

# Architecting Space Exploration Campaigns: A Decision-Analytic Approach <sup>1,2</sup>

Prof. Erin Baker  
edbaker@ecs.umass.edu  
(413) 545-0670

University of Massachusetts, Amherst  
Engineering Laboratory, 160 Governor's Drive, Amherst MA 01003

Elisabeth L. Morse  
elisabeth.l.morse@jpl.nasa.gov  
(818) 393-6933

Dr Andrew Gray  
andrew.a.gray@jpl.nasa.gov  
(818) 354-4906

Jet Propulsion Laboratory  
4800 Oak Grove Drive, Pasadena CA 91109

Robert Easter  
robert.w.easter@jpl.nasa.gov  
(818) 354-3552

*Abstract*—Under its Vision for Space Exploration, NASA is moving from designing single space missions to architecting whole exploration campaigns. In comparing campaign options, the flexibility to respond to discoveries and adapt to uncertainties is critical. This paper shows the benefits of Decision Analysis techniques for campaign design and evaluation. Important concepts of decision analysis are reviewed through the lens of designing a campaign to find exploitable equatorial water on Mars. The method developed herein is general to any *search* campaign. The paper concludes with a discussion of the challenges and opportunities in applying similar techniques to other types of campaigns.

## 1. INTRODUCTION - BACKGROUND

A number of papers have highlighted the importance of decision analysis and real options theory when evaluating space systems that have flexibility to respond to uncertainty [Iamassoure][Saleh][Shishko]. While recognized as valuable, these approaches have rarely been used for NASA missions, whose flexibility after launch is limited. But under the new Vision for Space Exploration, NASA is moving from designing single space missions to architecting whole exploration campaigns. A campaign includes links between subsequent missions, but also *options* to change future missions based on earlier outcomes, and on the resolution of technological and programmatic uncertainties.

### *Shortcoming of the traditional approach*

The traditional approach to space program planning has several shortcomings for campaign evaluation. First, even when designing a series of missions, the typical approach optimizes one mission at a time. This ignores the fact that the optimal choice for the first mission depends on the options available for subsequent missions. Second, the approach typically uses one of two comparison techniques: either setting requirements and comparing the cost of various design options; or setting a cost cap, and comparing the performance of various mission options. But even when designed to meet the same requirements, several

## TABLE OF CONTENTS

1. INTRODUCTION - BACKGROUND
2. PROBLEM DESCRIPTION
3. PROBLEM FORMULATION WITH INFLUENCE DIAGRAMS
4. MODELS FOR DECISION ANALYSIS INPUTS
5. PROBLEM SOLVING WITH DECISION TREES
6. CONCLUSIONS
7. REFERENCES
8. ACKNOWLEDGEMENTS
9. BIOGRAPHIES

<sup>1</sup> TBD/\$20.00© 2006 IEEE

<sup>2</sup> IEEEAC paper #2176, Version 0, Updated 11/2/2005

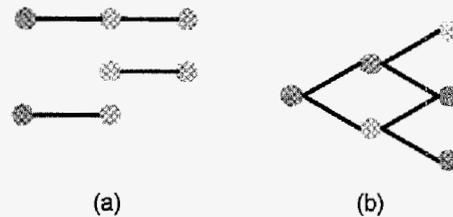
mission options often have different performance characteristics. More importantly, they often open up different options for the next mission in the campaign (for example, through technology feed-forward).

Finally, the traditional approach does not explicitly take uncertainty into account. When considering design options for the second mission in the campaign, it assumes a deterministic set of results from the first mission. In reality however, the results of each mission are probabilistic. The mission could fail; it could be moderately successful; it could discover the unexpected, etc. What the decision maker needs to select is not one option for the subsequent mission, but instead a *strategy*, i.e. a mission option for each possible outcome of the first mission. The probability of each of these possible outcomes should influence the decision.

Effective methods for designing space missions and estimating their cost have been used for years [teamx]. Recent improvements in concurrent engineering are expanding point designs into sizing models that capture the local trade space, making possible the rapid generation of cost/performance curves on a mission-per-mission basis [morse]. In order to properly evaluate alternate campaign design options, three elements need to be added to the space program planning methods: (1) evaluating cost-benefit metrics instead of cost alone, (2) taking uncertainty into account and (3) considering *strategies* for whole series of missions, instead of missions one-at-a-time.

#### *Benefits of a Decision-Analytic approach*

In order to address these shortcomings, we propose applying a framework of sequential decision making under uncertainty to campaign planning. This paper illustrates the use of Decision Analysis tools on the example of planning a campaign for finding exploitable equatorial water on Mars. After introducing the case study, we introduce Influence Diagrams, high level visual representations of decision problems, and illustrate their usefulness in structuring complex problems. We use probabilistic modeling to model the uncertainty about the location and amount of water on Mars. We use a decision tree to establish the best campaign, for given assumptions, and then apply extensive sensitivity analysis in order to determine the robustness of our results, and the value of gathering further information. The first benefit of this approach is that it explicitly incorporates uncertainty, rather than assuming that certain results will happen, or using averages. This allows for a well developed solution to the problem and helps to avoid surprises. The second benefit is that it allows us to incorporate the value of flexibility [Dixit\_Pyndyck]. Finally, it helps define which information requires further study, thus saving significant study time and money.



**Figure 1.** Conceptual difference between Traditional (a) and Decision-Analytic (b) Approaches.

## 2. PROBLEM DESCRIPTION

When faced with the task of determining the best campaign to search for exploitable equatorial water on Mars, JPL's "Eureka" Team found Decision Analysis concepts and tools particularly powerful [gray]. This "water-search" campaign is the example that will be followed in this paper, for its relevance and its illustrative qualities.

### *The importance of searching for water on Mars*

Under its new Vision for Space Exploration, NASA intends to send humans to the Moon and prepare for the eventual human exploration of Mars. Launch mass from Earth is a key problem and a cost driver when designing a human mission to Mars. The use of in-situ resource utilization (ISRU) has been advocated as a method to dramatically reduce the mass to be delivered at Mars, thus reducing the launch mass at Earth, and potentially doing away with the need for any in-space assembly [isru\_1]. If and where present, water would be, by far, the most interesting indigenous resource for use by a human mission to Mars [isru\_rapp]. Water is already thought to exist at the poles, in the form of water ice. But equatorial water would be particularly interesting, as other engineering considerations favor equatorial landing for the first human mission. Knowledge of the presence, form, and amount of any exploitable equatorial water would therefore have significant value when designing a human mission to Mars.

What is the best way to develop this knowledge? While remote sensing missions offer coverage, in-situ measurements provide the only unambiguous proof of the presence of water. Intermediate options include aerial platforms, rovers, and networks of small probes. To maximize returns, a several-mission approach might be interesting. What combination of missions would maximize the probability of finding exploitable water, while minimizing total cost?

### *Typical answer with the traditional approach*

The traditional approach to mission optimization would select a few deterministic series of missions, and compare their cost and performance. In this case, the candidates for comparison would likely include one stationary Lander

concept and one Rover concept, with the possibility to precede either mission with a remote-sensing Orbiter. The decision would be driven by one key constraint coming from above, such as cost or schedule. A logical conclusion would be as follows:

- If only one mission opportunity is allowed to search for water and the budget is low, only one stationary Lander is possible.
- If only one opportunity is allowed, but the budget makes it possible, send a Rover to take advantage of the additional range, improving the likelihood of a find and a characterization of the water extent.
- If two opportunities are available, send first a remote-sensing Orbiter, so as to improve the site selection for the in-situ mission. Then choose the in-situ mission (Lander or Rover) based on the available budget (as above).
- If three opportunities are available, with sufficient budget, send an Orbiter first, then refine the search with a long-range Rover. Finally send a stationary Lander equipped with an in-situ research utilization (ISRU) technology demonstration, at the site where the most resource was found.

#### Limitations of the traditional approach

This approach has the advantage of simplicity and rapidity, but it does not answer all the decision makers' questions. In particular, how are the ultimate decision makers to decide how many mission opportunities to allow for the search for water? Is it worth investing in remote sensing before sending an in-situ mission? Should the choice of a second mission depend on the outcome of the first mission, making it necessary to develop several designs in parallel? What is the value of the range provided by a Rover, compared to a stationary Lander, is it worth the cost? Can this value change after we learn more about the distribution of water on Mars? Should any funds be invested in developing new concepts such as aerial platforms and networks of small probes? etc.

### 3. PROBLEM FORMULATION WITH INFLUENCE DIAGRAMS

A decision-analytic approach addresses the shortfalls of the traditional approach. To make this point, this section will describe the proposed decision-analytic framework, illustrating the concepts at each step with their application to the search-for-water campaign.

#### General approach summary

Figure 2 summarizes the decision analysis process as this paper proposes to apply it to campaign architecture. The process starts with a definition of the campaign goals,

constraints, and figures of merit, i.e. the ultimate parameters of interest to the decision makers. The second and most important step is the formulation of the problem in terms of uncertainties, decisions, and relationships between them. Influence Diagrams (IDs) are useful tools in this problem formulation. Once formulated, the problem needs to be quantified. Mission studies, databases, or expert knowledge provide ranges of estimates for all required mission information. Expert knowledge provides the basis for modeling of the uncertainties. An influence diagram, once well documented with mission information and uncertainties models, provides the basis for generation of a decision tree (DT). "Rolling back" the decision tree solves for the optimal campaign strategy under nominal input values. Sensitivity analysis throughout the ranges of input values provides insight into the robustness of the results, and the need for more detailed information. The analysis process is iterated until a robust answer is reached.

