

SRS-4101 OCT-APP



Computers and Electronics in Agriculture
27 (2000) 355–375

Computers
and electronics
in agriculture

www.elsevier.com/locate/compag

Goals and goal orientation in decision support systems for ecosystem management

D. Nute^{a,b,*}, G. Rosenberg^b, S. Nath^{b,c}, B. Verma^{b,c},
H.M. Rauscher^d, M.J. Twery^e, M. Grove^{e,f}

^a Department of Philosophy, The University of Georgia, Athens, GA 30602, USA

^b Artificial Intelligence Center, The University of Georgia, Athens, GA 30602, USA

^c Department of Biological and Agricultural Engineering, The University of Georgia, Athens, GA 30602, USA

^d Southeast Forest Experiment Station, USDA Forest Service, Asheville, NC, USA

^e Northeast Forest Experiment Station, USDA Forest Service, Burlington, VT, USA

^f University of Vermont, Burlington, VT, USA

Abstract

We explore a goal-oriented as opposed to a problem-oriented approach to DSS development for ecosystem management. Ecosystem management ordinarily is guided by a set of goals that may conflict in various ways. Problems are perceived obstacles to realizing goals. Identifying and resolving conflicts between goals, testing current or projected situations for goal satisfaction, and problem identification all require a robust model of the goal structure for the intended domain. The lowest level of this goal structure must be represented as *desirable future conditions* consisting of proposed values for observable indicators. A model of the causal, legal, and other institutional relations between these desirable future conditions is also needed. Two projects based on a goal-oriented approach to DSS development are described. The first project has produced an initial prototype that incorporates goals for forest management in rules representing three tiers: management unit goals, stand-level goals, and desirable future conditions. The second, at an initial knowledge acquisition stage, is an attempt to develop a participatory decision-making methodology for socially and environmentally sensitive economic development in Central America. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Decision support systems; Goals; Ecosystem management; Forest management; Knowledge representation; Knowledge acquisition

* Corresponding author. Tel.: +1-706-5422823; fax: +1-706-5422839.

E-mail address: dnute@ai.uga.edu (D. Nute)

0168-1699/00/\$ - see front matter © 2000 Elsevier Science B.V. All rights reserved.

PII: S0168-1699(00)00092-2

1. Introduction

Many decision support systems (DSSs) ask the user to identify a goal and then proceed directly to the process of finding recommendations for achieving the selected goal. For example, a DSS for managing a single-species even-aged stand of trees might elicit a timber objective, then base a treatment recommendation on the objective selected and information about the stand. But not all decision making can be characterized in this way. Where an individual or group is engaged in managing a complex enterprise (such as forest management or the management of economic development for a watershed), there will typically be several goals relevant to any decisions being made. Where there are multiple goals, DSS development takes on new dimensions. What kind of knowledge about the goal structure for a domain does a DSS need to help users select a compatible set of goals, identify problems that arise in trying to satisfy that set of goals, and generate and implement solutions to those problems?

Three issues that arise in the development of any knowledge-based system are knowledge acquisition, representation, and utilization. How do we learn the goal structure for a domain? How do we represent this structure? And how is the goal structure used to improve decisions? In this paper, we will not be concerned with knowledge acquisition. Knowledge acquisition is certainly important and difficult; it is often the most difficult part of any knowledge-based system development project. But before we can attempt to construct the goal model for a particular domain, we need to ask ourselves whether there are any common features of the ways goals can be represented for domains where multiple goals are involved. We will begin by first looking at some of the ways we expect to use goals in decision making. This should help us determine what kind of model we will need for the goal structure of a particular domain. Just as in designing a hammer, form will follow function. Although we will avoid issues of knowledge acquisition wherever we can, we will need examples of partial goal structures to illustrate some of the points in the paper. Our focus will be on ecosystem management, but our conclusions should apply to management of other complex enterprises.

We identify three essential tasks where a model for the goal structure of a domain will be used by a mature DSS providing aid for a decision process involving multiple goals. First the DSS must test proposed sets of goals for compatibility. Second, the system must help users resolve any conflicts in a proposed goal set. Third, the DSS must estimate how successful a particular recommendation will be in achieving a set of goals. To perform these tasks, goals must be modeled in a hierarchical structure where the lowest level goals are directly measurable. Furthermore, the DSS must have a robust knowledge of causal and other important relationships among at least the lowest level goals in the system. An additional concept integral to providing meaningful decision support to the ecosystem manager is that of adaptive management (Holling, 1978; Walters and Holling, 1990; Bormann et al., 1994). Adaptive management can be described best as an explicit, structured, and systematic process for learning from one's experience through a cycle of planning, acting, monitoring, and evaluating. The process of refining both

goals and proposed sets of actions to improve the success of management is central to both adaptive management and the process we describe here. In the rest of this paper, we will expand on each of these issues individually.

Finally, we describe two projects aimed at developing goal-oriented decision support systems for ecosystem management. The USDA Forest Service has produced an initial prototype, NED-1, that incorporates goals for forest management in rules representing three tiers: management unit goals, stand-level goals, and desirable future conditions. A second project of Centro Internacional por Agricultura Tropical is at an initial knowledge acquisition stage. This project represents an attempt to develop a participatory decision-making methodology for socially and environmentally sensitive economic development in Central America.

2. Problems and goals

In our research, we are exploring both goal-driven and problem-driven approaches to decision making. What is the difference between goals and problems as we envision them? A *goal* is a *desirable condition*, a situation that someone is willing to allocate resources (time, effort, money, etc.) to achieve. The situation we associate with a goal may be one we want to bring about (as when our goal is to get a particular candidate elected to some public office) or one we want to sustain (as when our goal is to sustain economic growth.) We identify a *problem* when we recognize some obstacle to achieving or maintaining a goal. The existence of an obstacle implies that either a goal state has not been reached or some intervention is required to maintain a goal state that has been reached. Establishing a goal may not imply a need to intervene; if a desirable goal state has been reached, then the most that may be required is that the situation be monitored to make sure it does not change. But identifying a problem invariably implies the need for intervention. For example, a forest manager might say that a healthy forest is required and plan to extract a specific quantity of timber each year. If the forest is already healthy and there is no immediate threat to its health, then no intervention may be required to sustain that situation. All that is required is to monitor the forest to make sure no threats arise to its continued health. However, if the manager has not met the timber goal for the current year, the problem is to select the harvest method and the trees to harvest. Or if a new threat to forest health arises, the problem is to make an appropriate response. Recognizing the problem amounts to recognizing that some action needs to be taken to achieve or maintain a goal state even though the manager may not yet understand what actions might be appropriate. Notice another important difference in the goals used in this example. Having a healthy forest and extracting a specific quantity of timber each year are situations the manager wishes to bring about (if not already present) and to sustain once they exist. Harvesting the desired amount of timber during the current year is a situation the manager wants to bring about, but this is not a sustainable situation. To distinguish between these two types of goals, we will call single harvest quantity a *one-time* goal, and call maintaining forest health and continued timber harvest

sustainable goals. A sustainable goal may or may not require intervention under current circumstances.

