Using DPL
for
ISS Trade Studies

Dave Greeson
Dynacs Engineering Co., Inc.

International Space Station Program Office
Vehicle Effectiveness & Trades
Analysis & Integration Team

14 November 1995

Abstract

This white paper describes the methodology and process for applying DPL (Decision Programming Language) to complex engineering trade studies on the ISS (International Space Station) program. The benefits and rationale for using DPL to conduct analytical trade studies and risk assessments are also discussed. The paper has been distributed to ISS team leaders to make them aware of the capabilities and potential for DPL to support their efforts. DPL is an open-ended tool that can be applied to virtually any complex decision analysis problem where the major variables can be quantified with a reasonable degree of accuracy. DPL can also be used to conduct risk analyses in situations with uncertainty. Typical decision analysis trade-offs involve a combination of multiple attributes such as technical performance, cost, schedule, and risk. A simplified example is used to illustrate the concepts of the influence diagram, the decision tree, the objective function, and sensitivity analyses. Four examples of actual trades performed on the ISS program are also described and discussed to illustrate how DPL has been used to date.

Using DPL for ISS Trade Studies

Dave Greeson
Dynacs Engineering Co., Inc.

International Space Station Program Office
Vehicle Effectiveness & Trades
Analysis & Integration Team

November 14, 1995

Contents

Introduction and Overview

The Vehicle Effectiveness & Trades Analysis & Integration Team (VE&T AIT) is tasked with supporting and performing trade studies for the International Space Station (ISS) Program Office. The trade study charter of the VE&T AIT includes helping other teams to perform complex trades within their own areas of responsibility and also whenever they encounter issues that cross functional boundaries and affect multiple subsystems. The support provided by the VE&T AIT to other teams includes ROM (Rough Order of Magnitude) cost estimates, issue definition, presentation support, risk analysis, and decision analysis. The subject of this white paper is how the VE&T AIT uses DPL (Decision Programming Language) for complex trade studies that require a quantitative decision analysis. DPL has been used on approximately a dozen trade studies to date, ranging from relatively simple cases to extremely complex decision analysis problems that would have been very difficult to solve without DPL. In all instances where DPL has been used for a trade study, the primary objective has been to determine what course of action provides the most cost effective "bang for the buck" for the International Space Station program. DPL is used to make decisions that maximize technical performance while simultaneously minimizing cost, schedule, and risk. Figure 1 shows the general concept underlying a DPL analysis, where the goal is to reach certain cost, schedule, and performance targets with an acceptable margin of risk for each.

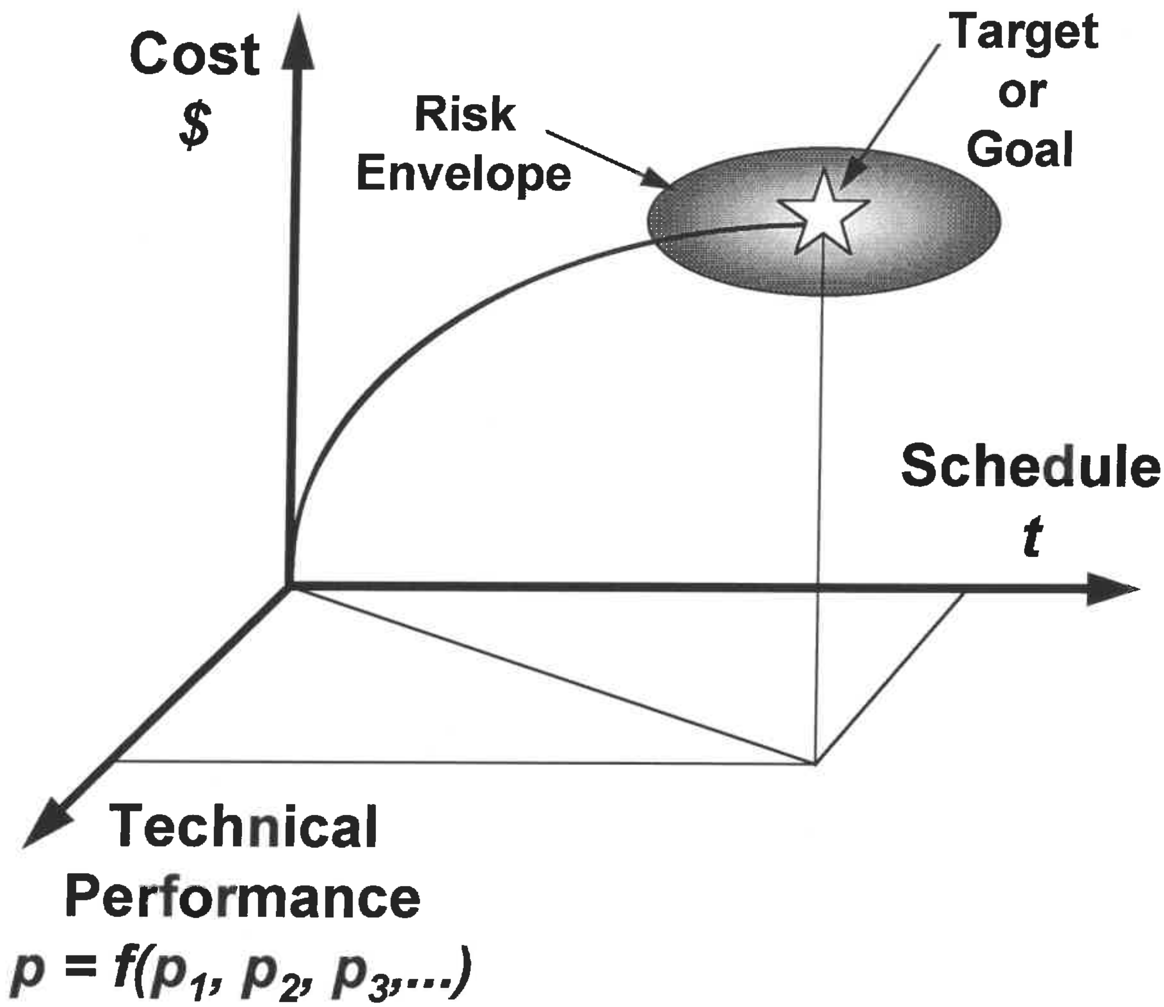

Figure 1, DPL Is Used To Achieve Cost, Schedule, and Performance Goals

DPL (Advanced Version) is used for complex trade studies that involve several technical options and uncertainties in cost, schedule, performance, and risk. DPL is particularly useful when there is no clear-cut choice between trade study options or one option is not obviously better than all the other options. DPL is an IBM PC Windows application developed by ADA Decision Systems, Menlo Park, California. DPL is a decision analysis tool that allows the user to graphically program complex influence diagrams and decision trees. DPL is a very flexible and powerful analytical tool that can be used to explore dozens of feasible alternatives and arrive at a decision based on a balance of cost, schedule, performance, and risk.

Several general categories of trade study analyses can be performed with DPL. The one that occurs most often is choosing between several feasible technical solutions. Another type of analysis that occurs is to use DPL's computational capabilities to generate the entire range of possible options, based on the decisions that have to be made. A third type of analysis that can be done is a time-phased risk analysis to determine the tradeoffs involved in delaying a key decision. Examples of all three types of analyses are presented below, in the section on actual ISS trade study examples. Other types of analyses are certainly possible with DPL, and it may be used in new ways as the need arises. Another logical use of the tool is to prioritize among many different tasks to achieve the most favorable cost vs. risk tradeoff.

An overview of the application of DPL in the trade study process follows. An example trade study which illustrates the analysis process is discussed below. Influence diagrams are used to capture the total scope of the decision analysis and enter the data for all the various cost, schedule, performance, and risk factors for each technical design option. Decision trees are constructed using the data entered in the influence diagram. Once the decision tree has been constructed, GET/PAY expressions for performance, cost, schedule, and risk attributes are applied to each branch of the tree as appropriate. An objective function can then be formulated to normalize these attributes to a common scale and apply weighting factors to each attribute. Weighting factors are determined through discussions with the trade study customer to derive the relative importance of the different attributes. A decision analysis is then run with the multi-attribute objective function. The output of the decision analysis is a decision policy tree showing the various options, which option is the optimal choice, and the expected utility for each option. The decision analysis also outputs a cumulative probability distribution that shows the range of expected utility. DPL can also be used to conduct parametric sensitivity analyses on uncertain variables, such as the probability that an event will or will not occur. In some cases, a sensitivity analysis will show that the optimal decision policy changes over the range of possible values for the uncertain variable.

Applying DPL to a Typical Trade Study

To illustrate the process of using DPL for a trade study, a simplified fictitious example has been generated. The example is a trade between Option A and Option B, which could be any two real-world technical design options. Option A has cost, schedule, performance, and operational risk factors to consider. Option B has cost, schedule, and performance factors to consider. In addition, the performance factor for Option B is dependent upon a chance event, whose probability is uncertain. Based solely on the numbers, it is difficult for the average person to make a choice between options, particularly as the number of different factors increases. Note that DPL could handle many more options (up to 20 branches on any decision or chance node). The example shown here has only two options for simplicity and clarity. The number of factors associated with each option is also not limited to four - there could be many more. The DPL software is capable of handling models with literally millions of endpoints - probably more than the typical PC hardware can handle (a large model slows down processing time in direct proportion to the number of nodes and branches).

The Influence Diagram

Figure 2 shows the influence diagram for the example trade study. In Figure 2, the rectangular node is a decision node, the elliptical nodes are chance nodes, and the rounded rectangular node is a value node. The arcs between nodes are known as influence arcs, and are used to define conditioning relationships between different nodes. Decision nodes are used to represent choices to be made by the decision-maker. Chance nodes are used to represent factors that are uncertain or events that the decision-maker has no control over. Value nodes are used to create variables, constants, or numerical combinations of other nodes. When the state branches of decision and chance nodes have values assigned, DPL automatically creates a variable using the node name. Note that the chance nodes representing uncertainties follow the decision, which reflects the typical real-world situation where a decision must be made without knowing the exact outcome of any choice. The implication is that all choices have some level of uncertainty and risk, which is consistent with common sense and everyday experience ("risk-free" is an invention of the advertising industry, not real-world systems engineering).

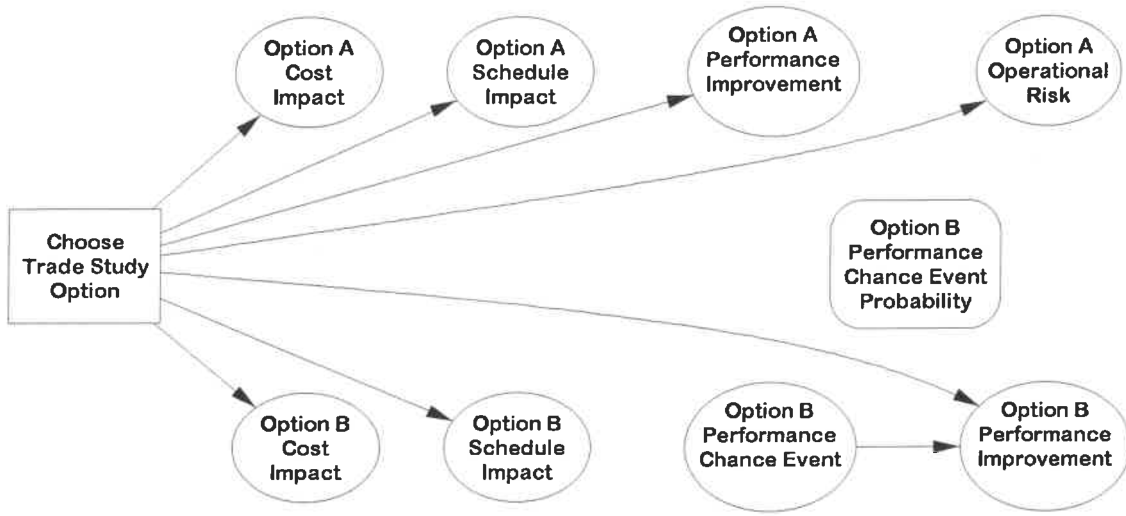

Figure 2, DPL Trade Study Example Influence Diagram

Figure 3 shows the underlying state branches of the decision node. Values could be assigned to the branches, but this is not necessary for the example at hand since the decision depends on the chance nodes associated with the option. DPL prompts the user to choose between maximizing or minimizing each decision node that is created. By maximizing or minimizing a decision, the user expresses a preferred outcome, such as minimum cost, or maximum utility. In a trade study where different attributes such as cost, schedule, performance, and risk are combined using an objective function, DPL will maximize or minimize the expected utility.

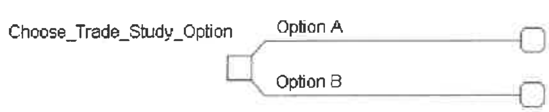

Figure 3, DPL Decision Node

Figure 4 shows the underlying state branches for several typical chance nodes from the example. In a chance node, probabilities of occurrence must be assigned to each state branch, and the sum of the probabilities must equal 1.0. Variables may also be used for the probabilities, as will be shown later. Values may also be assigned to each state branch to represent a numerical outcome. The number to the left on each state branch is the probability, the number to the right is the value associated with that branch. The numerical outcome can be in any units the user wishes, such as dollars, man-months, pounds, EVA hours, or any other parameter. The different classes or categories of numerical outcomes are referred to as 'attributes' in DPL. For a chance node representing a risk factor, subjective judgments based on expert opinion and experience can be captured and ranked numerically. Applying numerical rankings to subjective judgments is helpful when it is difficult or impossible to quantify a risk in terms of some common scale such as dollars, EVA time, or pounds. It is often impractical to "dollarize" all the risks, especially when all the possible outcomes of a chance node are not well defined or understood ("unknown unknowns"). When the different attributes are combined in an objective function, everything is normalized to a common arithmetic scale. The objective function allows an "apples to oranges" balance to be achieved. One commonsense guideline for combining different types of attributes is that the user should be consistent in his definition of each attribute's units. For example cost should be consistently expressed as dollars or man-hours, but not a combination. A combination could be used, but this would unnecessarily complicate the objective function.

As mentioned previously, each node may have up to 20 state branches. Fortunately, most chance events usually have a binary character - they either happen or they don't, so only two state branches are usually required to represent a chance event. Using more state branches to represent a chance event is advantageous when there is a wide range of expected outcomes. In any case, the sum of the probabilities assigned to the different state branches must add up to 1.0.

