

MADeM: a multi-modal decision making for social MAS

Francisco Grimaldo
Computer Science Dept.
University of Valencia
Av. Vicent Andrés Estellés s/n
46100 Burjassot, Spain
francisco.grimaldo@uv.es

Miguel Lozano
Computer Science Dept.
University of Valencia
Av. Vicent Andrés Estellés s/n
46100 Burjassot, Spain
miguel.lozano@uv.es

Fernando Barber
Computer Science Dept.
University of Valencia
Av. Vicent Andrés Estellés s/n
46100 Burjassot, Spain
Fernando.Barber@uv.es

ABSTRACT

This paper presents MADeM, a multi-modal agent decision making to provide virtual agents with socially acceptable decisions. We consider multi-modal decisions as those that are able to merge multiple information sources received from a MAS. MADeM performs social decisions since it relies on auctions, a well known market-based coordination mechanism. Our social agents express their preferences for the different solutions considered for a specific decision problem, using utility functions. Therefore, coordinated social behaviors such as task passing or planned meetings can be evaluated to finally obtain socially acceptable behaviors. Additionally, MADeM is able to simulate different kinds of societies (e.g. elitist, utilitarian, etc), as well as social attitudes of their members such as, egoism, altruism, indifference or reciprocity. MADeM agents have been successfully verified in a 3D dynamic environment while simulating a virtual university bar, where different types of waiters (eg. coordinated, social, egalitarian) and customers (e.g. social, lazy) interact to finally animate complex social scenes.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence - *Multiagent systems / Intelligent Agents*

General Terms

Design, Economics, Experimentation

Keywords

Social Reasoning, Multiagent Resource Allocation, Welfare Economics

1. INTRODUCTION AND RELATED WORK

Decision making is the cognitive process leading to the selection of a course of action among variations. There are several factors that influence this process, although probably the most important could be the amount of information the agent manages when deciding its actions. This is specially important in social MAS simulations, since our aim is

Cite as: MADeM: a multi-modal decision making for social MAS, Francisco Grimaldo, Miguel Lozano, Fernando Barber, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 183-190.

Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

to produce socially intelligent agents capable of displaying acceptable decisions for a specific society model.

From virtual agents community, the behavioral animation problem points to the construction of an intelligent system which is able to integrate the different techniques required for the realistic simulation of virtual humans behavior. Among them, we can include perception, motion control, goal selection, action execution, communication between agents, interaction with the environment, etc. Traditionally, designers have sought to make their agents rational so that they can behave efficiently (i.e. the shorter the plan the better). Therefore, social simulations have incorporated group coordination due to the fact that self interested agents (i.e. agents devoted to accomplish a set of goals) easily come into conflicts in a resource bounded environment even though their goals are compatible.

Socially intelligent agents are autonomous problem solvers that have to achieve their goals by interacting with other similarly autonomous entities [12]. Bearing this in mind, multi-agent systems are normally referred to as societies of agents, and provide an elegant and formal framework to design social behaviors for autonomous agents.

Much research has been done in behavioral animation of virtual agents for the last few years [1, 17, 19]. The pioneer work of Dimitri Terzopoulos [20] showed how to design a natural ecosystems animation framework with minimal input from the animator. He simulated *Artificial fishes* in virtual underwater worlds. However, human behavior is clearly different and more complex to emulate. In [16, 19] the goal is to design agents with a high degree of autonomy without losing control. These agents are an extension of the BDI architecture described in [15], and they include internal states such as emotions, reliability, trust and others. Emotional architectures have been also applied to virtual agents (animals and humans) to manage sociability and rationality and to produce believable groups of synthetic characters [8, 10, 13]. MAS-SOC [2] aims at creating a platform for multi-agent based social simulations, which is similar to our purposes. In this context, there is ongoing work in order to incorporate social-reasoning mechanisms based on exchange values [18].

Behavioral animation has also been tackled from the field of coordinated multi-agent systems. For example in Generalized Partial Global Planning (GPGP) [7], agents merge the meta-plans describing their operational procedures and figure out the better action in order to maximize the global utility. Another example is Multi-agent Planning Language (MAPL) [4], which assigns the control over each resource to a unique agent and uses speech-acts to synchronize planned

tasks. Collaboration is supported in the RETSINA system [9] thanks to the use of communicative acts that synchronize tasks and occasionally manage conflicts. Team formation and task coordination for heuristic search planning characters is presented in [11] to adapt better to the dynamism of shared environments.

Several approaches show realistic task-oriented behaviors, but autonomous virtual humans should also display social behaviors (e.g. interchanging information with their partners or grouping and chatting with their friends). This kind of social agents is required in many complex simulation environments: military/civil simulations, social pedestrians in virtual cities, games, etc. We consider that socially acceptable agents need to evaluate the social impact of their actions, so they could decide what to do in accordance with the society being simulated.

The purpose of MADeM is to provide a MAS simulation framework with agents managing multi-modal social decisions. We introduce this kind of decisions as those able to consider different focuses of attention coming from different sources. On the one hand, different focuses of attention (i.e. points of view) are required to provide more informed self-interested decisions. For instance, when deciding their actions, humans easily balance several aspects such as: efficiency, tiredness, skillfulness, mood, etc. On the other hand, a MADeM agent integrates social influence by asking external agents about the points of view the former is interested in. This feedback represents the preferences of the others for a specific situation, that basically includes certain resource and task allocations (see section 2.1). Internally, MADeM is based on the MARA theory [5] and it uses auctions as a basic procedure to provide the social feedback mentioned (see [6] for a good set of social welfare examples). However, we let the agents manage more than one auction and allocation associated to each decision, so they can simulate multi-modal social decisions in a MAS.

