

## **J-MADeM, a market-based model for complex decision problems\***

**Francisco Grimaldo<sup>1</sup>, Miguel Lozano<sup>1</sup> and Fernando Barber<sup>1</sup>**

<sup>1</sup> *Computer Science Department, University of Valencia,  
Dr. Moliner 50, (Burjassot) Valencia, Spain*

emails: francisco.grimaldo@uv.es, miguel.lozano@uv.es,  
fernando.barber@uv.es

### **Abstract**

This paper presents J-MADeM, a multi-modal decision making mechanism to provide agents in a Multi-Agent Systems (MAS) with a market-based model for complex decision problems. J-MADeM is now available as an open source library fully integrated into Jason, the successful interpreter for the AgentSpeak programming language. The aim of this work is to improve Jason by incorporating an agent decision-making module able to merge multiple information sources received from the rest of the agents. This information is modeled as a set of utility functions expressing the preferences of the agents for a specific problem. Then, J-MADeM agents use one-round sealed-bid combinatorial auctions as the main procedure to choose among different solutions to a decision problem. Thus, coordinated behaviors based on task/object passing can be evaluated to finally obtain acceptable allocations within a synthetic society. J-MADeM is able to simulate different kinds of societies (e.g. elitist, utilitarian, etc), as well as social attitudes of their members (e.g. egoism, altruism, reciprocity, etc). To evaluate the performance obtained by J-MADeM, we analyze two problems: (i) the Gold Miners problem, where a new multi-agent organization is proposed to adapt better to different gold distributions; (ii) a virtual university bar simulation, as an example of complex environment to verify the quality of the social behaviors achieved.

## **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

---

\*This work has been jointly supported by the Spanish MEC and European Commission FEDER funds under grants Consolider Ingenio-2010 CSD2006-00046 and TIN2009-14475-C04-04

deciding its actions. This is specially important in MAS simulations, where collective intelligence appears as the ability of the agents to achieve their goals by interacting with other similarly autonomous entities [8].

Social decision making models are being studied under the scope of MAS in order to regulate the autonomy of self-interested agents. Nowadays, the performance of a MAS is determined not only by the degree of deliberativeness but also by the degree of sociability. In this sense, sociability points to the ability to communicate, cooperate, collaborate, form alliances, coalitions and teams. Social reasoning has been extensively studied in multi-agent systems in order to incorporate social actions to cognitive agents [4]. As a result of these works, agent interaction models have evolved to social networks that try to imitate the social structures found in real life [7]. Social dependence networks allow agents to cooperate or to perform social exchanges attending to their dependence relations (i.e. social dependence/power [11]). Trust networks can define different delegation strategies by means of representing the attitude towards the others through the use of some kind of trust model (e.g. reputation [5]). Lastly, agents in preference networks express their preferences normally using utility functions so that personal attitudes can be represented by the differential utilitarian importance they place on the others' utilities.

Following this preferential approach, the MADeM (Multi-modal Agent Decision Making) model [6] is a market-based mechanism for social decision making, capable of simulating different kinds of social welfares (e.g. elitist, utilitarian, etc.), as well as social attitudes of their members (e.g. egoism, altruism, etc.). 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. MADeM is based on the MARA theory [3] and it uses auctions as a basic procedure to provide the social feedback mentioned.

In this paper we describe how the MADeM model has been integrated into an agent programming language to make socially acceptable decisions available to agents eventually part of an organization. Among several languages for agent programming, we have chosen the AgentSpeak language [10] and its open source interpreter Jason [2] to program this kind of social agents. This choice was made because the language is based on the well known BDI architecture and the interpreter can be easily customised to include the MADeM support. The coupling of MADeM with Jason is inspired in other extensions of Jason, in particular J-MOISE+ [9] and hence the name J-MADeM, as it joins Jason and MADeM.

The rest of the paper is organized as follows: The next section reviews the definition of the MADeM model and describes its decision making procedure. Section 3 explains the J-MADeM architecture that provides Jason agents with the built-in feature of performing MADeM decisions. In section 4 we analyze the performance obtained by J-MADeM in two application examples: (i) the Gold Miners problem, where a new multi-agent organization is proposed to adapt better to different gold distributions; (ii) a virtual university bar simulation, as an example of complex environment to verify the quality of the social behaviors achieved. Finally, in section 5 we state the conclusions of this work.

