Multi-Agent Patient Representation in Primary Care

Tracking Number: C0162

Abstract. Though multi-agent systems have been explored in a wide variety of medical settings, their role at the primary care level has been relatively little investigated. In this paper, we present a system that is currently being piloted for future rollout in Scotland that employs an industrial strength multi-agent platform to tackle both technical and sociological challenges within primary care. In particular, the work is motivated by several specific issues: (i) the need to widen mechanisms for access to primary care; (ii) the need to harness technical solutions to reduce load not only for general practitioners, but also for practice nurses and administrators; (iii) the need to design and deploy technical solutions in such a way that they fit in to existing professional activity, rather than demanding changes in current practice. With direct representation of individuals in health care relationships implemented in a multi-agent system (with one multi-functional agents representing each patient, doctor, nurse, pharmacist, etc.) it becomes straightforward first to model and then to integrate with existing practice. It is for this reason that the system described here successfully widens access for patients (by opening up novel communication channels of email and SMS texting) and reduces load on the practice (by streamlining communications and semi-automating appointment arrangement). It does this by ensuring that the solution is not imposed on, but rather, integrated with what currently goes on in primary care. Furthermore, with agents responsible for maintaining audit trails for the patients they represent, it becomes possible to see elements of the electronic patient record (EPR) emerging under agent control. This EPR can be extended through structured interaction with the practice system (here, we examine the GPASS system, the market leader in Scotland), to allow rich agent-agent and agent-human interactions. By using multi-agent design and implementation techniques, we have been able to build a solution that integrates both with individuals and extant software to successfully tackle real problems in primary care.

1 Introduction

Increasingly, research in multi-agent systems is exploring the advantages offered by the emerging multi-agent software engineering paradigm (see, for example [13]). By focusing on the relationship between an entity in the real world (an individual, a relationship, an organisation, a datum) and its corresponding entity in the "electronic world" (an agent, a relationship between agents, a set of agents, a datum), design can become easier and quicker [4]. The process of teasing out complex and intricate stakeholder relationships in real world domains is not made harder by the restrictions and assumptions of the implementation paradigm. As a result, interesting multi-agent based models of complex social structures have started to emerge [22].

The health care system is a perfect example. It offers extremely rich interdependent sets of relationships, stakeholders, rights, requirements and agendas that, though interesting for the modeller, have proved an enormous challenge for the deployment and uptake of practical IT systems [15]. It is an environment subject to constant change of workforce, 'clients', and infrastructure, often with a zero tolerance of error. The challenge for multi-agent systems presented by the complexities of health care has been well documented (see e.g. [10, 12, 16]) and has provided a rich field for research, but primary care - and its own unique set of issues - is often omitted from these investigations. The work described here harnesses agentive representation and techniques from agent oriented programming in tackling some of these features of primary care.

2 Background

Patients, politicians and health care professionals all agree that Patient Centred Care is a good thing, but translation of concept to reality has yet to be achieved [23]. The fields of health and computing share the common jargon of 'user friendly', accessible, and flexibility and so it is unsurprising that agents have been touted for patient-centered health care.

One of the earliest examples of work examining the role of multi-agent systems in health care is offered by [10]. The focus of the work presented there, and of the broader context in which it was conducted (*viz*., the DILEMMA project), is upon appropriate theorem proving in decision support systems that have to deal with complex, incomplete, inconsistent and potentially conflicting data. The agent component is designed to support the distribution of tasks amongst players in the system, in a manner similar to the much earlier Contract Net protocol for automated task distribution [24]. A prototype of the AADCare system as a whole was implemented for the management of cancer patients in the UK NHS system, though the extent of deployment and its subsequent success is unclear. Crucially, access for the patient to their medical record, and the unification of record components across different health services for an individual patient was not a focus for AADCare in either implementation or theory, and therefore mostly side-steps issues of patient-centered health care.

The Guardian Angel project [8] represents a "manifesto" developed since 1994 that tackles the patient-centered approach head on. Some elements of the manifesto have lead to implementation, of which the earliest was the Personal Internetworked Notary and Guardian project, PING [21]. Although that work mentions "agents" in passing, its focus is upon implementing basic security mechanisms for (conceptually) centralised data stored using XML. It makes no use of the agent oriented, peer-to-peer approach in either design or implementation. The motivating concepts, described in [14] however, are precisely those addressed in the current work: the need to balance patient access and security; the need to reduce fragmentation in medical records; the need for IT infrastructure to be interoperable; etc.

