

Formalizing an Electronic Institution for the distribution of Human Tissues

J.Vázquez-Salceda ^a, J.A.Padget ^b, U.Cortés ^a, A. López-Navidad ^c,
F. Caballero ^c

^a*Departament de Llenguatges i Sistemes Informàtics. Universitat Politècnica de Catalunya. c/ Jordi Girona 1-3. E08034 Barcelona, Spain.*

^b*Department of Computer Science. University of Bath. Bath, BA2 7AY. United Kingdom.*

^c*Banc de Teixits. Hospital de la Santa Creu i Sant Pau. c/ St. Antoni M. Claret, 167. E00025 Barcelona, Spain.*

Abstract

The use of Multi-Agent Systems (MAS) in health-care domains is increasing. Such Agent-mediated Medical Systems can manage complex tasks and have the potential to adapt gracefully to unexpected events. However, in these kinds of systems the issues of privacy, security and trust are particularly sensitive in relation to matters such as agents' access to patient records, what is acceptable behaviour for an agent in a particular role and the development of trust both between (heterogeneous) agents and between users and agents.

To address these issues we propose a formal normative framework, deriving from and developing the notion of an electronic institution. Such institutions provide a framework to define and police norms that guide, control and regulate the behaviour of the heterogeneous agents that participate in the institution. These norms define the acceptable actions that each agent may perform depending on the role or roles it is playing, and clearly specifies the data it may access and/or modify in playing those roles.

In this paper we present the formalization of Carrel, a virtual organization for the procurement of organs and tissues for transplantation purposes, as an electronic institution using the ISLANDER [15] institution specification language as formalizing languages. We demonstrate aspects of the formalization of such an institution, example fragments in the language used for the textual specification, and how such formalization can be used as blueprint in the implementation of the final agent architecture, through techniques such as skeleton generation from institution specifications described in [39] and [37].

Key words: Autonomous Agents, Real-Time Systems, Multi-Agent Systems, Agent Mediated Institutions, Negotiation.

1 Introduction

Organ and tissue transplantation is a widely-used therapy against life-threatening diseases. But there are two issues that make transplantation management a very complex issue: (i) *scarcity* of donors, so it is important to try to maximize the number of successful transplants (ii) *donor/recipient matching*, because of the diversity and multiplicity of genetic factors involved in the response to the transplant. In this paper we propose an agent-based architecture for the tasks involved in managing the vast amount of data to be processed in carrying out

- recipient selection (e.g. from patient waiting lists and patient records),
- organ/tissue allocation (based on organ and tissue records),
- ensuring adherence to legislation,
- following approved protocols and
- preparing delivery plans (e.g. using, say, train and airline schedules).

The relative scarcity of donors has led to the creation of international coalitions of transplant organizations. This new, more geographically distributed, environment makes an even stronger case for the application of distributed software systems, but with the added complication of the necessity to accommodate a complex set of, in some cases conflicting, national and international regulations, legislation and protocols governing the exchange of organs and tissues. It is the last point which underpins our case for the use of so-called *electronic institutions* whose purpose is to provide over-arching frameworks for interaction in the same way as institutions, or equivalently, collections of social norms, do in the physical world.

As we will explain later, electronic institutions work with explicit norms[12, 4, 7, 11]. Expressing all the regulations and protocols in the form of computable norms—instead of hard-coding them so they are scattered throughout the logic of a program—not only admits a readier verification, both informal and formal, of adherence but also it gives the system the added flexibility of behaviour that it can be adapted in the light of regulatory changes (an event that is not uncommon).

1.1 Institutions

The first question to which we should provide an answer is "what is an institution?", while a second question is "why are institutions useful?".

Most human interactions are governed by conventions or rules of some sort, hav-

Email addresses: jvazquez@lsi.upc.es (J.Vázquez-Salceda),
jap@cs.bath.ac.uk (J.A.Padget), ia@lsi.upc.es (U.Cortés),
alopeznavidad@hsp.santpau.es (A. López-Navidad).

ing their origin in society (emergent) or the laws (codification of emergent rules) that society has developed. Thus we find that all human societies, even the least developed ones, have some kind of social constraints upon their members in order to structure their relations and simplify their interactions. Some of these constraints are quite informal (taboos, customs, traditions) while some others are formally defined (written laws, constitutions). Individuals absorb these rules into their behaviour so they become second nature, which is how the customs of one culture can create problems in the context of another culture. However, social anthropology is not the subject of this paper, we are simply drawing attention to the fact that collections of expected behavioural patterns are all around us, in shops, banks, conversations, lectures, clubs etc..

The economist and Nobel laureate Douglas North has analyzed the effect of this corpora of constraints, that he refers as to *institutions*, on the behavior of human organizations (including human societies). North states in [29] that institutional constraints ease human interaction (reducing the cost of this interaction), shaping choices and making outcomes foreseeable. By the creation of these constraints, either the organizations and the interactions they require can grow in complexity while interaction costs can even be reduced. Having established these institutional constraints, every competent participant in the institution will be able to act—and expect others to act—according to a list of rights, duties, and protocols of interaction.

As norms are in fact the elements that characterize institutions they are not only useful as norms to be followed, but also they enable people to recognize an organization (such as a particular company) as being an instance of a particular kind of institution, and then use these information to predict other norms that could be applicable.

Thus, the answers to the two questions with which this section began are that (i) an institution is a set of possibly interdependent norms (ii) using institutions simplifies interaction and creates confidence.

1.2 *Electronic Institutions*

Not surprisingly, an electronic institution is the modelling of an institution through the specification of its norms in some suitable formalism(s). The statement is simple, but the practice is less so, partly because of the wide variety of norms, so that the form of expression (deontic logic, temporal logic, modal logic, classical predicate logic, types, for example) for one level is not necessarily suitable or even usable for another and partly because this stratification then leads to a need to demonstrate coverage and consistency between the levels (see [31] for a further discussion of these ideas). Nevertheless, the fact is that the essence of an institution, through

its norms and protocols *can* be captured in a precise machine processable form and this key idea forms the core of the nascent topic of institutional modelling.

As we observed earlier, typically, many human institutions have evolved through millennia of interaction, but we are now starting to see the invention of new organizations (typically as business models) which are, in effect, human experiments with institutional structures. This practice is now being taken up, but with considerably more freedom of design, in the multi-agent systems community, where it is software agents that participate in the electronic institutions and which gain from a simplification of the interaction protocols and a confidence that a transaction will be properly concluded according to the norms of the institution in which they are participating.

It is the capacity of multi-agent systems (MAS) to distribute the work among several autonomous, rational and pro-active entities (the software agents) that has lead this technology to be considered the best option for the implementation of open distributed systems in complex domains such as electronic commerce or virtual organizations. But building a multi-agent systems is a very complex process, as its aim is to establish a society of autonomous agents that interact among themselves in order to achieve one or several goals. One perceived way to tame that complexity, which is largely the result of unconstrained interaction, is to create virtual organizations—instances of institutions—following the principle put forward by North that the performance of human organizations is improved through the creation of institutional frameworks.

We have perhaps been a little free with the use of the terms *organization* and *institution* and so to clarify our position, we depict our view of their relationship in figure 1. This expresses the idea that there is a collection of norms, which may be fairly abstract desires or wishes such as “fairness” or “equity”, implemented through policies, such that particular sets of norms characterize institutions and then that organizations may either be instances of institutions or a combination of a bespoke set of norms. Such an organizational model may then be copied and expanded or simplified, while another organization may establish new norms which may be added into the general collection.

In conclusion, we believe institutions as norms are advantageous because they:

- reduce uncertainty about other agents’ behavior inside the institution,
- reduce misunderstanding, as there is a common corpus of norms governing agent interaction,
- permit agents to foresee the outcome of a certain interaction among participants
- simplify the decision-making process inside each agent, by decreasing the number of possible actions.