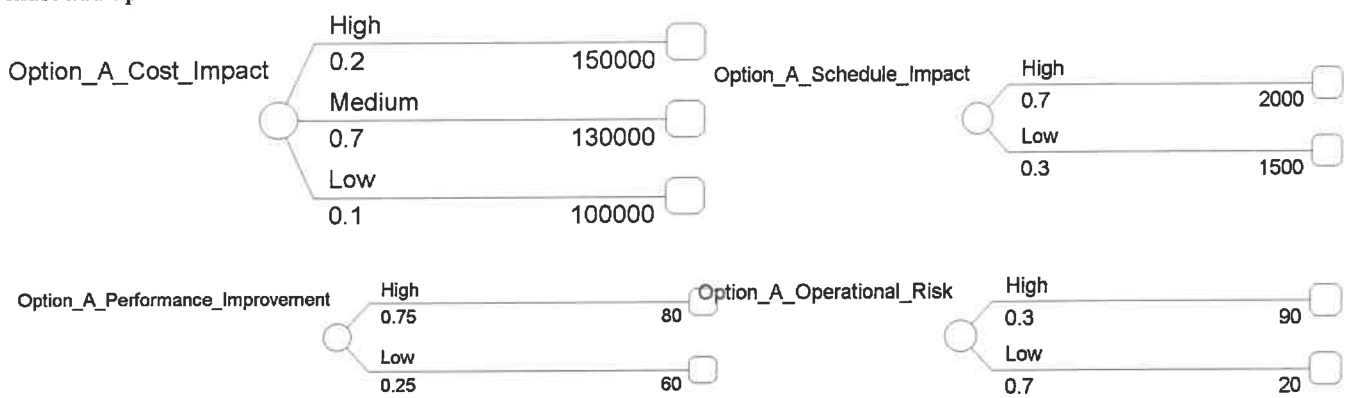

Figure 4, DPL Chance Nodes

The Decision Tree

Figure 5 shows how the various nodes in the example influence diagram have been combined to construct a decision tree. From the figure, it can be seen that the operational risk impact for Option A applies only in the case of a low performance improvement. For Option B, the performance improvement is dependent upon a conditioning chance event. If the event occurs, there will be some performance improvement whose value can be expressed probabilistically. If the chance event does not occur, there is no performance improvement for Option B.

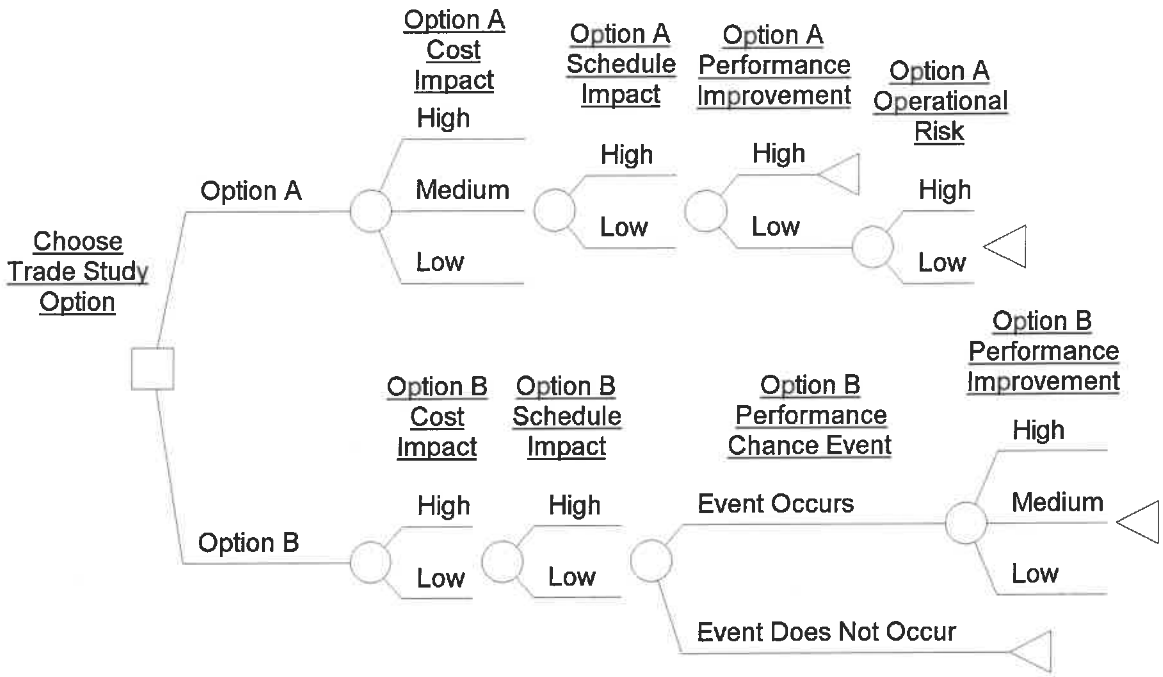

Figure 5, DPL Trade Study Example Decision Tree

Each branch of the decision tree has a GET/PAY expression assigned to it. For this example, four attributes are being carried - cost, schedule, performance, and risk. Cost, schedule, and risk branches in the decision tree have PAY expressions assigned to them. Performance branches have GET expressions assigned to them. The GET/PAY expression is a digital switching array, which selects the attribute to contribute to for any particular branch. The first attribute is cost, then schedule, then performance, then risk. Zeroes are used as placeholders for the attributes that do not apply to that particular node.

For example, the GET/PAY expressions for a cost node are shown below in Figure 6. The DPL variable Option_A_Cost_Impact is automatically created when the user creates the corresponding chance node in the influence diagram.

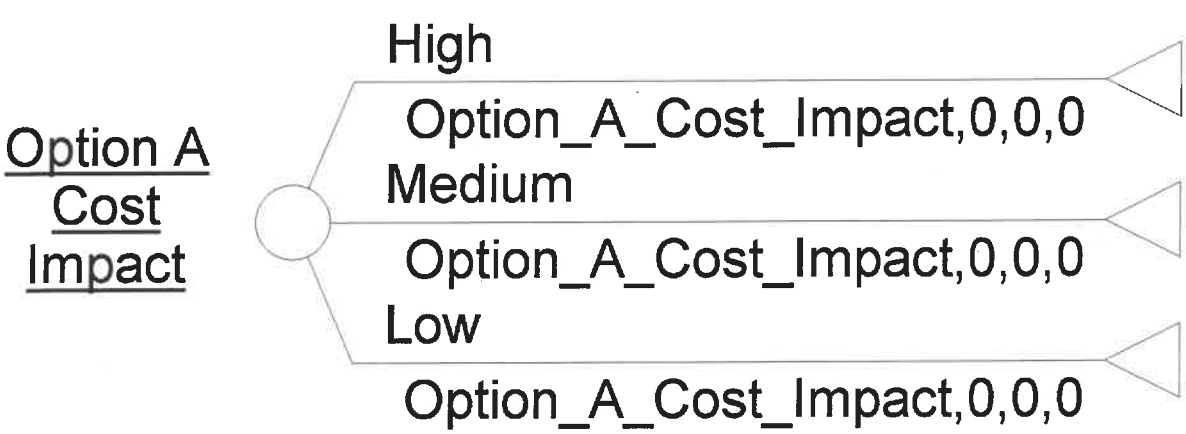

Figure 6, Chance Node GET/PAY Expressions

The Objective Function

The objective function is used to combine the various attributes being tracked through the decision tree. For this example, the objective function is as follows:

$$C_{sf}*C_W*\$1 + S_{sf}*S_W*\$2 + P_{sf}*P_W*\$3 + R_{sf}*R_W*\$4$$

where

C_{sf} = cost normalization scale factor = 100/150,000 ~ 0.00067
C_W = cost weighting factor = 0.5
$\$1$ = cost attribute (DPL placeholder variable)
S_{sf} = schedule normalization scale factor = 100/2000 ~ 0.05
S_W = schedule weighting factor = 0.05
$\$2$ = schedule attribute (DPL placeholder variable)
P_{sf} = performance normalization scale factor = 100/80 ~ 1.25
P_W = performance weighting factor = 0.35
$\$3$ = performance attribute (DPL placeholder variable)
R_{sf} = risk normalization scale factor = 100/100 = 1.0 (not required in the equation)
R_W = risk weighting factor = 0.1
$\$4$ = risk attribute (DPL placeholder variable)

The normalization scale factors are required to prevent the magnitude of the numbers for the different attributes from overwhelming the weighting factors. All the attributes must be normalized to a common scale. In this example, a scale of 0 to 100 was used. Any arbitrary scale can be used, such as -100 to +100, or 0.5 to 17.38. The scale selected makes no difference mathematically to DPL, it is strictly a matter of convenience for the analyst. Using a scale from 0 to 100 has proven convenient for trade study purposes, because people tend to think in terms of percentages and gradations from 0 to 100. When ranking subjective factors such as risk relative to one another, a scale from 0 to 100 is easily understood and no risk normalization scale factor is needed. A scale from 0 to 100 also provides adequate granularity to rank options relative to each other. The largest value in any of the branches for a particular attribute is then divided into the scale size to derive a scale factor for that attribute. In the example shown above, the cost attribute scale factor is 0.00067, because the largest cost number in the entire diagram was $150,000 and 100 divided by 150,000 equals 0.00067.

The weighting factors are required to achieve a linear combination of the attributes. The sum of the weighting factors should equal 1.0. Weights are usually assigned by the trade study customer, based on his particular priorities. In this example, cost is the most important attribute to the overall decision, with a weighting factor of 0.5. The next most important attribute is performance, at 0.35; then risk, at 0.1; and finally schedule, at 0.05. If there is a desire to perform a sensitivity analysis on the relative weights, value nodes can be created in the influence diagram and the weighting factors can be varied over a specified range using DPL's Value Sensitivity Analysis function.

The attributes, $1, $2, $3, $4, are simply placeholder variables for DPL to sum up the GET/PAY expressions for the particular branch being evaluated. In full enumeration mode, DPL calculates every possible permutation and combination of branches through the overall decision tree. DPL can handle up to 64 different attributes, if required. Most trade studies can be accomplished with three or four attributes, however. Every trade study eventually boils down to a balance between cost, schedule, performance, and risk. In some cases it may be necessary to distinguish between different aspects of a particular attribute, such as development costs versus operational costs, where there may be a desire to weight development costs differently than operational costs. In such a situation, another attribute is simply added to the objective function and then scaled and weighted appropriately.

Running the Decision Analysis

When a decision analysis is run with DPL, the optimal decision policy is output, as shown in Figure 7 below. The bolded line on Option B shows that it is the preferred option. The numbers in parentheses indicate the expected utility for that option. Expected utility is the most likely combination of the cost, schedule, performance, and risk attributes for that option. The expected utility is computed using the scaling factors and weighting factors in the objective function.

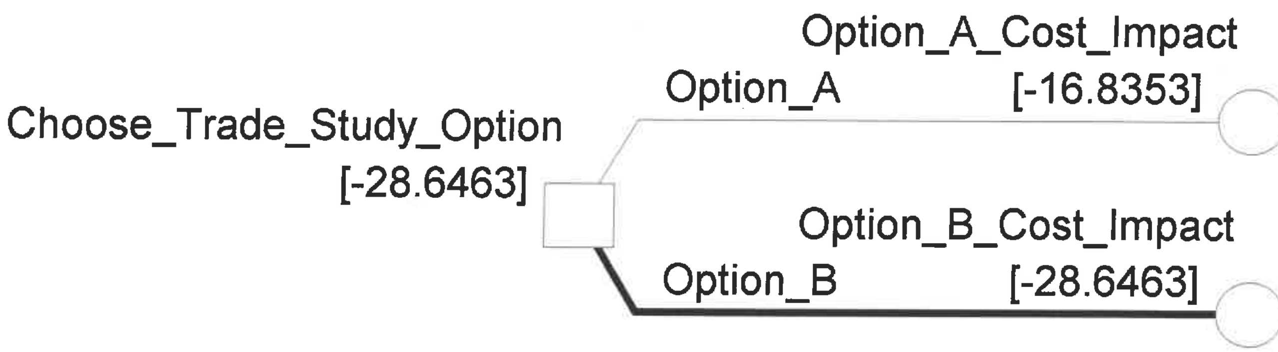

Figure 7, DPL Optimal Decision Policy

Figure 8 shows another output of the decision analysis, the risk profile, a cumulative probability distribution of outcomes (in this case, expected utility) for the optimal decision policy. This graph is useful for seeing where the decision policy changes from one option to another as a function of the expected utility, and also for learning what the potential maximum and minimum ranges are for the expected value. This result gives the analyst a good idea of what the full range of possible outcomes are and how likely they are to occur. From the graph, it can be seen that the range of possible expected values is from -90 to +5, and the optimal decision policy changes from Option A to Option B when the expected value is zero. The reason there is a switch in the decision policy is because of the chance node for Option B Performance Improvement. For Option B, the performance improvement is programmed as a GET expression, which makes a positive contribution to the expected utility. Negative impacts (cost, schedule, risk) are modeled as PAY expressions. The end result is that Option B maximizes the expected utility because it's positive GET contributions outweigh the negative PAY contributions.

