DTIC FILE COPY

AD-A955 947

The Counterplanning Process: A Model of Decision-Making in Adverse Situations.

Jaime G. Carbonell

February, 1979

DEPARTMENT
of
COMPUTER SCIENCE

DTIC
ELECTE
JUL 17 1990
S E D

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

Carnegie-Mellon University

90 07 16 455

DO NOT REMOVE
ZDAAAAAA1470202

The Counterplanning Process: A Model of Decision-Making in Adverse Situations.

Jaime G. Carbonell

Carnegie-Mellon University
Computer Science Department
Pittsburgh, Pennsylvania 15213

February, 1979

Accession For
NTIS GRA&I ☒
DTIC TAB ☐
Unannounced ☐
Justification
By
Distribution/
Availability Codes
Dist | Avail and/or Special
A-1

UNANNOUNCED

DTIC COPY INSPECTED 5

APPROVED FOR PUBLIC RELEASE
DISTRIBUTION UNLIMITED

[1]This Research was sponsored in part by the Defense Advanced Research Projects Agency (DOD), 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.

Table of Contents

1. Introduction. 1
2. The POLITICS system. 2
3. The Counterplanning Control Process. 5
4. The structure of the counterplanning strategies. 9
5. Diversionary counterplanning strategies. 10
6. Predictive vs. Explanatory Reasoning. 13
7. Counterplanning strategies based on goal compromise. 15
8. Counterplanning strategies based on goal substitution. 21
9. Understanding the reasons for plan conflicts. 24
    9.1. Accidentally vs. Intentionally caused conflict situations. 26
    9.2. Inconvenience vs. total goal blockage. 28
10. The structure of a plan conflict. 30
    10.1. Time-dependent resource limitations. 32
    10.2. Human-assistance resource: Help from a third party. 34
    10.3. Resources as measures of cost. 36
11. Counterplanning against blocked preconditions. 37
12. Concluding remark. 39
13. References. 41

The Counterplanning Process: A Model of Decision-Making in Adverse Situations.

Jaime G. Carbonell
Carnegie-Mellon University
Department of Computer Science
January, 1979

Abstract

A heuristic model of human reasoning is discussed where the reasoner must contend with dynamically-changing goals and actions of other actors in the world. 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 goal in spite of attempts by others to block his plans. The model has been implemented as part of the POLITICS system, a computer program that understands simple natural language accounts of international political conflicts.

KEY WORDS AND TOPICS: planning, counterplanning, natural language, heuristics, inference, rule-based systems.

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 human decision processes. We focused our research on situations where the decision maker has to contend with other parties that may be continuously trying to thwart his efforts. In such circumstances, the formulation of plans of action and the decisions of how and when to implement the plans become much more complex 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 part of potential adversaries, formulating alliances with other parties that have common goals, and compromising with one's adversaries when necessary. We discuss our theory of counterplanning and its implementation in POLITICS, 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 (e.g., Newell and Simon [1972], Sacerdoti [1975]) and plan interpretation in story understanding tasks (e.g., Schank and Abelson [1977], Wilensky [1978], Schmidt and Sridharan [1978]). 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 (Carbonell [1979]) 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 - Schank and Abelson [1977], Cullingford [1977]), 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, 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 from the point of view of a conservative American policy maker, and answers some 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 VIOLATEO:
 *US* MILITARY STRENGTH DECREASEO.

 *US* GOAL VIOLATED:
 ACTOR (#COUNTRY TYPE (*COMMUNIST*))
   <=> (*MTRANS*)
  MOBJ ((ACTOR (*US*) TOWARO (*STRENGTH* TYPE (*MILIT*))
                                          INC (-2)))

*RUSSIA* GOAL ACTIVATEO.
ACTOR (*RUSSIA*) TOWARO (*SCONT* OBJ (*PANAMA*))
                                     INC (*POS*)

 INTEPRETATION COMPLETEO, 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 UNITEO STATES WOULO 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 UNITEO 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 SHOULO 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 the 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 to model possible US decisions. These heuristic strategies create the more probable hypothetical scenarios (as illustrated in the answer to question 5), and subsequently decide on means to bring about or prevent the hypothetical scenario. In answering question 7, POLITICS is trying 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 for 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 an integrated understanding system, it can be conceptually divided into several modules, including: 1) Natural Language understanding and generation in a semantically-rich domain. 2) A goal hierarchy representation of the political ideologies that focus the inference process. 3) Script application and situational inference rules. 4) A reasoning 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 been heretofore directly addressed in Artificial Intelligence.

