The Ideal Engineering Decision Support System

Dr. David G. Ullman
Robust Decisions Inc.
ullman@robustdecisions.com

Abstract

This paper has four goals: The first is to develop the importance of decision thinking. Decision thinking suggests that looking at design or business problem-solving as decision-making is a useful approach to understanding problems and generating a good solution to them. The second goal is to present support for this position based on literature from psychology, sociology, decision theoretics, business and engineering. Third, an ideal system to support engineering design decisions is defined. This ideal system is described by sixteen measures. These are based on the literature and on the study of existing methods. Forth and finally, five existing decision support methods are compared to the ideal. These include IBIS, MAUT, AHP, Pugh's decision matrix and *Accord*. *Accord* is a new system developed by the author and his colleagues with the ideal in mind.

1. INTRODUCTION

What were you taught about decision-making? Probably what you learned was that if you got an answer that matched the one in the book, you got the right answer. Matching the right answer makes decision-making easy. Yet, when you went to buy a car, to design a product, or to make a business decision, there was no right answer. In fact, most issues have multiple, alternative, satisfactory solutions. These options evolve as you work our way through the problem. The criteria used to evaluate these potential solutions also evolve as more is learned about the issue. The bigger the problems and the more people involved in making the decision, the harder it is to manage the alternatives, criteria and evaluation. If you think of the alternatives, and the criteria you use to judge the alternatives, such as information, then in general:

> Problem solving is generating and refining information punctuated by decision-making.

Observe the dynamics in a meeting and see if you can find any methods or tools that help the team manage the information. You probably won't see any. You may see people trying to develop and share information that supports their favorite alternative. However, these activities are, for the most part, ad-hoc. They offer little help to insure the knowledge and abilities of the participants is well utilized. In fact, this ad-hoc nature results in unresolved issues, dissatisfaction with the results and a partial explanation of why people dislike meetings so much; especially those meetings that are poorly structured and not clearly focused.

It is important to put heavy emphasis on decision-making because:

A decision is a commitment to use resources.

Whenever a decision is made, one of the proposed alternatives is chosen. Future activity focused on the chosen alternative uses time, money and other resources, and excludes any effort on the alternatives rejected. Thus, if a poor choice was made and it is later decided to revise the decision, all the intervening time is lost and expenditures are, for the most part, unrecoverable. Thus, all decisions commit the decision-makers and others to further effort. In fact, part of decision-making is determining how much commitment each alternative will require to bring it to fruition.

To help develop systems that support decision making and are useful in daily work this paper has four goals: First is to develop the importance of decision thinking. The second goal is to present support for this position based on literature from psychology, sociology, decision theoretics, business and engineering. Third, an ideal system to support engineering design decisions is defined. Forth and finally, five existing decision support methods are compared to the ideal. These include IBIS, MAUT, AHP, Pugh's decision matrix and *Accord*. *Accord* is a new system developed by the author and his colleagues with the ideal in mind.

2. THE IMPORTANCE OF DECISIONS

2.1 The Decision Pyramid

Not all types of information are of equal value in solving a problem. In Figure1, four classes of information are shown. The most basic form of information is raw data. Raw data is numbers, textual clauses or other descriptive information about some object or idea. Models are a form of information that represent the relationships between data. These relationships may be mental pictures of a situation, math equations, full sentences or paragraphs, or graphic images that relate basic data resulting in a richer form. These models are static relationships among the data and are key to evaluation. During evaluation, if you find that an alternative does not meet a criterion, this discovery itself does not tell you how to change the alternative to better fit the need. To gain knowledge necessary to make an informed change, the behavior of models must be understood and interpreted. It is the knowledge gained during evaluation that we depend on to refine the alternatives and criteria. Finally, when knowledge is sufficient, decisions using judgment based on this knowledge can be made. Thus, according to this argument, the most valuable type of information is a decision, as it is based on all the less valuable information types. In other words, decision-making support requires the management of data, models and knowledge, and the associated judgment on which decisions are based.

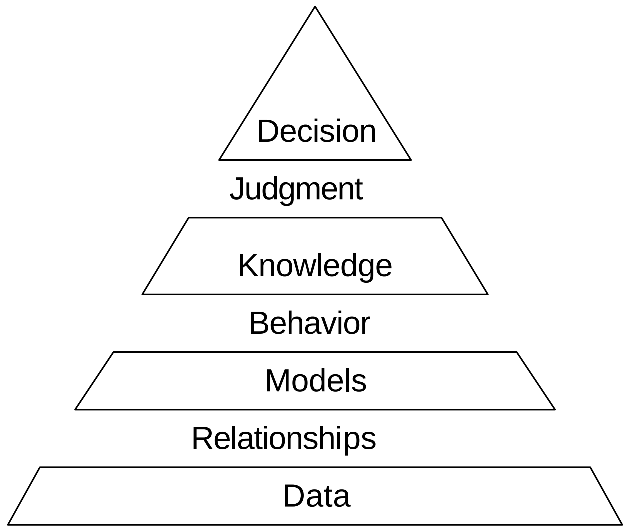

Figure 1: The Value of Information

If for example, you are interested in buying a new computer system, you look on-line or in a catalog and find all kinds of data on processor clock speed and memory size for each alternative computer you want to consider. If you know the relationships between these data, you have a model for how a potential computer might perform. In fact, some of the computer magazines generate measures based on such models. If you have worked with computers enough (i.e. you have enough knowledge of computer systems), you can use the data and models to actually predict the performance of the computers you are considering. Furthermore, your knowledge helps determine the criteria for selecting your new computer from the alternatives. Based on this knowledge, the data and the models, you use your judgment to make a decision about which computer is best to buy.

As shown in Figure 1, decisions are dependent on the weaker types of information. In the statement made at the beginning of the paper, “Problem solving is generating and refining information punctuated by decision-making”, both generation and refinement use data, models and knowledge coupled through relationships and behaviors.

2.2 Decision Thinking Background

The path to understanding the importance of decisions began for the author over fifteen years ago. In 1985 he and colleagues in Artificial Intelligence began a study to understand the structure and operations used by engineering designers [Stauffer 87, Ullman 1988]. To understand design engineer performance it was necessary to observe engineers engaged in the design process. This was accomplished in this research through a technique called protocol analysis. Protocol analysis involved the video and audio taping of engineers as they solved specific design problems. While developing their solutions they were asked to "remain verbal" [Newell 1972]. The engineers' verbalizations were then transcribed. Reading the transcripts, and viewing the designer's gestures and drawing actions on videotape many types of protocol data were extracted from the data.

A total of five separate design protocols were recorded. Two experienced design engineers solved a problem concerned with the packaging of small batteries in a computer and three others solved the problem of designing a machine to coat both sides of an aluminum plate with a thin film.

Results from this study led to a model of how engineers process information that is still referenced today [Ullman 1988]. Additionally, it provided data for understanding the information requests of engineers when redesigning a device [Kuffner 1991] and a deeper understanding of the evolution of criteria during the design process [McGinnis 1992].

Through this research, concepts were developed for a system that recorded the history of the design process for later review. Efforts to build such a system began with focus on the assemblies, parts and features that evolve during the design process. This was not successful as these objects are the results of the activity and thus are impossible to track early in the process and change during the process. Thus, focus turned to tracking the issues addressed [Nagy 1992]. This view was based on the IBIS model of information.

Rittel [Rittel 1973] describes IBIS, a model for organizing the deliberation process that occurs during complex decision-making. The IBIS model organizes the deliberation process into a network of three data elements, *Issues, Positions,* and *Arguments.* An *issue* is an identified problem to be resolved by deliberation. Each issue can have many *Positions* that are proposed solutions developed to resolve the issue. Each position can have any number of *Arguments* that support or oppose that position. The model presented in the next section is based on IBIS.

