Modeling and Simulation in the Design of a Lunar Rover Suspension with Extensibility to Martian Applications

Stephanie C. Thompson*
Systems Realization Laboratory, Georgia Institute of Technology, Atlanta, GA, 30332

Nathan Young†
Systems Realization Laboratory, Georgia Institute of Technology, Savannah, GA, 31407

Robert Thiets‡,
Systems Realization Laboratory, Georgia Institute of Technology, Atlanta, GA, 30332

Thomas M. Groshans§
Georgia Institute of Technology, Atlanta, GA, 30332

and

Christiaan J. J. Paredis**
Systems Realization Laboratory, Georgia Institute of Technology, Atlanta, GA, 30332

In this paper, we present a simulation-based study of the design of the suspension system for a manned lunar rover based on utility theory. The study includes the use of both analysis and decision models. The design objectives are modeled using a fundamental objectives hierarchy and a means objectives network. Our fundamental objective is to maximize the research value of the rover via the suspension system by maximizing the safety of passengers and equipment. The means objectives which contribute to maximizing the safety of equipment are the focus of this study and include maximizing energy dissipation, minimizing vibrations caused by the terrain, and minimizing the vertical travel of the payload. The scope of this study is limited to the suspension parameters and does not include the design of the geometric configuration of the system. The measures of effectiveness for these means objectives include the vertical displacement, the upward and downward acceleration, and the settling time of the rover body.

An influence diagram is developed to organize the decisions, chance events, and computational outcomes that influence the achievement of the design objective. In this design study, the primary decision is to identify the preferred values of the effective spring rate and damping coefficient for each wheel. An energy-based model of the suspension system is developed using the software package Dymola and the modeling language Modelica. The energy-based model is then used to simulate the dynamic response of the suspension to obtain values of the measures of effectiveness. The suspension system is designed for extensibility to Martian applications by modeling the environment (moon or Mars) of the rover as an uncertain chance event that influences the design decision. Additional chance events include the mass (of the rover, occupants, and payload), the terrain, the flexibility of the rover body, and the length and width of the rover. The chance

---

* NSF Graduate Research Fellow, G. W. Woodruff School of Mechanical Engineering, AIAA Student Member.
† Graduate Research Assistant, G. W. Woodruff School of Mechanical Engineering - Savannah, AIAA Student Member.
‡ Graduate Student, G. W. Woodruff School of Mechanical Engineering.
§ Graduate Research Assistant, G. W. Woodruff School of Mechanical Engineering.
** Assistant Professor, G. W. Woodruff School of Mechanical Engineering, chris.paredis@gatech.edu.

events are modeled using statistical parameters, and Latin Hypercube Sampling with common random numbers is used to estimate the expected utility of each decision.

The final phase of the study includes preference modeling and optimization. The preferences of the designer are modeled with a multi-attribute utility function. The goal is then to maximize the expected utility under the uncertainty of the chance events by adjusting the design variables of spring rate and damping coefficient. An optimizer is employed in the ModelCenter simulation integration framework to maximize the expected utility estimated via Latin Hypercube Sampling. To reduce the computational complexity, the expected utility estimator samples an adaptive Kriging model rather than the dynamic simulation model directly. This results in a speed-up of several orders of magnitude.

Nomenclature

A	=	acceleration of rover center of mass [m/s^2]
AD	=	upward acceleration of rover center of mass [m/s^2]
AU	=	downward acceleration of rover center of mass [m/s^2]
k_i	=	multi-attribute utility function scaling constant pertaining to attribute i
ST	=	settling time of the rover center of mass [s]
T	=	maximum travel of the rover center of mass [m]
$u(x,y,z)$	=	utility as a function of attributes x, y, z

I. Introduction

IN the Vision for Space Exploration set forth by President Bush in January 2004, a major objective is to return to the moon and establish a lunar base to gain experience for future manned missions to Mars.[1] To achieve this goal, the members of the Exploration Systems Architecture Study group recommend Mars-forward testing, in which the systems that will be used on Mars missions are first developed for and tested on the moon.[2] Rovers are needed on both the moon and Mars to provide access to a large area of the surface, and these rovers must be designed to accommodate not only variations in the environment between the moon and Mars, but also variations in use and loading conditions. In this paper, we present a simulation-based study of the design of the suspension system for a manned lunar rover with extensibility to Martian terrain. The scope of this study is limited to the suspension parameters (spring rate and damping coefficient) and does not include the design of the geometric configuration of the system. It is assumed that the geometry of the rover will be considered in more detail later in the design process when the decisions concerning component sizing and materials are made.

Decision-based design research recognizes the pivotal role that decisions play in all phases of a design process.[3,4] In this context, we focus on modeling the design decision involved in selecting spring rates and damping coefficients for the suspension system of a manned rover in this paper. In accordance with Normative Decision Theory, a rational decision maker should choose the decision alternative which has the highest expected utility.[5] To evaluate the expected utility of each decision alternative represented by values of spring rate and damping coefficient, a combination of energy-based system models, uncertainty models, and preference models are employed.

The model of the decision situation is explained in Section II. A fundamental objectives hierarchy as well as a means objective network are generated and used to define measurable and comprehensive attributes for the decision. The design alternatives are synthesized through systematic decomposition and an influence diagram is created to represent the decisions, chance events, and measures of effectiveness. The energy-based model of the suspension system is developed using the software package Dymola and the modeling language Modelica. An overview of this model is given in Section III. The energy-based model is then used to simulate the dynamic response of the suspension as it travels over lunar and Martian terrain in order to obtain values for the measures of effectiveness. Chance events are modeled using statistical parameters, and Latin Hypercube Sampling with common random numbers is used to estimate the expected utility of each decision. A detailed explanation of the uncertainty models is given in Section IV.

The final phase of the simulation-based design study includes preference modeling and optimization. Preferences are modeled with a multi-attribute utility function. The optimization objective is to maximize the expected utility under the uncertainty of the chance events by adjusting the design variables. An optimizer is employed in

ModelCenter to maximize the expected utility calculated via Latin Hypercube Sampling. Detailed explanations of preference modeling and optimization are located in Section V and Section VI, respectively.

II. Decision Modeling

The rover suspension decision is modeled using objectives hierarchies and an influence diagram. These are discussed in this section.

A. Objectives Hierarchies

Generating fundamental objectives hierarchies and means objectives networks assists the design team in identifying the objectives involved in the decision.[6] The fundamental objective of rover suspension is to maximize the research value of the rover via the suspension system by maximizing the safety of passengers and equipment. This objective is decomposed into specific components which do not contain any overlap, specifically: to maximize safety of payload, minimize launch cost, and to minimize build cost. The process is repeated as long as the objectives remain solution independent. The resulting fundamental objectives hierarchy for this design study is shown in Figure 1. In this design study, we focus on the fundamental objectives which maximize the safety of the rover payload via maximizing energy dissipation, minimizing vibrations, and minimizing vertical travel of payload. Means objectives are elicited by determining possible methods for achieving fundamental objectives. A means objective network showing the influence and interaction between the fundamental and means objective is shown in Figure 2.

