

# Multi-agent negotiation to support an economy for online help and tutoring

Chhaya Mudgal and Julita Vassileva

University of Saskatchewan  
Computer Science Department  
{chm906, jiv}@cs.usask.ca

## ABSTRACT

We are designing a computational architecture for a "learning economy" based on personal software agents who represent users in a virtual society and assist them in finding learning resources and peer help. In order to motivate users to participate, to share their experience, offer help and create on-line learning resources, payment is involved in virtual currency and the agents negotiate for services and prices, as in a free market. We model negotiation among personal agents by means of an influence diagram, a decision theoretic tool. In addition, agents create models of their opponents<sup>1</sup> during negotiation to predict opponent actions. Simulations have been carried out to test the effectiveness of the negotiation mechanism. An attempt has been made to test the benefits of the proposed economy as a basis for the peer help environment, I-Help.

**Keywords** multi-agent architecture, negotiation, peer-help system, distributed educational systems, learning economy, personal agents, motivation

## 1. INTRODUCTION

The Internet provides a variety of options for on-line training, tutoring and help, from access to FAQs and multi-media teaching materials in the form of web-pages, to more interactive forms like discussion forums, on-line tutoring, collaboration or peer-help sessions. The creation of high quality teaching materials is associated with significant costs, which usually have to be paid by those who benefit directly from them, i.e. the learners. There is a potential for a rapidly growing market of on-line training and there has been a significant increase in the number of commercial vendors in this area. A number of universities are already offering on-line degrees, and charge significant fees (still, somewhat lower than the costs of traditional university education).

However, still the most on-line training materials appear informally, collaboration and help happen spontaneously. University lecturers post their course outlines, lecture notes and course readings / materials on-line as an additional source of information for their students. People facing problems in a certain area search for a newsgroup related to the area and send their question there, hoping for someone more competent than themselves to read it and answer it. People ask their colleagues, personal acquaintances and friends for help. This is a huge pool of knowledge and

expertise, which is not formally valued in organizational or commercial form and which is used only randomly, occasionally and scarcely. Our goal is to provide an infrastructure that motivates the usage of this knowledge. We hope to achieve this by creating a marketplace for learning resources, i.e. an e-commerce environment for trading with intangible goods (advice, help, teaching or tutoring). This economy encompasses information exchange, which happens both in asynchronous and in a synchronous way. For example, the use of on-line resources like web-pages, FAQ entries, or the use of e-mail to ask a question and provide advice can be viewed as asynchronous information exchange, since they don't imply that both sides (the learner and the helper/ tutor are present and involved in interaction in the same time). Synchronous information exchange involves both sides in a real-time, live contact -- for example, in an on-line help session via some chat tool, telephone line, or collaboration environment.

The basic assumption in the design of a learning economy model is that resources like effort and time spent to provide help or to create teaching material have inherent costs. To take them into account, these resources should be made tradable. Thus paying the helper/tutor may motivate a user to get online and help another user. In this paper we focus on a synchronous information exchange since it is related with more immediate motivational need. However, the approach encompasses asynchronous information exchange too.

Maes et al. [6] proposed to help consumers in e-commerce applications in the search of goods, price comparison and negotiation or bidding by providing them with personal agents / assistants. We believe that this is even more important in trading with knowledge resources, since users have to be able to concentrate on their work or learning rather than thinking about how to get a better deal. The free market infrastructure for learning resources that we propose is based on personal agents representing individual users in a distributed (web-based) learning environment. The personal agents form an economic society designed to motivate the students who are knowledgeable to help their fellow students by receiving payment in a cyber pseudo currency.

## 2. MULTI-AGENT BASED LEARNING ECONOMY

I-Help provides a student of a university course with a matchmaking service to find a peer-student online who can

---

<sup>1</sup> We will use the word "opponent" to denote the other agent in negotiation, though we don't imply necessarily an adversary or strongly competitive negotiation.

help with a given question/problem [3,4]. The most recent implementation of I-Help is based on Multi AGent Architecture for Adaptive Learning Environment (MAGALE<sup>2</sup>), described in [11,12], which ensures an economic infrastructure for trading with help. MAGALE is a society of agents trading with knowledge-resources. The users who possess knowledge resources become sellers and the users who seek for help or advice, tutoring or teaching materials on a specific topic become buyers. The buyer is ready to pay some amount of virtual (or real) currency in order to achieve the goal of getting knowledge while the seller of the resources is ready to give advice in exchange for money, thus achieving the goal of accumulating currency. Like any market system, in MAGALE (and respectively in its implementation, I-Help) the price of a good depends on the demand and the importance of that good to the buyer. A more detailed description about the economic model in MAGALE and the requirements for the economy in I-Help can be found in [5].