One important reason for beginning with goals is that some of the problems that arise in trying to achieve one goal may be the result of the way we pursue other goals we have adopted. Consider again a forest manager who wants to harvest a certain amount of timber from a healthy forest. Suppose that certain trees are diseased and that the only way to stop the spread of the disease is to remove and destroy the diseased trees. The manager may have the means to accomplish either of these goals. The problem is that if resources are allocated to removing the diseased trees, the manager does not have sufficient resources left to harvest the desired timber. The problem may be obvious in this simple case, but it may be far from obvious in more complex cases. Moreover, when several constituencies are involved in and affected by a decision, they may not be fully aware of each other's goals and the way they affect each other. Worse yet, the different stakeholders in a decision may *misunderstand* each other's goals in ways that lead them to think there are conflicts (and problems) where there are none.

Of course, it does not make sense to build a DSS to assist in achieving just any possible set of goals. The goal of a school soccer coach might be to improve the school's soccer field while the goal of a Minister of Agriculture might be to improve the country's maize production. Building a DSS to allow us to consider these two goals as part of a single decision process would strike us as silly. But it would be reasonable to group the soccer coach's goals with a goal of the school's chemistry department to equip a new chemistry lab or to group the Minister of Agriculture's goal with the goal of the Minister of Trade to increase production of manufactured goods in the country. These groupings are more reasonable because the soccer coach and the chemistry department likely depend on the same resources to meet their goals and their decisions involve many of the same stakeholders. The same is true for the Ministers of Agriculture and Trade. So we will want to group together into the same decision making process or DSS only goals which concern the same *issues*. The notion of issue we are using here is not well defined and we will not try to define it better. For present purposes, it is important simply to note that there must be some relationship between the goals in a system that combines them into a single issue and makes it appropriate to consider them together in making decisions. As we use the term, 'issue' is closely related to the equally vague term 'domain'.

Our work suggests that whether or not we consciously approach decision making in the manner outlined above, it is difficult to identify problems until we know what our goals are. Indeed, *there can be no problems without goals*, although a goal does not necessarily imply a problem. A desirable situation may already exist and may be sustainable with little or no effort. On the other hand, some problems may be associated not with a single goal, but with the relationship among several goals. When we employ a single purpose, problem-driven expert or decision support system, we have bypassed the goal setting and problem-identification stages of decision making and proceeded directly to the problem-solving stage. There is certainly nothing wrong with systems that serve this kind of purpose, but our

attention in this paper is focused on the special features of decision making that arise out of the desire to satisfy a multiplicity of goals. For such situations, we believe that a *goal-driven* rather than a *problem-driven* approach is more appropriate. Our thinking is consistent with recent work on participatory approaches to sustainable natural resource management (Allen et al., 1996).

3. Measuring goal satisfaction

Whatever goals we are trying to achieve, there are at least four situations in which we will need to measure how close we are to achieving those goals. First, we will want to evaluate our initial situation to see how far we have to go to realize all of our goals. This is a crucial step in problem identification. Second, we will want to develop and evaluate alternative courses of action (i.e. decisions) expected to help us achieve our goals. Third, a course of action needs to be selected from the alternatives evaluated (typically one that is expected to achieve best the desired goals within any constraints imposed upon decision makers). Finally, after we have selected a course of action and begun to implement it, we will want to monitor our progress towards the stated goals.

Leaving aside for the moment the question of compatibility of goals, we envision the decision process as involving the following steps. First, the decision makers will provide information about the current situation and propose goals to be achieved. Next, the decision makers will determine how success in achieving each of the proposed goals is to be measured. Third, the decision makers will evaluate the current situation to see how well it satisfies the goals. Fourth, the decision makers will reach decisions intended to achieve all the goals proposed or at least to reduce the distance between the current situation and one in which all the goals would be satisfied. We believe there is a role for a software system that will assist the decision makers in some or all of these steps. Our concern is with how to represent the goal structure in such a system.

To measure goal satisfaction, we must *operationalize* our goals. Most goals will not be measurable in the form they are proposed. The decision makers will have to determine specific criteria to be used for measuring the goals. These criteria become more specific goals that we will call *desirable future conditions* or *DFCs*. Of course, all goals are desirable conditions. We use the term 'desirable future condition' and the acronym 'DFC' as technical terms to signify goals that are not analyzed in terms of other goals in our goal structure. We have also noticed that some goals may involve situations that already exist. If the goal is a one-time goal, we will drop it from our goal set as soon as we know that it has been satisfied. A goal will remain in our goal set only if it is either a one-time goal that has not yet been satisfied or a sustainable goal that we want to achieve or maintain. Even though a sustainable DFC might represent a situation that already exists, we will still call it a desirable *future condition* since it is a condition we want to maintain in the future. DFCs can also be thought of as a subset of indicators of any natural resource program, where a complete set of indicators may include variables for monitoring the system as well.

4. Determining goal compatibility and resolving conflicts

Decision makers may propose sets of goals that cannot be simultaneously satisfied. The incompatibility of a set of goals may be very deep so that the goals could not be simultaneously satisfied under any conditions whatsoever, or the incompatibility could be contingent upon the context of the decision. To give an example from forest management, suppose one goal is to produce a certain volume of timber and another goal is to manage the forest for certain forest-dwelling wildlife species. These goals are not logically incompatible. Given a large enough forest, we should be able to produce any amount of timber and also provide habitat for any species that can live in that region of the world. But if the management unit is small, it may be possible to meet the desired timber goals or to provide habitat for the desired species, but it may not be possible to satisfy both goals at the same time. Of course, it may also turn out that neither goal can be met. So even a set containing a single goal may turn out to be 'incompatible' under actual conditions, such as providing polar bear habitat in Georgia. But determining that a single goal cannot be met seems in principle easier than determining that some larger set of goals cannot all be met at the same time in a given context. Our goal structure must include considerable information about causal and other relationships among goals to identify conflicts and incompatibilities.

If any conflicts among goals have been identified, we will want the DSS to provide methods for developing a compatible goal set from the current incompatible set of goals. Of course, this will depend essentially on the priorities of the users. But the representation of the goal structure for the domain can play a critical role in how this is accomplished, or even whether it can be accomplished. Negotiation systems have been developed to help resolve conflicts among multiple parties in order to arrive at an acceptable decision (Jelassi and Foughi, 1989; Thiessen and Loucks, 1992; Fang et al., 1993; Kilgour et al., 1995; Hipel et al., 1996; Faber et al., 1997). Multi-criterial optimization and multi-objective systems (Cohon, 1978; Yeh, 1985; Janssen, 1992; van Keulen, 1992; Yakowitz et al., 1993; van Huylenbroeck, 1996; Yakowitz, 1998) provide support for users to specify goals, constraints and priorities. Essentially, such tools search for solutions that best match (i.e. optimize) user specified criteria. We are looking for something different. We want to explore the possibility of a DSS that incorporates enough knowledge about the domain to propose specific intelligent alternatives to challenged goals as part of the conflict resolution process. We will need a representation of the goal structure for a domain that allows the system to navigate easily from a challenged goal to a possible alternative. Potentially such systems can supplement existing tools such as negotiation and multi-criteria optimization tools.

