In the Proceedings of the 1990 AAAI Spring Symposium workshop on Planning. These proceedings also appear as University of Maryland Systems Research Center Technical Report 90-45.

Reactive Planning using a "Situation Space"

Stacy C. Marsella and Charles F. Schmidt

Department of Computer Science
Rutgers University
New Brunswick, NJ

Introduction

Consider the problem of how a single agent, the actor, can plan rationally in a world populated with multiple independent intelligent agents. Each agent can bring about changes in the actor's world which could frustrate or facilitate some course of action that the actor might entertain as part of a possible plan for achieving a goal. In the classical planning paradigm the actor is the sole source of change. This assumption, coupled with an assumption that the actor possesses a complete and correct model of the world, allows us to adopt a fairly straightforward definition of rationality for such a planning system. The system behaves rationally if it creates correct plans which are in a certain sense also minimal. A plan is correct if the goal can be proved to follow from the initial state and the plan. It satisfies a sense of minimality if no part of the plan can be removed and the proof remain valid. This is an intuitively reasonable criterion for rationality in this classical paradigm. However, research within this paradigm has led to a deeper understanding of the difficulties that stand in the way of creating an efficient computational model of planning that is guaranteed to satisfy this notion of rationality when the world is even of moderate complexity.

More recently, there has been a good deal of research in AI planning on the problem of how to extend the investigation of planning to contexts where the assumptions of the classical model are relaxed (e.g. Hendler and Sanborn, 1987; Firby, 1987; Schmidt et al., 1989). The relaxation ranges from withdrawing the assumption that guarantees success of action execution to the withdrawal of the assumption that the planner is the only source of change in the planning environment. Not surprisingly, the notion of rationality assumed in classical planning is difficult to realistically apply to most non-classical planning environments. If complete and certain knowledge is either lacking or of a complexity that precludes its effective use, then it may be the case that the success of any plan cannot be guaranteed.

In this paper we will tentatively put forward a hypothesis about a way in which to implicitly structure the space of states associated with a domain such that a weaker form of rationality is enforced when a planner uses this structure to guide its planning activity. We will term this structuring of the space of states a *situation space*.

An Example Domain - PWorld

It is useful to define an example domain within which the idea of a situation space can be motivated and the intuitions behind this structure discussed. The domain that is used for this purpose is referred to as PWorld. It was loosely defined by analogy to the kind of goals that might govern the movement of persons congregrated at a party. In this world there is a set of agents and a set of cells which are spatially structured as a grid. At any point in time each agent is located at some cell and no two agents may occupy the same cell.

The primitive actions available to each actor are of two types. Either staying in the current location or movement to a rectilinearly adjacent cell. Staying in a location can always be carried out. The rules of movement are basically those of rectilinear motion with constraints which depend upon the location of other agents. First, an agent cannot move to a cell occupied by another agent since this would violate the constraint that only one agent can occupy any cell. More interestingly, when two agents are adjacent only to one another and no other agent then neither agent can move to a cell that would cause the other agent to no longer be adjacent to some agent. Here adjacency holds between two agents if they occupy cells which are immediately rectilinearly or diagonally adjacent. Call this the *politeness constraint*. Note that if the agent is adjacent to more than one other agent, then it is possible for an agent to move to a cell that breaks this adjacency relation and possibly causes the agent not to be adjacent to any other agent. This politeness constraint is sufficient to yield interesting "group" behavior when a sufficient number of agents are "packed" into a fixed set of cells.

These minimally interesting agents can have goals of being adjacent to other agents, not being adjacent to other agents, or some conjunct of these possibilities. We assume that there is no communication about goals among the agents and that no agent's goal involves intentionally blocking or facilitating the goal achievement of another agent. It is a benign though possibly chaotic world. We further assume that the agents' goals are sufficiently diverse to make it unlikely that the true situation is a competitive one. For example, where more than eight agents had the goal of being adjacent to the same other agent.

The agents are assumed to plan and act asynchronously. Whenever, two or more agents attempt to move to the same free location at the same time, we assume a world monitor that indifferently decides which agent's action succeeds. In the case where an agent attempts to move to a cell occupied by an agent who is executing the action of staying in that cell, then the world monitor always rules in favor of the agent that is staying in the cell.

Now let us distinguish one particular agent, called the actor, and consider the problems encountered by the actor in attempting to achieve the goal of being aligned with some specific other agent, the attractor. One agent is said to be aligned with another agent if the agent occupies a cell that is immediately north, south, east, or west of the other.