Planning and Counterplanning require self-knowledge of the goals that one strives for, 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, instrumentally, 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 the ideological goals.

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 <u>obstructive</u> <u>counterplanning</u>. The second type of counterplanning scenario is essentially the dual of the first type: An actor X is trying to achieve his goal G(X) in spite of attemps by Y to prevent G(X) and to block X's plans for pursuing G(X). This process, called <u>constructive</u> <u>counterplanning</u>, differs from obstructive counterplanning only in terms of the subjective perspective of the counterplanner. The perspective shift causes the counterplanner to apply different strategies at different times, as the application of strategies is goal driven. Figure 1 is a control-flow diagram for the obstructive counterplanning process.

To illustrate the obstructive counterplanning mechanism, consider a simple example. Prison guard X wants to prevent prisoner Y from escaping. Hence, G(Y) is Y being free outside the prison. We enter figure 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, if any, plan the prisoner formulated. This plan determination may itself involve some planning: Should the guard ask the prisoner?, Should he threaten him? (See Schank and Abelson [1977] for a discussion of social planning units). Let us again assume that the guard fails. At this point we enter the third box in figure 1. The guard can ask himself "what would I do if I was trying to escape? If he finds a reasonable plan, he should assume that this may be the prisoner's plan and he should apply the obstructive counterplanning strategies to blocking the plan. For instance, the guard may find that stealing the key is a reasonable plan, In which case he would take 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

