

A Formalization of Equilibria for Multiagent Planning

Michael Bowling, Rune Jensen, and Manuela Veloso

Computer Science Department
Carnegie Mellon University
Pittsburgh PA, 15213-3891

Abstract

Planning has traditionally focused on single agent systems. Although planning domain languages have been extended to multiagent domains, solution concepts have not. Previous solution concepts either focus on planning for teams of agents with a single goal, or on a single agent in an environment containing other agents with unknown and unpredictable behavior. In reality other agents are usually acting to achieve their own explicit goals, which may not be the same or even related. In game theory the notion of an equilibria provides a framework for thinking about self-interested, utility maximizing agents. We define a formalization of the multiagent nondeterministic planning problem and introduce a notion of equilibria inspired by the game theoretic concept. As far as we know, this is the first solution framework that explicitly accounts for the various goals of *all* agents. In addition to the formalization we also demonstrate how this framework applies in a number of different domains.

Introduction

Traditionally, planning assumes a single agent for which a planner needs to find a sequence of actions that can transform some initial state into some state where a given goal statement is satisfied. But in general, “planning” can be viewed as being concerned with the general action selection problem. The planning framework has extended the classical deterministic state-action plan generation focus to many other dimensions, in particular nondeterministic actions. With the introduction of nondeterministic actions, the presence of an environment and other agents becomes a consideration. In fact, actions may have nondeterministic effects not only because of the uncertainty of their own execution, but also due to the uncertainty of the model of the environment where the actions are executed, or due to uncertainty in the actions of other agents. The possible presence of other agents as executors is viewed as the substrate of “multiagent planning.” The interest in this area has been steadily increasing and many issues remain open.

Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

We note that appropriate languages have been contributed for the representation of domains with nondeterministic actions including an explicit statement of the existence of other agents (Jensen & Veloso 2000). Other efforts offer modular architectures to represent multiagent planning domains (Wilkins & Myers 1998). However, we believe that there has not been a formal discussion of the space of multiagent plans or *solutions*. In this work, we do not focus on the problem of plan generation for multiagent planning. Instead, we focus on the interesting question of analyzing and comparing solutions for multiagent planning. Our motivation comes from making an analogy with game theory and the notion of equilibria (Owen 1995). In simple terms, in game theory an equilibrium is a *joint* solution for all the agents, such that there is no reason for any agent to change their own choice of actions given their desire to maximize some real-valued utility function.

Inspired by game theory and extending previous formal definitions of single-agent planning (Cimatti *et al.* 2000), in this paper, we introduce a formal definition of *equilibria* for multiagent planning. We begin motivating through an example that multiagent planning must address domains where agents have different goals. This is a rich space of scenarios including planning for teammates, agents with similar goals, adversaries, and the large spectrum in between. We then present a formalization of the *multiagent planning problem* that includes all multiagent goal situations. We also define the concept of *multiagent planning equilibria* as a unifying solution concept for these varied scenarios. Throughout the paper, the theoretical definitions are illustrated in a number of carefully designed examples. Finally, we discuss some questions raised by this work along with future directions and conclude.

Multiagent Plans Depend on Goals

From the beginning, plans have been contingent upon the agent’s goals. Plans are usually evaluated as to whether they achieve the goals; sometimes considering how quickly, with what probability, or from what initial states. In addition, goals are often the driving mechanism for finding good plans through Means-Ends Anal-




Figure 1: A soccer-like domain.

ysis (Newell & Simon 1963). In multiagent domains plans are of course still contingent on goals. There is an additional dependence, though. Good plans also depend on the plans of the other agents, which as we have stated, depends heavily on their goals. We will illustrate this with an example to demonstrate the importance of taking into account the other agents' goals.

