OPTIMIZING TECHNOLOGY INVESTMENTS: A BROAD MISSION MODEL APPROACH

Robert Shishko, Ph.D.*
Jet Propulsion Laboratory, California Institute of Technology
Pasadena, CA 91109

ABSTRACT

A long-standing problem in NASA is how to allocate scarce technology development resources across advanced technologies in order to best support a large set of future potential missions. Within NASA, two (orthogonal) paradigms have received attention in recent years: the real-options approach and the broad mission model approach. This paper focuses on the latter. Two broad mission models are developed for Mars Science Laboratory (MSL)-type missions—a large mobile rover/laboratory versus a fixed laboratory with a small "fetch" rover. Two technologies that are critical to the amount of science returned make up the technology portfolio. Within each mission model, the technology program manager (TPM) maximizes the science return by allocating a technology development budget and controlling reserves across the two technologies. The TPM must ultimately choose between a higher science return and a higher probability of development success for the technology portfolio. The paper concludes with prospects for implementing the broad mission model approach.

INTRODUCTION

The work reported in this paper is part of NASA's effort to improve the allocation of technology development resources. The focus here is on Mars robotic exploration technologies that have achieved the proof-of-concept stage of development, but have not yet been demonstrated in a space environment. (Those familiar with NASA's Technology Readiness Level (TRL) scale will recognize this range as TRL 3 through TRL 6.) The overall objective of this work is to create a consistent, rigorous, and risk-based approach to guide the Mars Exploration Program in selecting its technology investment portfolio (TIP).

Within NASA, two (orthogonal) paradigms for TIP selection have received attention in recent years: the real-options approach and the broad mission model approach. Since the real-options approach has been documented elsewhere[1,2], this paper demonstrates the broad mission model approach. The real-options approach focuses on one technology investment and its application to missions that potentially might use it. The question the real-options approach addresses is: "How much is it worth for the right to undertake this technology development?" The broad mission model approach focuses on one broad mission concept and all the technology investments that could potentially be made. The question the broad mission model approach addresses is: "How should limited technology development resources be applied to improve that mission's outcome?" Clearly, the two approaches answer different questions, yet both could be used for TIP selection.

This paper applies the broad mission model approach to two Mars Science Laboratory[†] (MSL)-type missions—a large mobile rover/laboratory versus a large fixed laboratory with a small "fetch" rover to gather rocks and other samples. One conclusion is that even a modest change in mission concept like this requires a completely different mission model.

OPTIMIZATION IN MODEL 1

In the first model, a long-range rover/ laboratory acquires samples and makes measurements on them at a number of widely separated science investigation sites on Mars. There are two technologies that are candidates for further investment—a sample acquisition technology that decreases acquisition and handling time, and a roving technology that increases the average velocity over Martian terrain.

The objective function to be maximized is the science measurement rate, $\dot{M}$, which depends on the performance achieved by these two technologies. This objective function, given in

[*] Principal Systems Engineer/Economist

[†] MSL is a NASA mission, slated for launch in 2009, to land laboratory instruments on Mars with capabilities to enable a significant breakthrough in the area of astrobiology.

Eq.(1), was developed by Professor Dave Miller and his student Julien Lamamy at M.I.T. The maximization is subject to three constraints that I have added. The first two are mission constraints: reliability and site diversity; the third is the technology development budget constraint, i.e., the proposed cost of developing the two technologies should not exceed a fixed budget.

The reliability constraint is stated in terms of total distance travelled, including travel within a site prior to final approach to a sample. The site diversity requirement is stated in terms of an area over which the search for scientifically interesting sites must be conducted. The reliability constraint keeps the rover from exceeding its design capabilities, while the diversity constraint acts as a countervailing condition that forces the rover to travel further in its search for diverse sites. Other expressions could be used for the site diversity requirement; the one in this paper is based on the idea that the rate of change of site diversity with traverse distance (per site) is a power function of the intersite traverse distance.

$$\text{Max } \dot{M} = \frac{M_{site}}{\left(\dfrac{D_{trav}+\beta}{v} + T_{recon} + M_{site}(T_{acq}+T_{sa})\right)} \qquad (1)$$

subject to

$$D_{trav}N_{sites} + \beta(N_{sites}+1) \le \varepsilon D^{*} \qquad (2)$$

$$\Psi_0 \ge \left(\frac{\theta}{2}\right) D_{trav}^2 N_{sites} \qquad (3)$$

$$C(v, T_{acq}) \le B \qquad (4)$$

In Model 1, the endogeneous variables are:

