

Organ Transplantation Management

Swiss Federal Institute of Technology Lausanne (EPFL)
IN-Ecublens, 1015 Lausanne, Switzerland

Student: Rachele Fuzzati
Professor: Boi Faltings
Co-Supervisor: Monique Calisti

EPFL Technical Report IC/2005/022

Table of Contents

1. Introduction	3
2. Healthcare Context	5
2.1 Requirements	5
2.2 Innovative Approach	6
3. Organ Transplantation Management	8
3.1 Problems	11
3.2 Ideal Characteristic of Solution	11
4. Agent Based Solution	13
4.1 State of the Art	14
4.2 The System Proposed	15
5. The Problem of Establishing the Winner	17
5.1 An Example in the Organ Transplantation Domain	17
5.2 Our Idea	17
5.3 Expected Scenario	19
5.4 Problems Arisen	19
5.5 Solutions Adopted	19
6. The System Realized	22
6.1 Conventions in the Future Descriptions	22
6.2 Framework	22
6.3 Target	23
6.4 System Implemented for the Simulations	23
6.5 Variables	26
7. Obtained Results	27
7.1 First Set of Results	27
7.2 Interpretation of the First Set of Results	30
7.3 Suggestion	31
7.4 New Results	31
7.5 Interpretation of the New Results	36
8. Conclusions and Future Developments	38
8.1 Conclusions	38
8.2 Future Developments	38
Bibliography	40

Organ Transplantation Management

Graduate School Project, Artificial Intelligence Laboratory
Swiss Federal Institute of Technology Lausanne (EPFL)
IN-Ecublens, 1015 Lausanne, Switzerland

Student: Rachele Fuzzati
Professor: Boi Faltings
Co-Supervisor: Monique Calisti

Final Report, July 2002

1. Introduction

In the healthcare context, organ transplantation has raised to great importance in the last years. Improvements in medical techniques and pharmacological anti-reject therapies have made transplantation a powerful and valid way to treat diseases. Thanks to this, and to the relevance that mass media put on it, the number of donors is constantly increasing all over Europe. The decision to assign an organ from a particular donor to a particular recipient is a very complex process which can be decomposed into the following activities:

- Gather, store and manage a mandatory set of personal and medical information about each recipient and each donor (ex: blood group, weight, height, tissues characteristics, ...)
- In presence of a donor, find a group of potential recipients which are compatible with the donor with respect to the mandatory sets of information stored
- Among the group of potential recipients, find the one that best fits with the donor. This decision is taken not just on the base of medical parameters (such as the current health state of the patient), but also on the base of logistical considerations (such as the possibility to transport the organ from the donor's to the recipient's hospital and the availability of medical teams to perform the operation)

On top of this, there is a very strict constraint in time. Donors are usually persons who have been involved in serious accidents and are kept artificially alive though cerebral death has been stated. In such patients is extremely difficult to maintain acceptable vital parameters for a long time, and any variation in them can lead to the loss of organs. Moreover, after the organs have been taken from the body, they can be stored for very few hours and the transplant must take place in the shortest period possible.

Currently, the activities described above are almost entirely performed by human beings, with weak supports to process large amounts of data and to coordinate with each other. The organization and storage of information in compact, re-usable ways and the introduction of software support systems in the process of matching can accelerate and simplify the assignment of organs.

In this project, the organ transplantation management problem is mapped into a multiagent system. In the system, donors, hospitals representing potential recipients, information coordinators and medical experts are all represented by agents. Donor and hospital agents register to the system by revealing their information to the information coordinator agent, who is in charge of storing and

managing it in a private and confidential way. Donor agents offer organs, hospital agents compete among each other to get the organ they need, medical expert agents cooperate with information coordination agents to interpret information and guarantee a correct and fair allocation of the organs.

For simplicity, the logistical aspects relative to organ transportation and medical teams scheduling are not taken into account in this first realization. The focus of the project is on the operation of matchmaking. For the characteristics that it presents (agents offering a “good” and agents interested in getting the “good”), the system strongly resembles a marketplace, with the connected problems of ensuring the fair play of agents. However, most of the techniques applied to existent marketplaces to force agents to a fair behaviour involve money payment and economic fees and are not directly applicable in the case of organ transplantation management since the law bans the use of money in this field. Therefore, it is necessary to find other methods to perform payments and achieve the fairness result.

An overview of the healthcare context is given in Section 2. Then, the focus is restricted to the organ transplantation domain with the description, in Section 3, of the activities performed, the related problems and the ideal solution. Section 4 justifies why multi-agent systems are considered suitable to solve this kind of problems, then presents the state of the art in the domain and the suggested solution to which this project gives contribution. Section 5 enters the main topic of the project with the presentation of our idea on how to perform matchmaking and establish the effective recipient of and organ, the problems arisen in developing this idea and the solutions applied to solve them. The system realized and used for the simulations is described in Section 6, while Section 7 reports the results obtained during the simulations and tries to give explanation to some unexpected behaviours of the system. Section 8, Conclusions and Future Developments, draws a summary of the main aspects emerged during the realization of this project and suggest open points of future development.

2. Healthcare Context

Healthcare is a context in which, sooner or later, everybody could be involved. Many steps forward have been done in this domain in the last years. Medical studies and improvement in techniques have led to unimaginable capabilities. Unfortunately, the more the possibilities and knowledge grow, the more is difficult for the current healthcare system to guarantee adequate performances. New medical techniques involve the composition of different medical services and the cooperation between different medical structures, with the exchange of large amounts of complex information. Healthcare systems are still related to the past. They are not flexible and do not yet have the right tools to cope with the problems introduced by new technologies. Difficulties in composition of services and managing of information lead to a not complete use of the possible therapies, to longer hospitalisation times, to a waste of resources. This is negative both for the patients, who spend more time being treated, and for the healthcare structures, who do not use their resources at the best.

2.1 Requirements

Among the domain, it is possible to distinguish two different kinds of problems. A more general/theoretical kind, which is related with the representation and organization of information and its reuse for medical purposes. And a kind related to possible application, which can help both patients to better afford their diseases and medical personnel to better intervene on patients. In both cases, the aim is the improvement of the existent services offered by hospitals and healthcare systems and the creation of new ones, to guarantee more efficient medical assistance to the citizens.

Global problems can be distinguished into:

- Dependencies and correlation of medical information and related services. Medical personnel should be able to access complete sets of information about patients and build personalized services (such as therapies or treatment prescriptions) on top of that. For example, whenever a patient is checked by a doctor or is treated for some accident or disease, the complete situation of the patient has to be reported. This has to be put in relation with the precedent health-history of the patient, as it might reveal important correlations (such as the existence of an allergy, epilepsy and so on). Indeed, knowing the complete health-picture of a patient and being able to retrieve it quickly can help to save time in case of banal injuries, but can help to save life in case of an emergency.
- Need for coordination. It is possible that there is the need to combine distinct services within a single hospital or organization. In this case is very important to carefully coordinate the parts which perform the different services to avoid wasting of time and resources. For example, each day, in a hospital, physicians request certain tests and or therapy to be performed as a part of the diagnosis and treatment of a patient. Often, tests and treatments are performed by independent departments in the hospital (ex: patient residing in intensive care who need X-ray service and orthopaedic service). It is necessary to carefully coordinate the activities and the personnel of the different departments involved and to plan the order in which the activities have to be performed (ex: the patient can't be treated by orthopaedic service if he hasn't passed through X-ray service first). When the hospitals/organizations involved in the composition of services are more than one, coordination becomes an even more stringent requirement and the problem is complicated by the introduction of confidentiality, privacy and security aspects. For example, if the hospital (say A) in which the patient is residing can not guarantee X-ray service, hospital B has to be contacted for it. Hospital B may be willing to perform the service, but may want to protect its information and refuse to reveal the X-ray timetable for the day. In this case the coordination becomes

hard since some information must be revealed between the two hospitals while privacy and security aspects have to be issued.

Application driven problems are:

- **Monitoring.** Some patients (diabetics or elderly people, for example) do not need to be hospitalised, but they need to be constantly monitored. It is necessary to control they take the prescript drugs at the scheduled time and maybe it would be useful to vary each time the quantity of medicine to assume (insulin, for example) according to the current physical conditions.
- **Planning.** It is useful to plan checkups basing on the real needing of the patients. In this way, hospitals can give a better quality of service to patients, concentrating resources and personnel just on cases of real need. On the other side, healthy patients feel less oppressed by their disease, having a higher quality of life.
- **Scheduling.** Nurses and doctors have precise timetables for their work, but it is possible that they have to substitute each other or change their turns. In such cases it is extremely important to have a clear schedule of the tasks to be performed and the therapies to be executed on the patients.

2.2 Innovative Approach

Most of the aspects discussed above require high costs in terms of effort and time spent by healthcare systems and medical personnel to solve them. Things are made even more complex by the fact that, in healthcare domain, time is always relevant and constraints quickly change with it (ex: the conditions of a patient can get suddenly worse, requiring a quick and unexpected intervention). However, some of the tasks (such as planning or scheduling) presented in the examples are not just a prerogative of humans and could be left to software applications designed to well take into account the constraints of the problems. Given the particular field, one can not expect to completely substitute human beings, but the appropriate technical support could lead to:

- better use of human resources,
- faster and more efficient composition of medical services,
- more coordinated cooperation between medical structures

and, consequently, to:

- better conditions for the patients
- saving of money/resources for the healthcare systems.