In the 1980s IBIS was applied and extended to support software design and information capture [Yakemovic 1989, Yakemovic 1990]. Since that time, this basic model has been used and further extended by others. When applied to tracking the evolution of design information, the use of IBIS meant organizing the system about the issues addressed rather than the objects developed. However, this to proved limiting, as, throughout the design process new issues are introduced, old issues are revisited, and some issues abandoned. What became evident was that it was not the issues themselves that were important to track, but their resolution – the decisions to resolve an issue, generate new issues, postpone issue resolution, or abandon an issue. Thus, evolved the decisioncentric view taken in this paper and the statement; “Problem solving is generating and refining information punctuated by decision-making.”

3. THE INFORMATION, ACTIVITIES AND STRATEGIES OF DECISION-MAKING

The following sections summarize what is known about the information managed during decision-making, the activities used to develop and refine the information, and strategies used to manage the activities. The models presented are not the only ones in the literature. However, they summarize the predominant view of how individuals and teams mange the decision-making process.

3.1 Decision-Making Information

The default model for discussing design as decision-making is based the IBIS model introduced in Section 2.2. The types of information used in decision-making are shown in Figure 2. Each of the classes of information and their relationships shown in the figure are defined in the text below.

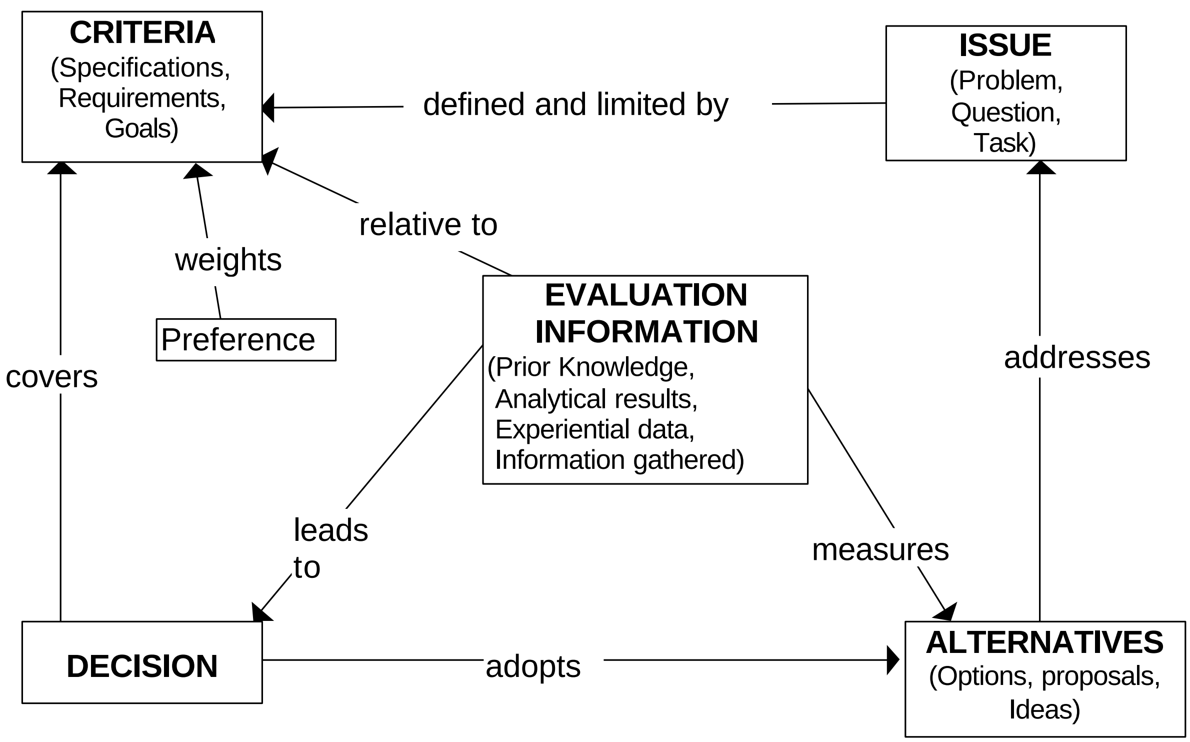

Figure 2: Decision-Making Information

Issue: *An Issue is* a call for action to resolve some question or a problem. An issue is defined and limited by the criteria used to measure its resolution. *Issues are generally expressed as the desire to change, design, redesign, create, fix, develop, or choose an object which meets a number of stated and unstated criteria.* The term "object" in the previous definition can mean any technical or business system, assembly, part, or feature. It can refer to hardware, electric device or software. It can refer to the form or function of the object.

Criteria: *Criteria limit solutions raised by an issue.* There are two major parts to a criterion, the attribute of the alternative measured and a target value for the attribute. The term "criterion" is used synonymously with "requirement", "goal" or "specification" as all limit the space of acceptable solutions for the issue. *Criteria are developed by the issue stakeholders, those individuals responsible for or affected by the resolution of the issue.* Each stakeholder has **preference** for how important s/he feels that each criterion is to the successful resolution of the issue. The combination of the criteria and the preference for them is often called the *value model* because their combination is used to measure or place a value on the alternatives.

Alternatives: *An Alternative is an option generated to address or respond to a particular issue. The goal of the decision-making is to find an alternative that the decision-makers agree to adopt.* Alternatives are often called "options", "ideas", "proposals", or "positions." Any number of alternatives may be developed to resolve a design issue.

Evaluation: *Evaluation information comprises the results of determining how well the alternatives resolve the issue.* Evaluation is the activity of argumentation supported by information developed through prior knowledge, analysis, experimentation, or information gathering (e.g. expert advice). *An argument is the rationale for either supporting or opposing a particular alternative.* Argumentation measures alternatives with respect to criteria, and these arguments lead to agreements.

Decision: *A decision is the agreement to adopt an alternative(s) to resolve the issue.* Decisions are dynamic; they may later be changed as criteria and preferences change, and as new alternatives are generated.

3.2 Decision-Making Activities

The activities that generate and manage the various types of information have been well studied by the author [Ullman 1988, Stauffer 1991, Nagy 1992, McGinnis 1992, and Herling 1997] and many others [Hales 1987 and Blessing 1994]. Recent papers by Girod [Girod 2000b, c] do a good job of summarizing these activities. The listing below is based on his work, but has been condensed to better serve this document.

- **Issues**
    - Generating issues
    - Organizing issues to be worked on
- **Criteria**
    - Identifying criteria
    - Refining criteria to insure understanding
    - Weighting criteria (establishing preference)
- **Alternatives**
    - Identifying alternatives
    - Clarifying the alternatives' working principles
    - Clarifying the alternatives' environment
- **Evaluation**
    - Establishing alternative performance relative to a particular criterion
    - Gathering external information
    - Generating analytical or experimental results
- **Decision**
    - Choosing the best alternative
    - Decision what to do next
- **The Process**
    - Controlling the decision-making process

Among Girod's findings, addressed later in this paper, the time spent on the different activities was more dependent on process approach than on group member's level of professional experience [Girod 2000b]. Thus, the final item on the list above "controlling the decision-making process" is the topic of the next two sections. The first focuses on individual decision-makers and the second on teams. In these sections, the action of controlling the process will be referred to as the decision-making strategy.

3.3 Strategies of Individual Decision-makers

In *The Adaptive Decision Maker*, [Payne 1993] the authors show that an individual's decision-making strategy is contingent on the problem (task variables, context variables), the person (cognitive ability, prior knowledge) and social context (accountability, group membership). The basic thesis in this book is "*that the use of various decision strategies is an adaptive response of a limited-capacity information processor* [i.e. the human] *to the demands of complex task environments*". Here, a decision strategy is a sequence of activities used to transform initial information to a point where the decision maker views the problem solved. A fundamental assumption in this work is that "*individuals decide how to decide by considering both the cognitive effort and the accuracy of various strategies*". In other words, they do a cost-benefit analysis in choosing a strategy.