Some Observations about PWorld

What are some of the aspects of PWorld that make planning particularly difficult? First, simply consider some estimates of aspects of the state space associated with a PWorld problem. Assume there are 15 agents at a party. Recall that there are five primitive actions that an agent might attempt. A reasonable estimate of the average number of these actions that an agent can actually perform at a point in time is three. With 15 agents there are roughly 3^{15} or over 14 million possible next states. Even if each agent could only choose between two actions, then the average branching factor of the state space would be 2^{15} or over 33 thousand. Consequently, reasoning through this space of possible next states to find a sequence of actions that guarantees goal achievement regardless of the actions of the other agents would be a formidable search task. The problem is particularly difficult if the plan involves many actions.

Another possibility would be to attempt to accurately model each agent's plan. However, if each agent's plan depends on modelling the planning of other agents, the combinatorics of these models of models, ..., becomes prohibitive, if not actually inconsistent. What seems to be required is some means of forming abstractions or reformulations of the primitive state space useful to the planner. An obvious candidate is to aggregate agents that are adjacent to each other into groups based on the transitive closure of the adjacency relationship. With 15 agents we might have states which involve one group of all agents, states which involve 2 groups of 1 and 14 agents, or 2 and 13, and so on, to the state where there are 15 groups, that is, each agent is isolated. From the rules of motion for individuals, it is possible to determine the rules of group motion and the possibilities of group formation and dissolution for groups of different size and spatial distribution. For example, a two agent group can undergo an identity transformation, a rigid translation in any of the four directions, a clockwise or counter-clockwise circling of one agent about the other, and so on. A three agent group has a set of group transformations unique to it such as deformation from horizontally aligned to a triangular configuration, and so on, together with the two agent group transformations possible when two agents actions are joined and the third allowed to vary independently, and so on. The advantage of such a reformulation accrues from the ability to collapse states based on the symmetries discovered. Unfortunately, our actor would have to be rather well-versed in the mathematical theory of groups to accomplish this reformulation and there is no guarantee that such a reformulation would simplify the planning for achievement of the actor's goal.

These observations about PWorld are consistent with the following conclusions. First, a classical planning formulation of this type of domain is very unlikely to be of practical use. Second, it is unlikely that a coherent formulation that allows the planner to reason from "correct" models of each agent to deterministic predictions of each agent's actions is possible. And, even if possible, it is unlikely that the use of such models is tractable. Third, even if a reformulated space describing aggregate motion could be derived from the laws of individual motion and states, the resulting state would still be highly complex and carry no guarantee that it could be usefully employed to significantly simplify the planning problem. Thus, we seem to find ourselves in the world of problems which require the paradoxical notion of "reactive planning."

A Plan-Derived Situation Space

One suggestion for dealing with unpredictable environments is to create a universal plan which "specifies appropriate reactions to every possible situation within a given domain..." (Schoppers, 1987). This suggestion is less than helpful in a domain such as PWorld. Agre and Chapman (1987) have studied the problem of selecting actions in an arcade game, Pengo, that shares some of the characteristics of our PWorld. We share their belief that a "state collapsing" mechanism to control the selection of action in domains such as these is required. Their hope is that such a mechanism emerges from an appropriate coupling or "situated" interplay between a "vision" system and concrete activity. In constrast, the hypothesis advanced here is that the traditional representations available in planning systems may provide a basis for usefully collapsing the states associated with a problem to yield a kind of "situated planner."

We refer to this basis for controlling planning as a situation space. The intent is that a situation space provide an efficient representation of an abstract partial plan that can be used to monitor the changing state of the world, provide the planner information about the appropriate goal to pursue in the "current situation", and enforce a goal-oriented coherence to a planning system that typically attempts to achieve and maintain subgoals in a local fashion.

The problem in realizing this intent is that of determining how to identify an abstract partial plan for a problem without exhaustively planning in the primitive problem space. Our proposed solution involves two basic moves. First, note that for a complex domain such as PWorld there will often be many differing plans that might yield a solution. The representation of an abstraction of this "OR-space" of plans can yield a situation space that itself would be difficult to use and will make the identification of a construction procedure more problematic. What is desired is the imposition of a simple structure on the abstract partial plan. This will result in curtailing the possible plans that a planner controlled by this space can produce. And, it can mean that a plan is not found for some specific problem despite the fact that one exists. The second move is to construct this abstract partial plan in as egocentric a fashion as possible. That is, the idea is to exploit the simplification that results by considering the affect that the world can have on the actor's ability to produce actions within a local or egocentrically defined portion of the space rather than to emphasize reasoning about the global properties of the space. An additional problem that must be faced by any constructive procedure is that of the choice of parameters for the basic objects of the space. Clearly, a situation space for a PWorld with a suitably fixed number of cells and involving 15 agents will differ and be largely irrelevant for one that involves only four agents.

The current ideas concerning a plausible method for constructing such a situation space will be presented by sketching a possible situation space for our PWorld problem. The basic strategy is inductive in form but the induction is controlled to yield partial state descriptions that collapse the space into action and goal related classes.