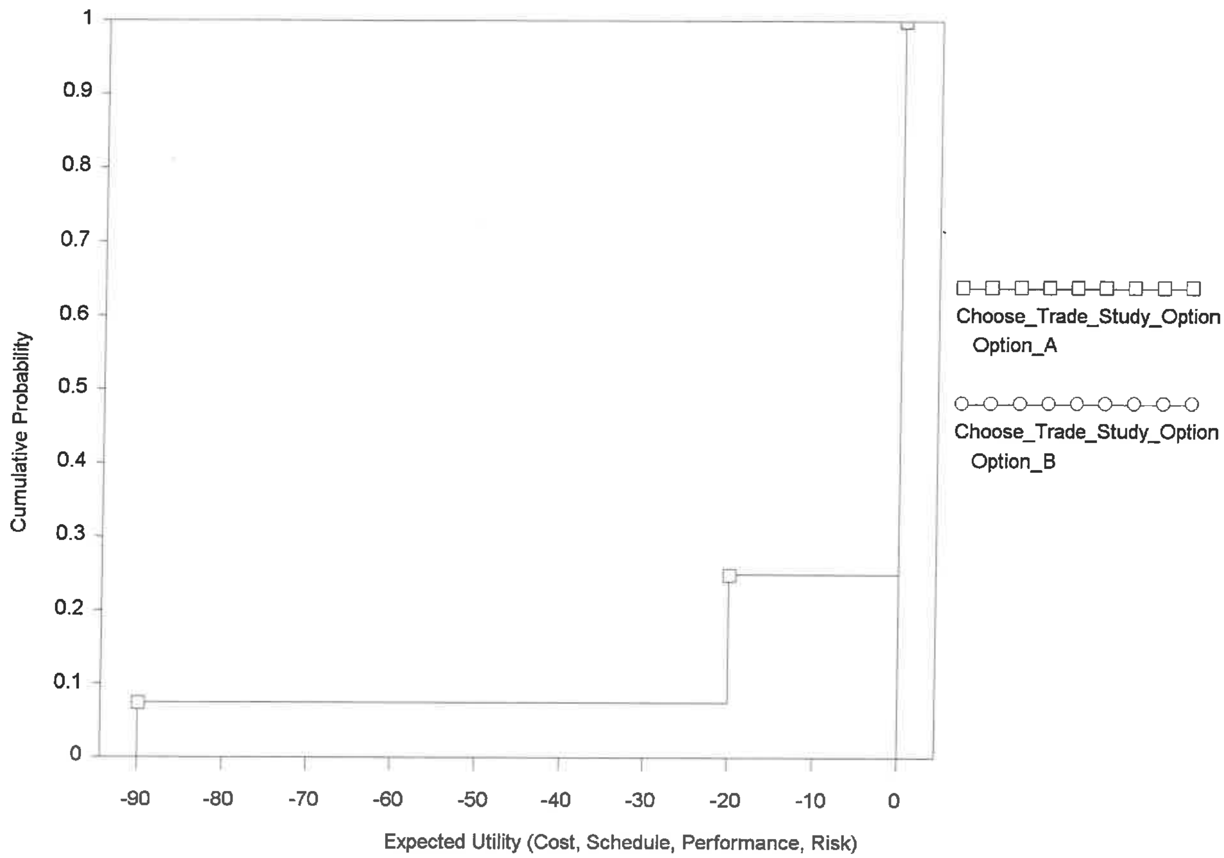

Figure 8, DPL Decision Policy Cumulative Distribution

Running Sensitivity Analyses

In most complex trade studies, there are uncertainties that cannot be specified exactly. In most cases, the probability of an event can never be known exactly. Other times , the resulting value of an event is not known precisely. In such situations, DPL provides the capability to perform three types of parametric sensitivity analyses. A parametric study allows the analyst to vary the specified parameter over a range of values to see what effect it has on the overall outcome. By varying the parameter or parameters and observing the resulting effects, the key drivers in the decision become clear and the break-even points where one option becomes more favorable than another can be determined. Such insights are usually key to making sure the problem is understood completely and nothing is "hidden" in the data. Sensitivity analyses provide insights into how sensitive the results produced by the model are to changes in the numbers in the model. In most decision analysis problems, an iterative approach is best. A preliminary model is constructed with preliminary estimates for probabilities and values, then the model and the data contained within it is refined until the analyst is confident the model is robust and accurately represents the real world situation. The three types of sensitivity analyses DPL provides are the rainbow diagram value sensitivity analysis, the tornado diagram value sensitivity analysis, and the tornado diagram event sensitivity analysis. Each type of sensitivity analysis will be discussed below, continuing with the fictitious trade study example.

Value Sensitivity Analysis (Rainbow Diagram)

Figure 9 below shows a rainbow diagram where the probability of occurrence for a chance event was varied from 0.01% to 99.9%. In other words, from a virtual impossibility to a virtual certainty. The diagram shows how the expected value (utility) and the decision policy changes as the probability of the chance event is varied. The variable, "Option B Performance Chance Event Probability", was entered into the model as a value node on the influence diagram (see Figure 1). At approximately 0.62% probability, the slope of the curve changes, as does the shading under the curve. The meaning of this is that the optimal decision policy changes as the probability is varied. When comparing this diagram to Figure 8, another insight into the problem is that even if the probability is varied across the full range, the expected value changes across a much smaller range than in the cumulative probability distribution.

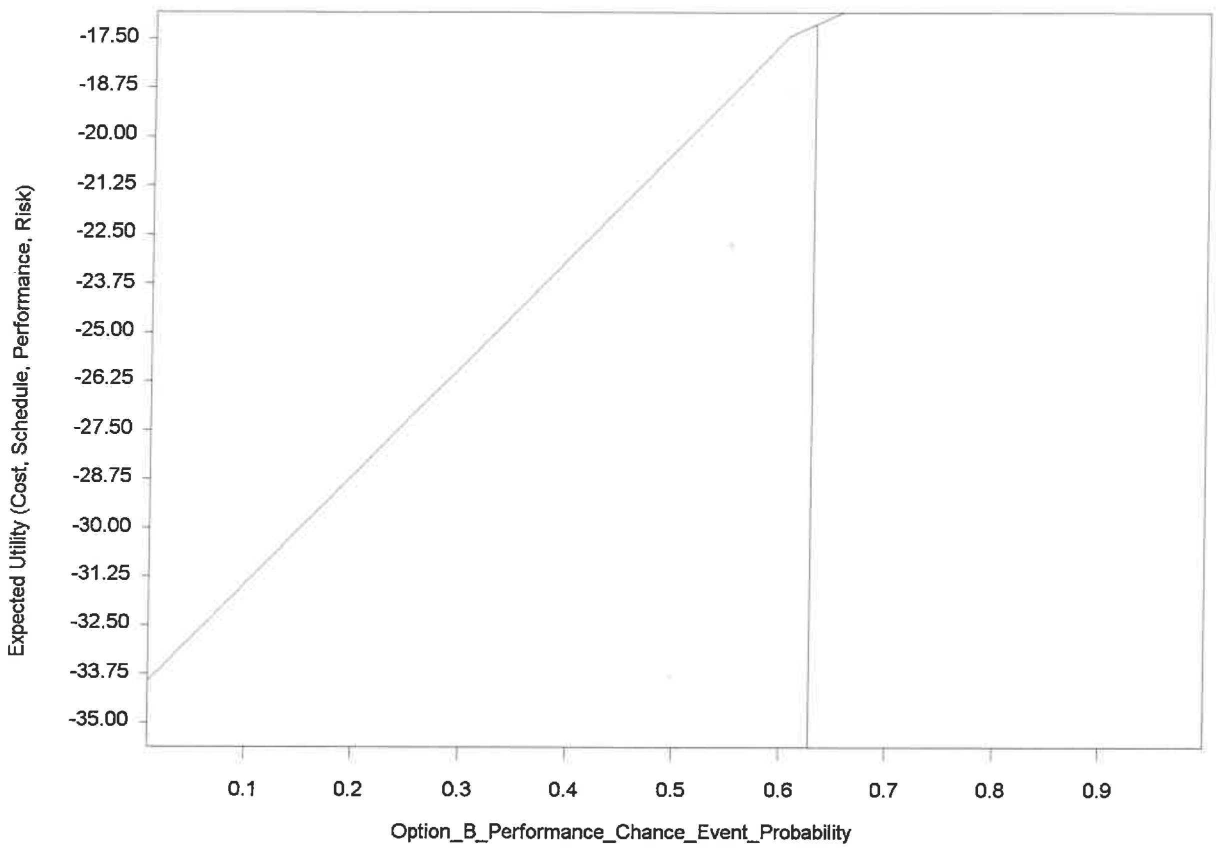

Figure 9, Value Sensitivity Analysis (Rainbow Diagram)

The main value of having a diagram such as Figure 9 is that it provides an insight into what leverage is available to the decision-maker to affect the outcome. If, for instance, the probability of occurrence of an event can be maximized or minimized through extra work or funding, it is more likely that a desired outcome can be achieved. On the other hand, if a probability is related to some event the decision-maker has no control over, such as a natural phenomenon or a political event, the rainbow diagram provides an insight into how vulnerable the desired outcome is to external circumstances. In some cases, the difference between two competing decision policies may turn on a key probability, whose value will be an assumption or educated guess. In the early stages of a complex study or analysis, there are usually many uncertainties, which can only be resolved through detailed engineering analyses or definitive cost estimates. For trade studies, however, a decision is usually required before precise information becomes available. The strength of DPL sensitivity analysis is that uncertain parameters can be modeled and varied to see what effect they will have on the decision.

In Figure 10, the chance event is shown, with the variable probability shown on one branch. The other branch, "Event Does Not Occur", has no probability assigned to it. When the probability variable is designated as the parameter for a rainbow diagram value sensitivity analysis, it is varied over the range of values the analyst specifies. The analyst inputs the starting and ending value, and the number of increments. When DPL evaluates the model, it computes a new expected value for each increment of the parameter across the specified range. In this case, the value is a probability on one state branch of a chance node. DPL automatically computes the probability for the other state branch as 1.0 - Current Parameter Value, so no value or variable is required to be entered for that state branch. Another important point about Figure 9 is that the chance node shown here has no outcome values specified, only probabilities. The chance node is being used to create a conditional branching scenario, based on a chance event. This capability is useful in modeling decisions and downstream courses of action that depend on downstream events. The usual situation, as the example model illustrates, is that a decision must be made before all the facts are available. This usually forces trade study decision trees to have a decision on the left side, followed by a series of chance nodes (refer back to Figure 4). This principle will also be shown in the Actual ISS Trade Study Examples section that follows.

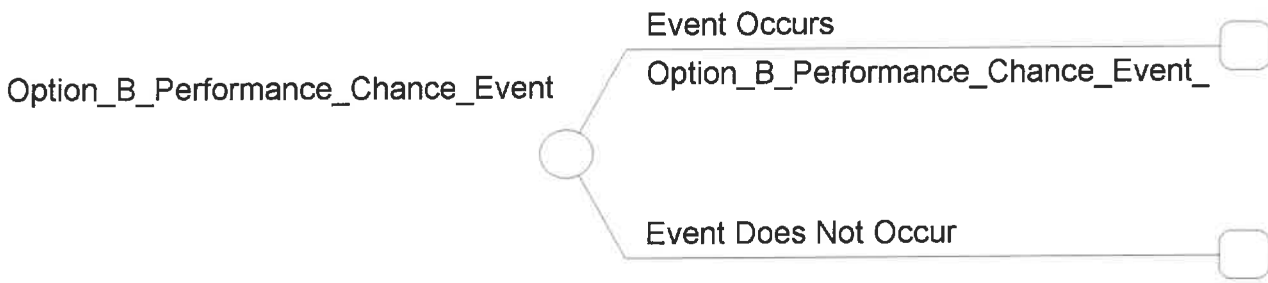

Figure 10, Chance Event With Variable Probability

Value Sensitivity Comparison (Tornado Diagram)

Figure 11 below shows a tornado diagram for a value sensitivity comparison of the example trade study. The diagram shows the relative importance of the model's variables as the values are varied over a user-specified range. The diagram can be tailored to show any desired grouping of variables in the DPL model. The diagram automatically ranks the selected variables from greatest to least influence, placing those with the greatest influence at the top of the diagram. The wider the bar for a variable is, the more influence it has over the decision analysis outcome. Where a bar changes color indicates a decision policy change, similar to the way the rainbow diagram indicates policy changes. The analyst must specify the starting and ending value for each variable to be added to the diagram. The solid vertical line indicates the expected value or expected utility for the nominal model, without variations. The numbers below the ends of each bar indicate first the value of the variable and then the resulting expected value. As can be seen from the figure, the variable "Option B Cost Impact" has the greatest relative effect on the decision, while "Option A Operational Risk has the smallest relative effect. The expected values and decision policy changes that are computed with the tornado diagram take full account of the weighting and scaling factors defined in the objective function.

It is important to note that the high and low values for each sensitivity variable are chosen by the analyst. The analyst should choose ranges for each variable that represent roughly the same degree of uncertainty so the comparison will be meaningful. For example, the uncertainty range for one variable might be specified as plus or minus 80%, while another may be specified as plus or minus 10%. The 80% uncertainty would likely overwhelm the 10% uncertainty and skew the results of the tornado diagram. A consistent uncertainty range for the variables is preferable, but not absolutely required. In some cases, the inherent uncertainty of some variables may be naturally higher than for others. The tornado diagram allows the analyst to examine the relative effects of all the uncertainties on a single diagram.

The significance of Figure 11 and the Value Sensitivity Comparison for trade studies is that it allows the analyst to vary all the model variables simultaneously and observe the relative effects of each. This capability is very useful when a model contains multiple uncertainties whose values may vary over a wide range. Another description of the value tornado diagram is "multivariate sensitivity comparison". The tornado diagram gives the decision-maker a quantitative insight into which variables are driving the decision and which ones are relatively unimportant. Sometimes, people involved in a decision may focus on factors that seem hugely important to themselves, yet play a relatively small role in the big picture. The tornado diagram can be used illustrate what factors are really driving a decision and clarify what leverage is available from each factor. Using the tornado diagram, we can see where additional effort or funding will really make a difference and where it will have only a small effect.

