7º Simposio Argentino de Inteligencia Artificial - ASAI2005
Rosario, 29-30 de Agosto de 2005

A Multi-Agent System to Support Ambulance Coordination in Time-Critical Patient Treatment

Beatriz López, Bianca Innocenti, Silvana Aciar, and Isabel Cuevas

University of Girona
Campus Montilivi, edifice P4, 17071 Girona, Spain
{blopez, bianca, saciar,icuevas}@eia.udg.es

Abstract. Stroke is the third highest cause of mortality and the first cause of disabled people in western countries. A significant number of the people who survive live with serious physical and psychological disabilities and require permanent assistance in their dairy activities. When detected, there is a limited time in which to take effective treatment measures. In this paper we present a multi-agent system, MASICTUS, with the aim of supporting the diagnosis of acute stroke diseases while coordinating ambulance services and expert neurologists to attend the patient in time. In particular, we propose using an auction mechanism based on trust to coordinate the ambulances. To assure that the patient is treated in time, the system has reactive behavior that is able to deal with incidents that could occur when the ambulance is traveling to the patient's location or the hospital.

1 Introduction

Stroke is a cerebrovascular disease which affects the blood vessels that provide blood to the brain. It is also called an acute cerebrovascular accident (*ictus*), emboli or thrombosis. As a consequence of a stroke neural cells in the affected area do not receive oxygen and therefore cannot work, dying within minutes. There are two main kinds of strokes: ischemic and hemorrhagic. In the former, blood vessels are internally obstructed, while in the latter, the blood explodes in the brain [4].

Stroke is the third highest cause of mortality and the first cause of disabled people in western countries [1]. A significant number of people who survive live with serious physical and psychological disabilities and require permanent assistance in their dairy activities. Mortality rates that have been descending in the last decades, have currently increased due to the large proportion of elderly people in the population, who have a higher risk of having a stroke.Therefore in the future, in addition to personal, family, social, and labor consequences, acute stroke will cause a significant health and economic burden for health systems. WHO (World Health Organization) has also stated the importance of the illness in Europe and has arrived at a set of principles aimed at providing the best stroke practice (Helsinborg declaration).

Acute strokes are medical emergencies, because they arise acutely and unexpectedly (but not unpredictably), and either the patient or their family request quick attention for the neurological fault. Emergency treatment is particularly important because around the affected area, the ischemic penumbra, there is a critical therapeutic time window. That is, there is a limited amount of time in which the treatment given will be effective. This time window is not crisp, but evidence has shown that it is no longer than 6 hours, and for the best results it is 3 hours [3].

Recent studies in a given region [5], have shown that the lack of expert neurologists in every health center means that strokes are detected outside the therapeutic time window. Expert neurologists are placed in large hospitals, usually located in big cities. Moreover, the ambulance teams responsible for moving patients from their original location to the large hospitals are private services interested on maximizing their benefits. So conflicts of interest can arise between patient treatment and transportation.

In this context, those in charge of the large hospitals have considered the possibility of developing a computing support system to coordinate expert neurologists and ambulance services. From our experience, we have proposed and developed a multi-agent system, MASICTUS, which supports the medical stroke protocol that assures that the appropriate treatment arrives to the patient. Our aim is to show to the health care authorities how a multi-agent approach can improve the current ambulance service.

Two main features of the system are outlined in this paper: ambulance coordination, and ambulance reactive behavior for dealing with incidents that could occur when the ambulance is travelling to the patient's location or the hospital. The paper is organized as follows: First, in Section 2 we describe the multi-agent architecture. Details about ambulance agents are provided in Section 3 and the ambulance coordination method and reactive behavior is explained in Section 4 and 5 correspondingly. The first implementation results are shown in Section 6, and we end with some discussion and conclusions in Section 7.

2 MASICTUS Architecture

In order to support the stroke protocol, we have designed a multi-agent architecture in which two main kinds of agents are distinguished: agents related to the health care service (patients, health centers, stroke protocol and ambulance teams), and supporting agents (expert systems and trust agents) (see Fig. 1). The health-care agents are in charge of assuring that the medical protocol is carried out correctly for ischemia stroke treatment. In this protocol an ambulance is required to transport the patient to the corresponding health center. On the other hand, supporting agents help in the process of diagnosing the disease and assess ambulance reliability.