One of the fundamental steps to do in this direction is to find a better way of representing information. In healthcare context, each patient has an enormous amount of information related to himself. This information, which can come from different doctors and different hospitals, is usually a non structured collection of medical reports and documents. Therefore, it is often very difficult, even for real medical personnel, to evaluate the entire clinical history of the patient. To improve the readability of information and to allow the creation of software tools capable to aid medical personnel in its job, it is fundamental to well organize the available information and to represent it in compact ways that make it reusable.

The use of ontologies is a way through which is possible to reach this purpose [7, 10]. The term “ontology” refers to a shared understand of the domain of interest. It is a unifying framework which embodies objects and concepts, their definitions and the relationships between them. Ontologies provide consistent vocabularies and world representation necessary for clear communication within knowledge domains. An example of ontology in the medical domain is UMLS [24]. UMLS is a long term project of the National Library of Medicine and is specifically created to enable new

information technologies to take advantage of controlled medical vocabularies. However, the development of ontologies is still a young discipline, it does not yet embrace all the possible branches of medicine and healthcare, and most of the work still has to be done.

Once the information is compact, structured and well represented, it is possible to provide solutions which use and manipulate it guaranteeing:

- Flexibility. Every hospital/healthcare system has a different organization. Every patient has a different history. The tools realized should adapt to these characteristics, being able to find new solutions whenever changes in the state or in the constraints arise. For example, in the case the X-ray machine of a section of a hospital goes out of order, a tool could find another section with the same machine working, contact it and send the patients to the new section.
- Automation. Basing on the information available about the patients, the medical constraints and the healthcare systems, the software solution tools should be able to take autonomous decisions regarding scheduling, planning, matching. For example, a tool could coordinate nurses and hospital departments by scheduling tests and therapies for patients [3, 4].
- Personalization. Given that they deal with real patients, who are different one from the others in terms of needing and reaction to therapies, software solution tools should be able to automatically take into account the single characteristics and personalize the treatments. For example, in case of a rising of the glucose level in the blood, a young diabetic patient may be suggested to do more physical exercise and see if the problem reduces, while an old patient may be suggested to go immediately to the doctor.

3. Organ Transplantation Management

Organ transplantation is a very complex problem that summarizes most of the aspects examined until now: knowledge management, planning/scheduling, coordination, monitoring. Moreover, on top of this, there is also a very strict constraint in time due to the intrinsic nature of the problem (organs have a short time life outside bodies and can be stored for very few hours). All these characteristics, together with the wide diffusion of this practice all over the world, make organ transplantation management an exciting, challenging problem to be solved.

Organ transplantation has raised to a great importance in the last years. Transplantation is more and more seen as a valid way to treat diseases and no longer as just the last option therapy. Thanks to frequent campaigns for citizens information, the number of donors is constantly increasing all over Europe. In Switzerland, there is a unique national centre of coordination, Swisstransplant [26], which is in charge of registering patients, analyse organs compatibilities, set up logistic for transplants. Currently, due to strict time constraints and legal aspects, the origin of transplanted organs is mainly local. However, most neighbouring nations share reciprocal agreements between their transplant organizations: in the event no matching is found among the recipients of a country, the organ is advertised to the neighbour countries (in the case of Switzerland: Italy and France). This gives better chances to find an appropriate recipient for each organ and avoid wasting. Also, an international European organization, Eurotransplant [25], has been recently created with the aim of increase and better coordinate the transplants all over the continent. The Eurotransplant International Foundation is responsible for the mediation and allocation of organ donation procedures in Austria, Belgium, Germany, Luxemburg, the Netherlands and Slovenia. In this international collaborative framework, the participants include all transplant hospitals, tissue-typing laboratories and hospitals where organ donations take place. The Eurotransplant region numbers well over 118 million inhabitants and it is open to new countries to join.

Organ transplantation requires a delicate balance between fairness and medicine to decide in each case who will be the effective recipient.

From a medical point of view, there is a set of information which has to be provided by both potential recipients and donors and on which the first operations of choice are performed. The main information among this set is about:

- blood group
- tissue characteristics (HLA groups)
- weight
- height
- age

Some of these characteristics (ex: blood group) determine a definite incompatibility, most of the others just give indications on how suitable the organ is (ex: weight) or on how much the donor matches with the recipient (ex: tissue characteristics). The latter aspect is the most important: in transplantation, a rejection against an organ occurs when the recipient's immune system recognizes the donor organ as alien; so, the more the tissue characteristics match those of the recipient, the less chance there is that a rejection reaction will occur.

On the other side, from a fairness point of view, there are other important factors that influence the final assignation of an organ. Information of key importance for the definition of the effective recipient are:

- the age of the patient being under sixteen
- the relative position in the waiting list
- the presence of 0-emergency cases

For example, doctors might decide it is better to assign an organ to a patient who has a worst compatibility but is in danger of life if he does not receive a new organ in a short time (0-emergency), instead of giving the organ to the patient with just the best compatibility. In this cases, it is extremely important to verify that the real conditions of the potential recipients are fairly

declared and correctly reported, as some doctors might try to present the situation of their patients worse than how it actually is, in order to gain the organ.

The organ transplantation process is mainly composed by two different phases: the procurement phase and the surgery phase (Figure 1). The procurement phase is about finding the most appropriate recipient and organizing the logistic solutions for the transplant. The surgery phase comes later and is almost exclusively related to medicine: surgical operations, anti-rejection treatments and subsequent therapies. The focus of this project will be on the procurement phase. This can be further decomposed in a set of activities that are depicted in Figure 2:

- **Matchmaking.** All the patients waiting for an organ have to register to a waiting list and reveal to the hospitals or to the competent institutions the required information about themselves (age, weight, blood group, etc). When there is the presence of a donor, all the necessary and useful information about him are forwarded to the competent institutions. In these institutions, medical experts analyse all the potential recipients and eventually come up with the one that best matches with the donor accordingly to the criteria described above.
- **Transport Routes Planning.** After a suitable recipient has been chosen, the organ must be explanted and sent from the donor hospital to the recipient hospital. In some cases (ex: Switzerland), the law requires that a team from the recipient hospital assists to (or directly performs) the explant of the organ. In this case, it is necessary to arrange both the incoming and the outgoing journeys of the medical team.
- **Medical Teams Scheduling.** When a suitable recipient is found, the organ must be explanted from the donor and implanted in the recipient. Both the explant and the implant operation require the reservation of operating theatres and the arrangement of the presence of medical personnel (sometimes from both the donor and the recipient hospitals contemporarily).




Figure 1

Procurement Phase

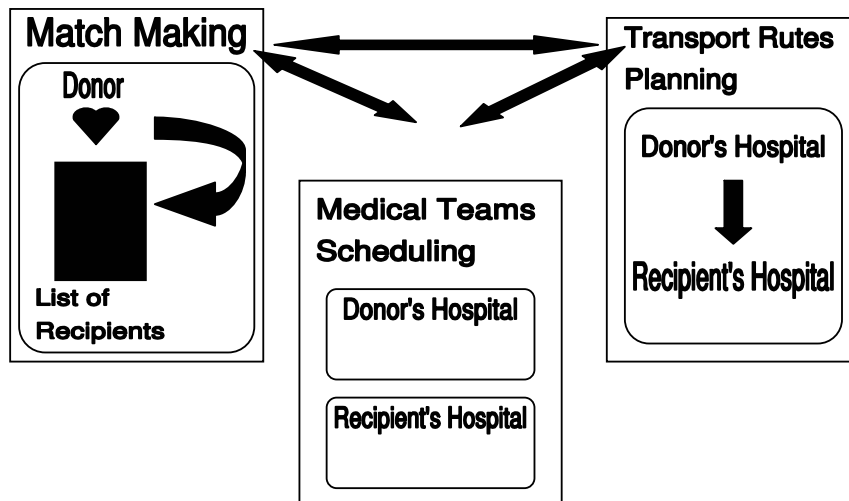


Figure 2

3.1 Problems

On top of the activities described above, there is a very strict constraint in time. Donors are usually persons who have been involved in serious accidents and are kept artificially alive though cerebral death has been stated. In such patients is extremely difficult to maintain acceptable vital parameters for a long time, and any variation in them can lead to the loss of organs. Moreover, after the organs have been taken from the body, they can be stored for very few hours (four for heart, six for lungs, sixteen-twenty four for liver, thirty six for kidney [26]) and the transplant must take place in the shortest period possible.

Currently, the way in which the procurement phase tasks are performed is still largely non-automated and non-coordinated. The main pitfalls in the transplant process are:

- Medical experts have to consider one by one all the possible receivers and evaluate the matching with weak supports to process large amount of data.
- Information is usually not stored in a compact, re-usable way, therefore the coordination between medical experts and surgeons has to pass through telephone and facsimile.
- Finding the best route involves looking up several timetables of means of transport (such as trains or planes) and making spatial and temporal reasoning to provide the most efficient solution.
- Scheduling the medical teams involves looking up the timetables of operating theatres and medical personnel to find solutions which are available at the required times.

