Agent-based careflow using CPGs

David ISERN , Antonio MORENO
Departament d'Enginyeria Informàtica i Matemàtiques. Universitat Rovira i Virgili.
Avda. Països Catalans, 26. 43007 Tarragona, Catalunya
{disern,amoreno}@etse.urv.es

Abstract. The management and coordination of the human and material resources needed in the careflow of a hospital is a difficult task. In this paper we propose to use an agent-based system to tackle this problem. This system is an improvement of our previous work on the multi-agent system *HeCaSe*, that provided a basic set of secure and personalised medical services. An important novelty of our approach is the use of agents to help doctors to follow *Clinical Practice Guidelines (CPG)* associated to different pathologies. All the actors of the system are modelled using agents, and CPGs are used to mediate the negotiation of appointments between patients and services with the doctor's supervision.

Keywords. multi-agent systems, intelligent autonomous agents, coordination in distributed systems, clinical practice guidelines, careflow management.

1. Introduction

A medical centre has a large number of resources that have to be managed in the *careflow* process (*i.e.* the workflow process involved in the provision of care, [1,2]). There are *human resources* (many different kinds of professionals, such as doctors, nurses, administrators, receptionists) and *material resources* (*e.g.* X-ray, gym, rehabilitation unit). All of them play a specific role within the medical centre organisation, and they must coordinate their activities to provide the best possible care to patients. Any computer system designed to work in a medical centre has to take into account different issues; in [3] it was argued that most of the following ones suggest the use of *agent technology* [4] in the health care area:

- *Heterogeneous data*: this kind of organisation generates a lot of data from different sources (an X-ray image, a blood test, the result of a medical visit, an electronic medical record, etc.) and it is necessary to integrate them smoothly (*e.g.* a doctor should be able to add easily the data of a clinical examination –symptoms, diagnosis, treatment- into the patient's medical record, or the result of a blood test could be automatically incorporated to the medical record and sent to the responsible doctor).
- *Autonomy*: services, departments, medical practitioners and patients are autonomous entities with their own knowledge, beliefs and goals.
- *Distributed data*: the data related to a patient is usually distributed among different units of the hospital.
- *Complex coordination*: there are several kind of interactions to coordinate: human-resource, resource-resource and human-human. Also, this coordination has temporal constraints which must be considered.

A *clinical practice guideline (CPG)* indicates the protocol to be followed when a patient is diagnosed a certain illness (e.g. which medical tests have to be performed on the patient to get further data, or what steps have to be taken according to the results of the tests). Therefore, they provide very detailed information concerning the resources needed in the treatment of the patient. In this paper we propose the design of an agent-based system, based on our previous work on the *HeCaSe* system [5-7], in which CPGs are automatically incorporated into the workflow of doctors and hospital services; in that way, care is improved at least in two ways:

- Doctors are reminded about the steps that should be followed in the treatment of a certain disease, and that reduces the possibility of making errors or forgetting tasks to be done.
- Agents representing patients, doctors and hospital services can automatically coordinate their activities to provide a fast care (*e.g.* they can arrange the dates for different tests to be performed on the patient easily, without the patient having to visit personally different units of the hospital to arrange those tests).

The rest of the paper is organised as follows. First of all, the next section comments the main features of medical guidelines and the current available tools to create and store them. Section 3 describes an agent-based system called *HeCaSe2*, that combines agents and medical guidelines in order to improve collaboration and coordination between services and humans. Finally, section 4 gives some conclusions and some future lines of research.

2. PRO*forma*: a language for describing CPGs

A *Clinical Practice Guideline (CPG)* is a systematically developed statement designed to assist clinicians in making decisions about appropriate health care for specific clinical circumstances. The vast majority of CPGs are published as text, and typically include criteria describing their applicability to particular groups of patients, the recommended processes of care and appropriate use of materials and procedures, as well as providing ancillary information such as supporting evidence. CPGs are generally concerned with two main areas that directly affect patient care: the quality of clinical decision-making, and the correct and timely management of clinical tasks [6]. From an AI perspective there is an alternative to publishing guidelines solely in human readable form such as text, tables or flow diagrams, which is to formalise the medical knowledge contained in these guidelines in a machine-readable format that a computer can use to support clinicians in their routine work. Many researchers have tackled this problem in the last years and there are different approaches to represent CPGs in a structured/formal way. Different languages such as Asbru, EON, GLIF, PRO*forma*, PRODIGY and GLARE were analysed and compared in [8,9]. Basically, all these languages use task-based representations with differences in the basic elements.