Or, to summarize and take a more robust stance, the use of norms makes agents

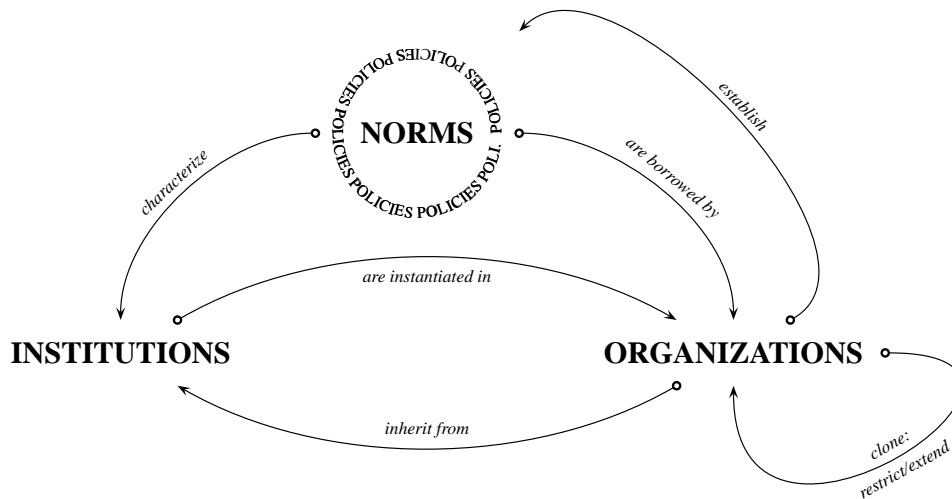


Fig. 1. Relationship between norms, institutions and organizations

more successful in the achievement of their goals despite their having bounded rationality.

1.3 Plan of this paper

In §2 we introduce the reader to the organ and tissue allocation problem, present an analysis of the problem and then sketch the design of the multi-agent system we have prototyped to support the allocation process.

The core of the paper is §3, where we show the process of formalizing the institution, by grouping interactions together to make so-called *scenes*, linking the scenes to create the *performative structure* and defining the norms that carry obligations across scenes, resulting in a precise and comprehensive description. §3.4 presents some illustrative fragments of the Carrel specification in the textual representation of ISLANDER and following this in §3.5 we present the resulting multi-agent system architecture. A prototype has been implemented in JADE, but we are also investigating other agent toolkits and the use of emerging tools for the automatic verification and generation of agent platforms [40, 39].

In §4 we appraise the related literature, present some conclusions and point to future research issues.

2 A case study in transplant organization

Modelling the transplant process cannot be done effectively as a *gedanken* exercise: it must be rooted in reality because there are so many actors, so many regulations

and so many factors to take into account. It is also essential to talk to the people involved to get a proper picture of the activities that take place.

We have been fortunate to work with the Spanish Organización Nacional de Transplantes¹ (ONT) and the Catalan Organització CATalana de Transplantaments² (OCATT) transplant organizations, which are statistically among the most effective and demonstrate the highest global volume of transplants per head of population, thanks to the creation and implementation of a network of well-coordinated hospitals and tissue banks, coupled to the definition of clear procedures for the distribution of organs and tissues and a high level of citizen awareness of the value of donation.

Since 1980 the number of requests for the application of transplant techniques has risen so much³ that the human coordinators—the people at the hospitals who act as the interfaces between the surgeons internally and the organ and tissue banks externally—are facing significant problems in dealing with the volume of work involved in the management of requests and piece⁴ assignment and distribution. Transplant-based therapies are the subject of much investigation and increasing application, such that demand for pieces may well rise rapidly in the near future. A review of the coordinator's role and the difficulties faced is presented by López-Navidad [23].

Consideration of the factors we have summarized above leads to the question of whether some kind of automation of the assignation process is desirable and if so, whether it is possible. In the design and construction of our prototype, we have applied agent technology to develop software which may assist specialists in making some of these decisions and an allocation of pieces that is acceptable according to the legislative requirements and other procedures governing the the process. To achieve this, the agents have to be given domain-specific knowledge (kinds of pieces, attributes to describe them, etc.) so they can act *rationally*⁵, and the rules they should follow, such as which actions can be done when, what information can be accessed or given out, etc..

In this kind of medical application the use of rules in order to guide the agents' behaviour is mandatory, as a mistake can lead to an unsuccessful transplant and potentially the death of a patient, as well as the waste of piece which might have better benefited someone else. A further complication is that all the agent actions must respect legislation on the distribution and use of pieces for transplantation. This need

¹ National Transplant Organization

² Catalan Transplant Organization

³ The continuous raise in requests is due, among other factors, to the introduction of new immunosuppressors which have significantly decreased rejection in recipients' clinical evolution.

⁴ From now on we will use the term *piece* instead of *tissue* or *organ*

⁵ In the sense of Simon and Newells' Bounded Rationality[27].

for regulation is one of the reasons why we have modelled the Carrel Multi-Agent System as an Electronic Institution [8]. And one of our hopes is that this choice will help to clarify the interaction among the different norms and regulations from the very basic, at hospital level, up to national and European levels.

2.1 *The organ and tissue allocation process*

We freely admit that our experience of practice is limited (to date) to the Spanish and Catalan organizations, but there is little doubt that they are both examples of best practice, and so constitute viable physical institutions on which to base electronic ones. The Spanish organizational model has two levels of action:

- *intra-hospital*: where the role of Hospital Transplant Coordinator was created to improve the coordination of all the people working at any step of the donor procurement, allocation and transplantation process
- *inter-hospital*: where an intermediary organization—OCATT for Catalonia, ONT for the whole of Spain—was created to improve the communication and coordination of all the participating health-care organizations, namely hospitals and tissue banks.

2.1.1 *Procurement*

In the case of organs the process starts at the *procurement stage*, when the members of the coordination team inside a certain hospital are made aware of a potential donor. A donor alarm is then sent to the ONT—except in Catalonia, when it is sent to OCATT. This alarm is signalled by telephone, and a human member of the staff has to list the basic attributes of the donor, including the results of clinical analysis, and a first evaluation of the organs and tissues that could be extracted is carried out. This first call is done as early as possible, usually when brain death of the potential donor is diagnosed.

2.1.2 *Search*

The next step is to search for suitable recipients. The intermediary organization (ONT, OCATT) carries out a recipient search for each organ that may be available by calling all hospitals with information about the pieces. To speed up this search process, protocols have been defined such as, looking first for recipients in the same hospital as the donor, then to hospitals nearby, then the same city, the same area, and lastly anywhere in the country. In the situation that there is no suitable recipient in the country, then an organ is offered to third countries, either inside or outside the EU.

2.1.3 *Emergencies*

In Spain, the ONT maintains a list of urgent cases containing all the recipients whose condition is life-threatening. If there is a suitable recipient for an organ in this list then it is accorded higher priority in the assignation process over all other recipients.

2.1.4 *Acceptance and delivery*

Once the recipient for a certain organ has been identified, a formal offer is then sent to the surgery team, that will implant the organ via the Hospital Transplant Coordinator. The team can analyze all the information available about the organ and decide whether they want to accept the organ or whether they think it is not well enough suited to their patient. If they accept the organ then the *delivery stage* starts, planning all the vehicles that will be needed (ambulance or helicopter from one hospital to one nearby, to the airport or to a train station; plane and/or train from one city to the other). Thus, delivery plans must take into account transportation system schedules.

2.1.5 *Time scales*

In the case of organs, time is one of the key issues, since they can only be kept outside a human body for a short time-span (hours), as all the preservation methods only can delay, but not stop, the decay process arising from their not receiving enough blood (ischaemia).

In the case of tissues, such as corneas, skin or bones, they can be preserved for longer periods—days, even weeks. Such relative resilience permits not only an effective preservation process, but also allows time for assessing the quality of the tissues and establishing the absence of bacteria and viruses. The consequence is a search stage that works in the opposite way to that for organs. With organs, the process is triggered when a donor appears, whereas in the case of tissues, it is the appearance of a new recipient that triggers the search through the tissue banks for a suitable piece for transplant. This search too is done by means of several time-consuming telephone calls from hospitals to tissue banks.

2.2 *The Carrel system*

For computer systems, medicine is one of the most difficult fields to manage, as it is extremely difficult to foresee all the conditions that may occur, leading to unexpected side effects when certain decisions or actions are performed in an unanticipated situation. As Fox & Das argue in [18], these are the kind of ill-defined

fields which have historically been the concern of Artificial Intelligence. They also identify software agents as having many characteristics which are appropriate for the implementation of medical applications.

An *Agent* is a computer program capable of taking its own decisions with no external control (*autonomy*), based on its perceptions of the environment and the objectives it aims to satisfy [41]. An agent may take action in response to changes in the environment (*reactivity*) and also take initiatives (*pro-activity*). A further attribute of agents is their ability to communicate with other agents (*social ability*), not only to share information but also to coordinate actions in order to achieve goals for which agents do not have plans they can fulfill on their own, solving even more complex problems. Agents can also exhibit other useful attributes such as *learning*, *rationality*, *mobility*, *benevolence* and *veracity*. A group of agents that interact among themselves and share some common objectives is called *agency*.

