PLAN RECOGNITION FOR AIRBORNE TACTICAL DECISION MAKING

Jerome Azarewicz
Glenn Fala
Ralph Fink
Christof Heithecker

Naval Air Development Center
Warminster, PA 18974

ABSTRACT

Airborne tactical decision making is degraded as a result of sophisticated threat capabilities, high data rates and uncertainties, and the necessity for timely response. Under investigation at the Naval Air Development Center is the concept of a plan recognition model to assist the tactical decision maker in interpreting and predicting the activities of enemy platforms.* On-going work in the field of plan recognition was surveyed, knowledge acquisition conducted, and a prototype plan recognition model has emerged. The model is a hierarchical, blackboard based adaptation of a more general architecture of cognition. The model attempts to overcome some of the perceived shortfalls of other approaches relative to the complexities of the tactical situation. Extensions to accommodate uncertain events and elusive goals in multi-hypothesis situations are the focus of current activities.

I INTRODUCTION

Command and Control decision makers aboard Naval aircraft face difficult tasks in assessing and acting upon the activities of enemy platforms (e.g., aircraft) in a tactical environment. Enemy activity is monitored via on-board and remote sensor systems. In high-threat situations, large volumes of sensor data must be analyzed and correlated in real time in order to construct an accurate representation of the situation as it unfolds. The data arrives quickly and may be incomplete, inaccurate, and ambiguous as a result of sensor limitations, threat deception, and other factors. A tremendous burden is placed on the decision maker, who must absorb and assimilate this data to make time-critical tactical decisions on which the survival of the task force may depend.

A key factor in intelligent tactical decision making is the correct interpretation of the tactical situation. The interpretation process can be cast as a form of plan recognition which asserts that the tactical observer interprets the activity of enemy platforms by hypothesizing their goals and inferring the plans that are being carried out in order to achieve the goals. An automated, on-line plan recognition model would serve to assist the decision maker in real-time tactical situation interpretation, alerting him to significant events and trends. Feasibility of this concept is under investigation at the Naval Air Development Center through the development of a prototype Plan Recognition Model (PRM).

Previous work in plan recognition has focused on a number of problem domains, (Litman and Allen, 1984; Wilensky, 1983; Schmidt, 1976; and Carver, Lesser, and McCue, 1984). Although current models provide considerable insight into the plan recognition problem, they have not fully addressed the complexities encountered in domains such as that of tactical situation assessment. If one compares the present status of plan recognition problems against what is required for the tactical problem, several shortfalls can be identified as follows:

Present Plan Recognition	Tactical Problem
Single Agent	Multiple Agents-Independent or Interacting
Cooperative Situations	Adversarial Situations
Well Defined Goals	Elusive Goals
Known Events	Uncertain Events
Limited Hypotheses Set	Large Hypotheses Set
Time Factor is Negligible	Time Factor is Critical

It is the intent of the current work at NADC to overcome some of these limitations in current plan recognition models, especially in the areas of increasing the hypothesis set and analyzing adversarial situations with uncertain events. The remainder of the paper gives an overview of the design and operation of the single agent/multiple hypotheses PRM architecture.

---

*The work described in this paper has been supported by Naval Air Development Center Independent Research, the Office of Naval Technology, and the Naval Air Systems Command.

II PRM: COGNITIVE ARCHITECTURE OVERVIEW

The design of the PRM was largely derived from the work done by Anderson on the Architecture of Cognition (Anderson, 1983). Figure 1 is an overview of our interpretation of three major components of this theory: Long Term Memory, Procedural Memory, and Short Term Memory.

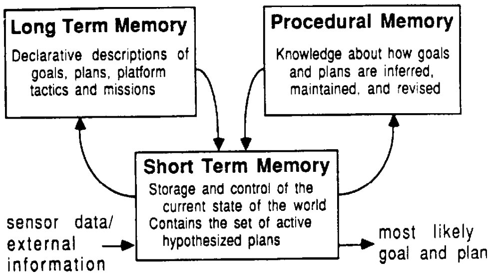

FIGURE 1. PLAN RECOGNITION MODEL ARCHITECTURE

Long term memory (LTM) consists of a declarative description of the set of missions which can be held by an agent under observation. Each mission is represented by a class of plans which can be performed to achieve a particular goal. The structure of these plan descriptions is hierarchical. Plans can be decomposed into a sequence of events, which, in turn, can be decomposed into a set of parameters (Figure 2). Constraints, the necessary conditions for the occurrence of an object in the LTM, are imposed at each level (plan, event, or parameters) of the plan hierarchy. If a constraint has been violated, then the LTM object to which the constraint is attached cannot have occurred.

