio selections would be indifferent to a change in the specific value of a particular cost or utility. In other words, it revealed whether a given uncertainty in a cost or utility might lead to a portfolio different from the one computed as optimal. In Figure 1, for example, the green bars represent the range in which the expected utility can vary for selected capabilities without triggering a change in this particular optimal portfolio. The red bars denote the indifference range for the remaining (non-selected) capabilities. The capabilities are sorted based on the expected utility/cost ratio in descending order.

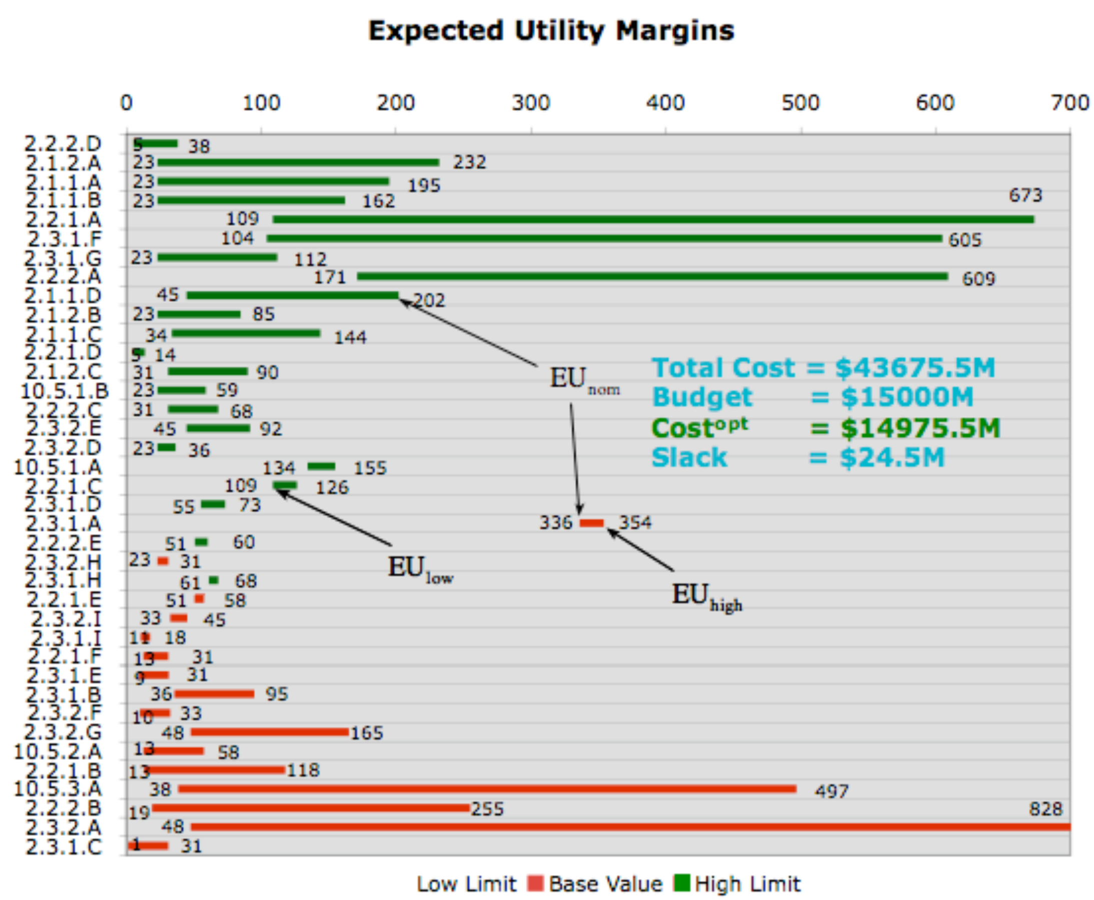

Figure 1. Expected utility ranges for which the given optimal portfolio remains unchanged. The green bars indicate selected capabilities, and the red bars indicate those that were not selected.