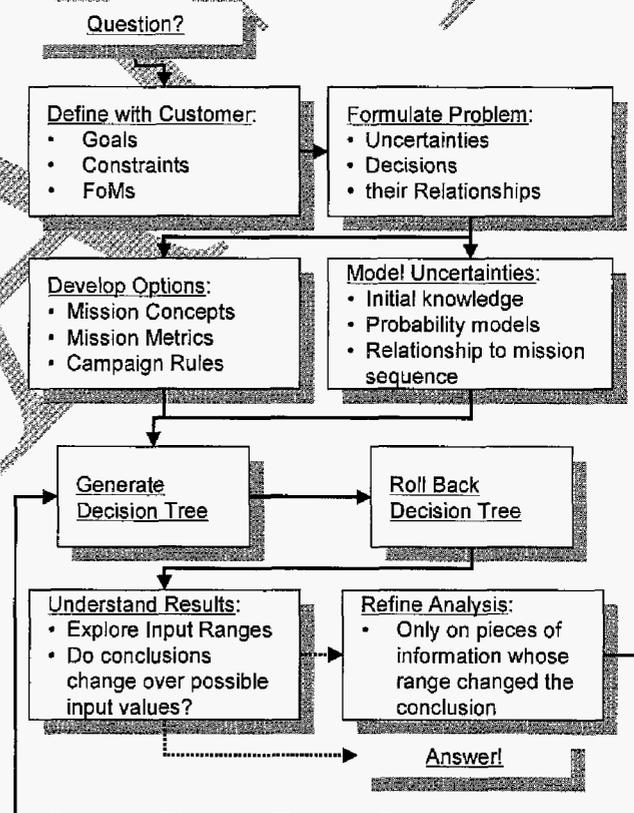


Figure 2. General Campaign Analysis Approach

#### Defining campaign goals

**Campaign goals-** It is important to start the analysis with a clear understanding of the campaign goals, against which alternative strategies will be evaluated. A campaign can have primary and secondary goals.

In this case, the primary goal is to locate near-surface (within the upper 5-m) water in the equatorial regions of Mars (+/- 30 deg latitude), and characterize its form and extent (3-dimensional concentration maps). Other human precursor measurements can constitute secondary goals.

*Campaign constraints-* When they exist, strict constraints on the campaign design must be identified upfront.

The search for water on Mars would be subject to a constraint on total budget, yearly budget, total schedule, and/or number of mission disasters. Only campaign strategies that meet all constraints will be considered in the analysis. The analysis below does not include numerical values for these constraints, which were unknown for the study; instead, these constraints were kept for sensitivity analysis.

*Formulating the problem*

Once the goals are clear, three key elements define any decision problem: values, alternatives, and uncertainties.

*Figures of Merit-* Figures of Merit (FoMs), or *Values* are what the decision maker (DM) uses to determine which outcomes are better than others. Values are principles used to evaluate the “actual or potential consequences of action and inaction” [Keeney]. They must measure how well the campaign meets its goals and constraints.

In the search for water, the following FoMs are of prime interest: total cost of the campaign (to minimize); total time to complete the campaign (to minimize); and likelihood of finding an equatorial site with sufficiently high water concentration to carry out ISRU (to maximize). Secondary FoMs include the coverage of the planet (to maximize), as greater coverage provides more information, and thus more options, to human mission planners; the number of additional human-precursor measurements carried out over the course of the campaign (to maximize); and the number of mission disasters (to minimize), which would strongly impact the public perception and the likelihood of securing further funds for the campaign.

In order to determine an *optimal strategy* for the campaign, it is necessary to aggregate the FoMs into one “value” metric. This aggregation can be an equation of any form. For simplicity, in this study the aggregate value was taken as the weighted sum of all FoMs. The choice of weights is highly dependent on the decision maker. Sensitivity analysis with respect to FoM weights is important. Table 1 lists the FoMs and their baseline weight as used in the study.

Cost (in Millions \$)	-1
Time to Complete (in months)	-10
Number of Mission Disasters	-1000
Likelihood of “H” find at the site level	10,000
Likelihood of “M” find at the site level	1000
Likelihood of “H” find at the area level	10
Likelihood of “M” find at the area level	1
Coverage, per 100 regions	5
Coverage, per 100 areas	5
Coverage, per 100 sites	5
Additional Measurements	300

Table 5 contains the baseline values for the FoM weights used to generate the aggregate value metric. The weights can be interpreted as the dollar amount assigned to the outcome. For example, consider the weight of -10 assigned to *Time to Complete*. This implies that the DM would be willing to pay up to \$10M for every month that could be shaved off the completion time; or conversely, would want to save more than \$10M to justify a one month delay. Similarly, the value of 10,000 assigned to *Likelihood of “H” find at the site level*, implies that the DM would be willing to pay up to \$10B on a campaign that was certain to find “H” at the area level.

*Alternatives-* Alternatives are actions that the DM can choose between. The availability of future alternatives to respond to the resolution of uncertainties is what defines a decision problem.

In our example, the alternatives are mission options for each Mars launch opportunity.

The study assumes a discovery-driven program, where the choice of each mission is dependent on the results of the previous mission. This can be achieved either by spacing the missions sufficiently so as to allow for development time (i.e., launch a mission to Mars every other opportunity), or by developing several mission options in parallel (thus being able to launch at every opportunity). The former approach will be assumed in the paper as a baseline. The two approaches are easy to compare using the framework proposed herein.

*Information-* Information is what connects *alternatives* to *values*. Some information is well known, or can be known with varying levels of accuracy by carrying out studies at various levels of detail. This is the case, for example, of mission cost, or likelihood of a mission disaster. For the sake of the study, this information can be considered deterministic. If studies have not yet estimated the information with accuracy, a range of values can be used; sensitivity analysis will determine the importance of refining the estimate.

**Table 1. Baseline FoM Weights**

FoM	Weight
-----	--------

But other key information is often unknown, and cannot be known before exercising one of the *alternatives* available in the decision problem. This is the case, in our example, of the distribution of water on Mars, and of the amount of water that a given mission will find. These pieces of information correspond to the key *uncertainties* of the problem.

*Uncertainties*- Some of the uncertainties will get resolved as the campaign progresses, and their resolution will influence future decisions. Each uncertainty has a list of possible outcomes. For example one uncertainty is “What is the highest concentration of water within a given area on Mars?”. The possible outcomes range from 0% to 100%.

*Influence Diagrams*

Influence Diagrams (IDs) are useful tools in formulating decision problems [ID\_citations]. Influence Diagrams are a visual representation of a decision problem. They are “graphic representations of formal mathematical models” [owens shacter nease], which lay out the relationships between decisions, uncertainties and values.

Figure 3 illustrates the Influence Diagram for the water campaign. In one simple drawing, the influence diagram summarizes all the relationships between the decisions to make in the campaign, the uncertainties that will get resolved over time, and how they relate to the ultimate campaign values. For this reason, it is a powerful tool for problem formulation, as well as for communication.

*Boxes*- The three key elements of a decision problem are represented in Influence Diagrams. Values are represented

as diamonds; alternatives as squares; and uncertainties as ovals. Each node is associated with tables or trees of information. For example, each decision is associated with its alternatives. For this example, the alternatives in Mission 1 are Orbiter, Aerial, Lander, Network, and Rover. Each uncertainty is associated with a list of outcomes, and the probabilities over those outcomes. The final node (“FoMs”) represents the aggregate value of the campaign. It aggregates all the FoMs into one value metric, and aggregates all values for each possible outcome of the campaign.

*Arcs*- The other key elements in IDs are the arrows or arcs. They illustrate how the various elements in the decision problem are related.

An arrow between uncertainty nodes indicates that a probabilistic relationship may exist. This occurs when the outcome of one node impacts the probability of another node. For example, the distribution of water on Mars impacts the probability that any given mission will find a site with high water concentration. Note that the existence of an arrow is a weak assertion – a relationship may exist. The absence of an arrow is a strong assertion: the events are conditionally independent, given all other relationships on the ID. For example, the cost after the first mission does not influence the distribution of water on Mars.

An arrow going from one decision node into another implies that the first decision is made and known before the second decision must be taken. Similarly, an arrow going from an uncertainty node into a decision node implies that the outcome of that uncertainty node will be known before the decision must be taken. Thus, our diagram indicates that we will not implement mission 2 until the results of mission 1

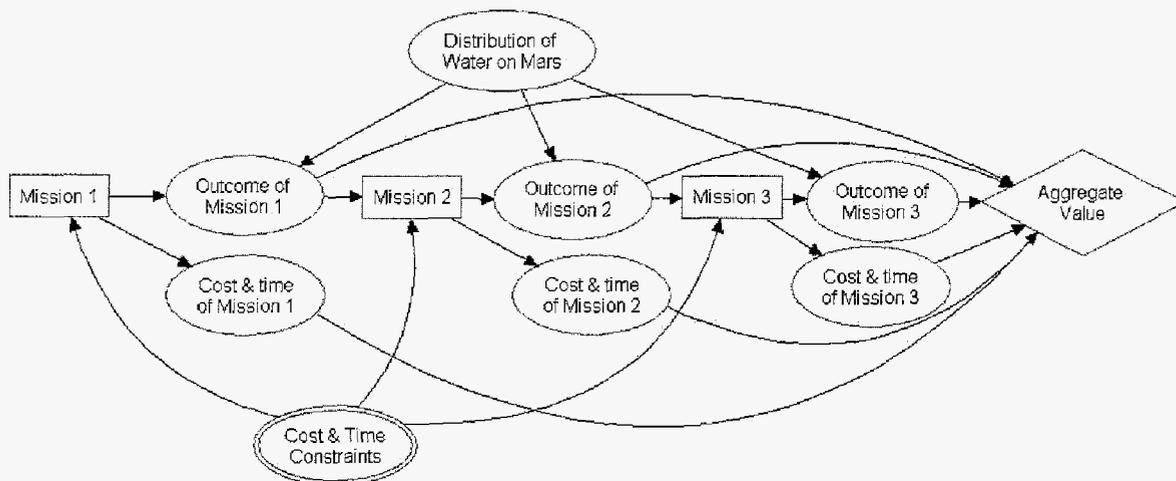


Figure 3. Influence Diagram for the Water-Search Campaign

are known.

