

General Disclaimer

One or more of the Following Statements may affect this Document

- This document has been reproduced from the best copy furnished by the organizational source. It is being released in the interest of making available as much information as possible.
- This document may contain data, which exceeds the sheet parameters. It was furnished in this condition by the organizational source and is the best copy available.
- This document may contain tone-on-tone or color graphs, charts and/or pictures, which have been reproduced in black and white.
- This document is paginated as submitted by the original source.
- Portions of this document are not fully legible due to the historical nature of some of the material. However, it is the best reproduction available from the original submission.

Influence Diagram Use With Respect to Technology Planning and Investment

Daniel J. H. Levack¹

Pratt & Whitney Rocketdyne, Canoga Park, California, 91309

Bryan DeHoff²

SynGenics Corp, Delaware, Ohio, 43015

Russel E. Rhodes³

NASA, Kennedy Space Center, Florida, 32899

Influence diagrams are relatively simple, but powerful, tools for assessing the impact of choices or resource allocations on goals or requirements. They are very general and can be used on a wide range of problems. They can be used for any problem that has defined goals, a set of factors that influence the goals or the other factors, and a set of inputs.

Influence diagrams show the relationship among a set of results and the attributes that influence them and the inputs that influence the attributes. If the results are goals or requirements of a program, then the influence diagram can be used to examine how the requirements are affected by changes to technology investment.

This paper uses an example to show how to construct and interpret influence diagrams, how to assign weights to the inputs and attributes, how to assign weights to the transfer functions (influences), and how to calculate the resulting influences of the inputs on the results.

A study is also presented as an example of how using influence diagrams can help in technology planning and investment.

The Space Propulsion Synergy Team (SPST) used this technique to examine the impact of R&D spending on the Life Cycle Cost (LCC) of a space transportation system. The question addressed was the effect on the recurring and the non-recurring portions of LCC of the proportion of R&D resources spent to impact technology objectives versus the proportion spent to impact operational dependability objectives.

The goals, attributes, and the inputs were established. All of the linkages (influences) were determined. The weighting of each of the attributes and each of the linkages was determined. Finally the inputs were varied and the impacts on the LCC determined and are presented. The paper discusses how each of these was accomplished both for credibility and as an example for future studies using influence diagrams for technology planning and investment planning.

¹ Program Manager, Advanced Programs, Pratt & Whitney Rocketdyne, P.O. Box 7922, 6633 Canoga Ave, Canoga Park, CA 91309/ MS RFB19, AIAA Member.

² Principal Aerospace Engineer, SynGenics Corp, 5190 Olentangy, River Road, Delaware, OH 43015, AIAA Senior Member.

³ AST, Technical Management, Engineering Directorate Design & Development Eng Div Sys Engineering & Integration Br, Kennedy Space Center, Florida/NE-D2, AIAA Senior Member.

Nomenclature

<i>DDT&E</i>	= Design, Development, Test and Evaluation
<i>LCC</i>	= Life Cycle Cost
<i>MSFC</i>	= Marshall Space Flight Center
<i>NASA</i>	= National Aeronautics and Space Administration
<i>OPS</i>	= Operations
<i>QFD</i>	= Quality Function Deployment
<i>R&D</i>	= Research and Development
<i>SPST</i>	= Space Propulsion Synergy Team

I. Introduction

INFLUENCE diagrams are relatively simple, but powerful, tools for assessing the impact of choices or resource allocations on goals or requirements. They are very general and can be used on a wide range of problems. They can be used for any problem that has defined goals, a set of factors that influence the goals or the other factors, and a set of inputs.

Influence diagrams show the relationship among a set of results, the attributes that influence them, and the inputs that influence the attributes. If the results are goals or requirements of a program and the inputs are different ways to use R&D, then the influence diagram can be used to examine how the requirements are affected by changes to technology investment.

This paper will discuss how to construct and interpret influence diagrams, address how to assign weights to the inputs and attributes, address how to assign weights to the transfer functions (influences), and how to calculate the resulting influences of the inputs on the results.

II. Constructing an Influence Diagram

The first step in establishing an influence diagram is to define the goals. The second step is to define all of the attributes that have any impact on the goals. The next step after that is to establish any attributes that affect the attributes established in step two. Continue this process until the attributes are actually inputs – items that the user has the ability to control.

In other words, define what it is that needs to be achieved (the “goal” or “goals”), then determine all the processes that have an impact on the goal or goals (the “attributes”), and lastly determine what inputs are available to affect the processes. Once all of these have been defined, arrange all of the processes in a block flow diagram showing the interactions between the different processes, and the methods by which the inputs affect specific processes.

Influence diagrams generally flow from right to left and no loops are allowed in the diagrams.

There are three types of boxes in an influence diagram: inputs, attributes, and goals. Goals are boxes that only have influences going into them. They have no outputs – they are the end results of the influence diagram. Attributes are boxes that have influences coming into them and influences leaving from them, but they have no independent value, they are simply processes that act as transfer boxes and switches. Most inputs are boxes that have influences leaving from them but have no influences coming into them – they are normally starting points for a chain of influences. They have independent values that normally represent some parameter that can be controlled to ultimately influence some goal. Although an input does not have to have influences coming into it, it may also be affected by other inputs or other attributes. Consequently, an input is allowed to have influences going into it. The key difference between an input and an attribute is that the input has an independent value.