In the context of UK health care, Pouloudi and Reed [20] offer a relatively early example of using multi-agent systems to represent and model interactions in the NHS in an attempt to build a realistic foundation for integrative systems. The work combines intra-agent representational concepts with inter-agent communication and relationship structures in modelling the interactions between stakeholder relationships in patient data. The model there was theoretical and unimplemented.

More recently, the Advanced Computational Lab at Cancer Research UK has built multi-agent models of the same sorts of complex relationships specifically in the context of cancer, from initial patient contact with their GP through various stages of care and maintenance [2; 7]. They employ the mature COGENT and PRO*forma* tools and the tried-and-tested Domino model of agent architecture, but still the focus is squarely upon the interacting agencies of the health system, for which the patient is simply a customer.

Moreno *et al*. [16] describes a system that moves closer to the ideal of patient centered involvement and access. In their HeCaSe system, patients have an interface that supports appointment booking and various static configuration parameters. It is, however, focused only on the interaction between patients and initial, primary care consultations, and is run from PC clients. Crucially from a deployment point of view, it requires doctors to switch to a new system.

Finally, it is worth noting that perhaps the largest impact of multi-agent systems in health care to date has been in specifically targeted applications that focus on particular functions of the health system. So, for example, there are prototypes and demonstrators of multi-agent system applications in areas such as organ transplant [6; 17], antibiotic prescription [9], pharmacy in general [3], protocol monitoring [1], proactive information provision in anaesthesia [11], data flow in Leukemia management [12] and others (such as those in the special issue, volume 27 issue 3, of *AI in Medicine*). These examples are clinician centred attempts to streamline existing processes. Thus multi-agent systems as a tool is having an impact in many areas. But this is peripheral to the argument that we hope to make here, namely, that multi-agent systems as a paradigm fits the goal of patient-centered health care perfectly not only at a conceptual level, but also in implementation and deployment.

3 Multi-Agent Systems for Patient-Centered Health Care

The concept of **agentive representation** is implicit in very many agent-based models of real world structures, and even entire agent based methodologies such as Gaia [25]. The idea is simply that one component in the real world is represented by a single corresponding agent in the system. This idea is now also starting to gain traction in the commercial world [4]. In the medical domain, agentive representation means agents representing general practitioners, consultants, pharmacists, and, of course, patients.

To build systems that are to be deployed in real health care situations, it is important that the infrastructure meets and exceeds a range of basic expectations of the users with respect to various aspects such as security, scaling and reliability. In the work described here, we have selected the JUDE platform [5] for reasons of flexibility and robustness. The architecture of agents in JUDE is simple in that each agent is equipped with a set of generic functionality (such as basic reasoning and communication) that can then be augmented with additional modular functionality as needed.

The system implements the patient-centered approach by equipping agents representing patients with all the functionality they require to represent their corresponding patient in the electronic health care world. So for example, the patient agent can access data on that patient held in different locations and by different parties. The patient agent can communicate with the local surgery to organise appointments. The patient agent can access information on pharmacy location and availability. And, of course, the patient agent can communicate with, and be contacted by, the patient themselves. This communication can make use of whatever channels may happen to be available at a given moment – from web to SMS. But every time, and in every case, the patient is simply communicating with their own, persistent, agent.

Detailing the implementation in full is beyond the scope of this paper, but it is useful to offer depth in a subset of the functionality and go on to show how this functionality is deployed and fitted in to existing primary care processes.

3.1 An Example: Reducing DNA rates

In this currently live trial, patient agents are equipped with the ability to interact with agents representing individuals in a GP surgery, including the receptionist and GP. One agent-mediated interaction is the process of booking and confirming appointments. The appointments process is a current area of interest in practice management, since a substantial proportion of valuable GP time is wasted as a result of people who book appointments but then subsequently do not attend (DNA). Reducing the DNA rate offers practical and substantial advantages to GPs and practices in the UK.

The current model is to allow patients to ring the practice receptionist and negotiate verbally to arrange an appointment time convenient for both GP and patient. In some cases, as described above, some or all of this process may be conducted by email instead of over the telephone.