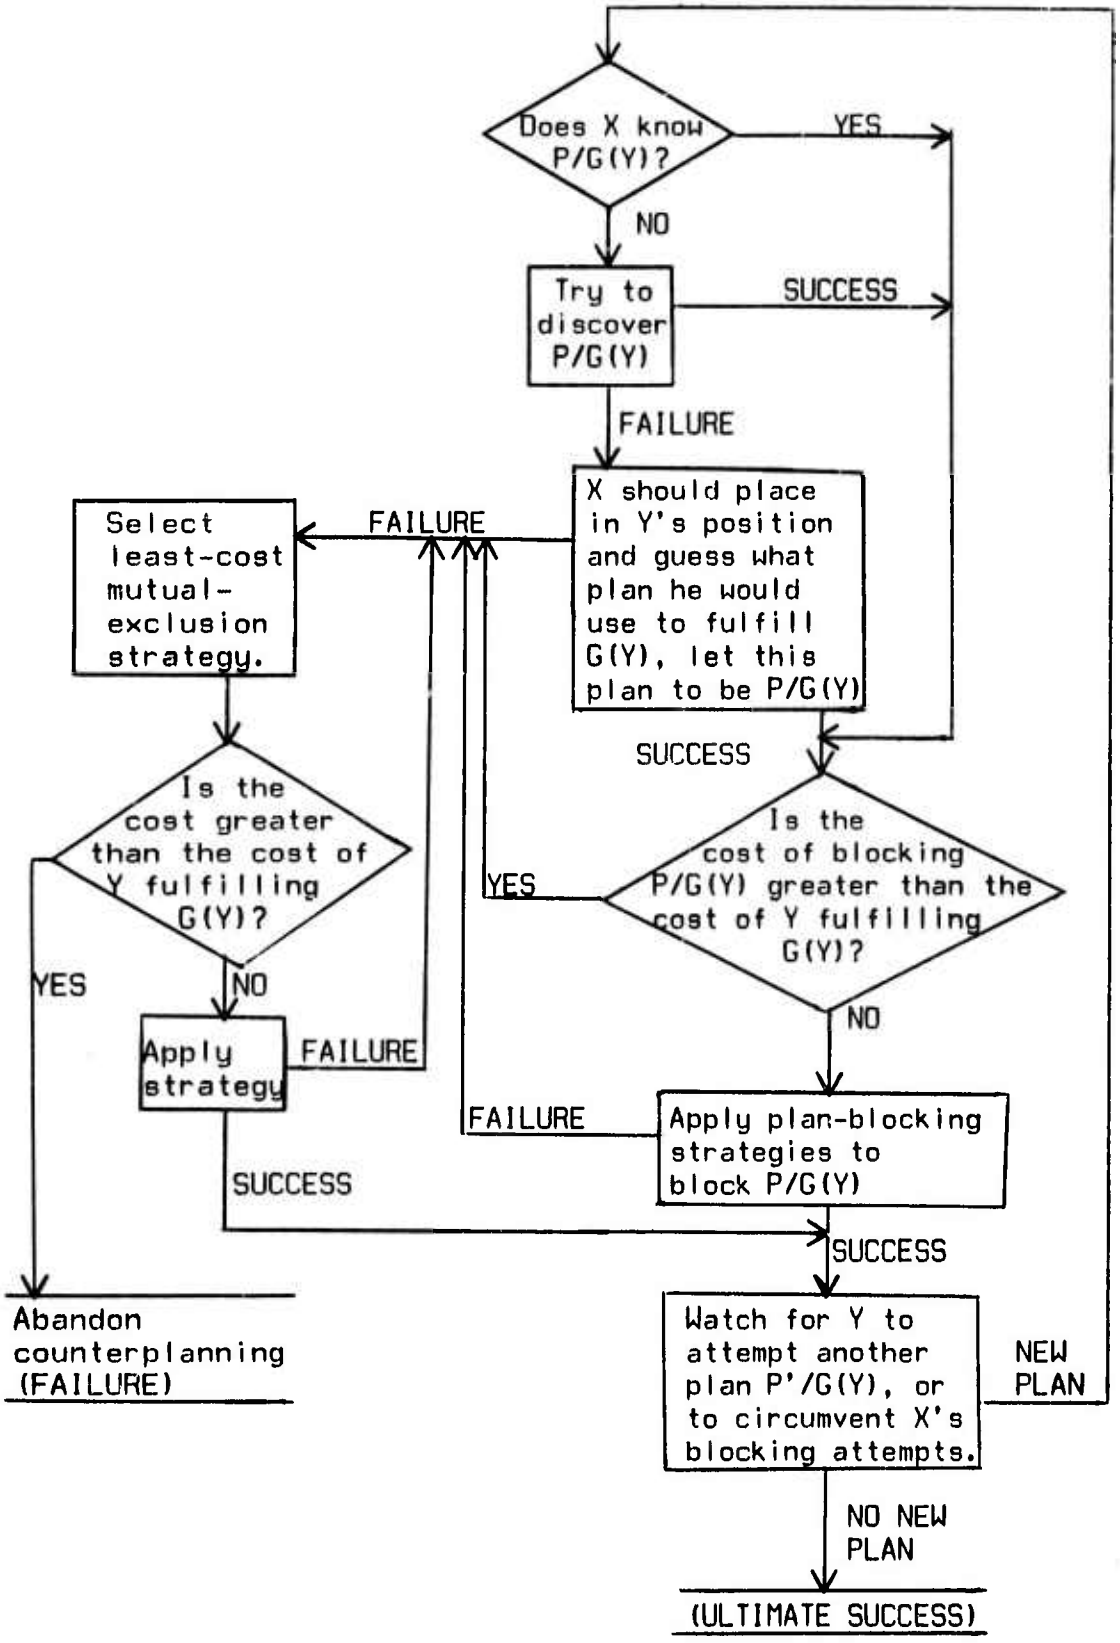

Figure 1: Control structure of the Counterplanning process (obstructive phase). Actor X attempts to block actor Y from achieving his goal G(Y). P/G(Y) is Y's plan for fulfilling G(Y).

(i.e., apply mutual exclusion strategies) such as pointing a gun at the prisoner and informing him that he either remains put, or he will be dead. Thus escaping and staylng 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 risk of being blown up is more costly than his goal of thwarting the prisoner's goal. 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 briefly discussed in a later section.

Figure 2 is the control-flow for constructive counterplanning. Let us consider the same situation from the point of view of the prisoner. We enter the first box in figure 2: Is the escape goal directly blocked by the guard? If this is the case, then 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, then 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, then 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. thereby expecting the guard to pursue his higher-level goal of staying alive and giving the prisoner a chance to escape.

Finally, there is another set of counterplanning strategies that the prisoner may pursue. 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 (e.g., keeping a constant eye on the prisoner, so he has no time to dig his escape tunnel), such as time, materials, or outside assistance. There are heuristic strategies for attempting to overcome each type of blockln action. Constructive counterplanning also has exit conditions based on the relative cost of abandoning or continuing the effort.