Soccer Domain

Consider the simple soccer-like domain simplified from Littman's two-player grid game (1994), diagrammed in Figure 1. There are two agents A and B each of which solely occupies one of the four unshaded squares on the field. Agent A begins in possession of the ball. The shaded squares are out-of-bounds for our simplification of the domain. The agents have operators or actions associated with each of the compass directions (N, S, E, and W) or can wait, holding its position (H). The two agents in this domain select their actions simultaneously, but in execution there is an undetermined delay before these actions are carried out. So, the execution is serial but it is nondeterministic as to which agent's action is carried out first. The effect of an action is to simply move the agent in the specified direction or remain in its position as long as the target square is unoccupied. If the target is occupied then the agent does not move, but if the agent was carrying the ball it loses the ball. For agent A, losing the ball terminates execution as a failure. The goal for agent A is to move into either of the two labeled goal squares, which also terminates execution.

Figure 2 shows two examples of how the agents operators affect the state. From the initial state, if both agents choose their south operator (S,S) they will both simply move south. But if agent A selects south and agent B selects East (S,E), then there are two possible outcomes depending on their order of execution: (i) agent A moves south first, and then agent B moves east into the now unoccupied square; or (ii) agent B bumps into agent A first causing no position change, and then agent A moves south.

Possible Plans

There are a number of different possible plans for agent A to achieve its goals in this domain. The various plans depend on what actions we expect agent B to perform. We describe these plans without rigorous definitions or proofs, appealing to the reader's intuition as to what




Figure 2: Two example effects of joint actions in the soccer domain. The ordered pairs of actions represent the actions of agents A and B, respectively.

constitutes a "good" plan. The concept of a "good" plan will be rigorously formalized later in the paper.

Nondeterministic. One very simple case is if we believe agent B will select its actions nondeterministically, i.e., randomly. A sensible plan would be to hold position until the agent's actions carry it into the bottom right state. From this position regardless of the action of the other agent, a plan of two consecutive west actions is guaranteed to reach the goal. Since agent B selects actions nondeterministically it will eventually enter the bottom right state and so this plan is guaranteed to reach the goal. Other plans risk agent B's random actions causing it to move in the way, resulting in the loss of the ball and failure. Although this plan guarantees reaching the goal, it does not necessarily guarantee achievement in a finite time, as it requires waiting a possibly infinite number of cycles before agent B moves to the bottom right square.

Teammate. As we assume that most agents have goals, and often these goals are known, we go beyond the assumption that the other agent selects actions randomly. Their actions therefore are not likely to be nondeterministic but rather planned carefully to reach their own goals. Consider the case that agent B is actually agent A's teammate and therefore they have identical set of goal states. Then there is a much more efficient plan. Agent A should simply hold at the initial state while its teammate moves south out of its way. Then move west into the goal, without fear that the other agent will move in front of it. As a teammate it can be sure that the agent will comply with these assumptions since its goals are identical.¹ This plan is guaranteed to reach the goal in a finite number of steps, as opposed to the plan for the nondeterministic agent.

Adversary. Neither of these plans though have any guarantees if agent B is in fact planning to stop agent A from succeeding. In this case agent B's goal states

¹This admittedly does not address crucial issues of how this team compliance can be achieved, and may require planned communication and coordination strategies for distributed execution.

are all the states that are not goal states for agent A. If agent B simply holds its ground, then neither the nondeterministic nor teammate plan for agent A would ever reach its goal. In this situation an adversarial plan (Jensen, Veloso, & Bowling 2001) that can provide worst-case guarantees is more appropriate. One such plan is to nondeterministically select between holding and moving north or south until the other agent is not in front of it. Then, move west into the goal hoping its action gets executed first. This plan has no guarantee of success as the opponent may still move in front of it while it advances toward the goal causing the ball to be lost. It does, though, have some possibility of success. In fact, against a good plan by the opponent this is all that can be guaranteed.

Overlapping Goals. A whole new situation arises if we believe that agent B is not quite an opponent but not quite a teammate. Suppose its goal is to have agent A score, but only across the south square. In this case moving south from the initial state and then west would reach the goal without its interference. This plan, like the teammate plan, is guaranteed to succeed in a finite number of steps. Notice that the teammate-based plan and the nondeterministic-based plan would both fail in this case as both agents would hold indefinitely.