Figure 1 presents an example of a situation space for PWorld. The construction begins with a tentative set of situation predicates, **Free**, **Paired**, and **Surrounded**. This tentative set is based on a classification of partial state configurations that differentially constrain the actor's actions. Then we show how these set of situation predicates might be modified when fit into a situation space.

Consider first the Surrounded predicate, which returns true when there are other agents in the four cells that are vertically and horizontally adjacent to the actor's cell location. The effect of being in such a state is to make the actions of the actor conditional on those four agents. In particular, if the surrounding agents do not move, the actor cannot move. The Paired predicate is when there is exactly one other agent adjacent to the actor (i.e. in just one of the 8 vertically, horizontally, or diagonally adjacent cells) and none of the actor's immediate legal moves bring additional agents into the group. The effect is to place the actor in a two person group, whereby the *politeness constraint* restricts the actor's actions. If the other agent does not move, the actor can take at most 2 different moves and if the other agent continues not to move over some sequence of the actor's actions there will be a a cyclic structure to that action sequence. Free is true when the actor is not adjacent to any other agents or alternatively is in a group of more than 2 persons and is not surrounded. Thus the actor is either not in a group and can thus immediately take any of four different actions (walls permitting) or is capable of leaving the group and thus eventually will be able to take any of four different actions, again assuming the other agents do not move.

These three situation predicates coincide with different restrictions on the Actor's view of its ability to compose actions. This view is egocentric in the sense that it ignores the other agents' actions. Furthermore, it is shortsighted because it ignores the actor's goals. Thus the task remains to coalesce this action based view of what an actor can do with a goal based view of what the actor wants to do.

Consider now a fourth situation, **Aligned**, which is when the actor is vertically or horizontally adjacent to the Attractor. States satisfying this predicate are the desired goal states. Being Aligned does not coincide with a unique or uniform restriction on action. Aligned can restrict the Actor's actions, and can do so to different degrees depending on whether there are other members in the group. For instance, one can be aligned in a two or three person group. The distinction here is that when alignment has been achieved restriction on action is no longer a major hindrance vis-a-vis goal achievement and can even be beneficial.

Clearly, Aligned is important for goal achievement, yet it is not simply defined in terms of restrictions on actions. This is why situations cannot only express restrictions on the ability of the actor to act. A situation space must also overlay onto or otherwise distinguish situations based on some incomplete overall plan or plans to achieve the goal and how subgoals in those plans are associated with situations. So, whereas situations are in part "locally" defined in terms of the actor's capabilities they must also fit into a situation space that coherently relates a situation to the overall goal. This suggests certain modifications to the situation predicates.

For instance, the most important goal is to achieve and maintain alignment. As we have seen, the various states in which the goal is true do not uniquely or uniformly satisfy any of the three situations based on action restriction. So an Aligned situation predicate is defined and associated with the overall goal of achieving and maintaining alignment. However, in order to distinguish the Aligned situation from Surrounded, we further constrain Surrounded so that none of the other, surrounding, agents is the Attractor. Likewise, we constrain Paired so that the other agent is not the Attractor. See Figure 2 for examples that satisfy these modified situation predicates.

Now that we have the situations we must fit them into a situation space that guides the actor's planning and action execution. This involves adding additional structure. First, associated with each situation predicate is a situationally appropriate goal to pursue when the predicate is true. The goal associated with Surrounded is to achieve a state that satisfies Free. The goal associated with Paired is also Free. When in the situation Free, the goal is either "Approach Attractor and Maintain Free" or, failing that just "Maintain Free". This disjunctive goal is due to possibility of approach paths being blocked by other agents.

In addition to the situation predicates and goals, Figure 1 represents possible transitions from one situation to another as arcs connecting situations. Although not explicitly represented in Figure 1, each of these arcs has an associated predicate to be monitored that determines when the transition occurs. There are two kinds of transition arcs, solid and dotted. The solid arcs represent transitions that are associated with the successful achievement of the goal associated with the situation at the tail end of the arc. In contrast, the dotted arcs mark unexpected transitions. A path along solid arcs encodes a possible sequence of situations that represent various "expected" paths for the actor in the situation space. Any such sequence from some current situation to the Aligned situation is a decomposition into a collection of planning islands with the following characteristics:

- there is an order on the planning for the islands and on the execution of the solutions as noted by the arcs,
- an island's goal ignores future islands,
- an island's goal includes some maintenance goals.

This situation space denotes the limited rationality of the Actor. Due to the instability of PWorld it doesn't make sense to make complete plans that plan across future islands. Nevertheless the situation space itself insures that the planning for the present situation coincides with at least one formulation of an abstract decomposition structure for the goal. Roughly speaking, we can characterize this as get Free, Approach Attractor, and Align. Fortuitous transitions simply bypass part of the structure. The maintenance goals serve to guard against catastrophic transitions that shift the actor into a less advantageous situation in terms of the situation space's partial order.