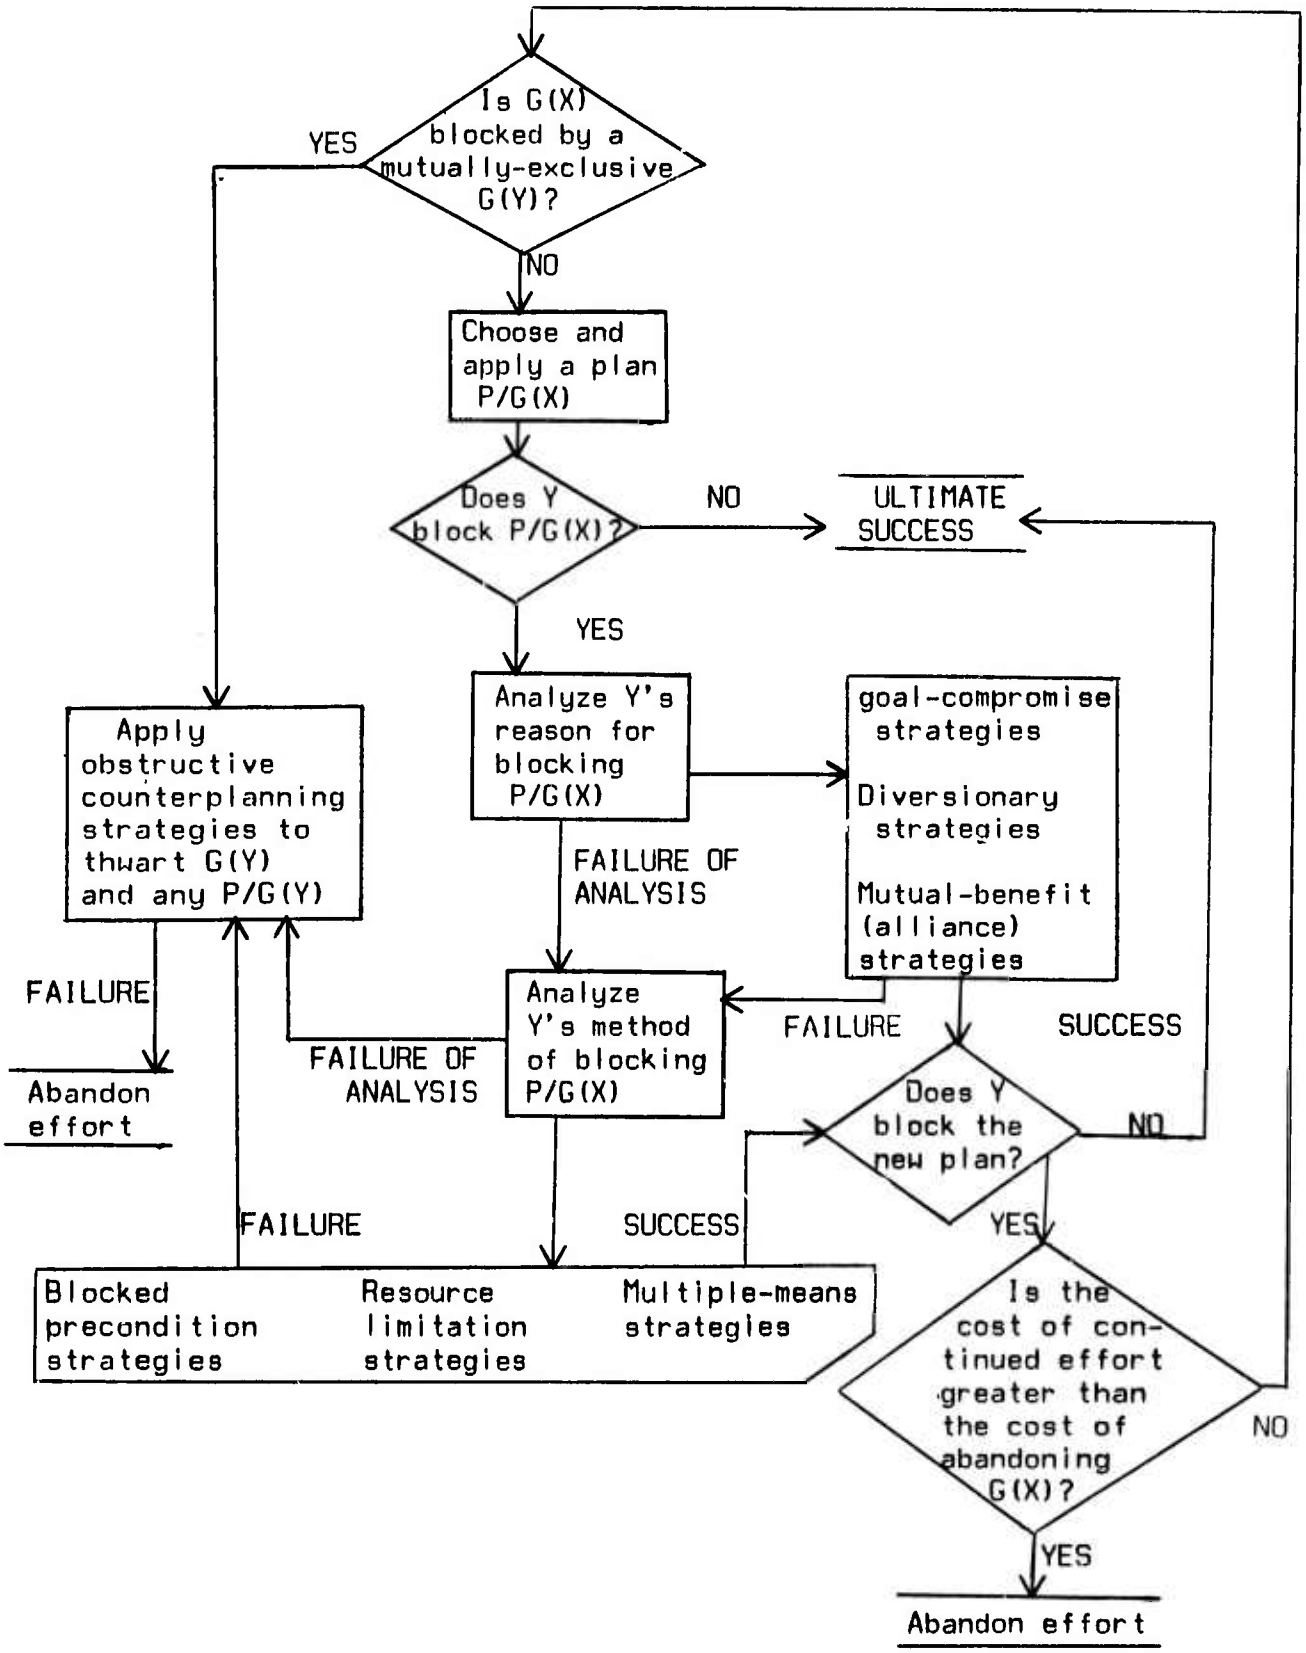

Figure 2: Control structure of the counterplanning process (constructive phase.) Actor X wants to achieve G(X) in spite of Y's attempts to thwart his plans or prevent his ultimate goal.

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 that 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 in figure 3:

```
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, ...>
```

Figure 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 that the process model must spend searching for an applicable strategy. All the trigger conditions are "inexpensive" tests: that is, they are tests that can be applied directly to the situation without requiring further inference or complicated matching. Furthermore, 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 tests are performed. These tests may be arbitrarily complex, requiring further inference and possibly requiring other counterplanning strategies to be invoked. However, the trigger conditions usually restrict the set of applicable counterplanning strategies to a small number in any given situation (typically three or four). Furthermore, counterplanning strategies reflect common-sense reasoning about how to deal with adverse situations. As such, each strategy is sufficiently general to apply across most reasonable human conflict domains. 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 counterplanning strategies suffice to model the counterplanning actions in most situations and most goal conflicts.

The action part of each strategy (denoted by the word "THEN") 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 after the sequence of actions is performed. The refinement contains one or more additional rules that usually provide further detail to the counterplanning situation. These rules are truly subrules to the counterplanning strategies because 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, Barstow [1977]) than to more constrained production systems (e.g., Newell's [1973] PSG system).

The following sections analyze some of the more significant types of counterplanning strategies. A much more complete set is discussed in (Carbonell [1979]). We also discuss the way in which counterplanning strategies are used in text-understanding and question-answering tasks.