If the expected utility for a selected capability were reduced below the lowest value in the range, the capability would be rejected, possibly making room in the budget for the selection of one of the currently unselected capabilities represented by a red bar. Similarly, if the expected utility of an unselected capability is increased beyond the limits of its indifference range it would become selected, possibly knocking one or more of the previously selected capabilities out of the portfolio. Changes in cost can be even more unpredictable. For example, raising the cost of a selected capability beyond the limits of its range could cause it to become unselected—or it could be retained at the sacrifice of a different, less-valuable capability.

The above-mentioned procedure not only produces the indifference ranges for each capability, but also their individual tendencies to remain, enter or exit the portfolio. This information is the result of tracking and cumulating the observed changes in the portfolio at the edge of the indifference range. Figure 2 depicts the cumulated tendencies for each capability during the parametric screening on the expected utility. The negative numbers represent exits from the original portfolio composition, while the positive numbers reflect entries.

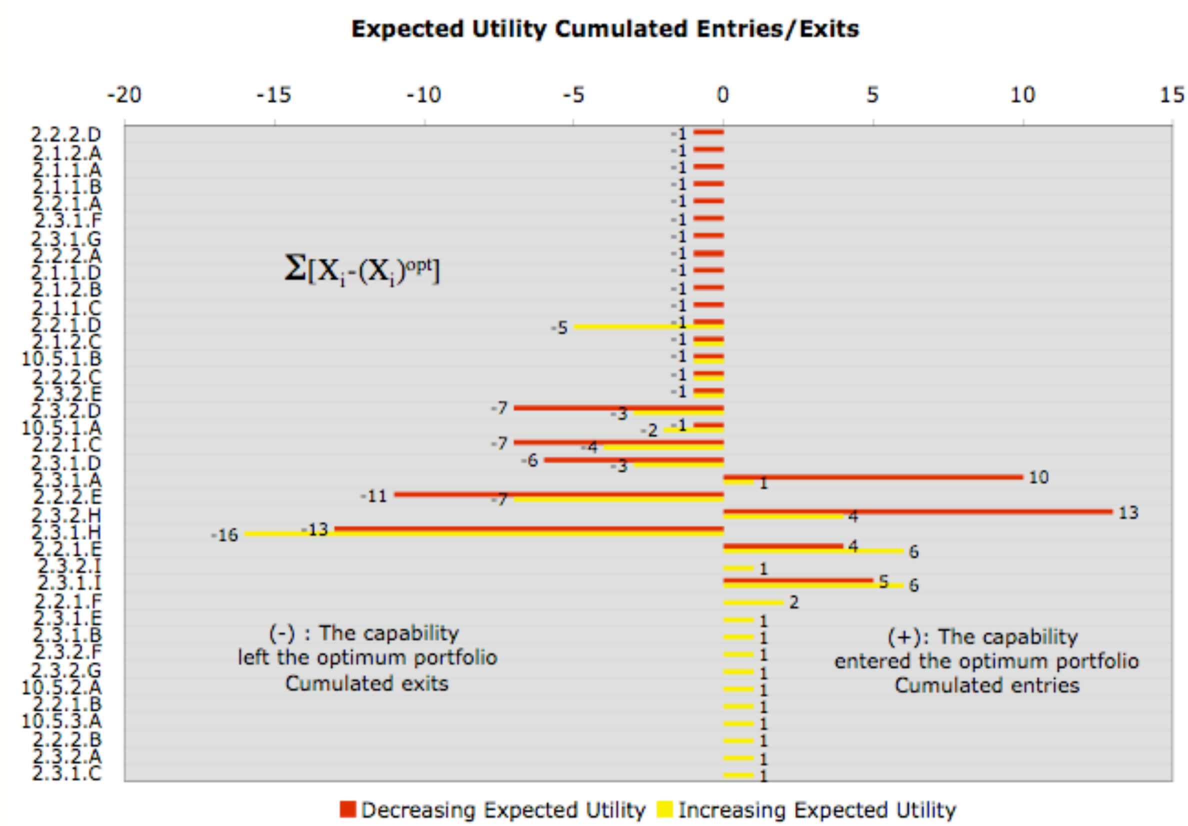

Figure 2. Individual tendencies of each capability during the parametric sweep. The red/yellow bars denote the cumulated moves while decreasing/increasing the expected utility, respectively.