In the case of unavailability of medical personnel/operating theatres in the hospitals involved or in the case of routes which require a too long time, the designed recipient has to be discarded and a new one, for which is possible to come up with appropriate solutions, has to be found. All the process has to start again from scratch: find another receiver and organizing medical team and transportation. This leads to undesired delays and sometimes causes the waste of the organ.

The situation is even more complex if we take into account the fact that usually, when there is a donor, the amount of organs available for a transplant is greater than one, and the described procedure has to be repeated in parallel for each one of the organs. To deal with all these difficulties, often medical experts stop the matching procedure when they reach the first fit, instead of looking for the best fit (best compatibility, better chances of life...) or they tend to assign more than one organ to the same hospital in order to organize just one delivery for multiple organs.

3.2 Ideal Characteristics of Solution

The management of organ transplantation processes is a very complex and delicate task. However, software support systems could help in speeding up and simplifying some of the operations of the procurement phase, guaranteeing a better use of resources and rising the chances of success.

Organizing information in compact, re-usable ways for both patients and medical personnel is fundamental and allows the use of software to perform matching and planning. Possible ideal solutions for the three different aspects of the problem are:

- Matchmaking. A software tool could be in charge of performing the matching operation between donor and recipients. Accordingly to the desired level of involvement of human beings, there can be two possibilities. The first one is about the realization of a decision

support system that evaluates all the possible receivers, following some given criteria, discards all the patients that have a level of compatibility lower than a certain threshold and passes to the medical expert just the few patients that could be really compatible. In this way, the medical experts could concentrate just on a restricted number of receivers, making sure they choose the best fit match. The second possibility is to allow a completely automatic mechanism of decision. Such system would rely on a trustable mechanism of decision which, following the given constraints, is able to come up with just one recipient, the one considered most suitable as the result of the elaboration. The role of the human medical expert in this case would be marginal. Once the mechanism has proved to be safe and fair, and the system has proved to be secure, the human task will be reduced to verify the effective compatibility of the donor with the designated recipient.

- **Transport Routes Planning.** A software tool could perform the timetable looking up, combine the data to fulfil the requirements and suggest the possible best route. Human beings should just verify the correctness of the route. The tool could also be in charge of booking places, make reservations, communicate the priority of the fare to the transport companies.
- **Medical Teams Scheduling.** A software tool could look up for operating theatre timetables and medical personnel schedules, combine the information to fulfil the requirements and suggest times for the operation, available operating theatres and medical teams. Human beings should just verify the correctness of the results. The tool could also be in charge of booking operating theatres and communicating the scheduling of the operation to the medical personnel.

With such tools, the research for the best receiver could become quicker. Moreover, the event of not finding a suitable medical team for the operation or a convenient route for the transportation would no longer represent a dramatic problem. In fact, with the matchmaking tool producing a list of potential recipients, the experts would have now to deal with a restricted number of possible receivers and undertaking another matching process would be easier. For example, they could pick another recipient from the list of suitable ones. With the matchmaking tool producing just one result, the experts will simply have to run the process of matchmaking again. If the system is well designed, the process should not require too long time to be performed and the repetition would not affect human resources, which would be left free to concentrate on other tasks. Moreover, the system could be designed in such a way to exchange information with the other tools and take into account their results, producing a result which is already suitable under the three points of view. Even in the case of multiple organs to be assigned, the process would turn out to be more efficient, since different software tools (one for each organ/medical team/route) could work in parallel substituting human beings and producing results with less cost in term of human resources involved.

4. An Agent Based Solution

Summarizing the description made in the previous Section, the scenario in which the solution has to be drawn has the following characteristics:

- Constantly changing environment. Constraints on the system can be added or removed (for example by surgeons), new recipients can register to the waiting list, new organs can become suddenly available.
- Possibly conflicting tasks to perform. The best recipient for an organ accordingly to the physical characteristics could be a certain patient, but the presence of a 0- emergency case could force the result to be different.
- Many different possible courses of actions available. The goal to reach is the selection of a recipient, but the actions through which is possible to get to the goal are multiple and the appropriate one must be chosen every time. For example, with a rare blood group donor, the first criteria of selection (in order to speed up the decision process) could be the blood group, in other cases another criteria could be more efficient.
- Necessity to make prediction about the future. The possible consequences of an action must always be taken into account and influence the decision. For example, in selecting the best recipient, the system has to predict the patients' probabilities of life and use them as a parameter for the result.
- Necessity to take decisions in a timely fashion. Time constraints in organ transplantation management process are very strict and the selection of the recipient has to be taken in the shortest time possible.

It is clear that the problems are very complex and the traditional software tools are not adequate to the purpose. The solution should be able to provide great flexibility and simulate human behaviour in taking decisions. A possible suggestion when dealing with such problems is the use of intelligent agents and multi-agent systems.

In literature, the properties of agents are defined as [23]:

- autonomy: agents operate without the direct intervention of humans and they have control over their actions and state
- social ability: agents are able to cooperate and communicate with humans or other agents in order to achieve their tasks
- reactivity: agents perceive their environment and respond in timely fashion to changes that occur in it
- pro-activeness: agents do not simply act in response to their environment, they are able to exhibit goal-directed behaviour by "taking the initiative"

And also, multi-agent systems are depicted as ideally suited to representing problems that have multiple solving methods, multiple perspectives and/or multiple solving entities [6]. Such systems have the traditional advantage of sophisticated patterns of interactions. Examples of common types of interactions include:

- cooperation: working together towards a common aim
- coordination: organising problem solving activity so that harmful interactions are avoided or beneficial interactions are exploited
- negotiation: coming to an agreement which is acceptable to all parties involved

Therefore, a multi-agent system can be defined as a loosely coupled network of problem solvers that work together to solve problems that are beyond the individual capabilities or knowledge of each problem solver. These problem solvers – agents – are autonomous and may be heterogeneous in nature.

Thus, we conclude that, considering the procurement phase of the organ transplantation management, software agents and multi-agent systems turn out to be the more appropriate solution as they can evaluate and optimise the utility of specific actions, gather and process huge amounts of information, follow specific strategies and interaction protocols more efficiently and more rapidly than humans or traditional mechanisms can.

4.1 State of the Art

Organ transplantation management is a new field to which artificial intelligence has turned, therefore very few solutions have been proposed by now. The most important work has been done by the University of Catalonia, together with the University of Rovira i Virgili. The tools proposed are three multi-agent systems that cover the three phases of organ transplantation management: matchmaking, transport routes planning and medical teams scheduling. Many assumptions have been made and many constraints have been relaxed to make the creation of the tools possible. Moreover, the model of organ transplantation management examined was the Spanish one, which is simpler than the Swiss under some aspects (for example, the Swiss law requires the presence of recipient hospital medical team during the explant operation, this increases the complexity of the transport organization).

- **Matchmaking.** Carrel [19, 20] is an Agent Mediate Institution capable to regulate and speed up the organ assignation process. It is an intelligent resource management service designed to share, among the hospitals that are members, all the information stored in the different Banks of Tissues. Through a special kind of negotiation, it assigns the “best” piece (organ, tissue or bone) to the recipient that shows the “best” match with it. The concept of best is related to a set of rules, which are imposed by both the policy of the hospital transplant unit and the surgeon who is treating the potential recipient. The coordination among Banks of Tissues can be done at regional, national or international level, taking advantage of the single negotiation protocol and of the standardization of information. A delicate problem, however, is the respect of all the national and international laws.
- **Transport Routes Planning.** The multi-agent system described in [14] is capable of finding the best ways to send an organ from one city to another. The concept of “best” takes into account both the time needed for the transfer and the expenses in economic terms. As means of transport, just trains and planes are considered, but the research could be extended to other means (ex: busses). The weak points of this system are some too relaxed constraints applied in it. The first one is that the problem has been limited to the transport of an organ between two Spanish province capitals (which reduces the number of cities to 50). This is against the idea of having a European organization for transplants. The second one is that the time needed to change from one mean of transport to another has not been taken into account, even if it is determinant in reality.
- **Medical Teams Scheduling.** The multi-agent system described in [12] is capable of finding, at the recipient hospital, an operating theatre in which the transplant operation can be performed and a medical team (doctors, nurses, anaesthetist) that will be assigned to perform the operation. The structure of the system is similar to the previous case.

