A Game-theoretic Approach to the Socialization of Utility-based Agents

Alcino Cunha
alcino@di.uminho.pt

José Neves
jneves@di-ia.uminho.pt

Departamento de Informática, Universidade do Minho
Largo do Paço, 4709 Braga Codex, PORTUGAL

Abstract

This paper presents a formal framework in which to study the socialization processes evolving among utility-based agents. These agents are self-interested, being their different social attitudes (cooperativeness, competitiveness or indifference) a consequence of this behavior. The dynamics of the socialization process are captured by a relation that measures the similarities between the desires of two groups of agents. This similitude relation is derived from the system's model, defined as a probabilistic transition system and a set of individual preference relations. Game-theoretic concepts are used in order to determine the rational (or expected) transitions of the system.

1. Introduction

The work presented in this paper concerns the study of the socialization process among utility-based agents. Finding the results of this process consists in determining the social attitudes that each agent or group of agents exhibits towards the other elements of the society. Traditionally, software agents are designed to play a specific role in a predefined society. Under this perspective, the socialization problem is of no interest since the social attitude of an agent is fixed and it is a direct consequence of its role. In the design of a multi-agent system, the existence of a common problem to be solved usually results in a *benevolence assumption*, that states that all the system's agents are equally committed in solving a problem the best they can. This traditional way of viewing multi-agent systems is gradually disappearing, in particular due to the fact that the software world is becoming extremely dynamic and diverse; software agents should be prepared to inter-operate in a changing and heterogeneous environment, and their social attitudes should change according to the problems being solved and the involved parties.

According to their social attitudes, agents are usually divided into categories [3]: *cooperative*, if they try to maximize the welfare of a group of agents; *self-interested*, if they only care about its own interests and welfare; and *hostile* or *competitive*, if the maximization of its welfare implies the decrease of the others' satisfaction. An utility-based agent acts according to the principle of expected utility maximization, and, consequently, it is always a self-interested entity. Cooperativeness or competitiveness may still exist, but only as an indirect consequence of this self-interested behavior. One's view of the agent's socialization problem relies on this assumption, and the algorithm to be described attempts to determine and analyze such indirect social attitude.

2. Modeling societies of utility-based agents

The proposed model of utility-based multi-agent systems is based on classical concepts of game [2] and concurrency theory [4]. Transition systems are used as a starting point for the model. The agents' beliefs are represented by *information partitions*, a concept widely used in game-theory to represent the *beliefs* of the players. The *desires* of an agent are assessed by a simple real-valued utility function defined over states. Some agents may prefer to reach a state s within a short period of time, than a state s' with higher utility within a larger one. To capture this kind of time preference, *discount factors* are also incorporated into the model. Nondeterminism is modeled by assigning a lottery to the possible outcomes of executing an action. Lot_S will denote the set of all possible lotteries defined over the set of states S.

Definition 1 *A model of a utility-based multi-agent system is an octet*

$$M = (A, S, i, L, t, u, \delta, B)$$

where $A = \{A_1, \ldots, A_n\}$ *is an ordered set of agent names;* S *is a set of world states, being* i *the initial state;* $L = \prod_{a \in A}(L_a \cup \{*\})$ *is a set of possible action labels, where* L_a *represents the set of the possible actions of agent* a *and* $*$ *represents the null action;* $t : S \times L \hookrightarrow Lot_S$ *defines* *the transition relation over the state space (it is a partial function, once it may be impossible for a given group of agents to execute a particular action in a given state); $u = \{u_a\}_{a \in A}$ is an indexed set with the utility functions, where $u_a : S \cup Lot_S \to I\!R$; $\delta = \{\delta_a\}_{a \in A}$ is an indexed set with the discount factors; $B = \{B_a\}_{a \in A}$ is an indexed set of information partitions.*

Since that in each state all the agents choose their actions simultaneously, the states of the transition system will be modeled by the normal-form representation of a static game. The solution concept that will be used in order to determine the strategies that will be followed in each of these games is that of the Nash equilibrium [1].