Various pricing models have been incorporated in e-commerce systems. The most common are "post and charge", "pay-per-use" and "auction". "Post and charge" is applied in I-Help for paying for asynchronous resources, such as web materials, FAQ items, or answers in a discussion forum. One can post an answer to a question in I-Help's discussion forum and people who read it would be charged to pay a certain price. A similar model is implemented in the Marketplace for Java Technology Support [10], a community where people buy and sell technical support (the forum is operated by HotDispatch, Inc).

The "pay-per-use" model implies paying a certain rate for a unit of usage time of the resource, for example paying for a telephone call. This can be an appropriate mechanism when the duration of the service is connected with costs and it can not be fixed or agreed upon in advance. This is an appropriate model of payment for the various forms of synchronous knowledge transfer that are supported in I-Help (chat, phone-communication or collaboration). The duration of a help session implies costs to the helper who is asked to interrupt some current task. It is hard to say in advance what duration will be required since it depends on the question, on the ability of the helper to explain, and on the helpee's ability to understand. Therefore, it is appropriate to deploy this payment method in synchronous help allowing both sides to interrupt the session when they feel that it doesn't make sense for them to continue.

The "auction" model, where several agents are bidding for goods [6] is appropriate when there is a big demand and short supply. It allows the resource to be allocated to a consumer who values it most. This could be an appropriate model in the case where synchronous information exchange (e.g. help request) is required by many users and there are

few knowledgeable users on-line to provide help. This model has not been applied in I-Help yet, but it could be.

The auction model is, in fact, a way of collective negotiation of the price for a resource, where the main factors that determine the price are the demand and the supply. The other two models don't imply per se a mechanism for determining the price - they assume that there is a price that is agreed upon in advance. The price can be established centrally by a component that analyses the state of the market at the moment or it can be negotiated between the agents who participate in the deal [13]. The advantage of negotiation is that it allows for including multiple factors to dynamically calculate the price depending on the specific buyer and seller, i.e. the agents can settle on the most suitable price for both parties.

The price of a learning resource depends on many factors. Of course, the supply and demand (e.g. how many competent helpers are currently on line and how many people are requesting help) play a major role. However, many other factors can play a role, for example, whether the help is urgently needed or not, whether the potential helper minds being interrupted, whether the helper and the person asking for help (the helpee) have some social connection. For example, the helper might not want to be interrupted in principle, but would make an exception for a friend. Therefore, a negotiation mechanism is appropriate as a way to dynamically determine the price, especially for synchronous information exchange.

We have proposed a negotiation mechanism for the personal agents in MAGALE that determines the price for synchronous information exchange (e.g. on-line peer help in I-Help) using the "pay per use" payment model. This mechanism mimics the process of human negotiation in a buyer-seller situation, by representing it as an iterative decision making process. It also allows the negotiator to anticipate the opposing party's actions and takes into account the personal risk attitude towards money of the user represented by the agent. The purpose of negotiation is to find the best deal for the user independently on whether the user requires help or is playing the role of a helper.

### 3. NEGOTIATION MECHANISM

The MAGALE architecture underlying I-Help consists of personal agents representing the users/ students. The agents maintain user models containing information about the user's goals, knowledge and preferences [3]. When the students in the class need help their agents contact a centralized matchmaker who knows which other agents are online. These agents negotiate with each other about the price (in our case this is the payment rate per unit of help time) and when a deal is made they inform their user. If the user agrees to help, a chat window opens for both sides and the help session starts. Thus, in I-Help the agents make decisions on behalf of their users about the price to offer

---

<sup>2</sup> The name MAGALE is introduced to distinguish the more general architecture from I-Help, which is an application.

and how to increase or decrease the price to strike a better deal depending on user specified constraints, such as the urgency of the user's current work, importance of money to the user and the user's risk behavior.

### 3.1 Decision Theoretic Approach to Negotiation

We have developed a novel negotiation approach, using influence diagrams and based on decision theory and on modelling the opponent agent. Negotiation in a buyer-seller context can be viewed an iterative process in which the agents make offers and counteroffers based on their preferences. Modelling negotiation as iterative decision making supports the dynamics of the situation, e.g. it allows the negotiating agents to change their preferences and their beliefs about the likelihood of uncertainties.

