

Counterplanning: A Strategy-Based Model of Adversary Planning in Real-World Situations*

Jaime G. Carbonell

*Computer Science Department, Carnegie-Mellon University,
Pittsburgh, PA 15213, U.S.A.*

Recommended by N.S. Sridharan

ABSTRACT

Plan formulation and plan comprehension in adversary situations constitute integral components of human reasoning. A process model based on heuristic strategies is presented for decision making in obstructive and constructive counterplanning situations. The former situation is characterized by an actor striving to thwart the goals and plans of a second actor. The latter is the dual situation; it provides general means for an actor to pursue his goals in spite of attempts by others to block his initial plans. The model has been implemented as part of the POLITICS system, a computer program that understands brief natural language accounts of international political conflicts.

1. Introduction

Computer simulation of human decision-making strategies provides a useful means of modeling, and thereby coming to a better understanding of the mechanisms underlying this central aspect of human cognition. We focused our research on situations where the decision maker must contend with other parties who may be continuously trying to thwart his efforts. In such circumstances, formulating plans and deciding how and when to implement one's plans may become much more complex processes than in simple one-actor planning situations. The primary problems that must be addressed in such situations include the non-deterministic nature of planning under uncertain outcomes and changing circumstances, predicting the most likely actions on the

* This research was sponsored in part by the Defense Advanced Research Projects Agency (DOD), Monitored by the Office of Naval Research under contract N00014-75-C-1111, and in part by ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory under contract F33615-78-C-1511. The views and conclusions contained in this document are those of the author, and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

part of potential adversaries, formulating alliances with other parties sharing common goals, and compromising with one's adversaries when necessary. In order to make the simulation problem more tractable we assume that the counterplanner has acquired internal models of his potential adversaries encoding their primary motivations. We discuss our theory of counterplanning and its implementation within the **POLITICS** framework. **POLITICS** is a computer program that models ideologically-oriented decision-making tasks in the domain of international diplomacy and political confrontations.

Computer systems that model various facets of human reasoning abound in the Artificial Intelligence literature. Models of human decision-making, however, have been mostly confined to the study of problem-solving situations (see, e.g., [10, 12]) and plan interpretation in story understanding tasks (see, e.g., [14, 15, 17]). These research efforts did not focus on simulating decision-making in situations where there are several parties actively scheming to thwart each other's goals and block each other's plans while simultaneously furthering their own goals. We developed the counterplanning strategies precisely for the purpose of modeling such complex interactions among the plans and goals of several competing parties.

2. The **POLITICS** System

The **POLITICS** system [3, 2] understands a natural language description of a conflict situation by abstracting its meaning in terms of contextual knowledge of similar conflict situations (encoded as scripts [7, 14]), determining how the situation affects the goals of each political actor, and later applying counterplanning strategies to model the decisions open to the political actors under the circumstances defined by the input conflict situation. **POLITICS** evaluates possible decisions in terms of their effects upon the goals of the political actor, and their likelihood for success (e.g., whether the planned decision can be implemented, or whether it is likely to be blocked by an opposing political actor).

In the following example the **POLITICS** program analyzes a political event from the point of view of a conservative American policy maker. Later, it answers several questions to illustrate its goal-oriented decision making process.

```
*(INTERPRET US-CONSERVATIVE)
INCORPORATING US-CONSERVATIVE IDEOLOGY GOAL TREES...
INITIALIZING PARSER... DONE.
```

```
INPUT STORY: +The US Congress is expected to approve the
               Panama Canal treaty.
```

```
PARSING... COMPLETED.
```

```
INSTANTIATING SCRIPT: $PARLIAMENT-PROC
```

EXPECTED TRACK:

⇔((\$VOTE VOTERS (#ORG NAME (CONGRESS) PARTOF
 (*US*))
 ISSUE (#TREATY NAME (PANAMA CANAL)
 REF (TREATY03))
 SIGNEES (*US* *PANAMA*))
 OUTCOME (*POS*)))

CANAL-ZONE FROM (*CONTROL* VAL (*US*))
 TOWARD (*CONTROL* VAL (*PANAMA*))

US GOAL VIOLATED:

US MILITARY STRENGTH DECREASED.

US GOAL VIOLATED:

ACTOR (#COUNTRY TYPE (*COMMUNIST*))

⇔ (*MTRANS*)

MOBJ ((ACTOR (*US*) TOWARD (*STRENGTH* TYPE
 (*MILIT*))
 INC (-2)))

RUSSIA GOAL ACTIVATED.

ACTOR (*RUSSIA*) TOWARD (*SCONT* OBJ (*PANAMA*))
 INC (*POS*)

INTERPRETATION COMPLETED, READY TO ACCEPT
 QUESTIONS.

Q1: Should the US approve the treaty?

A1: NO, THE TREATY IS BAD FOR THE UNITED STATES.

Q2: Why is the treaty bad for the US?

A2: THE UNITED STATES WOULD LOSE THE CANAL TO
 PANAMA.

Q3: What might happen if the United States loses the Canal?

A3: RUSSIA WILL TRY TO CONTROL THE CANAL.

Q4: Why would Russia try to take over the Canal?

A4: RUSSIA WANTS TO EXPAND ITS MILITARY CONTROL.
 THE PANAMA CANAL HAS HIGH MILITARY VALUE.

Q5: How might this happen?

A5: THE UNITED STATES MAY NOT BE ABLE TO STOP ANYONE TAKING OVER THE PANAMA CANAL. TORRIJOS IS A COMMUNIST. TORRIJOS WILL LET RUSSIA CONTROL PANAMA.

Q6: Why should the US worry about Russia taking over the Canal?

A6: BECAUSE RUSSIA WANTS TO CONTROL THE WORLD.

Q7: What should the US do to stop Russia from taking over the Canal?

A7: CONGRESS SHOULD VOTE AGAINST THE PANAMA CANAL TREATY.

In the above example, *POLITICS* inferred that congressional approval means voting in favor of a resolution, which in turn means that the provisions of the resolution will come to pass. Since the key provisions of the Panama Canal treaty are stored in memory, *POLITICS* focuses its attention on their effect upon the goals of the US (as defined by the US conservative ideology). *POLITICS* concludes that two important US goals may be violated: Military strength will decrease, and ability to deter Russian aggression will suffer.

In the question-answering phase, counterplanning strategies focus on the affected goals in order to model possible US decisions. These strategies create probable hypothetical scenarios (as illustrated in the answer to question 5), and subsequently decide on means to bring about or prevent the hypothetical scenario. For instance, in answering question 7, *POLITICS* attempted to prevent the possible Russian takeover of Panama, because such a takeover violates the goals of the US. *POLITICS* has no absolute measure of the likelihood that any hypothetical situation will become reality. The counterplanning process suggests plans of action that make the realization of possible situations more (or less) probable. Hence, there is a notion of relative likelihood of hypothetical outcomes, mediated by the plans of the various actors.

The questions answered by the *POLITICS* system require an inference process that relates the input event to factual information contained in memory, as well as a decision-making process whose focus is to further the goals of the United States as perceived by a conservative American policy maker. Although the *POLITICS* system is essentially an integrated understanding system, it can be divided conceptually into several modules, including: (1) Natural language understanding and generation in a semantically-rich domain. (2) A process to focus attention and thereby constrain inferences based on the goal-hierarchy representation of political ideologies. (3) A script and situational inference rule applier. (4) A planning system based on our heuristic model of counterplanning. This paper focuses only on the counterplanning process, a significant aspect of human reasoning that has not heretofore been directly addressed in Artificial Intelligence.

Planning and counterplanning require self-knowledge of one's motivations, and at least partial knowledge of the goals of other actors with whom one interacts. In *POLITICS*, the goals of all relevant actors form the bulk of the ideological-belief model. For instance, a model of a conservative American asserts that the primary Soviet goal is world domination, and that the primary US goals are communist containment and a strong national defense. The *POLITICS* model of a US liberal, however, states that the primary US and Soviet goals are the maintainment of world peace, and that humanitarian goals are more important than a strong defense. Because of the varying sets of goals and the different priorities among the goals, *POLITICS* is able to model different political ideologies using the same reasoning process. This reasoning process consists primarily of the counterplanning heuristics focused on fulfilling predefined ideological goals in novel situations.

3. The Counterplanning Control Process

The counterplanning process is encoded as a set of heuristic strategies applicable to general conflict situations, and a control-flow algorithm that determines when to apply the various classes of counterplanning strategies. We first discuss the control-flow algorithm; later we analyze the structure and content of the heuristic strategies.

There are two general types of conflict situations where an actor may apply the counterplanning process. The first situation is characterized by an actor (X) trying to thwart another actor (Y) from achieving his goal G(Y). X may prevent Y from achieving G(Y) by directly making the goal state impossible to achieve, or by repeatedly blocking Y's plans to fulfill G(Y). We call this process *obstructive counterplanning*. The second type of counterplanning scenario is essentially the dual of the first: An actor X may try to achieve his goal G(X) in spite of potential attempts by Y to nullify the goal state G(X) or block X's plans for pursuing G(X). This process, called *constructive counterplanning*, differs from obstructive counterplanning only in terms of the subjective viewpoint of the counterplanner. The perspective shift causes the counterplanner to apply different strategies at different times, as the application of strategies is goal driven. Fig. 1 is a control-flow diagram for the obstructive counterplanning process.

