

Planning a Large-Scale Progression of R&D – a Pilot Study in the Aerospace Domain

Andrew A. Shapiro
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91009
818 393-7311
Andrew.A.Shapiro@jpl.nasa.gov

Steven L. Cornford
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91009
818 354-1701
Steven.L.Cornford@jpl.nasa.gov

Martin S. Feather
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91009
818 354-1194
Martin.S.Feather@jpl.nasa.gov

George Price
NASA Headquarters 300 E ST SW
Washington DC 20546-0001
Mail Suite: 6J39-B
202 358-1223
george.price@nasa.gov

Yuri O. Gawdiak
NASA Headquarters 300 E ST SW
Washington DC 20546-0001
Mail Suite: 4K70-A
202 358-1853
yuri.o.gawdiak@nasa.gov

Wendell R. Ricks
NASA Langley Research Center
Hampton, VA 23681-2199
Mail Stop 449
757 864-6733
Wendell.R.Ricks@nasa.gov

Abstract—^{1,2} A prior case study reported in “Programmatic Risk Balancing” [D.M. Tralli, IEEE Aerospace Conference 200] established the suitability of adopting a lifecycle risk management decision-support tool to the planning of application projects across NASA’s Earth Science Enterprise. Here we report on a pilot study to gauge the suitability of this same approach for large-scale planning of a progression of research and development efforts in the Aerospace domain. The purpose of the study was to assess feasibility and utility, and to prototype adaptations to the approach as and when such adaptations were found to be needed.

The novel challenges posed by this domain included: scale – the overall goal is to plan and monitor \$1.5B worth of R&D spread over several years; scope – information spans task-level, project-level and program-level concerns; distributed expertise – the information on which to base decisions requires combining inputs from multiple geographically dispersed, busy people (i.e., they won’t be available to all meet concurrently, even via a teleconference); and novel problem domain aspects – for example the world continues to evolve as the multi-year R&D efforts take place, so that what might be desirable solutions to aim for this year may be rendered obsolete and unnecessary a few years hence as other capabilities mature, or alternately, may continue to be necessary but less sufficient as demands increase.

The net result was promising: the approach worked, and a number of interesting observations could indeed be drawn from the accumulated information. Overall it also pointed to

several possible avenues to scale-up the approach, together with some remaining key problems.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. BACKGROUND	2
3. PILOT STUDY BASELINE	5
4. APPROACH.....	5
5. ANALYSIS AND OBSERVATIONS	8
6. CONCLUSIONS	11
ACKNOWLEDGEMENTS	11
REFERENCES	11
BIOGRAPHIES	12

1. INTRODUCTION

The purpose of this pilot study was to determine if the use of Defect Detection and Prevention (DDP), a tool developed by Steve Cornford and Martin Feather for NASA at JPL, initially intended for lifecycle risk management decision support analysis [1,2,3], can be used for the intelligent management of an R&D portfolio. In this case, we consider the evaluation of a number of ongoing research tasks and their weighted relevance to meeting the objectives of a narrowly defined set of goals. The purpose of this project was to demonstrate and evaluate the potential of the Defect Detection and Prevention (DDP) method to measure progress toward research objectives, specifically toward system-level capabilities enabled by NASA aeronautical technology projects.

The pilot study took as a representative overall objective of the Aerospace domain that to increase by 50% the

¹  © 2005 IEEE

² IEEEAC paper #001, Version 0, Updated September 25, 2005

throughput of the US airspace over the next several years, with no diminution in safety levels. Gauging feasibility included determinations of what inputs would be needed, how they might be gathered from multiple experts and thereafter reconciled, and how might they be represented within the risk-centric model. Gauging utility included application of the decision-support tool's capabilities to reveal, via a variety of graphical presentations, the results of its various computations over the accumulated dataset; the insights gleaned from doing so would be indicative of the potential value this approach would have in the overall R&D planning efforts to come.

The approach followed was to gather representative data from a pair of domain experts, doing so in a deliberately loosely coordinated manner to mirror (on a small scale) the anticipated challenges of information gathering. Preliminary capabilities of the decision-support tool were exercised to assist in identification and reconciliation of areas of contention in the separately collected but overlapping

datasets. Also stressed was the use of the decision-support tool's capability to represent distinct stages of activities, in this case to capture the multiple years of R&D efforts, and simultaneously the increasing challenges posed by the evolution of the "status quo" as airspace throughput demands are anticipated to increase over several years, and the influences of other, non R&D, factors (e.g., sociological considerations) in this same domain.

2. BACKGROUND

DDP was developed by JPL as a tool for life-cycle risk management. It comprises software and methods for eliciting the necessary input and status data. Applications to date have focused on assessing and optimizing risk for projects at various levels. Figure 1 shows the history of the development of the DDP tool, as well as important case studies.