An agent-based solution offers a technical improvement whilst integrating with existing practice to minimise barriers to use. A patient's agent is responsible for intervening in some or all of the communication between the patient and the practice, and is responsible for reminding the patient of upcoming appointments. A patient is allocated an agent in the system following consent agreement. At that point, the patient can send a text message or an email to their agent, which, in either case, then communicates with the agent representing the practice receptionist. The receptionist's agent then communicates with the receptionist through the most appropriate means – at the moment, that is email. The negotiation is conducted in this way between patient and receptionist until agreement on appointment time is met. (It is unreasonable to expect patients to be using an electronic diary, and thus it is not possible to automate the patient end of the negotiation process. Similarly, it is important to keep the receptionist in the loop, and so automating that end is also counterproductive.) When agreement is reached, the receptionist confirms the appointment through a web interface provided by the receptionist's agent. That agent then informs the patient's agent of the confirmed appointment time. At 24 hours before the appointment, the patient's agent sends a reminder (currently by SMS). At two hours before the appointment is due, the patient's agent sends a second reminder. If the patient does not reply to that second reminder, thereby failing to confirm that they still intend to keep the appointment, the patient's agent will inform the receptionist's agent of a problem. In this case, the receptionist's agent will take some default action, which is currently to email the receptionist suggesting that the appointment be cancelled and the time freed up.

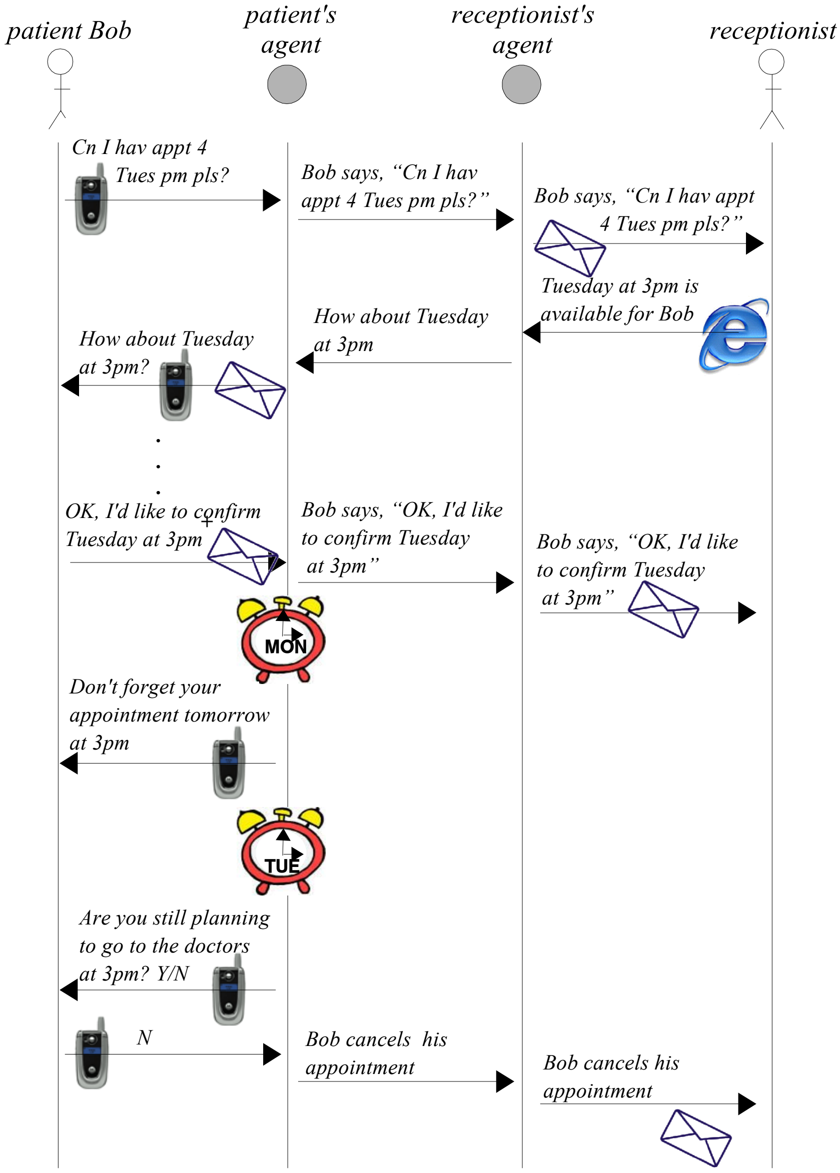

Fig. 1. Sample Interaction

Figure 1, above, summarises an example interaction, demonstrating the various communication mechanisms (text, email and web, indicated by the icons on the interaction arrows), the simple proactive work of the patient's agent, and the involvement of the receptionist and patient in the system.