## 2 The MADeM model

The MADeM model provides agents with a general mechanism to make socially acceptable decisions. In this kind of decisions, the members of an organization are required to express their preferences with regard to the different solutions for a specific decision problem. The whole model is based on the MARA (Multi-Agent Resource Allocation) theory [3], therefore, it represents each one of these solutions as a set of resource allocations. Thus, the definition domain of MADeM is composed by the following elements:

- A set of agents  $\mathcal{A} = \{a_1, \dots, a_n\}$  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.
- A set of resources to be allocated by the agents  $R = \{r_1, \dots, r_m\}$ . MADeM represents resources in the form of  $task(slot)$ , where the  $slot$  is a parameter that needs to be assigned in order to execute the  $task$ . Then, it identifies each one of the solutions for a specific decision problem as an allocation  $P$  of elements (either agents or objects) to task-slots as follows:

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

- 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. 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.
- A set of collective utility functions  $Cuf \in \{elitist, egalitarian, utilitarian, nash\}$ , representing the social welfare of the simulated society, that is, the type of society where agents are located.

MADeM uses a market-based winner determination problem to merge the different preferences being collected according to the kind of agent or society simulated. The details of the whole decision making procedure are explained in the following subsection.

### 2.1 Decision making procedure

MADeM uses one-round sealed-bid combinatorial auctions to choose among different solutions to a decision problem. 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. On the other hand, agents receiving the auction would bid their utility values provided that they were interested in the task being auctioned. Thus, MADeM lies in between centralized and distributed market-based allocation.

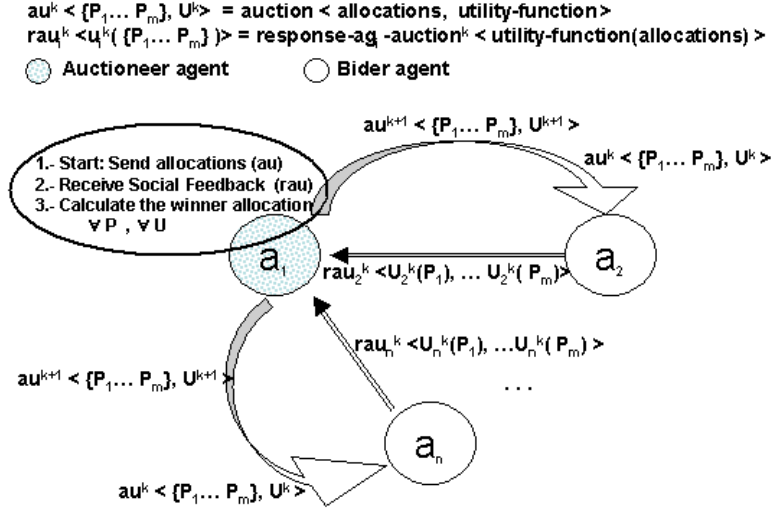


Figure 1: MADeM Procedure

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

1. *Auctioning phase*: This phase is carried out by a single 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 ( $< P_1, P_2, \dots, P_m >$ ). 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 ( $au^k(< P_1, P_2, \dots, P_m >, U^k)$ ). 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  $au^q$ ).
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 ( $rau_i^k = < U_i^k(P_1), \dots, U_i^k(P_m) >$ ).
3. *Winner determination phase*: In this phase, the auctioneer selects a winner allocation for each launched auction. To do this, he uses a classical winner determination problem. Afterwards, he chooses one final winner allocation among these auction winners using a multi-modal decision making process. Thus, the final winner allocation will represent an acceptable decision for the society being simulated. The details of these calculations are fully described in section 2.2.

## 2.2 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 1 groups these utility values in a set of vectors, one for each allocation.