It is with these points in mind that we propose to use Agent technologies to automate part of the allocation process, finding suitable pieces for transplant and giving support to the decision making steps in that process. The over-arching goal is to improve the process of selection and procurement. A key principle of our solution is to deploy software agents to create, negotiate and coordinate plans for the extraction, transfer and implantation of pieces. To do so we follow the same scheme as was developed in Spanish, creating solutions for *inter-hospital* and *intra-hospital* levels. In this paper we focus solely on the tissue allocation procedure⁶ nor do we address any issues related to planning.

2.2.1 The inter-hospital level

At the inter-hospital level we have created the *Carrel Institution*, an agent platform whose behaviour could be briefly described as an *agency* that receives a piece request from one hospital and then tries to allocate the *best* piece available from all the known tissue banks. In this agency different entities (the agents) play different roles that are determined by their goals, rights and duties.

Figure 2 depicts all the entities that interact with the Carrel system, where TB denotes a tissue bank and UCTx denotes a transplant coordination unit. There are a) the hospitals that create the piece requests, b) the Tissue Banks, and c) observers such as the ONT [30] and OCATT. We make it a requirement in our model that all hospitals, even the ones that own a tissue bank, will make their requests through Carrel in order to ensure an acceptable distribution of pieces and to ease the tracking of all pieces from extraction to implant, in the same manner as ONT and OCATT require for organs.

The role of the Carrel Institution can be summarized in terms of following tasks:

⁶ Organ allocation is integrated into Carrel in [38].

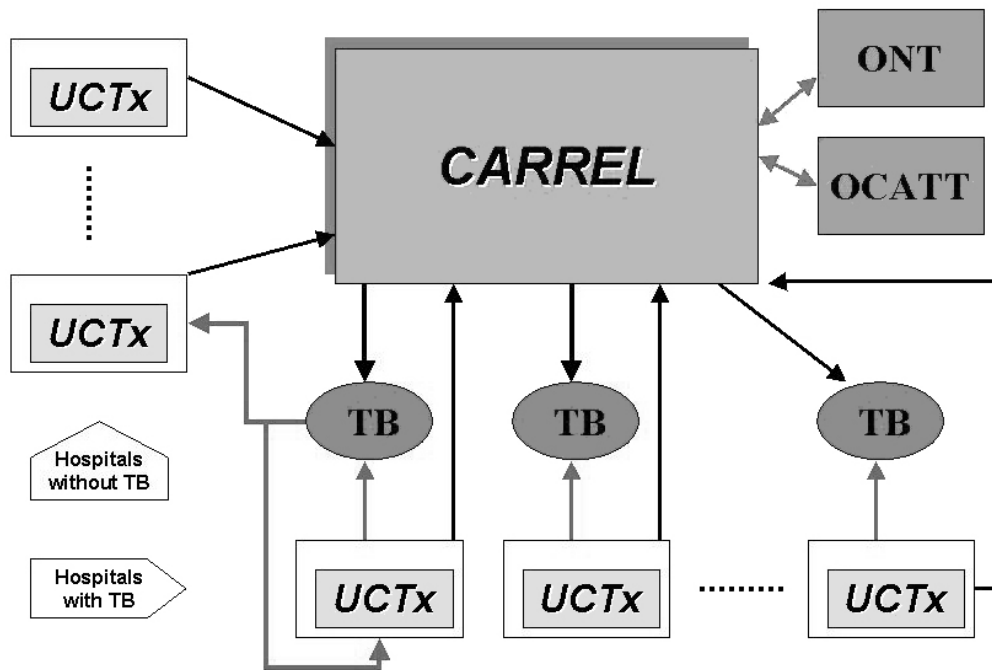


Fig. 2. The Carrel System environment: the Carrel Institution and the UCTx

- T1 ensure that all the agents which enter the institution behave properly—which is to say they follow the behavioural norms
- T2 acquire timely information on all the available pieces in the tissue banks
- T3 verify that all the hospitals and tissue banks fulfill all the criteria specified for interaction with Carrel
- T4 provide for external monitoring of the fulfillment of the commitments made within Carrel
- T5 coordinate piece delivery from tissue banks to hospitals
- T6 register all incidents related to a piece

A hospital becomes a member of the Carrel institution in order to make use of the services provided. In doing so, they are agreeing to accept the outcomes of the negotiation (assignment) process. Also, the agents representing the hospitals must respect the institutional norms. A hospital is represented in Carrel by the Transplant Coordination Unit (*UCTx: Unidad de Coordinación de Transplantes*). This agent serves as interface between the surgeons and Carrel. When a surgeon needs a piece he makes his request through the UCTx system, which analyzes the information entered by the surgeon, adds the information about the recipient and, finally, creates a *Finder Agent*, that is, the agent that goes to Carrel looking for a suitable piece.

The information required by the *Finder Agent* to look for a piece in *Carrel* is held in an electronic *Sealed Envelope*. The information contained in the envelope is summarized in Table 1. The *Selection Function* is the part of the information contained in the *Sealed Envelope* that allows the *Finder Agent* to perform a negotiation. It is composed of a set of rules, each one a constraint on the piece to be

Table 1
The envelope contents

<i>Urgency level</i>	that works as electronic postage stamp and sets the urgency level of the request (in Spain: normal, urgency-1 or urgency-0)
<i>Hospital identification</i>	a certificate issued by the Certification Authority associated with the Carrel institution [8], to allow the institution to authenticate the sender of each request and ensure that only <i>Finder Agents</i> with requests from authorized hospitals can enter and negotiate inside Carrel.
<i>Piece information</i>	type, parameters, etc.
<i>recipient data</i>	age, sex, laboratory analysis, etc.
<i>selection function</i>	see Table 2

Table 2
The selection function predicates

<i>piece</i>	predicates that describe the constraints the selected piece has to satisfy, such as the age of the donor or the dimensions of the piece itself.
<i>origin</i>	predicates that can set constraints about the tissue bank(s) preferred by the surgeon or the hospital
<i>cost</i>	predicates about the cost of the piece, such as price. Note: the cost is just that of extraction and preservation. Settlement is managed via a clearing house.

selected. Some of these rules may originate from the policy of the whole transplant unit of the hospital, but the others are introduced by the surgeon, who can set the constraints associated with a given recipient. Table 2 lists the kind of predicates a *Selection Function* can include.

2.2.2 *The intra-hospital level*

The functioning of the Transplant Coordination Unit may also benefit from the use of agents to help coordinate all the people in the hospital related to a particular case. Hence, our modelling of the UCT represents not only the surgeons, but also:

- the human transplant coordinator, who is the overall coordinator and who must coordinate all the tasks to be done and who must also be informed of any transplant related event
- any member of the hospital staff who plays a role at any step of the transplant process, from the moment when the piece arrives to the hospital to the moment when the piece has been implanted in the recipient.

The complete architecture is shown in figure 3. The role of the *Finder Agent* has already been described in the previous section.

