



JOSEPH SALEH
JIN KIM

ANTONIO LEAL
JUDEA PEARL

COGNITIVE SYSTEMS LABORATORY

PROGRESS TOWARD A GOAL-DIRECTED DECISION SUPPORT SYSTEM

This work was supported by the
Engineering Psychology Programs, Office of Naval Research
Contract N00014-78-C-0372, Work Unit Number NR 197-049

Approved for Public Release; Distribution Unlimited.
Reproduction in whole or in part is permitted for
any purpose of the United States Government.

PROGRESS TOWARD A GOAL-DIRECTED DECISION SUPPORT SYSTEM

Joseph Saleh
Antonio Leal
Jin Kim
Judea Pearl

Technical Report covering the period 5/1/78 to 10/31/79
Work performed at Cognitive Systems Laboratory
School of Engineering and Applied Science
University of California, Los Angeles
Professor Judea Pearl, Principal Investigator

This work was supported by the
Engineering Psychology Programs, Office of Naval Research
Contract N00014-78-C-0372, Work Unit Number NR 197-049

Approved for Public Release; Distribution Unlimited.
Reproduction in whole or in part is permitted for
any purpose of the United States Government.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) PROGRESS TOWARD A GOAL-DIRECTED DECISION SUPPORT SYSTEM		5. TYPE OF REPORT & PERIOD COVERED Technical Report 5/1/78 - 10/31/79
		6. PERFORMING ORG. REPORT NUMBER UCLA-ENG-CSL-7973
7. AUTHOR(s) Joseph Saleh Antonio Leal Jin Kim Judea Pearl		8. CONTRACT OR GRANT NUMBER(s) N00014-78-C-0372
9. PERFORMING ORGANIZATION NAME AND ADDRESS University of California, Los Angeles School of Engineering and Applied Science Los Angeles, California 90024		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 197-049
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research 800 N. Quincy Street Arlington, Virginia 22217		12. REPORT DATE October 31, 1979
		13. NUMBER OF PAGES 127 pages
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public Release; Distribution Unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Decision Analysis, Decision Support Systems, Means-Ends Analysis, Knowledge-Based Systems, AND/OR Graphs		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report describes a preliminary operational version of a computerized, domain-independent, decision-support system which is based on a novel, goal-directed structure for representing decision problems. The structure allows the user to state relations among aspects, effects, conditions, and goals, in addition to actions and states which are the basic components of the traditional decision-tree approach. The program		

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

interacts with the user in a stylized English-like dialogue, starting with the stated objectives and proceeding to unravel the more detailed means by which these objectives can be realized. At any point in time, the program focuses the user's attention on the issues which are most crucial to the problem at hand. The structure used is more compatible with the way people encode knowledge about problems and actions and, therefore, promises to offer the following advantages: (1) Judgments and beliefs issued by the user would constitute a more valid representation of the user's experience, and (2) The user may be guided toward the discovery of action alternatives he otherwise would not have identified.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

TABLE OF CONTENTS

	<u>Page</u>
1.0. INTRODUCTION	1
1.1. Background	1
1.2. Deficiencies of Decision-Tree Representation	4
1.3. A Goal-Directed Approach	7
2.0. ORGANIZATIONAL DESIGN FOR A GOAL-DIRECTED DECISION SUPPORT SYSTEM	13
2.1. Model Structure	13
2.2. The Major Goal	15
2.3. Subgoals	15
2.4. Actions	16
2.5. Preconditions	19
2.6. Side Effects of Actions on Subgoals	24
3.0. OPTIMIZATION	29
3.1. Motivation	29
3.2. Local Optimization	30
3.3. Optimization Based on First-Order Approximation of Side Effects	35
4.0. DIALOGUE MANAGEMENT	37
4.1. Methodology	37
4.2. Calculation of Expected Value-of-Analysis	40
4.3. An Alternative to the Expected Value-of-Analysis Criterion	45
4.4. Dialogue Management Algorithm	47
5.0. STRUCTURE OF UNCERTAIN EVENTS	49
5.1. Motivation	49
5.2. Uncertain Events	49

TABLE OF CONTENTS (Continued)

	<u>Page</u>
5.3. Propagating the Impact of Uncertain Events	51
5.4. Optimization of Goal-Directed Graphs Containing Uncertain Events	53
6.0. IMPLEMENTATION PROGRESS	55
6.1. Elicitation	55
6.2. Computations and Heuristic Optimization	57
6.3. System Features	58
6.4. Programmer Features	60
6.5. Program Structure	61
6.6. Data Structure	61
7.0. SAMPLE PROTOCOL	69
7.1. Scenario	69
7.2. Goal-Directed System Interview	70
7.3. User Comments	84
7.4. Resulting Goal-Directed Graph Structure	90
7.5. System Recommendation	93
8.0. PRELIMINARY EVALUATION AND FUTURE PROSPECTS	95
APPENDIX A. SIMULATION RESULTS FOR LOCAL OPTIMIZATION	101
APPENDIX B. SENSITIVITY ANALYSIS	115
B.1. Calculation of Δ	115
B.2. Vector Representation of Propagation Functions	119
B.3. Calculation of the Expected Value-of-Analysis (EVA)	125
REFERENCES	127

LIST OF FIGURES

	<u>Page</u>
Figure 1-1. Decision-Tree Representation of Terrorist Attack Problem	9
Figure 1-2. Goal-Directed Representation of Terrorist Attack Problem	10
Figure 2-1. Model Structure	14
Figure 2-2. Goal Structure	17
Figure 2-3. Subgoal Structure	18
Figure 2-4. Action Strategy Structure	20
Figure 2-5. Action Mode Structure	21
Figure 2-6. Criticality Function	23
Figure 2-7. Precondition Structure	25
Figure 2-8. Side Effects of Actions on Subgoals	26
Figure 3-1. A Sample Goal-Directed Graph	32
Figure 3-2. Structure of the Sample Goal-Directed Graph	33
Figure 3-3. First-Order Approximation Model for Side Effects	34
Figure 4-1. An Incomplete Goal-Directed Graph	39
Figure 4-2. Parametric Representation of the Canonical Functions	42
Figure 4-3. Graphical Representation for the EVA Formula	43
Figure 4-4. Graphical Representation for the EVA Formula	44
Figure 5-1. Long-Term Industrial Planning	50
Figure 5-2. Example of a Goal-Directed Graph Containing Uncertain Event	52
Figure 6-1. Program Structure	62
Figure 6-2. Main Data Matrix	63
Figure 6-3. Side Effects Matrix	64

LIST OF FIGURES (Continued)

	<u>Page</u>
Figure 6-4. Name Matrix	64
Figure 6-5. Successor Matrix	66
Figure 6-6. Elicitation Management Matrix	66
Figure 7-1. Structure of Goal-Directed Graph Elicited from the Dialogue	92
Figure A-1. Representative Actions	102
Figure A-2. Simulation Result for the First Representative Action Selection Method	106
Figure A-3. Simulation Result for the Second Representative Action Selection Method	107
Figure A-4. Simulation Result for the Third Representative Action Selection Method	108
Figure A-5. Simulation Result for the Fourth Representative Action Selection Method	109
Figure A-6. A Sample Two-Level Goal-Directed Graph	110
Figure A-7. Goal-Directed Graph Reduction for Identifying Local Optimal Action Under S_1	111
Figure A-8. Graph Compression Process	112
Figure B-1. Propagation of Piecewise Linear Functions in a Goal-Directed Graph	116
Figure B-2. Two Typical Propagation Functions, y_1 and y_2 , Representing Different Subgoals of Goal-Directed Graphs	118
Figure B-3. Cascade of the Two Propagation Functions, y_1 and y_2	120
Figure B-4. Cascade of the Two Propagation Functions, y_1 and y_2	121
Figure B-5. Derivation of Δ From Cascaded Propagation Function	122
Figure B-6. Parametric Representation of the Canonical Functions	123

1.0. INTRODUCTION

This report summarizes the work performed toward the design and implementation of a Goal-Directed Decision-Support System during the period 5/1/78 to 10/31/79. The project was conducted under research contract N00014-78-C-0372 funded by the Engineering Psychology Programs Division of the Office of Naval Research.

The ultimate objective of this project has been to develop and evaluate a computerized decision-support system based on a new, more effective, representational structure which promises to offer several advantages over the traditional decision-tree representation. The scope of this report is confined to the description of the design and implementation of the current operational version of the support system. Other efforts, such as the construction of tools for evaluating the merit of this approach (Kim, 1979) and experiments designed to resolve some theoretical issues (Burns and Pearl, 1979), have been documented elsewhere and will not be included here.

Chapter 1 describes the background, rationale, and motivation for the pursued approach. Chapters 2, 3, 4, 5, and 6 discuss various design issues and the implementation forms currently adapted. Chapter 7 contains a protocol of an actual man-machine dialogue between a user and the current version of the decision-support program. Chapter 8 outlines directions for future research.

1.1. Background

Decision Support Systems (DSS) can be classified into two major categories: Knowledge-Based Systems and Situation-Based Systems. Knowledge-based systems store and employ a large data-base which contains the features and constraints specific to a given problem environment (e.g., they may employ a large medical or legal library) and enable the user to obtain an immediate access to factual information from the problem environment. It is the user's task, then, to

mentally incorporate this information with additional inputs regarding the specific problem situation and come up with a decision strategy. Situation-based systems are domain-independent. They rely on the user carrying most of the background knowledge and expertise and only map into the machine that section of knowledge which the user perceives as relevant to the problem at hand. In this mode the machine acts as a sophisticated friendly 'sounding board'; it does not provide information of its own, but it assists the user in structuring and searching his own knowledge and provides advice on alternative courses of action.

Decision-Analytic technology employs situation-based support. Decision analysts who are called upon to assist in the solution of a given planning problem usually possess less specific knowledge about the problem domain than their customers. The benefit of their services stems primarily from their familiarity with a skeleton structure (i.e., a decision tree) common to all problems, and their ability to represent all problems within the confines of this structure and to draw optimal conclusions from the formal structure once it solidifies. While the optimization process is usually performed on electronic computers, the formalization phase has been accomplished manually, using lengthy interviews with persons intimately familiar with the problem domain.

In the early part of 1975 a project was initiated at UCLA to automate the formalization phase using an interactive computer system which would guide the decision maker through a structured English-like dialogue and construct a decision-tree from his responses. The objectives of this work were three-fold: 1) to provide the decision analysis industry with a practical automated tool for eliciting decision trees in cases where manual elicitation techniques are either infeasible or non-economical, 2) to cast the decision analysts' behavior into a formal framework in order to examine the principles governing the elicitation

procedure and gain a deeper understanding of the dialogue process itself, and 3) to provide experimental psychologists with a standard automated research tool for comparing subjects' behaviors under various conditions and under different support techniques.

From a practical viewpoint, though, the major drawback of manual interviews is their length and cost. Since real-time analysis of decision-trees is beyond the limitation of human computational capability, it invariably happens that many hours of interviews are spent on eliciting portions of the decision tree which do not have decisive bearing on the problem(s) at hand. This fact can only be discovered at a later stage once the problem structure is formalized, and a sensitivity analysis has been conducted on an electronic computer. During the interview itself, however, it is impossible for the analyst to process the entire information obtained by him up to that point, and to select the optimum course for conducting his future inquiries.

A direct man-machine interface could provide three distinct advantages. First, it offers the capability of real-time sensitivity analysis, which in turn could be used to guide the growth of the decision-tree in only the more promising directions. Second, it provides an inexpensive means of updating the program with new knowledge, even by the non-technical decision maker. Finally it opens the way to computerized real-time Delphi methods for aggregating opinions of several remotely located experts.

This project was pursued by A. Leal and was completed in 1976 (Leal, 1976). It culminated in "An Interactive Program for Dynamic Elicitation of Decision Structures" demonstrating the feasibility of constructing a computerized system which interacts with a person in pseudo-natural English and provides assistance in structuring his problem perception, making plan recommendations and communicating the structure to others (Leal and Pearl, 1977). The program's main

techniques were borrowed from both artificial intelligence (AI) and decision-analysis (DA). DA provided a formal structure of knowledge representation in the form of a decision-tree quantified with probability and value assessments. AI provided techniques for heuristic search of game trees and, to a lesser degree, some capabilities for natural languages processing.

Since the completion of Leal's program, the feasibility of automating the process of tree elicitation has attracted the interest of several other laboratories. Merkhofer et al. (1977) at SRI describe a tree structuring support system for command and control applications. A. Leal et al. (1978) at Perceptronics describe an interactive computer aiding system for group decision making designed to support crisis management situations.

1.2. Deficiencies of Decision-Tree Representation

Our experience with the operation of Leal's program confirmed earlier hopes that due to the structural simplicity of decision-trees, only very primitive levels of language-understanding would be sufficient to conduct natural, English-like dialogues. However, the lack of sophisticated language understanding features, aside from accounting for the simplicity of the program, also resulted in several deficiencies. The most serious deficiency arises from the constraint of representing knowledge in tree form.

In many real-world applications, the decision maker may not perceive a problem in the form of a time sequence of decision alternatives and event outcomes, but rather as a static network of influences surrounding issues and factors. Consider, for example, our perception of the environmental pollution problem. The issues of capital investment, energy needs, energy supply, unemployment, public health, etc., all seem to be tightly interwoven in a network of cause and effect relationships. The first step in attacking such a problem should be to explicate the underlying causal network rather than to

hypothesize and evaluate various action/event scenarios.

When a person confronts such a complex problem he is rarely aware of the set of relevant alternative actions available to him at the onset. In fact he usually hopes the analyst would help him identify those alternatives on the basis of certain things he desires to achieve and others he wishes to prevent.

Imagine how awkward it must sound for a person planning the long range economic policy of the U.S. to be asked:

Computer: "What seems to be your problem?"

Planner: "Our long-range economic policy."

Computer: "List the alternatives available to you."

A much more natural and useful question would be:

Computer: "List the effects you would like to see accomplished."

or, Computer: "List the disturbing characteristics of the present situation."

The user may become aware of his immediate options only after unraveling the processes which influence the desired and undesired effects, the preparations needed to make these processes more or less effective, and the conditions which should prevail before an action becomes applicable.

The major difference in the formal representation required for such problems and the one handled by decision-trees is that the atomic entities admitted by the latter representation are restricted to be descriptions of 'world states' or decision 'situations'. The decision maker can express relations among these situations but is unable to express relations between their constituents. For example, when a decision maker is asked to assess the value of a situation resulting from a given event/action sequence, he is presented with the entire sequence and is forced to aggregate the effects of all the event/action components by mental manipulations. He cannot, for example, explicitly

express the belief that raising taxes is a positive contributor to unemployment regardless of other situational factors such as air pollution or the energy embargo. Likewise, he is unable to state explicitly that increased employment (a situational factor) may enhance tax payer's willingness to support more public transportation systems. Instead, he would be required to globally assess the likelihood of obtaining tax payers' support given prior actions and situations.

Decision Analysis is founded on the paradigm that the reliability of human judgments increases when the format of these judgments are made more compatible with the internal format used by people to encode experience. In fact, the sole rationale of problem-decomposition 'divide and conquer' approaches is to reformulate a given problem statement in terms of many, so-called more 'elementary', problem statements to which reliable judgments can be assigned. The reason that one expects these elementary judgments to be more reliable than those involving global considerations is only that the former are more likely to match the format in which human experience is encoded. The decomposition affected by decision-tree analyses only offers the first step toward a structural match between the external and the internal codes. The fragmentation, however, remains too crude to allow the user to express beliefs in a natural and, therefore, reliable manner.

The main objective of the current research project is to devise a richer structure for eliciting knowledge about decision problems, a structure in which aspects, issues, and conditions are represented as independent entities. On the basis of such a structure, it should be feasible to construct a decision-support program that, starting with the stated objectives, should guide the decision maker toward the discovery of action alternatives he otherwise would not have identified.



1.3. A Goal-Directed Approach

To facilitate an 'issue-oriented' problem elicitation program, the internal machine representation of problem situations could be based on the methodology known in the artificial intelligence literature as 'problem reduction' or 'means-ends analysis' (Nilsson, 1971). Each node in this structure represents a subproblem or a subgoal rather than a state description. The task of resolving each separate 'issue', which the decision maker perceives as part of his problem, is regarded as a reduction of the global problem into several components. These can be further reduced to their constituencies, and so on.

A 'means-ends analysis' was first employed in the General-Problem-Solver (GPS) program developed in the late 1960's (Ernst and Newell, 1969). The program is controlled by 'differences': a set of features which make the goal different from the current state. The programmer had to specify along what dimensions these differences are measured, which differences are easier to remove, what are the operators available for the reduction of the differences, and under what condition each reduction operator is applicable. A successful planning program, called STRIPS, based on the same principles was implemented in SRI to plan the actions of an object-manipulating robot (Fikes et al., 1971). In STRIPS too, actions are brought up for consideration by virtue of their potential for reducing the differences (mismatched logical assertions) standing between the desired goal and the current state. When the current state does not possess the conditions necessary for enacting a desired, difference reducing operation, a subgoal is created to generate the missing conditions. The structure underlying this form of reasoning is no longer a tree but an AND/OR graph. The OR nodes represent various types of actions one can employ in attempting to achieve a given subgoal, and the AND nodes represent the remaining subgoals (differences)

all of which should be resolved before a solution is reached. These latter sets of subproblems are of two types; the first contains a set of preconditions that must be realized before the enactment of a previously identified desirable action could be feasible, the second contains a set of adverse effects (additional differences) introduced by such an action.

We have selected the AND/OR Graph as the basic representation for structuring decision problems and, since at each level of expansion the content of deeper levels is determined by the available set of subgoals, we call it a goal-directed structure (Pearl, 1978).

To demonstrate the difference between this structure and the traditional decision-tree, consider two possible conceptualizations of the problem of handling a terrorist attack. Figure 1-1 represents a possible beginning of a decision-tree describing the crisis, while Figure 1-2 represents a goal-directed structure for the same problem. The two basic entities in the latter structure are actions (in  boxes) and subgoals or issues (in  boxes). The root of the graph labeled TERRORIST ATTACK is recognized as involving two main issues: securing the hostages' safety and discouraging future attacks. These are connected by an AND arc to indicate that both issues must be dealt with simultaneously. At this point the natural question for the computer to ask would be, "Could you think of an action which would serve the hostages' safety and at the same time would deter future attacks?" The possibility of 'ATTACK TO RESCUE' immediately comes to one's mind, and the various aspects of this suggestion are explicated. Other actions, intended to resolve each subgoal separately, are then elicited. Each action is characterized by two lists: 1) a preconditions list and 2) an effects list. Any one of the preconditions which is not yet satisfied generates a subgoal (e.g., the condition 'terrorists must agree to postpone deadline' generates the subgoal 'provide terrorists with incentive

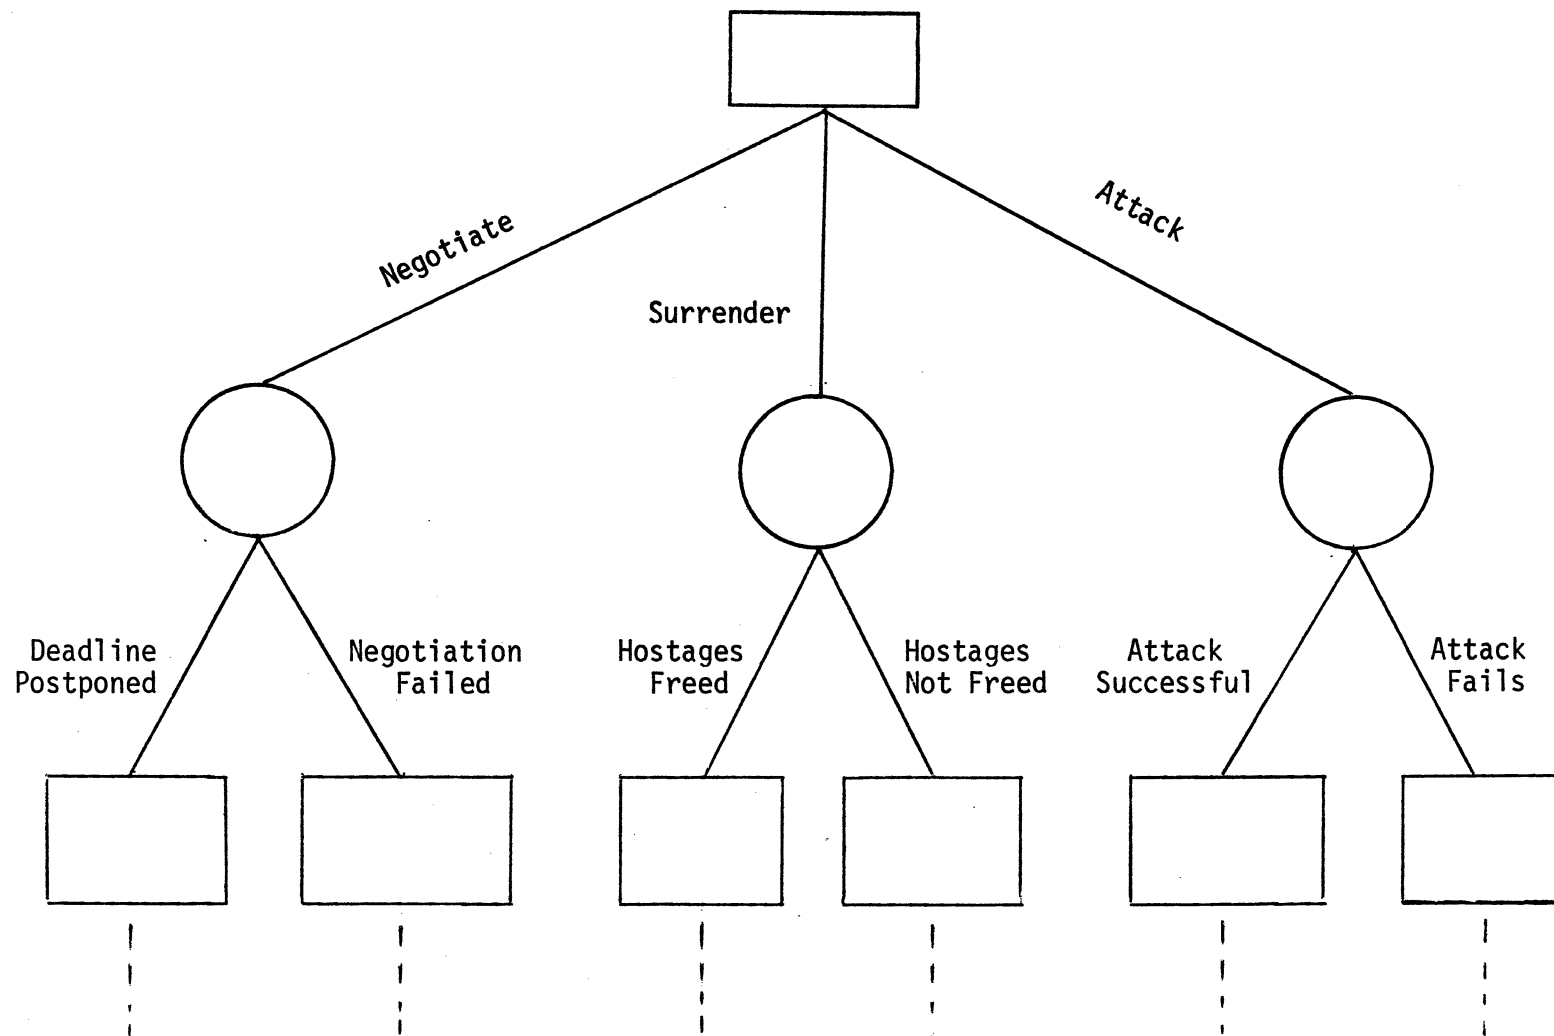


Figure 1-1. Decision-Tree Representation of Terrorist Attack Problem.