Problems in which a choice must be made can be categorized as static or dynamic. In static choice problems, the alternatives and the criteria for choosing among them are known and itemizable. Static problems with small sets of discrete choices are well studied by researchers in psychology and marketing. Static problems with continuous choices or a mix of continuous and discrete are routinely solved in engineering using optimization methods. However, if the alternatives and criteria are evolving during the solution of the problem, then it is a dynamic choice problem. Dynamic characteristics are common in many design problems, procurement situations, medical diagnosis, and other business and technical situations.

There are many strategies which people use in making choices in static problems. Static strategies are commonly divided into two broad categories, noncompensatory and compensatory. Noncompensatory strategies do not allow very good performance relative to one criterion to make up for poor performance on another. In other words, no matter how good an alternative is, if it fails on one evaluative criterion, it is eliminated from consideration. One static noncompensatory strategy that is well documented is called the lexicographic strategy. The lexicographic process is to first select all the alternatives that perform above some minimum on the most important criteria. If more than one alternative is above the minimum performance level, then move to second most important criterion and so on until one alternative remains. This is like looking up a word in a dictionary, looking up the first letter eliminates all the words that start with another letter, then you move to the second letter and so on.

Static compensatory strategies allow the decision-makers to balance the good features of an alternative with its weaker features. Compensatory strategies are characterized by summing the judgments across all evaluative criteria. Common compensatory methods are the decision matrix, also known as Pugh's method [Pugh 1990, Ullman 1997b] and utility theory methods [Keeney 1993 and Von Winterfeldt 1986].

When faced with making a decision, people are very adaptive in selecting the strategy they use. In general, conscious decisions on how to decide are not often made. For most simple decision problems, people assess the situation and unconsciously decide on a strategy to be used.

The choice of strategy is sensitive to the number of alternatives since cognitive effort (attention) is a scarce resource. For problems with two alternatives, people use compensatory strategies. If more than two, they use noncompensatory methods. Additionally, time pressure also forces strategies that reduce the problem size. It is common knowledge that, in many engineering situations, the decision is made when the time runs out. In other words, the engineers are still trying to learn more about the alternatives through simulation, testing or other information gathering method when they need to make a decision and move on to the next problem.

All else being equal, decision-makers prefer more accurate and less effortful choices. Since these desires are conflicting, strategy selection is a compromise. One study of students selecting a brand of car based on seven criteria showed that 60% used the lexicographic strategy, 32% used a compensatory strategy (Weighted Sum or Equal Weight), and 8% used a combination of strategies [Payne 1993]. Thus, most of the students traded accuracy for time. The compensatory methods give the greatest accuracy but the noncompensatory methods take the least time. The choice is based on assessment of the importance of an accurate solution, the time available, personal knowledge of strategies and other individual differences.

In an industrial experiment [Naude 1997] (detailed in Section 3.4) a team of managers noted their intuitive choice on a critical, static problem. They then used a structured decision support process and found that their intuitive choice was the least satisfactory. On reflection, they (the subjects not the researchers) concluded that they gave unduly heavy weighting to a very few attributes. In other words, they used a noncompensatory strategy and arrived at a poor choice.

What can be concluded from the above discussion is that, **a method that gives the accuracy of the compensatory strategy with the effort of the noncompensatory strategy would add value to human decision-making activities.**

All of the above discussion is limited to an individual solving a static decision problem. However, many problems are dynamic with the alternatives and criteria developing as the knowledge is gained. This is especially true in engineering and business.

Dynamic choice strategies require the designer to search for new information. One of the first experiments aimed at understanding human information processing during design tasks [Stauffer 1991] showed that over two thirds of the strategies used by the design engineers were searches through design space. In these searches, the designer collected more evidence to refine the alternatives and criteria, and to better evaluate the alternatives relative to the criteria. The searches ended when some satisfactory choice had been made in the time available.

Where the static methods described above have been well studied by psychologists and marketing researchers; dynamic problems, although common, are only beginning to be addressed. In these types of problems, the decision-maker may not be sure enough about the alternatives or the relative importance of the criteria to use the noncompensatory methods. Thus, compensatory methods are used until either all that can be learned is learned, time runs out, or someone demands a decision; then the problem is frozen - becomes static- and a noncompensatory method is used to make a final decision. Thus, **a method that supports compensatory strategies as information about the**

alternatives and criteria evolves would add value to human decision making activities.

3.4 Team Strategies

There has been extensive research in the psychological and sociological communities on how individuals make decisions. There has been substantially less research on teams reaching agreement. A large amount of this research has come from the engineering and business communities.

Olson [Olson 1992, 1996] and his colleagues studied 10 software design groups of 3-7 people to find the structure of team deliberation. They observed experienced teams as they worked on large projects at Anderson Consulting and Microelectronics and Computer Technology Corporation (MCC). The time of observation totaled over 17 hours of different phases of four very different software design projects. They found that they could easily code the meetings using the IBIS model with the addition of categories for project and meeting management, summarizing and other secondary activities. The activity of "Summarizing" is where the team restates the issues, alternatives and criteria in a simple list like fashion.

On reducing the data, they found the only major difference between teams when considering the time allocations for the information categories was for project management. If they factored that time out, they found "we were struck with the overall similarity of the pattern of time use across these 10 meetings." Olson summarized his results in a diagram that showed the distribution of time spent in developing and working on issues, alternatives and criteria, and the transitions between them. A simplified version of this diagram is shown in Figure 3. Here the area of the circles corresponds to the amount of time spent in that category and the width of the arrow corresponds to the frequency of the transitions between the categories.

A limitation of this data is that it is not clear what part of the data corresponds to argumentation – the evaluation of the alternatives relative to the criteria. From the definitions in the references, it appears that most of this activity was lumped with "criteria".

Olson found that 46% of all transitions between categories were in or out of criteria (as implied by the arrows in Figure 3). Similarly, 41% were in or out of alternatives. He also found that 2 or more alternatives were discussed in most issues, 40% of the issues had 3 or more. Of the alternatives proposed 63% of were evaluated verbally and the remainder eliminated without audible evaluation..

What can be concluded from Olson's data is that a model, such as that in Figure 2, can well represent the information in meetings as it includes the major data types in Figure 3. Further, **any support system must allow easy transitions between the major information types.**