The rest of the paper is organized as follows: The next section presents the MADeM information domain, that basically includes the types of resources being used (section 2.1) and the agent preferences representation (section 2.2). Then, we explain the decision making procedure, including the auctioning process (section 3) and the winner determination problems faced (sections 3.1, 3.2). Section 4 shows the MAS framework designed to integrate MADeM agents and an application example created to verify the agent social skills. Finally, we show the social behaviors obtained by different types of agents (eg. waiters, customers) according to the main MADeM parameters.

2. MADEM: MULTI-MODAL AGENT DECISION MAKING

The multi-modal agent decision making presented in this paper (MADeM) is based on the MARA theory, thus, it shares a similar domain of definition but making some additions to it. We summarize them as follows:

- A set of agents $A = \{a_1, \dots, a_n\}$ where each a_i represents a particular agent involved in the decision. A vector of weights $\vec{w} = \langle w_1, \dots, w_n \rangle$ is associated to each agent representing the internal attitude of the agent towards other individuals. This information will be used to weigh the information received from other agents (section 3.1).

- A set of resources to be allocated by the agents $R = \{r_1, \dots, r_m\}$. The definition of resources we use is different from the classical definitions found in the literature and it will be explained in detail in section 2.1.
- Instead of having only one utility function as in classical MARA problems, each agent will have a set of utility functions $\{U^1, U^2, \dots, U^q\}$. These utility functions will be used to evaluate the allocations from different points of view. We better explain the utility functions for our agents in section 2.2. Additionally, each agent will have a vector of utility weights $\vec{w}_u = \langle w_{u_1}, \dots, w_{u_q} \rangle$ representing the importance given to each point of view in the multi-modal agent decision making.

MADeM also has a winner determination procedure to decide the winner allocation of an auction and consequently, simulate socially acceptable decisions of agents. But as complex decisions may need to consider more than one point of view, the auctioneer may require from each agent to express its preferences using several utility functions. Therefore the auctioneer will receive several utility values for the proposed allocations and, as a result, will have to execute several winner determination procedures, one for each point of view, and a final multi-modal winner determination procedure to decide which point of view has the winner allocation. The details of the whole winner determination procedure will be explained in sections 3.1 and 3.2.

An example illustrating the functionality of MADeM could be a customer entering a virtual bar and deciding which waiter to place his order to. Different points of view could be considered: to ask the nearer waiter (i.e. tiredness), to ask the less occupied waiter (i.e. utilitarianism) or to ask a waiter who is a friend of him (i.e. sociability). Auctions are used to select the better candidates for each category. In this case, utility weights would express personal tendencies such as laziness, impatience or sociability. Looking at them, MADeM would be able to choose among these three possibilities.

2.1 Types of resources

In order to obtain social and intelligent behaviors, we provide the agents with the ability to ask for opinion or social feedback about different solutions for a specific decision problem (e.g. whether to pass the execution of a task to another agent or not). Agents do this by following an auctioning mechanism in which the resources being auctioned are tasks. Task auctioning has been widely used in the MAS community but its application to social virtual agents is not so common and needs to take new issues into account. The novelty of our approach is that we do not auction only the execution of a task as found in the literature. Instead, the resource to be auctioned is what we term *task slots*. We consider task slots as slots that need to be assigned in order to execute a task, thus, they can be considered as parameters of the action. When considering any kind of task, we differentiate two main types of task slots: *agent slots* which correspond to agents that play different roles in the task execution, as for example the executant of the task or the beneficiary of the task; and *object slots*, that correspond to objects needed to perform the action.

We also differentiate slots depending on the role they play in the task. There is a slot present in every kind of task

which is the agent who carries it out (Ag_e). Besides, we consider as general slots two additional slots present in many tasks: the main object of the action (Obj_m) and the beneficiary or destinatory in any sense of the action (Ag_d). For example, in the *Give* action, Ag_e is the agent who initially has the object, Obj_m is the object being given and Ag_d is the agent who receives the object. In any case, it is always possible for a particular problem to consider more task slots if necessary.

It can be seen that the classical task auctioning is subsumed within our approach as it corresponds to the auctioning of the Ag_e slot. Moreover, the auctioning of objects is also possible by auctioning the appropriate task slot of the considered action (e.g. the Obj_m slot of an action *Own*).

According to this, we represent resources or task slots as $r = task(slot)$, where we consider the slot as a typed parameter according to an ontology defined for the problem. An assignment of an element (agent or object) to a slot is then represented as $task(slot) \leftarrow element$, and an allocation P of elements to task slots has the following representation:

$$P = \{t_1(s_1) \leftarrow e_1, t_1(s_2) \leftarrow e_2, \dots, t_1(s_n) \leftarrow e_n, t_2(s_1) \leftarrow e_{n+1}, \dots, t_m(s_n) \leftarrow e_{n+m}\}$$

Therefore, we represent each one of the solutions considered for a decision problem as an allocation of this type.

