A game-theoretic approach to cooperation in Multi-Agent Systems *

Matei Popovici
POLITEHNICA University of Bucharest
Splaiul Independentei nr. 313, sector 6
Bucharest, Romania
matei.popovici@cs.pub.ro

Ciprian Dobre
POLITEHNICA University of Bucharest
Splaiul Independentei nr. 313, sector 6
Bucharest, Romania
ciprian.dobre@cs.pub.ro

ABSTRACT

The development of semantic agent frameworks is strongly influenced by the means for representing and reasoning about knowledge. Nevertheless, when faced with decision-making in multi-agent environments, current approaches are limited in their ability to model strategic interactions. As a result, ontological approaches such as those based on OWL, offer only a reactive behavior for such interactions. We argue that a game-theoretic approach is more suitable when faced with modeling cooperation in Multi-Agent Systems. We introduce a new type of games, Boolean Games with Currency, which combine features of Boolean Games, with the transferable payoff setting. We also give a computational characterisation for a solution concept for coalitional stability: the core. We show that the core of a Boolean Game with Currency is always non-empty, and we prove the core membership problem to be co-NP complete.

Categories and Subject Descriptors

H.1 [**Models and principles**]: General; F.2 [**Analysis of algorithms and problem complexity**]: General—*complexity measures*

General Terms

Theory

1. INTRODUCTION

The success of a Multi-Agent System (MAS) is strongly affected by the way decision-making is approached. In traditional semantic systems, agents are treated as inanimate entities that interact with the environment in a strictly reactive way, by responding to external stimuli with particular actions, in order to achieve some goals. In frameworks that support agent cooperation such as [1], the emphasis falls on the knowledge representation means, and on developing a suitable agent platform. Semantic learning methods such as [8] or [7], focus on achieving a certain more-or-less fixed behavior, using machine learning techniques that do not require an in-depth understanding of the underlying processes. All these approaches ignore the fact that, in many situations, the decision to act in a certain way is influenced by the decisions of other agents. To illustrate this, consider the following example: agents 1, 2 and 3 are travelling together, and must decide on the transportation means. The possible options are (a) airplane, (b) bus, and (t) train. The decision is made by a simple voting procedure. If there is a tie, the fastest form of transportation wins. Agents have the following preferences (we use $x \succ_i y$ to denote that agent i prefers x to y): $a \succ_1 b \succ_1 t$, $b \succ_2 a \succ_2 t$, $t \succ_3 b \succ_3 a$. If each agent votes for his most preferred candidate, then the winner would be a, since there is a tie and it is assumed that airplanes are faster than buses or trains. This outcome is highly undesirable for agent 3, who least prefers travelling by air. Nevertheless, if 3 would observe the preferences of 1 and 2, then, he could change his vote from t to b, and therefore change the outcome of the voting procedure from a to b. Even if b is not 3's most desirable travelling means, for him it is better than a. In this latter case, 3 behaves in a *strategic* way: his decision is based both on his goal, as well as on the others' decisions.

By *strategic interaction*, we refer to any situation where: the state of an agent is dependent both on the environment as well as on the states of other agents, each agent is aware of this fact and exploits it to his own advantage. In the following, we are interested in Multi-Agent Systems where any interactions are done in accordance with these properties. We also make the following assumptions on the characteristics of agents: (i) they are rational, (ii) self-interested and (iii) have limited resources.

Assumption (i) refers to agents that are fully aware of the actions available for them as well as their consequences and have well-defined preferences over all possible states-of-the world. By (ii) we understand that each agent has a certain preferred state-of-the world, and would take any possible action to achieve this state, regardless of the interests of other agents. Assumption (iii) states that any agent can only use some finite computational effort, in order to derive the proper actions he should undertake.

Game Theory (GT) is a powerful tool in the design of any kind of strategic interaction between agents in a MAS. It uses the following elements: a set of *actions* available for each agent, a set of *consequences* that result from performing some actions, a *consequence function* that, for each action, associates a proper consequence, and finally, for each agent, a preference relation over the set of consequences. GT distinguishes between two modeling approaches: non-cooperative, and cooperative. In the non-cooperative setting the emphasis falls on the actions available to the *individual*, as well as his preferences. The cooperative setting, still preservers the preferences of individuals but the emphasis now falls on what agents can *jointly* achieve.