N_{sites}	= number of science sites investigated
D_{trav}	= typical intersite distance
T_{miss}	= surface mission time
v^*	= optimal surface velocity
T^*_{acq}	= optimal sample acquisition and handling time
M	= total science measurements
$\Delta M/\Delta T_{miss}$	= science measurements per unit of surface mission time
λ	= marginal science measurements per technology investment dollar

The exogeneously supplied variables in Model 1 are:

D_{site}	= site characteristic size
M_{site}	= number of science measurements per site
D^*	= reliability requirement
Ψ_0	= site diversity requirement
ε	= traverse efficiency = 1/odometer multiplier
θ	= constant of proportionality
B	= technology development budget
R	= technology development reserve
T_{recon}	= site reconnaissance time
T_{sa}	= sample analysis time
β	= typical intrasite distance = $(M_{site})^{1/2} D_{site}$

Irrespective of whether the objective function is to maximize the science measurement rate, Eq. (1), or to maximize the total science measurements subject to an additional constraint on total surface mission time, or to minimize total surface mission time subject to an additional constraint on science measurements, the first-order conditions are the same! These are:

$$-\left(\frac{D_{trav}+\beta}{v^2 M_{site}}\right) = \frac{\partial C/\partial v}{\partial C/\partial T_{acq}} \qquad (5)$$

$$N_{sites} = \left(\frac{\varepsilon D^{*} - \beta}{D_{trav}+\beta}\right) \qquad (6)$$

$$\Psi_0 = \left(\frac{\theta}{2}\right) D_{trav}^2 N_{sites} \qquad (7)$$

$$C(v, T_{acq}) = B \qquad (8)$$

$$\lambda = \frac{-\dot{M}^2}{\partial C/\partial T_{acq}}\Big|_{(v^*,\, T^*_{acq})} \qquad (9)$$

An examination of the objective function reveals that the shortest possible D_{trav} (i.e., D_{trav} = 0) maximizes Eq.(1). However, this degeneracy is avoided by the diversity constraint. The feasible region, shown in Figure 1, is the area below the reliability constraint curve and above the site diversity constraint curve. In fact, the reliability and site diversity contraints in Eqs. (6) and (7) interact in a stable manner to determine unique values for N_{sites} and D_{trav} given by the intersection of the two curves in Figure 1.

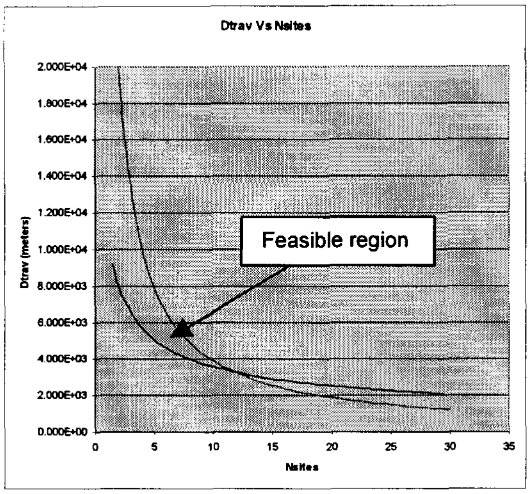

Figure 1— Solutions for N_{sites} and D_{trav} for Typical Model I Parameters

The Cost Functions

To determine the optimal mix of technology investments, Eqs. (5) and (8) must be solved together. Eq. (5) simply states that the optimum lies where the ratio of marginal payoffs in terms of the science measurement rate equals the ratio of marginal technology development costs. Knowledge of each technology development cost function is therefore critical to obtaining a quantitative solution.

In this paper, the combined technology development cost function is shown as Eq. (10).

$$\begin{aligned} C(v,T_{acq}) &= C_1(v)+C_2(T_{acq}) \\ &= -c_0 + v_0 c e^{(v/v_0)} + (a - bT_{acq}) \end{aligned} \tag{10}$$

where a, b, c, c_0 and v_0 are all positive constants and are chosen so as to reflect current technology when C_1 and C_2 = 0. The linear form for C_2 results in a significant computational benefit, since the marginal cost of decreasing sample acquisition and handling time is constant.