These four plans all are completely different, despite the fact that the conditions that generated the plans, the domain rules, and the agent’s goal did not change from situation to situation. This demonstrates that multiagent planning solutions need to take into account the goals of the other agent. Under certain specific circumstances, such as nondeterministic agents or adversarial agents, plans may and have been rigorously specified. On the other hand, there is no unifying framework for defining these sorts of planning solutions. In the next section, we define precisely a solution concept for multiagent planning that accounts for the other agents’ goals.

Planning Equilibria

In this section, we formalize the concept of a multiagent planning equilibrium. In order to help make these concepts clear we will first describe an example simpler than the soccer domain. We will use this example to make our formal concepts more concrete, and will return to the soccer domain later.

A Simple Example — The Narrow Doorway

Consider a two agent robot domain where both agents are in a hallway and want to move into the same room through a single doorway. The agents have an operator to go through the door (G) that only succeeds if the other agent is not also trying to go through the door. They also have the choice of waiting (W). Each agent’s goal is simply to be in the room.

The Formalization

We first begin by formalizing some planning related concepts. The definitions parallel closely with Cimatti and colleagues’ single-agent formalization (2000). We extend their definitions of planning domains, problems, and solutions to encompass multiple agents. We then follow this formal framework with a definition of multiagent planning equilibrium. Notice that the definitions and concepts presented are not bound to any particular planning algorithm or language.

We start by defining a multiagent planning domain.

Definition 1 (Multiagent Planning Domain)

A multiagent planning domain D is a tuple $\langle \mathcal{P}, \mathcal{S}, n, \mathcal{A}_{i=1\dots n}, \mathcal{R} \rangle$ where,

- \mathcal{P} is the finite set of propositions,
- $\mathcal{S} \subseteq 2^{\mathcal{P}}$ is the set of valid states,
- n is the number of agents,
- \mathcal{A}_i is agent i ’s finite set of actions, and
- $\mathcal{R} \subseteq \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ is a nondeterministic transition relation where $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ and must satisfy the following condition. If $\langle s, a, s' \rangle \in \mathcal{R}$ and $\langle s, b, s'' \rangle \in \mathcal{R}$ then, $\forall i$ there exists $s''' \in \mathcal{S}$,

$$\langle s, \langle a_1, \dots, a_{i-1}, b_i, a_{i+1}, \dots, a_n \rangle, s''' \rangle \in \mathcal{R}.$$

I.e., each agent’s set of actions that can be executed from a state are independent.

In addition, let $\text{ACT}_i(s) \subseteq \mathcal{A}_i$ be the set of applicable or executable actions in state s . Formally,

$$\text{ACT}_i(s) = \{a_i \in \mathcal{A}_i \mid \exists \langle s, \langle \dots, a_i, \dots \rangle, \cdot \rangle \in \mathcal{R}\}.$$

The additional condition in the planning domain definition on \mathcal{R} requires that each agent be capable of selecting actions independently. Formally this amounts to the following. For all states s and executable actions for the agents $a_i \in \text{ACT}_i(s)$ there exists some transition $\langle s, \langle a_{i=1\dots n} \rangle, s' \rangle$ that is in \mathcal{R} .

In the doorway domain, \mathcal{P} contains two propositions, *A-in-room* and *B-in-room*. The set of states \mathcal{S} corresponds to all four possible subsets of \mathcal{P} , since all combinations of propositions are valid in this domain. n is two and $\mathcal{A}_{A,B}$ is the set of actions $\{G, W\}$. The transition relation \mathcal{R} is defined by the rules described above. The complete enumeration of states and transitions is shown in Figure 3. The figure also numbers the states so they can be referred to in an abbreviated form. Note that this domain satisfies the independent action condition on \mathcal{R} .

Definition 2 (Multiagent Planning Problem)

Let $\mathcal{D} = \langle \mathcal{P}, \mathcal{S}, n, \mathcal{A}_{i=1\dots n}, \mathcal{R} \rangle$ be a multiagent planning domain. A multiagent planning problem P for \mathcal{D} is a tuple $\langle \mathcal{D}, \mathcal{I}, \mathcal{G}_{i=1\dots n} \rangle$, where $\mathcal{I} \subseteq \mathcal{S}$ is the set of possible initial states and $\mathcal{G}_i \subseteq \mathcal{S}$ is the set of goal states for agent i .