For example, in the bar environment presented in section 4 we have modelled the task *Give(Waiter, Object, Customer)* as the action used by waiters to give a product to a customer. In this task we have 2 slots that may be auctioned: the Ag_e , which is the waiter that gives, and the Ag_d , which is the customer that receives the product. These slots would be represented as *Give(Ag_e)* and *Give(Ag_d)*. Hence, a possible allocation for these slots could be $\langle Give(Ag_e) \leftarrow a_1, Give(Ag_d) \leftarrow a_5 \rangle$, given that a_1 is a waiter and a_5 is a customer.

2.2 Agent preferences

In the auctioning process, the auctioneer asks the agents to bid for one or more task slot allocations, each bidder then uses its utility functions to evaluate and score the different allocations being considered.

MADeM agents use non-negative 2-additive utility functions of the form:

$$U(P) = \sum_{p \in P} u(p) + \sum_{p_1, p_2 \in P} u(p_1, p_2) \quad (1)$$

where $u(p)$ is a utility function for a unique slot assignment and $u(p_1, p_2)$ is the extra utility value given to the situation in which both assignments are done. In order to obtain the utility functions ($U(P)$) normalized in the range $[0, 1]$, the utility functions of all agents must be divided by the same constant, which will depend on the particular problem.

Despite the fact that the types of resources being auctioned are tasks, utility functions express benefit. Therefore, agents would aim at maximizing their utility. Using cost functions could also be possible but agents would aim at minimizing their costs.

3. DECISION MAKING PROCEDURE

MADeM uses one-round sealed-bid combinatorial auctions to choose among different solutions to a decision problem.

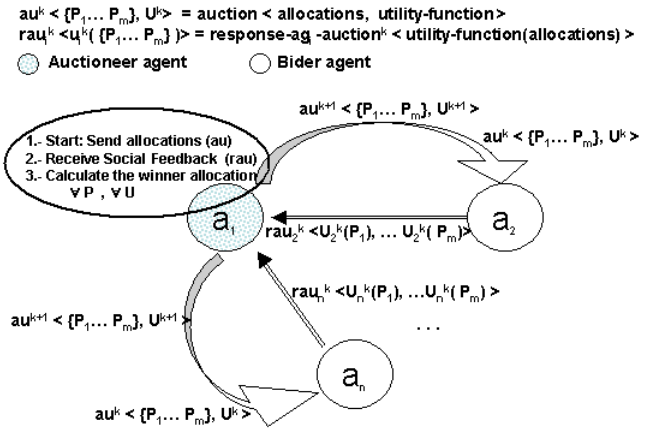


Figure 1: MADeM Procedure

As mentioned above, a solution is represented by an allocation of one or many task slots. Auctioneer and bidder roles are not played by fixed agents throughout the simulation. Instead of that, every agent can dynamically adopt each role depending on his/her needs or interests. For example, an agent would be the auctioneer when he wanted to pass a task to another agent following a social behavior. On the other hand, agents receiving the auction would bid their utility values provided that they were interested in the task slots being auctioned. Our model allows more than one auction to be running at the same time. Similarly, agents can participate in several auctions simultaneously. Thus, this approach lies in between centralized and distributed market-based allocation.

An overview of the multi-modal decision making procedure followed by the agents to generate socially acceptable decisions is shown in figure 1. This procedure is mainly based on the following steps:

1. *Auctioning phase*: This phase is carried out by an agent (a_1) who wants to socially solve a decision problem (e.g. where to sit). This agent then constructs the set of allocations representing all the possible solutions for the problem ($\langle P_1, P_2, \dots, P_m \rangle$). These allocations have the form of task slots assignments such as *SitAt(Obj_m)* $\leftarrow table_1$. Next, he auctions them to a particular group of agents, that we call the target agents. Each auction also includes a single type of utility function that the agent is interested in evaluating from the others ($au^k(\langle P_1, P_2, \dots, P_m \rangle, U^k)$). As complex decisions require to take into consideration more than one point of view, the auctioneer agent can start different auctions for the same set of allocations (au^1 through au^g).

The process to select the target agents for an auction varies depending on the kinds of task slots being allocated. When dealing with agent slots, the target agents can be extracted from the type of the task slot being auctioned. For example, in our bar environment, the task slot *Make_Coffee(Ag_e)* should only be auctioned to agents of the class *Waiter*. On the other hand, when allocating object slots, the target agents could be those agents that are somehow related to objects of the same type as the task slot being auctioned.

For instance, the task slot $SitAt(Obj_m)$, where Obj_m corresponds with the table where to sit at, could be auctioned to all the agents related to a table in the environment. When there is no type system available, the target agents would be the whole set of agents.

2. *Bidding phase*: Since the auctioneer informs about both the task slot allocations and the utility functions being considered, bidders simply have to compute the requested utility functions and return the values corresponding to each auction back to the auctioneer ($ra_i^k = \langle U_i^k(P_1), \dots, U_i^k(P_m) \rangle$).
3. *Winner determination phase*: In this phase, the auctioneer selects a winner allocation for each launched auction, that is, for each point of view being considered. To do this, he uses a classical winner determination problem, as explained in section 3.1. Afterwards, he chooses one final winner allocation among these auction winners using a multi-modal decision making process. The details of this calculation are fully described in section 3.2. Thus, the final winner allocation will represent an acceptable decision for the society being simulated.