In the system described in this paper CPGs are represented in a language called PRO*forma*, which was developed by the Advanced Computation Laboratory at the Imperial Cancer Research Fund (now Cancer Research UK) under the EU 4th Framework PROMPT project (1996-1999) [10]. Fig. 1(a) is a screenshot of TALLIS, a tool developed to manage CPGs in that language. Fig. 1(b) shows the four basic elements defined in the language: *Decisions* represent the choice points in a guideline, *Actions* represent an activity performed by an external agent (*e.g.* an injection), *Enquiries* represent the acquisition of data from some external source (*e.g.* from the patient's medical record), and *Plans* allow tasks to be composed in order to design complex structured protocols [10,11].

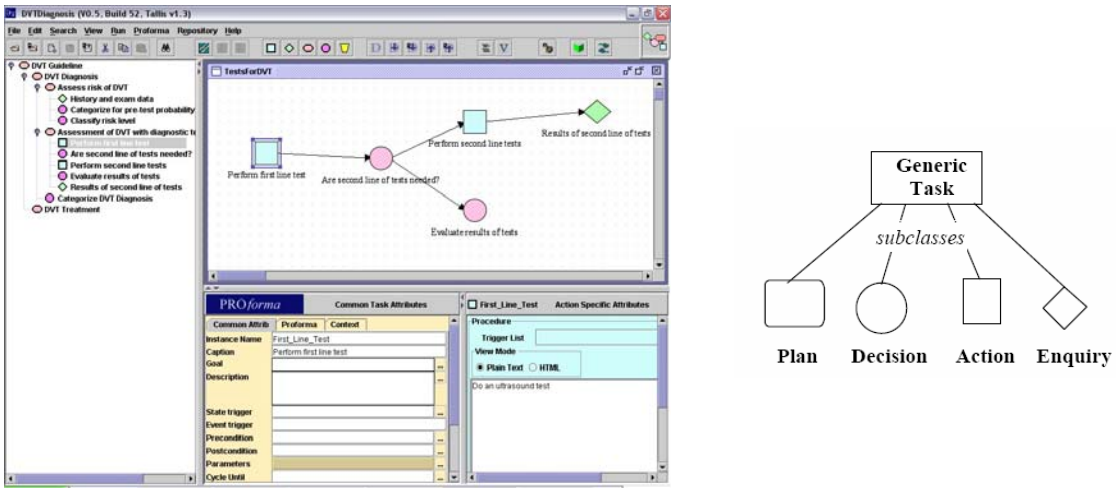

(a) CPG retrieved from TALLIS repository (b) Tasks defined in PRO*forma*

Figure 1: Development of CPGs using PRO*forma*

Certain properties are common to all tasks (preconditions, postconditions, goal and a textual description), but each subclass contains specific attributes. Each *Decision* has a list of options that can be chosen called *candidates*, and a choice mode of those candidates that can be single or multiple. Any *Plan* has a list of components, a set of scheduling constraints about the whole plan or several components, a set of abort conditions, and a set of termination conditions. An *Enquiry* contains a list of the names of the data items to be considered. Finally, an *Action* only contains the procedure to be performed that can be free text or a SQL statement. Data definitions enable PRO*forma* to communicate with external elements (databases, agents, web services).

3. HeCaSe2: Providing Distributed Guideline-Based Health Care Services

First of all, we shall summarise the main characteristics of our previous system *HeCaSe*. After that we will show the architecture of the new system, *HeCaSe2*, and the way in which CPGs are managed to improve the care given to patients and the coordination between the hospital staff and services.

3.1 Features of HeCaSe

Health Care Services (*HeCaSe*) is an agent-based prototype that models the medical institutions of Catalonia. It features the following main characteristics (full details of security measures and service personalisation in [5-7]):

- The patient may request information about the medical centres available in a geographical area, specifying a set of search constraints.
- It is possible to book a visit with a doctor, selecting a proposed appointment from a list of possible ones suggested by the system.
- The patient has a secure access to his own medical record.
- Each patient has an agenda to manage his appointments that is taken into consideration in the booking process (personalised proposals).
- He can also rank the possible doctors, medical centres or cities in order to teach the system which ones he prefers for future booking negotiations (learning capabilities).
- Health care personnel can manage easily their working hours and patient's appointments.
- During an examination, doctors automatically access the patient's medical record and update it with the result of the visit.

