H.J. Antonisse, J.W. Benolt, &
B.G. Silverman (Eds.) Proceedings
of the Annual AI Systems in
Government Conference, March 27-31,
1989, Washington, D.C. Computer
Society Press.

REACTIVE PLANNING USING A "SITUATION SPACE"

Charles F. Schmidt[1], John L. Goodson[2]
Stacy C. Marsella[1], John L. Bresina[1]

1. Lab. for Computer Science Research
Rutgers University
New Brunswick, NJ 08903

2. GE Advanced Technology Laboratories
Route 38, Moorestown Corporate Center
Moorestown, NJ 08057

Abstract

The problem of how to plan in tactical situations where planning must be responsive to events, or other agents' actions, which lie outside the predictive capability of the planner is addressed. The basic difficulty presented by this type of planning problem is how to engage in actions that are coherently related to the achievement of the overall goal despite the fact that often the planner can not develop a complete plan for achieving that overall goal. In order to overcome this difficulty we have developed a planning model within which the planner is controlled by knowledge organized into what we have termed a situation space. The situation space guides the selection of goals and the construction of complete subplans which are appropriate to situations that arise and are coherently related to the overall goal. The situation space supports the principled generation of contingency plans and thereby softens the impact of plan failure and replanning in a reactive environment. The model is presented using an example of competitive multiagent interaction.

Introduction

Within AI, a plan is typically defined as some partially ordered sequence of "primitive actions" that solve a problem. The problem consists of initial state and goal state descriptions, both of which are partial descriptions of world states. Execution of the plan transforms a state implied by the initial state description into a state implied by the goal state description.

Standard AI models of planning have typically employed what might be termed a predictability assumption; that is, it is assumed that the planner's model of the world as well as of the effect of its actions on that world is complete and correct (e.g., STRIPS [3] and NOAH [6]). This assumption is violated if either an action fails to achieve its intended effect and/or plan relevant events lie outside the bounds of the planner's ability to control and/or predict.

Even assuming the planner's ability to predict relevant aspects of the world, the planner still may lack the effectiveness required to achieve the goal. There may not exist a sequence of the planner's primitive actions that achieve the goal state from the initial state. In such a case, classical planning models are typically designed to fail; that is, they provide no plan at all. This condition on planning success will be called the effectiveness condition.

There are many types of problems which require some form of planning to achieve an overall goal where either the predictability assumption or the effectiveness condition are violated. We will refer to the problem of how to plan in contexts where at least one of these is violated as the reactive planning problem. Note that when predictability and/or effectiveness are suspect then plan execution must be monitored in order to detect plan failure and trigger some form of replanning. Such replanning is now interleaved with execution and in many contexts this places real-time constraints on the time available for replanning. Consequently, any approach to reactive planning must be sensitive to the issue of how to bound the time required for replanning. This is a particularly difficult problem because there is no guarantee that the failed plan can be revised nor even that there exists a solution to the problem from the state entered on plan failure.

There have been several recent approaches to the reactive planning problem (cf. [7, 5, 1, 4]). Many of these involve identifying and more or less judiciously selecting a plan or action sequence for each state in the problem space that may be encountered. There may be many plans for some of these states and no plan for others. If there are many possible plans for a particular state, then the idea is to select "the best". Thus, if there are m encounterable states and n states (where n is less than or equal to m) from which the goal is reachable, then n such plans are created and indexed to the state to which they apply. Obviously, these n plans in turn affect the number of states that are reachable via failure of some aspect of one of the plans. In particular, selection of the n plans affects p, the number of states reachable via plan failure from which no recovery exists. A judicious selection of these n plans might be one which yields a compact representation of the n plans and/or one which minimizes p. The creation of this "space of plans" is typically carried out prior to execution and thus no time is required for replanning. Replanning simply involves retrieval of the plan indexed on the current state.