As illustrated on Figure 3, the search for water on Mars is driven by a key uncertainty: the distribution of near-surface equatorial water on Mars. This distribution, which is highly uncertain at the beginning of the campaign, influences the outcomes of each mission. In turn, after each mission we can update the state of knowledge about the water distribution using Bayesian updating (discussed in detail in the next section). Each mission also increases the total campaign cost and time to complete, and influences all other FoMs. The final campaign value is a function of these aggregate mission FoMs, together with the final state of knowledge on the water distribution.

#### 4. MODELS FOR DECISION ANALYSIS INPUTS

Once the problem formulation has been laid out, quantitative estimates of uncertainties and alternatives must be developed as numerical inputs to the decision analysis problem. The required amount of modeling is very dependent on the campaign problem. As an illustrative case, this section will describe the models used for the search-for-water case.

##### Modeling uncertainties

As illustrated in the ID (Figure 3), the key uncertainty is the distribution of equatorial water on Mars. Current knowledge provides for averages over large regions only. Combining data from the Mars Odyssey and Mars Global Surveyor (MGS) orbiters, full coverage is available for the water concentration, induced from hydrogen concentration, in the top 1-m of the surface, averaged over pixels of size 5 deg longitude by 5 deg latitude [ref\_watermap]. In some areas, information is available on pixels of size 2 deg x 2 deg [water\_rapp]. Each such pixel corresponds to 120 km x 120 km region of Mars. Whether the distribution of water inside each region is uniform, implying a small concentration everywhere, or widely varying, with the presence of subsurface “lakes”, is important for ISRU. It is also important for discriminating between stationary and mobile platforms when searching for Martian water.

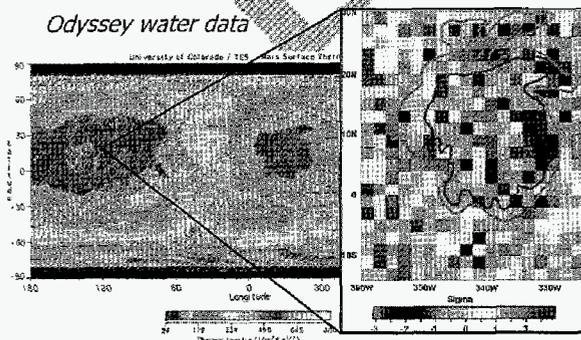


Figure 4. Initial State of Knowledge [water\_rapp]

*Geographical extents*- The relationship between the water distribution and the range of exploration of various mission platforms can be captured by defining several “geographical extents”. The chosen definitions are summarized in Table 2.

Table 2. Geographical Extents (“pixels”)

Pixel	Side Size (deg)	Side Size (m)
Region	2	120,000
Area	0.2	12,000
Site	0.002	120

*Concentration levels* -The second dimension of water distribution is the *concentration* of water within the top 5-m of the subsurface in each “pixel”. For simplicity, this dimension is quantized as well within the problem. Three levels of water concentration are defined: low (ISRU not possible), medium (ISRU possible), and high (very interesting for ISRU). Baseline threshold values for these levels are summarized in Table 3.

Table 3. Water Concentration Levels

Level	Definition	Ref. average
Low (L)	<5%	2.5%
Medium (M)	>5%, <20%	10%
High (H)	>20%	25%

In this model, the state of knowledge about water distribution can be described with two parameters: the water distribution already found; and the percentage of the planet already covered.

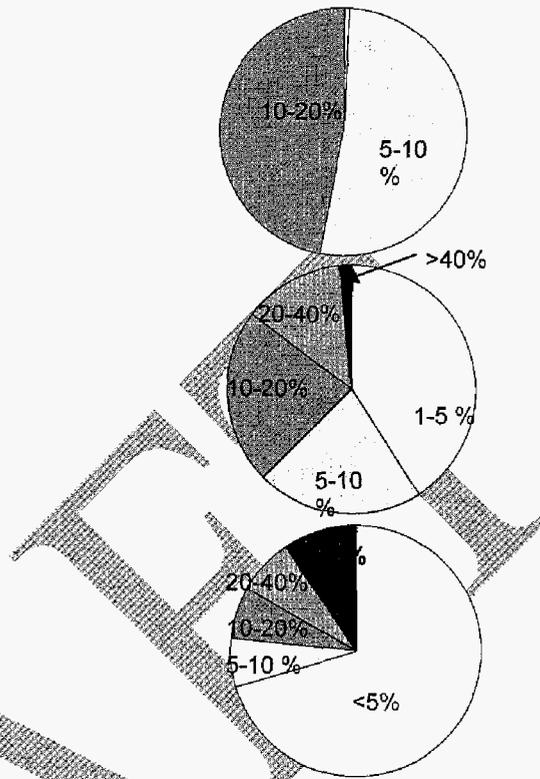
The water distribution already found is a vector of size 3 (number of “pixel sizes” in Table 2), where each element can take 4 values (per Table 3):

- “NS” if no search was carried out at that level of resolution (such is the case at the “area” and “site” level at the beginning of the campaign),
- “L” if all explored pixels of that size have been found to have low water concentration,
- “M” if at least one pixel has been found with medium concentration, but none with high concentration (such is the case at the “region” level at the beginning of the campaign),
- and “H” if at least one pixel with high water concentration was found.

The coverage parameter is a matrix of size 3 x 3 per the number of geographical extents. Element (i,j) in the matrix is the percentage of all pixels of size (i) on the planet that have been explored already, with resolution (j). At the beginning of the campaign, the whole planet has been explored at the “region” level.

*Water distribution-* Once the average water concentration is known over a pixel, the *distribution* of water within the pixel can be modeled as a probability density function; this is a convenient way to capture the likelihood that a randomly chosen “small pixel” within the “large pixel” would have a certain water concentration. In this study, a beta distribution was chosen to model water distribution. This distribution has the benefits of allowing only values between 0% and 100%; and of taking any possible shape by changing its parameters. With a beta distribution, the water distribution is characterized by the average and the standard deviation of the water concentration. The reference average values chosen for each level (L, M, H) are summarized in Table 2. The standard deviation is an unknown of the problem. Therefore, standard deviation must be one of the uncertainties captured in the state of knowledge. Figure 5 illustrates 3 representative water distributions, corresponding to 3 representative standard deviations of the water concentration. Lacking further information, baseline inputs assume these three distributions have equal likelihood.

*Correlation-* Another important parameter to estimate the value of mobility is the *correlation* between adjacent pixels. If a site has low water concentration, how likely are neighboring sites to also score “low”? This effect is captured by a correlation factor between adjacent sites. The factor is unknown at the beginning of the campaign. It can depend on the geographical extent, since sites are more likely correlated than regions. It can also depend on the standard deviation of the water distribution, i.e. its overall shape (uniform or “lakes”). Therefore, a distribution/correlation matrix is the third and last parameter describing the state of knowledge about water distribution. For simplicity, the model captures only three possibilities for the standard deviations (low, medium, and high as illustrated in Figure 6) and two possibilities for correlation (low and high). The matrix is of size 6 (for the number of distributions captured) x 3 (for the number of “pixel” sizes). Element (i,j) in the matrix is the likelihood of distribution (i) over geographical extent (j), given the current state of knowledge.



**Figure 5.** Representative Water Distributions with (a) low, (b) medium and (c) high standard deviation

**Table 4.** Initial Stage State of Knowledge assumption on water distribution. For each “pixel size”, the table shows 6 possible standard deviation/correlation pairs, with their probability occurrence.

Region-level						
St Dev.	25% (l)		50% (m)		100% (h)	
Correl.	50%	40%	30%	70%	60%	50%
Proba.	17%	17%	17%	17%	17%	17%
Area-level						
St Dev.	25%		50%		100%	
Correl.	70%	60%	60%	90%	80%	70%
Proba.	17%	17%	17%	17%	17%	17%
Site-level						
St Dev.	25%		50%		100%	
Correl.	95%	85%	75%	99%	95%	85%
Proba.	17%	17%	17%	17%	17%	17%

*Mission outcomes-* Each mission has four possible outcomes: disaster (no science return), “L” (over the resolution searched, only pixels with low concentration are found), “M” (at least one pixel with medium concentration) and “H” (at least one pixel with high concentration). Given no mission disaster the likelihood of “L”, “M” and “H” is calculated in three steps.

The first step calculates the likelihood of each outcome at the first pixel explored as a function of the average water concentration over the region explored, the standard deviation of this water concentration, and the percentage of the region that has already been explored (assuming that we don't go twice to the same area). The average water concentration and its standard deviation determine the number of "L", "M", and "H" pixels over the region explored by the mission, by using the beta distribution. The likelihood of each find is therefore simply a "find 1 out of N" problem.

The second step takes correlation into account to determine the likelihood of each find at each following pixel. A Markov chain provides a good model of this problem, as Figure 7 illustrates. In this figure, Pc is the correlation factor, and Pi is the likelihood that any random site has a concentration i (where i is L, M or H).

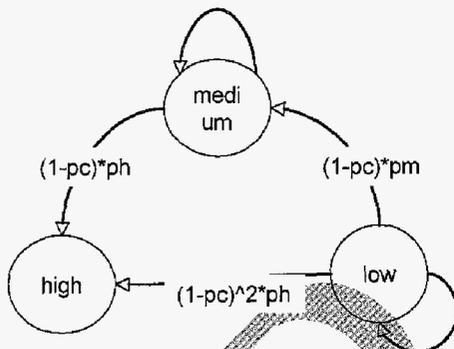


Figure 6. Markov Chain for pixel correlation

After visiting N sites, the vector of outcome probabilities is the product of the vector of outcome probabilities after the first sites, by the Nth power of the Markov transition matrix.