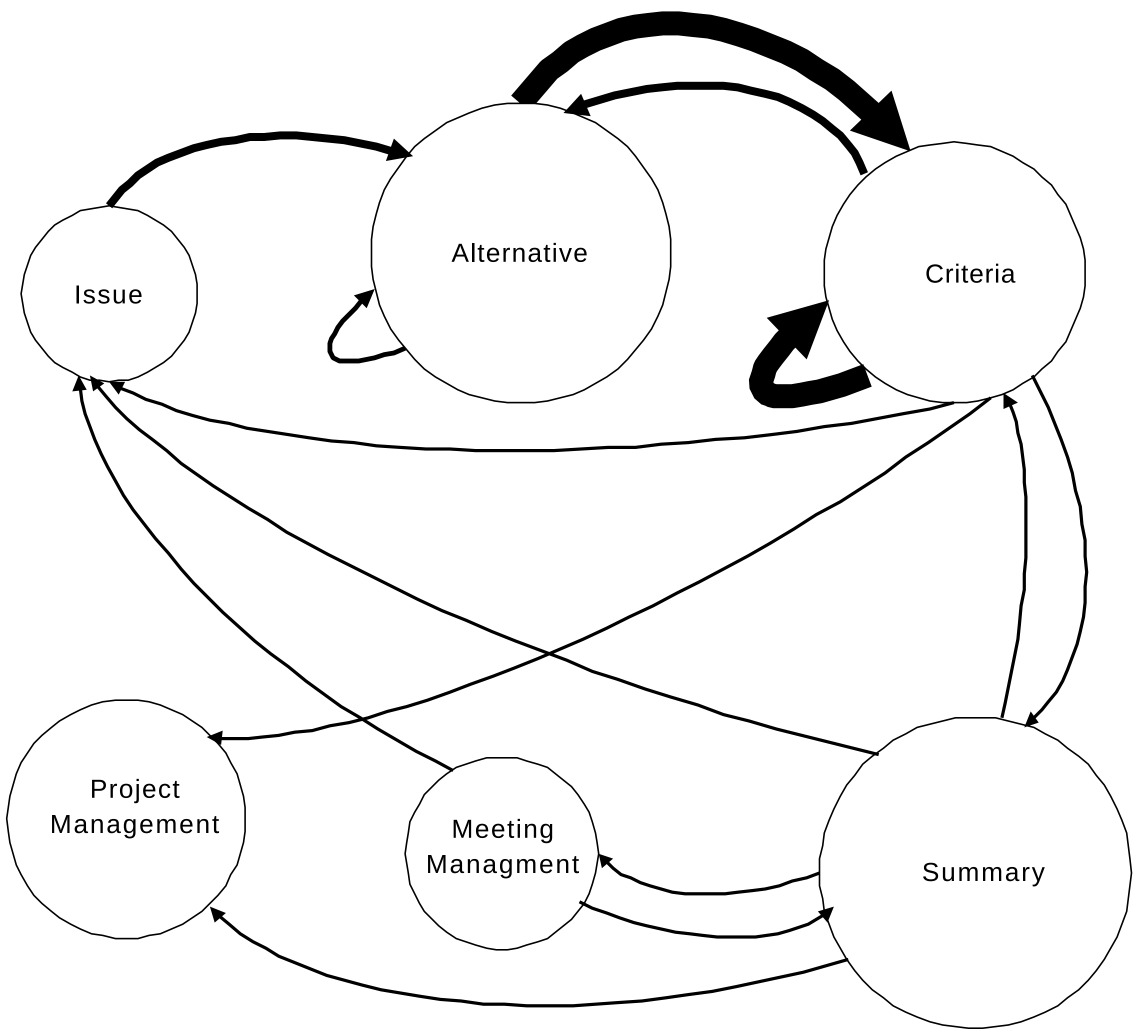

Figure 3: Simplified Diagram From Olson

Another study [Naude 1997] was performed on a team within a brake liner manufacturing company. The researchers introduced a multi-criteria decision making (MCDM) support system. This system basically forced the team to itemize the alternatives and the attributes, weight the attributes and evaluate the alternatives relative to attributes, and then it automatically ranked the alternatives. The team used the system to solve two static, interdependent problems. One had five alternatives and nine attributes (45 relationships). The other was smaller. The team was initially shocked at the results they obtaining using the system.   After reflection they accepted the results as an honest reflection of their perceptions and realized that they were initially forced to make sub-optimal decisions due to cognitive limitations.

4. THE IDEAL DECISION-SUPPORT SYSTEM

The previous sections presented the state of understanding about how individuals and teams make decisions. The goal of this section is to translate this knowledge, and additional evidence, into description of an ideal decision support tool. This philosophy of working from the needs of the user not only follows common design practice, but also is based on a quote from Tversky [Tversky 1988], a psychologist troubled by the poor match between the needs and the tools, "Rather than abandoning decision analysis, we should try to make it responsive to the complexities and limitations of the human mind."

In this section, the attributes of an ideal system are listed and then discussed. The section following this will compare current methods to this ideal. It is assumed that the ideal system will support a team of decision-makers. However, all but two of the attributes (4 and 10) apply equally to individuals. This list builds on earlier efforts to define an ideal [Cook 1993 and Ullman 1995]. The attributes of an ideal decision support system are:

1. Support **inconsistent** decision-making information
2. Support **incomplete** decision-making information
3. Support **uncertain** decision-making information
4. Support **evolving** decision-making information
5. Support the building of a **shared vision**
6. Calculate alternative **ranking, rating and risk**
7. Suggest direction for additional work, **what to do next**
8. Require low **cognitive load**
9. Support a **rational strategy**
10. Leave a **traceable logic** trail
11. Support a **distributed** team

Each of these will now be developed in detail. Before discussing them, an omission from this list must be noted. Many commercially available "decision support" tools actually provide analysis that models or simulates the performance of an alternative. This type of analysis helps the user understand the behavior of the proposed alternative to build knowledge (see Figure 1) and support evaluation (see Figure 2). However, beyond providing information on alternative performance for specified criteria, product analysis does not directly support decision-making. Thus, it has not been included in this listing. A good overview of the use of product analysis in design is in [Fertig 1997] [1].

[1] Work on product analysis in this reference completed by the author of this paper.

The first four attributes of an ideal system focus on the information used in making decisions.

4.1 Support inconsistent decision-making information

Information inconsistency can occur in three ways. First, in a team situation, there can be many different viewpoints about what is important. Viewpoint inconsistency occurs because different people on a team represent different corporate functions or stakeholders. Support systems can either average[2] the viewpoints to generate a single, consistent viewpoint about what is important, or alternatively, honor the diversity. As Naude points out, "The more difficult route is to accept that there are many decision-makers involved and to solicit their individual weights as a preliminary step toward reaching consensus" [Naude 1997].

An ideal decision support tool honors **viewpoint inconsistency**.

The second type of inconsistency is, evaluation inconsistency. This occurs when the evaluation of an alternative relative to a criterion is different across team members. Variation across team members often occurs when information is not well refined, but also happens when there is good experimental or analytical data. This variation naturally occurs because all evaluation is tempered by prior knowledge and level of risk adversity of the team member. Typically, deterministic systems need consistent evaluation information, whereas stochastic methods can support a distribution of evaluation results.

An ideal decision support tool can manage stochastic information representing the **data variation or inconsistency** across team members.

A third type of decision information inconsistency is abstraction inconsistency, the natural mix of qualitative and quantitative information found in most problems. Although engineers like to quantify all the features in a problem, many design constraints are not based on physical laws and cannot be easily represented quantitatively. As stated by Ehrlenspiel and Lenk [Ehrlenspiel 1993] "evaluations are usually expressed in vague, imprecise terms." Thus, the ideal system can exploit both the qualitative and quantitative nature of most problems.

An ideal decision support tool can manage different levels of **abstraction inconsistency** ranging across quantitative and qualitative information.

---

[2] Averaging is dangerous as it violates Arrow's Theorem, which states that averaging can lead to a spurious result.

4.2 Support incomplete decision-making information

Often the information used to make a decision is incomplete. There are two categories of information that may be only partially known or articulated.

Information describing the alternatives and criteria may be incomplete. If all the alternatives are known and all the criteria for evaluation can be itemized (i.e. fixed), then the problem is considered complete. In most problems the alternatives and the criteria for their evaluation evolve as the discussion progresses (see Section 4.4). Team members seldom itemize the entire set of potential alternatives and even when using a system such as Quality Function Deployment (QFD) [Ullman 1997b] they are never assured that they have addressed all the criteria. Thus, most problems are open to new alternatives and criteria and an ideal support system should allow easy addition of new alternatives and criteria.

> An ideal decision support tool should manage **incomplete alternatives and criteria generation** and allow their addition throughout the decision-making

Secondly, in most engineering decision making problems everyone on the design team does not evaluate all the alternatives against all the criteria. This is especially true if the team is multi-disciplinary. Where the completeness in the previous paragraph refers to the number of alternatives and criteria, this characteristic focuses on the completeness of the team evaluation of them.

> An ideal decision support tool should manage **incomplete evaluation** of alternatives relative to criteria.

4.3 Support uncertain decision-making information

Regardless of the all other factors, decision-making is always based on uncertain information. As information evolves (see next item) the uncertainty usually decreases. But even in the most refined engineering models there are inaccuracies, variations and noise that cannot be ignored [Ullman 1997b].