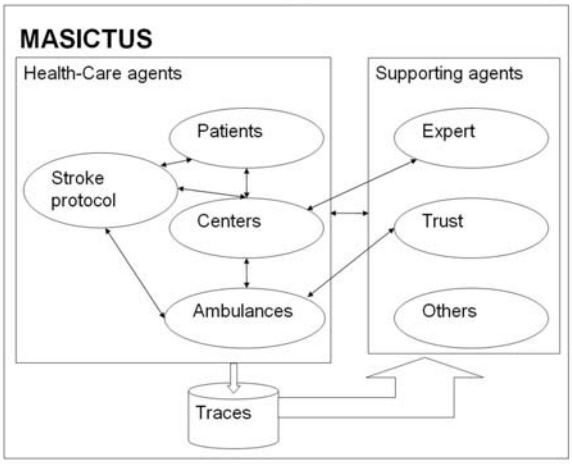

Fig. 1. MASICTUS architecture

2.1 Health-Care Agents

There are four main kinds of health-care agents: patients, health-care centers, stroke protocol and ambulance teams. First, patient agents deal with all the information related to patients. When a citizen suffering from a stroke either arrives at a health center or calls an emergency phone number, a patient agent is created. This kind of agent keeps the records of the patient until he/she is finally admitted to the hospital.

Second, there is a health care agent for every health care center involved in the stroke attention: local centers (primary attention), zone hospitals, main hospitals and the emergency phone centers (061 phone calls). A citizen can arrive at any of these centers with an acute stroke and they all should be able to detect the illness and apply the appropriate stroke medical protocol.

Each health care agent interacts with an expert agent that helps in the process of diagnosing the patient by following a fuzzy logic approach. The outcome of the expert agent determines the kind of center to which the patient should be transported: zonal or main hospital. Then, the health care agents activate the stroke protocol by interacting with the corresponding stroke protocol agent.

Third, the stroke protocol agent starts interacting with the main hospital in order to alert expert neurologists about the new patient. In addition, it requests an ambulance from the ambulance agent in order to transport the patient to the destination hospital.

Finally, ambulance agents enter the scene by providing a service to the patient. The aim of the ambulance agent is to provide the health care center with the requested ambulance on time. To achieve these objectives, two tasks should be addressed. First, selecting the closest ambulance to the health care center to fulfill the time constraints. And second monitoring the ambulance arrival, so they react on time to any incident that may occur. Due to its complexity, the

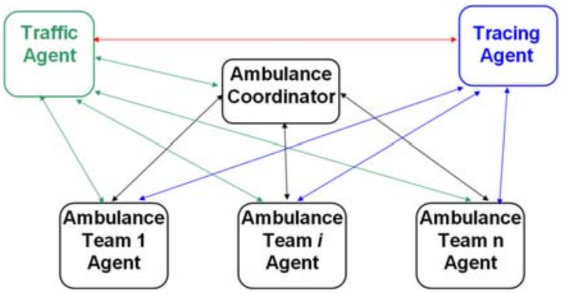

Fig. 2. Ambulance multi-agent system

ambulance agent has been conceived as an abstract agent in the form of a multi-agent system. We explain this in detail in the next section. In section 4 and 6 the different selection and monitoring activities are explained, respectively.

2.2 Supporting Agents

There are two supporting agents: the expert system agent and trust agent. The expert system agent implements the decision diagnostic procedure in order to find out if the patient is suffering an acute stroke. This agent is quite important since the drugs to heal patients are relatively new and most of the health care community is not aware of the corresponding medical protocol. For the sake of simplicity we will not discuss the details of the agent here, and we refer the reader to [4].