In open multi-agent systems (i.e. the systems in which new agents dynamically enter or leave) there is a high degree of uncertainty about the current state of the market (i.e. the demand/supply ratio), or the preferences of the opponent. An influence diagram is a graphical structure for modelling uncertain variables and decisions. It explicitly reveals probabilistic dependence and the flow of information [8].

An influence diagram is a directed acyclic graph with three different kinds of nodes: decision nodes, chance nodes and a value node. These nodes are represented as squares, circles, and diamonds respectively. The decision nodes represent choices available to the user, the chance nodes carry probabilistic information corresponding to the uncertainty about the environment and the opponent, and the value node represents the utility, which the agent wants to maximize. Arcs into random variables indicate probabilistic dependence and the arcs into a decision node specify the information available at the time of making decision. Evaluating the diagram gives an optimal solution for the problem. Influence diagrams provide a means to capture the nature of the problem, identify important objectives, and generate alternative courses of action. A decision model based on an influence diagram can deal with multiple objectives and allows tradeoffs of benefits in one area against costs in another. A good introduction to influence diagrams and methods to evaluate them can be found in [8,9].

The negotiation protocol is based on decision theory and is a straightforward iterative process of making offers and counteroffers. So, during negotiation the agent can be in *Offer* or *Counter-offer* state repeatedly. The final state will be *Accept* or *Reject*. Similar to [13], we use "negotiation strategy" to denote the actions an agent takes in every iteration depending on its preference model. In our model once the agent is in a final state, it cannot retreat back from it. The negotiation mechanism takes into account the preferences of the user, which usually depend in the domain of the negotiation context. The preferences include:

- the *maximum price* of the buyer (i.e. how much the helpee is willing to pay),

- the *urgency* of the current goal (to get help for the buyer, or the seller's current work, which she has to interrupt in order to help),
- the *importance* that either agent attaches to *money*, and
- the user's risk behavior (risk-averse or a risk-seeking person).

We have incorporated utility to model the way in which the decision-maker values different outcomes and objectives. Each agent in I-Help can be a buyer or seller of help. The utility for the buyer (helpee) and the seller (helper) for the actions *accept*, *reject* and *counter-propose* vary according to their risk behavior.

It is important to note that money importance and risk-behavior are two different entities and they are set by the user in the user preference model. The risk behavior of the user instructs the personal agent about the increase or decrease in the price offers to be made. A risk-seeking person will try to counter-propose an offer rather than accepting. A risk-averse person will accept whatever minimum price he/she is offered and will refrain from counter proposing in fear of losing. The agent calculates the utility values of the action alternatives that it has at any time during negotiation. The utility of actions depends upon the money that the seller gets and the buyer has to pay. It also varies with the specified risk behavior of the user. For instance, as shown in the Figure 1 the utility of accepting an offer for a risk-averse buyer increases much slower as the difference between the offered price and the preferred price decreases. That means that as long as the offer price of the seller comes closer to the preferred price of the agent (buyer), it will be more willing to accept the offer, since there is not significant growth in utility if it continues to counter-propose. For a risk-seeking agent, the utility continues to grow fast in this case, since it is willing to take the risk of counter-proposing, hoping to get a price even lower than the preferred price.




Figure 1: Variation of  $U_{accept}$  for a **buyer**

Risk behavior also affects the increment and the decrement of the buyer and the seller. For a risk-averse buyer, if the urgency of the current task is very high and the importance of money is also high, it will start by offering a price, which is near to the maximum price it is willing to pay. A risk-seeking buyer will start from a very low price and will try to get the lowest price possible. For a risk-seeking seller the utility of accepting an offer increases if it gets more

money than its minimum price. The functions the agents use to increase or decrease their offers and counteroffers as a buyer and as a seller are defined as follows:

**For Buyers**

If  $\text{max\_price} > \text{std\_price}$  then  
 Offered price =  $\text{std\_price} - \Delta$   
 Else  
 Offered price =  $\text{max\_price} - \Delta$

**For Sellers**

If  $\text{min\_price} > \text{std\_price}$  then  
 Offered price =  $\text{min\_price} + \Delta$   
 Else  
 Offered price =  $\text{std\_price} + \Delta$

where  $\text{std\_price}$  is the market price provided by the matchmaker. It is calculated based on the current situation of the market of help on this topic and on the difficulty of the topic, thus providing some measure for actual worth of the resource.