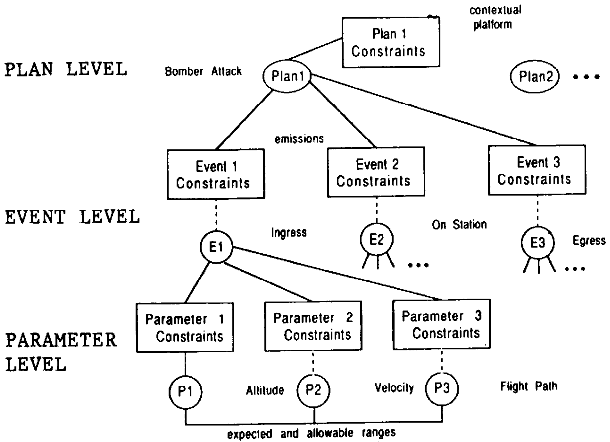

Figure 2. Plan Structure Organization

Procedural Memory (PM) contains knowledge about how to reason about the state of the world. Using the knowledge resident in both short term memory and long term memory, procedural memory is the vehicle for matching observations of the environment to plan structures, postulating goals, weighing hypotheses and revising plans. Components of PM include search strategies and utilities for matching objects stored in LTM.

Short term memory (STM) is the storage space for the intermediate results of the plan recognition process. Knowledge about the current state of the world resides here, such as measurements of the environment, the history of observed events, the set of active hypothesized plans, and the measures of belief and disbelief associated with these plans. STM is the medium through which procedural memory interacts with LTM; procedural memory will only act on those plans stored in LTM which have been retrieved and posted on STM.

III PRM COMPONENT DESCRIPTIONS

In the following sections, the PRM is discussed in the domain-specific context of airborne tactical decision making. Knowledge acquisition activities were carried out to extract and represent the declarative and procedural knowledge the decision maker brings to bear in the plan recognition process. Work has focused on achieving a PRM that can generate and maintain multiple hypotheses to interpret and predict the behavior of a single threat platform. Extensions to accommodate multiple platforms are currently under study but will not be discussed below.

A. Long Term Memory

LTM consists of a declarative description of a set of possible missions held by an observed platform. Each mission represents a class of plans whose successful execution results in the achievement of the mission goal. For example, the goal of the mission class ATTACK can be achieved by invoking one of several available plans or variations of plans in the ATTACK class. Several different platforms may have the same mission goal but use different plans to achieve it. On the other hand, platform limitations restrict the set of plans that may be executed to achieve a particular mission. The plan description must capture the significant features of the ways a mission goal may be achieved. The following is a discussion of the two components of each plan structure in LTM: the plan hierarchy and the Deterministic Finite Automata (DFA).

The hierarchical component defines a plan as tiered structure (Figure 2) consisting of a plan name, a set of events, various parameters pertaining to the events, and constraints at each tier in the hierarchy.

Constraints in the plan hierarchy form a nested set of necessary conditions which must be satisfied for entry to a tier. The matching of environment measurements against objects in the plan structure involves checking the constraints within the structure of the plan, event, and parameter levels. If a measurement of the environment satisfies the plan constraints, the plan is plausible and the event constraints may be checked; if these are satisfied for a particular event, then that event's parameter constraints may be checked. The constraints serve to reduce the search spaces of both the plans possibly held by the platform and the event being performed within that plan.