Note that the wider bars represent capabilities that entered or exited the portfolio more frequently in this study. Such behaviour characterizes the marginal groups, whose performance-cost ratios made them expendable or marginally acceptable, subject to the vagaries of their own cost and performance expectations and those of the other capability groups. Although a univariate analysis such as this represents an "ideal case" in which only one parameter is uncertain, it does provide the decision maker with essential information regarding their technology portfolio. For example, some capabilities would require at least a doubling of their expected utility to get selected and consequently are definitely not a contender.

Monte Carlo Analysis. In addition to this procedure, which dealt with the effect on a portfolio of only one variation in only one capability group at a time, we also performed a Monte Carlo simulation in which variations were applied to all capability groups simultaneously. In this study, the portfolio optimizations were run 1000 times with the cost and expected utility of each capability group varied randomly each time up to a 10% increase or decrease relative to its initially assigned value. Then an additional 1000 runs were performed with variations up to 25%.

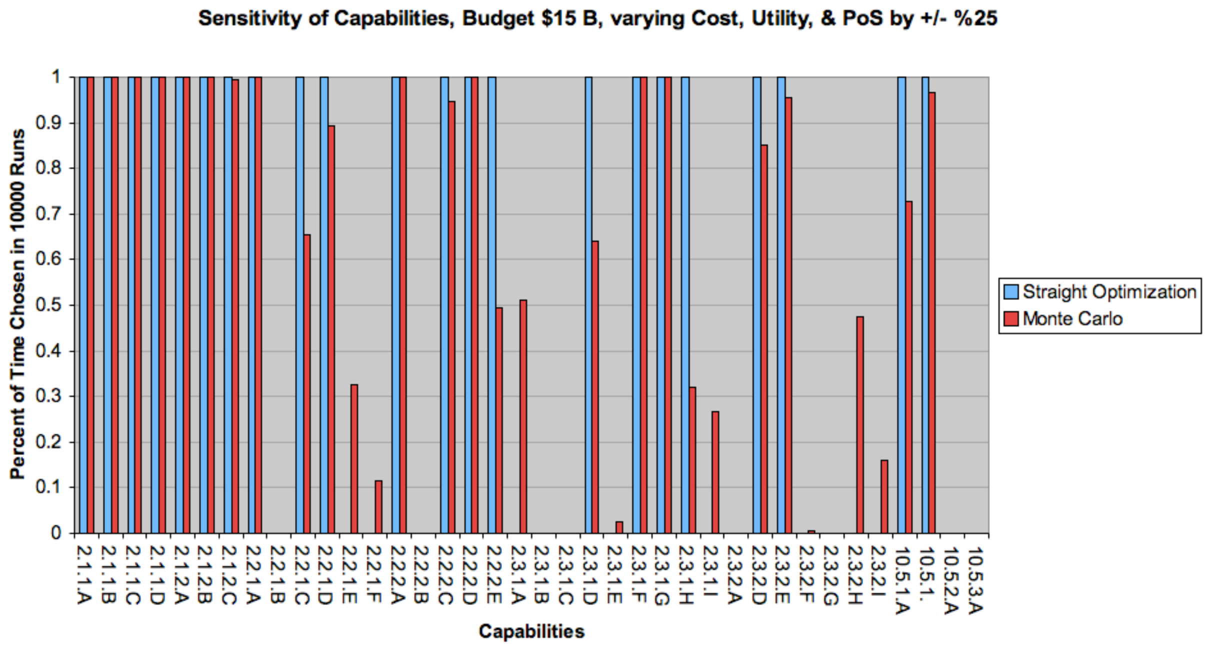

Figure 3. Selection frequency for each capability in the Monte Carlo analysis. The red bars indicate the percentage selection under varying parameters. The blue bars represent the nominal optimization.

