Towards a Multi-Agent Representation of Stakeholder Interests

Athanasia Pouloudi and Chris Reed
Department of Information Systems and Computing
Brunel University
Uxbridge UB8 3PH
{Nancy.Pouloudi, Chris.Reed}@brunel.ac.uk

Abstract
Stakeholder identification and analysis plays a key role in the design and management of complex interorganisational systems, creating a vast catalogue of interrelated vested interests. This rich source of information, however, can quickly become unmanageable and opaque. This paper presents an agent-theoretic approach to modelling the interests of, and communication between, stakeholders, and explores the relationships between the two concepts. The agent model is founded on notions of argumentative support which then form the basis of the communication protocol. Through analysis of an extended case study the benefits of this approach are clearly demonstrated and the complex social structure of stakeholder interrelations is examined. It is posited that this approach facilitates both the initial stage of identifying stakeholders in a system and also the subsequent process in which the stakeholders reshape their interests and arguments as the negotiation process develops.

1. Introduction

This paper sets out to explore the relation between the concepts of stakeholder and agent. In particular, an examination is presented of the way in which multi-agent technology can be used to represent stakeholders, their interests and their interrelations. The ADEPT project (Jennings *et al.*, 1996; Sierra *et al.*, 1997) has demonstrated that agent technology can be successfully applied in modelling business processes. As a key component of business process management, stakeholder analysis and its effects might also be suspected to be amenable to modelling in agent-theoretic terms.

In order to explore the utility of the representation in a complex domain, the paper uses the example of a network implemented in the healthcare context in the United Kingdom, where multiple parties with diverse interests are involved. The paper uses this case study to illustrate how these stakeholders, their beliefs and the issues emerging from their conflicting interests can be represented in a consistent and systematic way. Furthermore, we explore the benefits and implications of this modelling for understanding stakeholder behaviour.

The next section introduces an interpretive approach to stakeholder analysis that has previously been applied in information systems research. We then use this approach to describe our case study, the NHSnet, in terms of some key stakeholders that can affect or be affected by the network and are for different reasons 'interested' in its state of development. Section four models the implicit and explicit stakeholder beliefs and gives an example of how conflicting issues can emerge and evolve over time as a result of the communication between different types of agents. The paper concludes by outlining the contribution resulting from the interaction between the two research areas and by making recommendations regarding further research directions.

2. Stakeholder analysis for interorganisational systems

Computer systems and networks that transcend the boundaries of a single organisation, i.e., interorganisational systems (Cash & Konsynski, 1985), affect and are affected by a large number of groups and organisations; these are termed *stakeholders*. An interpretive stakeholder analysis approach can be used for identifying these stakeholders as well as their interests, their perceptions of interorganisational systems and the way in which they present their interests and interact with other stakeholders.

Most stakeholder analysis approaches that have been used in both a strategic management and an information systems context have failed to give guidance for the identification of stakeholders assuming that stakeholders can readily be identified (e.g., Freeman, 1984; Richardson & Richardson, 1992; Ruohonen, 1991). We would argue that this might not be so particularly when issues such as the exchange of information and its electronic support are at stake. In such cases stakeholder identification is not necessarily a straightforward process. It is useful however as it can help to unveil the broader context within which interorganisational systems are used. Recent research in interorganisational systems (Pouloudi & Whitley, 1997) has suggested a number of principles underlying stakeholder behaviour. These can guide the analyst in the identification, and eventually in the analysis of the stakeholders' perspectives that seem relevant. These principles and their analytical implications are summarised in Table 1.

Principles of stakeholder behaviour	**Implications for interorganisational systems stakeholders' identification and analysis**
stakeholders depend on the particular context and change over time	generic (and static) stakeholder lists are inadequate; stakeholder identification should reflect the context and be reviewed over time
stakeholders do not exist in isolation; they interact and exchange information	stakeholders can be identified through a progressive, iterative process that follows up direct and implicit stakeholder interactions (including those in informal fora)
information systems stakes may change over time	stakeholder analysis should look into current as well as previous perceptions of stakeholders and investigate how these may evolve (cf. Pouloudi & Whitley, 1996)
some stakeholders' wishes may not be realised	conflicting stakeholders' expectations of an interorganisational systems should be investigated: they will probably influence the future of the system

Table 1 Guidelines for the identification and analysis of interorganisational systems stakeholders (Adapted from Pouloudi & Whitley, 1997)

The implication of these principles is that the future of an interorganisational system does not only rely on its technical feasibility and the availability of resources, but also on the perceived benefits or problems that it brings to its stakeholders. Therefore, stakeholder analysis as a method for the investigation of interorganisational systems development and use cannot stop at the identification of a broad number of stakeholders. Examining the interactions between the stakeholders identified and understanding the reasons why different stakeholders may have different ideas and feelings about the evolution and the future of an interorganisational system is equally important; it allows understanding of which issues matter for different stakeholders as well as what the explicit and implicit conflicts on these issues are. Of course, stakeholder identification and understanding the stakeholders' perceptions are not independent activities. This is evident from the previous presentation of principles of stakeholder behaviour and implications for stakeholder identification and analysis. For example, some stakeholders may be unable to identify further interested parties unless they are given an opportunity to raise their concerns or discuss what they believe are important issues relating to the evolution, current state and future of the interorganisational system under investigation.