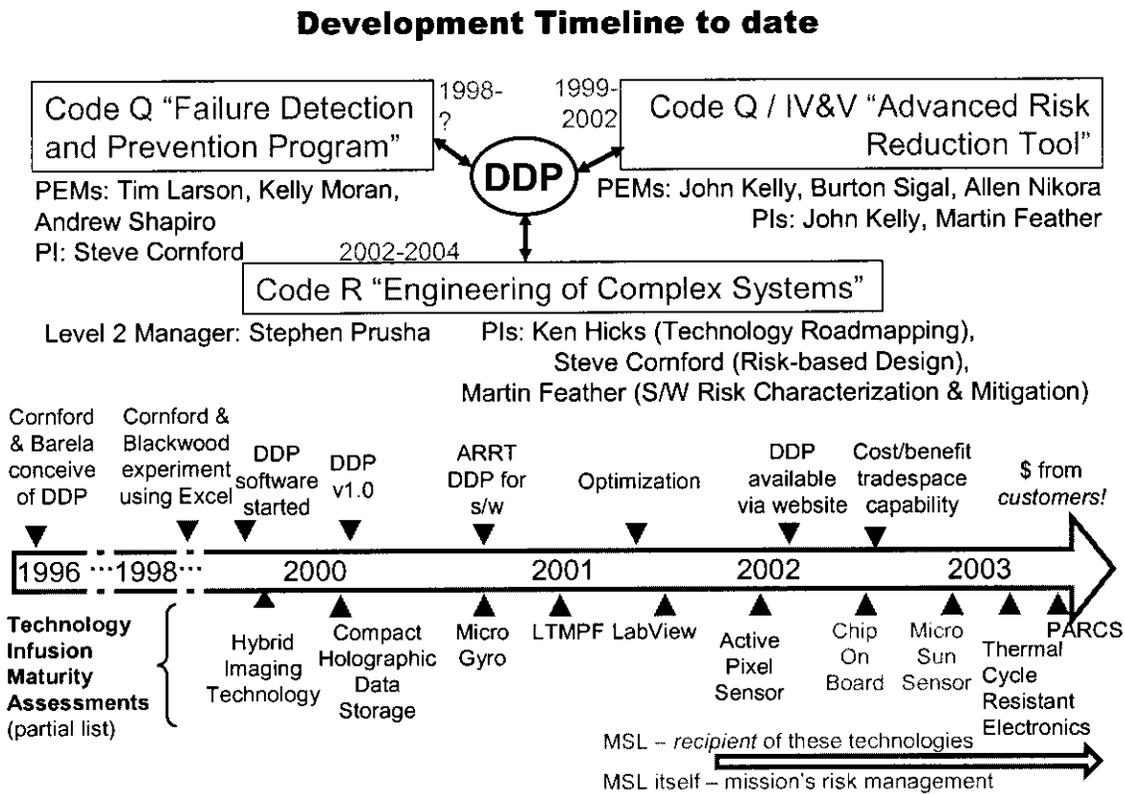


Figure 1. An historical timeline of the development of the DDP tool.

The current effort is an attempt to extend this method to measuring progress toward research objectives. Measuring such progress for aeronautical technologies has been a difficult task for several reasons: the relationship between

research results and NASA objectives for improving the air transportation system is complex and often not explicit; research projects can represent alternative, parallel, or dependent approaches to the desired outcomes; and measuring progress of research tasks is itself a challenging

activity. DDP has the potential to address these issues largely because DDP is a risk-centric method. Since risk is expressed in terms of probabilities, the mathematics involved in summing the results of alternative, parallel or dependent activities are relatively straightforward.

NASA's Aeronautics Research Mission Directorate was also interested in use of the DDP approach during the planning phase. DDP not only provides information about risks, it also affords a method of optimizing portfolio investments by planning for maximum risk abatement for a fixed budget or, alternatively, determining the minimum budget and portfolio of activities required to achieve a specified level of risk. This application would be iterative and dynamic, updating the best investment strategy as a program progresses, anticipated risks are abated, and new risks are identified. DDP also enables program activities to be considered as real options, focusing on investment strategies to develop information that will enable optimal downstream decisions. To avoid assessments based on perceived, rather than real, changes in risk as a program progresses, these applications imply a need for an objective, relatively detailed, explicit method of evaluating risk. DDP is such a method.

Using a diverse team of selected experts, the methodology of defect detection and prevention has three major steps. The first is a listing of the system objectives portrayed by the initiating team. In an orthogonal row, relevant risks are identified by the team of experts and the initiating team based on their collective experience. These are then ranked in terms of the impact of a particular risk on all of the objectives, see Figure 2. The sigma on the top represents the sum of an individual risk's effect across all of the objectives and the sum at the bottom is a sum downward of all risks on a particular objective.

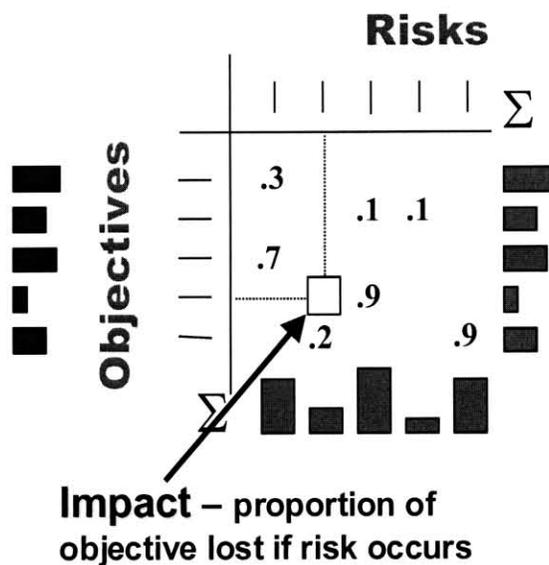


Figure 2. The first DDP matrix, Objectives vs Risks and their impact.

Second, in a separate matrix, a list of mitigations strategies are generated and their effects scored against all of the identified risks, shown in Figure 3. The pi represents the product of the risk reductions delivered by each mitigation. The risk product is the effect of a particular mitigation on all risks.

Third, costs are associated with mitigations and trades are performed to select combinations of mitigations are selected to cost effectively reduce risk. Tradeoffs may be made manually, using the expertise of the team. Alternatively, artificial intelligence (AI) techniques may be used, such as simulated annealing, to give graphic representations of analysis. This technique can give a clear representation of design/cost trades that will allow the project portfolio management to make informed decisions. [2,3] It is the use of this third tool, DDP, which is recommended for this application of enabling portfolio effectiveness.