Table 1: Recommended Portfolio Composition for the Next Generation Air Transportation System at a budget level of $15B.

Robust Selection	
Less than 10 exits for a selected capability in the deterministic analysis. *Greater than 85% selection record for a selected capability in the Monte Carlo.*	
2.1.1.A	Protect/Prevent Abnormal Operations & System Failures
2.1.1.B	Detect & Mitigate Natural Hazards
2.1.1.C	Prevent Breakdown of Human/Machine Interface
2.1.1.D	Integrity & Efficiency of Accepting Advanced Software Systems
2.1.2.A	Detect & Inform Potential System Vulnerabilities
2.1.2.B	Mitigate Consequences from Intentional Attack
2.1.2.C	Detect & Contain Diseases & Bio/Chem Agents
2.2.1.A	Low emission subsonic vehicles
2.2.1.D	Low emission personal air vehicles
2.2.2.A	Low noise subsonic vehicles
2.2.2.C	Low noise ESTOL vehicles
2.2.2.D	Low noise personal air vehicles
2.3.1.F	Increase Arrival/Landing Rates at Commercial Airports
2.3.1.G	Commercial Operations from Small/Underused Airport
2.3.2.E	Efficient all-weather rotorcraft
10.5.1.B	Conduct Routine UAV in NAS

Not Recommended	
Less than 10 entries for a non-selected capability in the deterministic analysis. *Less than 15% selection record for a non-selected capability in the Monte Carlo*	
2.2.1.B	Low emission supersonic vehicles
2.2.1.F	Low emission UAVs
2.2.2.B	Low noise supersonic vehicles
2.3.1.B	General Aviation During Peak Demand
2.3.1.C	Public Service Aircraft During Peak Demand
2.3.1.E	Globally Harmonized Equipage & Operations
2.3.2.A	Efficient subsonic vehicles
2.3.2.F	Complete Decision Information to All in NAS
2.3.2.G	Low Cost Vehicles for Bulk Cargo
2.3.2.I	Minimum Impediments of Mode Change
10.5.2.A	Extended Autonomous Flight in Mars Atmosphere
10.5.3.A	Incorporating Hypersonic Air-Breathing Propulsion

Trade Candidate	
2.2.1.C	Low emission ESTOL vehicles
2.2.1.E	Low emission rotorcraft
2.2.2.E	Low noise rotorcraft
2.3.1.A	Capacity En-Route Commercial Operations in NAS
2.3.1.D	Minimize System-Wide Disruptions
2.3.1.H	Commercial Operations with short/no Runways
2.3.1.I	Incorporate Full Spectrum of Aircraft to NAS
2.3.2.D	Efficient easy-to-operate personal air vehicles
2.3.2.H	Increased Speed & Range for Pedestrian Travel
10.5.1.A	Autonomous high altitude long-endurance flight

The status (in or out) of each capability is accumulated from each run such that a selection frequency is computed from this stage of the parametric screening. Figure 3 shows the selection frequency chart for the Monte Carlo runs with +/-25% variations (with capabilities sorted alphabetically).

The results from the two parametric approaches are mutually calibrated in order to issue a common categorization of the projects sets as “robustly selected”, “robustly rejected”, and “trade candidates”. For example, we found that in this study that “robust selection” translates into “less than 10 exits for a selected capability” in the deterministic analysis and “greater than 85% selection record” in the Monte Carlo analysis. Illustrative results of this procedure applied to the NGATS 2004 system are shown in Table 1.

Furthermore, these results can be summarized in a graph referred to as a “frontier plot”, where the capabilities are represented in the phase space (expected cost vs. expected utility) based on their selection and sensitivity state. The symbolic frontier (green dotted) line separates the selected set from the unselected set and the closed two curves, green and red, identify the robustly selected set and the robustly rejected capabilities, respectively. For a given set of capabilities the frontier moves down to the right with increasing budget.

