Coordinating Decision Making In Large Organizations

William P. Birmingham, Joseph G. D'Ambrosio, Tim Darr, and Ed Durfee
The University of Michigan
Ann Arbor, MI 48109
wpb@eecs.umich.edu
313-936-1590

Abstract

We believe that separation of economic decisions from feasibility decisions is a necessity for achieving organizational goals in a timely manner. In this paper we describe an approach that coordinates both causal and economic decisions of an organization. Concurrency is enhanced as economic decisions can be made without causal information, and the results of these decisions can be used to guide the causal decision makers. The causal decision makers can proceed, for the most part, in parallel since economic interdependencies have been eliminated.

1.0 Introduction

The goal of any organizational structure is to enhance the abilities of its constituents for gainful activity. While organizational structures vary, the need for coordinating constituent decision making transcends any particular structure. We have been developing techniques to support decision making in large-scale organizations, particularly for the design of complex systems (*e.g.*, automobiles). Our approach is based on separating decisions about design feasibility from those regarding the economic value of particular designs.

We have developed a representation and set of algorithms based on constraint-satisfaction techniques to reason over entire design spaces of very large size (hundreds of variables). This representation allows individual agents (constituents in an organization) to make decisions about the feasibility of the design space concurrently, and in some cases to eliminate iteration through this space. This dramatically reduces design time. These algorithms deal with design causality; they describe what is physically possible given some set of component and manufacturing technologies. They do not, however, effectively consider economic impact.

We have noticed that the design of large systems involves a great deal of economic decision making that is entirely separate from causal decision making. The objective here is to assign economic value to points or regions in the feasible design space. Traditional optimization processes, for example, confuse these two things by the very nature of the way a problem is specified. Through well-founded statements of preference (not constraints) about design alternatives and explicit statements about the value of time, we can effectively capture the economic decision-making process. Furthermore, the majority of interactions among causal decision makers stems from economic interdependencies (*e.g.,* all agents must share common budgets such as dollar cost).

In this paper, we describe an approach that coordinates both causal and economic decisions of an organization. Section 2 describes a decision-theoretic approach for distributed economic decision making. Section 3 provides an overview of the representation and algorithms for casual decision making. Finally, Section 4 provides a discussion of the merits of our approach.

2.0 Distributed Economic Decision Making

The field of decision theory [Lindley71] provides tools for rational decision making. Optimality is defined in terms of preference statements made by the decision maker. Specifying economic preferences between alternatives provides simple means for capturing economic goals and is well understood by decision makers. Furthermore, it is mathematically well founded.

Solving a problem using decision theory proceeds as follows. All decision alternatives are identified along with their respective consequences. The desirability of each consequence is determined using statements of preference from the decision maker. Probability is used to measure the likelihood of a consequence and a utility function is used to measure

desirability of an alternative/consequence pair. Using this formulation, the alternative that provides the highest expected utility is chosen.

Multi-attribute decision making allows the evaluation of alternatives based on a number of attributes. In general, attributes cannot be directly combined into a single objective function since they are non-commensurate. For example, the attribute profit is quantitative and has units of dollars, while the attribute corporate image is qualitative and has no standard set of units. Fortunately, Multi-Attribute Utility Theory (MAUT) [Fishburn70] provides means to evaluate the desirability of multi-attribute consequences, and thus, facilitates multi-attribute decision making using a decision theoretic approach. For mutually preferentially independent attributes, the multi-attribute utility function is expressed as a weighted summation of attribute utility functions.

A major limitation of a utility function is that all feasible alternatives must be enumerated and evaluated in order to specify the utility function. Although this requirement prohibits the use of a utility function for many problems, an imprecise utility function [White84], which can be specified based on a subset of alternatives, can be used. An imprecise utility function specifies a partial order on the set of feasible alternatives, and can be used to identify a non-dominated set of alternatives.