> An ideal decision support tool should manage **uncertain decision-making information.**

4.4 Support evolving decision-making information

"Decisions are not made after gathering all the facts, but rather constructed through an incremental process of planning by successive refinement" [Kuipers 1988]. This statement implies that decision problems are not static, but dynamic (see Section 3.3). The decision-makers are learning about the alternatives and refining the criteria as the choice is being made.

Knowledge changes during problem solution. The more effort, the more is learned about the alternatives, the criteria to evaluate them, and the results of evaluations. These changes occur no matter the level of effort at the beginning of the project to fully define everything. This maturing of the information was documented in a study on the evolution of constraints [McGinnis 1992] based on the same data as used by Stauffer and Ullman discussed in Section 1.1 [Stauffer 1991 and Ullman 1988]. In this study, 227 product features of a simple, seven-part product were studied to determine their origins. Results showed that 70% of the constraints on which conceptual design decisions were based were given in the problem description at the beginning of the effort. Most of the remaining constraints were created by earlier decisions. In the study these were called "derived constraints" as they were derived from the particular process and decisions made by the designer. Later in the design process, during detail design, 84% of the constraints were derived from earlier decisions. Only 7% were based on the given problem description and the remaining 8% of the constraints came from the designers personal knowledge about the function and availability of mechanical components. What is important here is that fully 7/8ths of the constraints in detail evolved as decisions were made. These results led the author to the statement "Problem solving is the evolution of information punctuated by decisions".

Further support for the need to manage evolving information comes from recent studies at Cambridge [Ahmed 2000]. In this study, thirty-four engineering designers from an aerospace company were divided into two groups, experienced and novice. Those categorized as experienced had an average of 19.5 years in industry and those considered novice had an average of 1.5 years. Observations were made in each designer's own environment with additional data taken from discourses and interviews. The most significant result showed that experienced designer preformed more intermediate evaluations than novices. They found that novices generate alternatives, implement them, evaluate how well the alternatives meet the requirements and then iterate back to generating new or refined alternatives. Experienced designers on the other hand, perform preliminary evaluation before implementation, short cutting the cost and time for implementation and final evaluation.

The implications of these results are that it must be easy to change the data used by support systems as the problem is solved. Alternatives must be easily added and deleted, criteria must be refined and the results of new evaluations must be reflected in the decision analysis.

An ideal decision support tool should allow for the **evolution of information.**

4.5 Support the building of a shared vision

The quality of the decision reached is dependent on the amount the team discusses previously unshared information. In an experiment by Winquist [Winquist 98] three-person groups decided which of 2 professors should be nominated for a teaching award. When evaluated along traditional academic measures with all the information, one of the candidates was clearly better than the other. In the experiment however, prior to team discussion, half of the information available for the decision was given to all group members (shared information), and the other half was evenly divided and each portion given to only one member (unshared information). Further, the information was distributed in such a way that the best choice was not obvious to members prior to discussion. Results showed that discussion focused more on members' shared information rather than their unshared information. However, decision quality was affected only by the amount of unshared information discussed. The amount of shared information discussed did not affect decision quality.

Thus, a decision support system should be a mechanism to foster the sharing of information. This activity is often referred to as building a shared vision of the decision problem information. One result of the IBIS studies (Section 2.1) is that even encouraging the team to list the alternatives, criteria and evaluations, appears to improve the decision quality.

An ideal decision support tool should support the building of a **shared vision**.

4.6 Provide alternative ranking, rating and risk

One goal of decision support systems is to help the users choose the best from a list of alternatives. This ranking puts the alternatives in an ordinal listing and is the simplest of the possible results.

An ideal decision support tool should calculate **alternative's rank**.

More informative is to not only rank the alternatives, but to also rate them versus some ideal level of satisfaction. In such systems, typically an alternative that rates 100%, fully satisfies the users. What makes this more informative is that an alternative that is ranked best might only have a rating of 40%, for example, implying that all the alternatives are weak and new ones are needed.

An ideal decision support tool should calculate **alternative's rating**.

Finally, even more information is included if the system can provide an indication of risk involved in selecting an alternative. Even though an alternative ranks best and has a high rating, there still may be substantial risk in choosing it because, for example, the choice may be based on inconsistent (Section 4.1), incomplete (Section 4.2) or evolving information (Section 4.3).

An ideal decision support tool should calculate the **alternative selection risk.**

4.7 Provide direction for additional work, what to do next

There is usually not enough time or other resources to gather the needed information to make decisions with complete, consistent and fully evolved information. One challenge faced by decision makers is to decide which alternatives to eliminate from consideration, and which attribute(s) of the remaining alternative(s) to refine and evaluate more effectively. Support for evaluation can come in terms of developing analytical or physical models, obtaining previously developed information or hiring consultants to supply the needed information. Regardless of source, this need for information creates a sub-problem within each problem. Namely, under the constraints of time, current knowledge and resources to develop increased knowledge, what research should be undertaken to develop the information needed to make a more informed decision. In other words, a sub problem is to decide what to do next [Ullman 1997a].

Support for what to do next usually comes in the form of a sensitivity analysis. This analysis potentially permits exploration of the impact of changes in the alternatives, criteria or preferences (e.g. tradeoff studies) among the criteria. It can tell team members how changes will affect the ranking, rating and risk for each alternative. Based on these changes, the system or the team can determine what to do next.

Note that where the previous six attributes of an ideal decision support system focused on information used in making a decision, this one is directed at managing the decision-making process.

An ideal decision support tool should give guidance on **what to do next.**

4.8 Require low cognitive load

In the one of the best books on how individuals make decisions [Payne 1993], the authors state, “One way we can encourage decision makers to be more normative (more precisely reflect their values and beliefs) is by reducing the cognitive effort demands of the task.” However, it appears that reducing the cognitive load is a very challenging undertaking. Consider the following quotation: “Two decades of research have emphasized the short comings of human judgment and decision-making processes. A focus on errors and biases in decision-making has led naturally to a concern with approaches for aiding and improving decision-making. One such approach is to train decision makers to recognize and avoid errors. Another approach is to construct a decision aid, usually based on a mathematical representation of the problem, to support the error prone decision maker. Although the literature shows that both of these approaches are capable of improving on intuitive processes, decision makers are often reluctant to rely on them' [Kleinmuntz 1993]. In other words, in spite of the best efforts to date, methods and tools are still not well accepted. It is hypothesized that much of this nonacceptance is due to additional cognitive load imposed by these methods.

An ideal decision support system should; help a team reach a better decision than they would without its use (see Naude's research results in Section 3.4), and not require an increase in cognitive load. An ideal system should provide sufficient value added to the team that they want to use it.

An ideal decision support tool should require **low cognitive load.**

4.9 Support a rational strategy

Dylla [Dylla 1989,1991] in an experiment consisting of 6 mechanical engineers individually designing a simple mounting system found that there was a positive correlation between the time spent on goal analysis and the quality of the solution, Figure 4. By "goal analysis" Dylla was referring to understanding the criteria that define the ideal solution to the problem. In his experiments he measured quality by using a group of experts to judge the final designs. One aspect of their judgment was how well the solution met the original criteria and so it is not surprising that the subjects who spent more time had better results. This finding experimentally supports anecdotal evidence that says that work early to understand the problem pays off with shorter time to market and higher product quality [Ullman 1997b].

Dylla also found that the more complete the consideration of design space, the higher the quality of the final product, Figure 4. In the experiment he identified 18 partial solutions to the problem. Combinations of these partial solutions were used by the subjects to build complete devices. He found that the higher number of the partial solutions the better the product quality. This supports the anecdotal evidence that "If you develop one idea, it will probably be a poor idea; if you generate twenty ideas, you might have one good idea" [Ullman 97b].