**Table 5. Markov Transition Matrix**

1	(1-Pc)*Ph	(1-Pc) <sup>2</sup> *Ph
0	1-(1-Pc)*Ph	(1-Pc)*Pm
0	0	1-(1-Pc)*Pm-(1-Pc) <sup>2</sup> *Ph

The third step repeats this probability calculation over each possible water distribution (standard deviation and correlation pair). The likelihood of each pair given the current state of knowledge is used to calculate the "expected value" of each outcome probability.

*Bayesian updating*

A key component of uncertainty modeling for "search" campaigns is to update the state of knowledge based on each mission find. For example, if an Orbiter mission finds all sites to be "medium", this suggests that the standard deviation of the water distribution is lower than initially expected; this needs to be taken into account before

evaluating the next mission, as it will diminish the value of mobility and favor stationary platforms. In order to update the probability distribution over the distribution of water on Mars, Bayes' rule is used. Denoting Pc as the correlation factor, S the standard deviation, j the index of each possible standard deviation/correlation pair, and K the outcome observed, and noting that P(x|y) means the probability of outcome x given outcome y:

$$P(S_j, P_{c_j} | K) = P(K | S_j, P_{c_j}) * P(S_j, P_{c_j}) / P(K)$$

where the *a-priori* probability of outcome K was, as explained above is:

$$P(K) = \sum_j \{ P(S_j, P_{c_j}) * P(K | S_j, P_{c_j}) \}$$

The other two components of the state of knowledge are easily updated. The coverage of the planet is the sum of the previous coverage at that resolution, and the coverage of the current mission. The "best" find at each level of resolution is a simple comparison between the previous find, and the current mission outcome. For example, if only low concentrations were previously found at the area level, but the latest mission explored with area resolution and found "high", then the area find is updated to "high".

By changing the numerical values, and the number of quantification "bins", the uncertainty model described in the above is general to any "search" campaign. The same model could be used as a start to consider searching for resources on the Moon, or for life on Mars, etc. But a campaign with totally different goal, such as emplacing a power and communications infrastructure to prepare for a human visit, would be faced with a very different set of uncertainties. It would require the development of new and different uncertainty models.

*Modeling alternatives*

In a campaign problem, the alternatives are various mission options. There are three components in modeling the alternatives: (1) defining the set of available options, (2) defining the information required on each mission option and (3) filling out the matrix of mission information for each option with numerical values.

*Defining options-* The set of mission options is typically determined by expert knowledge, consultation of mission study databases, or brainstorming sessions. In order to keep the problem manageable and to understand the results, the initial set of options should be small. If necessary, options can be initially grouped in families.

In this case, we define five initial families: Orbiter, Lander, Rover, Aerial (platform) and Network (of small landers). Only those families that will prove interesting after a first round of analysis will be worth breaking down into several design options.

**Table 6. Initial Mission Information (cost not shown)**

Mission	Orbiter	Aerial	Lander	Network	Rover
Explores	Planet	Region	Area	Area	Area
#platforms	1	1	1	3	1
Resolution	Region	Area	Site	Site	Site
Coverage	50%	0.5%	0.01%	0.01%	0.8%
Error Factor	1.5	3	1	1	2
P(disaster)	5%	20%	10%	10%	10%
Error Factor	1.5	3	2	2	1.5
Add <sup>al</sup> Meas.	2	2	1	1	3
Time (months)	52	26	26	26	52

*Defining information-* The information to capture for each mission option flows from the campaign goals and constraints on the one hand, and the uncertainty model on the other hand.

The information necessary to calculate campaign FoMs flows directly from these FoMs. In our case it consists of: mission cost (for total campaign cost), time to complete the mission (for time to complete the campaign), mission coverage and resolution (for campaign coverage at each resolution), and number of additional human precursor measurements.

The information required for the uncertainty model includes the likelihood of a mission disaster (launch, cruise, arrival, or payload failure), and the metrics of performance of the mission: coverage, resolution, and number of resolution units explored (a measure of mobility).

*Numerical estimates-* Whether mission options are all included or grouped in families, coming up with one numerical estimate for each piece of mission information is a challenge. In reality, such estimates always include an error bar. Instead of running the analysis using one “expected” value, it is best for the first round of analysis to keep numerical ranges of values for all mission information. Sensitivity analysis results will indicate which ranges are too large to conclude on a best campaign strategy.

The initial mission information for the water-search campaign is summarized in Table 4 (except for the cost information). The “error factor” applies to the baseline value for estimation of the upper (multiplication) and lower (division) bounds of the “error bar”. In some cases, such as the aerial platform, the error bar is a-priori very large. This reflects a new concept, for which no detailed design study has been carried out yet.

This set of mission information is typical of any campaign. Cost, time to complete, and probability of mission disaster will always be important. In addition, every campaign will

have campaign-specific metrics, such as the number of additional human precursor measurements in this case. These metrics need to be defined on a case-by-case basis, with a careful consideration of the campaign goals.

## 5. PROBLEM SOLVING WITH DECISION TREES

### *Decision Trees*

The power of Influence Diagrams for problem formulation was described in section 3. Decision trees are another representation of decision problems. They are a diagrammatic representation of all the possible realizations of all the possible strategies. They provide a straightforward way to calculate the best strategy through dynamic programming. Decision trees illustrate a sequence of decisions and uncertainties, presented in the tree in the order that the decisions will be made and the outcomes will be discovered.

Decision analysis uses the same conventions for the nodes in the tree that in the influence diagram: squares represent decisions, ovals represent uncertainties, and a polygon, in this case a triangle, represents the values or FoMs.

*Decision nodes-* The branches extending from a decision node represent each of the alternatives available at that decision point. For example, in Figure 8, there are five alternatives for the first mission: Orbiter, Aerial, Lander, Network, or Rover. Each alternative may be associated with a value. In this case, each alternative mission is associated with its cost. The value is presented in a box just below the name of the alternative. We use variables rather than fixed values to increase the flexibility of the tree. Thus, below “Orbiter” you see a box with “Cost\_Orbiter”.

*Uncertainty nodes-* The branches extending from an uncertainty node represent the possible outcomes of the uncertainty. For example, in Figure 8, there are four possible outcomes resulting from sending an orbiter: Disaster, Low, Medium, and High. Each outcome has a probability, which is placed below the name of the outcome. Again, the probabilities are represented as variables here, with

P\_Dis\_Orbiter representing the probability of a disaster with an orbiter.

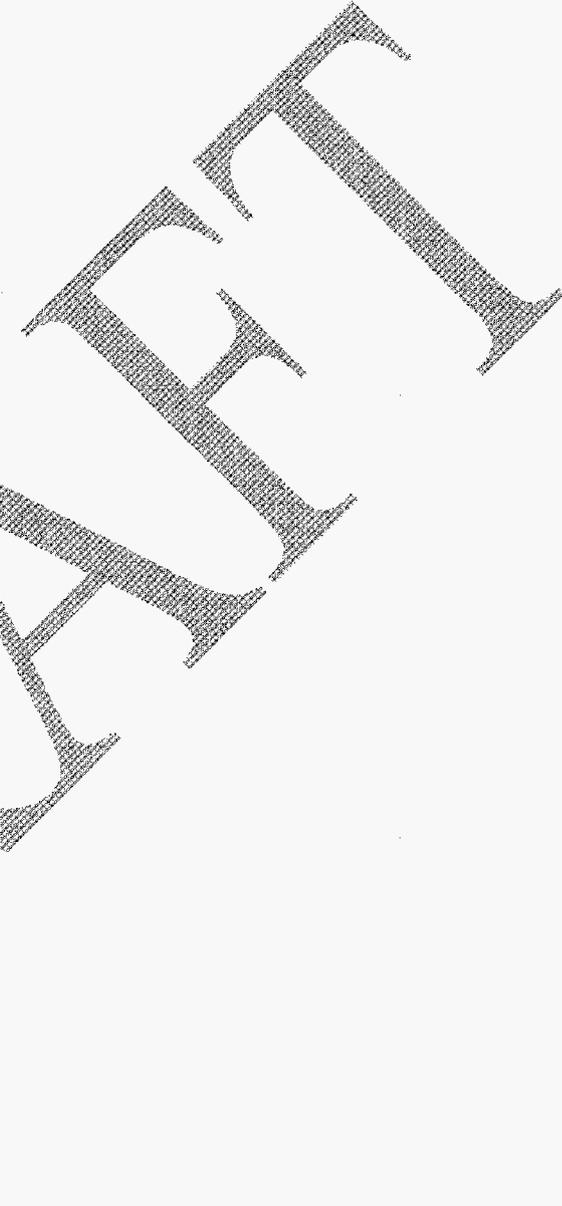
*Terminal nodes-* At the end of each terminal branch there is a terminal value node. Here the aggregate value is equal to the value of the amount of water that has been found minus the sum of the costs of all missions that have been performed to get to that point. For example, the value associated with the second terminal node is equal to the value of finding a Low site minus the cost of the orbiter minus the cost of the lander. In this figure, the formula has been shortened for display clarity. The real value includes all the FoMs with their associated weights as summarized in Table 1.

Note that the Bayesian updating and other probability calculations take place outside the tree. The probability that a Lander finds a high site depends on the current set of probabilities over the distribution of water, as well as on the mission information on resolution and coverage.

Given the number of possible mission options and the number of possible outcomes for each option, the decision tree for a campaign problem rapidly “explodes”. As an illustration, Figure 7 shows only a small extract of the decision tree for our case study.