Figure 1 shows examples of goals, attributes, and inputs. Box “A” is a classic example of an input – it has an independent value (shown in blue above the box), it has two influences leaving from it, and it has no influences going into it. Box “B” is also an input since it has an independent value (also shown in blue above the box). But it has influences coming into it and leaving from it. It is an input and not an attribute because it has an independent, controllable, value. Boxes “C” and “D” are attributes – they have influences coming into them and influences leaving from them, they transfer and direct influences, but they have no independent value. Box “E” is a goal – it has influences coming into it, it has no influences leaving from it, and it has no independent value.




Figure 1. Example of Boxes Defined as Inputs, Attributes, and Goals.

This paper presents an example of an influence diagram. The purpose of this example is to have a case complex enough to show all the types of interactions that would be present in any case of any complexity. But at the same time, to have a case simple enough that all the numbers and interactions can be easily seen and followed.

III. Detailed Construction of Example Influence Diagram

The example that has been chosen is determining what influences a diner's satisfaction in a restaurant. Not every single possible attribute will be examined, but enough attributes will be examined that it will be a rational case that makes sense.

The first job is to determine the goal. The goal in this case is a diner's satisfaction. The second job is to establish what attributes directly impact the goal of a diner's satisfaction. For this example only three attributes have been chosen. They are: food presentation, food taste, and food price. Other attributes such as quality of service could also have been included but would only complicate the example without adding new information about influence diagram construction. Figure 2 shows the influence diagram as defined up to this point.




Figure 2. Definition of Goal and Directly Affecting Attributes.

The food presentation, food taste, and food price are not characteristics that can be directly controlled, i.e. they are not inputs. The things that can be done to directly impact these attributes are the method by which the food is prepared, i.e. the food preparation, the recipe used, the quality of the ingredients used, and the ability of the cook. Figure 3 shows the influence diagram with these items added. These items are all things that can be directly controlled, and thus they are the inputs for this example. In other words, how resources are spread across these inputs will determine the value of the attributes and, consequently, the diner's satisfaction.

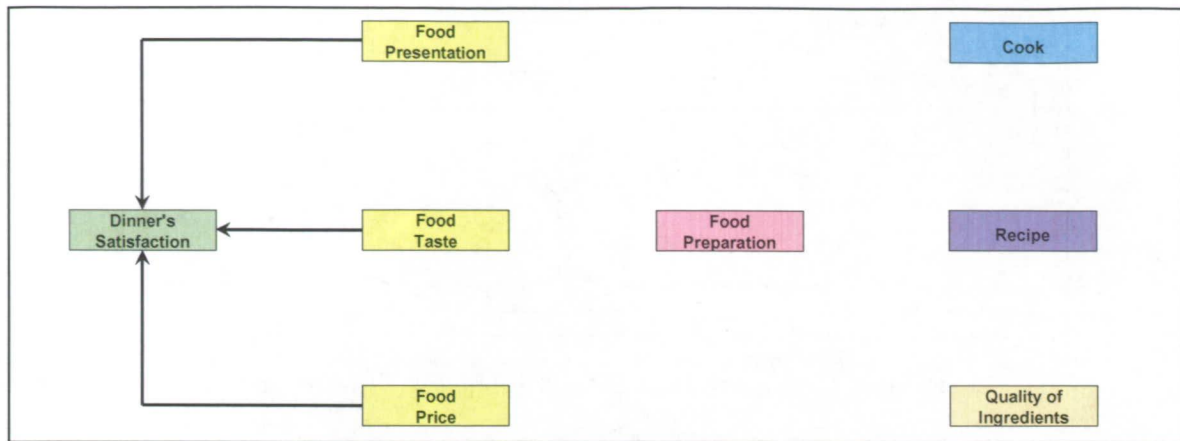


Figure 3. Definition of Goal, Attributes, and Inputs.

At this point the linkages between the attributes and the goal have been established, and the inputs and attributes themselves have been established. But the linkages and interactions between the inputs and the attributes have not yet been established. The food preparation input is actually affected by the other inputs. It is affected by the cook, it is affected by the recipe used, and it is affected by the quality of the ingredients. The other three inputs do not directly influence each other. The cook has a direct influence on the food presentation attribute and, through the cost of the cook, on the food price. The cook also has an influence on the food preparation. The recipe has a direct influence on the food preparation and on the food presentation. The quality of ingredients has a direct influence on the food preparation, it has a direct influence on the food taste, and certainly an influence on the food price. The food preparation input directly affects the food price, and directly affects both the food presentation and the food taste. Assuming these relationships are valid, then the influence diagram now looks like Figure 4.