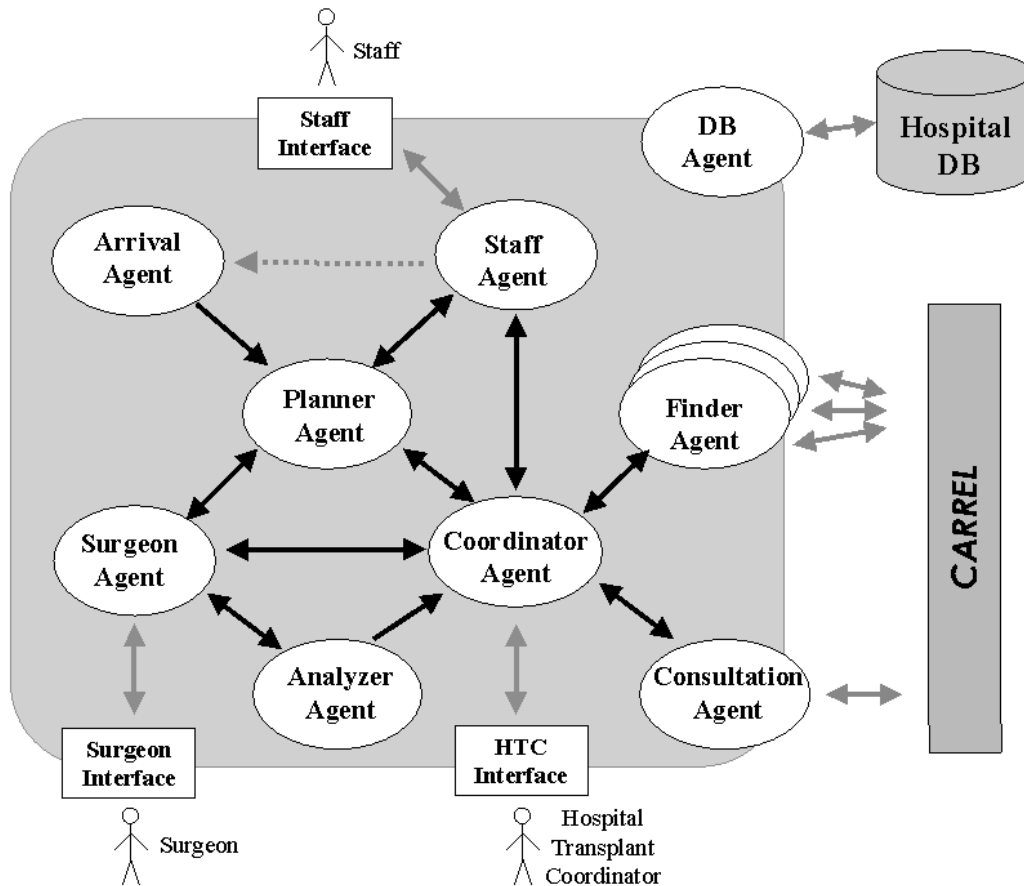


Fig. 3. The Transplant Coordination Unit Agency

Staff Agent: is a role for the representation of any member of the staff involved in transplants, with the exception of surgeons and the Hospital Transplant Coordinator, which are more specialized. The *Staff Agent* helps the people to plan their daily tasks, acting as their electronic diary and as a message receiver for communications from the Hospital Transplant Coordinator through the *Coordinator Agent*. The agent implementing this role also gives, for each person, access to the clinical information of patients in their care.

Surgeon Agent: is a role extending the *Staff Agent* functionalities for the purpose of use by surgeons. It communicates with the surgeons through the *Surgeon Interface* to collect and format the requests for pieces for transplant. Each request, as mentioned in the previous section, has to include the relevant information about the patient, the required piece, optional medical and economic restrictions and the *Selection Function*. The *Surgeon Agent* sends the request to the *Analyzer Agent*, which then checks if all the characteristics needed have been supplied, and if the values are consistent following a given protocol.

Coordinator Agent: responsible for the distribution and coordination of the different tasks that make up the whole process. The *Hospital Transplant Coordinator* in person can communicate with this Agent through the *Coordinator Interface*,

and control the behavior of the agent.

DB Agent: verifies that each access made by agents in *UCTx* to the clinical data in the Hospital Database is authorized, following a role-based access policy [21].

Planner Agent: is responsible for creating the transplant plan, that is, finding an operating theatre to match the arrival time of the piece and the surgeon's available schedule. The *Planner Agent* can send several proposals to the surgeon and the rest of the staff involved and tries to reach an agreement.

Arrival Agent: is responsible for updating the *Planner Agent* about events that can change the delivery plan, events that can occur while the transportation of the tissue is made from the tissue bank to the Hospital.

Consultation Agent: acts as the interface to the institution's database, and serves to process the different types of queries sent by the *Surgeon Agent*, the *Planner Agent* or the *Coordinator Agent*. This is done, as with Finder Agents, by going to the *Carrel* institution.

For more details about the *UCTx* system, including a more detailed description of each agent and the specification of the messages used, see [9].

2.3 Security considerations

Transplant information is considered high-risk data as it includes sensitive information about people (donors and recipients). So the *UCTx* and *Carrel* systems have to observe the local, national and European Union legislation on transplants (see the reports of the ONT in [10] and the recommendations of the Transplant Experts Committee in [24]). It also follows the European directives and the Spanish law on personal data protection [22] [13] [32]. In particular, both have to ensure confidentiality, privacy and integrity of patient and donor data. This is a long-standing issue in health care that acquires new facets with the use of Electronic Medical Records (EMR). One of the benefits of using electronic records is that it assures access for authorized and authenticated users as well as tracking access as demanded by law. It is important for patient and donor trust that it can be demonstrated that the information about them may not be used for any purpose beyond that for which it was collected.

3 Formalizing the Carrel Institution

We will focus now on the *Carrel* Institution, as it is the one that performs the mediation among hospitals and tissue banks in the transplant process.

To give a formal description of the *Carrel* system we will follow the ISLANDER formalism described in [15]. This formalism views an agent-based electronic insti-

tution as a type of *dialogical system* where all the interactions inside the institution are a composition of multiple dialogic activities (message exchanges). These interactions (called *illocutions* [28]) are structured through agent group meetings called *scenes* that follow well-defined protocols. A second key element of the formalism is the notion of agent *role*. Each agent can be associated to one or more roles, and these roles define the scenes the agent can enter and the protocols it should follow. The *scene protocols* are defined as multi-role conversational patterns.

3.1 *The performative structure*

The scenes and the connections between them constitutes the *performative structure*. This is a network of scenes that defines the possible paths that each agent role may take through the institution. In accordance with its role, an agent may or may not be permitted to follow a particular path through the performative structure, and ultimately, may be required to leave the institution. Because of the need to model the system accurately and verifiably, we use formal specification techniques which then permit the automatic verification of properties by means of theorem-proving (future work) and model-checking (work in progress) and finally lead to automatic generation of validated agent skeletons, representation of ontologies and the actual institutional structure.

In the Carrel System, the interaction among the agents and the institution is structured using the following scenes:

Reception Room: where all the external agents should identify themselves and where they are assigned the initial roles they are authorized to play. Any piece requests carried by arriving agents are validated at this stage.

Exchange Room: where agents carrying piece requests start negotiation with the institution for the assignation of a piece.

Confirmation Room: where the provisional assignments made in the *Exchange Room* are confirmed or withdrawn because of the arrival of another request with higher priority. Subsequently, in case of the confirmed ones, a delivery plan is prepared.

Consultation Room: where the institution is updated about any event or incident related to a piece. Agents coming from tissue banks inform the institution about piece availability, while agents coming from hospitals inform the institution about the reception of the piece, the transplant and the progress of the recipient.

The next step in the formalization process is to define the roles that may be played in the Carrel system. There are two kinds of roles: the *external roles* (roles for incoming agents) and the *institutional roles* (roles for agents that carry out the management of the institution). The external roles are the following:

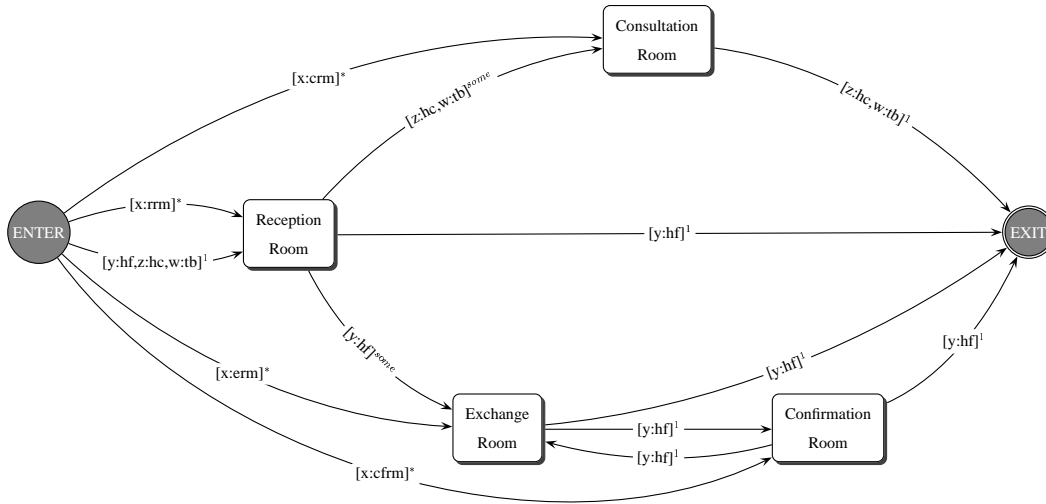


Fig. 4. The Carrel Institution performative structure

Hospital Finder Agent (hf): agents sent by hospitals with piece requests to be served by the Carrel system.

Hospital Consultation Agent (hc): agents sent by hospitals to keep the Carrel system updated about any event related to a piece. They can also perform queries on the Carrel database by means of the DB agent (see §3.5).