$\Delta$  is determined as follows:

**For Buyers**

If  $\text{urgency} == \text{very urgent}$  then  
 If  $\text{risk\_behavior} == \text{risk seeking}$  then  
 $\Delta = 1 - e^{-x/R}$       $x > R$   
 If  $\text{risk\_behavior} == \text{risk averse}$  then  
 $\Delta = 1 - e^{-x/R}$       $x < R$

Where  $R$  is the preferred price and  $x$  is the offered price.

**For Sellers**

If  $\text{urgency} == \text{very urgent}$  then  
 If  $\text{risk\_behavior} == \text{risk seeking}$  then  
 $\Delta = \sqrt{\text{min\_price}}$   
 If  $\text{risk\_behavior} == \text{risk averse}$  then  
 $\Delta = \log(\text{min\_price})$

For both the buyer and the seller the values of  $\Delta$  should not exceed the preferred prices.

We use an influence diagram that has a conditional node representing the uncertainty about the other party (see Figure 2). The outcomes of this node are the probabilities that an opponent can be in any of the states *accept*, *reject* and *counter-offer*, because at every step the agents have to choose between these three actions. They do so by calculating the maximum expected utility for the actions, which are represented as the possible choices for the decision node in the influence diagram. In any practical application of negotiation there are multiple objectives involved and there has to be tradeoff between one over the other. Before the decision is made the factors that are already known to affect the decision (deterministic nodes) are taken into account as they affect the actions to be made. The node corresponding to the opponent's action can be considered conditional since nothing is known about the opponent at the beginning of the negotiation. We can either treat the outcomes of the opponent node as equally likely or replace the equal likelihood of the opponent's actions with the outcome of a model of the opponent using a probabilistic influence diagram.




Figure 2: Influence Diagram for the decision model

**3.2 Modeling the Opponent**

One of the basic ingredients of a negotiation process is the correct anticipation of the other side's actions. In a dynamic environment like a market place where the situation is changing all the time and new buyers and sellers keep on entering and leaving the system, it is very costly for agents to create and maintain models of the other participants in the environment. In the I-Help system the environment is dynamic and since the agents represent real users, it is hard to predict the actions of the opponent agent on the basis of its past behavior (since the user's preferences which participate in the agent's negotiation strategy can change in the meantime). It is unlikely that the user will be willing to share preferences with other users (or their agents) before or during the negotiation process. However, it is useful for an agent to model the opponent's behavior during the negotiation session, since this can help predict the opponent's reaction. It is important to note that we are not doing recursive or nested agent modeling. Agents initially have no knowledge about each other. After the first round of offers made the agent starts using the opponent's response to infer a model of the opponent's preferences and to predict the possible reaction of the opponent to the counteroffer that the agent is about to make.




Figure 3: Probabilistic influence diagram representing the opponent's model

An appropriate tool for this purpose is a probabilistic influence diagram. Figure 3 shows the model of the opponent represented as a probabilistic influence diagram. The oval nodes are conditional and the double-circled node

is deterministic. The conditional probability distribution of the conditional nodes over the outcomes is assessed on the basis of the first offer. The probability distribution for the "Opponent's action" node can be calculated by performing reductions over the nodes. For instance, performing arc reversal from the "Money Importance" node to the "Opponent's Action" node makes "Money Importance" a barren node. Hence, it can be removed from the diagram and a new conditional probability distribution is calculated. Conditional predecessors of the nodes (if any) are inherited. In a similar way the diagram can again be simplified by using arc reversal and barren node removal, which finally gives the probability distribution for the Opponent's Action node. If the next move of the opponent does not match with the predicted action, Bayes' update rule is used to update the probability distributions. For more information about probabilistic influence diagrams refer to [9].

#### **4. EVALUATION**

We evaluated the proposed negotiation mechanism in an environment, where agents represented only themselves, i.e. no real users were involved. In this way we were free to vary the negotiation parameters and generate a lot of experimental data. The purpose was to evaluate the results of the decision method only.

The results [7] showed that the proposed negotiation approach achieves a better deal for the agent that uses it compared to other negotiation approaches, for example, one based on decreasing (for seller) / increasing (for buyer) of the offered price. We carried out further experiment, which showed that if the agents are bluffing, i.e. offering help at much higher price than their preferred price, the acceptance percentage of their negotiation is low. Agents who are more reasonable get a good deal maximum number of times.