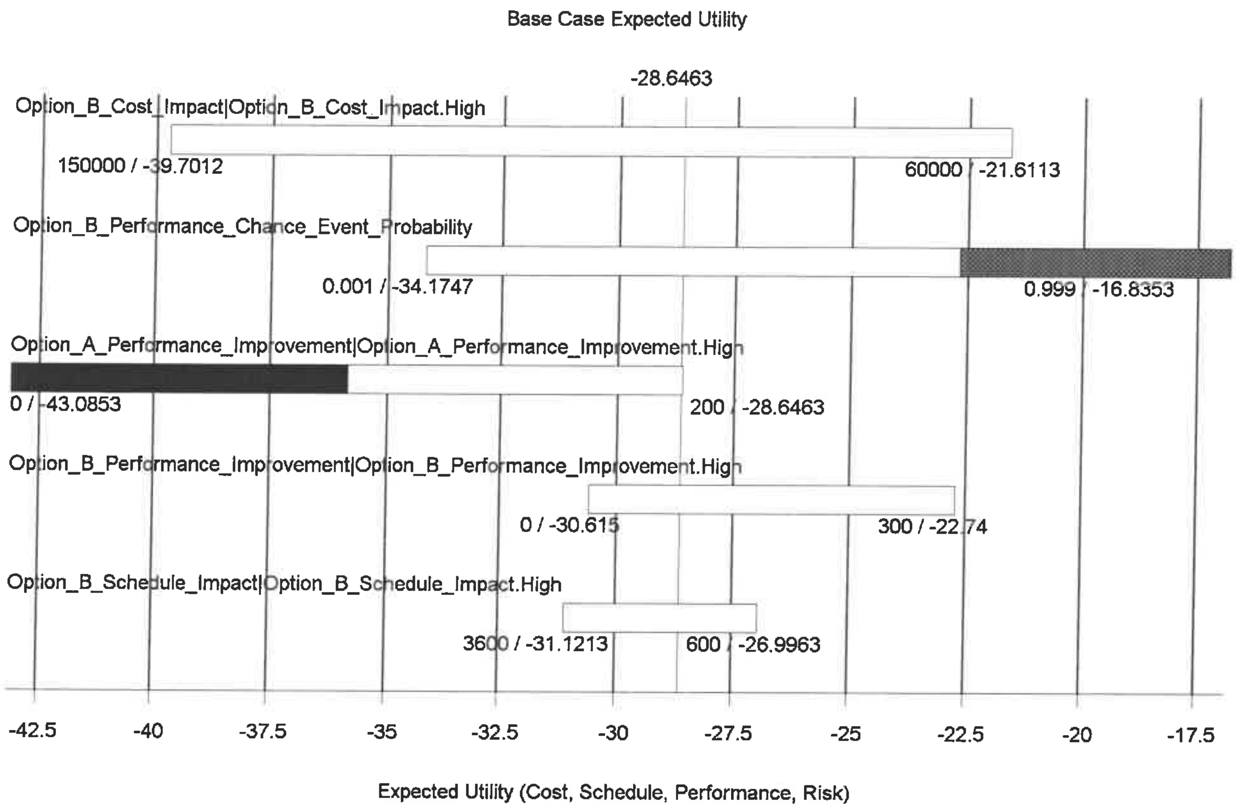

Figure 11, Value Sensitivity Comparison (Tornado Diagram)

Event Sensitivity Comparison (Tornado Diagram)

Figure 12 below shows a tornado diagram for an event sensitivity comparison of the example trade study. The diagram compares the relative effects of the uncertainties in the model. The diagram shows all the events (variables) in the model and their relative effects on the decision analysis, using the probabilities and values that were originally entered with the influence diagram. The event sensitivity comparison differs from the value sensitivity comparison in that the range of variation for the event variables is not controllable. As in the value sensitivity comparison tornado diagram, a change in the color or shading of a bar indicates a decision policy change. The event sensitivity comparison can be run several different ways - nominal, deterministic, combination, and probabilistic. Figure 12 below shows a nominal comparison. The nominal method is a deterministic sensitivity analysis that sets all uncertainties to their nominal states, then varies each one from a low state to a high state, keeping all other events at their nominal states. The solid vertical line is the *base case* policy line. The base case is the result when all events are set to their nominal value, and is not the expected value. From Figure 12, it can be seen that the variability of "Option A Performance Improvement" has the greatest effect on the end result, while "Option B Schedule Impact" has the least effect.

The other types of event sensitivity comparisons are useful in various situations. They each change how the base case bar is calculated and how the probabilities are handled. See the DPL Advanced Version User Manual for additional details on Event Sensitivity Comparisons.

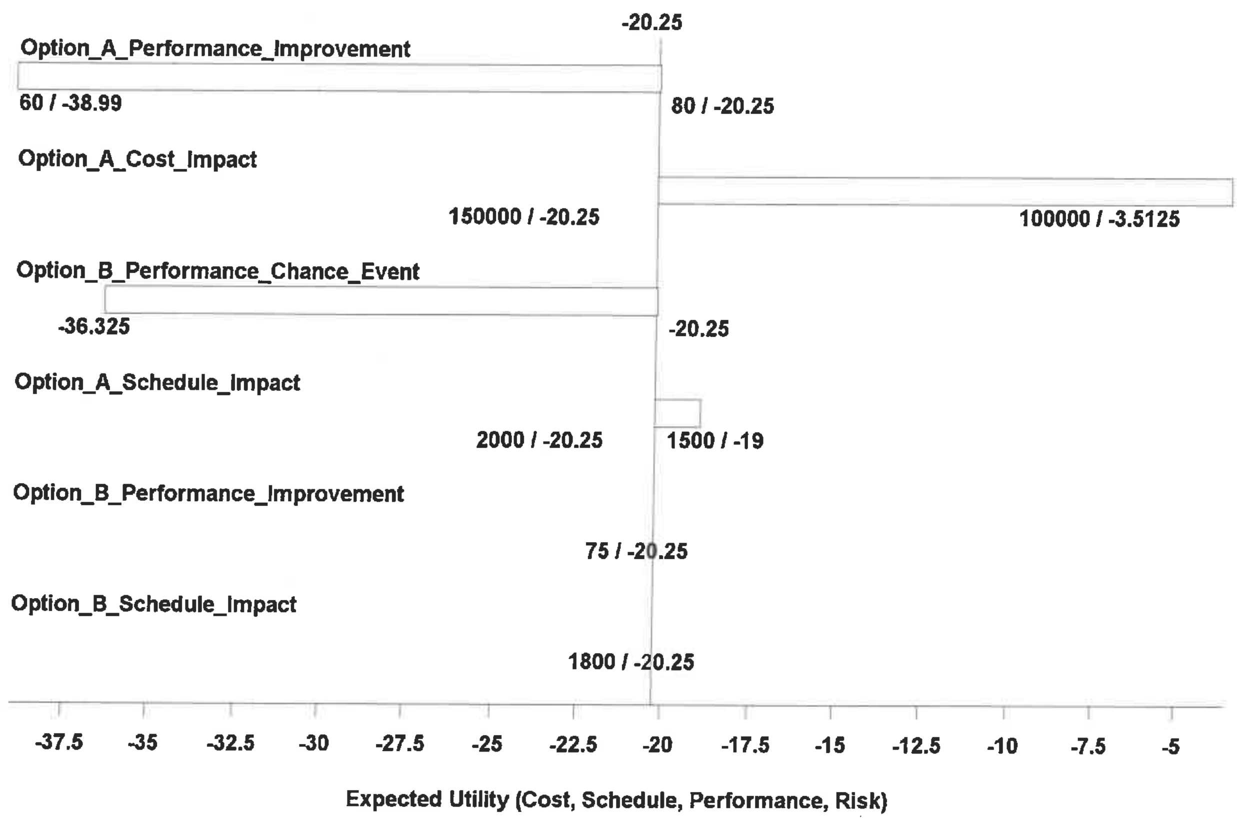

Figure 12, Event Sensitivity Comparison (Tornado Diagram)

Actual ISS Trade Study Examples

DPL has been used to support over a dozen different International Space Station (ISS) trade studies to date. It has been used in several different ways, as explained in the introduction and overview. There are undoubtedly other ways DPL could be used to support the ISS program. The Program Integrated Risk Management group has ordered four copies of DPL Advanced for use in risk analysis. DPL is flexible and adaptable to a broad range of uses, just like any other high-level programming language such as Ada or C^{++}. DPL is an object-oriented system with built-in optimization and Monte Carlo capabilities. DPL's real strength is that it is tailored to formulating and solving complex decision analysis problems. While there are other decision analysis tools available on the open market such as DATA, DEMOS, Criterium, and others, only DPL provides both the Influence Diagram and Decision Tree views. DPL Advanced is also particularly well-suited for multiple attribute decision problems, where a tradeoff between several attributes such as cost, schedule, performance, and risk is required. Nearly all real-world problems will involve some combination of these factors. The representative examples on the following pages illustrate how DPL has been used to date on the ISS program and what the benefits of using DPL are for complex decision analysis problems.

GPS RPU (Global Positioning System Receiver/Processor Unit)

This trade study was performed at the request of the Guidance, Navigation, and Control (GN&C) AIT team leader. The issue that prompted the trade study was that the GPS RPU does not meet certain requirements for functioning in a depressurized environment. The GPS RPU is critical for the end-to-end functionality of the GN&C subsystem, since it is the only source of attitude data on the United States On Orbit Segment (USOS). This trade study resulted in the most complex DPL model attempted to date, mainly because of the lack of definitive data. Nearly every aspect of the problem was uncertain, requiring multiple chance nodes to represent all the different variables. Several iterations of the model were made, each more complex than the last. The final model had nine different technical options, four decision nodes, and sixty chance nodes. Figure 13 below shows the influence diagram for this trade study. In this DPL model, the influence diagram was used as a problem definition and data entry tool, not to solve the problem. To actually select the appropriate design option, a customized decision tree was created. See Figure 14 for the overall GPS decision tree structure. The decision nodes are on the left, branching into the various options, each of which has several chance nodes associated with it. This trade study involved cost, schedule, and risk. Technical performance was not a driving factor in the analysis.

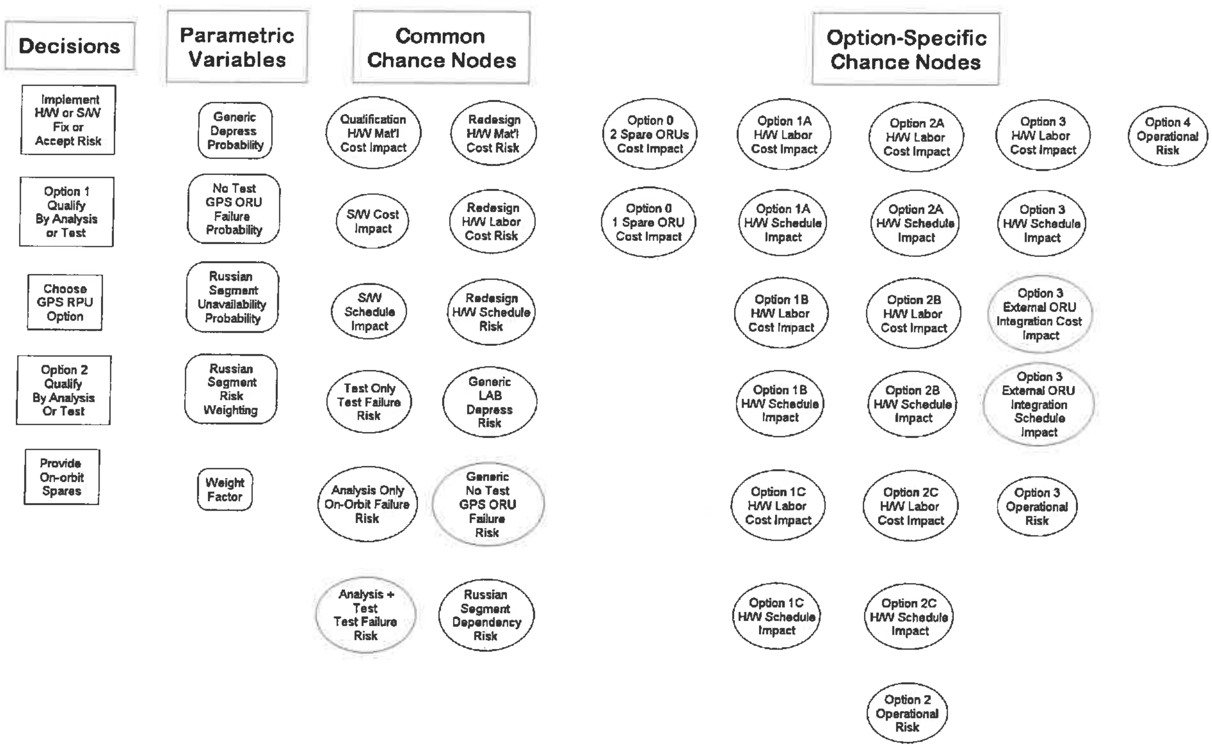

Figure 13, GPS RPU Trade Study Influence Diagram

Figure 15 below shows a zoomed-in view from a small section of the overall GPS decision tree of Figure 14. The highlighted section is from the middle right of the overall diagram. The branches that are shown in the figure were for Options 2A, 2B, and 2C which had uncertainties in hardware cost, hardware labor, hardware schedule, software labor, software schedule, operational risk, Russian segment dependency risk, test failure risk, lab depressurization risk, on-orbit failure test, redesign cost risk and redesign schedule risk. As the figure shows, most of the uncertainties are modeled using three branch chance nodes. This technique is analogous to using a three point estimate for risk analysis (high, medium, and low or optimistic, expected, and pessimistic). More than three branches can be used, but usually add neglible value or further insight into the problem. Chance nodes with only two branches are used to model conditional branching events, such as failures.