The trust agent keeps information about ambulance service reliabilities. That is, past ambulance behavior is kept in a file (see Fig. 1). With this information, a trust value of each service is computed by the trust agent according to the methodology defined in [18].This trust value is defined within the [0,1] interval; 0 means the agent is untrustworthy, while 1 indicates a completely reliable agent. The trust value is used in the ambulance coordination method described in section 4.

3 Ambulance Abstract Agent

As stated above, the ambulance agent in the MASICTUS architecture is an abstract agent composed of several agents, namely: the ambulance coordinator, ambulance teams, the traffic agent, and the tracing agent (see Fig. 2). There is an ambulance team for each real ambulance, while there is a single ambulance coordinator, traffic agent and tracing agent.

3.1 Traffic Agent

The traffic agent is in permanent contact with the national traffic central in order to get information related to traffic obstructions, accidents and temporally closed streets and roads. Thus, it is possible to locate points on the map where ambulances cannot pass. This agent proactively updates when a traffic incident is detected by de national traffic central, broadcasting information about the incident to all the ambulance team agents. Likewise, it informs the tracing agent about traffic situations, since it has to know whether there are changes in ambulance trajectories and why.

3.2 Tracing Agent

The aim of the tracing agent is to record information about the history of the ambulance's past activity. In particular, if the service provided by an ambulance team has been successful, that is, if the ambulance has arrived at the center where the patient is placed in the appropriate time. If it has not, the reasons for the failure are recorded: whether the driver followed the best route or not, if there was a problem in the trajectory and why, etc. This information is stored in a tracing file (see Fig. 1) and can be used by the best agent to modify ambulance reliability.

This agent also sends results of previous similar trajectories to the ambulance team agents when requested. The past experiences help ambulance team agents to make decisions.

3.3 Ambulance Coordinator Agent

This agent decides which ambulance team will go to pick up the patient from their current location and take him/her to the hospital. In order to choose the ambulance an auction process is applied, which is explained in section 4.1.

3.4 Ambulance Team Agent

This agent represents a physical vehicle and it has two main roles: Bidder and Monitor. As a bidder participates in the auction process coordinated by the ambulance coordinator in order to be selected to pick up the patient. As a monitor, once a patient has been assigned to the ambulance team, the trajectory followed by the physical ambulance is monitored in order to detect possible deviations from the estimated arrival time. Both, the bidding and monitoring role are based on the facilities provided by GPS and path planning techniques. Each team agent is composed of three modules: the GPS module, the Trajectory module, and the decision module (see Fig. 3).

The GPS module obtains the global position of the ambulance and locates the emergency vehicle on the map. It uses an electronic device attached to the vehicle that constantly sends information to the ambulance agent about the position of the vehicle.

The trajectory module, based on information from the GPS module, the Traffic Agent and the time given by the Ambulance Coordinator Agent (maximum time), calculates the best trajectory to the patient location. To calculate

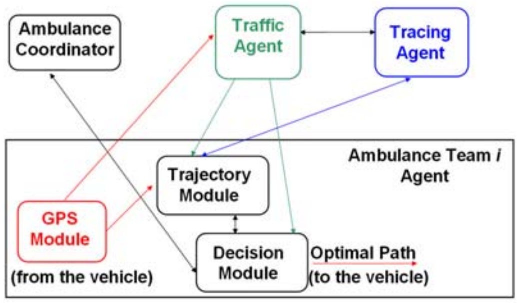

Fig. 3. Component of Ambulance Agent and links with the other agents of the system

the path, two constraints are taken into account. Firstly the time needed to execute the path should be less than the time given by the ambulance coordinator; otherwise the ambulance arriving to the patient will not actually help the patient. Secondly the trajectory must be free of obstructions. It is also possible to use information about previous trips in order to modify current paths that are similar to past ones that did not succeed. This information is available in the tracing file (see Fig. 1). The outcome of this module is the estimated time of arrival according to the best path found.