Tissue Bank Notifier (tb): agents sent by tissue banks in order to update Carrel about piece availability.

The institutional roles consist of one agent to coordinate all the scene relationships, the **Institution Manager (im)** and one agent to manage each scene: **Reception Room Manager (rrm)**, the **Exchange Room Manager (erm)**, the **Confirmation Room Manager (cfrm)** and the **Consultation Room Manager (crm)**

With all the scenes and roles identified, the performative structure can be drawn (see figure 4). Nodes are the scenes listed above plus *enter* and *exit* nodes in order to define beginning and ending points in the diagram. Arcs are labelled with tags of the form $[variable:role]^{access}$, where *variable* is an agent_{*i*}, *role* is one of the identified roles for the Carrel system and *access* is either the number of instances of agents playing such role that can enter together into the scene, or * meaning that the agent is the one that creates the scene. The diagram in figure 4 shows, for instance, that scene managers go directly from the *enter* point to the scene they should manage (the * means, as mentioned before, that they are the ones creating the scene), while all the external agents must proceed through the *Reception Room* scene in order to be registered and then be directed to the proper scene according to their roles.

The performative structure is also useful as a blueprint for the final agent architecture (see section §3.5).

3.2 Describing the scene protocols

Once the performative structure is defined, the next step is to define the interactions between the agents within the scenes. We do this by means of a *scene protocol*, which defines the accepted sequences of messages that two or more agents can utter within an scene. The protocol is represented by means of a directed graph, where each node is an step or state of the conversation, and arcs are utterances. For each illocution there is an illocution scheme, which defines precisely the nature of the utterance, the roles of the sender and the receiver(s) of the utterance and the information that is exchanged. Examples of such representations can be seen in figures 6 to 12. Within the conversation graph, the greyed out nodes denote points at which agents may join the conversation, those with a double circle, where an agent may leave the conversation, while the rest are intermediate nodes.

Reception Room As explained above, in the *Reception Room* agents enter and are registered with the platform and according to their roles are directed to the proper room. The protocol for this scene is shown in figure 6: an agent_{*i*} makes a request for admission (1) that can be accepted (messages 3a, 3b, 3c) or refused (message 2, exit state w_1). According to the role of the incoming agent_{*i*} it is directed to the *Consultation Room* (exit state w_2) or its request is validated (messages 4 and 5, exit state w_3). Then agent_{*i*} waits until an instance of *Exchange Room* is available in which to negotiate over the kind of pieces agent_{*i*} is requesting (messages 6 and 7, exit state w_4).

Exchange Room The *Exchange Room* is the place where negotiation over pieces is performed. The protocol of this scene is shown in figure 8: agent_{*i*}, a hospital *Finder Agent*, asks the scene manager for piece offers which match the requirements given in their petition. Then the scene manager gives a list of available pieces (message 9) that is evaluated by the external agent_{*i*} (message 10). With this information the scene manager can make a provisional assignment and resolve the situation of when two agents are interested in the same piece. When this provisional assignment is delivered (message 11) then agent_{*i*} exits the scene to go to the *Confirmation Room* represented by state w_2 . There is an alternative path for the case when there are no available pieces matching the requirements described in the petition (message 9 with null list). In this case agent_{*i*} requests an exit permission from the institution (message 12, exit state w_1), including the reason for leaving. The reason provided is recorded in the institution logs to form an audit trail for the relevant authorities to inspect. For further information about this negotiation process see [8].

Confirmation Room In this scene, the provisional assignments made in the *Exchange Room* are either confirmed or withdrawn. Figure 10 shows the protocol of this scene: the agent_{*i*} can analyze the assigned piece data and then accept or refuse it (message 13). If the agent_{*i*} accepts the piece and no higher-priority requests appear during a certain time window then the provisional assignment is confirmed and a delivery plan is given to the agent_{*i*} (message 14), and then it ex-




Fig. 5. The conversation graph for the Reception room

Msg#	Illocution
1	(request (?x hf hc tb) (?y rrm) (admission ?id_agent ?role ?hospital_certificate))
2	(deny (!y rrm) (!x hf hc tb) (deny ?deny_reason))
3a	(accept (!y rrm) (!x hf) (accept_hf))
3b	(accept (!y rrm) (!x hc) (accept_hc))
3c	(accept (!y rrm) (!x tb) (accept_tb))
4	(inform (?x hf) (?y rrm) (petition ?id_hospital ?urgency_level ?time_to_deliver ?piece_type (?piece_parameters) (?info_recipient)))
5	(inform (!y rrm) (!x hf) (petition_state ?id_ppetition ok faulty))
6	(inform (?y rrm) (?x hf) (init_exchange ?piece_type ?id_exchange_room))
7	(request (?x hf) (?y rrm) (exchange_entrance_request !id_exchange_room))

Fig. 6. The illocutions for the Reception room

its the Carrel system (exit state w_2). When there is a request with higher priority that needs the piece provisionally assigned to agent_{*i*}, a conflict arises. To resolve the conflict the scene manager notifies the agent_{*i*} that the assignment has been withdrawn (message 15) and that it is then entitled to a fresh request for another piece (message 16) to be negotiated again in the *Exchange Room*.

Consultation Room This scene allows agents coming from hospitals or tissue banks to inform Carrel of the location and status of known and new pieces. Note that we assume that a tissue bank only participates in one instance of Carrel at the same time. The protocol of this scene is shown in figure 12. The incoming agents can perform notifications (messages 17, 18, 19) and are informed if the notification is successful (message 20). The agents coming from hospitals representing the Hospital Transplant Coordinator can also perform queries (message 21) about




Fig. 7. The conversation graph for the Exchange room

Msg#	Illocution
8	<code>(query-if (?x hf) (?y erm) (offer_list ?id_petition))</code>
9	<code>(inform (!y erm) (!x hf) (offer_list !id_petition (list (?id_piecel ?info_piecel) ... (?id_piecen ?info_piecen)))</code>
10	<code>(inform (?x hf) (?y erm) (weighted_list !id_petition (list (!id_piecel ?weight) ... (!id_piecen ?weight)))</code>
11	<code>(query-if (?y erm) (?x hf) (piece_offer (?id_petition ?id_piece ?cost_estimation) void))</code>
12	<code>(request (?x hf) (y im) (exit ?exit_reason))</code>

Fig. 8. The illocutions for the Exchange room

historical information, such as statistics on, say, successful cornea transplantations over a certain period. The queries are answered (message 22) with the level of detail that is permitted for a certain role, where the permissions are verified according to a *Role-Based Access Model* [21]. When the incoming agents have performed all the queries and notifications, they exit the Carrel system (message 23).

3.3 Normative rules

Although we as humans function in the context of norms in many different situations every day, we have become so used to them that we are not necessarily aware of them. The premise behind our use of electronic institutions is that we are trying to define and construct similar mechanisms to help control and simplify interaction for agents in virtual worlds. And the objectives and outcomes remain the same: to govern the behaviour of agents, establish rights and obligations and above all, to create an atmosphere of trust.

At a high level norms appear to say very a little and convey no hint of how they

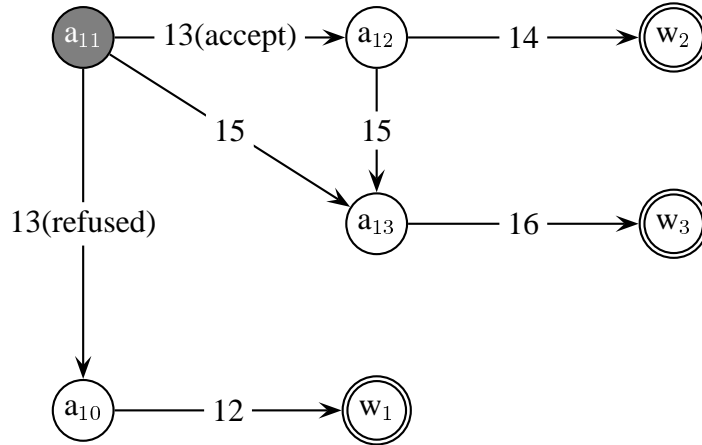


Fig. 9. Conversation graph for the Confirmation room

Msg#	Illocution
12	(request (?x hf) (y im) (exit ?exit_reason))
13	(inform (?x hf) (?y cfrm) (piece_eval ?id_petition ?id_piece accepted refused))
14	(inform (?y cfrm) (?x hf) (piece_delivery ?id_petition ?id_hospital ?id_tissue_bank ?delivery_plan))
15	(inform (?y cfrm) (?x hf) (piece_reassigned_exception ?id_petition ?id_piece ?reassignment_reason))
16	(query-if (?x hf) (?y cfrm) (another_offer_list ?id_petition))

Fig. 10. Illocutions for the Confirmation room

may be implemented: we might express desires for “fairness” or “equity”, but such norms are useless—instead we have to invent rules of behaviour, restrictions, constraints, requirements and obligations which satisfy us that they *support* the more abstract norms. And eventually, this norm refinement process can be seen to bring us right down to the level of individual actions, such as leaving a scene, or individual expressions, such as the protocols defined by the conversation graphs. Thus, we can also find norms defined implicitly in the performative structure:

- *inter-scene constraints*: connections among scenes are, in fact, constraints about the accepted paths for each agent role inside the scene network.
- *intra-scene constraints*: the scene protocols define the accepted interactions for each agent role (what can be said, by whom, to whom and when).