3.1. Obstructive counterplanning

To illustrate the obstructive counterplanning mechanism, consider a simple example. Prison guard X wants to prevent prisoner Y from escaping. Hence, G(Y) is the state "Y is free outside the prison". We enter Fig. 1 at the top. Does the guard know the prisoner's escape plan? Let us assume that he does not. His next step is to determine what plan, if any, the prisoner formulated. This plan determination may itself involve some planning: Should the guard

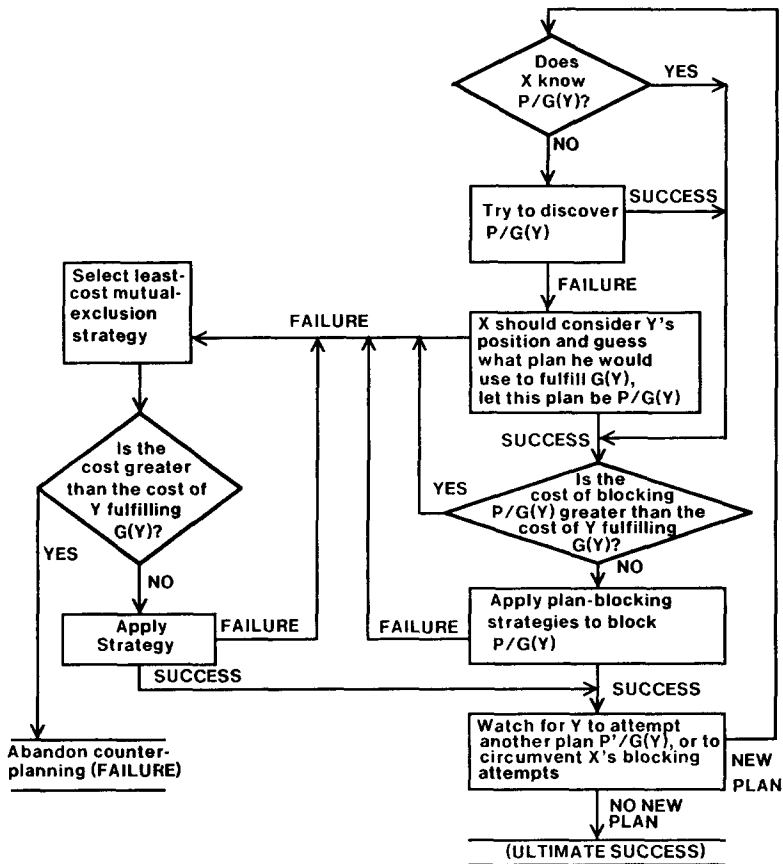


FIG. 1. Control structure of the counterplanning process (obstructive phase). Actor X attempts to block actor Y from achieving his goal $G(Y)$. $P/G(Y)$ is Y's plan for fulfilling $G(Y)$.

ask the prisoner? Should he threaten him? (See [14] for a discussion of social planning units). Let us again assume that the guard fails. At this point we enter the third box in Fig. 1. The guard can ask himself "what would I do if I was trying to escape?" If he finds a reasonable plan, he may hypothesize that this is the prisoner's plan and he would apply the obstructive counterplanning strategies to block that plan. For instance, the guard may find that stealing the key is a reasonable plan, in which case he may apply the violate-necessary-precondition strategies (discussed later) to conclude that he should keep the keys away from the prisoner.

If no plan presents itself, the only option open to the guard is to take general precautions (i.e., apply mutual exclusion strategies) such as pointing a gun at the prisoner and informing him that either he remains put, or he will be dead.

Thus escaping and staying alive become mutually-exclusive states in the prisoner's mind.

There are exit conditions in the counterplanning algorithm. For instance, if the prisoner's plan is to blow up the prison (and he has the means to do so), the guard may decide that the risk of being blown up is more costly than his goal of thwarting the prisoner's escape. Similarly, if he has to keep a gun trained on the prisoner for the length of the prison sentence, the guard may decide that this is more costly than letting the prisoner escape. Cost measures in counterplanning are discussed briefly in a later section.

3.2. Constructive counterplanning

Fig. 2 is the control-flow for constructive counterplanning. Let us consider the

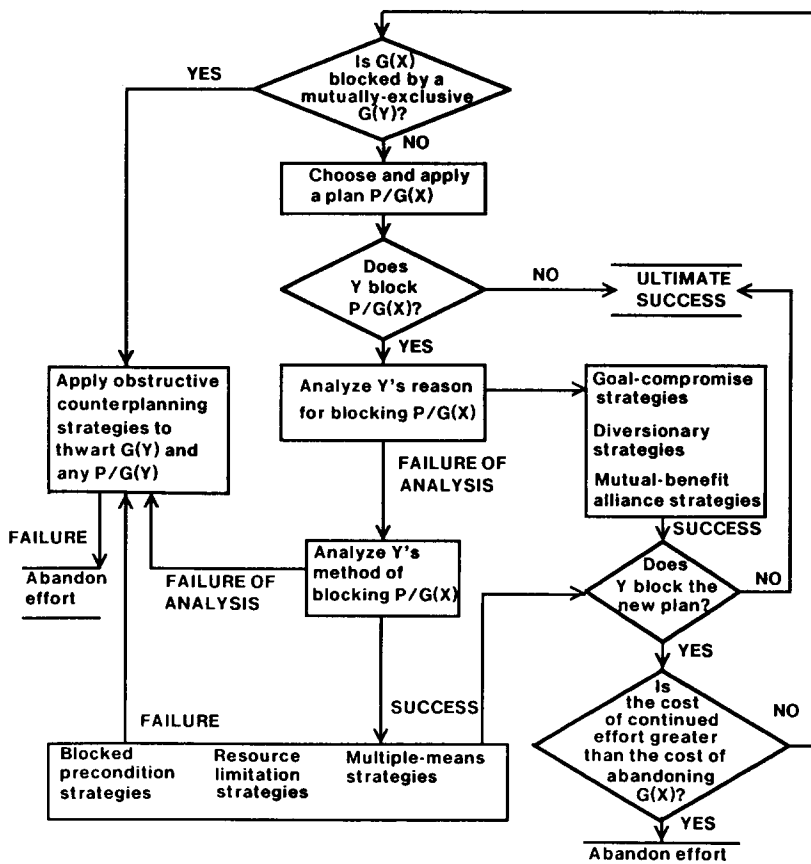


FIG. 2. Control structure of the counterplanning process (constructive phase). Actor X wants to achieve G(X) in spite of Y's attempts to thwart his plans or prevent his ultimate goal.

same situation from the point of view of the prisoner. We enter the first box in Fig. 2: Is the prisoner's escape goal directly blocked by the guard? If this is the case, the only recourse open to the prisoner is obstructive counterplanning against the guard's goal. For instance, if the guard has a gun trained on his head, the prisoner must address this problem before formulating any specific escape plan.

If his goal is not directly blocked by a mutual-exclusion state, the prisoner can then address the problem of escaping. If he is able to formulate a workable plan and the guard does not counterplan, the prisoner succeeds. Otherwise he must address the guard's obstructive counterplanning. First, the prisoner can analyze the reasons for the guard's counterplanning. For instance, if the reason is that the guard needs his job to make money, the prisoner may try bribing the guard (a mutual-benefit strategy, in effect forming a temporary alliance with the guard). If he discovers a more important goal that the guard pursues and he is able to threaten this goal, the guard may be diverted from his counterplanning efforts. For instance, the prisoner may convince the guard (truthfully or not) that he is about to blow up the prison and thereby induce the guard to pursue his higher-level goal of staying alive by fleeing.

Finally, the prisoner may pursue a different set of counterplanning strategies. These strategies depend upon his analysis of the method used by the guard to counterplan against him. For instance, the guard may be blocking a precondition (e.g., keeping the keys away from the prisoner) or limiting an essential resource such as time, materials, or outside assistance. There are heuristic strategies for attempting to overcome each type of blocking action. Constructive counterplanning also has exit conditions based on the relative cost of abandoning or continuing the effort.

4. The Structure of the Counterplanning Strategies

Our model of counterplanning is essentially a goal-directed rule-based system. Each strategy is a rule that tests the goal conflict state between the two actors, the goal trees of each actor, and the plan which each actor is pursuing. If the test of a rule is true, the action part of the rule suggests a counterplanning method that is likely to succeed. The action part of the rule may have additional conditions or refinements. The structure of a counterplanning strategy is illustrated below.

```

STRATEGY <rule#k>(title)
  TRIGGER Conjunction of <test1, test2, . . . , testn>
    IF Conjunction of <testn + 1, testn + 2, . . . >
      THEN Sequence of <action1, action2, . . . >
  REFINEMENT <subrule#k1, subrule#k2, . . . >

```

FIG. 3. The structure of a counterplanning strategy.

The test part of each rule is divided into the 'TRIGGER' clause and an additional 'IF' clause. In order for a rule to apply, both clauses must be true in the counterplanning situation. The reason for our division of the test clause in such a manner is to reduce the time the process model spends searching for an applicable strategy. All the trigger conditions are 'inexpensive' tests; these tests can be applied directly to the situation without requiring further inference or complicated matching. The trigger conditions are compiled into a discrimination network. This organization allows the addition of new strategies without a corresponding linear increase in the search time required to test all the applicable counterplanning strategies.