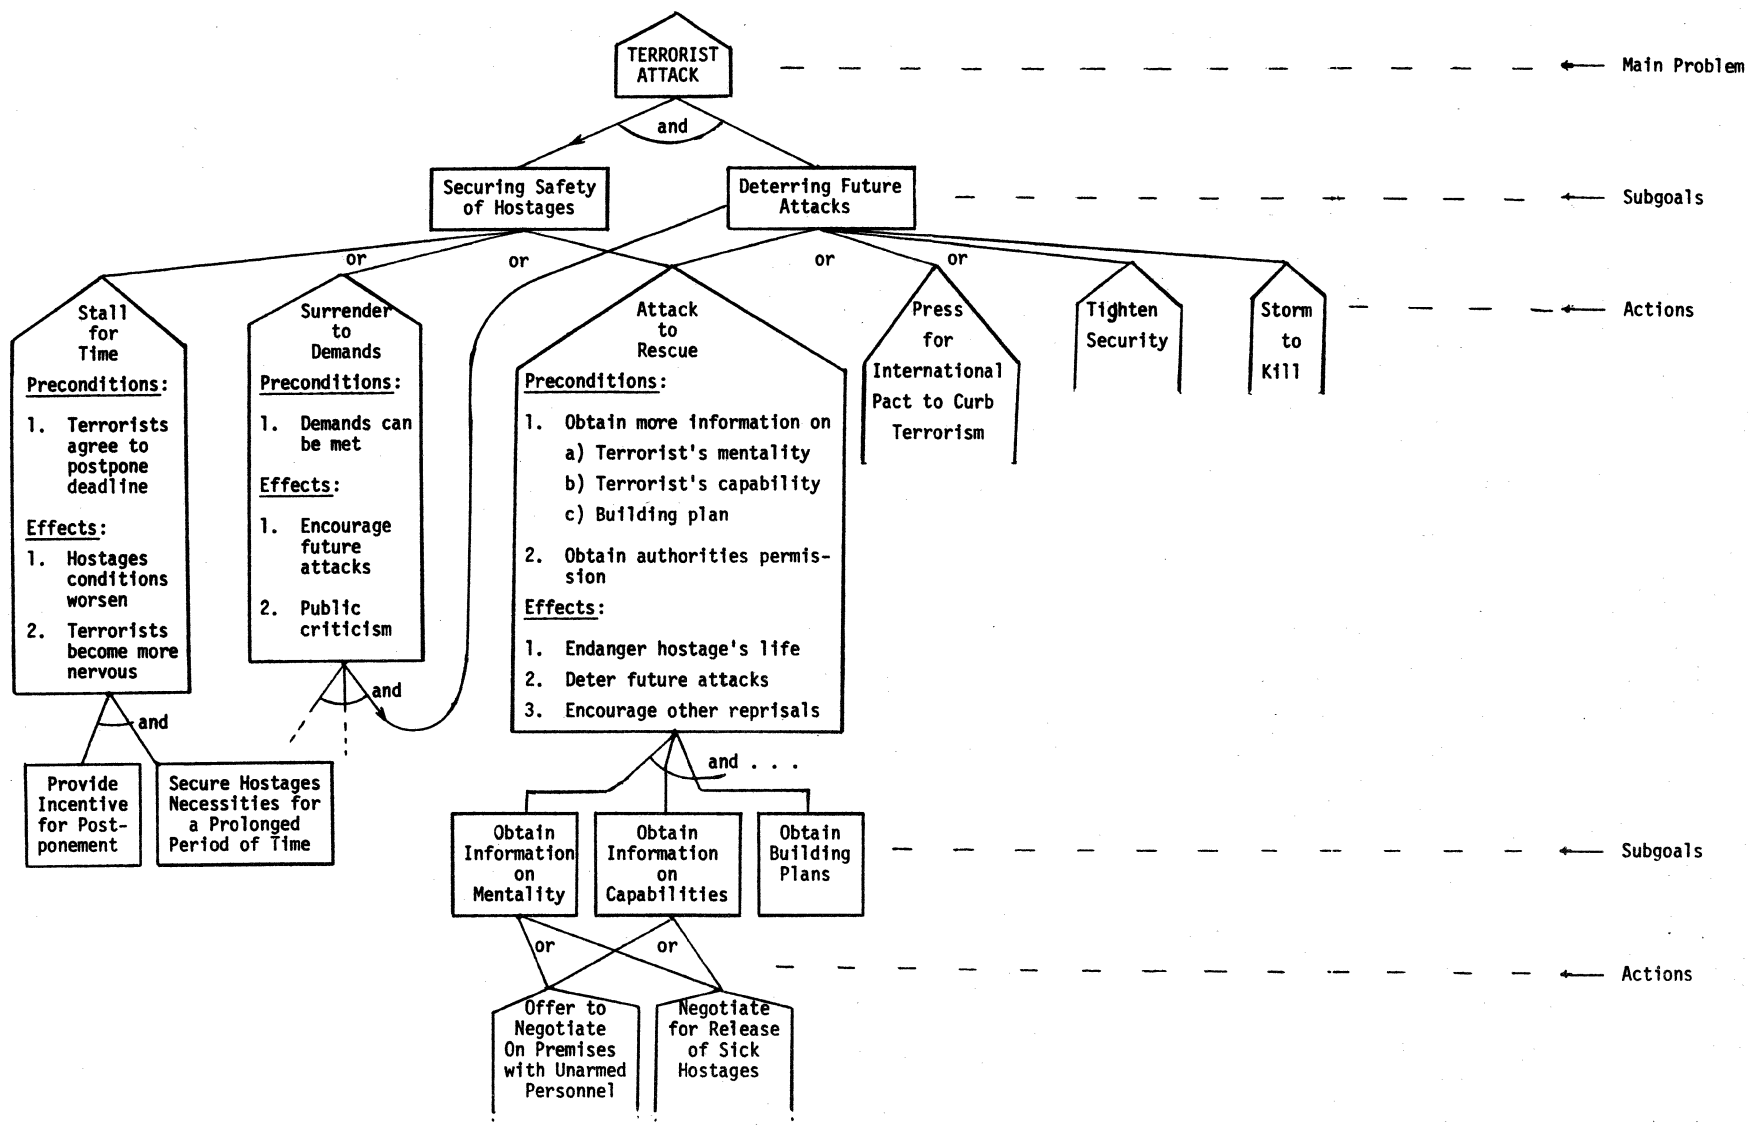


Figure 1-2. Goal-Directed Representation of Terrorist Attack Problem.

for postponement'.

Some arcs of the graph may point back toward higher levels in the structure (e.g., one of the effects of 'surrender to demands' is found to be 'encourage future attacks' which generates, since it is an adverse effect, a subgoal of eliminating this effect, namely the subgoal 'detering future attacks' which is already listed in the first level).

The main advantage of this structure is that the intent of each action is spelled out explicitly prior to naming the action. The analysis proceeds from the ends toward the means which encourages the user to discover novel alternatives. For example, the alternative 'negotiate for release of sick hostages' only came to mind after drawing the subgoals 'obtain information on terrorist's mentality and capabilities'. Clearly, similar goals may also implicitly influence one's thoughts during a decision-tree elicitation. For example, the alternative 'negotiate' in Figure 1-1 may have been identified for the purpose of obtaining additional information about the terrorist mentality. However, not having such purposes spelled out formally may cause the user to neglect exploring a large set of alternatives which can make up a workable solution plan.

In formal problem-solving, such as theorem proving or robot planning, problems are said to be solved when a sequence of operators is found which removes all differences between the desired and the current state. In real-life problems, such as the terrorist problem above, issues seldom get 'resolved'. They are, at best, alleviated, or controlled within acceptable ranges. For example, one has no guarantee that meeting the terrorists' demands would result in the hostages' safety. The latter is only a plausible expectation. Similarly, one cannot be sure of the degree to which storming the building would deter future terrorist attacks. Such estimates must be assessed using educated guesses and quantified using a formal structure. One of the tasks accomplished during

this research project has been to equip the goal-directed structure of Figure 1-2 with capabilities of handling uncertain and value-driven relationships. The descriptions of the actions should contain information on the degree to which each of the preconditions contributes to the realization of each subgoal. For example, the action 'attack to rescue' should specify how obtaining the various information would influence the hostages' safety during the attack. Similarly, a value judgment must be attached to each of the mentioned subgoals in order to determine both the relative merit of candidate solution plans and the direction of future elicitation queries.

It is interesting to note that the structure depicted in Figure 1-2 could also constitute a 'frame' (or template) for representing the generic aspects of terrorist-attack problems. Once elicited in detail, such a structure could be pre-stored as an 'expert' on terrorist confrontations, and be consulted when a particular crisis develops. The advantage of pre-storing the 'frame' is that during the crisis, only the problem-specific parameters need be explored in detail. On the basis of these parameters, the program could also suggest pre-stored contingency plans for consideration by the user, provide explanation for its suggestions, and, to some degree, be able to understand queries posed to it in English.

2.0. ORGANIZATIONAL DESIGN FOR A GOAL-DIRECTED DECISION SUPPORT SYSTEM

2.1. Model Structure

The decision support system aids the decision maker through computer interaction by structuring his decision problem and eliciting informational values associated with his perception of the problem environment. The system is 'goal-directed' (Pearl, 1978) in the sense that, unlike the decision-tree structure, the information supplied by the user is evoked and guided by explicit references to the objectives which need be accomplished and the means by which they may be attained. Figure 2-1 shows the model structure and the components relevant to attaining the goal. The main components are the following:

- (1) Goal - the major objective of the decision maker.
- (2) Subgoals - the goal 'dimensions', 'attributes', or detailed items that combine to form the overall goal.
- (3) Actions - the major action strategies that are open to the decision maker for accomplishing a particular subgoal.
- (4) Modes - the possible implementation methods of performing each action.
- (5) Preconditions - those states of nature or the environment that must exist before a particular (action) mode can be implemented effectively.

The structure should be thought of as a 'tree'. Thus, the goal is divided into many subgoals, each subgoal has a number of possible actions that could accomplish it, each action has a number of ways (modes) it can be performed, and each mode has a number of preconditions that must be completed. Once the preconditions are specified, they lead directly to new subgoals. That is, the subgoal of completing the specific precondition that allows actions to be taken, etc. If the realization of a precondition is beyond the direct control of the user and is, instead, perceived to depend on externally controlled eventualities, then that precondition is treated as an uncertain event node quantified by

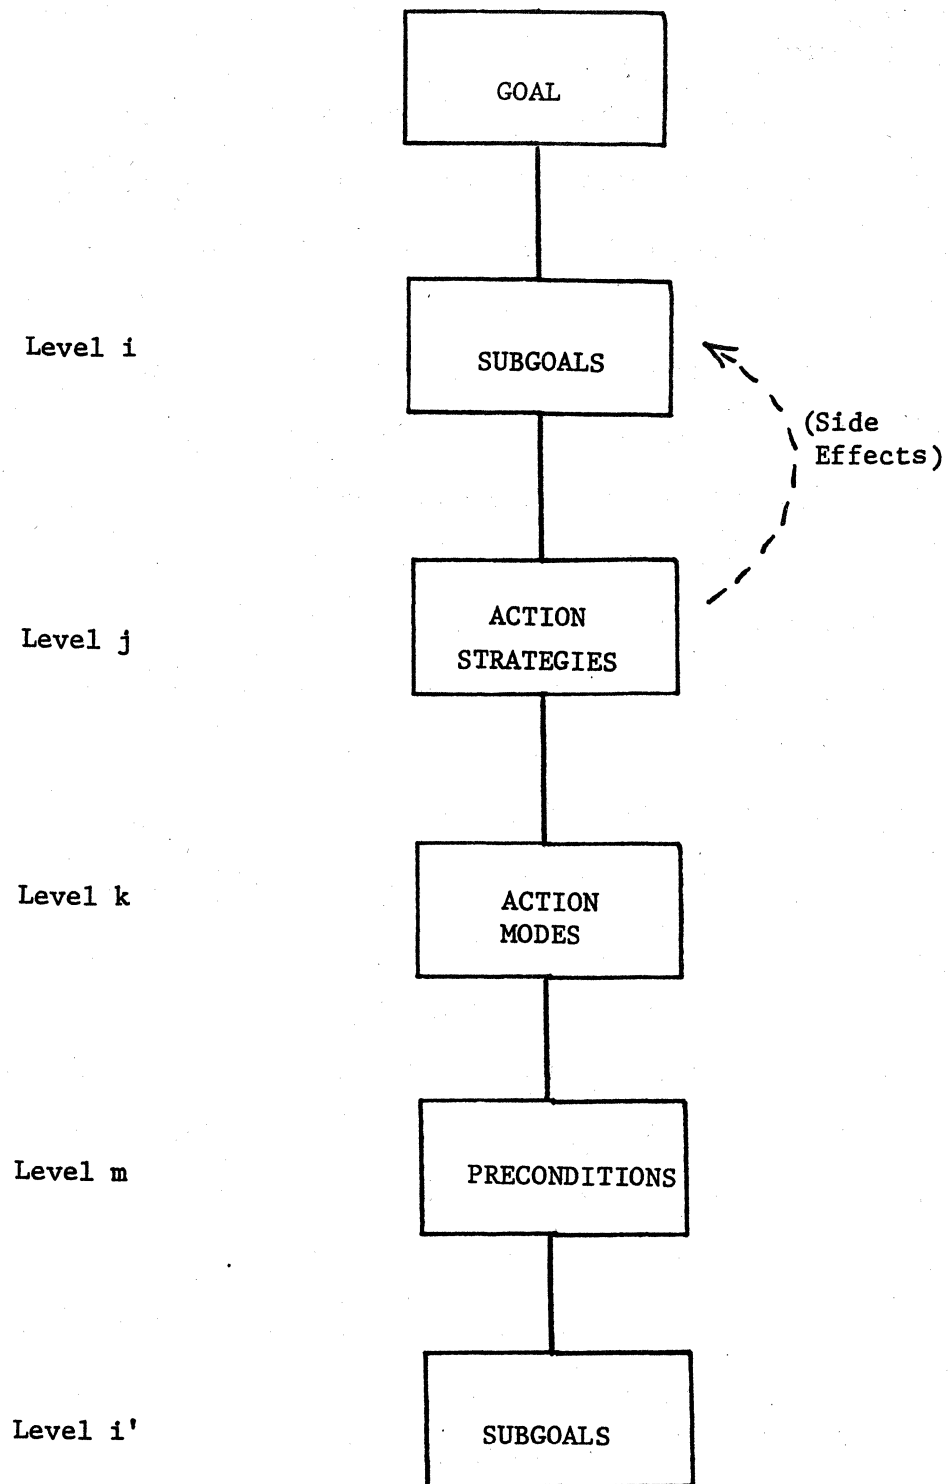


Figure 2-1. Model Structure

likelihood estimates. The structure can then be repeated.

Cross-relationships can also exist in the structure. For example, it is possible for one action to do 'double duty'. A single action may have a beneficial or adverse effect on a subgoal to which it is not directly connected. (Thus, the structure is more properly called a 'graph' rather than a tree.) These cross-relationships must be elicited from the decision maker and will affect the recommendations given by the system.

The following sections outline each of the above components in detail including the required information values and algorithms for aggregating them. The structure 'levels' have been indexed for purposes of referencing the various values and parameters.

2.2. The Major Goal

The decision maker will usually state the major goal in terms of a particular state of affairs that he desires. The overall goal has an associated value G ($0 \leq G \leq 1$) that captures the decision maker's accomplishment of the goal to different levels of satisfaction. It is the task of the decision support system to maximize this value. The value of G need not be 0 at the beginning of the elicitation session, that is, a portion of the goal may already be attained. The value $G=0$ will reflect a pessimistic state of affairs and $G=1$ an optimistic situation, conveniently chosen by the user for references purposes.

2.3. Subgoals

With the major goal stated, it is now necessary to explore in detail the goal dimensions, or 'subgoals'. The subgoals are simply the components that combine to form the overall goal. The subgoals may either be desired dimensions or adverse dimensions (hopes and concerns). Adverse dimensions are those whose elimination supports goal attainment. The subgoals should completely describe the major goal in such a way that if all desired dimensions were fulfilled to

their utmost level and all adverse dimensions were reduced to their lowest possible level, the goal would be achieved completely.

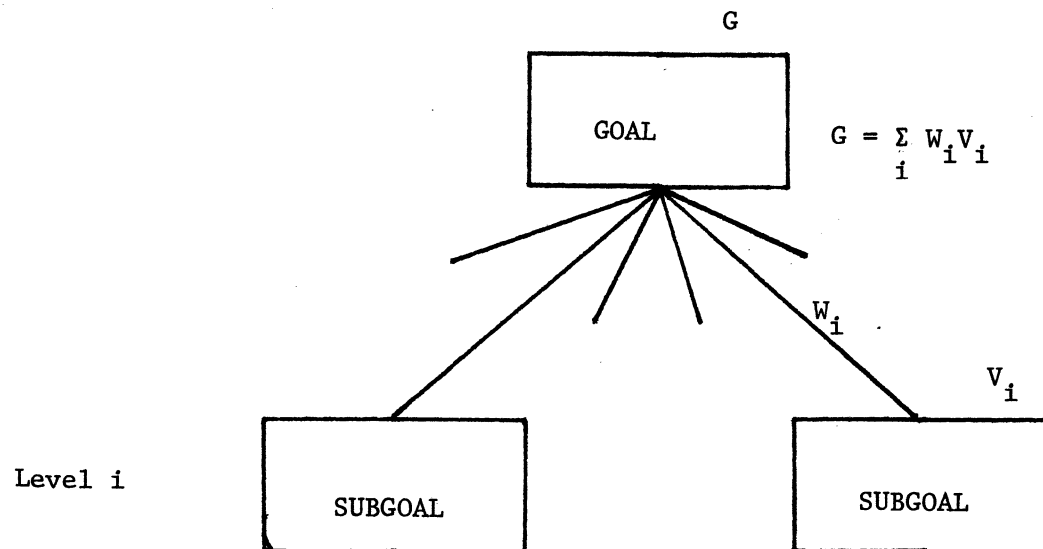
Figure 2-2 shows the structure of the major goal in terms of its subgoals. The relation between the major goal and the subgoals is represented by two numbers associated with each subgoal: value and weight. The value V_i ($0 \leq V_i \leq 1$) of subgoal i is the degree to which it has been achieved. (This parallels to the value G for the goal.) The weight W_i ($0 \leq W_i \leq 1$) for subgoal i is a measure of its importance relative to the other subgoals. The decision maker is instructed to estimate the degree to which the accomplishment of the subgoal adds to the satisfaction of the major goal. The weights are constrained to sum to 1 ($\sum_i W_i = 1$).

The goal value G is obtained from the subgoal values and weights by a linear combination ($G = \sum_i W_i V_i$). Thus, the subgoal structure corresponds to a linear multi-attribute model.

2.4. Actions

After the list of specific subgoals has been established, the decision support system begins elicitation of actions. For each subgoal, the decision maker is asked to think of possible actions that would cause the desired subgoals to be achieved (attained) or the adverse subgoals to be eliminated (reduced). More than one action may be listed. However, each action should have the capacity, by itself, to affect the subgoal, and they should be mutually exclusive.

Actions are divided into two levels: action 'strategies' and action 'modes'. An action strategy is a statement of a plan or a short description of what is to be done. An action mode is a more detailed specification of the method for accomplishing the action strategy. An action mode may be thought of as simply a 'subaction'. Figure 2-3 shows a subgoal with a number of supporting action strategies. Each action strategy has an associated 'effectiveness'



G Goal Value $0 \leq G \leq 1$

W_i Subgoal Weight (importance) $0 < W_i < 1$ $\sum_i W_i = 1$

V_i Subgoal Value (Level of Attainment) $0 \leq V_i \leq 1$

Figure 2-2. Goal Structure

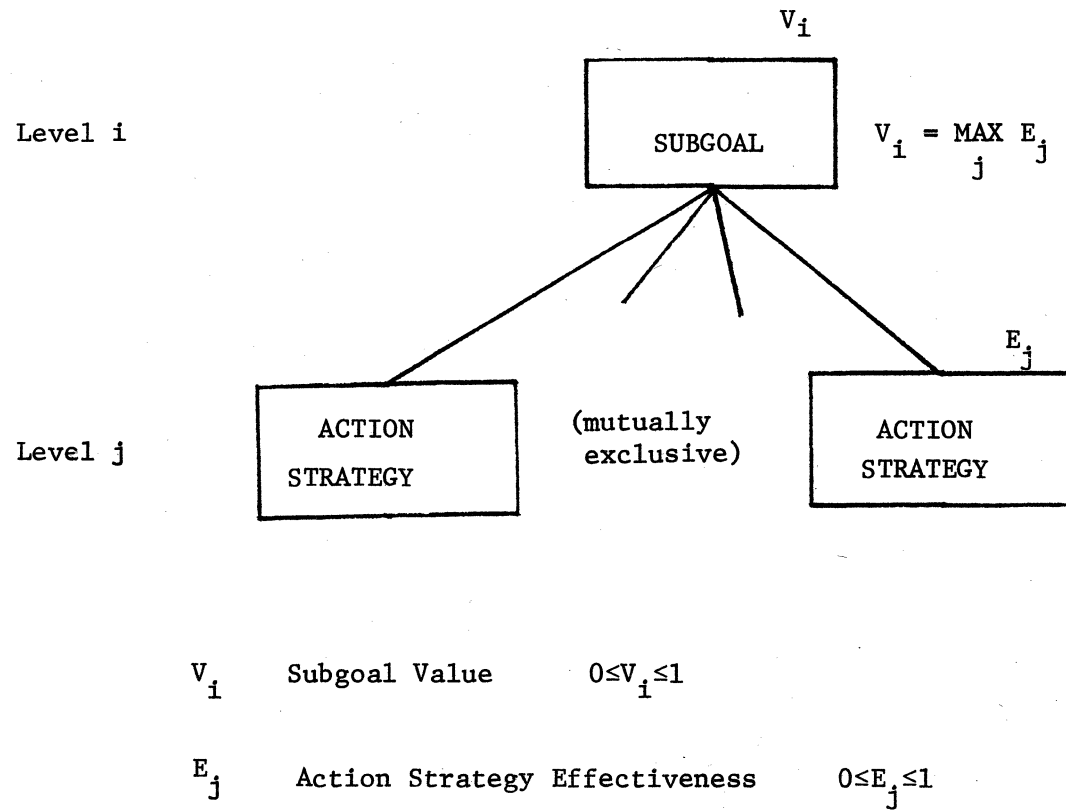


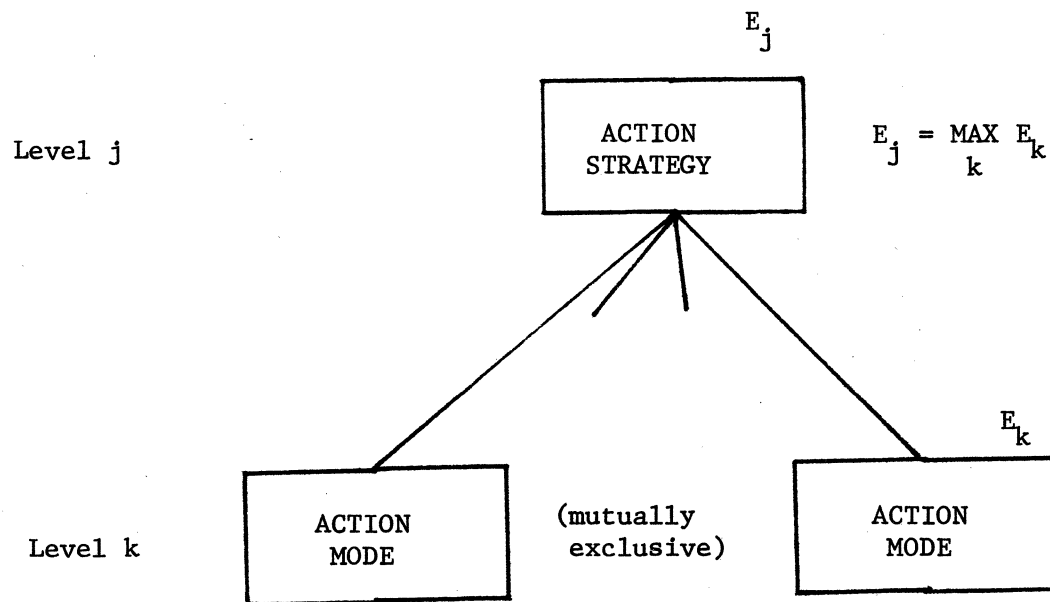
Figure 2-3. Subgoal Structure

measure E_j ($0 \leq E_j \leq 1$), which indicates the level of subgoal attainment to be expected if action strategy j were executed. The value V_i of the subgoal is the maximum of the supporting action strategies ($V_i = \text{Max}_j E_j$) representing the option of selecting that action strategy which produces the highest subgoal attainment.

Figure 2-4 shows the action strategy structure composed of supporting action modes. The action mode effectiveness E_k ($0 \leq E_k \leq 1$) is the amount that the corresponding mode affects the success of the action strategy. Again, the (locally) best action mode is the one with the highest value ($E_j = \text{MAX}_k E_k$). It may not always be necessary to subdivide every action strategy into action modes. If the action strategy can only be implemented in one way, the mode level is not needed. The benefit of characterizing actions by a two-level structure lies in the fact that many properties of the various modes (e.g., preconditions) would be identical to all modes of a particular strategy. This would enable us to store these common sets of properties in the description of the parent strategy, thus saving the storage and elicitation time otherwise consumed by duplication.

2.5. Preconditions

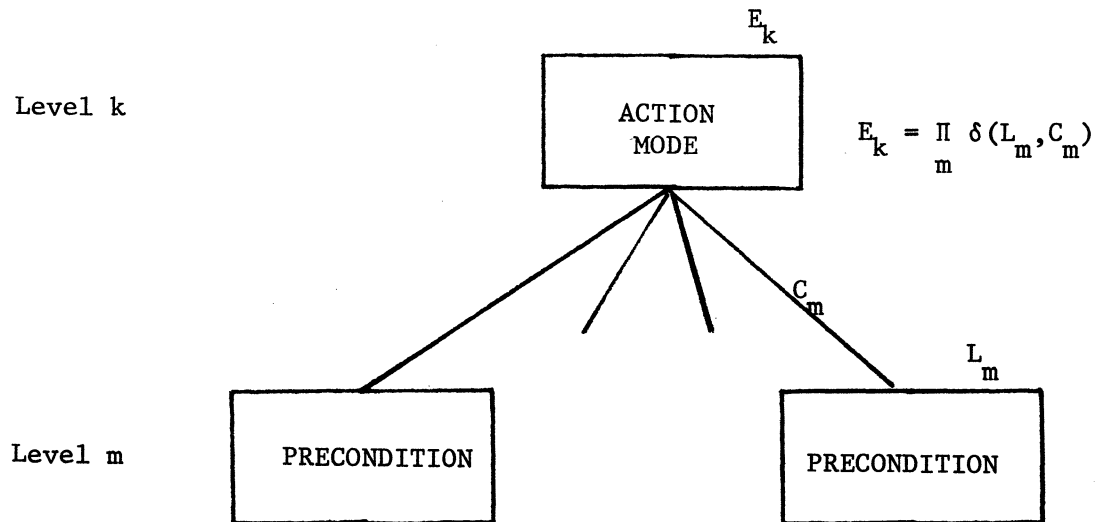
A 'precondition' is a state of nature or the environment that must exist before an action mode (or strategy) can be implemented effectively. Precondition satisfaction need not be an 'all or nothing' requirement. The effectiveness E_k of an action mode may be proportional in some way to the level of attainment (completion) of the precondition state. Each precondition is characterized by a measure of completion L_m ($0 \leq L_m \leq 1$) (Figure 2-5). This measure is used to calculate the effectiveness of the corresponding action mode. Since all of the preconditions should be completed before the effectiveness of the action mode can be fully realized, the mode effectiveness should be based on the normalized product of the supporting precondition completion levels.



E_j Action Strategy Effectiveness $0 \leq E_j \leq 1$

E_k Action Mode Effectiveness $0 \leq E_k \leq 1$

Figure 2-4. Action Strategy Structure



E_k Action Mode Effectiveness $0 \leq E_k \leq 1$

L_m Precondition Attainment (Completion) Level $0 \leq L_m \leq 1$

C_m Criticality Threshold $0 \leq C_m \leq 1$

$\delta(L_m, C_m)$ Criticality Function $0 \leq \delta \leq 1$


Figure 2-5. Action Mode Structure

A second measure associated with each precondition is its 'criticality' threshold C_m ($0 \leq C_m \leq 1$). The criticality is a threshold on the completion level of the precondition below which the effectiveness of the corresponding action mode is nullified. A threshold of 0 means that the action mode can be executed (to some degree of effectiveness) even if the precondition exists at its minimum level of attainment. A threshold of 1 means that the mode cannot be implemented (or has 0 effectiveness) unless the precondition is fully satisfied.

There can be a number of possible relationships between the precondition completion level and the effectiveness of the supported mode. If this relationship is called δ , the 'criticality function', it may be expressed in terms of a graph within the interval $[0,1]$. Figure 2-6 shows the chosen δ functions. It expresses the relationship between the precondition completion level L_m , its criticality C_m , and the degree to which it enables the action mode to be executed. Once the function is determined, the overall effectiveness of the action mode can be obtained by taking the normalized product of the criticality functions of each of the connecting preconditions as shown in Figure 2-5.

There are two types of preconditions: controllable and uncontrollable. A 'controllable' precondition means that the level of its completion is either known or can be controlled directly. An 'uncontrollable' precondition is one whose current level of attainment is both uncertain and not directly adjustable. For example, in the context of business decision making, the user may consider the action mode 'lower prices by 10 percent' as a potential action for achieving the subgoal 'capture a larger share of the market'. The effectiveness of this action depends (among other factors) on the variables 'competitor's prices' and 'buyers' price awareness'. The latter may be controlled via advertisement while the former must be treated as an uncertain variable not subject to one's direct control or scrutiny. A more detailed description of structuring uncertain events

$$E = \delta(L, C)$$



$$\delta = \begin{cases} 0 & \text{if } L < C \\ \frac{L-C}{1-C} & \text{if } L \geq C \end{cases}$$

L = Precondition Completion Level

C = Criticality Threshold


Figure 2-6. Criticality Function

is given in Chapter 5.