*Rolling back the tree: general principle*

A basic Decision Analysis method, “rolling back the decision tree” is the operation by which the information input into a DT can be translated into an optimal campaign strategy. There are two essential operations to rolling back the tree: calculating expected values and choosing among options. The operation starts from the end points and works backwards, calculating expected values for all uncertainty nodes and choosing the option with the highest expected value at all decision nodes. Figure 8a illustrates the principles in a very simplified decision tree with example values and probabilities.



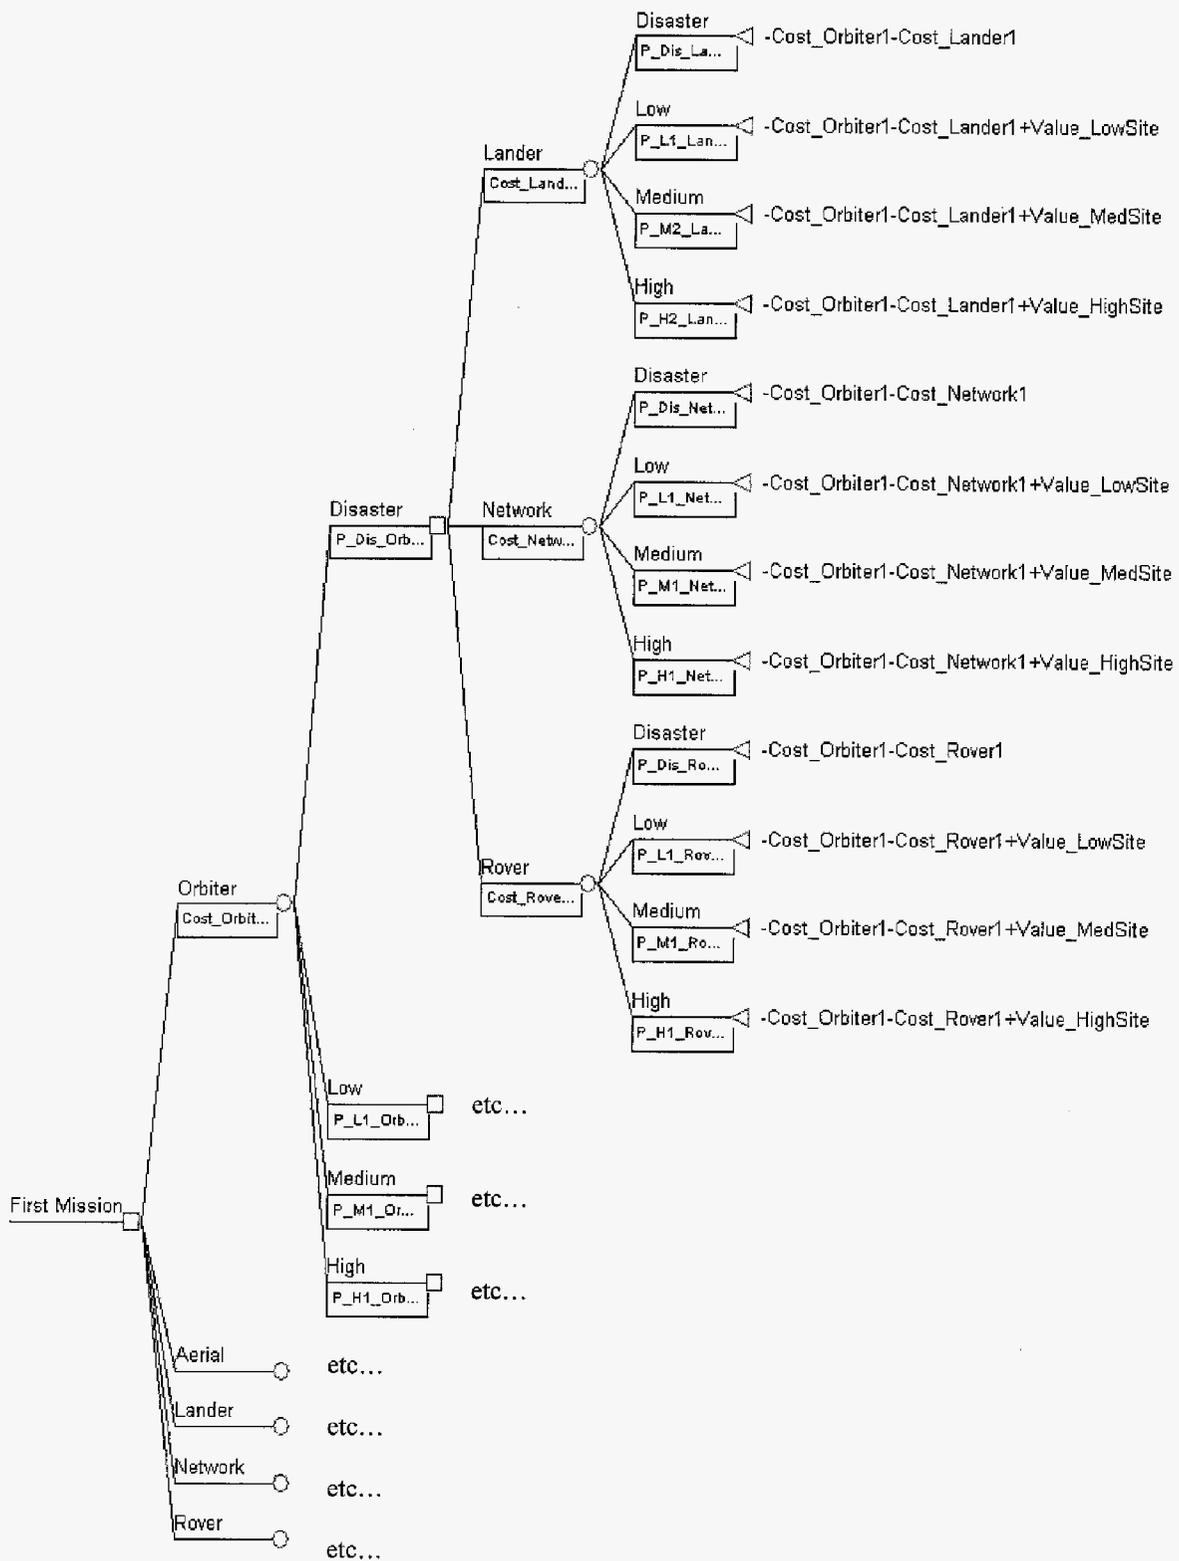


Figure 7. Decision Tree Extract

*Calculating Expected Outcomes-* The first step calculates the expected values at each of the uncertainty nodes that precede the terminal value nodes. The expected value after each uncertainty node is calculated by taking the product of the probability of an outcome with the aggregated value for that outcome, and then summing these products over all outcomes of the uncertainty node. For example, the expected value of the Lander following an Orbiter disaster is  $0.8*(-2) + 0.2*(8) = 0$ . The expected values are shown in the shaded boxes below the probabilities.

*Choosing best decisions-* The next step is to choose the best alternative at each decision node preceding the uncertainty nodes. For example, following an Orbiter Disaster, we would choose the alternative (Lander, Network, or Rover) with the highest expected value - in this example, we would choose Rover with an expected value of 2. This value becomes the expected value of a Disaster. The tree can then be rolled back another level. For example, the expected value over all the outcomes of the initial Orbiter mission equals  $0.2*2 + \dots + 0.1*14 = 7.1$ .

*Identifying the best strategy-* The rolling back continues

until we end up with a single *strategy*. The strategy will consist of a first decision, and then later decisions for each possible outcome. An example strategy is presented in Figure 8b. In this strategy an Orbiter is followed by a Network if it does not find a medium or high water site. The program stops after the network, except in one case. If the Orbiter has a disaster and the Network finds low water, then a Rover is sent (these examples are for illustrative value; results of the actual study follow).

### Automatic Tree Generation

Given the size of the decision tree, and the need for extensive sensitivity analysis on all input parameters, manual entry of the decision tree is not practical.

If the same mission options were available at each decision point in the Influence Diagram (Figure 3), and if each mission outcome had the same probability no matter what the preceding missions, then available off-the-shelf software could translate our simple ID into a decision tree. This might be the case for some campaign analyzes.

For our example, however, both the available options and the likelihood of a mission outcome take unique values