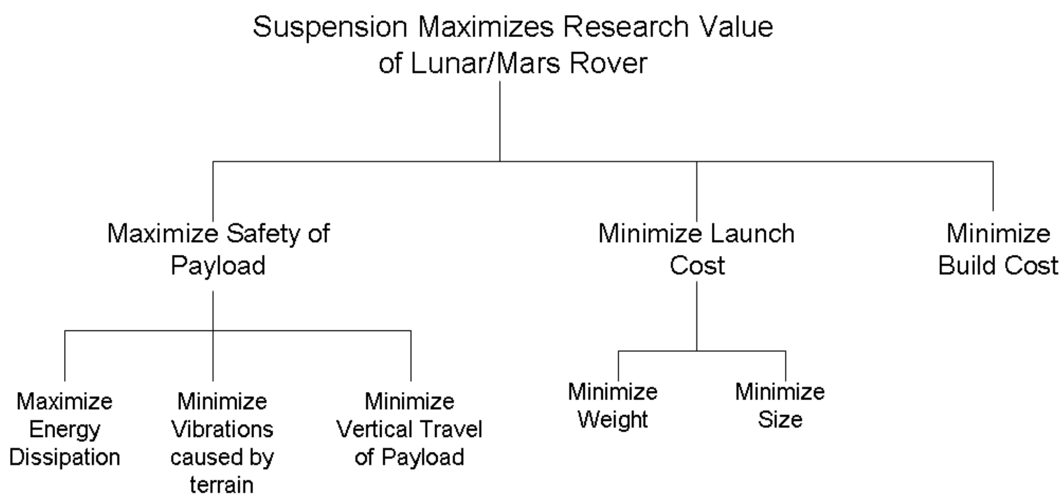

Figure 1. Fundamental Objectives Hierarchy

B. Influence Diagram

An influence diagram, shown in Figure 3, is used to model the structure of the design problem and is composed of design decisions, chance events, and computation outcomes.[7] Design decisions are marked by a square box in the influence diagram and represent key milestones in the design process. The effective spring rate represents a lumped spring rate including contributions from a spring, the tire, and deflection of suspension structure. The damping ratio is the rate at which energy is dissipated in the suspension as it is deflected. Each wheel of the rover is assumed to have identical spring rates and damping coefficients.

The chance events considered in this design study, marked by ovals in the influence diagram, are payload mass, rover location, terrain amplitude, angle of approach, torsional stiffness, torsional damping, axle length and wheelbase. A detailed description of each chance event can be found in Section IV.

Intermediate computational outcomes are measurable attributes which have been derived from the means objectives. The intermediate computational outcomes are marked with a rounded box in the influence diagram and include the settling time, acceleration, and travel of the center of mass of the rover. Settling time is the time required for a response to reach, and stay within 2% of its final value. In terms of the suspension design, this represents the required time to damp vibration energy to an acceptable value. The maximum acceleration seen by the rover's center of mass represents the ability of the suspension to maintain correct ride height, dissipate impact energy, and maintain traction. The total distance traveled by the rover's center of mass represents the ability of the suspension to maintain correct ride height and allows for design of adequate ground clearance.