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

Remember that every agent had an associated vector of weights representing its attitude towards the other individuals ( $\vec{w} = \langle w_1, \dots, w_n \rangle$ ). According to it, the auctioneer weighs the utility vectors in equation 1 doing a component by component multiplication with the attitude vector as shown in equation 2.

$$\overrightarrow{U_w(P_j)} = \overrightarrow{U(P_j)} * \vec{w} \quad \forall j \in [1..m] \quad (2)$$

Attitude weights are used to model the social behavior of the auctioneer agent. For example, a whole range of behaviors between egoism and altruism can be modelled using the vector of equation (3), 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.

$$\begin{aligned} \text{Egoism} - \text{Altruism} : \quad \vec{w} &= \langle p, \dots, p, 1 - p, p, \dots, p \rangle \\ w_i &= 1 - p, w_{j \neq i} = p, i = \text{Myself} \end{aligned} \quad (3)$$

It is also possible to model reciprocal attitudes by means of the vector  $\vec{w}$ . A simple example is shown in equation 4, where weights are based on the interchange of favors between agents.

$$\text{Reciprocity} : \quad w_i = \frac{\text{Favors\_from}(i)}{\text{Favors\_to}(i)} \quad (4)$$

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 (equation 5).

$$\text{Winner} = P_w \longleftrightarrow sw(P_w) = \max_{j \in [1..m]} sw(P_j) \quad (5)$$

To compute the social welfare of an allocation, the auctioneer uses Collective Utility Functions (CUFs) and the weighted utilities defined in equation 2 (as shown in equation 6). 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 .

$$sw(P) = \text{cuf}(\overrightarrow{U_w(P)}) \quad \begin{cases} \text{cuf}_{utilitarian} = \Sigma u_w(i) \\ \text{cuf}_{egalitarian} = \min\{u_w(i)\} \\ \text{cuf}_{elitist} = \max\{u_w(i)\} \\ \text{cuf}_{nash} = \Pi u_w(i) \end{cases} \quad (6)$$

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 7).

$$au_1(\{P_1, \dots, P_m\}, U^1) \rightarrow (P_{w1}, sw(P_{w1})), \dots, au_k(\{P_1, \dots, P_m\}, U^k) \rightarrow (P_{wk}, sw(P_{wk})) \quad (7)$$

Lastly, the MADeM final winner allocation is that which maximizes the welfare of the society after having multiplied it by the corresponding utility weight  $\bar{w}_u$  (see equation 8),

$$P_w = P_{wi} \longleftrightarrow sw(P_{wi}) = \max_{i \in [1..k]} w_u(i) * sw(P_{wi}) \quad (8)$$

### 3 J-MADeM architecture

This section describes how the MADeM model [6] has been integrated into Jason [2] as an open source library named J-MADeM. The J-MADeM is built upon the *Jason Communication Infrastructure*, thus extending the communication level options available in Jason with a set of modules that provide the agents with the built-in feature of performing MADeM decisions. Figure 3a illustrates how these components are integrated into Jason. The J-MADeM basically offers to the AgentSpeak programmer: (i) an agent architecture that Jason agents can use to carry out their own MADeM decisions, (ii) an interface to develop utility functions that can be used along with the MADeM model and (iii) a set of internal actions to manage the parameters of these kinds of decisions.

The *J-MADeM Agent Architecture* extends the *Jason Agent Architecture* in order to incorporate all the necessary modules that allow MADeM decisions to be automatically carried out. The main components of the *J-MADeM Agent Architecture* are shown in the figure 3b, where we can identify the following elements:

- *MADeM Parameters* : This data storage contains the MADeM context currently defined for the agent. Essentially, it stores the personal weights, the utility weights, the collective utility function and the bid timeout to be used in future MADeM decisions.
- *Decision Launcher* : This module starts the MADeM process for a particular decision. Firstly, it stores the MADeM context for this decision into the *Decision Data* storage, thus allowing other decisions to be concurrently performed with different MADeM parameters. Secondly, it auctions each of the allocations being considered as solutions to the target agents.