Once the trigger conditions for a rule have been met, the additional 'IF' tests are performed. These tests may be arbitrarily complex, perhaps requiring further inference and the invocation of other counterplanning strategies. The trigger conditions usually restrict the set of applicable counterplanning strategies to a small number in any given situation (typically three or four). Moreover, counterplanning strategies reflect common-sense reasoning about how to deal with adverse situations. As such, each strategy is sufficiently general to apply across many reasonable types of human conflict. This means that the total number of counterplanning strategies is relatively small compared to the total number of rules and information contained in the situation-specific scripts. We have found that approximately forty to fifty counterplanning strategies suffice to model the counterplanning actions in the situations and goal conflicts we have observed.

The action part of each strategy (the 'THEN' clause) is a sequence of counterplanning methods to be applied in the current context by one of the actors. If a strategy is applied, the 'REFINEMENT' field is checked before the sequence of actions is performed. The refinement contains one or more additional rules that usually provide further information useful to the counterplanning process. These rules are truly subrules to the counterplanning strategies, as they are invoked only in the case that all the tests of the strategy are true and, in addition, the test clause of each subrule is also true. In structure, our strategies are much closer to the rules in expert systems (such as PECOS [1] and MYCIN [16]) than to pure production systems (e.g. PSG [11] or OPS [8]).

The following sections analyze some of the more significant types of counterplanning strategies. A more complete set is discussed in [3]. We also discuss the way in which counterplanning strategies are used in text-understanding and question-answering tasks.

5. Diversionary Counterplanning Strategies

Consider some general principles underlying counterplanning strategies applicable to mutual-exclusion goal conflicts. A frequently-encountered set of strategies operate on the principle of diverting the efforts of an actor involved

in the goal conflict away from direct pursuit of his goal. There are essentially three classes of diversionary strategies, all relying on the fact that if an actor must divert his efforts to other matters, he is less likely to succeed at his original task. We group the diversionary strategies into the three categories listed below:

Diversionary strategies. (1) Threaten higher-level goals of one's opponent.

(2) Dissipation of effort—Threaten opponent's other goals.

(3) Deception—Misleadingly convince one's opponent that other goals are threatened.

Let us look at one important diversionary strategy and its deceptive option. The process of convincing one's opponent that the actions of a counterplanning strategy have been carried out, without this being the case, is called the *deceptive option* of that strategy.

Strategy 1. (Threaten higher-level goal)

Trigger Mutual-exclusion goal-conflict situation between $G(Y)$ and $G(X)$.

If X can find a goal $G'(Y)$ to block, where $G'(Y)$ is a goal of higher importance to Y than $G(Y)$,

Then X should try to block $G'(Y)$. X can expect Y to pursue $G'(Y)$ and abandon $G(Y)$.

Refinement X should choose $G'(Y)$ such that

- (1) Y cannot pursue both $G'(Y)$ and $G(Y)$ simultaneously, or
- (2) Y can accomplish $G'(Y)$ only if the blocking action is abandoned.

Strategy 1 means that in a mutual-exclusion goal-conflict situation, one actor (X) may threaten a higher-level goal of the second actor (Y), in order to make Y divert his efforts to preserve that higher-level goal. Since an actor's attention, time, and material resources are limited, Y may not be able to protect his higher-level goal and pursue the goal that conflicts with X 's goal simultaneously. Thus X will be in a better position to resolve the conflict situation in his favor. The 'refinement' field of the rule gives some advice to help X make his choice on which one of Y 's higher-importance goals to threaten. X knows which goals Y may consider more important, by examining Y 's goal tree. Our model assumes that the various actors know about each other's primary motivations. X should threaten a goal that requires Y 's full attention, time, or material resources to protect. Alternatively, X can threaten an important goal that Y cannot protect. This gives X a bargaining position to tell Y that he will stop his threat only if Y abandons the (presumably less important) goal that

conflicts with X's goal. Bargaining strategies are discussed in greater detail later in this paper.

Let us see the deceptive version of Strategy 1. This strategy is based on the same principle of diverting the attention and efforts of the opposing counter-planner away from the original goal conflict. However, the method used need not correlate with reality. As such, one should expect that a deceptive strategy is more likely to fail (e.g., if the opposing actor discovers the truth).

Strategy 2 (Deception: False threat)

Same as Strategy 1 but either

- (1) X only appears to threaten the higher level goal $G'(Y)$, or
- (2) X threatens to block $G'(Y)$ but has falsely convinced Y that $G'(Y)$ should be one of Y's high-level goals.

Strategy 2 is rather ubiquitous in everyday life. Bluffing and making threats with no intent to carry them out are some common applications of counter-planning by means of invoking the deceptive option. Strategy 2 is also useful in understanding some novel situations such as the following story.

Event 1. John and Bill were competing for Mary's attentions at a party. John, noticing that Bill was more successful, went to the telephone next door and called the host to tell him that Bill's house was on fire. Regretably, Bill had to leave the party in great haste.

In our society, John's and Bill's goals of courting Mary are mutually exclusive; therefore the counterplanning strategies, including Strategy 2, are applicable. Bill's higher-level goal of preserving his house, belongings, and possibly his family overrides further immediate consideration of his previous goal. The fact that Bill's house was not on fire is relevant only to the likelihood of success of Strategy 2. It is possible that Bill may see through the ruse by, for instance, calling his house or calling his neighbors.

There are other diversionary strategies (listed in [3]); however, they all share a common principle: An actor cannot overtax his attention and resources by simultaneously pursuing multiple courses of action. Awareness of this simple principle applied to other actors (as well as the counterplanner) guides our formulation of the diversionary counterplanning strategies. Each strategy is based on a different method of causing an actor to worry about more than one goal at a given time. We classify the limitations on the simultaneous pursuit of multiple courses of action into the following categories.

- (1) Limitations on available time.
- (2) Mental and physical limitations on the number of actions which can be performed simultaneously.
- (3) Limitations of material resources instrumental to one's plan.

- (4) Limitations of physical or mental ability.
- (5) Goal of avoiding certain consequences of one's actions.
- (6) Interactions between different courses of actions.

6. Dual Processes: Plan Interpretation and Plan Generation

Before continuing our analysis of specific counterplanning strategies and the principles on which they are based, let us step back and consider the general counterplanning and plan interpretation processes. Here we present a global view of the process-model encoded in the entire *POLITICS* system, rather than the more specific method-oriented discussion of Section 3.

Individual counterplanning strategies may be used for plan interpretation as well as plan generation. In our discussion thus far, we assumed the latter task: Given the goals of the planner, those of his adversaries, and a potential conflict situation, formulate a plan to achieve the planner's goals in light of expected adversity. The plan interpretation *process* differs from the plan generation process, but not the *knowledge base* of strategies, goal trees and situational scripts. The plan interpreter is given a sequence of actions taken by one or more actors in a particular conflict situation. However, the interpreter may have only *partial* knowledge of the conflict situation, and must *infer* which goals (among those in the actors' goal trees) are being pursued.

The interpretation of a sequence of actions as a coherent plan could proceed by generating potential plans from partial knowledge, matching these plans against the given sequence of actions, and trying anew upon failure of match. However, this application of the generate-and-test method is neither computationally efficient nor a reasonable model of human plan interpretation. Instead, *POLITICS* applies the counterplanning strategies to generate a small number of *general predictions*, coupled with a wait-and-see heuristic. Some general predictions may be confirmed and instantiated into specific plans by subsequent information. The key idea is to generate few predictions at the appropriate level of abstraction. This general paradigm for plan interpretation is not new to *POLITICS*. (E.g. the PAM system [17], SAM [7], the 'hypothesize-and-revise' method [15], and frame-driven comprehension [6] contain variations on the common theme: predict a general structure, and depending upon later evidence, instantiate, revise, or discard the general prediction.) *POLITICS* extends the general plan-interpretation process to multi-party adversary-planning scenarios.

In essence, plan interpretation consists of discovering the goals which motivate the observed actions of purposeful actors by means of inferring relevant intermediate stages in their planning process. *POLITICS* takes advantage of the counterplanning strategies and goal trees for the dual processes of planning and plan interpretation. The knowledge encoded in the strategies and goal trees is essential to constrain search in the multi-actor adversary-planning

space. Hence much of the complexity introduced in extending planning and plan interpretation from a single-actor situation to the more complex adversary-planning scenario is balanced by the introduction of our new knowledge sources.

Let us turn to an overview of the planning and plan interpretation processes in the context of the *POLITICS* system, and give an illustration of how the plan interpretation process applies to a conflict situation.

6.1. Synopsis of the plan generation process

The following steps give an overview of the plan generation process.

(1) Interpret the natural language description of the conflict situation¹ instantiating scripts and querying memory to flesh out information left implicit in the original description.

(2) Identify the actors involved in the conflict situation, selecting their activated or threatened goals from their goal trees.²

(3) Call the plan interpretation process, if necessary, to infer active goals and goal interactions among adversary actors not stated in the input description.

(4) Apply constructive and/or obstructive counterplanning methods (discussed in Section 3) to invoke the appropriate counterplanning strategies, and formulate a viable plan of action.