3.2 Description of the architecture of HECASE2

The basic architecture of the proposed MAS is depicted in the following figure:

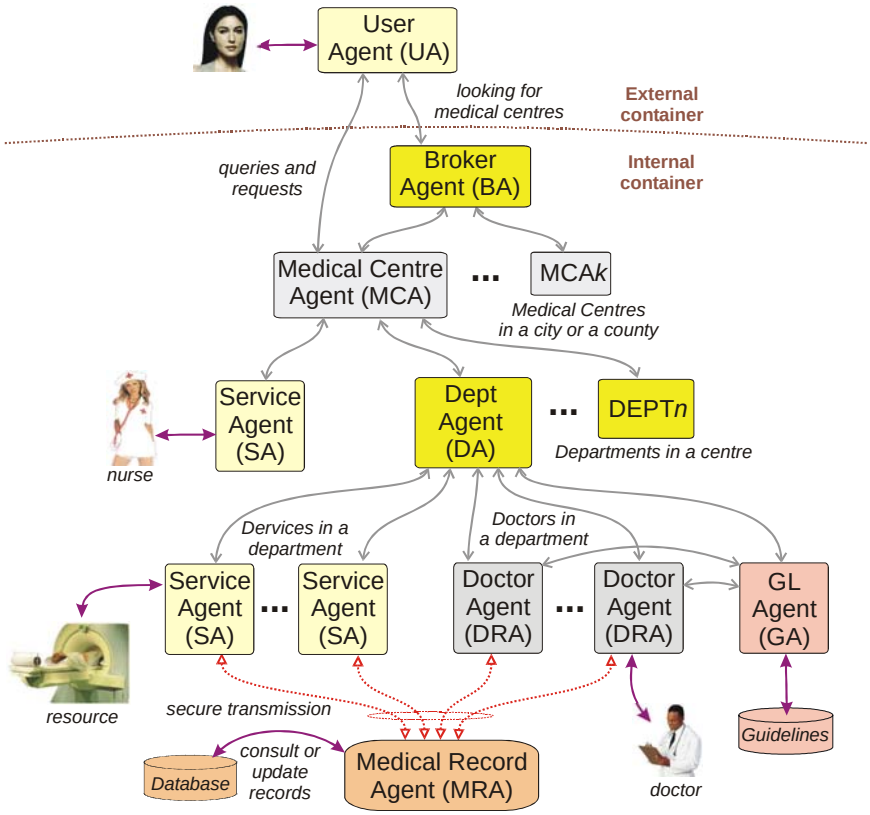

Figure 2. Architecture of the multi-agent system

At the top of the architecture is placed the user, who interacts with the system through his *User Agent* (UA). This agent stores *static data* related to the user (*e.g.* national healthcare number, name, address, access information -login, password, and keys-) and *dynamic data* (the timetable and the preferences of the user, [7]). The *Broker Agent* (BA) is an agent that knows about all the medical centres located in a certain area. A *Medical Centre Agent* (MCA) centralises and monitors the outsiders accesses to the agents that manage the information of a medical centre. A MCA monitors all of its departments, represented by *Department Agents* (DAs), and a set of general services linked to human or physical resources, represented by *Service Agents* (SAs) (*e.g.* a blood test service). Each department has a staff of several doctors, modelled through *Doctor Agents* (DRAs), and offers more specific services, also modelled as SAs (*e.g.*, a nurse that can take different observations *in situ*). Both MCAs and DAs are aware of the services they can provide (when a SA enters the system, it sends a message detailing its services to the associated MCA or DA). In addition, each department contains a *Guideline Agent* (GA) that performs all actions involved with guidelines (*e.g.* it can retrieve the CPG associated to a specific illness). This GA contains only CPGs related to the department where it is located. At the bottom of the architecture, a *Medical Record Agent* (MRA) controls the access to a database that stores all medical records of the patients of the medical centre. Appropriate security measures have been taken to ensure that only properly authenticated and authorised agents may access and update the medical records ([5]).