In the following we introduce a new type of game, suitable for representing and reasoning about compromise. To our knowledge, there are no actual applications of game theory in semantic agent frameworks, although game theory is briefly mentioned in papers such as [5] or [4]. The rest of the paper is structured as follows: In Section 2 we introduce a new type of game, Boolean Games with Currency, which combines the features of Cooperative Boolean Games, with the transferable payoff setting. In Section 3 we give some complexity bounds for the computation of the core and finally, in Section 4 we make some final remarks and sketch future directions concerning our approach.

2. GAMES WITH CURRENCY

We propose a new type of cooperative games inspired from [2], where agents are able to specify their goals as propositional formulae. The difference between our approach and Cooperative Boolean Games (CBG) [2] is that we consider situations in which agents are able to exchange payoff. As a result, our setting combines features of CBG with that of cooperative games with transferable payoff [6].

Let $\mathcal{V} = \{p, q, r, \ldots\}$ be a finite set of propositional variables, and $\mathcal{L}_\mathcal{V}$ be the standard propositional logic consisting of formulas built using variables from $\mathcal{V}$, the negation operator $\neg$ and the connectives $\wedge$ and $\vee$. In a *Cooperative Boolean Game*, each agent i controls a subset of variables $\mathcal{V}_i \in \mathcal{V}$. $\mathcal{V}_1, \mathcal{V}_2, \ldots, \mathcal{V}_n$ is a proper partition of $\mathcal{V}$ such that no two agents can share control over a variable. The possible actions available to agents consist of setting truth values to the variables they control. Each action has a certain cost. The agent i's objective is expressed as a boolean formula $\gamma_i \in \mathcal{L}_\mathcal{V}$.

DEFINITION 1. *A boolean game with currency (BGC) is a tuple* $\mathfrak{B} = \langle N, \mathcal{V}, (\mathcal{V}_i)_{i\in N}, (\gamma_i)_{i\in N}, \mu, c\rangle$ *where:*

- $N = \{1, 2, \ldots, n\}$ *is the set of agents;*
- $\mathcal{V}$ *is a finite set of propositional variables;*
- $\mathcal{V}_1, \mathcal{V}_2, \ldots, \mathcal{V}_n$ *is a partition of* $\mathcal{V}$*. Each* $\mathcal{V}_i$ *denotes the set of variables controlled by agent* i*;*
- $\mu : N \to \mathbb{R}_{\geq 0}$*:* μ *assigns a value to an agent's goal formula; in the following, we use* μ_i *as a shorthand for the value* $\mu(i)$.
- $\gamma_1, \gamma_2, \ldots, \gamma_n \in \mathcal{L}_\mathcal{V}$ *are the goals of each agent;*
- $c : N \times \mathcal{V} \times \mathbb{B} \to \mathbb{R}_{\geq 0}$ *is a function assigning for each agent, variable and truth value, a cost of setting that particular truth value. We use* $c_i(w, t)$ *as a shorthand for* $c(i, w, t)$*;*

EXAMPLE 1 (BGC). *Consider the following scenario: Jim, Tom and Harry are colleagues and, occasionally, they go out together. They can choose between a restaurant and a pub. Jim would like to eat at the restaurant, but he is looking forward to spend more time with Harry. Harry is only interested to go to the pub and have some beers, and he is not interested in company. Like Jim, Tom is also indifferent about the location, and doesn't want to go out alone. But Jim finds Tom to be a really boring person, and prefers the company of Harry. This scenario is modeled by a BGC having:*

- $N = \{1, 2, 3\}$ *- the agents, Jim, Harry and Tom*
- $\mathcal{V} = \{p_1, p_2, p_3, r_1, r_2, r_3\}$ *and the following partition: for all* $i \in \{1, 2, 3\}$ $\mathcal{V}_i = \{p_i, r_i\}$ *model the agent* i*'s possible choices:* p_i *and* r_i *models agent* i*'s choice: going to the pub, or restaurant, respectively.*
- $c = 0$ *is a constant cost function;*
- $\mu_1 = 6$, $\mu_2 = \mu_3 = 2$*;*
- $\gamma_1 = (p_1 \wedge p_2 \wedge \neg p_3) \vee (r_1 \wedge r_2 \wedge \neg r_3)$*;*
- $\gamma_2 = p_2 \wedge \neg r_2$*;*
- $\gamma_3 = (p_3 \wedge p_2) \vee (r_3 \wedge r_2)$*;*

We now look at the possible *outcomes* of BGC and, given a certain outcome, we define a measure of fulfillment for each agent, i.e. a *utility*. For an arbitrary set of variables $V \in 2^\mathcal{V}$, let $\xi = (\xi_w)_{w\in V}$ be a V-valuation giving truth assignments for variables in V. If w is a variable in V, then either $\xi_w = \top$ (true) or $\xi_w = \bot$ (false). For all agents i, we define a function testing the entailment of i's goal:

$$1_i(\xi) = \begin{cases} 1 & \text{if } \xi \models \gamma_i \\ 0 & \text{otherwise} \end{cases}$$