These implicit norms are made explicit by the agents while they are following the




Fig. 11. Conversation graph for the Consultation room

Msg#	Illocution
17	(inform (?x hc) (?y crm) (piece_arrival ?id_hospital ?id_tissue_bank ?id_piece (?state)))
18	(inform (?x hc) (?y crm) (transplantation_eval ?id_piece ?id_recipient ?date (?info_transplantation)))
19	(inform (?x tb) (?y crm) (tissue_bank_update ?id_tissue_bank ?id_piece (?specifications)))
20	(inform (!y crm) (!x hc tb) (notification_ack !id_piece ok error))
21	(query-if (?x hc) (?y crm) (?query))
22	(inform (!y crm) (!x hc) (query_results (?results)))
23	(request (?x hc tb) (?y im) (end))

Fig. 12. Illocutions for the Consultation room

paths.

But there are other norms that are not yet defined. These are norms where actions taken by an agent_i inside a certain scene may have consequences for its future interactions outside that scene. An instance of this kind of situation in Carrel is the arrival of a high-priority request in the Reception Room scene has effects in the Confirmation Room scene, as any provisional assignment cannot be confirmed until the high-priority request has been served. This can be expressed informally by a rule such as the following:

*if a request has arrived in the reception room,
and the request is high-priority,
and the request is well-formed,*

*and no piece has yet been assigned to the request,
then there can be no assignment confirmation in the Confirmation room.*

Figure 13 shows this rule split into its elements (preconditions and postconditions) along with a formal specification of each element.

3.4 *The textual specification*

A topic of on-going work is the use of model-checking to establish properties of the scene protocols and of the performative structure. For this reason, there is a complementary textual presentation of the illocutions, scenes, dialogic framework and performative structure, which is amenable to transformation and analysis for exactly this purpose. It must be emphasized that the textual representation is designed to capture the institutional information, rather than being directly suitable for model-checking. On the other hand, it is relatively straightforward to extract the transition networks and the arc labels from the specification and then generate textual input for model checkers such as Spin [20] or nuSMV[5]. Figure 14 shows part of the textual description of the Carrel system.

3.5 *The agent architecture*

Once all the above framework is in place, we are then able to describe the whole multi-agent system and the relations among the agents.

It is straightforward to get the agent architecture (figure 15) from the Performative Structure. First of all, there is at least one staff agent for each institutional role (identified in §3.1): the **RR Agent** as a *reception room manager*; the **CR Agent** as a *consultation room manager*; an **ER Agent** for each exchange room, acting as an *exchange room manager*; the **CfR Agent** as a *confirmation room manager*; and the **IM Agent** as the *institution manager*, coordinating the other agents.

In order to assist the agents mentioned above, two specialized agents are added: the **Planner Agent** which is a specialized agent tasked with carrying out the planning needed in the Confirmation Room, and the **DB Agent** (an agent specialized in the access control to the Database, also depicted in the figure).

Finally we have included an interface connected to the *Institution Manager*, to allow a human administrator to access the system.

A prototype of this architecture has been implemented in JADE [2] and it is thereby (mostly!) FIPA compliant [16].

<pre>done((inform (?x hf) (?y rrm) (petition ?id_hospital ?urgency_level ?time_to_deliver ?piece_type (?piece_parameters) (?info_recipient))), r-room)</pre>	<p>if a request has arrived in the reception room</p>
\wedge	AND
<pre>!urgency_level = urgency_0</pre>	<p>the request is high-priority</p>
\wedge	AND
<pre>done((inform (!y rrm) (!x hf) (petition_state ?id_petition ok))), r-room)</pre>	<p>the request is well-formed</p>
\wedge	AND
<pre>\neg done((inform (?z cfrm) (!x hf) (piece_delivery !id_petition !id_hospital ?id_tissue_bank ?delivery_plan)), cf-room)</pre>	<p>no piece has yet been assigned to this request</p>
\Rightarrow	THEN
<pre>\forall (w hf): (w \neq x) : \neg done(((inform (!z cfrm) (!w hf) (piece_delivery ?id_petition ?id_hospital ?id_tissue_bank ?delivery_plan))), cf-room)</pre>	<p>there can be no assignment confirmation in the Confirmation room</p>

Fig. 13. A norm linking actions in one scene with consequences in another

3.6 Formalizing the institutional agents

At this point all the interactions among institutional agents and external agents have been modelled. In fact such interactions have been specified in a very formal way by means of *norms*. The last step to complete the definition of the institution is to describe the behavior of the institutional agents and the interactions between them.

Following [31] we are working with the Ambient calculus (a process algebra with locations [3]) in order to obtain a precise description of the interaction of the institu-

```

(define-performative-structure carr-performative-structure
  scenes:
    ((root root-scene)
     (r-room reception-room-scene)
     (e-room exchange-room-scene)
     (cf-room confirmation-room-scene)
     (cs-room consultation-room-scene)
     (output output-scene)
    )
  connections:
    ((root r-room consulter ...))
)
; declare the
; scenes comprising
; the institution

(define-dialogic-framework fm-dialogic-framework
  ontology: carr-ontology
  representation-language: first-order-logic
  illocutionary-particles: (... ..)
  external-roles: (hf, hc, tb)
  internal-roles: (im, xrm, erm, cfrm, crm)
  role-hierarchy: ((im xrm) (im erm)
                  (im cfrm) (im crm))
  role-compatibility: (incompatible: (hf tb) (hc tb) (hf hc))
)
; list all the roles
; that participating
; agents may play and
; how the roles are related
; to one another

(define scene e-room
  roles: (erm hf)
  dialogic-framework: e-room-df
  states: (a6 a7 a8 a9 a10 w1 w2)
  initial-state: a6
  final-states: (w1 w2)
  access-states: ((erm (a6)) (hf (a6 a7)))
  exit-states: ((erm (w2)) (hf (w1 w2)))
  agents-per-role: ((erm 1 1) (hf 0 n))
  transitions:
    ((a6 a7 (query-if (?x hf) (?y erm) (offer_list ?id_petition) ))
     (a7 a8 (inform (!y erm) (!x hf) (offer_list !id_petition
              (list (?id_piecel ?info_piecel) ... (?id_piecen ?info_piecen))))))
     ...
     (a10 w1 (request (?x hf) (?y im) (exit ?exit_reason) ))
    )
  constraints: ()
)
; describe the conversation graph
; contained in a particular scene
; specifying which roles may participate
; the rest is just a textual
; representation of a finite state
; machine

...

(define-norm urgency0-arrival
  antecedent:
    ((r-room
     (inform (?x hf) (?y rrm)
      (petition ?id_hospital ?urgency_level ?time_to_deliver ?piece_type
       (?piece_parameters) (?info_recipient) ) ) )
     (!urgency_level = urgency_0))
     (r-room
      (inform (!y rrm) (!x hf) (petition_state ?id_petition ok) ) )
     (nil)))
  defeasible-antecedent:
    ((cf-room
     (inform (?z cfrm) (!x hf)
      (piece_delivery !id_petition !id_hospital ?id_tissue_bank ?delivery_plan))))
  consequent:
    (forall (w hf) (not (= w x))
     (obliged !z (not (inform (!z cfrm) (!w hf)
      (piece_delivery ?id_petition ?id_hospital ?id_tissue_bank ?delivery_plan) ), cf-room)))
)
; define a norm in three parts

```

Fig. 14. Part of the Carrel specification in text format

tional agents. We do not present any of this aspect of the formalization here, partly because of its mathematical nature and partly because of the amount of introductory material that would be appropriate to assist in its comprehension, but primarily because we regard the focus of this paper not to be process algebraic modelling, but rather the demonstration of the use of norms—implicit and explicit—to specify sound systems in medical domains.