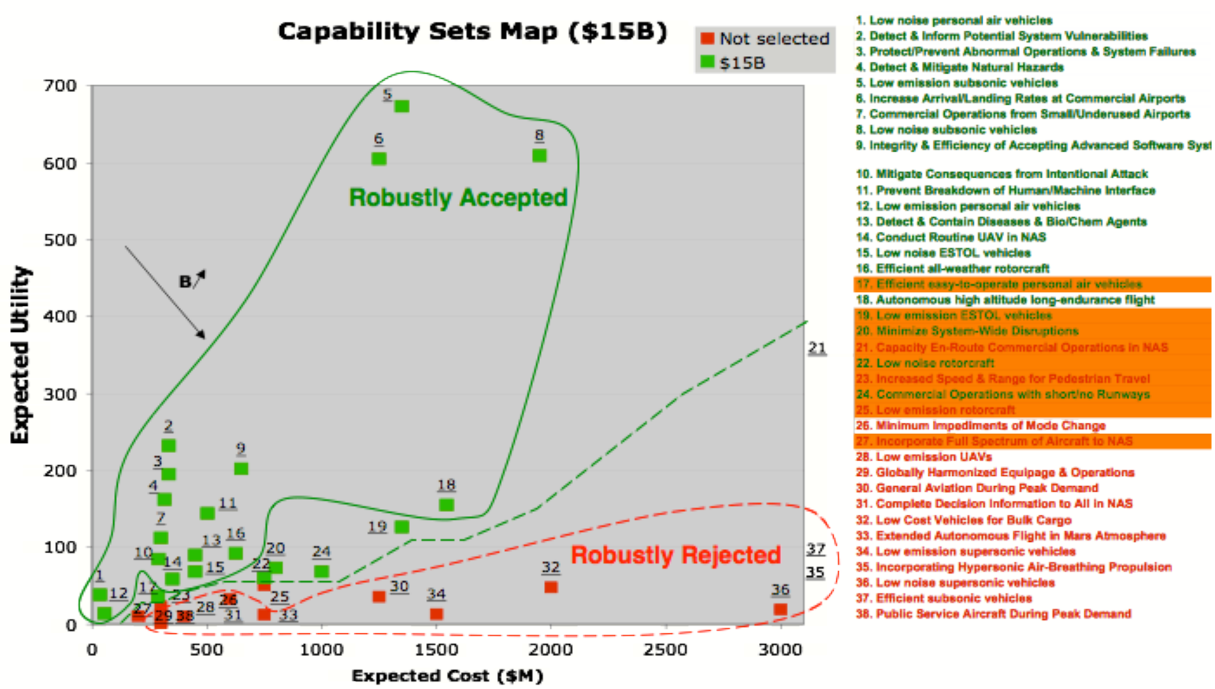

Figure 4. Optimal portfolio for a total investment of $15B. The capabilities shown in green on the left plot and on the table to the right are those that have been robustly chosen. The capabilities in red have been robustly rejected. The remaining capabilities highlighted in orange have been accepted or rejected with higher uncertainty, and could be subject to further assessment by the decision-maker.

Figure 4 depicts the “frontier plot” for the case under consideration. The stable portfolio components (selected or not) are farther from the frontier than the trade candidates. This chart can be used not only as visual aid in the decision making process, but also for verification of the optimization runs. Furthermore, the trade candidates could be subject to further scrutiny by the decision-maker.

K-best Analysis. The k-best sets analysis (Guikema and Milke, 2003) offers the "k" suboptimal portfolios closest to the optimal recommendation for a given budget level. Based on the k-best sets the decision-maker can take into account aspects of the problem that are not easily modelled quantitatively, as well as additional constraints important to the decision.

When finding the k-best sets with the base case input parameters and then comparing the values of these sets over the entire range of possible values for the input parameters, competitor portfolios can be proposed. The intersection of the k-best portfolios with the optimal portfolio produces a set of project selections deemed as "persistent."

Figure 5 shows the relative positioning of the five closest competitive portfolios with respect to the optimal recommendation in an aggregated expected utility/total cost mapping. From the placement and composition of the suboptimal portfolios the decision-maker can fulfil supplementary requirements. For example, if the extra constraint is to spend most of the available budget, KB3 is the close to the optimal portfolio, but in addition it minimizes the budget slack.

