

A Structured Approach to Strategic Decision Making for NASA's Technology Development

J. Derleth, C. Weisbin, G. Rodríguez, J. Mrozinski

Jet Propulsion Laboratory, California Institute of Technology

MS 301-180

4800 Oak Grove Dr.

Pasadena, CA 91109

Jason.E.Derleth@jpl.nasa.gov

Abstract

A method for collecting quantitative technology development information and matching it with capability needs of future NASA missions is described. Three quantitative analyses are performed, and results presented in terms of optimal portfolios, one at each of many varying budget levels. This is a tool that a decision maker can use to assess trends and better understand the underlying value of each capability area. These results show that quantitative decision making practices are possible in complex systems such as the NASA technology development program.

Introduction

The National Aeronautics and Space Agency (NASA) through its Office of the Space Architect¹ organized an effort to enable a systematic approach to technology

investment in support of space exploration. The Capabilities, Requirements, Analysis and Integration Team (CRAI) was formed as an inter-center effort to capture, validate, and analyze the relevant information.

Information on capability requirements spanning multiple missions of interest to the agency was sought, including mission name, quantified capability need with uncertainty range (and associated units), and date required for each of the capabilities. The team selected a mission set for a proposed Lunar/Mars campaign consistent with implementing the Presidential Vision for Space Exploration.

Technologies which aim to satisfy one or more of the requirements were characterized within a structure which included state-of-the-art and projected quantified performance, cost, schedule, and associated uncertainties, along with proposed milestones.

Decision-maker preferences are important for any analysis. One possibility for guiding investment decisions would be the “mission-enabling” approach: select sets of technologies that adequately enable particular missions. An alternative is a “democratic” approach: select technologies based only upon relative performance gain and applicability to the

¹ A recent NASA reorganization has placed the Space Architect functions within a new Advanced Planning and Integration Office

mission set. This study took the democratic approach, as it is most applicable for long-term planning studies where focused mission designs have not as yet been selected.

A second study was completed at the Jet Propulsion Laboratory (JPL) under separate sponsorship, in parallel with the CRAI effort. This study utilized the same analysis structure, and its results are presented here with the title “Exploration Mission Analysis.”

Methodology

The CRAI team employed a methodology described in Reference 4, and in the web site: <http://start1.jpl.nasa.gov>. This method has been developed over several years at JPL and several papers have been published.^{1 2 3 4}

CRAI and the NASA Space Architect’s Architecture Group were to work closely together to develop requirements. Unfortunately, the President’s announcement in February 2004 of a new exploration focus for NASA made the immediate acquisition of specific requirements difficult. Therefore, CRAI decided to use three pre-existing architectures as a focus for their work. These architectures are:

1. 1998 Mars Reference Mission⁵
2. OASIS mission⁶
3. JSC Architecture #1⁷

These three architectures were chosen because they were available, substantial, and somewhat quantitative. The CRAI block leads extracted quantitative requirement data from these reports.

Technology Data Collection

A data template was created early in the process to meet the needs of the analysis. Intended to capture quantitative data provided by hundreds of technologists, it was designed

to be as simple as possible. The key data that it collects are: contact information, technology name, linkage to the capability-breakdown structure, expected cost, probability of success (PoS) if fully funded, metrics used, state of the art (SOA), expected increase over state of the art, and budgetary and development plans.

CRAI Analysis

The CRAI effort supplied a total of 440 technology data sheets, categorized into the following Technology Areas:

- 2.1 Communications & Info Systems
- 2.2 Space Utilities and Power
- 2.3 Human Support Systems
- 2.4 Automation and Robotics
- 2.5 In-Space Transportation
- 2.6 Scientific Instruments and Sensors
- 2.7 Structures and Materials
- 2.8 Crew Mobility
- 2.9 Launch Access

Of the 440 data sheets, 102 were unusable due to incompleteness in some of the essential fields. The remaining 338 data sheets were matched to requirements extracted from the three architectures. A technology was considered matching only if its metrics were the same as the requirements, i.e., the technology “Ka-Band Travelling-Wave Tube (TWT) 100 to 250W” was considered to be a needed technology because its metric, watts of power in a travelling-wave tube, was matched by a requirement of 100W of power for a TWT in the Mars Reference Mission. This strict matching of requirements to technologies reduced the number of technologies eligible for analysis to 47, and only technology areas 2.1-2.5 had matches.