5. Desirable future conditions

There is little point in setting a goal if we have no way of knowing when that goal has been achieved. Being able to tell whether a goal has been achieved is not

the same thing as knowing how to achieve it. If we do not know how to achieve a goal, we have a problem. But if we do not know what would count as having achieved a goal, we have nothing.

We have already introduced a special kind of goal that we call a *desirable future condition* or *DFC*. To elaborate, a DFC will be associated with a single variable corresponding to some observable situation in the world. This variable might take various kinds of values including both quantitative and qualitative values. The simplest kind of DFC simply specifies that the associated variable has a value within a certain range or set of values. For example, one DFC might be that canopy closure in a stand be at least 80% (canopy closure is defined as the proportion of the sky blocked by leaves or branches when viewed from the ground), or that a particular reservoir contain at least 10 000 m³ of water. These DFCs concern variables that have numerical values. A qualitative DFC might be that there is a pond on a particular piece of land.

Some DFCs, including our example of canopy closure for a stand of trees, involve thresholds or boundary values. When establishing a DFC, we may not know the exact boundary we wish to set. In that case, we may establish a *fuzzy* DFC with *fuzzy* thresholds. Depending on our purposes, we can interpret these thresholds in various ways. We could decide that when the value of the associated variable is equal to or greater than a threshold, then the DFC is satisfied, and when the variable has a value that is at least 5% less than the threshold, then the DFC is not met. Then we can use a linear function to represent the degree to which the DFC is met for the range between 5% below the threshold and the actual threshold. In our example, the extent to which a stand of trees with 77% canopy closure meets our DFC would be represented as 0.4. Where it is difficult or impossible to completely satisfy all of a set of non-fuzzy DFCs, we might use fuzzy DFCs to try to maximize partial fulfilment of the set.

Our first proposal for the goal structure for a domain is that every goal should be reducible to a set of DFCs. This reduction might be complex, involving various alternative groupings of DFCs. But unless some such mechanism is available for unpacking goals into sets of values for observable variables, determining when a goal is satisfied will remain a mystery. Once our model of the goals for a domain provides a reduction of each goal to some (possibly alternative) sets of DFCs, we can evaluate goal satisfaction. Given any situation, actual or projected, from which we can infer values for all the variables associated with the DFCs that define the set of selected goals, we can determine which of the goals in the set are satisfied. This is a necessary step in problem satisfaction and in testing proposed solutions to problems. Analysis of DFCs also provides a means for identifying information gaps even before decision alternatives have been developed and evaluated.

The notion of a DFC is relative to our purposes and resources. For example, one of our primary goals for a watershed might be drinkable water that does not need to be filtered. An intermediate-level goal might be that the turbidity of the water be low at certain locations where human activity affects the turbidity. Maintaining this goal will require monitoring. If we have the resources to measure settling times and amounts of sediment, then we could establish DFCs in these terms. But if we do not

have the resources for these measurements, we could simply set the goal that the water is 'clear'. As another example, it might be ideal in managing a large ecosystem such as a national forest to set DFCs that can only be monitored properly if we have extensive information about each of a vast number of stands. But these data may not be available. In this case, we might be forced to set our DFCs in terms of properties easily observable from aerial surveys.

6. Relationships among goals

To check the compatibility of a set of goals, we need to see whether there is some set of DFCs corresponding to the goal set that can all be satisfied at the same time. Since there may be alternative ways to satisfy a goal, there may be different sets of DFCs corresponding to the entire goal set that must be checked to determine compatibility. Let's use a forestry example again. Suppose one of our goals is that at least half of the stands in a forest have a closed canopy and another is that at least one-third of the stands in the same forest have an open canopy. Assuming that open and closed canopy are defined in terms of the percentage of canopy closure, we will stipulate that a stand has a closed canopy if its canopy closure is at least 80% while it has an open canopy if its canopy closure is no greater than 20%. Notice that we are concerned with two distinct DFCs even though the two DFCs involve the same variable. Our two DFCs, that canopy closure be at least 80% and that canopy closure be no more than 20%, cannot both be satisfied by the same stand at the same time. But our goals do not require this. With two stands or more, we can satisfy both goals. There are two ways distinct sets of DFCs correspond to our goal with only two stands A and B: canopy closure is at least 80% in A and no more than 20% in B, or canopy closure is at least 80% in B and no more than 20% in A. There are many more possibilities with more stands.

In this example, the relationship between the two DFCs is logical. Other DFCs might be related causally. Continuing our forestry example, suppose one of our goals is to manage for a certain wildlife species and that one of the requirements for the species is sufficient hard mast (acorns, nuts, etc.) to serve as food for the species. We might measure hard mast production in bushels per acre. One of our DFCs might be that we have a stand that produces 50 bushels of hard mast per acre per year. This might be biologically impossible if canopy closure is less than 20% since hard mast is produced by mature trees and there could not be enough mature trees to produce the required hard mast in a stand where canopy closure is less than 20%. Here the relationship between the two DFCs, again one of incompatibility, is a causal one, or more specifically a biological one, rather than a logical one.

Besides logical and causal relationships, DFCs might stand in legal, cultural or other relationships to each other. In addition to the model of the goal structure of a domain that includes a reduction of each goal in the structure to alternative sets of DFCs, we will also need appropriate causal, legal, or other models of the relations between the variables for which DFCs are specified. In many cases this will allow evaluation of the compatibility of sets of DFCs and, thus, of any sets of

goals. In some cases, we will be able to determine that two DFCs cannot be satisfied at the same time without additional information about the particular situation in which decisions are being made. But in other cases, the incompatibility will depend upon actual circumstances. It is possible to have a closed canopy on half the stands in a forest and an open canopy on one-third of the stands in the same forest provided the forest includes more than one stand. This is not possible if we are talking about a forest consisting of a single stand.

While it may be enough in many cases to have appropriate knowledge of the relationships among DFCs, it will often be more efficient if we have some direct knowledge of relationships among at least some higher-level goals independent of any disaggregation of those goals into sets of DFCs. For example, we may know that certain agricultural practices cause water pollution. If one of our goals is to encourage these agricultural practices and another is to reduce water pollution, we do not have to analyze our goals to the level of DFCs to recognize a conflict even though the conflict should certainly show up at the DFC level. In other cases, we may find that certain relationships among higher-level goals become obscured at the DFC level. For example, one goal might be to increase agricultural production and another might be to increase the amount of forested land in the region. A DFC for the first goal might be to convert 100 ha of forested land to maize production per year for the next 10 years, and a DFC for the second might be to plant 100 ha of agricultural land with trees each year for the next 10 years. There is no incompatibility in converting 100 ha of forested land to agricultural use at the same time we convert 100 ha of agricultural land to forest; so the DFCs are compatible. But the first DFC conflicts with the second goal and the second DFC conflicts with the first goal. This conflict will go undiscovered if we only look for conflicts at the DFC level.