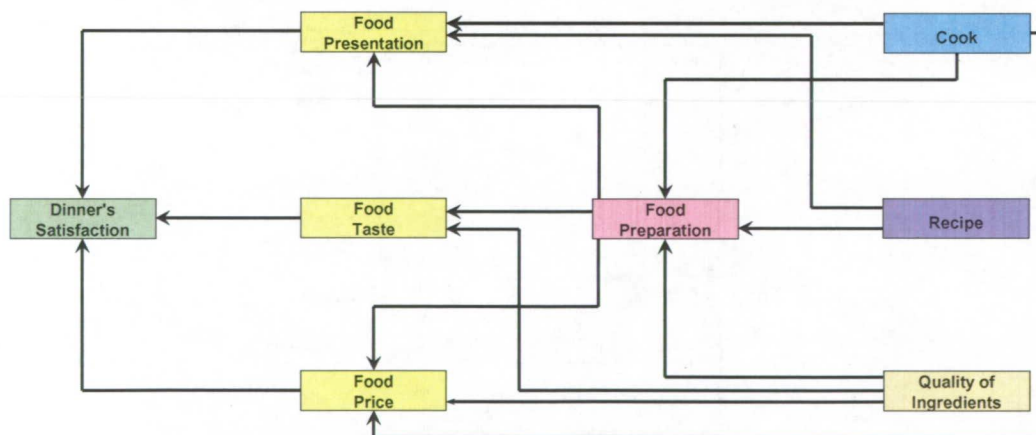


Figure 4. Influence Diagram with All Influences Shown.

The influence diagram shown in Figure 4 now defines the relationships between the inputs, the attributes, and the goal, but it is not yet useful for numerical analysis of the impact of the inputs. The next step, and perhaps one of the most important steps, is to find the strength of the relationships between each input, each attribute, and the goal. Another way of looking at the strength of these relationships is that they are influence coefficients or transfer functions. The method used to determine the strength of the relationships among the attributes and the inputs is often different from the method used to determine the strength of the relationships between the final attributes and the goal.

First consider the relationships between the final attributes and the goal. In this example the food presentation might have a high value, perhaps 100%, at a high-end restaurant, while the food price has relatively low value, perhaps 50%. On the other hand, at a taco stand the food presentation would be expected to have a low value,

perhaps 10%, and the food price a very high value, perhaps 100%. The point is that the value for the strength of the relationships depends upon the case that is being examined even when the relationships are the same. For example, in an aerospace case, the relative strengths of the relationships between cost and performance would differ between a nationally critical military mission and a commercial airliner.

The strength of the relationships between the final attributes and the goal should be determined either from historical knowledge of the customer or of the political and practical constraints the customer must work under, or from interviewing the customer directly.

For this diner's satisfaction example, the strengths of the relationships between the final attributes and the goal have been assumed to be equal and set at 100%.

Next consider the relationships among the attributes and the inputs. The strengths of these relationships are probably best established from historical data or using estimates from people who are very familiar with the processes that the attributes, and their interactions with the inputs, represent. The strength of the relationships used for the diner's satisfaction example will be shown in a subsequent figure.

Influence strengths can be positive or negative and their strengths can range from zero percent (no influence) to well above 100 percent.

The correct establishment of the links and the strength of the relationships is critical to the utility of influence diagrams.

Once the influence diagram and the links have been established, the strengths of the relationships can be varied to determine the sensitivity of the goal to the value of each of the relationship strengths. Because it can require significant resources to establish a value for each of the relationship strengths, it is valuable to exercise the influence diagram with variable relationship strengths before expending the effort to determine the actual strengths. Then effort can best be spent only on those relationships where the goal is particularly sensitive to the actual value of the relationship strength.

The last step is to assign values to the inputs. A common practice is to assign all the inputs as a proportion of the available resources, perhaps adding up to 100%, or as fractions adding up to one. An alternate method is to assign the inputs as dollars. Somewhat arbitrarily assigning the values of 30 to the cook, 20 to the recipe, 25 to the quality of ingredients, and 25 to the food preparation, the final version of the example influence diagram results. It is shown in Figure 5.

Examination of Figure 5 shows how an influence diagram really works. The figure is embellished and color-coded to show all of the flows of the inputs through all of the attributes to the goal. The three pure inputs and the one combination input/attribute have been color-coded and then every arrow into every attribute has been given four boxes that correspond to the colors of the inputs. Thus the contribution of each input along each path of influence can be easily seen. The value of all influences from all sources is shown in black just above the left top of each box. The value that starts from each input is shown in red just above the right top of each input box. The strength of each influence path is shown in red as a percentage above each black arrow path.

First look at the combination input/attribute "Food Preparation". The first influence flowing into it is from "Cook". This influence has a value of 30 (the input value of 30 from "Cook" multiplied by the influence factor of 100%). The second influence flowing into "Food Preparation" is from "Recipe". This influence has a value of 20 (the input value of 20 from "Recipe" multiplied by the influence factor of 100%). The last influence flowing into "Food Preparation" is from "Quality of Ingredients". This influence has a value of 15 (the input value 25 multiplied by the influence factor of 60%). Since "Food Preparation" is itself an input with a value of 25, the value of the box is 90 ($90 = 30$ from "Cook" + 20 from "Recipe" + 15 from "Quality of Ingredients" + 25 from itself). All the influences emanating from "Food Preparation" will have this value of 90 multiplied by the influence factor of the specific out-going influence flow path.