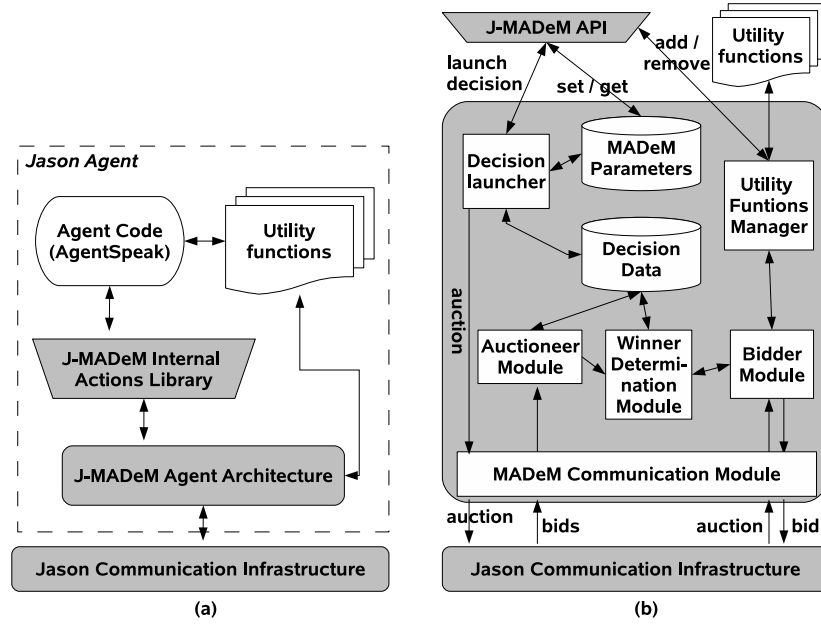


Figure 2: (a) Overview of the J-MADeM architecture and (b) detailed view of the J-MADeM Agent Architecture.

- *Decision Data* : This data storage holds all the information related to the MADeM decisions still in process. Therefore, it contains their MADeM context, their considered allocations and the preferences received for each of them.
- *MADeM Communication Module* : This module extends the Jason agent communication module in order to deal with MADeM messages. When it receives a MADeM auction, it invokes the *Bidder Module* to get the agent's preferences over the considered allocations. On the other hand, when it receives a MADeM bid, it informs the *Auctioneer Module* about the received preferences.
- *Bidder Module* : This module manages the reception of a MADeM auction. It extracts the considered allocations and bids for them according to the agent's preferences. To express these preferences it relies on the utility values provided by the *Utility Functions Manager*.
- *Utility Functions Manager* : This component acts as an interface between the built-in MADeM mechanism and the user defined *Utility Functions*. Thus, it is in charge of locating and invoking them in order to calculate the agents' utilities for the set of considered allocations.
- *Auctioneer Module* : This module manages the reception of MADeM bids. It extracts the sender's preferences and stores them into the *Decision Data*. As soon as the preferences from all the target agents have been received, it calls the *Winner Determination Module* to solve the decision.

- *Winner Determination Module* : This module solves the MADeM winner determination problem using the information stored into the *Decision Data* for the decision being resolved (i.e. considered allocations, agents' preferences, personal weights, utility weights, social welfare,...). Once resolved, it notifies the agent about the winner solution.

For a complete description of the utility function interface and the set of internal actions refer to the J-MADeM documentation available at the Jason website [1].

## 4 Experiments and Results

This section summarizes the results obtained in two application examples developed to test J-MADeM agents. First, we revisit the Gold Miners problem [2], a classical simulation scenario where agents must compete for the resources (gold) located at the environment. This example allows us to evaluate the efficiency of the auction-based method proposed and to experiment with dynamic organizations. Second, in order to test the sociability features provided by J-MADeM, we have created a virtual university bar simulation. In this scenario waiter agents serve the orders placed by customer agents. According to the model parameters and the society being simulated, waiters are able to combine social behaviors (i.e. chatting) and efficiency at work.

### 4.1 The Gold Miners

In this example a team of gold-mining agents has to find a set of chunks of gold, randomly scattered in a grid-like territory, in order to carry them to a depot. The original Jason team consists of a leader that assigns each miner to a quadrant of the grid. The miners then explore and pick up the pieces of gold they find in the environment. To be more efficient, miners can use an auction-based model to inform the others about new gold locations.

Following the original Jason implementation, we have created an equivalent MAS with J-MADeM agents to be able to compare both. We have implemented an utility function, equivalent to the distance function used by the original miners, so that J-MADeM miners can express their preferences with regard to the gold units. Then, any agent can launch a J-MADeM decision as follows:

```
jmadem.launch_decision(LAgents, A1, [goldDistance], DecisionId);
```

where **LAgents** is the list containing all the agents apart from the announcer, since the original Jason implementation uses a broadcast auction model. A full version of this example can be downloaded from the Jason website [1].