This approach to stakeholder analysis is interpretive in that it accommodates the different perceptions of stakeholders concerning the way in which an interorganisational system is working and progressing. Indeed, different stakeholders may have different views about why a system is progressing (or not progressing) in a given direction. All conflicting views need to be seen as legitimate, representing the background of particular stakeholders, but also their distinct interest in participating in (or abstaining from) the interorganisational system. Certainly, some stakeholders will only present what they want others to perceive as their interest. For this reason it is interesting to investigate what ultimate motives other stakeholders attribute to their behaviour and perspectives. An interpretive approach means each of these opinions, however controversial, is respected, even though the interpretations of the researcher in the presentation of the stakeholders' views (including the choice of which views are relevant) and in pointing to some issues as more important for the future of the interorganisational system under scrutiny are also present. It is this interpretive approach which forms one of the key features of the agent-theoretic analysis presented in section 4.

To illustrate the value of this analysis we have taken the example of an interorganisational network recently implemented in the healthcare sector in Britain. The next section describes the project and focuses on some of the problematic issues for its implementation. This is then followed by a discussion focusing on the stakeholders of the network and their perceptions concerning its current and future use.

3. The NHSnet

The Information Management Group of the NHS Executive, the body responsible for the execution of health care policy in Britain (NHS Executive, 1994b), launched the NHS-wide networking project in 1993, as "an integrated approach to interorganisational communications within the NHS" (NHS Executive, 1994a p. 6). The objective of this network has been to enhance communication and information exchange between various health care providers and administrators. Thus, the NHSnet is expected to support data communications that cover a variety of information flows across different levels. Its infrastructure is expected to cover a variety of business areas, including patient related service delivery, patient related administration, commissioning and contracting, information services, management related flows and supplies of NHS organisations (NHS Executive, 1995).

The NHSnet is available since 1996. Yet, despite the technological success of the project, and in particular its completion within schedule, its implementation has suffered from the lack of acceptance by the medical profession. Doctors remain sceptical mainly of the security that the network has to offer. These concerns have been overtly voiced, mainly by the British Medical Association (BMA), the national professional body of physicians in the United Kingdom, but also by computer security consultants. These parties fear that patient data may be misused by both NHS members (referred to as "insiders") and external parties (Willcox, 1995).

As a result of their concern, doctors, again through the voice of the BMA, threatened not to participate in the electronic exchange of data unless they can be convinced that patient privacy is safeguarded. On the other hand, the NHS Executive have stated that the proposed system will be better than the previous: data confidentiality was quoted as one of the shortcomings of the previous situation and one that the NHS-wide networking infrastructure would safeguard (NHS Executive, 1994a). A recent conference in Healthcare Computing (18-20 March 1996, Harrogate, United Kingdom) provided the opportunity for a direct confrontation of the two sides on the matter:

> When you create a large database with information on millions of people you create an extremely valuable resource. British medical records as held in GP [general practice] organisations could be worth £2bn if they can be mined and resold to insurance companies, pharmaceutical companies and the like. (Dr Ross Anderson, Security Advisor, BMA)

> The measures we have put in place are to stop anybody who is unauthorised getting at data form, and via, the [NHS-wide networking] system and one of the key parts of that system is a strong authentication challenge. (Ray Rogers, Executive Director, NHS Information Management Group)

(Both quotations are reported in the British Journal of Healthcare Computing and Information Management, vol. 13, no. 3, 1996, p. 6).

These views are typical of the concerns voiced by stakeholders. It is worth noting that the concerns on confidentiality and patient-identifiable information and the debates about alternative solutions are ongoing (e.g., Barber, 1998; Turner, 1998), particularly as Britain has to conform to the European Union Directive regarding the protection of personal data (95/46/EC). The next section provides a summary of the interests of four major stakeholders in the NHSnet case, which will serve as the bases for the development of an agent model in section 4.

3.1 A complex picture of stakeholder views

The application of the stakeholder analysis approach outlined in the second section reveals a complex picture of stakeholders and interests (see Table 2 for a summary of all the stakeholders that have been identified). For the purposes of this paper we will concentrate on a subset of those NHSnet stakeholders that have been more actively involved in or affected by the debate concerning the future use of the network (for a more detailed account of the NHSnet stakeholders and their interests see Pouloudi, 1997). These include the doctors and their representative organisations, the patients, the security consultants and the NHS Executive.

