Negotiation in State-Oriented Domains with Incomplete Information over Goals

Shlomit Bergman and **Elan Pavlov** and **Jeffrey S. Rosenschein**[1]

Abstract.

State Oriented Domains (SODs) are domains where agents are concerned with moving the world from an initial state into one of a set of target states. Negotiation in this environment was explored by Rosenschein and Zlotkin [9], who provided an analysis of incentive compatible mechanisms over a variety of two-agent, single-encounter types. Their model included the concept of an agent's worth (the agent's benefit from achieving its goal), using it as a baseline for utility calculation of a negotiation's outcome. One scenario left unexamined, however, was the case where agents know one another's worths, but not one another's goals. This situation creates the possibility of agents' lying to one another solely about goals, to influence the outcome of a negotiation.

In this paper, we explore this specific case of known worths and unknown goals in two-agent State Oriented Domains, in a variety of encounter types. Through analysis and examples, it is shown that an agent can benefit from declaring less costly goals, but that there are certain limits to the lies an agent can beneficially declare. We also analyze the connection of this work to classic game theory results, including general work on incentive compatible mechanisms and the revelation principle.

1 INTRODUCTION AND BACKGROUND

Negotiation in Multiagent Systems (MAS) is often used as a technique for agents to agree on a specific goal, on a plan of action, or on an allocation of tasks among themselves. Models of negotiation in MAS have been built for specific types of domains, one of the most popular being so-called "State Oriented Domains" (SODs) [9]. In State Oriented Domains, agents are concerned with moving the world from an initial state into one of a set of target states. The set of target states is called the agent's *goal*. This formalization implies that goals cannot be partially achieved; a state is either in the goal set, or it is not.

In order to reach one of the goal states, an agent carries out a sequence of actions (a *plan*), that terminates in a goal state. This state-oriented description of an agent's environment thus reflects the classical approach of artificial intelligence systems [10], and specifically the state-space search of artificial intelligence planning [4].

In a multiagent SOD, the goals of individual agents may or may not overlap. These multiple agents, to avoid harmful interference and leverage helpful synergy, may agree on a *deal* that specifies each agent's plan (the actions that each one should execute), bringing the world into a state that is in the intersection of their goal sets.

An agent's plan to achieve a goal has a *cost*, reflecting (in some sense) the amount of work to be done by the plan. Assume that each agent in a multiagent SOD has some upper bound on the cost that he is willing to pay to reach a goal state. We call this upper bound the *worth* of the agent's goal (following [9]).

The *utility* for an agent of the plan that a deal prescribes for him will be assumed to be the difference between the worth of his goal and the cost of his role in that plan. We further assume that an agent will only attempt to reach a goal using a specific plan if its utility is greater than 0 (negative utilities imply that an agent loses more than it gains, while 0 utility implies indifference between accepting or rejecting the deal).

The definition of State Oriented Domains permits agent actions to have side effects: agents can achieve one another's goals (even when such achievement is not the reason actions are being carried out), or inhibit one another from achieving their goals. However, there can also be a conflict between agents' goals, and there is no guarantee that the goal states of agents overlap. The possibility of agents reaching a deal (i.e., an assignment of plans) that is mutually acceptable depends both on the existence of a goal state that satisfies them both, as well as the allocation of effort, which needs to provide them both with positive utility.[2]

1.1 Incomplete Information

An environment in which agents have complete information about other agents' worth and goal set is relatively straightforward. Protocols may be implemented such that both agents will agree on a deal that results in equal utility for participants, with the resulting maximization of the product of agent utilities. This solution has certain attributes that an environment designer might find desirable to achieve.

What about situations of incomplete information? Given our interaction model, an agent might have incomplete information about other agents' worth and/or goals. There are three possible cases:

Worth	*Goals*
unknown	known
unknown	unknown
known	**unknown**

In order to cope with the lack of information (again, following [9]), we add a "-1-phase" where agents simultaneously declare private information before the beginning of the negotiation. The negotiation then proceeds as if the revealed information were true.