At this point, we will say a few words about *constraints*. Our decisions are guided both by our goals and by external constraints. The difference is that we might not choose the constraints under which we must decide and act, but we must nevertheless conform to them. On our model, constraints are represented in part by the causal, legal, and other relationships among goals. Constraints will also apply to the alternative solutions we consider to the problems we identify later in the decision process. We do not currently have a separate method for representing constraints in our goal-driven decision model, but we may find this is necessary later.

The amount of knowledge about relationships among goals that is needed to thoroughly evaluate a set of goals for compatibility is daunting. Automating this analysis for complex domains may require more knowledge than we can ever obtain. In implementing goal-driven decision making for complex domains, it will probably always be necessary to depend on the users for some of this knowledge. But the more knowledge of causal and other relationships between goals incorporated into a DSS, the more help the DSS can provide in goal conflict analysis and the better it will be able to recommend acceptable methods for achieving the goals that are set.

7. Hierarchical structure of goals

We have concluded that our goals must be reducible to alternative sets of DFCs if we are to determine whether a set of goals is satisfied in an actual or projected situation. We have also concluded that we must have a model of essential causal, legal, or other relationships among the DFCs in the domain if we are to determine whether a set of goals can be simultaneously satisfied in a given situation. This suggests a simple, two-tiered goal structure with DFCs in the bottom level and all other goals in the top level. This is in fact the kind of structure used in NED-1, the DSS recently released by the USDA Forest (see below.)

Although a two-tiered goal structure together with some representation of the relationships among at least the DFCs in the structure may be adequate for both goal satisfaction and goal compatibility, a more complex structure will serve us better when we discover incompatibilities in a goal set. If we find that we cannot satisfy all of the selected goals, we must re-evaluate and try to find an acceptable revision of our goal set that might be simultaneously satisfied. The re-evaluation and revision of a goal set is an example of adaptive management and illustrates the necessity of such a flexible approach to the complex task of ecosystem management. The only way to revise a goal set in a two-tiered goal structure is to eliminate one or more of the goals in the top level. Deciding which goals to eliminate will depend on the priorities attached to the different goals. But another revision strategy is possible when we have a richer goal structure (i.e. one that has multiple tiers) because this lets us replace some goals with other goals that on examination may be equally acceptable.

To illustrate this strategy, suppose an initial goal set for a family includes sending the oldest daughter to Old Ivy College. It becomes evident upon analysis that the family resources will not permit the daughter to attend Old Ivy and still satisfy other goals which have a high priority for the family. From the available information, we can infer that a goal of the family is for the oldest daughter to acquire a college education. This is a more general goal and sending the daughter to Old Ivy is just one way the family can achieve it. It may be the preferred way to achieve the goal, but there are certainly others such as sending the daughter to State University where costs are lower. Whether the family makes other sacrifices to send the daughter to Old Ivy or revises its initial goal set by modifying the college the daughter plans to attend depends on the priorities the family attaches to the different goals involved. But whatever the priorities may dictate in the actual case, this example points to a two-step process for revising a goal set. The first step is to *ascend* from a more specific to a more general goal and then *descend* from this general goal to a different specific way to achieve the general goal.

Goal ascent and descent require a rich goal structure. When trying to revise an incompatible goal set to produce an acceptable, compatible goal set, we will want to find as many places as possible where an ascent/descent alternative is available. Whatever priorities we may be able to apply to the goal structure will dictate which of the possible revisions are more attractive. We can then test the resulting revised goal sets for compatibility. Assuming a set of priorities that partially orders the

alternative goal sets, we have a strategy for discovering the most attractive set of goals that has a reasonable chance to be satisfied simultaneously.

These considerations suggest that we should strive to find more abstract goals under which other goals in our developing goal structure fall. We should decompose goals into other goals at the highest level of abstraction or generality possible, finally reducing the lowest level of abstract goals into DFCs only when further decomposition into goals more abstract than DFCs appears unlikely to provide useful means for goal set revision.

8. Goals and problem identification

Earlier we characterized a problem as some obstacle that stands in the way of achieving a goal, an obstacle that requires intervention. The two most likely problems that come to mind are that we either don't know how to achieve the goal or we don't have some resource we need to achieve the goal. A needed resource might be time, labor, money, or almost anything else, but one important resource is information. And it may be impossible to figure out how to solve a particular problem until we have more information. In some cases, we may not be able to figure out how to solve one problem until we have formed a plan for solving another problem.

When we identify a problem, we often identify an additional goal. This is obvious in the case where the solution to a problem is to find additional resources. The new goal, then, is exactly that: to find the additional resources needed to achieve our original goal. Thus, our goal set grows and new problems may be recognized as a result. When we add a new goal to our set, new possibilities for incompatibility arise. If we do not keep all of our goals in sight as we do this, constantly reviewing the set of goals as it develops, we can defeat our own purposes. Consider the example of the O'Henry story, 'The Gift of the Magi'. In the story, a husband sells his cherished watch to buy combs for his wife's hair while the wife cuts and sells her beautiful hair to buy a chain for her husband's watch. Each person had a goal, the obstacle to achieving the goal in each case was a lack of resources, and each came up with a solution to that problem that thwarted the goal. Here the decision makers deliberately prevent communication about and coordination of their decisions because they want their gifts to be a surprise for the other person, but this kind of situation can happen whenever decision makers do not take the broad view for whatever reason.

Besides enlarging our goal set by identifying problems, we may also enlarge our goal set by refining our goals. In O'Henry's story, the initial goal for each of the characters was to buy a suitable Christmas gift for the other. This became refined to the more specific goal of buying fancy combs in one case and to the more specific goal of buying a watch chain in the other. Both the initial goals and the refined goals are in principle compatible. Each identifies money as a problem; so each adopts as a new goal raising sufficient money to buy the chosen present. There is still no incompatibility in the total set of goals for husband and wife. Now each

refines the goal set further, the man deciding to sell his watch and the woman deciding to sell her hair. At this point, the goals become incompatible and the irony arises in the story because the method used to obtain each gift makes the other gift unsuitable. We belabor the example to point out how problem identification can lead to new goals and how new goals may require further iterations of problem identification. Whenever new goals are added to the goal set, incompatibilities may arise. If this happens, we must retreat to a previous stage and revise the goal set. Goal selection, goal conflict analysis, and problem identification are components in an iterative process that proceeds until we reach a point where we can begin to implement some of our decisions.

9. Goals and problem solution

The interdependence of goals discussed in the last section also has consequences for how we generate and choose solutions to the problems we identify. Suppose that we have several goals, some of which can be achieved by sustaining an existing situation and others that can only be achieved by bringing about new situations. We will call the first kind of goal a *currently satisfied* goal and the second an *unsatisfied* goal although these terms are somewhat misleading. Even though a currently satisfied goal is met by some existing situation, it might be possible to meet the goal by interrupting the situation that now satisfies the goal and bringing about a new situation that will also satisfy the goal. If this change in the way a goal is satisfied requires resources, our natural tendency is to leave matters as they are; so we tend to concentrate on the goals that are not currently being met.