At the point below a precondition, the decision structure repeats with new subgoals (Figure 2-7). The new subgoal refers to the possible methods which can be used to satisfy the corresponding precondition. This may involve another entire structure including action strategies, action modes, further preconditions, etc.

2.6. Side Effects of Actions on Subgoals

It often happens that the execution of a particular action has a beneficial or adverse effect on a subgoal to which it is not directly connected. The decision system asks if any such effects are present and elicits an impact measure for the relationship. Figure 2-8 shows the side effects of an action k on a subgoal i . The 'remote' impact I_j^r must be elicited and is considered to be the amount ($0 < I_j^r \leq 1$) that the action, if implemented, will increase or decrease the level of attainment of the affected subgoal. Assume that the subgoal attainment level has already reached a value V_i (between 0 and 1) as a result of action j directly connected to it. The remote action k will modify V_i in one of two different ways. If the remote action has a beneficial effect of degree I_i^k , V_i will increase to $V_i' = V_i + (1 - V_i) I_i^k$, thus mapping the current value into the range ($V_i \leq V_i' \leq 1$).



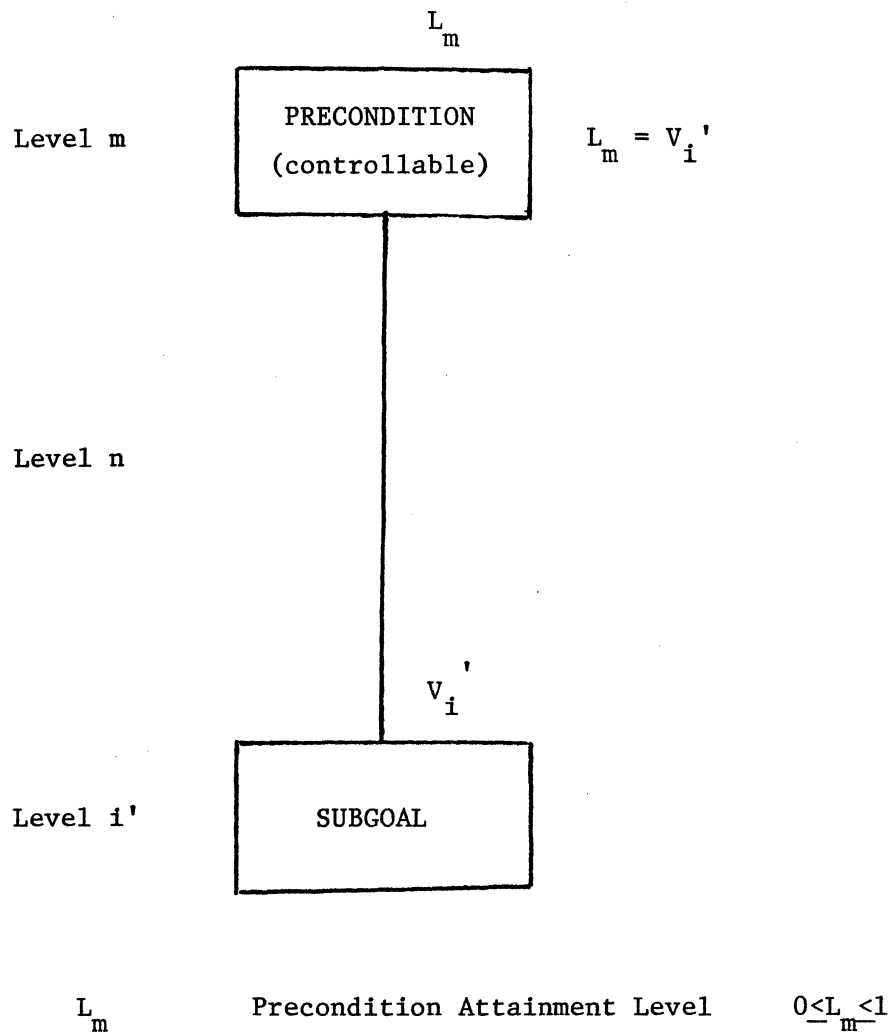


Figure 2-7. Precondition Structure

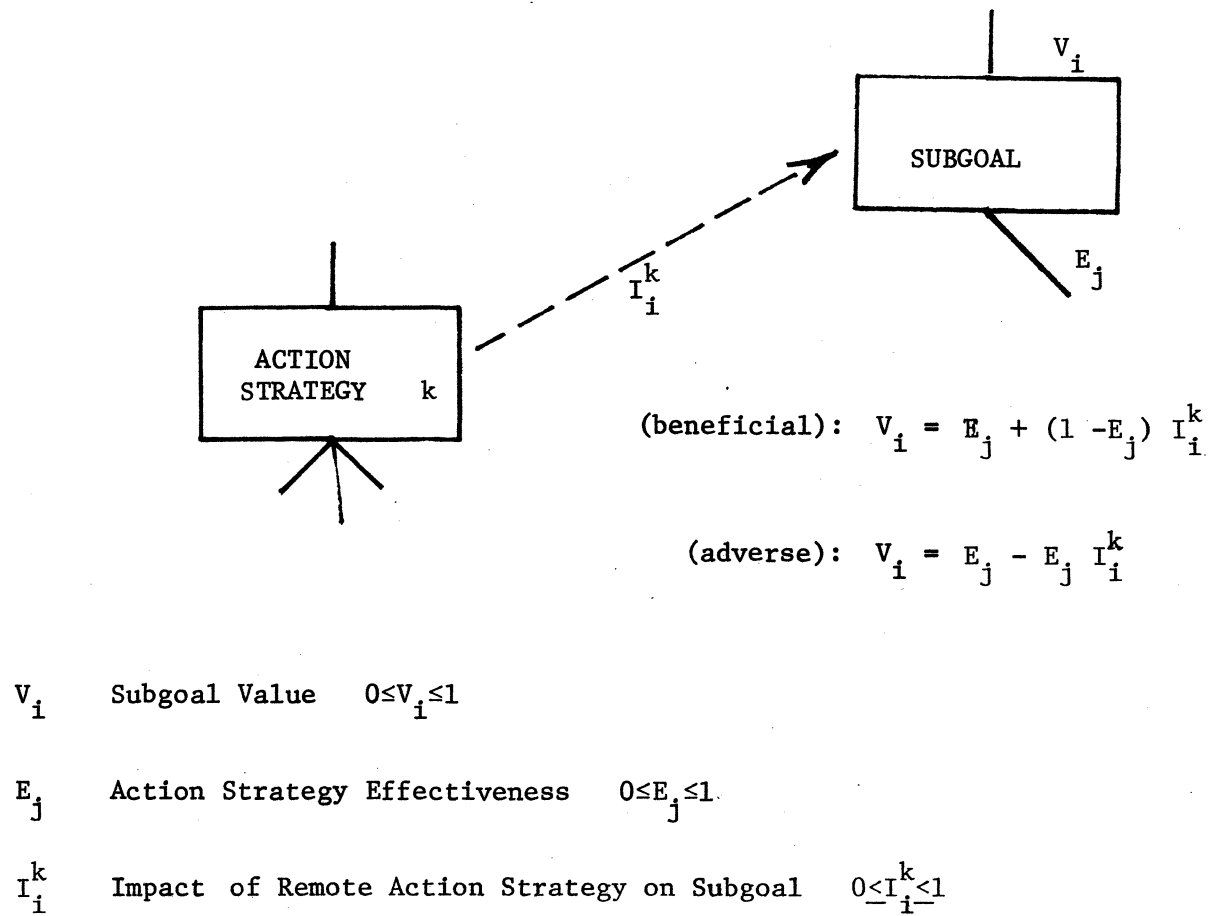



Figure 2-8. Side Effects of Actions on Subgoals

In the extremes, a value of $I_i^k = 0$ will not change V_i ; a value of $I_i^k = 1$ will cause the subgoal to reach its full attainment level, $V_i = 1$.

If the remote action has an adverse effect on the subgoal, it will lower the subgoal attainment level V_i to $V_i' = V_i - V_i I_i^k$.



A remote adverse impact of degree $I_i^k = 0$ will not change the current subgoal value V_i , whereas $I_i^k = 1$ will reduce V_i to 0.

When several actions affect the same subgoal remotely, their cumulative impact is computed by applying each individual action in succession (in any order).

3.0. OPTIMIZATION

3.1. Motivation

The main result required from a goal-directed decision-support system is the identification of the optimal feasible action plan whose implementation results in the highest level of attainment for the major goal; a feasible action plan is a sequence of action combinations, one for each level of the goal-directed graph. The identification process requires assessment of all feasible action combinations. Had the assessment of a plausible action combination been derivable from a linear aggregation of individual actions composing each action combination, the complexity of the optimization effort of a goal-directed graph with n actions would have been of the same order as n . However, since side-effects are combined multiplicatively, the effect of each individual action depends on the action combination in which it appears. Therefore, rather than independent actions, feasible action combinations must be assessed in their entirety resulting in a much higher complexity. For instance, for a graph with n actions clustered in k classes of m mutually exclusive actions $(a_1^1 \dots a_m^1) \dots (a_1^k \dots a_m^k)$, where $n = mk$, the number of plausible action combinations subject to optimization is $(m)^k$. If the value of G had to be calculated separately for each action combination, the evaluation effort would be prohibitively long.

During the use of the system, the optimization procedure is executed frequently. Each time a recommendation is requested the optimization must be performed. Additionally, it should be invoked each time a node is considered for expansion. As the number of actions in a goal-directed graph increases, the time required for optimization increases rapidly. Even during early stages of the elicitation phase, the optimization process may take a major portion of the processing time. In a larger goal-directed graph, the optimization process may take minutes. Therefore, employment of the exact and exhaustive optimization

process will be unacceptable. A local optimization procedure has been proposed and simulated to investigate the effect of different conditions on the accuracy of the results. An optimization procedure based on first order approximation of side effects has also been developed, studied, and finally chosen for adaptation. These subjects are summarized in the following section leaving the more detailed elaboration to Appendix A.

3.2. Local Optimization

The local optimization procedure consists of iteratively selecting a single action from each (mutually exclusive) action set. The 'action set' is the collection of actions supporting a common subgoal. In each stage of the local optimization procedure, a criterion is optimized with respect to only one action set, while the other sets are kept at a fixed level. While the complexity of exhaustive optimization increases exponentially with n , the complexity of each iteration cycle of the local optimization is of the same order as n .

Due to the multiplicative nature of the side-effect combination rule, the ultimate effect of an individual action on the major goal cannot be assessed in isolation. The entire action plan, composed of a single action from each action set, must be known for an exact valuation. Therefore, prior to the selection of individual actions, it is necessary to represent each action set by a 'representative' action which has roughly the same overall effects as that action which would be ultimately chosen by a global optimization (Figure A-1). The representative action may not necessarily be a member of the given set. Rather it can be 'virtual', producing a mixture of direct and side effects typical of those produced by the actual set.

The collection of all representative actions constitutes the 'initial action combination'. In the first cycle of the optimization procedure, the

and low inflation (Figure 3-1). Among other actions, 'subsidizing businesses' and 'reducing business and inventory taxes' are proposed as alternative means to cause an increase in employment, while 'increasing interest rate' and 'reducing public expenditures' are viewed as ways to lower inflation. According to this model, although 'increasing interest rate' enhances the level of attainment of 'low inflation', its implementation has an adverse effect on the other subgoal 'high employment'. Here, the degree of the adverse effect is more naturally expressed by the percentage reduction in the effectiveness of the subgoal. In the example of Figure 3-1, the side effect produced by 'increasing interest rate' amounts to a reduction of 30 percent in the level that the subgoal 'high employment' would otherwise attain.*

According to this model, if the degree of effectiveness of the action selected under the first subgoal is E_1 and the adverse impact of the action selected under the second subgoal with respect to the first subgoal is E_2^1 (Figure 3-2), then the level of attainment of the first subgoal (V_1) will be:

$$V_1 = E_1 (1 - E_2^1)$$

Since the optimization complexity in such multiplicative models is very high, the possibility of reducing it by employing a first-order additive approximation has been investigated.

In Figure 3-3, the line L represents the multiplicative model $V_1 = E_1 (1 - E_2^1)$ while the dotted line represents the additive approximation. It is parallel to the line $V_1 = E_1$ drawn from a representative point (\bar{E}_1) on L. The additive model is expressed by:

$$V_1 = E_1 - \bar{E}_1 E_2^1$$

which does not contain the product term $E_1 E_2^1$.

* The level of 'high employment' (as well as any other subgoal) is not measured on the normal physical scale of population proportions. Rather, it reflects the subjective scale of gravity perceived by the user.

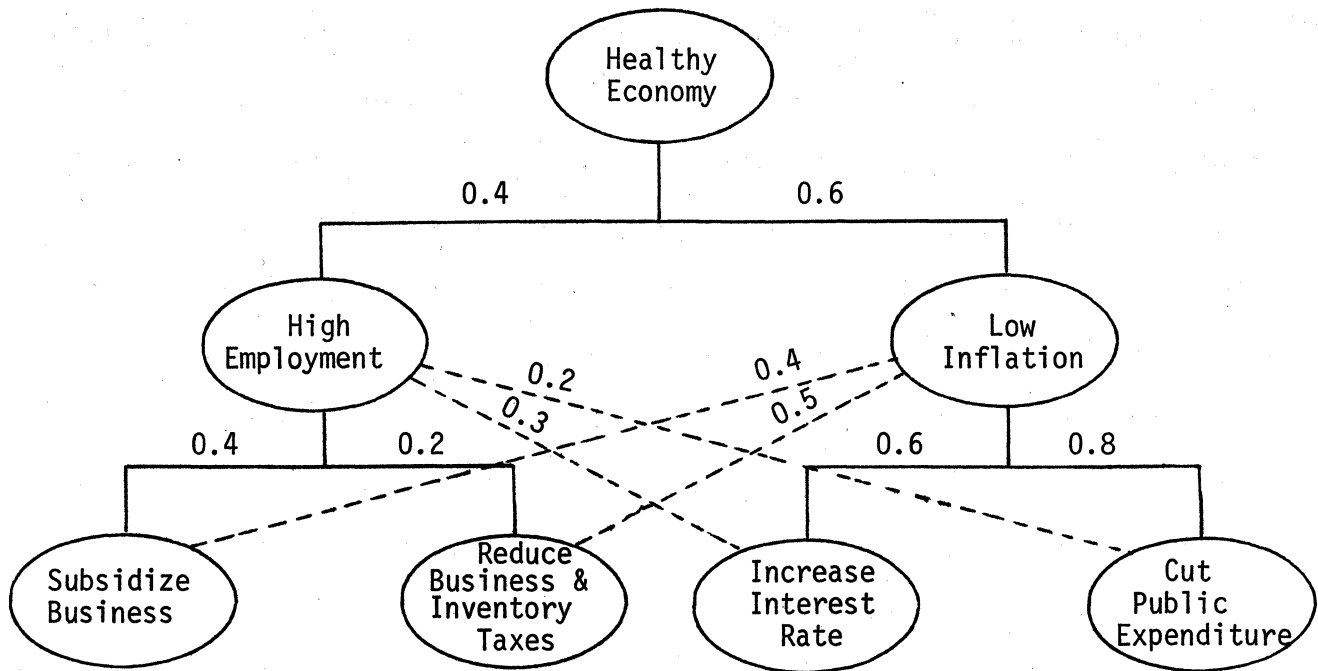


Figure 3-1. A Sample Goal-Directed Graph




Figure 3-2. Structure of the Sample Goal-Directed Graph




Figure 3-3. First-Order Approximation Model for Side Effects

virtual actions are replaced, one at a time, by real actions taken from the corresponding sets. The process can be iterated any desired number of times. The resultant action combination may not always converge to the globally optimum and the degree of suboptimality depends on the choice of the initial action combination.

A simulation program has been written to identify effective choices for initial action combination. The results of that simulation are presented in Appendix A. It was found that by equipping each representative action with a direct effect equal to the maximum of the actual direct effects in its corresponding set, a close to optimal solution plan was obtained after a single iteration. In 143 out of 216 randomly selected cases, the globally optimal plan was found. The accuracy of the optimization procedure may be enhanced by applying an increasing number of iterations, however, as in the most nonlinear problems of this kind, there is no guarantee that the process will ultimately converge.

The simulation studies were conducted on a two-level graph with 16 terminal nodes. Not having the assurance that the quality of this procedure would remain consistently acceptable for graphs of much larger dimensions, an alternative procedure was chosen, based on a first-order approximation of side effects.

3.3. Optimization Based on First-Order Approximation of Side Effects

The complexity required for global optimization would be drastically reduced if the impact of side effects of actions on the degree of attainment of other subgoals would combine additively. Our choice of the multiplicative model instead was based on the belief that people perceive the side effect of an action to depend on the degree of effectiveness of other actions taken simultaneously. For example, the major goal of realizing a healthy economy may be perceived, in a particular case, to be attainable by maintaining high employment

The accuracy of the model depends on the selection of the representative point \bar{E}_1 , and would increase when \bar{E}_1 is chosen close to the effectiveness of the action which will be ultimately chosen.

Since the chosen action remains unknown at the time the additive model is constructed, we chose \bar{E}_1 to be the weighted average of the effectiveness of the actions under the given subgoal. Since the likelihood of eventually choosing an action increases with its effectiveness, we chose the normalized effectiveness as the weight associated with this action.

The accuracy of the model also depends on the magnitude of the side effects. For sufficiently small side effects, the model provides an adequate representation. If E_2^1 and E_3^1 are side effects of actions emerging from other subgoals, then according to the exact model:

$$\begin{aligned} V_1 &= E_1 (1 - E_2^1) (1 - E_3^1) \\ &= E_1 (1 - E_2^1 - E_3^1 + E_2^1 E_3^1) \end{aligned}$$

With E_2^1 and E_3^1 sufficiently small we have:

$$V_1 \approx E_1 (1 - E_2^1 - E_3^1)$$

which can be linearized by $V_1 \approx E_1 - \bar{E}_1 (E_2^1 + E_3^1)$.

The additive model results in a substantial reduction of the required optimization effort since, for a single level graph, the optimal action plan can be identified in one single iteration. However, the impact of this approximation on the quality of the plan found has not been tested yet.

4.0. DIALOGUE MANAGEMENT

One of the main advantages that computerized decision-support systems offer over manual methods is the ability to compute (in real-time) which areas of the problem graph deserve further exploration and guide the dialogue in such a way that at any stage the user would focus attention on the most crucial issues. The procedure for identifying the crucial issues and presenting them for further analysis is denoted by dialogue management.

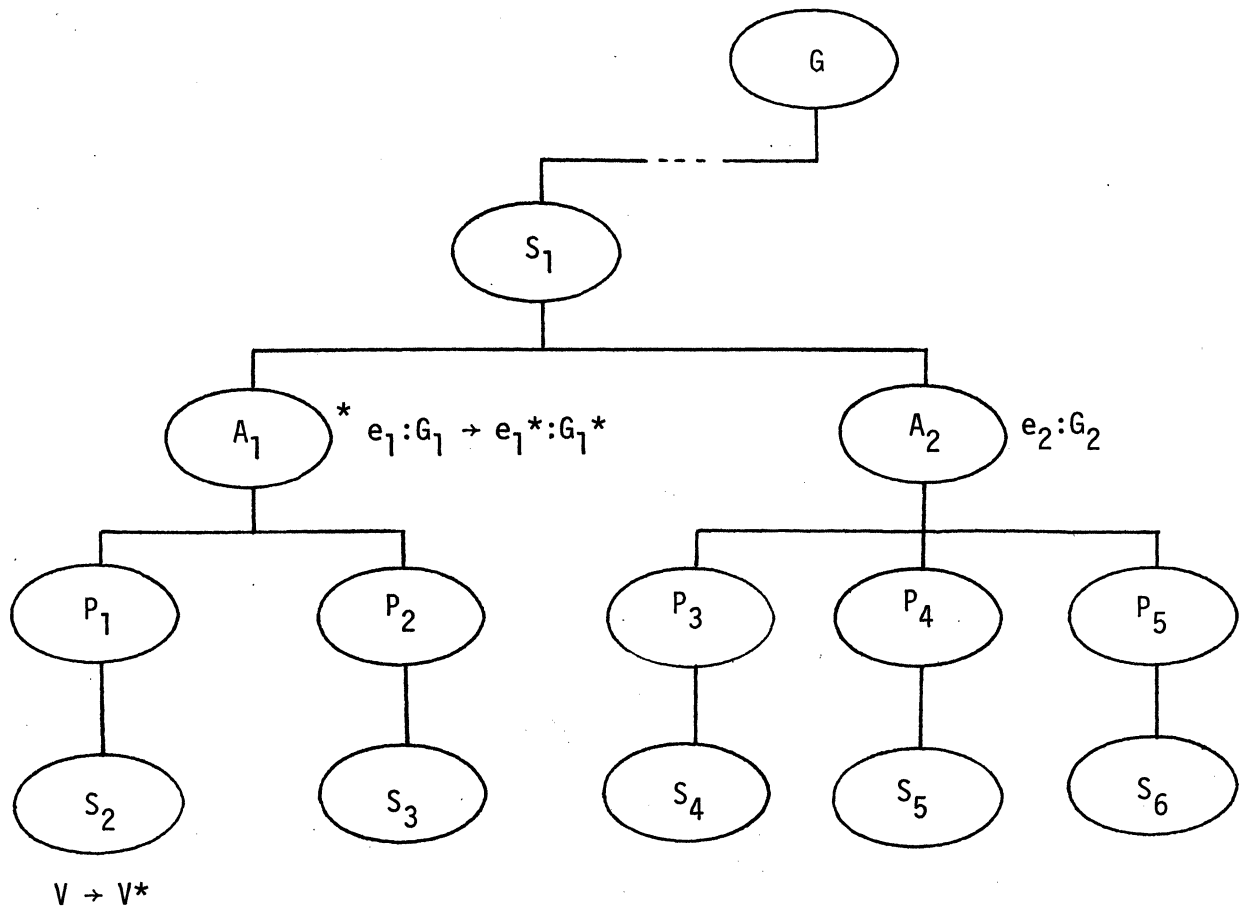
4.1. Methodology

The purpose of constructing a goal-directed graph is to provide a means for estimating the relative value of alternative ways of accomplishing the major goal, while providing aiding in generation of those alternatives. The degree of attainment of the major goal provided by each alternative action plan (strategy) is computed by assessing values of different graph elements and then rolling these values back through the graph. The rollback value that is obtained for degree of attainment of the major goal provided by each alternative action plan (strategy) is then a measure of the utility of that action plan (strategy). Beside the potential of aiding in generating richer sets of alternative plans, the time consuming effort of a goal-directed graph construction is justified by the belief that the utilities computed by rolling back the goal-directed graph are more valid estimates than those which would be produced through direct elicitation of the utilities of each alternative action plan. This difference in reliability, which constitutes the driving force behind the necessity to ask the user for more detailed information, was chosen, therefore, as a criterion for determining the future direction of the dialogue.

Suppose that at each subgoal node the decision maker is asked, prior to the construction of the complete goal-directed graph, to provide assessments of the maximum degree of attainment a specific subgoal would achieve if the complete

goal-directed graph were developed and solved. This assessment is expected to be generated on the basis of the knowledge and experience of the decision maker concerning similar issues. However, since he is uncertain about these degrees and is incapable of mentally manipulating the knowledge he possesses, the assessments could not be provided as an exact point estimate; they rather have to be viewed as random variables in the form of probability distributions. The value of resolving the uncertainty regarding the rollback value of each alternative action plan could then be computed using a method similar to that of calculating the value of information in standard decision-tree analysis (Merkhofer et al., 1977). Since constructing the complete goal-directed graph would be instrumental in resolving this uncertainty, the value computed for resolving uncertainty on the rollback value for a given subgoal can be interpreted as the expected value of further analysis along the path emerging from that subgoal. The value-of-analysis concept can then be used to guide the expansion of a goal-directed graph structure by selecting for expansion the node with the highest value-of-analysis.

A value-of-analysis measure can be assigned to each terminal subgoal node in a goal-directed graph structure which reflects the worth of focusing attention on the problem of planning the implementation of this subgoal. As an example, consider the incomplete goal-directed graph presented in Figure 4-1. This graph is considered incomplete since it is recognized by the user that the levels of attainment of the terminal subgoals (i.e., S_2 , S_3 , S_4 , and S_5) are uncertain, thus requiring further analysis. Therefore, the values assigned to these nodes are only approximations to the complete graph. The decision maker approximates the degrees of attainment of terminal subgoals by mentally manipulating alternative plans which may lead to its realization. Although the decision maker is uncertain about what the rollback values corresponding to the terminal subgoals



- v : the provisional value of the subgoal
- v^* : the true value of the subgoal would be resulted by complete analysis of the subgoal
- y : denotes already satisfied preconditions

Figure 4-1. An Incomplete Goal-Directed Graph

would become, were the graph expanded completely, he may be able to quantify his uncertainty by specifying a range or a distribution for that value. The terminal subgoals in the simplified graph that have the greatest need for expansion are those nodes for which the expected value of resolving the current uncertainty on true rollback value is the highest, thus, providing a criterion for recommending nodes for expansion.

The process of structuring a goal-directed graph proceeds as follows. First, an estimate for the uncertainty in the provisional value of each terminal subgoal node is elicited from the decision maker. Then, based on the elicited estimates, the value-of-analysis for each terminal subgoal is calculated. Finally, the subgoal with the highest value-of-analysis will be selected for expansion.

4.2. Calculation of Expected Value-of-Analysis

When the decision maker estimates the provisional value (V) of a specific subgoal (e.g, terminal subgoal S_2 in Figure 4-1), he obtains this value by crudely considering a set of relevant actions and their associated modes and preconditions emanating from this subgoal. Therefore, in absence of a deeper analysis of the subgoal, the decision maker believes that he will be able to implement the set of actions leading to a level of attainment roughly equal to V , for the subgoal. But, since this value may not coincide with that obtained by a more complete analysis, the potential increase in utility is given by:

$$\Delta G = G(V^*, V^*) - G(V, V^*)$$

and its expectation:

$$EVA = E_{V^*} [G(V^*, V^*) - G(V, V^*)]$$

where $E_{V^*}(\cdot)$ stands for the expected value with respect to the random variable V^* , and V and V^* designate the provisional value and the true value of the level of attainment of the subgoal, respectively; $G(V_1, V_2)$ denotes the level of attainment of the major goal assuming that actions are selected believing that V_1 is the degree of attainment of the subgoal, while in reality V_2 is its true value.