3.3 Automatic guideline analysis

When a patient has an appointment to be visited by a doctor, his medical record is automatically sent to the DRA by the MRA (this feature of *HeCaSe* was described in [6]). Thus, the doctor can take into account all the data of the medical record –previous visits, tests, illnesses, treatments, allergies, etc.- in the new visit. The doctor, during the medical examination, makes the diagnosis that the patient has a certain disease, *d*. He may provide this information to the system through the graphical interface of the DRA, along with other data such as the drugs to be taken, as shown in the following figure:

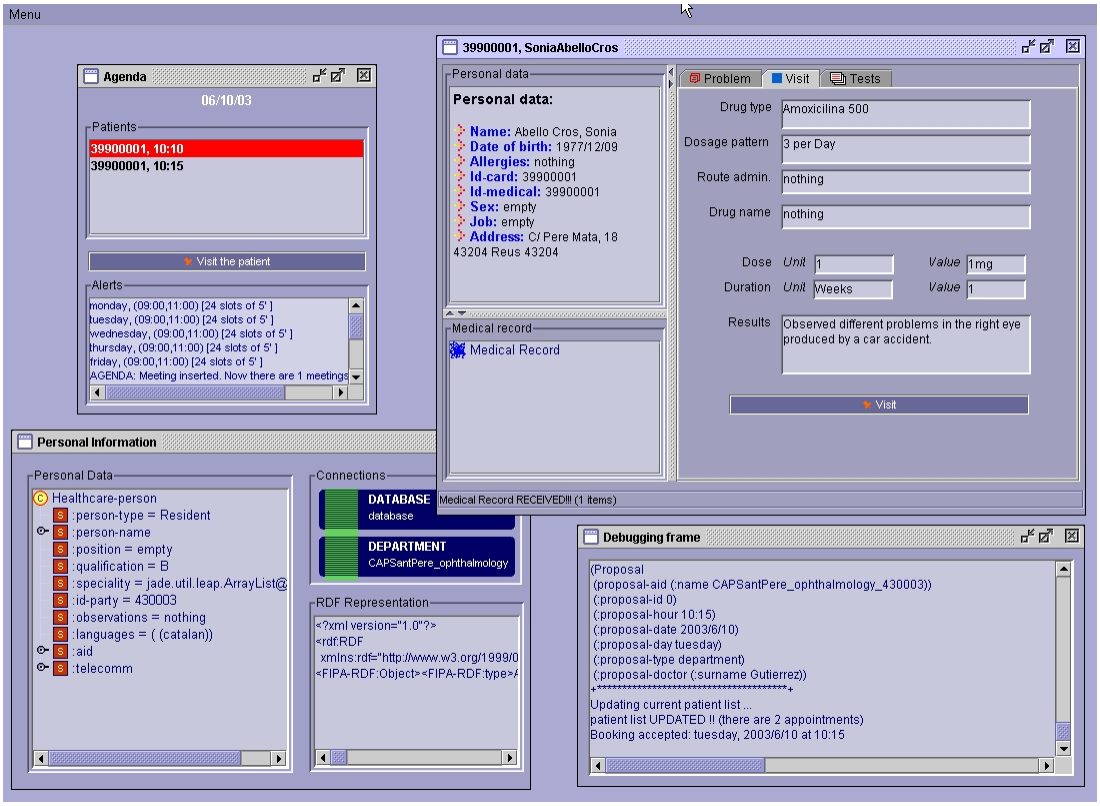

Figure 3. Screenshot of the DRA

This information is sent by the DRA to the GA of the department, which sends to the DRA the guideline associated with the disease *d*, CPG_d (represented in the PRO*forma* language commented on section 2). That guideline specifies the recommended clinical protocol associated with the illness, *i.e.* the steps to be followed to give an appropriate care to a patient with that pathology.

The DRA starts to analyse the tasks contained in the guideline, and it performs the following actions (recall the four kinds of tasks defined in PRO*forma* on fig. 1b):

- If a *plan* is found, it has to analyse its components (a *plan* is just a way of structuring the tasks to be performed in a CPG, so that they can be represented in different levels of abstraction or detail).
- If an *action* is found, it has to be executed *in situ* by a human, external to the system. An *action* may involve a question to the patient or a direct action on the patient (such as giving an injection or measuring some clinical data -*e.g.* blood pressure- by the doctor or a nurse).
- If a *decision* is found, that means that in previous steps of the CPG the data involved in the decision have been requested (in direct *actions* or in *enquiries*).