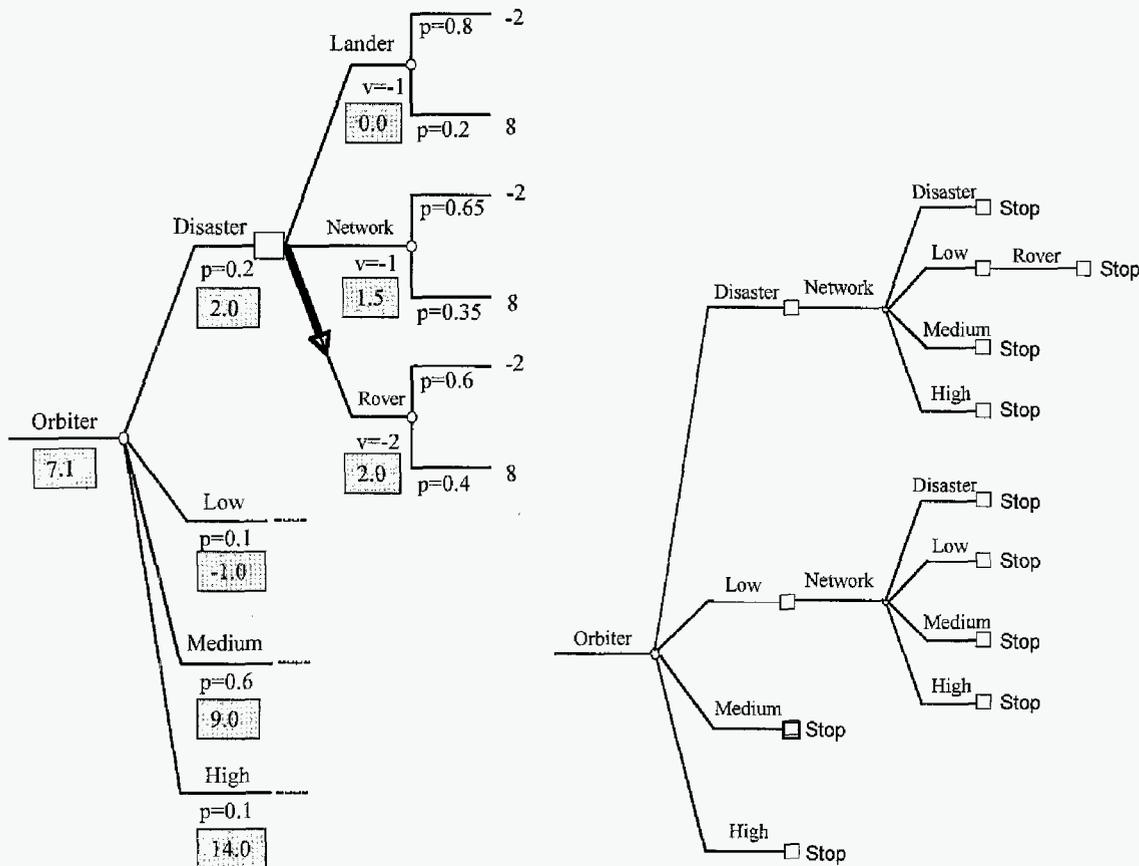


Figure 8. Example of "Rolling back" a DT: (a) Tree (b) Optimal Strategy

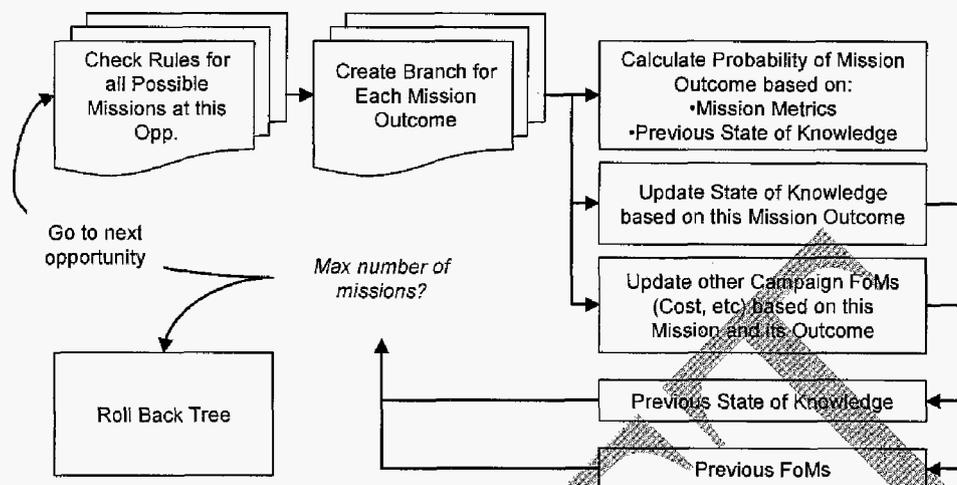


Figure 9. Automatic Tree Generation

throughout the tree. At each mission opportunity, constraints apply to restrict the available mission options: science constraints (no remote sensing mission can be last), cost constraints (if limited by a total budget or a budget profile), and programmatic constraints (such as maximum acceptable number of mission disasters, or whether the same mission would be attempted again after a disaster). In addition, the probability of each mission outcome depends on the previous missions flown and their outcomes, as described in the uncertainty modeling section. Bayesian updating needs to happen after each mission outcome.

For this application, a new routine was written for automatic generation of the decision tree based on simple rules. The routine uses an iterative algorithm at each mission opportunity. The steps of this algorithm, illustrated on Figure 9, can be summarized as follows:

1. For the mission opportunity at hand, given the cost and time already spent, and the latest find, check the campaign constraints for the available mission options.
2. For each available mission option, create an *alternative* branch. Read the mission information database to determine the possible mission outcomes.
3. For each possible mission outcome, create an *uncertainty* branch. Calculate the probability of the mission outcome based on the previous state of knowledge, and the mission performance information (as described in the uncertainty modeling section).
4. For each mission outcome, read mission information to update the campaign FoMs and use Bayesian updating to calculate the new state of knowledge.
5. Repeat for next mission opportunity, using the new values of campaign FoMs and state of knowledge as a start.
6. When no mission is available based on the tree building rules, start rolling back the tree.

The general structure of the automatic tree generation algorithm is generic to any campaign. Parts to be tailored for each campaign are the functions calculating the probability of each mission outcome; and, if different from this case, the functions updating the unique campaign FoMs based on mission information.

#### Baseline results

Table 7 summarizes the baseline results. Several conclusions can already be reached.

*Number of mission opportunities-* The first conclusion is that the overall value of the campaign increases as the number of mission opportunities increases. For the baseline FoM weights, the increased likelihood of a find, together with the increased coverage, are worth the increased cost and time.

*Ending the search-* The second conclusion is that the search should end as soon as a site with high water concentration is found. The additional coverage alone is not worth additional missions at that point. Compared to a traditional planning approach, which would assume a set number of missions, the decision-analytic approach thus saves the cost of a 3<sup>rd</sup> mission for some cases, increasing the total value of the campaign. This reflects the value of the “option to launch or not launch a third mission”.

Table 7. Baseline Results

# missions	Best strategy	Value (\$M-eq.)
1	Rover	Expected: 3460
2	Rover until find “H”	Expected: 5253 Worst exp for last mission: 1510

		Best: 9185 (39% chance, 1 <sup>st</sup> finds)
3	Rover until find "H"	Expected: 5317 Worst: -440 (2 disasters) Best: 9185 (39% chance, 1 <sup>st</sup> finds)

*Dominating mission-* The third conclusion is that for the baseline set of information and FoM weights, the Rover concept dominates all other missions. No matter what the number of mission opportunities, the optimal strategy is to send a Rover until a high-water-concentration site is found. Given that regions exist with an average medium water concentration, and given the range of a Rover, the likelihood that at least one of the sites visited by the Rover will be "H" is 39%. This is sufficient to outweigh the value of remote sensing for site selection.

*Importance of sensitivity analysis*

Sensitivity analysis is always important to understand the context of the results, as well as their robustness. In the case of campaign design, sensitivity analysis with respect to all mission information "with error bars" is particularly important. If the conclusions of the analysis change when varying parameters within their error bar, then the analysis is inconclusive; smaller error bars are required.

Figure 10 offers another snapshot of the baseline results. It compares the value of the campaign as a function of mission opportunity, and as a function of the choice of mission for the first opportunity, assuming the optimal choice of mission is made thereafter (in this case, all subsequent missions are Rovers).

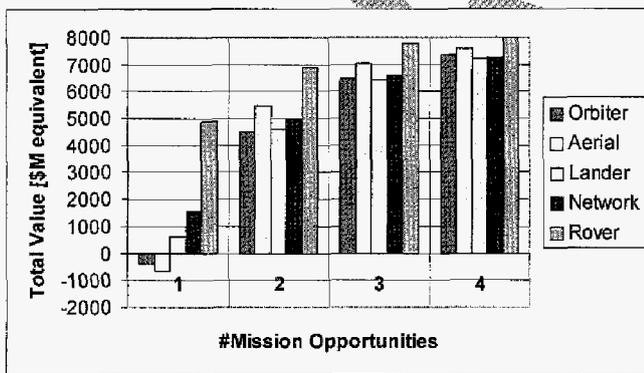


Figure 10. Baseline Results per 1<sup>st</sup> mission

The figure shows how the value of the Rover relative to the other alternatives diminishes with increasing number of mission opportunities. The closest contender to the Rover mission appears to be the Aerial platform, closely followed by the Orbiter. This suggests that the value of remote sensing as a first mission might depend on the numerical assumptions. While the Aerial platform is a new concept, the

results suggest that there might be value in developing designs and technologies for such a mission.

*Aerial versus Rover-* Figure 11 shows the results of sensitivity analysis with respect to all mission information for the Aerial platform and the Rover concepts, for the case of 3 mission opportunities. The figure represents the space of Rover information improvement versus Aerial platform information improvement. For example, +100% on the horizontal axis represents the case where all Rover metrics were increased by 100% of their error bar, i.e. the "best" end of the Rover error bar (smallest cost, highest coverage, smallest probability of disaster). The vertical axis similarly represents the error bar of the Aerial concept. For all points above the blue "Baseline" curve, the Aerial platform concept dominates as the first mission; the Rover is the optimal first mission for all points below the "Baseline" curve. The boundary curve is very close to the baseline values (0%/0%), showing that the range of possible input values is equally split between the two possible decisions. Refined mission information is required to conclude on the optimal first mission.

*Sensitivity to preferences-* Figure 11 also shows how this boundary between the two concepts changes when the FoM weights are changed. The red dashed "FOMs x10" curve corresponds to the case where weights that favor an aerial platform (coverage, area find) are multiplied by 10, and weights that favor a Rover (site find, disaster) divided by 10. This curve is very close to the baseline curve, showing that the FoM weights have little influence on the conclusion. In this case, the decision maker preferences influence the results less than the mission cost and performance information.