The 'connected' NHSnet stakeholders	The 'unconnected' NHSnet stakeholders
doctors (**GPs** and hospital doctors)	**patients**
local medical committees (LMCs)	Data Protection Registrar
British Medical Association (**BMA**)	legal organizations; pressure groups; media
health authorities	Members of Parliament; Secretary of Health
hospital management	Insurance companies
central and local communication management groups	pharmacies
	local pharmaceutical committees; the PSNC
national and local user representative groups	pharmaceutical companies
Prescription Pricing Authority	Medicines Control Agency
Department of Health	IT & telecommunications suppliers
NHS Executive (**NHS-Ex**)	**computer security consultants (SecCon)**
Information Management Group (of the NHS-Ex)	GCHQ
	researchers

Table 2 An Overview of the NHSnet stakeholders (adopted from Pouloudi, 1997)

Doctors in primary care (GPs) and their representative organisations

In the United Kingdom, patients register with a general practitioner (GP), who is responsible for delivering primary care and who acts as a 'gatekeeper' between primary and secondary care. Patient information is therefore in most cases collected and maintained by the GP, who is increasingly assisted in this task by computer systems. The existing variant level of computerisation in a practice (cf. Gillies, 1995; Hayes, 1997) affects the attitude towards participation in the NHSnet. Although some GPs can already exchange administrative information with health authorities using this network, more benefits are expected if they can also exchange clinical patient information with other healthcare providers. This need emerges primarily when a patient is referred to specialists or hospital doctors.

Doctors and the BMA are particularly concerned with the issue of confidentiality over such information exchange, as is evident from their active involvement in the NHSnet debate (Creasey, 1996). It is worth noting that not all GPs were equally aware of (and therefore concerned about) the confidentiality issues that may result from the use of the network. However, since the BMA raised the issue of the confidentiality they realised that the patients' privacy may be compromised unless the information is exchanged across a network that is safe from interceptions and received only by the professionals who need this information to deliver care.

Some of those stakeholders who question the GPs' concerns (e.g., employees of the NHS Executive) argue that the doctors use data confidentiality also as a pretext for controlling the exchange of patient information, and thus protecting and enhancing their role in healthcare. Other stakeholders, including patients, justify the attitude of doctors on the basis that such information they hold has been disclosed as a consequence of the patients' trust in GPs. Making this information more widely available would damage the doctor-patient relationship. The doctors fear that patients may then refuse to disclose sensitive information to the doctor, with unpredictable effects for diagnosis and hence the provision of appropriate care.

The NHS Executive (NHS-Ex)

The NHS Executive (NHS-Ex) is responsible for executing health policies as set by the Department of Health in the UK. In this capacity they are resented by most of the other stakeholders for failing to satisfy their expectations. Yet, most stakeholders recognise their proactive role in putting forward ambitious policies for the use of information technology in healthcare and in seeing that the projects are implemented within the set guidelines. It is worth noting that they normally involve users in piloting new projects, as implementations are typically more easy and successful if they result from co-operation with willing parties. The NHSnet is a characteristic example; following the confidentiality debate, a number of the suggestions of the BMA have been taken on board, and this has been a key factor enabling the progress of the project.

Patients

Although the confidentiality of patient data is at the heart of the NHSnet debate, the patients, either as individuals or through associations representing their interests – those of patients in general or of specific patient groups – seem absent from the debate and the negotiations about the security of the network. Rather, they rely on the initiatives of the medical profession. According to the doctors, this reflects the trust of patients in their GPs. Not all patients accept this explanation, but at the same time they may be unaware of any mechanisms that would allow them to formally phrase their concerns; the British Medical Association is perceived as a more powerful body for negotiating with the government.

Computer security consultants

Given the complexity of the NHSnet debate on security, it is not surprising that computer security consultants became actively involved, particularly in terms of consulting the BMA on security matters and often representing them in the conflict with the NHS Executive. Other security consultants collaborate with the NHS Executive. Thus, Zergo Limited has been commissioned by the Information Management Group of the NHS Executive to “undertake a study looking at the ramifications of using encryption and related services across the NHS-Wide Network” (NHS Executive, 1996). Although ‘the Zergo report’ was interpreted as a willingness of the government to take the doctors’ concerns over confidentiality seriously, it also generated new issues regarding the appropriate use of encryption and access to encryption mechanisms. Finally, security and privacy specialists have also become involved in the debate to create awareness in the patient population about the dangers of the electronic exchange of healthcare data (Anderson, 1996; Bywater & Wilkins, 1996; Cohen, 1996; Davies, 1996).

As expected, the identification of stakeholders and the discussion of their perceptions in relation to the current and future use of the NHSnet reveal a very complicated picture, as each stakeholder has different concerns and thus raises different issues. As a result of its richness, the NHSnet provides an excellent setting for investigating the potential to represent stakeholders through agents. However, the complexity of the stakeholder interrelations and the scope of this paper, which is primarily aiming at exploring the possibility to represent stakeholders using agents makes it necessary to confine our study to a limited number of stakeholders. Thus we will concentrate on a representation of those stakeholders printed in bold in Table 2 (*viz*. GPs, BMA, NHS-Ex and SecCon). As demonstrated below, this small subset of the broader stakeholders group provides ample opportunity for raising issues that are of interest to the application of agent technology.

4. Agents and inter-agent communication

The interests of a stakeholder represent a complex network of beliefs of various types, including facts about the world, values, opinions, hypothetical scenarios, and so on. Furthermore, each stakeholder also holds *second-order beliefs*, i.e. beliefs about their own opinions, the opinions of others, and how these inter-relate. Finally stakeholder analysis also offers an insight into the 'intentions' and 'desires' (to use agent rather than stakeholder terminology) of each stakeholder. Such a characterisation is directly amenable to agent-theoretic representation, whereby a stakeholder (or prototypical stakeholder) is modelled as an individual agent. Indeed, the notion of a stakeholder ("someone that affects or is affected by a system") is closely compatible with the more common definitions of an agent – although aspects of the latter are often left implicit in the definition of the former (self-interest and proactivity, for example, form intrinsic features of the notion of a stakeholder).