Therefore, the DRA may evaluate the conditions involved in the decision and follow the candidate option that satisfies them.

- If an *enquiry* is found, that means that a set of values is requested (so that they can be used in posterior decisions). Some of these values may already be available in the medical record (*e.g.* the allergy to penicillin), other values may be obtained *in situ* with direct actions (*e.g.* present blood pressure), and other values may only be obtained through tests (*e.g.* a count of white corpuscles implies the need of a blood analysis). In this latter case, the DRA has to contact the agents that provide medical services (SAs, that may depend on a department or a medical centre) so that the patient may take that analysis. This process is commented on the following section.

3.4 Service coordination

When the DRA reaches an enquiry task, it may be asked a number of data that it ignores (*e.g.* the number of white corpuscles and the level of glucose in urine). Using a specific domain ontology, it finds out the services that provide those data (*e.g.* the first item can be discovered with a blood analysis, whereas the second one requires a urine test). The domain knowledge also indicates the temporal constraints among differents kinds of tests (*e.g.* there are no constraints between blood analysis and urine test).
At that moment the DRA must find SAs that perform the needed medical services, and make bookings on behalf of the patient. The coordination process is the following:

- The DRA sends a message to its associated DA, requesting the performance of a blood analysis and a urine test on the patient.
- The DA, as commented in the description of the architecture, knows about the services available in the department. If there are SAs in the department that offer those services, a negotiation between those agents and the DRA is made to schedule the performance of the services (the process is similar to the booking process in *HeCaSe*, described in detail in [7], in which the preferences of the user and its timetable are taken into account to provide a service as personalised as possible). The result could be the following: the patient may take the blood test on May 25th at 11:00, and the urine test on June 1st at 17:00; the results of the tests will be available on May 30th and June 7th, respectively. The SAs send the results of the tests to the MRA when they are available, so that they are automatically stored in the patient's medical record. After setting up the tests to be performed, the DRA and the UA may also schedule the date of the next visit (from June 8th, date in which the results of both tests will be available in the electronic medical record).
- If there is any test that is not performed by any SA of the department, the DA sends the request to the MCA, which knows the hospital-wide transversal services. We assume that a doctor of one department will only need services specific to the department or transversal to all the medical centre, but not services offered in other departments of the medical centre, which will be very specific to other kinds of pathologies, nor services of other medical centres. The MCA will tell to the DRA which agents provide the required services, and the same negotiation process (see [7]) will be followed to set up the dates of the tests and of the next visit of the patient.

4. Conclusion and future work

The system proposed in this paper is another step towards a further use of agent technology in health care, which has experimented an important growth in the last years ([3]). This system is based on our previous on the *HeCaSe* multi-agent system ([5-7]), which is being improved in two main ways:

- The proactiovity of the DRAs helps the doctors to follow the CPG associated with a disease, by reminding them the steps to be followed, the actions and decisions to be taken and the tests to be made.
- Doctors and services negotiate automatically in order to find out the dates in which the patient must take the clinical tests suggested in the CPG.

The automated coordination of the services available in a hospital may provide many benefits to patients and doctors:

- The patient knows, during the visit, when and where he must go to have the tests taken. He does not have to spend time visiting different services of the department or the medical centre to schedule the performance of the tests.
- The automatisation of the booking procedure also gives more free time to doctors and nurses to care for patients.
- The preferences of the patient and his previous commitments are taken into account, so that the tests are set up in moments which are appropriate for the patient.
- The medical data of the patient is automatically considered, and he does not have to spend time during the visit to answer routine questions (*e.g*. allergies).
- The doctor also sets up the date of the next visit knowing that the results of the tests will already be available in the patient's medical record.
- If a certain data is already available in the medical record, it is considered at that moment and no test is scheduled to obtain it needlessly.
- The CPG associated with the diagnosed disease is cautiously followed, and the doctor has less possibilities of making errors or forgetting essential tests to be performed.