Rosenschein and Zlotkin analyzed the two cases where worth is not known. In this paper, we analyze a situation that was not dealt with in the original treatment, namely the case where worth is known

[1] School of Engineering and Computer Science, The Hebrew University of Jerusalem, Jerusalem, Israel; email: {bshlomit, elan, jeff}@cs.huji.ac.il

[2] In [9] probabilistic deals are also explored, where even if no mutually acceptable goal exists, the agents might agree on a deal that leaves them both with positive *expected* utility—though only one reaches its goal.

and goals are not known. The general conclusion of Rosenschein and Zlotkin was that, in situations of partial information, a strategic player can benefit by pretending that its worth is lower than it actually is, either by declaring lowered worth or by declaring a cheaper goal. In our case, when the worth is known, an agent cannot lower its worth, but as we will see, it can get more utility by declaring a cheaper goal.

2 DEFINITIONS

We adopt the definitions of [9].

Definition 1 (State Oriented Domain) *A State Oriented Domain (SOD) is a tuple* $< \mathcal{S}, \mathcal{A}, \mathcal{J}, c >$ *where:*

1. $\mathcal{S}$ *is the set of all possible world states.*
2. $\mathcal{A} = \{A_1, A_2, ...A_n\}$ *is an ordered list of agents.*
3. $\mathcal{J}$ *is the set of all possible joint plans. A joint plan* $J \in \mathcal{J}$ *moves the world from one state in* $\mathcal{S}$ *to another. The actions taken by agent* k *are called* k*'s* role *in* J*, and will be written as* J_k*. We can also write* J *as* $(J_1, ..., J_n)$*.*
4. c *is a function* $c : \mathcal{J} \rightarrow (\mathbb{R}^+)^n$*. For each joint plan* J *in* $\mathcal{J}$*,* $c(J)$ *is a vector of* n *positive real numbers, the cost of each agent's role in the joint plan. If an agent plays no role in* J*, his cost is 0.*

Definition 2 (Encounter) *An* Encounter *within an SOD* $\langle \mathcal{S}, \mathcal{A}, \mathcal{J}, c\rangle$ *is a tuple* $\langle s, (G_1, ..., G_n)\rangle$ *such that* $s \in \mathcal{S}$ *is the initial state of the world, and for all* $k \in \{1...n\}$*,* G_k *is the set of all acceptable final world states from* $\mathcal{S}$ *for agent* A_k*.* G_k *will also be called* A_k*'s* goal*.*

Definition 3 (Deals) *Given an encounter in a two-agent SOD* $\langle s, (G_1, G_2)\rangle$*:*

- *We define a* Pure Deal *to be a joint plan* $J \in \mathcal{J}$ *that moves the world from state* s *to a state in* $G_1 \cap G_2$*.*
- *We define a* Deal *to be a mixed joint plan* $(J : p)$*;* $0 \leq p \leq 1 \in R$ *such that* J *is a Pure Deal.*

The semantics of a Deal are that the agents will perform the joint plan (J_1, J_2) with probability p, or (J_2, J_1) with probability $1 - p$.

Definition 4 (Cost) *If* $\delta = (J : p)$ *is a Deal then* $Cost_i(\delta)$ *is defined to be* $pc(J)_i + (1 - p)c(J)_k$ *(where* k *is* i*'s opponent).* $c(J)_i$ *is the* i*-th element of the cost vector, i.e., it is the cost of the* i*-th role in* J*.*

Definition 5 (Worth of a goal) *The* worth *of goal* G_i *to agent* i*, denoted* w_i*, is defined as the maximum expected cost that agent* i *is willing to pay in order to achieve his goal* G_i*.*

Definition 6 *The* utility *of a deal* δ *for agent* A_i *is defined as* $u_i(\delta) = w_i - Cost_i(\delta)$*.*

Definition 7 *A deal* δ *is* individually rational *if, for all* i*,* $u_i(\delta) \geq 0$

Definition 8 *A Deal* δ *is* Pareto optimal *if there does not exist another Deal that dominates it—there does not exist another Deal that is better for one of the agents and not worse for the other.*

Definition 9 *The* negotiation set *NS is the set of all the deals that are both individual rational and pareto optimal.*

We also introduce several additional definitions, standard in the game theory literature, that were not originally used in [9].

Definition 10 (Budget-Balance) *A negotiation is said to be budget balanced when the total payments received during the process* from *agents is at least as great as the total payments made during the process* to *agents [2, 7].*

The concept of budget balance assumes that payments to and from agents may be part of the negotiation mechanism.