In order to clearly motivate the agent-theoretic approach, it is important to emphasise the aims of the model. Stakeholder analysis, as exemplified by the previous section, can be a lengthy, complex process involving a large number of individuals or representatives. As a result, the information drawn from the process is also an intricate, complex web of vested interests, bias, and opinions. It is clear that both the process and the result could benefit from a clear model – stakeholder analysis could then use the model to guide the design of interview material and further elicitation of potential stakeholders, while a model of the results could used to test the effects of particular decisions or developments within the interorganisational system. To perform such tasks, however, it is necessary to support integration of and communication between the interests of the stakeholders represented by the individual agents.

In order to filter and view subsets of stakeholders' interests, the model distinguishes between *stakeholder agents* (SA's), which model the belief structure of a stakeholder, and *issue agents* (IA's) which represent dynamic, proactive, information-seeking tasks for formulating and amalgamating views on particular issues. It is these IA's which effect the primary co-ordination of communication: an IA elicits views from the SA's, integrates those views into its current model of the issue, determines what further information is required, distributes drafts of the issue summary, and so on.

A sample scenario is given in Figure 1, below. The set of stakeholders holding a stake in a given issue is a subset of the stakeholders of the system as a whole: determining this set represents a preliminary phase of the work of an IA. The process of stakeholder identification (whereby stakeholders themselves may indicate other stakeholders) is mirrored in the work of the IA in determining its issue's stakeholders: results of the first round approximation may give rise to additional stakeholders being requested at the second round, and so on. (Clearly, the notion of a 'round' implies a temporal inflexibility which, although applicable to the slow, real-world communication in the stakeholder analysis process, is an undesirable feature of a system which does not suffer from such restrictions: in fact, there is no need to enforce this synchrony, since the IA can dynamically update its set of stakeholders whenever it receives information that such a change is warranted). Thus at a given stage in the process, an IA will have active lines of communication with identified stakeholders, and will be unaware of additional lines of communication which will become active at a later stage ('potential lines of communication').

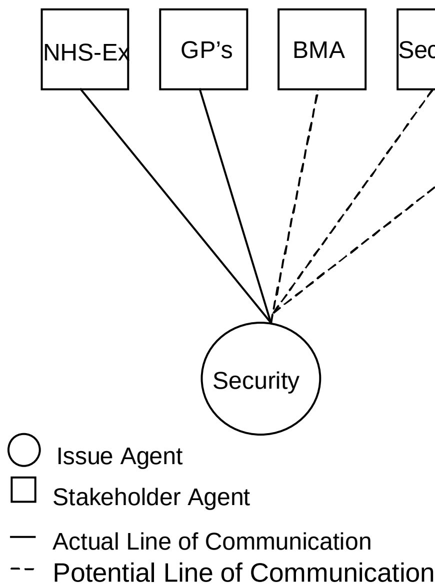

Fig. 1. Sample Scenario.

4.1 Stakeholder agents

The views and opinions of stakeholders can be construed as a network of beliefs with various support relationships holding between them: in a sense, it is appropriate to see these structures as *arguments*, albeit that these arguments are not voiced. Thus one belief may support, attack, or have no relation with another. In addition, the models must admit second-order beliefs, such that one belief may support, attack, be supported by or be attacked by another belief thought to be held by a separate agent. For example, the NHS-Ex agent has beliefs expressing (i) that the security of the system is better than before; (ii) that NHSnet possesses the ability to safeguard information; (iii) that GP's believe that NHSnet does not possess the ability to safeguard information. There is then a clear attacking relationship between (ii) and (iii). Finally, each belief is tagged with a number of issue markers – thus the three NHS-Ex beliefs mentioned above would be tagged with the *security* marker (amongst others – the first may also be tagged for *efficiency* or *cost*, for example). This argument-based means of representing information offers a number of benefits, including the ability for individual agents and the system as a whole to tackle the issues of uncertain and incomplete information which are characteristic of real world domains (see, for example Parsons and Jennings, 1996; Reed *et al.*, 1997; Elvang-Gøransson *et al.*, 1993; etc). In addition, the model is amenable to a number of well-understood extensions to increase its flexibility and expressiveness (see, for example, Reed and Long 1997a) – although the notion of argumentation employed here is simplified, it is easily extensible to cope with further complexity and subtleties in belief modelling (such as the notion of mutual belief, qualified belief and the absence of belief in other agents), and to enrich the communication process (Parsons and Jennings, 1996; Reed, 1998).

The partial characterisation of four stakeholders in the *security* issue is shown below in Figure 2. The NHS-Ex agent makes a case for the security of NHSnet, the conclusion of which is attacked by beliefs that NHS-Ex knows GPs to hold. In addition, the NHS-Ex agent also models the belief that GP's have an ulterior motive in safeguarding information (as discussed above in §3.1), which is attacked by NHS-Ex's own belief of the inappropriateness of the argument. The GPs agent has a variety of reasons for coming to its conclusion, including the conclusion of the BMA agent, which again has a number of sources for its argument, as does the BMA, which employs arguments from both SecCon and GPs. SecCon itself has two separate arguments, one concerning the relationship between NHS "insiders" (i.e. employees with the potential to abuse the system) and the presence of an authentication challenge, and one relating the worth of the data to the potential for misuse.