In Appendix B it will be shown that under the propagation rules defined in Chapter 2 the function $G(V_1, V_2)$ would be a piecewise linear function of both arguments, like the one depicted in Figure 4-2. Such functions are completely specified by a list (vector) of three parameters a, b, α which determine the level, break point, and slope of the function, respectively. The parameters (a, b, α) could be calculated recursively, top down, and be stored as a characteristic vector for each node of the graph.

In Appendix B, we also show that once the characteristic vector of a terminal subgoal is calculated, its EVA value is given by:

$$EVA = \alpha p |V + \Delta - m|$$

where: α is the characteristic slope of the subgoal,

p is the probability that V^* will deviate from V by an amount exceeding Δ ,

Δ is the change in the level of the subgoal required to affect a switch in action strategy, and

$$m = \begin{cases} E_{V^*} (V^* \mid V^* < V + \Delta) \text{ for a node in the current optimal} \\ \text{plan } (\Delta < 0) \\ E_{V^*} (V^* \mid V^* > V + \Delta) \text{ for a node not in the current optimal} \\ \text{plan } (\Delta > 0). \end{cases}$$

A graphic representation of the formula for EVA is given in Figures 4-3 and 4-4.




Figure 4-2. Parametric Representation of the Canonical Functions





Figure 4-3. Graphical Representation for the EVA Formula
 (Case 1: The Subgoal is Part of the Current Optimal Plan)





Figure 4-4. Graphical Representation for the EVA Formula
 (Case 2: The Subgoal is Not a Part of the Current Optimal Plan)

α and Δ are characteristics of the path leading to each terminal subgoal, whereas p and m can be calculated on the basis of distribution function of V^* which the user estimates for each subgoal.

The simplest way of eliciting the user's uncertainty regarding the true value of V^* is for the user to specify a range of $[V_{\max}, V_{\min}]$ where V^* is most likely to be found. If we assume that V^* is uniformly distributed over this range, we get:

$$V = \frac{V_{\max} + V_{\min}}{2}$$

and

$$p = \begin{cases} \frac{V + \Delta - V_{\min}}{V_{\max} - V_{\min}} & \text{for a node in the optimal plan} \\ \frac{V_{\max} - V - \Delta}{V_{\max} - V_{\min}} & \text{for a node not in the optimal plan} \end{cases}$$

4.3. An Alternative to the Expected Value-of-Analysis Criterion

In trying to utilize the EVA as a criterion for node selection we have discovered a practical difficulty stemming from a basic flaw in the method of calculating the EVA. After one or two levels of expansion the EVA of all terminal nodes became identically equal to zero. Since the impacts of subgoals are quantified by numbers smaller than one (weight factors) at each junction, the value of G becomes less sensitive to variations in the levels of subgoals remote from the root. Consequently, as the subgoal appears in deeper levels of the graph, its calculated α decreases and Δ , the deviation in V necessary to result in a decision switch, becomes substantially greater than 1. Eventually none of the terminal nodes will be capable of causing a decision switch by its own variation and so, as is shown in Figures 4-3 and 4-4, all terminal nodes will be assigned a zero EVA value.

Although in some rare cases the utility of analysis for all terminal subgoals would indeed be zero, and should be interpreted as a signal to stop further analysis, in the more common case, this phenomenon is an artifact of the approximations used to calculate EVA. The formula for the EVA measure has been derived under the assumption that the levels of attainment of all subgoals except the one in question remain fixed at their most likely value. However, since the value of other subgoals are also subject to uncertainty within their range of variability, the level of attainment of every subgoal should also be treated as a random variable. Consequently, Δ (the change in V required to result in a decision switch) must also be considered as a random variable rather than a constant. The correct expression of the expected value-of-analysis will then be:

$$EVA = E_{\Delta} E_{V^*} \Delta G$$

Thus, whenever the distribution of V^* overlaps that of $V + \Delta$, there would be a finite probability for a decision switch which, in turn, would result in a finite value for the utility of analysis.

Lacking sufficient data to estimate the distribution of Δ , we decided to drive the dialogue by an approximate criterion related to the likelihood of obtaining a decision switch rather than the exact value of EVA. The criterion chosen is α times the ratio between the potential effective range of variation and the amount of change required for a decision switch.

If the subgoal subject to analysis emerges from the presently most effective plan, a switch will occur if the value of the random variable V^* falls below the value of the random variable $V + \Delta$. The potential range of variation effective for a switch, in this case, is $V - V_{\min}$ which positively contributes to the occurrence of a switch. The likelihood of a switch is also inversely

affected by $|\Delta|$. Therefore, the likelihood of a switch will be increasing with:

$$LS_1 \triangleq \frac{\Delta \text{ Potential Effective Range of Variation}}{\text{Mean of Change Required for a Switch}} = \frac{V - V_{\min}}{|\Delta'|}$$

where Δ' is calculated, as explained in Appendix B, on the basis of the provisional values of all other subgoals.

In the second case, where the subgoal subject to analysis is not in the optimal plan, a switch will occur if the value of the random variable V^* exceeds the value of the random variable $V + \Delta$. Thus, the potential range of variation effective for a switch is $V_{\max} - V$ which is a positive factor for causing a switch. Here also the likelihood of a switch is inversely affected by Δ , and so, the likelihood of a switch should be measured by:

$$LS_2 \triangleq \frac{\Delta \text{ Potential Effective Range of Variation}}{\text{Mean of Change Required for a Switch}} = \frac{V_{\max} - V}{|\Delta'|}$$

α , of course, measures the sensitivity of G to variations in V^* once a decision switch occurs, and should, therefore, multiply LS to yield a criterion approximating EVA.

4.4. Dialogue Management Algorithm

The dialogue management algorithm finally chosen is based on the ratio criterion described in the previous section requesting the user to provide only V_{\min} and V_{\max} for every subgoal and taking the provisional value as the midpoint $(V_{\max} + V_{\min})/2$. Therefore:

$$LS_1 = \frac{V - V_{\min}}{|\Delta'|} = \frac{\text{Range of } V}{2|\Delta'|}$$

$$LS_2 = \frac{V_{\max} - V}{|\Delta'|} = \frac{\text{Range of } V}{2|\Delta'|} = LS_1$$

The algorithm is as follows:

- (1) Do the following for all subgoals (S) subject to expansion.
 - (a) Elicit the Range of $V(S)$, $[V_{\max}(S), V_{\min}(S)]$.
 - (b) Calculate $\Delta'(S_1)$ and $\alpha(S)$ assuming all other subgoals attaining the value $(V_{\max} + V_{\min})/2$ as described in the algorithm of Appendix B.
 - (c) Calculate $\frac{\text{Range of } V}{2 \Delta'(S)} \cdot \alpha(S)$.
- (2) Choose the subgoal with the largest criterion (c) for further analysis.

In order to maintain a reasonable continuity of attention we have chosen to treat the entire subgraph residing between a given subgoal and its descendent subgoals as a single atomic unit for dialogue management purposes. Thus, once a subgoal is selected for expansion, the entire subgraph supporting it (i.e., actions/modes/subgoals) becomes committed for analysis in a breadth-first fashion.

5.0. STRUCTURE OF UNCERTAIN EVENTS

5.1. Motivation

The primitive skeleton of the goal-directed graph structure, as described in Figure 2-1, is not capable of modeling preconditions where the level of attainment is both uncertain and not directly adjustable. For example (see Figure 5-1), in the context of long-range industrial planning, at a certain point during the planning process, the decision maker may realize the necessity of generating 'additional revenue'. The goal-directed system will guide the user by asking him to list alternative actions such as 'Expand Sales to New Customers'. Suppose that the user indicates that the effectiveness of this action depends on three preconditions, 1) 'New Products Available', 2) 'Sufficient Market Interest', and 3) 'Competitors' Prices Higher'. Whereas, the level of attainment of the first precondition can be controlled via actions such as 'Allocate Sufficient R&D Budget' and 'Expand Production Capacity', the latter two must be treated as uncertain preconditions not subject to one's direct control or scrutiny. The establishment of subgoals directed toward the attainment of such preconditions will only introduce unpursuable objectives into the structure. Thus, it is necessary to augment the goal-directed structure by incorporating a special treatment of uncontrollable preconditions.

5.2. Uncertain Events

Although the level of completion of an uncontrollable precondition cannot be directly adjusted by a user's actions, the user may be able to implement actions to enhance the likelihood of some events which, in turn, would increase the expectation of reaching a higher level of completion for the desired precondition. For example, consider the uncontrollable precondition 'Sufficient Market Interest' in Figure 5-1. One of two possible uncertain events may occur: 1) there is a sufficient market interest, or 2) the market interest is

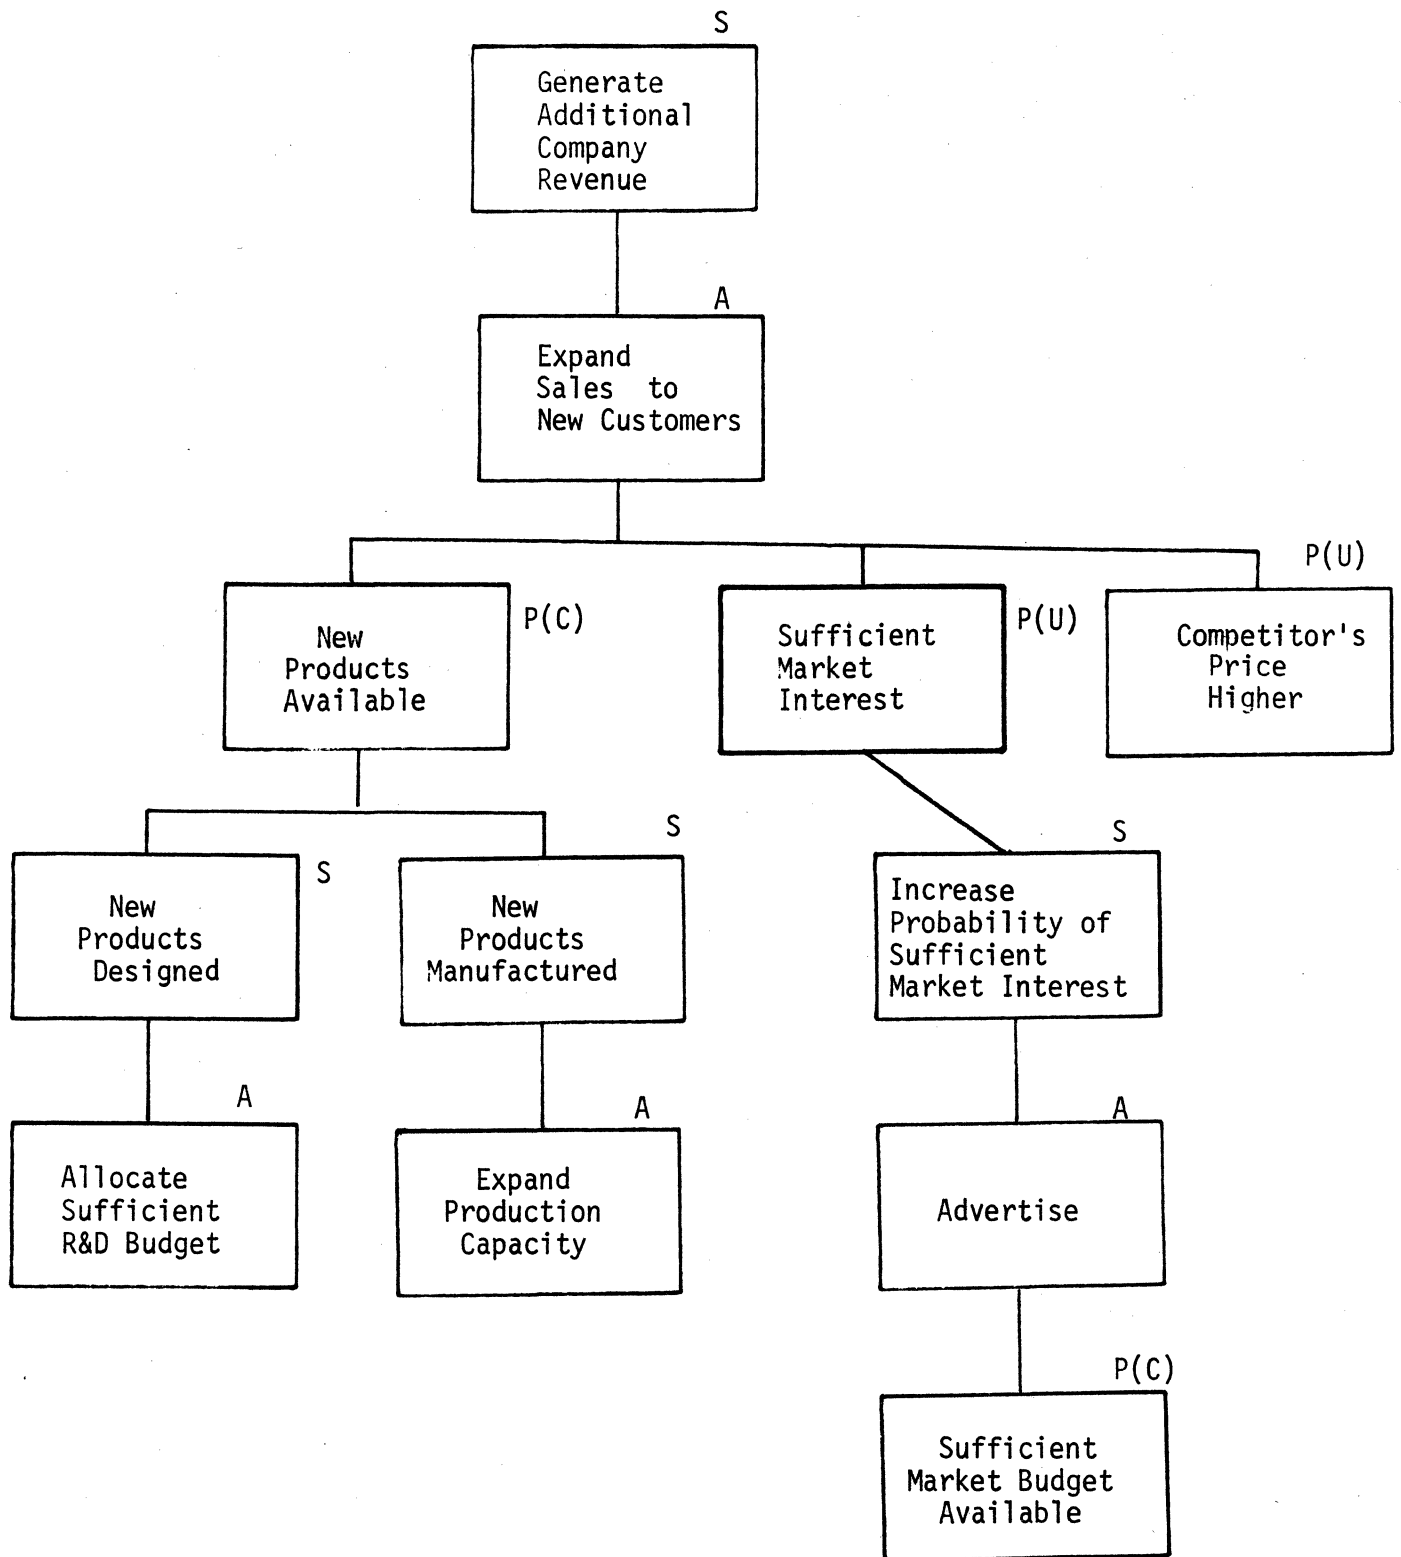


Figure 5-1. Long-Term Industrial Planning

insufficient. These events are 'uncertain' since their occurrence one way or the other is not perceived to be controlled by specific user actions but rather, must be represented by a probabilistic model. In such a case, rather than pursuing subgoals that directly satisfy the precondition, a more useful approach is to attempt to establish subgoals which are directed toward increasing the probability of the favorable events and decreasing the probability of unfavorable events. For instance, in the example of Figure 5-1, the subgoal of 'Higher Probability of Sufficient Market Interest' can be presented to the user, which may evoke actions, such as 'Advertising'. These actions may, in turn, require the satisfaction of certain controllable preconditions (e.g., availability of sufficient marketing budget) which is compatible with the structure described in the previous chapter.

5.3. Propagating the Impact of Uncertain Events

Two parameter vectors are associated with each uncontrollable precondition (Figure 5-2). The first vector $[p(t_1), p(t_2), p(t_3)]$ contains the probability of occurrence of each uncertain outcome. The second vector $[(L|t_1), (L|t_2), (L|t_3)]$ contains the level of completion of the precondition, given the occurrence of the corresponding uncertain outcome (Figure 5-2). Once the vectors are elicited, the system compares the elements of the second vector and selects, for further elicitation, the outcome with the highest level, i.e., the most desirable outcome. A new subgoal can then be created which is aimed at increasing the probability of the selected outcome. Usually, a strategy aimed at increasing the likelihood of the desired outcome is also helpful in avoiding its undesirable counterparts. Therefore, it is not necessary to setup a separate subgoal for each outcome.

As the expansion of the new subgoal continues, the probability vector is updated. At each point in time, during the elicitation process, the expected

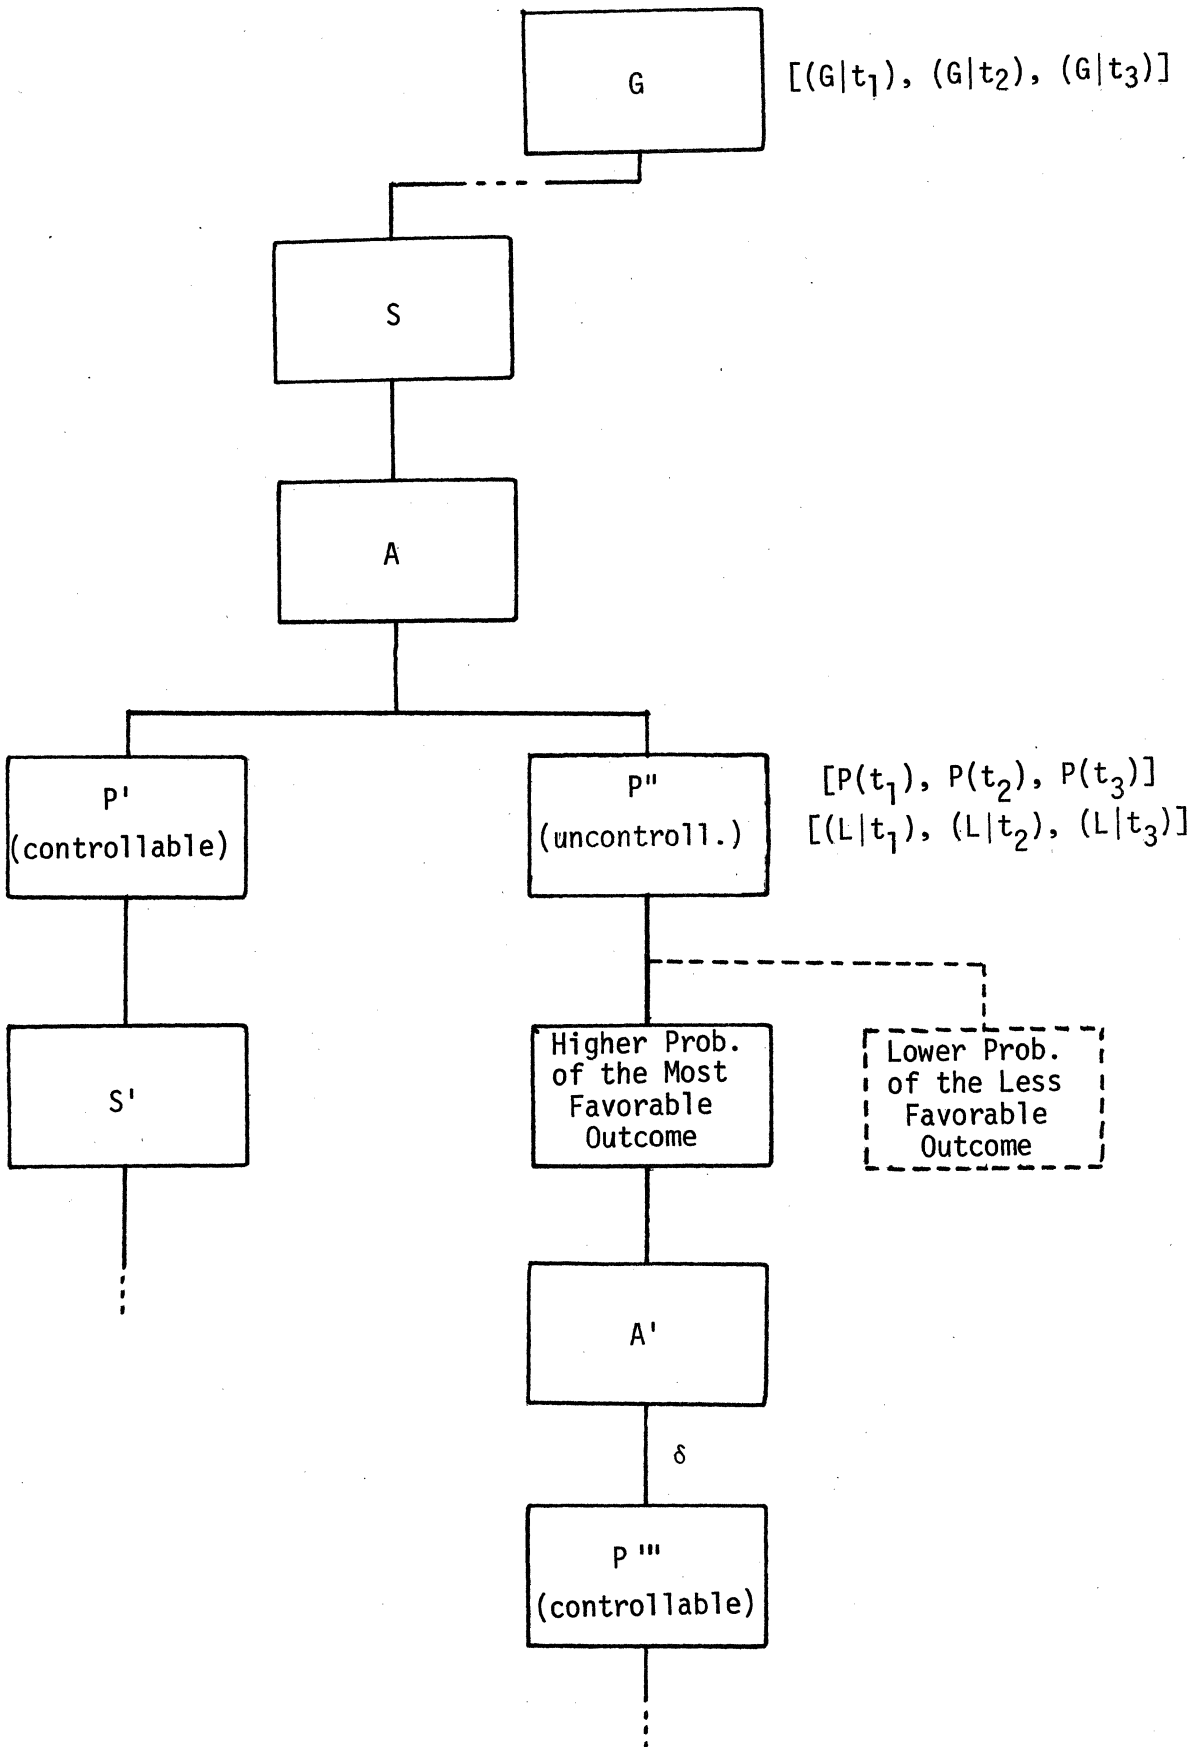


Figure 5-2. Example of a Goal-Directed Graph Containing Uncertain Event

level of attainment of the major goal can be calculated based on the two vectors associated with the uncontrollable precondition. The process begins by calculating the level of attainment of the major goal given the occurrence of each uncertain outcome individually. The result will be a new vector of conditional levels of attainment of the major goal, $[(G|t_1), (G|t_2), (G|t_3)]$ in Figure 5-2. In other words, $(G|t_i)$ is the value of the major goal assuming that outcome t_i has occurred. Then, the expected level of attainment of the major goal can be calculated:

$$G = (G|t_1) \cdot p(t_1) + (G|t_2) \cdot p(t_2) + (G|t_3) \cdot p(t_3)$$

Since the values of the probability vector $p(t_1)$, $p(t_2)$, $p(t_3)$ depends on the actions ultimately emerging from the uncontrollable precondition, the calculated value of G provides a selection criteria for the emerging action plans.

5.4. Optimization of Goal-Directed Graphs Containing Uncertain Events

The process of optimizing goal-directed graphs with uncertain events is similar to the process discussed in Chapter 3. However, dealing with a number of uncontrollable preconditions with multiple uncertain events, can lead to an exponential explosion in the number of optimization steps required. As an example, consider a section of a graph where an action A requires two uncontrollable preconditions each depending on three uncertain outcomes. Since the effectiveness of the action A depends on the level of completion of both preconditions, each possible combination of uncertain outcomes must be considered separately. This would require nine separate calculations in the example above. Thus, as the number of uncontrollable preconditions across the graph increases, the number of required calculations proliferates.

To overcome this complexity problem, a heuristic rule has been employed. In order to select the best action from a given set, the rule computes the

expectation of G by fixing the impacts (i.e., the δ 's) of all uncontrollable preconditions in the graph at their local expected value (expect those supporting the action under consideration). Thus, for all other preconditions (p_k) with uncertain outcomes ($t_1, \dots, t_i, \dots, t_n$) we assume:

$$\delta_k = \sum_{i=1}^n p(t_i) \cdot \delta_k (L|t_i)$$

That amounts to replacement of the uncertain events by a 'virtual event' producing an equivalent local impact. This heuristic rule circumvents the problem of exponential complexity and would not introduce significant inaccuracy beyond the approximation already used in consideration of multiple subgoals. In this case, also, it is assumed that all subgoals except the one under consideration are fixed at the mid-point of their range of variation.

6.0. IMPLEMENTATION PROGRESS

6.1. Elicitation

The program is currently able to elicit from the decision maker all of the required information for the construction of a goal-directed decision graph. The four major node types elicited are (1) objectives, (2) action strategies, (3) action modes, and (4) preconditions. Specific information requested for each type is the following:

(1) Objectives (subgoals)

- (a) Names - the name of each objective that describes the overall goal.
- (b) Weights - the relative importance of each objective.
- (c) Levels - a gross estimate of the expected level of achievement of each objective.

(2) Action Strategies

- (a) Names - the name of every action strategy (including 'no action' if applicable) that could be taken to accomplish the corresponding objective. This list should be mutually exclusive and complete.
- (b) Levels - an estimate of the level of achievement of the corresponding objective if each action were to be implemented.
- (c) Side Effects - an estimate of the amount of decrease in the achievement level of non-related objectives. The program currently elicits adverse side-effects only.

(3) Action Modes

- (a) Names - a list of the different methods that could be used to implement the corresponding action strategy.
- (b) Levels - a revised estimate of the effectiveness of the corresponding action strategy due to the implementation of the particular action mode.