Definition 11 (Efficiency) *The negotiation is efficient when it maximizes the total (sum) of utilities of all agents.*

We are interested in goals with *(weakly) dominant strategies*. The agents will choose their best possible strategy, and this strategy will be independent of what the other agents do. Furthermore, all of our strategies will be *incentive compatible* (IC) [2, 3], meaning that the agents will tell the truth during the negotiation process.

3 IMPOSSIBILITY AND HARDNESS RESULTS

In order to show the impossibility of solving the negotiation problem (i.e., finding a suitable deal), as formalized above, we show its reduction to the multicast problem. The multicast problem is a particular case of a bilateral trading mechanism, in which several parties attempt to bargain to achieve an agreed-upon goal. A well known result [6] states that it is impossible to achieve budget balance (BB) and individual rationality (IR) while maximizing utility in these types of problems. We will relate to this impossibility result below in Section 5, where we design a mechanism for our scenario that exhibits some of these attributes.

In fact, even if we are willing to forgo budget balance and merely find the most efficient IC mechanism, the solution to the problem is still NP-hard. This is due to a result [5] which states that the multicast problem is NP-hard to solve, and another result [1] which states that it is NP-hard even to approximate. We first define the multicast problem, and then show a simple reduction from the multicast problem to the problem of negotiation, which demonstrates that it is impossible to achieve BB and IR while maximizing welfare. This means that we have to compromise on some of these properties.

Multicast transmission is defined as follows (we utilize the same notation as [1, 5]). There is a user population P residing at a set of nodes N that are connected by links L. There is a privileged node α that will be called the source. We assume that there is a *universal* tree $T(P)$. Given any set of receivers $R \subseteq P$, the multicast tree $T(R)$ is the minimal tree $T(R) \subseteq T(P) \cap R$. Each node has a valuation v_i if it is in R.

We now show the reduction between the SOD negotiation scenario and the multicast problem. Let us assume for a contradiction that there is a mechanism for the negotiation scenario that is IR, BB, and maximizes welfare. We then show that there is such a mechanism for the multicast problem. Given any multicast problem we define a world in which there are P blocks (each corresponding to a different player) that are in P different spaces and such that there is an empty space to the right of each block. We define a set of P agents who negotiate: agent i wishes that the blocks corresponding to a path to node i in the multicast problem be moved to the right. The worth for agent i of any state of the world that has these blocks moved is the valuation in the multicast problem.

Note that any solution to this blocks world problem results in a solution to the multicast problem, where R consists of blocks that have been moved. By the definition of the states of the world, this set R is connected along links in the multicast problem.

Since the multicast problem has no general solution that is IR, BB and maximizes welfare, the blocks world similarly (a classic State Oriented Domain) has no such general solution.

To summarize, by [6] it is impossible to find an IR+BB welfare-maximizing mechanism for this problem (even in the case in which there are no conflicts) no matter how much computational power we invest. By [5] it is hard to find the most efficient allocation with IC agents in polynomial time, and by [1] it is hard to even approximate the most efficient solution in polynomial time. On the other hand, it is possible to achieve some of the desired properties. For example, IR plus BB can be achieved by the empty set. IR plus maximized welfare can be achieved by the complete set (in the case where there are no conflicts). BB plus maximizing welfare can be achieved by charging one player for all of the cost incurred by the protocol.

4 INCOMPLETE INFORMATION ABOUT GOALS

In this section, we consider agents negotiating with incomplete information about goals, but knowing one another's worth, in a variety of encounter types: cooperative encounters using mixed deals, conflict encounters using semi-cooperative deals (where there is a common subgoal), and conflict encounters using multi-plan deals (where there is no common subgoal).

We will use examples from the slotted blocks domain, a State Oriented Domain, to illustrate agent interactions. In the slotted blocks domain, there is a table with a bounded number of slots. The agent operations are PickUp and PutDown; each operation has a cost of one. An agent can only hold one block at a time. We assume that the worth of both agents is publically known, but their goals are private information.

Agents are assumed not to agree on a deal that is not individual rational, which means they will not agree on a deal where either one's utility is negative. We tighten this assumption a little, for the purposes of our analysis: we assume that if $u_i(\delta) = 0$, an agent will prefer the conflict deal, rather than δ, even though the two options are the same for him. The rationale for this is that the agent with utility 0 would like to be able to proclaim to the other agent, "If I do not get more than 0 from our deal, there will be no deal."