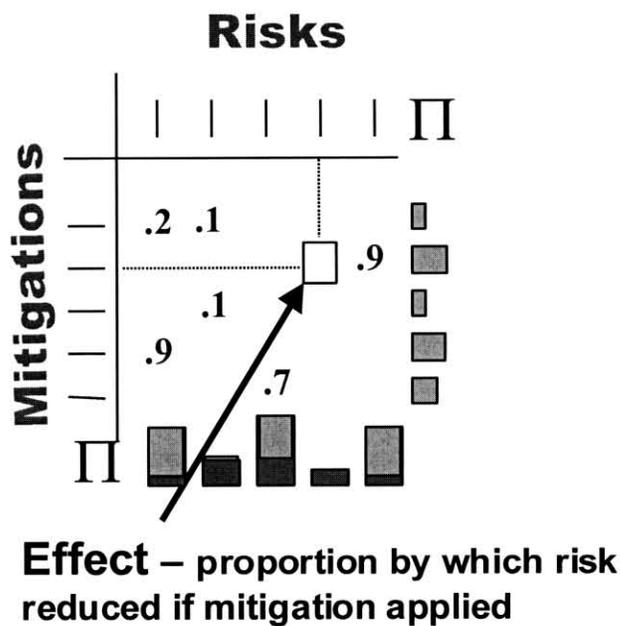


Figure 3. The second DDP matrix, Mitigations vs Risks and their effectiveness.

Figures 4-6 show several examples of DDP output views. The view in Figure 4 is a product of a process called simulated annealing where a number of combinations of mitigations and their costs are examined and plotted for their cumulative effectiveness. In this case it is there is a clear "knee" in the curve. To the left of the knee, effectiveness drops rapidly with small changes in cost. To the right of the knee additional cost provides very little benefit. The implication is that there is a critical cost/mitigation threshold, in this case right around \$1M.

Figure 5 shows another sample view of a DDP output. The upper bar chart (red and purple) shows a sorted list of risks for this particular program (from the risks affecting the most objectives to the ones affecting the least weighted by their potential impact). These risks are also color coded by the type of risk, in this case new technology versus engineering risk. Other possible categories may also be represented in the DDP tool such as schedule risks, etc.

The lower bar chart in Figure 5 shows a scenario with an estimate of what would result if a particular set of mitigations has been implemented (green bars). The resulting plot shows where mitigations have been predicted to be very effective for some tasks (nearly all green) or slightly effective (little green).

Figure 6 shows yet another useful DDP output view. On the left of this figure, a “thumbnail” chart of all of the risks sorted by category (general, new technology, engineering,

and programmatic) and then rank. Just to the right is a second “thumbnail” showing the effectiveness of a suggested first year’s mitigation program in each of the four risk areas.

The larger plot in Figure 6 shows the risks and suggested mitigations for the first year’s program with colors representing schedule, cost, technology and general risks (grey, yellow, orange and red respectively).

Inset on the right of Figure 6 is a version of the standard “traffic light” plot. This plot is usually given as a 5x5 matrix, but in the case of DDP, the granularity of choices is so high, an integrated or continuous contour version of the chart is generated. In this view, the threshold levels for high medium or low risk may be set by the user. Corresponding lines are shown on the bar chart.

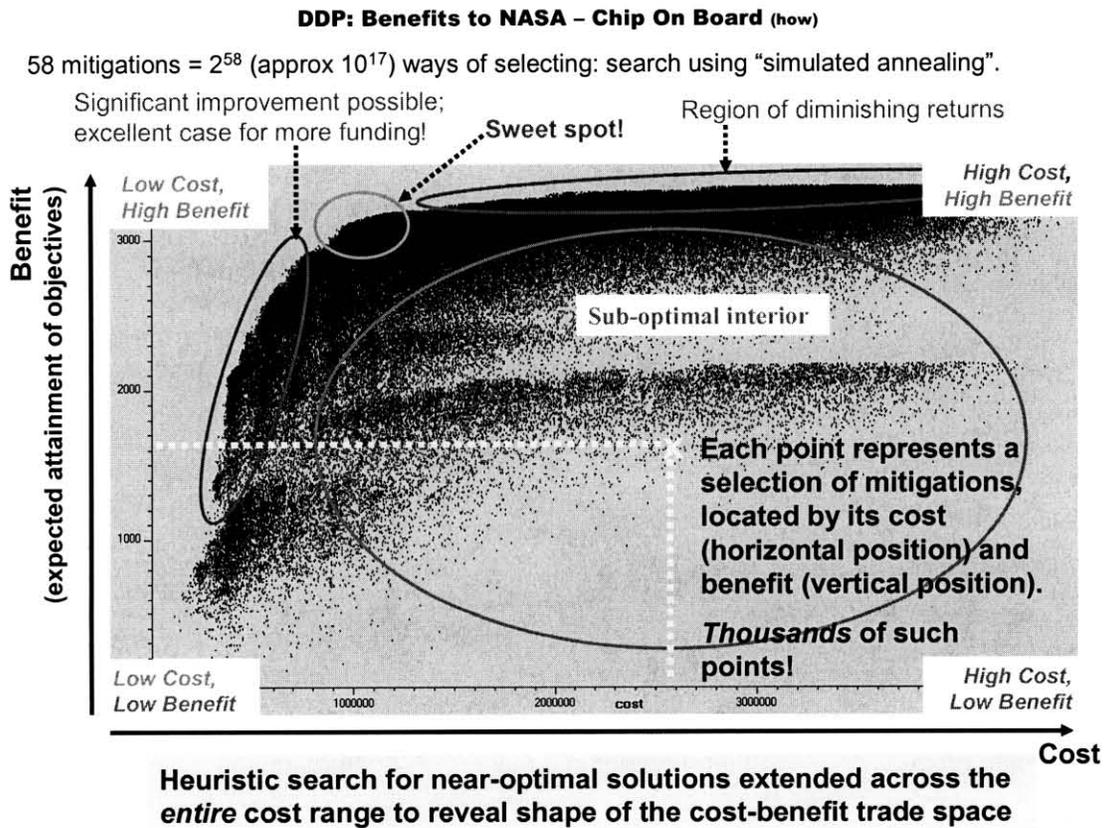


Figure 4. An example of a simulated annealing DDP output. Steep slopes indicate significant added value for small marginal expenditures. Shallow slopes indicate little gain for additional expenditures.