Note that the viability of this approach also hinges on its own kind of predictability assumption, i.e. that the planner can identify all of the possible states that are reachable on failure of some aspect of each of the n plans selected. Obviously, the size of n can become quite large even for relatively trivial problems. Further, unless the world is "cooperative" there may often be no principled way to select the **n** plans that avoids a quite large **p**. Achieving complex conjunctive goals that exhibit a high degree of dependency between (sub)goals can yield particularly "uncooperative" worlds. Consequently, the generality of this approach to reactive planning may be quite limited.

The approach to planning developed in this paper was motivated from our consideration of planning within the domain of tactical anti-submarine warfare (ASW). In this problem domain the planner's goals and subgoals are always defined in relation to usually uncertain knowledge about the knowledge, goals, and actions of other hostile agents. Consequently, the predictability assumption as well as the effectiveness condition of standard AI planning are violated. The competitive nature of the task precludes the possibility of simply waiting for an opportunity to arise. The overall ideal goal of approaching undetected within a particular distance and geometry of one or more hostile agents typically cannot be achieved without considerable planning.

The uncertainty of the knowledge about the characteristics and present/future actions of other hostile agents coupled with the strong dependency among the planner's subgoals precluded the adoption of the reactive planning framework discussed above. The space of possible plans under even a heuristically based strategy of closure is simply too large to be usefully pursued. This led us to develop what we refer to as a method of reactive planning using a "situation space" or in a more militaristic sense a "tactical space".

Characteristics of the Problem

The tactical ASW problem is a complicated one which will be simplified and idealized in our discussion. Consideration will be limited to the case where there is a single agent on whose planning we will focus and a single hostile agent. Each agent's overall goal is to attack and destroy the other.

There are three basic types of goals that the planner attempts to achieve or maintain. The first type are preservation goals. These include avoiding detection by other agents, avoiding collision with other agents, and avoiding attack and destruction by other agents. The second type are information goals. Knowing the presence and type of ship of other agents, their location, speed and direction are the basic information goals. Finally, there is the overall goal of destroying the other agent through attack. There are three basic actions which the planner can use to achieve these goals. These are setting the speed and direction of own ship and firing weapons. The achievement of any of these goals ultimately depends upon the relative distance between the planner's ship and the other agents' or target ships over some period of time.

There are three things to note about these goals. First, each is defined relative to the target agents. Second, the goals are not independent. The preservation goals are necessary to the normal achievement of the information goals as well as the overall goal. And third, the information goals must be true in order for the overall goal of destroying the other to be achieved. The relational character of the goals coupled with the uncertainty about the other's present or future course of action create the essential reactive characteristics of this problem. The fact that the achievement of the overall goal depends on the achievement and maintenance of other goals dictates the necessity for a planning approach rather than a purely reactive one. Finally, the planner's goal is not necessarily defined by a single fixed conjunction of these goals. Rather, differing conjunctive combinations of these goals may be pursued by the planner at particular points in the problem. Further, the planner does not have control over the conditions that give rise to these changes to the goal.

Planning begins whenever the planner believes that a hostile agent is in its area of operation. Initially the planner may not know the type of other ship, its location, speed or direction. The ability to plan in an interesting way begins when a ship has been detected. However, planning at this point does not involve an attempt to achieve the overall goal. Usually, too little is known with certainty about the other agent to support a rational approach to planning for the overall goal. In fact, the goal of current planning and execution can change considerably over the course of the interaction with another agent. The "active" goal is selected in a situationally reactive fashion. However, the achievement of the active goal is attempted not through the initiation of a precompiled reaction, but through real-time planning and execution. Further, because the success of any current plan depends upon some temporal projection of the future actions of the target, plan failure is common. However, not all failures have the same impact on the planner. The failure to achieve or maintain an information goal can lead to replanning for the achievement of this type of goal. Failure of a preservation or destruction goal changes the goal of the planner in a very significant way. Such basic observations of the characteristics of this problem motivated our attempt to identify the type of knowledge and a method for its use that could give rise to a reactive planning model that was applicable to this problem.