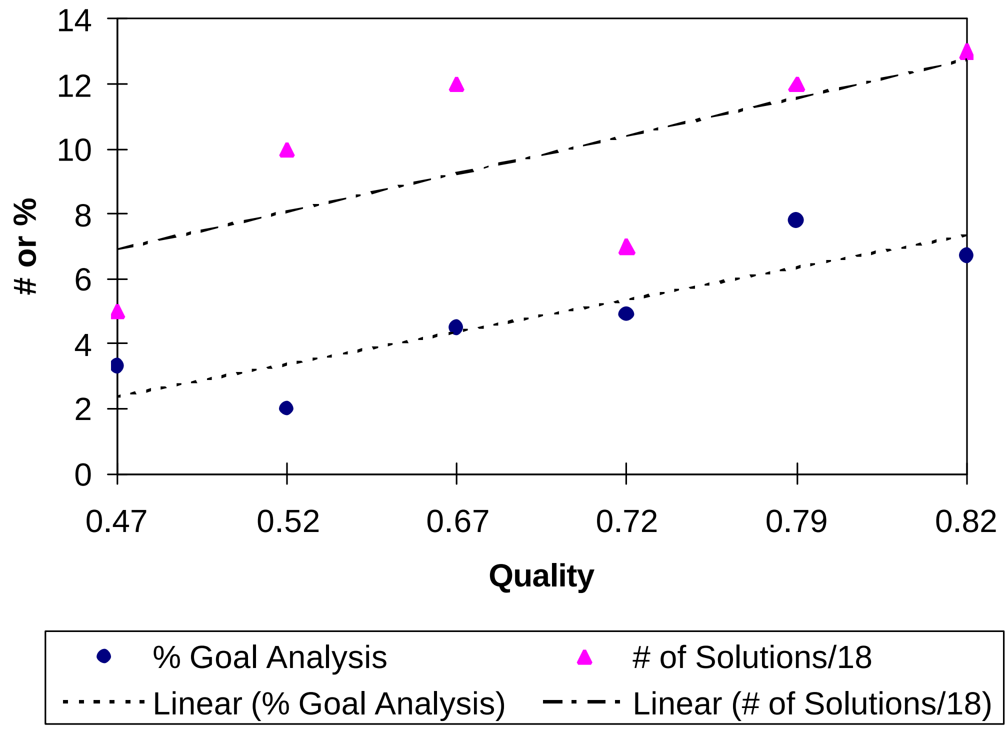

Figure 4. Dylla's research results

An experiment performed on student design teams [Merderer 1997][3] had as its goal to understand team strategies in a limited design space. Four teams of students were studied while playing the Delta World Design Negotiation Game [Bucciarelli 1994]. This design problem is played on a board representing a two-dimensional Delta World. Four participants -- a project manager, an architect, a structural engineer, and a thermal engineer -- are responsible for seeing that the performance specifications of their specialty are met in the design of a residence for their fellow citizens in this planar world. The specifications (i.e. criteria) are spelled out in detail.

The whole delta world environment is designed so that very little individual knowledge can be used in its solution. This is done be setting up artificial laws of gravity, and two-dimensional people who inhabit this world. This artificially reduces the effect of individual knowledge differences and affords focus on the problem solving strategy.

Merderer analyzed the tapes of four teams of University of Washington students[4] solving the Delta World problem. He compared their activities and strategies with

---

[3] This work was done while the author (Merderer) was a visiting scholar under the direction of current author.

[4] Tapes were complements of Professor Robert P. Smith, University of Washington.

how well the team's resulting design met the given criteria. Two of his conclusions directly supports Dylla's findings and extends them to teams. These are:

- Time spent developing criteria was positively correlated to constraint satisfaction.
- The creation of several alternatives leads to a better result.

Additionally, he found:

- A coherent and carefully conducted evaluation of an alternative leads to a better result.

In fact, he found that if changes in alternatives were made early and often, before careful evaluation of the existing alternatives, the changes did not improve the result and, in some cases, weakened it.

Although much of this paper addresses decision-making strategies, there is not a single strategy that fits all problems. As is evident from the title of the book "The Adaptive Decision Maker" [Payne 1993], individuals and teams adapt their strategy to meet the needs of the problem. Thus, it can be concluded from the two studies above that an ideal system should encourage the development of multiple alternatives and criteria in a flexible manner. Further, it should encourage the careful evaluation of these alternatives against the criteria, allowing the decision-maker(s) to work with the information in a manner that is both rational and flexible.

An ideal decision support tool should **support a rational strategy.**

4.10 Leave a traceable logic trail and information record for justification and reuse

There is extensive literature on the importance of capturing and querying the decision trail for reuse. Traditionally, most decisions in industry are either totally unrecorded or at best only the conclusion is noted in a memo. However, the logic behind the decision, the alternatives considered, the criteria used, and the arguments made are all lost. In engineering research this flow of decision-making is often called design rationale and is still the topic of much research [Ullman 1994].

Besides leaving a history of the decision-making process for later use, there is often need to show the customer, management and outside reviewers a record of what was done and what was not done to reach a decision. A recorded logical methodology for reaching a decision is important in these cases.

Finally, the importance of capturing decision information can be appreciated by looking at the results from a simple experiment [Plous 1993]. Researchers recorded a discussion unbeknownst to the participants. Two weeks later the participants were contacted and asked to write down everything they remembered about the discussion. When compared to the recording, they omitted 90% of the specific points. This is understandable as many of the points were probably not very important. However, of the points recalled nearly half were substantially incorrect. They remembered comments never made. They transformed casual remarks into lengthy orations. And, they converted implicit meanings into explicit comments. This experiment reinforces the point that memory is reconstructive and can not be relied on to explain how a decision was reached or for reuse of a decision-making process in future decisions.

An ideal decision support tool should **traceable logic trail** and information record for justification and reuse.

4.11 Support a distributed team

An ideal system needs to support a team of people, complete with their inconsistent (Section 4.1), incomplete (Section 4.2), uncertain (Section 4.3 and evolving (Section 4.4) input as they build a shared vision (Section 4.5). Additionally, decision-making is becoming increasingly distributed, thus, it is mandatory that a system be able to support people separated by time and distance.

An ideal decision support tool should **support a distributed team**.

5. HOW CURRENT DECISION-MAKING SUPPORT METHODS MEET THE IDEAL

In this section five current methods are compared to the ideal. First each of the methods will be briefly introduced. The length of this paper limits discussion on each method, so numerous references are given to supply more detail. In Table 1, the five methods are then rated relative to the 16 ideals developed in the previous section.

Three caveats are in order. First, the ideals are not easily measurable and are subjective. Well-formed criteria should be measurable and these are not. Second, the ratings were made only by the author and were based on his varying familiarity with them. Third, the final method, *Accord,* is relatively new and was developed by the author. The newness implies that it has not stood the test of time, as have the others. The self-evaluation has a good side and a bad side. On the positive, *Accord* was developed to meet the ideal. On the negative, it is self-evaluation. The evaluation presented here is a refinement of those done in [Girod 2000a] and [Ullman 1995].

IBIS

The IBIS model was introduced in Section 2.2 and further discussed in Section 3.4. MCC implemented IBIS on the computer for ease of posting issues, alternatives and arguments (a combination of criteria and evaluation). They tested IBIS on a software development project at NCR to support the design of a software system [Yakemovic 1989 and 1990]. The entire project addressed 2260 issues and some of the results are:

- IBIS provided a shared memory for the design team. The history of decisions made and recorded was easily reviewed and utilized by members of the design team and management.
- The use of IBIS helped the design team detect design issues that had ``fallen through the cracks''. In fact, it was estimated the resulting savings were 3-6 times greater than the cost of not using the tools.
- Team meetings seemed more productive. The tool helped structure the information, provide a shared vision and helped establish and focus the agenda.
- The team found that the use of the tools supported communication with other levels of the organization.