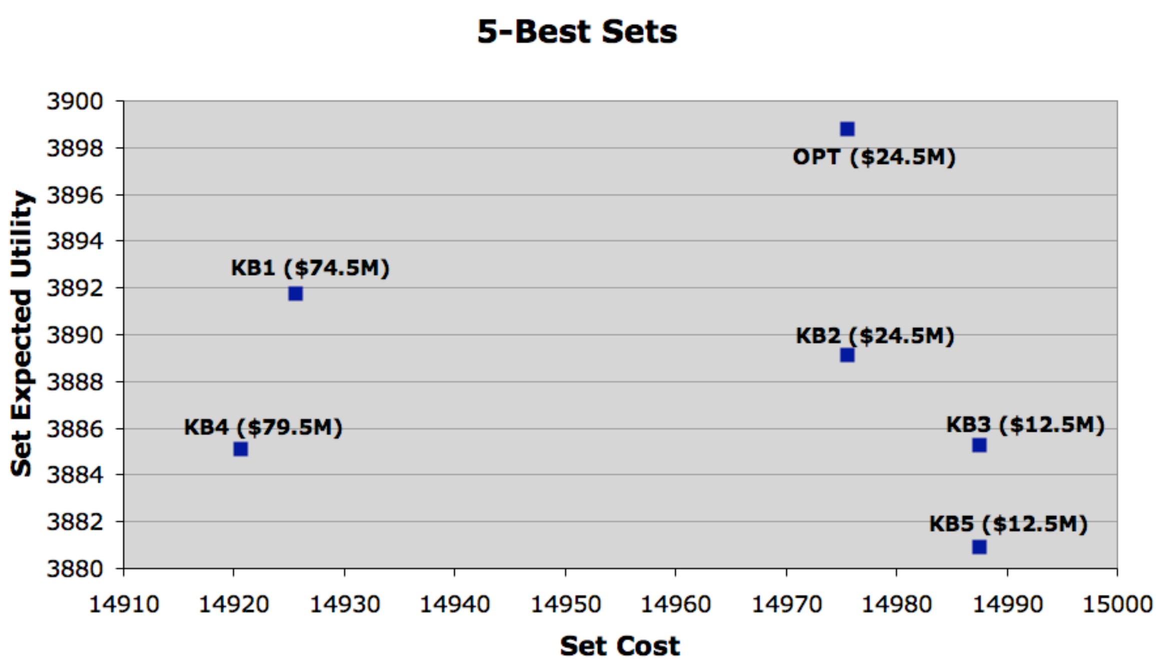

Figure 5. "5-best" portfolios mapping in the aggregated expected utility/ total cost. The figures in parenthesis denote the budget slack.

Table 2 identifies the actual 5-best portfolios and the categorization of the capabilities by their overall percent presence in the suboptimal portfolios (including the "persistent" set displayed in green colour). The colouring convention is similar to the one utilized in the parametric screening analysis. The colour green denotes the stable set, while the orange cells represent the trade candidates. One final observation can be made at this point: the parametric sensitivity analysis and the k-best analysis generate consistent choices of "robust" and "persistent" recommendations. With two exceptions the robust recommendations from the sensitivity analysis are the same as the ones suggested by the k-best analysis. This information can be used to increase the level of confidence in the decision-making process and to provide valuable insights and choices to the results of the optimality analysis.