The lesson we learn from O'Henry's story is that we need to guard against achieving one goal in a way that makes it impossible to achieve another goal. In particular, we do not want to inadvertently interrupt a situation that is satisfying one of our goals while producing a new situation that will meet another goal. And even though we would normally expect the most economical decisions to be one that finds ways to meet unsatisfied goals without affecting currently satisfied goals, this will not always be the case. Assuming that one criterion we use in evaluating possible decisions is conservation of resources, a thorough analysis may show that the best way to do this is to take an action that will result in some of our currently satisfied goals being satisfied in new ways. Resource limitations may even force us to find such a solution in some cases.

As an example, suppose a young man who is attending State University has two primary goals: to obtain a college education and to provide care for a recently disabled parent. By attending State University, the young man is already satisfying the first goal. He simply needs to find a way to satisfy the second goal without making it impossible to continue to satisfy the first. But assume that State University is far from the home of the parent who requires care, relocating the parent is not an option, and there is not enough money to pay someone to care for the parent. By transferring to Community College near his parent's home, the young man will be able to provide the necessary care personally. This may be the

optimal decision in the circumstances, but it requires that he terminate the situation which is currently satisfying one of his goals.

If the first lesson is that we must be sure to keep all goals in mind to avoid disrupting efforts to satisfy one goal by the way we pursue another, the second lesson must be that we should not become too attached to one particular way of satisfying a goal. What is certainly true is that better decisions are possible when we have a rich model of the goal structure for our domain and of the causal, legal, and other important relationships among the goals in this structure.

10. Goals and implementing decisions

As we mentioned before, we envision an iterative process involving goal selection, conflict analysis, problem identification, solution generation, and solution selection. The process is illustrated in Fig. 1. On this model, decision making begins with neither goals nor problems but with an issue or set of issues. A DSS would be designed around an issue such as forest management or management of economic development in a watershed. What constitutes an issue is a complex question that we will not discuss in detail here, but issues often arise through recognition of a set of problems and of their interrelatedness. Eventually, an issue becomes represented as a set of goals, constraints, and problems. Within our model, the issue is resolved into a goal structure together with the knowledge of the causal, legal, and other relationships among these goals.

Conflict analysis returns us immediately to goal selection if an incompatibility is detected. Otherwise, we proceed to problem identification. When we have identified one or more problems, considered alternative methods for resolving the problems, and selected one of these methods, we may in this process have established new, more specific goals that we must again analyze for conflicts. Thus, a single iteration of our cycle will likely produce a refinement of our original goals. This process is repeated until we have reached a sufficiently specific set of goals that we can proceed to implementation. At any stage in this process, we might detect a conflict in our current goal set. Whenever this happens, we must retrace our steps. This

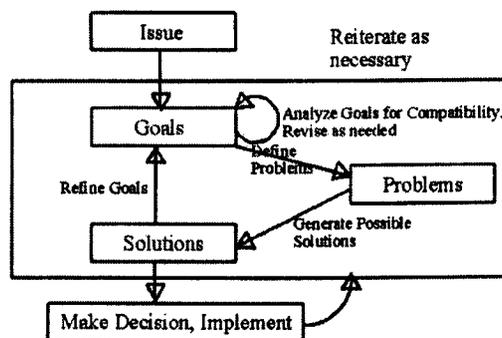


Fig. 1. Process diagram for goal-driven decision making.

might mean selecting a different method for solving a problem, or it might mean retreating all the way back to an earlier goal set to resolve conflicts not detected earlier.

Once we have a set of decisions sufficiently detailed that we can implement them, we will continue to determine goal satisfaction and evaluate for goal conflicts during the implementation phase. Most domains involve a changing situation with factors contributing to change that we cannot control. Since goal conflicts may depend upon actual circumstances, a change in circumstances may bring about conflicts that did not exist at the time of the original decisions. If this happens, we must revisit our decisions. We might do this by beginning the process afresh or by backtracking as far as necessary into our original decision process.

There is another source of conflict in the implementation phase that might force us to reconsider our decisions. A change in circumstances might not prevent us from achieving all of our original goals, but they might put us in a situation where we can only achieve our original goals in certain ways that are unacceptable. When we recognize that we can only proceed with our implementation plan through unacceptable means, we are in effect adding a constraint that we did not originally have to consider. Satisfying this new constraint becomes a new goal, one that is not compatible in the circumstances with our original goal set. For example, suppose a forest manager has decided to meet a long-term timber objective by clearcutting certain stands. Before this plan can be implemented, a residential area is developed from which the area to be clearcut is quite visible. The forest manager may now be unwilling to pursue timber objectives in a manner that will offend the aesthetic sensibilities of the new neighbors. The manager has recognized a new goal, one that could not have been considered in the original decision because residential development of the adjoining land was not foreseen. This new goal is incompatible with the manager's goal of clearcutting the areas originally planned to clearcut although it may not be incompatible with long-term timber objectives when considered more abstractly. The manager will need to abandon specific goals and, through goal ascent and descent, arrive at new specific goals to meet higher-level timber objectives. Indeed, this approach is very compatible with the growing recognition among natural resource personnel that ecosystem management is essentially an adaptive process requiring continual refinement and monitoring — in other words, we learn by doing through what is termed adaptive management (Walters and Hilborn, 1989). Putting such a process into place requires a long-term commitment to collecting and analyzing information from an ecosystem under management, plus the resolve to follow through on changes identified as necessary by evaluation of monitoring efforts. Efforts to implement adaptive management on a practical scale are in place on National Forests throughout the United States and will yield a rich base of information on how well the process has been accepted. Further analysis of whether a clear definition of a goal set early in the planning process leads to successful management may clarify the utility of emphasizing goal definition.