Proposed Model of Situationally Reactive Planning

The basic difficulty presented by this type of planning problem is how to engage in actions that are coherently related to the achievement of the overall goal despite the fact that often the planner can not develop a complete plan for achieving that overall goal. In order to overcome this difficulty we have developed a planning model within which the planner is controlled by knowledge organized into what we have termed a **situation space.**

Syntactically a situation space is specified as a collection of situations each of which includes the following information:

- situation name
- world state

- goal expression
- characterization
- transitions

The *world state* is an expression of the information that must be gleaned from the present state of the world in order for planning to ensue. The *goal expression* is a conjunctive expression describing the appropriate planning goal for this situation. Both the world situation and goal expression are partial state descriptions. *Characterization* is an indicator of relative advantage; its value is plus if in this situation the planner holds an advantage and minus otherwise. *Transitions* is a set of pairs, one for each of the possible next situations. The first member of a pair is a situation name and the second member is an expression defining the conditions for transition to the named situation. Hence, transitions define the connectivity of the situation space.

Note that this syntax allows a transition to depend on aspects of the current state that are independent of the agent's actions and thus a situation monitoring task is presupposed. In our problem domain, these expressions typically involve expressions that depend upon the success or failure of the plan as well as actions of the other agent.

Note that the situation space abstractly specifies the set of goals that might be pursued and the transition between these goals. It does not specify possible plans but only the context in which to plan. The situation space implicitly provides a functional abstraction of the kind of interaction histories that the planning agent can have in its pursuit of some overall goal. It is a useful abstraction when: (a) there are situations where achievement of the overall goal can not presently be planned; (b) there is a situationally appropriate goal that can be planned; and (c) plans can fail. The situation space provides at the meta-planning level a global view of the history and possible futures in the agent's pursuit of its overall goal. This global view allows the meta-level planner to select a situationally appropriate goal and object-level planner, and allows it to anticipate possible contingency planning based on its knowledge of the transition graph.

Planning in a Situation Space

Planning within a situation space is quite straightforward. No overall plan is created. Whenever a transition is made in the situation space the world state and goal expressions associated with the situation provide a basis for creating a situationally appropriate subplan. The goal expression associated with the situation is instantiated relative to the state of the world that holds when the situation is entered. Thus, each entry into a situation defines an island within the overall sequence of planning and execution. The state associated with the situation serves as the current starting state for the planner. Since the overall plan is broken into subplans, each plan is projected over a limited temporal (and indirectly spatial) span which serves in this problem domain to drastically limit the uncertainty about the actions of the other actor during the course of the subplan. For example, a simple default projection of the other agent's course and speed based on that agent's current course and speed can often provide an accurate basis for planning over a short time span but be quite useless over long time spans. This creation of subplans allows planning to proceed over spans where the predictability assumption may often be viable. Thus, the correctness of a plan can be evaluated, and, because of the limited scope of the plan, the real-time requirements of interleaving planning and execution can be better accommodated. Of course, plan execution must be monitored and replanning may be required. However, again, because of the limited scope of the plan, replanning time is more likely to fall within the time bounds allowed to support the necessity to initiate action.

Figure 1 presents an example of what such a situation space might look like for a simplified tactical ASW problem. This example will provide a basis for discussing the intuitive ideas concerning the role of such a situation space in this style of reactive planning.

In Figure 1 each node represents a situation and the labels on the nodes characterize the type of goal that is the major distinctive focus of plans created within that situation. Those nodes outlined in black are situations in which the planner enjoys an advantage over the other agent. In general, the planner holds the advantage if all of the preservation goals have been maintained. The nodes outlined in grey, labelled *Evade* and *Reattack* in Figure 1, represent situations where these preservation goals have failed and the planner is at a relative disadvantage. The arcs connecting situations represent possible transitions from one situation to another. The planner typically begins in the situation labelled *Search*. The transitions from *Search* to *Assess* to *Localize & Approach* correspond to the achievement of the information goals mentioned earlier. This sequence of situations leads to a transition to the *Attack* situation. The overall plan represented by this situation sequence from *Search* to *Attack* corresponds to the preferred attack plan. If the attack plan is successful, then a transition is made into the situation labelled *Success* and the interaction with the other agent is terminated. The transition to the situation labelled *Evade* represents a case where the planner has been detected or attacked. In this situation, the main goal is to simply evade destruction and re-establish the goal of not being detected and localized by the other. The node labelled *Quick Reaction Attack* is a case where the goal is to attack but in this case under the conditions where the information goals are partially achieved and the goal of remaining undetected may have been violated.