3.1 Winner Determination Problem

Once bidders have answered to an auction call (no answering means no preference, therefore, utility zero) the auctioneer agent has the utility values ($U_i(P_j)$) given by each bidder ($i \in A$) to every allocation being evaluated (P_j). Equation 2 groups these utility values in a set of vectors, one for each allocation.

$$\overrightarrow{U(P_j)} = \langle U_1(P_j), \dots, U_n(P_j) \rangle \quad \forall j \in [1..m] \quad (2)$$

Remember that every agent had an associated vector of weights representing its attitude towards the other individuals (see equation 3). According to it, the auctioneer weighs the utility vectors in equation 2 doing a component by component multiplication with the attitude vector as shown in equation 4.

$$\overrightarrow{w} = \langle w_1, \dots, w_n \rangle \quad (3)$$

$$\overrightarrow{U_w(P_j)} = \overrightarrow{U(P_j)} * \overrightarrow{w} \quad \forall j \in [1..m] \quad (4)$$

Attitude weights are used to model the social behavior of the auctioneer agent. For example, egoism is modelled giving weight 1 to himself and 0 to all other agents. In fact, a whole range of behaviors between egoism and altruism can be modelled using the vector of equation (5), where $p = 0$ represents the previous behavior, $p = 1$ represents total altruism and $p = 0.5$ represents an egalitarian behavior or indifference between oneself and the rest of the agents.

$$\begin{aligned} \text{Egoism} - \text{Altruism} : \quad \overrightarrow{w} = \langle p, \dots, p, 1 - p, p, \dots, p \rangle \\ w_i = 1 - p, w_{j \neq i} = p, i = \text{Myself} \end{aligned} \quad (5)$$

It is also possible to model reciprocal attitudes by means of the vector \overrightarrow{w} . A simple example is shown in equation 6, where weights are based on the interchange of favors between agents.

$$\text{Reciprocity} : \quad w_i = \frac{\text{Favors_from}(i)}{\text{Favors_to}(i)} \quad (6)$$

In order to behave socially, the auctioneer agent attends to the social welfare value when selecting the winner allocation of an auction. Therefore, the winner determination problem chooses the allocation that maximizes the welfare of the society (see equation 7).

$$\text{Winner} = P_w \longleftrightarrow sw(P_w) = \max_{j \in [1..m]} sw(P_j) \quad (7)$$

To compute the social welfare of an allocation, the auctioneer uses Collective Utility Functions (CUFs) and the weighted utilities defined in equation 4 (as shown in equation 8). MADeM allows to select among different CUFs when evaluating the social welfare of an allocation. At the moment, four CUFs have been integrated in MADeM, each one related to a kind of society: utilitarian, egalitarian, elitist and nash .

$$sw(P) = \text{cuf}(\overrightarrow{U_w(P)}) \quad (8)$$

$$\begin{aligned} \text{cuf}_{\text{utilitarian}} &= \Sigma u_w(i) \\ \text{cuf}_{\text{egalitarian}} &= \min\{u_w(i)\} \\ \text{cuf}_{\text{elitist}} &= \max\{u_w(i)\} \\ \text{cuf}_{\text{nash}} &= \Pi u_w(i) \end{aligned} \quad (9)$$

3.2 Multi-modal decision making

An agent can ask other agents about different points of view (e.g. efficiency, tiredness, etc). In order to do this, he performs several auctions with different types of utility functions (see parameter U^k in figure 1). Once all these auctions have been resolved, the auctioneer has the winner allocation for each point of view and the social welfare obtained provided that allocation is adopted (see equation 10).

$$\begin{aligned} \text{auction}_1(\{P_1, P_2, \dots, P_m\}, U^1) &\longrightarrow (P_{w1}, sw(P_{w1})) \\ &\dots \\ \text{auction}_k(\{P_1, P_2, \dots, P_m\}, U^k) &\longrightarrow (P_{wk}, sw(P_{wk})) \end{aligned} \quad (10)$$

The final winner allocation is then chosen using the vector of utility weights $\overrightarrow{w_u}$. MADeM allows two different ways to decide the final winner:

- The first possibility is to take the values in $\overrightarrow{w_u}$ as weights assigned to each utility function. Hence, the final winner allocation is that which maximizes the welfare of the society after having multiplied it by the corresponding utility weight (see equation 11),

$$P_w = P_{wi} \longleftrightarrow sw(P_{wi}) = \max_{i \in [1..k]} w_u(i) * sw(P_{wi}) \quad (11)$$

- The second possibility is to consider the values in $\overrightarrow{w_u}$ as the probabilities for the auctioneer to choose the corresponding point of view. In this case, the final winner allocation is chosen probabilistically.

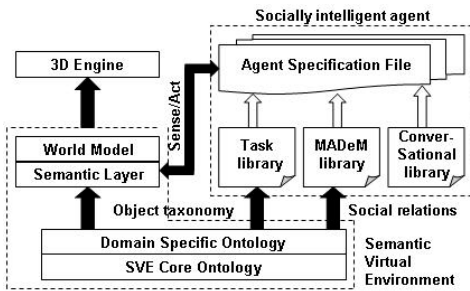


Figure 2: Multi-agent simulation framework.