(4) Preconditions

- (a) Names - the preconditions that must be satisfied before the corresponding action mode can be taken.
- (b) Criticality Thresholds - the completion level of each precondition in which the corresponding action mode just begins to become effective. That is, the precondition completion level below which the action mode is completely ineffective.
- (c) Completion levels - an estimate of the completion level of each precondition.

For uncontrollable preconditions, the following information is required:

- (a) Names - the names of the states of nature (events) that describe the corresponding precondition.
- (b) Probabilities - the probability of occurrence of each state.
- (c) Levels - the completion level that the corresponding precondition would receive if the particular state of nature were to occur.

During the elicitation process, the user is prompted by explanation of the type of information he is to provide. If he does not understand, or wishes more explanation, he may enter a question mark on the computer terminal and an alternate, more detailed explanation will be printed. The user is also kept aware of his current 'location' in the decision structure by constant reference to nodes along the path to the goal.

At the completion of each node expansion, the user may exit the elicitation mode by entering the key word 'terminate'. If he chooses not to continue the program exits from the elicitation procedure and enters a "system mode" in which the user can exercise various program options that process the current structure. To resume elicitation, the user simply types 'continue' and the elicitation procedure is re-entered at the proper point. This feature allows

the user to gain information about the current structure whenever he wishes, without resorting to a great number of questions.

6.2. Computations and Heuristic Optimization

The level of each node is calculated from the information available on all of its successors. The exact formulas for node level determination were explained in the previous sections. The following is a summary of these formulas and a listing of the currently operational computations in the program:

- (1) Goal - the level of the goal is a weighted average of its successor objectives (sub-goals).
- (2) Objectives - the level of each objective is the maximum of the levels of its successor action strategies modified by any side effects.
- (3) Action Strategies - the level of each action strategy is the maximum of the levels of its successor action modes.
- (4) Action Modes - the level of each action mode is the product of its successor preconditions. The preconditions may be individually modified by a criticality threshold.
- (5) Preconditions - each precondition is characterized by a vector of values corresponding to its successor states. The overall effect of a precondition on its parent action is given by the expected value of the latter with respect to the successor states.
- (6) Events - the level of each event is a weighted average of its successor objectives.

Currently, a single number is requested for each level estimate. Future implementation will provide an opportunity to input a minimum and a maximum value. Side effects are only allowed on action strategies and not on action modes. Further, the side effect is assumed to be adverse. Side effects that benefit related objectives is planned as a future task.

6.3. System Features

Once a user has entered system mode , he may exercise a number of options that will provide information about the decision structure that he has built so far. All of the system options require only a single keyword to be typed into the computer terminal. These keywords operate only in system mode and are not effective during the elicitation procedure. The following is a listing of the currently available options and an explanation of their function.

6.3.1. HELP

By typing HELP, the user can get a listing of all of the available option keywords as well as an explanation of how to get back into elicitation mode.

6.3.2. START

The START option clears out the memory containing the decision structure and begins a new elicitation session.

6.3.3. CONTINUE

The CONTINUE keyword allows the user to re-enter elicitation mode and continue building the decision structure from the point it was previously terminated.

6.3.4. ROLLBACK

This feature allows the user to find the current value of the goal as well as the best actions to take in order to maximize the goal value. The value of each node is computed from the value of its successors for all possible combinations of action strategies. That combination which produces the highest goal value is chosen for recommendation along with the expected attainment level of each objective.

6.3.5. GOALCHECK

This option allows the user to calculate the goal value with any desired

action subset. He inputs the combination of actions he wishes (currently by node number) and the program will calculate the value of the goal with that particular set. It does not disturb the current values.

6.3.6. STRUCTURE

The STRUCTURE option prints out the complete decision graph along with levels and modifiers (i.e., weights, criticalities, probabilities). The structure is printed in outline form with successor relationships shown as indentations.

6.3.7. ACTIONS

This option prints only the action subsets and their predecessors. It is helpful for the GOALCHECK option which requires a subset of actions by node number.

6.3.8. GOAL

The current value of the goal can always be checked by using this option.

6.3.9. INFO

This option allows the user to obtain all of the information about a particular node. This information includes:

- (1) The name.
- (2) The value.
- (3) The modifier value.
- (4) The predecessor.
- (5) The successors.
- (6) The side effects.
- (7) Whether or not it has been expanded.
- (8) Whether or not it is on the best path.

The node is specified by number which can be obtained using the STRUCTURE option.

6.3.10. TRACE

This option allows the user to 'watch' the rollback function as it calculates

the goal value for all possible combinations of action strategies looking for the one that maximizes the goal. Each and every action combination will be printed along with the corresponding goal value.

6.3.11. LOAD and SAVE

This command causes a copy of the program to be loaded from disc or to be saved on disc. Thus, a user can stop the session and continue at another time with all information saves.

6.3.12. PATH

By specifying any internal node, this option will print the path from that node to the goal.

6.3.13. CHANGE

This option allows the user to change the name or value of any node. (Future options will allow restructuring.)

6.3.14. EXPAND

If the user does not wish to expand the node recommended by the elicitation management algorithm, he may choose any terminal node he wishes and expand it.

6.4. Programmer Features

The following features were implemented for the purpose of facilitating programmer changes to the elicitation messages and phrases.

6.4.1. MESSAGE

By using a specific message code, the programmer can print any of the 240 message lines.

6.4.2. MESSAGEMAP

Allows the programmer to peer inside the message file and determine which messages are stored for various elicitation procedures.

6.4.3. EDIT

Allows easy editing of specific message lines. This includes options to

break the line if a small CRT terminal is in use or the option of inserting node names directly into the text.

6.5. Program Structure

Figure 6-1 shows a flow diagram of the major program structure. The 'Start' module begins a new elicitation session by evoking the 'Initialize' module. All data matrices are initialized and storage is reserved for node information. The 'Select' module is responsible for selecting the particular node to expand. If necessary, it can request a 'Rollback' to obtain the best action combinations as well as invoking the 'Elicitation Management' module to determine the best elicitation strategy.

The 'Expand' module is responsible for expanding a single node. It uses pre-structured messages and phrases stored in the 'Message File'. If the user exits elicitation mode, he can return through the 'Continue' module which continues the elicitation from the point where it left off. All modules access the data files to get information on the nodes.

6.6. Data Structure

The data structure for the program consists of five major matrices which store all of the required information about each node. The Main Data Matrix (Figure 6-2) contains most of the information and is three-dimensional measuring $4 \times 9 \times N$. The first dimension indicates node types: (1) Subgoals, (2) Actions, (3) Modes, and (4) Preconditions. The third dimension signifies the node number. There is no relation between the node number and its location in the decision graph. The nodes are simply stored sequentially as they are defined.

The second dimension of the Main Data Matrix contains the data items associated with each node. The 'Name Code' is a number that identifies a particular row in the Name Matrix (Figure 6-4). Since the name of a node contains more than one character, it must be referenced indirectly by code number.

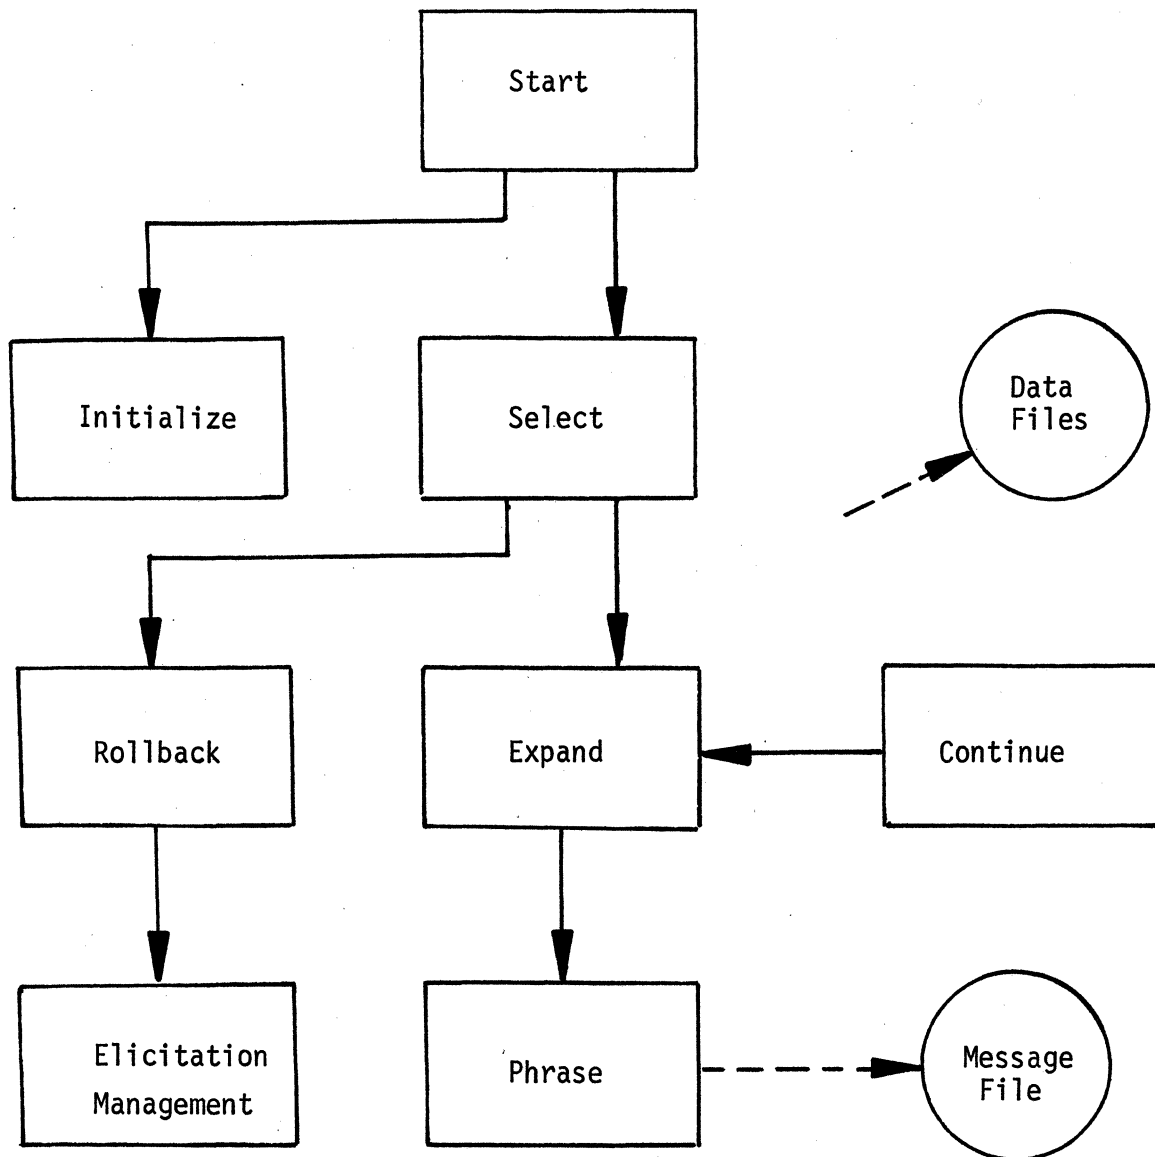


Figure 6-1. Program Structure

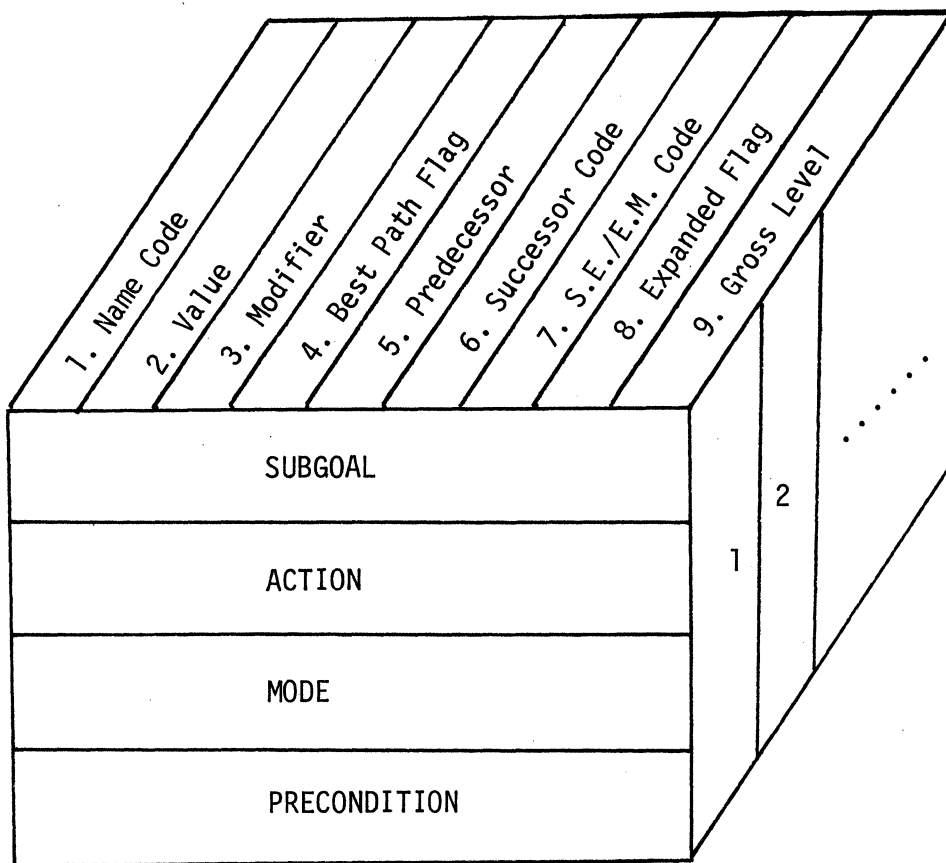


Figure 6-2. Main Data Matrix




Figure 6-3. Side Effects Matrix




Figure 6-4. Name Matrix

The numbers are simply assigned sequentially as nodes are defined.

The 'Value' is the level of attainment of the particular node and is stored directly in the Main Data Matrix. In the case of subgoals, the value is the average of the minimum value and the maximum value given by the user. The 'Modifier' is the term for the numeric value appearing on the arc immediately above the node. In the case of a subgoal, the modifier is the 'weight'. In the case of a precondition, the modifier is the 'criticality'. Actions and modes have no modifier and, thus, this data item is not used, but set to 1 as a default value.

The 'Best Path Flag' is a binary value which is set to 1 if the node is on the optimal path. The 'Predecessor' is a specific node number (on the third dimension) which corresponds to the predecessor node. Since nodes have just one predecessor, and since their type is completely determined, a single number is sufficient. For example, if Action Node #3 has a Subgoal #7 as its predecessor, a '7' is stored in the predecessor location for Action Node #3. Since actions must have subgoals as predecessors, the node is uniquely determined.

The 'Successor Code' is a reference number that accesses rows in the Successor Matrix (Figure 6-5). Unlike predecessors, nodes may have more than one successor. Thus, each row in the Successor Matrix lists the node numbers of the successors. The data item stored in the Main Data Matrix is this row number.

The 'S.E./E.M. Code' is used for two separate purposes depending on the node type. If the node is an action, the stored number is a code that refers to the Side Effects Matrix (Figure 6-3). Each action may have more than one side effect and each side effect is characterized by the subgoal to which it applies and the level of adverse effect. Thus, the Side Effects Matrix is three-dimensional and the S.E. Code Number accesses a particular 'plane' in the

Successor Code	1					
	2					
	3					
	4					
	5					

Figure 6-5. Successor Matrix

E.M. Code								
	1	2	3	4	5	6	7	8
	V_{\min}	V_{\max}	a	b		LS	sub-goal	pred. subg.

Figure 6-6. Elicitation Management Matrix

matrix. If the node is a subgoal, the Code Number refers to a row in the Elicitation Management Matrix (Figure 6-6). This matrix contains all of the necessary information for calculating the value-of-analysis and choosing the next node to expand. The first two entries are the minimum and maximum values of the subgoal. The following three entries are the vector values (a , b , α) which characterize the canonical function. The sixth entry is the 'likelihood switch' value. The last two entries contain redundant information about the next highest subgoals. This information saves processing time.

The eighth item in the Main Data Matrix is the Expanded Flag Code. This binary value is set to 1 if the node has been expanded and to 0 if the node has not yet been expanded. The last item is a gross level indicator which is simply a count of the number of subgoals from the node to the goal along the direct path.

7.0. SAMPLE PROTOCOL

In this chapter a sample protocol, resulting from an actual goal-directed system application, is presented. In this stage of the implementation, elicitation management is not active and nodes are expanded in breadth-first order (left-to-right). Evaluation of the major goal is accomplished by considering all possible combinations of actions -- a process which does not take a great deal of computer time with a small decision structure. Uncertain events are assumed to be binary.

7.1. Scenario

The user is a man in his mid-thirties with a Ph.D. in an engineering field of high demand. He is currently employed in a high technology research and development company. The company's incentive and payroll system is organized such that the total pay and qualification for more overtime are directly related to the quality of work. The user owns a house in the city and a piece of undeveloped land in the suburban area. He has never owned or managed his own business before. However, he believes that with the proper assistance he will be capable of conducting a personally owned business. He is married with two children both of preschool age. At the present time, his expenses are covered by his present salary, but he is unable to save substantial amounts of money. He considers his problem to be that of planning a strategy for improving his financial status without reducing the standard of living of his family.

The user responses are preceded by an asterisk (*). User comments, referenced by the numbers in the right margins, are shown in Section 7.3.

7.2. Goal-Directed System Interview

start

Goal-Directed Decision Support System

Do you wish to begin a new problem?

* yes

What is your problem?

* Improve my financial status

We will assume that your ultimate goal is to reach a satisfactory solution of the problem:

"Improve my financial status"

List all major objectives which, if realized, would contribute to the fulfillment of the goal.

1:

* ?

Let's analyze ways of achieving the main goal by considering its major components. List the major objectives which must be achieved before the goal can be attained.

1:

* Have large assets

2:

* Have high steady income

3:

*

Estimate the importance (0 to 10) of the objectives above.
objective 1: "Have large assets"

* ?

Assume that all objectives are partially satisfied and that you have a choice of increasing the level by which a particular objective will be attained. Rate the importance of such an improvement on each of the listed objectives. Use a scale of 0 to 10 for assessing the degree of importance: the higher the assessment value, the more important the objective.
objective 1: "Have large assets"

* 6

objective 2: "Have high steady income"

* 4

What level (0 to 10) is the following objective expected to reach in a well-planned strategy?

objective 1: "Have large assets"

* ?

Recall your previous experiences in this problem area.

Imagine similar situations in the past and assume that you are given a sufficient amount of time to plan your actions judiciously.

Estimate a reasonable level (0 to 10) of attainment for each objective which reflects

both environmental constraints and your available resources.

objective 1: "Have large assets"

* 6

objective 2: "Have high steady income"

* 7

C-1

C-2

C-3

C-4

Let's consider the objective "Have large assets"
List all possible action strategies that you can take
toward the fulfillment of "Have large assets".
Make sure that the list contains all
physically feasible approaches (including possibly
"doing nothing") and that no two of them can be combined.

1:
* Develop my land
2:
* Invest in stock market
3:
* Invest in real estate
4:
* Speculate in precious metals
5:
*

Suppose that these actions are taken. What is your
assessment of the resulting level of attainment
of "Have large assets" under such conditions?
action strategy 1: "Develop my land"

* 9 C-5
action strategy 2: "Invest in stock market"

* 6 C-6
action strategy 3: "Invest in real estate"

* 7 C-7
action strategy 4: "Speculate in precious metals"

* 6 C-8
Consider the following objectives"
1 Have large assets
2 Have high steady income
Does the action strategy "Develop my land"
have any adverse effect on any of the above objectives?
* yes
List the affected objectives by number (separated by spaces)
* 2 C-9
Suppose that the objective "Have high steady income"
would otherwise attain a level of (for example) 8.
What would its new level be assuming that
"Develop my land" was enacted?
* 5
Consider the following objectives:
1 Have large assets
2 Have high steady income
Does the action strategy "Invest in stock market"
have any adverse effect on any of the above objectives?
* no C-10
Consider the following objectives:
1 Have large assets
2 Have high steady income
Does the action strategy "Invest in real estate"
have any adverse effect on any of the above objectives?
* no C-11
Consider the following objectives:
1 Have large assets
2 Have high steady income

Does the action strategy "Speculate in precious metals" have any adverse effect on any of the above objectives?

C-12

* yes

List the affected objectives by number (separated by spaces)

* 2

Suppose that the objective "Have high steady income" would otherwise attain a level of (for example) 8.

What would its new level be assuming that "Speculate in precious metals" was enacted?

* 6

Let's consider the objective "Have high steady income"

List all possible action strategies that you can take toward the fulfillment of "Have high steady income".

Make sure that the list contains all physically feasible approaches (including possibly "doing nothing") and that no two of them can be combined.

1:

* Open a business

2:

* Find a second job

3:

* Increase my over-time

4:

*

Suppose that these actions are taken. What is your assessment of the resulting level of attainment of "Have high steady income" under such conditions? action strategy 1: "Open a business"

* 8

C-13

action strategy 2: "Find a second job"

* 4

C-14

action strategy 3: "Increase my over-time"

* 4

C-15

Consider the following objectives:

1 Have large assets

2 Have high steady income

Does the action strategy "Open a business"

have any adverse effect on any of the above objectives?

* yes

List the affected objectives by number (separated by spaces)

* 1

C-16

Suppose that the objective "Have large assets" would otherwise attain a level of (for example) 8.

What would its new level be assuming that "Open a business" was enacted?

* 3

Consider the following objectives:

1 Have large assets

2 Have high steady income

Does the action strategy "Find a second job"

have any adverse effect on any of the above objectives?

* yes

C-17

List the affected objectives by number (separated by spaces)

* 1

Suppose that the objective "Have large assets" would otherwise attain a level of (for example) 8. What would its new level be assuming that "Find a second job" was enacted?

* 3

Consider the following objectives"

- 1 Have large assets
- 2 Have high steady income

Does the action strategy "Increase my over-time" have any adverse effect on any of the above objectives?

* yes

C-18

List the affected objectives by number (separated by spaces)

* 1

Suppose that the objective "Have large assets" would otherwise attain a level of (for example) 8. What would its new level be assuming that "Increase my over-time" was enacted?

* 4

Let's consider the action strategy "Develop my land" List the names for various detailed methods by which "Develop my land" can be executed.

1:

* ?

Normally, approaches such as "Develop my land" can be executed in a variety of different ways (methods) in varying degrees of detail. Think about the most promising execution plans available and list their names. (Make sure that the list consists of mutually exclusive modes of operation).

1:

* Build a house

C-19

2:

*

For each mode, re-estimate the level (0 to 10) that "Have large assets" would attain as a result of enacting: action mode 1: "Develop my land - Build a house"

* ?

You previously estimated that "Develop my land" would have a certain effectiveness on a scale from 0 to 10. Normally, this value depends on the selected mode of implementation. For each listed mode, enter the revised level that objective "Have large assets" would assume, given that the mode was actually implemented.

action mode 1: "Develop my land - Build a house"

* 9

Let's consider the action strategy "Invest in stock market" List the names for various detailed methods by which "Invest in stock market" can be executed.

1:

* Utilities

C-20

2:

* Energy

3:

* Electronics

4:

*

For each mode, re-estimate the level (0 to 10) that
 "Have large assets" would attain as a result of enacting:
 action mode 1: "Invest in stock market - Utilities" C-21

* 5
 action mode 2: "Invest in stock market - Energy" C-22

* 7
 action mode 3: "Invest in stock market - Electronics" C-23

* 8
 Let's consider the action strategy "Invest in real estate"
 List the names for various detailed methods by which
 "Invest in real estate" can be executed.
 1:
 * Apartment building
 2:
 * Undeveloped land
 3:
 *

For each mode, re-estimate the level (0 to 10) that
 "Have large assets" would attain as a result of enacting:
 action mode 1: "Invest in real estate - Apartment building" C-24

* 7
 action mode 2: "Invest in real estate - Undeveloped land" C-25

* 5
 Let's consider the action strategy "Speculate in precious metals"
 List the names for various detailed methods by which
 "Speculate in precious metals" can be executed.
 1:
 * Gold
 2:
 * Silver
 3:
 *

For each mode, re-estimate the level (0 to 10) that
 "Have large assets" would attain as a resulting of enacting:
 action mode 1: "Speculate in precious metals - Gold" C-26

* 7
 action mode 2: "Speculate in precious metals - Silver" C-27

* 5
 Let's consider the action strategy "Open a business"
 List the names for various detailed methods by which
 "Open a business" can be executed.
 1:
 * Franchise C-28
 2:
 * Joint venture
 3:
 *

For each mode, re-estimate the level (0 to 10) that
 "Have high steady income" would attain as a result of enacting:
 action mode 1: "Open a business - Franchise" C-29

* 3
 action mode 2: "Open a business - Joint venture"

* 8

Let's consider the action strategy "Find a second job"
List the names for various detailed methods by which
"Find a second job" can be executed.

1:

* none

C-30

Let's consider the action strategy "Increase my over-time"
List the names for various detailed methods by which
"Increase my over-time" can be executed.

1:

* none

C-31

Let's consider the action mode "Develop my land - Build a house"
List any preconditions that must be established
before mode "Develop my land - Build a house" can be implemented.

1:

* ?

C-32

Normally, some preparation is usually required to
secure an effective execution of "Develop my land - Build a house".
List the conditions that must materialize
(be satisfied) or must be maintained (required effort
to guard against deterioration) so that the
effectiveness of "Develop my land - Build a house" (in terms of accomplishing
objective "Have large assets") is enhanced or secured.

1:

* Have sufficient funds

2:

*

Let's explore the importance of the realization
of these preconditions with respect to securing the
effectiveness of "Develop my land - Build a house". Estimate
(on a scale from 0 to 10) the level of completion
of the preconditions below which "Develop my land - Build a house"
is totally ineffective.

precondition 1: "Have sufficient funds"

* ?

C-33

It may be helpful to think about two extreme physical
quantities reflecting the completion level of
the preconditions to follow. Mentally assign 10 to
the highest imagined quantity and 0 to the
lowest. Keeping these quantities in mind, estimate
(on a scale from 0 to 10) where the following
precondition should be placed so that "Develop my land - Build a house"
just begins to become effective.

precondition 1: "Have sufficient funds"

* 8

C-34

What level (0 to 10) would you expect the following
precondition to reach in a well-planned
strategy aimed at realizing the precondition in question
and at the same time not having a
severe adverse effect on your other objectives?