(5) When applying counterplanning strategies, use the discrimination net of static 'TRIGGER' tests to focus on the few strategies relevant to the current situation. Then apply the first strategy whose 'test' field is satisfied.

6.2. A synopsis of the plan interpretation process

An overview of the plan interpretation process is given by:

(1) Interpret the initial description of the conflict situation, in the same manner as in the plan formulation process.

(2) Identify the actors involved and retrieve their goal trees. Check whether their actions match scripts where the goals of the instantiated script roles are known. If so, and if the scriptal goals are present in the actor's goal tree, activate those goals.

(3) Apply the counterplanning-strategy discrimination net to the partially formulated conflict situation in order to focus on the strategies which may have been applied by the actors to formulate their plans (whose instantiations consist of the observed actions).

(4) Match the general *action fields* of the activated strategies to the observed

¹ This process is described in [3, 4], where other aspects of the *POLITICS* system are discussed, including integrated parsing and ideological reasoning.

² See [3, 5] for a discussion of goal trees and goal identification.

actions of the actors. If the latter can be interpreted as instantiations of the actions suggested by the counterplanning strategies (i.e., if the match succeeds), hypothesize that the corresponding *test* fields of the matched strategies were true. These are the general predictions.

5. If the predictions are consistent with the actor's goal tree, and if they are borne out and further instantiated by subsequent actions, then the plans of the actors involved have been interpreted (and their goals discovered, as the test clauses of strategies relate high-level goals to general action frameworks).

6.3. An example of counterplan interpretation

The processes outlined above illustrate the dual use of counterplanning strategies: (1) Plan formulation by focusing on the strategies relevant to the conflict situation, matching on the test fields of these strategies, and suggesting the contents of the corresponding action fields as general plans. (2) Plan interpretation by, once again, focusing on the relevant strategies, but now matching on the action fields, and suggesting the contents of the test fields as a rationalization for the observed actions (as long as each rationalization is consistent with the known goals of the actors involved).

Now, we turn to an example of the plan interpretation process. Consider the following scenario.

Event 2. The United States supported Israel in the 1973 Arab-Israeli war. Subsequently the Arabs imposed an oil embargo on the United States, and its allies.

Question. Why did the Arabs impose the oil embargo?

POLITICS interprets the first sentence, identifying the US, Israel, and Arab Nations as the primary actors. Moreover, in instantiating the war script and the military-aid script, POLITICS infers that the Arabs and Israelis have an active goal of winning the war, and the US has the goal of Israel winning the war. Many of the inferred goals are pairwise mutually-exclusive. Therefore Strategy 1 is activated as possibly relevant to the scenario (along with a few other strategies not yet discussed). When the second sentence is encountered, POLITICS tries to see if it can be interpreted as part of the action field of Strategy 1. In fact, the action with refinement (2) matches the type of interaction between the Arab plan and the US goal. The Arabs are threatening a high-level US preservation goal by cutting off oil supplies, and the US cannot do anything to directly remedy the situation. Now the Arabs are in a position to bargain for to end the embargo in return for the US ending aid to Israel. Therefore, the understander can establish the Arab counterplan by applying Strategy 1 in an interpretive manner. The *result* of the Arab actions is matched to the *action*

part of the strategy. This match, suggested by our previous expectation that mutual-exclusion counterplanning strategies may be relevant, allows the understander to infer that the Arab plan was formulated by invoking counterplanning Strategy 1. Therefore, we can say that the *reason for the Arab action is the test clause of the strategy*.

It would be unreasonable to expect a human understander to predict an Arab embargo, given only the first sentence in Event 2, plus general world-knowledge. In fact, before and during the 1973 Arab-Israeli war, few people foresaw an Arab oil embargo. When the embargo came, however, the reasons for the Arab action became clear. Therefore, the plan interpretation process must cope with unforeseen actions, just as a planner must be capable of formulating effective plans in specific unforeseen conflict situations.

Now, having shown the duality of the general planning and plan interpretation processes, let us resume our discussion of individual counterplanning strategies, which enable these processes to proceed in an effective manner.

7. Counterplanning Strategies Based on Goal Compromise

A significant class of counterplanning strategies in mutual-exclusion goal conflicts involves various forms of accommodation with one's adversaries. These strategies are characterized by actors in mutual opposition who are willing to compromise. Willingness to compromise is mediated by many factors such as the two adversaries' perceptions of whether compromise is possible, necessary, or desirable.

There are two basic classes of bargaining strategies. Each class is characterized by the type of compromise the counterplanning actors are willing to accept. Let $G(X)$ be X 's goal and $G(Y)$ be Y 's goal. As before, suppose these two goals define a mutual-exclusion conflict.

Bargaining strategies

- (1) Partial fulfillment of $G(X)$ and/or $G(Y)$.
- (2) Goal substitution of $G(X)$ and/or $G(Y)$ by another goal not involved in the conflict.

The first class assumes the possibility of either goal in the goal conflict being partially fulfilled. Let us define 'partial fulfillment' of a goal with the aid of some illustrative examples.

Event 3. John was a very ambitious salesman. He wanted to become president of General Petroleum Corporation.

Case 1. John was fired and blacklisted for unethical business practices.

Case 2. John tried very hard and was eventually appointed vice-president of General Petroleum.

Case 3. John's meteoric rise in the corporate structure culminated in the presidency of the company.

In Case 1 John failed to achieve his goal. In Case 3 John clearly fulfilled his goal. What can we say about Case 2? Strictly speaking, John did not achieve his goal of becoming president. However, John did not totally fail in fulfilling his ambition. We classify Case 2 as partial fulfillment of his goal. John achieved something less than the presidency of the company, but, by any reasonable measure of success, he succeeded in obtaining some of the power, prestige, and wealth associated with the presidency of a company.

The following examples are instances where there can be no partial fulfillment of a goal.

Event 4. The New York Yankees wanted to beat the Boston Red Sox. The Yankees were leading until the bottom of the ninth when Rice won the game for Boston with a grand slam.

Event 5. Hubert Humphrey came very close to winning the 1968 presidential elections, but lost to Nixon.

In Event 4 the Yankees failed to fulfill their goal. The fact that they almost won cannot be considered partial fulfillment, as there are no intermediate states between winning and losing a baseball game. The same argument applies to Event 5. There was no real partial fulfillment of Humphrey's goal, no matter how close he came to winning the election.

What is the crucial difference between Event 3 and the two following examples? Becoming president of a company is a complex goal subsuming several simpler goals such as achieving social respect, power, and wealth. (Wilensky [17] discusses the phenomenon of goal subsumption.) Success on each one of these simpler goals is measured on a continuum rather than on an all-or-nothing criterion. Therefore there are two measures of partial success in achieving a goal. The first measure applies to a complex goal that subsumes several other goals. We define success to be the case where all the subsumed goals are fulfilled, and failure to be the case where none of the subsumed goals are fulfilled. Partial fulfillment is defined in the obvious manner: the goal is partially fulfilled if some of the subsumed goals are fulfilled. For instance consider the following case.

Event 6. John wanted to marry Susan. They decided to live together instead.

Did John achieve his goal? Strictly speaking the answer is 'no'. But, if we understand that marriage is a complex goal, we realize that for all intents and purposes John fulfilled most of the goals subsumed by marriage, such as

achieving companionship, periodic satisfaction of the sex drive, etc. Since marriage also involves a change in social and legal status not necessarily associated with living together, we say that John partially fulfilled his goal. Therefore, by this measure, partial fulfillment of a goal is the case where the specific goal sought is not fulfilled, but a significant fraction of the underlying reasons for pursuing the goal is fulfilled.

The second way in which partial goal fulfillment is measured applies to goal states that range over a continuum of values. Acquisition of money and achievement of social stature are examples of continuous-valued goals: there are virtually infinite degrees of social stature and of the amount of money a person can acquire. We define success differently for preservation goals than for achievement goals (called 'P-goals' and 'A-goals' respectively—Schank and Abelson [14] define the goal taxonomy). Let us call the *initial* value of the goal state 'I', the *desired* value of the goal state 'D', and its *resultant* value at the time when we must decide whether the goal has been attained, 'R'. For P-goals it is usually the case that $I = D$, and for A-goals $I < D$. The success and failure conditions of continuous-valued goals are given by the following table:

TABLE 1. Partial fulfillment of goal states

	Failure	Partial success	Success
A-Goal	$I = R < D$	$I < R < D$	$I < R = D$
P-Goal	$R < I = D$	$R \text{ 'slightly' } < I = D$	$R = I = D$

Let us consider how the partial-success table applies to an example of related A-goals and P-goals.

Event 7. Somalia wanted to conquer the Ogaden region in Ethiopia. The Somalis launched an invasion.

Case 1. The Somalis conquered the entire Ogaden.

Case 2. The Ethiopian army checked the Somali invasion, yielding only a border sector and several villages.

Case 3. Ethiopia beat back the Somali attack, defeating the invading army.