4.2 The System Proposed

The solution we propose, and to which this project gives contribution, is based on the observation of the procedures currently followed by medical personnel to organize a transplant. The aim is to substitute the tasks that could be performed automatically with multi-agent systems capable of reproducing the behaviour of humans. For simplicity, as a first realization, we consider to work just on a national (Swiss) level. This will allow us not to take into account international rules and laws and will make the matching easier. The process is mainly divided into five parts (Figure 3):

- Registration of recipients. Every hospital which has at least one patient in its waiting list and wants to participate to the assigning procedure has to register to the system. The registration has both the function of subscription and of information storing in the organ transplantation mechanism database.
- Registration of donors. All the information about donors that are necessary to perform the matching must be available in the system. The registration of donor in this case has the function of submitting information to the system.
- Automatic Matching. When a donor registers to the system, the process of matching starts. This is a delicate step in which the best matching between donor and recipient has to be found. The system is provided with mechanisms to follow, rules to respect and constraints to satisfy, all mapped from the human decision criteria. In this way the system can come up with a solution, which can be either the “best” recipient given the conditions imposed or a list of possible recipients to be analysed by a human medical expert.
- Supported Decision. Organ transplantation management is a delicate problem and one can not expect to completely avoid the involvement of human beings. Even with good software tools, the solution reported after the matchmaking phase should always be checked by a medical expert who can agree and give his opinion on the choice.
- Supported Coordination. When choosing a recipient, the logistical phase of transportation and medical teams arrangement has to be taken into account. Also in this case, software tools can come up with a solution that a human coordinator will verify before proceeding.

The most complex part of this solution, both from the point of view of the theory behind it and from the point of view of its practical realization, is the process of matchmaking. There are different possible ways of performing matchmaking operations in a multi-agent system. The choice has to be made with regard to the desired level of involvement of human beings and on the basis of the correctness and reliability properties of the system. As said before, the two main possible approaches are:

- have a matchmaker agent that verifies, one by one, the matching of the donor with all the potential recipients. This is undoubtedly a good solution if applied in systems with a small amount of potential recipients, but turns out to be not scalable when trying to apply to bigger environments (international case).
- let the agents who represent the hospitals (recipients) negotiate among each other to find out who can gain the organ. This solution raises a great number of problems such as how to regulate negotiation and how to ensure a fair behaviour of the agents.

In this project, we have decided to investigate the second aspect described above. Our belief is that the best results can be achieved leaving the hospitals free of competing to get the organ. If the

competition is well organized, all the requirements of fairness can be respected. Moreover, this scenario strongly resembles the actual reality, where doctors present their candidate and try to prove that they deserve to get the organ more than the others. Also, with this approach, we exclude the presence of an entity like the matchmaker, which in the first solution is charged with too much power and too much responsibility to be left alone to decide without any control.

The following chapters will describe in detail the mechanism we have thought about and the results we got working on it. They will also explain the theoretical problems we had to solve and the solutions we decided to adopt. Eventually, in the last Section, a description of the obtained results and their possible interpretation is provided.

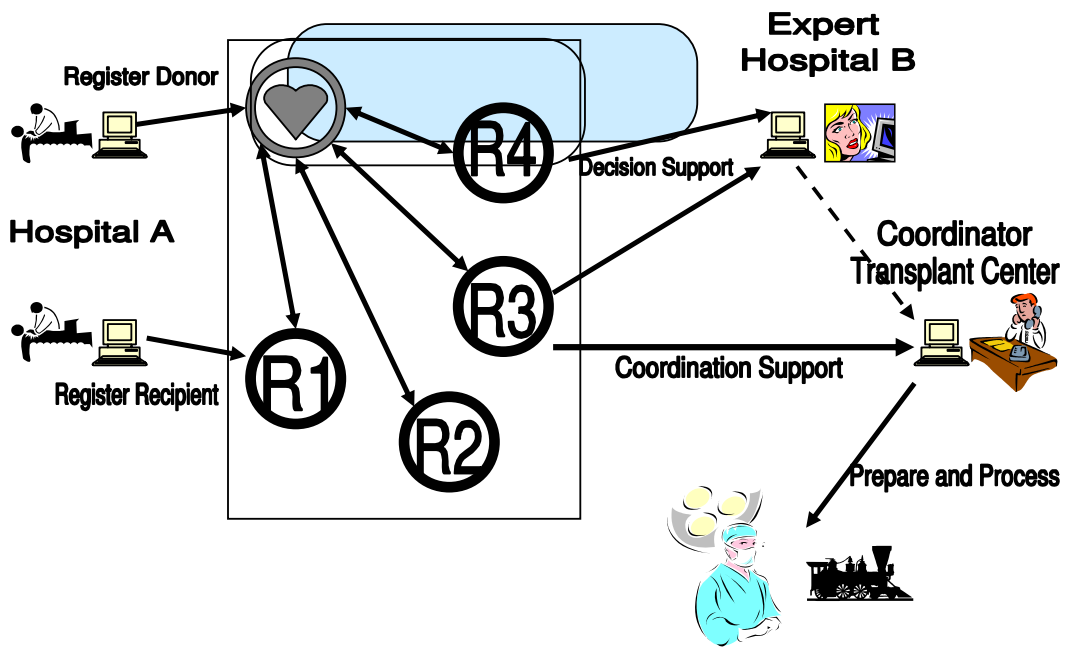


Figure 3

5. The Problem of Establishing the Winner

After estimating the level of matching of the organ and the possibilities and costs of the procurement phase, it is necessary to establish who, among the agents, will eventually gain the organ. This is the most delicate phase of all the process, since a wrong decision could lead to the waste of the organ and to the loss of human lives.

The decision to automate this part of the process must be taken very carefully. It is of extreme importance to establish a mechanism which guarantees the best results and the best level of fairness in the vast majority of the cases.

5.1 An Example in the Organ Transplantation Domain

At the present time, the only example in the domain which can be found in literature is relative to the Spanish model Carrel [19, 20], already described in Section 4.1. Although the description of the process of assignment is quite obscure and the most relevant parts concerning the decision are missing, we try here to illustrate it.

In Carrel, the negotiation between the agents takes place in what they call “Negotiation Room”. The manager agent of the Negotiation Room analyses each agent’s organ query and assigns a subset of the pieces (organs) offered in a Central Data Base that matches its query. Each piece offered to an agent has associated both an estimated procurement cost and an estimated distribution cost, so that the agent can use these data when deliberating. When an agent has received an offer list, it can start deliberating about it, evaluating each organ accordingly to the matching of its needs and the organ characteristics using its own selection function. The agent is free to reject all the pieces it dislikes. Once the agent has performed its deliberation, the assignment process takes place. This process uses a weighted graph that relates each agent with its pieces. The evaluation provided by the agent for each one of its pieces is used as the weight of each arc. The claimed goal of the assignment process is to achieve a maximum satisfaction with a minimum cost for each agent. However, this goal and the way it is achieved are not fully clear just following the description provided. Too many details are missing, like what happens if more than one agent gives the same evaluation for the same organ, or on which base the manager agent decides which subset of the available organs to propose the querying agents.

Moreover, this example suggests a model extremely dependent on a centralized part (the Negotiation Room and its manager agent). This can be a disadvantage because it requires communication and strong cooperation between the parties (querying agents and manager agent). Also, the matching operations are performed more than once, first by the manager agent and subsequently by the requiring agent itself. This can lead to a relevant waste of time and can be determinant for the result of the entire process.

5.2 Our Idea

The solution we present for solving the problem of establishing the eventual recipient of an organ moves towards a different horizon from the previous example. Our idea is to diminish the influence of third parties and let the agents compete among themselves to gain the organs. This could lead to a better use of the resources, since querying agents do not have to communicate frequently with the centralized part, and to the saving of precious time, since every agent performs the matching evaluation among its patients, the matching procedure is executed just once for every organ, and the singular evaluations can be done in parallel by different hospitals.

A possible realization of these concepts is through the use of a certain appropriate kind of auction. All the agents representing the hospitals which would like to participate in the assignment of organs register to an entity, which we call Auction House. Whenever there is a possible donor, the agent

representing the donor registers as well to the Auction House and communicates all its characteristics. Subsequently, the Auction House advertises the organs which have become available. At this point, each hospital evaluates the matching of the characteristics of the organs with the characteristic requirements of the patients, which are inserted in its private waiting list. If a hospital finds a patient who can be compatible with the donor, it tries to evaluate the possibilities and the costs of transport routes planning and medical teams scheduling (see Section 3). Finally, taking into account the results obtained, it decides how much it is worth to bid to get that organ during the auction.

When all the subscribed hospitals are done with the evaluation and matching parts, the auction can take place. The hospital which thinks the most it is worth for its patient to be transplanted wins the organ and is allowed to proceed with the operation. All the process is done under the supervision of an Observer, which is in charge of checking that the basic medical issues relative to the matching between the donor and the recipient are respected. After some time from the transplant operation, the Observer evaluates the result of the transplant. If the recipient is good and there have been no complications, it assigns to the hospital a positive comment and a reward. If there are complications, the patient is sick or, even worse, dead, it assigns to the hospital a negative comment, and of course no reward.

A graphical description of the system is given in Figure 4.