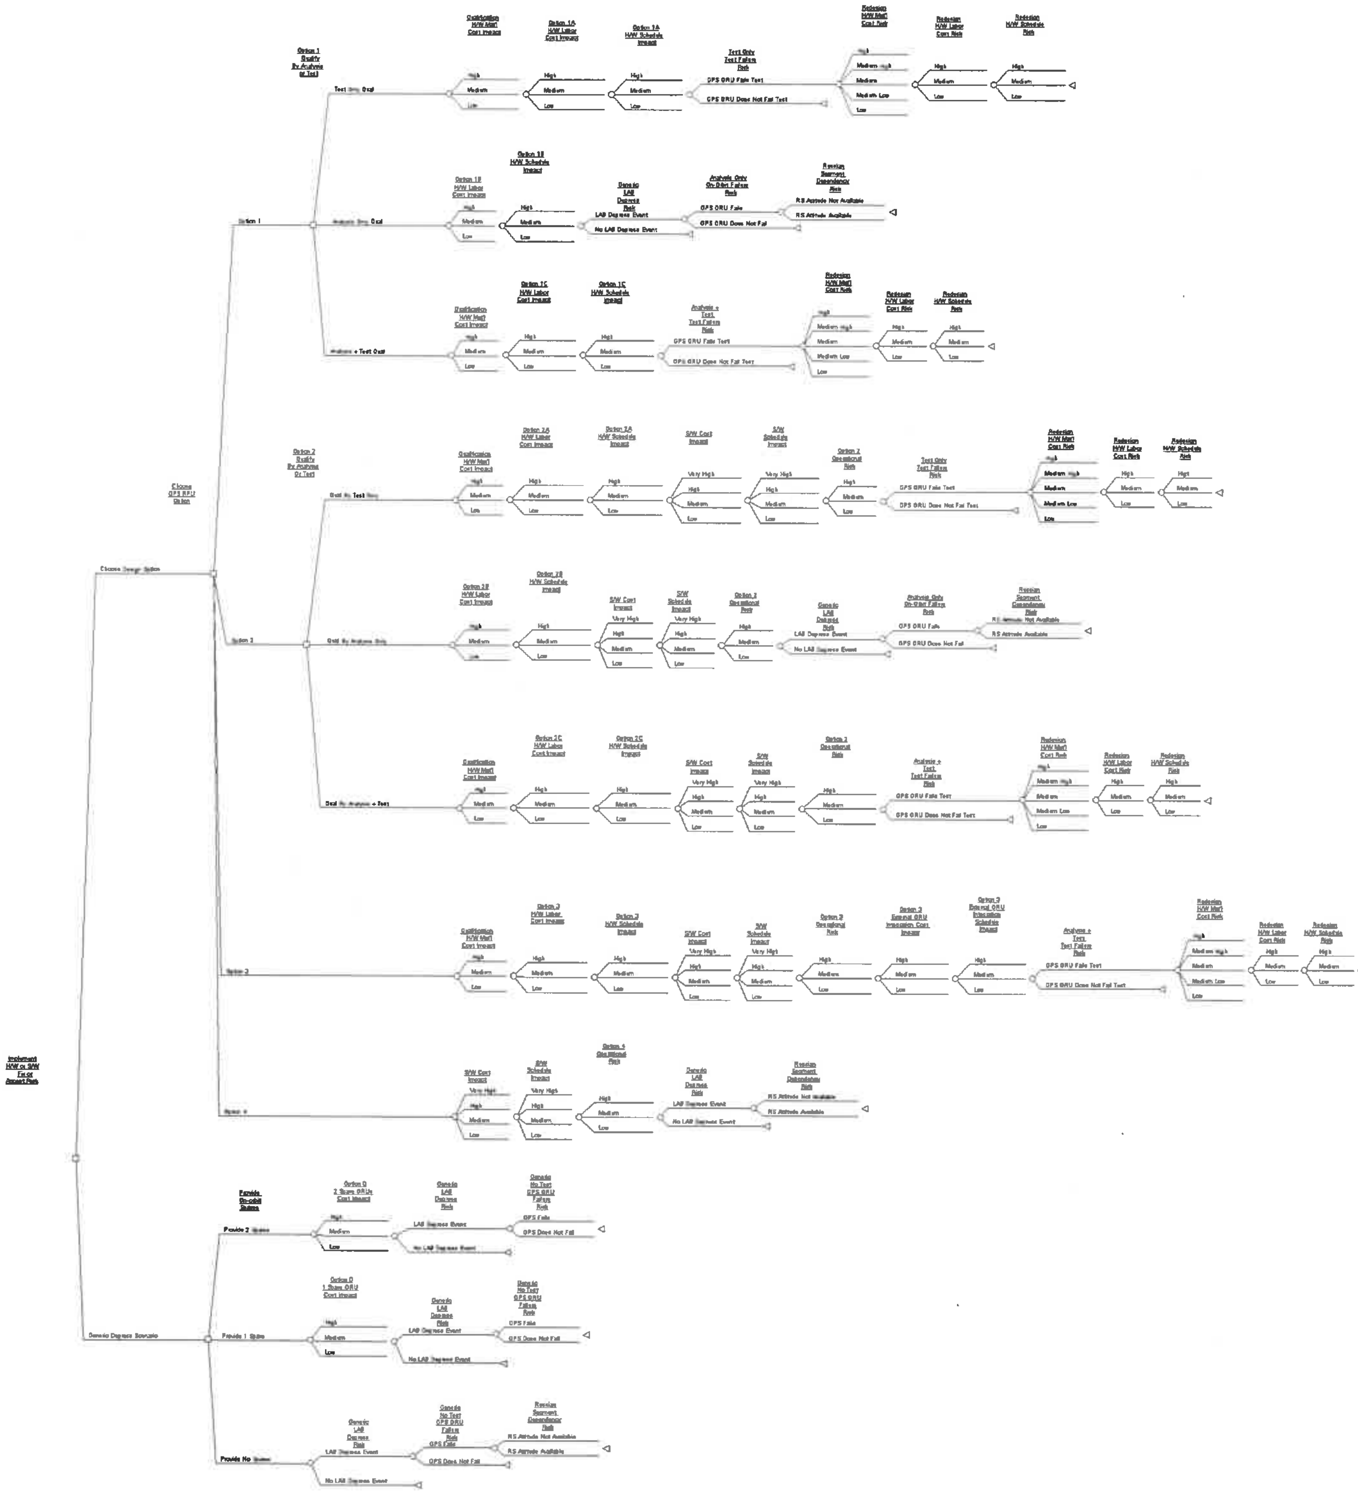

Figure 14, GPS RPU Trade Study Decision Tree

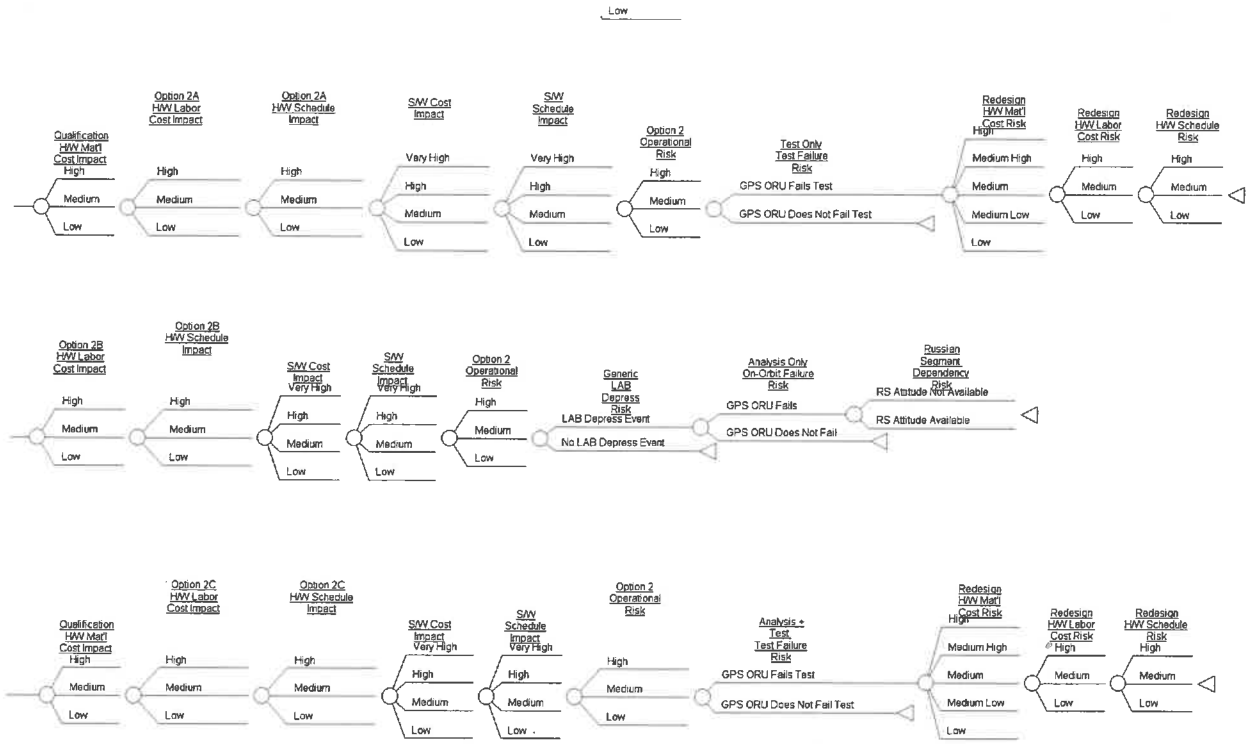

Figure 15, Subsection of GPS Trade Study Decision Tree

The "answer" provided by the DPL model was a selection of the best option, based on the preferences of the trade study customer. The customer's desired weighting factors were 40% cost, 20% schedule, and 40% risk. Those weighting factors were used in the objective function, along with the appropriate scaling factors. Figure 16 below shows a sample decision policy from the DPL model. The bold path is the optimal decision policy, and the numbers in brackets are the expected utility for each decision path. As the figure shows, the optimal policy was to accept the risk of a depressurization and make no changes to the GPS. DPL chose this option based on the extremely low probability of a depressurization and the relatively low weighting put on risk vs. cost and schedule. The next suboptimal decision policy was to implement Option 4, which involved changing the software to accomodate a failure. The next two suboptimal policies were implementation of Option 1A or Option 2A, both of which involved hardware changes. Option 4 was the recommended option to alleviate safety concerns, regardless of the inherently low risk involved. DPL provided a quantitative analysis of the problem that had been lacking before. Initially, it was very unclear which option was best from an objective systems engineering point of view.

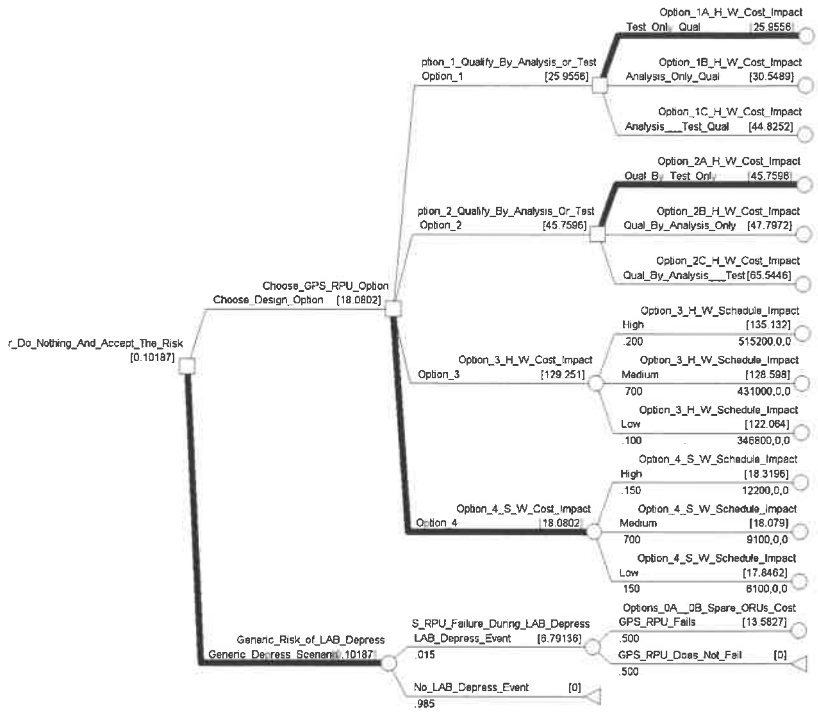

Figure 16, GPS RPU Trade Study Decision Policy

To gain further insight into the problem, several sensitivity analyses were run. One analysis replaced the two identified scenarios that resulted in a laboratory depressurization with a single generic event. The probability of that event occurring was varied to see what effect, if any, it would have on the final outcome. Figure 17 shows the resulting rainbow diagram of the value sensitivity analysis that was run. As can be seen in the figure, the optimal decision policy did not change when the probability was varied from 0.001 to 0.999. If the optimal policy had changed, there would be a change of slope and a change of shading or pattern in the diagram.

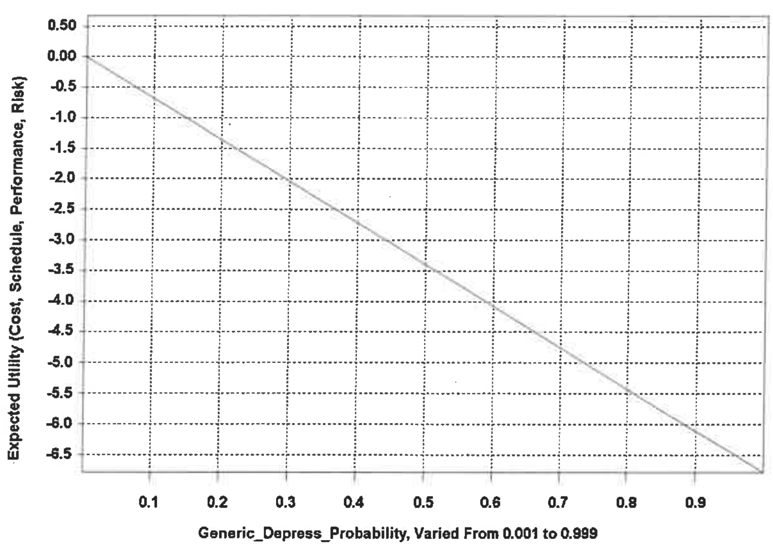

Figure 17, GPS RPU Trade Study Sensitivity Analysis

EEATCS (Early External Active Thermal Control System)

Another complex trade study that was performed looked at the cost vs. risk tradeoffs involved in selecting a design option for the EEATCS. The trade study was inititated because of concerns that the EEATCS ammonia fluid lines could not withstand the expected micrometeoroid and debris environment. If the lines were puntured or otherwise damaged, the ammonia could leak out and reduce the cooling capability of the EEATCS. A reduction of the EEATCS cooling capacity would reduce the amount of electrical power available to the Station, thus resulting in other systems having to be shut down or cycled on and off. Nine different technical options were identified, including choices between aluminum, nextel, and no shielding. The other major trade decision was whether or not to buy fluid line repair and reservicing capability. An objective function was used to give relative weighting to the cost and the risk. The trade study DPL decision policy appears below in Figure 18, with the optimal policy and suboptimal policy lines in bold.