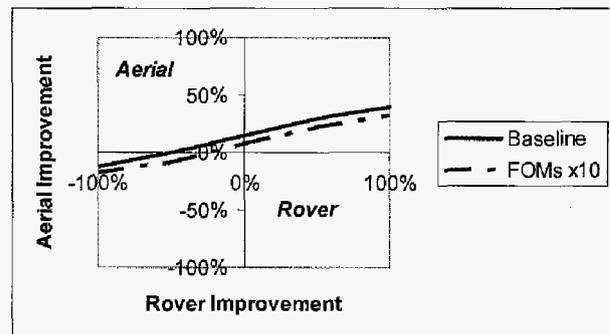


Figure 11. Aerial vs Rover: Information Space. Each axis represents changes from the baseline within the error bar of mission information. Curves represent the points above which an Aerial platform would dominate as 1<sup>st</sup> mission.

*Sensitivity to initial knowledge-* To conclude the Rover versus Aerial platform comparison, Figure 12 compares the effect of various inputs on the conclusion. The plot compares the total value of a 3-opportunity campaign

starting with an Aerial platform to that starting with a Rover, under different conditions. The “bound of mission data” conditions correspond to the “best” end of the Aerial platform mission information error bar, and the “worst” end of the Rover mission information error bar. The “bound of water distribution” conditions correspond to a case where areas are known to be highly uncorrelated, but sites are known to be very correlated; in such a case, an aerial platform is very valuable to pinpoint the right area before sending a surface mission. The figure shows that varying the water distribution (likelihood of each standard deviation/correlation pair) is the most significant driver of the comparison. While the water distribution is unknown, its impact on the best campaign strategy is actually more important than that of mission cost and performance. Any study or expert interview that would make possible a refined model of water distribution would therefore have significant value.

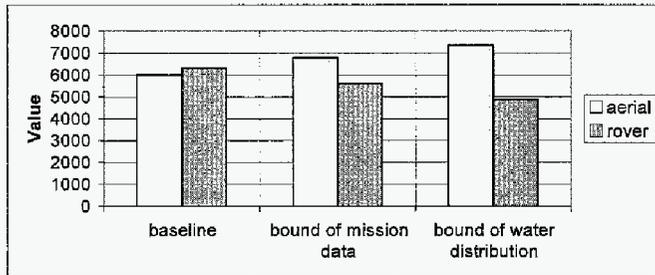


Figure 12. Aerial vs Rover Comparison Drivers

**Model refinement**

Sensitivity analysis results such as illustrated above help determine where it is worth carrying out detailed study. For example, the results just described help conclude that:

- No Customer interview to refine FoM weights is required,
- A better model of the current knowledge about Martian water distribution would have value (at least determining the believed likelihood of various standard deviation/correlation models),
- More detailed information on both mission concepts is required before concluding on the Rover or Aerial platform as a first mission, and
- After the first mission, a Rover mission is always the best strategy. Therefore, there might be value in refining the selection of possible Rover designs.

Table 8. Refinement of Rover options (cost missing)

Mission	1 solar Rover	2 Solar Rovers	1 RPS Rover
Explores	Area	Area	Area
#platforms	1	2	1
Resolution	Site	Site	Site

Coverage	2.1%	2.1% (each)	6.9%
Error Factor	2	2	2
P(disaster)	10%	10%	15%
Error Factor	1.5	1.5	1.5
Add <sup>al</sup> Meas.	3	3	3
Time (months)	52	52	78

Out of this list of study refinements, consider for example the breaking down of Rover options. Figure 12 breaks down the Rover family into three concepts with smaller error bars: 1 solar Rover, 2 solar Rovers, or 1 RPS (radioisotope power source) Rover. Compared to 1 solar Rover, the 2-Rover concept offers an increased total range, and an increased likelihood of a find through the visit to two different areas. The RPS Rover offers a higher total range than both the 1-solar-Rover and the 2-solar-Rover concept, but at the cost of increased mission duration; that mission visits only one area.

**Baseline results:** Table 6 includes the new results with this list of Rover options. All three options score better than the Aerial platform. The results suggest two interesting conclusions.

First, solar Rovers dominate over an RPS rover. For the baseline uncertainty on water distribution, the additional range provided by an RPS Rover is not worth its additional time and risk (due to the novelty of this technology for a Rover).

Second, when only one or two mission opportunities are available, a second solar Rover sent to a different area increases the probability of finding water to a degree that it is worth the cost. But if three or more mission opportunities are available, then an incremental investment is preferred. The optimal strategy becomes to send only 1 solar Rover at a time, until a high-concentration site is found. Only at the last mission opportunity does it become worthwhile sending a 2-Rover mission.

Table 9. Baseline Results

# missions	Best strategy	Value (\$M-eq.)
1	2 solar Rovers	Expected: 4182
2	2 solar Rovers until find “H”	Expected: 6142 Worst exp for last mission: 2232 Best: 9185 (42% chance, 1 <sup>st</sup> finds)
3	1 Solar Rover until find “H”, or 2 solar Rovers for last mission	Expected: 7163 Worst: 845 (2 disasters) Best: 9635 (36% chance, 1 <sup>st</sup> finds)

*Non-trivial strategies-* For 100% improvement in the Aerial platform information, the Aerial platform is the optimal first mission. In this case, the optimal strategy becomes interesting. The model results recommend repeating Aerial platforms until a high-concentration area is found, then send 1 solar Rover to that area. If no high-concentration area is found before the last mission opportunity, then 2 solar Rovers should be sent to the highest-concentration area found so far. As in the baseline results, the added value of the second Rover is worth the cost only in some circumstances. The ease with which such conclusions can be reached and quantified shows the power of a decision-analytic approach.

*Understand design drivers-* Figure \*\* shows the total campaign value as a function of number of mission opportunities, and as a function of the mission chosen for the first opportunity (solar Rovers being optimal thereafter). The figure shows that the value of the three Rover concepts is actually very close. This suggests the importance of sensitivity analysis with respect to Rover information, so as to understand the drivers of Rover value. Also, using our baseline values, it appears that sending 2 Solar Rovers as a first mission is robust to the number of mission opportunities: it is not the optimal first choice if there are three opportunities, but the value is extremely close to that of the optimal. Thus, if there is some uncertainty over the number of mission opportunities that will be available, sending 2 Solar Rovers would be the best response.

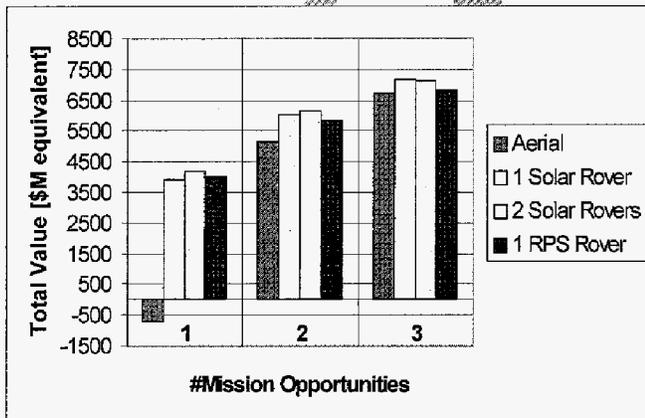
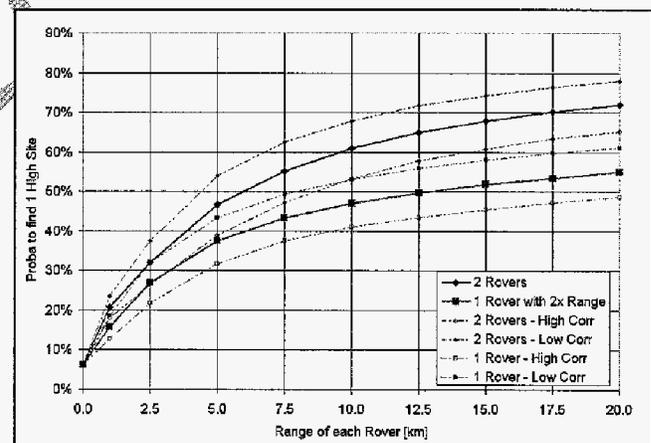


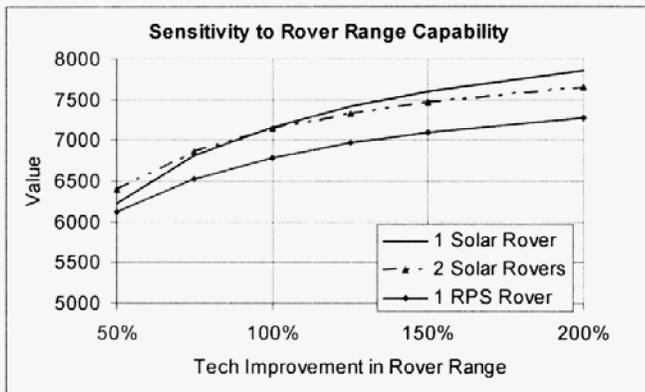
Figure 13. Rover Options Comparison

*Impact of Range on Performance-* As an example study of design features that drive mission value, Figure 14 shows the impact of Rover range. Based on the probability model described in Section 3, the first plot shows how the probability of a “high” find varies with Rover range capability. Several curves compare a 1-Rover concept with a 2-Rover concept where each Rover would have only ½ the range (same total range). Dashed curves show how the results change when sites have, respectively, low or high

correlation. From this plot, it becomes apparent that additional Rover range offers diminishing returns beyond 15-20 km. This explains the relative value of solar Rovers compared to the RPS Rover concept. The plots also show that this threshold range is insensitive to number of Rovers, and to level of site correlation. Finally, the figure illustrates the value of “breaking down the range” between two separate Rovers, instead of increasing the range of one Rover. For the same total range, the 2-Rover concept always offers a better likelihood of finding water, because it visits two different areas.