From a Somali perspective, Event 7 is an attempt to fulfill an A-goal, achieving military control of the Ogaden region, previously under Ethiopian control. Hence the present state: I = no control over the Ogaden, and the desired state: D = controlling the entire Ogaden region. In Case 1 the Somalis achieved a military control state over the Ogaden, resulting in state $R = D > I$. Therefore the Somalis fully succeeded. In Case 2 the achieved Somali goal is military control of a (probably small) part of the Ogaden region. Thus $I < R < D$:

the Somalis partially succeeded. In Case 3 Somalia achieved nothing relevant to their goal. Therefore $I = R < D$: total failure.

We can interpret Event 7 from an Ethiopian perspective. The primary Ethiopian goal is to preserve their sovereignty over the Ogaden region: a P-goal with $I = D$ = Ethiopian control of the Ogaden. In Case 1 the Somali control the Ogaden: a mutually exclusive state with Ethiopian control of the same region. Thus $R < I = D$, meaning that the Ethiopian P-goal failed. In Case 2 Ethiopia controls most but not all of the Ogaden; therefore R is slightly less than I and D : our condition for partial success. In Case 3 Ethiopia fulfills its goal of maintaining sovereignty over the Ogaden: $R = I = D$.

Let us turn our attention to partial-goal-fulfillment and goal-substitution bargaining strategies to determine how the nature of the goal conflict affects their application.

Strategy 3 (Goal compromise)

Trigger $G(X)$ and $G(Y)$ are mutually exclusive and may be fulfilled partially.

If X cannot achieve $G(X)$ by other counterplanning strategies,

Then X should try to partially fulfill $G(X)$ by bargaining with Y to compromise mutually on partially fulfilling their respective goals.

Metaphorically, Strategy 3 states that if one cannot have the entire pie, then one should try to bargain for at least a slice of the pie. The following example illustrates partial fulfillment of goals as a result of a compromise.

Event 8. Professors Smith and Jones kept blocking each other's efforts to become chairman of their department. Eventually Smith agreed to support Jones in the next election, with the stipulation that Jones support Smith in the following election.

Strategy 3 is invoked when other measures fail; one need not compromise if one can fulfill one's initial goal. To understand Event 8, one must realize that the very existence of a goal conflict is preventing either actor from making further progress towards his goal. In this case, the understander should expect compromise as the only reasonable course of action. The compromise in Event 8 is that each actor fulfills his acquisition of power/prestige goal, but only for a limited time.

Strategy 4 is more cooperative in nature than Strategy 3, but involves the same principle of compromise on partially attainable goals.

Strategy 4 (Cooperation by mutual need)

Trigger $G(X)$ and $G(Y)$ are mutually exclusive and may be partially fulfilled.

If neither X nor Y can independently achieve their respective goals,
but can succeed only by pooling their efforts,
Then X and Y should compromise on partially fulfilling G(X) and G(Y),
and plan jointly for their fulfillment.

Strategy 4 states that cooperation may be a necessary course of action in spite of conflicting goals. Event 9 illustrates this point.

Event 9. Jesse James and Bill Morgan joined forces to heist the payroll train.

We infer that both actors had the A-goal of acquiring the money in the payroll train. These goals are mutually exclusive but may be partially fulfilled. Suppose we had to answer the question: "Why did Jesse James and Bill Morgan join forces?" The most reasonable answer is: "Probably because neither could heist the train by himself." This answer cannot be inferred from the goal conflict itself; the existence of a goal conflict suggests competitive rather than cooperative actions. Therefore the understander has to be aware of strategy 4—cooperation between actors with conflicting goals is reasonable if neither actor can otherwise fulfill his goal. Strategy 4 predicts that Jesse and Bill will partially fulfill their A-goals; i.e., they will split the take. This prediction accords with our intuition of what normally happens in this type of situation.

The following event is another example of forced cooperation. Without the compromise of sharing government power neither party would achieve its goal of governing Italy.

Event 10. In Italy neither the Christian-Democrats nor the Communist Party were able to form a majority government. They formed an unprecedented coalition to govern the country.

Strategies 3 and 4 have their deceptive options, defined similarly to the deceptive option of Strategy 1.

Strategy 5. (Deception: False compromise)

Same as Strategy 3, but X can pretend to agree on compromise (to divert Y's efforts) while pursuing some other means of counterplanning.

Strategy 6. (Deception: False cooperation)

Same as Strategy 4, but after X and Y have mutually fulfilled their common goal, X can counterplan against Y with the purpose of blocking G(Y) and totally fulfilling G(X).

The deception can be applied to Events 8 and 9 as follows.

Event 8 (possible continuation). Jones supported Smith's successful bid for the chairmanship, but prior to the next election Smith bad-mouthed Jones.

Event 9 (possible continuation). After the successful train heist, Jesse James shot Bill Morgan, and kept the entire loot.

The interpretation process for Event 8 assumes that both Jones and Smith invoke Strategy 3. The use of the word 'but' in the continuation of Event 8 suggest that an expectation has been violated. (Schank [13] describes the 'but test': a heuristic method to determine what is inferred in understanding a text). Since bad-mouthing Jones violates Smith's part of the bargain, the violated expectation can be interpreted as a signal to the understander that Smith has invoked Strategy 5. The bargain was only a ruse to prevent Jones from blocking Smith's A-goal.

Understanding the continuation of Event 9 also requires one to realize that deception has been invoked. The result of the cooperation between Jesse and Bill was partial fulfillment of their respective goals, but then Jesse counterplanned (unbeknownst to Bill) to fulfill completely his A-money goal at Bill's expense. It is interesting to note that if the understander of Event 9 knows about Jesse James's goals and subgoals (i.e., has a reasonably detailed goal-subgoal importance hierarchy of the type discussed in [3, 5]), the understander may expect treachery on Jesse's part. Hence the goal hierarchies of the actors help to determine which counterplanning strategy is more likely to be invoked by the actors in a goal-conflict situation.

8. Counterplanning Strategies Based on Goal Substitution

In some goal-conflict situations, a useful class of bargaining strategies is based on goal substitution. There are different types of goal substitution, as discussed in [14]. Here we are concerned with how goal substitution is invoked in bargaining strategies. Let us consider an example of goal substitution.

Event 11. Johnny and Billy were arguing over who would get to ride their new bicycle. Johnny said, "If you let me ride it, I'll give you my candy bar." Billy promptly agreed.

In Event 11 Billy substituted eating a candy bar for his goal of riding the bicycle. Moreover, Johnny was aware that Billy might be amenable to this goal substitution. Johnny used his knowledge of Billy's goals to propose the bargain which resolved their goal conflict. The general bargaining strategy invoked by Johnny is the following.

Strategy 7 (Goal substitution)

Trigger $G(X)$ and $G(Y)$ are competing goals and X knows about Y 's other goals.

If X can bring about $G'(Y)$, one of Y 's goals which is at least as important as $G(Y)$,

Then X should bargain with Y to substitute $G'(Y)$ for $G(Y)$ as Y 's actively pursued goal. (This leaves X free to pursue $G(X)$.)

Refinement Apply this strategy only if $G(Y)$ is not a P-goal.

Strategy 7 is usually more applicable if Y 's goal in the conflict is an A-goal, rather than a P-goal, for two reasons: First, P-goals are usually more important; thereby making X 's task of finding a $G'(Y)$ of at least equal importance more difficult. Second, it is more difficult for one to sacrifice something one already has achieved than to abandon an A-goal, even in the case where the latter may be of more importance. For instance, it is usually not the case that a person will abandon his job in order to spend his time applying for another, possibly higher-paying job. It is more likely that he will retain his current job (P-goal) and, as time permits, apply for a better job (A-goal).

Strategy 7, being asymmetrical with respect to X and Y , has two deceptive options:

Strategy 8 (Deception: False substitution)

Same as Strategy 7, but X does not bring about $G'(Y)$ after Y abandons $G(Y)$.

Strategy 9 (Deception: False acceptance)

Same as Strategy 7, but Y re-establishes $G(Y)$ (thereby blocking $G(X)$) after X brings about $G'(Y)$.

Another type of a goal-substitution bargaining strategy involves mutual abandonment of both $G(X)$ and $G(Y)$. This strategy is invoked in the case the continued goal conflict itself violates more important goals for both X and Y than the original conflict goals.

Strategy 10 (Mutual goal abandonment)

Trigger $G(X)$ and $G(Y)$ are mutually exclusive and X also has the goal of $\text{NOT}(G(Y))$.

If continued counterplanning efforts may violate higher-level goals than $G(X)$ and $G(Y)$,

Then X should negotiate with Y to abandon mutually $G(X)$ and $G(Y)$. (Failing that X may simply choose to abandon pursuit of $G(X)$.)

Strategy 10 plays an important role in the POLITICS domain. For instance,

consider the following event processed by *POLITICS* from a US-liberal perspective.

Event 12. The United States Congress voted to fund the Trident submarine project.

Question: What might Russia do next?

Answer: Russia will probably build more submarines.

Question: What should the US do if Russia also builds nuclear submarines?

Answer: The United States should negotiate with Russia to stop the arms race.

Strategy 10 is invoked to answer the second question. From a US-liberal perspective both the United States and the Soviet Union have the high-importance goal of preserving world peace. Continued goal conflict on which country is militarily stronger (i.e., the arms race) could violate the higher-level P-goal of preserving peace. Since both countries also want to keep the other country from becoming the dominating world power, Strategy 10 suggests negotiation to abandon bilaterally the goals of attaining military superiority.