3. PILOT STUDY BASELINE

The baseline of this particular study was provided by our sponsors at NASA. [4] Risks that affect achievement of NASA aeronautics technology goals include technical performance, resource availability, schedule, implementation, and political (i.e., effect of policy decisions). This evaluation will concentrate on technical, cost, and implementation risk.

The project team has selected for evaluation two system-level capabilities included in the Pathfinder³ framework. Each of these capabilities is achieved by successful completion of specific projects, as identified in the framework. The two capabilities and their associated projects are as follows:

Capability 1

Ability, through optimization and integration improvements, to increase peak takeoff and landings per hour by 50% relative to a 1997 baseline (as measured at the 35 FAA benchmark airports).

Rationale: This capability represents an improvement over the previously-established objective for 2004 (enable a 35% increase in aviation system throughput in the terminal area based on 1997 National Airspace System (NAS) capacities).

Associated projects:

- AATT
- Efficient Flight Path Management
- Transformational Operations
- Advanced Information Capabilities
- System Evaluation & Engineering
- Quiet Aircraft Technology

Capability 2

Ability, through vehicle and operation improvements, to routinely service (35 takeoffs or landings per hour) transport aircraft (cruise speed > .8 mach, 2100 mi. range, passengers > 120) on runways < 5000 ft.

³ Pathfinder is a relational database of the elements, assumptions, and substantiations that link technology or project development alternatives, system requirements, strategies, and strategic objectives. The Pathfinder database provides input to the ARMD portfolio and project management analytical framework.

Rationale: this capability is an operating tempo that represents today's capabilities for regular airports and range, speed, and payload for average commercial transports, extended to runway lengths less than 5,000 ft. The runway length was determined from an LMI study that showed that stub runways or real estate could be available for such lengths at about half of major airports and hub airports. It should be understood that this applies to instrument as well as visual meteorological conditions.

Associated projects:

- Transformational Operations
- Integrated Tailored Aerostructures
- Autonomous Robust

4. APPROACH

The team looked at the specific objectives outlined in the previous section and then identified risks to achieving those objectives. The way that risks were identified were two-fold, first we looked at a typical flight profile from clearance delivery to the final unloading of passengers. A cartoon of a typical profile is shown in Figure 7. The second source of risk identification was a series of interviews with NASA experts. Figure 8 shows some of the risks identified in the exercise with some indented hierarchy.

The second step in the process was to apply the lists of ongoing related tasks provided by the aeronautical research team to the DDP matrix. A more detailed description of the process for implementing DDP has been described previously [5]. In this case, unlike most DDP exercises, the mitigations were already in place. For the pilot study, the existing FAA tasks appear as entities that reduce/retire these risks. The tasks may impact many risks (and vice versa) either positively or negatively. Figure 9 shows these lists with portions expanded (the entire list is too long to show conveniently).

The risks, in this case, were looked-at in terms of the likelihood of additional incidents, or a decrease in safety, encountered along the typical flight profile given in Figure 7 while attempting to meet the given objectives. The primary focus was on the many hand-offs and how they might be impacted by the increased volume requirements outlined in the objectives.

We then interviewed experts (with varying degrees of buy-in) and had them review and improved the risk model as well as correlate risks to mitigation tasks. We then scored risks versus mitigations with internal experts as

well as some outside inputs. The scoring was between zero and one with zero having no effect and one having a most significant impact. Figure 10 shows a portion of the scored matrix. These scores were only “representative” and should not be construed as data of sufficient fidelity to make decisions. In a formal DDP exercise, it is best to have all of the experts make the evaluation together or at the very least to have more formal interviews.

and the views of the tool to be demonstrated in a more familiar context, to resolve any disparities, add interesting alternatives and to visualize and interpret the results.

The scores for the pilot study were to allow the features

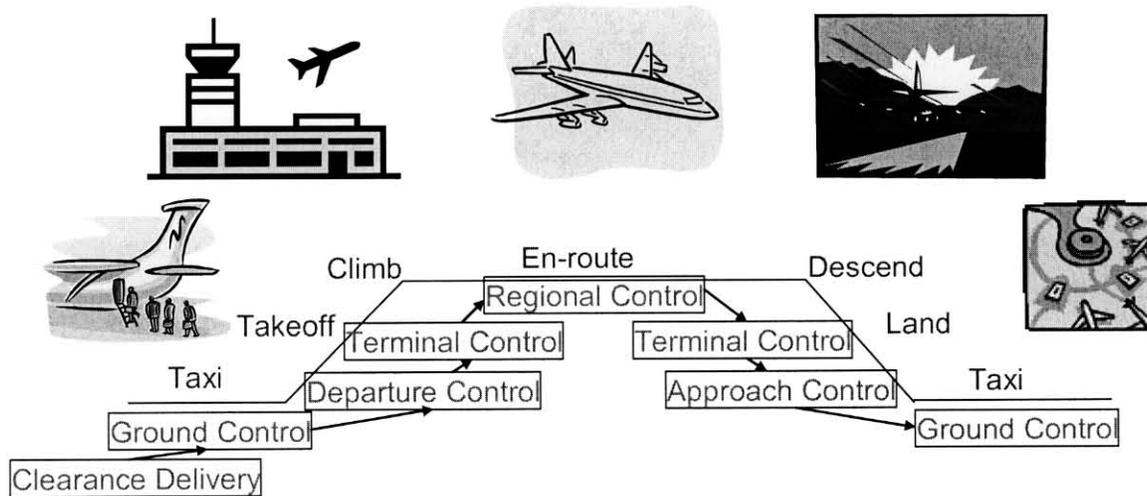


Figure 7. Handoff sequence used for pilot analysis evaluating risk mitigation interactions. [6]