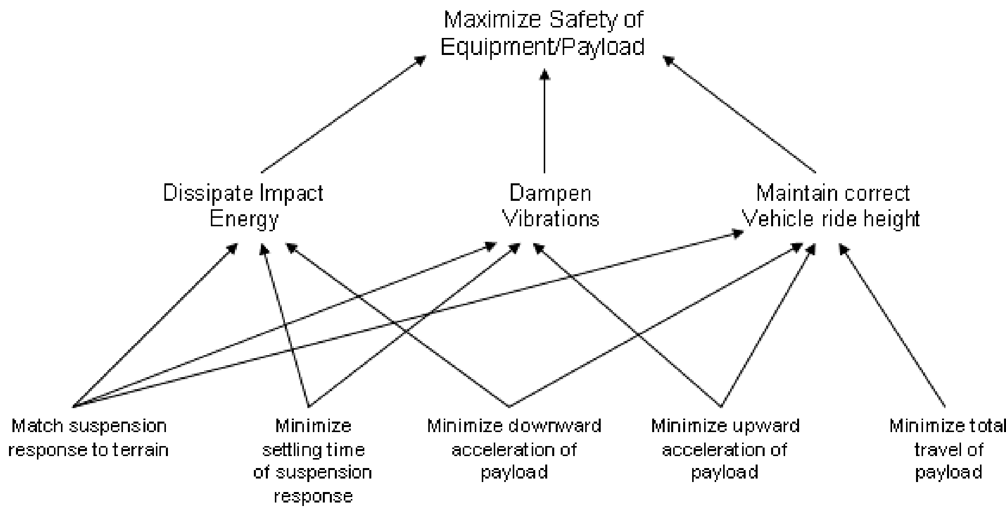

Figure 2. Means Objectives Network

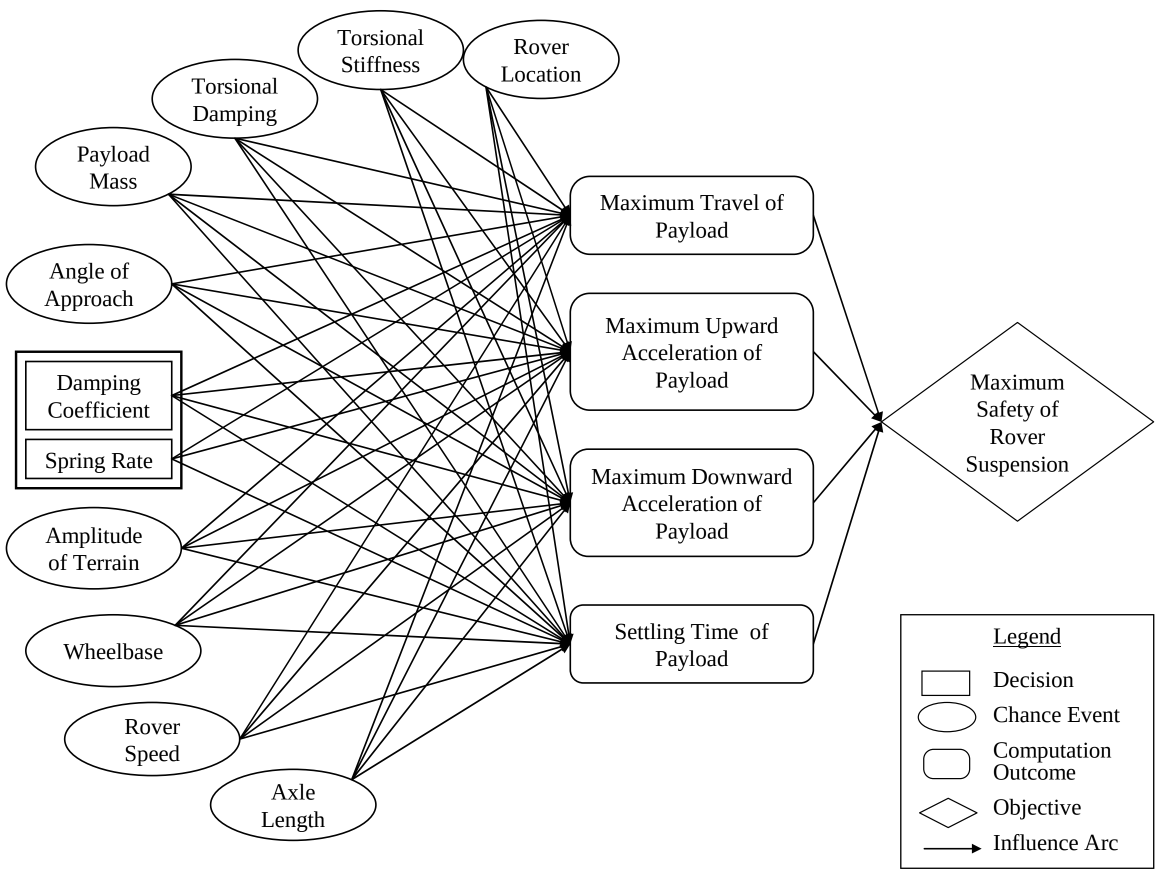

Figure 3. Influence Diagram

III. Systems Modeling

An energy-based model of the suspension system is developed using the software package Dymola and the modeling language Modelica using components from the Multibody Library.[8,9] The model is divided into two main components; one component simulates the rover suspension and the other simulates the terrain.

The rover is modeled as two axles connected by a rigid frame using revolute joints. These revolute joints allow control over the chassis flexibility. A point mass located in the center of the frame allows for control over the payload mass. Adjustable spring/dampers allow for control over the suspension parameters. A test frame model is used to simulate the movement of the rover over the terrain. A time-varying displacement signal is sent to four prismatic actuators, which convert the signal into displacements in the wheels of the rover model. One of the actuators is fixed and the other three are free to move in the horizontal plane. This configuration simulates the dynamic response of the rover in the vertical direction only.

The following assumptions and relationships are necessary to describe the interaction between the rover suspension and terrain components and to reduce the complexity of the model:

1) The suspension response is simulated by a time-varying change in the vertical position of the contact points between the rover suspension and terrain components
2) The rover never breaks contact with the terrain
3) The rover has perfectly rigid and mass-less wheels/tires
4) The rover travels at a constant velocity regardless of the terrain
5) The terrain only influences the vertical position of the contact points between the models
6) There is no friction between the terrain and the suspension
7) The rover is acted on by either lunar or Martian gravity
8) The suspension has four contact points with the terrain
9) Each contact point is independent
10) The spring force-deflection relationship is linear
11) All members are perfectly rigid
12) Flex due to the frame is approximated by a stiff spring
13) Friction in the joints is negligible

IV. Uncertainty Modeling

As shown in the influence diagram in Figure 3, the suspension decision is influenced by several chance events. These chance events represent the uncertainty associated with the behavior of the suspension, and these uncertainties must be captured using an appropriate modeling approach. Uncertainty modeling is a characterization of uncertain parameters which affect the realization of decision model outcomes. Characteristics of good uncertainty models involve representation, inference, and decision making. A good representation captures all available information and knowledge and has a clear, unambiguous operational definition. A good uncertainty model also allows the decision maker to infer information relevant to the decision in a computationally inexpensive and internally consistent manner. Furthermore, a good uncertainty model promotes decision making which yields a coherent decision with no sure loss. These are the characteristics of an ideal uncertainty model; however, no formalism completely satisfies these ideals. We implement probability theory to represent uncertainty in our design, which is the uncertainty formalism of Normative Decision Theory.[5]

A. Subjective Probability

Probabilities are often introduced from a frequentist perspective whereby the probability of an event is the limit value of frequency for an outcome over a long run such as the probability of getting heads in a coin toss. This interpretation breaks down for certain events, such as the probability of rain tomorrow. In this work, we implement subjective probabilities which express a willingness to bet on a particular outcome in a gamble in which one receives $1 if the event occurs and $0 otherwise. Subjective probabilities are related to frequentist probabilities in that one's beliefs expressed as a subjective probability should be consistent with scientific, factual information. Therefore, frequentist probabilities inform our beliefs about the outcome of an event and thus impact the formulation of our subjective probabilities about that event. The advantages of using this type of approach include an unambiguous meaning and operational definition which are both important in the support of decision making. The elicitation of beliefs about the uncertainties in the suspension decision is discussed in the next section to demonstrate the development of subjective probabilities.

B. Eliciting Beliefs

Uncertain events are represented by either discrete or continuous probability distributions. A discrete distribution can only take on certain values; however, the number of values within the discrete distribution can be infinite or finite. A coin toss is an example of a discrete distribution because the outcomes can only be heads or tails. Similarly, in the suspension decision, the location of the rover is represented by a discrete distribution because the only possible outcomes are the moon or Mars. In a continuous probability distribution, the outcome can take on any value in a range. This range can be bounded on one or both sides or be unbounded, and the number of possible values in the range is infinite. With the exception of the rover location, all of the chance events influencing the suspension decision are represented by continuous distributions.

To elicit the subjective probability for a discrete variable, questions are asked based upon the possible outcomes. To begin the process, a single question is posed in the following manner for an arbitrary uncertain variable. "How much would you be willing to bet that outcome *x* occurs?" For *n* possible outcomes, *n-1* questions are asked to elicit appropriate probabilities for each outcome because the probabilities for each outcome must sum to one. To elicit our beliefs for the rover operating environment, the question is "How much would you be willing to bet that the rover is operating on the moon?" Our answer of $0.70 reflects our beliefs that the rover is more likely to be operating on the moon because the rover suspension may go through a redesign or the program may be cancelled prior to the rover's use on Mars. Since there are two possible outcomes, only one question must be asked.