4 Realities of using a Multi-Agent System in Primary Care

The published work on MAS applications in health care understandably concentrates on the computing methodology and on the need to design software around complex health care problems. It will become increasingly important to involve the people using the system – clinicians and patients - in future developments. Research into any new technology, including MAS, must look at subtle patient centred issues. Innovations need to be based around what people actually want to use. Multi-agent systems have the potential to support an ongoing dialogue between doctor and patient. Research must therefore include an ongoing dialogue to plan and implement findings.

The "Do Not Attend" component of the system described in the previous section is adopted as an exemplar by which to follow through this process. By drawing on previous experience of integrating new technology into the patient-facing side of primary care [18; 19], the design of pilot trials and ethical approval for those trials has been achievable. The current deployment is based around 'as needed' creation of agents for up to 100 patients who volunteer to take part in the trial. Text messaging based appointment booking is offered as an extra service.

For the receptionist, all interaction with the system, and with patients using the system is through a simple web interface and email, both of which are familiar and non-threatening. The interface between the system and patients is text messaging, with which all who sign up are very familiar. Finally, the interface between the system and additional components (such as GPs) is built around email, again fitting in with existing practices. As a result, introduction of the service has been straightforward and has hit no major problems.

Clearly, the system, its deployment and its uptake are in very early days, but the current status marks an important milestone. The key step that has been taken is to provide each individual patient with an agent that interacts, on the patient's behalf, with agents representing health care professionals. With this scenario engineered and deployed, and with the dynamic upgrade facilities in JUDE (whereby new functionality can be deployed to existing agents without needing to "reboot" either the systems or those agents that are upgraded), it is relatively easy to map out a programme of rolling development and deployment of patient-centered services.

With appointment booking and reminders in place, the next step is to integrate another burdensome practice task: repeat prescription ordering. Ordering repeat prescriptions is, in the UK and other health systems, a task initiated by a patient who has a long-term need for prescription drugs. It requires a GP to sign off, which takes both GP time, and a physical appointment, involving both travel and time for the patient, as well as practice overhead in the form of receptionist time for booking. With agents representing all the parties, it becomes relatively simple for the request to be routed from the patient to the GP, confirmed (or otherwise) by the GP, and then forwarded to the appropriate pharmacist, from where the patient can collect their prescription. Of course this is far from the first time that a proposal has been tabled for streamlining this process, nor is it the only IT-based solution. The novelty is that it is exactly the same system as is currently used for appointment booking, made available for repeat prescription ordering by virtue of the direct representation of the people involved and the relationships between them. One of the additional benefits that arises “for free” with the approach is an audit trail at the individual patient level, so that health care professionals and patients can track where a request is in the system so cutting out unnecessary phone calls to the practice or pharmacist. Also, with the advent of electronic prescribing, our agent system represents a solution anchored to existing practice systems.

5 An EPR under Agent Control?

In Scotland, well over 80% of primary care practices have adopted the GPASS system to manage patient data and other practice functions[1]. With a wide variety of proposals for electronic patient record (EPR) management currently under discussion and development, integration with primary care systems is a key concern. GPASS, like many of its competitors, does provide programmatic access to the data it stores, via mediated SQL queries. Such a clean interface makes it a good starting point for investigating the alignment of the notion of an EPR with that of one agent for every patient.

The EPR can be analysed to yield a key subset of patient medical information which has general utility for both patients and health-care professionals. This subset contains information regarding a patient's significant clinical events, current prescribed medication (both acute, and repeat), key indicators such as tobacco and alcohol history, and their age, height, weight and body mass index. Together with primary personal information – name, address, date-of-birth, and CHI (unique health service number) – this constitutes what we call the Core Clinical Summary (CCS).

The information which makes up the CCS is stored in the MS SQL database which acts as the data store for the GPASS clinical system. In order to access the information required to construct the CCS, a JUDE module was developed for use by the appropriate agents to log into and make the relevant queries to GPASS, from which those agents might then construct the CCS accordingly. Figure 3 sumamrises the mechanisms by which this information is extracted in response to a patient initiated request.