The utility of an agent can be defined in the spirit of [2]. If ξ is a V-valuation, then:

$$u_i(\xi) = 1_i(\xi) * \mu_i - \sum_{w\in V\cap\mathcal{V}_i} c_i(w, \xi_w)$$

Notice that, as opposed to CBG [2], in our setting it might be the case that the utility of an agent is not positive when his goal is satisfied. Such a condition could be achieved only by adding restrictions to the value and cost functions μ and c. Therefore, it is not always the case that an agent prefers an outcome where his goal is satisfied.

PROPOSITION 1. *Cooperative Boolean Games are a special case of Boolean Games with Currency.*

PROOF. Consider a game $\mathfrak{B} = \langle N, \mathcal{V}, (\mathcal{V}_i)_{i\in N}, (\gamma_i)_{i\in N}, \mu, c\rangle$, with the following restrictions on μ and c:

- $\forall w \in \mathcal{V}, c_i(w, \bot) = 0$; only setting variables to true inflicts a cost on agents;
- $\forall w \in \mathcal{V}, c_i(w, \top) = c_j(w, \top)$; the costs for setting a variable to true are the same for all agents;
- $\mu_i > \sum_{w\in\mathcal{V}} c_i(w, \top)$; the value of each agent's goal is strictly larger than the sum of all the costs for variables he controls.

The last restriction ensures positive values for utilities u_i, if agents satisfy their goals. Then u_i can induce a preference relationship over valuations ξ such that: (i) between a valuation that does satisfy his goal and another that does not, a agent will always choose the former, and (ii) between two valuations that either both satisfy or do not satisfy his goal, a agent will chose the one minimising costs. □

3. THE CORE OF A BGC

In Example 1, it is obvious that goal satisfaction is dependent on a certain deal of cooperation between agents. In the following we formally define how cooperation between agents occurs.

For a game $\mathfrak{B} = \langle N, \mathcal{V}, (\mathcal{V}_i)_{i\in N}, (\gamma_i)_{i\in N}, \mu, c\rangle$, and a coalition $S \subseteq N$, the *endowment* of each agent $i \in S$ is:

$$1_i(\xi)\mu_i - \sum_{w\in\mathcal{V}\cap\mathcal{V}_i} c_i(w, \xi_w) \qquad (1)$$

The endowment is computed by taking the value of the satisfied goal, if it is indeed satisfied, and subtracting the involved costs. Based on the endowment, the characteristic function v is defined as follows:

$$v(S) = \max_{\xi=(\xi_i)_{i\in V}} \sum_{i\in S} \left\{ 1_i(\xi)\mu_i - \sum_{w\in V\cap\mathcal{V}_i} c_i(w, \xi_w) \right\} \qquad (2)$$

Equation 2 describes the worth of a coalition S as being the maximum sum of all endowments, obtained by some valuation ξ_i.

The **cf form** $\langle N, v\rangle$ of a BGC $\mathfrak{B}$ is obtained by taking N to be the number of agents of $\mathfrak{B}$, and by computing the characteristic function v according to Equation 2. The **core** of a BGC expressed in cf form $\langle N, v\rangle$ is the set of all feasible payoff profiles $(x_i)_{i\in N}$, for which there is no coalition S such that $v(S) > x(S)$. We say that $(x_i)_{i\in N}$ is group rational: no other coalition S has an incentive to deviate.

EXAMPLE 2 (CORE). *Going back to Example 1, it is straightforward that there is no valuation able to satisfy all goals. A stable valuation is $\xi = \{p_2, p_3\}$, which leaves agent Jim unsatisfied, but here he cannot object, since there is no coalition and ξ' that can guarantee $\neg p_3$ for him. But, since Jim's goal satisfaction is measured by a particular real value μ_1, and since this value can be transferred between agents, then he could change the outcome of the game, by* tipping off *Tom and thus determining him not to come to the bar. This would be possible under the following assumption: Jim's happiness when going with Harry is much larger than Tom's, and therefore he can afford to* tip off *Tom, and make him earn more by staying at home (or going to the restaurant), than by coming to the pub to see Harry.*

This gives us the intuition that, in settings such as ours, agents might attempt to maximise their utilities by means other than satisfying goals. One such mean is by participating in coalitions where some agents have goals with great values. Agents contribute to these goals instead of their own, and, in certain settings, can achieve more value. This is the case for Tom, who could participate in forming the grand coalition by fulfilling Jim's goal instead of his own, under the condition that he can obtain a higher value this way.