4. APPLICATION EXAMPLE

In this section we show how we have integrated MADeM into a multi-agent framework oriented to simulate socially intelligent characters in 3D virtual environments. This framework is developed over Jason [3], which allows the definition of BDI agents using an extended version of AgentSpeak(L) [14]. Figure 2 depicts the architecture of the system, which can be basically divided into two parts:

- The *Semantic Virtual Environment* uses ontologies to define the world knowledge base (i.e. the object taxonomy and the object interrelations) as well as the set of all possible relations among the agents within an artificial society. The environment is handled by the *Semantic Layer*, which acts as an interface between the agent and the world, thus sensing and executing the actions requested by the agents. The animation system – virtual characters, motion tables, etc. – is located at the *3D Engine* that can extract graphical information from the *World Model* database thus performing visualization.
- *Socially intelligent agents* receive sensorial information from the *Semantic Layer* and calculate the appropriate sequence of actions in order to achieve their goals. The agent’s finite state machine is defined in the *Agent Specification File*. It calls the following libraries to enrich agent behavior: the *Task Library* contains the operators that sequence the actions needed to animate a task; the *MADeM Library* provides the agents with the mechanisms to make social decisions; finally, the *Conversational Library* contains the set of plans that handle the animation of the interactions between characters (e.g. ask someone a favor, planned meetings, chats between friends...).

In order to test MADeM, we have created a virtual university bar where waiters take orders placed by customers (see figure 3). The typical objects in a bar, such as a juice machine, behave like dispensers that have an associated time of use to supply their products (e.g. 2 minutes to make an orange juice) and they can only be occupied by one agent at a time. Therefore, waiters should coordinate to avoid conflicts. Agents can be socially linked using the concepts defined in the ontology. For example, waiters and customers create social relationships with their friends and this defined social network is used when deciding whether to do favors, to promote social meetings, etc.

Both waiters and customers carry out tasks in the virtual bar and they use MADeM to decide among different task

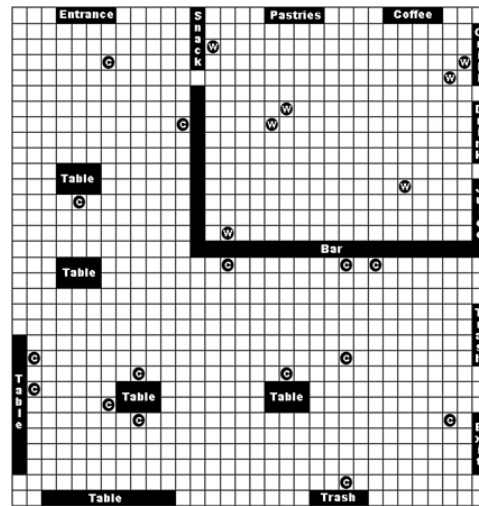


Figure 3: 2D virtual university bar environment

Tasks/Slots	Ag_e	Obj_m	Ag_d
Use	Waiter	ProductDispenser	-
Give	Waiter	Product	Customer
SitAt	Customer	SeatingPlace	-

Table 1: Tasks and slots being considered.

slot assignments. Table 1 summarizes the tasks and the slots being considered for each task. Waiters perform tasks *Use* and *Give* and they try to allocate them to the best candidate. Hence, the slot being auctioned will be the executor of the task (Ag_e) for both of them. However, since the decision problem of the customers is to choose among different *SeatingPlaces* to consume, they auction the slot Obj_m of the action *SitAt* that corresponds to the place where to sit. The classes of objects used to describe task slots are extracted from object taxonomy defined in the ontology. Therefore, objects such as a table or the bar would be implemented as instances of the class *SeatingPlace* and they could be considered as places to sit at.

Waiters serve orders basically in two steps: first, using the corresponding dispenser (e.g. the grill to produce a sandwich); and second, giving the product to the customer. Tasks are always auctioned using MADeM before their execution in order to find good social allocations. The set of allocations among which to decide is represented in equation 12. That is, for each task, a waiter evaluates whether to carry out the task against the chance to pass it to another waiter and perform his next task.

$$\begin{cases} P_0 = \{t(Ag_e) \leftarrow Myself\} \\ P_i = \{t(Ag_e) \leftarrow a_i, t_{next}(Ag_e) \leftarrow Myself\} \quad \forall a_i \in A \end{cases} \quad (12)$$

Waiters take into account three points of view when calling MADeM: *performance*, *chatting* and *tiredness*. Equations 13 and 14 define the utility values returned by the performance utility function for the tasks *Use* and *Give*. This function aims at maximizing the number of tasks being performed at the same time and represents the waiters’ willingness to serve orders as fast as possible. Social be-

haviors defined for a waiter are oriented to animate chats among his friends at work. Therefore, waiters implement the social utility function detailed in equations 15 and 16, where *Near* computes the distance between the agents while they are executing a pair of tasks. These functions evaluate social interest as the chance to meet a friend in the near future, thus performing a planned meeting. Finally, equation 17 defines the tiredness utility function for a waiter. This later function implements the basic principle of minimum energy, widely applied by humans at work. The type of society being simulated for waiters is elitist, thus, waiters will choose those allocations that maximize the utility functions previously defined.