The situation space provides a way in which to structure knowledge needed to choose an appropriate goal and strategy for planning. The overall plan is divided into separate planning episodes which allows for the interleaving of planning and execution in a controlled and evaluable fashion. The situation space itself represents a kind of functional abstraction of the states that can be reached during an attempt to achieve the overall goal. Note that it includes states

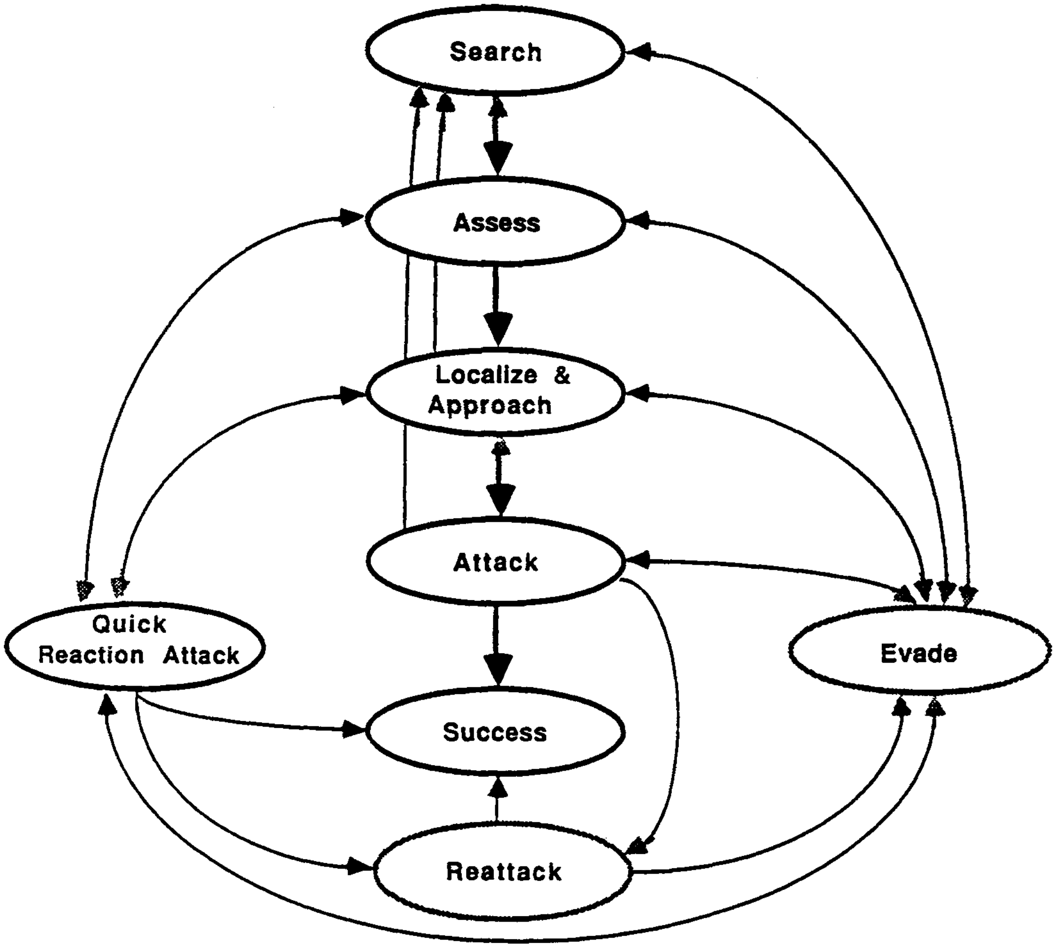

Figure 1. Situation space for a simplified tactical problem

that result from subgoal failure within a situation as well as subgoal success. The space itself provides information concerning what situations can possibly follow a current situation. A sequential record of the situations visited in this situation space can provide information concerning the history of the interaction with the other agent in the dynamic environment.