To elicit the subjective probability for a continuous variable, a series of lottery questions are asked in order to populate a cumulative distribution function (CDF). The questions are stated in the following manner for an arbitrary value *x* of random variable, *X*: "How much would you be willing to bet that *X* is less than or equal to *x*?" By answering such lottery questions for several different *x* values of random variable *X*, one can elicit points on a cumulative probability distribution. These points are then used to find an appropriate parametric probability distribution using Maximum Likelihood Estimation (MLE). A parametric probability distribution is needed for implementation in ModelCenter using Latin Hypercube Sampling; however, in general, any appropriate curve fitted to the elicited data can be used.

The rover payload is a continuous variable which can take on any real value in some range. To elicit points on the CDF for the rover payload, the following question is asked for several values of *x*: "How much would you be willing to bet that the mass of the rover payload is less than or equal to *x*?" The answers to these questions are the points on the CDF for the payload as shown in Figure 4. The other uncertain variables (terrain amplitude, rover speed, angle of approach, torsional stiffness, torsional damping, wheelbase, and axle length) are also continuous variables. To elicit the CDF for each variable, the process is carried out the same way as for the payload of the rover. From these elicited CDF's, approximating probability distributions are found through MLE. In Table 1, each uncertain variable is shown with its corresponding approximating distribution.

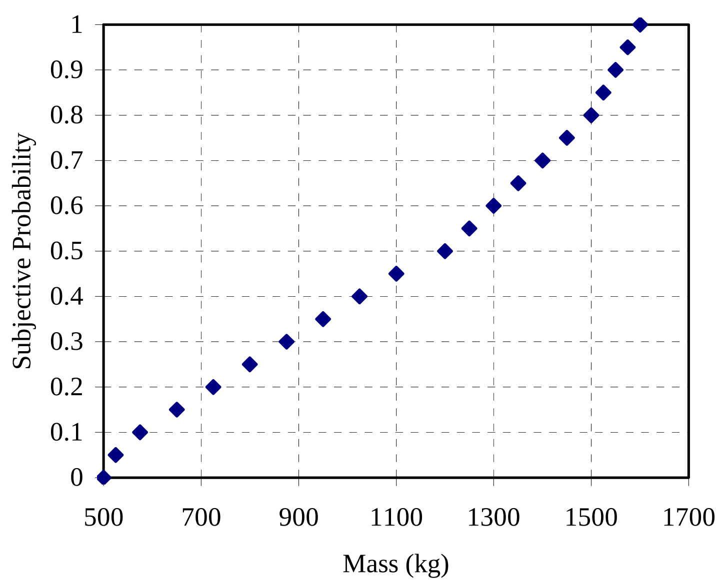

Figure 4. Cumulative Distribution Function of Payload

Table 1. Approximating Probability Distributions for the Uncertain Parameters

Parameter	Meaning	Distribution
Payload Mass	Mass in kg of the rover body, occupants, cargo, etc.	Uniform(500, 1600)
Rover Location	Subjective probability that the rover is on the moon (as opposed to Mars)	0.70
Terrain Amplitude	Integrator gain of the terrain signal, which is related to the height of terrain features.	Uniform(0.04, 0.71)
Rover Speed	Speed of the rover in mph.	Uniform(3, 10)
Angle of Approach	Angle between the front axle and the longitudinal axis of the rover body.	Exponential(9.49) – 0.001
Torsional Stiffness	Stiffness of the rover body in twist, in N/m.	Triangular(7500, 8000, 67500)
Torsional Damping	Damping of the rover frame in twist, in N/m/s.	Uniform(7500, 75000)
Axle Length	Length between the left and right wheels in meters.	Normal(1.32, 0.247)
Wheelbase	Length between the front and back wheels in meters.	Uniform(1.5, 4)

V. Preference Modeling Using Multi-Attribute Utility Theory

A utility function assigns appropriate utility values to a set of possible attribute alternatives. In practical terms, this allows the designer to rank order a given set of attribute values such as acceleration with respect to utility.[5] Utility functions are preferred to value functions in design decisions because utility functions allow decision makers to express preferences under uncertainty and all design decisions involve some degree of uncertainty.[10] In this design study, there are four attributes: upward acceleration, downward acceleration, travel, and settling time. To develop a multi-attribute utility function involving these four attributes, there are three steps: verify mutual utility independence, assess conditional utilities for each attribute, and determine the tradeoffs between the attributes.

A. Verifying Mutual Utility Independence

Given some number of attributes, one must formulate a utility function that predicts the overall utility given values of these attributes. In general this is a very complex problem, but for cases in which the attributes are utility independent, a multi-linear utility function can be defined. Utility independence means that the conditional preferences for lotteries on x given a value of attribute y do not depend on the value of y. In the context of the suspension decision, this means that our preferences for the acceleration of the rover should be independent of the value of the settling time and the value of travel, our preferences for the settling time should be independent of the value of acceleration and the value of travel, and our preferences for the travel should be independent of the values of acceleration and settling time. This is in fact the case. Therefore, a multi-linear utility function can be defined.

B. Elicitation of Conditional Utilities for Individual Attributes

The process for determining conditional utility functions for individual attributes is comprised of four steps. These steps include identifying the best and worst case attribute values, determining the certainty equivalent corresponding to a 50/50 gamble between the best and worst case values, repeating the identification of certainty equivalents for various points until curve is smooth, and fitting a curve to the elicited points. This process is demonstrated below for the upward acceleration of the rover center of mass. The results of the elicitation of conditional utilities for all four attributes are shown in Table 2.

The best case value of upward acceleration is 0 m/s^2 meaning that the suspension system transmits no upward acceleration to the rover center of mass. This best case value corresponds to a utility of one. The worst case acceleration value is 9.8 m/s^2, which corresponds to a utility of zero. This value is chosen because a large upward acceleration relative to the local gravity could lift the rover off of the surface, possibly resulting in a loss of crew or equipment. To elicit the intermediate values, one must identify certainty equivalents. The certainty equivalent is the value at which one is indifferent between receiving that value for certain and a 50/50 gamble of two values. The certainty equivalent then has the average utility of the two values involved in the 50/50 gamble. To identify the first certainty equivalent (corresponding to 0.5 utility), the following question is asked: "At what upward acceleration would I be indifferent between the chosen amount for certain and a 50/50 gamble of 9.8 m/s^2 or 0 m/s^2?" The answer of 2.8 m/s^2 reflects a risk-seeking attitude because it means that the decision maker must receive an acceleration value better than the expected value of the gamble to take the certain option over the gamble. To identify another certainty equivalent (corresponding to 0.75 utility), the following question is asked: "At what upward acceleration would I be indifferent between the chosen amount for certain and a 50/50 gamble of 2.8 m/s^2 or 0 m/s^2?" This elicitation process is continued until the utility function for the upward acceleration is completely defined. The elicited certainty equivalents for the upward acceleration are shown in Figure 5. We see from the shape of the elicited points that the decision maker is risk-seeking as utility approaches zero but risk-averse as the utility approaches one. A 2D function finder is then used to arrive at a function for the conditional utility of the upward acceleration.[11] In general, elicited preferences often exhibit the S-shape shown in Figure 5 and are well approximated by sigmoid functions.