$$U^{perf}(\text{Use}(Ag_e) \leftarrow a) = \begin{cases} 0 & \text{if } a \neq \text{Myself} \\ 1 & \text{if } (a = \text{Myself}) \wedge \\ & \text{IsUsing}(\text{Myself}, \text{Obj}_m) \wedge \\ & \text{not}(\text{IsComplete}(\text{Obj}_m)) \\ \frac{1}{t_{to be free} + t_{queue}} & \text{Otherwise} \end{cases} \quad (13)$$

$$U^{perf}(\text{Give}(Ag_e) \leftarrow a) = \begin{cases} 0 & \text{if } a \neq \text{Myself} \\ 1 & \text{if } (a = \text{Myself}) \wedge \\ & \text{CurrentTask} = \text{'Give'} \wedge \\ & \text{not}(\text{HandsBusy}(\text{Myself}) < 2) \\ \frac{1}{t_{to be free}} & \text{Otherwise} \end{cases} \quad (14)$$

$$U^{soc}(t_1(Ag_e) \leftarrow a_1, t_2(Ag_e) \leftarrow a_2) = \begin{cases} 0 & \text{if } a_1 \neq \text{Myself} \vee a_2 \neq \text{Auctioneer} \\ 1 & \text{if } (a_1 = \text{Myself}) \wedge a_2 = \text{Auctioneer} \wedge \\ & \text{IsFriend}(a_1, a_2) \wedge \text{Near}(\text{Pos}(t_1), \text{Pos}(t_2)) \wedge \\ & \text{ExecTime}(t_2) > \text{RemainTime}(\text{CurrentTask}) \\ 0 & \text{Otherwise} \end{cases} \quad (15)$$

$$U^{soc}(t(Ag_e) \leftarrow a) = \begin{cases} 0 & \text{if } a \neq \text{Auctioneer} \\ 1 & \text{if } a = \text{Auctioneer} \wedge \text{IsFriend}(a, \text{Myself}) \\ & \wedge \text{Near}(\text{Pos}(\text{CurrentTask}), \text{Pos}(t)) \wedge \\ & \wedge \text{TimeToStart}(t, a) = \text{Now} \\ 0 & \text{Otherwise} \end{cases} \quad (16)$$

$$U^{tir}(t(Ag_e) \leftarrow a) = \begin{cases} 1 - \frac{\text{tasks_done}}{\text{total_tasks}} & \text{if } a = \text{Myself} \\ 0 & \text{Otherwise} \end{cases} \quad (17)$$

On the other hand, customers place orders and consume them when served. Now, we are interested in animating interactions between customers that are consuming with their friends. Therefore, customers call MADeM to solve the problem of *where to sit*. The set of allocations among which to decide is defined in equation 18. In this case, the task slot being auctioned is the place where to sit and the candidates being evaluated are all the tables in the environment as well as the bar.

$$\begin{cases} P_0 = \{\text{SitAt}(\text{Obj}_m) \leftarrow \text{bar}\} \\ P_j = \{\text{SitAt}(\text{Obj}_m) \leftarrow \text{table}\} \quad \forall \text{table} \in \text{Tables} \end{cases} \quad (18)$$

Customers consider two points of view when calling MADeM: *sociability* and *laziness*. Equations 19 and 20 correspond to the social utility function defined for customers.

This function assigns a maximum value to a table provided that there is a friend sitting on it. To consume standing up at the bar is not considered of social interest at all, hence, its utility value is defined as zero. The laziness utility function is represented by equations 21 and 22. In this function, tables are evaluated according to their distance to the customer and, opposite to sociability, standing at the bar is now considered the best option. The type of society being simulated for customers is utilitarian, therefore, customers will choose those allocations that maximize the addition of the utility values previously defined.

$$U^{soc}(\text{SitAt}(\text{Obj}_m) \leftarrow \text{table}) = \begin{cases} 1 & \text{if } \text{AtTable}(\text{Myself}, \text{table}) \wedge \\ & \text{IsFriend}(\text{Myself}, \text{Auctioneer}) \wedge \\ & \text{AvailablePlace}(\text{table}) \\ 0 & \text{Otherwise} \end{cases} \quad (19)$$

$$U^{soc}(\text{SitAt}(\text{Obj}_m) \leftarrow \text{bar}) = 0 \quad (20)$$

$$U^{laz}(\text{SitAt}(\text{Obj}_m) \leftarrow \text{table}) = \begin{cases} \frac{\min_{tbl \in \text{Tables}} \{\text{DistanceTo}(tbl)\}}{\text{DistanceTo}(\text{table})} & \text{if } \text{Auctioneer} = \text{Myself} \wedge \\ & \text{AvailablePlace}(\text{table}) \\ 0 & \text{Otherwise} \end{cases} \quad (21)$$

$$U^{laz}(\text{SitAt}(\text{Obj}_m) \leftarrow \text{bar}) = \begin{cases} 1 & \text{if } \text{Auctioneer} = \text{Myself} \wedge \\ & \text{AvailablePlace}(\text{bar}) \\ 0 & \text{Otherwise} \end{cases} \quad (22)$$

As a final discussion, we would like to point out some implementation details that we have adopted in order to lower the complexity of the winner determination problem. MADeM considers the vector of utility weights w_u as the probability of each point of view, thus, it probabilistically selects which utility function will be finally taken into account and only performs one auction with this utility function. Moreover, the utility functions have been defined so that agents only express a preference different to zero under certain conditions of the allocations. Since a zero utility is similar to not voting, these utilities are not sent back to the auctioneer and the maximum number of responses a bidder gives to any auction is two (e.g. when waiters evaluate the social utility function for the two situations in which either the auctioneer or themselves perform the task). In fact, no auction at all is necessary when customers decide to behave following their laziness utility function, as it equals to zero for all bidders except for the auctioneer.

