S31-63

4P.

188113

Mission Planning for Autonomous Systems

G. Pearson
FMC Corporation
Santa Clara. CA 95052

FY 563409

1. Abstract

Planning is a necessary task for intelligent, adaptive systems operating independently of human controllers. ~~This paper reports—on~~ a mission planning system that performs task planning by decomposing a high-level mission objective into subtasks and synthesizing a plan for those tasks at varying levels of abstraction. We use a blackboard architecture to partition the search space and direct the focus of attention of the planner. Using advanced planning techniques, we can control plan synthesis for the complex planning tasks involved in mission planning.

Research

2. Introduction

Planning is a necessary task for intelligent, adaptive systems operating independently of human controllers. Autonomous systems need to plan their actions and adapt themselves to environmental changes for survival. Given a high level mission specification, the mission planning module needs to synthesize a sequence of actions to achieve mission goals. This requires advanced techniques that reason about constraints, granularity of search, spatial configurations, levels of abstraction and temporal orderings. Robotic systems benefit from these planning techniques by increasing their independence from mission control. As a result, expanded missions with reduced supervisory control can be conceived and executed[1].

Various approaches are being used in planning systems. STRIPS uses means-ends analysis in robot problem solving by identifying differences from a state description to a goal and selecting forward production rules to reduce them[2]. NOAH uses levels of abstraction and a least-commitment strategy to generate parallel plans hierarchically for an indoor mobile robot[3]. OPM uses opportunistic reasoning to solve the errand planning problem by combining both top-down and bottom-up planning activities[4]. MOLGEN uses skeletal refinement to plan experiments in molecular biology by instantiating specific operations from a sequence of generalized plan steps[5]. PROTEAN uses a blackboard architecture to configure protein structures by reasoning about the solution to the problem as well as the problem-solving process[6]. Our mission planner builds on these earlier systems and uses key aspects from each. It is implemented within the BB1 blackboard architecture[7] and uses the features offered by BB1.

3. Mission Planning

The overall control of autonomous systems requires the management of multiple subsystems cooperating to achieve mission goals. One such subsystem is mission planning. This paper reports on a mission planning system built for the autonomous operation of FMC's autonomous land vehicle, an M113 tracked vehicle.

Mission planning is the process of synthesizing a sequence of actions to satisfy goals and constraints posed by the mission manager. Mission plans are specified at varying levels of abstraction, with mission profiles at the higher levels and command sequences at the lower levels. Command sequences are fixed to perform specific tasks given the vehicle's operating characteristics; as such, they are best sequenced by computer programs that perform table lookups. Contention among these commands can usually be resolved at the higher levels, thus interaction among commands is minimal. At the other end of the planning hierarchy, mission profiles specify objectives and time tables for accomplishing the objectives. The vehicle achieves these objectives by executing command sequences downloaded by Mission Control at the appropriate times specified in the mission profile. Mission profiles lend themselves to template or script planning because they are specified at a level of detail higher in the abstraction hierarchy where interaction among the objectives is minimal.

Tasks, on the other hand, are synthesized into plans by considering the current state of the mission. Tasks consist of command sequences an autonomous vehicle executes to achieve some part of the overall mission. The planner performs task planning by decomposing the high-level mission objective into subtasks and synthesizing a plan for those tasks at varying levels of abstraction. Intermediate tasks must be selected and sequenced in such a way that subsequent goals can be achieved. An exemplary task performed by the planner is to develop a plan for conducting reconnaissance in a particular area specified by the mission commander. For example, a mission to conduct area reconnaissance is necessary when the commander desires specific information about certain locations or facilities within a defined area. To accomplish this mission, the planner must find overwatch positions for reconnoitering the targets, establish routes sequencing these positions, retain stealth of the operation and report all information rapidly and accurately[3].

Tasks refer to the intermediate abstraction levels in the planning hierarchy, those levels where interaction among planning decisions is the highest. Interaction among tasks either by sequencing tasks with prerequisite and postrequisite conditions or by decomposing tasks into subtasks makes up the interesting planning problem. These interactions occur in a dynamically changing environment and create a combinatorial explosion of the planning space; the search through this dynamic planning space is a key issue for mission planners.

4. Blackboard Architecture

We are implementing our mission planner using BEDS, a Blackboard Event Driven System[8], a version of the BB1

PAGE 302 INTENTIONALLY BLANK

blackboard architecture[7]. As such, it defines problem-solving knowledge sources for synthesizing plan steps, a multi-level solution blackboard for recording partial plans and a flexible control structure for controlling the expansion of the planning space.