The decision module is the core of the ambulance team agent. This module is aware of the current role of the agent: either bidder or monitor. In the first case, it uses the estimated arrival time computed by the trajectory module to bid in the auction process governed by the ambulance coordinator. In the latter, the module sends the desired trajectory to an electronic device attached to the vehicle in order to inform the driver about which path to follow. If an unexpected incident occurs, as for example a car accident occurs in a calculated path, which can delay the ambulance, it must calculate an alternative path in order to get to the health center on time. In the case that the new path takes more than the specified time, then this module notifies the Ambulance Coordinator of the situation, which can decide to assign a new ambulance team.

4 Ambulance Coordination

Coordination in the MASICTUS system is simple at the higher level: patient agents interact with center agents in a predefined way. That is, a patient agent interacts with the center agent in which the citizen has physically arrived. No choice is made at this level.

The key issue here is coordinating the ambulance teams. As stated in the introduction, ambulances depend on private services which are paid for the number of services they perform. In this situation, different ambulance companies are competing for patients. Therefore, a more sophisticated coordination mechanism is required. In particular we propose an auction mechanism based on trust. Auctions assure that the cheapest ambulance in terms of time is acquired, which for our problem is crucial. In addition, trust provides a mechanism to control the truth of the information provided by the ambulance teams in the auction process. As stated in 2.2, the trust model is explained in [18]. There you can find how trust is defined and updated according to the successes and failures of previous services assigned to the ambulances teams.

4.1 Auction Model

At a given point in time, there are several patients that require an ambulance in a given health region. Regarding the stroke medical protocol, however, it should be noted that there are three main emergency situations:

Case 1: If the acute stroke has occurred in less than 6 hours, the patient should be taken to the main hospital.
Case 2: If the time window is in [6-24] and the patient fulfils exclusion criteria (coma, epilepsy, etc.), then he/she should also be taken to the main hospital.
Case 3: Otherwise, the patient should be taken to the zonal hospital.

In the first case, the stroke protocol clearly defines that a patient suffering from an acute stroke has maximum priority and transporting the patient to the main hospital requires maximum attention. In the second and third case, the transportation priority is the same as any other patient suffering from a heart attack, traffic accident, etc.

Consistently, in the first case, we should assign an ambulance to a patient; while in the second case there is a set of patients to be transported with different ambulances. Since there are two different situations, two different ambulance allocation processes are distinguished. For the maximum priority case, we propose an inverse auction, while for the second case a combinatorial auction is the appropriated technique.

On one hand, in an inverse auction (also called contract net), the auctioneer proposes certain tasks to be performed with some conditions [6]. In our problem, the ambulance coordinator gives the ambulance teams the task of arriving to the health care center where the patient is currently located and gives a time window as the condition to be fulfilled. This time window is the result of subtracting from the treatment time window (provided by the patient agent) the time estimated to transport the patient from their current location to the hospital destination. Therefore, the bidders (ambulance teams) that could perform the task in the given time reply to the ambulance coordination with a bid, containing the estimated arrival time that has been computed according to the bidding policy. Then, the ambulance coordinator decides which ambulance to allocate the patient to according to the winner determination algorithm.

On the other hand, in a combinatorial auction, several patient locations are auctioned at the same time [19]. This system uses the same bidding policies but a different winner determination algorithm. In particular, we have applied CASS [2], a combinatorial auction algorithm. To reduce the length of the text, we only describe in this paper the methods we have developed for inverse auctions.