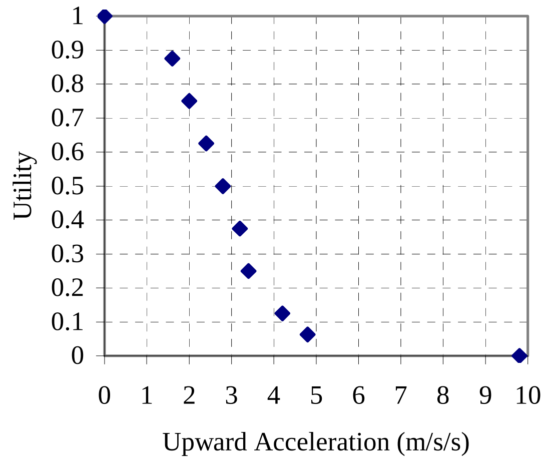

Figure 5. Elicited Certainty Equivalents for Upward Acceleration

Table 2. Conditional Utility Functions for Each Attribute

Attribute	**Best Case** $u(x) = 1$	**Worst Case** $u(x) = 0$	**1st Certainty Equivalent** $u(x) = 0.5$	**Fitted Function**	**Coefficients**
Upward Acceleration (m/s^2)	0	9.8	2.8	$y = \frac{a}{\left(1+be^{-cx}\right)} + d$	*a* = -1.02 *b* = 0.521 *c* = 1.45 *d* = 1.02
Downward Acceleration (m/s^2)	5.2	21	14	$y = ae^{\left(-e^{(b-cx)}\right)}$	*a* = 1.01 *b* = -5.45 *c* = -0.357
Settling Time (s)	0	5	1	$y = ae^{\left(-e^{(b-cx)}\right)} + d$	*a* = 0.997 *b* = 1.88 *c* = 2.00 *d* = 1.00
Travel (m)	0	0.3	0.1	$y = ae^{\left(-e^{(b-cx)}\right)} + d$	*a* = -1.00 *b* = 2.29 *c* = 26.5 *d* = 1.00

C. Combining Utilities into a Multi-Attribute Function

To create a multi-attribute utility function for the suspension decision, the elicited conditional utility functions for acceleration, settling time, and travel must be combined in a multi-linear utility function by assessing the tradeoffs between each attribute. The multi-linear utility equation for the three attributes is shown in Equation 1, where k_T, k_A, k_{ST}, $k_{T,A}$, $k_{T,ST}$, $k_{A,ST}$, and $k_{T,A,ST}$ represent the scaling constants and $u(T)$, and $u(ST)$ represent the conditional utility values for travel and settling time, respectively.

$$\begin{aligned} u(T,A,ST) &= k_T u(T) + k_A u(A) + k_{ST} u(ST) + k_{T,A} u(T)u(A) + \\ & k_{T,ST} u(T)u(ST) + k_{A,ST} u(A)u(ST) + k_{T,A,ST} u(T)u(A)u(ST) \end{aligned} \tag{1}$$

The term $u(A)$ is the value of a multi-linear utility function combining the conditional utilities for upward and downward acceleration. The multi-linear utility equation for these two attributes is shown in Equation 2 where k_{AU}, k_{AD}, and $k_{AU,AD}$ represent the scaling constants and $u(AU)$ and $u(AD)$ represent the conditional utilities of the upward and downward accelerations.

$$u(A) = k_{AU}u(AU) + k_{AD}u(AD) + k_{AU,AD}u(AU)u(AD) \quad (2)$$

To determine the scaling constants for a multi-linear utility function, a series of tradeoff questions are posed. In general, for a utility function with *n* attributes there are 2^n-1 scaling constants. The maximum overall utility is achieved when all the conditional utilities are equal to one. By setting the sum of the scaling constants equal to one, this maximum value is normalized to yield an overall utility of one. Due to this normalization, only 2^n-2 questions must be posed to obtain a set of 2^n-1 equations with the k_j's as the unknowns. The equations can be generated from certainty considerations, probabilistic considerations, or a combination; here we demonstrate the generation of equations from certainty considerations.

An equation is formed by identifying sets of attributes values between which the decision maker is indifferent. For example, consider a situation in which the rover suspension has the best case upward acceleration and the worst case downward acceleration. The decision maker compares that situation to another situation in which the conditional utility of the upward acceleration is equal to 0.5. The objective is to identify the value of downward acceleration in the second situation for which the decision maker is indifferent between the two situations. If the decision maker is indifferent between the two situations, then the utility for the two situations must be equal. One equation is formed by equating the utility for these two situations. Additional equations are created in a similar manner, and the resulting set of equations is solved to determine the values of the scaling constants for the two-attribute acceleration utility function.

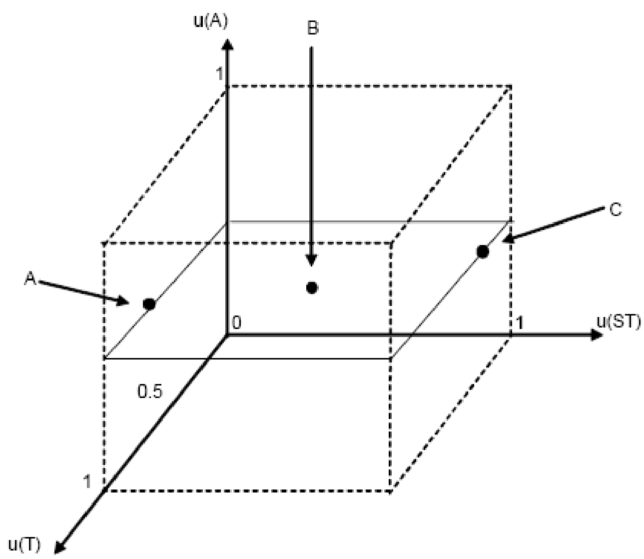

Figure 6. Graphical Representation the Utility Space

To identify the scaling constants in the three-attribute case, equations are generated in the same manner; however one of the attributes is held constant for the generation of each equation. Figure 6 is a graphical representation of the utility space for the suspension design decision in which the three axes *u(A)*, *u(ST)*, and *u(T)* represent the utilities of the acceleration, settling time, and travel, respectively. Although not shown, a three-dimensional utility landscape resides in this space. Points A, B, and C are used to compare travel and settling time with acceleration held constant at 0.5 utility. Point A represents an unknown travel utility with a settling time utility of 0. Point C represents an unknown travel utility at a settling time utility of 1. Point B represents the control point against which each case, A and C, is compared. A summary of the comparison of Points A and B is shown in Table 3, in which the bold entry indicates the value that is elicited in this comparison.