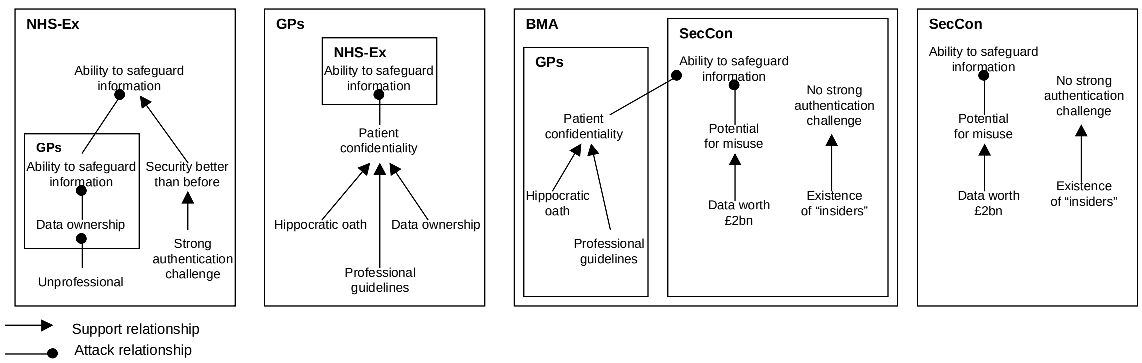

Fig. 2. Stakeholder agent representations.

4.2 Issue agents

The issue agents start with an initial (potentially singleton) set of stakeholders, from which the primary aim is to collate information and generate a summary of views. An IA does not aim to resolve any potential conflict it identifies – apart from the problems in trying to automate this process, there are a number of benefits in carefully maintaining the conflicting information: see, for example, Haggith's FORA system, (Haggith, 1996). In the current work – as in FORA – the emphasis is on representing the information in the debate, and thus clarifying the issues and views (and is thus comparable to work on disputation support, such as Gordon and Karacapilidis, 1997 and Jackson, 1997).

Thus an IA builds up a complete argument structure of units, annotated with the holder of the views. This argument structure is not necessarily composed entirely of units tagged by the marker for the IA's primary issue (i.e. the result of the *security* agent's work is not necessarily composed only of claims tagged with the *security* marker), since claims classified as being of one type may be supported by claims of quite another. An example of the structure created by the *security* agent on the basis of the small example in Figure 2 is given below in Figure 3. There are a number of potential techniques for presenting this argument structure in an appropriate way, though consideration of such techniques is beyond the scope of the current work (see Reed and Long, 1997b; Reed, 1997 for further exploration of this problem in relation to natural language, and Reed and Long, 1997a; Gordon and Karacapilidis, 1997 in relation to diagrammatic presentation).

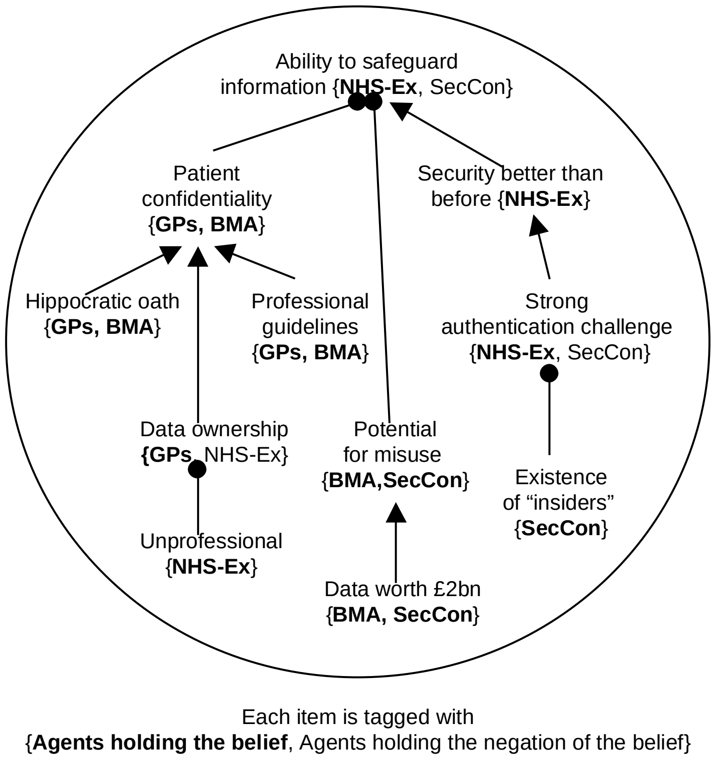

Fig. 3. The annotated argument structure generated by the IA.

Importantly, it is not possible to create this structure in a single pass; rather, it is necessary for the IA to initiate a number of dialogues with the stakeholders, during the course of which the various stakeholders are identified, and their beliefs are elicited and conflated. This process is discussed in more detail below.