Now look at the "Food Presentation" attribute. It has three flows of influence into it: two from inputs ("Cook" and "Recipe") and one from the combination input/attribute "Food Preparation". The first influence flowing into it is from "Cook". This influence has a value of 30 (the input value of 30 from "Cook" multiplied by the influence factor of 100%). The second influence flowing into "Food Presentation" is from "Recipe". This influence has a value of 10 (the input value of 20 from "Recipe" multiplied by the influence factor of 50%). The last influence is more interesting. The influence on "Food Presentation" from "Food Preparation" has a value of 39.6 (90, the value of the box "Food Preparation", multiplied by the influence factor of 44%). But it can be looked at as having the value of 13.2 from "Cook", 8.8 from "Recipe", 6.6 from "Quality of ingredients", and 11 from the input "Food Preparation" (not from the overall box "Food Preparation"). This second way of looking at the influence flows shows the impact of each of the inputs on every attribute and goal, and is what is highlighted in the colored boxes on each influence flow path in Figure 5.

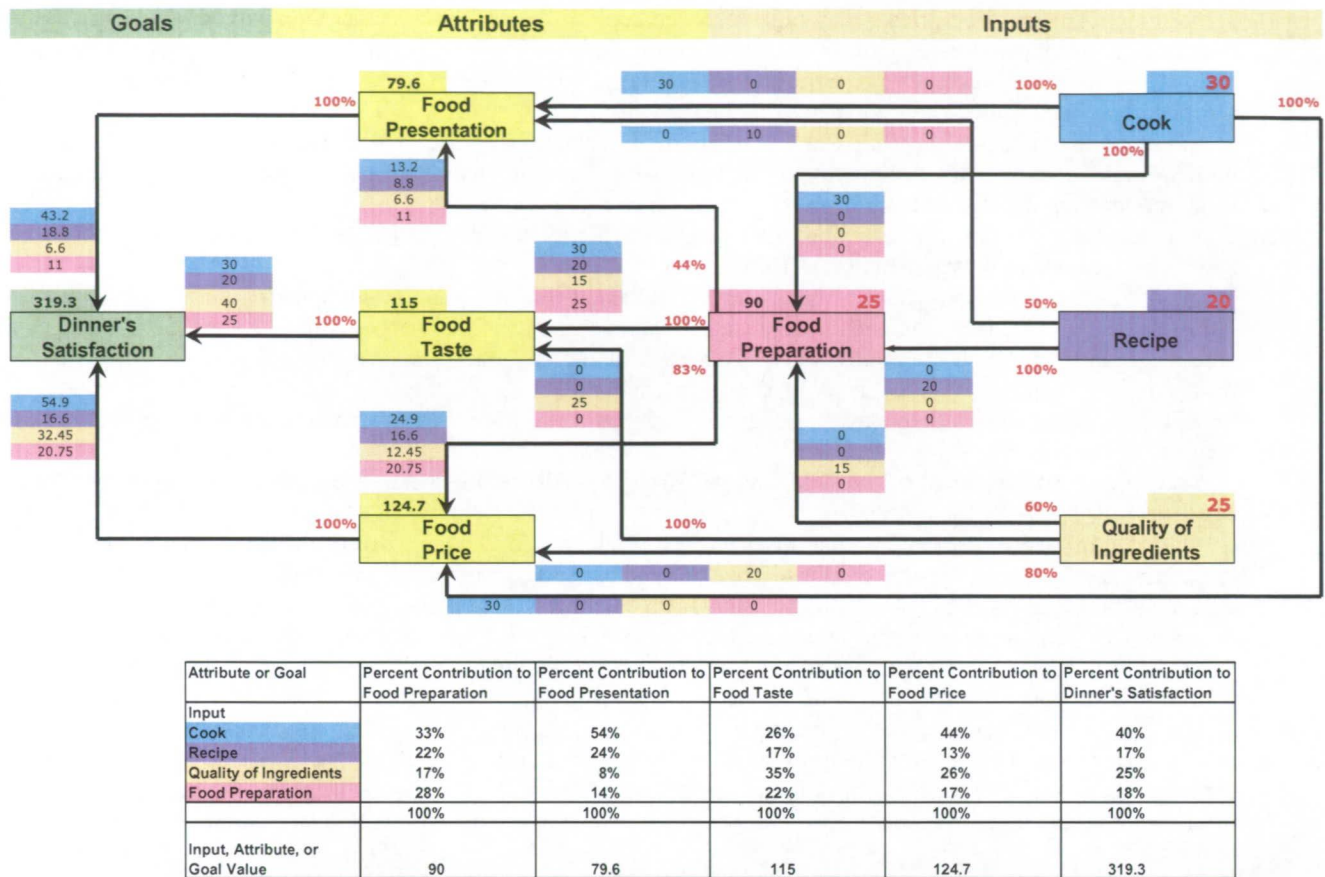


Figure 5. Fully Defined Diner's Satisfaction Influence Diagram.