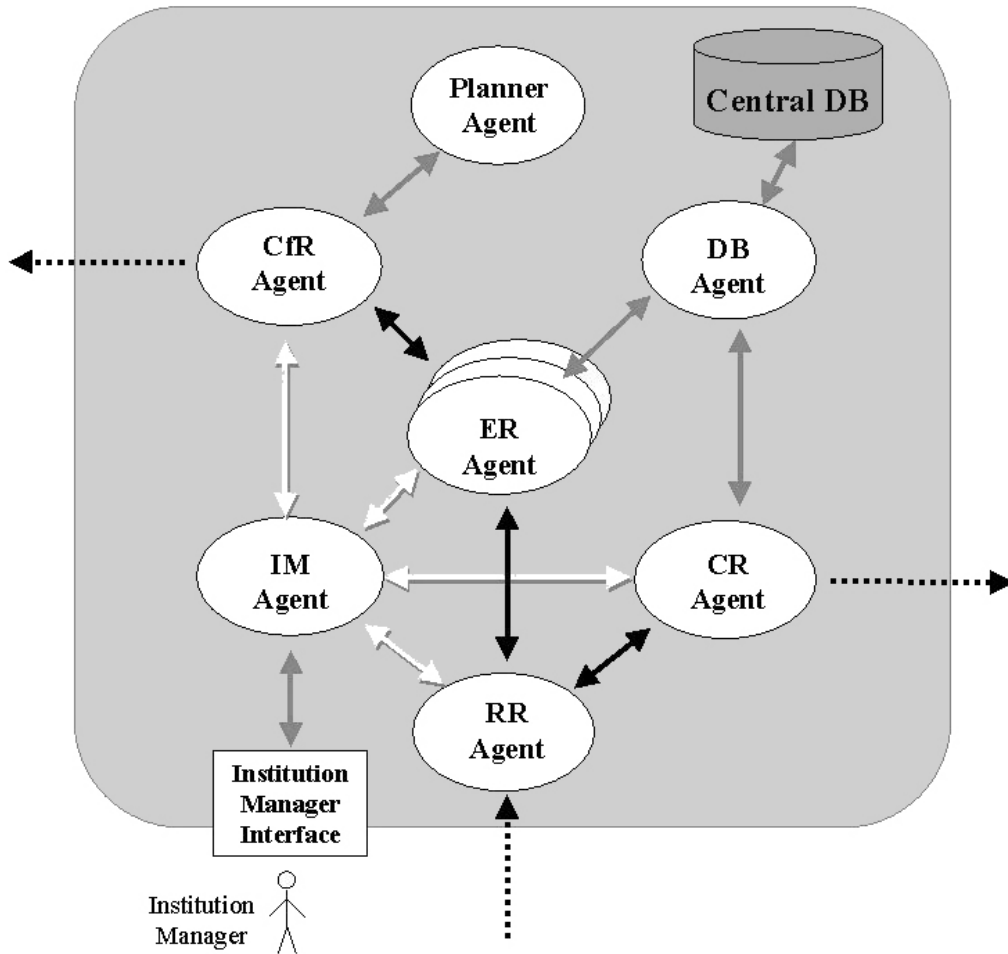


Fig. 15. The Carrel Institution Agency

4 Conclusions

The use of software solutions to assist the medical staff in their daily tasks is not a new topic. In fact there is a significant sub-area of software engineering called *Medical Informatics* or *Electronic Health-Care* aiming to find new ways to use information technologies to:

- model and apply medical knowledge
- create, store and retrieve patient records
- define, select or improve diagnosis and treatment guidelines
- improve communication among health professionals, and also between health professionals and patients
- (re-)organize and improve any health care service

In section §2.2 we presented the utility of *agents* in medical applications. As noted in [18, 35] the strengths of Agent Technologies (mainly *pro-activity* and *autonomy*)

make such technologies well-suited for medical applications. We will summarize here only some of their arguments:

- the capability of agents to anticipate pro-actively the information needs of users,
- their support of synchronous and asynchronous communication among parties,
- their suitability to support distributed decision making
- their ability to adapt to unpredicted situations
- their capability to adapt the Health Care services to the patient needs

As a result of these abilities, the use of agents for medical applications is rising. [35] provides a comprehensive review of the use of Agents in Health Care.

A central aspects of our work has been formalization. In [18] it is stated there is a need to develop methodologies and, if possible, formal design methods when applying agents to the medical domain. In their work they present the *PROforma* formalism, which permits the definition of guidelines to be followed by each agent in the system. For our purposes and from the point of view of distributed systems, the problem is that the formalism does not cover the communication acts and the dependencies between guidelines and hence how the behaviour of one agent may affect another.

We have presented here the suitability of a general formalism for multi-agent systems such as ISLANDER [15], based on the idea of a *dialogical framework*, applied in the medical domain. We have applied such formalism to the problem of transplant piece allocation, and by doing so we have been able to build a first prototype of the Carrel system.

Most of the work in the field of transplant allocation (such as EU projects RE-TRANSPLANT [33]⁷, TECN [34]) is devoted to the creation of a) standard formats to store and exchange information about pieces, donors and recipients among organizations, b) telematic networks, or c) distributed databases. The projected outcomes of this project are potentially extremely important and useful because the development of standard ontologies to define the concepts of the domain of discourse for agents is an essential stage on the way to the effective deployment of this technology. Project ESCULAPE [14] uses conventional software to help in tissue histocompatibility by developing HLA⁸ referencing computer systems and software packages to be used by hospitals and laboratories as a human tissue typing tool.

⁷ Unfortunately, there is no information available about the practical results of those projects other than the project URL.

⁸ *Human Leucocyte Antigens*.

There are very few references in the literature about the use of agents in the transplant domain. [36] describes an agent that uses multi-criteria decision techniques in the selection of the best receiver in a transplant, providing the Hospital Transplant Coordinator with a result according to the weights the user assigned to each criteria. [26] describes an hierarchical multi-agent system where the agent on the root node plans transport routes between hospitals using the information obtained from the other agents in the hierarchy, removing routes that will exceed the maximum available time for transportation and avoiding potential fatal delays due to mistakes in coordination of different means of transport. [25] proposes a multi-agent system architecture to coordinate hospital teams for organ transplants. Coordination is achieved through agents that keep track of the personnel schedules and the availability of the facilities (both described as time-tables divided into slots of thirty minutes). Finally [1] presents an alternative design for a multi-agent architecture for the Spanish organ allocation process. It identifies the agents needed to solve the problem and organizes them in four levels (*Hospital Level*, *Regional Level*, *Zonal Level* and *National Level*). However, no formalism is used in the development of the architecture.

As far as we know, CARREL is the first software system applied to the transplant allocation problem that joins the strengths of *agents* with the advantages of formal specifications.

As part of our future work we aim to expand the formalism presented here in order to describe other issues (such as the delegation of all non-fulfilled commitments of the incoming agents to the hospitals they belong to, which is not currently expressible by the formalism), and the creation of tools to build agent platforms automatically from institution specifications. A preliminary report on the generation of institutions and the agents to populate them appears in [39], while preliminary reports on the use of model-checking of institutional models appears in [40, 6].

We also aim to explore the extension of the *PROforma* formalism [17, 19] to include agent communication, so that it may be used to model systems with characteristics like Carrel.