In the doorway example, the goal states for agent A are $\{1, 3\}$ and for agent B are $\{2, 3\}$. The initial state set is the singular set $\{0\}$. With this definition of domain and problem, we can now formalize a notion of a plan.

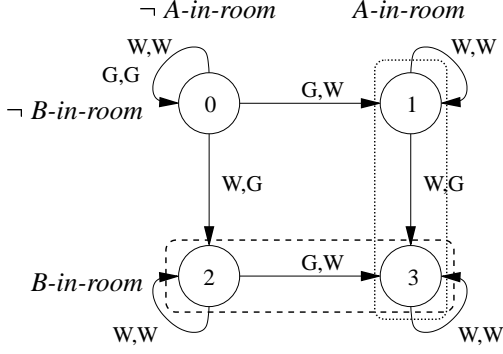


Figure 3: Doorway domain. \circ and \square represent A's and B's goal states, \mathcal{G}_A , and \mathcal{G}_B , respectively.

Definition 3 (State-Action Table)

A state-action table π_i for agent i in domain \mathcal{D} is a set of pairs $\{\langle s, a_i \rangle \mid s \in \mathcal{S}, a_i \in \text{ACT}_i(s)\}$. A joint state-action table π constructed from state-action tables for each agent $\pi_{i=1..n}$ is the set of pairs

$$\{\langle s, \langle a_1, \dots, a_n \rangle \rangle \mid s \in \mathcal{S}, \langle s, a_i \rangle \in \pi_i\}$$

A (joint) state-action table is complete if and only if for any $s \in \mathcal{S}$ there exists some pair $\langle s, \cdot \rangle$ in the state-action table.

For the doorway domain, a state-action table (or plan) for each agent might be,

$$\begin{aligned} \pi_A &= \{\langle 0, G \rangle, \langle 1, W \rangle, \langle 2, G \rangle, \langle 2, W \rangle, \langle 3, W \rangle\}, \\ \pi_B &= \{\langle 0, G \rangle, \langle 0, W \rangle, \langle 1, G \rangle, \langle 2, W \rangle, \langle 3, W \rangle\}. \end{aligned}$$

These are also *complete* state-action tables since they specify at least one action for each state. We can combine these tables into a complete *joint* state-action table. In general, a joint state-action table together with a multiagent planning domain determines the entire execution of the system. In order to define what it means for a plan to be a solution to a planning problem we need to formalize the notion of reachability and paths of execution. We will do this by first defining the execution structure of the multiagent system.

Definition 4 (Induced Execution Structure)

Let π be a joint state-action table of a multiagent planning domain $\mathcal{D} = \langle \mathcal{P}, \mathcal{S}, n, \mathcal{A}_i, \mathcal{R} \rangle$. The execution structure induced by π from the set of initial states $\mathcal{I} \subseteq \mathcal{S}$ is a tuple $K = \langle Q, T \rangle$ with $Q \subseteq \mathcal{S}$ and $T \subseteq \mathcal{S} \times \mathcal{S}$ inductively defined as follows:

- if $s \in \mathcal{I}$, then $s \in Q$, and
- if $s \in Q$ and there exists a state-action pair $\langle s, a \rangle \in \pi$ and transition $\langle s, a, s' \rangle \in \mathcal{R}$, then $s' \in Q$ and $\langle s, s' \rangle \in T$.

A state $s \in Q$ is a terminal state of K if and only if there is no $s' \in Q$ such that $\langle s, s' \rangle \in T$.

Intuitively, Q is the set of states that the system could reach during execution of the plan π , and T is the set

of transitions that the system could cross during execution. For our doorway domain the execution structure induced by our example joint state-action table is,

$$\begin{aligned} Q &= \{0, 1, 3\}, \\ T &= \{\langle 0, 1 \rangle, \langle 0, 0 \rangle, \langle 1, 3 \rangle, \langle 1, 1 \rangle, \langle 3, 3 \rangle\}. \end{aligned}$$

We can now formalize an execution path.