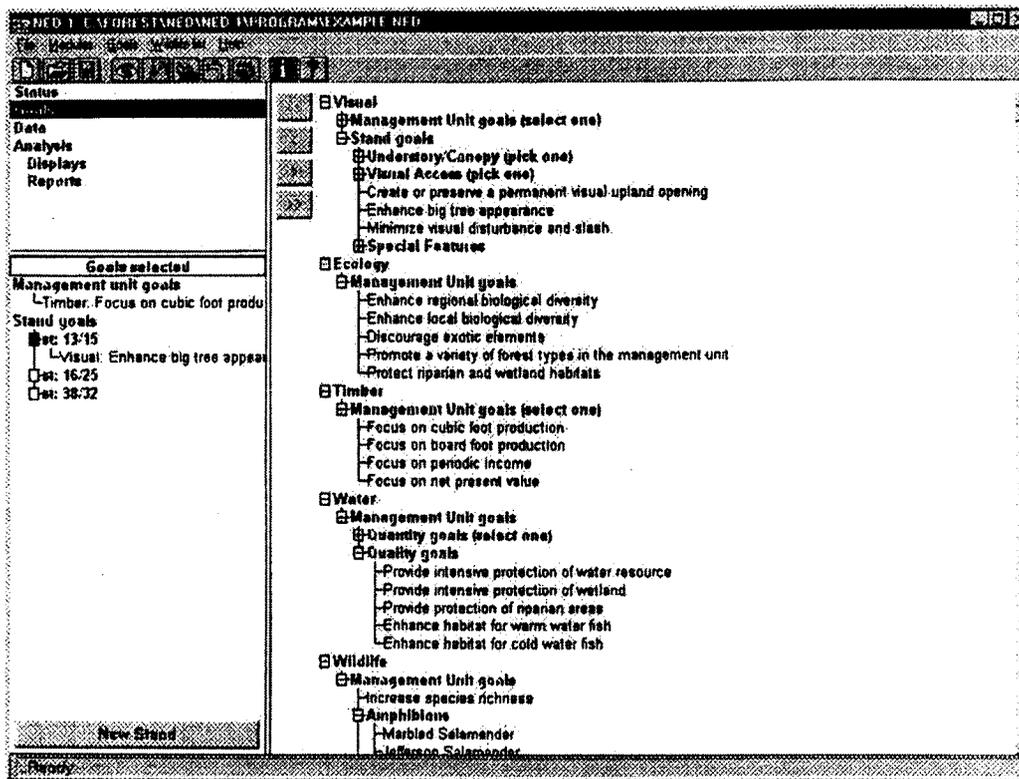


Fig. 2. NED-1 screen for selecting goals.

11. Knowledge acquisition and implementation experience

The conclusions and many of the examples in this paper derive from two projects to develop goal-driven decision support systems for ecosystem management, one sponsored by the USDA Forest Service and the other sponsored by Centro Internacional por Agricultura Tropical (CIAT).

The objective of the USDA Forest Service NED project is to develop a decision model for managing forests for a combination of timber, water, ecological, wildlife, and aesthetic goals (Rauscher et al., 1997; Twery et al., 1997, 1998). The first version of this system, NED-1, allows the user to select from among alternative goals in different categories (Fig. 2). NED-1 then generates a set of DFCs from the user's goal set and evaluates how well current forest conditions satisfy these DFCs. The goal structure for NED-1 is relatively simple. Most management unit goals are defined in terms of combinations of stand level goals, while stand level goals are defined in terms of DFCs for the stand. For example, a management unit goal to focus on cubic foot production as a timber objective requires that at least 65% of the acreage in the management unit is in stands that have a structure suitable for focusing on cubic foot production as a timber objective. A stand has a structure suitable for focusing on cubic foot production if the stand balanced size classes and

the total basal area and minimum acceptable growing stock for the stand fall in a certain range. A stand has balanced size classes if about 5–10% of the plots in the stand are in regeneration, 35–45% are in saplings and poles, 25–35% are in small saw timber, and 10–15% are in large saw timber. An aspen/birch stand satisfies the additional criteria if the total basal area is between 60 and 140 sq. ft/acre and minimum acceptable growing stock is less than or equal to 30 sq. ft/acre. The total basal area and minimum acceptable growing stock requirements vary for different stand types. NED-1 determines whether the goals for both the management unit and individual stands are satisfied and reports the results in a chart showing how well different goals are met (Fig. 3). This gives us a three-tiered goal structure. Management unit goals are at the highest level, stand level goals are at the middle level, and DFCs are at the lowest level.

NED-1 does not perform goal conflict analysis, provide any facilities for resolving goal conflicts, or prescribe either silvicultural methods or specific treatments. A beta version of NED-1 was evaluated at a workshop for government and private forest managers held in November 1998 at the Bent Creek Experimental Forest in Asheville, NC.

The CIAT-sponsored project focuses on socially and environmentally sensitive economic development in Central America. In the CIAT project we are conducting

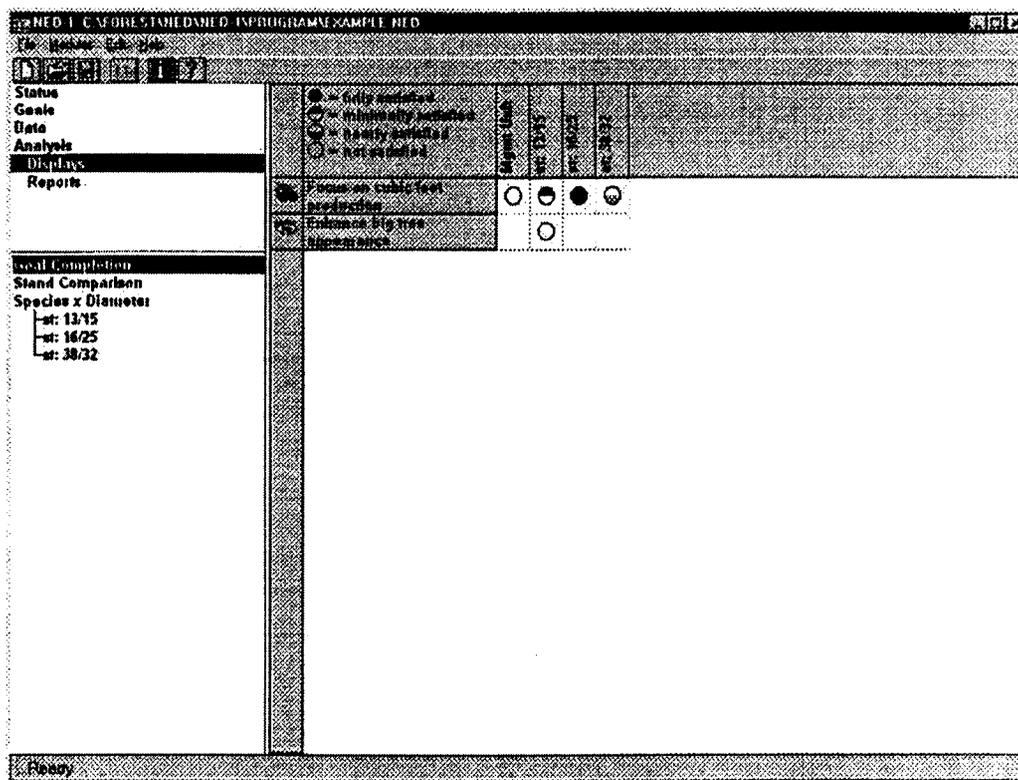


Fig. 3. NED-1 screen showing level of goal satisfaction by current forest conditions.

Figs. 4–6 are examples of some of the dialog screens that have been designed for this prototype system. The first of these screens is used to identify the different stakeholders affected by decisions made regarding the issue under consideration. The second screen lists goals that conflict with some other goal that the decision makers have at least tentatively set. This screen also permits the user to begin prioritizing goals. Similar but simpler screens are provided for goals shared by more than one member of the decision-making group or goals of individual decision makers that do not conflict with the goals of other decision makers. The third screen is an example of a screen of this type. Of course, all of these screens were translated into Spanish for use in our workshops.