4.3 Communication between IA's and SA's

Summary generation is – in the model as in the real world – an iterative process, wherein the IA produces a draft of the summary on incomplete information, and then offers that draft for consideration by the identified SA's, which can then offer modifications on an incremental basis. As mentioned above, this process relies on the IA soliciting two forms of information from the SA's – details of further SA's holding stakes in the IA's issue, and arguments pertaining to the IA's issue. This information can be requested and returned in tandem, with the former lying implicitly in the latter, where beliefs in the argument are second-order. For example, in the scenario illustrated above in Figure 2 and 3, the communication process might be described at an abstract level thus:

(i) The *security* IA, is initiated with two stakeholders identified, GPs and BMA, and polls GPs and BMA for information on the security issue (i.e. arguments for units tagged with the *security* marker).
(ii) The GPs returns the sub-argument against the NHS-Ex's "ability to safeguard information" claim. This subargument comprises a single claim, "patient confidentiality" supported by two further "the Hippocratic oath" and "professional guidelines". Notice that, as discussed below, an agent is not obliged to reveal its entire belief structure.
(iii) The BMA agent responds with the two arguments against the "ability to safeguard information" claim: that offered by the GPs, and that offered by the SecCon agent. The latter consists of the claim of "potential for misuse" supported by the "data worth £2bn" premises.
(iv) On receiving these replies, IA integrates the arguments of the BMA and GPs agents, and updates its list of relevant stakeholders to include NHS-Ex (from the GPs response) and SecCon (from the BMA response). It then sends requests to the NHS-Ex and SecCon agents.
(v) SecCon replies with the argument mentioned in (iii), as it has no more to add.
(vi) NHS-Ex replies with the subargument supporting the "ability to safeguard information" claim, and also the GPs argument with a counter to the GPs premise of "data ownership".
(vii) At this stage, the IA has polled all the stakeholders of whom it is aware, and can therefore generate a draft summary in the form of a complex, annotated argument structure. It then sends this draft to each of the stakeholders requesting comments.
(viii) The GP and NHS-Ex agents return the draft intact, as they possess no further information on any of its components. The BMA agent, however, on detecting the role of the "strong authentication challenge" unit in the NHS-Ex's argument, returns additional information to the effect that NHSnet doesn't offer protection in regard to "insiders".
(ix) The IA incorporates the BMA agent's new counterargument forming the complete structure in Figure 3; it sends this second draft to all relevant stakeholders and on receiving no alterations, terminates with its result.

At the final stage, another important change takes place – namely, that the NHS-Ex agent updates its model to incorporate the new-found views of the BMA; this dynamic nature of the SA's is closely allied to the analogous real-world process where stakeholders do not respond simply by querying some static set of beliefs, but rather alter their beliefs in the light of the opinions of other stakeholders (which may only become clear through the elicitation process). This belief update, however, does not represent the only dynamic aspect to SA's, for there is also important social and communicative structure holding between the SA's themselves.

4.4 Communication and social structure amongst SA's

From Figure 3 above, it can be seen that the BMA agent is rather different from the other SA's depicted, in that all it beliefs are second-order, belonging either to the security consultant SA, or to the GPs SA. This is unsurprising, since the remit of the BMA is to represent British GP's (and note that in a sense, the BMA is also representing the views of the security consultant). The BMA, then, acts as spokesperson for its various members, and as a result, the BMA agent can be seen to be actively requesting information from its member SA's in order to build a coherent argument. In building an argument structure composed of beliefs attributed to other agents, the BMA and other *association agents* (i.e. agents possessing only second order beliefs) are performing a similar task to that of the IA. There is, however, an important difference, since unlike the IA, the BMA agent is trying to produce a coherent (i.e. conflict free) argument – rather than simply summarise all the potentially disparate arguments of its members. The process by which it achieves such conflict resolution is beyond the scope of this paper, but key issues are addressed in Reed *et al*. (1997), Parsons and Jennings (1996), *inter alia*.

In addition, the process of association agents communicating with SA's to construct a coherent argument is dynamic, running concurrently with the IA/SA communication (as would be expected). In particular it explains why in the example above, the BMA agent does not at the first request return the "no strong authentication challenge" argument suggested by SecCon (as it would be expected to do, since the argument is tagged with the *security* marker). For at stage (vi), the BMA is *unaware* of this argument: not until the IA has offered the BMA it's first draft can the BMA pass on that draft to the SecCon agent, which then returns it with the new counterargument – which in turn, the BMA then passes back to the IA.

Furthermore, SA's may be unwilling to pass on some part of their argument structure to either IA's or other SA's. In the example above, the GPs agent only passes to the BMA two supports for the "patient confidentiality" claim: the third (that violating patient confidentiality impinges on the GP's notion of 'ownership' of patient data) is kept private (as it is liable to attack on the grounds of being unprofessional – as happens in the NHS-Ex model).

Finally, there remains one additional issue regarding one agent being the spokesperson of several others – the issue of abstraction. For the GPs agent itself is representing regional groups of GPs, which represent local practices, which ultimately represent individual GP's. And of course, individual GP's may also be represented through different routes (committees of the BMA, for example). This implies an intricate social structure amongst the SA's, with highly complex demands placed upon communication protocol. These problems represent an interesting direction for future formalisation, though the current approach is well suited to capturing such notions.