Definition 5 (Execution Path)

Let $K = \langle Q, T \rangle$ be the execution structure induced by a state-action table π from \mathcal{I} . An execution path of K from $s_0 \in \mathcal{I}$ is a possibly infinite sequence s_0, s_1, s_2, \dots of states in Q such that, for all states s_i in the sequence:

- either s_i is the last state of the sequence, in which case s_i is a terminal state of K , or
- $\langle s_i, s_{i+1} \rangle \in T$.

A state s' is reachable from a state s if and only if there is an execution path with $s_0 = s$ and $s_i = s'$.

For our doorway domain and example joint state-action table one execution path from the initial state is,

$$0, 0, 0, 0, 1, 1, \dots$$

We are now in a position to define what it means for our plan to solve a planning problem. We actually define multiple concepts increasing in strength. These concepts formalize some of the intuitive discussion from the previous section about whether a plan has one or more of the following properties:

- the possibility of reaching the goal,
- a guarantee of reaching the goal, and
- a guarantee of reaching the goal in a finite number of steps.

These concepts and their formalization are inspired and highly related to Cimatti and colleagues' single-agent solution concepts (Cimatti *et al.* 2000).

Definition 6 (Multiagent Planning Solutions)

Let \mathcal{D} be a multiagent planning domain and $P = \langle \mathcal{D}, \mathcal{I}, \mathcal{G}_{i=1..n} \rangle$ be a multiagent planning problem. Let π be a complete joint state-action table for \mathcal{D} . Let $K = \langle Q, T \rangle$ be the execution structure induced by π from \mathcal{I} . The following is an ordered list of solution concepts increasing in strength.

1. π is a weak solution for agent i if and only if for any state in \mathcal{I} some state in \mathcal{G}_i is reachable.
2. π is a strong cyclic solution for agent i if and only if from any state in Q some state in \mathcal{G}_i is reachable.
3. π is a strong solution for agent i if and only if all execution paths, including infinite length paths, from a state in Q contain a state in \mathcal{G}_i .
4. π is a perfect solution for agent i if and only if for all execution paths s_0, s_1, s_2, \dots from a state in Q there exists some $n \geq 0$ such that $\forall i \geq n, s_i \in \mathcal{G}_i$.

A state-action table's strength $\text{STRENGTH}(\mathcal{D}, \mathcal{P}, i, \pi)$ is the largest number whose condition above applies for agent i . If no conditions apply then $\text{STRENGTH}(\mathcal{D}, \mathcal{P}, i, \pi) = 0$.

For our doorway domain, the joint state-action table is a strong cyclic solution for both agents but not strong (i.e., it has a strength of 2 for both agents). This means that there is a path to the goal from any reachable state. But there are also paths that do not include either agents' goal states, and so it is not a strong solution for either agent.

The plans from the soccer domain can also be described under this solution framework. The plan that handles the nondeterministic agent B is a strong cyclic solution since a goal state is always reachable but there are infinite execution paths where agent A does not reach the goal (e.g., if agent B holds indefinitely). For the teammate case, the plan is a perfect solution since it is guaranteed to reach the goal in three steps and remain there. The same is true for the situation where agent B's goal is to have the ball scored in the southern square. In the adversarial case, the plan is only weak since some execution paths result in losing the ball and failing.

These solutions define what it means for one agent to be successful given a joint state-action table. The goal of planning from one agent's perspective is to find a plan that has the highest strength given the plans of the other agents. But the other agents' selection of a plan is equally contingent upon the first agent's plan. This recursive dependency leads to our main contribution of the paper: planning equilibria.