Scoring Algorithm

To compare disparate technologies with completely different goals is a difficult task.

Either the level of abstraction has to be so high as to remove any individuality of the technology, or a method has to be determined that allows the technology to be represented by a unitless score. Our scoring algorithm begins by converting the projected technology metrics into unitless scores.

For clarity, we will illustrate a simplified example: Disposal of Human Waste (a perhaps distasteful, but essential component to any human exploration). It has two metrics. Each metric is scored separately and the results are averaged.

We begin with the metrics that the technologist provided:

Table 1: Disposal of Human Waste Technology

Metric	SOA	Need
Time Waste Contained	0.25 years	200 years
Stored Volume Density	70 kg/m ³	700 kg/m ³

To make the score unitless, we divide the need by the state of the art:

Table 2: Calculations #1

Metric	Division	Result
Time Waste Contained	200/0.25	800
Stored Volume Density	700/70	10

There is often a large difference in scores, some of them being so large that they are unintelligible. To gain a better physical intuition of these make these results more intuitive, we take the log₂ of the score. This is then a measure of how many times a technology’s performance needs to double to reach the need:

Table 3: Calculations #2

Metric	Log	Result
Time Waste	log ₂ (800)	9.64

Contained		
Stored Volume Density	log ₂ (10)	3.32

The scores are then averaged. Averages are taken instead of sums to keep technologies with many metrics from scoring arbitrarily high.

Finally, the score is multiplied by the percentage chance of success to arrive at an expected value for the score:

Table 4: Final Calculation

Average Score	% Success	Expected Value:
6.48	90%	5.83

Now the unitless score can be seen as an estimate of the expected benefit of this technology to a single mission. If this technology were needed for another mission, a score would be calculated for that mission’s needs, and the two scores added. The average technology score for the 47 CRAI technologies using this method is approximately 1.56 points.

With a score and a cost, a benefit-cost ratio can be obtained. This ratio is calculated for each technology under consideration. A sample budgetary level for technology development across NASA is assumed, and a simple “grab-bag” optimization is performed. The sample budget level is then increased and the optimization ran again. By looking across a wide range of budgets, trends can be seen.

Exploration Mission Analysis

Similar to the CRAI analysis, the Exploration Mission analysis performed at JPL focused on the technology needs of a set of missions. Here, we examine these missions:

- 1) Mars Sample Return (MSR)
- 2) Astrobiology Field Laboratory (AFL)
- 3) Mars Science Laboratory (MSL)

4) Lunar Precursor Mission (LPM)

Other missions were considered in addition to these four, including the Jupiter Icy Moon Orbiter (JIMO) and a Lunar Sample Return mission. However, the JIMO office was still performing technology trades and thus could not provide information, and the Lunar Sample Return mission required only a few minor new technologies, making it less interesting for this study.

Technology Data Collection

Technology data was collected through interviews with the study managers, chief technologists and technologists for particular technologies for these four missions. Data collected was similar to the CRAI effort.

Each mission required several new technologies. MSR required 15, AFL 17, MSL 12 and LPM required 8, for a total number of 52 technologies to analyze.

Exploration Analysis

The same scoring algorithm used to analyze the CRAI data was used here with the one modifier that probabilities of success were not assessable for all technologies and therefore it was assumed that each technology was comparable regarding this parameter.

CRAI Results

Figure 1 presents technology-investment recommendations as a function of total resources available. The decision maker can choose to fund at any level and examine the portfolio for that budget. For instance, if the technology budget for the next 10 years were \$300 million, approximately \$60 million should be invested in Comm and Info Systems, approximately \$30 million in Space Utilities and Power, approximately \$110 million in Human Support Systems,

approximately \$70 million in Automation and Robotics, and approximately \$20 million in In-Space Transportation for a goal of maximizing total technological improvement while moving towards enabling the three CRAI design reference missions (The other areas do not appear in this graph because of lack of data).