In the remaining portion of this section, we describe an approach for distributed economic decision making based on the use of an imprecise utility function for capturing economic preferences. The approach is based on a hierarchy of attributes, and on the distribution of coordinator preferences to subordinates.

2.1 Objective/Attribute Hierarchy

For a given design problem, an attribute/objective hierarchy is identified. This hierarchy is similar to that involved in any MAUT based technique. Objectives are refined to the point where they are directly measurable in physical terms. Figure 2.1 provides an example of refinement of objectives. Measurable objectives are called attributes and are the basis for quantitative decision making. In the example organizational structure given in Figure 2.2, agent 1 would identify relevant attributes for agents 2-5, while agent 5 would refine its assigned attributes into those relevant to agents 6 and 7. The total objective/attribute hierarchy forms a top-down link to coordinate economic decision making in an organization.
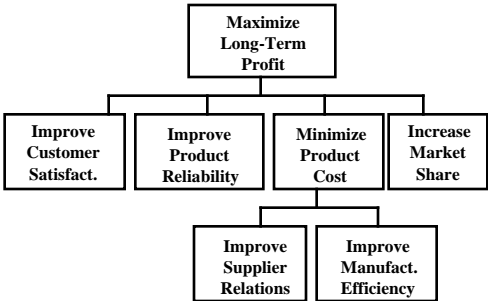

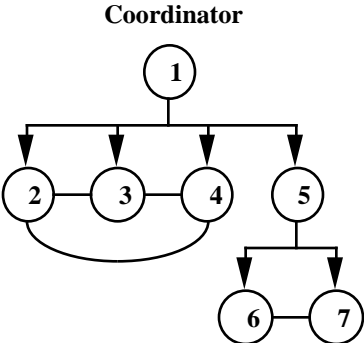**Figure 2.1:** Example objective hierarchy

Figure 2.2: Example organizational structure

2

In the ideal case, all attributes used for decision making are identified in the objective/attribute hierarchy. Unfortunately, it is difficult for large organizations, operating in a dynamic environment, to define and maintain such a hierarchy on a day-to-day basis. Thus other attributes are introduced into the decision making process, especially at the lower levels of the organizational hierarchy.

Consider the example given in Figure 2.3. Three different scenarios of attribute intersection are given for agents 5, 6, and 7 from Figure 2.2. In the first, the obj./attr. hierarchy of agent 5 produces two, mutually exclusive sets of attributes, one each for agents 6 and 7. For both agents, the total set of attributes used for decision making includes those identified by agent 5 as well as local attributes. Note that agents 6 and 7 do not share any common attributes, thus there is no interaction between these agents for economic decision making. In the other two scenarios, agents 6 and 7 share attributes, and may have to work together to achieve global goals.
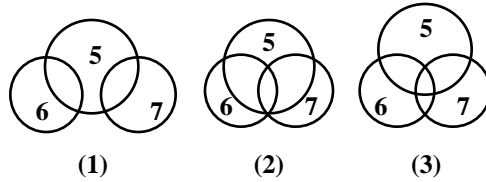
Figure 2.3: Three Venn diagrams illustrating the intersection of attribute sets.

2.2 Distributing Preferences

With the attribute hierarchy defined, statements of economic preference, in terms of relevant attributes, are made by the coordinating agent. These statements are generated by performing pair-wise comparison between a random sample of feasible design alternatives. Appropriate preference statements are then communicated to subordinate agents. For example, in Figure 2.2, statements of economic preference by agent 1 are distributed to agents 2-5. Each of the agents 2-5 combine these preference statements with their own local preferences to create a imprecise utility function to guide their decision making. Agent 5 would also provide appropriate preference statements to agents 6 and 7.