1 See http://www.gpass.co.uk

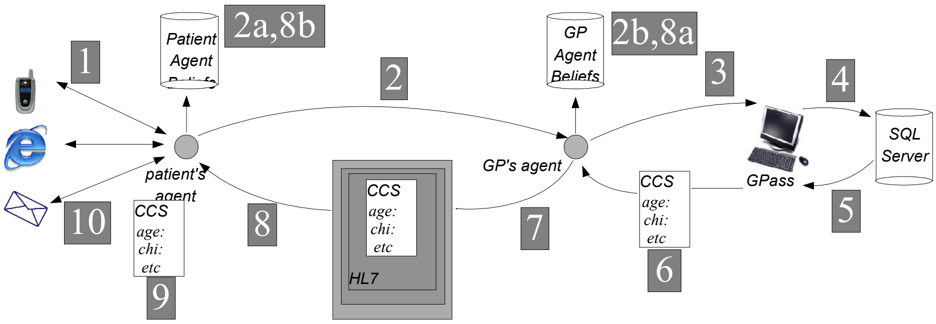

Fig. 3. Sample Patient-Agent-GP-GPASS interaction

In this example, at (1), the patient requests some part of CCS data via the web, email or their mobile. Then, at (2), the patient's agent requests CCS from the GP's agent, and by the implemented semantics of the inter-agent communication language, the Patient's agent (2a) and GP's agent (2b) update their respective belief sets to reflect that the request was made. At (3) the GP's agent logs in to GPASS clinical system and then makes set of requests to GPASS for patient information. In turn, GPASS calls into the SQL database (4) which returns the appropriate patient data (5). At (6) the GP's agent then uses that data to construct the CCS which it then embeds in the appropriate nested wrappers of a well formed HL7 message for subsequent transmission (7). At (8), the GP's agent informs the Patient's agent with the CCS, again leading to updates in the belief sets of the GP's agent (8a) and Patient's agent (8b).Finally, at (9) the Patients agent extracts CCS from HL7 wrappers and at (10) transmits CCS data to appropriate device of patient (where what is appropriate is determined through rich contextual reasoning). This same architecture also fits snugly into a more radical, long-term and ambitious picture in which patient agents have ultimate responsibility for maintaining an up-to-date picture of the EPR.

The ability of a patient to have to hand their own medical information via entirely ordinary devices, promotes the position of patient as genuine stakeholder in their own health care. The information available to patients via the CCS allows them to accurately inform their lifestyle using straightforward mechanisms, and on their terms. So, for example, a patient can easily match their recollection of past clinical events with the corresponding sections of their CCS. Further, in situations where clinical systems are not available (e.g., an accident in a remote location) the patient themselves has the means to provide key information (such as allergies, or current medications) to assist the health care professional.

Putting the patient in a position where they can easily interact proactively with their representation in the health care system also supports more interactive relationships between health care providers and patients. For GP's (and other health care professionals), the MAS based primary care solution provides all parties from all aspects of health care with a single point of contact to the patient, *viz.*, the patients agent. Whilst providing GPS with mobile access to CCS records may be convenient (immediately prior to house calls, for example), in positions where urgent access to medical information is vital, such as an emergency situation with an incapacitated patient outwith the health care infrastructure, the ability for the patient's agent to be accessed for information in the CCS (such as current medication regimes or illnesses) offers the potential for significant advantages.

6 Conclusions and Directions

The focus of this work is on the design, implementation, deployment and evaluation of initially very simple solutions to very real problems that despite their simplicity nevertheless use a full, mature, industrial-strength multi-agent system that can not only tackle the simple problems as they scale numerically, but also incrementally take on an ever more significant role, tackling ever more complex problems that demand ever more of the underlying technology.

Multi-agent systems have great potential to transform the process, and possibly even the outcome, of medical care. It is important that innovation goes hand in hand with evaluation. Our own next steps are to evaluate the outcomes of the pilot trial in designing larger scale trials which are based upon enriched functionality within individual agents. It is a logical extension to offer a MAS-based text message/email/web service including appointment booking, repeat prescription ordering and provision of clinical advice. Qualitative work will explore the views, aspirations and experiences of people who chose to use this service, with the first component of this work – health care staff interviewing – already almost complete. Quantitative work will measure when and how often people use the service, and the impact that the service has on the functioning of the practice. It will be important to look at the language and exchange of information in any patient – doctor text dialogue. Preliminary technical work and evaluation will then allow a larger cluster randomised trial to proceed in several practices involving hundreds of patients.

Multi-agent systems are likely to transform health care within the next decade. At a basic level the very nature of a consultation between a patient and a health care professional will need to be re-defined. At a more sophisticated level, patients will have the opportunity to integrate health-related activity and decision making into everyday behaviour. Professionals will be able to support and advise patients in a completely new and different way. The effect of these changes on health outcome needs to be addressed. The reactions of people using and working within the health service need to be explored. The major barrier to the implementation of multi-agent systems in health care may not be technical, but attitudinal, and it is this aspect that needs to be brought in to every stage of the development and deployment lifecycle.