- precondition 1: "Have sufficient funds"
- * ?
- Recall your previous experiences in this problem area.
 Imagine similar situations in the past and
 Assume that you are given a sufficient amount of
 time to plan your actions judiciously.
 Estimate a reasonable level (0 to 10) of attainment
 for each precondition to follow
 which reflects both environmental constraints and
 your available resources.
- precondition 1: "Have sufficient funds"
- * 9 C-35
- Let's consider the action mode "Invest in stock market - Utilities"
 List any preconditions that must be established
 before mode "Invest in stock market - Utilities" can be implemented.
 1:
- * none C-36
- Let's consider the action mode "Invest in stock market - Energy"
 List any preconditions that must be established
 before mode "Invest in stock market - Energy" can be implemented.
 1:
- * none C-37
- Let's consider the action mode "Invest in stock market - Electronics"
 List any preconditions that must be established
 before mode "Invest in stock market - Electronics" can be implemented.
 1:
- * none C-38
- Let's consider the action mode "Invest in real estate - Apartment building"
 List any preconditions that must be established
 before mode "Invest in real estate - Apartment building" can be implemented.
 1:
- * Secure apartment loan
- 2:
- * Let's explore the importance of the realization
 of these preconditions with respect to securing the
 effectiveness of "Invest in real estate - Apartment building". Estimate
 (on a scale from 0 to 10) the level of completion
 of the preconditions below which "Invest in real estate - Apartment building"
 is totally ineffective.
- precondition 1: "Secure apartment loan"
- * 1 C-39
- What level (0 to 10) would you expect the following
 precondition to reach in a well-planned
 strategy aimed at realizing the precondition in question
 and at the same time not having a
 severe adverse effect on your other objectives?

- precondition 1: "Secure apartment loan" C-40
- * 8
Let's consider the action mode "Invest in real estate - Undeveloped land"
List any preconditions that must be established
before mode "Invest in real estate - Undeveloped land" can be implemented.
1:
- * Secure land loan C-41
- 2:
- *
Let's explore the importance of the realization
of these preconditions with respect to securing the
effectiveness of "Invest in real estate - Undeveloped land". Estimate
(on a scale of 0 to 10) the level of completion
of the preconditions below which "Invest in real estate - Undeveloped land"
is totally ineffective.
precondition 1: "Secure land loan" C-42
- * 1
What level (0 to 10) would you expect the following
precondition to reach in a well-planned
strategy aimed at realizing the precondition in question
and at the same time not having a
severe adverse effect on your other objectives?
precondition 1: "Secure land loan" C-43
- * 8
Let's consider the action mode "Speculate in precious metals - Gold"
List any preconditions that must be established
before mode "Speculate in precious metals - Gold" can be implemented.
1:
- * none C-44
- Let's consider the action mode "Speculate in precious metals - Silver"
List any preconditions that must be established
before mode "Speculate in precious metals - Silver" can be implemented.
1:
- * none C-45
- Let's consider the action mode "Open a business - Franchise"
List any preconditions that must be established
before mode "Open a business - Franchise" can be implemented.
1:
- * none C-46
- Let's consider the action mode "Open a business - Joint venture"
List any preconditions that must be established
before mode "Open a business - Joint venture" can be implemented.
1:
- * Having management assistance C-47
- 2:
- *

Let's explore the importance of the realization of these preconditions with respect to securing the effectiveness of "Open a business - Joint venture". Estimate (on a scale from 0 to 10) the level of completion of the preconditions below which "Open a business - Joint venture" is totally ineffective.

precondition 1: "Having management assistance"

* 2

C-47

What level (0 to 10) would you expect the following precondition to reach in a well-planned strategy aimed at realizing the precondition in question and at the same time not having a severe adverse effect on your other objectives?

precondition 1: "Having management assistance"

* 6

C-48

Let's consider the precondition "Have sufficient funds" List conditions over which you have no direct control but which may, if they hold true, significantly impair or enhance the degree of effectiveness of "Develop my land - Build a house".

1:

* ?

Consider the effect of "Develop my land - Build a house" on "Have large assets". There may be uncertain factors (possible events, unexpected developments, unknown conditions, etc.) whose occurrence will significantly influence the effectiveness of "Develop my land - Build a house". Name such events.

1:

* Funds available

2:

* Funds not available

3:

*

What is the probability that the following event will hold true? (e.g. 0.6)

state 1: "Funds available"

* .7

C-49

state 2: "Funds not available"

* .3

Previously, you estimated that enacting "Develop my land - Build a house" would result in level 9 for objective "Have large assets". Now, suppose that the following event occurs.

Re-estimate the new level objective

"Have large assets" would reach.

state 1: "Funds available"

* 8

state 2: "Funds not available"

* 2

C-50

Let's consider the precondition "Secure apartment loan"

List conditions over which you have no direct control

but which may, if they hold true,

significantly impair or enhance the degree of

effectiveness of "Invest in real estate - Apartment building".

- 1:
- * Apartment loan approved
- 2:
- * Apartment loan denied
- 3:
- *

What is the probability that the following event will hold true? (e.g. 0.6)

- state 1: "Apartment loan approved"
- * .8
- state 2: "Apartment loan denied"
- * .2

C-51

Previously, you estimated that enacting "Invest in real estate - Apartment building" would result in level 6 for objective "Have large assets".

Now, suppose that the following event occurs.

Re-estimate the new level objective

"Have large assets" would reach.

- state 1: "Apartment loan approved"
- * 10
- state 2: "Apartment loan denied"
- * 0

C-52

Let's consider the precondition "Secure land loan"

List conditions over which you have no direct control

but which may, if they hold true,

significantly impair or enhance the degree of

effectiveness of "Invest in real estate - Undeveloped land".

- 1:
- * Land loan approved
- 2:
- * Land loan denied
- 3:
- *

What is the probability that the following event will hold true? (e.g. 0.6)

- state 1: "Land loan approved"
- * .6
- state 2: "Land loan denied"
- * .4

C-53

Previously, you estimated that enacting "Invest in real estate - Undeveloped land" would result in level 6 for objective "Have large assets".

Now, suppose that the following event occurs.

Re-estimate the new level objective

"Have large assets" would reach.

- state 1: "Land loan approved"
- * 10
- state 2: "Land loan denied"
- * 0

C-54

Let's consider the precondition "Having management assistance"

List conditions over which you have no direct control

but which may, if they hold true,

significantly impair or enhance the degree of

effectiveness of "Open a business - Joint venture".

1:
 * Assistance available
 2:
 * Assistance not available
 3:
 *

What is the probability that the following event will hold true? (e.g. 0.6)
 state 1: "Assistance available"
 * .5 C-55
 state 2: "Assistance not available"
 * .5

Previously, you estimated that enacting "Open a business - Joint venture" would result in level 7 for objective "Have high steady income".
 Now, suppose that the following event occurs.
 Re-estimate the new level objective
 "Have high steady income" would reach.
 state 1: "Assistance available"
 * 4 C-56
 state 2: "Assistance not available"
 * 7

Let's consider the state "Funds available"
 Consider this state as a new goal.
 List all major objectives which, if realized, would contribute to the fulfillment of the goal.
 1:
 * Get funds
 2:
 *

What level (0 to 10) is the following objective expected to reach in a well-planned strategy"
 objective 1: "Get funds"
 * 7 C-57

Let's consider the state "Apartment loan approved"
 Consider this state as a new goal.
 List all major objectives which, if realized, would contribute to the fulfillment of the goal.
 1:
 * none C-58

Let's consider the state "Land loan approved"
 Consider this state as a new goal.
 List all major objectives which, if realized, would contribute to the fulfillment of the goal.
 1:
 * none

Let's consider the state "Assistance available"
 Consider this state as a new goal.
 List all major objectives which, if realized, would contribute to the fulfillment of the goal.

- 1:
 * Find management assistance C-59
 2:
 *
- What level (0 to 10) is the following objective expected to reach in a well-planned strategy?
 objective 1: "Find management assistance"
 * 7 C-60
- Let's consider the objective "Get funds"
 List all possible action strategies that you can take toward the fulfillment of "Get funds".
 Make sure that the list contains all physically feasible approaches (including possibly "doing nothing") and that no two of them can be combined.
- 1:
 * Get building loan C-61
 2:
 * Refinance my house
 3:
 * Join with another investor
 4:
 *
- Suppose that these actions are taken. What is your assessment of the resulting level of attainment of "Get funds" under such conditions?
 action strategy 1: "Get building loan"
 * 8 C-62
 action strategy 2: "Refinance my house"
 * 7
 action strategy 3: "Join with another investor"
 * 8
- Consider the following objectives:
 1 Have large assets
 2 Have high steady income
 3 Get funds
 4 Find management assistance
- Does the action strategy "Get building loan" have any adverse effect on any of the above objectives"
 * yes C-63
 List the affected objectives by number (separated by spaces)
 * 2
- Suppose that the objective "Have high steady income" would otherwise attain a level of (for example) 8. What would its new level be assuming that "Get building loan" was enacted?
 * 6 C-64
- Consider the following objectives:
 1 Have large assets
 2 Have high steady income
 3 Get funds
 4 Find management assistance

Does the action strategy "Refinance my house" have any adverse effect on any of the above objectives?

* yes C-65
List the affected objectives by number (separated by spaces)

* 2
Suppose that the objective "Have high steady income" would otherwise attain a level of (for example) 8. What would its new level be assuming that "Refinance my house" was enacted?

* 6
Consider the following objectives:
1 Have large assets
2 Have high steady income
3 Get funds
4 Find management assistance
Does the action strategy "Join with another investor" have any adverse effect on any of the above objectives?

* no C-66
Let's consider the objective "Find management assistance" List all possible action strategies that you can take toward the fulfillment of "Find management assistance". Make sure that the list contains all physically feasible approaches (including possibly "doing nothing") and that no two of them can be combined.

1:
* Get a partner
2:
* Hire a manager C-67
3:
* Join an existing business
4:
*
Suppose that these actions are taken. What is your assessment of the resulting level of attainment of "Find management assistance" under such conditions?

action strategy 1: "Get a partner"

* 8 C-68
action strategy 2: "Hire a manager"

* 7
action strategy 3: "Join an existing business"

* 9
Consider the following objectives:
1 Have large assets
2 Have high steady income
3 Get funds
4 Find management assistance
Does the action strategy "Get a partner" have any adverse effect on any of the above objectives?

* no
Consider the following objectives:
1 Have large assets
2 Have high steady income
3 Get funds
4 Find management assistance

Does the action strategy "Hire a manager"
have any adverse effect on any of the above objectives?

* no

Consider the following objectives:

- 1 Have large assets
- 2 Have high steady income
- 3 Get funds
- 4 Find management assistance

Does the action strategy "Join an existing business"
have any adverse effect on any of the above objectives?

* yes

List the affected objectives by number (separated by spaces)

* 1

Suppose that the objective "Have large assets"
would otherwise attain a level of (for example) 8.

What would its new level be assuming that

"Join an existing business" was enacted?

* 6

Let's consider the action strategy "Get building loan"

List the names for various detailed methods by which

"Get building loan" can be executed.

terminate

Exit elicitation mode.

Enter system mode.

C-69

7.3. User Comments

- C-1: I did not understand what is meant by importance.
- C-2: Although both objectives, 'Have large assets' and 'Have high steady income', are instrumental for improving my financial status, I perceive having large assets to be more crucial to providing improvement in my financial status.
- C-3: Considering my present status and potential, it is possible to plan an investment program that would result in assets of about \$70k in three years.
- C-4: Considering my upcoming promotion and the possibility of increasing my overtime, there is a very high probability of increasing my steady income by about 25 percent in the near future. Besides, there is a possibility of finding a second job, therefore increasing my income by another factor of about 30 percent at the expense of reducing some of my overtime work at the present job.
- C-5: Current expected profit margin in land development is extremely high.
- C-6: Although I am familiar with some stocks with good expectation, the present economical situation makes investment in the stock market somewhat risky.
- C-7: Although the expected recession may decrease the market demand, real estate prices are tied to the inflation rate which is still rising. However, the option of investing in real estate will not be as profitable as developing my land.
- C-8: Since I consider buying only futures on a margin, although the price of precious metals rise with inflation, local fluctuations may get me out of the game, thus, loosing even my original investment.
- C-9: My present work environment is organized such that compensation, bonuses, and promotions are directly related to my effort. Considering that developing my land will make demands on my time and effort, it may

- actually decrease my potential for increasing my steady income.
- C-10: Since my investment in the stock market would be through my broker, the required time and effort would be negligible, permitting me to pursue other activities.
- C-11: The situation would be the same as investing in the stock market.
- C-12: Although speculation in precious metals would also be done through a broker, it requires continuous monitoring and analysis of market behavior, and thus demanding a considerable portion of my time and effort (however, less than the amount required for developing my land).
- C-13: Considering my potential level of investment in a business and my degree of capability in running it, I am pretty confident that I can develop a business with a net profit of at least \$25k per year.
- C-14: I expect finding a second job will increase my steady income by about \$17k per year.
- C-15: The amount of increase in my overtime at my present job is limited. I expect to be able to increase my income by about \$10k per year through extra overtime at my present job.
- C-16: Opening a business will occupy so much of my time that I will hardly be able to pursue any active investment at all.
- C-17: Having two jobs at the same time takes almost all my time.
- C-18: Although not as much as opening a business or having a second job, increasing my overtime sufficiently also decreases my available time and energy for pursuing an active investment.
- C-19: Considering the R1 zoning of my land, the only feasible development would be to build a single family residential unit.
- C-20: Considering the economic situation and its effect on different stocks, I consider the only three stocks with a promising future to be utilities,

energy, and electronic stocks.

- C-21: Although the profit margins of utility companies are increasing, the rise in fuel cost is very likely to slow down the profit margin rate of increase, thus negatively affecting the rate of increase in stocks.
- C-22: Energy stocks, especially the alternative energy stocks, are very promising in this period of energy shortage.
- C-23: Due to the rapid growth of the industry, resulting from the innovative technology, electronic stocks are probably the most promising stock today.
- C-24: Due to the high inflation rate, the price of building materials and construction workers' salaries is rising so rapidly that the rate of increase in the price of the building itself seems to be higher than the rate of increase in the price of undeveloped land.
- C-25: Although high inflation is always an insurance for an increase in real estate prices, the potential forthcoming recession will slow down the building activities, thus decreasing the demand for undeveloped land, which in turn will lower the rate of increase in undeveloped land price.
- C-26: Due to the historical importance of gold, a high inflation rate will cause the price of gold to increase greatly.
- C-27: Although the price of silver also rises according to the inflation rate, since the major use of silver is its industrial application, the potential forthcoming recession will have a negative effect on rising silver prices.
- C-28: Since I do not have sufficient know-how in running a business independently. Opening an independent business is not feasible.
- C-29: Although major help is offered in areas such as management and advertising by the parent company, a considerable portion of the profit will be indirectly transferred to the parent company.
- C-30: There are no different ways of finding a second job.

- C-31: There is no other way of increasing my overtime.
- C-32: What is meant by precondition?
- C-33: I need more explanation.
- C-34: Having sufficient funds is very critical. I cannot complete building the house unless initially I have at least 80 percent of the sufficient funds required to build the house.
- C-35: Considering available resources, I am very confident that I can acquire the sufficient funds.
- C-36: Although there are different utility companies that I can invest in, the nature of the investment in all these prospects would be the same.
- C-37: The same as in utility stocks.
- C-38: The same as in utility stocks.
- C-39: I cannot invest in an apartment building unless I acquire a mortgage loan.
- C-40: Considering my credit record, I am pretty confident that I can acquire a mortgage loan.
- C-41: The same as in the case of investing in an apartment building.
- C-42: Again, acquiring a mortgage loan is absolutely critical.
- C-43: The only feasible way to speculate in gold is to buy different gold futures on margin, which are basically the same.
- C-44: The same as speculating in gold.
- C-45: The nature of all feasible franchises are sufficiently similar.
- C-46: Not having sufficient background in managing a business, I need to have some management assistance.
- C-47: If I have 50 percent assistance in management, I can still operate a joint venture effectively.
- C-48: I think I can find management assistance at least sufficient to effectively run the business.

- C-49: Since I have different sources, other than getting a construction loan, to provide funds (such as a second mortgage on my present house), the probability of having funds available is pretty high.
- C-50: Even if sufficient funds for completing the building project were not available, I can still increase the value of my land by grading it using some savings that I already have.
- C-51: Considering my credit history and the fact that the apartment will create some further income, I will have a good chance of getting my loan approved.
- C-52: I can purchase the apartment only if my loan is approved.
- C-53: Since the land does not provide any form of income, the chance of approval of a mortgage loan application for purchasing land is less than that for purchasing an apartment.
- C-54: Again, I cannot purchase the land unless my loan is approved.
- C-55: At this point, I believe that I have a 50-50 chance of finding assistance.
- C-56: Without assistance, I can still run the business, but considerably less effectively.
- C-57: There is a good chance of getting sufficient funds.
- C-58: There is none.
- C-59: If there was no assistance available, I can search for other sources (such as finding a partner or employing a manager) for management assistance.
- C-60: I have a good chance of finding management assistance for the kind of business I have in mind.
- C-61: Besides getting a construction loan or refinancing my present house, I realize that I can provide sufficient funds by joining another investor.
- C-62: I believe the capital gain on my house is sufficient. However, I am more confident in the other two alternative ways of acquiring funds.

- C-63: Since I have to repay the mortgage loan monthly, my net steady income will decrease.
- C-64: Considering the monthly payments and my present income.
- C-65: The same as in getting a construction loan.
- C-66: Since I offer the land and he provides the money for construction, there will not be any monthly payments from one to the other.
- C-67: Now I realize that besides getting a partner or hiring a manager, I can acquire management assistance also by joining an existing business rather than opening my own.
- C-68: A partner may also not be a good manager and hired assistance may be good in management but unfamiliar with this business. The partners in an existing business, however, have proven to be good managers and also familiar with the business.
- C-69: Although it has many advantages, joining an existing business does not provide me with as much capital investment as in the case of opening my own business. In other words, I am paying something for the convenience of having the management assistance and working business.

7.4. Resulting Goal-Directed Graph Structure

```

structure
(Node number, Node name, Value, Modifier.)
Improve my financial status 0
  objective (level 1)
- 1. Have large assets 6 6
    action strategy (level 2)
- - 1. Develop my land 9 0
      action mode (level 3)
- - - 1. Develop my land - Build a house 9 0
        precondition (level 4)
- - - - 1. Have sufficient funds 9 8
          state (level 5)
- - - - - 1. Funds available 8 0.7
            objective (level 1)
- - - - - - 3. Get funds 7 10
              action strategy (level 2)
- - - - - - - 8. Get building loan 8 0
- - - - - - - 9. Refinance my house 7 0
- - - - - - - 10. Join with another investor 8 0
- - - - 2. Funds not available 2 0.3
- - 2. Invest in stock market 6 0
    action mode (level 3)
- - - 2. Invest in stock market - Utilities 5 0
- - - 3. Invest in stock market - Energy 7 0
- - - 4. Invest in stock market - Electronics 8 0
- - 3. Invest in real estate 7 0
    action mode (level 3)
- - - 5. Invest in real estate - Apartment building 7 0
        precondition (level 4)
- - - - 2. Secure apartment loan 8 1
          state (level 5)
- - - - - 3. Apartment loan approved 10 0.8
- - - - - 4. Apartment loan denied 0 0.2
- - - 6. Invest in real estate - Undeveloped land 5 0
        precondition (level 4)
- - - - 3. Secure land loan 8 1
          state (level 5)
- - - - - 5. Land loan approved 10 0.6
- - - - - 6. Land loan denied 0 0.4
- - 4. Speculate in precious metals 6 0
    action mode (level 3)
- - - 7. Speculate in precious metals - Gold 7 0
- - - 8. Speculate in precious metals - Silver 5 0

```

- 2. Have high steady income 7 4
action strategy (level 2)
- - 5. Open a business 8 0
action mode (level 3)
- - - 9. Open a business - Franchise 3 0
- - - 10. Open a business - Joint venture 8 0
precondition (level 4)
- - - - 4. Having management assistance 6 2
state (level 5)
- - - - - 7. Assistance available 4 0.5
objective (level 1)
- - - - - - 4. Find management assistance 7 10
action strategy (level 2)
- - - - - - - 11. Get a partner 8 0
- - - - - - - 12. Hire a manager 7 0
- - - - - - - 13. Join an existing business 9 0
- - - - - 8. Assistance not available 7 0.5
- - 6. Find a second job 4 0
- - 7. Increase my over-time 4 0

Side effects:

(Action number, Affected objective number, Adverse effect)

1 2 0.625
4 2 0.75
5 1 0.375
6 1 0.375
7 1 0.5
8 2 0.75
9 2 0.75
13 1 0.75

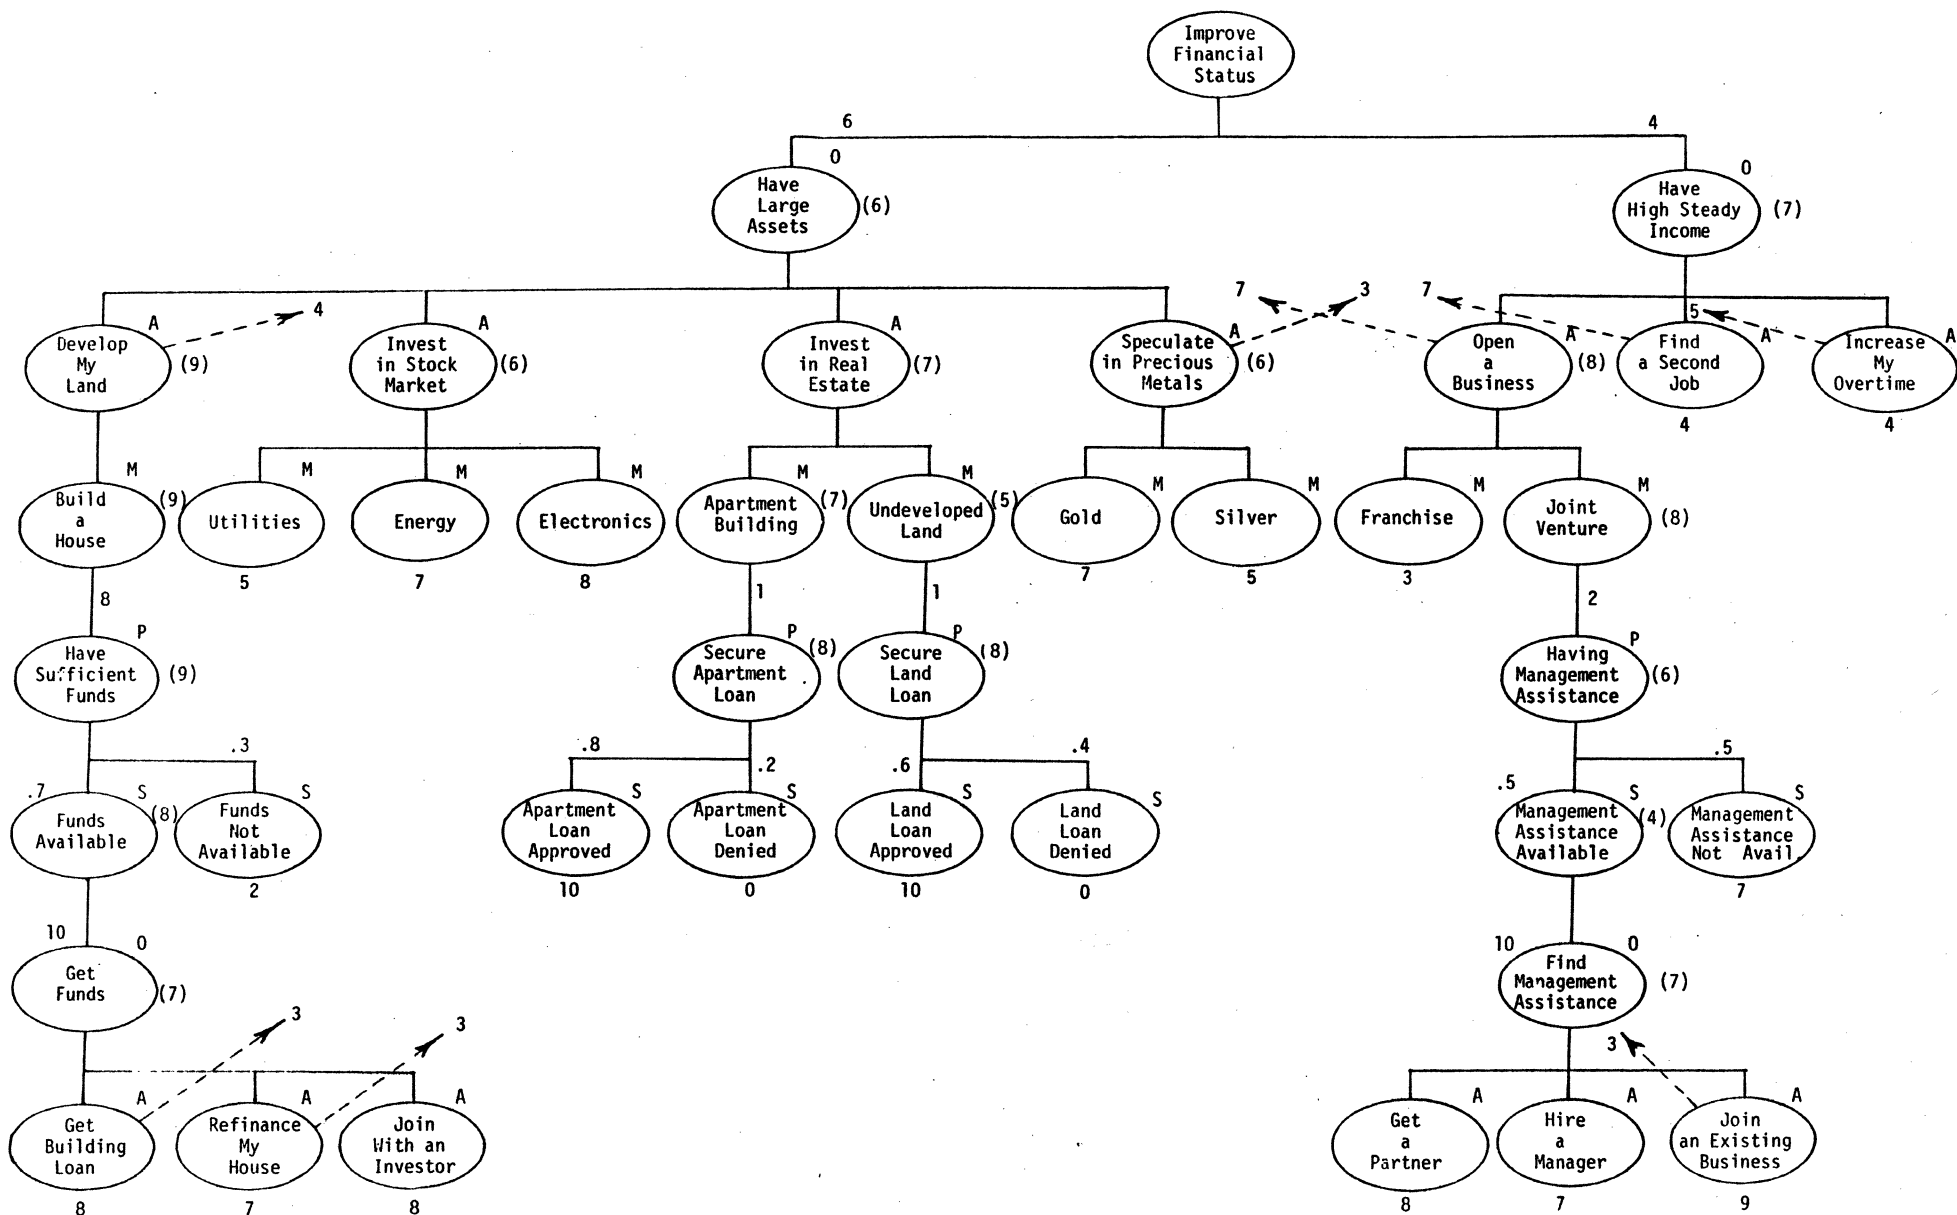


Figure 7-1. Structure of Goal-Directed Graph Elicited from the Dialogue

7.5. System Recommendation

The goal of "Improve my financial status" can be attained to level of 4.55 if the following actions are taken:

Implement "Invest in stock market - Electronics" toward the objective "Have large assets".