For the comparison between Points A and B, a travel value is elicited for Point A. This is possible because the utility is fully defined for Point B and the single missing utility for Point A is the travel. The travel value of 0.05 m is elicited based upon the preferences of the decision maker. Since the utilities for each point are fully defined, the utility values for Point A can be equated by using Equation 1. This results in Equation 3.

Table 3. A Summary of the Comparison of Points A and B

Point A	Attribute value	Attribute utility	Point B	Attribute value	Attribute utility
T'	**0.05**	0.926	T_0.5	0.1011	0.490
A_0.5		0.495	A_0.5		0.495
ST_0	5	0.004	ST_0.5	1.15	0.485

$$\begin{aligned}
u(\text{Point A}) = \quad & k_T(0.926) + k_A(0.495) + k_{ST}(0.004) + \\
& k_{T,A}(0.926)(0.495) + \\
& k_{T,ST}(0.926)(0.004) + \\
& k_{A,ST}(0.495)(0.004) + \\
& k_{A,ST,T}(0.926)(0.495)(0.004) \\
u(\text{Point B}) = \quad & k_T(0.490) + k_A(0.495) + k_{ST}(0.485) + \\
& k_{T,A}(0.490)(0.495) + \\
& k_{T,ST}(0.490)(0.485) + \\
& k_{A,ST}(0.495)(0.485) + \\
& k_{A,ST,T}(0.490)(0.495)(0.482) \\
u(\text{Point A}) = \quad & u(\text{Point B})
\end{aligned} \tag{3}$$

This process is repeated for the comparison between points B and C which result in another elicited value of travel and another equation. Additional sets of equations are generated by identifying similar points with equal total utility on other planes of constant conditional utility to make up at least six equations so that the unknown scaling constants can be found. Although only 2^n-1 equations are required, we have achieved better results by solving a set of more than 2^n-1 equations in a least-squares sense. The final multi-attribute utility function for the suspension decision is shown in Equation 4.

$$\begin{aligned}
u(T,A,ST) = \quad & (0.343)u(T) + (0.480)u(A) + (0.413)u(ST) + \\
& (-0.306)u(T)u(A) + \\
& (-0.217)u(T)u(ST) + \\
& (-0.386)u(A)u(ST) + \\
& (0.768)u(T)u(A)u(ST) \\
\textit{where} & \\
u(A) = \quad & (0.996)u(AU) + (0.966)u(AD) + (-0.962)u(AU)u(AD)
\end{aligned} \tag{4}$$

VI. Optimization

Having modeled all aspects of the design decision including the objectives, uncertainty, and preferences, the next phase of the design study is to optimize the expected utility to identify the preferred decision alternative. In theory, an optimizer can be wrapped around the energy-based system model and the Latin Hypercube Sampler to optimize the expected utility; however, multiple LHS runs compound with lengthy simulation times (about four minutes) to significantly increase the computation time to calculate the expected utility at a single point. Furthermore, most optimization algorithms require the evaluation of the objective function at several trial points in order to determine in which direction to move. Thus, there is motivation to reduce the simulation time of the suspension model in order to find a design solution in a reasonable amount of time. To reduce the complexity of the Latin Hypercube

Sampling, a screening experiment is conducted to determine the most significant of the uncertain factors influencing the suspension decision. This experiment is discussed in Section A. In addition, an adaptive kriging metamodel of the suspension model is implemented to reduce the simulation time, and this model is discussed in Section B. The optimization of the resulting simplified decision is presented in Section C.

A. Morris screening experiment and conclusions

To determine the most significant of the uncertain factors in the rover decision, a screening experiment is performed. Specifically, the method of Morris is implemented in this work to identify the uncertain factors which have significant effects on the utility.[13] 120 runs of the energy-based model are performed, and the utility of each run is calculated. The results are then summarized by plotting the effective standard deviation of each effect versus the effective mean of each effect. The resulting plot is shown in Figure 7.

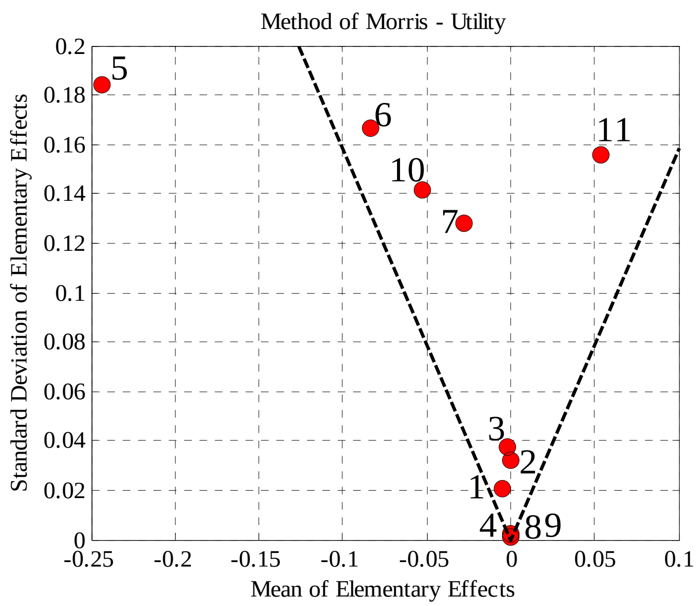

1	Spring Rate
2	Damping Coefficient
3	Payload Mass
4	Rover Location
5	Terrain Amplitude
6	Rover Speed
7	Angle of Approach
8	Torsional Stiffness
9	Torsional Damping
10	Axle Length
11	Wheelbase

Figure 7. Morris Plot of Suspension Utility

Large effective mean values indicate that the factor has a significant linear effect on the utility, whereas large effective standard deviation values indicate that the factor has either a significant nonlinear effect or a significant interaction effect. Points which lie within the “V” shape indicate that the effective mean value for that particular factor is not statistically distinct from zero. Based on these results the speed, angle, axle length, and wheelbase are found to be significant sources of uncertainty in the suspension decision due to the large effective standard deviations. Also, the terrain is found to be significant due to both a large effective standard deviation and a large magnitude effective mean. Therefore, the effects of these five uncertain parameters are modeled in the suspension decision using Latin Hypercube Sampling. The remaining four uncertain parameters (mass, location, torsional stiffness, and torsional damping) are thus held constant at their expected values for all samples. It is interesting to note that the mass and location factors are not significant as compared to some of the other uncertain factors, so these two factors can be ignored. Thus, the differences in gravitational forces on the two surfaces are not as significant to the suspension decision as the terrain, driving habits (speed and angle), and geometry of the rover (axle length and wheel base).