Solving for v^* and T^*_{acq} in Model 1

Using typical values for these cost function parameters, Figure 2 plots the marginal science payoff ratio and marginal cost ratio against v; the optimal v occurs at the intersection of the two curves.

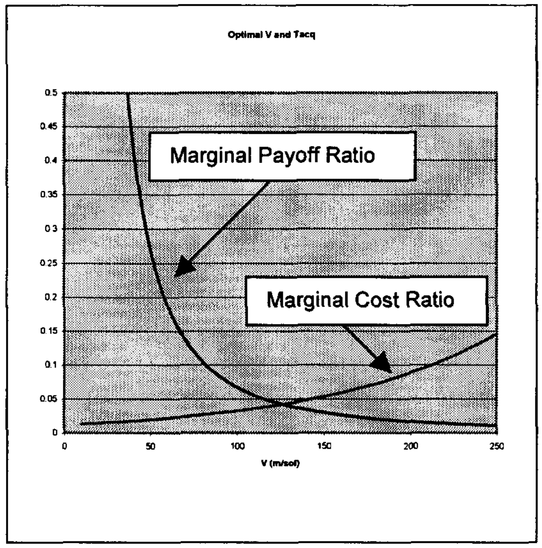

Figure 2—Solutions for v^* and T^*_{acq} for Typical Model 1 Parameters

With the convenience assumption of a constant marginal cost for decreasing T_{acq}, note that Eq. (5) depends only on v, holding the other Model 1 input parameters fixed. Consequently, with these technology development cost functions, an increase in the technology development budget leaves v^* unchanged, and only serves to decrease T^*_{acq}. Consequently, the proportion of the fixed technology development budget spent on velocity improvements decreases as the technology budget increases, with the proportion spent on sample acquisition and handling gaining accordingly. In general, one would expect that both v^* and T^*_{acq} would change with alternative budget levels.

The objective function, the science measurement rate, is an indirect function of the technology development budget in Model 1. That indirect function is shown in Figure 3.

Model 1 can be "deepened"–that is, expanded to include other technologies and multiple kinds of science measurements. Adding another technology adds another unknown variable to be optimized, but also adds another equation similar to Eq. (5) among the first-order conditions.‡

‡ It can be shown that the second-order conditions for a maximum are also satisfied with typical values for Model 1 parameters.

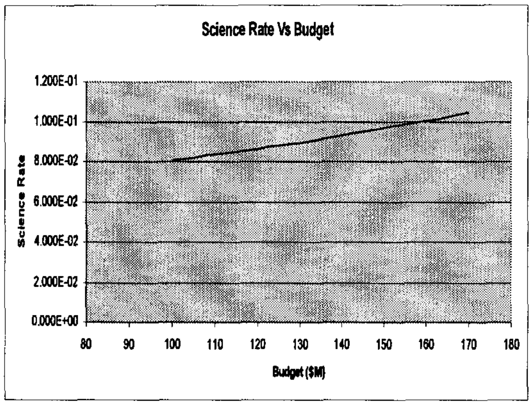

Figure 3—Science Rate Versus Budget for Typical Model 1 Parameters

Technology Development Uncertainty

So far technology development has been treated as a deterministic process–a given proposed cost yields a certain performance outcome. For a variety of reasons most of us are familiar with, technology development is anything but certain. I take the view that each technology development cost functions, $C_i(x)$, in this paper is a response function to a demand for performance x. Both the demander (usually a technology program manager), and the supplier (a technology developer) recognize that the actual performance outcome is highly uncertain, and that to achieve a given level of performance may require considerably more resources (and rarely less) than the proposed cost.