Table 2: Identification of the common set of capabilities in the "5-best" portfolios and their percent overall presence

Capability	Opt	KB1	KB2	KB3	KB4	KB5	Overall
2.2.2.D	1	1	1	1	1	1	100.00%
2.1.2.A	1	1	1	1	1	1	100.00%
2.1.1.A	1	1	1	1	1	1	100.00%
2.1.1.B	1	1	1	1	1	1	100.00%
2.2.1.A	1	1	1	1	1	1	100.00%
2.3.1.F	1	1	1	1	1	1	100.00%
2.3.1.G	1	1	1	1	1	1	100.00%
2.2.2.A	1	1	1	1	1	1	100.00%
2.1.1.D	1	1	1	1	1	1	100.00%
2.1.2.B	1	1	1	1	1	1	100.00%
2.1.1.C	1	1	1	1	1	1	100.00%
2.2.1.D	1	1	1	1	0	1	83.33%
2.1.2.C	1	1	1	1	1	1	100.00%
10.5.1.B	1	1	1	1	1	1	100.00%
2.2.2.C	1	1	1	1	1	1	100.00%
2.3.2.E	1	1	1	1	1	1	100.00%
2.3.2.D	1	1	1	0	1	0	66.67%
10.5.1.A	1	1	1	1	1	1	100.00%
2.2.1.C	1	1	1	1	1	0	83.33%
2.3.1.D	1	1	1	1	1	0	83.33%
2.3.1.A	0	0	0	0	0	1	16.67%
2.2.2.E	1	1	0	1	1	0	66.67%
2.3.2.H	0	0	0	1	0	0	16.67%
2.3.1.H	1	0	1	1	1	0	66.67%
2.2.1.E	0	1	1	0	0	0	33.33%
2.3.2.I	0	0	0	0	0	0	0.00%
2.3.1.I	0	1	0	0	0	1	33.33%
2.2.1.F	0	0	0	0	0	0	0.00%
2.3.1.E	0	0	0	0	0	0	0.00%
2.3.1.B	0	0	0	0	0	0	0.00%
2.3.2.F	0	0	0	0	0	0	0.00%
2.3.2.G	0	0	0	0	0	0	0.00%
10.5.2.A	0	0	0	0	0	0	0.00%
2.2.1.B	0	0	0	0	0	0	0.00%
10.5.3.A	0	0	0	0	0	0	0.00%
2.2.2.B	0	0	0	0	0	0	0.00%
2.3.2.A	0	0	0	0	0	0	0.00%
2.3.1.C	0	0	0	0	0	0	0.00%

Conclusions

We presented an approach, based on two complementary methods - parametric sensitivity analysis and k-best sets analysis, for qualifying optimal technology portfolios. The parametric sensitivity analysis relies on two types of evaluation procedures: deterministic and statistical (Monte Carlo).

The deterministic sampling yields the range within which the portfolio selections are invariant to changing cost or for the given budget. The statistical sampling expands the search domain with consideration of joint variation in capability input parameters. The change events are recorded and accumulated over the two parametric samplings. By performing a mutual calibration between the accumulated activities, the sets of projects "robustly selected", "robustly rejected", and "trade candidates" are identified.

The k-best sets analysis offers the "k" suboptimal portfolios closest to the optimal recommendation for a given budget level. The intersection of the k-best portfolios with the optimal portfolio produces a set of project selections deemed as "persistent." Although the two above approaches are complementary, their results are consistent, in that the "persistent" set is similar in composition to the "robust" set.

The goal of the postoptimality study is to enhance and improve the decision-making process by providing additional qualifications and substitutes to the optimal solution. The methodology proposed here is demonstrated on a NASA technology project selection. The results highlight the importance and the usefulness of the postoptimality analysis in providing a higher level of confidence to the technology portfolio recommendations.

Acknowledgments

The research described here was conducted at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration and would not have been possible without the encouragement of a number of outstanding and technically knowledgeable sponsors. We want to thank those particularly involved in the case studies presented in this paper, including Douglas A. Craig, Directorate Integration Office, NASA HQ; Jay Falker, Technology Assessment Lead, Office of Program Analysis and Evaluation, NASA HQ; Yuri O. Gawdiak, Program Integration Manager, Aeronautics Research Mission Directorate, NASA HQ; Terrence J. Hertz, Director, Aeronautics Technology, Aeronautics Research Mission Directorate, NASA HQ; and George Price, Special Assistant, Aeronautics Research Mission Directorate, NASA HQ. We would also like to express our appreciation for the support of our Program Office led by Stephen L. Prusha, Program Manager, Strategic Systems Technology Program Office, JPL/ California Institute of Technology.