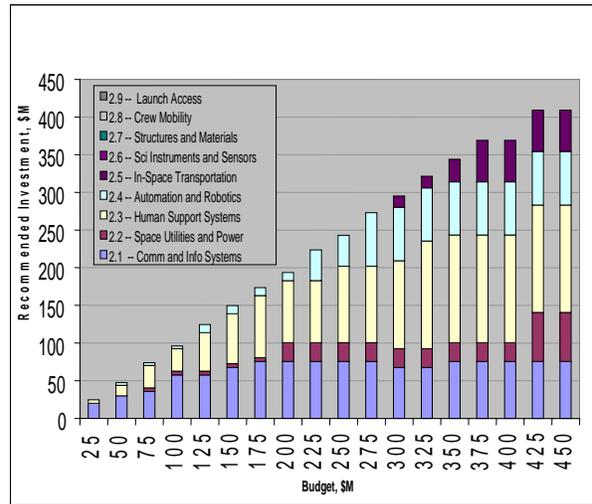


Figure 1 -- Investment Recommendations for Different Technology Areas as a Function of Resources Available

One of the more interesting results of this analysis is demonstrated clearly by the “Comm and Info Systems” (referred to as 2.1 from now on) portion of this graph. As the potential budget increases from \$25 million to \$100 million, the optimal portfolio increases the budget of 2.1. However, at \$125 million, the amount of money given to 2.1 actually decreases slightly. This displacement is due to a more expensive technology, which has a higher benefit/cost score, entering the portfolio. In other words, the budget has increased to a size where a technology that costs more, but has a higher benefit/cost ratio, can enter the optimal portfolio. This drop in funding in the 2.1 area occurs again at higher budget levels, and occurs in other areas at various times.

Ordering

Other factors beyond the particular budget investments were also of interest. These included: 1) the order that technology areas came into the portfolio, and 2) the order the technology areas “saturated,” i.e., when every technology that could be chosen was chosen.

The order in which technology areas first come into the portfolio shows us where the “low hanging fruit” is, where the highest return-per-dollar technologies are. The order of the five areas that had technology-requirement matches was:

- 1) 2.1 – Comm and Info Systems
- 2) 2.3 – Human Support Systems
- 3) 2.4 – Automation and Robotics
- 4) 2.2 – Space Utilities and Power
- 5) 2.5 – In-space Transportation

Primarily, this order is because there are medium-scoring, more affordable technologies in the Comm and Info Systems and Human Support Systems categories. For instance, High and Low Data Rate Coding Systems and Data Compression,” has a cost of only \$2M, so it is funded early on. Its score of 3.3 places it well above the average of 1.56.

The order of saturation is different, however. An area will saturate when there is enough budget to enable every technology within the area, i.e., when the technology that has the worst benefit to cost ratio is funded. The order of saturation is:

- 1) 2.4 – Automation and Robotics
- 2) 2.1 – Comm and Info Systems
- 3) 2.5 – In-space Transportation
- 4) 2.3 – Human Support Systems
- 5) 2.2 – Space Utilities and Power

The limiting technology for Space Utilities and Power is “Regenerative Fuel Cell Systems”, a technology that costs \$40M and has a score of only 0.08. (This low score is

likely due to an error of communication between the technologist and the analyst, since what the architecture states as a need is actually met by the state-of-the-art system!):

Table 5: Negative Score

Regenerative Fuel Cell Systems	SOA	Architecture #1	OASIS
W-h/KG	1880	400	1000

Since the need is significantly less than the SOA, the scoring algorithm produces a negative score. When combined with the positive score of the technology’s other metric, this gives an average score of 0.08. It seems clear that only the positive-scoring metric should have been submitted. This demonstrates the importance of iteration and communication in such a large project.

CRAI and Supplemental Data

As the data collected was sparse, we decided to supplement it with the JPL Exploration study (noted here on as the Pilot). Eleven more technology areas were added, with 156 new technologies, for a total of 203 technologies arranged into 16 areas. The new technologies were matched to requirements from future NASA missions.

Investment Strategies

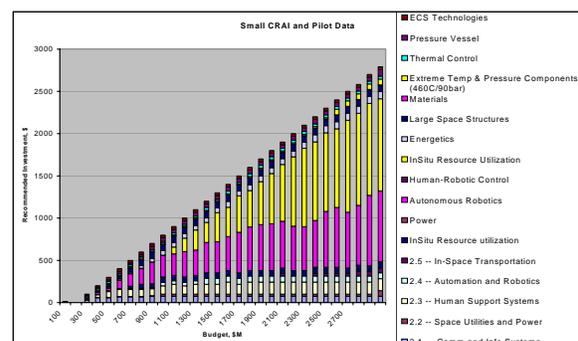


Figure 2 -- CRAI and Pilot Data Portfolio Recommendations

As can be seen in Figure 2, the total possible investment has increased. As before, the decision maker can choose to fund at any level and examine the optimal portfolio for that budget. Here, the “InSitu Resource Utilization” and the “Autonomous Robotics” categories have the largest possible budgets. This is due to more—and more expensive—technologies in these two categories.