In this paper, I assume that the technology program manager holds reserves, R, that can be allocated across several technology developments to increase the probability of success in achieving the originally sought performance in any one of them. The technology program manager deals with the cost risk problem by allocating these reserves so as to maximize the joint probability of technology development success. Mathematically,

$$\text{Max } F(x \mid B_{v^*}) G(y \mid B_{T^*_{acq}})$$

subject to

$$x + y \leq R$$

where $F(\ )$ and $G(\ )$ are the cumulative distribution functions (cdf) for the cost of meeting the perfomance outcomes for rover velocity and sample acquisition and handling, respectively. Note that both distributions are conditioned on the proposed budgets (and schedules) derived from the technology development cost functions. How these cdf's can be determined is described more fully in Fox, et al.[3]

The first-order conditions to the above maximization are straightforward when stochastic independence is assumed, and are shown as Eqs. (11) and (12) (for an arbitrary number of technologies).

$$\frac{F'(x)}{F(x)} = \frac{G'(y)}{G(y)} = \cdots = \zeta \tag{11}$$

$$x + y + \ldots = R \tag{12}$$

Again, to obtain a quantitative solution, I assumed a lognormal cdf for $F(\ )$ and a Weibull cdf for $G(\ )$. The specific parameters used in what follows produced the graphs in Figures 4 and 5.

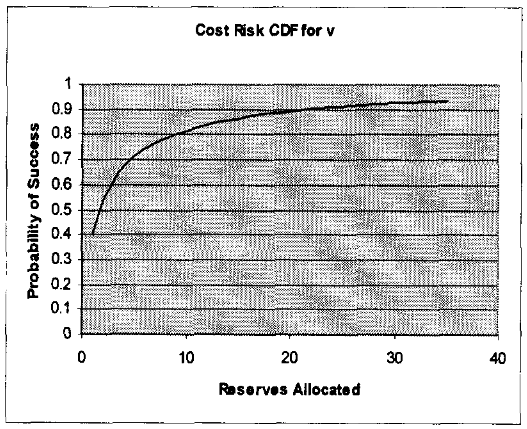

Figure 4—Cost Risk CDF for Rover Velocity Improvements

Applying Eqs. (11) and (12) to these two cost risk distributions permits the construction of the relationship between the maximized joint probability of success and the level of reserves, R. A graph of this relationship is shown as Figure 6 for typical model values.

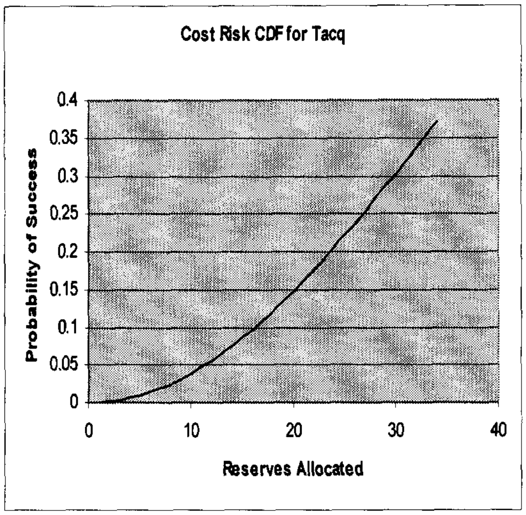

Figure 5—Cost Risk CDF for Sample Acquisition and Handling Improvements

At this point the technology program manager faces a significant tradeoff. The broad mission model approach *cannot* relieve the technology program manager from having to choose between buying more technology performance (and reaping a higher science measurement rate) and buying down the risk that one or more of the technology developments will fail to deliver. In other words, the technology program manager needs to choose the right values for B and R within the context of total programmatic resources, presumably $B + R$.

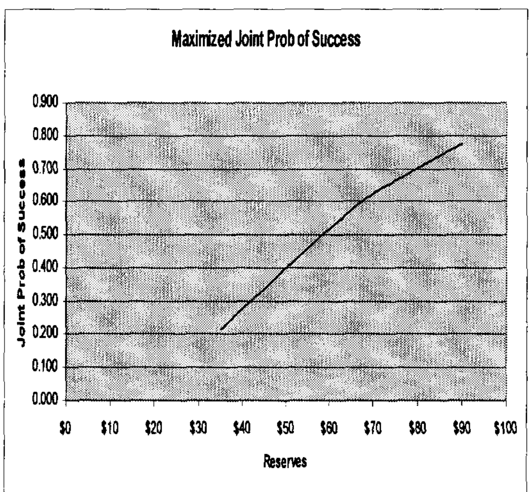

Figure 6—Relationship Between Reserve Level and Maximized Joint Probability of Success

One canonical approach to this choice problem is to postulate a utility function reflecting the technology development manager's preference for science rate versus risk, $U(\dot{M}(B), p_S(R))$, where $\dot{M}(B)$ is the maximized science measurement rate of Figure 3 and $p_S(R)$ is the maximized joint probability of technology development success of Figure 6. The utility function is assumed to have indifference curves reflecting a declining marginal rate of substitution between science rate and risk.