Using a blackboard, a hierarchy of abstraction levels where each level represents a partial state description of the world at some time, we can partition the search space and direct the focus of attention of the planner. We map the problem space onto the blackboard by specifying abstraction levels in the plan hierarchy. These levels represent both spatial and conceptual abstractions for the mission planning problem. For the mission of area reconnaissance, we generate a visibility map by creating boundary regions that contain locations visible to the target—a spatial abstraction. For the path planning task of the mission, we generate one type of nontrafficable region by creating water boundaries—a conceptual abstraction. Data abstractions help control the exponential search process required in planning by establishing planning islands; areas where local search can find plan anchors for attaching the remainder of the plan. The more independent the planning islands, the better the planner controls its planning space by relying on local search. The blackboard structures the planning space in the problem domain. To this structure, we apply the problem-solving strategy of skeletal plan refinement.

When a mission is specified, the planner chooses a general design. We specify a design with only the essential detail necessary to direct the initial search of the plan. This least-commitment strategy is maintained throughout the plan refinement process. The design specifies spatial configurations for plans and partitions the planning space into plan segments. Once these segments are found, the planner successively refines its plan by instantiating plan steps at the lower levels. Plan instantiations occur by creating planning elements using the correct data abstraction with the current plan abstraction. At the design level, the planner cannot use low-level data to form decisions. Instead, it uses high-level symbolic objects that represent the relationships between the tasks that make up the mission.

For example, consider a plan for a reconnaissance mission that synthesizes a sequence of tasks in both time and space such that the final configuration satisfies overriding objectives. A good design specifies the spatial orientation for each of the tasks. Finding this design depends on the reconnaissance tasks involved and their relationships to each other. At this level of abstraction, the planner reasons about the target location, the type of reconnaissance mission, visibility maps, nontrafficable regions, military strategy and communication requirements. Only after refinement of the design can plan steps involving exact task locations be instantiated using pixel data represented as coordinate triples.

4.1. Controlling Plan Synthesis

Plan synthesis occurs when knowledge sources instantiate plan steps recorded in the blackboard hierarchy. Without controlling plan synthesis, the planning system would exhaustively create the solution space of possible plans. While this works for simple planning problems, in mission planning, as the complexity of the mission increases, the number of tasks grows and the number of potential plans grows exponentially. We use a three-tiered structure for varying control over the execution of knowledge sources in the mission planning system that consists of *establishing focus decisions*, *executing strategies* and *ranking knowledge sources*. During problem solving, knowledge sources create decision elements in the plan hierarchy—as planning proceeds, more knowledge sources are activated and become available for execution. A controller rates these knowledge sources using focus decisions, strategies and rankings, and a scheduler selects a knowledge source to execute by choosing the one with the highest rating.

Focus decisions represent collections of heuristics against which knowledge sources are rated[6]. These decisions establish criteria used to evaluate the utility of knowledge sources. For each knowledge source the controller calculates a utility value by summing together, for each focus decision, the products of a focus weight representing the value of a focus decision and a satisfaction level, the degree to which a knowledge source satisfies a focus decision. This calculation results in ratings that prioritize the knowledge sources so a scheduler can select the knowledge source with the highest rating. Focus decisions are created during problem solving in response to changes in planning and reflect the general behavior of the system. They add high-level control decisions that the controller uses to direct the generation of plan steps.

Strategies provide a rigid control structure that directly controls the execution of knowledge sources. They permit the execution of a strict sequence of knowledge sources. A strategy represents a procedure for achieving a particular goal and consists of a goal, a status, a rationale and a list of strategies and tactics. The goal denotes what the strategy will accomplish when its status becomes operative, and the rationale describes what the strategy accomplishes. The ordered list of strategies and tactics defines the specific subgoals that make up the procedure. When strategies are operative, knowledge sources that achieve the same goals of the operative strategies receive higher priorities than ones that achieve different goals. Focus decisions are used to differentiate between knowledge sources with the same goals.

```
STRATEGY: FIND-LOCATION
   goal = FIND-LOCATION
   status = OPERATIVE
   rationale = "Instantiate best location for
                  performing a task"
   strategy&tactic = (INSTANTIATE-LOCATION
                      RATE-LOCATION
                      CHOOSE-LOCATION)

STRATEGY: FIND-R1
   goal = FIND-R1
   status = OPERATIVE
   rationale = "Controls search for R1"
   strategy&tactic = (INSTANTIATE-LOCAL-AREA-R1
                      FIND-LOCATION)
```

Figure 4-1: FIND-LOCATION and FIND-R1 strategies.