- 1: Insufficient safety
- 2: Failure to achieve capab1 (50% TO/L)
 - 2.1: Due to ground inadeq
 - 2.1.1: taxiways too crowded, not enough taxiways, too many runway crossings.
 - 2.1.2: gate: not enough (linkage to loading), too far.
 - 2.1.3: not enough ground controllers, not enough radar scopes, not enough planes/controller (bad for safety)
 - 2.1.4: not enough weather equipment (de-icers, people, trucks, fluid, etc.)
 - 2.1.5: too slow weather operations (de-icing que, runway clearing)
 - 2.2: Due to runway effects
 - 2.2.1: not enough runways
 - 2.2.2: spacing not small enough
 - 2.2.3: not enough highspeed turnoffs
 - 2.3: due to departure inadeq
 - 2.3.1: departure airspace too dense, not enough departure routes (insuff navig aids, noise abatement)
 - 2.3.2: not enough route density, too much traffic (busy airport)
 - 2.3.3: not enough departure controllers, not enough radar scopes, not enough planes/controller (bad for safety)
 - 2.4: due to approach inadeq
 - 2.4.1: Approach airspace too dense, not enough departure routes (insuff navig aids, noise abatement)
 - 2.4.2: not enough route density, too much traffic (busy airport)
 - 2.4.3: not enough approach controllers, not enough radar scopes, not enough planes/controller (bad for safety)
 - 2.5: shut down by weather
 - 2.5.1: tornados, microbursts, wind shear, storm from hell
 - 2.6: inadequate "orchestration"
 - 2.6.1: clearance delivery, routing]
 - 2.7: Other reasons
 - 2.7.1: Violate 'constant safety'
 - 2.7.2: Lack of buy-in, etc.
 - 2.8: FAA Will not implement mitigations
 - 2.9: Controllers will not accept automation mitigations
 - 2.10: Mitigation not cost effective
 - 2.11: Pilots will not implement

Figure 8. Example of risk inputs for pilot study.

5. ANALYSIS AND OBSERVATIONS

Direct results of the DDP pilot

Implications of the DDP Pilot for ARD

We wanted to pilot the DDP tool on an ARMD case using narrowly defined pilot task objective. For this task, we developed a relatively simple risk model which mapped tasks against the risks they retired. We held interviews with task/project managers and subject area experts to identify linkages (not score) results.

The results were that we identified some missing pieces necessary for infusion. For instance, issues/risks relating to the nature and attitudes of the FAA, air traffic controllers and the airlines which are not addressed in any of the identified related tasks.

Additionally, for the most of the issues: “everyone is working on everything”. This can be seen by a linkage map generated by the DDP tool in Figure 11. This shows how the different research tasks are related to the various risks. By the high density of lines, it appears that many of the tasks (upper lines) appear to be addressing many of the risks (lower lines).

Figure 12 also shows a similar effect from the simulated annealing with a very steep portion of the curve and a significantly extended flat portion. This plot indicates that a large number of tasks could be redirected to not consume resources addressing these particular risks.

For this pilot study, we had no overall budget information, so we rated the cost as purely the number of tasks. This essentially assumes that all tasks at the lowest work breakdown level are funded to about the same level. This is, of course, not true, but sufficient for the objectives of the pilot demonstration.

Scenarios were also run to view the effect of eliminating the effects on these particular narrow objectives of eliminating entire programs as well as time progressive scenarios with and without mitigations (the thought being that the traffic and crowding issues outlined in the objectives would get worse if no action is taken). Figure 13 shows the risk profile comparison of all mitigations “turned-on” with redirection of an existing program and a time sequence with business as usual (no mitigations implemented).

Topology view: Risks (top rows) vs Tasks (bottom rows)