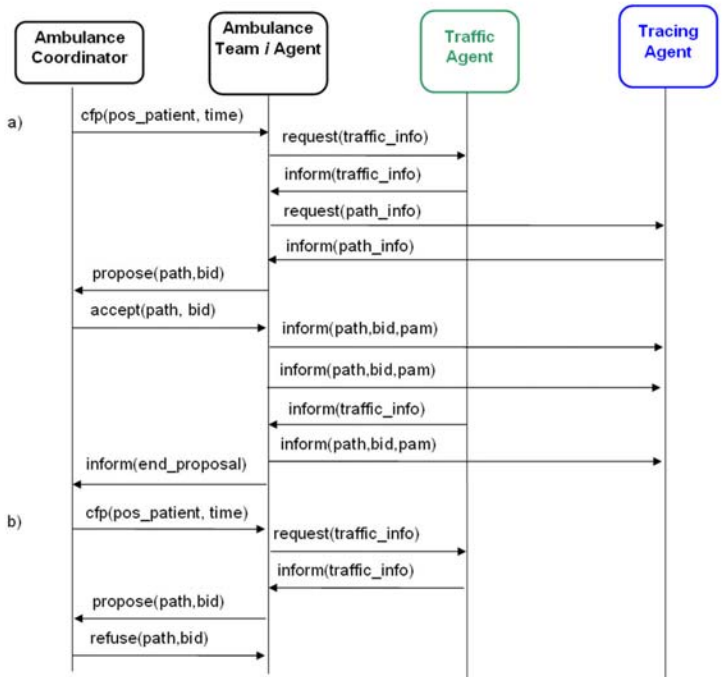

Fig. 4. Communication process among agents

4.2 Bidding Policy

As explained in section 3.4, the trajectory module of the ambulance agent team calculates the best path in terms of time and distance, to pick up the patient and transport them to the hospital (see also Fig. 4). The estimated arrival time is used as a bid in the auction process.

It is important to know that bidders have no incentives to deviate from the desired behavior, since an unappropriate behavior would be penalized with a decrease in the agent's trust value. With a lower trust value, an agent will have less opportunities to get new services on a near future, as explained in the following section.

4.3 Winner Determination Algorithm

The winner determination algorithm is applied by the ambulance coordinator to select the best proposal. This process has two parameters: the bid proposals, that is, the estimated arrival time of the ambulances, and their trust. The trust degree is defined in [0,1] and is computed by the trust agent. It does not necessarily hold that the ambulance with the best estimated time is the winner, it also needs to have a good trust degree. We have used fuzzy filters to filter the information provided by the ambulance teams according to their trust.

Fuzzy filters are good models for determining the degree to which the agents' assertions can be trusted in a competitive scenario [7].A fuzzy filter is a Mandami inference system in which the rules have the following form:

If A_1 *is* S_1 *and . and* A_n *is* S_n *then F is* L_1

where, A_i and F are fuzzy variables, S_j and L_1 are fuzzy labels. A_i are called the side variables, and F the filtered variable. There is a fuzzy filter for each agent, so the assertions of each agent, represented by the side variables, are then used to infer the filtered information.

For our purpose, the side variables are the estimated time (ET) and trust (t), and the filtered variable the increasing time (IT) to be added to the estimated time. The estimated time is defined in the universe of discourse [0,TTW], where TTW is the treatment time window; trust is defined in [0,1] (see section 2.2, trust agent) and the increasing time in [0,TTW-ET]. Observe, then, that for each agent the discourse domain of the increasing time will vary according to the estimated time provided by the agent. Then, with the increasing time the outcome of the fuzzy filter is proportional to the ET.

For each fuzzy variable, the following fuzzy labels are defined: ET: *very short, short, medium, long, very long*; t: *very low, low, medium, high, very high*; and IT: *very short, short, medium, long, very long*.

Note that the definition is dynamic, depending on TTW and ET. The fuzzy system consists in fuzzy rules such as:

R1: If ET is short and Trust is low then IT is very high
R22: If ET is short and Trust is very high then IT is very short

After applying the fuzzy filter, an increasing time is computed for each agent according to its trust. This increasing time is added to the original estimated time provided by each agent, so a new set of estimated times $ET_1^{'}$, $ET_2^{'}$, $ET_n^{'}$ is obtained. Then, the ambulance coordinator determines which agent has the lowest new time, and this agent is the winner of the auction process.

Our winner determination method then tries to be optimal regarding the preferences of the trusted individual agents in the system. If each individual preference is measured by the estimated time (cost), the concept of social welfare is then the sum of the individual utilities and can be used to measure the quality of the allocation from the viewpoint of the system as a whole according to [19].