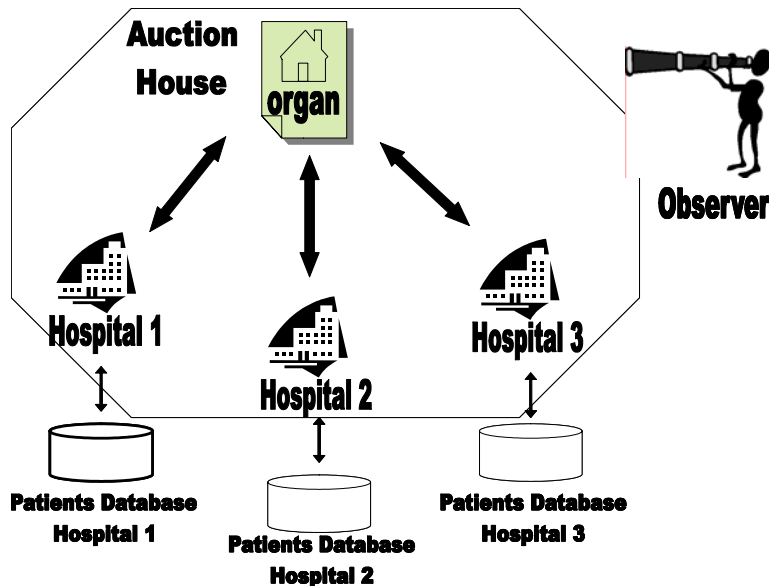


Figure 4

5.3 Expected Scenario

What we imagine as a scenario for what we described is that, like everywhere in real life, in our system there will be hospitals which are able to evaluate correctly and realistically the probabilities of performing a good transplant, and hospitals which, for corruption or real incapacity, will not be able to do so. What we expect from the proposed mechanism of payments and rewards is a sort of “natural selection”. The hospitals which can precisely evaluate their possibilities of succeeding will bid the appropriate amount; therefore, if they win the auction and get the organ, they will most likely succeed in the operation and get the reward. Just in very few, unlucky cases they will have to pay the penalty for not succeeding. For the hospitals which are not precise or correct, the situation will be exactly the opposite. If they will win the auction overbidding on an organ, while having low probabilities of good results, it is likely that they will perform a bad operation and pay for the failure. In this way, what we expect to happen is the precise hospitals to gain more and more, while the other hospitals lose more and more. With respect to this, we can imagine that the evaluation function of the bid amount will take also into account the current “capital” of the hospital, so that a hospital can not bid more than the “capital” it has and can not end up with a negative balance.

5.4 Problems Arisen

The depicted scenario presents two major problems:

- **Auction.** One of the targets of the system proposed in the last paragraph is the possibility to separate hospitals with good capacities of success probability evaluation from the others. To succeed in this target, it is fundamental that the hospitals use their real evaluation when placing a bid. In fact, just if they express their true beliefs the mechanism of payment and rewards can be effective and put in evidence the “good” hospitals. Therefore, it is determinant to find an auction mechanism which ensures truth telling as a dominant strategy.
- **Currency.** Organ transplantation deals with life and death of human beings. It is against every moral and every law to involve money in this process. Therefore, if we want the agents to be able to compete in the auction, and if we want the mechanism of payments and rewards to be effective, it is mandatory to find a “currency” which is as powerful as money in determining the behaviour of the agents and which can motivate them to compete.

5.5 Solutions Adopted

As we have seen in the previous paragraphs, we want to use an auction mechanism to allow (the agents representing) the hospitals to compete and gain organs. The problem which rises is how to force the hospitals to evaluate correctly their probabilities of success and compete in the auction in a fair way which guarantees the safety and the respect of the patients.

For this purpose, let us examine the most common kinds of auction mechanisms and decide which one is more appropriate to reach our target [9, 15].

- **First Price Auction:** each participant submits his bid in a sealed envelop; the agent with the highest bid wins the object and pays his bid, the others pay nothing.
- **Second Price Auction:** each participant submits his bid in a sealed envelop; the agent with the highest bid wins the object and pays the second highest bid, the others pay nothing.
- **English Auction:** there is an initial reservation price and at each time every player can increase the bid publicly; the auction is over if for a certain time period (fixed in advance) no one increases the bid.
- **Dutch Auction:** the auctioneer initially calls for a very high price and then continuously lowers the price until some bidder stops the auction and claims the good for that price.

Usually, the decision regarding the most appropriate auction method is left to the auctioneer, who chooses the auction mechanism that maximizes his revenue in equilibrium [11]. In our case, the auctioneer does not have any direct interest and it is possible to say that what has to be maximized is the social welfare of the society and the patients. In fact, the Auction House is a kind of “super partes” institution which does not gain anything from running the auction mechanism. We can imagine the Auction House as being an organization like Swisstransplant or Eurotransplant. The real interest in organ transplantation management is in treating patients and giving them a concrete possibility of life. Therefore, the auction mechanism must be such that the winner agent is the one representing the hospital which has the highest probability of performing a successful transplant. To reach this target, it is fundamental that the (agents representing the) hospitals truly express their estimated probabilities of succeeding, so that the organ can be given to the most suitable one.

In such a need, it is mandatory to adopt an auction mechanism which ensures incentive compatibility. An auction is called incentive compatible if all the participants are rationally motivated to reveal the truth about their valuations. Specifically, when considering the auction as a game, the truth is the dominant strategy, and this also frees bidders from strategic considerations. This is important since the hospitals do not have to worry about other hospitals’ possible behaviour and do not have to waste their time in game theory concerns [22]. Although incentive compatibility is nowadays one of the main topics in computerized settings [16], the first one to argue about its importance was Vickrey [21], who suggested for this purpose the use of the second price auction in his paper of 1961. Since it is easy to prove that the second price auction is actually the mechanism which ensures truth telling and incentive compatibility, we have chosen to adopt it in our solution.

The other problem raised in the previous paragraph is the fact that in the system it is not possible to use money. It is therefore necessary to find an alternative “currency” that allows the hospitals to compete in the auction, but at the same time forces them to behave correctly.

One of the first ways of payment and reward we thought about was a sort of adaptation of the Clarke Tax concept [5], and implied the exchange of services between the hospitals. For example, if a hospital wins an organ but it does not succeed in the transplant, it has to give free services (i.e. free x-rays) to some of the other hospitals (i.e. the second highest bidder) as a penalty for its mistaken evaluation and performance. However, this method would have still involved, even though in a more hidden way, the concept of money and economical interests. In the delicate field of organ transplantation, it is necessary to find a method which can be completely disjoint from economy, but still as much powerful as money for the results it can obtain.

Nowadays, one of the emergent topics in Artificial Intelligence and in multi-agent systems is the concept of reputation. Given that it is a common habit to think about and represent agents as humans, scientist have started applying them characteristics which are typical of the human world. One of these characteristics is precisely reputation. Every agent has, gains or loses a certain amount of reputation accordingly to its behaviour in the system. If an agent respects the rules of the system and the commitments it has with other agents or with external entities, it gains reputation and it is seen by the other agents as somebody who is worth to trust and cooperate with. On the contrary, if an agent does not respect the system rules or its commitments, it loses reputation and it is seen as somebody who is better to avoid. [1, 8, 13, 17, 18]

Following the new way of thinking, we have thought about using these concepts in our system also. Reputation is therefore the “currency” used by (the agents representing) the hospitals to bid, win and pay the organs in the auction mechanism. Reputation is also the way to give reward to the hospitals which have performed good and successful transplants. Thus, reputation becomes the key parameter to classify the hospitals in the system. A hospital with a high reputation is one which is likely to perform successful transplants, therefore patients think it is worth to subscribe to and be represented by. A hospital with a low reputation is one which, for a lot of different reasons, does not usually succeed in the transplants; therefore patients prefer not to be treated in such hospitals. As we can see, at the end economic aspects are always present, since hospitals receive financial contributions, from the government or from the insurances, for every patient they treat. However, in this framework, the economic aspects are far from the organ transplantation itself, and are

introduced just as a consequence of the good behaviour of the hospitals. We can therefore conclude that it is good, since hospitals need to get somehow an income for their performances.

6. The System Realized

After having described, in the previous chapters, the general structure of our idea and the choices made in it, this chapter will illustrate the system realized to represent and to simulate it. The simulations were carried out to verify if the general behaviour we expected from the system was correct. The results, which will be reported and commented in Chapter 7, gave us some surprises and actually show that the mechanism works just under particular conditions.

6.1 Conventions in the Future Descriptions

As it is common use, when speaking about agents, to address them as real persons, in this section we define the conventional names with which we refer to the different agents involved in our system.

- **Auction House:** agent in charge of organizing and managing the auction, registering the participants, collecting the payments;
- **Hospital:** agent that represents a hospital interested in participating to the transplant auction;
- **Donor:** agent that represents an organ donor or its hospital;
- **Transplant Observer:** agent which controls the compatibility between winner recipient and donor, declares the success or the failure of a performed transplant and gives rewards.

6.2 Framework

In a first phase, all the Hospitals which have at least one patient in their organ transplant waiting list subscribe to the Auction House. Whenever there is a new Donor, it subscribes to the Auction House and reveals all the information about the organs that it intends to donate. Let us suppose here, for simplicity of description and representation, that the Donor gives just one organ. In the case of multiple organs, the description that follows is repeated in parallel as many times as the number of organs available. After registering the Donor, the Auction House advertises the organ which has become available and communicates the information about it to the subscribed Hospitals.