4.1 Cooperative—Mixed Deals

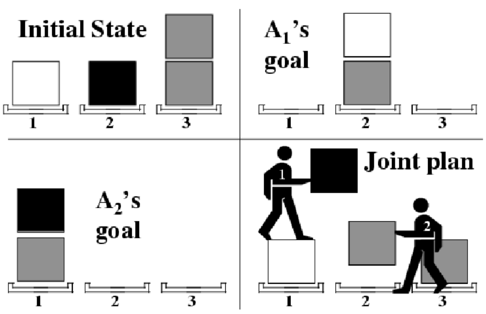

Figure 1. Cooperative Situation

Suppose that the initial state is as in Figure 1. Agent A_1's goal is that the white block is on a gray block at slot 2; he does not care where the other blocks are. Agent A_2's goal is that the black block is on a gray block at slot 1, also not caring where the other blocks are. As stated above, the goals are private information, but worth is public. Assume that w_1=12, and w_2=12. If each agent is alone in the world, it will cost him 8 to reach his goal alone, which will result in utility of 4 for agent A_1 $(12 - 8)$ and 4 for A_2 $(12 - 8)$. There is no contradiction between the two goals since there exists a state that satisfies both agents, and there is a joint plan to reach this state with a total cost of 8.

The optimal plan has two roles, one requiring six operations and one requiring two operations. One agent will lift the black (or white) block, while the other rearranges all the other blocks appropriately. The first agent will then put down the black (or white) block. When they both tell the truth, they will agree on $J : \frac{1}{2}$ where $J = (J_1, J_2)$ (which cost 2 and 6, respectively) giving each utility of 8.

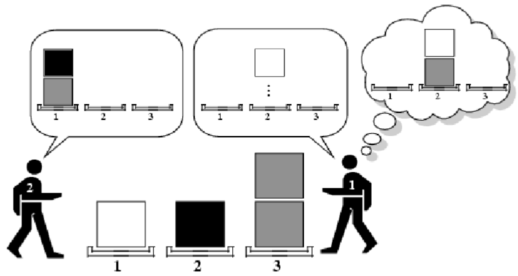

Figure 2. Lying in a Cooperative Situation

Suppose that A_1 decides to lie and instead of declaring his actual goal, says that his goal is for the white block to be at slot 2 but not on the table, as we can see in Figure 2. He would then not agree to carry out a plan of cost greater than two, and the agents will agree on $J = (J_1, J_2)$ as a joint plan ($p = 1$). This means that A_1 will always assume the cheaper role, giving A_1 utility of 10, and A_2 utility of 6. So by declaring a cheaper goal, A_1 gets more utility; A_1 is able to reduce the apparent cost for him of achieving his goal alone, while not compromising on the achievement of his real goal. This happens because there is only one state that satisfies both A_2's goal and A_1's lie, and this state also satisfies A_1's real goal.

4.2 Semi-Cooperative Deals

Beneficial lies can also exist in conflict situations, that is, when there does not exist a state that satisfies both agents' goals. In these situations, there can be a beneficial lie if there is a common subgoal for the agents.

Consider, for example, the situation depicted in Figure 3. Agent A_1 wants the black block to be at slot 1, whereas Agent A_2 wants it to be at slot 3. In addition, both want to swap the blocks at slot 4. The cost for each agent to achieve his goal alone is 10. An agent can move the gray block to slot 3, then the white block to slot 2, move the gray back to slot 4, then put the white on the gray, and finally move the black to slot 1 (or slot 3, depending on the agent).

Assume that $w_1 = w_2 = 12$. If both agents tell the truth, they will agree to do the swap cooperatively (which costs 2 for each agent, instead of 8), and then flip a fair coin to decide who will continue with his goal. This brings them an expected utility of 3: each agent will do 2 operations for swapping the blocks at slot 4, and with probability $\frac{1}{2}$ will achieve his goal (worth 12) with a cost of 2 more ($u_1 = u_2 = \frac{1}{2}(12-2) - 2 = 3$).