The table included at the bottom of Figure 5 shows the percentage contribution of each of the inputs to each of the attributes and to the goal. As an example look at the last column on the right of the table, the percent contribution to the diner's satisfaction due to the cook. From Figure 5 the value of the box "Diner's Satisfaction" is 319.3 (the value of all the influences going into the box). The influence from "Cook" is shown in the blue boxes along each of the three inputs to the "Diner's Satisfaction" box. They are: $43.2 + 30 + 54.9 = 128.1$. So the contribution due to "Cook" is $128.1/319.3 = 0.401$ or 40%. The other influences from the other inputs are calculated similarly.

The calculation of the influences of the inputs on intermediate attributes is similar. For example, the influence of "Cook" on "Food Presentation" is $30 + 0 + 13.2 = 43.2$. The overall value of "Food Presentation" is 79.6. So the contribution of "Cook" to "Food Presentation" is $43.2/79.6 = 0.543$ or 54%.

The table could have shown absolute numerical values instead of percentages, but it is often easier to perceive the relative effects if the values are converted into percentages, i.e., are normalized.

Once an influence diagram such as that shown in Figure 5 has been prepared, there are two common ways to use it. The most common way is to vary the inputs (generally the proportions of available resources) to determine what split of inputs maximizes the numerical value of the goal. This is often combined with a set of minimum proportions for each of the inputs. The second common way to use an influence diagram is to vary the influence factors to represent different scenarios. For example: how would the optimum resource allocation change if "Food Price" were much more important than "Food Presentation", or if "Food Presentation" were much more important than "Food Price"?

Understanding the elements that were used in this example will allow the construction of very complex influence diagrams. They simply have many more attributes and many more influences (transfer functions). Unfortunately, they can then become very hard to follow throughout the intermediate attributes. Consequently, it is important to understand the previous example. As long as the complex influence diagram is constructed consistent with the rules shown, then the results should be creditable. It is important to remember that those rules include not only the

mechanical rules of diagram construction, but especially the rules of getting good information about the influence paths and the strengths of the influence factors.

IV. Aerospace Influence Diagram Example

The Space Propulsion Synergy Team (an industry/government/academia group established in 1991) was tasked by NASA Marshall Space Flight Center (MSFC) to assess the impact of R&D expenditures on the life cycle cost of a space transportation system. The system to be assessed was the SpaceLiner 100. The system was looked at in two different ways: first, as a R&D program producing a single prototype, and second, as an R&D program followed by the production of a commercial fleet. The first case is representative of any case where the final fielded fleet would consist of only one or two vehicles. The question was what proportion of the R&D should be spent on efforts to improve recurring costs versus on efforts to improve non-recurring costs. An influence diagram was the tool chosen to address the question.

The work started from a determination of the primary cost drivers for a space transportation system. The SPST had conducted an exhaustive QFD study in the late 1990s to produce this set of primary cost drivers. The results were compiled in a design guide⁽¹⁾. Figure 6 shows the cost drivers – “attributes” – determined along with their weighting for a reusable space transportation system.

SPST ATTRIBUTES DEVELOPMENT WITH WEIGHTS

DATA REF: SL 100 Design Criteria Matrix (1-27-99).xls Space Des Crit (ETO Reusable)	Column A "Customer Rank" (J.Mankins & U.Hueter)	Column B Technology Improvement "Now"	Column C Technology Improvement "Plan"	Column D "Improvement Ratio" (C/B)	Column E "Sales Points"	Column F Score (A*D*E)	Weighted Score F/Sum(F) (Customer) %
Attributes							
Min. Cost Impact of Payloads on Launch Sys.	4	2	4	2.00	1	8.00	2.431
Low Recurring Cost							
Low Cost Sens. to Flt. Growth*	4	3	4	1.33	1	5.33	1.621
Operation and Support	5	1	5	5.00	1	25.00	7.597
Initial Acquisition	0	1	3	3.00	1	0.00	0.000
Vehicle/System Replacement	3	1	3	3.00	1	9.00	2.735
Highly Reliable	5	2	5	2.50	1	12.50	3.798
Intact Vehicle Recovery	5	3	5	1.67	1	8.33	2.532
Mission Success	3	4	3	0.75	1	2.25	0.684
Operate on Command	5	1	5	5.00	1	25.00	7.597
Robustness	5	2	5	2.50	1	12.50	3.798
Design Certainty	5	2	5	2.50	1	12.50	3.798
Flexible	2	1	2	2.00	1	4.00	1.215
Capacity	4	4	4	1.00	1	4.00	1.215
Operable							
Process Verification	5	3	5	1.67	1	8.33	2.532
Auto. Sys. Health Verification	5	1	5	5.00	1	25.00	7.597
Auto. Sys. Corrective Action	5	1	5	5.00	1	25.00	7.597
Ease of Vehicle/System Integration	2	1	2	2.00	1	4.00	1.215
Maintainable	4	1	4	4.00	1	16.00	4.862
Simple	5	1	5	5.00	1	25.00	7.597
Launch on Demand	2	1	2	2.00	1	4.00	1.215
Easily Supportable	5	1	5	5.00	1	25.00	7.597
Resiliency	3	1	3	3.00	1	9.00	2.735
Vehicle Safety	5	3	5	1.67	1	8.33	2.532
Personnel Safety	5	3	5	1.67	1	8.33	2.532
Public Safety	5	3	5	1.67	1	8.33	2.532
Equipment and Facility Safety	5	3	5	1.67	1	8.33	2.532
Minimum Impact on Space Environ.	4	2	4	2.00	1	8.00	2.431
Minimum Effect on Atmosphere	3	1	3	3.00	1	9.00	2.735
Minimum Environ. Impact all Sites	3	1	3	3.00	1	9.00	2.735
Benefit GNP	0	1	4	4.00	1	0.00	0.000
Social Perception	0	2	4	2.00	1	0.00	0.000
TOTALS=						329.08	100.000