Simultaneously, implement "Get a partner" toward "Find management assistance" which will eventually facilitate implementation of "Open a business - Joint venture" leading to attainment of "Have high steady income".

8.0. PRELIMINARY EVALUATION AND FUTURE PROSPECTS

A brief examination of the dialogue presented in the last chapter reveals the main strengths and weaknesses of this support-system. The most striking negative features of this dialogue are its apparent length and repetitiveness. These weaknesses can be attributed to several factors; some are basic to all fully-computerized situation-based support-systems and others are due to the incomplete state of the current version of the program.

The lack of a sophisticated dialogue-management procedure in this version of the program is responsible for many of its shortcomings. The breadth-first scheme used for controlling the dialogue reported gives the user no sense of direction. In many instances the user can predict the next question to be asked, which induces boredom and mistrust in the program's ability to adequately comprehend the information given to it. In addition, the elicitation of provisional value judgements, whose sole purpose is to support the elicitation-management algorithm, now seems totally unnecessary, while adding significantly to the length of the dialogue. We hope that these deficiencies will be rectified in the very near future, when the dialogue-management scheme described in Chapter 5 becomes operational.

It should also be mentioned that the repetitive, mechanical style displayed by the present form of the program is partly due to the fact that in this phase of the research, we have made no special effort to equip the program's queries with a more natural 'flair'. A significantly more human style of conversation can, for example, be obtained by a random selection of synonymous phrases to avoid repetition (see Leal and Pearl, 1977) and by exposing the queries' purpose: e.g., "It is crucial that we first examine ways of achieving 'X'" or "I am trying to find out whether you foresee any special difficulties in executing 'Y'", etc. Simple language-analysis features such as syntactic transformations, word

matching, and key-word control would also greatly enhance the natural flavor of the dialogue style.

Several observers of the goal-directed support system have also commented that they sometimes feel uncomfortable assigning numerical values to the judgments requested, and that they occasionally feel unsure of what these numerical values represent or how to calculate them. The current system is equipped with several instructional features which can provide, upon request, a more detailed explanation of the nature of the assessment requested. Part of the 'assessment discomfort' can be alleviated by improving these features, and part would be remedied when the dialogue-management program is installed and the user is asked to provide not a single number but a range of possible values.

However, we attribute the basic difficulty connected with assessing levels of attainment and strengths of influences to the fact that in everyday discourse, these same concepts and relations are communicated in qualitative, non-numeric castings. Not too long ago, before the general public became accustomed to numerical broadcasting of weather-predictions and accident statistics, the quantification of likelihood judgments (i.e., probability) met with similar resistance and uneasiness. We also found that after several days of working with the system, users saw no difficulty in interpreting and performing the assessments required. Consequently, we hope that the decision-makers who could benefit from frequent consultation with such support systems will quickly become familiar with its somewhat non-traditional parameters.

For the occasional, inexperienced, and non-technical users, we are currently examining a more drastic, but more promising, solution: disposing with numerical estimates altogether. Most of human knowledge and skills are acquired via non-numerical media. Most training manuals and committee's reports convey useful information in purely linguistic terms. We read a newspaper article and

feel very comfortable with statements such as "This vote by Congress would substantially impair the President's bargaining power". Although phrased qualitatively, we do acknowledge that such a statement conveys important and useful factual information without insisting on numerical explication of the degree of impairment. Similarly, it would be more natural and comfortable for the common decision-maker to respond to queries such as:

Computer: "Is this condition absolutely necessary for action X or just desirable?"

or

Computer: "Is it very likely or just probable? Choose the most appropriate term:

remotely possible, possible, probable, quite probable, likely, very likely, almost sure, ... "

Behind the scenes, the program can map the user's linguistic response onto an appropriate numerical scale and propagate the resulting value through the graph by the methods described in Chapter 2. The user, however, will be spared the labor of quantifying inherently linguistic variables and the guilt associated with issuing uncertain estimates.

This approach will undoubtedly raise objections of the traditional analysts who may view the reliance on linguistic, rather than numerical, inputs as a backward regress toward the prescientific era of speculative alchemy and 'seat of the pants' decision-making. However, the ultimate objective of decision-analysis is (see Chapter 1) to produce a formal and valid representation of the decision-maker's experience. Forcing a person to produce numbers would not, by itself, make the representation more valid, especially when one's experience is encoded qualitatively. A more reasonable approach would be to incorporate into the formal model as many of these qualitative relations as possible, so as to

make the end results insensitive to the exact magnitude assigned to each relation. We believe the goal-directed structure is a step in this direction; it is made up of many detailed and cognitively clear relationships which render the exact quantification of each component less critical. We feel, for instance, that the statement "noise level and safety are two factors of 'roughly equal' importance" conveys more reliable information than any reasonable numerical response to the query: "How many people seriously injured or killed per year, call that number x , makes you indifferent between the option: [x injured or killed and 2500 persons subjected to high noise levels] and the option: [one person injured or killed and 1,500,000 subjected to high noise levels]?" (Slovic et al., 1977, quotation from Keeney's analysis of 'The Mexico City Airport').

Succinctly, our basic position on this issue can be summarized by the belief that qualitative relationships of many cognitively meaningful concepts can be made to produce more accurate results than numerical quantification of few cognitively unmanageable relationships.

Although we have not performed systematic experiments for evaluating the merit of the goal-directed decision-support system (such experiments will be the main focus of our research during the next year), it appears that the goal-directed structure offers several advantages over the traditional decision-tree approach. Our personal experiences with the two types of decision support systems confirm earlier expectations that the goal-directed approach would offer superiority in both clarity and purposefulness.

We find it clear, natural, and pleasing to talk about one's need to obtain a loan in order to build a house, to quantify the degree of this need, or to express directly the fact that refinancing one's house would diminish one's spendable income. These options of expression are simply not provided by the decision-tree approach, where only action-sequences and would-states are

considered, while aspects, issues, and factors remain inexplicitly assumed.

Similarly, we have on several occasions noticed that the explicit mention of an objective by the program focuses the attention of the user on a host of related experiences and evokes a number of unconventional alternatives capable of realizing that objective. For example, the idea of refinancing one's house and using the funds to develop one's land is very common to anyone with a little experience in real estate. However, to a user with no previous exposure to real estate maneuvers, this possibility either may not occur or, in the more common case, the possibility may be discarded from conscious attention by virtue of emotional barriers or unpleasant associations it may carry. The goal-directed method weakens the impact of such barriers by focusing on a single objective at any given time and instructing the user to ignore, for the moment, all side effects. It should be very hard for a user responding to the query: "List all possible action strategies that you can take toward the fulfillment of 'Get funds'" not to mention the possibility 'Refinance my house', regardless of the adverse implications that such an alternative may carry.

Based on these preliminary results and observations, we cannot rule out the prospect that the goal-directed structure described in this report will develop into the standard architecture for next generation decision-support systems. It offers the capability of continuously sweeping the spectrum between situation-based and knowledge-based systems (depending on the scope and level of details required). It is capable of operating as a fully computerized system as well as in an 'analyst's apprentice' capacity. Finally, it is conceptually appealing and permits both systematic and directional formalization of knowledge.

APPENDIX A. SIMULATION RESULTS FOR LOCAL OPTIMIZATION

The local optimization procedure is based on the selection of individual actions from each (mutually exclusive) action set. The 'action set' is the collection of actions supporting a common subgoal. In each stage of the local optimization procedure, a criterion is optimized with respect to only one action set, while the other sets are kept at a fixed level. However, due to the multiplicative nature of the side-effect combination rule, the ultimate effect of an individual action on the major goal cannot be assessed in isolation. The entire action plan, composed of a single action from each action set, must be known for an exact solution. Therefore, prior to the selection of individual actions, it is beneficial to represent each action set by a fixed 'representative' action which has roughly the same overall effects as that action which would be ultimately chosen by global optimization (Figure A-1). The representative action may not necessarily be a member of the given set. Rather, it can be 'virtual', producing a combination of direct and side effects in accordance with some compression algorithm.

The collection of all representative actions constitutes the 'initial action combination'. In the first cycle of the optimization procedure, the virtual actions are replaced, one at a time, by real actions taken from the corresponding sets. The process can be iterated any desired number of times. The resultant action combination may not converge to the globally optimum and the degree of suboptimality depends on the choice of the initial action combination.

This appendix summarizes the results of a simulation performed to aid in the selection of an initial action combination which produces a close-to-optimal solution. For this simulation, four alternative methods for choosing representative actions were examined:

- (1) Setting the representative direct effect equal to the average of

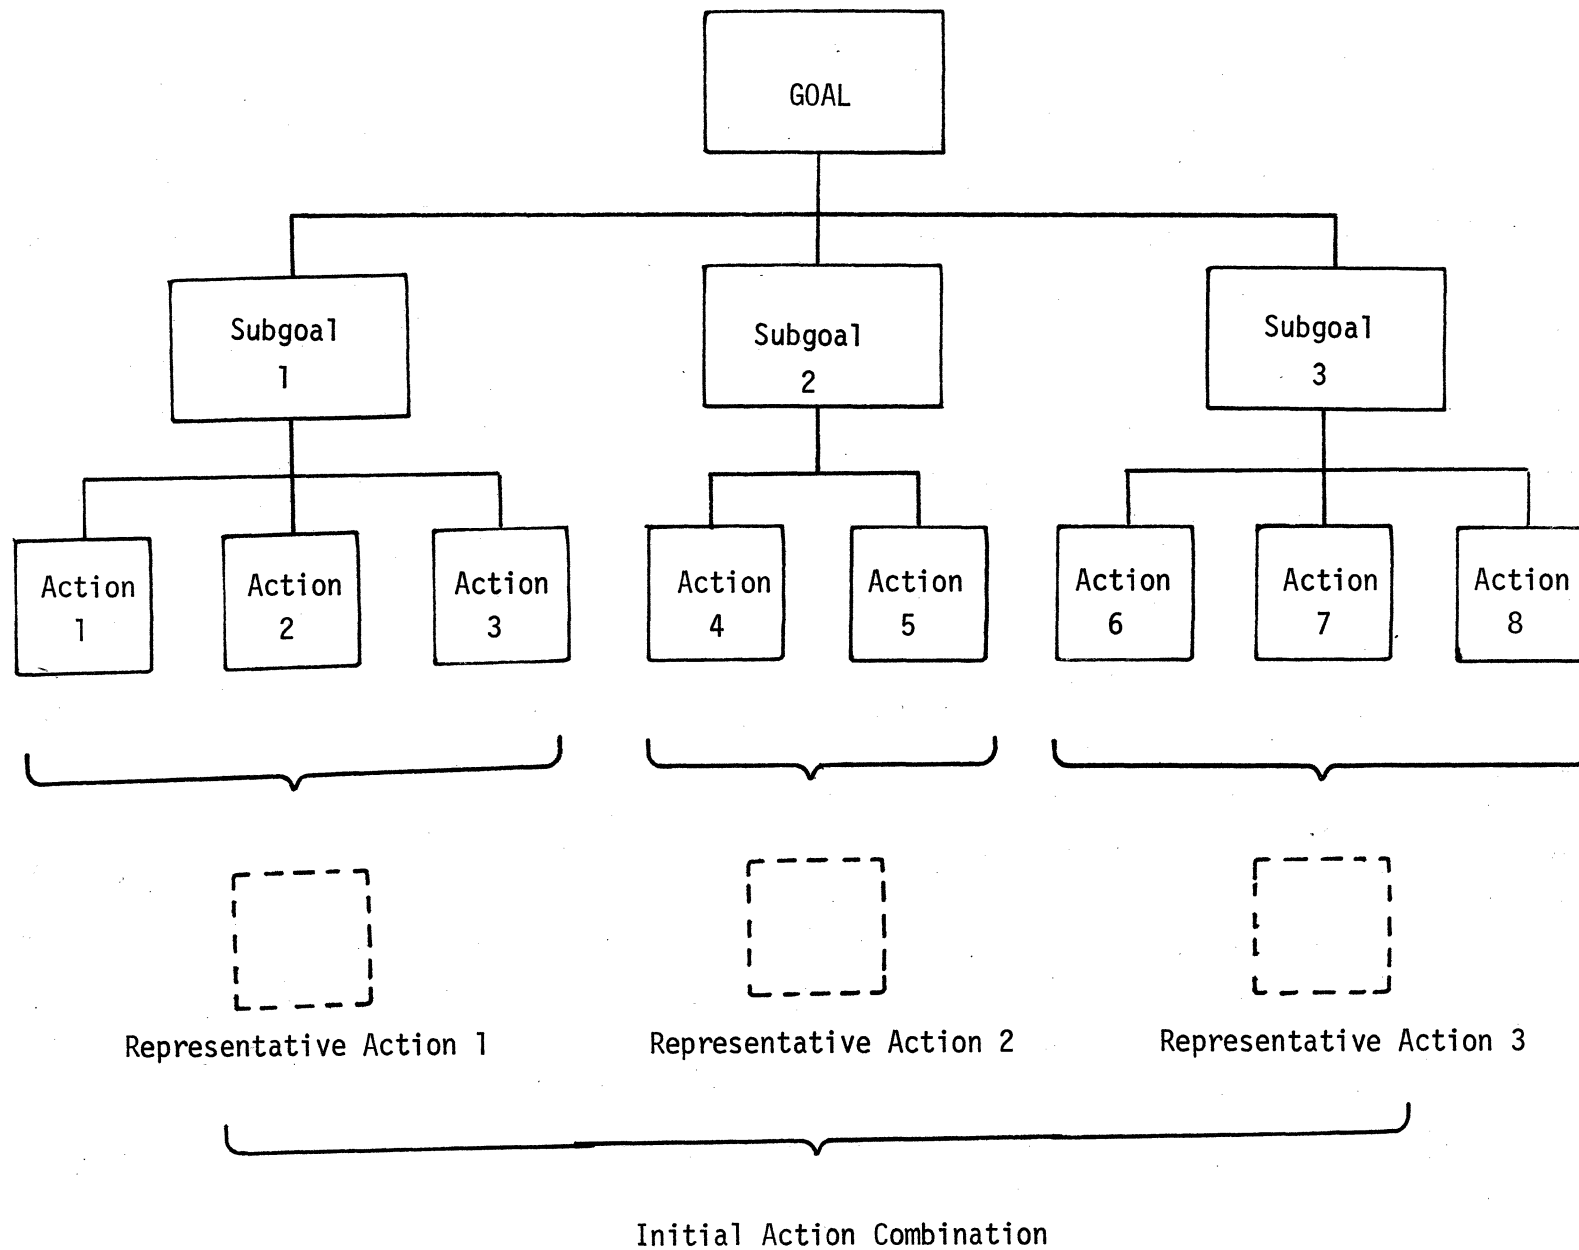


Figure A-1. Representative Actions

the actual direct effects of all actions supporting a common subgoal and setting the representative side effects equal to zero.

- (2) Setting the representative direct effect equal to the average of the actual direct effects and setting the representative side effects to the average of the actual side effects initiated from the set (directed towards the same target subgoals).
- (3) Setting the representative direct effect equal to the maximum of the actual direct effects and setting the representative side effects equal to zero.
- (4) Setting the representative direct effect equal to the maximum of the actual direct effects and setting the representative side effects equal to the side effects of the actual action with largest direct effect.

Among the four methods, only (4) employs representative actions directly from the actual actions suggested by the user. The other methods start the optimization using virtual effects as a substitution for the actual actions in the graph.

A simulation was performed to test these four alternative methods. A single-level goal-directed graph with variable parameters was structured. The weights, direct effects, and side effects were selected randomly, between zero and one, in such a way that all possible combinations of small, medium, and large direct effects and side effects would be generated (small: 0-0.33, medium: 0.33-0.66, large: 0.66-1).

At each simulation step, the optimal action combination was first calculated by an exhaustive search, and then compared to each of the methods above in terms of the quality of the resulting plans.

Let $\underline{\alpha} = (\alpha_1, \dots, \alpha_k)$ be an action plan which consists of k actions, where k equals the number of subgoals. According to the rollback formula, the level of

attainment of the major goal, $V_G(\underline{\alpha})$ becomes:

$$V_G(\underline{\alpha}) = \sum_{i=1}^k W_i E(\alpha_i) \prod_{\substack{j=1 \\ j \neq i}}^k [1 - S(\alpha_i, j)]$$

where

$E(\alpha_i)$ = direct effect of action α_i .

$S(\alpha_i, j)$ = side effect of action α_i on subgoal j .

With this convention, the selection criteria becomes:

$$\underline{\alpha}_{opt} = \underset{\substack{\text{all feasible} \\ \text{action combinations} \\ \underline{\alpha}_i}}{\text{Max}^{-1}} [V_G(\underline{\alpha}_i)]$$

The selection of the optimal action plan, in each simulation step, was immediately followed by selection of the action plans using the local optimization procedure. For each subgoal, a representative action was constructed according to each of the four alternative methods. Then, the local optimal action L_i under each subgoal i was calculated as:

$$L_i = \underset{j}{\text{Max}^{-1}} [V_G(H_1, \dots, H_{i-1}, \alpha_{ij}, H_{i+1}, \dots, H_k)]$$

where the α_{ij} 's are actions under subgoal i , and H_i is the representative action for subgoal i . The local optimal action plan is then:

$$\hat{\underline{\alpha}}_{opt} = (L_1, \dots, L_k)$$

The degree of merit of each representative action choice was calculated according to the following criterion:

$$d(\underline{\alpha}) = \frac{V_G(\underline{\alpha}) - V_G(\text{Worst Action Combination})}{V_G(\text{Optimal Action Combination}) - V_G(\text{Worst Action Combination})}$$

thus, $1 - d(\underline{\alpha})$ is proportional to the degree of suboptimality of the action plan. The value of $[1 - d(\underline{\alpha})]$ ranges between zero and one, where zero represents an optimal plan.

The simulation results are portrayed in Figures A-2, A-3, A-4, and A-5 for the alternative methods (1), (2), (3), and (4), respectively. Comparison of the results indicates that method (3) is superior to the others.

To investigate the performance of the local optimization procedure in multi-level goal-directed graphs, a simulation of a graph with sixteen actions appearing in two different levels was performed. The same three ranges (small: 0-0.33, medium: 0.33-0.66, and large: 0.66-1) were used for the value of the direct effects and side effects. To calculate the effects, it is not only necessary to define virtual actions, as in the single-level case, but also to define virtual subgoals. The virtual subgoals are used to reduce some subgraphs to linear structures, thus fixing the subgraph values. For example, the subgraph emerging from subgoal S_2 in Figure A-6 is reduced to the corresponding linear structure in Figure A-7. With this reduction, the values of the actions under subgoal S_1 are calculated under the assumption that virtual actions H_2 and H_{22} occur under S_2 .

The process of deriving Figure A-7 from Figure A-6 is outlined in detail in Figure A-8. First, virtual actions (H_i 's) are constructed for the bottom level actions, as described for the one-level graph as shown in Figure A-6(a) to A-6(b). Then, using the calculated values of the virtual actions, the lower level subgoals values are found followed by the calculation of the effectiveness of the next higher actions (A_3 and A_4). At this point, it is necessary to calculate the value for the level of attainment of the lower level subgoals (S_6 through S_9). The values of the compressed subgoals (HS_{67} and HS_{89}), emerging from a common action, are set to a value equal to the product of the values of the




Figure A-2. Simulation Result for the First Representative Action Selection Method




Figure A-3. Simulation Result for the Second Representative Action Selection Method




Figure A-4. Simulation Result for the Third Representative Action Selection Method




Figure A-5. Simulation Result for the Fourth Representative Action Selection Method

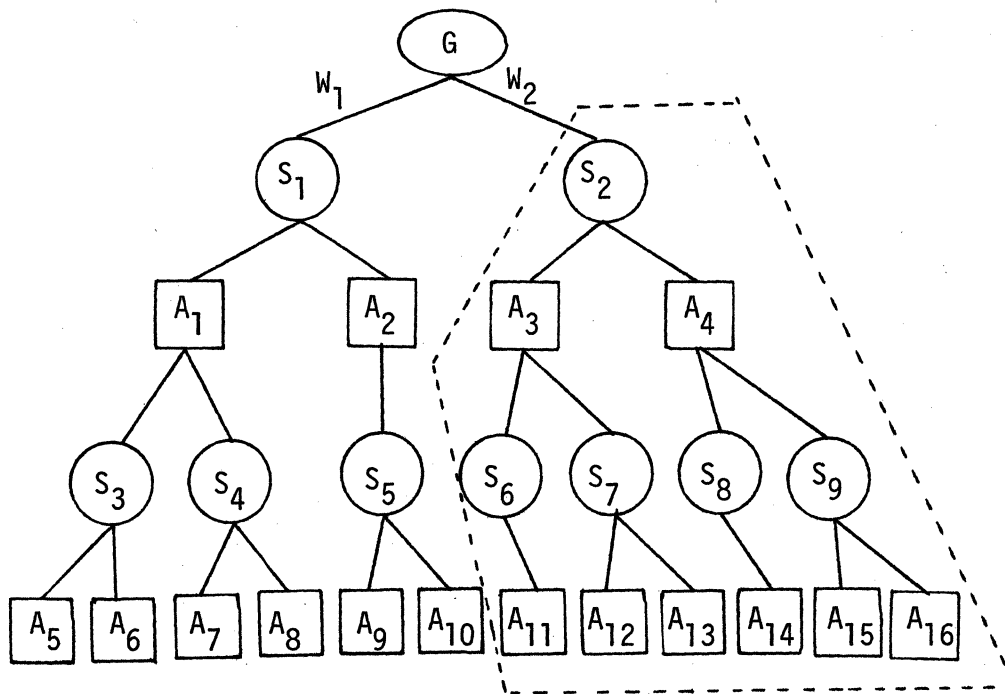


Figure A-6. A Sample Two-Level Goal-Directed Graph
(Side Effects Are Not Shown on the Graph for Simplicity)

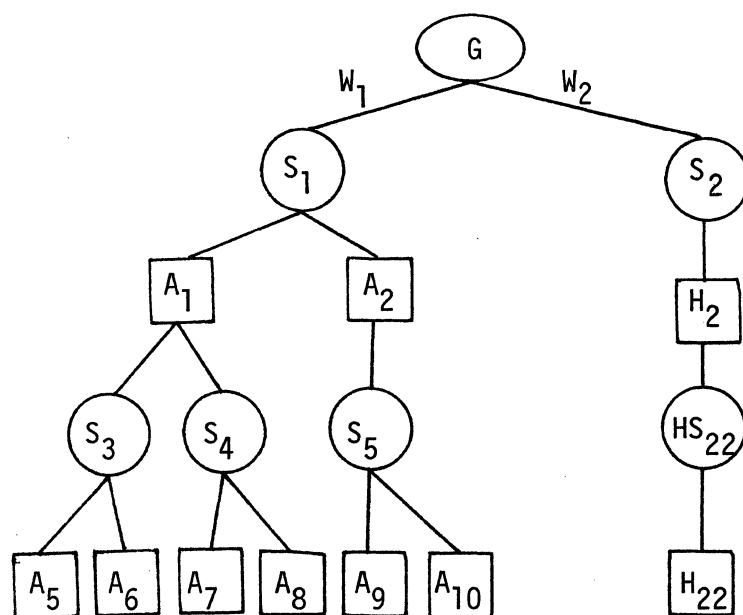


Figure A-7. Goal-Directed Graph Reduction for Identifying
Local Optimal Action Under S_1