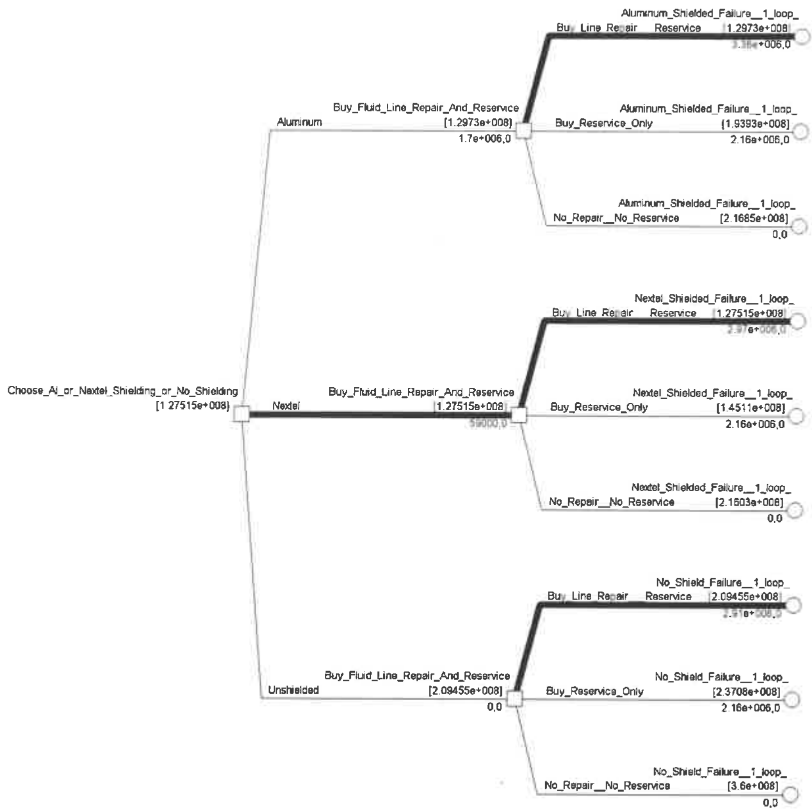

Figure 18, EEATCS Trade Study Decision Policy

The influence diagram for this trade study is shown in Figure 19. The diagram illustrates the overall problem structure. In this case, there were two decision nodes, three parametric sensitivity variables, and twelve chance nodes. The influence diagram was used more for organizing the problem structure and entering the data than for running the analysis. Not that there are no influence arcs in the diagram. The decision tree was used to actually solve the problem. The influence diagram view provides a convenient way to enter the data and create variables in DPL.

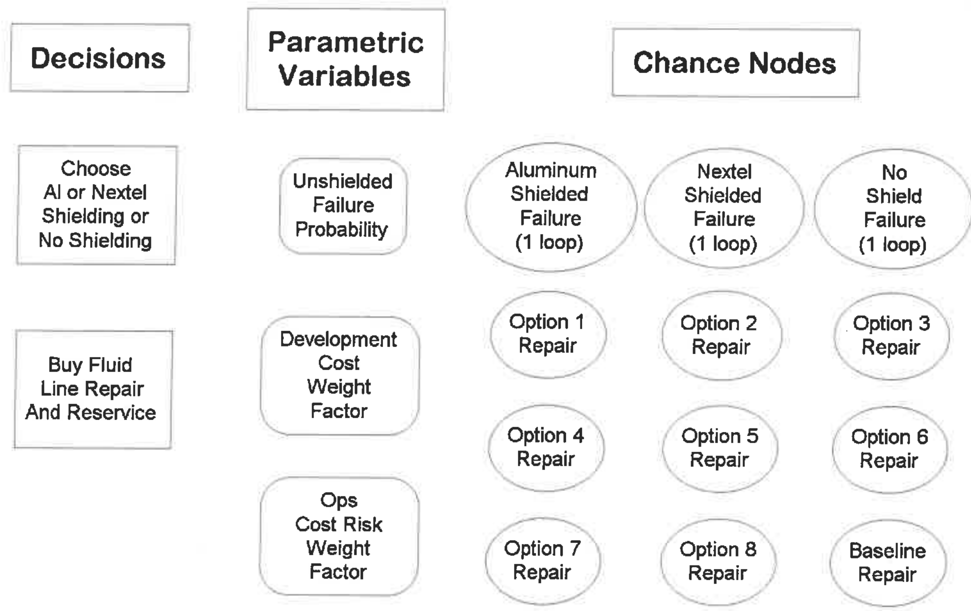

Figure 19, EEATCS Trade Study Influence Diagram

Figure 20 shows a section of the decision tree for the EEATCS trade study. The branches of the decision tree were loaded with 2-attribute GET/PAY expressions. The first attribute was the upfront development cost of each branch, the second was the associated dollarized operations-era risk for that branch. A zero for an attribute on a branch indicates that it is not applicable for that branch. As can be seen from the figure, the risk values are much greater than the cost values. The GET/PAY expressions allowed a unique set of cost and risk values to be assigned to each branch. In this trade study, the risks were dollarized to assign a dollar value to the potential losses. All attribute units are in dollars for this decision tree. The risk numbers are really ROM estimates, since it is typically impossible to assign an exact dollar value to a risk. In this particular trade study, the ultimate consequences of some branches could result in a major loss of function on the Station, requiring a major program restructuring to recover from the failures. The chance nodes in the decision tree have probabilities that were determined through a reliability and maintainability analysis.

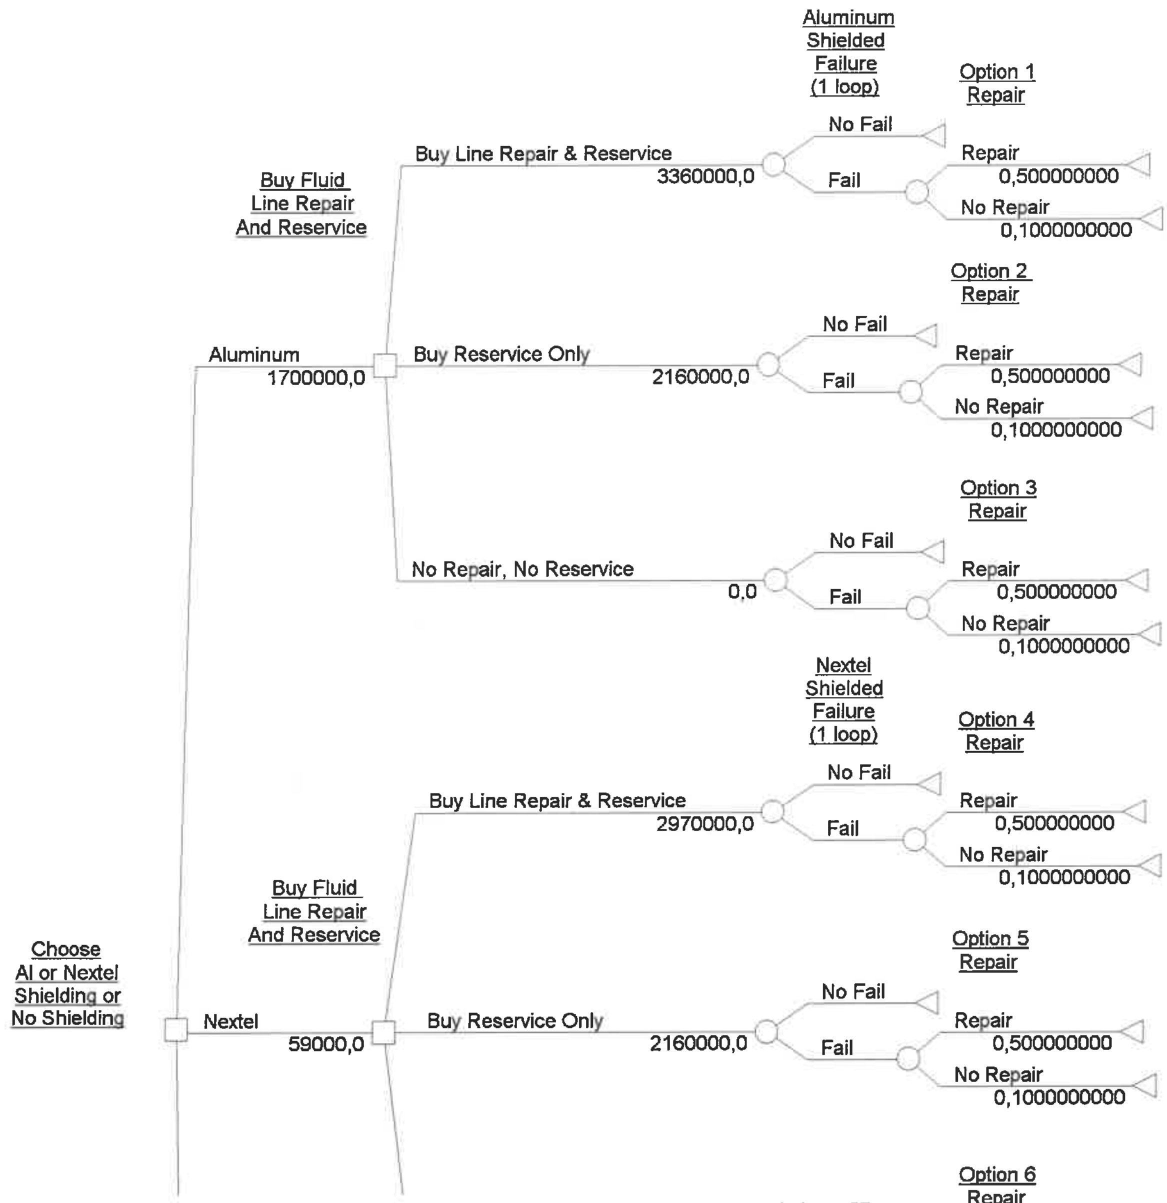

Figure 20, EEATCS Trade Study Decision Tree

For the EEATCS trade study, the balance between cost and risk drove the results of the trade study. The weighting factors for cost and risk were varied in the objective function by designating them as value nodes in the influence diagram (see Figure 19). The crux of the trade study was achieving a balance between near-term development costs and long-term operational risks. Since the total value of the two weighting factors had to add up to 1.0, an algebraic expression was used for one factor in terms of the other. A value sensitivity analysis was run on the weighting factors and is shown in Figure 21. The proper interpretation of the figure is that the decision policy does not change unless the weighting assigned to development cost is 90% or greater. The reason for this is that the downside risks in the operational era are several orders of magnitude greater than the up-front development costs. From a life cycle cost perspective, it makes much more sense to spend money up front and provide shielding, repair, and reservice capability to the EEATCS than it does to accept the risk of a major cascade failure during the operations era. The key decision to be made is whether or not to invest in additional survivability, repairability, and serviceability during the development phase or accept the risk of much greater costs and safety hazards during the operations phase.

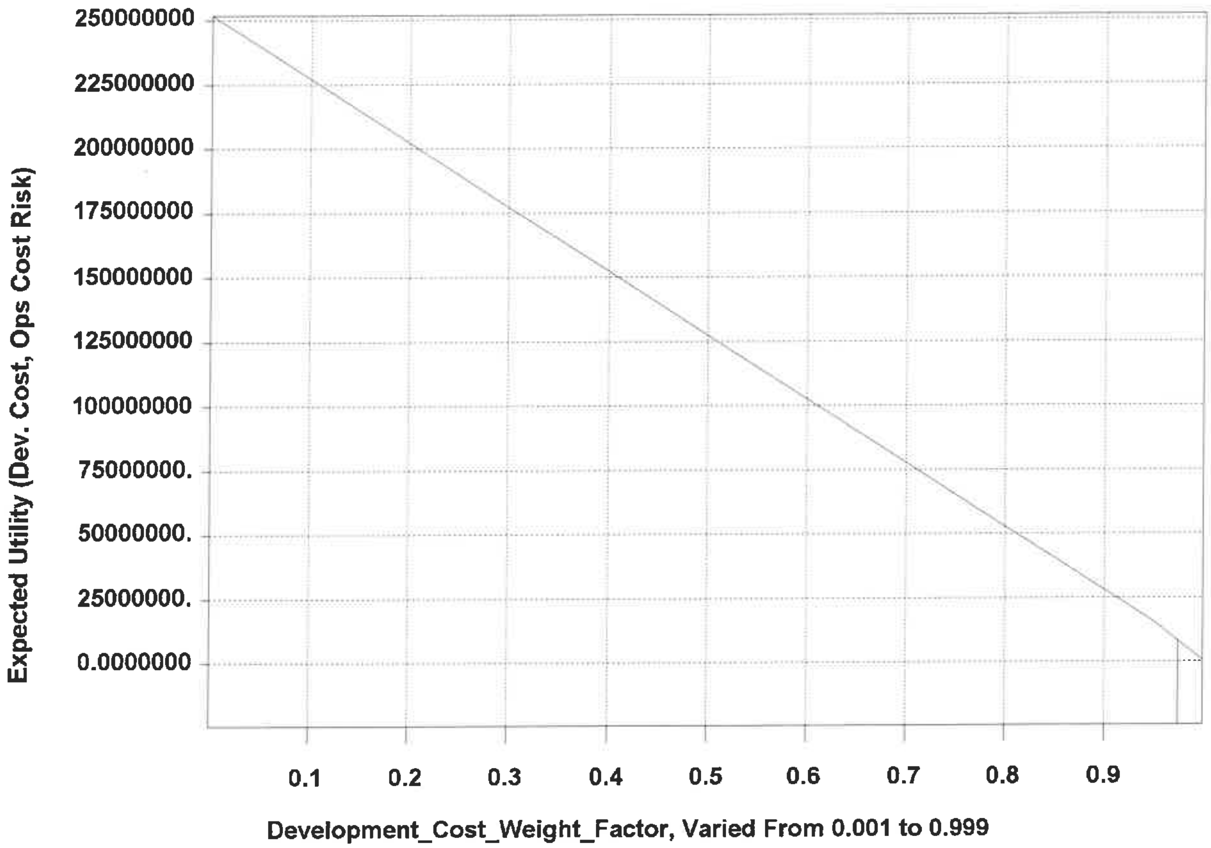

Figure 21, EEATCS Trade Study Cost vs. Risk Sensitivity Analysis

Another parametric sensitivity analysis was run for the EEATCS to see how the micrometeoroid and orbital debris (MMOD) environment affected the optimal decision policy. In this analysis, the probability that the unshielded system would spring an ammonia fluid leak was varied from 0.001 to 0.999, as shown in Figure 22 below. The major motivation for performing this analysis resulted from different MMOD environment predictions between Boeing Prime and Rocketdyne, the subcontractor producing the EEATCS. The low end of the probability scale (0.001) correlates to an orbital environment totally free of debris and micrometeoroids. The high end of the scale corresponds to an environment where the EEATCS is intensely bombarded by debris, virtually guaranteeing a leak. The real world lies somewhere between the two extremes. Results of the sensitivity analysis show that the decision policy is sensitive to the MMOD environment, with decision policy changes occurring at about 0.025 and 0.47 probability. The decision policy changes are indicated in Figure 22 by changes in shading and slope. The way this analysis relates to the trade study options is that the 0.47 leak probability corresponds to the first suboptimal policy with aluminum shielding and repair/ reservice capability. The 0.025 leak probability policy shift corresponds to the second suboptimal policy with no shielding and repair/reservice capability. The inherent reliability of the EEATCS was about 0.5 at the time the trade study was conducted, so any MMOD flux at all would push the leak probability higher without shielding. The inherent reliability is the reliability of the EEATCS due to equipment failures alone, in the absence of any MMOD punctures at all. This sensitivity analysis establishes two key insights: 1) the trade study options are highly sensitive to the predicted MMOD flux environment and 2) the inherent reliability of the EEATCS automatically limits which options are really viable.