Figure 6. Cost Driving Attributes for Reusable Space Transportation System.

The attributes were then arranged into an influence diagram with the influences determined by an interaction with the SpaceLiner 100 customer and refinement by the SPST. The influence diagram was based on an “algorithm” developed by Garry Lyles at MSFC for a “systems approach to safety, reliability and cost”. The algorithm was expanded into a flow diagram of the various attributes, and the flow of their influences, using the attributes and weightings from Figure 6. The resulting influence diagram is shown in Figure 7.

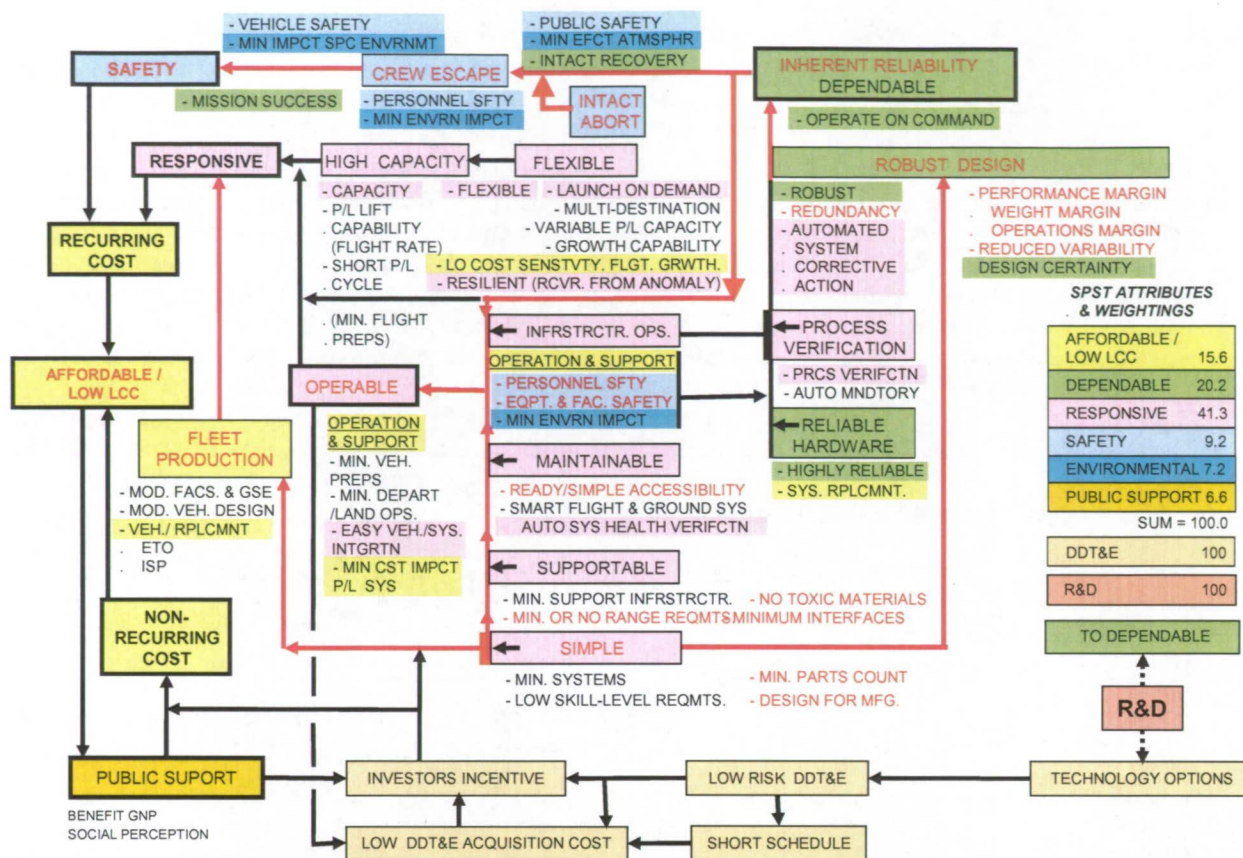


Figure 7. Detailed Influence Diagram for R&D Life Cycle Cost Goal.