The implementation of the proposed MAS is currently being made using a FIPA-compliant tool called JADE[1]. At the present stage, all agents have already been fully developed except the GA and the SAs. The medical service ontology, which expands the one implemented in the first version of the system [6], is currently under development; it basically describes the data that may be gathered in each kind of medical test. The TALLIS environment, kindly provided to us by Cancer Research UK, is being used to manage CPGs written in PRO*forma*. Some ideas for future work are the following:

- We are considering the implementation of UAs and DRAs in personal mobile devices such as PDAs or mobile phones, using the JADE-LEAP plug-in.
- The coordination procedures may be made more general, so that a doctor may request services that are located on other departments of his hospital or in other medical centres.
- The DRA now follows the CPG step by step, and it stops when it finds the first enquiry in which unknown data is requested. It could be possible to make a full

[1] *JADE (Java Agent DEvelopment framework)* is a set of Java libraries that facilitates the development of FIPA-compliant MAS. It can be downloaded from http://sharon.cselt.it/projects/jade.

analysis of the CPG, so that the DRA may for instance make pre-bookings for a number of medical tests that the patient might have to take in the future, depending on some intermediate decisions pending to be made. When the necessary data has been gathered and the decisions made, pre-bookings might be confirmed or refused.

- Doctors may choose not to follow exactly all the steps involved in a medical protocol [12], and the system should support them in these cases (*e.g.* by helping them to easily document why a certain step is not performed on a particular patient, so that the decision may be included in the patient's medical record).
- We plan to test the system with real users.

Acknowledgements

The authors would like to thank Dr. John Fox, from the Cancer Research Lab. (UK), for allowing them the use of TALLIS. The authors also want to acknowledge the rest of colleagues in the department for their helpful comments, especially D. Sánchez, J. Bocio and A. Valls.

References

[1] S. Quaglini, M. Stefanelli, A. Cavallini, G. Micieli, C. Fassino, C. Mossa, "Guideline-based careflow system", *Artificial Intelligence in Medicine* 20 (1), 5-22, 2000.

[2] S. Quaglini, M. Stefanelli, G. Lanzola, V. Caporusso, S. Panzarasa, "Flexible guideline-based patient careflow systems", *Artificial Intelligence in Medicine* 22 (1), 65-80, 2001.

[3] J. L. Nealon, A. Moreno, "Agent-Based Applications in Health Care," in *Applications of Software Agent Technology in the Health Care Domain*, Birkhäuser Verlag, pp. 3-18, 2003.

[4] M. Wooldridge, *An Introduction to Multiagent Systems*, J. Wiley Ltd, 2002.

[5] A. Moreno, D. Sánchez, D. Isern, "Security Measures in a Medical Multi-Agent System," in *Artificial Intelligence Research and Development*, IOS Press, pp. 244-255, 2003.

[6] A. Moreno, D. Isern, D. Sánchez, "Provision of agent-based health care services," *AI Communications*, 16 (3): 167-178, 2003.

[7] D. Isern, D. Sánchez, A. Moreno, A. Valls, "Personalisation of Medical Services," in *e-Health: Application of Computing Science in Medicine and Health Care, Research on Computer Science*, Instituto Politécnico Nacional, pp. 18-31, 2003.

[8] M. Peleg, S. Tu, J. Bury, P. Ciccarese, J. Fox, R. A. Greenes, R. Hall, P. D. Johnson, N. Jones, A. Kumar, S. Miksch, S. Quaglini, A. Seyfang, E. H. Shortliffe, M. Stefanelli, "Comparing Computer-Interpretable Guideline Models: A Case-Study Approach," *JAMIA*, 10 (1): 52-68, 2003.

[9] D. A. Wang, M. Peleg, S. Tu, E. H. Shortliffe, R. A. Greenes, "Representation of Clinical Practice Guidelines," *Medinfo*, London, UK, September 2-5, 2001.

[10] J. Bury, J. Fox, D. Sutton, "The PROforma guideline specification language: progress and prospects," *In EWGLP 2000*, Leipzig, Germany, 2000.

[11] A. Vollebregt, A. t. Teije, F. v. Harmelen, J. v. d. Lei, M. Mosseveld, "A study of PROforma, a development methodology for clinical procedures," *Artificial Intelligence in Medicine*, 17: 195-221, 1999.

[12] S. Quaglini, M. Stefanelli, V. Caporusso, S. Panzarasa, "Managing non-compliance in guideline-based careflow systems". AMIA 2000 Symposium, USA, 2000.