This automated template system by no means represents the goal-oriented DSS we hope eventually to develop for CIAT. Here we are only testing various dialogs to see how easily decision makers or facilitators can use them. The system will also

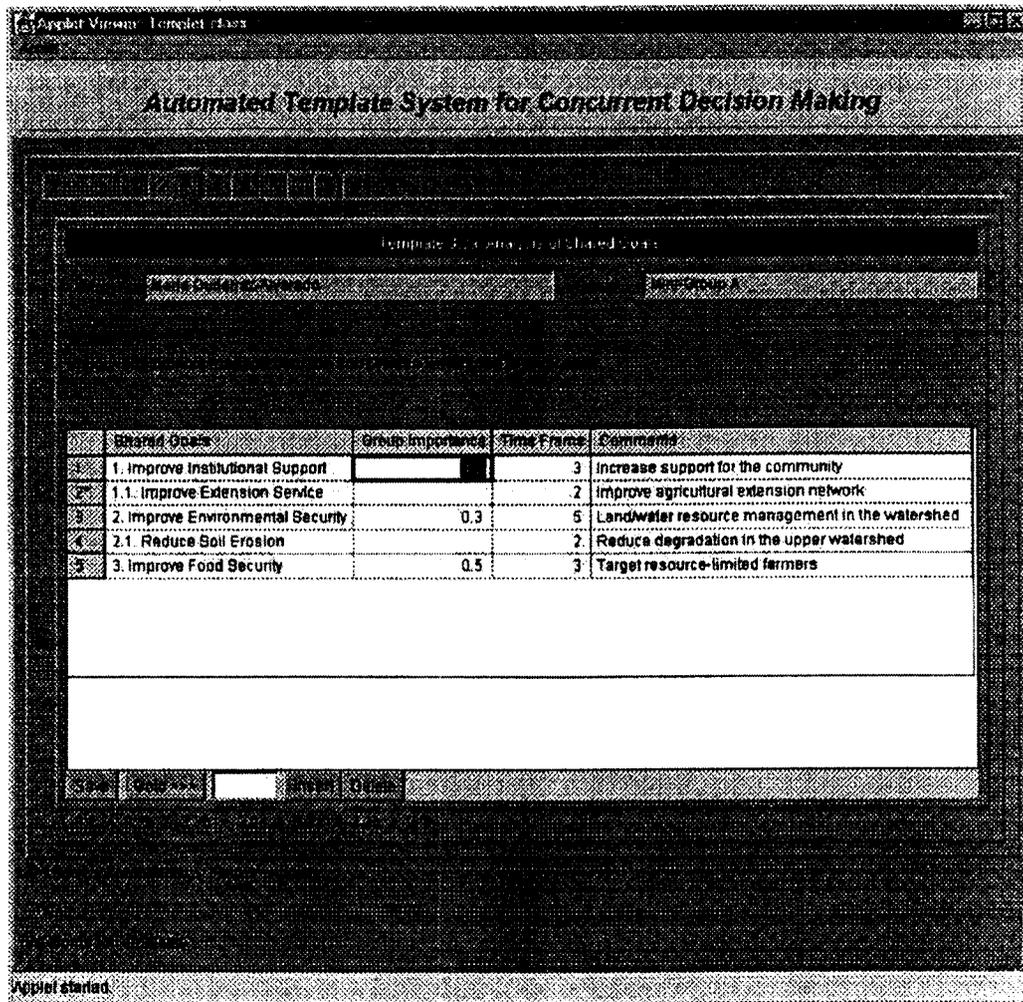


Fig. 5. CIAT project screen for identifying goals shared by all stakeholders.

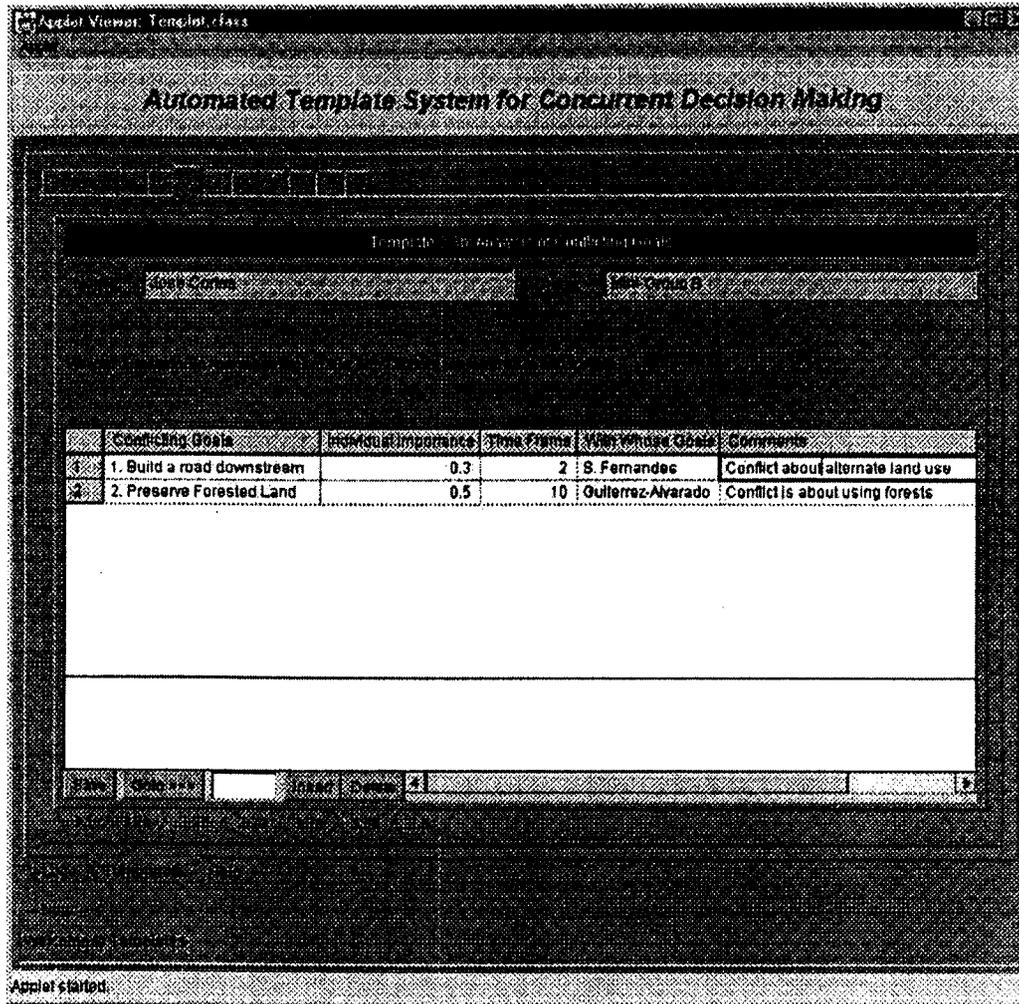


Fig. 6. CIAT project screen for identifying conflicting goals of stakeholders.

store all inputs for later analysis. In this way, they will play an important role in knowledge acquisition. From this information, we hope to glean valuable knowledge about the goal structure for issues we explore in the workshops, knowledge that will be used in later systems that can also identify conflicts in goal sets and recommend methods for resolving these conflicts.