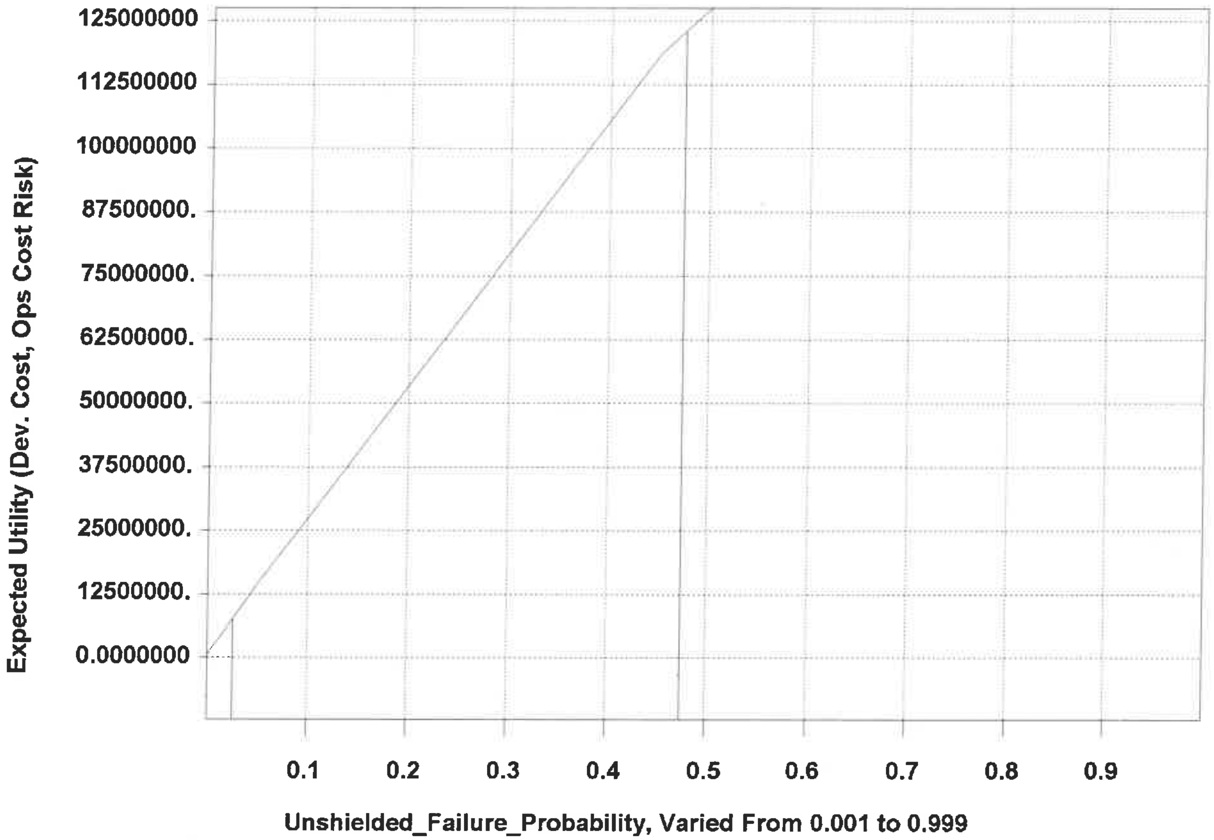

Figure 22, EEATCS Trade Study MMOD Sensitivity Analysis

To gauge the relative effects of the cost vs. risk sensitivity and the MMOD sensitivity, a tornado diagram value sensitivity comparison was performed, as shown in Figure 23. The tornado diagram shows that the relative impact of the two parametric variables is quite different. The Development_Cost_Weight_Factor variable influences the expected utility over a much wider range, from $490,941 to $251,751,000 than the Unshielded_Failure_Probability variable. It is also interesting to note that the direction of influence is reversed for the two variables, with the high and low probability ends switched on the two bars. Both bars show a policy shift to the left of the base case expected utility line, where the bar changes patterns. A shift on the right side of a bar indicates a shift on the high end of the variable's specifed range, while a shift on the left side of a bar indicates a shift on the low end. The policy shifts are placed halfway between the vertical base case line and the end of the bar, and are only used to indicate a policy shift occurs, not the exact probability it would occur at. The two previous sensitivity analyses would be used to find the precise probability where a policy shift occurs, if that were an important consideration.

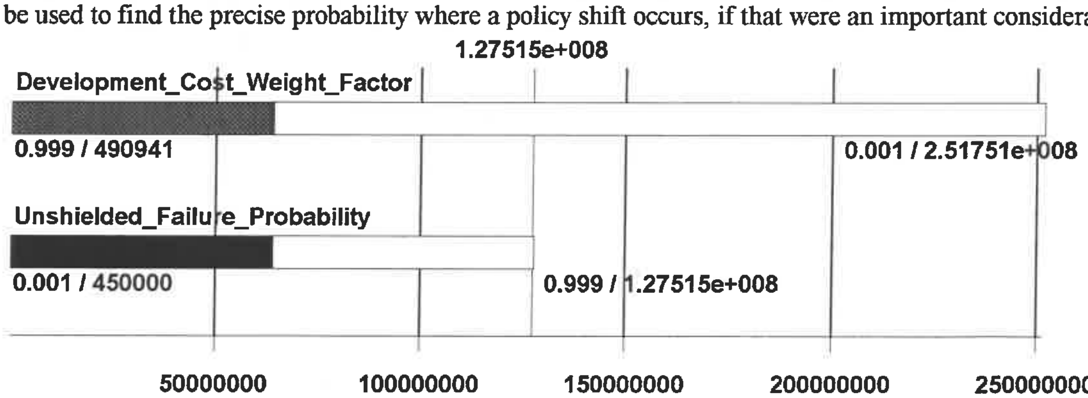

Figure 23, EEATCS Trade Study Sensitivity Comparison

Stowage Volume Shortfall Mitigation

One of the first ISS trade studies that utilized DPL involved an analysis of the options for stowing equipment and supplies onboard the Space Station. The problem was analogous to packing the trunk of a car - there are many different ways to put all the suitcases in the car, but only a few satisfy all the constraints. This trade study used DPL in a way that has not been shown previously, namely to generate the complete set of technical options based on the decisions and uncertainties involved. The influence diagram was used as a starting point, as shown in Figure 24 below. As can be seen from the diagram, there are six decision nodes, one chance node, and six value nodes. The end result of the trade is the value node in the center of the diagram "Total User Stowage at Assembly Complete", which is expressed mathematically as a summation of the values generated from all the other nodes in the diagram. The influence diagram was the key to solving this trade study, since the different decisions and constraints all had to be gathered together and their interdependencies identified. The decision to use the Centrifuge stowage volume, in the lower right of the diagram, drove the value nodes for centrifuge ambient stowage and refrigerator/freezer stowage.

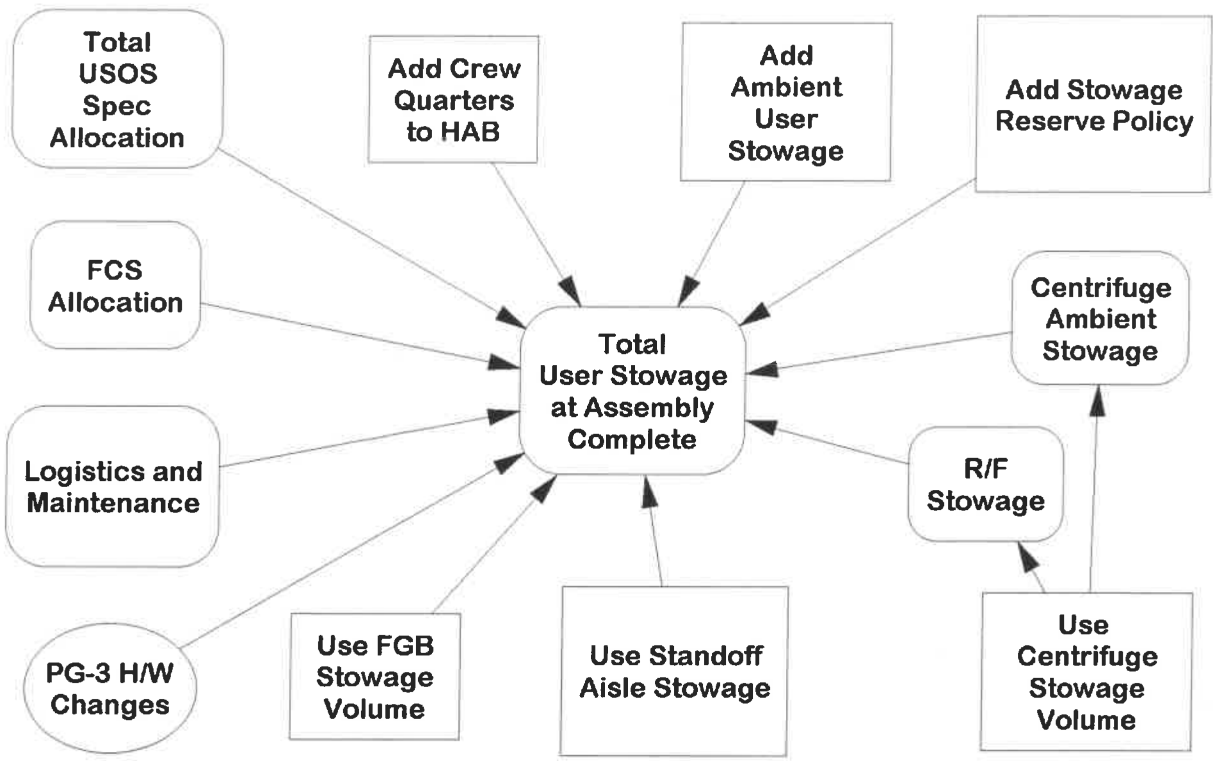

Figure 24, Stowage Volume Trade Study Influence Diagram

The major finding of the stowage volume trade study was an identification of all the possible combinations that could result from the six decisions and the chance node, and the resulting stowage volume associated with each option. Figure 25 shows the decision policy tree that had 64 possible endpoints, each of which represents a potential design option. The diagram is so large it has to be plotted to be legible, so Figure 25 is included only to show the symmetric decision tree structure that was produced. The trade study customer identified 7 options out of the 64 that were worthy of further consideration, based on the ISS program constraints and the stated desires of various stakeholders. For instance, the Astronaut Office position was that crew quarters were highly desirable, even mandatory. This constraint eliminated options that did not include crew quarters. By using DPL to generate the range of possible options, every combination of decision and chance outcomes could be quickly and easily generated without having to go through a laborious analysis process. The trade study customer was able to present the results of the analysis to management and show that he had considered all the potential options and had several viable backup options to his recommended solution.

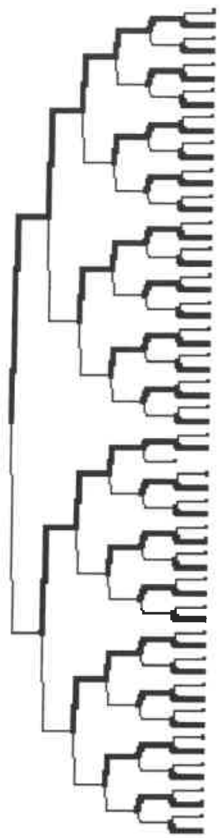

Figure 25, Stowage Volume Trade Study Decision Policy Tree

A noteworthy aspect of Figure 25 is that the decision policy tree is completely symmetrical. This is due to the fact that the tree was generated using the influence diagram. DPL will always draw a symmetrical tree from a problem entered with the influence diagram. Most decision problems are inherently asymmetric, however. In the case of a asymmetric problem, the decision tree generated by DPL can be modified or discarded and replaced with an asymmetric tree. Figures 4, 12, 13, and 18 show decision trees that were drawn manually, after the problem variables and values were entered using the influence diagram view. DPL provides the analyst with complete freedom to choose how to represent the problem and construct the model to achieve the best result. A background in decision analysis and probability theory is a basic prerequisite to good problem structuring, but once learned, the techniques and principles become easy to apply. Figure 26 below shows the decision tree that was automatically generated by DPL for the Stowage Volume Trade Study. As can be seen from the diagram, the decision tree structure is completely symmetric, since no modifications were made to the tree DPL generated.

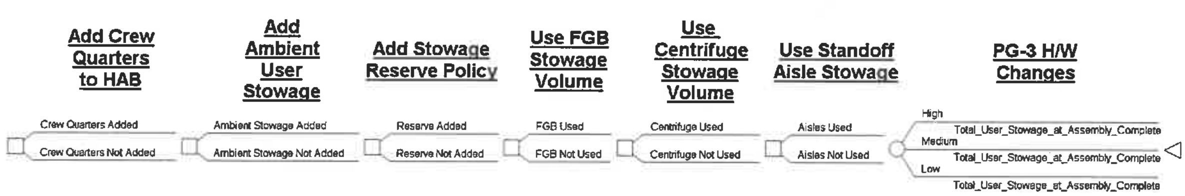

Figure 26, Stowage Volume Trade Study Decision Tree Generated By DPL

Scarring the US Habitat for a Crew Rescue Vehicle

The final example is a trade study that was really a risk analysis. The trade arose because of a desire by the United States to have a backup capability to the Russian Soyuz vehicles for assured crew return from the Space Station. An alternative vehicle to the Soyuz was the European Space Agency Crew Rescue Vehicle (CRV). The US desired the capability to dock the CRV to a port on the US side of the Station, rather than the Russian side. The US side needed some minor design changes (scars) to the US Habitat module to accommodate the CRV. Due to budgetary constraints, the decision to implement the scars needed to be delayed as long as possible. A DPL decision tree was constructed as shown in Figure 27. The tree is really a time-sequenced series of decisions to scar the Hab interleaved with chance nodes for the potential Russian backout from the ISS program. The decision nodes correspond to major programmatic milestones from the time the trade study was conducted (March, 1995) until the time the Hab is on dock at the Kennedy Space Center (November, 2001), ready to be launched into orbit on the Space Shuttle.