Definition 7 (Multiagent Planning Equilibria)

Let \mathcal{D} be a multiagent planning domain and $\mathcal{P} = \langle \mathcal{D}, \mathcal{I}, \mathcal{G}_{i=1..n} \rangle$ be a multiagent planning problem. Let π be a complete joint state-action table for \mathcal{D} . Let $K = \langle \mathcal{Q}, \mathcal{T} \rangle$ be the execution structure induced by π from \mathcal{I} . π is an equilibrium solution to \mathcal{P} if and only if for all agents i and for any complete joint state-action table π' such that $\pi'_{j \neq i} = \pi_j$,

$$\text{STRENGTH}(\mathcal{D}, \mathcal{P}, i, \pi) \geq \text{STRENGTH}(\mathcal{D}, \mathcal{P}, i, \pi').$$

I.e., each agent's state-action table attains the strongest solution concept possible given the state-action tables of the other agents.

Note that our example joint state-action table for the doorway domain is *not* an equilibrium. Both agents A and B currently have strength 2, but B can achieve a strength of 4 by choosing a different state-action table. Specifically, B should select the wait (W) action from the initial state and the go (G) action in state 1.

Examples

To make the concept of planning equilibria clearer, we will examine it in a number of illustrative domains. We return to the doorway domain. We then consider a domain representation of Rock-Paper-Scissors, and finally we reexamine the various plans in the soccer domain.

Doorway Domain

We gave above an example joint state-action table that is not a multiagent planning equilibria for this domain. An equilibria is the following state-action tables:

$$\begin{aligned} \pi_A &= \{ \langle 0, G \rangle, \langle 1, W \rangle, \langle 2, G \rangle, \langle 3, W \rangle \}, \\ \pi_B &= \{ \langle 0, W \rangle, \langle 1, G \rangle, \langle 2, W \rangle, \langle 3, W \rangle \}. \end{aligned}$$

In this case agent A goes through the door while agent B waits and then follows through the door. This is a perfect plan for both agents and so obviously no agent can achieve a higher strength with a different state-action table. Similarly, the symmetric tables where agent B goes through the door while agent A waits is also an equilibrium. There is an additional equilibrium,

$$\begin{aligned} \pi_A &= \{ \langle 0, G \rangle, \langle 0, W \rangle, \langle 1, W \rangle, \langle 2, G \rangle, \langle 3, W \rangle \}, \\ \pi_B &= \{ \langle 0, G \rangle, \langle 0, W \rangle, \langle 1, G \rangle, \langle 2, W \rangle, \langle 3, W \rangle \}. \end{aligned}$$

Here both agents nondeterministically decide between going through the door and waiting. This results in a strong cyclic solution for both agents, but given this state-action table for the other agent no strong or perfect plan exists for either agent. So this is also an equilibrium although obviously inferior to the other equilibria where both agents have higher strength plans. In game theory, such a joint strategy is called Pareto dominated.

Collision variation. Consider a variation on this domain where collisions (when both agents choose G) result in the robots becoming damaged and unable to move. In this case, the first two state-action tables above remain equilibria, but the third inferior table no longer is an equilibrium. The joint plan is now only a weak solution for both agents since there is a possibility of never achieving the goal. Each agent can also change to a different plan where it waits for the other agent to get through the door thus achieving a strong cyclic plan and a higher strength.

Door closing variation. Finally, consider that one agent entering the room sometimes causes the door to close behind it. Once the door is closed it cannot be opened and the doorway cannot be used. In this case, the same two joint plans are again an equilibrium but now they have different strengths for the different agents. The first joint state-action table is a strong plan for agent A, but only a weak plan for agent B, though it can do no better. The second is just a symmetry of this.

Rock-Paper-Scissors

Consider a planning domain representation of the children's game Rock-Paper-Scissors. Each agent simultaneously chooses one of rock (R), paper (P), or scissors (S). The winner is determined by a cyclic rule: rock loses to paper, paper loses to scissors, scissors loses to rock. Figure 4 gives the enumeration of states, transitions, and goals for this planning problem. In this case, there is a unique planning equilibrium where each agent's state-action table contains every action. This joint plan is a weak solution (strength 1) for both agents




Figure 4: Rock-Paper-Scissors as multiagent planning.

and neither agent can switch to a different plan and get a higher strength. This plan is analogous to the game's game theoretic equilibrium which randomizes evenly between all three actions (Fudenberg & Levine 1999).