Ordering

The order of the CRAI technology categories entering the optimal portfolio does not change. The amount of budget necessary for the CRAI areas to enter increases as they compete against the pilot study’s 11 areas.

What is interesting, however, is that the order of saturation does change. The order is as follows for both studies:

Table 6 -- Order of Saturation

CRAI only	CRAI and Pilot
2.4 – A & R	2.1 – Comm
2.1 – Comm	2.4 – A & R
2.5 – In-space Trans.	2.3 – Human Support
2.3 – Human Support	2.5 – In-space Trans.
2.2 – Space Utilities	2.2 – Space Utilities

What changes the order of saturation is not well understood. In fact, the last technology enabled in every area is identical between the CRAI-only and the CRAI-and-Pilot studies. The order may change because the addition of different technologies changes what mix of technologies is optimal. This has little effect in the beginning, where the funding is extremely limited and the best inexpensive technologies are chosen. Further on, when the budgetary levels are large enough to enable many of the average technologies, this larger set of technologies makes possible the “dropping” of average technologies to select technologies with better benefit/cost ratios. It is thought that because of this larger budget,

and because of the complex interaction of the technologies, the order of saturation changes.

Another possibility is the step size within the budget. If the technologies are of a lower order of magnitude than the step size, then a set of less expensive technologies are enabled all at once. Having a larger number of technologies increases the number of technologies that are an order of magnitude less than the step size, which can affect the complex interaction in the optimal portfolio.

Exploration Missions Results

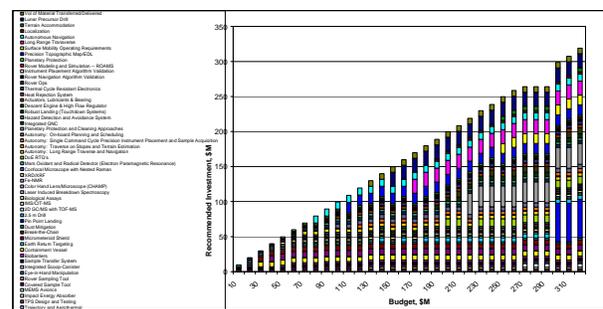


Figure 3 – Investment Recommendations for Different Technologies as a Function of Resources Available

Figure 3 gives the optimum technology choices at various budget levels, with the final budget level funding all technologies. These recommendations are the result of the democratic method, which allows for the greatest increase in technology without taking into account the constraints of having to satisfy any mission’s entire technology suite.

In Figure 4, we combine the technologies in Figure 3 working for the same mission together at each budget level to see how close the democratic method gets us to enabling any of the four missions.

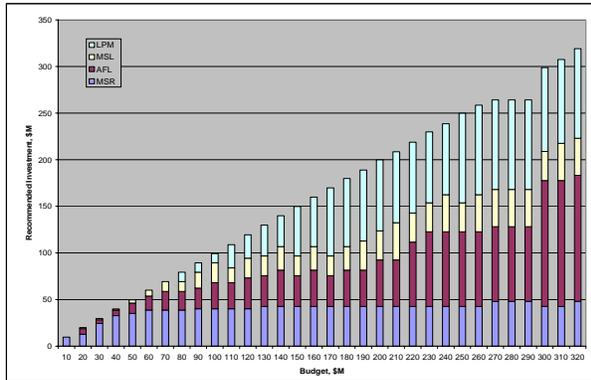


Figure 4 – Recommended Technologies Combined by Mission

From Figure 4 we begin to see that at some budget levels the democratic approach is not meeting any mission’s full technology needs. For example, it is clear that AFL does not have its full technology suite until large budgets. Filtering technologies which are not enabling a mission from Figure 4, we see that the democratic approach maximizes overall technology capabilities, but does not maximize the number of missions enabled for lower budgets (see Figure 5 below).