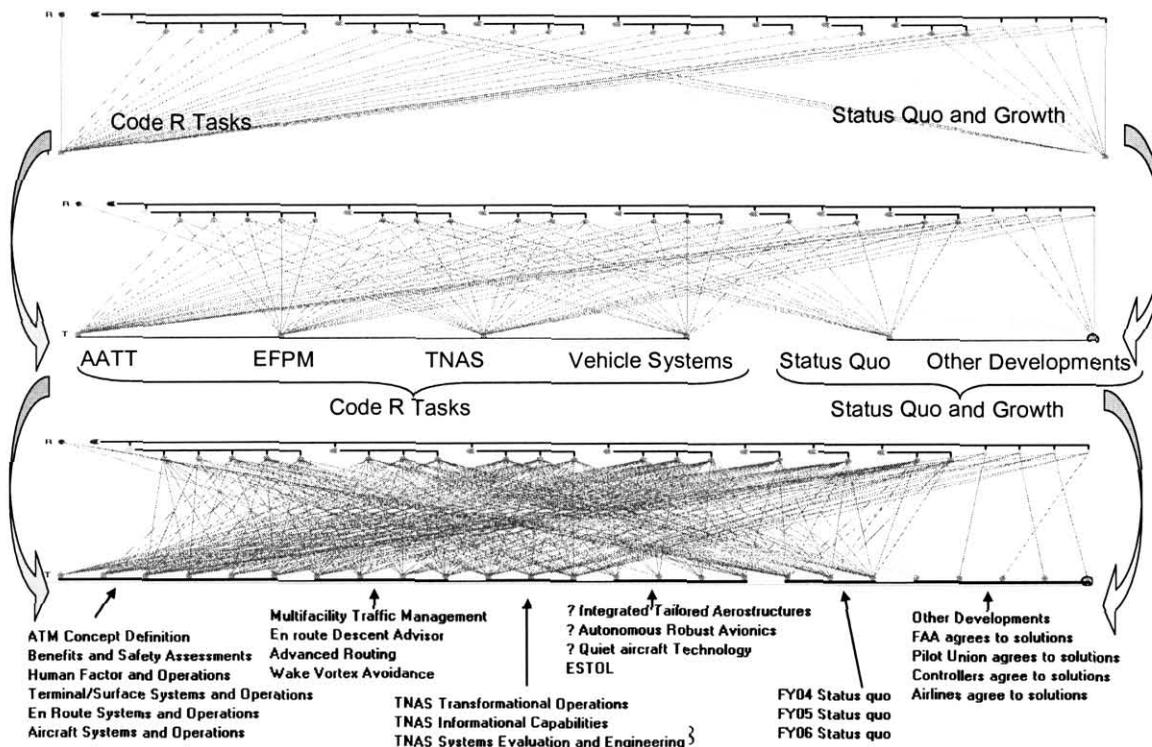


Figure 11. DDP output view showing the interactions between mitigation tasks and risks.

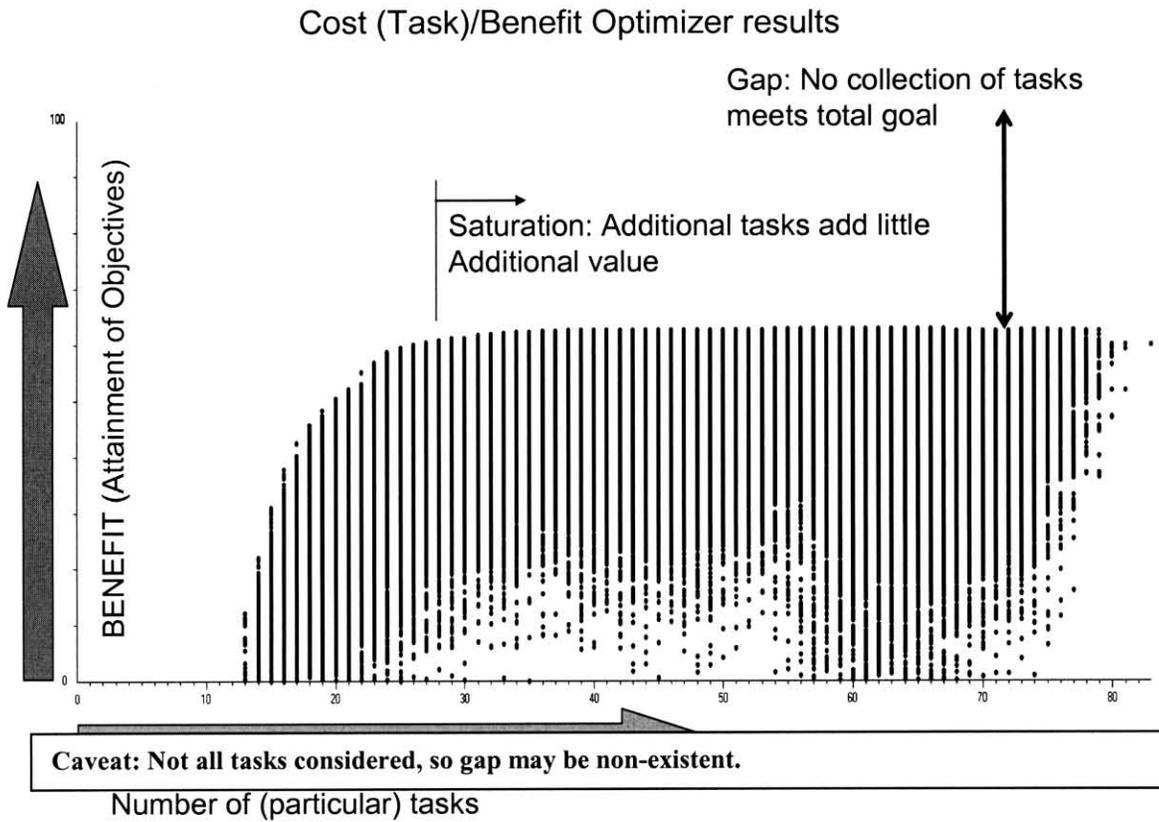


Figure 12. Mitigation tasks versus benefits for the pilot study.