The advertised organ matches differently with each of the patients of different Hospitals. Hospitals can compute the level of matching of the organ with their own patients. We can suppose that each Hospital decides fairly the priority level of the patients in its own waiting list, therefore each Hospital draws a ranking of candidates from all its possible recipients. Following these rankings, the Hospitals try to organize and estimate the costs of transport routes planning and medical team scheduling. Eventually, each Hospital comes up with its “best candidate”, the potential recipient for whom there is the highest possibility of succeeding in organizing and performing the transplant. Some of the Hospitals are very precise and can really evaluate the possibilities of performing a successful transplant, while others are disturbed by the presence of noise in the values. In our implementation, all the procedure of matchmaking and cost evaluation inside each Hospital is hidden (not represented) and for our study we suppose that each Hospital comes out already with the probability of success associated to its best candidate.

The Hospitals (each one representing a patient) compete in a second price auction, ruled by the Auction House, to gain the organ. The winner Hospital performs the operation on its patient under the control of the Transplant Observer, which subsequently monitors the health of the patient and eventually decides if the transplant was successful or not. The decision is of course taken on the basis of the real probability of success of the Hospital in that particular case.

If the transplant is successful, the Hospital gets back the amount paid plus a reward. If the transplant is not successful, the Hospital loses the paid amount. However, in order to help the Hospitals which

have been unlucky (correct evaluation, but unsuccessful transplant) to recover and get back into the mechanism, a small refund is given to all the Hospitals after every round. Given that Hospitals deal with life and death of human beings, it is not allowed to use money in the system. Hospitals pay the organs with reputation, which is therefore the official “currency”.

6.3 Target

The target is to investigate how the mechanism works. The ideal result would be to reach a sort of “natural selection”: after a certain number of organ assignments, the hospitals with a good evaluation of the probability of success should have high reputation, while the others should have low reputation.

6.4 System Implemented for the Simulations

A defined number n of Hospitals is created. Each Hospital has an initial amount of *reputation* (equal for all the Hospitals) and an associated *noise* (randomly generated), which denotes its possibility to correctly evaluate the probability of success. Hospitals are then grouped on the basis of their amount of *noise*. Each group contains the same number of Hospitals. A number of organs to be assigned sequentially (rounds) is selected and the simulation begins.

For every round:

1. The Auction House advertises an organ. For each Hospital, the Transplant Observer randomly defines a *probability of success* associated to that particular organ. Then each Hospital computes the *noisy probability of success*, which is the probability of success that the Hospital thinks its best candidate has, given the advertised organ. The *noisy probability of success* is computed choosing a random value of a Gaussian curve with mean *probability of success* and variance *noise*. If the value of the *noisy probability of success* is smaller than zero, it is assigned the value zero. If the value of the *noisy probability of success* is bigger than one, it is assigned the value one.
2. Each Hospital places its bid. The bids are computed with respect to the *noisy probability of success* and the amount of *reputation* that the Hospital has in the specific round. The function to compute the bid is the same for all the Hospitals and it is represented in Figure 5. The idea is to make the Hospitals be risk-averse. With the same *noisy probability of success*, if the Hospital is gaining reputation or has exactly the amount it is supposed to have at that round, the bid will be a linear function of the *noisy probability of success*. If the Hospital is losing reputation, the bid will be lower. The mathematical representation of this function is:

$$bid_{agent_i} = MaxR * \left(NoisyProbSucc_{agent_i} \right)^{reputation_{i,k}}$$

where $reputation_{(i,k)}$ is equal to:

$$reputation_{(i,k)} = \begin{cases} \frac{reputationAtRound_k}{reputation_{agent_i}} & \text{if } reputation_{agent_i} < reputationAtRound_k \\ 1 & \text{if } reputation_{agent_i} \geq reputationAtRound_k \end{cases}$$

The value of bid should be in the interval [0,1]. However, given that after every round each Hospital gains a refund, the maximum value that a Hospital is willing to bid is

$$MaxR = 1 - reputation_{refund}$$

If the computed value of *bid* is bigger than the amount of *reputation* that the Hospital has, *bid* is set equal to the available *reputation*.

3. The Auction House collects the bids and, following the second price auction mechanism, evaluates the winner and the amount of reputation to be paid. In case of draw of the winners, the Hospital with the lower value of *noise* is considered to be the fastest and therefore is chosen. In this case, the second price is the same as the first one, so the amount of reputation to be paid corresponds to the bid of the winner.
4. The *reputation refund* is given to the all Hospitals (including the winner of the round).
5. The Transplant Observer randomly evaluates if the transplant was successful or not, according to the *probability of success* of the winner Hospital. For example, if the *probability of success* of the winner Hospital was estimated to be 70%, there are seventy possibilities out of one hundred that the Transplant Observer will decide for a positive outcome. If the transplant is successful, the winner Hospital is given back the paid amount and gains a reward equal to the paid amount. If the transplant is not successful, the winner Hospital loses the paid amount.




Figure 5

6.5 Variables

The variables in the system are:

1. Number of rounds (organ assignments);
2. Number of Hospitals;
3. Noise in the evaluation of the probability of success;
4. Groups;
5. Maximal probability of success;
6. Initial amount of reputation of each Hospital;
7. Reward given for a successful transplant;
8. Payment due to gain an organ.

7. Obtained Results

The following results have been obtained running the simulation program. After each round, all the reputation values are collected, and then the average of reputation for each group is computed.

The diagrams reported show on the *x-axis* the number of rounds, on the *y-axis* the amount of reputation.

The division in groups is made on the basis of the value of noise of each Hospital. Hospitals in Group 1 have a value of noise between 0.0 and 0.1, in Group 2 between 0.1 and 0.2 and so on:

- Group 1 [0.0; 0.1],
- Group 2 [0.1; 0.2],
- Group 3 [0.2; 0.3],
- Group 4 [0.3; 0.4].

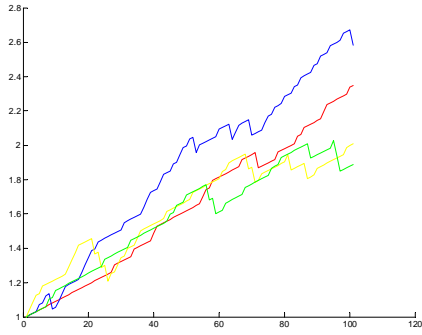
The mapping of the colours is:

- Group 1 **red**
- Group 2 **blue**
- Group 3 **yellow**
- Group 4 **green**


The values of the parameters used to compute each result are shown under the correspondent diagram. The study was initially made varying the Number of Rounds, the Number of Hospitals and the Maximal Probability of Success. Then, examining the obtained outcomes, it was decided to intervene also on the payment due to gain an organ and on the reward given for a successful transplant. The description of the reasons and the meaning of this decision is given later in Section 7.3.

7.1 First Set of Results

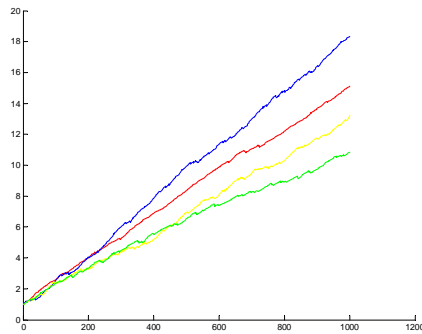
The first set of result was obtained for a number of Hospitals of 40 (10 Hospitals into each group) and 400 (100 Hospitals into each group). The number of rounds was 100, 1000 and 10000. The Maximal Probability of Success was 1.0 and 0.8.




Number of Rounds: 100
 Number of Hospitals: 40
 Probability of Success: 1.0




Number of Rounds: 100
 Number of Hospitals: 400
 Probability of Success: 1.0




Number of Rounds: 1000
 Number of Hospitals: 40
 Probability of Success: 1.0




Number of Rounds: 1000
 Number of Hospitals: 400
 Probability of Success: 1.0




Number of Rounds: 10000
 Number of Hospitals: 40
 Probability of Success: 1.0




Number of Rounds: 10000
 Number of Hospitals: 400
 Probability of Success: 1.0




Number of Rounds: 100
 Number of Hospitals: 40
 Probability of Success: 0.8



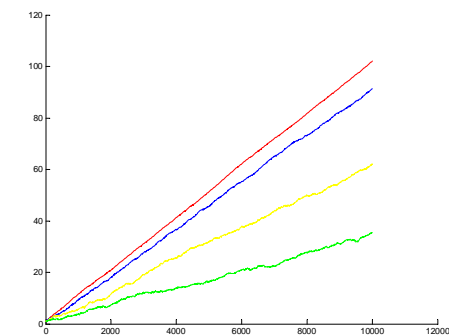
Number of Rounds: 100
 Number of Hospitals: 400
 Probability of Success: 0.8




Number of Rounds: 1000
 Number of Hospitals: 40
 Probability of Success: 0.8



Number of Rounds: 1000
 Number of Hospitals: 400
 Probability of Success: 0.8



Number of Rounds: 10000
 Number of Hospitals: 40
 Probability of Success: 0.8



Number of Rounds: 10000
 Number of Hospitals: 400
 Probability of Success: 0.8

7.2 Interpretation of the First Set of Results

Unfortunately, the results obtained are very far from the expected and desired results. In most of the cases, the differences in reputation between the groups are negligible and sometimes Hospitals with

high values of noise have more reputation than others with smaller values. One interpretation of these results can be the following.