Reading the graphs of Figures 3 and 6, one can construct the combinations of maximized science rate and maximized joint probability of success that are consistent with a fixed level of programmatic resources. Figure 7 shows the result for $225, consistent with previous graphs, and overlays some representative indifference curves of the postulated utility function.

The optimal combination occurs at the tangency point shown where:

$$MRS = \frac{\partial U / \partial \dot{M}}{\partial U / \partial p_S} = -dp_S / d\dot{M} = MRT \quad (13)^{\S}$$

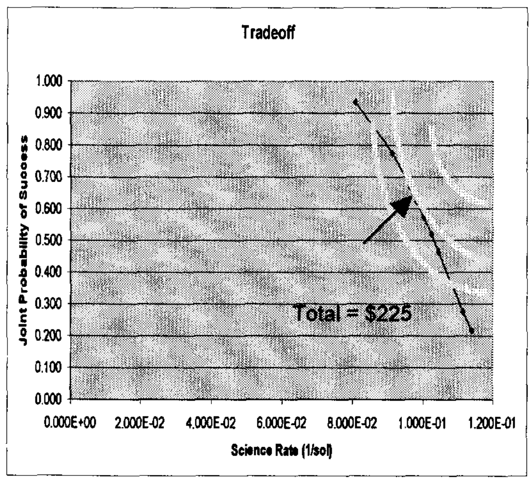

Figure 7—Optimal Combination of Science Rate and Risk for Model 1

OPTIMIZATION IN MODEL 2

In Model 2, I retain the same technology development cost functions, cost risk cdfs, and technology program manager utility function as

§ MRS is the marginal rate of substitution and MRT is the marginal rate of transformation in Model 1. The second model is based on a concept of operations in which a large fixed laboratory is landed on Mars that carries a deployable short-range "fetch" rover. In this broad mission model, samples of Martian material are obtained by the fetch rover one-at-a-time, and are brought to the laboratory where they are analyzed by its advanced instruments.

One question for modeling this concept concerns how big the "waiting room" is on the laboratory. With no waiting room for samples, the rover must wait until the laboratory has finished the last sample analysis before depositing the most recently collected sample into the first instrument.** A waiting room would allow the rover to temporarily place samples in the laboratory until it is ready to begin the analysis process. The rover is then free to search out other interesting samples.

In a purely deterministic world, the rate of science measurements would be the smaller of the (fixed) arrival rate of samples and the (also fixed) laboratory servicing rate of samples. This is independent of whether or not there is a waiting room. A better model treats both arrivals and servicing as stochastic processes with the arrival and servicing rates as parameters of a queueing process.

In Model 2, the objective function to be maximized, given in Eq. (14), is the *expected* rate of science measurements. The maximization is subject to a technology development cost constraint, Eq. (15).

$$
\begin{aligned}
\operatorname{Max} \mathrm{E}(\dot{M}) &= \mu + (\lambda - \mu) P_0 \\
&= \mu + (\lambda - \mu)\left[\frac{1-(\lambda/\mu)}{1-(\lambda/\mu)^{N+2}}\right] \qquad (14)
\end{aligned}
$$

$$
\begin{aligned}
\mu &= 1/T_{sa} \\
\lambda &= 1/(T_{acq} + D_{site}/2v)
\end{aligned}
$$

$$C(v, T_{acq}) \le B \qquad (15)$$

** Each sample is processed serially through the laboratory's instruments.

In Model 2, the new endogenous variables are:

$\mathrm{E}(M)$	= total expected science measurements
π	= marginal expected science measurements per technology investment dollar
λ	= arrival rate for samples
P_0	= probability that no sample is in or awaiting laboratory analysis

Model 2 also requires new exogeneous variables:

μ	= servicing rate for samples
N	= maximum number of samples that can be awaiting analysis (waiting room size)

The servicing rate is the reciprocal of the sample analysis time, which is a attribute of the laboratory's instruments and is exogenously given. The arrival rate is the reciprocal of the sample acquisition and handling time plus the *average* round-trip time between the sample and the laboratory. For the greatest simplicity here, I assumed that suitable samples are uniformly distributed and randomly selected over the science site. Obviously, if samples were selected from the periphery of the site, then the round-trip time would be double that shown.