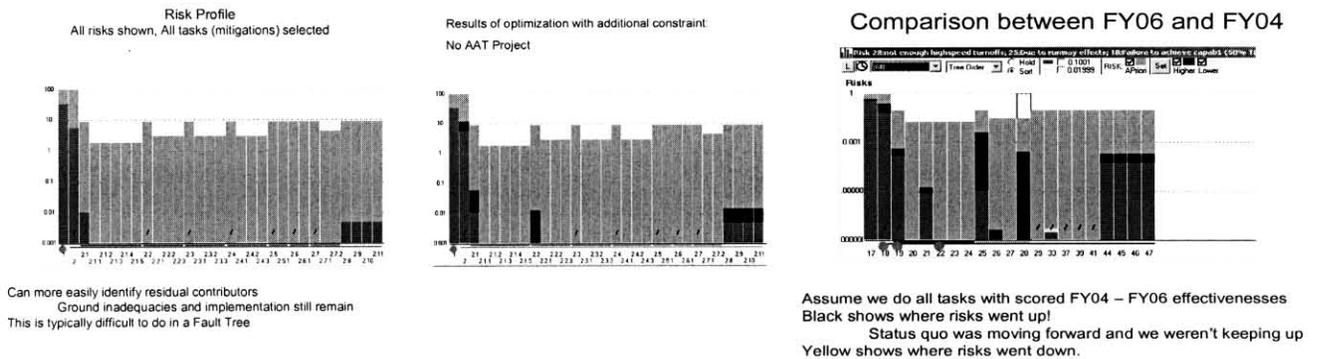


Figure 13. Comparison of risk mitigation scenarios. The left shows all mitigations activated, the center shows the impact of redirection of one of the related programs and the right view shows changes due to time with no action.

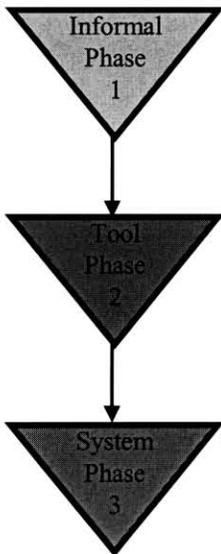


Figure 14. Tool implementation phases.

The impact of the DDP tool to the Aeronautics Research Mission Directorate (ARMD) from a management perspective has passed 2 classical phases. Hopefully, eventually, it will reach the desired third phase.

The first phase (Figure 14.) was the “business as usual” scenario which develops evaluation methods on-the-fly, does not have a systematic quality checking method, and minimum, to no collaboration tools for complex group processing. The good features of this phase are:

1. It’s infinitely compressible. Time is not a problem. No matter how little time you have you can configure the process to fit. By definition.
2. Plausible Deniability. Since it isn’t systematic, lacks built in quality tracking capabilities no one can be directly accused of poor evaluation management.

This phase has several very bad features which helps most organizations to start phase 2:

1. The outcome quality is usually very poor. The evaluators, the external review committees,

and most applicants often roundly denounce the outcomes.

2. It is not repeatable. You cannot evaluate future options with the same process because the techniques, individuals, and processes are not formalized and documented.
3. No way to improve neither the processes nor the outcomes. Because these evaluations are done on-the-fly there’s no corporate knowledge build up, no baselines to track and improve from, no method for isolating process/outcome failures.

This phase was also evident in a recent call for proposals for advanced technology development programs in another mission directorate. The review of proposals were done under incredibly short time frames with insufficient staff and under developed evaluation methods. It was impossible to a top down assessment nor trades across complex, disparate proposals.

The second phase, the tool phase, or as I like to call it “Kill the Messenger” phase is the worst phase to be in. The tool phase is where an organization jumps to in order to try and avoid the negative outcomes of phase 1. However, things have to get worse before they can get better. They actually don’t get worse - it’s just that the organization, through the use of a tool like DDP, realizes really how bad Phase 1 was and how much work they need to do. In particular the tool phase has these sobering impacts:

1. No more plausible deniability. The DDP tool documents the level, and quality of inputs and outputs. Pedigree becomes a key vocabulary word. The gaps become very exposed.
2. Technologists start to circle the wagons. DDP forces, at the very least, relative assessments to be made about key facets of a proposal/technology. Linkage to requirements, concept maturity, lab capabilities, facility availability, and funding trades are now suddenly visible and they often are major weak points in a given technology proposal.
3. The first round of DDP usage caused significant battles between technologists and evaluation managers/analysts. These battles occurred mainly because of poor communications and semantics that were glossed over or ignored in phase 1, but now had to be dealt with head-on through the

processes DDP required the parties to go through.

4. General panic. The parties involved now fear they've lost the ability to manipulate the system to meet their near term evaluation requirements. They realize that the on-the-fly method cannot survive long and will have to dedicate significantly more resources to systematic development and evaluations of technologies. The tool itself becomes a threat and accusations focus on the non feasibility of a systematic technology description and evaluation process.

There are however some positive aspects of phase 2:

1. A line has been drawn in the sand. This baseline, however tenuous, and whatever quality becomes an anchor point for improvement.
2. Hierarchy structure and strategic linkages. Just the very initial usage of the DDP tool provides a framework that helps gives a top down perspective of an organization's R&D portfolio.
3. Organizational situational awareness is significantly improved. The distance required to travel in order to improve the quality has become clearer.
4. Even if the organization continues to do on-the-fly evaluations the exposure to DDP like processes, visualizations, and outcomes will allow more rigorous critique of the phase 1 type processes and add additional fuel to the fire to not only permanently leave phase 1, but to also transition from phase 2 to 3.

6. CONCLUSIONS

A risk model pilot study was generated that allows tasks to be thought of as risk reduction activities. We can now ask what risks are reduced and how much by when. We can also identify areas for which additional data is desired and step back to a bird's eye view of the task portfolio.

This high level view allows us to identify where risks are being over or under addressed, provides a baseline against which to compare products, and allows us to more easily evaluate what if scenarios.

By running this pilot study, we are explicitly able to see the consequences of various changes in: program funding, task scope, task progress and the status quo.

Our vision of the future was demonstrated by addressing two (very hard) problems to solve. We have demonstrated a number of analysis scenarios and capabilities, including:

- How to pick an optimal portfolio
- Traceability of progress using available information
- Strategic planning
- Identifying levels of information fidelity
- How to validate the sum of the efforts to get what we want
- Task management
- How to use risk retirement as a measure of progress
- Predicting likelihood of success
- Identifying mitigation/adjustment/fall-back options
- Identifying driving requirements/objectives
- Identifying tall-poles
- How to verify the data from individual tasks.