The design variables (spring rate and damping coefficient) are also included in the experiment to determine the relative significance of the design variables with respect to the uncertain parameters. It is seen that the design variables have an effective mean near zero and a small effective standard deviation. This means that the design variables are not significant relative to some of the uncertain factors. Because of this, it may be appropriate to choose values of design variables which are robust to this uncertainty rather than value of design variables that maximize the expected utility. Although it is possible to pursue a robust design using utility theory, it requires a reformulation of the multi-attribute utility function to include attributes relating to the variation of each attribute. This approach is not pursued here, but is left for future work.

B. Kriging model

To reduce the computation time, an adaptive kriging model is implemented. The kriging model approximates the multi-attribute utility of the rover as a function of all the design and uncertain variables. The kriging algorithm is an adaptation of the algorithm available in the DACE kriging toolbox for Matlab developed by Lophaven, Nielsen and Søndergaard.[13] Typically, such a kriging model is generated once based on a dataset obtained through a space-filling design of experiments, such as a Latin Hypercube Sample. Based on this initial sample of the design space, a surrogate kriging model is created, and from that point forward only the surrogate model is considered. However, by sampling the entire design space evenly, a lot of resources are wasted characterizing the utility function in regions of the design space that are far removed from the optimal solution. Instead, we have implemented a modification of the kriging algorithm in which the design space is sampled more densely in areas of interest, as needed, in an iterative fashion, based on the predicted mean-square error of the interpolating Gaussian process. If the predicted error is larger than desired, then the utility function is evaluated based on the underlying high-fidelity Dymola simulation, and the kriging model is regenerated. Each time, the size of the dataset on which the kriging model is based increases by one data point. When the design space is large or when the desired accuracy of the kriging model is very small, the number of data points needed to reach the desired accuracy can be large. This poses a problem because the time required to generate the model is approximately proportional to the number of data points to the 4th power, meaning that computing the kriging model may take longer than performing the underlying simulation as the number of data points increases. To overcome this problem we have modified to kriging algorithm to include only the N nearest data points. If N is sufficiently small (e.g., N=200), then the cost of occasionally recomputing a small kriging model is smaller than the cost of computing one very large kriging model. Based on our experience with using this adaptive kriging algorithm in the nested loop of optimization and sampling-based estimation of the expected utility, the speed-up over direct sampling of the underlying simulation can be several orders of magnitude. This allows us to use large sample sizes for a better estimation of the expected utility.

We have implemented this adaptive kriging algorithm as a driver module in ModelCenter.[14] The algorithm is set up such that it is completely transparent to the user as to when the surrogate model or the high-fidelity underlying model is used. The driver module first generates a coarse, space-filling dataset and computes an initial kriging model. This may take some time and can best be performed in parallel, on a computer cluster. Thereafter, anytime an optimizer or uncertainty quantification algorithm requests a utility value, the kriging module determines whether the predicted error is within the desired accuracy. If the predicted error is not within the desired accuracy, the kriging module regenerates a local kriging model, possibly after first generating an additional data point. The kriging module incorporates and stores all these additional simulation runs as it proceeds, allowing the optimization to resume immediately where it left off if for any reason there is an interruption in the computation process.

C. Optimization

The expected utility of the suspension is optimized both deterministically and non-deterministically to gauge the effect of the uncertainty on the design solutions. But first, the design space is explored to find a reasonable starting point for the optimization. A full factorial experiment is conducted over ten levels of the two design variables while holding all the uncertain factors at their expected values. A surface plot of the data is then consulted to determine the shape of the utility surface and the location of the maximum. Two more full factorial experiments are conducted at the upper and lower limits of the terrain distribution, because the terrain is the uncertain parameter that is found to have the most significant effect on the utility in the screening experiment. By viewing the three surface plots together, the effect of the uncertainty in the terrain on the height and shape of the utility surface can be seen. The three surface plots are shown in Figure 8. The highest surface corresponds to the lowest terrain value, while the middle surface corresponds to the fully deterministic experiment and the lowest surface corresponds to the highest terrain value. It is seen from these plots that the terrain has a strong negative effect on the magnitude of the utility surface as well as some effect on the shape of the utility surface. The overall shape of the utility surface at all three terrain levels is mostly flat, with a peak occurring near the lower bounds for both design variables. Therefore, the starting point for the optimization runs is chosen as the lower bound for both design variables. It is expected that this point is the true maximum for the deterministic optimization; however, due to the effect of the uncertainty, the maximum expected utility of the non-deterministic utility surface may be different, although it is expected to be in the neighborhood of this initial point.

The deterministic optimization is performed by holding all the uncertain factors constant at their expected values and optimizing the utility. A gradient optimizer is used with the starting point at the lower bound of each design variable. The optimizer converges to the maximum utility at a spring constant of 12000 N/m and a damping coefficient of 5000 N/m/s. This point is the same as the optimization starting point, but additional starting points are tested and the results are consistent. This finding is consistent with the expected results based on the deterministic surface plot in Figure 8.