9. Understanding the Cause of a Plan Conflict

Thus far, we have discussed strategies for blocking goals and for circumventing blocked goals. The execution of a plan can also be blocked, intentionally or accidentally. In this section we analyze the reasons underlying plan conflicts and the types of counterplanning strategies this analysis suggests.

Let us categorize the knowledge an understander should look for in a plan-conflict situation in order to suggest or understand possible counterplanning measures. There are essentially six significant dimensions in most plan conflicts. These essentially orthogonal dimensions may be encoded as numerical scales, since we require them as measures of comparison (e.g., among the expected level of cooperation between actors in different circumstances, and the relative importance of the goals underlying a conflict). We propose ranking plan-conflict situations along the following six dimensions.

Dimension 1. Accidental vs. intentional plan conflict.

Dimension 2. Competitive vs. cooperative predispositions toward resolving the conflict.

Dimension 3. Interference as a mere inconvenience vs. total goal blockage.

Dimension 4. Mutual-exclusion goal conflict underlies the plan conflict vs. ultimate goal agreement but conflict on the means used to bring about the goal.

Dimension 5. Externally imposed conflicts vs. internally motivated ones.

Dimension 6. Very important goal threatened as a result of the plan conflict vs. insignificant goal threatened.

The classification of plan-interference conflicts along these dimensions has a

two-fold effect on the understanding process. First, this classification facilitates the selection of counterplanning strategies applicable to understanding a given conflict event. Second, the effectiveness of a given counterplanning strategy can be evaluated by the understander if the conflict situation has been ranked along the above dimensions. Thus the understander is better able to predict the probable outcome of a counterplanning strategy.

To illustrate our dimensional classification, we rank four events along our six dimensions. Later we see how counterplanning rules are triggered from this ranking.

Event 13. John was going to drive to the bowling alley, but his wife wanted the car to go shopping.

Event 14. The public works department was installing a new sewer main. John did not want them to dig up his front lawn. He made an appointment with the town planning board.

Event 15. The Seabrook nuclear power plant construction was repeatedly delayed by the anti-nuclear Clamshell alliance.

Event 16. The two gladiators were thrown in the arena. Cassius tried for a quick victory with a surprise double feint, but his Nubian opponent parried and counterattacked.

9.1. Accidental vs. intentional conflict

Let us consider what the dimension-rating table tells us about appropriate counterplanning strategies in each of the four events. We start by analyzing the counterplanning expectations generated from the values along the first dimension. If the plans of two actors interfere accidentally, it is often the case that low-order strategies are appropriate. A low-order strategy is a plan of action

TABLE 2. Dimensional analysis of plan conflicts

Dimension		Event 13		Event 14		Event 15		Event 16
(Negative)	(Positive)	both sides	Public works	John	Sea-brook	Clam-shell	both sides	
Accidental	Intentional	- 10	0	- 5	+ 10	+ 10	+ 10	
Competitive	Cooperative	+ 10	0	0	- 10	- 10	- 10	
Inconvenience	Total blockage	- 8	- 8	5	- 5	+ 10	+ 10	
Means conflict	Goal conflict	- 8	- 5	- 5	- 5	+ 10	+ 10	
External cause	Internal cause	+ 5	- 8	- 8	0	+ 5	- 10	
Low imp. goals	High imp. goals	- 5	- 8	+ 2	+ 10	+ 10	+ 10	

that takes very little effort to carry out, and usually has no adverse consequences or side-effects. For instance, a low-order strategy for obtaining some information is simply to ask. A higher-order strategy, applied to the same situation, is to threaten the person withholding the desired information. (Schank and Abelson [14] and Meehan [9] discuss the ranking of planning units, called planboxes, such as ASK and THREATEN.)

If the conflict between the two actors was intentionally brought about by one actor, the other party is not likely to succeed by applying low-order strategies. For instance, the Seabrook builders will not succeed in ending the Clamshell interference by simply asking, "Could you please let us continue building our nuclear power plant?" The strategies below suggest appropriate courses of action based on the accidental-intentional dimension. Let N be the value along this dimension, X be the counterplanning party, and Y be the other party in the conflict.

Strategy 11 (Accidental conflict)

Trigger The source of the conflict is known (or can be inferred) by the understander.

If the conflict is purely accidental ($N = -10$),

Then X should try first:

INFORM-REASON(conflict, Y), followed by

ASK(end of conflict, Y), or possibly

INVOKE-THEME(end of conflict, Y , theme).

Strategy 12 (Intentional conflict)

Trigger The source of the conflict is known (or can be inferred) by the understander.

If the conflict is purely intentional ($N = +10$),

Then X should abandon low-order counterplanning.

We apply these strategies to the dimensional rating of the four events. In Event 13, $N = -10$: the conflict situation between John and his wife is purely accidental. Strategy 11 tells us that John or his wife may only need to inform each other of the plan interference in order for the other party to help rectify the conflict situation. If awareness of the conflict is insufficient, John can ask his wife if it is all right with her for him to take the car. Alternatively, one spouse can draw upon their mutual relationship in order to request a favor; e.g., "Please drop me off at the mall on your way to the bowling alley." Schank and Abelson call this type of planning unit INVOKE-THEME. The existence of a social relationship, such as marriage, allows a person to request a favor and expect the other party to comply, within the scope of behavior defined by the social relationship.

In Event 14, John views his conflict with the public works department as

accidental in nature. He does not blame them with willful intent to destroy his front lawn. Therefore, it is likely that he may apply Strategy 11, informing or asking the planning board to reconsider their decision of installing the sewer line under his front lawn. There is less of an expectation that John can succeed in Event 14 than in Event 13. The reason for the diminished expectation is twofold: It is possible that the planning board already considered the side-effect of their plan (harming lawns) in making their decision; thus, the conflict is not totally accidental. This is why the accidental vs. intentional scale has the value -5 as compared with -10 for the previous event. The second reason for the decreased likelihood of Strategy 11 succeeding in Event 14 is that John cannot apply INVOKE-THEME as he could in the conflict situation with his wife. There is no well-established relationship between John and the planning board.

All the dimensions index certain counterplanning strategies. Here we will only consider one other illustrative dimension: Inconvenience vs. total goal blockage. The reader is referred to [3] for a more complete discussion of the other dimensions.

9.2. Inconvenience vs. total goal blockage

The inconvenience vs. total-goal-blockage dimension helps us to determine the focus of the actions taken by each actor in response to the plan conflict. Strategies 13 and 14 determine when an actor should focus his attention on counterplanning actions, and when he should reformulate his existing plan. These strategies are triggered on the basis of the value along our third dimension. This dimension tells us whether the plan interference leads to an inevitable goal blockage, or whether there are other means of achieving the actors' goal. In the latter case the plan interference is more of an inconvenience than a real goal blockage. Strategy 13 applies if the value along the goal-blockage dimension is high (i.e., $N > +5$), otherwise Strategy 14 is applicable. We use the following notational conventions: $G(X)$ is X 's goal, and $P/G(X)$ is X 's plan to fulfill $G(X)$.

Strategy 13 (Plan interference means goal blockage)

Trigger Y is blocking X 's plan $P/G(X)$.

If $P/G(X)$ is the only means for X to achieve $G(X)$ (i.e., $N > +5$),

Then X should counterplan to thwart Y 's efforts at blocking $P/G(X)$.

Strategy 14 (Alternate means method)

Trigger Y is blocking X 's plan $P/G(X)$.

If There are other means of achieving $G(X)$ (i.e., $N < +5$),

Then X should either:

- (1) Ignore Y 's interference and formulate a new or modified plan $P'/G(X)$, or

(2) If the new plan proves to be more costly to carry out than the old one plus the counterplanning actions against Y, X should counterplan to thwart Y's interference.

Refinement If the blocked goals are of high importance (e.g., the value along the sixth dimension is $> +5$), both alternatives 1 and 2 may be pursued simultaneously.

Consider how these strategies apply to our example events. In Event 13, $N = -8$. Therefore, if John or his wife have alternate means of transportation, Strategy 14 suggests that one or the other should choose this option. In Event 15, the Seabrook builders perceive the Clamshell alliance as more of a nuisance than a real threat. Strategy 14 predicts that they will choose the least-cost option of either modifying their plan to circumvent the Clamshell's efforts, or counterplan directly against Clamshell to stop their blocking actions. Since building the nuclear power plant is a high-importance goal, the refinement in Strategy 14 tells us that both reformulating the construction plans and counterplanning against Clamshell should be expected. (Seabrook has, indeed, altered their construction plans in an attempt to placate environmentalists, while pursuing legal actions against the Clamshell alliance.)

From the subjective viewpoint of the Clamshell alliance, Event 15 signifies total goal blockage. Clamshell cannot fulfill its goal of stopping nuclear power if Seabrook builds the power plant. Therefore, Strategy 13 predicts that they will focus their efforts entirely on counterplanning against the Seabrook builders.

In Event 14, John views the actions of the public works department as blocking his goal of preserving his lawn. He has no alternative plans to fulfill the preservation goal other than counterplanning against the threatened action by the public works department. Strategy 13 predicts that he will try to block their actions. His visit to the planning board must be interpreted in this light. Upon reading Event 14, we infer that his visit to the planning board must be related to a means of stopping or altering the sewer construction plan of the public works department. Strategy 13 applies also to Event 16. The goal of each gladiator is to preserve his own life. The goal of one gladiator can only be achieved by preventing the other gladiator from carrying out his combat plan. Thus we understand why they focus their efforts on countering each other's plans in the fight.