It is necessary to categorize preference statements into two different classes:
1. Primary preferences
2. Secondary preferences

The ordering of alternatives created by primary preferences must always be obeyed. As previously discussed, decision making by a subordinate may involve attributes other than those in the top-down obj./attr. hierarchy. A primary preference statement mandates that local attributes and preferences be ignored in decision making, unless there is a need to break ties. Thus, if the imprecise utility function specified by primary preferences is unable to order two alternatives, then decision making may expand to include local attributes and local preferences.

In contrast, secondary preferences specify an ordering, assuming all other values of local attributes are equal. The decision-making process for the subordinate is based on both those relevant attributes in the obj./attr. hierarchy and local attributes. Merging local and secondary preferences may result in conflicts, and an appropriate conflict resolution strategy must be used.

2.3 Additional Coordination Requirements

As shown in scenarios 2 and 3 in Figure 2.3, the value of some attributes may be influenced by more than one decision maker. Once a feasible set of alternatives has been generated (Section 3.2), it can be shown that, for additive attributes and linear attribute utility functions, a hierarchical distribution of preferences (Sections 2.1 & 2.2) is adequate to achieve optimal decision making by an organization. Unfortunately, such is not the case for non-linear attribute utility functions. In the worst case, the partial order specified by an imprecise utility function defaults to that specified by Pareto preference [Nemhauser89]. In this case, too many alternatives will remain in the non-dominated set, and the time required to evaluate all possible choices will bring the productive activity of an organization to a halt.

For non-linear attribute utility functions, the impact of a local decision, based on a set of attributes, on the overall utility function, is dependent on all decisions relating to those attributes. Assuming a representation as described in Section 3, it is possible to reduce the uncertainty in global utility relating to a local decision. By calculating the range of values that an attribute can take, a *strong* partial order, which produces a small set of non-dominated alternatives, can still be achieved.

3.0 Distributed Causal Decision Making

To coordinate constituent decision-making for determining a set of decision alternatives, we cast the organizational structure as a constraint-satisfaction problem (CSP) [Mackworth77]. A CSP consists of a set of variables, a set of domain values that can be assigned to the variables, and a set of constraints that restrict the assignment of values to variables. A feasible assignment is an assignment of values to variables that satisfies all constraints. The problem of finding a solution to the CSP problem is NP-complete, but we have developed an algorithm and representation that makes the problem tractable in a large number of cases for configuration design problems [Darr94].

3.1 Organizational Representation

In our CSP model, constituent members of an organization are modeled as CSP variables. These constituents are completely described by a set of attribute values that include the capabilities or the functions that the constituent may perform for the organization. Communication paths among constituents are modeled as CSP constraints. Communication of decisions by constituent team members is done wholly through constraints. Constraints define the set of feasible combinations of decisions among team members, and can be used to facilitate the concurrent execution of decision making by the organization as a whole.

In a traditional CSP, all domain values are scalar; we extend this to allow intervals or ranges of values, which facilitates a least commitment strategy for decision-making: a value is fixed at the last possible moment. In a similar vein, our constraints may be hierarchical, and can be refined in much the same way as described in Section 2.1. For example, in Figure 2.2, the coordinator may specify a very high-level constraint that is not directly useful for coordination of decision-making. This constraint may be further specified by constituents 3, 4 and 5 which would allow effective problem-solving behavior among these agents to achieve a set of decision alternatives that can be used in some other part of the organization.

To facilitate organizational decision-making, we augment the traditional CSP formulation by allowing constraints and variables to change dynamically during the decision-making cycle. This allows organization constituents to form teams based on the problem to be solved. We seek ways to achieve the greatest degree of coordination possible by minimizing unnecessary communication. For example, it may be the case that a constituent C1 can provide services (S1, S2, S3) needed by some other constituent or group of constituents. This constituent would be preferred over a set of three constituents (C2, C3, C4), where C2 provides S1, C3 provides S2 and C4 provides S3. The reason for this is that the first constituent can provide (S1, S2, S3) with little, if any, communication, whereas (C2, C3, C4) may require significant communication in order to achieve (S1, S2, S3). This is illustrated in Figure 3.1.
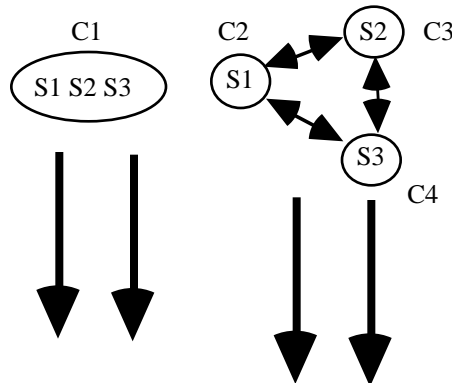