The events are temporarily ordered features which a tactical decision maker judges to be significant for hypothesizing of inferring a plan held by a platform. The events, in turn, consist of a set of observable parameters and their associated constraints. Parameters that define an event represent measures of the platform's kinematics (platform motion) and emissions (electromagnetic signals) behavior that are expected and allowable within that event. Each parameter that characterizes an event has an expected value or interval of values. To allow for variations in the unfolding of events, the constraints extend the expected parameter value range out to the maximum allowable range for the particular event. If the event constraints are satisfied by an observed measurement, then it is possible that the platform is executing this event in the hypothesized plan. Otherwise, this event cannot be held by the platform.

Event boundaries in the plan structures are defined in terms of significant transitions in expected parameter values. The events in Table 1 (E1, E2, etc.) are delineated with respect to a set of expected parameter values over an expected range. When there is a significant change in the expected parameter values, a transition from one event to the next occurs. For instance, the dramatic change in the altitude of the bomber in Table 1 indicates a transition from E1 to E2.

TABLE 1. PARTIAL DECLARATIVE DESCRIPTION BOMBER ATTACK PLAN

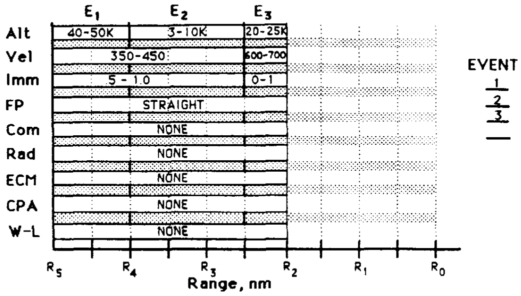

The event constraints serve to extend the expected event boundaries out to the maximum allowable parameter values for each event (Table 2). Because of these extended event parameter boundaries, events may overlap in time. The relaxation of expected parameter values through constraints provides for the gradual transition from one event to the next. Knowledge about to which event the platform will transition is embodied within the DPA. This concept is discussed at the end of this section.

TABLE 2. SAMPLE EVENT CONSTRAINTS FOR BOMBER ATTACK PLAN

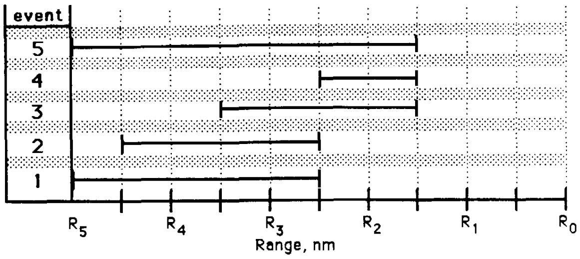

A measure of belief for the platform's events and state can be obtained by comparing the actual platform measurements against the expected/allowable parameter values. Rules attached to the parameter values in the plan structure are invoked to determine the degree of membership and associated belief in the event. Attached to each kinematic parameter are linear functions which are used to determine the degree-of-match (DOM) between an observed value for a parameter and the expected/allowable values for that parameter (Table 3). Utilizing kinematic information of the measurement, the linear functions return a value which falls into one of three classes: exact, partial, and none. The three results can be interpreted as follows:

Exact: The measurement falls within the expected range of values--the measurement is a member of the event-parameter set.

Partial: The measurement falls within the allowable range of values--the measurement is a plausible member of the event-parameter set.

None: The measurement falls outside of the allowable range of values--the event is removed from consideration.

TABLE 3. SAMPLE KINEMATIC PARAMETERS FOR EVENT 1 OF BOMBER ATTACK PLAN

Interpretation of Matcher Results

	Exact	Partial	None
Alt	range: 40-50K $b_{alt} = 1$	range: 25-40K $b_{alt} = \frac{1}{15000} alt - 5/3$	range: 0-25K, >50K $b_{alt} = 0 \rightarrow$ remove event
Vel	range: 350-450 $b_{vel} = 1$	range: 0-350 $b_{vel} = \frac{1}{350} vel$ range: 450-750 $b_{vel} = -\frac{1}{250} vel + 2.8$	range: > 700 $b_{vel} = 0 \rightarrow$ remove event
Imm	range: 0.5 - 1 $b_{imm} = 1$	range: 0 - .5 $b_{imm} = 2\ imm$	range: -1.0 - 0 $b_{imm} = 0 \rightarrow$ remove event

Knowledge about emission parameters represents the decision maker's heuristics used to support or refute a plan hypothesis (Table 4). This differs from the more computationally intensive kinematic match. A series of actions are initiated when an emission measurement satisfies certain preconditions; e.g., a message may be issued to increase or decrease the belief in the plan, or even remove a plan from consideration. For example if a bomber ID signal is observed, a message will be posted to remove all plans that cannot be held by a bomber. Values for the emissions parameters represent strong evidence to support or refute a hypothesized plan.