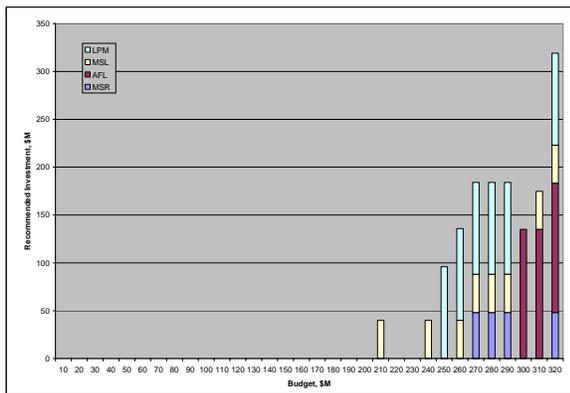


Figure 5 – Missions Enabled via the Democratic Method

From Figure 5, an odd conclusion is reached. At a technology investment budget level of \$210 M, only MSL is enabled – but to enable MSL’s entire technology suite is only on the order of \$40 M. In fact, for \$210 M, MSL plus MSR and LPM could all be

enabled. The democratic approach is not the method for optimizing enabled missions.

As the democratic approach has a poor time enabling missions at lower budget levels, we can see if we can do better with the mission enabling approach. Here, the analysis is run on each mission’s technology suite rather than each individual technology, where the score and the cost of the mission’s technology suite are simply the sum of each technology’s score and cost in that suite. The results of the mission enabling method can be seen in Figure 6 and can be compared with the results of the democratic method in Figure 5.

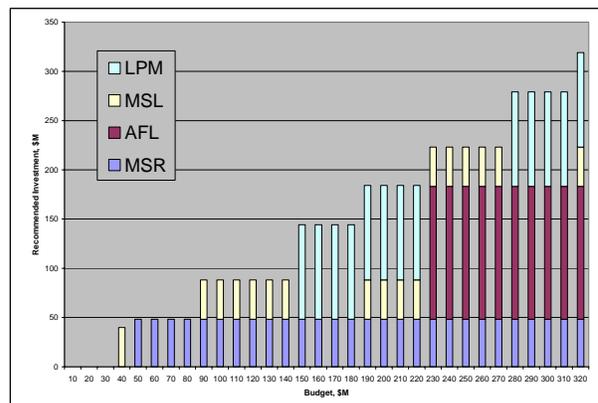


Figure 6 – Missions Enabled via the Mission Enabling Method

The biggest difference between Figure 5 and Figure 6 is that the mission enabling method allows for more missions at lower budget levels. This is due to the fact that each mission has those technologies that are considered good by the democratic method (those that have a high score and a low cost), and those that are not considered as good (those that have a lower score and a higher cost). In the democratic method, since each mission has some of these lower ranking technologies that it needs, and because these are always chosen last as they do not optimize the budget, the missions must wait until higher budgets are reached which can then allow in the lower ranking technologies in to complete the technology suite. This differs from the

mission enabling approach, where the better, cheaper technologies bring along the weaker, more expensive technologies at all times – the missions that are fully funded most frequently will be missions that have strong technologies to compensate for their weaker technologies.

It is seen that the lunar precursor mission is the third mission chosen in the mission enabling method, whereas in the democratic method, it is the second mission to be enabled. This is because those missions that have the smallest number of low ranking technologies tend to be enabled first in the democratic method, whereas in the mission enabling method, those missions with higher total scores and lower costs will be chosen first.

Competition Border

As was seen in the above analyses, as the budget level increases, some technologies fall out to allow in the more expensive (but higher scoring) technologies in. Figure 7 is a graphical representation showing the interaction of technologies as they come in and out of the Exploration Missions portfolio.

Figure 7 – Technology Competition Border

In Figure 7, the horizontal axis represents each of the 52 technologies. The technologies are ordered with the highest score/cost ratio

technologies on the left, decreasing towards the right as the trend line indicates. The main vertical axis indicates the total budget level. The color coding indicates each technologies funding status at each budget level.

Grey indicates that a technology has yet to be included at any budget level. Examining the rightmost technologies (those with the lowest score/cost ratio) remain grey through many budget increases.

Green indicates when the budget has increased enough to fund the technology.