We have illustrated (I hope) how the results of the risk model would be used in DDP and the value it would provide. For future work, we plan to go beyond representative data and do a more detailed case study. Additionally, we are responding to the question of "How would you do 100 of these a year?"

We have tried to paint a picture of how the task-level risk retirement approach that would be integrated with higher level portfolio work – top down meets bottom up in the middle and propose to take discrete steps

ACKNOWLEDGEMENTS

Much of 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. We wish to thank.....

REFERENCES

- [1] Martin S. Feather and Steven L. Cornford. "Quantitative Risk-based Requirements Reasoning", to appear in *Requirements Engineering Journal*; published online 25 February 2003, DOI 10.1007/s00766-002-0160-y
- [2] Martin S. Feather, K.A. Hicks, K.R. Johnson and S.L. Cornford; "Software Support for Improving Technology Infusion," Proceedings of the 1st International Conference on Space Mission Challenges for Information Technology; Pasadena, California, July 2003, pp. 359-367; JPL Publication 03-13A, Jet Propulsion Laboratory, California Institute of Technology.

[3] Martin S. Feather, S.L. Cornford and K. Moran, "Got Risk? A Risk-Centric Perspective for Spacecraft Technology Decision Making," *Fifth National Symposium on Space Systems Engineering and Risk Management*; 2004.

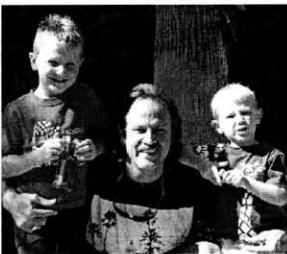
[4] George Price, NASA internal memo "Evaluation Plan Using Risk Assessment to Measure Project Progress." August 2004.

[5] M.S. Feather, S.L. Cornford, K.A. Hicks & K.R. Johnson Applications of tool support for risk-informed requirements reasoning;; *Computer Systems Science and Engineering (CRL Publishing Ltd)*; 20(1): 5-17, January 2005

[6] Some good reference to the airspace system with the handoffs, takeoff landing etc.

BIOGRAPHIES

Andrew A. Shapiro is the Division Lead Technologist for JPL's Enterprise Engineering Division and has been working in microelectronic interconnects for twenty years. He has worked as a member of the technical staff at Rockwell International and Hughes Aircraft, where he was responsible for the packaging of a number of phased array radars and ran their high density interconnect line. He was a Principal Scientist at Newport Communications/Broadcom, where made the first commercial polymer 10GHz Si packages, and he has designed and packaged 10 and 40GHz optoelectronic modules. He earned his BS in chemical engineering at U.C. Berkeley, his MS in Materials Science at UCLA and his Ph.D. in Materials Science at U.C. Irvine. He is on several national committees including NEMI optoelectronics roadmap, ECTC optoelectronics and, IMAPS education. Dr. Shapiro is also currently Assistant Adjunct Professor in Electrical Engineering at U.C. Irvine and is performing research in environmentally friendly manufacturing of electronics and optical and high frequency packaging.



Steven L. Cornford is a Senior Engineer in the Strategic Systems Technology Program Office at NASA's Jet Propulsion Laboratory. He graduated from UC Berkeley with undergraduate degrees in Mathematics and Physics and received his doctorate in

Physics from Texas A&M University in 1992. Since coming to JPL he focused his early efforts at JPL on establishing a quantitative basis for environmental test program selection and implementation. As Payload Reliability Assurance Program Element Manager, this evolved into establishing a quantitative basis for evaluating the effectiveness of overall reliability and test programs as well as performing residual

risk assessments of new technologies. This has resulted in the Defect Detection and Prevention (DDP) process is the motivation for this paper. He received the NASA Exceptional Service Medal in 1997 for his efforts to date. He has been an instrument system engineer, a test-bed Cognizant Engineer and is currently involved with improving JPL's technology infusion processes as well as the Principal Investigator for the development and implementation of the DDP software tool.



Martin S. Feather 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 early phase requirements engineering and risk management and of software validation (analysis, test automation, V&V techniques). 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. For further details, see <http://eis.jpl.nasa.gov/~mfeather>

George Price is the



Yuri O. Gawdiak is a program integration manager for NASA's Aeronautics Research Mission Directorate and has been working advanced information systems and strategic planning for twenty years. He has worked in industry, academia, and government covering space, aeronautics, and engineering research. At NASA he has worked at Johnson Space Flight Center on the Space Station program including principal investigator for the Wireless Network Experiment on STS-76/Mir-21. In aeronautics at Ames Research Center he was the build manager for the Surface Movement Advisor at Atlanta Hartsfield International Airport. At NASA headquarters Mr. Gawdiak was the program manager for the Engineering for Complex Systems program and has worked both in the Exploration Systems Mission Directorate as well as in the Aeronautics Research Mission Directorate. He earned his BS in information systems with a concentration in ergonomics at Carnegie Mellon University. For his work in Personal Satellite Assistant as well as other programs and project Mr. Gawdiak has been awarded both the NASA

Outstanding Leadership and Exceptional Achievement medals.

Wendell R. Ricks manages the NASA ARMD Aeronautical Systems Analysis Project and has been working in systems analysis for the past eleven years. He also manages the Revolutionary Systems Concepts for Aeronautics Project. Prior to working in systems analysis, he was a flight-systems research engineer for eleven years at NASA Langley Research Center in Hampton Virginia where he conducted simulator and flight tests to develop flight-deck design guidance that is better aligned to the manner by which commercial pilots manage flight information. He earned his BS at Old Dominion University and his MS at the Collage of William and Mary.