5 Reactive behavior

Due to the strict deadline imposed by an acute stroke's treatment time window, we have provided the ambulance teams with reactive behavior with the aim of dealing with different incidents. In the case that an incident occurs that can substantially effect the time needed to pick up the patient the ambulance team can contact the ambulance coordinator so that a new ambulance auction process is started. Additionally, if the tracing and traffic agent can provide information regarding any deviation of an ambulance team. So they can alert to the ambulance

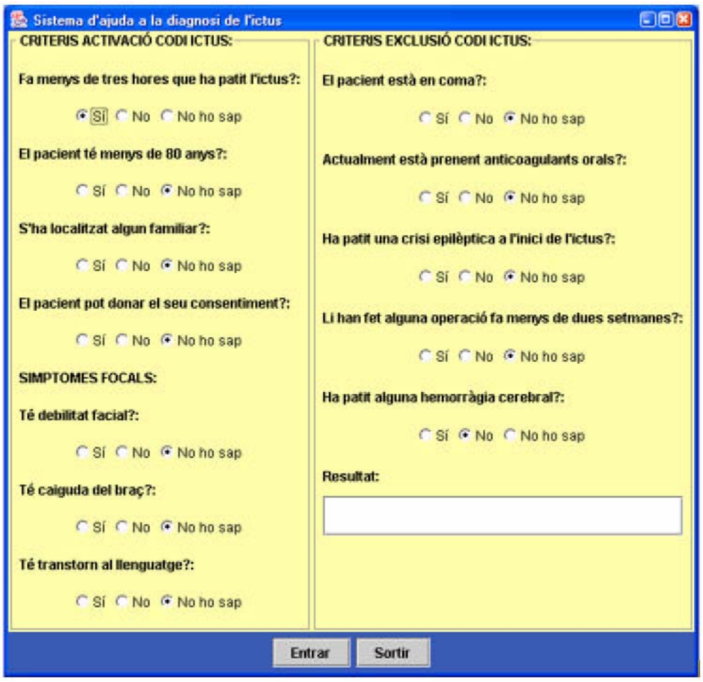

Fig. 5. Acute stroke diagnosis system interface.

coordination of any disfunction in the coordination process with the ambulance team.

6 Implementation

Currently we have a prototype of the system running in a JADE platform. Fig. 5 shows an interface of the system. Some of the functionalities of MASICTUS have already been completely deployed, namely, the patient, the center and the stroke protocol agents. Regarding the ambulance abstract agent, the inverse ambulance auction has been developed, that is, the functionality related to the highest priority patients.

7 Discussion and Conclusions

In this paper we have presented a multi-agent system with the aim of supporting the stroke medical protocol. Two main issues are addressed. First, we have provided a trust-based auction mechanism to deal with the decision making process to allocate an ambulance team to transport the patient. Then, we discussed reactive behavior in order to deal with incidents when carrying out the stroke medical protocol. There is a lot of work related to applying agents to the health care domain.

Regarding our problem, we would like to mention the research work of [13] on monitoring medical protocols. The authors propose a multi-agent system to assist and supervise the application of medical protocols in distributed hospital environments. The system is able to suggest actions and constraints (forbidden actions) to the medical staff. We believe that we can use the ideas in their work to improve our stroke protocol agent.

Regarding coordination of medical services, [14] propose an interesting multi-agent system for implementing medical guidelines, that is, sequences of actions, enquiries and decisions concerning a patient with a certain pathology. We think that this approach could be useful for implementing the decision support expert agent in MASICTUS. However, the stroke disease can be detected with very little information and simple tests, without time expensive tests such as blood analysis and other medical services. So coordinating different medical services in our case and regarding the patient pathology is limited to the patient record, the health center and the expert neurologist team in the main hospital.

Another interesting application of a multi-agent system for coordinating medical services is shown in [15]. Both, [15] and [14] replicate, to some extent, the existing human organization and authority structures in the multi-agent system. We also use this approach because we believe that changing the organization of authority in medical systems is unrealistic.

In [16] a multi-agent system is proposed for coordinating medical services provided by an ambulance team. Here, the approach has a different focus than ours. Instead of choosing an ambulance to transfer the patient, the ambulance proposes a hospital. The emphasis of this work is also on service coordination, but unlike our work, the authors focus on hospital service collaboration, while we emphasize competitive ambulance allocation.