5. RESULTS

In order to verify the social outcomes obtained with MADeM agents, we have simulated different types of waiters serving customers (see figure 4 for an snapshot of the running 3D virtual environment). The results shown in this section correspond to simulations where 10 waiters attend 100 customers.

As we have previously mentioned, we have modelled an elitist society of waiters within which agents attend to three points of view (i.e. performance, sociability and tiredness), each of them represented by its own utility function. In this context, utility weights can be adjusted to create different types of social waiters. For example, a *coordinated*




Figure 4: 3D virtual university bar environment.

Agent	σ_{NTasks}	$NChats$	Tasks/s
Coordinated	6.73	5	0.91
Social	4.37	29.4	0.65
Egalitarian	2.74	6.6	0.62
Self-interested	-	-	0.17

Table 2: Results for different types of waiters.

waiter could be an agent that chooses its decisions following performance 75% of the times and following sociability or tiredness in the rest of the situations. The vector of utility weights for a *coordinated* waiter would then be $\vec{w}_u = \langle 0.75, 0.125, 0.125 \rangle$, where each component represents the importance given to each utility function being evaluated. Similarly, we have defined *social* waiters as agents with the following vector of utility weights $\vec{w}_u = \langle 0.125, 0.75, 0.125 \rangle$ and *egalitarian* waiters as agents with $\vec{w}_u = \langle 0.125, 0.125, 0.75 \rangle$. Table 2 summarizes some results obtained with *coordinated*, *social* and *egalitarian* waiters against *self-interested* waiters with no social mechanism included. *Coordinated* waiters perform better (see column *Tasks/s*) since the majority of conflicts caused by the use of the same dispenser (e.g. the coffee machine) are resolved with specialization, that is, by passing the task to another waiter already using the dispenser. On the other hand, *social* waiters take more time to serve customers but animate a greater number of chats among friends (see the mean number of chats being animated in column *NChats*). *Egalitarian* waiters look at the tiredness utility function and try to allocate the task to the least tired waiter, therefore, the standard deviation in the number of tasks performed by each agent tends to zero (see column σ_{NTasks}). Finally, *self-interested* waiters demonstrate to perform worse than any kind of social waiter. As this agents are unable to do task passing nor chatting, columns σ_{NTasks} and *NChats* are not considered.

Besides the possibility to define the importance of each point of view through the vector of utility weights \vec{w}_u , MA-DeM allows for the definition of a vector of personal weights \vec{w} that models the attitude of an agent towards the other individuals. Table 3 shows the results obtained for the previously defined waiters using the models of attitude intro-

Coordinated	$\sigma_{Favours}$	<i>Favours</i>
Indifference	7.57	6.9
Reciprocity	1.15	8.8
Altruism	5.94	17
Egoism	1.41	0.7
Social	$\sigma_{Favours}$	<i>Favours</i>
Indifference	3.52	8.7
Reciprocity	1.76	7.8
Altruism	6.66	12.7
Egoism	0.81	0.4
Egalitarian	$\sigma_{Favours}$	<i>Favours</i>
Indifference	7.58	13.6
Reciprocity	2.4	15.5
Altruism	4.44	17.9
Egoism	0.47	0.1

Table 3: Results for different personal weights.

duced in section 3.1. Agents using *indifference* do not apply any modification over the utilities received, therefore, we consider the results of this attitude as the base values to compare with for each type of waiter. *Reciprocity* weights utilities attending to the ratio of favors already done between the agents. This attitude produces equilibrium in the number of favors exchanged as it can be seen in column $\sigma_{Favours}$. Altruism has been implemented in such a way that the weight given to oneself utilities is 0.25 whereas the weights for the rest of the agents is 0.75. As expected, altruist agents do more favors, since the importance given to the other’s opinions is three times the importance given to their own opinion (see high values for the mean number of favors exchanged *Favours*). On the other hand, egoism weights are 0.75 to oneself and 0.25 to the others, thus, agents rarely do favors (see low values in column *Favours*).

Agent’s preferences can sometimes go against personal attitudes. For example, whereas *reciprocity* tries to balance the number of favors, tiredness tends to assign tasks to the least tired waiter (see the greater $\sigma_{Favours}$ for egalitarian waiters). Another example is *egoism* applied to *egalitarian* waiters, in this case no task at all is passed among the agents ($\sigma_{Favours} = 0.1$). However, agent’s preferences can also empower personal attitudes. For instance, *altruism* applied to *coordinated* waiters produces a high level of specialization. This type of agents produces big values for $\sigma_{Favours}$ as the agents already using a dispenser (e.g. a juice machine) keep on getting products from the dispenser following both an altruist and a coordinated behavior that reduces collisions for the use of an exclusive resource. Despite that, personal weights have demonstrated to produce similar effects on the agents regardless of the kind of waiter being considered (i.e. *coordinated*, *social* or *egalitarian*).