5. Diversionary counterplanning strategies.

We turn to some of the more specific counterplanning strategies applicable to a mutual-exclusion goal conflict. A frequently-encountered set of strategies operate on the principle of diverting the efforts of an actor 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 has to 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 other goals.

3) Trick option - Convince one's counterplanning opponent that other important goals are threatened.

Let us look at one important diversionary strategy and its "trick option". A trick option consists of a pretense whose purpose is to convince other parties that one is pursuing a certain plan, while, in reality, this need not be the case.

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 G'(Y) should be chosen such that:
1) Y cannot pursue both G'(Y) and G(Y) simultaneously, or
2) Y can accomplish G'(Y) only if X stops his blocking action.

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 preserve that higher level goal. Since an actor's attention, time, and material resources are limited, Y may not be able to simultaneously protect his higher level goal and pursue his goal that conflicts with X's goal. Thus, X will be in a better position to win the conflict situation and achieve his own goal. The "refinement" part 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 querying 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 trick-option version of strategy 1. This strategy is based on the same principle of diverting the attention and efforts of the opposing counterplanner away from goal conflict. However, the method used need not correlate with reality. As such, one should expect that a trick-option strategy is more likely to fail (e.g., if the opposing actor discovers the truth.).

STRATEGY 2: TRICK OPTION: 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 then out are some common applications of counterplanning by means of invoking the trick-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 Carbonell [1979]); however they all share a common principle: An actor cannot simultaneously pursue 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 simultaneously have more than one goal to worry about. There are various kinds of limitations on the different items that an actor can focus his attention on at one 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 that can be performed simultaneously.

3) Limitations of material resources.

4) Limitations of ability.

5) Goal of avoiding certain consequences of one's actions.

6) Interactions between different courses of actions.

6. Predictive vs. Explanatory Reasoning.

For the task of generating hypothetical scenarios, one can only use counterplanning strategies in a predictive manner. However, in story understanding tasks, the counterplanning strategies are used to explain the actions of the actors. In such situations, it is better to apply the strategies in an explanatory manner. To see what we mean, consider event 2:

EVENT 2: The United States supported Israel in the 1973 Middle-East war. Subsequently the Arabs imposed an oil embargo on the United States and its allies.

QUESTION: Why did the Arabs impose the oil embargo?

Strategy 1 can be used In a predictive framework or in an explanatory manner. In understanding event 2 we need to pose and answer the question "Why did the Arabs impose an oil embargo?" Using knowledge about the goals of the Arabs and the goal of the United States to help Israel, the understander can establish the goal conflict between the Arabs and the United States. The mutual-exclusion goal conflict is between the US goal of aiding Israel and the Arab goal of preventing US aid to Israel.

Having interpreted the situation thus far, an understander can proceed in two different manners. The first manner is to predict all the possible Arab counterplanning actions to make the US stop aiding Israel. If, in interpreting the rest of the event, the understander matches an action with one of the previously predicted counterplanning actions, the understander can

conclude that indeed the Arabs were counterplanning against the US and their counterplan was the predicted course of action. Such a process would require the understander to generate vast numbers of hypotheses and subsequently test each hypothesis as a possible explanation of the situation. There is no evidence to suggest that people generate all possible inference paths in a given interpretation in order to discard all but one path that matches reality. It appears much more plausible that people only pursue a small number of relevant inference paths. Therefore, generating all possible plans of action is not a reasonable psychological model of human thinking, nor does it lend itself to reasonable constraints on the computational time that the system may require in order to generate and test all alternative actions on the part of the Arabs.

A more reasonable alternative to the generate and test process is the following: Given the existence of the mutual-exclusion goal conflict, we can predict that the two actors may counterplan against each other. No further predictive inferences are generated at this point. The rest of the event should be interpreted in light of the expectation that the two actors may counterplan to resolve their goal conflict.

When the understander learns of the Arab oil embargo, he tries to see if this is a reasonable course of action to take as a counterplan against the US goal of aiding Israel. Counterplanning strategy 1 (refinement 2) matches the type of interaction between the Arab plan and the US goal. The Arabs are threatening a higher level US preservation goal by cutting off oil supplies, and the US cannot do anything to directly remedy the situation. Now the Arabs can bargain to end the embargo in return for the end of US aid to Israel. Therefore, the understander can establish the Arab counterplanning actions by applying strategy 1 in an explanatory manner. The result of the Arab actions is matched to the action part of the strategy. This match, suggested by our previous expectation that a counterplanning action was likely, allows the understander to infer that the Arabs were invoking counterplanning strategy 1. Therefore, we can say that the reason for the Arab action is the test clause of the strategy.