References

Chiulli, R. M. , "Quantitative Analysis: An Introduction". CRC Press, 1998.

Elfes, A., Weisbin, C.R., Manvi, R., Adumitroaie, V., Lincoln, W.P. and Shelton, K., "Extending the START Framework: Computation of Optimal Capability Development Portfolios Using a Decision Theory Approach," Journal of the International Council on Systems Engineering, submitted for publication 2006.

Guikema, S. D. and Milke, M. W., “Sensitivity Analysis for Multi-Attribute Project Selection Problems,” Civil Engineering and Environmental Systems, Vol. 20, No. 3, pp. 143-162, 2003.

Manvi, R., Weisbin, C.R., Elfes, A. and Rodriguez, G., "Risk-Driven Interdependent Task Analysis Applied to a NASA R&D Program", IEA/AIE 2005 18th International Conference on Industrial & Engineering Applications (IEA) of Artificial Intelligence & Expert Systems (AIE), Bari, Italy, June 22-25, 2005. Published in the Springer-Verlag Lecture Notes on Artificial Intelligence series, Vol. 3533, 2005.

Martello, S. and Toth, P., Knapsack Problems - Algorithms and Computer Implementations, Wiley, 1990.

NGATS 2004, “Next Generation Air Transportation System”, Joint Program Development Office, December 2004. http://www.jpdo.aero.

Silberglitt, R. and Sherry, L.,”A Decision Framework for Prioritizing Industrial Materials Research and Development,” RAND Monograph Report MR-1558-NREL, 2002.

Weisbin, C.R., Rodriguez, G., Elfes, A. and Smith, J.H., "Toward a Systematic Approach for Selection of NASA Technology Portfolios," Systems Engineering Journal, Vol. 7, No. 4, pp. 285-302, 2004.

Weisbin, C.R., Manvi, R., Shelton, K., Adumitroaie, V., Derleth, J.E. and Elfes, A., “Capability-Development Return on Investment for the NASA Aeronautics Program”, 16th International Conference on Systems Engineering (ICSEng 2005), August 16-18, 2005, Las Vegas, NV.

Biography

Virgil Adumitroaie received his Ph.D. degree in Mechanical Engineering from the State University of New York at Buffalo in 1997. He served as a Senior Engineer at the CFD Research Corporation, Huntsville, AL prior to joining the Jet Propulsion Laboratory as a Senior Member of Technical Staff in 2004. Currently, his work in optimization methods and uncertainty analysis applied to decision problems supports tasks from the JPL's Strategic Technologies Program Office.

Kacie Shelton graduated from the California Institute of Technology in June 2000 with a B.S. in Physics. She has worked on the OnEarth project of making a web-based mosaic of Landsat images of the world at the PODAAC, and assisted in developing the Cassini image-processing pipeline. She now is a member of the START Team in technology assessment.

Alberto Elfes received his Ph.D. degree in Electrical and Computer Engineering, with concentration in Robotics, from Carnegie-Mellon University in 1989. Alberto Elfes is presently a Principal Member of Technical Staff at the Jet Propulsion Laboratory, where he leads projects in the areas of complex systems analysis, strategic technology and mission portfolio planning and design, autonomous airships, and supervisory control of teams of multiple heterogeneous robots.

Charles R. Weisbin, Eng. Sc. D. 1969 Columbia University. Charles R. Weisbin currently serves as Deputy Program Manager for the Strategic Systems Technology Office of the Chief Technologist at the Jet Propulsion Laboratory. He was the co-chairman of the NASA Telerobotics Intercenter Working Group for seven years and received the 1993 NASA Exceptional Service Medal for formulation and development of the NASA Telerobotics Program. He also received the award for Outstanding Leadership in Surface Robotics for the year 2000, a Lifetime Achievement Award from the World Automation Congress in 2004, and the award for Best Paper published in the System Engineering Journal in 2004.