The IBIS model allows for the representation of incomplete, dynamic information about interdependent issues associated with either the product or the process. Information is expressed informally, so that the design space can include quantitative or qualitative as well as deterministic or distributed data in support of the product or process. There is no formal structure to support belief or preference. Although IBIS can model complex decision-making information, it offers no automated support beyond the representation.

MAUT

Multi-attribute Utility Theory (MAUT) is a commonly used method to provide analytical support to the decision-making process. Multi-attribute utility theory underlies a set of methods for making these choices (in this paper, Pugh's method is a simplified MAUT and *Accord* is a probabilistic extension of it). All MAUT methods include:

1. Define the alternatives and relevant alternative attributes
2. Evaluate each alternative on each attribute.
3. Assign relative weights to the attributes to reflect preference.
4. Combine the attribute weights and evaluations to yield an overall satisfaction evaluation of each alternative.
5. Perform sensitivity analysis and make a decision.

Utility theory allows decision makers to give formalized preference to a space defined by the alternatives and criteria. For example, in one method, each alternative/criteria pair is given a score reflecting how well the alternative meets the criteria. The scores for each alternative are combined with measures of each criterion's importance (i.e. weight) to give a total utility for the alternative. Utility is a measure of preference for one alternative relative to another. A good introduction of this method is given in [Edwards 1994].

Traditional MAUT methods [Von Winterfelt 1986, Keeney 1993] were developed initially for individuals, but there are extensions for teams. A limitation of all MAUT methods is that information must be complete, consistent, certain and quantitative.

AHP

Thomas Saaty [Saaty 1995] developed this method around three principals: the principle of constructing hierarchies, the principle of establishing priorities, and the principle of logical consistency. The use of hierarchies helps to itemize the alternatives and attributes. Establishing priorities is based on pairwise comparisons between the alternatives, one criterion at a time. Thus a problem with 5 alternatives and 4 criteria requires 40 comparisons. He then reduces this data using a weighted average to find the ranking of the alternatives. The method allows for checking consistency, the third principal.

Pugh's Method

Pugh's method is the popular name for the decision matrix method [Pugh 1990, Ullman 1997b], a minimized, on-paper form of MAUT. Selection among itemized alternatives is accomplished by relative comparison to a set of criteria defined by the issue. Each alternative is weighed by its ability to meet each criterion. This method is used to support judgments about qualitative information. It results in an abstract satisfaction calculation for each alternative. Pugh's method supports an individual decision maker and thus uses consistent information.

Accord

Accord is a relatively new system that has been under development by the author for the last six years [Ullman 1997a]. *Accord* was designed to meet the ideal described in this paper and is a form of MAUT with probabilistic underpinnings. It is an extension of EDSS, which was reviewed in Girod [Girod 2000a].

Key points about *Accord* are:

1. It is an extension of MAUT, built on inference using Bayesian Networks [D'Ambrosio 1999]. This probabilistic methodology allows for calculating satisfaction and other results from inconsistent, incomplete, uncertain and evolving information.
2. *Accord* supports the management of knowledge and confidence, necessary components of uncertain and evolving information. Although MAUT allows a strong representation of preference, it has no model of belief. Thus, methods using MAUT focus on ways of modeling the objective or utility function, but have no formalized way to manage information concerning knowledge and confidence.
3. *Accord* is based on a natural model of team deliberation It directly supports the formalization and documentation of the problem elements

(i.e. issues, alternatives, criteria, criteria importance, knowledge and confidence).

4. *Accord* generates a series of team satisfaction values based on constructs of the input information. These values show individual satisfaction and combined team evaluations all based on a well accepted mathematical model [D’Ambrosio 1999]. Further, for each alternative it generated the probability that it is best.
5. Sensitivity analysis gives clear direction on what to do next with no additional information from the team members. This analysis shows the potential for increased (decreased) satisfaction with knowledge increased to the expert level.
6. Changes in the evaluation of the alternative/criterion pairs are recorded in a database which acts as a history of the decision making process. This history records the evolution of the decisions of the design team. Further, the PC instantiation of the method has a window for recording rationale with each alternative/criterion evaluation.
7. This methodology gives clear support for three questions decision makers repeatedly ask:

    A*What is the best alternative?”*
    A*Do we know enough to make a decision yet?”*
    A *What do we need to do next to feel confident about our decision?”*

In Table 1, these five methods are compared to the sixteen ideals. The comparisons are made as short statements and represent the typical use of the method. The caveats stated at the beginning of this section should be heeded when reading this table.

6. SUMMARY

This paper has been an effort to describe why it is important to understand and develop decision support systems. Although written by a design engineer, it has been more heavily influenced by business, psychology and sociology research. The importance of looking at design as a decision-making process has been defended and used as a basis to define the ideal decision support system. This definition has been influential in the development of a new methodology to support decision-making. This methodology along with others has compared to the ideal.

Just as the criteria for an ideal decision support system defined here is a refinement of the work of others, it is anticipated that others will refine what is proposed here. That is as it should be because; designing the criteria for an ideal system is evolutionary just like any other design problem.

			IBIS	MAUT	AHP	Pugh's	*Accord*
1	**Support Inconsistent Information**	**Viewpoint**	Does not represent	Designed for one weighting at a time	Designed for one weighting at a time	Designed for one weighting at a time	Can mange multiple viewpoints
2		**Data**	Can represent multiple arguments	Needs consistent information	Needs consistent information	Needs consistent information	Can manage inconsistent data
3		**Abstraction**	All information is text strings	Only quantitative	Only quantitative	Only qualitative	Both quantitative and qualitative
4	**Support Incomplete Information**	**Alternatives and criteria**	No concept of completeness	Results only as complete as data	Results only as complete as data	Results only as complete as data	Results only as complete as data
5		**Evaluation**	No evaluation supported	Needs complete evaluation	Needs complete evaluation	Needs complete evaluation	Can manage incomplete evaluation
6	**Support uncertain information**		No evaluation supported	Deterministic	Deterministic	Deterministic	Manages uncertainty
7	**Support evolving info.**		Weakly	No	No	No	Yes
8	**Support a shared vision**		Yes	Yes	Yes	Yes	Yes
9	**Generate alternative**	**Rank**	No	Yes	Yes	Yes	Yes
10		**Rate**	No	Yes	Yes	No	Yes
11		**Risk**	No	Limited	No	No	Yes
12	**Suggest what to do next**		No	No	No	No	Yes
13	**Require low cognitive load**		Small	High	High	Medium	Medium
14	**Provide rational strategy**		Yes	Yes	Yes	Yes	Yes
15	**Leave traceable logic trail**		Textual	Yes	Yes	Yes	Yes
16	**Support distributed team**		Limited	Single user	Single user	Single user	Asynchronous results sharing

Table 1.

7. REFERENCES

[Ahmed 2000] “Identifying the differences between novice and experienced designers” S Ahmed, K. Wallace, L. Blessing and M. Moss, in Design for Excellence, the proceedings of the Engineering Design Conference 2000, editors S Sivaloganathan and P.T.J. Andrews, Brunell University, June 2000, p97-106.

[Blessing, 94] Blessing, L., *A Process-Based Approach to Computer-Supported Engineering Design*. Black Bear Press, Ltd. Cambridge, 1994.

[Bucciarelli 1994] *Designing Engineers*. Bucciarelli, Louis L., The MIT Press: Cambridge, MA, 1994.

[Cook 1993] " An Empirical Investigation of Information Search Strategies with implications for Decision Support System Design" G.J.Cook, Decision Sciences, Vol 24, #3, 1993 pp683-697.

[D’Ambrosio 1999] “Inference in Bayesian Networks”, Bruce D’Ambrosio, AI Magazine, Summer 1999, Volume 20 #2, pp 21-36.