Figure 4-1 illustrates the structure of two strategies used by the Mission Planning System. The first strategy, FIND-LOCATION, consists of three tactics: INSTANTIATE-LOCATION, RATE-LOCATION and CHOOSE-BEST-LOCATION. This strategy finds a location by creating instances of locations, rating them and choosing the best one. The second strategy, FIND-R1, consists of the tactic INSTANTIATE-LOCAL-AREA-R1 and the strategy FIND-LOCATION. This recursive definition facilitates creating new strategies from existing ones. This strategy finds a position for performing reconnaissance by creating instances of local areas then finding locations within these areas.

The third level of control in this three-tiered structure, ranking knowledge sources, overlaps with the preceeding two. Ranking prioritizes knowledge sources that are grouped together because of similarities in function or strategy. During system design, knowledge sources are ranked to differentiate between subtleties in their performance characteristics. Usually, performance refers to processing speed with processing speed determining the granularity of search. Ranking gives the controller a discriminating factor when it chooses among knowledge sources with the same rating. Thus, ranking discriminates between knowledge sources that would otherwise be considered equal. Knowledge sources with the special rank of IMMEDIATE bypass the controller and execute immediately.

5. Constraint-based Reasoning

Another technique for controlling search is using constraints to limit the number of acceptable plans. Our planner uses constraints based on terrain feature information, resource limitations, vehicle limitations and military doctrine. In this way, the space of possible planning solutions is constrained by the specifications of the mission requirements. A mission must meet certain objectives while satisfying constraints that limit the success of an operation. The harder* the constraints, the less flexible the plan and the easier it is to confine the search. As constraints become softer** they contribute less to confining the space of possible plans. Our mission planner uses hard constraints to limit the number of acceptable plans by reasoning about terrain feature information, resource limitations, vehicle limitations and military doctrine during the planning process.

Mission planning is data intensive; therefore, search must be performed using different levels of granularity. We use a strategy that satisfies hard constraints before considering the soft constraints. Failure to satisfy any of the hard constraints results in the planner either terminating its search or backtracking by considering new potential solutions. Our planner performs simple backtracking by expanding its search through the data base. By continually increasing the resolution of the search, the planner increases the number of data points that it considers during planning. This technique allows it to make uniform cuts through the planning space as it refines plans top-down through the plan hierarchy. The planner performs more complicated backtracking by modifying constraints, thereby achieving the objective but with some loss of optimality. The planner relaxes constraints by propagating hard constraints. In our example of a reconnaissance mission, an original constraint is to use the reconnaissance technique of triangulation. However, when this constraint cannot be satisfied, the system relaxes it into one that allows straight reconnaissance. The constraint remains a hard constraint; it must be satisfied to complete the mission, but a plan allowing for straight reconnaissance is less desirable than one that uses triangulation.

Our planner works first from hard constraints to find a solution and backtracks only when necessary. It uses soft constraints, but considers them with less priority. Using constrained search, we confine the planning space to one that satisfies the hard constraints then find a solution that satisfies most of the soft constraints. As an example, consider the problem to place different sized objects in a leather pouch. A effectively simple strategy places the larger items in first, then squeezes in the smaller ones. This strategy works well because the planning space is characterized by the flexibility of the leather pouch. Mission planning has a similar flexible planning space because missions are defined to encompass changing conditions that occur during the operation of a mission. Thus, we use a strategy that satisfies the hard constraints, much like placing the larger items in the leather pouch first, then squeezes the soft constraints into place.

* Hard constraints refer to those constraints that must be satisfied when finding a valid solution.

** Soft constraints refer to those constraints that may or may not be satisfied when finding a valid solution.

6. Mission Planning Results

The Mission Planning System is written in ZetaLisp on a Symbolics Lisp Machine and interfaces to other software modules needed to control FMC's autonomous land vehicle. The planning system consists of 105 domain knowledge sources, 8 control knowledge sources, 27 strategies and 17 tactics.

Figure 6-1 shows the planning state of the Mission Planning System in the intermediate stages of finding a plan for conducting reconnaissance. In this example, the planner constructs a plan to reconnoiter an area using triangulation techniques that gather information about the target. It synthesizes a plan by executing knowledge sources and posting information on the blackboard. The blackboard levels are shown from the Strategy level down to the Route level. Strategy, Tactics and Foci make up the Control Blackboard, while Mission, Design, Region, Script, Lot, Location and Route make up the Domain Blackboard. Values are depicted in one of four states: underlined values represent operative criteria that the planner considers when making decisions; values shown in reverse-video denote the current activated subgoals in the planning space; boxed values are the selected positions and routes that make up the final plan; and shaded values (see figure 6-2) depict goals that have been accomplished.