It was noted that attributes relating to “dependable” flowed into other attributes and a simplified structure was developed to reflect this fact. Figure 8 shows the relationships, their weights, and, in the lower left corner of the figure, the resulting simplified diagram.

The simplifications that were shown in Figure 8 were incorporated into the influence diagram of Figure 7. The resulting influence diagram is shown in Figure 9. In the figure, the box “Fleet Production” has a dashed outline and the influences into and out of it are also dashed. This is being used as a convention to mean that when the single vehicle case is examined, then the influence into “Fleet Production” is zero percent and no influence is transferred through it to “Responsive”. However, when the commercial fleet case is examined, then the influence into “Fleet Production” is turned on, i.e., the influence into and out of “Fleet Production” is 100 percent. An Excel spreadsheet was developed to model the influence diagram of Figure 9.

Figures 10 and 11 show the results of exercising the influence diagram both for the single vehicle (prototype) case and the case where a commercial fleet is procured.

Figure 10 shows the results for the single vehicle or prototype case. The legend at the bottom of the chart shows the percentage of R&D resources used to address technology objectives (the “X percent” in Figure 9) over the percentage of the R&D resources used to address the operational dependability objectives (the “(100-X) percent” in Figure 9). The direction of “goodness” in these two figures is up. In other words, the cost, be it life cycle cost or individual attribute cost, goes down if the value in the figure goes up.

As the resource spread goes from 10 percent spent on technology options (the left side of the figure) to 90 percent spent on technology options (the right side of the figure), the overall life cycle cost improves, but the impact is small with only about a ten percent total change. The non-recurring costs are strongly impacted with about a 70 percent improvement (the schedule and testing is much shorter). But the recurring costs for operating the prototype is much worse at about two-to-one.

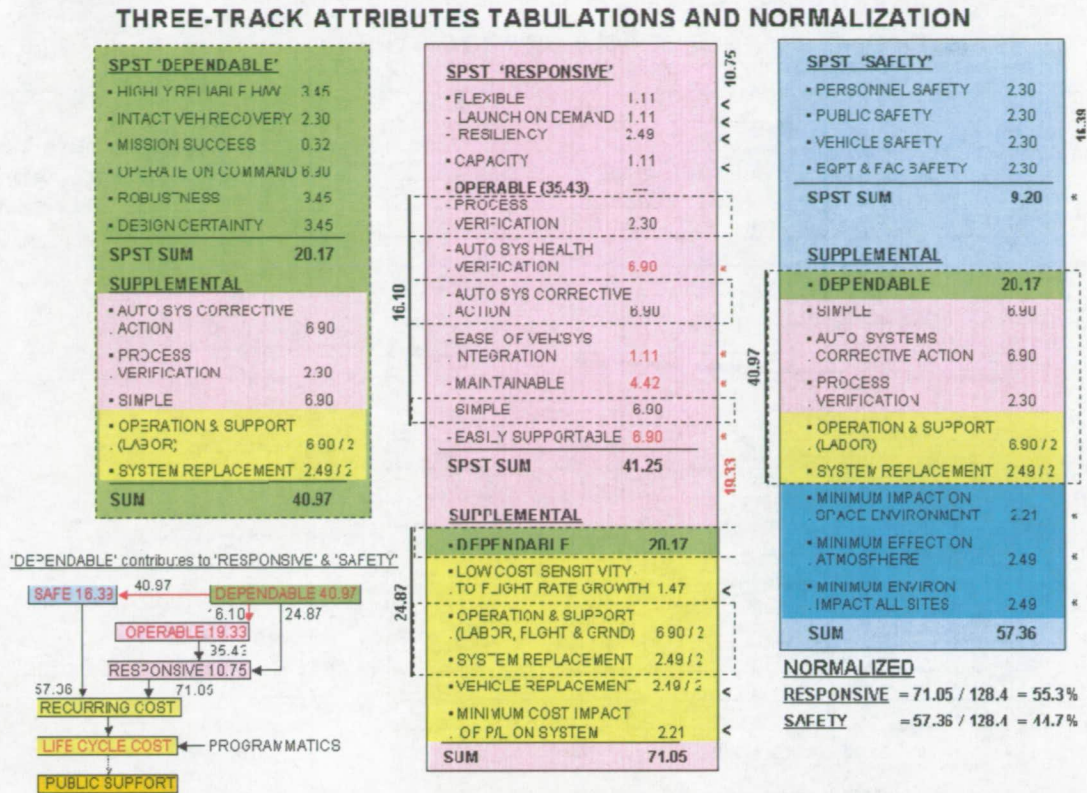


Figure 8. Combining Attributes to Produce Equivalent Attributes and Influences.