Soccer Domain

Let us reconsider the soccer-like domain. We presented three distinct planning problems where agent A's goals remained constant, but agent B's goals varied from having identical goals to A, opposing goals to A, and a subset of A's goals. The example plans described for these situations, if we add in the implied plan for agent B, are all equilibria to their respective multiagent planning problems. In the teammate case and the overlapping goal case, the equilibrium is a perfect solution for both agents. So, obviously, no agent can switch plans to improve on this solution. In the adversarial case, it is a weak solution for both agents, and neither agent can improve on this strength. This formalization of the planning equilibrium matches well with our intuitive notions of "good" plans in multiagent domains.

Discussion

Multiagent planning equilibria is a powerful concept both to understand the multiagent planning problem and as a viable solution that accounts for the other agents' goals. It also opens up many avenues for further research and understanding. We consider a couple important questions this work raises.

The first issue is the number of planning equilibria. The doorway domain illustrates that multiple equilibria may exist. Although some equilibria are obviously inferior to others, the equilibria framework does not select a single planning solution. For example, the two symmetric equilibria in the doorway domain are not equivalent, nor is one Pareto dominant. This calls for coordination or communication mechanisms to decide between competing equilibria.

Although some problems have many equilibria, other planning problems may have none. We believe that we have devised a (rather contrived) example of a domain with no equilibria (publication in preparation). Large classes of domains can be proven to have an equilibrium (e.g., team domains and adversarial domains.) Still, interesting questions are whether equilibria exist in most useful domains, or what are reasonable plans when equilibria do not exist.

Finally, this work presents a compelling framework and solution concept for multiagent planning, and gives

the challenge of devising methods to find a planning equilibrium. The equilibrium definition involves universal quantification over an agent's possible plans, which is exponential in the number of states, which in turn is exponential in the number of propositions. This is of course intractable for anything but domains with a handful of propositions. This opens up a new realm of interesting issues relating to efficiently finding equilibria under different planning frameworks, languages, or classes of domains. There are techniques for finding plans for teams of agents with identical goals as well as adversarial situations with competing goals (Jensen, Veloso, & Bowling 2001). These techniques, although limited to certain multiagent planning problems, provide evidence that multiagent planning equilibria can be both a theoretically and practically powerful concept.

Conclusion

We presented a formalization of multiagent planning and introduced the concept of a multiagent planning equilibrium. This is the first known solution concept that explicitly accounts for the goals of all the agents. This work provides a unifying framework for considering planning in multiagent domains with identical, competing, or overlapping goals. It also opens up many exciting questions related to practical algorithms for finding equilibria, the existence of equilibria, and the coordination of equilibria selection.

References

- Cimatti, A.; Pistore, M.; Roveri, M.; ; and Traverso, P. 2000. Weak, strong, and strong cyclic planning via symbolic model checking. IRST Technical report 0104-11, Istituto Trentino di Cultura.
- Fudenberg, D., and Levine, D. K. 1999. *The Theory of Learning in Games*. The MIT Press.
- Jensen, R., and Veloso, M. 2000. OBDD-based universal planning for synchronized agents in non-deterministic domains. *Journal of Artificial Intelligence Research* 13:189–226.
- Jensen, R. M.; Veloso, M. M.; and Bowling, M. H. 2001. OBDD-based optimistic and strong cyclic adversarial planning. In *Proceedings of the Sixth European Conference on Planning*.
- Littman, M. L. 1994. Markov games as a framework for multi-agent reinforcement learning. 157–163. Morgan Kaufman.
- Newell, A., and Simon, H. A. 1963. GPS, a program that simulates human thought. In Feigenbaum, E. A., and Feldman, J., eds., *Computers and Thought*. New York: McGraw-Hill. 279–296.
- Owen, G. 1995. *Game Theory*. Academic Press.
- Wilkins, D. E., and Myers, K. L. 1998. A multiagent planning architecture. In *Proceedings of the Fourth International Conference on Artificial Intelligence Planning Systems*, 154–162.