5. Conclusions and future research agenda

This paper has brought together two areas of research, that of stakeholder analysis and agent technology. Using the example of a network implemented in the healthcare context we have illustrated how stakeholders, their beliefs, and the issues emerging from their conflicting interests can be represented using two types of agents, stakeholder agents and issue agents. Due to its complexity, this example has been a useful pointer to both possibilities and limitations in this representation process.

The primary contribution of this paper has been to suggest a way of representing stakeholders and their complex interactions. Each representation is typically a 'snapshot' of relevant stakeholders for a given issue. Similarly, each issue agent represents the information that is available at a given point in time. Thus, it is possible to monitor the way in which the issues of interest are 'informed' by the stakeholders as well as how these shift over time. In this respect the agent representation offers an important tool for explaining and exploring the way in which the stakeholders' interests change over time, how these interests affect the issues of general importance, and how the stakeholders are affected by responses of other agents. This approach to representation addresses one of the shortcomings of the stakeholder analysis literature, namely the failure to represent stakeholders and their relations in a systematic, consistent and meaningful way. At the same time, it takes into account the principles of stakeholder behaviour and reflects the implications for their identification and analysis (cf. Table 1). For example, we have illustrated how agent representation can be used to expose implicit and explicit conflicts, and as a result, make it possible to follow and explain how different issues are reshaped and gain relative priority over time. This representation therefore shows how stakeholders alter their beliefs in response to the behaviour of other stakeholders. In a complex environment this framework of representation is of particular importance because it simplifies the interactions through consistent representation.

We have seen how an issue agent can accumulate and summarise similar or conflicting beliefs. Moreover, by creating this aggregate picture that is fed back to the stakeholder agents, the issue agent may generate further stakeholder reactions. For example, some stakeholders may respond to a problem situation by suggesting alternative solutions. This may generate new issue agents. We have seen how in our example the security issue has been reframed or led to the creation of an encryption issue, following the reactions of stakeholders. Such new issues may initially match the beliefs and interests of particular stakeholders but can also evolve due to reactions by other stakeholders. This will result in iteration of the stakeholder identification cycle followed by stakeholder and belief representation via new or updated stakeholder agents, summarisation in an issue agent that will generate further reactions and so on.

In practice, stakeholders will find this representation useful as it provides opportunities to find 'allies', i.e., stakeholder agents with similar interests and with which they can collaborate. At the same time, they may also find that they will need to make explicit different sets of beliefs, depending on the stakeholders with which they wish to collaborate and the issue in question. Furthermore, the possibility for similar representation is extremely important in a healthcare context similar to that presented in this paper. More specifically, the analysis of stakeholder interests through multi-agent representation presents a learning opportunity for those stakeholders that are likely to engage in similar debates. For example, the European Union has been pressing the Member States to harmonise their practices in relation to the way health information is exchanged. Thus, the approach suggested in this paper can facilitate the study of healthcare information exchange networks and thus offer important lessons for the development of similar projects even if these are to take place in a different environment.

A number of issues can be raised for further research. For example, it would be interesting to investigate whether interaction between stakeholder agents can occur through the intervention of issue stakeholders and what alternatives may be available. Another important point is the way in which stakeholders may choose to 'voice' their interests. It would be useful to monitor the way in which a 'summarised issue' agent can affect their future behaviour. One particularly interesting avenue is the possibility of new IAs being *spawned* by other agents (as opposed to being specified in the design phase) – an IA may recognise the importance of an issue (through the preponderance of markers of a particular type) and automatically create a new IA to prepare an appropriate summary on that issue; similarly SA's individually or in concert may be aware of relevant issues which should be addressed through the actions of a new IA (as in the security-encryption issues example mentioned above). The process of one agent spawning another may also have interesting applications in other areas of applied agent technology research, where the approach might afford greater flexibility through a reduction in the importance of the initial configuration produced in the system design phase.

Certainly, the agent representation is likely to place constraints on the richness of the actual stakeholder interaction. An interesting question that is open to future research is therefore a study of the extent to which agents can represent stakeholders and the issues that are of interest to them. This paper has illustrated the benefits of applying agent representation in a context of multiple stakeholders. Further work is required to continue analysis not only of the impact that agent based approaches have on modelling the interests of stakeholders in complex interorganisational systems, but also of the demands that such an application places on the underlying agent theory, and the enhancements to that theory required thereby.

References

95/46/EC European Union Directive. *On the Protection of Individuals with regard to the processing of personal data and on the free movement of such data*. 24 October 1995. OJ L281/31-50.

Anderson, R. J. (1996). Patient confidentiality - at risk from NHS-wide networking. In B. Richards (Ed.), *Current Perspectives in Healthcare Computing Conference*, (pp. 687-692). Harrogate: BJHC Ltd.

Barber, B. (1998). How should we treat personal medical data? *British Journal of Healthcare Computing and Information Management*, 13 (2), 10-11.

Bywater, M., & Wilkins, C. (1996). Mystic megabytes. *Journal of Healthcare Computing and Information Management*, 13 (2), 10-11.