References

1. Alsinet, T., Ansotegui, C., Bejar, R., Fernandez, C., Manya, F.: Automated monitoring of medical protocols: a secure and distributed architecture, Art. Int. in Med. 27(3) (2003) 367-392
2. Black, E.: Using Agent Technology to Model Complex Organisations AgentLink News 12 (2003)

3. Calabretto, J.P., Couper, D., Mulley, B., Nissen, M., Siow, S., Tuck, J., Warren, J.: Agent Support for Patients and Community Pharmacists, Proc. of the 35th HICSS, IEEE (2002)
4. Calico Jack: Multi-Agent Systems in Mobile Telecommunications Networks, Calico Jack White Paper #2004/01 (2004)
5. Calico Jack: "An Introduction to the Jackdaw University Development Environment", Calico Jack Developer Documentation (2004)
6. Calisti, M., Funk, P., Biellman, P., Bugnon, T.: Multi-Agent System for Organ Transplant Management, in (Moreno & Nealon, 2003) 199-212
7. Fox, J., Beveridge, M., Glasspool, D.: Understanding intelligent agents: analysis and synthesis, AI Communications 16 (3) (2003) 139-152
8. GA, http://www.ga.org Guardian Angel website, last updated 15 March 2002.
9. Godo, L., Puyol-Gruart, J., Sabater, J., Torra, V., Barrufet, P., Fabregas, X.: A multi-agent systems approach for monitoring the prescription of restricted use antibiotics, Art. Int. in Med. 27(3) (2003) 259-282.
10. Huang, J., Jennings, N.R., Fox, J.: An Agent-based Approach to Health Care Management. Applied Artificial Intelligence 9 (4) (1995) 401-420
11. Knublauch, H., Rose, T., Sedlmayr, M.: Towards a Multi-Agent System for Pro-Active Information Management in Anaesthesia. In Proc. of the Workshop on Autonomous Agents in Health Care at Agents 2000, Barcelona, Spain 2000.
12. Lanzola, G. Gatti, L. Falasconi, S., Stefanelli, M.: A Framework for Building Co-operative Software Agents in Medical Applications. Art. Int. in Med., 16(3) (1999) 223-249
13. Luck, M., Ashri, R. & D'Inverno, M.: *Agent-Based Software Development*, Artech 2004.
14. Mandl, K.D., Szolovits, P., Kohane, I.S.: Public Standards and Patients' Control: How to Keep Electronic Medical Records Accessible But Private. BMJ 322 (2001) 283-287.
15. Mark, A., Pencheon, D., Elliott, R.: Demanding Healthcare. Intl. J. of Health Planning & Mgmt. 15 (2000) 237-253
16. Moreno, A., Isern, D., Sanchez, D.: Provision of Agent-Based Healthcare Services, AI Comms. 16 (3) (2003) 167-178
17. Moreno, A., Valls, A., Bocio, J.: Management of Hospital Teams for Organ Transplant Coordination, Proc. of AIME 2001 (2001)
18. Neville R.G., Greene A.C., McLeod J., Tracy A., Surie J.: Mobile phone text messaging can help young people manage asthma. BMJ 2002 325 600
19. Neville R.G., McCowan C., Pagliari H.C., Mullen H., Fannin A.: E-mail Consultations in General Practice (in press)
20. Pouloudi A., Reed, C.A.: Towards a Multi-agent representation of stakeholder interests in NHSnet, Proc. PAAM98 (1998) 393-407
21 Riva, A., Mandl, K.D., Oh, D.H., Nigrin, D.J., Butte, A., Szolovits, P., Kohane, I.S.: The Personal Internetworked Notary and Guardian. Intl. J. of Med. Informatics 62 (2001) 27-40
22. Sabater, J., Sierra, C.: Reputation and Social Network Analysis in Multi-Agent Systems, Proc. AAMAS 2002, ACM Press.
23. White Paper – Patients as Partners in Care: A Plan for Action a Plan for Change HMSO (2003 The Scottish Executive
24. Smith, R.G.: The Contract Net Protocol: High Level Communication and Control in a Distributed Problem Solver. IEEE Trans. Computers 29 (12) (1980) 1104-1113
25. Wooldridge, M., Jennings, N.R., Kinny, D.: The Gaia Methodology for Agent-Oriented Analysis and Design. J. of Auton. Agents and Multi-Agent Sys. 3 (3) (2000) 285-312