References

- [1] A. Aldea, B. López, A. Moreno, D. Riaño, and A. Valls. A multi-agent system for organ transplant co-ordination. In S. Quaglini, P. Barahona, and S. Andreassen, editors, *Proceedings of the 8th. European Conference on Artificial Intelligence in Medicine, Portugal, 2001*, Lecture Notes in Artificial Intelligence 2101: Artificial Intelligence in Medicine, pages 413–416. Springer-Verlag, 2001.
- [2] F. Bellifemine, A. Poggi, and G. Rimassa. JADE : A FIPA compliant agent frame-work. In *Proceedings of Practical Application of Intelligent Agents*

- and Multi-agent Systems*, pages 97–108, London, 1999. Practical Application Company Ltd.
- [3] L. Cardelli. Mobility and security. In F. L. Bauer, editor, *Proceedings of the NATO Advanced Study Institute on Foundations of Secure Computation*, number ISBN 1 58603 015 9 in NATO Science Series, pages 3–37. IOS Press, 1999.
 - [4] C. Castelfranchi, F. Dignum, C. Jonker, and J. Treur. Deliberate normative agents: Principles and architecture. In *Proceedings of The Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Orlando, FL, Lecture Notes in Artificial Intelligence 1757. Springer-Verlag, 1999.
 - [5] A. Cimatti, E. M. Clarke, F. Giunchiglia, and M. Roveri. NuSMV: A new symbolic model checker. *International Journal on Software Tools for Technology Transfer*, 2(4):410–425, 2000.
 - [6] O. Cliffe and J. Padget. A framework for checking agent interaction within institutions. In *Proceedings of MOCHART workshop at ECAI’02*, Lyon, France, 2002.
 - [7] R. Conte and C. Castelfranchi. Are incentives good enough to achieve(info)social order? In C. Dellarocas and R. Conte, editors, *Workshop on Norms and Institutions in Multi-Agent Systems*. ACM-AAAI, ACM Press, 2000.
 - [8] U. Cortés, A. López-Navidad, J. Vázquez-Salceda, A. Vázquez, D. Busquets, M. Nicolás, S. Lopes, F. Vázquez, and F. Caballero. Carrel: An agent mediated institution for the exchange of human tissues among hospitals for transplantation. In *3^{er} Congrés Català d’Intel·ligència Artificial*, pages 15–22. ACIA, 2000.
 - [9] U. Cortés, J. Vázquez-Salceda, A. López-Navidad, and F. Caballero. UCTx: A multi-agent system to assist a transplant coordinator unit. In *Applied Intelligence: Machine Learning Policies*. 2002.
 - [10] O. N. de Transplantes, editor. *Informes y Documentos de Consenso promovidos por la Organización Nacional de Transplantes y la Comisión de Transplantes del Consejo Interterritorial del Sistema Nacional de Salud*. Editorial Complutense S. A., Madrid, Spain, 1st edition, 2000.
 - [11] C. Dellarocas and M. Klein. Contractual Agent Societies: Negotiated shared context and social control in open multi-agent systems. In C. Dellarocas and R. Conte, editors, *Workshop on Norms and Institutions in Multi-Agent Systems*, pages 1–11. ACM-AAAI, ACM Press, 2000.
 - [12] F. Dignum. Abstract norms and electronic institutions. In G. Lindemann, D. Moldt, M. Paolucci, and B. Yu, editors, *Int. Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA02)*, volume Mitteilung 318, pages 31–50. U. Hamburg, 2002.
 - [13] Directive 95/46/CE of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and of the free movement of such data, October 1995.
 - [14] ESCULAPE: Use of computer techniques for tissues matching and anal-

- ysis as an aid to human transplantation. http://dbs.cordis.lu/cordis-cgi/srchidadb?CALLER=EN_CORDIS&QZ_WEBSRCH=E%SCULAPE.
- [15] M. Esteva, J. Padget, and C. Sierra. Formalizing a language for institutions and norms. In J.-J. Meyer and M. Tambe, editors, *Intelligent Agents VIII*, volume 2333 of *Lecture Notes in Artificial Intelligence*, pages 348–366. Springer Verlag, 2001. ISBN 3-540-43858-0.
- [16] The Foundation for Intelligent Physical Agents, <http://www.fipa.org/repository/fipa2000.html>. *FIPA Specifications*, 2000.
- [17] J. Fox and S. Das. *Safe and Sound*. MIT Press, Boston, 1st edition, 1999.
- [18] J. Fox and S. Das. Guardian agents for safety-critical systems. In V. Shankararaman, editor, *Workshop on Autonomous Agents in Health Care*, pages 25–34, Barcelona, Spain, 2000.
- [19] J. Fox, N. Johns, A. Rahmzadeh, and R. Thompson. Proforma: A method and a language for specifying clinical guidelines and protocols. In J. Brender, J. P. Christensen, J.-R. Scherrer, and P. McNair, editors, *Medical Informatics Europe '96*, pages 516–520. IOS Press, 1996.
- [20] G. J. Holzmann. The Spin model checker. *IEEE Transactions on Software Engineering*, 23(5):279–95, May 1997.
- [21] A. Lin. Integrating policy-driven role-based access control with common data security architecture. Technical Report HPL-1999-59, HP Laboratories Bristol, 1999.
- [22] Ley Orgánica 15/1999 de protección de datos de carácter personal. Boletín Oficial del Estado 292, 14 de diciembre 1999.
- [23] A. López-Navidad. Professional characteristics of the transplant coordinator. *Transplantation Proceedings*, (23):1607–1613, 1997.
- [24] R. Matesanz. Meeting the organ shortage: Current status and strategies for improvement of organ donation. *Newsletter Transplant*, 4(1):5–17, 1999.
- [25] A. Moreno, A. Valls, and J. Bocio. Management of hospital teams for organ transplants using multi-agent systems. In S. Quaglini, P. Barahona, and S. Andreassen, editors, *Proceedings of the 8th. European Conference on Artificial Intelligence in Medicine, Portugal, 2001*, Lecture Notes in Artificial Intelligence 2101, pages 374–383. Springer-Verlag, 2001.
- [26] A. Moreno, A. Valls, and A. Ribes. Finding efficient organ transport routes using multi-agent systems. In *IEEE 3rd International Workshop on Enterprise Networking and Computing in Health Care Industry (Healthcom)*, L'Aquila, Italy, 2001.
- [27] A. Newell and H. Simon. The knowledge level. *Artificial Intelligence*, 18(1):87–127, 1982.
- [28] P. Noriega. *Agent-Mediated Auctions: The Fishmarket Metaphor*. Number 8 in IIIA Monograph Series. Institut d'Investigació en Intel·ligència Artificial (IIIA), Bellaterra, Spain, 1997. PhD Thesis.
- [29] D. North. *Institutions, Institutional Change and Economic Performance*. Cambridge Univ. Press, Cambridge, 1st edition, 1990.
- [30] Organización Nacional de Trasplantes. <http://www.msc.es/ont>.

- [31] J. Padget. Modelling simple market structures in process algebras with locations. *Artificial Intelligence and Simulation of Behaviour Journal*, 1(1):87–108, 2001. ISSN 1476-3036.
- [32] Real Decreto 994/1999 por el que se aprueba el reglamento de medidas de seguridad de los ficheros automatizados que contengan datos de carácter personal. Boletín Oficial del Estado 49, 26 de Febrero 2000.
- [33] REgional and International Integrated Telemedicine Network for Medical Assistance in End Stage Diseases and Organ TRANSPLANT. <http://retransplant.vitamib.com>.
- [34] TECN : Transplant Euro Computer Network. <http://www.ejeisa.com/nectar/t-book/html/health.htm\#TECN>.
- [35] T. P. V. Shankararaman, V. Ambrosiadou and B. Robinson. Agents in health care. In V. Shankararaman, editor, *Workshop on Autonomous Agents in Health Care*, pages 1–11, Barcelona, Spain, 2000.
- [36] A. Valls, A. Moreno, and D. Sánchez. A multi-criteria decision aid agent applied to the selection of the best receiver in a transplant. In J. Braz, M. Piattini, and J. Filipe, editors, *4th International Conference on Enterprise Information Systems (ICEIS-2002)*, Ciudad Real, Spain, 2002. ICEIS Press.
- [37] W. Vasconcelos, J. Sabater, C. Sierra, and J. Querol. Skeleton-based agent development for electronic institutions. In C. Castelfranchi and W. L. Johnson, editors, *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, volume 2, pages 696–703. ACM, ACM Press, July 2002.
- [38] J. Vázquez-Salceda, U. Cortés, A. López-Navidad, F. Caballero, and J. Padget. The organ allocation process: a natural extension of the Carrel agent mediated electronic institution. In *Proceedings of Agents in Healthcare workshop at ECAI'02*, 2002. Revised version to appear in *AI Communications* in 2002.
- [39] O. Vickers and J. Padget. Skeletal JADE Components for the Construction of Institutions. In J. Padget, D. Parkes, N. Sadeh, O. Shehory, and W. Walsh, editors, *Agent Mediated Electronic Commerce IV*, volume 2531 of *Lecture Notes in Artificial Intelligence*, pages 174–192. Springer Verlag, 2002. to appear.
- [40] M. Wooldridge, M. Fisher, M.-P. Huget, and S. Parsons. Model checking multiagent systems with mable. In C. Castelfranchi and W. L. Johnson, editors, *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, volume 2, pages 952–959. ACM, ACM Press, July 2002 2002.
- [41] M. Wooldridge and N. Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.