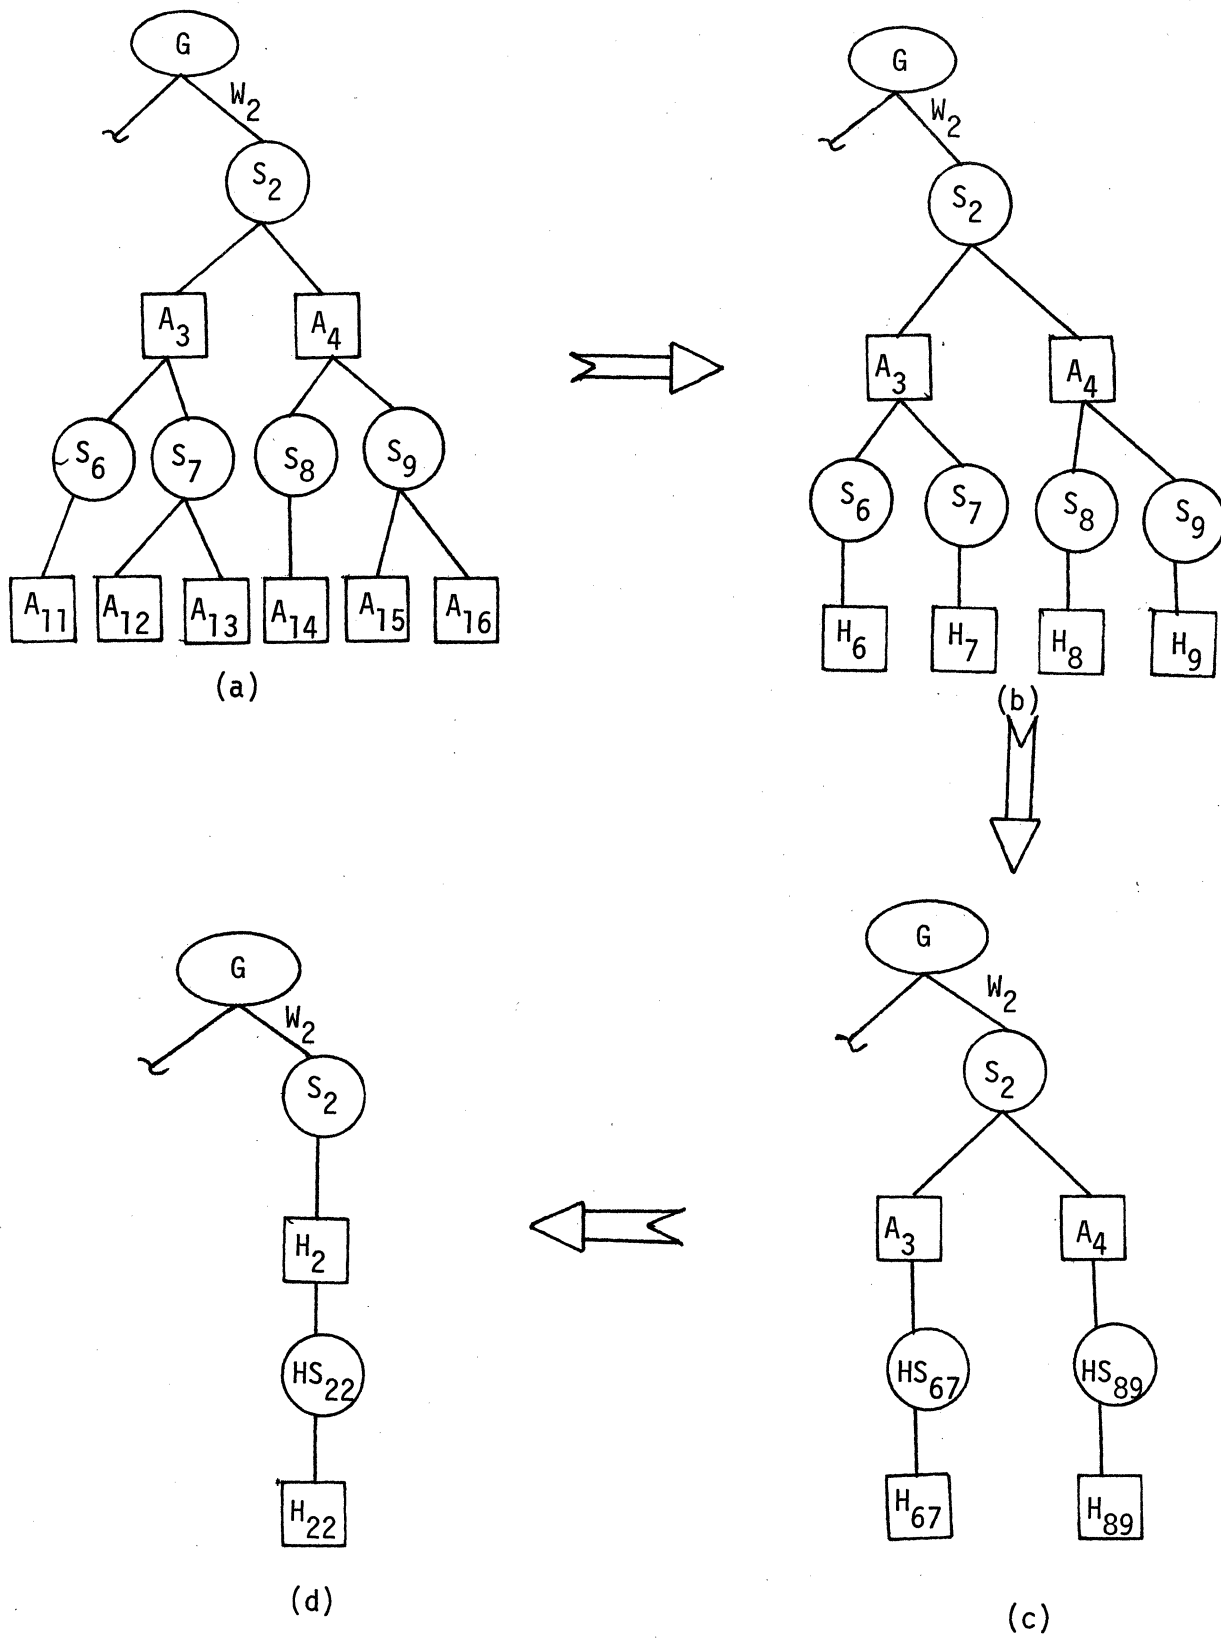


Figure A-8. Graph Compression Process

original subgoals (transition from Figure A-8(b) to A-8(c)). This is equivalent to the exact calculation under the assumption that δ -functions of intermediate preconditions are all straight lines emerging from the origin with a 45 degree slope. The process can be shown in the structure by replacing the subgoals with a virtual subgoal possessing the calculated value. Note that during the simulation, since the values of the higher level actions are being generated randomly, there is no need to calculate the value of the virtual subgoals because random values can be generated directly at the subgoal level. Therefore, the simulation was not based on assuming any particular δ -function. The virtual subgoals are used to facilitate aggregation of side effects. The side effects of the new virtual actions (H_{67} and H_{89} in Figure A-8) are calculated as follows (transition from Figure A-8(c) to A-8(d)).

$$S(H_N) = 1 - \prod_{i=1}^k [1 - S(H_{oi})]$$

where H_N is combined virtual action and the H_{oi} 's are virtual actions being combined.

By repeating the process, the structure in Figure A-6 is converted to the one shown in Figure A-7. Then, local optimal action plans are calculated in the manner defined earlier. As in the case of one-level graphs, the results also confirmed the superiority of the third representative action selection method.

APPENDIX B. SENSITIVITY ANALYSIS

B.1. Calculation of Δ

Due to the characteristics of the δ - functions used in the propagation of the impact of preconditions and the use of the Max-function in the calculation of the level of attainment of subgoals based on the effectiveness of actions, the exact relationship between the subgoals and the major goal can be described by piecewise linear functions. Examples of piecewise linear functions governing the propagation rules in a goal-directed graph are shown in Figure B-1.

In Figure B-1, the sensitivity of G with respect to terminal subgoal (S_4) is being examined. In the sensitivity analysis phase of the program we assume that the side effects of actions on the path leading to this terminal subgoal (i.e., out of A_3) are not affected by variations in its value. Side effects are activated only during the optimization phase when the corresponding action is chosen as a result of resulting in the highest utility over the action-set. Thus, the side effects of A_3 do not contribute to variations of $G(V^*)$.

With this understanding, the direct effect E_3 of action A_3 , as a function of all possible values V_4 of its successor subgoal S_4 , can be drawn as shown in Figure B-1(a). The resulting variation in the level of completion of precondition P_3 is the same as the variation in V_4 . In this case, action A_3 has only one precondition (P_3), so its effectiveness varies according to the corresponding δ -function (δ_3). The level of attainment V_3 of the predecessor subgoal S_3 is equal to the maximum of the direct effects of actions A_3 and A_4 (i.e., E_3 and E_4) multiplied by I_3 , the product of the side effects of all actions impacting V_3 . Since it is assumed that all other subgoals in the graph except S_4 remain constant, the direct effect of action A_4 is constant and equal to E_4 . Therefore, V_3 as a function of V_4 can be plotted as in Figure B-1(c). For small values of V_4 (which result in small values of E_3) the Max-function selects E_4 which, after




Figure B-1. Propagation of Piecewise Linear Functions in a Goal-Directed Graph

multiplication by I_3 (the product of side effects of remote actions on S_3), yields the value V_3 . As V_4 increases, V_3 remains constant until the value of V_4 becomes sufficiently large to cause $E_3 > E_4$. From this point on, the value of V_3 will be equal to $E_3 \cdot I_3$. Since E_3 with respect to V_4 has a linear relationship with slope α_3 , the relationship between V_3 and V_4 , from the angular point on, will be linear with slope $\alpha_3 \cdot I_3$.

The graph V_3 versus V_4 summarized the variations in the level of attainment of a subgoal with respect to the variations in the level of attainment of a successor subgoal located at one level lower. In the same manner, the variations in the level of attainment of the subgoal S_1 (V_1), with respect to variations in V_3 , can be calculated and drawn for one subgoal-level higher in the graph. The only difference is that the action A_1 has more than one precondition (P_1 and P_2). Therefore, the value of its direct effect is the result of the product of the δ -function corresponding to preconditions P_1 and P_2 ($\delta_1 \cdot \delta_2$). However, since the values of other subgoals remain constant, the level of completion L_2 of precondition P_2 remains constant. Therefore, $\delta_2(L_2)$ will be a constant number modifying the function δ_1 as a multiplicative factor. Thus, E_1 versus V_3 will be of the same form as δ_1 with only the slope changed from α_1 to $\alpha_1 \cdot \delta_2(L_2)$. The Max-function used in computing the value of V_1 , and the product of side effects of actions impacting S_1 , then have the same effects as described earlier, therefore, resulting in a V_1 function versus V_3 as shown in Figure B-1(f).

The propagation rule at each subgoal-level is in the form of canonical functions y_1 and y_2 shown in Figure B-2. Each terminal subgoal is thus related to the major goal through a path defined by the cascading series of functions of the same form. The calculation of impact propagation through the graph is substantially simplified by the fact that the cascade of two such function results in a third function of the same form. For example, the cascade of the two





Figure B-2. Two Typical Propagation Functions, y_1 and y_2 ,
Representing Different Subgoals
of Goal-Directed Graphs

functions, y_1 and y_2 in Figure B-2, gives the function of Figure B-3 (if

$a_2 \leq \frac{a_1}{\alpha_1} + c_1$) or the one in Figure B-4 (if $a_2 \geq \frac{a_1}{\alpha_1} + c_1$). Thus, after specifying

the propagation functions of all subgoal-levels on the path from the terminal subgoal in question to the major goal (G), the propagation functions along the path can be cascaded recursively, to express the remote relationship between the major goal and the terminal subgoal. At the highest level the resulting cascaded propagation function will be modified multiplicatively by the weight of the corresponding subgoal (e.g., w_1 in Figure B-1) and additively by the weighted impact of other subgoals (e.g., $w_2 v_2$).

The value of Δ is derived directly from the cascaded propagation function. The point on the abscissa (Figure B-5) corresponding to the function angular point b is equal to $V + \Delta$. Therefore, subtracting the provisional value V yields $\Delta = b - V$.

B.2. Vector Representation of Propagation Functions

The canonical function (Figure B-6) can be characterized by a three element vector, $\begin{pmatrix} a \\ b \\ \alpha \end{pmatrix}$, where a is equal to the constant level, $b = \frac{a + \alpha c}{\alpha}$ is the point on abscissa corresponding to the angular point, and α is the slope after the angular point. These vectors undergo four transformations when an influence is propagated through the graph. They are denoted by four operators on the characteristic vectors: (1) aggregation, (2) maximum, (3) cascade, and (4) summation.

The 'Aggregation' operator is used for three purposes: (1) to aggregate the impact of the preconditions on the effectiveness of the corresponding predecessor action, (2) to aggregate the impact of side effects and the direct effect of actions on the level of attainment of the corresponding predecessor subgoal, and (3) to aggregate the impact of a subgoal on the major goal with its corresponding weight. The aggregation operator, denoted by $\text{---}\rightarrow$, takes a canonical




Figure B-3. Cascade of the Two Propagation Functions,

$$y_1 \text{ and } y_2, \text{ if } a_2 \leq \frac{a_1}{\alpha_1} + c_1$$




Figure B-4. Cascade of the Two Propagation Functions,

$$y_1 \text{ and } y_2, \text{ if } a_2 \geq \frac{a_1}{\alpha_1} + c_1$$




Figure B-5. Derivation of Δ From Cascaded Propagation Function




Figure 4-2. Parametric Representation of the Canonical Functions

function and a constant and results in another canonical function:

$$\rightarrow \left[d, \begin{pmatrix} a \\ b \\ \alpha \end{pmatrix} \right] = \begin{pmatrix} a \cdot d \\ b \\ a \cdot \alpha \end{pmatrix}$$

It represents the multiplication of a canonical function by a constant.

The 'Maximum' operator is used to express the level of attainment of sub-goals based on the effectiveness of the corresponding actions. Since the effectiveness of all actions except one remain constant, the maximum operator works on a canonical function and a number of constants and results in another canonical function. The Maximum operator is denoted by $\rightarrow \text{Max} \leftarrow$:

$$\rightarrow \text{Max} \leftarrow \left[d, \begin{pmatrix} a \\ b \\ \alpha \end{pmatrix} \right] = \begin{cases} \begin{pmatrix} d \\ \frac{d + \alpha b - a}{\alpha} \\ \alpha \end{pmatrix} & \text{if } a < d \\ \begin{pmatrix} a \\ b \\ \alpha \end{pmatrix} & \text{if } a > d \end{cases}$$

It represents the maximum between a constant and a canonical function.

The 'Cascade' operator represents the cascading of functions along the path to the major goal. This binary operator takes two canonical functions and results in the cascade of the two functions which is also in canonical form. Since $f_1(f_2(x))$ is not equal to $f_2(f_1(x))$, the order of the arguments in a cascade operation is significant. If the vectors $\begin{pmatrix} a_1 \\ b_1 \\ \alpha_1 \end{pmatrix}$ and $\begin{pmatrix} a_2 \\ b_2 \\ \alpha_2 \end{pmatrix}$ represent the canonical functions $f_1(x)$ and $f_2(x)$, then, using \uparrow to denote the cascade operation $f_1(f_2(x))$, we have:

$$\begin{pmatrix} a_1 \\ b_1 \\ \alpha_1 \end{pmatrix} \uparrow \begin{pmatrix} a_1 \\ b_2 \\ \alpha_2 \end{pmatrix} = \begin{cases} \begin{pmatrix} a_1 \\ 1/\alpha_2 (b_1 - a_2) + b_2 \\ \alpha_1 \alpha_2 \end{pmatrix} & \text{if } a_2 \leq b_1 \\ \begin{pmatrix} \alpha_1 (a_2 - b_1) + a_1 \\ b_2 \\ \alpha_1 \alpha_2 \end{pmatrix} & \text{if } a_2 \geq b_1 \end{cases}$$

The 'Summation' operator is used at the highest level of the goal-directed graph to add the impact of the other first-level subgoals to the impact of the one under consideration with the varying level of attainment. (Since the level of attainment of all terminal subgoals, except one, is fixed, all other first-level subgoals can be represented by a constant). This operator takes a canonical function and a constant and results in another canonical function. The summation operator is denoted by Σ :

$$\Sigma \left[d, \begin{pmatrix} a \\ b \\ \alpha \end{pmatrix} \right] = \begin{pmatrix} a + d \\ b \\ \alpha \end{pmatrix}$$

The four operators are used to propagate the canonical functions represented in vector format. However, since the second element of the final resulting vector corresponds to the angular point and is equal to $V + \Delta$, Δ is readily obtained by the difference $V - b$.

B.3. Calculation of the Expected Value-of-Analysis (EVA)

If the relationship between the major goal and a given subgoal is characterized by a piecewise linear function such as the one depicted in Figure B-6, then the function ΔG would depend on which side of the angular point (b) V is found.

For $V < b$ (i.e., the subgoal is not part of the optimal plan) we obtain:

$$\Delta G(V^*) = \begin{cases} 0 & V^* < b \\ \alpha(V^* - b) & V^* > b \end{cases}$$

For $V > b$ (i.e., the subgoal is a part of the optimal plan) we have:

$$\Delta G(V^*) = \begin{cases} 0 & V^* > b \\ \alpha(b - V^*) & V^* < b \end{cases}$$

Consequently the EVA is given by:

$$EVA = E_{V^*} \Delta G(V^*) = \begin{cases} P(V^* > b) E(\Delta G | V^* > b) & \text{for } V < b \\ P(V^* < b) E(\Delta G | V^* < b) & \text{for } V > b \end{cases}$$

and, since $b = V + \Delta$, we have:

$$\begin{aligned} EVA &= \begin{cases} P(V^* > V + \Delta) \alpha E(V^* - V - \Delta | V^* > V + \Delta) & \text{for } \Delta > 0 \\ P(V^* < V + \Delta) \alpha E(V + \Delta - V^* | V^* < V + \Delta) & \text{for } \Delta < 0 \end{cases} \\ &= \begin{cases} \alpha p (m - V - \Delta) & \text{for } \Delta > 0 \\ \alpha p (V + \Delta - m) & \text{for } \Delta < 0 \end{cases} \end{aligned}$$

where p stands for the probability that a fluctuation in V^* would result in a decision switch, and $m = E(V^* | \text{a decision switch})$.

REFERENCES

- Burns, Michael and Judea Pearl, "Experiments in Cognitive Decomposition," UCLA-ENG-CSL-7951, School of Engineering and Applied Science, University of California, Los Angeles, August 1979.
- Ernst, George W. and Allen Newell, GPS: A Case Study in Generality and Problem Solving, Academic Press, New York, 1969.
- Fikes, Richard E. and Nils J. Nilsson, "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving," Artificial Intelligence, Vol. 2, pp. 184-208, 1971.
- Kim, Jin, "A Graphical System for Evaluating Decision Aids," UCLA-ENG-7915, School of Engineering and Applied Science, University of California, Los Angeles, March 1979.
- Leal, Antonio, "An Interactive Program for Conversational Elicitation of Decision Structures," UCLA-ENG-REP-7666, Ph.D. Dissertation, School of Engineering and Applied Science, University of California, Los Angeles, June 1976.
- Leal, Antonio and Judea Pearl, "An Interactive Program for Conversational Elicitation of Decision Structures," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-7, No. 5, pp. 368-376, May 1977.
- Leal, Antonio, Steven Levin, Steven Johnston, Marcy Agmon, and Gershon Weltman, "An Interactive Computer Aiding System for Group Decision Making," Technical Report PQTR-1046-78-2, Perceptronics, Woodland Hills, California, February 1978.
- Merkhofer, Miley W., Allen C. Miller, III, Burke E. Robinson, and Robert J. Korsan, "Decision Structuring Aid: Characterization and Preliminary Implementation," Stanford Research Institute, Menlo Park, California, September 1977.
- Nilsson, Nils J., Problem Solving Methods in Artificial Intelligence, McGraw-Hill, New York, 1971.
- Pearl, Judea, "A Goal-Directed Approach to Structuring Decision Problems," UCLA-ENG-7811, School of Engineering and Applied Science, University of California, Los Angeles, February 1978.
- Slovic, Paul, Baruch Fischhoff, and Sarah Lichtenstein, "Behavioral Decision Theory", Annual Review of Psychology, Vol. 28, p. 23, 1977.

DISTRIBUTION LIST

CDR Paul R. Chatelier
Office of the Deputy Under Secretary
of Defense
OUSDRE (E&LS)
Pentagon, Room 3D129
Washington, D.C. 20301

Director
Engineering Psychology Programs
Code 455
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217 (5 cys)

Director
Analysis and Support Division
Code 230
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Navy Analysis Programs
Code 431
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Operations Research Programs
Code 434
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Statistics and Probability Program
Code 436
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Information Systems Program
Code 437
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Special Assistant for Marine
Corps Matters
Code 100M
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Commanding Officer
ONR Branch Office
ATTN: Dr. J. Lester
Building 114, Section D
666 Summer Street
Boston, MA 02210

Commanding Officer
ONR Branch Office
ATTN: Dr. C. Davis
536 South Clark Street
Chicago, IL 60605

Commanding Officer
ONR Branch Office
ATTN: Mr. L. Lawson
1030 East Green Street
Pasadena, CA 91106

Commanding Officer
ONR Branch Office
ATTN: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Office of Naval Research
Scientific Liaison Group
American Embassy, Room A-407
APO San Francisco, CA 96503

Director
Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D.C. 20375 (6 cys)

Dr. Bruce Wald
Communications Sciences Division
Code 7500
Naval Research Laboratory
Washington, D.C. 20375

Dr. Robert G. Smith
Office of the Chief of Naval
Operations, OP987H
Personnel Logistics Plans
Washington, D.C. 20350

Naval Training Equipment Center
ATTN: Technical Library
Orlando, FL 32813

Human Factors Department
Code N215
Naval Training Equipment Center
Orlando, FL 32813

Dr. Gary Poock
Operations Research Department
Naval Postgraduate School
Monterey, CA 93940

Dean of Research Administration
Naval Postgraduate School
Monterey, CA 93940

Mr. Warren Lewis
Human Engineering Branch
Code 8231
Naval Ocean Systems Center
San Diego, CA 92152

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
Code RD-1
Washington, D.C. 20380

HQS, U.S. Marine Corps
ATTN: CCA40 (MAJOR Pennell)
Washington, D.C. 20380

Commanding Officer
MCTSSA
Marine Corps Base
Camp Pendleton, CA 92055

Chief, C³ Division
Development Center
MCDEC
Quantico, VA 22134

Mr. Arnold Rubinstein
Naval Material Command
NAVMAT 08T24
Washington, D.C. 20360

Commander
Naval Air Systems Command
Human Factors Programs
NAVAIR 340F
Washington, D.C. 20361

Commander
Naval Air Systems Command
Crew Station Design
NAVAIR 5313
Washington, D.C. 20361

Mr. Phillip Andrews
Naval Sea Systems Command
NAVSEA 0341
Washington, D.C. 20362

Naval Sea Systems Command
Personnel & Training Analyses Office
NAVSEA 074C1
Washington, D.C. 20362

Commander
Naval Electronics Systems Command
Human Factors Engineering Branch
Code 4701
Washington, D.C. 20360

Scientific Advisor to DCNO (MPT)
OP 01T (Dr. Marshall)
Washington, D.C. 20370

Dr. George Moeller
Human Factors Engineering Branch
Submarine Medical Research Lab
Naval Submarine Base
Groton, CT 06340

Navy Personnel Research and
Development Center
Manned Systems Design, Code 311
San Diego, CA 92152

Navy Personnel Research and
Development Center
Code 305
San Diego, CA 92152

Navy Personnel Research and
Development Center
Management Support Department
Code 210
San Diego, CA 92152

CDR P. M. Curran
Code 604
Human Factors Engineering Division
Naval Air Development Center
Warminster, PA 18974

Mr. Ronald A. Erickson
Human Factors Branch
Code 3175
Naval Weapons Center
China Lake, CA 93555

Human Factors Engineering Branch
Code 1226
Pacific Missile Test Center
Point Mugu, CA 93042

Dean of the Academic Departments
U.S. Navy Academy
Annapolis, MD 21402

LCDR W. Moroney
Code 55MP
Naval Postgraduate School
Monterey, CA 93940

Mr. Merlin Malehorn
Office of the Chief of Naval
Operations (OP 102X)
Washington, D.C. 20350

Mr. J. Barber
HQS, Department of the Army
DAPE-MBR
Washington, D.C. 20310

Dr. Joseph Zeidner
Technical Director
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Director, Organizations and
Systems Research Laboratory
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Edgar M. Johnson
Organizations and Systems Research
Laboratory
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Technical Director
U.S. Army Human Engineering Labs
Aberdeen Proving Ground, MD 21005

U.S. Army Aeromedical Research Lab
ATTN: CPT Gerald P. Krueger
Ft. Rucker, AL 36362

ARI Field Unit-USAREUR
ATTN: Library
C/O ODCSPER
HQ USAREUR & 7th Army
APO, NY 09403

U.S. Air Force Office of Scientific
Research
Life Sciences Directorate, NL
Bolling Air Force Base
Washington, D.C. 20332

Dr. Donald A. Topmiller
Chief, Systems Engineering Branch
Human Engineering Division
USAF AMRL/HES
Wright-Patterson AFB, OH 45433

Air University Library
Maxwell Air Force Base, AL 36112

Dr. Gordon Eckstrand
AFHRL/ASM
Wright-Patterson AFB, OH 45433

North East London Polytechnic
The Charles Myers Library
Livingstone Road
London E15 2LJ
ENGLAND

Professor Dr. Carl Graf Hoyos
Institute for Psychology
Technical University
8000 Munich
Arcisstr 21
FEDERAL REPUBLIC OF GERMANY

Dr. Kenneth Gardner
Applied Psychology Unit
Admiralty Marine Technology
Establishment
Teddington, Middlesex TW11 0LN
ENGLAND

Director, Human Factors Wing
Defence & Civil Institute of
Environmental Medicine
Post Office Box 2000
Downsview, Ontario M3M 3B9
CANADA

Dr. A. D. Baddeley
Director, Applied Psychology Unit
Medical Research Council
15 Chaucer Road
Cambridge, CB2 2EF
ENGLAND

Defense Documentation Center
Cameron Station, Bldg. 5
Alexandria, VA 22314 (12 cys)

Dr. Stephen J. Andriole
Director, Cybernetics Technology
Office
Defense Advanced Research Projects
Agency
1400 Wilson Blvd.
Arlington, VA 22209

Dr. Judith Daly
Cybernetics Technology Office
Defense Advanced Research Projects
Agency
1400 Wilson Blvd.
Arlington, VA 22209

Professor Douglas E. Hunter
Defense Intelligence School
Washington, D.C. 20374

Dr. Robert R. Mackie
Human Factors Research, Inc.
5775 Dawson Avenue
Goleta, CA 93017

Dr. Gary McClelland
Institute of Behavioral Sciences
University of Colorado
Boulder, CO 80309

Human Resources Research Office
300 N. Washington Street
Alexandria, VA 22314

Dr. Miley Merkhofer
Stanford Research Institute
Decision Analysis Group
Menlo Park, CA 94025

Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army-Navy Drive
Arlington, VA 22202

Professor Howard Raiffa
Graduate School of Business
Administration
Harvard University
Soldiers Field Road
Boston, MA 02163

Dr. Arthur I. Siegel
Applied Psychological Services, Inc.
404 East Lancaster Street
Wayne, PA 19087

Dr. Paul Slovic
Decision Research
1201 Oak Street
Eugene, OR 97401

Dr. Harry Snyder
Department of Industrial Engineering
Virginia Polytechnic Institute and
State University
Blacksburg, VA 24061

Dr. J. A. Swets
Bolt, Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Amos Tversky
Department of Psychology
Stanford University
Stanford, CA 94305

Dr. Gershon Weltman
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, CA 91364

Dr. Alphonse Chapanis
Department of Psychology
The Johns Hopkins University
Charles and 34th Streets
Baltimore, MD 21218

Dr. Meredith P. Crawford
American Psychological Association
Office of Educational Affairs
1200 17th Street, N.W.
Washington, D.C. 20036

Dr. Ward Edwards
Director, Social Science Research
Institute
University of Southern California
Los Angeles, CA 90007

Dr. Charles Gettys
Department of Psychology
University of Oklahoma
455 West Lindsey
Norman, OK 73069

Dr. Kenneth Hammond
Institute of Behavioral Science
University of Colorado
Room 201
Boulder, CO 80309

Dr. Ronald Howard
Department of Engineering-Economic
Systems
Stanford University
Stanford, CA 94305

Dr. William Howell
Department of Psychology
Rice University
Houston, TX 77001

Journal Supplement Abstract Service
American Psychological Association
1200 17th Street, N.W.
Washington, D.C. 20036 (3 cys)

Dr. Clinton Kelly
Decisions and Designs, Inc.,
8400 Westpark Drive, Suite 600
P.O. Box 907
McLean, VA 22101

Mr. Edward M. Connelly
Performance Measurement
Associates, Inc.
131 Park Street, N.W.
Vienna, VA 22180

Dr. Edward R. Jones
Chief, Human Factors Engineering
McDonnell-Douglas Astronautics
Company
St. Louis Division
Box 516
St. Louis, MO 63166