The original idea was to make the reputation of good Hospitals (Hospitals with a precise evaluation of the probabilities of success) grow high, while the reputation of bad Hospitals (Hospitals with a not precise evaluation of the probabilities of success or Hospitals who are trying to cheat) was supposed to decrease or, at most, increase very slowly. It was expected that the good Hospital could bid more conscientiously; avoiding losses and winning rewards in most of the cases, and so increase their reputation. While bad Hospitals were expected to bid indiscriminately and lose reputation by failing in performing transplants with low probability of success. The bidding function was chosen in order to represent a risk-averse behaviour of the Hospitals with respect to the amount of reputation. Such function was expected to enforce and speed up the mechanism, giving the possibility to the Hospitals with high reputation to bid more than the others.

This does not happen in reality, with the model realized and the parameters used in it. The explanation lays in the fact that, unlike what was expected, most of the times bad Hospitals win the auction against good Hospitals by bidding high values of reputation and then they perform a successful transplant, gaining the reward.

At the beginning, all the Hospitals have the same amount of initial reputation, so they all bid accordingly to the linear function with respect to their own *noisy probability of success*. Let us suppose (Figure 6) there are two Hospitals, one very good (red) and the other very bad (blue), which have the same *probability of success*, say 0.8. The good Hospital is really precise and sees a *noisy probability of success* equal to 0.85. The bad Hospital sees a *noisy probability of success* equal to 0.98. Therefore, in the auction phase, the bad Hospital will place a bid higher than the good Hospital and will win the organ. However, the probability that the bad Hospital will perform a successful transplant is pretty high and it is likely that the bad Hospital will gain reputation from its behaviour. This situation happens in the most of the cases during the simulation and prevents the good Hospitals from differentiating from the others.




Figure 6

Sometimes, slightly better results are obtained by lowering the maximal probability of success. In these cases, even if bad Hospitals win, the probability that they succeed is lower than what they expect and it is possible that they perform a not successful transplant, losing reputation. While good Hospitals can better evaluate the values and the risks of failure connected to these, choosing to bid accordingly to them. However, the value of maximal probability of success that is necessary to adopt in order to reach the desired results is fair too low (around 0.6, depending on the values of the other parameters) and does not really correspond to the values which can be observed in the real world.

The diagrams reported above show also another anomaly in the behaviour against the expectation. If the maximal probability of success is equal to 1.0, the system with 40 Hospitals produces results opposite to our desires and bad Hospitals overtake in reputation good Hospitals, while the system

with 400 Hospitals maintains the supremacy of good Hospitals on the others, but with a very short margin. If we lower the maximal probability of success to 0.8, the situation reverts and we notice a correct behaviour of the system with 40 Hospitals, even though with a very short margin, while in the system with 400 Hospitals all the Hospitals proceed together through the rounds without ever differentiating the ones from the others. This shows that the system is not scalable with respect to the number of agents involved and can be partially explained considering that, when the number of Hospitals is high, the probability of each Hospital to win the auction is low (too spread), therefore good Hospitals do not have enough possibilities to show their abilities. Moreover, lowering the maximal probability of success can cause good Hospitals to be victims of unlucky cases.

The reported results show clearly that, at the beginning, the system it is strongly instable, and the stability points it can reach during the evolution are multiple. This suggested us to modify some other parameters in order to push the system towards the desired behaviour.

7.3 Suggestion

For the purpose of differentiating the good Hospitals from the bad Hospitals, the first phase (the beginning) is the most critical one. In fact, bad Hospitals tend to overbid with respect to their real probability of success and they often win the auction against good Hospitals. However, if bad Hospitals are lucky and succeed in the transplant, they increase their amount of reputation, maintain a risk neutral behaviour and carry on overbidding. If bad Hospitals do not succeed in the transplant, the amount they pay is too low, this allows them to continue overbidding and therefore winning against good Hospitals. So, the problem is that good Hospitals often do not have the possibility to win the auction and prove their superiority. The aim is to make bad Hospitals lose reputation at the beginning in a way to force them to assume a risk averse behaviour. At this point, good Hospital can start winning the auctions and prove they are effectively better. This mechanism should auto-improve itself: once that good Hospitals start winning auctions, they gain more and more, keep a risk neutral behaviour, always bid the amount is worth to bid, succeed in most of the cases and gain reputation. In this scenario, bad Hospitals lose reputation and assume a risk averse behaviour, they bid less and, if sometimes they win an auction because they overestimate the probability of success, it is possible that they perform an unsuccessful transplant and lose more reputation.

An idea to implement the depicted situation could be to penalize the winning Hospital asking to pay an amount which is a multiple of the *secondPrice*. If the Hospital performs a good transplant, it receives the amount back plus a reward. If it doesn't succeed, it has lost what it had paid.

The mechanism could also be improved by diminishing the *reward* given for a successful transplant. In this way, bad Hospitals which win by overbidding and subsequently perform successful transplants will not have the possibility to differentiate too much from the other Hospitals.

Therefore, we define:

$$\begin{aligned} \text{PaidAmount} &= \text{PayCoefficient} * \text{secondPrice} \\ \text{RewardAmount} &= \text{RewardCoefficient} * \text{secondPrice} \end{aligned}$$

with


$$\begin{aligned} \text{PayCoefficient} &> 1 \\ \text{RewardCoefficient} &< 1 \end{aligned}$$

7.4 New Results


The graphs which follow compare the results of the simulations for 40 and 400 Hospitals made with:

- *Pay* coefficient equal to 1, *Reward* coefficient equal to 1 and a low maximal probability of success (50%);
- *Pay* coefficient equal to 1, *Reward* coefficient smaller than 1 and a low maximal probability of success (50%);
- *Pay* coefficient greater than 1, *Reward* coefficient equal to 1 and a higher maximal probability of success;
- *Pay* coefficient greater than 1, *Reward* coefficient smaller than 1 and a higher maximal probability of success.


The number of rounds considered in these cases was 100000.



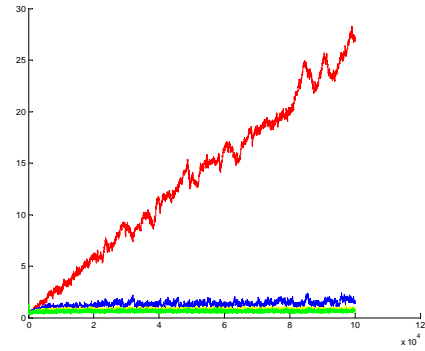
Number of Rounds: 100000
 Number of Hospitals: 40
 Probability of Success: 0.5
 Reward: 1.0
 Pay: 1.0




Number of Rounds: 100000
 Number of Hospitals: 40
 Probability of Success: 0.5
 Reward: 0.33
 Pay: 1.0




Number of Rounds: 100000
 Number of Hospitals: 40
 Probability of Success: 0.95
 Reward: 1.0
 Pay: 3.0




Number of Rounds: 100000
 Number of Hospitals: 40
 Probability of Success: 0.95
 Reward: 0.33
 Pay: 3.0




Number of Rounds: 100000
 Number of Hospitals: 400
 Probability of Success: 0.5
 Reward: 1.0
 Pay: 1.0




Number of Rounds: 100000
 Number of Hospitals: 400
 Probability of Success: 0.5
 Reward: 0.33
 Pay: 1.0




Number of Rounds: 100000
 Number of Hospitals: 400
 Probability of Success: 0.8
 Reward: 1.0
 Pay: 3.0




Number of Rounds: 100000
 Number of Hospitals: 400
 Probability of Success: 0.8
 Reward: 0.33
 Pay: 3.0




Number of Rounds: 100000
 Number of Hospitals: 400
 Probability of Success: 0.87
 Reward: 1.0
 Pay: 5.0



Number of Rounds: 100000
 Number of Hospitals: 400
 Probability of Success: 0.87
 Reward: 0.33
 Pay: 5.0



Number of Rounds: 100000
 Number of Hospitals: 400
 Probability of Success: 0.95
 Reward: 1.0
 Pay: 5.0



Number of Rounds: 100000
 Number of Hospitals: 400
 Probability of Success: 0.95
 Reward: 0.33
 Pay: 5.0

7.5 Interpretation of the New Results

From the diagrams shown in the previous paragraph it is easy to notice that the more the Hospitals have to pay for an error, the more good Hospitals differentiate from bad Hospitals.

If the *PayCoefficient* is equal to 1, the results obtained just lowering the maximal probability of success are not entirely satisfying. The margin between good and bad Hospitals is too narrow (40 Hospitals) and sometimes inexistent (400 Hospitals).

If the *PayCoefficient* is greater than 1, the results obtained are better. The *PayCoefficient* needed varies accordingly to the number of Hospitals we put in the simulation. For 40 Hospitals, a *PayCoefficient* of 3 is enough to obtain exactly the desired result, even with a very high maximal probability of success (95%). For 400 Hospitals, a *PayCoefficient* of 5 is needed to make the good Hospitals sharply differentiate from the bad Hospitals. However, if the maximal probability of success is very high (95%), not even the high price imposed prevents bad Hospitals from gaining more reputation than good Hospitals. In this case a solution is either to give a reward lower than the *secondPrice* (*RewardCoefficient* < 1) or to impose a lower maximal probability of success (87% is already enough).