The 'cost' measure mentioned in the second alternative of Strategy 14 is difficult to define precisely. We define a plan as being costly if it violates any of the goals on the relative-importance goal tree of the actor; the higher the violated goal is in the goal tree, the more costly is the plan. For instance, John having a prolonged argument with his wife may partially violate his goal of preserving a good marital relationship. By this measure, John's counterplanning may be costly, but finding alternative transportation is the least-cost plan, (i.e., it has

no goal-violation cost associated with it). Another measure of cost associated with carrying out a specific plan is the consumption of resources necessary to carry out other plans. (Resource limitations in counterplanning is discussed in the following section, see also [17].) Thus the true cost of pursuing a plan is a combination of both measures.

10. The Structure of a Plan Conflict

Understanding plan conflicts and subsequent counterplanning actions may require a deeper analysis of the plan than we presumed thus far. The dimensional ranking of plan conflicts provides us with one useful method of indexing plan-based counterplanning strategies. There is another productive source of information which can be tapped to help an understander discover applicable counterplanning strategies. The manner in which the plan-conflict state affects an actor's plan often suggests ways to nullify effects of the conflict. These ways of nullifying the effects of a conflict situation can be encoded as counterplanning strategies. To illustrate the necessity of analyzing the way in which the conflict situation blocks an actor's plan, consider the following event.

Event 17. Bill was standing in line for the shuttle flight to New York. The woman in front of him bought the last ticket. Bill offered to buy her ticket for \$20 more than she paid for it.

In order to understand the conflict situation in Event 17, we need to use the following information: Bill planned to take a plane to New York. Buying a ticket is a necessary precondition for Bill's plan. The woman buying the last ticket prevented Bill from fulfilling the necessary precondition. Moreover, we need to infer that Bill's offer to buy the ticket is an attempt to rectify the blocked precondition by means of a bargaining strategy. This reasoning process requires a more detailed analysis of Bill's plan, and knowledge of exactly *how* the plan was blocked, than the strategies presented in the previous section can give us. Therefore, in some cases the counterplanner should investigate the nature of the conflict situation with respect to the specific plan being pursued in order to formulate an effective counterplan. An understanding of the exact nature of the conflict complements the dimensional analysis of the circumstances surrounding the conflict situation. For instance, consider the following modification to Event 17.

Event 18. Bill was standing in line for the shuttle flight to New York. Pete, who intensely disliked Bill, bought the last ticket for the flight. Bill sought out another passenger and tried to purchase his ticket.

Understanding Event 18 proceeds in much the same way as understanding Event 17. We must realize that Bill is counterplanning to fulfill the blocked

necessary precondition to his plan. However, in order to answer the question “Why did Bill seek out another passenger?”, we must use information about Bill’s and Pete’s predispositions towards each other. Bill did not think that Pete would sell him the ticket. Why? Counterplanning Strategy 12, discussed in the previous section, tells us that if the actors have a hostile or competitive predisposition toward each other, bargaining strategies are likely to fail. Thus, in order to understand Event 18, we need to use information about the predisposition of the actors, encoded as one of our dimensions, as well as information about the precise nature of the plan interference.

Plan-conflict situations may be classified according to the aspect of the plan that prevents its realization. Most plan-interference events are rooted on violated preconditions or necessary resources denied to the planner. A necessary precondition is something that must be true in order for the plan to be carried out. Each type of plan-conflict situation in the taxonomy below suggests different counterplanning strategies.

- I. Both $P/G(X)$ and $P/G(Y)$ require the use of a common resource.
 - (a) The resource is time dependent.
 - (b) The resource is consumable and irreplaceable.
 - (c) The “resource” is help from a third party.
- II. Executing $P/G(Y)$ changes the state of the world so as to:
 - (a) Block a mediating precondition for $P/G(X)$.
 - (b) Block an uncontrollable precondition for $P/G(X)$.
 - (c) Block a controllable precondition for $P/G(X)$.
- III. $P/G(Y)$ makes executing $P/G(X)$ more troublesome (costly).

We analyze time-dependent, human-assistance resource limitations and blocked preconditions, presenting the counterplanning strategies triggered by each situation. The reader is referred to [3] for a discussion of the other plan-conflict situations in our taxonomy.

10.1. Time-dependent resource limitations

Consider some examples of counterplanning based on time-dependent resources.

Event 19. John wanted to drive to the bowling alley, but Mary took the car to go shopping.

Event 20. John was driving to Mary’s house. As he approached the Main Street intersection, he saw that the Thanksgiving parade was in full swing. John had to drive several miles out of his way.

Event 19 is our familiar example from the previous section. The plan interference, however, can also be characterized as two actors needing to use

the same resource, their car. Earlier, we discussed several possible counterplanning strategies for this situation, including plan abandonment by one actor and cooperative planning (such as one actor driving the other to his/her destination). Two additional courses of action that either actor can pursue are resource substitution and resource scheduling. Essentially, John can substitute for the car a resource which serves the same function of providing a means of transportation (e.g., ride his bicycle). Alternatively, John can reschedule the execution of his plan (e.g., wait until Mary returns). Each of these two courses of action corresponds to a counterplanning strategy ((16) and (17) respectively) based on the need for a shared time-dependent resource.

Strategy 15 (Control-resource subgoal)

Trigger $P/G(X)$ and $P/G(Y)$ need resource R where R cannot be used in more than one plan at a time.

If $G(X)$ and $G(Y)$ are in a mutual-exclusion conflict, or if Strategies 16 and 17 have failed,

Then X should create and pursue the high-priority goal $G'(X) = \text{DCONTROL}(X, R)$.

Strategy 16 (Resource substitution)

Trigger $P/G(X)$ and $P/G(Y)$ need resource R where R cannot be used in more than one plan at a time.

If X can find an alternative resource R' , such that $P/G(X)$ can use R' instead of R (possibly with minor modifications to P),

Then He should substitute R' for R and pursue his previous plan.

Strategy 17 (Resource scheduling)

Trigger $P/G(X)$ and $P/G(Y)$ need resource R where R cannot be used in more than one plan at a time.

If R is a time-dependent resource,

Then X should negotiate with Y to pursue $P/G(X)$ at time interval T_1 and pursue $P/G(Y)$ at time interval T_2 , where T_1 does not overlap with T_2 .

Refinement If one plan must be rescheduled, it should be the one that minimizes cost for both actors.

Strategies 16 and 17 are applicable when the two actors are cooperatively predisposed, and there is no underlying goal conflict to their plan interference. Strategy 15 should only be applied if X and Y cannot (or will not) cooperate. Hence, we see that the dimensional analysis proves useful in deciding the appropriateness of counterplanning strategies indexed by other means.

Event 20 can be understood in terms of Strategy 16, if one considers the

street to be a public resource. Main Street could not be used simultaneously by motorists and the Thanksgiving-day parade. Therefore, John chooses an alternate resource: He used different roads to pursue the same general plan of driving to Mary's house. Alternately, John could have waited for the parade to finish (an application of Strategy 17).

10.2. Human-assistance resource: Help from a third party

Sometimes two actors require assistance from the same third party. Some types of assistance, such as advice, can be given to many actors simultaneously, but other types of assistance require that the third party focus its attention on the needs of one actor to the exclusion of helping others. We classify this type of assistance in the same category as physical resources because it shares a large fraction of the counterplanning strategies with resource limitations. Consider the following event to which we may apply some of the previous resource-limitation counterplanning strategies.

Event 21. John and Bill wanted to take flying lessons. The flying instructor told them that his schedule was almost full. He could only accommodate one of them.

Continuation 21(1) John signed up and Bill decided to wait until the next set of lessons.

Continuation 21(2) John signed up and Bill went to another flying school.

Continuation 21(3) John convinced the instructor that Bill was suicidal. The instructor chose to teach John.

Continuation 21(4) John slipped a \$50 bill to the flying instructor. His lessons started that afternoon.

Event 21 describes a resource-limitation conflict where the resource is human assistance. Neither actor can carry out his plan of taking flying lessons without the active cooperation of the flying instructor. Since flying lessons require the full attention of the instructor on teaching a single student, John and Bill cannot simultaneously carry out their respective plans. Continuation 21(1) is an application by John and Bill of Strategy 17, resource scheduling. Bill applies resource substitution (Strategy 16) in Continuation 21(2).

The last two continuations illustrate a resource-limitation strategy that is only applicable if the resource is human assistance. Bill convinces or bargains with the instructor to secure his assistance over John's mutually exclusive wishes. The general strategy is presented below.

Strategy 18. (Securing human assistance)

Trigger $P/G(X)$ and $P/G(Y)$ need resource R where R cannot be used in more than one plan at a time.

If R is human assistance,
Then X has the following alternatives.

(1) X should use the PERSUADE package to convince R to help X instead of helping Y.

(2) If R is time limited, X should use the PERSUADE package to convince R to create more time (e.g., sacrifice other activities) so that R can assist both X and Y.

(3) X may ASK R if there is an R' who can perform the same function. If such an R' exists, either X or Y should invoke Strategy 16 (resource substitution) on R'.