The objective function here, Eq. (14), is a weighted average of the arrival and servicing rates, with the proportions determined by the probability that there is no sample in or awaiting laboratory analysis. The equation for that probability is a well-known result for a single-server exponential queueing system having a finite capacity of N+1.[4] The first-order conditions for the maximization are:

$$D_{site}/2v^2 = \frac{-\partial C/\partial v}{\partial C/\partial T_{acq}} \qquad (16)$$

$$C(v, T_{acq}) = B \qquad (17)$$

and when N = 0

$$\pi = \frac{-2}{\partial C/\partial T_{acq}}\left(\frac{\lambda}{(1+\lambda/\mu)}\right)^2 \Bigg|_{(v^*, T^*acq)} \qquad (18)$$

<u>Solving for v^* and T^*_{acq} in Model 2</u>

The convenience assumption of a constant marginal cost for decreasing T_{acq} means that Eq. (16) depends only on v, holding the other Model

2 input parameters fixed. This, of course, is not necessarily the case, but it highly simplifes these illustrative calculations. Again, with these technology development cost functions, an increase in the technology development budget leaves v^* unchanged, and only serves to decrease T^*_{acq}. Figure 8 shows the indirect expected science measurement rate function.

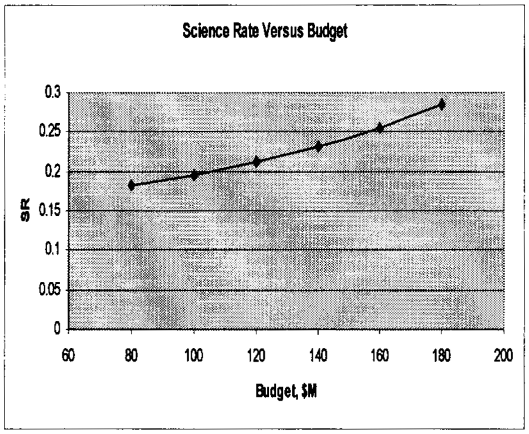

Figure 8—Expected Science Rate Versus Budget for Typical Model 2 Parameters

In Figure 8, the waiting room size is zero, which is consistent with the wishes of the MSL scientists to avoid any additional potential sources of sample contamination.

Using typical Model 2 parameter values and keeping parameters common to both Models 1 and 2 the same (especially the technology development budgets), one observes that v^* and T^*_{acq} are higher in Model 1 than in Model 2. This is exactly what one would expect since long-range roving consumes so much time that it pays to invest in technologies that can reduce that time when site diversity is a mission requirement. That science measurements are not be performed while roving means that the overall science rate in Model 1 is much lower than in Model 2. It is, however, the project scientists who must decide whether site diversity is important enough to forego the higher science rate.

Technology Development Uncertainty

The previous discussion and results regarding the maximized joint probability of technology development success apply to Model 2 as well. It is possible then to construct a new figure similar to Figure 7, which shows the optimal combination of science rate and technology development risk. Figure 9 shows that for Model 2.

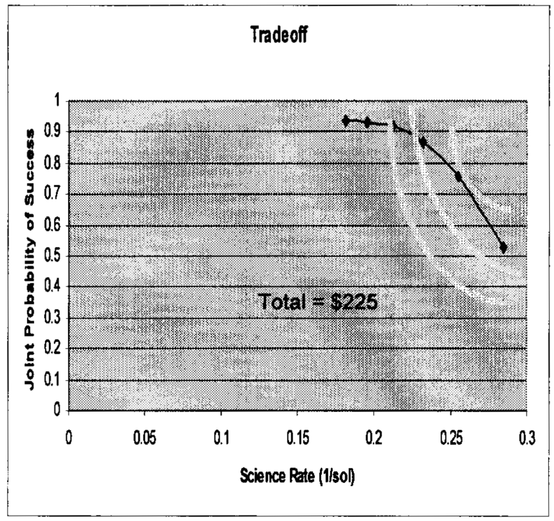

Figure 9— Optimal Combination of Science Rate and Risk for Model 2

The tangency condition in Eq. (13) applies in Model 2, and results in the optimal allocation of technology development resources to each technology, based on the demand for performance, and to reserves, based on portfolio risk.