TABLE 4. EMISSIONS PARAMETERS FOR EVENT 1 OF BOMBER ATTACK PLAN

	Signal Present	Signal Not Present
Com	decrease the belief in the plan	$\bar{d}$
Rad	if TAR see RAD Rule (3) below if other RAD then remove plan	$\bar{d}$
ID	see ID Rule (1) below	$\bar{d}$
ECM	remove plan	$\bar{d}$
W-L	see W-L Rule (2) below	$\bar{d}$

RULES:

(1) If ID equals Bomber then remove all non-Bomber plans from the Expectation board and decrease the belief in the plan
If ID does not equal Bomber then remove all non-ID plans

(2) If W-L equals yes then assert PREMATURE MISSILE LAUNCH

(3) If RAD equals TAR then assert PREMATURE RADAR ON

The second major component of the plan structures is captured by the DFA representation. The search algorithm of PM utilizes a description of a state transition diagram, which is the DFA stored in long term memory with each plan (Figure 3). The DFA of each plan specifies the set of legal sequences of events which must be performed in order to complete the plan. The state transition diagram indicates the various states a platform may be in and the significant events that are needed to transition to the next platform state. The diagram is very dependent upon the domain expert's characterization of a typical platform mission.

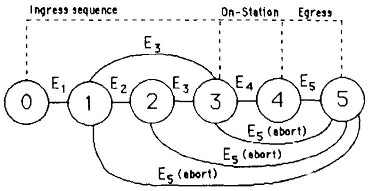

Event From To	Transition Condition
$E_1 E_2 E_3 E_4 E_5$	expected transition sequence - check for weapon-launch (W-L) in E_4 for successful ATTACK
$E_1 E_3 E_4 E_5$	optional transition sequence bypassing E_2, low flight - check for W-L as above

FIGURE 3. STATE TRANSITION DIAGRAM FOR THE EVENTS IN BOMBER ATTACK PLAN

In summary, LTM contains a set of plans which are represented as hierarchies and as deterministic finite automata. The hierarchies allow the portrayal of the tier structure of plans, events, and parameters. The DFAs are used to specify the sequences of events in which plans may unfold. The combination of these two structures is well-suited for representing knowledge about plans.

B. Procedural Memory

PM is the mechanism used to hypothesize the plan and sequence of events carried out by the platform given the declarative knowledge in LTM (plan hierarchies and DFA's), contextual information, and the history of observations of platform behavior. This discussion focuses upon the search and match strategies in PM used to estimate the current event within a hypothesized plan and the state of the platform within the DFA.

The goal of the search is to identify the current state of the platform. This is achieved through the recognition of a sequence of events which have been observed through measurements of the platform's behavior. This sequence of events is a partial instantiation of a hypothesized plan. When there is strong evidence that an event has occurred, the event is entered into the partial plan instantiation.

Uninformed strategies such as the depth-first and breadth-first searches would be inefficient in this application, since these strategies rely upon a cost function which allows for a large search space. The declarative knowledge about plans resident in LTM provides the capability to instead implement an informed search such as the best-first. The best-first algorithm allows us to utilize knowledge about the problem domain, an estimate of the state of the platform, a probable goal of the platform, and the information gathered by the search to determine the plan held by the platform (Pearl, 1984).

Since the state of a platform depends upon the event which is currently being performed by the platform, the searcher needs a heuristic evaluation function for determining the current event of the platform. This is the function of the event matcher. The event matcher consists of two subordinate matchers: the constraint matcher and parameter matcher. The constraint matcher checks if there are any violations of the event constraints; if the constraints are violated, then the event being matched cannot be occurring. Violations of the constraints narrow the search for a hypothesized plan and the event.