Cash, J. I., & Konsynski, B. R. (1985). IS Redraws Competitive Boundaries. *Harvard Business Review* (March-April), 134-142.

Cohen, N. (1996, 5 May 1996). UK: Your life, on the record. *Independent on Sunday*, p. 21.

Creasey, D. (1996). Personal health information: uncertainty ahead. *British Journal of Healthcare Computing and Information Management*, 13 (5), 6.

Davies, S. (1996). Dystopia on the Health Superhighway. *The Information Society*, 12, 89-93.

Elvang-Gøransson, M., Krause, P. & Fox, J. (1993) Dialectic reasoning with inconsistent information. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence* (UAI'93), pp114-121

Freeman, R. E. (1984). *Strategic Management: A Stakeholder Approach*. Cambridge, Mass.: Ballinger Publishing Co.

Gillies, A. (1995). On the computerization of general practice in the UK: the IT perspective. *Journal of Information Technology*, 10 (2), 75-85.

Gordon, T. F. & Karacapilidis, N. I. (1997). The Zeno Argumentation Framework. In *Proceedings of the International Conference on Artificial Intelligence and Law*, Melbourne.

Haggith, M. (1996) *A meta-level argumentation framework for representing and reasoning about disagreement*. PhD Thesis, University of Edinburgh

Hayes, G. M. (1997). How computers influence medical decision making: for good or evil? In *Informed drug prescribing: issues and advances in an electronic age*. London.

Jackson, S. A. (1997). Disputation by design. In *Proceedings of the Second Conference of the Ontario Society for the Study of Argument* (OSSA97). Ontario.

Jennings, N. R., Faratin, P., Johnson, M. J., O'Brien, P. & Wiegand, M. E. (1996). Using Intelligent Agents to Manage Business Processes. In *Proceedings of the First International Conference on The Practical Application of Intelligent Agents and Multi-Agents*. London.

NHS Executive (1994a). *A strategy for NHS-wide networking* (No. E5155). Information Management Group.

NHS Executive (1994b). *This is the IMG: A guide to the Information Management Group of the NHS Executive* (No. B2126). Information Management Group.

NHS Executive (1995). *NHS-wide networking: application requirements specification* (No. H8003). Information Management Group.

NHS Executive (1996). *The use of encryption and related services with the NHSnet: A report for the NHS Executive by Zergo Limited* (No. E5254). Information Management Group.

Parsons, S. & Jennings, N. R. (1996) Negotiation Through Argumentation - A Preliminary Report. In *Proceedings of the International Conference on Multi Agent Systems (ICMAS'96)*, Kyoto, Japan.

Pouloudi, A. (1997). Stakeholder Analysis in UK Health Interorganizational Systems. In K. V. Andersen (Eds.), *EDI and Data Networking in the Public Sector* (chapter 5). Boston: Kluwer Academic Publishers.

Pouloudi, A., & Whitley, E. A. (1996). Stakeholder analysis as a longitudinal approach to interorganizational systems analysis. In J. D. Coelho, T. Jelassi, W. Koening, H. Krcmar, R. O'Callaghan, & M. Saaksjarvi (Eds.), *Proceedings of the 4th European Conference on Information Systems* (pp. 33-44). Lisbon, Portugal.

Pouloudi, A., & Whitley, E. A. (1997). Stakeholder identification in interorganizational systems: gaining insights for drug use management systems. *European Journal of Information Systems*, 6 (1), 1-14.

Reed, C. A. (1997, in press). Representing and Applying Knowledge for Argumentation in a Social Context. *AI & Society*.

Reed, C. A. (1998, in press). Dialogue Frames in Agent Communication. *Proceedings of the 3rd International Conference on Multi Agent Systems (ICMAS'98)*, IEEE Computer Society Press.

Reed, C. A. & Long, D. P. (1997a). Multiple subarguments in logic, argumentation, rhetoric and text generation. In Gabbay, D., Kruse, R. & Ohlbach, H. (Eds.) *Qualitative and Quantitative Practical Reasoning* (pp. 496-510). Boston: Springer Verlag.

Reed, C. A. & Long, D. P. (1997b). Content Ordering in the Generation of Persuasive Discourse. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI'97)*. Nagoya, Japan.

Reed, C.A., Long, D.P., Fox, M. & Garagnani, M. (1997) Persuasion as a Form of Inter-Agent Negotiation. In Lukose, D., Zhang, C., (Eds.), *Multiagent Systems*. Springer Verlag.

Richardson, B., & Richardson, R. (1992). *Business Planning: an approach to strategic management* (2nd ed.). Pitman.

Ruohonen, M. (1991). Stakeholders of strategic information systems planning: theoretical concepts and empirical examples. *Journal of Strategic Information Systems*, 1 (1), 15-28.

Sierra, C., Jennings, N. R. Noriega, P. & Parsons, S. (1997). A framework for argumentation-based negotiation. In *Proceedings of the International Workshop on Agent Theories, Architectures and Languages (ATAL97)*. Providence, Rhode Island, USA.

Turner, R. (1998). The Caldicott Committee reports. *British Journal of Healthcare Computing and Information Management*, 13 (2), 10-11.

Willcox, D. (1995, 19 October). Health scare. *Computing*, p. 28-29.