PROPOSITION 2 (CORE MEMBERSHIP). *The core membership problem for BGC is co-NP complete.*

PROOF (SKETCH). Core membership is a decision problem $\mathrm{MEM}(\mathfrak{B}, x)$ which, given a Boolean Game with Currency $\mathfrak{B}$ and a vector x of size N, asks whether x is in the core of $\mathfrak{B}$. Recall the definition of core membership, which requests that there is no coalition $S \subseteq N$ and payoff profile y such that S prefers y over x. Now, consider the complement of core membership problem $\overline{\mathrm{MEM}}(\mathfrak{B}, x)$ which asks if there exists a coalition S, and a payoff y such that $y(S) > x(S)$. An equivalent definition for this problem, is asking whether there exists a coalition $S \subseteq N$ such that $v(S) > x(S)$. If such an S exists, then there is also an S-feasible payoff profile y, which agents in S will prefer over x.

In the following, we prove $\overline{\mathrm{MEM}}(\mathfrak{B}, x)$ to be NP-complete.

The membership $\overline{\mathrm{MEM}}(\mathfrak{B}, x) \in NP$ is straightforward: A procedure can build all coalitions S in nondeterministic polynomial time, and, for each S, checking whether $v(S) > x(S)$ holds can be done in deterministic polynomial time. The NP-hardness of $\overline{\mathrm{MEM}}(\mathfrak{B}, x)$ is due a reduction from the k-VERTEX-COVER(G) problem.

Let $G = (A, E)$ be a graph. k-VERTEX-COVER(G) asks if there is a subset $B \in A$ of vertices, with $|B| = k$, such that all edges from E are covered by at least one vertex. Starting from an instance of k-VERTEX-COVER(G), we build an instance of $\overline{\mathrm{MEM}}(\mathfrak{B}, x)$ in the following way:

- for each vertex $a \in A$, we create a agent $n_a \in N$. We add an additional agent n_o to N;
- for each edge $e = (a, b) \in E$, we create two variables $p_a, p_b \in \mathcal{V}$, and subsequently assign their control to the corresponding agents: $p_a \in \mathcal{V}_{n_a}$ and $p_b \in \mathcal{V}_{n_b}$ (each agent controls one "side" of an edge). For the additional agent, we have $\mathcal{V}_{n_o} = \emptyset$;
- we define the cost function c such that $\forall i$, $\forall p_e \in \mathcal{V}$, $c_i(p_e, \top) = 1$ and $c_i(p_e, \bot) = 0$;
- the goal of each agent n_a is $\gamma_{n_a} = \bigwedge_{e=(a,b)} (p_a \vee p_b)$ (n_a's goal is satisfied when at least one "side" of each of his

incident edges is set to $\top$). The goal for agent n_0 is $\gamma_{n_0} = \bigwedge_{e=(a,b)\in E} (p_a \vee p_b)$ (it is satisfied when all edges have at least one side set to $\top$);

- the goal value for each agent n_a is $\mu_{n_a} = 0$. The goal for agent n_0 is $\mu_{n_0} = 2|E|$;
- we build a payoff vector x such that $x_{n_o} = |E| - 1/2$ and $\forall n_a, x_{n_a} = \frac{1}{2(k+1/2)}$

We only show the first part of the double implication. The second part follows a similar argument. Suppose W is a vertex cover of size k. Then, the coalition $S_W = \{n_0\} \cup \{n_a | a \in W\}$ will achieve the maximum value under any valuation $\xi = (\xi_i)_{i\in S_W}$ such that, $\forall e = (a,b)$ $\xi_{p_a} = 1$ and $\xi_{p_b} = 0$ (exactly one "side" of each edge is set to true). Under any such valuations, all agents see their goals satisfied. Therefore:

$$v(S_W) = \sum_{i\in S_W} \mu_i - \sum_{i\in S_W} \sum_{w\in V\cap\nu_i} c_i(w,\top) = 2|E| - |E| = |E|$$

Since, according to the above construction, the payoff profile x associated to S_W yields:

$$x(S_W) = x_{n_0} + k * \frac{1}{2(k+1/2)} = \quad (3)$$

$$|E| - 1/2 + \frac{k}{2*k+1} = |E| - \frac{1}{2k+1} < |E| = v(S_W) \quad (4)$$

then x together with S_W is a credible deviation to the grand coalition. $\square$