We first analyze the original Jason gold-miners performance to estimate the importance of the original auctioning process and to perform a fair comparison. To evaluate the efficiency of the auction model, we have measured the number of gold chunks picked up as a result of an auction. We identify this situation as a successful auction, as it



Gold Density(total golds)	0.05(51)	0.1(102)	0.15(153)	0.2(204)
4 agents	9,3	23,5	33,2	37,2
8 agents	10,6	28,4	47,4	58,3
12 agents	10,8	26,2	41,2	56
16 agents	10,4	21,5	33,4	53,6

Table 1: Successful auctions in the original Jason Gold Miners implementation.

Gold Density(total golds)	0.05(51)	0.1(102)	0.15(153)	0.2(204)
4 agents	24,4	40,9	69,5	75,8
8 agents	31,5	67,9	108,6	109,3
12 agents	30,9	70,6	98,6	121,1
16 agents	33,6	71,7	107,2	131,2

Table 2: Successful auctions in the J-MADeM Gold Miners implementation.

is planned and completely executed. Table 1 shows the average of succesful auctions obtained over 10 simulations with different number of agents and gold densities. To calculate the number of gold chunks involved in each simulation we just multiply the gold density by the grid size (in these results wold size is 32x32). Suprisingly, table 1 indicates that only the 25% of the gold are captured as a result of an auctioning process. To the authors knowledge, the reason behind this behavior is that the centralized auction model implemented requires the leader to be able to manage parallel auctions during the simulation. However, the leader fails when dealing with concurrent auctions of the same pieces of gold.

On the other hand, table 2 shows the number of successful auctions obtained by the equivalent J-MADeM gold-miners implementation. J-MADeM agents offer a robust auction controller able to handle different auctions at the same time. Thus, this feature allows the J-MADeM miners to solve the same problem but with a higher number of succesful auctions, around 65%. We notice that some planned allocations can still not be completed, since new free agents might happen to pick up gold chunks previously auctioned.

To evaluate global performance estimators for this problem, we have compared the total number of steps taken by miners in both implementations. This experiment shows a similar efficiency for both approaches regardless of the number of agents and the gold density. This similarity also appears in other variables we have mesured, such as the simulation time or the standard deviation of the pieces of gold picked by the agents.

Finally, to test the flexibility of J-MADeM, we have changed the initial uniform gold distribution. Now, the gold chunks will be concentrated in a single quadrant of the grid, thus representing a non-uniform gold distribution from a global perspective. Besides, we have modelled a new dynamic multi-agent organization to face this new

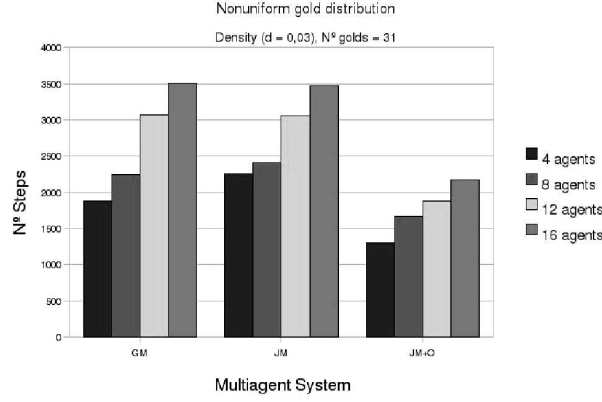


Figure 3: The J-Madem agent architecture.

situation while trying to increase the performance previously obtained. The new MAS has been organized in two roles: miners and bosses. Firstly, the new miner agents are similar to the original Jason miners but they modify the announcing process. That is, instead of broadcasting gold locations to the rest of the agents, they just inform their corresponding boss. Secondly, the bosses are new J-MADeM agents devoted to organize the work of a group of miners within a quadrant of the environment. The bosses are not allowed to directly pick up pieces of gold but they can allocate them to its subordinated miners. Initially a balanced miner-boss assignment is performed, so that all bosses manage the same number of miners. However, they can dynamically change the organization by borrowing miners from other bosses when the number of gold chunks found in the quadrant they control increases. The altruist bosses modelled will always prefer to lend one of their miners than to keep it idle (i.e. without utility). Therefore, the bosses launch MADeM decisions both to allocate golds to miners and to ask other bosses for additional miners. For this new social decision, the bosses will use an utility function that returns the number of free miners to express their preference, thus trying to keep the teams well balanced.