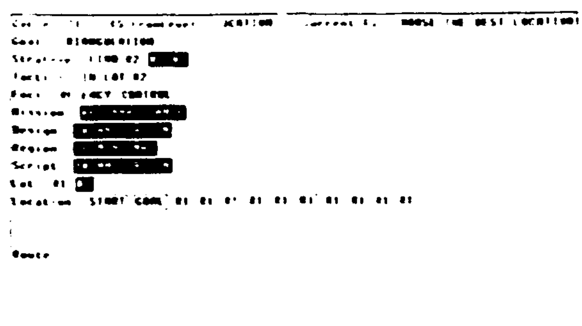

Figure 6-1: Mission Planning State on Cycle 21

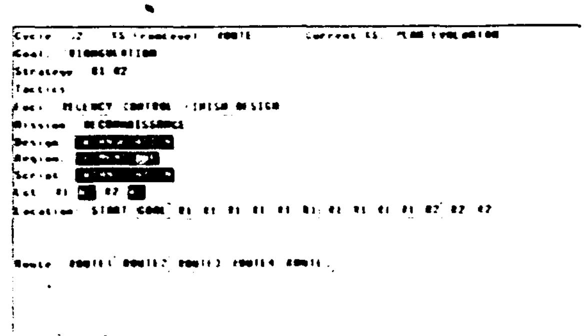

Figure 6-2: Mission Planning State on Cycle 32

After receiving a mission statement from the commander of the vehicle, the planning system begins to find a possible plan. It starts by creating focus decisions that provide top-down control. This behavior changes as the system operates and creates other control decisions. A *design* is selected based on the requirement for triangulation that identifies the tasks necessary to accomplish the mission. These tasks are confined to a *region* that is computed based on vehicle and terrain constraints. With the region computed, the *start* and *goal* locations are posted for reference by the planner. A *script* resulting in a reconnaissance mission using triangulation is selected and a *strategy* for instantiating the script is generated. Here, the planner implements the R1-R2 strategy.

Having generated a strategy, the planner can instantiate *lats*, local areas where the planner searches for reconnaissance locations. It then *finds locations* and *routes* and sequences them into the final plan.

Figure 6-2 shows the planning state after the planner has found one possible plan for performing area reconnaissance using triangulation. The final plan is represented as nodes at the Location and Route levels. In this final plan, the vehicle travels along ROUTE1 from its starting position to the first reconnaissance location, R1. After reconnoitering the target, it travels along ROUTE5 to the second reconnaissance location, R2. At this point, the vehicle triangulates data acquired from the first reconnaissance task and completes the mission by moving to its final destination along ROUTE2.

7. Conclusion

We have built a Mission Planning System capable of sequencing tasks to achieve higher level mission objectives. We have built this system using a blackboard architecture that defines knowledge sources, a multi-level blackboard and a flexible control structure. Using this architecture integrated with other planning techniques, we have some degree of control over the explosive search space inherent in mission planning problems.

8. Acknowledgements

I am indebted to Perry Thorndyke and Shawn Amirsardary for helpful comments on the content of this paper, with special thanks to Jolan Yao for her software support of the Mission Planning System. Barbara Hayes-Roth was responsible for the BB1 system and made it available to FMC.

References

1. Thorndyke, P., McTamaney, L., Kuan, D. & Pearson, G., "The Autonomous Land Vehicle", *Defense Science & Electronics*, Vol. 1 November 1985.

2. Nilsson, N., *Principles of Artificial Intelligence*, Tioga Publishing Co., 1980.

3. Sacerdoti, E. D., "A Structure for Plans and Behavior", Tech. report 109, SRI, AI Center, 1975.

4. Hayes-Roth, B. & Hayes-Roth, F., "A cognitive model of planning", *Cognitive Science*, Vol. 3 1979, pp. 275-310.

5. Friedland, P. & Iwasaki, Y., "The Concept and Implementation of Skeletal Plans", *Journal of Automated Reasoning*, Vol. 1 1985, pp. 161-208.

6. Hayes-Roth, B., et. al., "A Layered Environment for Reasoning about Action", Tech. report KSL 86-38, Stanford University, 1986.

7. Hayes-Roth, B., "A Blackboard Model of Control", Tech. report HPP 83-38, Stanford University, Computer Science Department, 1984.

8. Pearson, G. & Kuan, D., "Mission Planning System for an Autonomous Vehicle", *The Second Conference on Artificial Intelligence Applications*, 1985, pp. 162-167, Miami Beach, Fla.

9. Pearson G., "Mission Planning within the Framework of the Blackboard Model", *To appear in Proceedings from Expert Systems in Government Conference*, 1985, McLean, Va.