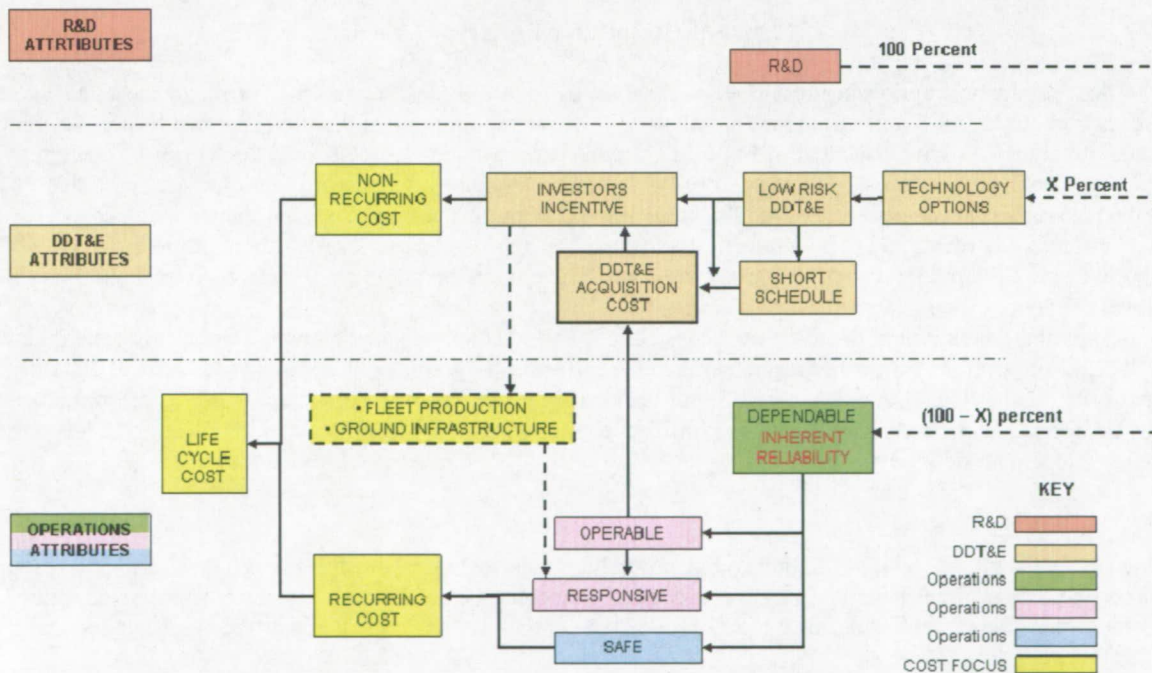


Figure 9. Final Version of Reusable Space Transportation System Influence Diagram.




Figure 10. Results for Prototype Case.

If the fleet production option in the influence diagram of Figure 9 is exercised (i.e., if the dashed line into the fleet production box is solid and influences are allowed to flow into and out of the box), the results are different in magnitude but similar in kind, as shown in Figure 11. In this case, as the resource spread goes from 10 percent spent on technology options (the left side of the figure) to 90 percent spent on technology options (the right side of the figure) the non-recurring costs are improved the same amount as in the previous case, but the recurring costs are also improved instead of harmed, and the overall life cycle costs are improved significantly (about 30 percent).

In both cases the trade favors R&D resources being spent on technology objectives over being spent on operational dependability objectives.

The design and development of this influence diagram, an Excel spreadsheet embodying the influence diagram, the results, and suggestions for further refinements of the influence diagram were presented to NASA in 2001⁽²⁾.

The results for the SpaceLiner 100 case are not necessarily general, but the objective is not to determine for all time how to spend R&D, but to show that using influence diagrams can be useful for making informed decisions for specific programs with defined goals.

V. Conclusions

Influence diagrams are relatively simple, but powerful, tools for assessing the impact of choices or resource allocations on goals or requirements. They are very general and can be used on a wide range of problems. They can be used for any problem that has defined goals, a set of factors that influence the goals or the other factors, and a set of inputs.




Figure 11. Results for Commercial Fleet Case.

As was shown in the “Diner’s Satisfaction” example, influence diagrams can not only be used to find the impact of changes in inputs on the value of the goals, but can also be used to find the individual impact of each input on each of the intermediate attributes. If intermediate attributes have constraints, this allows a means to ensure they are satisfied. It also allows the identification of attributes that, although thought to be important, actually have little effect.

The life cycle cost for a space transportation system example showed that influence diagrams can be used for resource planning for complex, interrelated, systems. However, that example also showed that influence diagrams are case specific. They must be tailored for the specific goals, and, especially, the influence factors that apply to a given customer or organization.

There are three keys to the successful use of influence diagrams. The first key is to define the goals, all of the attributes that can affect the goals and the other attributes, and the inputs. The second key is to define the influence paths. And the last key is to define the strengths of the influence (transfer functions).

References

- ¹ A Guide for the Design of Highly Reusable Space Transportation – Final Report, Space Propulsion Synergy Team (SPST), August 29, 1997.
- ² Rhodes, R. E., et al., SpaceLiner 100 Functional Requirements Subteam, “Assessing R&D and DDT&E Influence on Life Cycle Cost,” Briefing Supporting 3rd Generation RLV/SpaceLiner 100 Functional Requirements, Space Propulsion Synergy Team, 18 October 2001, KSC Archives, Loc. 37E-7/Folder #1 (unpublished).