Control Architecture that exploits the Situation Space

Situation spaces are used to control and delimit the planning that occurs while a particular situation holds. Minimally, this requires that the current situation be the basis for the invocation of planning for the goal specified by that situation. Whenever a transition to a next situation occurs, planning must be shifted to the goal specified by the new situation. Thus, situation based planning also presumes a monitoring task that determines whether a transition to a new situation has occurred.

A minimal control architecture must thereby perform two tasks, situation monitoring and situation based invocation of planning. The monitoring task determines whether the current situation holds or a transition to a next situation is necessary. This involves monitoring whether one of the next situation/expression pairs, in the set of transitions, has become true in the current state of the world.

Based on the current situation and that situation's goal, a planning task must be invoked to plan and execute a solution for the goal. This architecture allows differing methods of planning to be associated with particular situations. For example, if the *Quick Reaction Attack* is particularly time critical, a reactive planning method might be useful rather than a more traditional AI planning method.

Beyond the above two minimally required tasks of monitoring and planning, a situation based control architecture can also perform several other useful tasks. In particular, the situation space is very useful for determining whether to generate a contingency plan, that is, a plan appropriate for a situation other than the current one. Such a plan is contingent on transition to a new situation and serves to prepare the planner for that situation. Contingency planning is desirable to insure the real time response of the system to changing situations. However, it is not desirable for the system to expend resources on planning for contingencies that never occur. The situation space provides a principled basis for determining whether it is possible and worthwhile to form a contingency plan. In part, this determination is based on the probability of a transition to one of the next situations. This probability is estimated using information on the current state of the world, the expression defining the conditions for transition, and the situation sequence. The decision to form a contingency plan is also based on the availability of sufficient information (known or plausibly predictable) about the hypothesized next situation. The next situation's world state expression provides the description of this information by expressing what must be known for planning to ensue. Finally, the next situation's goal determines what goal is to be planned for in the contingency plan.

A meta level control architecture (e.g., REAPPR [2]) is used to control the three tasks of planning for the current situation's goal, monitoring for situation transition, and contingency planning. Given a new current situation, the meta level invokes the planning for the goal associated with that situation. As part of this invocation, the meta level provides the object level with planning strategies (i.e., a set of planning rules) appropriate to the situation and goal, as well as any (partial) contingency plans that were developed for this situation. This planning/execution task continues unabated until interrupted by the meta level due to a transition to a new situation. The meta level also invokes the monitoring task appropriate for this situation. Any changes are reported back to the meta level. Based on these reports, the meta level updates the contingency planning task. If these reports go so far as to indicate a new situation, all three tasks are re-invoked to reflect the change in current situation.

Finally, it should be noted that this description assumes that any dependencies between the three tasks are controlled by the situation space and how the meta level architecture employs it. So for instance, the situation monitoring task monitors those conditions which signal a transition from the current situation and the planning task can change the status of those conditions. However, as long as the current situation holds, the same conditions are monitored independent of the particular object level plan being developed and the status of its execution. In part, the successful simple performance of the two tasks depends on the degree to which such independence can be maintained,

Status of implementation

We are currently investigating the implementation of this approach. Much of the work to date has centered on the planning strategies that are specialized to and appropriate for a particular situation. For example, in the *Localize & Approach* situation, we have examined "critical island" strategies and variants of traditional path planning strategies. One so called approach planner generates a space of approach paths taking into account preservation and information goals relative to a set of hostile agents. The planner takes as starting state uncertain information about the past and current position, course, speed, and type of hostile agents within an operational area. Typically, a conjunctive goal specifies a primary target to be approached and secondary targets for which information and preservation goals are to be achieved.