Figure 3.1: Example Communication

3.2 Feasible Set Generation

Identification of a feasible set of alternatives is performed in a decentralized manner with communication between groups strictly through the constraints that link the activities of the two groups. The form of our constraints allows the decisions to be refined as the design progresses, which is consistent with the structures illustrated in Figures 2.1 and 2.2. Since constraints may be hierarchical, constraints can also from groups to minimize communication as illustrated in the previous section. For example, the objective to maximize gas mileage may be expressed in very high-level terms by the

coordinator, which would translate into a wide interval for the set of attributes that relate to that objective. Similarly, the sub-objectives would also be expressed with some "slack" that allows individual decision-makers a certain degree of leeway to make decisions. This is in contrast to an organizational structure where fixed target attribute values are given by the coordinator and teams must engage in bitter negotiation to gain possession of limited resources. Such a scheme would lead to many iterations until project management intervenes to force a resolution. Such an approach loses the opportunities for parallelism and concurrency that are possible with our approach.

4.0 Discussion

We believe that separation of economic decisions from feasibility decisions is a necessity for achieving goals in a timely manner. First, concurrency is enhanced as economic decisions can be made without causal information, and the results of these decisions can be used to guide the causal decision makers. The causal decision makers can proceed, for the most part, in parallel since economic interdependencies have been eliminated. Second, we can direct the proper information to the decision makers without overwhelming them with unnecessary information. Finally, we can provide an integrated, rational *decision-making process* that not only reduces design time, but makes the rationale behind design decisions much clearer.

This rationale can be used to guide organization re-engineering and analysis. For example, by studying past design processes, it is possible to dynamically create highly effective "tiger teams" for future design problems by bringing together only those participants who have proven capabilities relevant to a problem. Furthermore, the communication and decision-making needs of this group can be anticipated, so that a proper organizational structure can be developed to support them.

We have developed several prototype systems demonstrating the effectiveness of our approach. In one series of experiments, we show that the time for computer agents to make decisions is independent of the number of agents involved in the design process. In another series of experiments, we demonstrate how shared utility functions, even in hierarchical organizations, can be used to coordinate decision making in groups with goals that are not necessarily commensurate (at the individual agent level); we also demonstrate how such a utility function can be found using only preferences about decision alternatives.

References

[Darr94] T. P Darr and W. P. Birmingham, *An Attribute-Space Representation and Algorithm for Concurrent Engineering*, Technical Report, The University of Michigan Department of Electrical Engineering and Computer Science.

[Fishburn70] P.C. Fishburn, *Utility Theory for Decision Making*, John Wiley & Sons, New York, 1970.

[Lindley71] D.V. Lindley, *Making Decisions*, Wiley - Interscience: New York, 1971.

[Mackworth77] A. K. Mackworth, "Consistency in Networks of Relations", *Artificial Intelligence*, Vol. 8, 1977, pp. 99-118.

[Nemhauser89] G.L. Nemhauser *et al*., *Handbooks in Operations Research and Management Science: Optimization*, Vol. 1, North-Holland, New York, 1989.

[White84] C.C. White *et al,*. "A Model of Multi-attribute Decision Making and Trade-off Weight Determination Under Uncertainty,*" IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-14, No. 2, March/April, 1984, pp. 223-229.