Finally, [17] point out the importance of strokes as one of the major social diseases and propose the ADDHealth project to provide support to treatment comparisons. We find this proposal quite interesting since it opens the way to proving how computing based systems in general and agent-based systems in particular can improve health care treatment.

As a conclusion, we can say that there is a lot of work to do in relation to health care applications. We have tackled a particular approach to deal with critical-time patient treatment given that ambulance teams are managed by private companies. Currently we have a first prototype of MASICTUS. Our aim is that this prototype that simulates real ambulance coordination, convinces health care authorities about the utility of multi-agent supporting tools and integrate them into the health care system .

Acknowledgments

This research project has been partially funded by the Spanish MEC project TIN2004-06354-C02-02 and AECI INTERCAMPUS A/1562/04. Many thanks to Esteve del Acebo for his fruitful discussions about fuzzy filters.

References

1. Gil-Peralta, A.: Prevención primaria del Ictus. Revista Ictus, (1998).
2. Leyton-Brown, K.: Resource Allocation in Competitive Multiagent Systems. Ph.D. thesis, Stanford University (2003).
3. Fisher, M., Takano, K.: The penumbra, therapeutic time window and acute ischaemic stroke. Bailliees Clin Neurol Vol 4(2) (1995) 279–295.
4. Dávalos, A., Fisher, M.: Emerging therapies for cerbrovascular disorders. Stroke Vol 36(2) (Feb 2005) 208–210.
5. Dávalos, A. Personal communication.
6. Wellman, M. P. and Wurman, P. R.: Market-Aware Agents for a Multiagent World. Robotics and Autonomous Systems Vol 24 (1998) 115–125.
7. Del Acebo, E., de la Rosa, J.Ll.: A Fuzzy System Based Approach to Social Modeling in Multi-Agent Systems. AAMAS (2002).
8. Noriega, P., Sierra, C. and Rodríguez, J.A. : The Fishmarket Project. Reflections on Agent-mediated institutions for trustworthy E-Commerce. Workshop AMEC (1998).
9. Montaner, M., López, B., de la Rosa, J. Ll.: Developing Trust in Recommender Agents. In Proc. AAMAS (2002).
10. Yu, B. and Singh, M.P.: Searching Social Networks, Proceedings of Second International Joint Conference on Autonomous Agents and Multi-Agent Systems(2003) 65–72.
11. Yu, B. and Singh, M.P.: Towards a Probabilistic Model of Distributed Reputation Management, 4th Workshop on Deception, Fraud and Trust In Agent Societies Montreal (2002).
12. Patel, J., Teacy, W. T. L., Jennings, N. R. and Luck, M.: A Probabilistic Trust Model for Handling Inaccurate Reputation Sources. In Procc. Third International Conference on Trust Management (in press), (2005).
13. Alsinet, T., Anstegui, C., Bjar, R., Fernndez C., Many F.: Automated monitoring of medical protocols: a secure and distributed architecture. Artificial Intelligence in Medicine (2003).
14. Isern, D., Moreno, A.: Distributed Guideline-Based Health Care System. ISDA (2004).
15. Decker,K. and Li, J.: Coordinated hospital patient scheduling. In Proc. ICMAS (1998).
16. Ciampolini, A., Mello,A., Storari, S.: A multi-agent system for medical services synergy and coordination. ECAI workshop in Health Care (2004).
17. Greenwood, S., Raguin, T., Norring, C., Lan, T.: ADDHEalth. Agentlink III Technical Forum on Health Care. Rome (2004).
18. López, B., Aciar, S., Innocenti, B., Cuevas, I.: How multi-agent systems support acute stroke emergency treatment. IJCAI Health Care Workshop (2005).
19. Chevaleyre, Y. and Dunne, P.E. and Endris, U. and Lang, J. and Lemaître, M. and Maudet, N. and PAdget, J. and Phelps, S. and Rodríguez-Aguilar, J.A. and Sousa, P.: Issues in Multiatgent Resource Allocation. Agentlink III- Technical Forum Group on Multi-agent Resource Allocation. (2005).