If the event constraints are not violated, the parameter matcher is invoked to compare the current measurement of the environment with the allowable values for each parameter slot of the event. The parameter matcher will return a degree of membership of the observed measurements in the set of allowable values for each parameter. These degrees of membership are combined to obtain the overall belief in the occurrence of the event.

The event weights supplied by the heuristic evaluation function are used to detect a change in the platform state, i.e., a transition from one event to the next. These weights are stored in a linked tree structure for each plan (Figure 4). For backtracking purposes, attached to each node of the tree is information about the type of decision made by the searcher. At each observation the platform state is determined on the basis of these weights. The search path from the current estimate of the state to the initial guess of the platform state represents the partial instantiation of a plan.

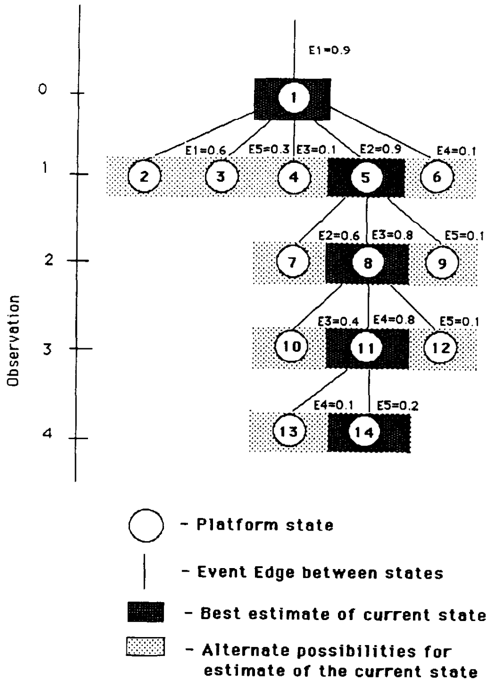

FIGURE 4. BEST FIRST SEARCH

PM uses STM as a means of accessing the hierarchical plan structures and the DFAs stored in LTM. We would not want to access LTM directly, since we are only maintaining a local search of the plan structures. It is for this reason that the knowledge about the current state of the hypothesized plans is maintained in STM. This is discussed in the next section.

C. Short Term Memory

STM is a blackboard used as a workspace for the intermediate results of the PRM process. As can be seen from Figure 5 there are five partitions of the blackboard workspace within STM.

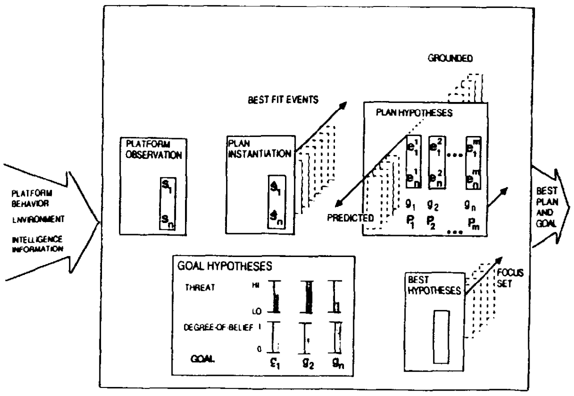

FIGURE 5. SHORT TERM MEMORY

The Platform Observation Board integrates the measurement of the observed platform behavior with contextual information. This is a mechanism to supplement the information observed by sensors to aid in the recognition of events. As the time into the mission increases, the emphasis on the observed platform behavior as a measure of an event increases, while the emphasis on contextual information decreases.

The best fit events are captured on the Plan Instantiation Board. For a given sequence of observations, the event closest to representing the platform's actions is selected. Conceptually, segments of various plans are being assembled to represent the platform's actions.

The Plan Hypotheses (PH) Board consists of a set of plans which the search and match processes indicate are likely to be held by the platform. Each hypothesized plan is a template retrieved from LTM. The system strives to maintain at least one plan on the board at all times. The plans on the PH board have the following characteristics: 1. Each hypothesized plan is assigned a weight which represents the accumulated evidence that the template fits the observations; 2. The events of each hypothesized plan are tagged as either "grounded", "current", or "unmarked". A grounded event was observed, matched at least partially with the measurement of the platform behavior, and is no longer occurring. The current event best matches the current observations. An unmarked event has just gotten underway or has yet to be observed.