The explanatory mode of reasoning is superior to the predictive mode because it does not require the generation and subsequent testing of an arbitrarily large set of possible courses of action. It is also a more reasonable model of human thought. 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 were clear to almost everyone.

7. Counterplanning strategies based on goal compromise.

A significant class of counterplanning strategies in mutual-exclusion goal conflicts is the set of bargaining strategies. These strategies are characterized by a willingness to compromise on the part of the disputing actors. Willingness to compromise is mediated by many factors such as whether compromise is possible, necessary, or desirable on the part of both disputing actors.

There are two basic classes of bargaining strategies, each class is characterized by the type of compromise that the counterplanning actors are willing to accept. Let G(X) be X's goal, and G(Y) be Y's goal. As before, 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 a) John was fired and blacklisted for unethical business practices.

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

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

In case (a), John failed to achieve his goal. In case (c), John clearly achieved his goal. What can we say about case (b)? Strictly speaking, John did not achieve his goal of becoming president. However, John did not totally fail in fulfilling his ambition. We classify case (b) 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 examples 3 and the two latter examples? Becoming president of a company is a complex goal subsuming several simpler goals such as achieving social respect, power and wealth. (Wilensky [1978] 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 outcome. 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 goal subsumption state, 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 are fulfilled.

The second way in which partial goal fulfillment is measured applies to goal states that can take a continuum of values. Acquisition of money and achievement of social stature are examples of continuum-valued goals; there are virtually infinite degrees of social stature and of the amount of money that 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 [1977] define the goal taxonomy.) Let us call the <u>initial</u> value of the goal state "I", the <u>desired</u> value of the goal state "D", and its <u>resultant</u> value at the time when we must decide whether the goal succeeded "R". For P-goals it is usually the case that I = D, and for A-goals I < D. The success and failure conditions of continuum-valued goals are given by the following table:

	FAILURE	PARTIAL SUCCESS	SUCCESS
A-GOAL	I = R < D	I < R < D	I < R = D
P-GOAL	R < I = D	R slightly < I = D	R = I = D

Table 1: Partial fulfillment of goal states.

Let us consider how the 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 = controling 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 Somalls 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 Somalis 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 partially fulfilled.

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 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 does not compromise if one can totally fulfill one's goal. In order 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 would predict 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 Communist Party were able to form a majority government. They formed an unprecedented coalition to govern the country.

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

STRATEGY 5: TRICK OPTION: 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: TRICK OPTION: 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 trick option can be applied to events 8 and 9 as follows:

EVENT 8 (possible continuation):
Jones supported Smith's successful bid for the chairmanship, but in the next election Smith badmouthed Jones.

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

The understanding 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 suggests than an expectation has been violated. (Schank [1975] describes the "but test", a heuristic method to determine what is inferred in understanding a text.) Since badmouthing 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 the trick option 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 totally fulfill 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 the Carbonell [1979]), then 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 Schank and Abelson [1977]. 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 his goal of riding the bicycle with the goal of eating a candy bar. 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 that ended 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 that 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, people rarely compromise on P-goals, regardless of their importance. 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, has two trick options:

STRATEGY 8: Trick option: FALSE SUBSTITUTION

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

STRATEGY 9: Trick option: 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 goal-substitution bargaining strategy involves mutual abandonment of both G(X) and G(Y). This strategy is invoked in the case that 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 other counterplanning efforts may violate higher level goals than G(X) and G(Y).

IF X also has the goal of NOT(G(Y)),

THEN he should negotiate with Y to mutually abandon G(X) and G(Y). (Otherwise 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 have the goal of the other country not being the dominating world power, strategy 10 suggests negotiation to mutually abandon the goals of attaining military superiority.

9. Understanding the reasons for plan conflicts.

Thus far, we have discussed strategies for direct goal blockage, 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 type of counterplanning strategies this analysis suggests.

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

1) Accidental vs. intentional plan conflict.

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

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

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

5) Externally imposed conflicts vs. internally motivated ones.

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, our 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.