Concluding Remarks

The further development and evaluation of these ideas of how to construct and use situation spaces in planning requires the pursuit of two directions of research. One is to find an appropriate formal characterization of the construction procedure in order to determine its appropriateness and generality. A second is to experimentally explore the behavior of possible situated planners that can be designed to use the situation space to control planning activity.

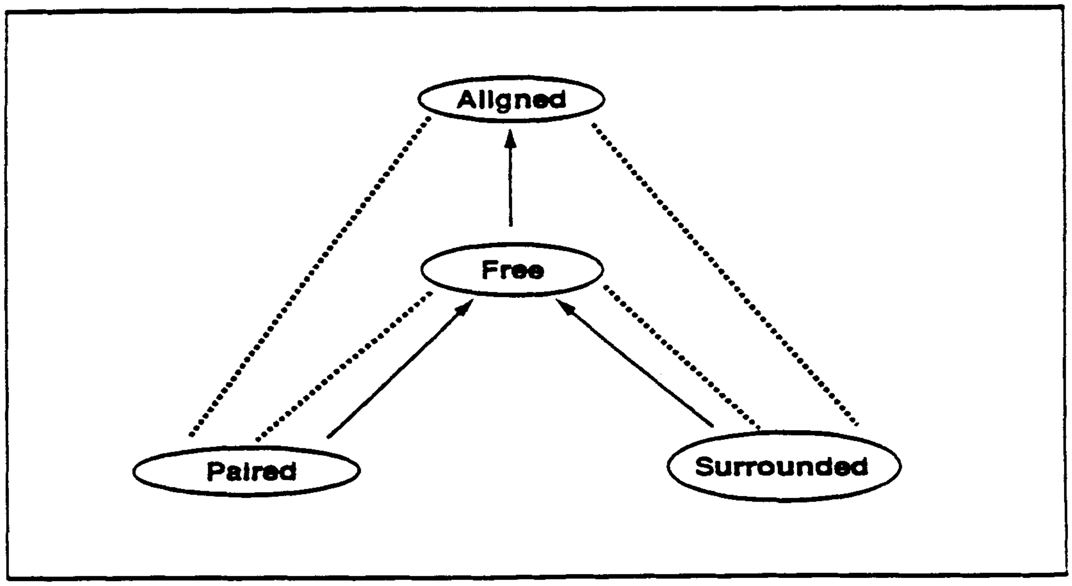

Figure 1: Example Situation Space

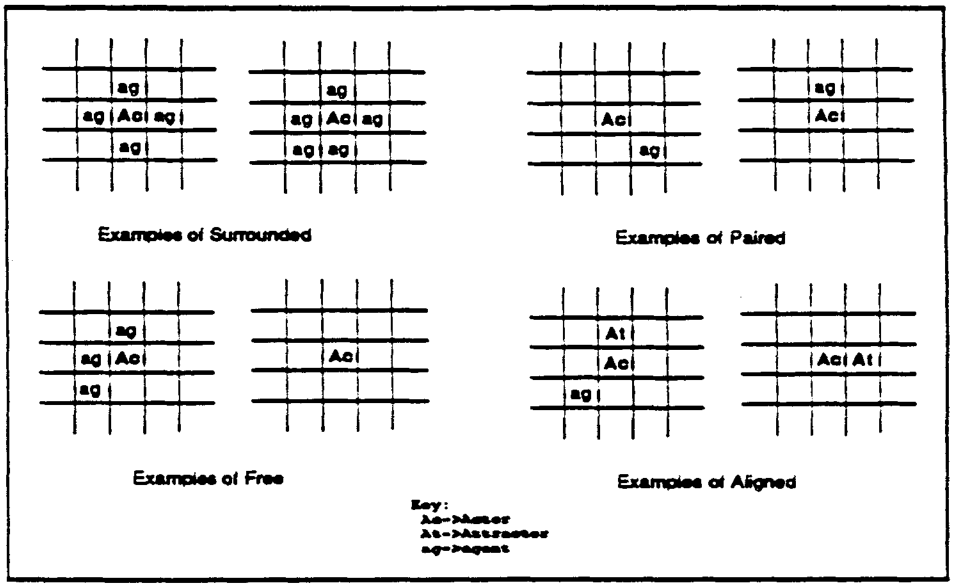

Figure 2: Example Situations

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, AAAI, 1987, 268-272.

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

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

[4] Schmidt, C.F., Goodson, J.L., Marsella, S.C., Bresina, J.L., Reactive Planning Using a Situation Space. In Proceedings of the 1989 AI Systems in Government Conference, IEEE, 1989.

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