CONCLUDING REMARKS

The work described in this paper was meant to demonstrate how a broad mission model can be developed and used to solve the TIP selection problem in a risk-based framework. Several issues can be raised in that regard.

Feasibility

In order to be truly useful to NASA's Mars Exploration Program, there is a need to broaden the mission model so that the optimization can include technology developments related to Mars precision landing and safe landing (i.e., landing hazard detection and avoidance). This can be done with a more complex mission model. Next, there is a related need to develop the means to precisely *translate* improvements in the technical attributes of hardware and software to improvements in the mission performance variables used in the broad mission model. In this regard, some good work has been done with respect to rover autonomy.[5]

One of the essential elements of analysis in the broad mission model approach to TIP selection is knowledge of the technology development costs, expressed as functions of the mission performance variables. Unfortunately, this is an area in which many unfounded assumptions are typically made (including in this paper). In connection with translating improvements in technical attributes into improvements in mission performance variables, there is a parallel need to translate the costs and cost risk information (often available only at the technical attribute level) to the level required by the optimization in the broad mission model approach.

Another essential element of analysis is the technology program manager's utility function. The tradeoff between portfolio risk and performance is essentially unavoidable, but different stakeholders may have different risk preferences. Whose utility function will be used is, as we know from Arrow's Impossibility Theorem, a thorny issue.

As I mentioned earlier, even a modest change in mission concept may require a completely different mission model. Once an appropriate mission concept has been selected and the corresponding mission model objective function and constraint equations developed, optimization should not be difficult given today's tools. In fact, all the calculations for this paper were done using a commonly available spreadsheet tool. (Appendix A shows the user input area and output area of the spreadsheet for Models 1 and 2. These are included to inform the reader as to the typical values of the inputs used and output variables obtained.)

Which Approach Is Better?

In my introduction, I alluded to two paradigms for TIP selection–the real-options approach and the broad mission model approach. A comparison of the two seems inevitable. While this topic certainly could be the subject of a separate paper, several points are worth making here.

First, the technology development cost functions and cost risk cdfs are essential elements of analysis in both approaches. Both also require some metric of value for performance improvements. In the broad mission model approach, the value is generally associated with the amount or rate of science measurements. This is a metric that is clear and meaningful to scientists, engineers, and project managers. In the real-options approach, the value of performance improvements in NASA flight projects is measured in dollars, which, while clear enough, is harder to quantify and to explain to the above groups.[††]

Second, schedules–the timing of both costs and benefits–are generally implicit and fixed in the broad mission model approach, while they are explicitly treated and of considerably greater importance in the real-options approach. Further, the way risks are modeled and propagated through time is substantially different in the two approaches.

Lastly, the broad mission model addresses the TIP selection problem in the context of a single mission that has a very high likelihood of occuring because it's the next one on the roadmap. As such it seems very approprate as a tool for focused technology developments tied to a specific roadmapped mission. Because NASA has mission roadmaps for many thematic programs and enterprises, many such broad mission models are needed for NASA as a whole. The limited context of the broad mission model approach leads me to conclude that it is more tactical in nature than the real-options approach, which takes a more strategic view of the TIP selection problem.

ACKNOWLEDGMENTS

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

OK

The author wishes to thank Steve Prusha and Alberto Elfes of JPL, who supported this work under NASA's Engineering of Complex Systems Program.

APPENDIX A

Models 1 and 2 user inputs and outputs follow:

[††] In actual practice, whenever the real-options approach requires the ability to quantify changes in the amount or rate of science measurements/data, a broad mission model is often helpful. The real-options approach also requires a credible way of translating that into dollars.