A possible explanation of the reason why it is necessary to introduce the *PayCoefficient* to push the behaviour of the system towards the desired direction can be found considering the expected value of the reward of a winning Hospital A at round t.

$$\begin{aligned} E\left[reputation_A(t) - reputation_A(t-1) | A \text{ wins at round } t\right] = \\ = MaxProbSucc * RewardCoefficient * secondPrice(t) - \\ - (1 - MaxProbSucc) PayCoefficient * secondPrice(t) + \\ + reputationRefund \end{aligned}$$

In order to be worth for the Hospital to participate, it should be:

$$E\left[reputation_A(t) - reputation_A(t-1) | A \text{ wins at round } t\right] > 0$$

$$MaxProbSucc * RewardCoefficient * secondPrice(t) - (1 - MaxProbSucc) PayCoefficient * secondPrice(t) + reputationRefund > 0$$

Therefore, the maximum *PayCoefficient* so that there is an expected return is

$$PayCoefficient < \frac{MaxProbSucc}{1 - MaxProbSucc} * RewardCoefficient + \frac{reputationRefund}{(1 - MaxProbSucc) * secondPrice(t)}$$

In our case (maximal probability of success = 95%, *RewardCoefficient* = 0.33, *reputationRefund* = 0.01).

$$PayCoefficient < \frac{0.95}{1.0 - 0.95} * 0.33 + \frac{0.01}{(1.0 - 0.95) * secondPrice(t)}$$

Which implies

$$PayCoefficient < 6.27 + \frac{0.2}{secondPrice(t)}$$

secondPrice is always included between 0.0 and 0.99, but usually the value is around 0.9. If we take such a value, we obtain

$PayCoefficient < 6.3$

Therefore, if the $PayCoefficient$ we impose is smaller than the computed value, the Hospitals have an incentive to participate, since the expected reward they get in case of victory in the auction is positive.

However, this does not justify why the value of the $PayCoefficient$ has to vary accordingly to the number of Hospitals in the system. A better explanation should take into account also the probability of Hospital A of winning the auction at round t (which is instead related to the number of Hospitals in the system).

$$\begin{aligned} E[reputation_A(t) - reputation_A(t-1)] &= \\ &= ProbWinAuction_A(t) * (MaxProbSucc * RewardCoefficient * secondPrice(t) - \\ &\quad - (1 - MaxProbSucc) * PayCoefficient * secondPrice(t)) + \\ &\quad + reputationRefund \end{aligned}$$

$$E[reputation_A(t) - reputation_A(t-1)] > 0$$

$$\begin{aligned} PayCoefficient < \frac{MaxProbSucc}{1 - MaxProbSucc} * RewardCoefficient + \\ + \frac{reputationRefund}{ProbWinAuction_A(t) * (1 - MaxProbSucc) * secondPrice(t)} \end{aligned}$$

This value is impossible to compute without involving game theory concepts and complex considerations about the behaviour of the system to evaluate $ProbWinAuction$ and $secondPrice$. However, this equation helps to explain the variation of the $PayCoefficient$ that it is necessary to introduce: $PayCoefficient$ is inversely dependent on $ProbWinAuction$, which is expected to decrease when the number of Hospitals increases. So, the $PayCoefficient$ for which the Hospitals have interest to bid (expected positive reward) increases with the number of participants to the system

8. Conclusions and Future Developments

8.1 Conclusions

The project has required, during its first part, a deep studying of the problem of healthcare and organ transplantation management. The current methods of organ transplantation management have been analyzed and hypotheses have been made on the possible ways to improve them with software support systems. Existent tools in this area have been examined and criticized, pointing out their weaknesses and suggesting possible alternative ways. The focus has then been restricted to the problem of finding the most suitable recipient for an organ among a set of possible recipients represented by different hospitals. The problem has been approached allowing the hospitals (each one representing a potential recipient) to compete among themselves in a second price auction to gain the organ. Since, in all the countries, laws absolutely ban the use of money in the ambit of organ transplantation, reputation has been chosen as the official currency in the system.

The idea was to represent a possible and (under some points of view) “realistic” environment in which there are hospitals which are able to evaluate correctly the probabilities of performing a good transplant, and hospitals which, for corruption or real incapacity, are not able to do so. A mechanism of payments for gaining the organ and rewards to successful transplants has been set out with the intent of performing a sort of “natural selection” between the hospitals. The hospitals which could precisely evaluate their possibilities of succeeding were supposed to almost constantly increase their amount of reputation, while the other hospitals were supposed to lose reputation more and more.

The simulations of such scenario were rich of unexpected results. We observed that actually, if the maximum probability of succeeding in the transplant was high, bad hospitals were winning reputation by overbidding. Even lowering the maximum probability of success did not give better results. Moreover, the system showed up to be instable and strongly dependent on the events happening at the start. Therefore, we decided to try to differentiate good and bad hospitals from the very beginning by penalizing unsuccessful transplants asking for a payment multiple of the second price and rewarding successful transplant with a reward smaller than the second price. This method gave exactly the expected results and, after a very short bootstrap phase, good hospitals increased their reputation while all the others were left with low values. However, the payment coefficient to be introduced in order to obtain good results is not unique and varies with respect to the number of hospitals in the system: the more the hospitals, the higher the value.

8.2 Future Developments

The possible future subjects of study, in the context of this project, are multiple:

- The influence of the introduced penalization on the bidding strategy of the hospitals has still to be investigated and could lead to different results and a deeper understanding of the behaviour of the entire system.
- The idea of using reputation as “the official currency” to substitute money in auctions or negotiations could be extended and applied to other systems where the use of money is not appropriate or not allowed. For example, in the context of the academia, some decisions could be taken basing on the amount of reputation of (the agent representing) a professor, rather than on the amount of money in his bank account which would not be representative of his research skills.
- The possible presence of antisocial agents could be taken into account [2] and new simulations could be carried out to verify the changes that such agents could introduce in the system. Also, Vickrey auction would become inappropriate in such cases and it would be opportune to find new solutions to the problem.

Bibliography

- [1] Abdul-Rahman A., Hailes S. *Supporting Trust in Virtual Communities* 2000
- [2] Brandt F., Weiss G. *Antisocial Agents and Vickrey Auctions*
- [3] Decker K., Li J. *Coordination Assistance for Mixed Human and Computational Agent Systems* UMass Computer Science Technical Report 95-31 1995
- [4] Decker K., Li J. *Coordinated Hospital Patient Scheduling* 1998
- [5] Ephrati E., Rosenschein J. *The Clarke Tax as a Consensus Mechanism Among Automated Agents* 1991
- [6] Jennings N., Sycara K., Wooldridge M. *A Roadmap of Agent Research and Development Autonomous Agents and Multi-Agent Systems I*, 275-306 Kluwer Academic Publishers, Boston 1998
- [7] Kamel Boulos MN, Roudsari AV, Carson ER *Towards a Semantic Medical Web: HealthCyberMap's Dublin Core Ontology in Protégé-2000* Fifth International Protégé Workshop, Sowerby Centre for Health Informatics Newcastle 2001
- [8] Kraus S., Sycara K., Evenchik A. *Reaching agreements through argumentation: a logical model and implementation* 1998
- [9] Kumar M., Feldman S.I. *Internet Auctions* Technical report IBM 1998
- [10] Leroy G., Chen H. *Meeting Medical Terminology Needs – The Ontology Enhanced Medical Concept Mapper* IEEE Transactions on Information Technology in Biomedicine, Vol 5, No 4, December 2001
- [11] Monderer D., Tennenholtz M. *Optimal Auctions Revisited* 1998
- [12] Moreno A., Valls A., Bocio J. *Management of Hospital Teams for Organ Transplants using Multi-Agent Systems*
- [13] Mui L., Mohtashemi M. *Rational Decision Making Using Social Information* 2002
- [14] Ribes A., Valls A., Moreno A. *Planning Organ Transport Routes*
- [15] Sandholm T. *Automated Negotiation* Communications of the ACM, 1999
- [16] Sandholm T., Ygge F. *On the gains and losses of speculation in equilibrium markets* In Proc. of the 15th Int Joint Conference on Artificial Intelligence (IJCAI), 1999
- [17] Schneider J., Kortuem G., Jager J., Fickas S., Segall Z. *Disseminating Trust Information in Wearable Communities* 2000
- [18] Sen S., Biswas A., Debnath S. *Believing Others: Pros and Cons*
- [19] Valls A., Moreno A., Alonso H., Sanchez D. *A Multi-Agent System to Support the Assignment of Transplantable Organs*
- [20] Von Neumann J., Morgenstern O. *Theory of Games and Economic Behaviour* 1947

- [21] Vazquez-Salceda J., Cortes U. *Using Agent Mediated Institutions for the Distribution of Human Tissues among Hospitals*
- [22] Vickrey, W. *Counterspeculations, auctions, and competitive sealed tenders* Journal of Finance, 1961
- [23] Wooldridge M. *Agent-Based Computing* Baltzer Journals 1997
- [24] <http://umlsks.nlm.nih.gov>
- [25] <http://www.eurotransplant.org>
- [26] <http://www.swisstransplant.org>