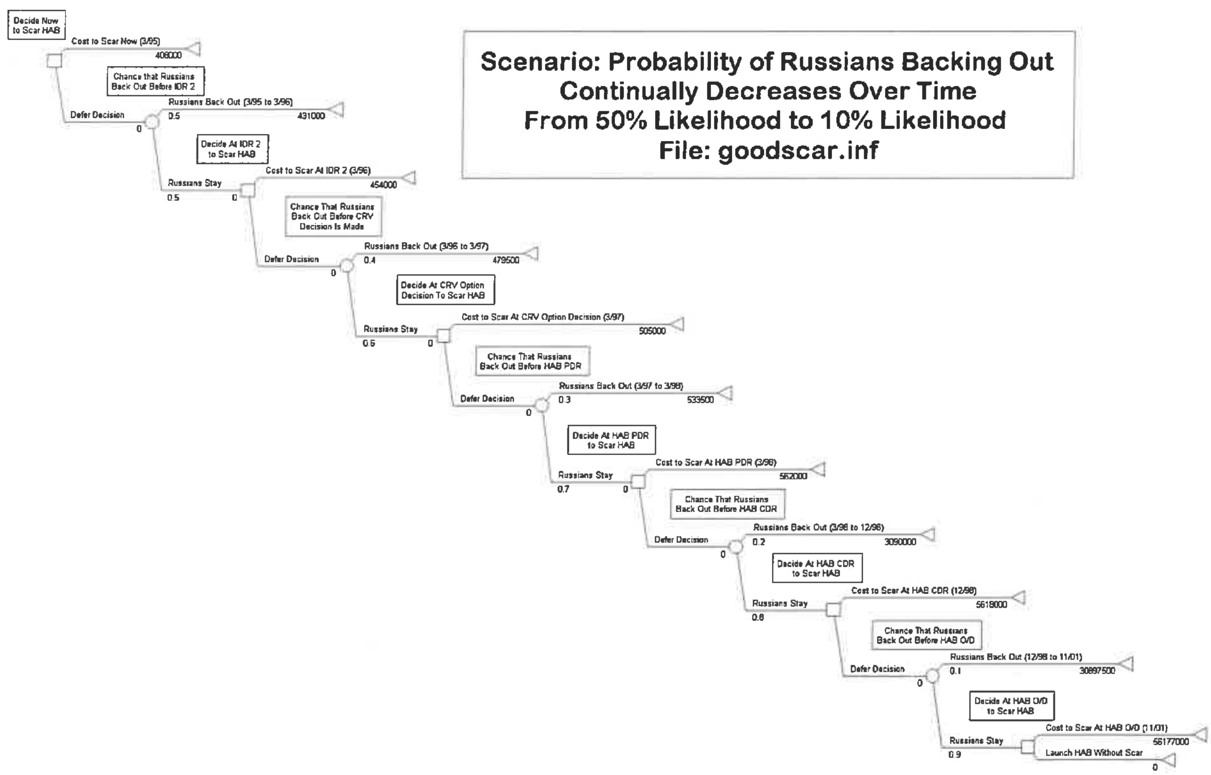

Figure 27, Hab Module CRV Scar Trade Study Decision Tree

A small section of the overall decision tree is shown below in Figure 28. The decision tree consists of a series of decision nodes and chance nodes. At each decision node, the choices are to implement the design changes or defer the decision again. If the design is changed, there is a cost impact, as shown by the GET/PAY expression on the upper branch of each decision node. Each chance node has two possible outcomes - the Russians back out of the ISS program or they stay. If the Russians back out, the design must be changed and the cost is absorbed at that time. If the Russians stay, there is no cost impact. The chance nodes have different probability numbers depending on which scenario is being examined (optimistic, pessimistic, and indeterminate). There is further discussion of the uncertainty analysis below, before Figure 30.

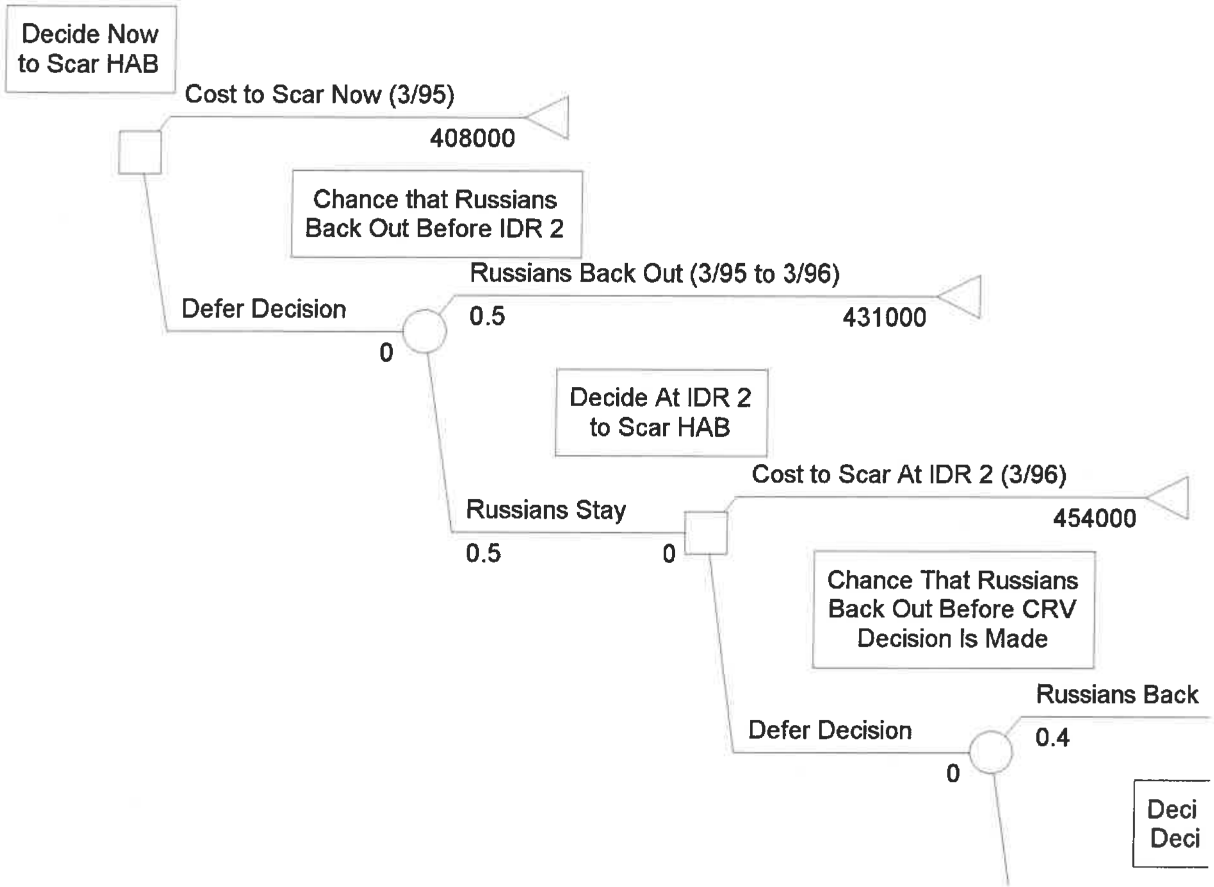

Figure 28, Hab Module CRV Scar Trade Study Decision Tree Subsection

The ISS program management could conceivably decide at any time between 1995 and 2001 to implement the Hab module scar. The problem with delaying the decision is that the cost of implemention rises as the Hab module proceeds through design, development, and test. The conclusion one can draw from the cost escalation is that the decision could be delayed to the year 1998 without experiencing major cost growth. After 1998, the Hab module enters the Preliminary Design Review (PDR) stage and the slope of the curve ticks upward. In 1999, the Hab Critical Design Review (CDR) occurs, and the slope of the curve makes a dramatic surge upward. The change in slope is due to the fact that actual hardware is being produced after PDR and CDR, and the cost to redesign existing hardware is much greater than the cost to change a paper design. Figure 29 below shows how the cost rises as a function of time. The cost increase was based on the Taguchi parametric rule of thumb that gives the cost of a engineering design change as a function of the program life cycle phasing.

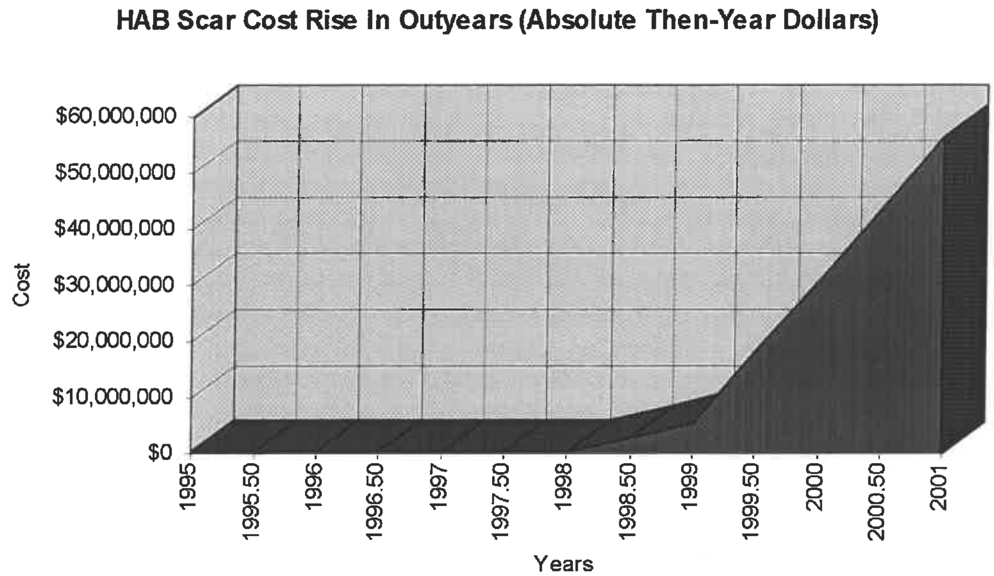

Figure 29, Hab Module CRV Scar Trade Study Cost Escalation

The cost escalation was one major influence in the trade study. The other driving factor in the analysis was the uncertainty associated with a Russian backout from the ISS program. The Russians could, for any number of reasons, drop out of the program at some time in the future. The NASA ISS program management would like to be prepared for such an eventuality. To address this issue, three versions of the decision tree in Figure 27 were created. The first version had the probability of a Russian backout continually decreasing over time, from 0.5 to 0.1. The second version had continually increasing probabilities, from 0.5 to 0.9. The third version had a constant backout probability of 0.5. The three cases represent the best case, worst case, and indeterminate case. The indeterminate case is where the probability of the Russians dropping out stays at 0.5 from 1995 to 2001 (the Russians are as likely to stay in the program as they are to drop out). In each scenario, the decision analysis was run to see where the optimal policy changed. As shown in Figure 30 below, for the best case scenario, at some time just before the Hab module CDR, it is no longer cost-effective to scar the Hab. At that point, it is more cost-effective to accept the risk of a Russian backout than to spend the additional money to scar the Hab. The worst case scenario decision policy recommended scarring the Hab right up until launch, while the indeterminate case had a later change in the decision policy. The conclusion that can be drawn from such decision policies are that it in all cases, it is most cost effective to scar the Hab module as soon as possible. The earlier the module is scarred for a CRV, the less it will cost and the less risk is involved.

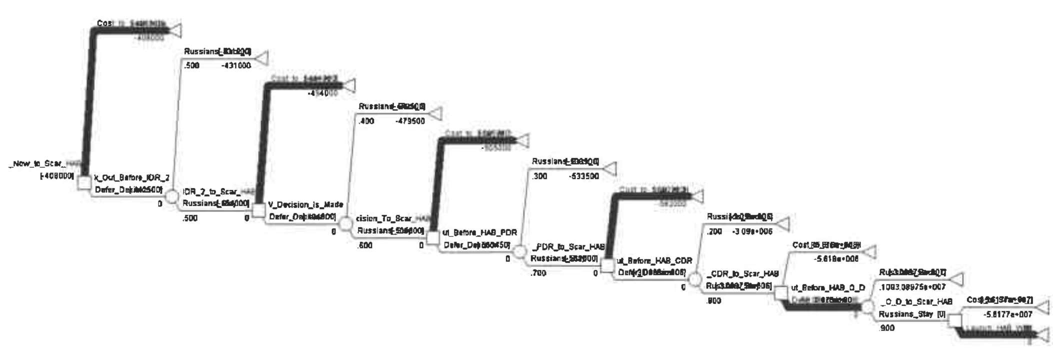

Figure 30, Hab Module CRV Scar Trade Study Decision Policy (Decreasing Backout Probability Case)

Other DPL Capabilities

DPL has other advanced features and capabilities that have not been used to date to support trade studies. DPL programs can be linked to Excel spreadsheets using dynamic data exchange, and spreadsheet models can be converted to DPL. DPL also has the capability to be used as a pure programming language, with expressions, values, constants, operators, series, arrays, and mathematical functions. The program supports many different functions and probability distributions. For decision analysis problems, DPL has capabilities that support analyzing the value of decision and control, risk attitude, constraints, and deterministic events. As time goes on, the other features of DPL will likely be used for ISS trade studies and risk analyses as the the need arises.

Conclusions

DPL is successfully being used to support complex engineering trade studies on the ISS program. The tool brings a wide range of capabilities to the Vehicle Effectiveness & Trades AIT for use in providing high quality analytical support to other teams throughout the program. DPL provides a new level of analysis capability that did not exist on the ISS program before. Through the use of DPL, the ISS program now has the tools and techniques to perform quantitative decision analyses and make consistently high quality decisions on important program issues.

Acronyms

ADA	Advanced Decision Analysis
AIT	Analysis and Integration Team
CDR	Critical Design Review
CRV	Crew Rescue Vehicle
DPL	Decision Programming Language
EEATCS	Early External Active Thermal Control System
EVA	Extravehicular Activity
FCS	Flight Crew Systems
FGB	Functional Energy Block (Russian acronym)
GPS	Global Positioning System
GN&C	Guidance, Navigation and Control
HAB	Habitat Module
IBM	International Business Machines
ISS	International Space Station
MMOD	Micrometeoroid / Orbital Debris
NASA	National Aeronautics and Space Administration
PC	Personal Computer
PDR	Preliminary Design Review
PG-3	Product Group 3
R/F	Refrigerator / Freezer
ROM	Rough Order of Magnitude
RPU	Receiver / Processor Unit
US	United States
USOS	United States On Orbit Segment
VE&T	Vehicle Effectiveness and Trades