Model 1

Input Parameter Name	Value	Units	Comments
Reliability	0.93		0.74 < R < 0.99
Min Potential Explorable Area, Psi	1.00E+08	m2	Diversity-driven area constraint
Potential Area Focus, Theta	1.570795	radians	Pr{x>0}~=density* psi
Msite	5		"Rock" measurements per site
Dsite	50	m	Site diameter
Trecon	4	sols	Time for site reconnaissance
Traverse Efficiency	0.667		
Beta	111.80	m	=SQRT(Msite)*Dsite
D*	6.11E+04	m	Reliability-driven total distance limit
Tsa	3	sols	Time for sample analysis per rock
Cost Functions and Budget			
C0_Tacq	100	$M	Linear cost fn intercept
MC_Tacq	-25	$M/sol	Linear cost fn slope
C	0.3		Exponential cost fn
V0	100	m/sol	Exponential cost fn
Budget	160	$M	
Reserves	65	$M	

Output Parameter Name	Value	Units	Comments
Dtrav	3.257E+03	m	
Nsites	12		
Mission Time	6.587E+02	sols	
Mission Time	1.85	Yrs (on Earth)	
Science Rate	9.108E-02	1/sol	Maximized objective
Potential Explorable Area	1.000E+08	m2	>=Psi
Total Distance	6.078E+04	m	<=Dmax
Total Euclidean Distance	4.054E+04	m	=Tot Dist * Trav Eff
Optimal V	125	m/sol	
Optimal Tacq	1.788	sols	Sample acq time
Gamma	27.942	sols	=Msite*(Tacq+Tsa)+Trecon
Optimal Percent Spent on V	65.4%		
Optimal Percent Spent on Tacq	34.6%		
Lambda	3.32E-04	1/sol-$M	Shadow price
Science (Mission Total)	60		
Percent Mission Time At Sites	52.5%		
Optimal Percent Reserve on V	7.7%		
Optimal Percent Reserve on Tacq	92.3%		
Joint Probability of Success	75.6%		Maximized joint prob

Model 2

Input Parameter Name	Value	Units	Comments
Waiting Room	0		on MSL
Max Rocks in System	1		N + 1
Discard Probability	0		currently not used
Dsite	500	m	site diameter
Tda	2	sols	
Service Time	3	sols	Tsa = Tda + Dcycle
Service Rate	0.333333	1/sol	1 / Tsa
Left-Hand Side Constant	9.94431		
Mission Time	180	sols	
Mission Decision Cycle	1	sols	Tdcycle
Cost Inputs	**Value**	**Units**	**Comments**
C	0.3	$M	
v0	100	m/sol	
MC_Tacq	-$25.0	$M/sol	
C0_Tacq	$100.0	$M	
Budget	160	$M	

Output Parameter Name		Value	Units
Expected Science Rate		0.31603738	"rocks"/sol
Total Expected Science		56.8867277	"rocks"
Optimal Tacq		0.59299219	sols
Optimal v		91.3952062	m/sol
Utilization from QTP		0.47405606	
E(number in system) from QTP		0.47405606	
E(time in waiting room) from QTP		0	
E(time in system) from QTP		3	
State Probabilities from QTP			
	0	0.52594394	
	1	0.47405606	
	2		
	3		
	4		
	5		
	6		
Interarrival Time		3.328365	sols
Arrival Rate		0.30044778	1/sol
Cost Outputs		**Value**	**Units**
C(v)		$74.82	
C(Tacq)		$85.18	
Total Expenditure		$160.00	
Fraction Spent on v		46.8%	
Lagrangian multiplier, pi		1.99759E-03	"rocks"/sol-$M

REFERENCES

1 Shishko, R., et al., "NASA Technology Assessment Using Real-Options Valuation", *Proceedings of the 12th Annual International INCOSE Symposium*, Las Vegas, NV, 2002

2 Saleh, J.H., Lamassoure, E., and Hastings, D.E., "Space Systems Flexibility Provided by On-Orbit Servicing: Part 1," *Journal of Spacecraft and Rockets*, Vol. 39, No. 4, 2002, pp. 551-560

3 Ebbeler, D., Fox, G., and Habib-agahi, H., "Dynamic Cost Risk Estimation and Budget Misspecification", *Proceedings of the AIAA SPACE 2003 Conference*, Long Beach, CA, September 2003

4 Ross, Sheldon M., *Introduction to Probability Models, 7th Ed.*, Harcourt Academic Press, San Diego, CA, 2000

5 Lincoln, W.P., Elfes, A., et al., "Relative Benefits of Potential Autonomy Technology Investments", *International Conference on Space Mission Challenges for Information Technology (SMC-IT)*, Pasadena, CA July 2003