*Value of Rover Range-* The second graph shows how these results translate in terms of total campaign value. For each Rover concept, it plots the campaign value as a function of changes in Rover range capability (percentage with respect to baseline value). In this plot, the effect of “diminishing returns” in performance are combined with the other campaign FoMs, in particular cost. Because its total range is smaller, the 1-solar-Rover concept would benefit more from range improvement than the 2-solar-Rover concept, which is already in the area of diminishing returns. No matter what the range capability, the 1-solar-Rover concept is a less expensive mission. As a result, there is a range beyond which the 1-Rover concept becomes optimal: the added performance of a second Rover does not justify its cost. This cross-over happens for about 100% improvement in Rover range capability. The value of this capability increase is several \$100M, given the baseline FoM weights and the baseline mission cost estimates.





**Figure 14.** Effect of Rover Range (a) on likelihood of find, (b) on campaign value

### Study Conclusions

*Direct conclusions-* Even with limited available information and large error bars on mission information, the decision-analytic approach helped gain a number of insights into the search for water on Mars. In particular, the following conclusions can be reached without further study:

- A Rover is the best platform to find the water.
- If the water find is valuable, it is worth spending several mission opportunities searching for water. Staggering the investment is less expensive than sending many Rovers on one opportunity.
- Depending (1) on the disparity in the distribution of water and (2) on mission cost/performance, an aerial platform can provide valuable remote sensing, pinpointing the location where it is best to send the Rover mission.
- Other missions are dominated over the whole range of possible input values. Orbiters are not very interesting because regional-level information is already available; only an Orbiter that could improve on the resolution available from MGS and Odyssey would add valuable information. Landers and Networks are dominated because of the large advantage of mobility; between 0 and 10-15 km, added mobility significantly increases the likelihood of a find.
- The value of additional Rover range diminishes beyond 15-20 km. There is value in technology development to increase solar rover range up to that point.

*Focusing future studies-* Furthermore, this high-level study helps reach non-trivial conclusions as to what information is worth more work. Fully detailing all possible mission concepts, and accurately estimating their cost and performance, would be a daunting task. Instead, the results help focus future work towards the following areas:

- Refined models are needed for the possible water distributions over different geographical scales. For each model, a thorough expert survey should be carried

out, so as to conclude on the consensus “likelihood” of each model. This is a particularly interesting insight, as uncertainty models are not a part of traditional space program planning. This analysis shows that they should be, since they drive the optimal exploration strategy.

- Mission information for Rovers (both solar and RPS concepts) and for Aerial platform are worth refinement. Together with a better water distribution model, they would help conclude on the optimal mission strategy. They would also help conclude on the value of technology investment into aerial concepts.
- For the baseline campaign FoMs, the time to achieve its full range appears to penalize the RPS-Rover concept. In reality, that Rover would already achieve a valuable range in a shorter time. In order to compare several Rover concepts fairly, the relationship between their lifetime and their performance needs to be taken into account. A simple Markov model of the Rover operational lifetime, including failure rate and daily range, would capture this effect adequately.
- The baseline results suggest that it is worth spending numerous mission opportunities on the search for water. In reality, significant Programmatic and Cost constraints would limit the possibilities. These constraints need to be well understood. For a smaller cost or schedule cap, a different campaign strategy would become optimal.

The modeling described herein was kept simple, so as to illustrate the various concepts of decision analysis. Before answering the search-for-water question, the study would benefit from improvements in a number of other areas. First, the form of water should be taken into account; subsurface ice would be more valuable than hydrous minerals. Second, a number of other mission concepts are possible and should be evaluated. Finally, the value of other technology investment (besides Rover range) should be considered.

### Applicability to other Campaigns

General and specific points of this analysis

## 6. CONCLUSIONS

Through the example of a robotic campaign to search for exploitable equatorial water on Mars, this paper demonstrated the benefits of decision analysis techniques to space program planning. The example is illustrative both of the additional work required beyond traditional space mission design, and of the benefits that can be gained from that work.

Applying decision analysis to space exploration campaigns presents opportunities to improve both space program planning and decision analysis theory.

Traditional space program planning deals primarily with requirements definition and cost estimation on a one-mission-at-a-time basis. These tools are insufficient when planning a discovery-driven space campaign, where future missions depend on the uncertain results of initial missions. With Influence Diagrams, the decision analytic approach forces the study team to explicitly formulate the campaign design problem in terms of decisions, alternatives, and uncertainties. This problem formulation step already generates valuable insights. ID solving tools or Decision Tree analysis then provide a tool to solve the decision problem in terms of an optimal *strategy*. Only initial estimates of mission information and uncertainties are required for the first round of analysis. Extensive sensitivity analysis helps eliminate a large portion of the trade space, and *guide* future studies. It identifies the pieces of mission information, and the parts of the uncertainty model that matter most to the campaign decisions. It also provides a basis to determine the required level of resolution for each model.

Decision Analysis is most applicable for high-impact decisions involving significant investment, high complexity and elements of uncertainty. In fully developing and applying the decision analysis to space exploration campaigns, we are opening up a very large opportunity to apply and improve decision analysis. Probably more than any other application, space exploration campaigns push the limits of current decision analysis theory in two directions. First, in the valuation method, with the challenges of valuing mission outcomes where mission are primarily science-based. Second, space exploration campaigns offer complex problems where there is a large uncertainty about every single piece of information in the model; there are interesting research opportunities into techniques to rigorously deal with this uncertainty without limiting or invalidating the insights gained from the analysis.

**Shortcomings / Caveats**

**Potential for Future Work – Other Fields of Application**

## 7. REFERENCES

[lamassoure] Elisabeth Lamassoure, Joseph Saleh, Daniel E. Hastings. "Space Systems Flexibility Provided by On-Orbit Servicing: Part 1" and "Part 2". *Journal of Spacecraft and Rockets*, Vol 39, No. 4, pp561-570, July-August 2002

**[saleh]\*\***

[shishko] Robert Shishko Donald H. Ebbeler and George Fox. "A Real-Options Approach for NASA Strategic Technology Selection". *Systems Engineering*, Volume 7, Issue 1, Pages 1 – 13, 2003.

[morse] Elisabeth L. Morse, Tracy Leavens, Babak Cohanim, Corey Harmon, Eric Mahr, Brian Lewis. "Next-Generation Concurrent Engineering: Building Models to Complement Point Designs". *IEEE Aerospace Conference 2006*, Big Sky, Montana, March 3-11, 2006.

**[Dixit, Pyndyck] Avinash K. Dixit and Robert Pyndyck  
Investment and Real Options**

**[mavr] Andrew S. Mavr "Real Options"**

[isru\_1] MEPAG (2005), "Mars Scientific Goals, Objectives, Investigations, and Priorities: 2005", 31 p. white paper posted August, 2005 by the Mars Exploration Program Analysis Group (MEPAG) at: <http://mepag.jpl.nasa.gov/reports/index.html>.

[isru\_rapp] Donald Rapp, Jason Andringa, Robert Easter, Jeffrey H. Smight, Thomas Wilson, Larry Clark and Kevin Payne. "Preliminary System Analysis of Mars ISRU Alternatives. FY04 Report". JPL report D-31341, October 2004.

[keeney] Ralph Keeney. Value-Focused Thinking. Harvard University Press, Cambridge Massachusetts. 1992\*\*

[ID\_citations] John Diffenbach. "Influence Diagrams for Complex Strategic Issues" in *Strategic Management Journal*, Vol 3No2 April-June 1982 \*\*

[owens shacter nease] Douglas K. Owens, MD, msc, Ross D. Shachter, phd, Robert F. Nease, Jr., phd, "Representation and Analysis of Medical Decision Problems with Influence Diagrams," in *Medical Decision Making*, Vol 17, No. 3, Jul-Sept 1997.

[ref\_watermap] W. C. Feldman, T. H. Prettyman, S. Maurice, J. J. Plaut, D. L. Bish, D. T. Vaniman, M. T. Mellon, A. E. Metzger, S. W. Squyres, S. Karunatillake, W. V. Boynton, R. C. Elphic, H. O. Funsten, D. J. Lawrence, and R. L. Tokar, "Global distribution of near-surface hydrogen on Mars", *Journal of Geophysical Research*, Vol. 109, E09006, doi:10.1029/2003JE002160, 2004.

[water\_rapp] Donald Rapp. "Accessible Water on Mars". JPL report 31343, December 2004.

## 8. ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration **Contract with Aeronautics**

Corp. Significant contributions were made by Don Rapp, Eureka Team, Matt Olfert, and others.

## 9. BIOGRAPHIES

**Prof. Erin Baker** is Assistant Professor of Industrial Engineering and Operations Research in the department of Mechanical and Industrial Engineering at University of Massachusetts, Amherst. Her research is on Decision Making Under Uncertainty and Learning, and has been applied primarily to the Climate Change Policy. She has published in journals such as *Operations Research*, *Resource and Energy Economics*, and *Climatic Change*. Prof. Baker was a Summer Faculty Fellow at the Jet Propulsion Lab during the summer of 2005. She received her Ph.D. from the department of Management Science and Engineering at Stanford University in 2002.



**Elisabeth Morse** (previously Elisabeth Lamassoure) is a member of the Mission Systems Concepts section at the Jet Propulsion Laboratory, specializing in collaborative engineering and pre-formulation phase studies. She is currently the lead Systems Engineer for JPL's Advanced Projects Design Team (Team X). Over her 4 years at JPL, Elisabeth was a systems engineer on more than 50 Team X studies, from Earth orbiters to Mars Sample Return missions. She also participated in Mars Scout and New Frontiers proposals, Mars Program concept studies with emphasis on human precursor missions, and probabilistic risk assessment (PRA) for the Mars Sample Return planetary protection program. Elisabeth has an M.S. equivalent from Ecole Polytechnique, France, with Physics major, and an M.S. in Aeronautics and Astronautics from the Massachusetts Institute of Technology (MIT).



**Dr. Andrew Gray** is a

\*\*Andrew pix-bio

**Robert Easter** is a

\*\*Bob pix-bio

End of File