A hypothesized plan contains the set of possible next events which describe future platform behavior and indicate a possible goal. The predicted events consist of a sequence of events or a path determined by the DFA.

The sequence of grounded events attached to a plan on the PH Board is a path through the DFA which was supported by measurements: there was sufficient evidence to state that these events occurred. With each plan is a degree of belief; i.e., a measure of how well the hypothesized plan fits the plan templates in LTM. The current event is the one best supported by evidence from the most recent observations.

The Goal Hypotheses Board contains the goals of each of the hypothesized plans, the threat level associated with the goal, and the degree of belief that the platform holds this goal. The threat level is an indication of the importance of maintaining a plan. For example, a plan and goal may have low degrees of belief, yet their high threat capability will merit their maintenance on the board.

Given a set of available focusing heuristics (Carver, Lesser, and McCue, 1984), the Best Hypotheses (BH) Board maintains a history of the most likely hypothesized plans and goals for the platform. The plan and goal hypotheses held on this board are the output of PRM.

To summarize, STM supplies a cache of memory for the intermediate results of the search and match processes, a workspace for the blackboards, and control information essential to the PRM process (described in the next section).

IV PRM PROCESS DESCRIPTION

The current PRM is being implemented in a frame-based blackboard architecture. In this representation, STM is a blackboard which consists of several contributing blackboards (Figure 6), each of which is associated with various experts which make up the PRM. These experts consist of a STM Manager (i.e., the executive), a Plan Hypotheses (PH) Manager, a Plan Expert, an Event Expert, and a Parameter Expert.

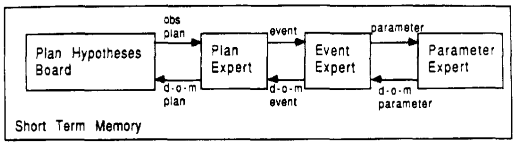

FIGURE 6. STM BLACKBOARD HIERARCHY

The STM Manager will oversee the operation of the PRM process on the STM blackboard. It functions as an executive directing top-level control of the PRM. STM thus is a workspace for the various knowledge sources (experts) which conduct the various facets of plan recognition. The STM Manager is responsible for handling messages sent from the user (primarily at start-up time) and messages passed back from the PH Manager that may have been generated by either the PH Manager or one of the other experts at a different level in the hierarchy of the process. In response to these messages, the STM Manager has duties pertaining to initializing the system, getting observations, and system restart/recovery.

When all processing is complete at the executive level, control is generally passed to the PH Manager. The Plan, Event, and Parameter Experts deposit and withdraw information from the PH workspace. On this workspace, the set of all possible plans, given the state of the system, are posted. The PH Manager is responsible for managing this workspace by passing messages to both the STM Manager and the Plan Expert.

The PH Manager maintains the set of all possible plans that may be held by the platform given the current state of the system. The PH Manager receives messages that refer to modifying the current state of the PH Board. These messages may be from the Plan Expert, such as a "Remove this plan" message, or from the STM Manager, such as "Instantiate plans that correspond to the current observation." The PH Manager also sends messages to the STM Manager and to the Plan Expert. These messages may be "PH Board empty, take appropriate action" or "Match the current observation to this plan."

Each plan on the PH board is a partial instantiation of some plan that exists in LTM. Each partially instantiated plan contains information referring to how well it corresponds to the observed platform behavior. Recall that a plan includes information regarding the platform's most likely current state and event, a list representing the history of observed (grounded) events, and the set of expected (extrapolated) states and events, all with respect to the plan in LTM. The maintenance of the plan on the PH board is the responsibility of the Plan Expert.