The planning process utilizes knowledge about the operational area, own ship resources, target behavior, and tactics. Approach paths are constructed in a symbolic abstraction space which permits a representative yet manageable number of the possible paths to be developed and critiqued. The grain or size of the solution space is controllable and can be varied to stay within the limits of planning resources. Approach plan solutions are annotated by the planning process to support comparison among alternatives and replanning during execution.

The planning process relies on simple models of target behavior within a given situation to project the course of events to a limited temporal horizon. These assumptions must be monitored to evaluate the viability of the plan chosen for execution. Of course, target behavior often deviates from these projections necessitating replanning. Currently, we are investigating the tradeoffs between local plan revision and strategies in which the current plan is abandoned and a new complete plan is generated.

Concluding Remarks

This model for situationally based control of planning and execution was motivated by two aspects of the tactical ASW problem. First, in this problem the achievement of the overall goal as well as the major subgoals depends not only on the actions of the planner, but also upon events outside the control or accurate prediction of the planner. Second, the planner can not simply wait for an opportune situation to arise. Rather, the planner must plan and act to achieve and maintain certain subgoals to create and take advantage of an opportune situation.

The model is useful for this type of problem because it allows alternative abstract subproblem sequences for the overall goal to be represented and used to control planning. The situation space provides the basis for local planning where the planning can be carried out with a locally limited version of the predictability assumption used in standard AI planning models. At the same time the situation space provides a global view of the planning and execution process upon which to base meta planning decisions such as contingency planning. Further, the movement through the situation space provides a basis for monitoring and summarizing the course of the planning effort. Such a situation sequence history can be useful in appropriately modifying the object level planning tasks.

From a purely syntactic point of view this situation planning model is quite general. The standard planning model can be embedded within it as a special case where only two situations exist in the space; namely a *Start* and a *Success* situation. This embedding is of little use. In general the usefulness of this situational planning model will depend upon the number of situations and the structure of their connectivity. Further, the model assumes that the time required to create a plan, or to replan, within a particular situation is bounded to allow planning to complete prior to the necessity to initiate action. In our example problem, the situation space exhibits a fairly high degree of connectivity, yet it is sparse enough to support meta level planning strategies which depend on the global view of the problem. If the situation space is completely connected, then the usefulness of this global view is less clear. Our hope is that a planning model of this type will prove to be useful for a variety of problems that require reactive planning because of the possibility of subplan failure and/or the inability to predictively model or control plan relevant events.

References

[1] Agre, P.E. and Chapman, D.
Pengi: An implementation of a theory of activity.
In *Proceedings of the National Conference on Artificial Intelligence*, pages 268-272.
AAAI, Seattle, WA, August, 1987.

[2] Bresina, J.L., Marsella, S.C. & Schmidt, C.F.
Predicting subproblem interactions.
Technical Report LCSR-TR-92, Laboratory for Computer Science Research, Rutgers University, February, 1987.

[3] Fikes, R.E., Hart, P.E. & Nilsson, N.J.
STRIPS: A new approach to the application of theorem proving to problem solving.
Artificial Intelligence 2:189-208, 1971.

[4] Firby, R.J.
An investigation into reactive planning in complex domains.
In *Proceedings of the National Conference on Artificial Intelligence*, pages 202-206.
AAAI, Seattle, WA, August, 1987.

[5] Hendler, J.A. and Sanborn, J.C.
A model of reaction for planning in dynamic environments.
In *Proceedings of the Knowledge-Based Planning Workshop*, pages 24-1 - 24-10.
DARPA, Austin, TX, December, 1987.

[6] Sacerdoti, E.D.
A Structure for Plans and Behavior.
Elsevier North-Holland, NY, 1977.
(Artificial Intelligence Series.).

[7] Schoppers, M.J.
Universal plans for reactive robots in unpredictable environments.
In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 1039-1042. IJCAI, Milan, Italy, August, 1987.