If after turning green a technology then turns blue, this indicates that the technology has continued to be funded. However, if a technology turns red, the technology has lost its status of being funded do to competition. The region defined by the red and green is referred to as the “Competition Border.”

Technologies which never turn red are the stronger technologies. Looking to the left of Figure 7, we see several technologies which enter at a low budget level, turn blue and never turn red – the competition is no match.

Technology #12, on the other hand, entered (turned green) at the first budget level, but at the first increase turns red. The reason for this is that at the first budget level, it was just cheap enough to fill in the remainder of the budget, and it was cheaper then the higher score/cost ratio technologies to its left. However, once the budget increased a step, the higher score/cost ratio technologies could enter the budget – provided that they also took away #12’s funding! Several other technologies can be seen entering the budget before their better neighbors to the left, only to be booted shortly thereafter.

Sensitivity/Robustness Analysis

Once all of these levels are run, there is an optimal portfolio for each budgetary level. This represents the recommended investment under the proposed conditions stated earlier. But how sensitive is a portfolio

recommendation with respect to uncertainties regarding the capability requirements and technology improvement characterization?

To measure the robustness of inclusion for each technology, a sensitivity-type analysis was run. For each budgetary level, every technology's score and cost were adjusted a random amount from -5 to 5% and the optimization was run. This was repeated 200 times for each budgetary level.

The Monte Carlo robustness/sensitivity runs showed several technologies at each budgetary level that did not switch in and out. These technologies were different at each budgetary level. This is understood when examining Figure 7—once a technology has passed a certain level, it is always selected.

The technologies that shift in and out of the optimal technology portfolio are more interesting. The following graph will be used for illustration:

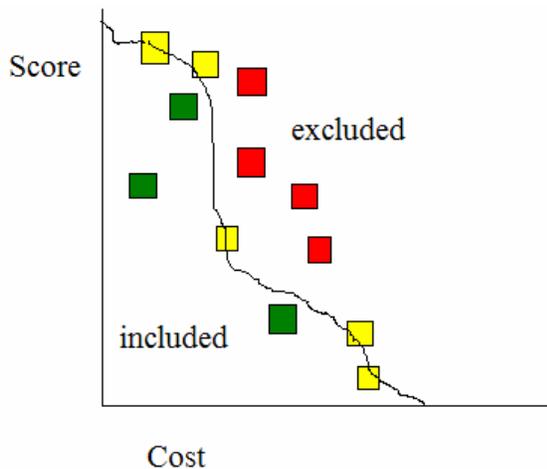


Figure 8 -- Estimation of Robustness

The boxes each represent a technology and its uncertainty bounds. The curved line is imagined to be the optimal portfolio border. Its shape is unknown and changes with each analysis run, but is drawn here for reference.

The red boxes have too high of a cost for their score, so they are not included in the optimal portfolio. The green boxes are within the border; they have a lower cost-benefit ratio

and are included. The yellow boxes are on the line. Depending on their score and cost, they will be included in the portfolio or not. By the number of times they are chosen in the 200 runs, we can estimate how close to the line the center of their box is, a measurement of how robust the technology is within that portfolio.

Conclusions

This work has demonstrated that a portfolio analysis approach applicable to strategic decision-making for the Exploration Mission is feasible. This process:

- Is objective, traceable, quantitative, and repeatable
- Is capable of enabling a strategy to tactics approach
- Provides a method to determine the robustness of results to alternative data sets and policy preferences through a systematic sensitivity analysis

Primarily due to the lack of data, and the lack of verification of the data we had, we are unable to say that the graphs above actually reflect optimal portfolios. However, there are several things that became clear through analysis of the data:

Certain **categories** were consistently chosen by the algorithm at lower budgets than others, even with the Monte Carlo analysis. An example of this is the 2.1 – Comm and Info Systems area. This was due to lower cost or higher value technologies being a part of the category. These technologies tend to be chosen first, thus the category received funding at a lower budgetary level.

Certain **categories** were consistently “saturated” first, meaning that all of the possible technologies were chosen by the optimization routine first. What stood out in these categories was that there were no low-scoring, high-cost technologies. An example of this is 2.4 – Automation & Robotics. This area's technologies are consistently used by

multiple missions, giving many of them relatively high scores.