What if agent A_1 lies and states that his goal is: “The black block is at slot 1 and the white block is on the gray block”? The cost for A_1 to achieve his apparent goal is 6 because he can build the reversed stack at slot 3 with cost of 4. With this lie, assuming that A_2 states his true goal in the -1-phase, when they do not reach an agreement,[3] A_1 would get alone an apparent utility of $3 = \frac{1}{2}(12-6)$. A_2 would get alone $1 = \frac{1}{2}(12-10)$.

If A_1 does not lie about his goal, the expected utilities of A_1 and A_2 are equal. However, if A_1 lies about his goal, it appears that A_1’s expected utility is greater than A_2’s. In this situation, A_1 can demand to use a weighted coin with the claim that “my expected utility is greater than yours”.

Even in this situation, it is better for them to cooperate, and “divide” the saved utility. However, the question arises as to how they should divide the additional utility. Let us take their utility without agreement, (3,1), i.e., (A_1,A_2), and keep the relation of their utilities, 3:1, which leads to:

$$p\,(12-4) + (1-p)(-2) = 3[(1-p)(12-4) + p(-2)]$$

meaning that the utility of A_1 will be three times as much as the utility of A_2. Both agents will do the swap together, which will cost 2 for each. With probability p, A_1 will win the coin toss and pay 2 more to achieve his goal, so we get $p\,(12-4) + (1-p)(-2)$. For A_2, the calculation is the same except that we switch $(1-p)$ and p, and in order to keep the relation of 3:1 we multiply it by 3.

¿From the equation above we get $p = \frac{13}{20}$, so A_1 will get $4\frac{1}{2}$ and A_2 will get $1\frac{1}{2}$. With his lie, A_1 got more utility. This happened because the cheapest way for the agents to cooperate still brought the world to A_1’s real goal state.

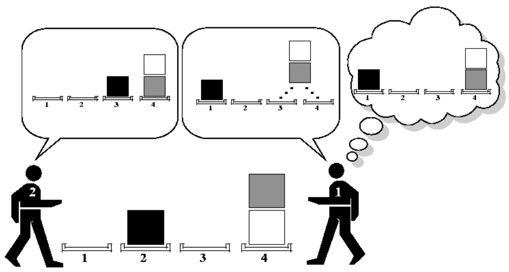

Figure 3. Common Subgoal

4.3 Multi-Plan Deals

There can also be beneficial lies in conflict situations with no common subgoal. Consider the initial interaction situation seen in Figure 4. The goal of A_1 is to reverse the blocks in slot 2, and to leave the blocks in slot 1 in their initial position. The goal of A_2 is to reverse the blocks in slot 1, and to leave the blocks in slot 2 in their initial position. The worth for each agent is known: $w_1 = w_2 = 10$. If an agent is alone in the world he can achieve his goal with a cost of at least 8.

The goals here contradict one another. There is no final state that satisfies both agents, and in addition there is no common subgoal that allows them to (partially) cooperate, as we had in the previous section. But they can still flip a coin and *then* cooperate: they flip a coin and then they work together to achieve the goal of the agent who won the coin toss. To achieve any goal together costs a total of 4 for the two agents—2 for each.[4]

If both agents tell the truth, they will agree on $(\delta_1, \delta_2) : \frac{1}{2}$, where δ_i brings the world to i’s goal state. The utility for each agent would be $3 = \frac{1}{2}(10-2) + \frac{1}{2}(-2)$, with probability $\frac{1}{2}$ that an agent gets his goal satisfied, and he “pays” 2, yielding $\frac{1}{2}(10-2)$. With probability $\frac{1}{2}$ his goal is not satisfied but he must pay 2 in order to achieve his opponent’s goal, yielding $\frac{1}{2}(-2)$.

Suppose agent A_1 lies and claims that his goal is to reverse the blocks in slot 2 and leave the blocks in slot 1 *OR* to have the white block in slot 2. To achieve his new goal alone would cost him 6 instead of 8. If they run into conflict, they will toss a coin and the winner will do his work alone, giving A_1 an apparent utility of 2 ($2 = \frac{1}{2}(10-6)$), and A_2 utility of 1 ($1 = \frac{1}{2}(10-8)$). When they cooperate they can do the work together with cost of 4 to achieve each goal. They can still toss a coin and then work together to achieve the winner’s goal.