PROPOSITION 3. *Given a BGC in cf-form, the characteristic function v is super-additive, that is for every $S_1, S_2 \in N$ such that $S_1 \cap S_2 = \emptyset$, the following holds: $v(S_1 \cup S_2) \geq v(S_1) + v(S_2)$.*

PROOF. Let S_1, S_2 be two disjoint coalitions and ξ_1, ξ_2 be the solutions of Equation 2, giving the values $v(S_1)$ and $v(S_2)$, respectively. Let $V_1 = \cup_{i\in S_1} \mathcal{V}_i$ and $V_2 = \cup_{i\in S_2} \mathcal{V}_i$. It is straightforward from the definition of a BGC that variable sets V_1 and V_2 are disjoint (since no two agents can have control over the same variable). Since ξ_1 and ξ_2 contain truth assignments for *every* variable in V_1 and V_2 respectively, then ξ_1 and ξ_2 cannot assign different values to the same variable. Then, for all goals γ_1 and γ_2 satisfied by ξ_1 and ξ_2 respectively, it is the case that $\xi_1 \cup \xi_2 \models \gamma_1$ and $\xi_1 \cup \xi_2 \models \gamma_2$. Therefore, each satisfied goal in S_1 or S_2 is satisfied also in $S_1 \cup S_2$ (possibly cheaper). As a result we have that $v(S_1 \cup S_2) \geq v(S_1) + v(S_2)$. $\square$

PROPOSITION 4. *The core of a BGC is non-empty.*

PROOF. This is a direct consequence of the property of super-additivity of BGC. First of all, notice that the definition of the core from Section 3 can be reformulated as:

$$core(\mathfrak{B}) = \{x \quad | \quad \forall S \subseteq N.x(S) \geq v(S)\}$$

Now, since $\mathfrak{B}$ is super-additive, it follows that $\forall S \subset N$, $v(N) \geq v(S)$, since $N = S \cup N \setminus S$ and $v(N) \geq v(S) + v(N \setminus S)$. Then, any feasible payoff profile that offers a division of $v(N)$ is in the core of $\mathfrak{B}$. $\square$

4. CONCLUSIONS AND FUTURE WORK

The game-theoretic setting proves to be a suitable direction for studying agent interaction and cooperation. The results we describe, the computational complexity for the core membership and core emptiness problems are inherently theoretic, but not without applicability. Although not quite optimistic (the core membership is shown to be co-NP-complete), the results suggest that, for small games, computing the core is possible, and could be implemented in systems that assist humans in decision-making. Also, iterative techniques, in the spirit of [2], for deriving a more efficient computation process should be researched. Following the line of [3], the goal value function from BGC's can be used to incentivise certain behavior within a group of agents. We are currently tackling these ideas, as well as attempting to produce an efficient implementation, based on our results.

5. ACKNOWLEDGMENTS

The research presented in this paper is supported by national project: "TRANSYS - Models and Techniques for Traffic Optimizing in Urban Environments", Contract No. 4/28.07.2010, Project CNCSIS-PN-II-RU-PD ID: 238.

6. REFERENCES

[1] N. Dragoni, M. Gaspari, and D. Guidi. An infrastructure to support cooperation of knowledge-level agents on the semantic grid. *Appl. Intell.*, 25(2):159–180, 2006.

[2] P. E. Dunne, W. van der Hoek, S. Kraus, and M. Wooldridge. Cooperative boolean games. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems - Volume 2*, AAMAS '08, pages 1015–1022, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems.

[3] U. Endriss, S. Kraus, J. Lang, and M. Wooldridge. Designing incentives for boolean games. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '11, pages 79–86, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.

[4] Jorge Ejarque and Raúl Sirvent and Rosa M. Badia. A multi-agent approach for semantic resource allocation. In *CloudCom*, pages 335–342, 2010.

[5] J. Lin, S. Sedigh, and A. Miller. A semantic agent framework for cyber-physical systems. In A. Elçi, M. Koné, and M. Orgun, editors, *Semantic Agent Systems*, volume 344 of *Studies in Computational Intelligence*, pages 189–213. Springer Berlin / Heidelberg, 2011.

[6] M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[7] S. Thiel, S. Dalakakis, and D. R. 0002. A learning agent for knowledge extraction from an active semantic network. In C. Ardil, editor, *IEC (Prague)*, pages 217–220. Enformatika, Çanakkale, Turkey, 2005.

[8] Y.-H. Wang, W.-N. Wang, and C.-C. Huang. An intelligent semantic agent for supervising chat rooms in e-learning system. In *ICDCS Workshops*, pages 680–685, 2005.