Unlike waiters, customers make decisions within an utilitarian society where they consider two points of view: *sociability* and *laziness*. Figure 5 shows the behavior obtained with different types of customers. We compare two metrics: the mean number of social meetings performed among the customers and mean distance covered to consume. Lazy customers, with low utility weights for sociability, most of the time choose to consume at the bar or to sit at a nearby table (see point $\vec{w}_u = \langle 0.3, 0.7 \rangle$). Therefore, the mean distance to the consuming place is short but only a few so-




Figure 5: Lazy vs social customers

cial meetings are animated. On the other hand, social customers, with higher utility weights for sociability, perform more social meetings but they also need to move longer distances to find their friends. Points $\vec{w}_u = \langle 0.6, 0.4 \rangle$ and $\vec{w}_u = \langle 0.7, 0.3 \rangle$ in figure 5 correspond to some examples of social customers.

6. CONCLUSIONS

The decision making approach followed in this paper aims at incorporating human style social reasoning for character animation. According to this, we have presented MADeM as a market based multi-modal agent decision making for social MAS that is able to obtain different kind of coordinated behaviors for the agents involved.

On the one hand, MADeM decisions allow the agents to manage several points of view related to its actions. This social feedback is modelled via utility functions that express the different preferences of each agent for the allocations received for each solution being considered. As an example, two groups of socially intelligent agents have been created that consider different points of view in their decision making: first, a team of waiters using performance, sociability and tiredness; and second, a model of customer that evaluates sociability and laziness.

On the other hand, and besides the possibility of weighting the points of view considered, MADeM provides the agents with the possibility of modelling specific attitude towards the others. The attitudes modelled include indifference (do not modify the other's utility), reciprocity (use the ratio of favors between the agents involved), altruism (the more work for the others the better) or egoism (the more work for me the better).

7. REFERENCES

- [1] N. Badler, C. Philips, and B. Webber. *Simulating Humans: Computer Graphics, Animation and Control*. Oxford University Press, 1993.
- [2] R. H. Bordini, A. C. da Rocha, J. F. Hübner, A. F. Moreira, F. Y. Okuyama, and R. Vieira. MAS-SOC: a social simulation platform based on agent-oriented programming. *Journal of Artificial Societies and Social Simulation*, 8(3), 2005.
- [3] R. H. Bordini and J. F. Hübner. Jason. Available at <http://jason.sourceforge.net/>, March 2007.

- [4] M. Brenner. A multi-agent planning language. In *Proc. of ICAPS'03 Workshop on PDDL*, 2003.
- [5] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaitre, N. Maudet, J. Padget, S. Phelps, J. A. Rodriguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. In *Informatica*, pages 3–31, 2006.
- [6] Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Welfare engineering in practice: On the variety of multiagent resource allocation problems. In *ESAW*, pages 335–347, 2004.
- [7] K. S. Decker and V. R. Lesser. *Readings in Agents*, chapter Designing a family of coordination algorithms. 1997.
- [8] C. Delgado-Mata and R. Aylett. Emotion and action selection: Regulating the collective behaviour of agents in virtual environments. In *AAMAS '04*, pages 1304–1305, 2004.
- [9] J. A. Giampapa and K. Sycara. Team-oriented agent coordination in the RETSINA multi-agent system. Tech. report CMU-RI-TR-02-34, Robotics Institute-Carnegie Mellon University, 2002.
- [10] P. J. Gmytrasiewicz and C. L. Lisetti. Emotions and personality in agent design. In *AAMAS '02*, pages 360–361, 2002.
- [11] F. Grimaldo, M. Lozano, F. Barber., and J. Orduña. Integrating social skills in task-oriented 3D IVA. In *IVA '05*. Springer-Verlag LNAI, 2005.
- [12] L. M. Hogg and N. Jennings. Socially intelligent reasoning for autonomous agents. *IEEE Transactions on System Man and Cybernetics*, 31(5):381–393, 2001.
- [13] R. Prada and A. Paiva. Believable groups of synthetic characters. In *AAMAS '05*, pages 37–43, 2005.
- [14] A. S. Rao. AgentSpeak(L): BDI agents speak out in a logical computable language. In S. Verlag, editor, *Proc. of MAAMAW'96*, number 1038 in LNAI, pages 42–55, 1996.
- [15] A. S. Rao and M. P. Georgeff. Modeling rational agents within a bdi-architecture. In *KR'91*, pages 473–484. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, 1991.
- [16] S. Raupp and D. Thalmann. Hierarchical model for real time simulation of virtual human crowds. *IEEE Transactions on Visualization and Computer Graphics*, 7(2):152–164, 2001.
- [17] C. Reynolds. Flocks, herds and schools: a distributed behavioral model. *Computer Graphics*, 21 (4):25–34, 1987.
- [18] M. Ribeiro, A. C. da Rocha, and R. H. Bordini. A system of exchange values to support social interactions in artificial societies. In ACM, editor, *AAMAS'03*, 2003.
- [19] D. Thalmann and J. Monzani. Behavioural animation of virtual humans: What kind of law and rules? In *Proceedings of Computer Animation*, pages 154–163. IEEE Computer Society Press, 2002.
- [20] X. Tu and D. Terzopoulos. Artificial fishes: physics, locomotion, perception, behavior. In *Proc. of SIGGRAPH'94*, pages 43–50. ACM, 1994.