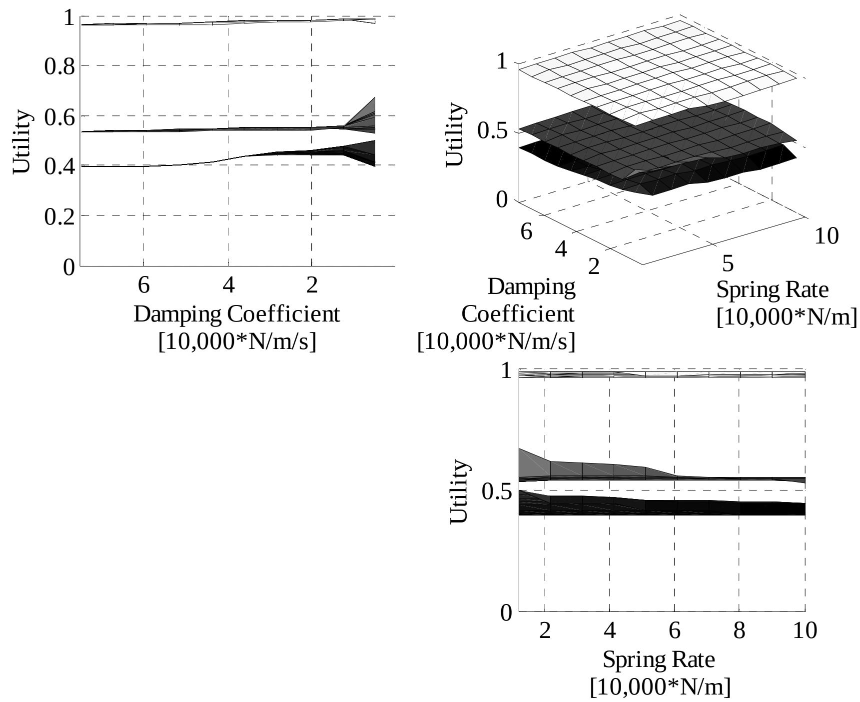

Figure 8. Utility Surface Plots at Three Values of Terrain Amplitude

The non-deterministic optimization is performed using Latin Hypercube Sampling to estimate the expected utility. A gradient optimizer is used again with the same starting point at the lower bounds of both design variables. It is important to note that the sampler uses common random numbers for each set of Latin Hypercube runs. This introduces a small bias in the estimate, but ensures that the gradient optimizer is able to compare points in the design space without noise due to the sampling. The optimizer again converges to the maximum expected utility at a spring constant of 12000 N/m and a damping coefficient of 5000 N/m/s, which is the same as the result from the deterministic optimization. Although the preferred decision alternative is the same in both the deterministic and non-deterministic case, the value of the expected utility differs due to the uncertainty in the problem. In the deterministic case, the utility at the preferred decision alternative is 0.671, while in the non-deterministic case the expected utility is 0.624 with a standard deviation of 0.157. Because the preferred decision alternative is the same for both the deterministic and non-deterministic cases, the decision is insensitive to the uncertainty in the problem.

In general, a utility optimization problem is not bounded; however, in this study, lower bounds are implemented for the spring rate and the damping coefficient. These bounds are imposed to prevent the Dymola simulation from reaching states that are not possible in a physical system. For some combinations of damping coefficient and spring rates, it is possible for the suspension model to reach an inverted state. This is because there is no collision detection between the various members of the suspension system incorporated in the Dymola simulation. Thus, the suspension members can pass through each other during the simulation, which of course is not possible in a physical system. If the spring rate and damping coefficient are large enough, this type of inversion does not occur in the model. In this study, the lower limits of spring rate and damping coefficient are identified by running simulations and viewing the resulting animation of the suspension behavior, and these limits are imposed as lower bounds. Another way to handle this problem is to incorporate a non-linearity in the spring rate to simulate reaching the solid length of the spring. This type of approach is more appropriate in the context of utility theory because bounds are not needed on the design variables.

VII. Closure

In this paper we have demonstrated the modeling and analysis of a design decision regarding the suspension parameters for a manned rover. The design decision itself is modeled using objectives hierarchies, which assist the design team in identifying the fundamental and means objectives involved in the decision, as well as an influence diagram, which helps the design team to identify the decisions, computations, and uncertain events that influence the achievement of the design objectives. In accordance with normative decision theory, the uncertain aspects of the decision are represented by subjective probability distributions, and the preferences of the design team are modeled with a multi-attribute utility function. When combined with a system model, these models enable the estimation of the expected utility of each design alternative through Latin Hypercube Sampling. An optimizer is then used to find the design alternative with the maximum expected utility.

In this study we found that the maximum expected utility occurs at a spring constant of 12000 N/m and a damping coefficient of 5000 N/m/s; however, this finding is based on the many assumptions made in the modeling of the system, the uncertain factors, and the preferences of the design team. Using subjective probabilities, we have captured our beliefs about the uncertain factors influencing this design decision, but subject matter experts will have different beliefs as a result of their expertise in the area, which will result in different subjective probability distributions and a different design solution. Similarly, we have captured our preferences concerning the performance attributes of the rover, but our preferences will differ from the preferences of other decision makers who have more knowledge of the working environment and requirements for a manned rover. Different preferences will again result in a different multi-attribute utility function and a different design solution. There are limitations, as well, related to the suspension system model. For simplicity, our model captures only vertical travel, acceleration, and settling time of the center of mass of the rover. Also, because the test-frame model induces wheel displacements to represent the terrain, our model does not allow the rover to separate from the surface and fall back to the surface under gravitational acceleration. As a result, the acceleration data captured in the simulation is not representative of the true acceleration experienced by an actual rover. Due to these limitations, the results we have presented in the paper should not be taken as conclusive. Rather, we wish to demonstrate the method of modeling a design decision in accordance with normative decision theory.

Acknowledgments

We thank Dynasim AB and Phoenix Integration for the use of their software packages and technical support. This work was completed within the context of our semester project for *ME 6105: Modeling and Simulation in Design*; we wish to thank our classmates for their input through class discussion. Stephanie C. Thompson gratefully acknowledges the support of a National Science Foundation Graduate Research Fellowship.

References

[1]*The Vision for Space Exploration*. 2004, National Aeronautics and Space Administration.

[2]*Exploration Systems Architecture Study Final Report*. 2005, National Aeronautics and Space Administration.

[3]Hazelrigg, G. A., *A Framework for Decision-Based Design.* Journal of Mechanical Design, 1998. **120**(4): p. 653-658.

[4]Marston, M., Allen, J. K., and Mistree, F., *The Decision Support Problem Technique: Integrating Descriptive and Normative Approaches.* Engineering Valuation and Cost Analysis, Special Issue on Decision-Based Design: Status and Promise, 2000. **3**: p. 107-129.

[5]von Neumann, J. and Morgenstern, O., *Theory of Games and Economic Behavior*, 3rd ed, Princeton University Press, Princeton, NJ, 1980.

[6]Keeney, R. L., *Creativity in Decision Making with Value-Focused Thinking.* Sloan Management Review, 1994. **35**(4): p. 33-41.

[7]Howard, R. A. and Matheson, J. E., *Influence Diagrams.* Decision Analysis, 2005. **2**(3): p. 127-143.

[8]Dymola, Multi-Engineering Modeling and Simulation, Software Package, Ver. 6, Dynasim AB, Lund, Sweden, 2007

[9]Frtizson, P., *Principles of Object-Oriented Modeling and Simulation with Modelica 2.1*. 2003, New York: Wiley-IEEE Computer Society Press.

[10]Keeney, R. L. and Raiffa, H., *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. 1976, New York, NY: John Wiley & Sons.

[11]Phillips, J. R., *ZunZun.com Interactive 2-Dimensional and 3-Dimensional Data Modeling*, [online software], URL: zunzun.com, [cited 2007 May].

[12]Morris, M. D., Factorial Sampling Plans for Preliminary Computational Experiments. Technometrics, 1991. **33**(2): p. 161-174.

[13]Lophaven, S. N., Nielsen, H. B., and Søndergaard, J., "*Dace: A MATLAB Kriging Toolbox, Version 2.0,*" Technical University of Denmark, Technical Report IMM-TR-2002-12, Kgs. Lyngby, Denmark, August 1, 2002.

[14]ModelCenter, Software Package, Ver. 7.1, Phoenix Integration, Inc., Wayne, PA, 2007.