In order to evaluate the principal usefulness of an economic model to motivate users a version of I-Help was developed, which was used as a basis for comparison for the negotiation approaches. This system was applied in a 3rd year undergraduate computer science class at the University of Saskatchewan. Initially there seemed to be an enthusiasm among the students about the system, however, consequently there turned out to be very little usage, which didn't allow us to draw any conclusions about the efficiency of the economy or the planned control measures. There were several different reasons for this failure. The most important one is the slow response time of the system, especially off campus, due to slow network connections during this period. It must be pointed out that the slow response was completely due to reasons independent on the implementation of the system or the negotiation mechanism. A second reason might have been an inappropriate interface design, which made interaction with the personal agent somewhat cumbersome. Finally, the lack of usage can also be attributed to the fact that the

3<sup>rd</sup> year students knew each other very well, had established multiple ways of interacting with one another in class and in the labs and hence they did not find any need to login to the system to get help. The reasons for this choice were purely pragmatic: time until semester begin was short and this implementation required the least adaptation effort, for the domain representation and student modelling were already developed.

Yet, our inability to obtain strong (whether positive or negative) evaluation results taught us a good lesson: that introducing such advanced mechanisms makes sense only when the basic technology works reliably (with respect to network speed, response time and user interface design). Another lesson we learned is that the right user group should be selected very carefully before trying to test and evaluate such system. We hope that if the proposed market economy model is utilized in distance learning or a very large first year class where students don't know each other and have no other incentives to be helpful to each other, it will prove to be successful. Our next test is in a large introductory computer science class.

This negative experience also shows that there are sometimes unexpected difficulties in testing such complex distributed multi-agent systems, due to very basic "low-level" problems, completely unrelated to the proposed technology. It seems that new evaluation methodologies are needed which would allow evaluation without the need of developing of stable nearly ready for marketing system.

#### **5. RELATED WORK AND DISCUSSION**

To our best knowledge, there is currently no other work in the area of market economy based distributed systems that support human learning. Developing a learning economy has been made [1], but it was based on the barter (exchange) model and has not been implemented. IBM has proposed an economy for trading information resources [2], however this proposal assumes that the resources are ready documents and it focuses mainly on pricing models that are appropriate for them. The most closely related work to ours is in the field of multi-agent negotiation in e-commerce [13]. In [13] negotiation and modelling the opponent is realized by using a Bayesian network where the agents have all the relevant information about each other, while in our approach negotiation is modelled as a decision process. In addition, our agents don't share information about each other's priorities and model each other to predict the actions of the opponents and thus to optimize their decisions.

Our approach opens some interesting research avenues in student / user modelling to be pursued further. There are multiple models about each user in the system. They are created by different agents, contain different (but also sometimes overlapping) information, are created under different circumstances.

First, there is a centralized representation of a collection of user models in the system, which consists of some social

features (eagerness, helpfulness), as well as the knowledge level of the user. The matchmaker uses these models to select a best helper for the topic of the request.

Second, to enable agents to make decisions according to the preferences of their users each agent consults a model of its user's negotiation preferences. These preferences include starting price, risk attitude, current goal urgency, list of friends / enemies. This "negotiation" user model is initialized by the user and can be updated by the user, after receiving feedback from the agent about the success rate in negotiation.

Third, each personal agent makes a model of its opponent agent i.e. the negotiation preferences of the user represented by the opponent agent

So at any time, three different types of models of each user exist, including multiple versions of the user's negotiation preferences, as developed by each agent who is negotiating with this user's agent. Therefore, the following questions arise:

1) Are these models consistent with each other? Is it desirable/necessary that they should be consistent? How to aggregate information from these models for different purposes?

2) How to handle discrepancies between the model of negotiation preferences and the qualities of the resources offered by the user, i.e. how to avoid people with weak abilities charging too much for their services?

3) Should the agents be able to access the model of the opponent as maintained by the system? How can the personal agent verify the preconditions for good-quality service for its user?

More research on these issues will help to find the benefits and pitfalls of distributed user modelling.