12. Discussion

Our focus has been on knowledge representation and utilization rather than on knowledge acquisition. Any effective DSS requires a model of the causal, legal, and other relationships between the entities in its domain, a model that is adequate for the purposes of the DSS. The goals in a goal structure, and particularly the DFCs,

help identify the entities in a domain and assist the developer in developing a domain model for the DSS. Nevertheless, the requirement that we develop a goal structure sufficiently rich to assist in resolution of goal conflicts significantly increases the difficulty of an already formidable knowledge acquisition task. In early goal-driven DSSs, and perhaps even in mature systems, we will have to provide ways for users to modify the goal structure for the domain. A problem with allowing users to modify the goal structure is that the DSS may not possess a causal, legal, or other model of the relationships between the new goals and goals already represented in the DSS. This particular problem could be avoided if users are required to analyze any new goals they enter in terms of other goals already in the DSS. Another potential problem is that allowing users to modify the existing goal hierarchy could compromise the DSS since part of the expertise incorporated into the DSS is the understanding of the relationships between higher-level goals and DFCs. In NED, for example, our domain experts' best scientific understanding of the habitat requirements for forest-dwelling wildlife species are incorporated into the relationships between the goals to manage for these species and the DFCs represented in the system. These are some of the challenges we face as we attempt to implement goal-driven DSSs.

References

- Allen, W.J., Bosch, O.J.H., Gibson, R.G., Jopp, A.J., 1996. Co-learning our way to sustainability: an integrated and community-based research approach to support natural resource management decision making. In: El-Swaify, S.A., Yakowitz, D. (Eds.), *Multiple objective decision making for land, water, and environmental management: Proceedings of the First International Conference on Multiple Objective Decision Support System (MODSS) for Land, Water, and Environmental Management*, Honolulu, HI, pp. 51–59.
- Bormann, B.T., Brookes, M.H., Ford, E.D., Kiester, A.R., Oliver, C.D., Weigand, J.F., 1994. A framework for sustainable ecosystem management. Gen. Tech. Rep. PNW-GTR-331. U.S. Department of Agriculture, Forest Service, Pacific Northwest Research Station, Portland, OR.
- Cohon, J.L., 1978. *Multiobjective Programming and Planning*. Academic Press, New York.
- Faber, B.G., Wallace, W., Croteau, K., Thomas, V., Small, L., 1997. Active response GIS: An architecture for interactive resource modeling. In: *Proceedings of the GIS'97 Annual Symposium on Geographic Information Systems*. Vancouver, BC, GIS World, Inc., pp. 296–301.
- Fang, L., Hipel, K.W., Kilgour, D.M., 1993. *Interactive Decision Making: The Graph Model for Conflict Resolution*. Wiley, New York.
- Hipel, K.W., Kilgour, D.M., Fang, L., Peng, X., 1996. Using the decision support system GMCR for resolving conflict in resource management. In: El-Swaify, S.A., Yakowitz, D. (Eds.), *Multiple objective decision making for land, water, and environmental management: Proceedings of the First International Conference on Multiple Objective Decision Support System (MODSS) for Land, Water, and Environmental Management*, Honolulu, HI, pp. 23–47.
- Holling, C.S. (Ed.), 1978. *Adaptive Environmental Assessment and Management*. Wiley, New York.
- van Huylbroeck, G., 1996. A multicriteria approach for trade-off analysis between economic and environmental objectives in rural planning. In: El-Swaify, S.A., Yakowitz, D. (Eds.), *Multiple objective decision making for land, water, and environmental management: Proceedings of the First International Conference on Multiple Objective Decision Support System (MODSS) for Land, Water, and Environmental Management*, Honolulu, HI, pp. 419–435.

- Janssen, R., 1992. *Multiobjective Decision Support for Environmental Management*. Kluwer, Dordrecht.
- Jelassi, M.T., Foroughi, A., 1989. Negotiation support systems: an overview of design issues and existing software. *Decision Support Syst.* 5, 167–181.
- van Keulen, H., 1992. Options for agricultural development: a new quantitative approach. In: *Systems Approaches for Agricultural Development, Vol. 2, Proceedings of the International Symposium for Agricultural Development, Bangkok, Thailand*, pp. 355–365.
- Kilgour, D.M., Fang, L., Hipel, K.W., 1995. GMCR in negotiations. *Negotiations J.* 11, 151–156.
- Rauscher, H.M., Kollasch, R.P., Thomasma, S.A., Nute, D.E., Chen, N., Twery, M.J., Bennett, D.J., Cleveland, H., 1997. NED-1: a goal-driven ecosystem management decision support system: technical description. *Integrating Spatial Information Technologies for Tomorrow: GIS 97 Conference Proceedings, Vancouver, BC, February 17–20, GIS World Inc.*, pp. 324–332.
- Thiessen, E.M., Loucks, D.P., 1992. Computer-assisted negotiation of multiobjective water resources conflicts. *Water Resour. Bull.* 28, 163–177.
- Twery, M.J., Bennett, D., Kollasch, R.P., Thomasma, S.A., Stout, S.L., deCalesta, D., Hornbeck, J., Steinman, J., Miller, G., Grove, M., Rauscher, H.M., Gustafson, Cleveland, H., Palmer, J., Hoffman, R., McGuinness, B., Chen, N., Nute, D.E., 1997. NED-1: an integrated decision support system for ecosystem management. In: *1997 ACSM/ASPRS: Annual Convention and Exposition Technical Papers, Vol. IV, American Society for Photogrammetry and Remote Sensing and American Congress on Surveying and Mapping, Bethesda, MD*, pp. 331–342.
- Twery, M.J., Stout, S.L., Loftis, D.L., 1998. Using desired future conditions to integrate multiple resource prescriptions: the Northeast decision model. In: El-Swaify, S.A., Yakowitz, D. (Eds.), *Multiple objective decision making for land, water, and environmental management, Proceedings of the First International Conference on Multiple Objective Decision Support Systems for Land, Water, and Environmental Management: Concepts, Approaches, and Applications, 23–27 July, 1995, Honolulu, HI*. CRC Press, Boca Raton, FL, pp. 197–203.
- Walters, C.J., Hilborn, R., 1989. Ecological optimization and adaptive management. *Annu. Rev. Ecol. Syst.* 9, 157–188.
- Walters, C.J., Holling, C.S., 1990. Large-scale management experiments and learning by doing. *Ecology* 71, 2060–2068.
- Yakowitz, D.S., 1998. A multiattribute tool for decision support: ranking a finite number of alternatives. In: El-Swaify, S.A., Yakowitz, D.S. (Eds.), *Multiple objective decision making for land, water, and environmental management, Proceedings of the First International Conference on Multiple Objective Decision Support Systems for Land, Water, and Environmental Management: Concepts, Approaches, and Applications, 23–27 July, 1995, Honolulu, HI*. CRC Press, Boca Raton, FL, pp. 205–215.
- Yakowitz, D.S., Stone, J.J., Lane, L.J., Heilman, P., Masterson, J., Aboit, J., Imam, B.A., 1993. Decision Support System for evaluating the effects of alternative farm management systems on water quality and economics. *Water Sci. Technol.* 28, 47–54.
- Yeh, W.W.-G., 1985. Reservoir management and operations models: a state-of-the-art review. *Water Resour. Res.* 21, 1797–1818.