Suppose that a given joint-strategy c was accorded to be executed by a group of agents. In each state s, $c(s)$ determines which joint action will be executed in state s. If the system is deterministic, and the initial state is i, the execution trace that would be induced by this strategy is

$$s_0 = i, s_1 = t(s_0, c(s_0)), s_2 = t(s_1, c(s_1)), \ldots$$

The expected utility for agent $a \in A$, if c is executed, would be the weighted sum

$$u_a(s_0) + \delta_a \times u_a(s_1) + \delta_a^2 \times u_a(s_2) + \ldots$$

If only the first n states are considered when calculating the expected utility of an agent, the percentage of the total weight that will be covered is

$$\alpha = \frac{\sum_{x=0}^{n} \delta^x}{1/(1-\delta)}$$

If the reverse process is adopted, that is, if the percentage of the total weight that is to be analyzed (the parameter α) is specified in the first place, then it is possible to determine the number of states n that should be looked ahead when determining the reward of an agent's decision. The parameter α can be viewed either as the confidence one has in the expected utility that is being calculated, or as the measure of the bounded rationality of the system's agents. A value near 0 characterizes a system where the agents have very limited resources, and a value near 1 characterizes a system where the agents are almost perfect reasoners. After determining n, the equilibrium strategies (and hence the rational transitions) should be evaluated by backward induction, starting at the states that are situated n actions ahead from the initial state i.

3. Agent socialization

The social attitudes that emerge in a multi-agent system will be represented through a similitude relation, which expresses the similarities between the agents *desires*.

Definition 2 *The similitude relation between two groups of agents, in a multi-agent system $M = (A, S, i, L, t, u, \delta, B)$, is expressed by the function*

$$sr : \mathcal{P}(A) \times \mathcal{P}(A) \times \alpha \to I\!R$$

which is to be understood as a measure of the similitude among the desires of both groups. A negative result means that the first group acts in a competitive or hostile manner towards the second group. A similitude of 0 means that the attitudes of the first group do not influence the satisfaction of the second. A positive result indicates a cooperative behavior. α is the result's confidence value.

Given a procedure to determine the expected or rational transitions of a system (such as the one presented in the previous section), the process of determining $sr(J, K, \alpha)$, under M, may now be stated as follows: (1) Determine the rational transitions of the multi-agent system M with confidence value α. (2) Determine the average expected utility u for all agents in K. (3) Define a model M', similar to M, taking into consideration only the actions that do not had active participation by agents in J; i.e. $M' = (A, S, i, L', t, u, \delta, B)$ with $L' = \{l | l \in L \wedge \forall j \in J \cdot \pi_j(l) = *\}$, where $\pi_j(l)$ denotes the j's contribution to the joint-action l. (4) Determine the rational transitions of the multi-agent system M' with confidence value α. (5) Determine the average expected utility u' for all agents in K, taking into account the rational traces of M'. (6) Evaluate $sr(J, K, \alpha) = u' - u$.

Consider the following example: imagine that one wants to build an agent r, by assembling two independent and simpler utility based agents named, respectively, h and v. The question that may arise is which of the agents should be used, such that the final agent will act as expected. The solution to this problem consists in finding h and v among the available instances such that $sr(\{h, v\}, \{r\}, \alpha)$ is maximized. If, given two models for the agents h and v, one is asked if h and v are a good choice to build the agent r, then one has to prove that $sr(\{h, v\}, \{r\}, \alpha) > 0$.

References

[1] J. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36:48–49, 1950.

[2] E. Rasmusen. *Games and Information: an Introduction to Game Theory*. Basil Blackwell, 1989.

[3] T. Sandholm and V. R. Lesser. On automated contracting in multi-enterprise manufacturing. In *Improving Manufacturing Performance in a Distributed Enterprise: Advanced Systems and Tools*, pages 33–42, Edimburgh, Scotland, 1995.

[4] G. Winskel and M. Nielsen. Models of concurrency. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 4, pages 1–148. Clarendon Press, Oxford, 1995.