Figure 4.1 compares the total number of steps followed by the three gold-miners implementations considered. The first group corresponds to the performance obtained by the original Jason implementation (GM) when facing environments with non-uniform gold distribution. Then we have included the equivalent J-MADeM implementation (JM) and finally the new dynamic organization (JM+O). The figure shows how the new organization obtains higher performance regardless of the number of miners being simulated. Similar results have been also obtained for the simulation time, where the new organization carries out the same work but reducing significantly the required time.

## 4.2 Virtual university bar

This social scenario represents a virtual university bar where waiters take orders placed by customers. Both waiters and customers carry out tasks in the virtual bar and they use J-MADeM to decide among different task slot assignments. Next, we summarize

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

Table 3: Results for different personal weights.

the main results obtained. For a full description of the simulated scenario see [6].

Table 3 shows the results obtained for three kinds of waiters (i.e. *coordinated*, *social* or *egalitarian*) using different models of attitude. For instance, 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  $\sigma_{Favors}$ . 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  $\overline{Favors}$ ). 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  $\overline{Favors}$ ).

Agent’s preferences can sometimes go against personal attitudes. For example, *egoism* applied to *egalitarian* waiters produces that no task at all is passed among the agents ( $\overline{Favors} = 0.1$ ). 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  $\sigma_{Favors}$  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. In spite of that, personal weights have demonstrated to produce similar effects on the agents regardless of the kind of waiter being considered.

## 5 Conclusions

This paper has presented J-MADeM, a new open source library oriented to create different types of social simulation agents. J-MADeM agents are able to merge several points of view received from other agents. This social feedback is modelled via utility functions that express the preferences of each agent for every solution being considered. The J-MADeM architecture and its integration into Jason have been reviewed. The application examples presented demonstrates the robustness and flexibility to create complex market-based social simulations. First, the Gold Miners example can

use J-MADEm to maximize the efficiency of the MAS. Second, the virtual university bar example offers a more complex environment where several utilities (personal and collective) and weights (personal and utility-based) allow to design elaborated social simulations.

## References

- [1] R. H. Bordini and J. F. Hübner. Jason. Available at <http://jason.sourceforge.net/>, March 2007.
- [2] R. H. Bordini, J. F. Hübner, and M. Wooldridge. *Programming Multi-Agent Systems in AgentSpeak using Jason*. Wiley, 2007.
- [3] 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. *Informatica*, 30:3–31, 2006.
- [4] R. Conte and C. Castelfranchi. *Cognitive and Social Action*. UCL Press, 1995.
- [5] R. Falcone, G. Pezzulo, C. Castelfranchi, and G. Calvi. Why a cognitive trustier performs better: Simulating trust-based contract nets. In *Proc. of AAMAS'04: Autonomous Agents and Multi-Agent Systems*, pages 1392–1393. ACM, 2004.
- [6] F. Grimaldo, M. Lozano, and F. Barber. MADEm: a multi-modal decision making for social MAS. In *Proc. of AAMAS'08: Autonomous Agents and Multi-Agent Systems*, pages 183–190. ACM, 2008.
- [7] H. Hexmoor. From inter-agents to groups. In *Proc. of ISAI'01: International Symposium on Artificial Intelligence*, 2001.
- [8] L. M. Hogg and N. Jennings. Socially intelligent reasoning for autonomous agents. *IEEE Transactions on System Man and Cybernetics*, 31(5):381–393, 2001.
- [9] J. F. Hübner, J. S. Sichman, and O. Boissier. Developing organised multi-agent systems using the Moise+ model: Programming issues at the system and agent levels. *Journal of Agent-Oriented Software Engineering*, 1(3/4):370–395, 2007.
- [10] A. S. Rao. AgentSpeak(L): BDI agents speak out in a logical computable language.
- [11] J. Sichman and Y. Demazeau. On social reasoning in multi-agent systems. *Revista Ibero-Americana de Inteligencia Artificial*, 13:68–84, 2001.