Certain **technologies** were always chosen, regardless of budgetary levels, technology mix, or sensitivity/robustness runs. These technologies are robust technologies and were most often technologies that were used by a number of missions. A few examples:

- a. 2.1: Ka Band TWT 100 to 250W – used by multiple missions, low cost
- b. 2.3: Water Recovery From Waste – very high score, fairly low cost
- c. 2.4: End-Effector placement – relatively low cost

Certain **technologies** were never chosen, most likely due to data collection errors. Many of these had negative scores, showing us that even simple data collection becomes difficult when the data collection effort is scaled to an agency as large as NASA. A few examples are provided here:

- a) 2.5: LOX/L-CH₄ OMS Engine – negative score, the SOA is sufficient
- b) 2.2: Stirling Reactor – negative score, the state of the art is sufficient

Lessons Learned

There are several recommendations that can be made to facilitate this process if attempted in the future:

Begin with a web-based data collection system. It is believed that this will make the data entry job easier for the technologists. This type of system also eases transfer of the data into a spreadsheet.

Block leads should spend the time up front to create a hierarchy showing all of the technologies that they hope to collect data for. This will show alternative technology pairs, technology pairs that should be enabled simultaneously, and also any dependencies that might exist.

Data collection should be done early to give enough time for a thorough analysis.

The analysis structure can be created early on with preliminary data, and all analysis routines can be constructed with preliminary data as well. Preliminary results, while not believable, can be used to verify that the structure has been written correctly and that the technologists understand how their data will be used.

The CRAI effort has shown that quantitative data collection and analytical decision making can, in fact, be done

References

1. G. Rodriguez and C.R. Weisbin, "A New Method to Evaluate Human-Robot System Performance", special issue of Journal on Autonomous Robots, Kluwer-Publishers, 14,165-178, March 2003.
2. G. Rodriguez and C. R. Weisbin, "A New Method for Evaluation of Human-Robot System Resiliency," JPL Publication 02-26, November 2002
3. J. P. Chase, A. Elfes, W.P. Lincoln, and C. R. Weisbin," Identifying Technology Investments for Future Space Missions", Space 2003 Conference and Exposition, AIAA, September 25, 2003, Long Beach, CA, U.S.A
4. C.R. Weisbin, G. Rodriguez, A. Elfes, and J. H. Smith, "Toward a Systematic Approach for Selection of NASA Technology Portfolios", Wiley InterScience, Sept 2004
5. "Mars Reference Mission", http://www_sn.jsc.nasa.gov/marsref/contents.html
6. "Orbital Aggregation & Space Infrastructure Systems", http://rasc.larc.nasa.gov/rasc_new/HPM/OASISEXEC_97.pdf
7. "JSC Architecture 1" -- NASA Exploration Team (NEXT) Design Reference Missions Summary, July 11, 2002

Author Biographies:

Jason Derleth is a systems engineer at the Jet Propulsion Laboratory (JPL) of the California Institute of Technology. He received an SM degree in Aeronautical and Astronautical Engineering from the Massachusetts Institute of Technology in 2003, and a BA in the Liberal Arts from St. John's College in Annapolis, Maryland in 2000, where he received awards for writing and excellence in the arts as demonstrated by his hand carved cello.

Jason's current research interests include new approaches to complex systems analysis, technology portfolio and mission design optimization; and incorporating such understanding into optimization tools.

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 1999 NASA Exceptional Service Medal for development and infusion of robotics from basic research into NASA flight missions and missions of other agencies, and a Lifetime Achievement Award from the World Automation Congress in 2004.

Guillermo Rodriguez is a research engineer at the Jet Propulsion Laboratory of the California Institute of Technology. He received a Ph.D. degree in system science from the University of California at Los Angeles in 1974, and has been at JPL since then. He has participated in the development

of autonomous control systems flown in planetary missions, including the Viking, the Voyager, and the Galileo missions.

Since the mid 1980's, he has been involved in research for NASA research and technology programs. His research interests include analysis in hierarchical systems, estimation theory, nonlinear multi-body system control, and control for autonomous systems.

Joe Mrozinski is a systems engineer at the NASA Jet Propulsion Laboratory, currently involved with technology assessment and mission trade models and tools. Joe attended the University of Michigan where he earned a BS in Aerospace Engineering in 2002, a BA in Philosophy in 2003 and a Masters in Space Systems Engineering in 2004.