The Plan Expert receives messages concerning which plans to deal with and in what order, given the current state of the system. The Plan Expert invokes the knowledge about searching, i.e., which events to process for each plan and in what order. This knowledge is represented as the DFA that reside in LTM. The Plan Expert sends a match message to the Event Expert to process prescribed events. As a result of this processing a DOM is assigned for each plan and a message to remove invalid plans may be sent. The Plan Expert only has knowledge about high-level plan maintenance such as transitioning to the next event of each knowledge sources invoked by the Plan Expert to perform the detailed steps of the plan maintenance process.

When the Event Expert receives the match message and the current observation from the Plan Expert, the event constraints are first checked. If an event's constraints are violated, the Event Expert passes a message to the Plan Expert to remove that event from consideration as the current event, otherwise, the Event Expert invokes the Parameter Expert to process each of the individual parameters corresponding to that event. The Parameter Expert does the actual matching of observed data against the parameter values that would be expected and are allowable. As a result of the matching phase, the degree to which a parameter matched is returned.

V STATUS AND FUTURE WORK

The PRM described above was developed and implemented on the Symbolics using the Flavors package. The plan library contains structures for nine threat scenarios gleaned from knowledge acquisition activities over a six month period. The simulation testbed system can provide the scenarios and variations of them to PRM in terms of input files of kinematic and emissions parameters for the threat of interest. Output to the user consists of a dynamic depiction of the emerging tactical scenario, as well as continual updates of the hypotheses being maintained and their associated belief measures.

The PRM is currently a single-threat/ multi-hypothesis model. Although it can maintain multiple hypotheses to explain the threat behavior, it is limited to an analysis of a single threat agent. Extensions to accommodate multiple threat agents are under investigation. Such agents could be pursuing multiple goals independently but more likely will be working in concert to effect a single goal. Effective management and pruning of the search space will become paramount in a multi-threat environment. A judicious mix of data driven and goal driven processing will need to be invoked. Emphasis must be placed on the key features and indicators that serve to discriminate among the different possible plans.

Also under scrutiny are richer representations for the heuristic evaluation function and the handling of uncertain data and knowledge. Extensions to multi-valued logic, fuzzy schemata, and work on the theory of endorsements are potential candidates. Ongoing efforts in machine learning and reasoning by analogy could find suitable application here.

Finally there is the issue of time-critical operation. In a fleet air defense setting, correct interpretation of threat actions and appropriate friendly force response formulation must occur in a matter of minutes or event seconds. In addition to opportune search strategies and the use of focusing, it will be necessary to exploit the parallelism inherent in a multi-threat/multi-hypothesis plan recognition environment. The issue of mapping the PRM architecture onto various parallel processing topologies has received considerable attention in ongoing work at NADC.

VI ACKNOWLEGDEMENTS

In the area of knowledge acquisition, Joe Alfano has made significant contributions to the PRM project. Raymond Kirsch of LaSalle University has been instrumental in addressing the exploitation of parallelism in machine architectures for AI computing in general and the PRM in particular.

REFERENCES

[1] Anderson, J., The Architecture of Cognition, Harvard University Press: Cambridge, MA, 1983.

[2] Carver, N., V. Lesser, and D. McCue, "Focusing in Plan Recognition", The National Conference on Artificial Intelligence, AAA1-84, Austin, Texas, August 1984.

[3] Litman, D., and J. Allen, "A Plan Recognition Model for Subdialogues in Conversations", Technical Report 141, Department of Computer Science, University of Rochester, Rochester, N.Y., November 1984.

[4] Pearl, J., Heuristics: Intelligent Search Strategies for Computer Problem Solving, Addison-Wesley: Reading, MA, 1984.

[5] Schank, R. and R. Abelson, Scripts, Plans, Goals, and Understanding - An Inquiry into Human Knowledge Structures, Lawrence Erlbaum Associates, 1977.

[6] Schmidt, C., N. Sridharan, and J. Goodson, The Plan Recognition Problem, An Intersection of Psychology and Artifical Intelligence", Artifical Intelligence, 11: 45-83, August 1978.

[7] Stenerson, R., "Integrating AI into the avionics engineering environment," IEEE Computer, February 1986.

[8] Wilensky, R., Planning and Understanding - A Computational Approach to Human Reasoning, Addison-Wesley, 1983.