On the other side, more research is needed on analyzing the global behavior of a system based on individual negotiations between agents, like ours. Especially in an educational system, it is very important to predict and be able to control the overall behaviour that emerges as a result of interaction of personal agents and users. An economic model should show how to effectively distribute learning resources. We have proposed an economic model [5], which provides a wide variety of options to control the economy from outside to stimulate desirable behavior of the system. For example, since the dynamics of the system depends very much on the availability of the resource at the time of need and the urgency of the resource demand, it is possible for a teacher to "inject" from outside more resources of the type that are in high demand (e.g. post a web-page with explanation on a newsgroup). Since the MAGALE architecture provides for application agents, (representing software applications) or individual information resources (like web-pages), it is possible to throw such resources on the market with application agents that are selling them for a cheaper price (within a "pay per use" paying model), thus satisfying the demand. This can also stimulate users to create their own web-resources. For

example, if a user has been asked several times to answer the same question, she might prefer to create a web-page with an explanation and attach to it an agent who is going to charge users reading the page some small amount, to benefit the user. However, it will be hard to design an experiment to test the benefit of these measures, since the system is very complex - so many factors come to play, that it is hard to attribute success or failure even to a group of factors. New methods, possibly borrowed from sociology, will be needed to evaluate such systems.

## 6. CONCLUSION

We have developed an original approach for negotiation among personal agents based on decision theory and influence diagrams. By use of probabilistic influence diagrams agents are able to model their opponents during the negotiation process and thus to predict better their actions. Experiments on a simulation showed the effectiveness of the proposed negotiation mechanism [7]. An attempt has been made to evaluate the benefits of the proposed economy as a basis for the peer help environment I-Help in a third level university class. Our experience showed that such experiments have to be designed very carefully to keep complexity and technical issues under control and in the same time to be able to answer some interesting research questions. Probably new evaluation methodologies for distributed agent based systems on the Internet will be necessary.

**Acknowledgement.** This work has evolved in discussions with our colleagues from the ARIES Group at the University of Saskatchewan. It has been partially funded by the Telelearning NCE under Project No. 6.28.

## 7. REFERENCES

- [1] Boyd, G. 1997. Providing Real Learning with Virtual Currency. Proceedings of the International Conference on Distance Education, Penn State University, June 1997.
- [2] Greenwald A. and J. Kephart 1999. Shopbots and Pricebots. in Proceedings of IJCAI '99. Stockholm, July 31 - August 6, 1999. also available on line at: <http://www.research.ibm.com/infoecon/researchpapers.html>
- [3] Greer, J., McCalla, G., Cook, J., Collins, J., Kumar, V., Bishop, A. and Vassileva, J. (1998) The Intelligent HelpDesk: Supporting Peer Help in a University Course, Proceedings ITS'98, San Antonio, Texas, 494-503.
- [4] Greer, J. McCalla G., Collins J., Kumar V., Meagher P., Vassileva J. (1998) Supporting Peer Help and Collaboration in Distributed Workplace Environments, *International Journal of AI and Education*, 9.
- [5] Kostuik, K., Vassileva, J. Free Market Control for a Multi-Agent Based Peer Help Environment. Workshop on Agents for Electronic Commerce and Managing the Internet-Enabled Supply Chain, Seattle, Autonomous Agents' 99, Washington, May 1, 1999

- [6] Maes, P., Guttman, R., Moukas, G., Agents that Buy and Sell. Communications of the ACM. March 1999- Volume 42, Number 3, 81-83.
- [7] Mudgal, C., Vassileva, J. (to appear) An Influence Diagram Model for Multi-Agent Negotiation, in Proceedings of International Conference on Multi-Agent Systems, ICMAS'2000, 7-12 July 2000, Boston, MA.
- [8] Shachter, R., Evaluating Influence Diagrams. Operations Research. Volume 34, No 36, 1986, 871-882.
- [9] Shachter, R., Probabilistic inference and influence diagrams. Operations Research. Volume 36, No.4 1988, 589-604.
- [10] Marketplace for Java™ Technology Support. on-line at <http://www.hotdispatch.com/sun>
- [11] Vassileva J., Deters R., Greer J., McCalla G., Kumar V., Mudgal C. (1998) A Multi-Agent Architecture for Peer-Help in a University Course, Proceedings of the Workshop on Pedagogical Agents at ITS'98, San Antonio, Texas, 64-68.
- [12] Vassileva J., Greer J., McCalla G., Deters R., Zapata D., Mudgal C., Grant S. A Multi-Agent Approach to the Design of Peer-Help Environments, in Proceedings of AIED'99, Le Mans, France, July, 1999, 38-45. also available on line at: <http://julita.usask.ca/homepage/Agents.html>
- [13] Zheng, D., and Sycara, K. Benefits of Learning in Negotiation in Proceedings of Fifteenth National Conference on Artificial Intelligence, 1997. 36-41.