[Dylla 91] N. Dylla, Thinking Methods and Procedures in Mechanical Design, Dissertation, Technical University of Munich, 1991, in German.

[Dylla 89] Dylla, Norbert. *Experimental Investigation of the Design Process*. International Conference on Design Engineering, 22-25 August 1989, Vol. 1 No. C377/090.

[Edwards 1994] "SMARTS and SMARTER: Improved Simple Methods for Multiattribute Utility Measurement", Edwards W. and F. H. Barron, Organizational Behavior and Human Decision Processes, #60, pp 306-325.

[Ehrlenspeil 1993] “Einflusse auf den Bewertungsprozess beim Konstruieren” Proceeding of the International Conference on Engineering Design, ICER 93, Ehrlenspeil K. and Lenk E The Hague, NL, pp449-456.

[Fertig 1997] Made Phase II Final Report, Flexible Environment For Conceptual Design, Geometric Modeling and Analysis, and Assembly Process Planning, Fertig, K.W., Ullman D.G. et.al, Contract Number – F33615-94-C-4426, Rockwell Science Center, Palo Alto Laboratory, January 27, 1997.

[Girod 2000a] Decision-making and design concept selection, M Girod, A.C. Elliott, I.C. Wright, in Design for Excellence, the proceedings of the Engineering Design Conference 2000, editors S Sivaloganathan and P.T.J. Andrews, Brunell University, June 2000, p659-666.

[Girod 2000b] Activities in collaborative concept selection processes for engineering design, M. Girod, A.C. Elliot, I.C. Wright and N.D. Burns, DETC2000/DTM-14548, The Design Theory and Methodology Conference, part of the ASM 2000 Design Engineering Technical Conferences, Baltimore Maryland, 9/2000.

[Girod 2000c] A Descriptive Model of Collaborative Concept Selection Process in Engineering Design”, M. Girod, A.C. Elliot, I.C. Wright and N.D. Burns, Proceedings of the International Conference CE2000, Lyon France, pp494-503.

[Hales 1987] Analysis of the Engineering Design Process in an Industrial Context, Dissertation, Universtiy of Cambridge, Gants Hill Publications, Hampshire, UK, 1987.

[Herling 1995] Herling, D., Ullman, D.G., and B. D'Ambrosio, "Engineering Decision Support System (EDSS)", ASME 1995 Design Engineering Technical Conferences, Vol 2, 1995, pp 619-626.

[Herling 1997] “An Engineering Decision Support System (EDSS) with Alternative-Criterion Pair Evaluations”, Herling, D.,Ph.D. Thesis, Department of Mechanical Engineering, Oregon State University, 1997.

[Keeney 1993] *Decisions with Multiple Objectives*, R. L. Keeney and H. Raiffa, Cambridge University Press, 1993.
[Kleinmuntz 1993] “Information Displays and Decision Processes”, Don Kleinmuntz and D A Schkade, Psychological Science, Vol 4, #4 July 1993.
[Kuffner 1991] Kuffner, T. and D.G. Ullman, “The Information Requests of Mechanical Design Engineers”, Design Studies, 12(1), pp 42-50, 1991.
[Kuipers 1988] "Critical Decisions under Uncertainty: Representation and Structure" Kuipers B. , A.J. Moskowitz and J.P. Kassirer, Cognitive Science #12, 177-210, 1988, pp177-211.
[McGinnis 1992] McGinnis, B., D.G. Ullman, "The Evolution of Commitments in the Design of a Component," Journal of Mechanical Design, Vol. 144, March 1992, pp. 1-7.
[Merderer 1997] “Evolution of Information in a Design Team”, C. Merderer, Diplom Aribeit, University of Munich, 1997.
[Nagy 1992] Nagy, R.L. and D.G. Ullman, "A Data Representation for Collaborative Mechanical Design," Research in Engineering Design, Vol. 3, No. 4, 1992, pp. 233-242.
[Naude 1997] (A Case study of Strategic Engineering Decision Making Using Judgmental Modeling and Psychological Profiling, Pate Naude, Geoff Lockett and Kevin Holms, Transactions on Engineering Management, Vol 44 #3, Aug 1997 pp 237-247.
[Newell 1972] *Human Problem Solving*, Alan Newell and Herbert A. Simon, Prentice Hall, Englewood Cliffs, New Jersey, 1972.
[Olson 1992] Olson, G.M. et. al., “Small Group Design Meetings: An analysis of Collaboration”, Human-Computer Interaction, 1992, Vol 7, pp 347-374.
[Olson 1996] Olson, G.M. et.al., “The Structure of Activity During Design Meetings”, in Design Rationale: Concepts, Techniques, and Use, Edited by Thomas P. Moran and John M. Carroll, Lawrence Erlbaum Associates, Publishers, Mahwah New Jersey, USA, 1996, pp 217-244.
[Payne 1993] *The Adaptive Decision Maker*, John Payne, James Bettman and Eric Johnson, Cambridge University Press, 1993.
[Plous 1993] *The Psychology of Decision Making*, Scott Plous, McGraw Hill, 1993.
[Pugh 1990] *Total Design*, Pugh S., Addison Wesley , 1990.
[Rittel 1973] “Dilemmas In a General Theory of Planning”, Rittel HW and Webber MM, Policy Sciences, 4, pp 155-169, 1973.
[Saaty 1995] *Decision Making for Leaders,* Thomas Saaty, RWS Oublications,1995.
[Stauffer 1987] "An Empirical Study on the Process of Mechanical Design", Larry Stauffer, PhD Dissertation, Oregon State University, September 1987.
[Stauffer 1991] Stauffer, L.A. and D.G. Ullman, "Fundamental Processes of Mechanical Designers Based on Empirical Data," Journal of Engineering Design, Vol. 2, No. 2, 1991, pp. 113-126.
[Tversky 1988] *Risk and Rationality: Can Normative and Descriptive Analysis Be Reconciled,* Tversky, A. and D.l Kahneman, Institute of Philosophy & Public Policy; April 1988.
[Ullman 1988] "A Model of the Mechanical Design Process Based on Empirical Data," Ullman, D.G., T.G. Dietterich, L. Stauffer, Academic Press, Artificial Intelligence in Engineering Design and Manufacturing, 2(1), 1988, pp. 33-52.
[Ullman 1994] Ullman, D. Issues Critical to the Development of Design History, Design Rationale and Design Intent Systems. Design Theory and Methodology 94, DE, 68:249-258.
[Ullman 1995] “A Taxonomy for Classifying Engineering Decision Problems and Support Systems”, Ullman, D and D’Ambrosio B., Artificial Intelligence for Engineering Design, Analysis, and Manufacturing, AI EDAM #9, pp. 427-438, 1995.

[Ullman 1997a] Ullman, D., Herling, D. and D'Ambrosio, B., "What to do Next: Using Problem Status to Determine Course of Action", Research in Engineering Design, #9, 1997, pp 214-227.
[Ullman 1997b] *The Mechanical Design Process*, Ullman, D. 2nd edition, McGraw Hill, 1997.
[Ullman 2001] *Twelve Steps to Robust Decisions: Building Consensus in Product Development and Business,* to be published in 2001.
[Von Winterfelt 1986] *Decision Analysis and Behavioral Research*, Von Winterfelt D. and W. Edwards, Cambridge University Press, 1986.
[Winquist 98] "Information pooling: when it impacts group decision making.", Jennifer R. Winquist and James R. Larson Jr. Journal of Personality and Social Psychology, Feb 1998 v74 n2 p371(7).
[Yakemovic 1989] Yakemovic, K.C.B. and Conklin J. 1989. The Capture of Design Rationale on an Industrial Development Project: Preliminary Report Technical Report. STP-279-89 MCC.
[Yakemovic1990] Yakemovic, K.C.B. and Conklin J. 1990. Report on a Development Project Use of an Issue-based Information System: Technical Report STP-279-90 MCC.