So they will agree that A_1 will get double the utility of A_2, and we get: $p(10-2) - (1-p)(2) = 2[(1-p)(10-2) - p(2)]$ which leads to $p = \frac{3}{5}$. When cooperating, they will bring the world to A_1’s real goal state because it is cheaper to accomplish this together than it is to bring the world to his false goal state. This gives A_1 utility of 4, and A_2 utility of 2, so A_1 benefited from his lie.

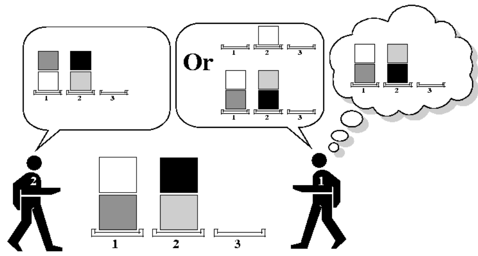

Figure 4. Conflict Situation

4.4 Analysis—Goal Fabrication, Worth is Known

As shown above, an agent can benefit by declaring a cheaper goal. Furthermore, assuming true goals are internal to the agent and hidden from the outside world, there is only a single, focused risk involved in lying. Even if *both* agents lie, they can still reach agreement *as long as the intersection of their declared goals is not empty and is included in their real goals*. Basically, the more an agent reduces the apparent cost of his goal, the more his expected gain increases. Note that the choice of the false goal is dependent on the other agent’s goal.

Consider, for example, the situation in Figure 5. Assume that $w_1 = w_2 = 12$. Agent A_1 lies and says that his goal is for the white block to be at slot 2 but not on the table. Agent A_2 lies and says that his goal is for the black block to be at slot 1 but not on the table. The combination of these two lies leads to a situation in which both agents stand to gain the same amount as in the original case, where neither agent lied.

[3] We assume here that if agents fail to reach an agreement, they flip a fair coin for who gets to do his work alone towards his own goal, and the agents *do not* leave the world in the initial state.

[4] We of course assume that agents can enter into a binding contract to cooperate, even if they lose the coin toss—or that there are sufficient disincentives to defection that an agent is motivated to cooperate even when he loses the toss.

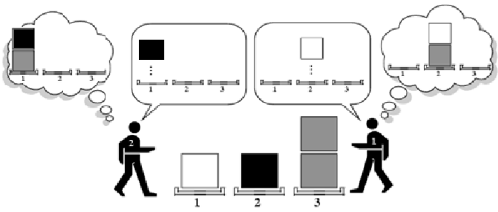

Figure 5. Two Agents Lying in a Cooperative Situation

They will agree on equal work for each, giving a utility of 8 to each agent. Although it is the case here that when both agents lie, they get the same utility as when both tell the truth, it is still "advisable" for the agents to lie. For example, agent A_1 cannot predict whether agent A_2 will lie, and if A_1 tells the truth and A_2 lies, A_1 stands to lose. Since when A_1 lies he does not risk his gain, and he might increase his gain, one may conclude that it is advisable for A_1 to lie. Clearly A_2 should also lie for symmetric reasons. The Nash equilibrium for the decision whether to lie or not is for both agents to lie.

If both agents declare a cheaper goal, they do not risk reaching an agreement, because according to our assumption worth is known, and of course higher than the new goal's cost. The incentive to lie without any risk is absolute — as long as the intersection of the agents' declared goals is not empty and is included in their real goals.

If we compare this to the two other worth/goal cases from [9] where worth is not known, we can see that when the goal is known, the strategy is to declare lower worth, but if both lower their worth too much they may not reach an agreement (depending on the strict/tolerant mechanisms explored in that research). When both goal and worth are unknown, the dominant strategy is to declare a cheaper goal, and worth is calculated from the stand-alone cost. Because agents will agree on a deal that maximizes the product of their utilities, by declaring a cheaper goal an agent can confidently reduce the amount of work he will do in the final deal.