DIMENSION		EV. 13	EVENT 14		EVENT 15		EV. 16
		both	Pub.	John	Seab	Clam	both
(Negative)	(Positive)	sides	works		rook	shell	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.	-5	-8	+2	+10	+10	+10

Table 2: Dimensional analysis of plan conflicts.

9.1. Accidentally vs. Intentionally caused conflict situations.

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 inappropriate. A low-order strategy is a plan of action 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 to simply ask for it. A higher order strategy, applied to the same situation, is to threaten the person withholding the desired information. (Schank and Abelson [1977] and Meehan [1976] 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 need only 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 that strategy 11 can succeed 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 refered to Carbonell [1979] 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 actor's 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 (e.g., 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) (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) (N < +5),

THEN X should either:

1) Ignore Y's interference and either modify or totally formulate a new 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 alternative 2 of strategy 14 is difficult to define precisely. We define a plan as being costly if it violates any of the goals on the 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, 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, also see Wilensky [1978]). 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 that 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. Furthermore, 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 that are denied to the planner. A necessary precondition is something that must be true of the world 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 irreplacable.
  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).
) 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 refered to Carbonell [1979] 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 that 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) = 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, X should agree with Y to have X pursue P/G(X) at time T1, and Y pursue P/G(Y) at time T2, where T1 does not equal T2.

REFINEMENT If one plan has to be rescheduled, it should be the one that causes the least cost problems 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 chose an alternate resource: He used different roads to pursue the same general plan of driving to Mary's house.

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 for which some of our previous resource-limitation counterplanning strategies are applicable.

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 15, 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 [1977] 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 the giving $50 is part of John's plan in his bargaining strategy to secure the services of the instructor. Therefore, some rule like 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 possible 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 can be considered a combination of two factors: 1) other goals that 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 former case, an actor precludes the pursuit of another plan he wished to pursue if the second plan required the same resource (Wilensky [1978]). 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 an actor's present actions.

The acquisition of some resources, such as money for people, and energy-producing substances for countries, becomes an important subsumption goal in itself. In a similar manner, cultivating friendships and having people owe favors become goal states, as these can be considered to be resources useful in 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 because it wastes human 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 [1977] 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 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

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 alternative 1 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

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 chairman of the department. The elections were cancelled that year by the dean. 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 chairman."

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

          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.
                    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).

                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.
```

12. Concluding remark.

We have seen how counterplanning is a necessary process for both decision making in the face of adversity, and understanding natural language events about conflict situations. The counterplanning strategies discussed in this paper do not form a complete set (primarily due to limitations of length, but see Carbonell [1979]); rather, they are meant as an illustrative set.

Counterplanning is a general inference mechanism for understanding human conflict situations. Since most interesting, "real world" stories and events involve some type of

conflict and attempts at resolving various aspects of the conflict, counterplanning is a necessary tool in the Artificial Intelligence repertoire.

13. References.

Barstow, D. R. (1977). Automatic Construction of Algorithms and Data Structures Using a Knowledge Base of Programming Rules. Ph.D. Thesis. Stanford University, Memo AIM-308.

Carbonell, J. G. (1979) Subjective Understanding: Computer Models of Belief Systems. PhD Thesis. Yale University research report #150.

Cullingford, R. (1977). Script Application: Computer Understanding of Newspaper Stories. Ph.D. Thesis. Research Report #116, Yale University.

Meehan, J. R. (1976). The Metanovel: Writing Stories by Computer. Ph.D. Thesis. Research Report #74, Yale University.

Newell A., and Simon, H. A. (1972). Human Problem Solving. Prentice-Hall, Inc., New Jersey.

Newell, A. (1973) Production Systems: Models of Control structures. In W. G. Chase (Ed.), Visual Information Processing. Academic Press. New York.

Sacerdoti, E. D. (1977). A Structure for Plans and Behavior. Elsevier North-Holland Pub. Co. Amsterdam.

Schank, R. C. (1975). Conceptual Information Processing. North Holland Pub. Co. Amsterdam.

Schank, R. C. and Abelson R. P. (1977). Scripts, Goals, Plans and Understanding. Lawrence Erlbaum. Hillside, NJ.

Schmidt, C., Sridharan N. and Goodson, J. (1978), The Plan Recognition Problem. Artificial Intelligence, Vol 11, No. 1,2. 45-83.

Wilensky, R. (1978). Understanding Goal-Based Stories. Ph.D. Thesis, Research report #140, Yale University.