(The 'PERSUADE package' refers to Schank and Abelson's [14] methods for one actor to convince another to follow a desired course of action.)

The first alternative of the above strategy was invoked by John in Continuation 21(3). If the understander is not aware that one of John's options is to persuade the instructor to accommodate him, it is very difficult to understand why John told him Bill was suicidal. The same alternative, using a different means of persuasion, was invoked by John in Continuation 21(4). We need to understand that giving \$50 is part of John's plan in his bargaining strategy to secure the services of the instructor. Therefore, the knowledge encoded in Strategy 18 (as well as knowledge about the self-preservation and acquisition of money goals that the instructor must have) is necessary to understand Event 21 followed by its third or fourth continuations.

The second and third alternatives of Strategy 18 are illustrated in the following two additional continuations of event 21.

Continuation 21(5) John offered to let the instructor use his yacht if he gave John lessons on Saturday.

Continuation 21(6) John asked to be referred to some other flying instructor, and was told of a pilot who gives private lessons. Bill went to see this pilot.

10.3. Resources as measures of cost

The cost associated with executing a plan (as defined earlier) can be considered as a combination of two factors: (1) other goals which may be violated by carrying out the plan, and (2) the consumption of resources not available for use in other plans. The second measure of cost applies to time-dependent resources and human-assistance resources. In the physical-resource case, an actor precludes the pursuit of another plan he wished to pursue if the second plan required the same resource [17]. In the latter case, the same problem may exist, but, more importantly, the person may not be willing to be of assistance more than a few times (or may require compensation in the form of another

useful resource). Thus, cost is measured in terms of what cannot be done in the future as a result of actor's present actions.

The acquisition of some resources, such as money for people and energy-producing substances for countries, may become important subsumption goals in themselves. In a similar manner, cultivating friendships and having people owe favors become goal states, as these can be considered to be useful resources in formulating future plans. For instance, having political connections is a useful resource-subsumption state for many different types of plans. Severing political or interpersonal ties, therefore, can have a high cost, as it wastes potential human-assistance resources.

11. Counterplanning Against Blocked Preconditions

In order to successfully execute a plan, there are usually some preconditions that must hold true. Schank and Abelson [14] classify preconditions into three categories: controllable, uncontrollable, and mediating preconditions. For instance, if John's plan is to ask Mary where the bank is located, the following preconditions must be met: (1) John must establish a communications link with Mary, such as telephoning her, or being in physical proximity. (2) Mary must know where the bank is located. (3) Mary must be willing to convey the information to John. The first precondition is a controllable precondition, because we assume that John can achieve it at will. The second is an uncontrollable precondition; John cannot bring this precondition about by any action on his part. The third precondition is called a mediating precondition. Mediating means that the planner can try to bring this condition about, but its final outcome rests on the actions or predispositions of another party. For instance, John can use persuasion strategies to convince Mary that she should tell him where the bank is, but there is no guarantee that John will succeed. In addition to persuasion, there are other methods to circumvent a blocked mediating precondition, listed in the following strategy.

Strategy 19 (Unblocking a mediating precondition)

Trigger Y interferes with X's plan $P/G(X)$.

If Y is blocking a mediating precondition for $P/G(X)$,

Then X should choose the least cost alternative among:

(1) Establish the goal $G1(X)$ = "Bring about the mediating precondition by some means independent of Y's actions." Suspend $P/G(X)$ until $G1(X)$ is achieved.

(2) Choose a new plan P' , such that $P'/G(X)$ does not require this mediating precondition.

(3) Set up the mutual-exclusion goal conflict $G(X)$ = "Y is not blocking the mediating precondition to $P/G(X)$." Invoke the mutual-exclusion counterplanning strategies.

Consider the case where one actor's plan blocks an uncontrollable precondition of another actor's plan. We cannot apply the first alternative of Strategy 19 because, by definition, an actor cannot bring about uncontrollable preconditions. The second and third alternatives are applicable to both mediating and uncontrollable preconditions. Another possibility suggests itself if the blockage is temporary: wait until the uncontrollable precondition is no longer blocked. For blocked mediating preconditions, the waiting option is subsumed under the first alternative in Strategy 19 if a more active means to re-establish the blocked mediating precondition cannot be applied. Here is the blocked uncontrollable-precondition counterplanning strategy.

Strategy 20 (Unblocking an uncontrollable precondition)

Trigger Y interferes with X's plan $P/G(X)$.

If Y is blocking an uncontrollable precondition for $P/G(X)$,

Then X should choose the least cost alternative among:

- (1) Choose a new plan P' such that $P'/G(X)$ does not require this uncontrollable precondition.
- (2) Invoke the mutual-exclusion counterplanning strategies to thwart Y's continued blocking of the uncontrollable precondition.
- (3) If Y's blockage of an uncontrollable precondition for $P/G(X)$ is temporary, X should suspend P until Y no longer blocks this precondition. Then X should resume the pursuit of $P/G(X)$.

To illustrate the third alternative of Strategy 20, consider the following event.

Event 22. Smith wanted to be elected president of the republic. The elections were cancelled that year by the military junta. Smith bided his time until the following year.

Question. What was Smith waiting for?

If the understander realizes that Event 22 is an instance of the third alternative of Strategy 20, it is clear how the question should be answered. Smith has not given up his goal; he merely suspended it until the temporary blockage of the uncontrollable precondition was removed in the following year. Hence an appropriate answer is: "Smith was waiting until the following year to run for president."

Counterplanning against blocked controllable preconditions is much simpler. Indeed, in many cases the actor whose plan's controllable precondition was blocked re-establishes the precondition without recourse to other counterplanning measures. The following is the blocked controllable precondition strategy.

Strategy 21 (Unblocking a controllable precondition)

Trigger Y interferes with X's plan P/G(X).

If Y is blocking a controllable precondition for P/G(X),

Then X should choose the least cost alternative among:

(1) Re-establish the blocked controllable precondition.

(2) Use the PERSUADE package to convince Y to abandon or change his plan in order that the controllable precondition to P/G(X) is no longer blocked.

Refinement X should choose the method for re-establishing this precondition in such a way that it will not be blocked once again by P/G(Y).

12. Concluding Remark

We have seen how counterplanning is a necessary process of both decision making in the face of adversity and for understanding natural language descriptions of such conflict situations. The counterplanning strategies discussed in this paper are illustrative rather than exhaustive. A more complete set of strategies is discussed in [3].

Counterplanning is a general inference mechanism for understanding human conflict situations. Since most interesting stories and events involve some type of goal conflict, we need to include counterplanning as a necessary tool in the Artificial Intelligence repertoire.

REFERENCES

1. Barstow, D.R., Automatic construction of algorithms and data structures using a knowledge base of programming rules, Ph.D. Dissertation, Stanford University, Stanford, (Nov. 1977).
2. Carbonell, J.G., POLITICS: Automated ideological reasoning, *Cognitive Sci.* 2(1) (1978) 27-51.
3. Carbonell, J.G., Subjective understanding: Computer models of belief systems., Ph.D. Dissertation, Yale University (Jan. 1979) Comput. Sci. Rept. No. 150.
4. Carbonell, J.G., POLITICS: An experiment in subjective understanding and integrated reasoning, in: Schank, R.C. and Riesbeck, C.K. (Eds.), *Inside Computer Understanding: Five Programs Plus Miniatures* (Erlbaum, Hillside, 1980).
5. Carbonell, J.G., Towards a process model of human personality traits, *Artificial Intelligence*, 15(1,2) (1980) 49-74.
6. Charniak, E., A framed painting: Representation of a commonsense knowledge fragment, *Cognitive Sci.* 1(4) (1977) 355-394.
7. Cullingford, R., Script application: Computer understanding of newspaper stories, Ph.D. Dissertation, Yale University, New Haven (Sept. 1977).
8. Forgy, C.L. and McDermott, J., OPS, a domain-independent production system language, *Proc. of the Fifth Internat. Joint Conf. on Artificial Intelligence* (1977) 933-939.
9. Meehan, J.R., The metanovel: Writing stories by computer, Ph.D. Dissertation, Yale University, New Haven (Sept. 1976).
10. Newell, A. and Simon, H.A., *Human Problem Solving* (Prentice-Hall, Englewood Cliffs, 1972).
11. Newell, A., Productions Systems: Models of control structures, in: Chase, W.G. (Ed.), *Visual Information Processing* (Academic Press, New York, 1973).

12. Sacerdoti, E.D., *A Structure for Plans and Behavior* (North-Holland, Amsterdam, 1977).
13. Schank, R.C., *Conceptual Information Processing* (North-Holland, Amsterdam, 1975).
14. Schank, R.C. and Abelson, R.P., *Scripts, Goals, Plans and Understanding* (Erlbaum, Hillside, 1977).
15. Schmidt, C., Sridharan, N. and Goodson, J., The plan recognition problem, *Artificial Intelligence* **11**(2) (1978) 45-83.
16. Shortliffe, E., *Computer Based Medical Consultations: MYCIN* (Elsevier, New York, 1976).
17. Wilensky, R., Understanding goal-based stories, Ph.D. Dissertation, Yale University, New Haven (Sept. 1978).

Received November 1971; revised version received November 1980