5 DESIGNING A MECHANISM

We would like to develop an incentive compatible mechanism for our given negotiation scenario; such a mechanism would mean that agents will gain nothing by lying. As noted in Section 3, it is easy to get an incentive compatible mechanism if we have no additional demands. In principle, we would like to have a welfare maximizing, IR, and BB mechanism, but unfortunately this is impossible. So in order to proceed, we relax our demands of the mechanism and require one that maximizes welfare and is IR; such a mechanism will have "truth telling" as a dominant strategy (in other words, agents will disclose their true goals to the mechanism). For such a mechanism to find a solution, it has to subsidize the agents such that the sum total budget of the mechanism is negative (the mechanism "loses money" to guarantee incentive compatibility).

For example, we can take the situation depicted in Figure 5. The best lie that an agent can tell is as depicted in the figure; for agent A_1, the best lie is "The white block is at slot 2 but not on the table" (similarly, for agent A_2). Now the solution is that each agent will pay only 2. One agent will lift the white block (with a cost of 2), the other agent will arrange all the other blocks (with a cost of 6). The mechanism will pay 4 to the one who did the role which cost 6, so we finish with an efficient and individual rational solution, and the mecanism subsidized it (by 4). Here we got individual rationality and efficiency, but not a budget balanced solution. In this situation, no agent can gain any benefit from lying, no matter what the other agents will do.

The first thing to note is that the mechanism has to subsidize all of the common subgoals. This is simply because each player wishes to impose the cost of this subgoal on another player. Secondly, there are some lies that are hard to avoid, since it is possible for a player to fabricate a goal that has a large subgoal with another player, but is close to his real goal. Therefore, the agents will pay the cost from their *best lie* to their goal.

This can be calculated by doing a breadth first search (BFS) from the initial state. The depth of the BFS will be bounded by the sum total of worths. For every possible goal for the agent i (holding all other goals to be the declared goal), an optimal path can be calculated. Agent i will then pay the minimum from his goal to the path that is closest to it. It is easy to see that this is IR, since the price that player i pays does not depend on the valuation of player i, and since the price is monotonic.

6 CONCLUSION

We have explored the case of known worths and unknown goals in two-agent State Oriented Domains, in cooperative encounters using mixed deals, conflict encounters using semi-cooperative deals (where there is a common subgoal), and conflict encounters using multi-plan deals (where there is no common subgoal). It was shown that agents stand to benefit by lying so as to reduce the apparent cost of their goals, even more than in other incomplete information scenarios; they do not risk reaching an agreement, as long as the intersection of their declared goals is not empty and is included in their real goals.

Along with the above analysis, we presented a reduction of our negotiation scenario to the multicast problem, establishing the relevance of impossibility and hardness results over our model. Although we cannot have a general mechanism that is budget balanced and individual rational while maximizing utility, we did present an incentive compatible mechanism that makes use of mechanism and agent payments to maximize welfare and remain individual rational, while forgoing budget balance.

REFERENCES

[1] A. Archer, J. Feigenbaum, A. Krishnamurthy, R. Sami, and S. Shenker, 'Approximation and collusion in multicast cost sharing', *Games and Economic Behavior*, (to appear).

[2] Kenneth Binmore, *Fun and Games: A Text on Game Theory*, Houghton Mifflin Company, 1992.

[3] R. K. Dash, N. R. Jennings, and D. C. Parkes, 'Computational mechanism design: A call to arms', *IEEE Intelligent Systems*, **18**(6), 40–47, (2003).

[4] T. L. Dean and M. P. Wellman, *Planning and Control*, Morgan Kaufmann, 1991.

[5] J. Feigenbaum, C. Papadimitriou, and S. Shenker, 'Sharing the cost of multicast transmissions', *Journal of Computer and System Sciences*, **63**, 21–41, (2001).

[6] R. B. Myerson and M. A. Satterthwaite, 'Efficient mechanisms for bilateral trading', *Journal of Economic Theory*, **29**, 265–281, (1983).

[7] D. C. Parkes, J. Kalgnanam, and M. Eso, 'Achieving budget-balance with VCG-based payment schemes in combinatorial exchanges', Technical Report Research Report RC 22218, IBM, (March 2002).

[8] J. S. Rosenschein and G. Zlotkin, 'Cooperation and conflict resolution via negotiation among autonomous agents in noncooperative domains', *IEEE Transactions on Systems, Man, and Cybernetics*, **21**(6), 1317–1324, (1991).

[9] J. S. Rosenschein and G. Zlotkin, *Rules of Encounter*, MIT Press, 1994.

[10] Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall, 2003.